

ICP-Flow: LiDAR Scene Flow Estimation with ICP

Yancong Lin and Holger Caesar
 Delft University of Technology, The Netherlands

Abstract

Scene flow characterizes the 3D motion between two LiDAR scans captured by an autonomous vehicle at nearby timesteps. Prevalent methods consider scene flow as point-wise unconstrained flow vectors that can be learned by either large-scale training beforehand or time-consuming optimization at inference. However, these methods do not take into account that objects in autonomous driving often move rigidly. We incorporate this rigid-motion assumption into our design, where the goal is to associate objects over scans and then estimate the locally rigid transformations. We propose ICP-Flow, a learning-free flow estimator. The core of our design is the conventional Iterative Closest Point (ICP) algorithm, which aligns the objects over time and outputs the corresponding rigid transformations. Crucially, to aid ICP, we propose a histogram-based initialization that discovers the most likely translation, thus providing a good starting point for ICP. The complete scene flow is then recovered from the rigid transformations. We outperform state-of-the-art baselines, including supervised models, on the Waymo dataset and perform competitively on Argoverse-v2 and nuScenes. Further, we train a feedforward neural network, supervised by the pseudo labels from our model, and achieve top performance among all models capable of real-time inference. We validate the advantage of our model on scene flow estimation with longer temporal gaps, up to 0.4 seconds where other models fail to deliver meaningful results.

1. Introduction

Motion is vital for visual perception, particularly for highly automated vehicles that operate in a dynamically changing 3D world, as motion facilitates the detection of dynamic objects around an autonomous vehicle. A popular task in motion prediction is scene flow estimation, which calculates point-wise motion from two temporarily adjacent LiDAR scans, *i.e.* a 3D vector that describes the displacement of a point [3, 12, 28, 43]. Scene flow lays the foundation for numerous high-level tasks in perception, particularly in scene understanding without relying on large amounts of annota-

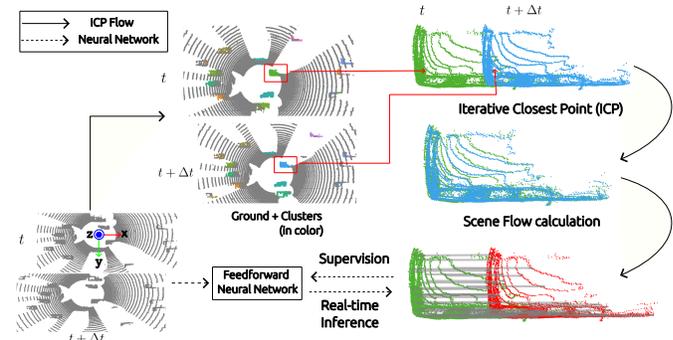


Figure 1. **ICP for scene flow.** Given two LiDAR scans, we remove the ground, cluster points, and align clusters using ICP, as objects move rigidly. We infer a rigid transformation for each pair of clusters, from which the scene flow can be recovered. Further, we train a feedforward network using the prediction from our model as supervision. The network runs in real-time with only marginal performance loss.

tions [14, 34, 51, 57]. As an example, [34] leverages scene flow to segment dynamic objects from a scene and associate them over multiple frames, from which one can create bounding boxes for training object detectors in an unsupervised manner. [51] builds on top of scene flow and takes advantage of motion cues to discover and track objects from a large volume of unlabeled data. These works treat scene flow estimation as a cornerstone and consider motion as a useful prior for temporal perception. The motion prior not only reduces the dependency on manual annotation but also scales up when ample data is available, particularly in autonomous driving where ample unlabeled data is relatively cheap to acquire. Thus, it is important to develop a reliable scene flow method for autonomous driving.

There has been a strong demand for unsupervised scene flow, as many works [14, 34, 51, 57] count on scene flow to extract object information, such as tracked bounding boxes over time, for free, *i.e.*, without human labeling. Recent works in [2, 33] have made great strides towards this direction, by utilizing the cycle consistency of forward and backward flows. Further, the work in [26, 27] proposes test-time optimization that conducts learning from scratch per sample, thus eliminating the need for training data. Un-

fortunately, these approaches are only able to predict free-form unconstrained scene flow, due to the lack of multi-body rigidity, *i.e.* a scene is composed of multiple rigidly-moving objects. As a consequence, the flow vectors from the same object, *e.g.*, a moving vehicle, may not agree in terms of their direction or magnitude. Although [15, 17] have done pilot work on incorporating the motion rigidity into scene flow, they rely on either partial or full annotation for model training. Recent work [25] achieves unsupervised learning without losing motion rigidity. We share the same spirit as [25] and further eliminate the need for large-scale data and lengthy training processes.

Another concern for scene flow is the inference cost, which is crucial for processing large volumes of data, particularly in autonomous driving. However, recent works [9, 26, 45], although being data-independent, suffer from significant inference latency. Processing a single sample can take more than a minute on a modern GPU [9, 26], making scene flow a time-consuming and resource-intensive task in real-world deployment. Other unsupervised work [25] is unable to process full LiDAR scan during inference and requires downsampling due to high demand on GPU memory.

We propose ICP-Flow, a learning-free model to overcome the reliance on data and the lack of motion rigidity. ICP-Flow also provides high-quality pseudo labels for training a neural network that runs in *real time* at inference. Our model, as shown in Fig. 1, builds on top of the Iterative Closest Point (ICP) [4] algorithm and is fully hand-crafted, thus demanding neither human annotation nor training data. Although simple, we are able to achieve competitive performance on common benchmarks, including Waymo [41], Argoverse-v2 [54] and nuScenes [5]. Further, we treat our predictions as pseudo-labels for supervising feedforward neural networks and achieve *real-time inference* with only a marginal performance decrease. Additionally, we extend our ICP-Flow to scene flow estimation over a longer temporal horizon of up to 0.4 seconds, where other models fail.

To summarize, our contributions are as follows:

- We introduce a learning-free LiDAR scene flow estimator that requires neither large datasets nor manual annotation.
- Our ICP-Flow incorporates the multi-body rigid-motion assumption by design and produces consistent scene flow per object. ICP-Flow is the top-performing model on Waymo [41] and nuScenes [5].
- Our ICP-Flow generates high-quality pseudo labels for supervising a feedforward neural network that performs on-par with the state-of-the-art, but with a considerably lower inference latency.

2. Related work

There have been numerous works that estimate scene flow from RGB or RGBD images [32, 47, 48, 53]. However, our focus is on scene flow from *point clouds*, particularly in

autonomous driving. We highlight works within this scope.

2.1. Scene flow from point clouds

Early work [12] on scene flow formulates the task as an energy minimization problem by assuming geometric constancy and motion smoothness. [43] converts point clouds into occupancy grids and computes a flow field by tracking the occupancy grids using expectation maximization. Recent works are mostly data-driven models that estimate scene flow in an end-to-end fashion [3, 8, 16, 17, 21, 28, 29, 36, 49, 52]. However, model training requires massive data labeled by human experts. In contrast, our model is free from explicit learning and costly annotation.

To remedy the need for manual labels, [1, 33, 42] take advantage of the cycle consistency and propose a self-supervised mechanism for model training. [55] achieves the same goal by minimizing the Chamfer distance between two point clouds after flow compensation, with smoothness constraint and Laplacian regularizer. Instead of self-supervised learning, [18] develops a synthetic dataset with ground truth annotations to aid learning. Another line of research focuses on knowledge distillation from imperfect pseudo labels [24, 44]. [44] supervises model training using predictions from [26] and is able to outperform the teacher model [26] when a sufficient amount of data is available. Although manual labels are no longer needed, these models still demand large amounts of data. In contrast, our model requires neither training data nor human labeling.

Recently, runtime optimization has gained popularity because of its independence on data. [9, 22, 26, 27] learn scene flow for each sample at test time by iteratively minimizing the Chamfer distances between two point clouds. Our model shares the same spirit and eliminates the need for data. However, our design is hand-crafted, thus free from the time-consuming test time optimization.

2.2. Motion rigidity in scene flow

Instead of predicting an unconstrained free-form flow vector per point, [12] uses the rigid body assumption in scene flow, *i.e.* objects do not deform, and predicts a rigid transformation per object. Similarly, [13, 15, 17, 25] also adopt the rigidity assumption by design. [45] considers rigidity as an additional regularizer and improves upon previous work [26], which only produces an unconstrained flow field. Inspired by this line of research, we also convert scene flow into rigid transformation estimation from which the complete scene flow can be recovered. Notably, previous work [25] has proposed using ICP to align LiDAR segments where the initial transformation of ICP is estimated from a deep network trained on large-scale datasets. In contrast, we eliminate the need for computationally expensive training of deep networks on large datasets by using a hand-crafted histogram-based scheme to aid ICP.

2.3. ICP

ICP [4] is a commonly used technique for registering 3D shapes or point clouds, based on point correspondences. There have been numerous works on extracting reliable correspondences, ranging from classic feature engineering [7, 35, 39, 40, 46, 56] to deep feature learning [10, 50, 58]. However, they are primarily designed to match scene-scale data, *e.g.*, full-size LiDAR scans, rather than individual segments from LiDAR data. We opt for a conventional ICP implementation [4] to match clustered LiDAR segments.

3. Method

3.1. Problem statement

Scene flow estimation takes as input a pair of LiDAR scans \mathbf{X}^t and $\mathbf{X}^{t+\Delta t}$, captured by an autonomous vehicle at two adjacent time steps t and $t + \Delta t$, where $\mathbf{X} \in \mathbb{R}^{3 \times L} = \{\mathbf{x}_l \in \mathbb{R}^3\}_{l=1}^L$ denotes a point cloud of length L . The goal is to estimate a flow field $\mathbf{F}^t \in \mathbb{R}^{3 \times L} = \{\mathbf{f}_l \in \mathbb{R}^3\}_{l=1}^L$ such that $\mathbf{X}^t + \mathbf{F}^t \approx \mathbf{X}^{t+\Delta t}$. Notably, the size of \mathbf{X}^t may differ from $\mathbf{X}^{t+\Delta t}$, while \mathbf{X}^t and \mathbf{F}^t are always of the same length.

Additionally, we assume that each LiDAR scan \mathbf{X} can be decomposed into *background* \mathbf{X}_{bg} that is static over time and *foreground* \mathbf{X}_{fg} consisting of K rigidly-moving objects that may or may not move at the given time step, denoted as $\mathbf{X}_{fg} = \{\mathbf{C}_k \in \mathbb{R}^{3 \times L_k}\}_{k=1}^K$, where \mathbf{C}_k is the k -th object, *i.e.*, a cluster of points that represents a particular object. Our aim is to estimate a rigid transformation $\mathbf{T}_k \in SE(3)$ for each object \mathbf{C}_k , from which we can recover its scene flow \mathbf{F}_k by

$$\mathbf{F}_k = \mathbf{T}_k \mathbf{T}_{ego} \circ \mathbf{C}_k - \mathbf{C}_k, \quad (1)$$

where \circ indicates applying a rigid transformation to a set of points. Matrix $\mathbf{T}_{ego} \in SE(3)$ is the ego-motion transformation at the corresponding time step. \mathbf{X}_{bg} is static over time and therefore its transformation T_{bg} is equivalent to an identity matrix $\mathbf{T}_{eye} \in SE(3)$. Thus, our end goal is to decompose a scene into a set of clusters $\{\mathbf{C}_k\}_{k=1}^{K+1}$ representing the objects, and then to estimate their transformations $\{\mathbf{T}_k\}_{k=1}^{K+1}$ between two scans \mathbf{X}^t and $\mathbf{X}^{t+\Delta t}$. For simplicity, we consider the *background* as an additional object that does not move over time.

3.2. Overview of ICP-Flow

Fig. 2 shows a sketch of ICP-Flow. Given two LiDAR scans, we first conduct ego-motion compensation to align them in the same coordinate system. Sequentially, we remove the ground points from each scan separately and fuse the remaining points for subsequent clustering, resulting in a set of clusters (*i.e.* clusters of points) at time t and $t + \Delta t$, respectively. We then employ the Iterative Closest Point (ICP) [4] algorithm to associate clusters over time. Notably,

rather than matching two LiDAR scans, we apply ICP to each pair of clusters over time and estimate a transformation matrix that minimizes the point-wise distance between paired clusters. Ultimately, we calculate a scene flow per cluster and assign it to the corresponding points in the original LiDAR scan. Crucially, we highlight that naive ICP does not deliver competitive results because it relies to a great extent on a good initial guess of the transformation. To overcome this issue, we design a simple yet effective histogram-based initialization for ICP.

Moreover, we train a feedforward neural network to further reduce the inference latency, using the pseudo-ground truth generated by our model.

3.3. Ego motion Compensation

Ego motion compensation can significantly reduce the difficulty in scene flow estimation, as the *background* and static objects no longer “move” after compensation. Further, ego motion is directly available in autonomous driving (from the IMU or other odometry) and in common benchmarks [5, 41, 54]. Thus, we take advantage of the given ego motion in our design. For benchmarks where ego motion is not available, we adopt KISS-ICP [46] to estimate the relative transformation between a pair of scans.

3.4. Ground removal and point clustering

We use Patchwork++ [23] to remove the ground from each LiDAR scan and feed the remaining points into HDBSCAN [6] for clustering. Instead of clustering each scan individually, we first fuse the non-ground points from both scans and then conduct HDBSCAN clustering. Afterward, we separate fused points by time and obtain a set of clusters $\{\mathbf{C}_m^t\}_{m=1}^M$ and $\{\mathbf{C}_n^{t+\Delta t}\}_{n=1}^N$, each of which has a timestamp and a cluster index as denoted by the superscript and subscript.

3.5. Cluster pairing

For each cluster at time t , we search for several candidate clusters at $t + \Delta t$ that are likely to match the given cluster. This is to reduce the search space and does *not* enforce one-to-one correspondence. We refer to this step as cluster pairing, after which all paired clusters are fed to ICP matching (Section 3.6) for further verification (Section 3.7). If successful, we save the transformation \mathbf{T}_m^t that best aligns each pair of clusters for scene flow calculation. Simply speaking, we pair each cluster at time t with its neighboring clusters at time $t + \Delta t$ that lie in a predefined area of $\tau_x \times \tau_y$ where τ_x and τ_y (in meters) are the maximal translations possible within Δt along the x and y dimension, respectively. Subsequently, we feed all pairs to ICP matching and cluster association if successful. This procedure can be further simplified by exploiting the clustering indices from HDBSCAN. We provide details in the supplementary material.

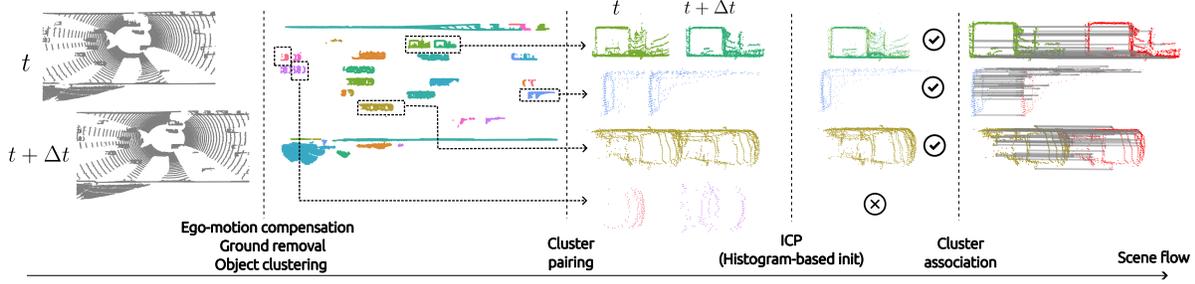


Figure 2. **Overview of ICP-Flow.** Given two full-size LiDAR scans as input, we first do ego-motion compensation and ground removal on each scan. Subsequently, we fuse the non-ground points from both scans and group them into a set of clusters. We pair clusters by spatial locality and feed them to ICP matching for further verification and transformation estimation. We then filter unreliable matches and associate clusters over time. The scene flow is recovered by using the rigid-motion assumption. Crucially, to aid ICP matching, we develop a histogram-based voting strategy for initialization, by exploring the motion rigidity.

3.6. ICP matching

Given a pair of clusters \mathbf{C}_m^t and $\mathbf{C}_n^{t+\Delta t}$ consisting of L_m and L_n points respectively, we leverage ICP [4] to estimate a transformation matrix \mathbf{T}_m^t and calculate quantitative metrics to measure the alignment. However, ICP requires a reasonably good initialization; otherwise, it produces a suboptimal estimation. We refer the readers to the ablation study in the supplementary material for an in-depth analysis.

Histogram-based initialization. We take advantage of the fact that objects rarely change directions sharply within a short temporal window Δt , where $\Delta t \leq 0.1$ seconds, leaving translation the major variable to infer during initialization. Moreover, we explicitly incorporate the rigid motion assumption, indicating that points from the same object share approximately the same translation within a time gap of Δt . Taking this inspiration, we compute a histogram for all translation vectors between a pair of clusters and then select the dominant translation shared by the majority. We construct the histogram \mathbf{H} by discretizing the maximal translation τ_x within Δt along the x dimension into equally spaced bins of 0.1 meters. The same also applies to the y and z dimensions. This results in a histogram \mathbf{H} of size $L_x \times L_y \times L_z$. To conduct voting, we calculate the *point-wise* translation vectors between \mathbf{C}_m^t and $\mathbf{C}_n^{t+\Delta t}$ ¹, namely $\{\mathbf{x}_i^t - \mathbf{x}_j^{t+\Delta t} \mid \mathbf{x}_i^t \in \mathbf{C}_m^t, \mathbf{x}_j^{t+\Delta t} \in \mathbf{C}_n^{t+\Delta t}\}$ where $i = 1, \dots, L_m$, and $j = 1, \dots, L_n$. Subsequently, we discretize each translation vector and cast a vote to its bin in \mathbf{H} . After voting, we localize the bin with the most votes and initialize \mathbf{T}_m^t with the associated translation.

ICP matching. We further refine the initial transformation using ICP [4]. We opt for a conventional ICP implementation rather than recent works that are primarily designed for large-scale point clouds and require advanced GPUs for optimization. We adopt the Pytorch3D implementation [37].

To evaluate the quality of ICP matching, we adopt two

metrics: average distance d between transformed point correspondences and ratio of inliers r , defined by $d = \frac{1}{L_m} \sum_{i=1}^{L_m} \|d_i\|$ and $r = \frac{\sum_{i=1}^{L_m} \mathbb{1}(d_i)}{L_m + L_n - \sum_{i=1}^{L_m} \mathbb{1}(d_i)}$, respectively. Here d_i indicates the distance between a point $\mathbf{x}_i^t \in \mathbf{C}_m^t$ and its nearest neighbor $\mathbf{x}_j^{t+\Delta t} \in \mathbf{C}_n^{t+\Delta t}$ after transformation. $\mathbb{1}(d_i)$ is an indicator function, defined by Eq. (2), that categorizes a pair of correspondences as an inlier if the distance does not exceed τ_{inlier} :

$$\mathbb{1}(d_i) = \begin{cases} 1 & \|d_i\| \leq \tau_{inlier} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

3.7. Cluster association

After ICP matching, we construct a distance matrix \mathbf{M}_d of size $L_m \times L_n$, where $\mathbf{M}_d(m, n)$ indicates the distance d between paired clusters \mathbf{C}_m^t and $\mathbf{C}_n^{t+\Delta t}$. Similarly, we also build an inlier-ratio matrix \mathbf{M}_r where $\mathbf{M}_r(m, n)$ indicates the inlier ratio r between paired clusters. For unpaired clusters, we assign ∞ in \mathbf{M}_d and 0 in \mathbf{M}_r .

Sequentially, for each query cluster at time t , we seek a paired cluster at time $t + \Delta t$ that has the smallest distance d . However, it might be possible that a pair of clusters differ substantially in size, such that the cluster with fewer points always matches well with the other. To prevent this, we reject a pair of clusters once $r < \tau_r$ where τ_r is a predefined threshold. Afterwards, we search for the best match for the query cluster. This is equivalent to an $\arg \min$ over columns for each row in \mathbf{M}_d . Similarly, we reject paired clusters once $d > \tau_d$, where τ_d is also a threshold. For clusters at time t that have no match at time $t + \Delta t$, we simply assign an identity transformation. Finally, we recover the scene flow using Eq. (1).

Alternatively, we can apply Hungarian matching [11] to associate the clusters from two scans, by enforcing one-to-one correspondence. In general, we find that $\arg \min$ matching works well on common benchmarks.

¹This is equivalent to a broadcasted matrix subtraction.

3.8. ICP-Flow pseudo labels as supervision

We also train a feedforward neural network for real-time inference. We supervise model training with the pseudo labels from ICP-Flow. We adopt the same setup as ZeroFlow [44], including both the model architecture and the loss function.

3.9. Implementation details

We adopt the default parameters in Patchwork++ [23] during ground removal. We use the default parameters in HDBSCAN [30, 31], except that *min_cluster_size* is set to 20, below which ICP matching becomes substantially harder. We take maximally 200 clusters after HDBSCAN sorted by the number of points for cluster pairing. For the other clusters, we simply set their transformations to be an identity matrix. Assuming $\Delta t = 0.1$ seconds, we set the maximal translation τ_x and τ_y to be 3.33 meters, which is equivalent to the distance that an agent travels at 120 km/h. Scalar τ_z is set to be 0.1 meters, as objects barely move up/downward. We set the inlier threshold τ_{inlier} during ICP [37, 59] to be 0.1 meters. We set the rejection threshold for cluster association to be $\tau_d = 0.2$ meters or $\tau_r = 0.2$.

Regarding neural network training, we use the Adam optimizer [20] with an initial learning rate of $2e^{-4}$, which is multiplied by 0.1 after 25 epochs. We train the model for 50 epochs on 4 Nvidia V100 GPUs and an Intel(R) Xeon(R) W-2245 CPU @ 3.90GHz. The entire training process takes approximately 5 days on the Waymo scene flow dataset [17, 41]. We conduct model inference on the same device. Our code is available at <https://github.com/yanconglin/ICP-Flow>.

4. Experiments

4.1. Datasets

We conduct experiments on the Waymo [41], nuScenes [5] and Argoverse-v2 [54] datasets, which are the largest datasets for scene flow in autonomous driving. We take full-size LiDAR scans as input without any downsampling.

Waymo. We use the *modified* Waymo dataset released by [17], where the ground truth is calculated from annotated 3D bounding boxes. There are 11,440/4,013/4,032 samples for training/validation/test, where each sample consists of 5 consecutive scans spanning 0.4 seconds, as the LiDAR frequency is 10Hz. The average number of points per scan is 177,000 [9]. We follow [17] to remove the ground points, by applying a threshold along the z axis.

nuScenes. Similar to Waymo, we also use the *modified* nuScenes dataset from [17]. There are 10,921/2,973/2,973 samples for training/validation/test, where each sample contains a sequence of 11 consecutive scans captured at 20Hz. Notably, nuScenes is sparser than Waymo due to the sensor difference (32 beams vs. 64 beams). The average number

of points per scan is 25,000 [9]. We also remove the ground points by thresholding along the z axis [2, 17].

We are also aware of the existence of other subsets for Waymo and nuScenes, such as the ones used in [2, 19, 26, 27]. We choose the subset from [17] because it has (1) abundant samples for test; and (2) paired scans over a longer temporal horizon.

Argoverse-v2. We adopt the recent Argoverse-v2 [54], captured by two roof-mounted 32-beam LiDARs. This dataset only contains paired LiDAR scans at two successive time steps with an interval of 0.1 seconds. We follow the exact preprocessing procedure as in [44] and conduct evaluation on the official validation split. The average number of points per scan is 83,000 [9]. The ground points are removed according to a rasterized heightmap.

4.2. Evaluation

We adopt three metrics for evaluation [9, 17], including (1) 3D end-point-error (EPE, in *meters*) which measures the average L_2 error of all flow vectors; (2) strict accuracy (Acc-S, in %), equivalent to the fraction of points with $EPE \leq 0.05m$ or relative EPE error (to ground truth norm) ≤ 0.05 ; (3) relaxed accuracy (Acc-R, in %), similar to Acc-S but with a threshold of $0.1m$ and 0.1 . We report these three metrics on *static foreground*, *static background*, and *dynamic foreground*² [9]. This provides a more comprehensive evaluation than an overall metric averaged over all points, as *static background* points are the majority in a scene. The evaluation is limited to points within a $64m \times 64m$ area surrounding the ego-car on Waymo and nuScenes [17]. On Argoverse-v2, the evaluation is extended to a $102.4m \times 102.4m$ area [44, 54].

4.3. Baselines

We compare our mode against 5 recent baselines, including RigidFlow [25], NSFP [26], FastNSF [27], FastFlow [19], ZeroFlow [44], PCA [17]. RigidFlow [25] shares the same strategy as ours except that the initial transformation for ICP is estimated by a pre-trained neural network. We use the released checkpoint by the authors (trained on KITTI_r [24]) and report its results. NSFP [26] and FastNSF [27] come from a family of work that employs test-time optimization, with FastNSF being substantially faster, as indicated by its name. Both methods require no training data or manual annotation. ZeroFlow [44] and FastFlow [19] are both data-driven methods. The major difference is that FastFlow is supervised by ground truth labels, while ZeroFlow learns from pseudo labels generated by NSFP [26]. PCA [17] is a fully-supervised data-driven approach that incorporates into its design the multi-body rigidity. For ZeroFlow, we directly use the pre-trained checkpoints released by the authors. We

²A point is considered as *dynamic* if its ground truth velocity is above $0.5m/s$.

Metrics	Label	Time	Dynamic Foreground			Static Foreground			Static Background		
			EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑
NSFP [26]	-	90s	0.0966	61.12	79.64	0.0162	92.26	98.74	0.0538	89.56	93.32
Chodosh <i>et al.</i> [9]	-	93s	0.1081	59.15	78.30	0.0156	95.54	98.80	0.0448	91.35	95.46
FastNSF [27]	-	0.5s	0.2983	32.06	46.84	0.0146	97.69	99.30	0.0402	94.39	96.73
ZeroFlow [44]	-	21ms	0.2229	15.20	36.56	0.0123	96.30	98.46	0.0198	96.75	97.84
RigidFlow [25]	-	0.6s	0.2575	15.99	38.29	0.1299	30.11	58.37	0.2517	18.08	40.15
FastFlow [19]	✓	21ms	0.1950	25.44	49.55	0.0170	89.17	96.19	0.0031	99.15	99.50
PCA [17]	✓	35ms	0.1083	61.93	83.41	0.0156	98.35	97.78	0.0199	98.12	99.56
Ours	-	3.3s	0.0799	76.10	88.53	0.0165	96.46	99.68	0.0270	96.83	98.51
Ours+FNN	-	21ms	0.1254	41.22	67.31	0.0044	98.91	99.45	0.0058	99.58	99.62

Table 1. **Comparison on Waymo dataset [17, 41].** This dataset contains paired LiDAR scans from successive time steps, *after ego motion compensation*. We evaluate all methods using EPE, Acc-S, and Acc-R, on dynamic foreground, static foreground, and static background separately. Overall, our model and its derivatives perform the best over multiple metrics. Notably, we are also able to outperform supervised baselines, particularly for the dynamic foreground. Among all methods, ZeroFlow, FastFlow, and Ours+FNN have identical model designs, thus having the same inference speed.

Metrics	Label	Time	Dynamic Foreground			Static Foreground			Static Background		
			EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑
NSFP [26]	-	26s	0.1328	48.84	70.97	0.0343	87.65	96.37	0.0371	89.19	95.92
Chodosh <i>et al.</i> [9]	-	35s	0.1311	49.40	71.78	0.0261	89.52	97.15	0.0213	91.37	96.51
FastNSF [27]	-	0.5s	0.3684	23.60	39.24	0.0227	94.65	99.02	0.0268	93.13	98.19
ZeroFlow [44]	-	29ms	0.2244	15.82	40.04	0.0205	94.67	97.80	0.0125	97.62	98.88
FastFlow [19]	✓	29ms	0.1836	22.52	50.67	0.0198	92.49	96.86	0.0064	98.38	99.09
Ours	-	3.5s	0.1653	48.61	70.70	0.0391	79.31	95.32	0.0320	86.68	95.68
Ours+FNN	-	29ms	0.1697	36.01	62.88	0.0189	94.76	99.22	0.0035	99.37	99.61

Table 2. **Comparison on Argoverse-v2 dataset [9, 44, 54].** The dataset contains pairs of successive LiDAR scans *after ego motion compensation*. NSFP [26] and Chodosh *et al.* are the state-of-the-art for dynamic foreground. However, they require significantly longer time than others for optimization, up to half a minute. In contrast, Ours+FNN, a feedforward neural network supervised by Ours, is capable of real-time inference without decreasing the performance. Although being less superior on dynamic foreground, Ours+FNN achieves top results on static foreground and background.

choose ZeroFlow-1X as it does not require training on external data [44]. For FastFlow, we directly use the checkpoints from [44], which has done extensive comparisons between FastFlow and ZeroFlow. For PCA [17], we take the official checkpoints released by the authors. For NSFP and FastNSF, we take the official implementation and adapt it to corresponding datasets. We keep the default parameters, except that weight decay is disabled [9]. For Chodosh *et al.* [9], we use a third-party implementation³, as no official code is available. We test baseline models on corresponding datasets *by ourselves* due to the lack of certain metrics.

4.4. Comparison to state-of-the-art

Waymo. Tab. 1 shows the result on the Waymo Open dataset [17, 41]. We compare the EPE, Acc-S, and Acc-R metrics on dynamic foreground, static foreground, and static background separately. Our method outperforms not only the unsupervised competitors but also supervised models trained with massive data and annotation, such as PCA [17] and FastFlow [19], especially on dynamic fore-

ground, *i.e.*, annotated objects that move faster than 0.5m/s. Our advantage over the best-performing baseline NSFP [26] is approximately 1.5cm per point in terms of EPE. Notably, our method not only excels in EPE but also improves Acc-S substantially by more than 10% and Acc-R by 5%. Regarding static objects, most models produce reasonably good results and the performance gap among baselines is marginal. Ours+FNN is a feedforward neural network that shares the same architecture as ZeroFlow [44]. The only difference is the source of supervision. Ours+FNN is supervised by the pseudo labels from *Ours*, while ZeroFlow uses the NSFP pseudo-labels. Although less competitive than Ours, Ours+FNN still outperforms ZeroFlow by a margin of 10 cm, which shows the value of pseudo-labels generated from our method. NSFP and Chodosh *et al.* [9] are also strong competitors in terms of performance, but they are dramatically slow during inference. A single inference takes more than 1 minute, preventing them from real-world deployment. In contrast, Ours only takes around 3 seconds, thus being approximately 30× faster. Ours+FNN further reduces the runtime by more than ×1000, without sacrificing much performance. It is worth mentioning that ZeroFlow

³<https://github.com/kylevedder/zeroflow/blob/master/models/chodosh.py>

training requires calculating NSFP pseudo labels beforehand. In this case, NSFP takes several months of GPU compute [44], while Ours reduces the effort to several days. We also compare Ours to RigidFlow [25] as both models follow the “clustering + ICP” design. The main difference is that RigidFlow requires a deep network for initial pose estimation, while Ours adopts histogram-based initialization. In Tab. 1, Ours achieves $\times 3$ better result on dynamic foreground and $\times 10$ on static part than RigidFlow in terms of EPE, indicating the usefulness of the proposed initialization. We provide an additional comparison on KITTI_o [24] in the supplementary material to further validate the advantage of our design.

Argoverse-v2. We also make comparisons on the recent Argoverse-v2 datasets [9, 54]. Chodosh *et al.* [9] and NSFP [26] are the leading methods and outperform others on the dynamic foreground by approximately 3 cm in EPE. However, as indicated by the running time, they are remarkably slower than others (up to $\times 1000$), due to the time-consuming runtime optimization. FastNSF alleviates this issue but suffers from observable performance drops. In spite of the inferiority on dynamic foreground, Ours+FNN, a feedforward neural network supervised by Ours, excels in static foreground and background. More importantly, it enables real-time inference, which is crucial for processing large volumes of data. When compared to ZeroFlow - another unsupervised model that shares the same architecture, Ours+FNN is able to outperform significantly on dynamic foreground (6cm in EPE and 10% in Acc-S and Acc-R). Overall, we are able to achieve the best result among all models that run in real time.

nuScenes. Additionally, we show a comparison on nuScenes [5], composed of paired scans every 0.05 seconds, after ego-motion compensation. Differing from Waymo and Argoverse-v2, the data is captured by a single 32-beam LiDAR, thus being much sparser. Generally, our model achieves top results in all three categories, compared to other unsupervised baselines, as shown in Tab. 3. Notably, Ours is only marginally worse than the supervised baseline PCA [17] in terms of EPE for the dynamic foreground.

4.5. Scene flow over a longer temporal horizon

So far we compared different methods on scene flow from two successive frames. We also test the capability of various methods on samples with a longer time difference. This is particularly useful when processing temporally downsampled data [38]. Thus we conduct experiments on scene flow estimation from clips of LiDAR scans, each of which contains 5 consecutive scans from Waymo, following [17]. We calculate scene flow between the first frame and the remaining frames, thus resulting in 4 pairs of LiDAR scans whose

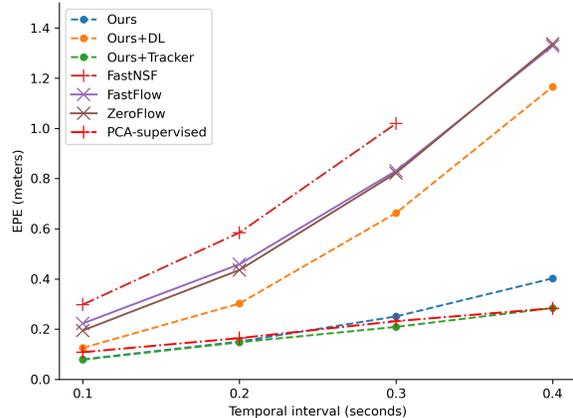


Figure 3. **Scene flow errors with increasing time gap.** We show the EPE values for dynamic foreground with respect to the time duration. As the time gap increases, Ours degrades gracefully and the gap to PCA [17], a supervised model designated for this task, is marginal till 0.3 seconds. In contrast, other methods fail to generalize to a longer duration. Ours+Tracker, an extension of Ours that does tracking over time, is able to achieve comparable results without relying on learning from costly annotation.

time difference gradually increases from 0.1 to 0.4 seconds.

We plot the EPE errors over time for the dynamic foreground in Fig. 3. With the increase of time, the performance of Ours decreases gracefully. The difference between Ours and PCA [17], a fully supervised data-driven approach on this task, is insignificant within a temporal window of 0.3 seconds. In comparison, FastFlow [19], ZeroFlow [44] and Ours+FNN fail to produce a reliable estimation, as the error becomes substantially large, since these methods are not trained for this scenario. FastNSF [27] does not produce reasonable predictions at $t = 0.4s$ and is thus absent from comparison. We exclude NSFP [26] here because its performance is highly unstable. We are unable to find a set of hyperparameters that work at all time steps. To conclude, our design not only works competitively for scene flow from successive scans but also generalizes to further-away scans within a temporal window of 0.4 seconds. Quantitative comparisons are available in Tab. 4.

We also extend Ours to a tracker, namely Ours+Tracker, which associates clusters over time, *i.e.*, over the entire clip. We provide its details in the supplementary material. In Tab. 4, *Ours+Tracker* is able to further improve the result over dynamic foreground by approximately 2cm in EPE as we no longer lose track over time. However, it deteriorates in Acc-S/R, since errors accumulate over time.

4.6. Limitations

Our design is a feature-engineering solution that only exploits geometric information during scene flow. However, it can fail when (1) ground removal and clustering do not

Metrics	Label	Dynamic Foreground			Static Foreground			Static Background		
		EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑
NSFP [26]	-	0.1527	38.82	59.46	0.0406	80.69	91.84	0.0812	66.25	79.92
Chodosh <i>et al.</i> [9]	-	0.1571	33.27	57.17	0.0404	75.82	90.49	0.0776	63.61	80.17
FastNSF [27]	-	0.1591	42.63	61.43	0.0418	80.09	91.84	0.0902	58.81	77.28
PCA [17]	✓	0.1340	31.89	64.91	0.0356	80.34	95.35	0.0514	64.16	88.95
Ours	-	0.1445	49.59	66.02	0.0298	87.26	96.22	0.0403	80.60	91.75
Ours+FNN	-	0.1850	21.52	45.35	0.0150	96.05	98.76	0.0090	98.10	98.81

Table 3. **Comparison on nuScenes dataset [5].** nuScenes dataset contains paired scans captured by a 32-beam LiDAR sensor. Ours outperforms the unsupervised baselines on all metrics, while getting close to the fully supervised baseline PCA [17].

Metrics	Label	Dynamic Foreground			Static Foreground			Static Background		
		EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑	EPE (m) ↓	Acc-S (%) ↑	Acc-R (%) ↑
NSFP [26]	-	0.8533	20.22	36.37	0.0257	95.12	98.25	0.0575	93.67	96.88
FastNSF [27]	-	4.1415	24.72	39.15	0.0360	96.10	98.33	0.6372	93.11	95.77
ZeroFlow [44]	-	0.7097	10.30	25.04	0.0231	94.51	97.14	0.0340	95.71	97.17
FastFlow [19]	-	0.6968	16.46	32.83	0.0241	88.68	95.43	0.0049	99.07	99.47
PCA [17]	✓	0.1970	53.31	77.49	0.0216	97.16	99.44	0.0289	97.16	99.44
Ours	-	0.2209	67.59	84.66	0.0272	96.08	99.16	0.0711	96.49	97.96
Ours+FNN	-	0.5636	30.76	51.73	0.0105	97.44	98.54	0.0085	99.43	99.51
Ours+Tracker	-	0.1799	58.98	80.98	0.0341	88.53	97.73	0.0722	93.74	97.46

Table 4. **Scene flow on Waymo dataset [41], over a longer temporal horizon (5 consecutive frames, up to 0.4 seconds).** Given a clip of 5 consecutive scans, we compute the flow between the first frame and the other frames, leading to 4 pairs per clip. The result is averaged over all points. Most methods fail to generalize to a longer temporal duration, while *Ours* still produces reasonably good results, compared to PCA [17] which is specifically designed for this task. Additionally, we also include Ours+Tracker, an extension of our design that utilizes intermediate scans and tracks clusters iteratively over time. It offers better results than the fully supervised PCA. It is worth mentioning that Ours+Tracker takes as input intermediate scans while others do not.

work decently, resulting in over/under segmentation; (2) there are multiple similar objects nearby in a scene; (3) an object is no longer in the perception range, particularly for fast-moving objects; (4) point density decreases as ICP struggles to match sparse clusters. Moreover, the naive matching strategy for cluster association (Section 3.7) does not take into consideration the one-to-one correspondence, such that a cluster might be matched multiple times. Further, the rigid body assumption may not always hold for deformable objects, such as bendy or articulated buses and trucks.

We include typical failure cases in the supplementary material for qualitative comparison.

5. Conclusion

We propose a learning-free framework for scene flow estimation, particularly for LiDAR-based perception in autonomous driving. Our design is inspired by motion rigidity that assumes objects in a scene move without any deformation. To trace motion, we adopt classic ICP matching which finds the optimal transformation that aligns two clusters. We then recover the scene flow per cluster from the transformation matrix. To aid ICP, we develop a histogram-based voting for translation initialization, enabling better

ICP matching. Further, we train a feedforward neural network that is capable of real-time inference using the pseudo labels from our model. We show quantitatively on Waymo, Argoverse-v2 and nuScenes the advantage of our model over other unsupervised baselines, not only from successive time steps but also from a longer time duration. Future work will fit our design into a neural network for exploiting both geometric and semantic features.

References

- [1] Stefan Baur, David Emmerichs, Frank Moosmann, Peter Pinggera, Bjorn Ommer, and Andreas Geiger. Slim: Self-supervised lidar scene flow and motion segmentation. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [2] 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. 1, 5
- [3] Aseem Behl, Despoina Paschalidou, Simon Donné, and Andreas Geiger. Pointflownet: Learning representations for rigid motion estimation from point clouds. In *Proceedings*

- of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7962–7971, 2019. 1, 2
- [4] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, pages 586–606. Spie, 1992. 2, 3, 4
- [5] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2, 3, 5, 7, 8
- [6] Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 160–172. Springer, 2013. 3
- [7] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992. 3
- [8] 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
- [9] Nathaniel Chodosh, Deva Ramanan, and Simon Lucey. Re-evaluating lidar scene flow for autonomous driving. *arXiv preprint arXiv:2304.02150*, 2023. 2, 5, 6, 7, 8
- [10] Christopher Choy, Wei Dong, and Vladlen Koltun. Deep global registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2514–2523, 2020. 3
- [11] David F Crouse. On implementing 2d rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4):1679–1696, 2016. 4
- [12] Ayush Dewan, Tim Caselitz, Gian Diego Tipaldi, and Wolfram Burgard. Rigid scene flow for 3d lidar scans. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1765–1770. IEEE, 2016. 1, 2
- [13] Guanting Dong, Yueyi Zhang, Hanlin Li, Xiaoyan Sun, and Zhiwei Xiong. Exploiting rigidity constraints for lidar scene flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12776–12785, 2022. 2
- [14] Emeç Erçelik, Ekim Yurtsever, Mingyu Liu, Zhijie Yang, Hanzhen Zhang, Pınar Topçam, Maximilian Listl, Yılmaz Kaan Caylı, and Alois Knoll. 3d object detection with a self-supervised lidar scene flow backbone. In *European Conference on Computer Vision*, pages 247–265. Springer, 2022. 1
- [15] Zan Gojcic, Or Litany, Andreas Wieser, Leonidas J Guibas, and Tolga Birdal. Weakly supervised learning of rigid 3d scene flow. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5692–5703, 2021. 2
- [16] Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, and Panqu Wang. Hplfflownet: Hierarchical permutohedral lattice fliownet for scene flow estimation on large-scale point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3254–3263, 2019. 2
- [17] Shengyu Huang, Zan Gojcic, Jiahui Huang, Andreas Wieser, and Konrad Schindler. Dynamic 3d scene analysis by point cloud accumulation. In *European Conference on Computer Vision*, pages 674–690. Springer, 2022. 2, 5, 6, 7, 8
- [18] Zhao Jin, Yinjie Lei, Naveed Akhtar, Haifeng Li, and Munawar Hayat. Deformation and correspondence aware unsupervised synthetic-to-real scene flow estimation for point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7233–7243, 2022. 2
- [19] Philipp Jund, Chris Sweeney, Nichola Abdo, Zhifeng Chen, and Jonathon Shlens. Scalable scene flow from point clouds in the real world. *IEEE Robotics and Automation Letters*, 7(2):1589–1596, 2021. 5, 6, 7, 8
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [21] Yair Kittenplon, Yonina C Eldar, and Dan Raviv. Flow-step3d: Model unrolling for self-supervised scene flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4114–4123, 2021. 2
- [22] 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. 2
- [23] Seungjae Lee, Hyungtae Lim, and Hyun Myung. Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3d point cloud. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13276–13283. IEEE, 2022. 3, 5
- [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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15577–15586, 2021. 2, 5, 7
- [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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16959–16968, 2022. 2, 5, 6, 7
- [26] Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural scene flow prior. *Advances in Neural Information Processing Systems*, 34:7838–7851, 2021. 1, 2, 5, 6, 7, 8
- [27] Xueqian Li, Jianqiao Zheng, Francesco Ferroni, Jhony Kaesemodel Pontes, and Simon Lucey. Fast neural scene flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9878–9890, 2023. 1, 2, 5, 6, 7, 8
- [28] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. Flownet3d: Learning scene flow in 3d point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 529–537, 2019. 1, 2
- [29] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. Meteor-net: Deep learning on dynamic 3d point cloud sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9246–9255, 2019. 2

- [30] Leland McInnes and John Healy. Accelerated hierarchical density based clustering. In *Data Mining Workshops (ICDMW), 2017 IEEE International Conference on*, pages 33–42. IEEE, 2017. 5
- [31] Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11):205, 2017. 5
- [32] 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. 2
- [33] Himangi Mittal, Brian Okorn, and David Held. Just go with the flow: Self-supervised scene flow estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11177–11185, 2020. 1, 2
- [34] 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 *European Conference on Computer Vision*, pages 424–443. Springer, 2022. 1
- [35] Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Colored point cloud registration revisited. In *Proceedings of the IEEE international conference on computer vision*, pages 143–152, 2017. 3
- [36] Gilles Puy, Alexandre Boulch, and Renaud Marlet. FLOT: Scene Flow on Point Clouds Guided by Optimal Transport. In *European Conference on Computer Vision*, 2020. 2
- [37] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 4, 5
- [38] Davis Rempe, Tolga Birdal, Yongheng Zhao, Zan Gojcic, Srinath Sridhar, and Leonidas J Guibas. Caspr: Learning canonical spatiotemporal point cloud representations. *Advances in neural information processing systems*, 33:13688–13701, 2020. 7
- [39] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *Proceedings third international conference on 3-D digital imaging and modeling*, pages 145–152. IEEE, 2001. 3
- [40] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE international conference on robotics and automation*, pages 3212–3217. IEEE, 2009. 3
- [41] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 2, 3, 5, 6, 8
- [42] Ivan Tishchenko, Sandro Lombardi, Martin R Oswald, and Marc Pollefeys. Self-supervised learning of non-rigid residual flow and ego-motion. In *2020 international conference on 3D vision (3DV)*, pages 150–159. IEEE, 2020. 2
- [43] Arash K Ushani, Ryan W Wolcott, Jeffrey M Walls, and Ryan M Eustice. A learning approach for real-time temporal scene flow estimation from lidar data. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5666–5673. IEEE, 2017. 1, 2
- [44] Kyle Vedder, Neehar Peri, Nathaniel Chodosh, Ishan Khatri, Eric Eaton, Dinesh Jayaraman, Yang Liu, Deva Ramanan, and James Hays. Zeroflow: Fast zero label scene flow via distillation. *arXiv preprint arXiv:2305.10424*, 2023. 2, 5, 6, 7, 8
- [45] Kavisha Vidanapathirana, Shin-Fang Chng, Xueqian Li, and Simon Lucey. Multi-body neural scene flow. *arXiv preprint arXiv:2310.10301*, 2023. 2
- [46] Ignacio Vizzo, Tiziano Guadagnino, Benedikt Mersch, Louis Wiesmann, Jens Behley, and Cyrill Stachniss. Kiss-icp: In defense of point-to-point icp—simple, accurate, and robust registration if done the right way. *IEEE Robotics and Automation Letters*, 8(2):1029–1036, 2023. 3
- [47] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a rigid motion prior. In *2011 International Conference on Computer Vision*, pages 1291–1298. IEEE, 2011. 2
- [48] Christoph Vogel, Konrad Schindler, and Stefan Roth. Piecewise rigid scene flow. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1377–1384, 2013. 2
- [49] 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. 2
- [50] Yue Wang and Justin M Solomon. Deep closest point: Learning representations for point cloud registration. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3523–3532, 2019. 3
- [51] Yuqi Wang, Yuntao Chen, and ZHAO-XIANG ZHANG. 4d unsupervised object discovery. *Advances in Neural Information Processing Systems*, 35:35563–35575, 2022. 1
- [52] Zirui Wang, Shuda Li, Henry Howard-Jenkins, Victor Prisacariu, and Min Chen. Flownet3d++: Geometric losses for deep scene flow estimation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 91–98, 2020. 2
- [53] Andreas Wedel, Thomas Brox, Tobi Vaudrey, Clemens Rabe, Uwe Franke, and Daniel Cremers. Stereoscopic scene flow computation for 3d motion understanding. *International Journal of Computer Vision*, 95:29–51, 2011. 2
- [54] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021. 2, 3, 5, 6, 7
- [55] 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

- [56] Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 37(2):314–333, 2020. 3
- [57] Guangyao Zhai, Xin Kong, Jinhao Cui, Yong Liu, and Zhen Yang. Flowmot: 3d multi-object tracking by scene flow association. *arXiv preprint arXiv:2012.07541*, 2020. 1
- [58] Xiyu Zhang, Jiaqi Yang, Shikun Zhang, and Yanning Zhang. 3d registration with maximal cliques. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17745–17754, 2023. 3
- [59] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3d: A modern library for 3d data processing. *arXiv preprint arXiv:1801.09847*, 2018. 5