

Re-Evaluating LiDAR Scene Flow

Nathaniel Chodosh
Carnegie Mellon University
nchodosh@andrew.cmu.edu

Deva Ramanan*
Carnegie Mellon University
deva@cs.cmu.edu

Simon Lucey*
University of Adelaide
simon.lucey@adelaide.edu.au

Abstract

Popular benchmarks for self-supervised LiDAR scene flow (stereoKITTI, and FlyingThings3D) have unrealistic rates of dynamic motion, unrealistic correspondences, and unrealistic sampling patterns. As a result, progress on these benchmarks is misleading and may cause researchers to focus on the wrong problems. We evaluate a suite of top methods on a suite of real-world datasets (Argoverse 2.0, Waymo, and NuScenes) and report several conclusions. First, we find that performance on stereoKITTI is negatively correlated with performance on real-world data. Second, we find that one of this task’s key components – removing the dominant ego-motion – is better solved by classic ICP than any tested method. Finally, we show that despite the emphasis placed on learning, most performance gains are caused by pre- and post-processing steps: piecewise-rigid refinement and ground removal. We demonstrate this through a baseline method that combines these processing steps with a learning-free test-time flow optimization. This baseline outperforms every evaluated method.

1. Introduction

Research is often guided by improving benchmark performance, but we show that popular scene flow benchmarks may be guiding research in the wrong direction. Scene flow, the 3D analog of optical flow [52], can help autonomous vehicles identify moving objects. Most autonomous vehicles use LiDAR sensors for 3D perception, creating interest in estimating the dynamic motion between successive scans [4, 6, 17, 27, 34, 40]. If this task can be accomplished without relying on labeled data, it can give autonomous vehicles awareness of moving objects outside the detection taxonomy [42]. The potential for label-free detection has led to significant interest in self-supervised scene flow methods [4, 26, 40, 58]. We focus on this self-supervised setting and show that the standard benchmarks have several fundamental flaws. When evaluated more realistically,

*Equal contribution.

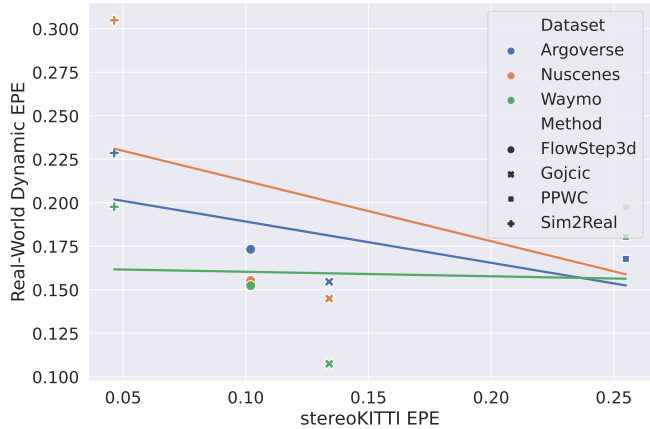


Figure 1. Recent self-supervised scene flow methods [17, 26, 29, 58] typically focus on performance on stereoKITTI [34]. We call attention to several problematic aspects of its semi-synthetic construction. When evaluated on real-world datasets (Waymo, Argoverse, NuScenes) we find a **negative correlation with stereoKITTI performance**. For each method and LiDAR dataset, we plot a point corresponding to the self-reported end-point error on stereoKITTI versus the end-point error on dynamic points of that method trained on real data. For each dataset, we plot the best-fit line to visualize the correlation.

apparent benchmark improvements correspond to stalled or even decreasing real-world performance (Fig. 1).

Most work follows the evaluation framework from FlowNet3D [34], which primarily measures the average end-point-error (EPE) across FlyingThings3D [36] and KITTI-SF [38] (now typically referred to as stereoKITTI). These datasets have several issues: (1) stereoKITTI samples dynamic objects with an artificial pattern that differs from the pattern on static objects, inadvertently providing part of the “answer” to learning-based approaches. (2) Both FlyingThings3D and stereoKITTI ensure successive point clouds are in one-to-one correspondence, which is never the case for actual scans. (3) Both datasets contain an unrealistically high percentage of dynamic points. In contrast, real-world data is dominated by the background.

Together these issues obfuscate the main challenges of

LiDAR scene flow: identifying the few non-static points and estimating their motions robustly given the lack of correspondences. We demonstrate the impact of these issues by evaluating a suite of top methods on several large-scale real-world datasets (Argoverse, NuScenes, Waymo), finding that performance is *negatively correlated* with performance on the standard benchmarks.

We also believe these benchmarks have led researchers to ignore the virtues of classic optimization-based approaches. Many learning-based methods have been proposed for first estimating and removing ego-motion [17,51], but we show that none outperform Iterative Closest Point (ICP). Furthermore, we show that the common pre- and post-processing steps of ground removal and enforcing piecewise-rigidity have a larger impact on performance than any learning strategy. We show this through a baseline method that combines these steps with a test-time optimization flow method. This baseline, without any learning, outperforms every method in our suite and outperforms all self-supervised methods on the self-reported NuScenes [36] and lidarKITTI [17] benchmarks. This leads to the conclusion that current learning methods fail to extract information not present in a single example despite using large amounts of data.

In summary, our main contributions are:

- An investigation of the weakness of current self-supervised scene flow evaluations.
- A *dataless* flow method which gives state-of-the-art results.
- A standard evaluation protocol and codebase along with a “model-zoo” of top methods as well as flow labels for Argoverse 2.0¹.

2. Related Work

Scene Flow: Scene flow was introduced by [53], who posed the problem in the stereo RGB setting and spawned a large body of subsequent work [3, 8, 19, 20, 22, 37, 46, 47, 55, 56]. A related problem is non-rigid registration [1, 10, 24, 31, 45], which is focused on fitting dense point clouds or meshes. We are interested in the setting without images, based purely on sparse LiDAR point clouds.

Optimization based LiDAR Scene Flow: Dataless LiDAR scene flow estimation was first proposed by Dewan *et al.* [11]. Inspired by a non-rigid registration method regularized by the graph Laplacian [13], Pontes *et al.* [48] created an improved method. Their results were further improved upon by Li *et al.* [33] with the implicit regularization of coordinate networks [2, 9, 39, 44, 50]. We use [33] as the

¹Code, weights, and outputs for all the evaluated methods will be released

backbone of our baseline and additionally employ coordinate networks for height-map estimation. Apart from these, the vast majority of recent methods have focused on deep learning based solutions.

Self and Weakly-Supervised LiDAR Scene Flow: “Just Go with the Flow” [40] demonstrated that using a combination of nearest-neighbor and cycle consistency losses was enough to train the FlowNet3D network, avoiding the reliance on labeled data. This led to many other works which adopted similar losses [4, 29, 51, 58]. Others made use of easier to acquire sources of supervision such as foreground/background segmentation masks [12, 17] or addressed the synthetic to real domain gap [26]. Of particular relevance are those methods which make use of ego-motion estimation and piecewise rigid representations [12, 17, 32]. We show that these steps are critical to good performance, but find that ICP vastly outperforms learned approaches and that piecewise rigidity is more effective as a post-processing step than as a loss regularizer.

Ground Segmentation: In robotics, ground segmentation has been studied as a subset of general dataless segmentation [41], traversable area identification [21], and as a pre-processing step for object detection, classification and tracking [25, 30, 43, 59]. Current scene flow methods do not make use of these sophisticated methods and instead rely on basic plane fitting [4, 18, 42]. In order to show how even small improvements to pre- and post- processing steps can drastically improve scene flow estimation, we propose a simple segmentation method based on coordinate networks.

3. Benchmark Issues

Most self-supervised flow estimation works inherit their evaluation protocol from [34], which is based on the synthetic FlyingThings3D [36] and stereoKITTI [16, 38]. Both were created for evaluating RGB-based flow methods and [34] extended them to point clouds by lifting the optical flow and depth annotations to 3D. These datasets and protocols suffer from three main deficiencies.

The one-to-one correspondence assumption removes the main challenge of working with LiDAR data. Since both FlyingThings3D and stereoKITTI are created by lifting optical flow annotations, the point clouds for each input pair are in one-to-one correspondence. For each 3D point \mathbf{p}_i and ground truth 3D flow \mathbf{f}_i in the first frame, there exists a 3D point \mathbf{q}_i in the second frame such that $\mathbf{p}_i + \mathbf{f}_i = \mathbf{q}_i$. Real-world LiDAR scans do not have this property. The presence of correspondences fundamentally changes the LiDAR scene flow problem into a point-matching problem. It has been claimed that randomly sub-sampling the input breaks this correspondence. However, given the total number of points in each scan (90k) and the number of sub-sampled points (8192), there are still an expected 745 corresponding points in the input. Since each example has only a

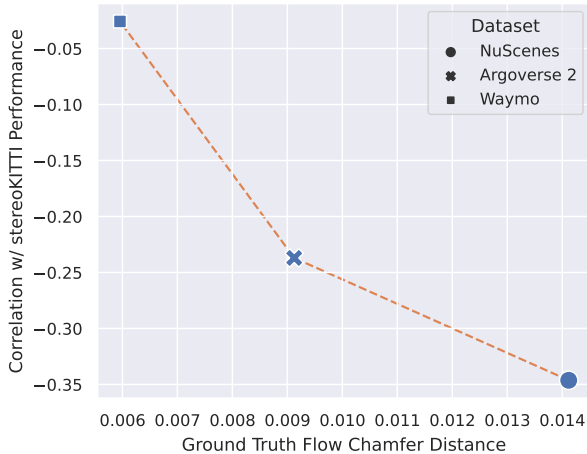


Figure 2. The correlation between performance on stereoKITTI and performance on real datasets as a function of how much those datasets violate the one-to-one correspondence assumption present in stereoKITTI. The more a dataset violates this assumption (higher chamfer distance due to sparser pointclouds), the worse the correlation with stereoKITTI performance.

handful of independent motions, these correspondences are enough to constrain the solution.

This problem is especially severe for self-supervised scene flow methods trained using the chamfer distance, a symmetric extension of the nearest neighbor distance [40]. Essentially, each method deforms the first point cloud using the predicted flows and then computes the distance to the second point cloud. The chamfer distance is an excellent self-supervised objective for point clouds with one-to-one correspondences since the ground truth flow achieves the minimum distance of 0. However, for real-world point clouds with no correspondences the chamfer distance becomes a weaker proxy, and even the ground truth flow will have a non-zero distance. We can quantify this effect by computing the chamfer distance of the ground truth flows for several real-world datasets. In Fig. 2 we show the relationship between that quantity and the slopes of the best-fit lines shown in Fig. 1. As can be seen, the more a dataset violates the one-to-one assumption, the more *good* performance on stereoKITTI indicates *poor* performance on that dataset.

Current benchmarks’ unrealistic rates of dynamic motion make estimating ego-motion seem harder than it is. In any scene, some portion of the measured points belong to the static background rather than dynamically moving objects. In real-world LiDAR scans such as those in NuScenes [7], Waymo [27], or Argoverse 2 [57], the percentage of static points is very high, ranging from 85-100% (Fig. 4). In contrast, FlyingThings3D has 0% static points since it consists of flying things, and the KITTI-SF dataset has approximately 75-85% static points. Dealing

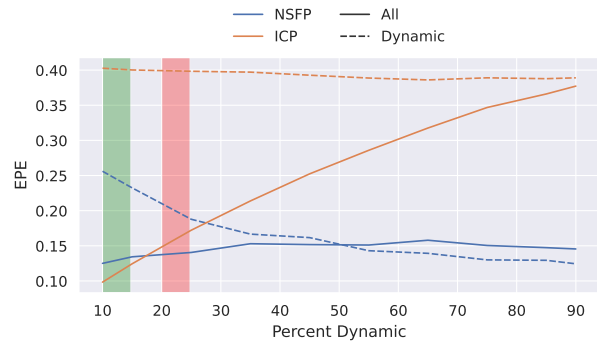


Figure 3. The performance of ICP and NSFP versus the ratio of dynamic points in each example. The green region indicates the ratio found in real data and the red region indicates the ratio found in the stereoKITTI dataset. The unrealistic dynamic ratio of stereoKITTI causes ICP to appear to perform worse than it does on real-world data such as Argoverse.




Figure 4. An analysis of the motion profile of points found in various autonomous driving datasets. The leftmost column is the percentage of points belonging to any tracked object. The right three columns show those points separated by their speed once ego-motion has been removed. All three datasets have a very low rate of dynamic foreground points, in contrast to the typical stereoKITTI benchmark.

with this data imbalance is one of the key difficulties of self-supervised flow estimation but it is not present in the popular benchmarks. This discrepancy also explains why ICP has been ignored as a technique for ego-motion estimation [34].

In order to understand why recent methods have neglected ICP, we need to understand how the amount of dynamic points in a scene affects the performance of ICP. We manually re-sampled each example in the Argoverse 2.0 validation dataset to take a specified number of dynamic and static points and then ran a test-time optimisation scene flow method [33] and ICP [8]. The results (Fig. 3) show that when looking at the total EPE averaged over the entire dataset, like most evaluations, the performance of ICP steadily degrades as the percentage of dynamic points increases. At 20% dynamic points, double the rate of real-world data but approximately the same ratio as KITTI-SF,




Figure 5. Comparison of sampling patterns from KITTI-SF (left) versus the corresponding real LiDAR scan (right). KITTI-SF uses dense CAD models for foreground objects, which makes it easier for learning-based methods to find them among the sparse LiDAR background points.




Figure 6. An example ground truth segmentation from KITTI-SF (left) compared with (right) an example prediction from our local model. Note how the model can even identify parked vs moving cars since only moving vehicles have the dense sampling pattern.

ICP’s performance has degraded to significantly below the performance of NSFP. This explains why early methods which did include ICP in their comparisons [34] were able to outperform it, leading future methods to leave it out even as more realistic data became available.

The sampling pattern of dynamic objects in stereoKITTI gives away the answer in the input. The ground truth flow and disparity of KITTI-SF were created by fitting 3D models to the LiDAR points and 2D annotations [37]. This gives the dynamic objects a distinctive dense sampling pattern that separates them from the background (see Fig. 5). However, determining which points belong to moving objects is one of the main challenges and the sampling pattern effectively gives away the answer in the input. We demonstrate that learning-based system can simply identify moving objects by the sampling pattern without learning anything about the motion.

Model: We use a model that takes a small amount of local context and *no information from the second frame*, to predict motion segmentation for a point \mathbf{p} . That is, the input to our network is the set of vectors $\{\mathbf{q} - \mathbf{p} \mid \forall \mathbf{q} \in \mathcal{N}_r(\mathbf{p})\}$ where \mathcal{N}_r is a ball of radius r centered at \mathbf{p} . We use $r = 0.05$ m in our experiment. Then we use a standard PointNet architecture to predict a binary label for the point.

Training: We train our model on the first 150 examples of KITTI-SF and test on the remaining 50. Note this is essentially the split used by FlowNet3D [34] and the recent RigidFlow [32] for fine-tuning. We train our network using a cross-entropy loss weighted by the inverse frequency of each label. For comparison, we perform the same experiment on real LiDAR scans from Argoverse. We sample 150 scans from the train set and 50 scans from the validation set.




Figure 7. Our proposed pipeline for LiDAR scene flow based largely on classical pre- and post-processing.

Results: Using only information in a 5 cm ball around each point we are able to segment KITTI-SF with high accuracy. We achieve a mean intersection over union on the foreground of 0.83 and 0.97 on the foreground. An example is visualized in Fig. 6. In comparison when we attempt this on real LiDAR scans that do not have a biased sampling pattern, we achieve foreground and background mean intersection over union scores of 0.16 and 0.81 respectively. This demonstrates that the sampling pattern of KITTI-SF makes it trivial to identify foreground points and should not be used for training or validation.

3.1. Real-World Flow Evaluation

Due to these deficiencies, recent works have made progress on improving results on these benchmarks without improving results on real-world data (Fig. 1). Some of the problems could perhaps be addressed through re-sampling, but doing so requires modifying the data based on the ground truth static-vs-dynamic labels. This runs counter to the goal of creating methods that operate on raw, unlabeled data.

Some works have evaluated on real LiDAR data by transferring synthetic labels [17], or more commonly by computing flows from object-level tracks [4, 26, 27, 33]. However, those evaluations have been presented as auxiliary results [4, 33] with limited comparisons to existing methods, or without comparison entirely [27]. As a result, KITTI-SF and FlyingThing3D remain the standard benchmarks. We argue that evaluating on real-data should be the “gold-standard” for scene flow as opposed to synthetic benchmarks.

To further emphasise how top learning-based methods fail when evaluated properly on real data, we also demonstrate that they can be outperformed by a simple test-time optimization baseline. We construct this baseline by making small improvements to standard pre- and post-processing steps and wrapping them around neural scene flow prior.

4. Baseline

Our pipeline (Fig. 7), consists mainly of pre- and post-processing steps around the test-time flow optimization method of [33].

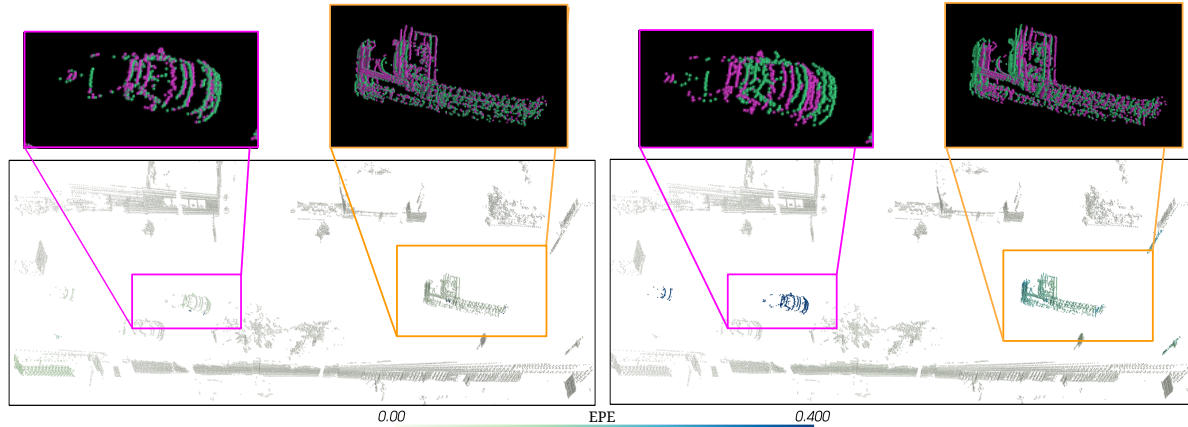


Figure 8. Comparison between NSFP with our proposed pre and post processing steps (**left**) and standard NSFP (**right**). In the (**bottom**) views points are color coded by EPE. The (**top**) detail views show the first and second frames aligned by the predicted flow. NSFP struggles to represent both foreground and background motion. We find that first using ICP to remove ego-motion greatly improves the estimates on dynamic points.

Motion Compensation: Since real-world scenes consist mostly of static background objects, removing the ego-motion of the sensor makes estimating the dynamic motion significantly easier (Fig. 8). For datasets such as Argoverse or Waymo, the ego-motion is provided and we transform the first scan into the coordinate frame of the second. NuScenes and lidarKITTI also have this information but previous work has opted to not use it. In this case, we use ICP [8, 54] to first estimate the ego-motion. We demonstrate experimentally that this is much more effective than learned approaches.

Ground Removal: The sampling pattern of a LiDAR sensor creates “swimming” artifacts as the sensor moves. These artifacts create the appearance of motion and present a large problem for the nearest neighbor loss function used by almost all self-supervised methods. The largest artifacts come from the ground, leading most methods to remove them by height thresholding or by fitting a plane [4, 34]. This can fail when the ground changes elevation either from hills or even sidewalks, see Fig. 9. We propose an improvement based on analyzing the assumptions behind the current approaches. These assumptions are:

1. The sensor has been calibrated such that the ground can be represented as a height map $h = f(x, y)$.
2. Except for a small number of noise returns, all measured points (x, y, z) satisfy $z \geq f(x, y)$
3. The ground function $f(x, y)$ is in some way “simple”. Existing methods make use of this assumption by requiring $f(x, y) = c$ for thresholding or $f(x, y) = ax + by + c$ for plane fitting.

The issue comes from an overly strict application of as-

sumption (3). Instead, we allow our height map to be *piecewise linear* rather than linear. To represent our piecewise linear height map we use a 3-layer coordinate network with ReLU activations and 64 hidden units per layer. To fit it we use assumption (2) to design a one-sided robust loss:

$$\mathcal{L}_{height}(h, z) = \begin{cases} (h - z)^2 & z < h \\ \text{Huber}(h, z) & h \leq z \end{cases} \quad (1)$$

The Huber function [23] allows the loss to ignore points high above the ground. For each point cloud the parameters of our height-map θ_h are found by optimizing

$$\min_{\theta_h} \sum_{i=1}^N \mathcal{L}_{height}(f_{\theta_h}(x_i, y_i), z_i). \quad (2)$$

Finally, any points which are less than 0.3m above our predicted ground are removed.

Scene Flow Estimation: Once the ego-motion and ground points have been removed, we estimate the scene flow using the method of Li *et al.* [33]. Briefly, this means that we represent the forwards and backward flow using two coordinate networks $f_{\theta_{fw}}, f_{\theta_{bw}} \in \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}^{N \times 3}$ which are optimized with gradient descent on the objective:

$$\min_{\theta_{fw}, \theta_{bw}} \mathcal{C}(g_{\theta_{fw}}(\mathbf{P}^t), \mathbf{P}^{t+\Delta}) + \mathcal{C}(g_{\theta_{bw}}(g_{\theta_{fw}}(\mathbf{P}^t)), \mathbf{P}^t). \quad (3)$$

For compactness, we have let $g(\mathbf{X}) = f(\mathbf{X}) + \mathbf{X}$ and \mathcal{C} be the truncated symmetric chamfer distance as described in [33]. We differ from the original method in that we do not use ℓ_2 regularization on the weights. We find that when combined with motion compensation that regularization leads to zero flow predictions everywhere.




Figure 9. An example of our ground removal technique on a Waymo [27] scene with a non-planar ground. In all images, color coding indicates the height of each point. From left to right: the input point cloud, the result of thresholding, our result.

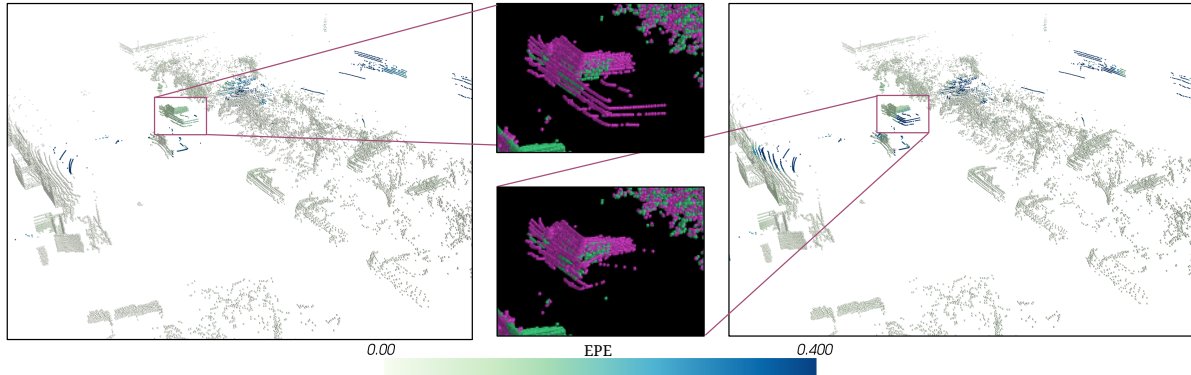


Figure 10. Optimizing for the standard nearest neighbor self-supervised loss can cause mis-predictions in the presence of strong occlusions. Here, the bed of the truck is collapsed into the cab by NSFP (left). Our RANSAC non-rigid refinement step can fix this type of error (right).

Rigid Refinement: Piecewise rigidity has been a common choice of prior for scene flow estimation and has been widely used for LiDAR scene flow [4, 12, 17]. Many learning-based methods require differentiable rigid refinement for loss functions, but we show that simply fitting rigid motion to the final predictions performs better. This is similar to the method of [17] but we use RANSAC to filter outliers. First we use DBSCAN [14] to produce a set of clusters $\{\mathcal{V}_j = \{\mathbf{p}_k\}_{k=1}^K \}_{j=1}^J$. Then for each cluster, we use RANSAC [15] to fit a rigid model to flow predictions for that cluster. That is for each RANSAC iteration we randomly sample 3 points $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$ (the minimum required to get a unique solution) and their associated flow vectors $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3$ and then fit rigid motion parameters by solving:

$$\min_{\mathbf{R} \in SO(3), \mathbf{t} \in \mathbb{R}^3} \sum_{i=1}^3 \|\mathbf{R}\mathbf{p}_i + \mathbf{t} - (\mathbf{p}_i + \mathbf{f}_i)\|_2^2, \quad (4)$$

with the Kabsch algorithm [28]. We then compute the norm of the difference between the raw and rigid flows, considering any below a threshold to be inliers. At the end of T iterations the rigid parameters with the highest number of inliers are selected and the parameters are recomputed with respect to the inlier set producing $\mathbf{R}^*, \mathbf{t}^*$. Since we know that with motion-compensated inputs the vast majority of flows should be zero, we further refine the flows by setting the rigid motion parameters to the identity transform if $\|\mathbf{t}^*\|_2$ is below a threshold. Then all points in the cluster are

Supervision	EPE			AccR	AccS		
	Avg	Dynamic	Static		Dynamic	Dynamic	
			FG	FG			BG
Gojicic [17]	Weak	0.083	0.155	0.064	0.032	0.650	0.368
EgoFlow [51]	Weak	0.205	0.447	0.079	0.090	0.111	0.018
Sim2Real [26]	Synth	0.157	0.229	0.106	0.137	0.565	0.254
PPWC [58]	Self	0.130	0.168	0.092	0.129	0.556	0.229
FlowStep3D [29]	Self	0.161	0.173	0.132	0.176	0.553	0.248
Odometry	None	0.198	0.583	0.010	0.000	0.108	0.002
ICP [8]	None	0.204	0.557	0.025	0.028	0.112	0.015
NSFP [33]	None	0.088	0.193	<u>0.033</u>	0.039	0.542	0.327
Ours	None	0.055	0.105	<u>0.033</u>	0.028	0.777	0.537

Table 1. Quantitative results on Argoverse 2. Our baseline predicts the motion of dynamic objects by 30% over [17], despite that method using ground truth foreground masks for training. We also see from the static background EPE that ICP outperforms learning-based methods at predicting ego-motion.

assigned the flow $\mathbf{f}_i = \mathbf{R}^*\mathbf{p}_i + \mathbf{t}^* - \mathbf{p}_i$. Points that were not assigned to any cluster by DBSCAN have their predictions unchanged. The effect of this step can be seen in Fig. 10.

5. Evaluation

Our central claim is that the use of popular benchmarks causes leading methods to degrade in quality when evaluated on real-world data. The main result in support of this is shown in Fig. 1; here we detail the procedure used to se-

Supervision		EPE			AccR	AccS	
		Avg	Dynamic	Static		Dynamic	Dynamic
		FG	FG	BG	FG	FG	
Gojic [17]	Weak	0.084	0.145	0.060	0.047	0.714	0.436
EgoFlow [51]	Weak	0.236	0.520	0.074	0.114	0.105	0.022
Sim2Real [26]	Synth	0.214	0.305	0.143	0.194	0.426	0.155
PPWC [58]	Self	0.156	0.197	0.097	0.173	0.468	0.183
FlowStep3D [29]	Self	0.156	0.155	0.090	0.224	0.660	0.370
ICP	None	0.176	0.450	0.033	0.046	0.151	0.047
NSFP	None	0.088	0.130	0.034	0.101	0.711	0.447
Ours	None	0.055	0.061	0.020	0.083	0.891	0.681

Table 2. Quantitative results on NuScenes. Our baseline halves the EPE on dynamic objects when compared to the next best method (NSFP). Since our baseline also uses NSFP as its backbone, this indicates that pre- and post-processing steps can have an enormous impact. Again we also see that ICP has the best ego-motion prediction.

Supervision		EPE			AccR	AccS	
		Avg	Dynamic	Static		Dynamic	Dynamic
		FG	FG	BG	FG	FG	
Gojic [17]	Weak	0.059	0.108	0.045	0.025	0.844	0.584
EgoFlow [51]	Weak	0.183	0.390	0.069	0.089	0.178	0.046
Sim2Real [26]	Synth	0.132	0.180	0.075	0.142	0.497	0.217
PPWC [58]	Self	0.132	0.180	0.075	0.142	0.497	0.217
FlowStep3D [29]	Self	0.169	0.152	0.123	0.232	0.708	0.405
ICP	None	0.192	0.498	<u>0.022</u>	0.055	0.172	0.047
NSFP	None	0.100	0.171	0.021	0.108	0.539	0.331
Ours	None	0.041	0.073	0.013	0.039	0.877	0.726

Table 3. Quantitative results on Waymo. Our baseline outperforms all tested methods but notably in this instance [17] performs better than ICP at predicting the background ego-motion.

	Moving		Static	50-50
	EPE	Accuracy Relax	EPE	EPE
Zero	0.6381	0.1632	0.5248	0.5814
ICP	0.2101	0.6151	0.0290	0.1196
PPWC	0.3539	0.2543	0.1974	0.2756
EgoFlow	0.7399	0.0000	0.0570	0.3985
SLIM (U)	0.1050	0.7365	0.0925	0.0987
Ours	0.0625	0.894	0.0660	<u>0.064</u>
SLIM (S)	<u>0.0702</u>	<u>0.8921</u>	<u>0.0499</u>	0.0600

Table 4. The results of our baseline on NuScenes as evaluated by Baur *et al.* [4]. U and S refer to the unsupervised and fully-supervised versions of their method. Our dataless baseline outperforms all self-supervised methods and even achieves comparable and superior results to the fully supervised network.

lect, train, and evaluate the tested methods (Sec. 5.1). We also discuss the performance of our baseline, which we find outperforms all the tested methods, as well as validate these results through comparison to other authors’s self-reported metrics on NuScenes and lidarKITTI (Sec. 5.2). Finally, we evaluate our ground segmentation method (Sec. 5.3).

	Supervision	EPE	AccS	AccR
PPWC [58]	Full	0.710	0.114	0.219
FLOT [49]	Full	0.773	0.084	0.177
MeteorNet [35]	Full	0.277	/	/
Gojic [17]	Weak	0.133	0.460	0.746
Gojic++	Weak	0.102	0.686	0.819
Dong (Waymo Open) [12]	Weak	0.077	0.812	0.906
Dong (Semantic KITTI)	Weak	<u>0.065</u>	<u>0.857</u>	<u>0.940</u>
Ours	None	0.061	0.917	0.962

Table 5. Results of our method on lidarKITTI w/ ground [17]. We outperform all existing methods without using any training data.

5.1. Training and Evaluation Protocol

Evaluation Metrics: We use a set of standard metrics:

- **EPE:** Average end-point-error *i.e.* the ℓ_2 norm of the difference between predicted and ground truth flow.
- **Accuracy Relax:** Ratio of predictions with absolute EPE less than 0.1m or relative error below 0.1.
- **Accuracy Strict:** Same as Accuracy Relax but with a threshold of 0.05.

Rather than averaging these metrics over the entire dataset, we break them into three classes of points: dynamic foreground, static foreground, and static background. Points are considered belonging to the foreground if they are contained in the bounding box of some tracked object and they are considered dynamic if they have a flow magnitude of at least 0.5 m s^{-1} . The choice of threshold is discussed in the supplement. Separating the metrics in this way is vital due to the low ratio of dynamic to static points, without it the metrics are dominated by performance on the static background. In order to produce a single number for ranking purposes we combine the EPE results on the classes into a three-way average similar to [4]. In order to save computation we evaluate on a subset on of each dataset’s validation split. We use the entire NuScenes validation set, and every 5th frame of the Argoverse and Waymo validation sets.

Method Selection: We chose to evaluate 6 methods to serve as a thorough exploration of recent research. Methods were chosen if they presented a self or weakly supervised method, appeared in a recent conference, and made available a working implementation of their method. This led us to choose: PointPWC Net [58], EgoFlow [51], FlowStep3D [29], NSFP [33], Sim2Real [26], and Gojic *et al.* [17]. Additionally, we include as baselines an off-the-shelf ICP [5] implementation [54]. Most of the baseline methods clip points to a depth of 35m [26, 29, 51, 58]. To make comparisons fair we only include points contained in a 70m square around the sensor as was done in [4].

Training Procedure: As much as possible we attempted to use the same training strategy as the original authors, but

we also wished to enforce standardization on the amount of computing resources allocated to each method. As a result, we chose to adopt the two-stage training regime from [58]. First, we train for seven days on a quarter of the whole training set followed by five days of fine-tuning on the entire dataset. Each method was trained using the largest batch size possible on a single NVIDIA T4 GPU (some methods were not set up for multi-GPU training), and using the authors’ optimizer and learning rate schedule. Methods that were able to include ground truth ego-motion in their loss formulation were also given that information. We also used ground truth foreground/background masks to train the weakly supervised method [17]. The only method that required substantial changes was [26] as its training is based on transferring information from a synthetic labeled dataset. The synthetic data comes from a virtual depth camera, so we clip the LiDAR points to match the field of view.

Results: Our results can be found in tables 1, 2, and 3. The main result in Fig. 1 was created by taking the dynamic foreground error for each learning method/dataset combination and plotting it against their self-reported stereoKITTI error. We use the dynamic foreground EPE since this is the hardest and most important component of the real-world scene flow problem. We remove EgoFlow from these plots since it is a large outlier in terms of both real-world dynamic EPE and stereoKITTI EPE.

We also claim that synthetic benchmarks are causing researchers to focus on the wrong problems, which we test by comparing the tested methods to our simple test-time optimization baseline. Validating this claim, we can observe that our baseline outperforms all tested methods despite not using any training data; even Gojcic *et al.* who use ground truth foreground masks. Each tested method claims some learning-based novelty as its main contribution, which is then shown to be effective through experiments on KITTI-SF and FlyingThings3D. However, none can outperform a carefully designed baseline when evaluated on real data. Further validating this claim is the fact that no method was able to match the performance of ICP at predicting the ego-motion. This is in spite of the fact that several of the methods [4, 17, 51] explicitly claim ego-motion predictions as a benefit of their architectures. As discussed in Sec. 3, this results from the unrealistic dynamic ratio found in the standard benchmarks. Gojcic *et al.* comes the closest to outperforming ICP and does so on Waymo. However, this is because it essentially incorporates ICP as a part of its final test-time refinement of the ego-motion.

5.2. Existing Benchmark Comparison

There may be some concern that the tested methods performed poorly compared to the baseline due to a lack of tuning of hyperparameters. To address this concern, we also compare our baseline to self-reported metrics on two real-

world datasets: NuScenes as evaluated in [4] and lidarKITTI which was introduced by the Gojcic *et al.* [17].

NuScenes: The authors of SLIM [4] also reported numbers on real-world data. As can be seen in Tab. 4 our method vastly outperforms SLIM and the other methods they tested on the main 50-50 EPE metric, dynamic EPE, and dynamic accuracy. Again, ICP performs far better than every other method on static points. We also include the SLIM results when trained in a supervised manner and show that we achieve comparable and some superior results, despite not using any training data.

lidarKITTI: lidarKITTI [17] was generated by associating KITTI-SF labels with the raw LiDAR scans. We use the version with ground points as ground removal is a component of our method. Ground point flow is set to the ego-motion. The results are shown in Tab. 5 and further confirm that our baseline performs well even compared to self-reported numbers. It substantially outperforms all but one of the existing self- and weakly-supervised methods.

5.3. Ground Segmentation

We close with an evaluation of our ground segmentation method. First, we present a qualitative example (Fig. 9) showing our method effectively handling a scene with a non-planar ground. To quantitatively analyze our method we need to consider what constitutes failure. The subtlety lies in the ill-defined nature of the ground, making it impossible to define normal precision and recall metrics. Given that our goal is motion estimation, the failure we are most concerned with is classifying a dynamic point as belonging to the ground. Therefore we look at the rate at which we make this error. On NuScenes we find that 99.3% of the time points which are classified as ground are in fact static and achieve a rate of 99.4% on Waymo.

6. Conclusion

We re-examined the evaluations of self- and weakly-supervised scene flow methods in the context of autonomous driving and found several deficiencies. We claimed that these deficiencies are impacting the quality and types of recently proposed methods. To provide evidence for this claim we evaluated a large number of top methods on several real-world datasets. We found a negative correlation between performance on the standard stereoKITTI benchmark and performance on all of the real-world datasets. Additionally, we proposed a dataless estimation technique that far outperformed the existing approaches, demonstrating that focus on the current benchmarks is causing researchers to ignore effective methods. Given that our baseline is based on pre- and post-processing techniques, we believe that other methods not based on test-time optimization will also benefit from them.

References

- [1] Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–8. IEEE, 2007. [2](#)
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2565–2574, 2020. [2](#)
- [3] Tali Basha, Yael Moses, and Nahum Kiryati. Multi-view scene flow estimation: A view centered variational approach. *Int. J. Comput. Vis.*, 101(1):6–21, 2013. [2](#)
- [4] 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 *Int. Conf. Comput. Vis.*, pages 13126–13136, 2021. [1, 2, 4, 5, 6, 7, 8](#)
- [5] Albert E Beaton and John W Tukey. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*, 16(2):147–185, 1974. [7](#)
- [6] Aseem Behl, Despoina Paschalidou, Simon Donné, and Andreas Geiger. Pointflownet: Learning representations for rigid motion estimation from point clouds. In *Int. Conf. Comput. Vis.*, pages 7962–7971, 2019. [1](#)
- [7] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multi-modal dataset for autonomous driving. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11621–11631, 2020. [3](#)
- [8] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Img. Vis. Comput.*, 10(3):145–155, 1992. [2, 3, 5, 6](#)
- [9] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5939–5948, 2019. [2](#)
- [10] Haili Chui and Anand Rangarajan. A new point matching algorithm for non-rigid registration. *Comput. Vis. Img. Und.*, 89(2-3):114–141, 2003. [2](#)
- [11] Ayush Dewan, Tim Caselitz, Gian Diego Tipaldi, and Wolfram Burgard. Rigid scene flow for 3d lidar scans. In *Int. Conf. Intel. Rob. Sys.*, pages 1765–1770. IEEE, 2016. [2](#)
- [12] Guanting Dong, Yueyi Zhang, Hanlin Li, Xiaoyan Sun, and Zhiwei Xiong. Exploiting rigidity constraints for lidar scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12776–12785, 2022. [2, 6, 7](#)
- [13] Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. Smooth shells: Multi-scale shape registration with functional maps. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12265–12274, 2020. [2](#)
- [14] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Int. Conf. Know. Disc. Data Min.*, KDD’96, page 226–231. AAAI Press, 1996. [6](#)
- [15] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Comm. ACM*, 24(6):381–395, 1981. [6](#)
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3354–3361. IEEE, 2012. [2](#)
- [17] Zan Gojcic, Or Litany, Andreas Wieser, Leonidas J Guibas, and Tolga Birdal. Weakly supervised learning of rigid 3d scene flow. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5692–5703, 2021. [1, 2, 4, 6, 7, 8](#)
- [18] Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, and Panqu Wang. Hplflownet: Hierarchical permutohedral lattice flownet for scene flow estimation on large-scale point clouds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3254–3263, 2019. [2](#)
- [19] Simon Hadfield and Richard Bowden. Kinecting the dots: Particle based scene flow from depth sensors. In *Int. Conf. Comput. Vis.*, pages 2290–2295. IEEE, 2011. [2](#)
- [20] Simon Hadfield and Richard Bowden. Scene particles: Unregularized particle-based scene flow estimation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(3):564–576, 2013. [2](#)
- [21] Michael Himmelsbach, Felix V Hundelshausen, and H-J Wuensche. Fast segmentation of 3d point clouds for ground vehicles. In *Intel. Veh. Symp. IV*, pages 560–565. IEEE, 2010. [2](#)
- [22] Michael Hornacek, Andrew Fitzgibbon, and Carsten Rother. Sphereflow: 6 dof scene flow from rgb-d pairs. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3526–3533, 2014. [2](#)
- [23] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992. [5](#)
- [24] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, and Andrew Fitzgibbon. Kinectfusion: Real-time 3d reconstruction and interaction using a moving depth camera. In *User Int. Soft. Tech.*, ’11, page 559–568, New York, NY, USA, 2011. Association for Computing Machinery. [2](#)
- [25] Víctor Jiménez, Jorge Godoy, Antonio Artuñedo, and Jorge Villagra. Ground segmentation algorithm for sloped terrain and sparse lidar point cloud. *IEEE Access*, 9:132914–132927, 2021. [2](#)
- [26] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7233–7243, 2022. [1, 2, 4, 6, 7, 8](#)
- [27] Philipp Jund, Chris Sweeney, Nichola Abdo, Zhifeng Chen, and Jonathon Shlens. Scalable scene flow from point clouds in the real world. *IEEE Rob. Aut. Letters*, 7(2):1589–1596, 2021. [1, 3, 4, 6](#)
- [28] Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976. [6](#)
- [29] Yair Kittenplon, Yonina C Eldar, and Dan Raviv. Flow-step3d: Model unrolling for self-supervised scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4114–4123, 2021. [1, 2, 6, 7](#)

- [30] Seungjae Lee, Hyungtae Lim, and Hyun Myung. Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3d point cloud. In *Int. Conf. Intel. Rob. Sys.*, pages 13276–13283. IEEE, 2022. 2
- [31] Hao Li, Robert W Sumner, and Mark Pauly. Global correspondence optimization for non-rigid registration of depth scans. In *Comput. Graph. For.*, volume 27, pages 1421–1430. Wiley Online Library, 2008. 2
- [32] Ruibo Li, Chi Zhang, Guosheng Lin, Zhe Wang, and Chunhua Shen. Rigidflow: Self-supervised scene flow learning on point clouds by local rigidity prior. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16959–16968, 2022. 2, 4
- [33] Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural scene flow prior. *Adv. Neural Inform. Process. Syst.*, 34:7838–7851, 2021. 2, 3, 4, 5, 6, 7
- [34] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. FlowNet3d: Learning scene flow in 3d point clouds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 529–537, 2019. 1, 2, 3, 4, 5
- [35] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. MeteorNet: Deep learning on dynamic 3d point cloud sequences. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9246–9255, 2019. 7
- [36] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4040–4048, 2016. 1, 2
- [37] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Int. Conf. Comput. Vis.*, pages 3061–3070, 2015. 2, 4
- [38] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3d estimation of vehicles and scene flow. *ISPRS Annals Photo. Rem. Sens. Spat. Inf. Sci.*, 2:427, 2015. 1, 2
- [39] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4460–4470, 2019. 2
- [40] Himangi Mittal, Brian Okorn, and David Held. Just go with the flow: Self-supervised scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2020. 1, 2, 3
- [41] Frank Moosmann, Oliver Pink, and Christoph Stiller. Segmentation of 3d lidar data in non-flat urban environments using a local convexity criterion. In *Intel. Veh. Symp. IV*, pages 215–220. IEEE, 2009. 2
- [42] 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 *Eur. Conf. Comput. Vis.*, pages 424–443. Springer, 2022. 1, 2
- [43] Patiphon Narksri, Eijiro Takeuchi, Yoshiki Ninomiya, Yoichi Morales, Naoki Akai, and Nobuo Kawaguchi. A slope-robust cascaded ground segmentation in 3d point cloud for autonomous vehicles. In *Int. Conf. Intel. Trans. Sys.*, pages 497–504. IEEE, 2018. 2
- [44] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 165–174, 2019. 2
- [45] Mark Pauly, Niloy J Mitra, Joachim Giesen, Markus H Gross, and Leonidas J Guibas. Example-based 3d scan completion. In *Symp. Geom. Proc.*, pages 23–32, 2005. 2
- [46] Jean-Philippe Pons, Renaud Keriven, and Olivier Faugeras. Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. *Int. J. Comput. Vis.*, 72(2):179–193, 2007. 2
- [47] J-P Pons, Renaud Keriven, O Faugeras, and Gerardo Hermosillo. Variational stereovision and 3d scene flow estimation with statistical similarity measures. In *Int. Conf. Comput. Vis.*, volume 2, pages 597–597. IEEE Computer Society, 2003. 2
- [48] Jhony Kaesemodel Pontes, James Hays, and Simon Lucey. Scene flow from point clouds with or without learning. In *Int. Conf. 3D Vis.*, pages 261–270. IEEE, 2020. 2
- [49] Gilles Puy, Alexandre Boulch, and Renaud Marlet. Flot: Scene flow on point clouds guided by optimal transport. In *Eur. Conf. Comput. Vis.*, pages 527–544. Springer, 2020. 7
- [50] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Adv. Neural Inform. Process. Syst.*, 33:7462–7473, 2020. 2
- [51] Ivan Tishchenko, Sandro Lombardi, Martin R Oswald, and Marc Pollefeys. Self-supervised learning of non-rigid residual flow and ego-motion. In *Int. Conf. 3D Vis.*, pages 150–159. IEEE, 2020. 2, 6, 7, 8
- [52] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *IEEE Conf. Comput. Vis. Pattern Recog.*, volume 2, pages 722–729. IEEE, 1999. 1
- [53] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *Int. Conf. Comput. Vis.*, volume 2, pages 722–729. IEEE, 1999. 2
- [54] 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 Rob. Aut. Letters*, 8(2):1–8, 2023. 5, 7
- [55] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a rigid motion prior. In *Int. Conf. Comput. Vis.*, pages 1291–1298. IEEE, 2011. 2
- [56] Christoph Vogel, Konrad Schindler, and Stefan Roth. Piecewise rigid scene flow. In *Int. Conf. Comput. Vis.*, pages 1377–1384, 2013. 2
- [57] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Neur. Inform. Process. Syst. Data. Bench. Track*, 2021. 3
- [58] 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 *Eur. Conf. Comput. Vis.*, pages 88–107. Springer, 2020. 1, 2, 6, 7, 8

- [59] Dimitris Zermas, Izzat Izzat, and Nikolaos Papanikolopoulos. Fast segmentation of 3d point clouds: A paradigm on lidar data for autonomous vehicle applications. In *Int. Conf. Rob. Aut.*, pages 5067–5073. IEEE, 2017. [2](#)