

# DifFlow3D: Toward Robust Uncertainty-Aware Scene Flow Estimation with Iterative Diffusion-Based Refinement

Jiuming Liu<sup>1</sup>, Guangming Wang<sup>2</sup>, Weicai Ye<sup>3</sup>, Chaokang Jiang<sup>4</sup>, Jinru Han<sup>1</sup>,  
 Zhe Liu<sup>5</sup>, Guofeng Zhang<sup>3</sup>, Dalong Du<sup>4</sup>, Hesheng Wang<sup>1\*</sup>

<sup>1</sup>Department of Automation, Shanghai Jiao Tong University    <sup>2</sup>University of Cambridge

<sup>3</sup>State Key Lab of CAD & CG, Zhejiang University    <sup>4</sup>PhiGent Robotics

<sup>5</sup> MoE Key Lab of Artificial Intelligence, Shanghai Jiao Tong University

{liujiuming, wangguangming, liuzhesjtu, wanghesheng}@sjtu.edu.cn

ts20060079a31@cumt.edu.cn    maikeyeweicai@gmail.com

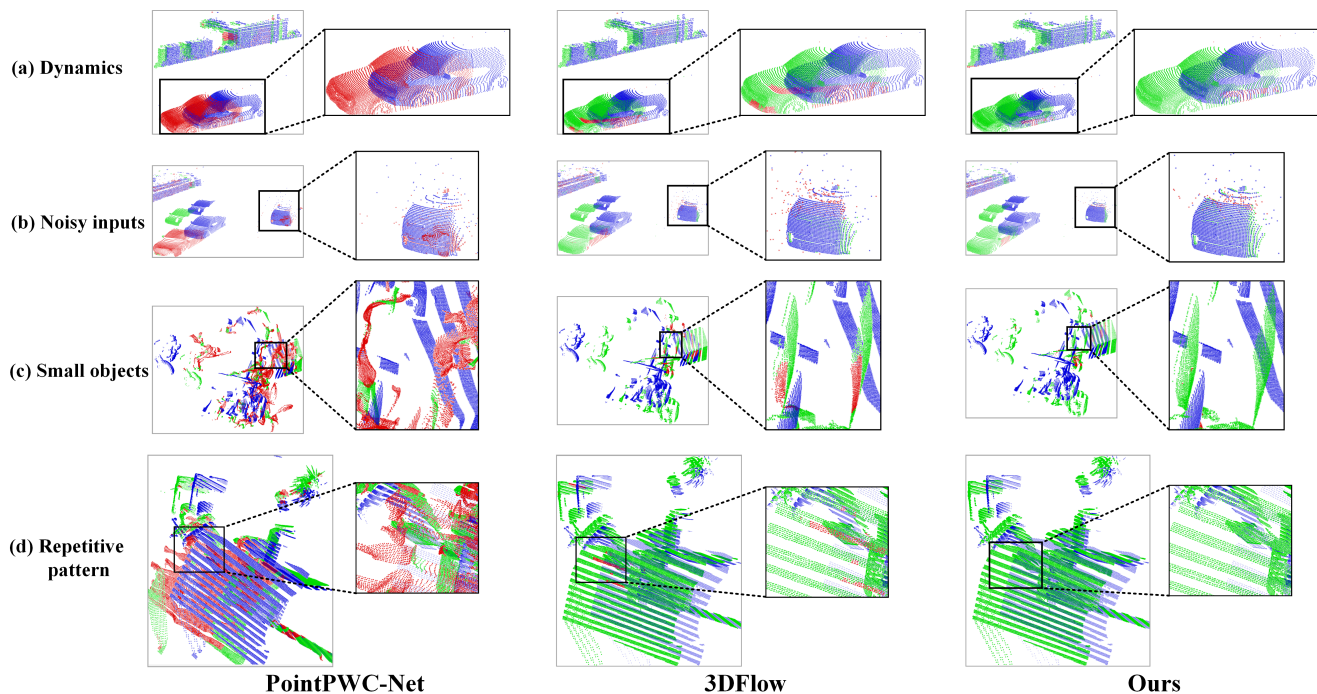


Figure 1. **Comparison on challenging cases.** DifFlow3D predicts uncertainty-aware scene flow with diffusion model, which has stronger robustness for: (a) dynamics, (b) noisy inputs, (c) small objects, and (d) repetitive patterns. Blue, green, red points respectively indicate the first frame  $PC_1$ , accurately estimated  $PC_2$  ( $PC_1$  warped by estimated flow), and inaccurately estimated  $PC_2$ .

## Abstract

Scene flow estimation, which aims to predict per-point 3D displacements of dynamic scenes, is a fundamental task in the computer vision field. However, previous works commonly suffer from unreliable correlation caused by locally constrained searching ranges, and struggle with accumulated inaccuracy arising from the coarse-to-fine structure. To alleviate these problems, we propose a novel uncertainty-aware scene flow estimation network

(DifFlow3D) with the diffusion probabilistic model. Iterative diffusion-based refinement is designed to enhance the correlation robustness and resilience to challenging cases, e.g. dynamics, noisy inputs, repetitive patterns, etc. To restrain the generation diversity, three key flow-related features are leveraged as conditions in our diffusion model. Furthermore, we also develop an uncertainty estimation module within diffusion to evaluate the reliability of estimated scene flow. Our DifFlow3D achieves state-of-the-art performance, with 24.0% and 29.1% EPE3D reduction respectively on FlyingThings3D and KITTI 2015 datasets.

<sup>1\*</sup>Corresponding Author.

Notably, our method achieves an unprecedented millimeter-level accuracy (0.0078m in EPE3D) on the KITTI dataset. Additionally, our diffusion-based refinement paradigm can be readily integrated as a plug-and-play module into existing scene flow networks, significantly increasing their estimation accuracy. Codes are released at <https://github.com/IRMVLab/DifFlow3D>.

## 1. Introduction

As a fundamental task in computer vision, scene flow refers to the 3D motion field estimated from consecutive images or point clouds. It provides the low-level perception of dynamic scenes and has various down-stream applications, such as autonomous driving [10, 27], SLAM [20, 52, 55], and motion segmentation [1, 6]. Early works focus on employing stereo [13] or RGB-D images [11] as input. With the increasing application of 3D sensors, e.g. LiDAR, recent works commonly take point clouds directly as input.

As a pioneering work, FlowNet3D [22] extracts hierarchical features with PointNet++ [32], and then regresses the scene flow iteratively. PointPWC [51] further improves it with the Pyramid, Warping, and Cost volume structure [39]. HALFlow [43] follows them and introduces the attention mechanism for better flow embedding. However, these regression-based works commonly suffer from unreliable correlation [21] and local optimum problems [23]. Reasons are mainly two folds: (1)  $K$  Nearest Neighbor (KNN) is utilized for searching point correspondences in their networks, which can not take into account correct yet distant point pairs and also has matching noise [8]. (2) Another potential problem arises from the coarse-to-fine structure widely used in previous works [22, 43, 44, 51]. Basically, an initial flow is estimated in the coarsest layer and then iteratively refined in higher resolutions. However, the performance of flow refinement highly relies on the reliability of initial coarse flow, since subsequent refinement is typically constrained within small spatial ranges around the initialization.

To address the unreliability issue, 3DFlow [44] designs an all-to-all point-gathering module with backward validation. Similarly, BPF [4] and its extension MSBRN [5] propose a bi-directional network with the forward-backward correlation. IHNet [47] leverages a recurrent network with high-resolution guidance and a re-sampling scheme. However, these networks mostly struggle with large computational costs due to their bi-directional association or recurrent iteration. In this paper, we discover that diffusion models can also enhance the correlation reliability and resilience to matching noise due to its denoising intrinsic (Fig. 1). Motivated by the findings in [38] that injecting random noise is helpful to jump out of local optimum, we reformulate the deterministic flow regression task with a probabilistic diffusion model as illustrated in Fig. 2. Moreover, our method can serve as a plug-and-play module in previ-




Figure 2. **An illustration of our diffusion for scene flow estimation.** During the forward process, we progressively add Gaussian noise on the ground truth flow residual ( $s_0$ ). A neural network  $M_\theta(\cdot, \cdot, \cdot)$  is trained to denoise the noisy flow residual  $s_t$  at time  $t$  based on condition information  $C$ .

ous scene flow networks, which is more generic and almost has no computational overhead (Section 4.5).

However, it is rather challenging to leverage a generative model in our task, due to the inherent generation diversity of diffusion models. Unlike the point cloud generation task requiring diverse output samples, scene flow prediction is a deterministic task that calculates precise per-point motion vectors. To tackle this problem, we leverage strong condition information to restrict the diversity and effectively control the generated flow. Specifically, a coarse sparse scene flow is first initialized, and then the flow residuals are iteratively generated with diffusion. In each diffusion-based refinement layer, coarse flow embedding, cost volume, and geometry encoding are all utilized as conditions. In this case, diffusion is applied to learn a probabilistic mapping from conditional inputs to the flow residual.

Furthermore, previous works rarely explore the confidence and reliability of the scene flow estimation. However, dense flow matching is prone to errors in the case of noise, dynamics, small objects, and repetitive patterns as in Fig. 1. Thus, it is significant to know whether each estimated point correspondence is trustworthy. Inspired by the recent success of uncertainty estimation in optical flow task [41], we propose a per-point uncertainty in the diffusion model to evaluate the reliability of our scene flow estimation.

Overall, the contributions of our paper are as follows:

- For robust scene flow estimation, we propose a novel plug-and-play diffusion-based refinement pipeline. To the best of our knowledge, this is the first work to leverage diffusion probabilistic model in the scene flow task.
- We design strong conditional guidance by combining coarse flow embeddings, geometry encoding, and cross-frame cost volume, to control the generation diversity.
- To evaluate the reliability of our estimated flow and identify inaccurate point matching, we also introduce a per-point uncertainty estimation within our diffusion model.
- Our method outperforms all existing methods on both FlyingThings3D and KITTI datasets. Especially, our DifFlow3D achieves the millimeter-level End-point-error (EPE3D) on KITTI dataset for the first time. Compared with previous works, our method has stronger robustness

for challenging cases, e.g., noisy inputs, dynamics, etc.

## 2. Related Work

**Scene Flow Estimation.** As the counterpart of optical flow in 3D vision, scene flow has witnessed remarkable advances in recent years. FlowNet3D [22] is a pioneering work that firstly extracts point features with PointNet++ [32], and then regresses the per-point motion vector hierarchically. PointPWC [51] further improves it with the Pyramid, Warping, and Cost volume structure [39] in a patch-to-patch manner. FLOT [31] proposes an optimal transport module to establish the cross-frame association. PV-RAFT [49] fuses the point and voxel representations for both local and global feature extraction. To tackle the unreliable correlation, 3DFlow, Bi-PointFlowNet, and MSBRN [4, 5, 44] design the bi-directional correlation layer. SFGAN [45] proposes a scene flow network with the generative adversarial network (GAN). DELFlow [30] proposes a dense efficient flow estimation network for large-scale points. PT-FlowNet [8] proposes a transformer-based pipeline to capture long-range dependencies. There are also some self-supervised methods [1, 14, 15, 17, 19, 35], without requiring ground truth flow annotations. In this paper, we rethink scene flow estimation with a different insight: *How to find a more generic method improving the robustness and reliability of previous regression-based methods?* On the one hand, we introduce a denoising diffusion model with strong conditions. On the other hand, uncertainty estimation is proposed to evaluate the reliability of dense flow matching.

**Diffusion models in 3D vision.** Recently, diffusion model [12, 36] has gained significant attention, which gradually removes noise from a Gaussian distribution to the sample data distribution by learning from its reverse process. It has yielded great advancements in image generation [33, 34], video synthesis [7, 53], and static point cloud generation and completion [24, 42]. DPM [24] first applies the diffusion model conditioned on a shape latent to point cloud generation. PVD [54] proposes a 3D shape generation and completion network through point-voxel diffusion. LION [42] combines VAE and two latent diffusion models (feature-based and point-based) to further enhance the generation abilities. *However, it remains unclear whether diffusion can model dynamic point clouds, i.e. scene flow.*

**Uncertainty estimation.** Uncertainty estimation has been widely explored in stereo matching, Multi-View Stereo (MVS), and optical flow fields. UCS-Net [3] proposes an uncertainty-aware cascaded network where uncertainty interval is inferred from depth probabilities to progressively sub-divide local depth ranges. SEDNet [2] designs a loss function for joint disparity and uncertainty estimation in deep stereo matching. ProbFlow [48] proposes to jointly estimate the optical flow and its uncertainty. Similarly, PDCNet [41] enhances ProbFlow with a new train-

ing strategy in terms of self-supervised training. *Nevertheless, the uncertainty of scene flow estimation has been rarely studied before.* In this paper, we design a per-point uncertainty to evaluate the reliability of our estimated flow.

## 3. Method

Given two consecutive point cloud frames  $PC_1 \in \mathbb{R}^{N \times 3}$  and  $PC_2 \in \mathbb{R}^{M \times 3}$ , scene flow  $sf \in \mathbb{R}^{N \times 3}$  indicates the 3D motion vector corresponding to each point in  $PC_1$ .

As illustrated in Fig. 3, our network first hierarchically extracts point features and generates the initial scene flow in the coarsest layer (Section 3.1). Then, we utilize a diffusion model with strong condition signals to refine the coarse scene flow by predicting dense flow residuals from coarse to fine (Section 3.2). To evaluate the per-point prediction reliability of scene flow, uncertainty is also estimated jointly (Section 3.3) and propagated iteratively (Section 3.4). Finally, the scene flow, flow residual, and uncertainty are all included for the network supervision (Section 3.5). We will discuss each module specifically in the following sections.

### 3.1. Feature Extraction and Flow Initialization

We adopt a hierarchical point feature extraction structure, where  $PC_1$  and  $PC_2$  are down-sampled through a series of set conv layers respectively. In each set conv layer, Farthest Point Sampling (FPS) is used to sample center points, and their neighboring point features are aggregated as:

$$f_i^G = \underset{m=1,2,\dots,M}{AGG} (MLP((x_i^m - x_i) \oplus f_i^m)), \quad (1)$$

where  $x_i$  is the  $i$ -th sampled center point.  $x_i^m$  and  $f_i^m$  represent its  $m$ -th neighbor point and feature respectively.  $f_i^G$  means the output feature.  $\oplus$  is the concatenation, and  $AGG$  indicates the feature aggregation operation. It is worth noting that our method can be adapted to different initialization manners. The feature aggregation can be adopted by the max pooling in PointNet++ [32] or Pointconv [50].

After extracting point features hierarchically, we then leverage the cross-frame correlation [5, 44] in the coarsest layer for the initialization of scene flow. However, the initialized scene flow  $sf^0$  is sparse and inaccurate due to the low resolution and mismatching. Inspired by recent success and denoising properties of the diffusion model [12], we design a novel diffusion-based scene flow refinement module to recover the denser scene flow progressively.

### 3.2. Diffusion-based Scene Flow Refinement

It is non-trivial to directly recover scene flow from pure noise due to the significant discrepancies between their data distributions. Since coarse scene flow has been already obtained in the above flow initialization, what we exactly need is the flow residual between the initialized scene flow and the ground truth. Therefore, we formulate the scene flow residual as the diffusion latent variable, and multi-scale flow

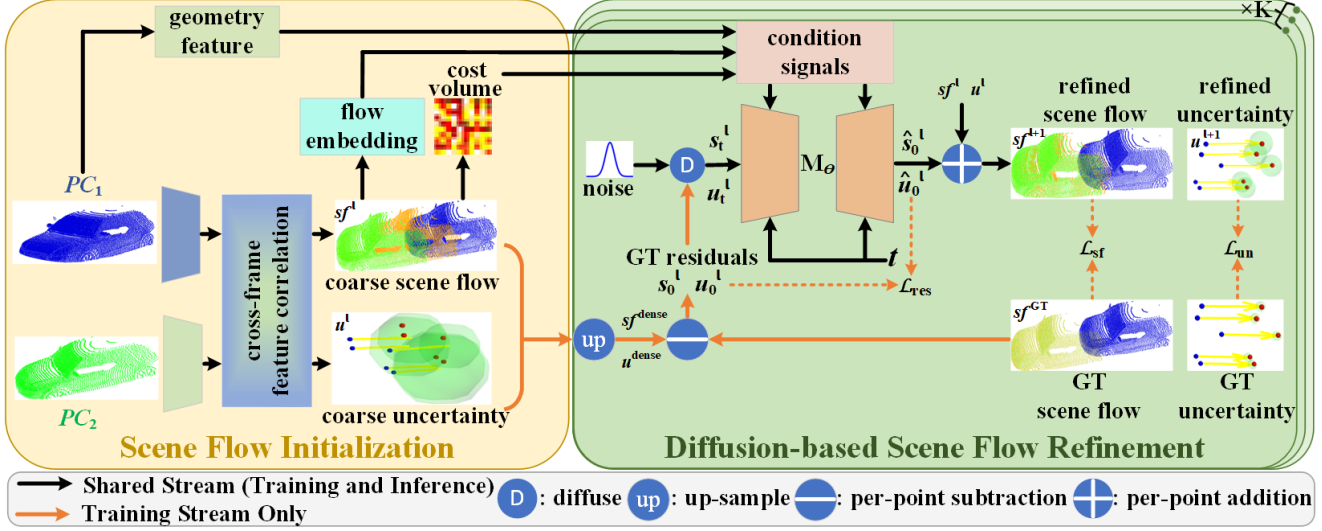


Figure 3. **The overall structure of DiffFlow3D.** We first initialize a coarse sparse scene flow in the bottom layer. Then, iterative diffusion-based refinement layers with flow-related condition signals are applied to recover the denser flow residuals. A per-point uncertainty is also predicted jointly with scene flow to evaluate the reliability of our estimated flow.

residuals are generated from the reverse process of diffusion model iteratively.

**Forward diffusion process.** While training, the forward diffusion process gradually adds Gaussian noise into the ground truth flow residual for  $T$  timesteps via a Markov chain:

$$q(s_{1:T}^i | s_0^i) = \prod_{t=1}^T q(s_t^i | s_{t-1}^i), \quad (2)$$

where  $s_0^i$  indicates the ground truth flow residual for  $i$ -th point in  $PC_1$ .  $s_t^i$  is the intermediate flow residual at timestamp  $t$ . Gaussian transition kernel  $q(s_t^i | s_{t-1}^i) = \mathcal{N}(s_t^i; \sqrt{1 - \beta_t} s_{t-1}^i, \beta_t I)$  is fixed by predefined hyperparameters  $\beta_t$ . The forward process progressively adds noise to the ground truth flow residual  $s_0^i$  and eventually turns it into a total Gaussian noise when  $T$  is large enough.

In practice, coarse sparse flow is first up-sampled by the Three-Nearest Neighbors (Three-NN) as in [5, 44]. The ground truth flow residual is then constructed through the subtraction between up-sampled coarse dense scene flow and the ground truth with the same resolution as in Fig. 3.

**Reverse denoising process.** To generate the flow residual with robustness to outliers, we resort to the reverse process of diffusion model. Basically, the reverse process can be represented as a parameterized Markov chain starting from a random noise  $p(s_T^i)$ :

$$p_\theta(s_{0:T}^i) = p(s_T^i) \prod_{t=1}^T p_\theta(s_{t-1}^i | s_t^i), \quad (3)$$

where the reverse transition kernel  $p_\theta(s_{t-1}^i | s_t^i)$  is approximated with a neural network. Here, we follow [40] to directly learn the latent variable - flow residual  $\hat{s}_0^i$  by the denoising network  $M_\theta(s_{t-1}^i | s_t^i, t)$ .

However, there still remain challenges since scene flow estimation requires precise values and directions, rather than diverse outputs as in the point cloud generation task.

Therefore, we constrain the generation diversity and control the reverse process by the powerful condition information  $C$ . In this way, the training objective of flow residual can be represented by:

$$\mathcal{L}_{res} := \mathbb{E}_{s_0^i, t} [\| s_0^i - M_\theta(s_t^i, t, C) \|_2^2], \quad (4)$$

where  $M_\theta(\cdot, \cdot, \cdot)$  is the denoising network.

**Design of condition signals.** 3DFlow [44] compares different design choices and investigates what really matters for 3D scene flow estimation. Inspired by their work, our diffusion condition signals include three key flow-related components: geometry feature  $p^i$  extracted from  $PC_1$ , cross-frame cost volume  $cv^i$ , and coarse dense flow embeddings  $e^i$ . Cross-frame cost volume [43] is constructed by warping the first point cloud  $PC_1$  and then attentively correlating with  $PC_2$ . The augmented version [5] with Gate Recurrent Unit (GRU) is used here for better cross-frame correlation. Coarse dense flow embeddings are generated through the set-upconv layer in [22] from coarse sparse flow embeddings. Finally, the condition information is the concatenation of these features:

$$c^i = p^i \oplus cv^i \oplus e^i. \quad (5)$$

Condition signals  $C = \{c^i | c^i \in \mathbb{R}^d, i = 1, \dots, n^l\}$  guide the reverse process to progressively produce the refined dense scene flow residual  $\hat{s}_0^i = M_\theta(s_t^i, t, C)$ . The superscript  $i$  in  $s_t^i$  will be omitted below for the simplicity.  $n^l$  is the point number of  $PC_1$  at layer  $l$ .

### 3.3. Uncertainty Estimation

Previous works rarely study the uncertainty and reliability of estimated scene flow. However, there exist point matching errors when associating two point cloud frames, especially for noisy inputs, dynamics, and repetitive patterns (Fig. 1). In these cases, estimated flow vectors from




Figure 4. **The visualization of uncertainty.** During the training process, our designed uncertainty intervals narrow progressively, which encourages predicted flow toward the ground truth.

the network are not equally reliable in terms of different points. Therefore, it is crucial to make the network aware of whether each estimated point correspondence is trustworthy during the training. For this purpose, we propose a per-point uncertainty jointly estimated with scene flow in our diffusion model to evaluate the estimation reliability.

**Estimation of uncertainty.** Cost volume represents the per-point matching degree between two point cloud frames. Therefore, in the coarsest layer, we output one additional channel as the initial uncertainty simultaneously with initial scene flow from cost volume as in Fig. 3. Then, in each refinement module, to control the flow residual generation more tightly, we also predict the uncertainty residuals using the same process as flow residuals through the denoising network of diffusion module:

$$\hat{s}_0^l, \hat{u}_0^l = M_\theta(s_t^l, u_t^l, t, C), \quad (6)$$

where  $\hat{u}_0^l$  represents the predicted uncertainty residual corresponding to each per-point flow residual  $\hat{s}_0^l$ . Finally, refined uncertainty in layer  $l$  is then propagated into layer  $l+1$  through the feature propagation process [22].

**Construction of the ground truth uncertainty.** The key issue is how to enable the network aware of uncertainty and how to supervise it toward the appropriate optimization direction. During the initial stages of training, the network has a relatively poor ability to estimate precise scene flow and tends to focus more on learning easier point correspondences. As the network converges, it has mastered the correspondence-finding prior to some extent for estimating more challenging cases. Therefore, our uncertainty intervals should also be narrowed dynamically throughout the training process. We obey this rule to construct the ground truth uncertainty as:

$$e_{ab} = \| sf^l - sf^{GT} \|_2, \quad (7)$$

$$e_{re} = \left| \frac{e_{ab}}{|sf^{GT}|} \right|, \quad (8)$$

$$u_0^l = \begin{cases} 0, & e_{ab} < E_1, e_{re} < E_2 \\ 1, & otherwise, \end{cases} \quad (9)$$

where  $u_0^l$  indicates the constructed uncertainty ground truth.  $E_1$  and  $E_2$  decay with the convergence of the network, which constrain the uncertainty intervals by absolute and relative estimation errors. As in Fig. 4, uncertainty intervals represent the reliability range of each estimated flow and progressively encourage the estimated flow to optimize toward the ground truth. The supervision of uncertainty residual is similar to flow residual in the diffusion:

$$\mathcal{L}_{un} := \mathbb{E}_{u_0^l, t} [\| u_0^l - \hat{u}_0^l \|_2^2]. \quad (10)$$

By incorporating uncertainty estimation into the training process, we enhance the network’s ability to capture and quantify the uncertainty associated with each estimated scene flow, leading to more reliable and accurate results.

### 3.4. Overall Architecture of DiffFlow3D

We adopt the coarse-to-fine structure to generate the scene flow residuals and uncertainty residuals at different scales. The inputs of our diffusion-based refinement at layer  $l$  are coarse sparse scene flow  $sf^l$ , coarse sparse uncertainty  $u^l$  and condition signals  $C^l$ . Firstly, coarse sparse scene flow  $sf^l$  is up-sampled to produce the coarse dense scene flow  $sf^{dense}$ . Similarly, coarse sparse uncertainty  $u^l$  is up-sampled as coarse dense uncertainty  $u^{dense}$ . Dense flow residual  $\hat{s}_0^l$ , uncertainty residual  $\hat{u}_0^l$  are then refined from the denoising network  $M_\theta$ . The refined dense scene flow (or uncertainty) at layer  $l$  is generated by the per-point addition between coarse dense flow (or uncertainty) and its corresponding residual. Finally, the refined dense flow and uncertainty will be the inputs of layer  $l+1$ :

$$sf^{l+1} = sf^{dense} + \hat{s}_0^l, \quad (11)$$

$$u^{l+1} = u^{dense} + \hat{u}_0^l. \quad (12)$$

### 3.5. Training Objective

We adopt a multi-scale supervision strategy in terms of scene flow, scene flow residuals, and uncertainty residuals. The scene flow loss is designed as:

$$\mathcal{L}_{sf}^l = \| sf^l - sf^{GT} \|_2, \quad (13)$$

where  $sf^l$  and  $sf^{GT}$  respectively indicate the estimated and ground truth scene flow at layer  $l$ .

The overall loss function is the mixture of three parts:

$$\mathcal{L}^l = \lambda_{sf} \times \mathcal{L}_{sf}^l + \lambda_{res} \times \mathcal{L}_{res}^l + \lambda_{un} \times \mathcal{L}_{un}^l, \quad (14)$$

$$\mathcal{L} = \sum_{l=0}^K \alpha^l \times \mathcal{L}^l, \quad (15)$$

where  $\alpha^l$  and  $\mathcal{L}^l$  indicate the weight and total loss at layer  $l$ .  $K$  is the number of our diffusion-based refinement layer.

Method	FT3D <sub>s</sub>						KITTI <sub>s</sub>					
	EPE3D↓	Acc3DS↑	Acc3DR↑	Outliers↓	EPE2D↓	Acc2D↑	EPE3D↓	Acc3DS↑	Acc3DR↑	Outliers↓	EPE2D↓	Acc2D↑
FlowNet3D [22]	0.1136	0.4125	0.7706	0.6016	5.9740	0.5692	0.1767	0.3738	0.6677	0.5271	7.2141	0.5093
HPLFlowNet [9]	0.0804	0.6144	0.8555	0.4287	4.6723	0.6764	0.1169	0.4783	0.7776	0.4103	4.8055	0.5938
PointPWC-Net [51]	0.0588	0.7379	0.9276	0.3424	3.2390	0.7994	0.0694	0.7281	0.8884	0.2648	3.0062	0.7673
HALFlow [43]	0.0492	0.7850	0.9468	0.3083	2.7555	0.8111	0.0622	0.7649	0.9026	0.2492	2.5140	0.8128
FLOT [31]	0.0520	0.7320	0.9270	0.3570	—	—	0.0560	0.7550	0.9080	0.2420	—	—
HCRF-Flow [18]	0.0488	0.8337	0.9507	0.2614	2.5652	0.8704	0.0531	0.8631	0.9444	0.1797	2.0700	0.8656
PV-RAFT [49]	0.0461	0.8169	0.9574	0.2924	—	—	0.0560	0.8226	0.9372	0.2163	—	—
FlowStep3D [16]	0.0455	0.8162	0.9614	0.2165	—	—	0.0546	0.8051	0.9254	0.1492	—	—
RCP [10]	0.0403	0.8567	0.9635	0.1976	—	—	0.0481	0.8491	0.9448	0.1228	—	—
3DFlow [44]	0.0281	0.9290	0.9817	0.1458	1.5229	0.9279	0.0309	0.9047	0.9580	0.1612	1.1285	0.9451
BPF [4]	0.0280	0.9180	0.9780	0.1430	1.5820	0.9290	0.0300	0.9200	0.9600	0.1410	1.0560	0.9490
PT-FlowNet [8]	0.0304	0.9142	0.9814	0.1735	1.6150	0.9312	0.0224	0.9551	0.9838	0.1186	0.9893	0.9667
IHNet [47]	0.0191	0.9601	0.9865	0.0715	1.0918	0.9563	0.0122	0.9779	0.9892	0.0913	0.4993	0.9862
MSBRN [5]	0.0150	0.9730	0.9920	0.0560	0.8330	0.9700	0.0110	0.9710	0.9890	0.0850	0.4430	0.9850
Ours	<b>0.0114</b>	<b>0.9836</b>	<b>0.9949</b>	<b>0.0350</b>	<b>0.6220</b>	<b>0.9824</b>	<b>0.0078</b>	<b>0.9817</b>	<b>0.9924</b>	<b>0.0795</b>	<b>0.2987</b>	<b>0.9932</b>

Table 1. Comparison results on FlyingThings3D and KITTI datasets without occlusion [9]. The best results are in bold. Our method has 24.0% and 29.1% EPE3D reduction respectively compared with previous works.

## 4. Experiment

### 4.1. Datasets and Implementation Details

We follow previous works [4, 9, 22, 31, 44, 49] to train our DifFlow3D on synthetic FlyingThings3D dataset [25], and evaluate the model on both FlyingThings3D and KITTI [26] datasets. All the experiments are conducted on a single RTX 3090 GPU. The Adam optimizer is adopted with  $\beta_1 = 0.9, \beta_2 = 0.999$ . The learning rate is initially set as 0.0001 and reduced by 0.8 every 10 epochs. We leverage DDIM [37] with the total timestamp  $T = 1000$ . Refinement layer number  $K$  is 3. Gaussian noise is set as  $\sigma = 1$ . We imitate the learning rate step decay strategy for  $E_1$  and  $E_2$ , which are initially set as 0.5, and decay with the rate 0.8 until 0.02. Our evaluation metrics include EPE3D (m), Acc3DS, Acc3DR, Outliers, EPE2D (px), and Acc2D, commonly used in previous works [4, 22, 44].

### 4.2. Comparison with State-of-the-Art Methods

To comprehensively demonstrate the superiority of our method, we show comparison results with two pre-process settings: without occlusion [9] and with occlusion [22].

**Comparison on point clouds without occlusion.** We compare our DifFlow3D with a series of state-of-the-art (SOTA) methods on both FlyingThings3D [25] and KITTI scene flow [26] datasets without occlusion. Table 1 demonstrates that our DifFlow3D outperforms all previous methods in terms of both 3D and 2D metrics on two datasets. Compared with the recent SOTA method MSBRN [5], our model reduces the End-point-error (EPE3D) by 24.0% and 29.1% respectively on FlyingThings3D and KITTI datasets. It is worth noting that DifFlow3D also has the best generalization capability. Only trained on the Flyingthings3D synthetic dataset, our method firstly achieves millimeter-level

Dataset	Method	Sup.	EPE3D↓	Acc3DS↑	Acc3DR↑	Outliers↓
FT3D <sub>o</sub>	FlowNet3D [22]	Full	0.169	0.254	0.579	0.789
	FLOT [31]	Full	0.156	0.343	0.643	0.700
	OGSFNet [29]	Full	0.163	0.551	0.776	0.518
	FESTA [46]	Full	0.111	0.431	0.744	—
	3DFlow [44]	Full	0.063	0.791	0.909	0.279
	BPF [4]	Full	0.073	0.791	0.896	0.274
	MSBRN [5]	Full	0.053	0.836	0.926	0.231
	Ours	Full	<b>0.043</b>	<b>0.891</b>	<b>0.944</b>	<b>0.133</b>
	KITTI <sub>o</sub>	FlowNet3D [22]	Full	0.173	0.276	0.609
FLOT [31]		Full	0.110	0.419	0.721	0.486
OGSFNet [29]		Full	0.075	0.706	0.869	0.327
FESTA [46]		Full	0.097	0.449	0.833	—
3DFlow [44]		Full	0.073	0.819	0.890	0.261
BPF [4]		Full	0.065	0.769	0.906	0.264
SCOOP [17]		Self	0.063	0.797	0.910	0.244
MSBRN [5]		Full	0.044	0.873	0.950	0.208
Ours		Full	<b>0.031</b>	<b>0.955</b>	<b>0.966</b>	<b>0.108</b>

Table 2. Comparison results on FlyingThings3D and KITTI datasets with occlusion [22]. “Self” and “Full” respectively mean self-supervised and fully-supervised training. Our method outperforms previous works by 18.9% and 29.5% in terms of EPE3D.

(0.0078m) EPE3D metric on the KITTI dataset.

**Comparison on point clouds with occlusion.** We also evaluate our model on the datasets with occlusion prepared by [22]. Experiment results in Table 2 show that our DifFlow3D still surpasses all previous methods on both two datasets. Notably, our evaluation results on the KITTI dataset outperform previous methods by a large margin. Only trained on FT3D<sub>o</sub>, our method reduces EPE3D by 29.5% and outliers by 48.1% on KITTI. This demonstrates the excellent generalization capability of our method.

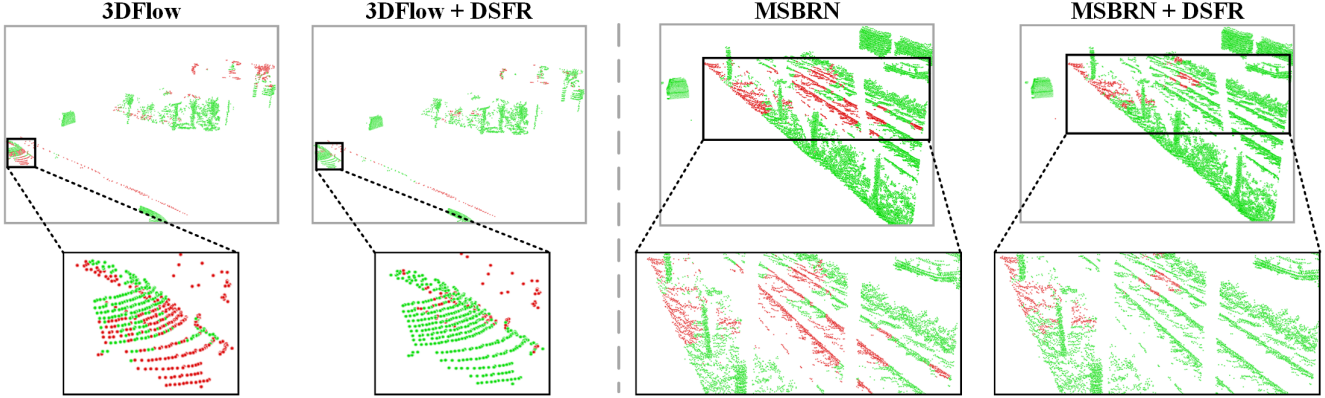


Figure 5. **Visualization results w/o or with our Diffusion-based Scene Flow Refinement (DSFR).** For better comparison, we only visualize the estimated  $PC_2$  by warping  $PC_1$  with estimated scene flow. **green, red** points respectively indicate **accurately estimated  $PC_2$**  and **inaccurately estimated  $PC_2$**  (measured by Acc3DR).

Method	Training	FT3D <sub>s</sub>						KITTI <sub>s</sub>					
		EPE3D↓	Acc3DS↑	Acc3DR↑	Outliers↓	EPE2D↓	Acc2D↑	EPE3D↓	Acc3DS↑	Acc3DR↑	Outliers↓	EPE2D↓	Acc2D↑
3DFlow [44]	Quarter	0.0317	0.9109	0.9757	0.1673	1.7436	0.9108	0.0332	0.8931	0.9528	0.1690	<b>1.2186</b>	0.9373
3DFlow+DSFR	Quarter	<b>0.0297</b> (↓ 6.3%)	<b>0.9207</b>	<b>0.9785</b>	<b>0.1548</b>	<b>1.6344</b>	<b>0.9188</b>	<b>0.0316</b> (↓ 4.8%)	<b>0.9028</b>	<b>0.9634</b>	<b>0.1604</b>	1.2247	<b>0.9413</b>
3DFlow [44]	Complete	0.0281	0.9290	0.9817	0.1458	1.5229	0.9279	0.0309	0.9047	0.9580	0.1612	1.1285	0.9451
3DFlow+DSFR	Complete	<b>0.0242</b> (↓ 13.9%)	<b>0.9494</b>	<b>0.9860</b>	<b>0.1166</b>	<b>1.3201</b>	<b>0.9459</b>	<b>0.0251</b> (↓ 18.8%)	<b>0.9398</b>	<b>0.9793</b>	<b>0.1302</b>	<b>0.9761</b>	<b>0.9686</b>
BPF [4]	Complete	0.0280	0.9180	0.9780	0.1430	1.5820	0.9290	0.0300	0.9200	0.9600	0.1410	1.0560	0.9490
BPF+DSFR	Complete	<b>0.0247</b> (↓ 11.8%)	<b>0.9390</b>	<b>0.9840</b>	<b>0.1192</b>	<b>1.3749</b>	<b>0.9468</b>	<b>0.0265</b> (↓ 11.7%)	<b>0.9355</b>	<b>0.9672</b>	<b>0.1290</b>	<b>1.0527</b>	<b>0.9637</b>
MSBRN [5]	Complete	0.0150	0.9730	0.9920	0.0560	0.8330	0.9700	0.0110	0.9710	0.9890	0.0850	0.4430	0.9850
MSBRN+DSFR	Complete	<b>0.0114</b> (↓ 24.0%)	<b>0.9836</b>	<b>0.9949</b>	<b>0.0350</b>	<b>0.6220</b>	<b>0.9824</b>	<b>0.0078</b> (↓ 29.1%)	<b>0.9817</b>	<b>0.9924</b>	<b>0.0795</b>	<b>0.2987</b>	<b>0.9932</b>

Table 3. **The plug-and-play capability of our methods.** Our Diffusion-based Scene Flow Refinement (DSFR) can effectively improve the accuracy introduced into recent methods on both FlyingThings3D and KITTI datasets. The best results are in bold.

### 4.3. Plug-and-Play on Previous Works

It is extremely encouraging that our diffusion-based refinement can serve as a plug-and-play module, improving the estimation accuracy of several recent SOTA works [4, 5, 44]. In Table 3, we replace their refinement modules with our diffusion-based one. 3DFlow [44] first trains their model on a quarter of the data, then finetunes results on the complete dataset. For the comprehensive comparison, we also show experiment results on both quarter and complete training sets. As shown in Table 3, our method respectively has 6.3% and 13.9% EPE3D reduction on the quarter and complete training set of FT3D<sub>s</sub> dataset. Similarly, introducing our diffusion refinement strategy into the GRU in BPF [4] and MSBRN [5] can also significantly improve their accuracy. Also, they innovate the cross-frame association with specific designs. Instead, our diffusion-based refinement is more generic in terms of different network designs or challenging cases. As in Fig. 5, estimations for dynamic cars or bushes with complex patterns are more accurate by adding our Diffusion-based Scene Flow Refinement (DSFR).

### 4.4. Ablation Study

Extensive experiments are conducted to validate the effectiveness of each proposed component.

**Designs of diffusion model.** We first remove the diffusion refinement module, remaining the coarsest initialization stage. As in Table 4 (a), there is a 0.14m EPE3D increase. We also compare different denoising network choices (MLP [44], point transformer [28], or GRU [5]). GRU has the best performance, partly because the recurrent network is more suitable for iterative refinement.

**Uncertainty.** As illustrated in Table 4 (b), without our proposed uncertainty, there is 28.9% higher EPE3D. Our uncertainty ground truth is designed as a binary value. We also conduct the experiment by substituting it with a continuous value ranging from 0 to 1. However, worse estimation results can be seen. To better guide the flow regression, we jointly estimate uncertainty and scene flow as in Section 3.3. If we predict uncertainty and scene flow by two separate fully connected layers from cost volume, there is a significant EPE3D increase. This is because uncertainty can not tightly control the flow due to the indirect connection and even mislead the flow learning in the wrong direction.

**Condition signals.** Conditional information is extremely significant in guiding the flow residual refinement module. We remove each of our designed condition signals to validate their importance as in Table 4 (c). Without cost volume as guidance, our model has obviously worse

Method	EPE3D↓	Acc3DS↑	Acc3DR↑	Outliers↓
<b>(a) Diffusion</b>				
w/o diffusion refinement	0.1559	0.0777	0.4462	0.8637
with transformer-based $M_\theta$ [28]	0.0374	0.8465	0.9638	0.1862
with MLP-based $M_\theta$ [44]	0.0244	0.9492	0.9858	0.1172
Ours (with GRU-based $M_\theta$ [5])	<b>0.0114</b>	<b>0.9836</b>	<b>0.9949</b>	<b>0.0350</b>
<b>(b) Uncertainty</b>				
w/o uncertainty	0.0147	0.9753	0.9929	0.0515
with continuous uncertainty GT	0.0178	0.9669	0.9906	0.0774
with separately estimated uncertainty	0.0170	0.9692	0.9914	0.0661
Ours (full)	<b>0.0114</b>	<b>0.9836</b>	<b>0.9949</b>	<b>0.0350</b>
<b>(c) Condition</b>				
w/o cost volume	0.0621	0.6821	0.9098	0.4038
w/o coarse flow embedding	0.0245	0.9415	0.9817	0.1058
w/o geometry feature	0.0150	0.9723	0.9921	0.0603
Ours (full, with all three conditions)	<b>0.0114</b>	<b>0.9836</b>	<b>0.9949</b>	<b>0.0350</b>

Table 4. Ablation studies on FlyingThings3D prepared by [9].

Method	Runtime	Method	Runtime
FLOT [5]	289.6ms	PV-RAFT [49]	781.1ms
FlowStep3D [16]	972.7ms	RCP [10]	2854.6ms
MSBRN [5]	221.1ms	Ours	228.3ms

Table 5. Runtime comparison. Compared with our baseline MSBRN [5], our proposed diffusion-based refinement only brings minimal computational overhead.

estimation (4.5 times larger EPE3D), since cost volume precisely represents the per-point correlation between two frames. Coarse flow embedding can also lead the flow generation because the coarse matching in the low-resolution layer has a more global understanding of cross-frame correlation. Thus, if we remove it, there is a significant accuracy drop. The geometry feature has low-level 3D structural information in dynamic scenes. 31.6% higher EPE3D can be seen without the geometry feature as guidance.

#### 4.5. Runtime Comparison

We compare our DifFlow3D with previous methods on efficiency. All results are evaluated on a single RTX 3090 GPU. Table 5 shows that DifFlow3D has highly competitive efficiency. Notice that our baseline here is based on MSBRN [5], where we combine its flow initialization with our diffusion-based refinement. Compared with MSBRN, our method has only a slightly more computational overhead (3.3%), but much higher estimation accuracy (24.0% and 29.1% on FT3D<sub>s</sub> and KITTI<sub>s</sub> as in Table 1).

### 5. Discussion

We discuss the motivations and reasons why our proposed diffusion-based refinement module works.

**Why diffusion?** As illustrated in Fig. 1, the introduction of diffusion model can improve the robustness for some challenging cases. We further show two quantitative comparisons by adding random noise in Fig. 6. Previous




Figure 6. Adding random Gaussian noise on input points. We add random Gaussian noise on input points, where the horizontal coordinate represents the standard deviation of the noise.




Figure 7. How uncertainty works during the training process. We visualize the uncertainty intervals respectively at different training stages (epoch=10 and epoch=100).

works all witness a dramatically increasing EPE3D, but our method has a relatively small fluctuation in accuracy.

**How uncertainty works?** Our uncertainty represents the reliability and confidence level of per-point predicted scene flow, which can progressively constrain the flow residual generation range. As in Fig. 7, there are large uncertainty intervals for all points at the initial stages of the training. After training for 100 epochs as in Fig. 7 (b), all uncertainty intervals are narrow obviously because our network has better estimation ability. Furthermore, uncertainty intervals also vary in terms of different objects. As shown in Fig. 7 (b), there are relatively larger uncertainty intervals for the dynamic car, which are relatively difficult to learn due to its inconsistent motion. But other static objects have almost no uncertainty estimations.

### 6. Conclusion

In this paper, we innovatively propose a diffusion-based scene flow refinement network with uncertainty awareness. Multi-scale diffusion refinement is adopted to generate fine-grained dense flow residuals. To improve the robustness of our estimation, we also introduce a per-point uncertainty jointly generated with scene flow. Extensive experiments demonstrate the superiority and generalization ability of our DifFlow3D. Notably, our diffusion-based refinement can serve as a plug-and-play module into previous works and shed new light on future works.

**Acknowledgement.** This work was supported in part by the Natural Science Foundation of China under Grant 62225309, 62073222, U21A20480, and 62361166632.



## References

- [1] Stefan Andreas Baur, David Josef Emmerichs, Frank Moosmann, Peter Pinggera, Björn Ommer, and Andreas Geiger. Slim: Self-supervised lidar scene flow and motion segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13126–13136, 2021. [2](#), [3](#)
- [2] Liyan Chen, Weihang Wang, and Philippos Mordohai. Learning the distribution of errors in stereo matching for joint disparity and uncertainty estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17235–17244, 2023. [3](#)
- [3] Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su. Deep stereo using adaptive thin volume representation with uncertainty awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2524–2534, 2020. [3](#)
- [4] Wencan Cheng and Jong Hwan Ko. Bi-pointflownet: Bidirectional learning for point cloud based scene flow estimation. In *European Conference on Computer Vision*, pages 108–124. Springer, 2022. [2](#), [3](#), [6](#), [7](#)
- [5] Wencan Cheng and Jong Hwan Ko. Multi-scale bidirectional recurrent network with hybrid correlation for point cloud based scene flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10041–10050, 2023. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [6] Tianchen Deng, Hongle Xie, Jingchuan Wang, and Weidong Chen. Long-term visual simultaneous localization and mapping: Using a bayesian persistence filter-based global map prediction. *IEEE Robotics & Automation Magazine*, 30(1): 36–49, 2023. [2](#)
- [7] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7346–7356, 2023. [3](#)
- [8] Jingyun Fu, Zhiyu Xiang, Chengyu Qiao, and Tingming Bai. Pt-flownet: Scene flow estimation on point clouds with point transformer. *IEEE Robotics and Automation Letters*, 8(5): 2566–2573, 2023. [2](#), [3](#), [6](#)
- [9] Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, and Panqu Wang. Hplflownet: Hierarchical permutohedral lattice flownet for scene flow estimation on large-scale point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3254–3263, 2019. [6](#), [8](#)
- [10] Xiaodong Gu, Chengzhou Tang, Weihao Yuan, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Rcp: recurrent closest point for point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8216–8226, 2022. [2](#), [6](#), [8](#)
- [11] Simon Hadfield and Richard Bowden. Kinecting the dots: Particle based scene flow from depth sensors. In *2011 International Conference on Computer Vision*, pages 2290–2295. IEEE, 2011. [2](#)
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [3](#)
- [13] Frédéric Huguet and Frédéric Devernay. A variational method for scene flow estimation from stereo sequences. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–7. IEEE, 2007. [2](#)
- [14] Chaokang Jiang, Guangming Wang, Yanzi Miao, and Hesheng Wang. 3d scene flow estimation on pseudo-lidar: Bridging the gap on estimating point motion. *IEEE Transactions on Industrial Informatics*, 2022. [3](#)
- [15] Chaokang Jiang, Guangming Wang, Jiuming Liu, Hesheng Wang, Zhuang Ma, Zhenqiang Liu, Zhujin Liang, Yi Shan, and Dalong Du. 3dsflabelling: Boosting 3d scene flow estimation by pseudo auto-labelling. *arXiv preprint arXiv:2402.18146*, 2024. [3](#)
- [16] Yair Kittenplon, Yonina C Eldar, and Dan Raviv. Flowstep3d: Model unrolling for self-supervised scene flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4114–4123, 2021. [6](#), [8](#)
- [17] Itai Lang, Dror Aiger, Forrester Cole, Shai Avidan, and Michael Rubinstein. Scoop: Self-supervised correspondence and optimization-based scene flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5281–5290, 2023. [3](#), [6](#)
- [18] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 364–373, 2021. [6](#)
- [19] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16959–16968, 2022. [3](#)
- [20] Jiuming Liu, Guangming Wang, Chaokang Jiang, Zhe Liu, and Hesheng Wang. Translo: A window-based masked point transformer framework for large-scale lidar odometry. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1683–1691, 2023. [2](#)
- [21] Jiuming Liu, Guangming Wang, Zhe Liu, Chaokang Jiang, Marc Pollefeys, and Hesheng Wang. Regformer: an efficient projection-aware transformer network for large-scale point cloud registration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8451–8460, 2023. [2](#)
- [22] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. Flownet3d: Learning scene flow in 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 529–537, 2019. [2](#), [3](#), [4](#), [5](#), [6](#)
- [23] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16014–16023, 2021. [2](#)
- [24] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021. [3](#)

- [25] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. 6
- [26] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. 6
- [27] Mahyar Najibi, Jingwei Ji, Yin Zhou, Charles R Qi, Xinchun Yan, Scott Ettinger, and Dragomir Anguelov. Motion inspired unsupervised perception and prediction in autonomous driving. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVIII*, pages 424–443. Springer, 2022. 2
- [28] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 7, 8
- [29] Bojun Ouyang and Dan Raviv. Occlusion guided scene flow estimation on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2805–2814, 2021. 6
- [30] Chensheng Peng, Guangming Wang, Xian Wan Lo, Xinrui Wu, Chenfeng Xu, Masayoshi Tomizuka, Wei Zhan, and Hesheng Wang. Delflow: Dense efficient learning of scene flow for large-scale point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16901–16910, 2023. 3
- [31] Gilles Puy, Alexandre Boulch, and Renaud Marlet. Flot: Scene flow on point clouds guided by optimal transport. In *European conference on computer vision*, pages 527–544. Springer, 2020. 3, 6
- [32] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3
- [34] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 3
- [35] Yaqi Shen, Le Hui, Jin Xie, and Jian Yang. Self-supervised 3d scene flow estimation guided by superpoints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5271–5280, 2023. 3
- [36] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [37] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 6
- [38] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2
- [39] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. 2, 3
- [40] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 4
- [41] Prune Truong, Martin Danelljan, Radu Timofte, and Luc Van Gool. Pdc-net+: Enhanced probabilistic dense correspondence network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2, 3
- [42] Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, Karsten Kreis, et al. Lion: Latent point diffusion models for 3d shape generation. *Advances in Neural Information Processing Systems*, 35:10021–10039, 2022. 3
- [43] Guangming Wang, Xinrui Wu, Zhe Liu, and Hesheng Wang. Hierarchical attention learning of scene flow in 3d point clouds. *IEEE Transactions on Image Processing*, 30:5168–5181, 2021. 2, 4, 6
- [44] Guangming Wang, Yunzhe Hu, Zhe Liu, Yiyang Zhou, Masayoshi Tomizuka, Wei Zhan, and Hesheng Wang. What matters for 3d scene flow network. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pages 38–55. Springer, 2022. 2, 3, 4, 6, 7, 8
- [45] Guangming Wang, Chaokang Jiang, Zehang Shen, Yanzi Miao, and Hesheng Wang. Sfgan: Unsupervised generative adversarial learning of 3d scene flow from the 3d scene self. *Advanced Intelligent Systems*, 4(4):2100197, 2022. 3
- [46] Haiyan Wang, Jiahao Pang, Muhammad A Lodhi, Yingli Tian, and Dong Tian. Festa: Flow estimation via spatial-temporal attention for scene point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14173–14182, 2021. 6
- [47] Yun Wang, Cheng Chi, Min Lin, and Xin Yang. Ihnet: Iterative hierarchical network guided by high-resolution estimated information for scene flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10073–10082, 2023. 2, 6
- [48] Anne S Wannenwetsch, Margret Keuper, and Stefan Roth. Proflow: Joint optical flow and uncertainty estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 1173–1182, 2017. 3
- [49] Yi Wei, Ziyi Wang, Yongming Rao, Jiwen Lu, and Jie Zhou. Pv-raft: Point-voxel correlation fields for scene flow estimation of point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6954–6963, 2021. 3, 6, 8

- [50] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 9621–9630, 2019. 3
- [51] 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 88–107. Springer, 2020. 2, 3, 6
- [52] Hongle Xie, Tianchen Deng, Jingchuan Wang, and Weidong Chen. Angular tracking consistency guided fast feature association for visual-inertial slam. *IEEE Transactions on Instrumentation and Measurement*, 2024. 2
- [53] Sihyun Yu, Kihyuk Sohn, Subin Kim, and Jinwoo Shin. Video probabilistic diffusion models in projected latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18456–18466, 2023. 3
- [54] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5826–5835, 2021. 3
- [55] Siting Zhu, Guangming Wang, Hermann Blum, Jiuming Liu, Liang Song, Marc Pollefeys, and Hesheng Wang. Sni-slam: Semantic neural implicit slam. *arXiv preprint arXiv:2311.11016*, 2023. 2