

# EgoFlowNet: Non-Rigid Scene Flow from Point Clouds with Ego-Motion Support

Ramy Batrawy<sup>1</sup>  
ramy.batrawy@dfki.de  
René Schuster<sup>1</sup>  
rene.schuster@dfki.de  
Didier Stricker<sup>1,2</sup>  
didier.stricker@dfki.de

<sup>1</sup> Augmented Vision  
German Research Center for  
Artificial Intelligence (DFKI)  
Kaiserslautern, Germany  
<sup>2</sup> Computer Science Department  
The University of  
Kaiserslautern-Landau (RPTU)  
Kaiserslautern, Germany

---

## Abstract

Recent weakly-supervised methods for scene flow estimation from LiDAR point clouds are limited to explicit reasoning on object-level. These methods perform multiple iterative optimizations for each rigid object, which makes them vulnerable to clustering robustness. In this paper, we propose our EgoFlowNet – a **point-level** scene flow estimation network trained in a weakly-supervised manner and without object-based abstraction. Our approach predicts a binary segmentation mask that implicitly drives two parallel branches for ego-motion and scene flow. Unlike previous methods, we provide both branches with all input points and carefully integrate the binary mask into the feature extraction and losses. We also use a shared cost volume with local refinement that is updated at multiple scales without explicit clustering or rigidity assumptions. On realistic KITTI scenes, we show that our EgoFlowNet performs better than state-of-the-art methods in the presence of ground surface points.

## 1 Introduction

Scene flow estimation is an important computer vision problem for navigation, planning, and autonomous driving systems. It provides a representation of the dynamic environment by estimating the 3D motion field relative to the observer. Until a few years ago, stereo images were used for joint disparity estimation and optical flow estimation to represent scene flow [6, 20, 28, 34, 35]. However, the two-view geometry used in self-driving cars has inherent limitations, such as inaccuracies in depth estimation in distant regions.

With the advent of LiDAR, many learning-based methods have been developed to estimate scene flow directly from point clouds in a fully-supervised manner [6, 15, 22, 24, 29]. They differ from each other in their basic feature extraction framework and the way they design their cost volume. Due to the lack of annotated data on realistic sequences, some methods train their end-to-end models with self-supervised losses [22, 24, 31, 39].




Figure 1: Our EgoFlowNet operates non-rigidly at the **point-level** and shows high accuracy for regions of varying local density (*e.g.*, the red, blue, and green rectangles).

Apart from the point-wise estimation of scene flow, some methods perform better when using self-supervised losses under conditions of rigidity [9, 25]. Other methods support scene flow estimation with ego-motion [2, 40, 43]. All of the above methods work well on ideal conditions (*e.g.*, no ground points, no occlusions, or with nearly direct correspondences between consecutive scenes).

A recent breakthrough has been achieved by WSLR [13], where a multi-task prediction network is designed and trained with real scenes in a weakly-supervised manner in the presence of ground points. This approach segments the scene into static parts (*i.e.* background (BG)) and moving agents (*i.e.* foreground (FG)). It then optimizes the initial estimate of ego-motion and scene flow via non-parametric object-based optimizations using explicit rigidity constraints. Towards learning-based optimization, ERC [10] uses the predicted segmentation mask from [13] and proposes a novel optimization method with an error-driven Gated Recurrent Unit and residual scene flow heads. These methods [10, 13] show impressive results for more difficult scenes (*e.g.* with ground points, outliers, occlusions, *etc.*). However, both methods rely on the DBSCAN clustering algorithm [10], which may limit their ability to work on low-density regions or under-sampled objects (*e.g.*, distant cars or small objects). In addition, they must perform iterative optimizations for each clustered region, which negatively impacts efficiency when the scene contains a large number of clusters.

Compared to these methods, our approach is far removed from any clustering strategy and instead predicts unconstrained scene flow at the point-level. To this end, we design our multi-task network to predict a binary FG/BG segmentation mask, which is then carefully used to estimate ego-motion and scene flow (*c.f.* Figure 1). Unlike object-based methods, we feed the ego-motion and scene flow branches with all input points, integrate our predicted mask into both branches and combine everything with point-based coarse-to-fine refinement to obtain accurate scene flow. For robust estimation in both branches, we also develop a hybrid feature extraction to provide both branches with well-suited features.

Our contributions are summarized as follows:

- We propose EgoFlowNet – a multi-task neural network architecture to estimate scene flow directly from *raw* point clouds that jointly estimates binary segmentation masks, ego-motion, and scene flow.
- We propose a hybrid feature extraction along with a hybrid warping layer and integrate the binary masks to obtain robust scene flow.
- We work with a point-level refinement of the scene flow, which is free of explicit clustering mechanisms or rigidity assumptions for dynamic objects.
- On difficult real LiDAR scenes (*i.e.*, with ground points, occlusions, and outliers), we show that our proposed approach outperforms recent clustering-based methods.

## 2 Related Work

3D scene flow was first introduced in the image domain using RGB-D [18, 19, 63, 67] for indoor scenarios and stereo images [6, 17, 20, 28, 65, 69, 41] for outdoor scenarios. However, learning scene flow directly from point clouds without relying on RGB images opens up a wide field of research [6, 8, 10, 13, 14, 15, 22, 23, 26, 32, 42, 46, 48, 49].

**GRU-based Scene Flow from Point Cloud:** The Gated Recurrent Unit (GRU) [7] is used to iteratively refine the global cost volume to provide an accurate estimate of the scene flow [69]. FlowStep3D [27] updates the cost volume locally using GRU with multiple reconstructions and iterative point cloud alignment. To encode a large correspondence set within the cost volume, PV-RAFT [48] combines a voxel representation with a point-wise cost volume. A point-wise optimization combined with a recurrent network regularization is proposed by RCP [14]. Our EgoFlowNet avoids strict iterative updates and works from coarse-to-fine, driven by a binary segmentation mask and jointly estimates scene flow and the ego-motion.

**Hierarchical Scene Flow from Point Cloud:** FlowNet3D [26] is the first work to introduce a cost volume layer from a point cloud with hierarchical refinement. However, it is limited to a single cost volume layer. To overcome this limitation, HPLFlowNet [15] introduces multi-scale correlation layers by projecting points into a permutohedral grid [68]. Moving away from the grid representation, PointPWC-Net [49] improves the direct estimation of scene flow from *raw* point clouds by constructing cost volumes at a range of scales from coarse-to-fine. Following the hierarchical point-based designs, intensive improvements are proposed in the development of cost volume using dual attentions as in [1, 47, 44, 45]. Our network is basically hierarchical, but integrates further multi-task estimates of ego-motion and segmentation. It operates in challenging outdoor scenes with typical occlusions and in dense scenes with ground points.

**Scene Flow from Point Cloud with Constraints:** Axiomatic concepts of rigidity assumptions are explored in [4, 25] along with cluster-based or object-level optimization. However, the above methods are not well explored with typical outdoor scenes in the presence of ground surface points. More recently, WSLR [13] has proposed pioneering weakly-supervised learning along with non-parametric optimization, and ERC [10] extends this to learning-based optimization. Both work well on challenging outdoor scenes where ground points are present. However, both require multiple optimization steps and work under object-level constraints using DBSCAN clustering [11]. Chodosh *et al.* [8] is a very recent conventional and cluster-based method that works by test time optimization using ICP [4, 66] and RANSAC to achieve appropriate piece-wise rigidity. In contrast to the above methods, we do not use clustering algorithms and work with point-level optimization, which allows us to estimate non-rigid motion and is more accurate and robust than state-of-the-art methods.

## 3 Network Design

Our EgoFlowNet estimates scene flow as translational vectors from two consecutive frames of point clouds, with no assumptions about object rigidity. Given Cartesian 3D point cloud frames  $P$  and  $Q$  at timestamps  $t$  and  $t + 1$ , our goal is to estimate point-wise 3D flow vectors  $\hat{S}$  for each point within  $P$ . Our network is designed to combine segmentation, ego-motion, and scene flow tasks at four scales  $\{L\}_{k=0}^3$ , where  $l_0$  is the full resolution of  $P$  and  $Q$ . An illustration of our hierarchical modules of our feature extraction, cost volume, ego-motion

and scene flow is given in [Figure 2](#), where the right part of the figure illustrates a single layer or scale  $l_k$  for each of these modules. The following sections describe the components of each module in detail.

### 3.1 Feature Extraction

Our feature extraction module consists of two networks: The first one is an encoder-decoder module, while the second one consists only of a context encoder. The backbone of our feature extraction is inspired by RandLA-Net [\[16\]](#).

**Encoder Module:** Each scale in the encoder module essentially consists of two layers, where Local-Feature-Aggregation (LFA) is applied to aggregate the features at the  $l_k$  scale, followed by a Down-Sampling (DS) layer to aggregate the features from the  $l_k$  level to  $l_{k+1}$ , resulting in a simultaneous decrease in resolution. The backbone of the LFA is inspired by RandLA-Net [\[16\]](#), which uses attentive pooling based on self-attention as in [\[50, 51\]](#). At all scales, we search for 16 neighbors in Euclidean space using K-Nearest-Neighbors (KNNs), then their weighted features are summed based on attentive pooling. The DS layer samples the points based on Farthest-Point-Sampling (FPS) to the defined resolution  $l_{k+1}$ , and aggregates 16 nearest neighbors in the higher resolution  $l_k$  for each selected sample simply by using Max-Pooling.

**Decoder Module:** The decoder module of the hourglass network consists of  $\{L\}_{k=0}^3$  layers for extracting the features up to the full (input) resolution  $l_0$  of  $P$  and  $Q$ , respectively. To up-sample from the  $l_{k+1}$  level to  $l_k$ , we simply assign the one nearest neighbor for each point of the higher resolution to the lower one, followed by a simple Multi-Layer Perceptron (MLP). To increase the quality of the features in the encoder-decoder network, lateral connections are added to each layer.

**Segmentation Features:** The encoder-decoder of the first network extracts the features  $F_{s,0}^P$  and  $F_{s,0}^Q$  at the input resolution, which are used to predict the binary segmentation masks ( $M_{fg}^P$  and  $M_{fg}^Q$ ) and/or ( $M_{bg}^P$  and  $M_{bg}^Q$ ) for  $P$  and  $Q$ , respectively.

**Hybrid Features:** The encoder module of the first network and the context network compute the features ( $F_{encoder,k}^P$ ,  $F_{encoder,k}^Q$ ) and ( $F_{context,l}^P$ ,  $F_{context,k}^Q$ ), at each scale level  $l_k$ . The encoder modules down-sample the input to the resolutions  $l_1 = 2048, l_2 = 512, l_3 = 128$  with feature dimensions  $C_0 = 32, C_1 = 128, C_2 = 256, C_3 = 512$ . The output features of the two encoders at each scale level  $l_k$  are merged using the predicted and down-sampled segmentation masks  $M_{fg,k}$  as follows:

$$HF_k = M_{fg,k} \cdot F_{context,k} + (1 - M_{fg,k}) \cdot \perp (F_{encoder,k}), \quad (1)$$

where  $M_{fg,k}$  refers to the binary mask of foreground points,  $(1 - M_{fg,k})$  refers to the background mask (*i.e.*,  $M_{bg,k}$ ) and  $\perp$  is an operator that sets the gradient of the operand to zero,  $\nabla(x) = 0$  (*i.e.*, stop gradient). Since the number of background points (BG) within a scene, including the ground points, is usually much higher than the number of foreground points (FG), using the stop gradient eliminates the negative effect of the ego-motion branch on the segmentation head and the scene flow branch. By merging the context encoders, the features of the FG points can be enhanced to provide an accurate estimate of scene flow for these points. We apply [Eq. \(1\)](#) at each scale level  $l_k$  using  $M_{fg,k}^P$  and  $M_{fg,k}^Q$ , resulting in  $HF_k^P$  and  $HF_k^Q$  for  $P$  and  $Q$ , respectively, which are then used for the shared cost volume.

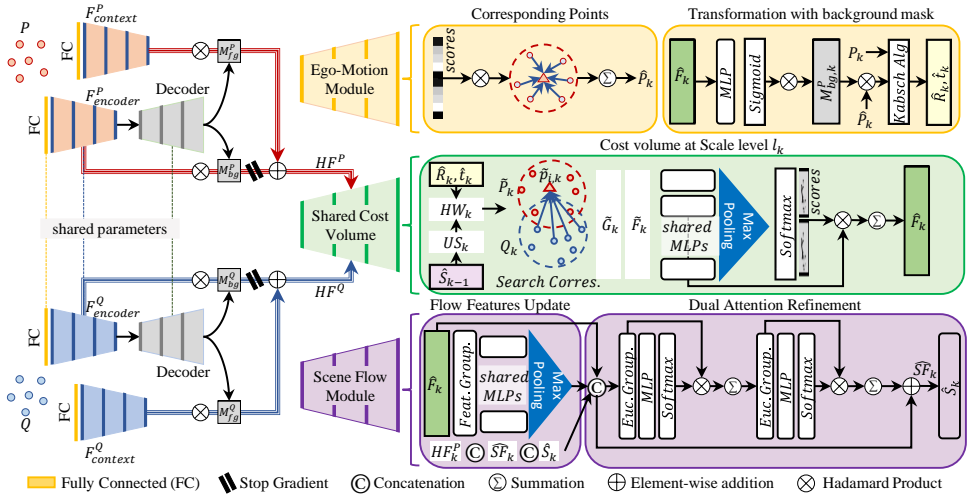


Figure 2: Our EgoFlowNet architecture predicts a binary segmentation masks ( $M_{fg}^P$  and  $M_{fg}^Q$ ) for foreground points ( $FG$ ) and ( $M_{bg}^P$  and  $M_{bg}^Q$ ) for background points ( $BG$ ). We use the binary mask to jointly estimate ego-motion and scene flow at the point-level. For this, we extract hybrid features ( $HF^P$  and  $HF^Q$ ) and hierarchically refine our point-wise scene flow.

### 3.2 Segmentation Head

We apply three layers of Multi-Layer Perceptions (MLPs) with 64, 32 and 1 output channels to the computed segmentation features  $F_{s,0}^P$  and  $F_{s,0}^Q$ . The output of the last layer provides the segmentation probabilities at full (input) resolution layer  $l_0$ , allowing us to define the binary segmentation masks  $M_{fg}^P$  and  $M_{fg}^Q$  for  $P$  and  $Q$ , respectively.

### 3.3 Shared Cost Volume

We learn the geometric and feature correlations based on the hybrid features  $HF_k^P$  and  $HF_k^Q$  (c.f. Eq. (1)).

**Searching for Correspondences:** As a first step, we need to find the correlation set in  $Q_k$  for each point in  $P_k$ . Since finding correlations based on Euclidean space may not be sufficient to capture distant correspondences, we use the feature space to find the correspondences at the coarsest scale  $l_3$ . This provides a high quality initial estimate of the scene flow  $\hat{S}_3$  and a high quality initial estimate of the ego-motion parameters represented by the rotation  $\hat{R}_3$  and translation  $\hat{t}_3$  components. For the scene flow, when searching for correspondences in feature space, the distant matches on the upper scales are approximated by our hybrid warping layer so that the warped point cloud  $\tilde{P}_k$  is close to its match in  $Q_k$ . With this initialization, it becomes worthwhile to search for the closest 16 matches in Euclidean space for the upper scales  $\{L\}_{k=0}^2$ .

After grouping the correspondence set with its geometric features  $\{q_{j,k}\}_{j=1}^{16}$  and hybrid features  $\{hf_{j,k}^q\}_{j=1}^{16}$ , we compute the differences to the point  $P_k$  and its hybrid feature  $hf_{i,k}^p$ , respectively. This yields the geometric and feature differences  $\tilde{G}_k$  and  $\tilde{F}_k$ , which are then concatenated. We apply Max-Pooling along the feature dimension to compute attentive weights

similar to HRegNet [27]. The geometric and feature differences are then smoothly weighted by the attentional weights and summed to obtain  $\hat{F}_k$ .

**Hybrid Warping Layer:** Our hybrid warping layer ( $HW_k$ ) is jointly driven by the ego-motion, the scene flow, and the predicted segmentation  $M_{fg,k}^P$  of frame  $P_k$ . After obtaining the initial ego-motion and the initial scene flow at the coarse scale  $l_3$ , we apply hybrid warping to refine the estimate at the upper scales  $\{L\}_{k=0}^2$ . For this purpose, we use the corresponding binary masks ( $M_{fg,k}^P$  and  $M_{bg,k}^P$ ), the up-sampled scene flow  $\hat{S}_k$  from the coarser scale through a simple Up-Sampling layer ( $US_k$ ), and the ego-motion transformation to warp the points in  $P_k$  towards the target  $Q_k$  and obtain  $\tilde{P}_k$ . For all upper scales, we apply the following equation:

$$\tilde{P}_k = M_{fg,k} \cdot (P_k + \hat{S}_k) + (M_{bg,k} \cdot ((\hat{R}_k P_k^T)^T + \hat{t}_k)) \quad (2)$$

### 3.4 Ego-Motion branch

We compute point correspondences and apply the Kabsch algorithm [28] to estimate the ego-motion parameters  $\hat{R}_k$  and  $\hat{t}_k$ .

**Corresponding Points:** Inspired by HRegNet [27], after obtaining the correspondence set  $\{q_{j,k}\}_{j=1}^{16}$  for each point in  $Q_k$  as described in the cost volume, we multiply the computed attentive weights with them and sum over the nearest neighbors to obtain the corresponding points  $\hat{P}_k$ .

**Optimal Transformation:** Multi-Layer Perceptrons (MLPs) are applied to the cost volume output  $\hat{F}_k$ , followed by a Sigmoid function to obtain confidence values inspired by [10, 13, 27]. However, unlike previous approaches [10, 13], we do not filter out the ( $FG$ ) points to feed the ego-motion branch with only ( $BG$ ) points. Instead, we feed this branch with all points and multiply the confidence values by  $M_{bg,k}^P$  to refine the corresponding points  $\hat{P}_k$  so that the transformation matrix can be computed according to [28] to obtain  $\hat{R}_k$  and  $\hat{t}_k$ .

### 3.5 Scene Flow branch

Across all scales, our scene flow branch consists of three refinement stages and four scene flow predictors with simple nearest-neighbor Up-Sampling (US). The total number of layers with attention-based refinement is inspired by RMS-FlowNet [10], which is designed to estimate scene flow only, but we add three feature updating units.

**Feature Updates:** With the obtained cost volume features  $\hat{F}_k$ , we search for the 16 nearest neighbors in the feature space, group them and then apply MLPs followed by Max-Pooling. This helps to capture similar features and implicitly extends features to semantic objects as inspired by DGCNN [27].

**Dual Attention Refinement:** We concatenate the updated features  $\hat{F}_k$  with  $HF_k^P$ ,  $\hat{S}_k$  and  $\hat{S}_k$ , where the latter two components are the scene flow features and the scene flow, respectively. Both are initialized to zero in the coarse layer  $l_3$  and are only used in the upper scales  $\{L\}_{k=1}^2$ . We use the defined nearest neighbors (16) in Euclidean space to group the concatenated features and we apply dual attentions to refine the corresponding features as performed in [10].

**Scene Flow Predictor:** Our EgoFlowNet predicts scene flow at multiple scales, inspired by [10, 29]. The scene flow estimation head takes the resulting scene flow features at each scale  $\hat{S}_k$  and applies three layers of MLPs with 64, 32 and 3 output channels. Then, the estimated scene flow  $\hat{S}_k$  and the features from the attention-based refinement are up-sampled to the next higher scale using a simple KNN search.

### 3.6 Scene Flow of BG Points

At the input point resolution  $l_0$ , we compute the scene flow of the background  $BG$  from the predicted rotation and translation (i.e.,  $\hat{R}_0$  and  $\hat{t}_0$ ) of the ego-motion branch. We use the binary segmentation mask  $M_{bg}^P$  to merge the scene flow of the  $BG$  points with the output of the scene flow  $\hat{S}_0$  obtained from the scene flow branch.

### 3.7 Loss Function

To guide the training of segmentation, ego-motion, and scene flow, we combine three losses:

$$\mathcal{L}_{total} = \mathcal{L}_{seg} + \mathcal{L}_{ego} + \mathcal{L}_{sf}, \quad (4)$$

**Segmentation Loss:** We use the Weighted Binary Cross-Entropy loss to overcome the severe imbalance of  $FG$  and  $BG$  classes as follows:

$$\mathcal{L}_{seg} = -\frac{1}{N} \sum_{i=1}^N \gamma \cdot y_i \cdot \log(\sigma(p(y_i))) + (1 - y_i) \cdot \log(1 - \sigma(p(y_i))), \quad (4)$$

where  $i$  is the index in  $P$  or  $Q$ ,  $y_i$  is the ground truth label,  $p(y_i)$  is the probability of the predictions and  $\gamma$  is the  $FG$  class weight, which is set to 20.

**Ego-Motion Loss:** Inspired by [27], the ego-motion loss is computed hierarchically. Given a four-scale estimate of the transformation parameters  $\hat{R}_k$  and  $\hat{t}_k$  and the ground truth  $R$  and  $t$ , we compute the final ego-motion loss as follows:

$$\mathcal{L}_{ego} = \frac{1}{4} \sum_{k=0}^3 \beta \|\hat{R}_k^T R - I\|_2 + \|\hat{t}_k - t\|_2 \quad (5)$$

where  $\|\cdot\|_2$  denotes the  $L_2$ -norm and  $\beta$  is set to 1.8.

**Scene Flow Losses:** To train the scene flow branch, we apply a bidirectional Chamfer loss  $\mathcal{L}_{cd,k}$  and Smoothness loss  $\mathcal{L}_{sm,k}$  per scale, both driven by  $M_{fg,k}^P$  and  $M_{fg,k}^Q$  as follows:

$$\mathcal{L}_{cd,k} = \sum_{\tilde{p}_k \in \tilde{P}_k} m_{fg,k}^P \cdot \min_{q_k \in Q_k} \|\tilde{p}_k - q_k\|_2 + \sum_{q_k \in Q_k} m_{fg,k}^Q \cdot \min_{\tilde{p}_k \in \tilde{P}_k} \|\tilde{p}_k - q_k\|_2 \quad (6)$$

$$\mathcal{L}_{sm,k} = \sum_{p_k^i \in P_k} m_{fg,k}^P \cdot \frac{1}{N_k(p_k^i)} \sum_{p_k^j \in N_k(p_k^i)} \|\hat{S}_k(p_k^j) - \hat{S}_k(p_k^i)\|_1 \quad (7)$$

where  $\|\cdot\|_1$  denotes the  $L_1$ -norm, and the number of neighborhood points are  $N_0 = 16, N_1 = 12, N_2 = 8, N_3 = 4$ . Both losses are then combined as follows:

$$\mathcal{L}_{sf} = \sum_{k=0}^3 \alpha_k (\mathcal{L}_{cd,k} + \mathcal{L}_{sm,k}), \quad (8)$$

and the weights per scale are  $\alpha_0 = 0.02, \alpha_1 = 0.04, \alpha_2 = 0.08, \alpha_3 = 0.16$ .

## 4 Experiments

First, we give a brief description of the data sets and metrics used for evaluation. We also demonstrate the accuracy of the method in comparison to state-of-the-art methods. Finally, there is a verification of our design choices.



## 4.1 Evaluation Metrics

Let  $\hat{S}$  denotes the predicted scene flow, and  $S$  denotes the ground truth scene flow. The evaluation metrics for the 3D motion are averaged over all points and computed as follows:

- **EPE3D [m]**: The 3D end-point error computed in meters as  $\|\hat{S} - S\|_2$ .
- **Acc3DS [%]**: The strict 3D accuracy which is the ratio of points whose EPE3D  $< 0.05 m$  or relative error  $< 5\%$ .
- **Acc3DR [%]**: The relaxed 3D accuracy which is the ratio of points whose EPE3D  $< 0.1 m$  or relative error  $< 10\%$ .
- **Out3D [%]**: The ratio of outliers whose EPE3D  $> 0.3 m$  or relative error  $> 10\%$ .

To evaluate the predicted ego-motion parameters (*i.e.*,  $\hat{R}$  and  $\hat{t}$ ), compared to the ground truth ( $R$  and  $t$ ), respectively, we report the following metrics averaged over all the consecutive scenes:

- **RAE [°]**: The relative angular error computed in degrees as:  $\arccos(\text{Tr}(\hat{R}^T R - 1)/2)$ .
- **RTE [m]**: The relative translation error computed in meters as  $\|\hat{t} - t\|_2$ .

## 4.2 Data Sets and Preprocessing

As with all related methods, the point clouds generated from the following data sets are randomly sub-sampled to be evaluated at a defined resolution (*e.g.*, 8192 points) and are shuffled in a random order to resolve possible correlations between consecutive point clouds. We evaluate all of the methods in the different versions of KITTI that are described below. All data include ground surface points.

**semKITTI [3]**: It contains semantic labels of point clouds and ego-motion ground truth, including many sequences of real-world autonomous driving scenes. WSLR [13] has created a preprocessed version of this data set, including large sequences for training and a test split.

**stereoKITTI [30]**: This is a real scene flow data set with scene flow labels. As with most LiDAR-based methods, it is preprocessed using HPLFlowNet [15]. This processing creates direct correlations across the consecutive scenes, and exhibits non-uniform point cloud density.

**lidarKITTI [12]**: Unlike stereoKITTI, the consecutive point clouds in this data set are not in direct correspondence and some points have typical occlusions. The scene flow vectors of the ground truth are obtained by mapping the points to the corresponding pixels in the stereoKITTI data set. The point clouds have a non-uniform density that mimics the sampling pattern of a typical LiDAR scan. We use exactly the preprocessed and published data from WSLR [13].

## 4.3 Comparison to State-of-the-Art

To demonstrate the accuracy of our model, we compare our segmentation, ego-motion estimation, and scene flow estimation with state-of-the-art methods.

**Segmentation and Ego-Motion:** We compare the predicted mask and ego-motion estimates of our EgoFlowNet with the pioneering work of WSLR [13], which is the first to jointly predict binary segmentation, ego-motion, and scene flow for point clouds in a single network. The comparison is shown in Table 1. The results shown by WSLR [13] are the best optimized results presented in their paper, obtained by pre-training on FlyingThings3D [29] and subsequent further training on semKITTI [3]. With the exception of  $FG$  precision, our model trained from scratch on the same training split of semKITTI outperforms WSLR on all



Table 1: The segmentation accuracy of our EgoFlowNet generalizes better to lidarKITTI and shows better results for ego-motion estimation (*i.e.*, RAE and RTE) on semKITTI.

Method	semKITTI [8]						lidarKITTI [12]			
	prec. FG ↑ [%]	rec. FG ↑ [%]	prec. BG ↑ [%]	rec. BG ↑ [%]	RAE ↓ [°]	RTE ↓ [m]	prec. FG ↑ [%]	rec. FG ↑ [%]	prec. BG ↑ [%]	rec. BG ↑ [%]
WSLR [12]	<b>0.950</b>	0.892	0.991	<b>0.996</b>	0.116	0.029	0.734	0.855	0.991	<b>0.980</b>
Ours	0.898	<b>0.922</b>	<b>0.997</b>	<b>0.996</b>	<b>0.097</b>	<b>0.024</b>	<b>0.797</b>	<b>0.887</b>	<b>0.992</b>	0.975

Table 2: We outperform point-wise models that are fully supervised and methods that optimize for rigid motion at the object-level [8, 10, 13].

Data Set	Method	Sup.	Rigid.	stereoKITTI [8]				lidarKITTI [12]			
				EPE3D ↓ [m]	Out3D ↓ [%]	Acc3DS ↑ [%]	Acc3DR ↑ [%]	EPE3D ↓ [m]	Out3D ↓ [%]	Acc3DS ↑ [%]	Acc3DR ↑ [%]
FT3D, [13]	PointPWC-Net [13]	full	✗	0.204	0.645	0.292	0.556	0.710	0.932	0.114	0.219
	FlowStep3D [13]	full	✗	0.109	0.391	0.577	0.765	0.797	0.929	0.087	0.184
	RMS-FlowNet [8]	full	✗	0.199	0.547	0.391	0.618	0.652	0.920	0.120	0.233
	WM3D [13]	full	✗	0.119	0.487	0.488	0.721	0.646	0.928	0.165	0.270
	Bi-PointFlowNet [8]	full	✗	0.135	0.439	0.578	0.760	0.686	0.905	0.179	0.268
Chodosh <i>et al.</i> [8]	None	✓		-	-	-	-	0.061	-	0.917	0.962
semKITTI [8]	WSLR [12]	Weak	✓	0.068	0.263	0.836	0.897	0.080	0.369	0.742	0.850
	ERC [10]	Weak	✓	0.053	0.269	0.858	0.917	0.065	0.290	0.857	0.940
	Ours	Weak	✗	<b>0.039</b>	<b>0.212</b>	<b>0.922</b>	<b>0.966</b>	<b>0.049</b>	<b>0.267</b>	<b>0.918</b>	<b>0.964</b>

other segmentation metrics. Since we predict point-wise scene flow for the  $FG$ , errors in the segmentation have less impact compared to other methods that predict rigid object motion. That said, our segmentation generalizes better to lidarKITTI and outperforms WSLR in the  $FG$  category. In addition, our ego-motion estimation (*i.e.*, rotation RAE, and translation RTE errors) on semKITTI surpasses that of WSLR [12].

**Scene Flow:** The ultimate goal of our model is to predict the scene flow for each input point in the scene with respect to point cloud  $P$ . To this end, we evaluate our final scene flow estimate against point-wise methods [10, 8, 13, 14, 15], which are state-of-the-art methods that perform best on stereoKITTI when ground points are omitted. However, the accuracy of these methods is severely limited in the presence of ground points on stereoKITTI and even worse on lidarKITTI, a data set that resembles real LiDAR scenes with occlusions and no direct correspondences between successive LiDAR scans. Our point-based scene flow is comparable to the latest conventional method [8] on lidarKITTI, which integrates the ego-motion and rigidity assumptions into the scene flow estimation, but our method performs significantly better with respect to EPE3D. We also outperform the object-based weakly supervised methods WSLR [12] and ERC [10] on both KITTI versions in all metrics (*c.f.* Table 2). We also visualize our qualitative results on lidarKITTI in Figure 3. Further qualitative results can be found in the supplementary material.

**Efficiency:** Our model contains about 16 million parameters. For 8192 input points, it requires 11 GFLOPs, which takes an average of 140ms for a pair of point clouds on a single NVIDIA Titan V.

## 4.4 Ablation Study

We conduct several experiments by training on semKITTI and evaluating on lidarKITTI to verify our design decisions. To do this, we evaluate  $FG$  and  $BG$  points separately in  $EPE3D_{fg}$  and  $EPE3D_{bg}$ , respectively. We build our baseline design by computing our features from the encoder of the first feature extraction network ( $F_{encoder}$ ), so that no context encoder module and hybrid features are applied, and we do not consider binary masks for any of our network branches (*i.e.*, ego-motion and scene flow) and our warping layer is computed based on the scene flow branch only. Then, we check the basic estimates of our model by applying only

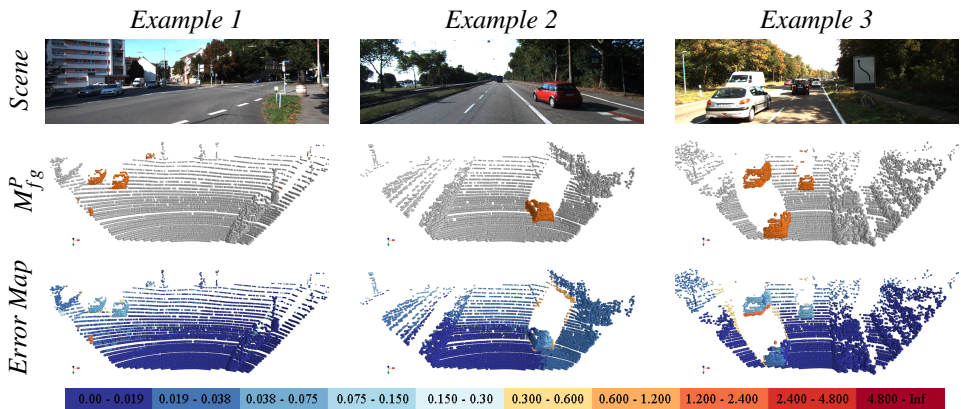


Figure 3: Three examples from lidarKITTI [12] show the qualitative results of our EgoFlowNet. For visual enhancement only, we show the RGB images of each scene. We visualize the predicted binary mask, where *BG* and *FG* points are encoded by gray and orange colors, respectively. The error map for each scene (third row) shows the end-point error in meters and is colored according to the map shown in the last row. Our EgoFlowNet shows low errors (dark blue) over a wide area in each scene, including *FG* and *BG* points.

the proposed losses (c.f. Eq. (3)), as shown in the 1<sup>st</sup> row in Table 3. In the 2<sup>nd</sup> row, we see the positive influence of integrating the predicted background mask  $M_{bg}^P$  into the ego-motion branch. In the 3<sup>rd</sup> row, we apply our hybrid warping layer as in Eq. (2), which further improves the results. Adding scene flow feature updating and dual attention refinement significantly improves the results for all metrics, as can be clearly seen in rows 4<sup>th</sup> and 5<sup>th</sup>. Without the *stop gradient* in Eq. (1), the results are the same as after using refinements, but applying it shows an improvement in  $EPE3D_{fg}$  and consequently in all other metrics. We provide further experiments in the supplementary material.

## 5 Conclusion

We propose EgoFlowNet, which predicts binary segmentation masks for dynamic and static LiDAR-based scenes and jointly estimates hierarchical ego-motion and scene flow. Our method works by estimating scene flow at the point-level rather than optimizing it at the object level. Our network is free of any clustering and uses point-level refinement, which produces better results than competing methods and allows for non-rigid object motions. Our approach outperforms recent approaches that rely on the object-level and shows robust accuracy in the presence of ground points.

Table 3: We investigate the contribution of each component in our design on lidarKITTI [12].

$M_{bg}^P$	Hybrid Warping	Feature Update	Attention Refinement	Hybrid Features	Hybrid Features $\perp$	$EPE3D \downarrow$ [m]	$EPE3D_{fg} \downarrow$ [m]	$EPE3D_{bg} \downarrow$ [m]	$Acc3DR \uparrow$ [%]
$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.168	0.375	0.154	0.689
$\checkmark$	$\times$	$\times$	$\times$	$\times$	$\times$	0.139	0.363	0.119	0.808
$\checkmark$	$\checkmark$	$\times$	$\times$	$\times$	$\times$	0.085	0.335	0.065	0.871
$\checkmark$	$\checkmark$	$\checkmark$	$\times$	$\times$	$\times$	0.067	0.326	0.048	0.931
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	$\times$	0.053	0.287	0.035	0.952
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	0.053	0.282	0.036	0.957
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	<b>0.049</b>	<b>0.267</b>	<b>0.033</b>	<b>0.964</b>

## Acknowledgment

This work was partially funded by the Federal Ministry of Education and Research Germany under the project DECODE (01IW21001) and partially in the funding program Photonics Research Germany under the project FUMOS (13N16302).

## References

- [1] Ramy Batraway, René Schuster, Mohammad-Ali Nikouei Mahani, and Didier Stricker. RMS-FlowNet: Efficient and Robust Multi-Scale Scene Flow Estimation for Large-Scale Point Clouds. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [2] Aseem Behl, Despoina Paschalidou, Simon Donn e, and Andreas Geiger. Point-FlowNet: Learning Representations for Rigid Motion Estimation from Point Clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [4] Paul J Besl and Neil D McKay. A Method for Registration of 3-D Shapes. In *Sensor fusion IV: control paradigms and data structures*, 1992.
- [5] Yuhua Chen, Luc Van Gool, Cordelia Schmid, and Cristian Sminchisescu. Consistency Guided Scene Flow Estimation. In *European Conference on Computer Vision (ECCV)*, 2020.
- [6] Wencan Cheng and Jong Hwan Ko. Bi-PointFlowNet: Bidirectional Learning for Point Cloud Based Scene Flow Estimation. In *European Conference on Computer Vision (ECCV)*, 2022.
- [7] Kyunghyun Cho, Bart Van Merri nboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [8] Nathaniel Chodosh, Deva Ramanan, and Simon Lucey. Re-Evaluating LiDAR Scene Flow for Autonomous Driving. *arXiv preprint arXiv:2304.02150*, 2023.
- [9] David Deng and Avideh Zakhor. RSF: Optimizing Rigid Scene Flow From 3D Point Clouds Without Labels. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2023.
- [10] Guanting Dong, Yueyi Zhang, Hanlin Li, Xiaoyan Sun, and Zhiwei Xiong. Exploiting Rigidity Constraints for LiDAR Scene Flow Estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [11] Martin Ester, Hans-Peter Kriegel, J rg Sander, Xiaowei Xu, et al. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *kdd*, 1996.

- [12] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [13] Zan Gojcic, Or Litany, Andreas Wieser, Leonidas J Guibas, and Tolga Birdal. Weakly Supervised Learning of Rigid 3D Scene Flow. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [14] Xiaodong Gu, Chengzhou Tang, Weihao Yuan, ZuoZhuo Dai, Siyu Zhu, and Pings Tan. RCP: Recurrent Closest Point for Scene Flow Estimation on 3D Point Clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [15] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [16] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [17] Eddy Ilg, Tonmoy Saikia, Margret Keuper, and Thomas Brox. Occlusions, Motion and Depth Boundaries with a Generic Network for Disparity, Optical Flow or Scene Flow Estimation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [18] Mariano Jaimez, Mohamed Souiai, Javier Gonzalez-Jimenez, and Daniel Cremers. A Primal-Dual Framework for Real-Time Dense RGB-D Scene Flow. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [19] Mariano Jaimez, Christian Kerl, Javier Gonzalez-Jimenez, and Daniel Cremers. Fast Odometry and Scene Flow from RGB-D Cameras based on Geometric Clustering. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [20] Huaizu Jiang, Deqing Sun, Varun Jampani, Zhaoyang Lv, Erik Learned-Miller, and Jan Kautz. SENSE: a Shared Encoder Network for Scene-flow Estimation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [21] Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 1976.
- [22] Yair Kittenplon, Yonina C Eldar, and Dan Raviv. FlowStep3D: Model Unrolling for Self-Supervised Scene Flow Estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [23] Ruibo Li, Guosheng Lin, Tong He, Fayao Liu, and Chunhua Shen. HCRF-Flow: Scene Flow from Point Clouds with Continuous High-order CRFs and Position-aware Flow Embedding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

- [24] Ruibo Li, Guosheng Lin, and Lihua Xie. Self-Point-Flow: Self-Supervised Scene Flow Estimation from Point Clouds with Optimal Transport and Random Walk. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [25] Ruibo Li, Chi Zhang, Guosheng Lin, Zhe Wang, and Chunhua Shen. RigidFlow: Self-Supervised Scene Flow Learning on Point Clouds by Local Rigidity Prior. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [26] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. FlowNet3D: Learning Scene Flow in 3D Point Clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [27] Fan Lu, Guang Chen, Yinlong Liu, Lijun Zhang, Sanqing Qu, Shu Liu, and Rongqi Gu. HRegNet: A Hierarchical Network for Large-scale Outdoor LiDAR Point Cloud Registration. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [28] Wei-Chiu Ma, Shenlong Wang, Rui Hu, Yuwen Xiong, and Raquel Urtasun. Deep Rigid Instance Scene Flow. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [29] 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 *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] Moritz Menze and Andreas Geiger. Object Scene Flow for Autonomous Vehicles. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [31] Himangi Mittal, Brian Okorn, and David Held. Just Go with the Flow: Self-Supervised Scene Flow Estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [32] Gilles Puy, Alexandre Boulch, and Renaud Marlet. FLOT: Scene Flow on Point Clouds Guided by Optimal Transport. In *European Conference on Computer Vision (ECCV)*, 2020.
- [33] Yi-Ling Qiao, Lin Gao, Yukun Lai, Fang-Lue Zhang, Ming-Ze Yuan, and Shihong Xia. SF-Net: Learning Scene Flow from RGB-D Images with CNNs. In *British Machine Vision Conference (BMVC)*, 2018.
- [34] Rohan Saxena, René Schuster, Oliver Wasenmüller, and Didier Stricker. PWOC-3D: Deep Occlusion-Aware End-to-End Scene Flow Estimation. *IEEE International Conference on Intelligent Vehicles Symposium (IV)*, 2019.
- [35] René Schuster, Oliver Wasenmüller, Christian Unger, Georg Kuschik, and Didier Stricker. SceneFlowFields++: Multi-frame Matching, Visibility Prediction, and Robust Interpolation for Scene Flow Estimation. *International Journal of Computer Vision (IJCV)*, 2020.
- [36] Aleksandr Segal, Dirk Haehnel, and Sebastian Thrun. Generalized-ICP. In *Robotics: Science and Systems*, 2009.

- [37] Lin Shao, Parth Shah, Vikranth Dwaracherla, and Jeannette Bohg. Motion-based Object Segmentation based on Dense RGB-D Scene Flow. *IEEE Robotics and Automation Letters (RA-L)*, 2018.
- [38] Hang Su, Varun Jampani, Deqing Sun, Subhansu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. SPLATNet: Sparse Lattice Networks for Point Cloud Processing. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [39] Zachary Teed and Jia Deng. RAFT-3D: Scene Flow using Rigid-Motion Embeddings. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [40] Ivan Tishchenko, Sandro Lombardi, Martin R Oswald, and Marc Pollefeys. Self-Supervised Learning of Non-Rigid Residual Flow and Ego-Motion. In *International Conference on 3D Vision (3DV)*, 2020.
- [41] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3D Scene Flow Estimation with a Piecewise Rigid Scene Model. *International Journal of Computer Vision (IJCV)*, 2015.
- [42] Guangming Wang, Xinrui Wu, Zhe Liu, and Hesheng Wang. Hierarchical Attention Learning of Scene Flow in 3D Point Clouds. *IEEE Transactions on Image Processing (TIP)*, 2021.
- [43] Guangming Wang, Zhiheng Feng, Chaokang Jiang, and Hesheng Wang. Unsupervised Learning of 3D Scene Flow with 3D Odometry Assistance. *arXiv preprint arXiv:2209.04945*, 2022.
- [44] Guangming Wang, Yunzhe Hu, Zhe Liu, Yiyang Zhou, Masayoshi Tomizuka, Wei Zhan, and Hesheng Wang. What Matters for 3D Scene Flow Network. In *European Conference on Computer Vision (ECCV)*, 2022.
- [45] Guangming Wang, Yunzhe Hu, Xinrui Wu, and Hesheng Wang. Residual 3D Scene Flow Learning with Context-Aware Feature Extraction. *IEEE Transactions on Instrumentation and Measurement (TIM)*, 2022.
- [46] Shenlong Wang, Simon Suo, Wei-Chiu Ma, Andrei Pokrovsky, and Raquel Urtasun. Deep Parametric Continuous Convolutional Neural Networks. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [47] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic Graph CNN for Learning on Point Clouds. *ACM Transactions on Graphics (ToG)*, 2019.
- [48] Yi Wei, Ziyi Wang, Yongming Rao, Jiwen Lu, and Jie Zhou. PV-RAFT: Point-Voxel Correlation Fields for Scene Flow Estimation of Point Clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [49] 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 *European Conference on Computer Vision (ECCV)*, 2020.

- 
- [50] Bo Yang, Sen Wang, Andrew Markham, and Niki Trigoni. Robust Attentional Aggregation of Deep Feature Sets for Multi-view 3D Reconstruction. *International Journal of Computer Vision (IJCV)*, 2020.
- [51] Wenxiao Zhang and Chunxia Xiao. PCAN: 3D Attention Map Learning Using Contextual Information for Point Cloud Based Retrieval. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.