

SCENE FLOW AS A PARTIAL DIFFERENTIAL EQUATION

Kyle Vedder^{1,2*} Neehar Peri^{2,3} Ishan Khatri³ Siyi Li¹ Eric Eaton¹
 Mehmet Kocamaz² Yue Wang² Zhiding Yu² Deva Ramanan³ Joachim Pehserl²
¹University of Pennsylvania ²NVIDIA ³Carnegie Mellon University

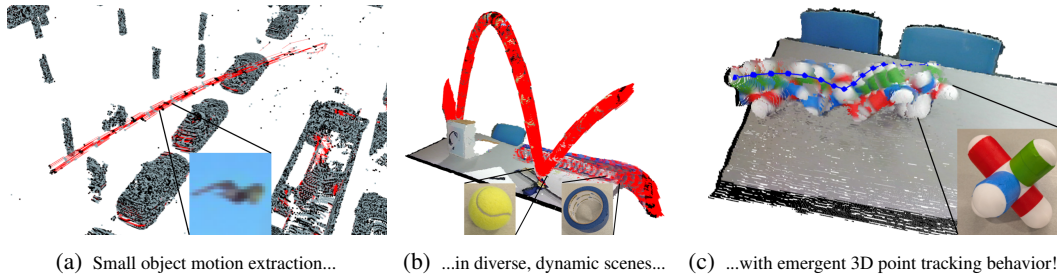


Figure 1: EulerFlow is able to capture the motion of small, fast moving objects with few lidar points, such a bird flying in front of an autonomous vehicle (Figure 1a). EulerFlow’s flexibility allows it to estimate scene flow for fast-moving table top objects *without additional hyperparameter tuning* (Figure 1b). EulerFlow’s PDE estimate exhibits emergent 3D point tracking behavior without explicit long-horizon supervision (Figure 1c). Note that point clouds are shown in color for visualization purposes only; RGB is not used during optimization.

Interactive scene visualizations at vedder.io/eulerflow

ABSTRACT

We reframe scene flow as the task of estimating a continuous space-time PDE that describes motion for an entire observation sequence, represented with a neural prior. Our method, *EulerFlow*, optimizes this neural prior estimate against several multi-observation reconstruction objectives, enabling high quality scene flow estimation via pure self-supervision on real-world data. EulerFlow works out-of-the-box without tuning across multiple domains, including large-scale autonomous driving scenes and dynamic tabletop settings. Remarkably, EulerFlow produces high quality flow estimates on small, fast moving objects like birds and tennis balls, and exhibits emergent 3D point tracking behavior by solving its estimated PDE over long-time horizons. On the Argoverse 2 2024 Scene Flow Challenge, EulerFlow outperforms *all* prior art, surpassing the next-best *unsupervised* method by more than 2.5 \times , and even exceeding the next-best *supervised* method by over 10%.

1 INTRODUCTION

Scene flow estimation is the task of describing motion with per-point 3D motion vectors between temporally successive point clouds (Dewan et al., 2016; Liu et al., 2019; Erçelik et al., 2022; Jund et al., 2021; Zhang et al., 2024b; Vedder et al., 2024; Khatri et al., 2024). Such per-point motion estimates are critical for autonomy in diverse environments, e.g. maneuvering around open-world objects like flying debris (Peri et al., 2022a) or folding deformable cloth (Weng et al., 2022). Importantly, scene flow estimation requires not only an understanding of object *geometry*, but also its *motion*. However, scene flow methods broadly do not work on smaller objects (Khatri et al., 2024). For example, in the autonomous vehicles domain, Khatri et al. highlight that even supervised methods struggle to describe the majority of pedestrian motion, with unsupervised methods failing dramatically. Scene flow promises to be a powerful primitive for understanding the dynamic world, but such failures

*Corresponding email: kvedder@seas.upenn.edu

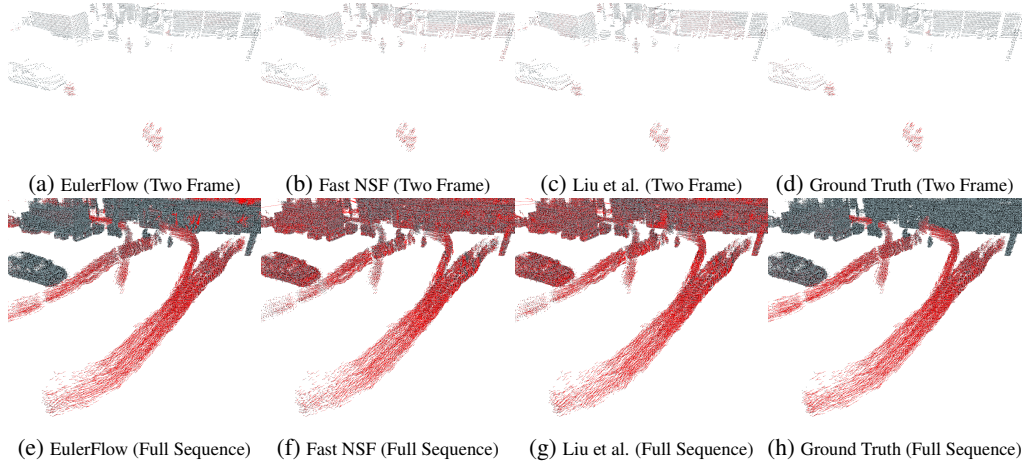


Figure 2: We visualize an example of five pedestrians crossing the street in front of a stopped car, cherrypicked to have unusually high density lidar returns, making it particularly easy to estimate flow. Figures 2a–2d depict a two-frame flow visualization of EulerFlow and several strong baselines. Notably, only visualizing flow over two frames makes it difficult to distinguish flow quality. In contrast, Figures 2e–2h depict flow vectors over the full sequence, making differences in quality clear; for example, EulerFlow is the only one without artifacts on the stopped car.

explain why it has limited adoption in downstream applications like tracking (Zhai et al., 2020) or open-world object extraction (Najibi et al., 2022).

Scene Flow via PDE. In Figure 2, visual assessment of scene flow quality is much easier in an accumulated global frame; while incomplete due to an implicit time axis, these accumulated flow vectors allow viewers to imagine how positions in 3D space evolve over *many* timesteps, and compare that to predicted flows. This imagination of scene flow as continuous motion over large time intervals motivates us to model scene flow as a Partial Differential Equation (PDE) that describes the scene’s instantaneous motion across continuous position and time; this is the same formalism used to represent fluid flow or heat flow in classical mechanics. Scene flow estimation then becomes the task of estimating this PDE. We can straightforwardly represent this PDE estimate with a neural prior (Li et al., 2021b) and optimize it against scene flow surrogate objectives, both over single frame pairs and extended *across arbitrary time intervals*, unlocking new optimization objectives that produce better quality estimates. We formalize this in Section 3 and propose the *Scene Flow via PDE* framework.

EulerFlow. We instantiate Scene Flow via PDE with standard point cloud distance objectives like Chamfer Distance to create *EulerFlow*. Notably, EulerFlow outperforms *all* prior art, supervised or unsupervised, on the Argoverse 2 2024 Scene Flow Challenge and Waymo Open Scene Flow benchmark. It outperforms prior *unsupervised* methods by a large margin ($> 2.5\times$ mean dynamic error reduction), and is able to capture small, fast moving objects, including those outside of labeled taxonomies (e.g. the flying bird in Figure 1a). Due to its simplicity, EulerFlow is able to provide high quality scene flow out-of-the-box on real-world data for other important domains such as dynamic tabletop settings (Figure 1b) *without* domain-specific tuning. Finally, we show that simple PDE solving techniques like Euler integration can be used to extract 3D point tracking behaviors (Figure 1c), which serves as both an exciting emergent behavior as well as a mechanism for visualizing and interpreting the quality of the continuous PDE estimate.

We present four primary contributions:

- We propose *Scene Flow via PDE* (SFvPDE), a reframing of scene flow estimation as the task of fitting a PDE that describes the change of continuous positions over continuous time, unlocking a new class of surrogate objectives that enable better scene flow estimates.
- We instantiate SFvPDE with *EulerFlow*, a flexible **unsupervised** scene flow method that achieves **state-of-the-art** performance on the Argoverse 2 2024 Scene Flow Challenge, **beating all prior supervised and unsupervised methods**.
- We study EulerFlow and show its strong performance is derived from its ability to optimize its PDE estimate against the full sequence of observations over arbitrary time horizons.
- We show that EulerFlow’s simple, flexible formulation allows it to run unmodified on a variety of domains, with emergent capabilities like 3D point tracking behavior.

2 BACKGROUND AND RELATED WORK

Evaluation. Dewan et al. formalized scene flow for point clouds as the task of estimating motion between point cloud P_t at time t and point cloud P_{t+1} at $t + 1$ by estimating the true flow $F_{t,t+1}^*$, i.e. true residual vectors for every point in P_t that describe its movement to its associated position at $t + 1$. Error is computed by measuring the per-point endpoint distance between estimated and ground truth vectors. Historically, these errors are reported with a per-point average (*Average EPE*); however, as Chodosh et al. show, Average EPE is dominated by background points, preventing meaningful measurement of non-ego object motion descriptions. Khatri et al. address this shortcoming with *Bucket Normalized EPE*, which reports per-class performance normalized by speed, allowing for direct comparisons across classes with very different average speeds (e.g. pedestrians and cars). Bucket Normalized EPE is the basis for the *Argoverse 2024 Scene Flow Challenge*¹, where methods are ranked by the mean error of their motion descriptions (*mean Dynamic Normalized EPE*).

Input / Output Formulation. Dewan et al.’s choice to formulate scene flow using *only* two input frames is arbitrary; it’s the minimal information needed to extract rigid motion, but there are not real-world problems constrained to *only* have access to two frames. Indeed, using five or ten frames of past observations is standard practice in the 3D detection literature (Zhu et al., 2019; Vedder & Eaton, 2022; Peri et al., 2022b; 2023; Nalty et al., 2022), and multi-frame formulations have started to appear in the scene flow literature: Liu et al. (2024) and Flow4D (Kim et al., 2024) use three (P_{t-1}, P_t, P_{t+1}) and five input frames (P_{t-3}, \dots, P_{t+1}) respectively to predict $\hat{F}_{t,t+1}$. As we discuss in Section 3, rather than just using more observations to estimate flow for a single frame pair, we formulate scene flow as a joint estimation problem: given the full observation sequence (P_0, \dots, P_N), we estimate *all* flows $\hat{F}_{0,1}, \dots, \hat{F}_{N-1,N}$ between *all* adjacent observations.

Feedforward Methods. Feedforward networks are a popular class of scene flow methods due to their fast inference speed (Liu et al., 2019; Behl et al., 2019; Tishchenko et al., 2020; Kittenplon et al., 2021; Wu et al., 2020; Puy et al., 2020; Li et al., 2021a; Jund et al., 2021; Gu et al., 2019; Batrawy et al., 2022; Wang et al., 2022b; Kim et al., 2024; Zhang et al., 2024a). While they are often trained with supervised labels, recent work has developed distillation pipelines that leverage unsupervised pseudolabelers (Vedder et al., 2024; Zhang et al., 2024b; Lin & Caesar, 2024).

Neural Scene Flow Prior. Li et al. (2021b) propose Neural Scene Flow Prior (NSFP), a widely adopted unsupervised scene flow approach. NSFP uses the inductive bias of the smooth, restricted learnable function class of two ReLU MLP coordinate networks (8 hidden layers of 128 neurons; one to estimate forward flow \hat{F}^+ and to estimate reverse flow \hat{F}^-) to minimize

$$\text{TruncatedChamfer}(P_t + \hat{F}^+, P_{t+1}) + \left\| P_t + \hat{F}^+ + \hat{F}^- - P_t \right\|_2, \quad (1)$$

where TruncatedChamfer is defined as the standard L_2 Chamfer distance, but with per-point distances above 2 meters set to zero in order to reduce the influence of outliers. NSFP is optimized for at most 1000 steps with early stopping.

Motion Beyond Two Frames. Wang et al. (2022a) tackles the adjacent problem of estimating 3D point *trajectories* over 25 frames with Neural Trajectory Prior (NTP) by jointly optimizing three separate ReLU MLP neural priors: 1) a sinusoidal embedded, time conditioned, 25 frame trajectory basis estimator (embed(t) \mapsto $256 \times 25 \times 3$ tensor, where 256 is the dimension of the trajectory basis), 2) a coordinate network bottleneck encoder, and 3) a bottleneck decoder to estimate a per-point linear combination over the learned trajectories. Trajectories are optimized for both a one-frame lookahead L_2 Chamfer loss and a cyclic consistency loss over the entire velocity space trajectory.

Deformation in Reconstruction. Nerfies (Park et al., 2021) and DynamicFusion (Newcombe et al., 2015) estimate a deformation field to warp a canonical frame to explain the observed frame. While capable of describing small motions, these methods require a canonical frame that contains all of the relevant geometry to deform; however, in large, highly dynamic scenes like autonomous driving, there is often no frame that contains all moving objects. By comparison, Scene Flow via PDE does not assume the existence of a canonical frame, instead only describing how the scene changes.

¹<https://www.argoverse.org/sceneflow>

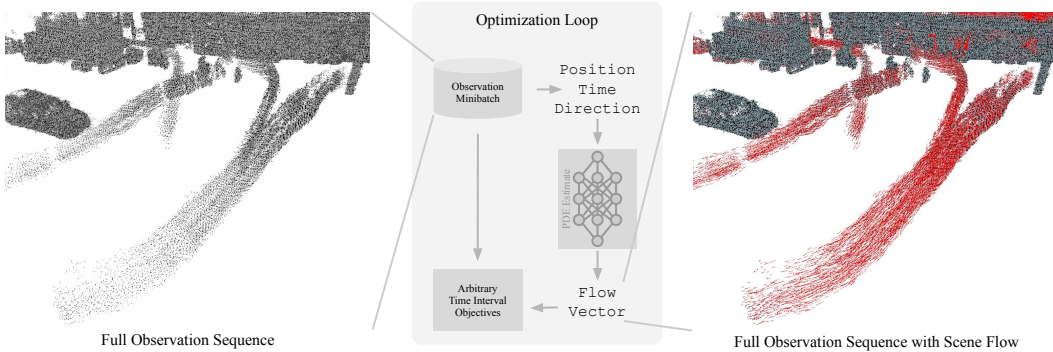


Figure 3: Overview of our *Scene Flow via PDE* framework, which estimates a position PDE across the entire observation sequence by optimizing against multi-frame objectives. This PDE estimate is represented with a neural prior (Li et al., 2021b), providing a flexible, general representation for describing position-time motion.

3 SCENE FLOW VIA PDE

Rather than consuming an arbitrary history of observations to estimate the true flow for a single adjacent observation pair, we pose the problem of jointly estimating the true flow $F_{t,t+1}^*$ for all adjacent point clouds P_t, P_{t+1} in the entire sequence (P_0, \dots, P_N) . To do this, rather than describing scene flow as positional change over a fixed interval ($F_{t,t+1}^*$ are residual vectors over the interval t to $t + 1$) as we did in Section 2, we can instead express these changes as a differential equation that describes positional change over *continuous* time.

Given some scene, let $L^*(\mathbf{x}; t)$ be the continuous function that describes the true position of coordinate $\mathbf{x} \in \mathbb{R}^3$ from a canonical frame at some arbitrary time $t \in \mathbb{R}$. The function describing how locations change over time would be L^* 's partial derivative, the Partial Differential Equation $\frac{\partial L^*}{\partial t}$. This formulation allows us to recover flow vectors over any arbitrary continuous time interval; for example, solving the PDE over the time interval t to $t + 1$ with $\vec{0}$ initial conditions and querying that function with the points of P_t will recover $F_{t,t+1}^*$. Additionally, it also avoids the important limitation of requiring a canonical frame; the canonical frame we describe defining L^* disappears during differentiation, as $\frac{\partial L^*}{\partial t}$ merely describes its change.

While we do not have access to the true L^* or its PDE $\frac{\partial L^*}{\partial t}$, we can estimate it by optimizing a single neural prior θ against surrogate objectives. This framing, which we formalize into the *Scene Flow via PDE* framework (SFvPDE; Figure 3), allows the representation of θ to benefit from constructive interference between objectives, as well as formulate objectives over arbitrarily long time horizons, enabling an optimization driven estimate of $\frac{\partial L^*}{\partial t}$ (Equation 2).

$$\frac{\partial L^*}{\partial t} \simeq \text{SFvPDE}(\theta, t) \tag{2}$$

In addition to not requiring a canonical frame definition, this formulation also benefits from the use of ODE solvers to extract residual vectors rather than regressing them directly. For example, in the deformation fields described in Section 2, the deformation estimator must directly output highly non-linear warps of the canonical frame across arbitrary time. In contrast, SFvPDE only requires that methods estimate local derivatives, with significant non-linear expressivity provided by the ODE solver's numerical integration.

4 EULERFLOW

We straightforwardly instantiate Scene Flow via PDE with *EulerFlow*, a point cloud only scene flow method². EulerFlow is far simpler than prior art, requiring only a single optimization loop over a

²Visualizations shown in color for better viewing. EulerFlow can also use monodepth estimates (Supp. A.2)

single neural prior. Our neural prior θ is a straightforward extension to NSFP’s coordinate network prior³; however, instead of taking a 3D space vector (positions $X, Y, Z \in \mathbb{R}$) as input, we encode a 5D space-time-direction vector: positions $X, Y, Z, \in \mathbb{R}$, sequence normalized time $t \in [-1, 1]$ (i.e. the point cloud time scaled to this range), and direction $d \in \{\text{BWD} = -1, \text{FWD} = 1\}$. This simple encoding scheme enables description of arbitrary regions of the PDE, allowing for the PDE to be queried at frequencies different from the sensor frame rate. Euler integration enables simple implementation of multi-step forward, backward, and cyclic consistency losses without extra bells and whistles. For efficiency, we use Euler integration with Δt set as the time between observations for our PDE solver, enabling support for arbitrary sensor frame rates. We optimize θ as formulated in Equation 3, featuring bi-directional three step Euler integration and one step cycle consistency, over 1000 epochs using Adam optimizer with a learning rate $8e^{-5}$ with five frame minibatches and NSFP’s early stopping schedule. Formally, for P_t we optimize

$$\arg \min_{\theta} \sum \begin{matrix} \forall k \in \{1, 2, 3\} : \text{TruncatedChamfer}(\text{Euler}_{\theta}(P_t, \text{FWD}, k), P_{t+k}) \\ \forall k \in \{1, 2, 3\} : \text{TruncatedChamfer}(\text{Euler}_{\theta}(P_t, \text{BWD}, k), P_{t-k}) \\ \|\text{Euler}_{\theta}(\text{Euler}_{\theta}(P_t, \text{FWD}, 1), \text{BWD}, 1) - P_t\|_2 \times 0.01 \end{matrix} . \quad (3)$$

Like with NSFP, TruncatedChamfer is defined as the standard L_2 Chamfer distance with per-point distances below 2 meters. Additionally, we introduce $\text{Euler}_{\theta}(P_t, d, k)$, which is k step Euler integration from t to $t \pm k$ depending on the direction d , with the given point cloud P_t as the initial conditions. As we show in Section 5, EulerFlow’s simple PDE estimation formulation produces high quality flow and solving this PDE over arbitrary time spans unlocks emergent point tracking behavior.

5 EXPERIMENTS

In order to validate EulerFlow’s construction and better understand the impact of its design choices, we perform extensive experiments on the Argoverse 2 (Wilson et al., 2021) and Waymo Open (Sun et al., 2020) autonomous vehicle datasets. We compare against open source implementations of FastNSF (Li et al., 2023), Liu et al., NSFP (Li et al., 2021b), FastFlow3D (Jund et al., 2021), and variants of ZeroFlow (Vedder et al., 2024) provided by the ZeroFlow model zoo⁴, a third-party implementation of NTP (Wang et al., 2022a) from Vidanapathirana et al., and Argoverse 2 2024 Scene Flow Challenge leaderboard submission results from the authors of Flow4D (Kim et al., 2024), TrackFlow (Khatri et al., 2024), DeFlow++/DeFlow (Zhang et al., 2024a), ICP Flow (Lin & Caesar, 2024), and SeFlow (Zhang et al., 2024b). As discussed in Khatri et al. and used in the Argoverse 2 2024 Scene Flow Challenge, methods are ranked by their speed normalized *mean Dynamic Normalized EPE*.

Implementation Details. To showcase the flexibility of EulerFlow without hyperparameter tuning, for all quantitative experiments we run with a neural prior of depth 8 (NSFP’s default depth), except for our submission to the Argoverse 2 2024 Scene Flow Challenge (Section 5.1) where, based on our depth ablation study on the val split (Section 5.2.3), we set the depth of the neural prior to 18. As discussed in NTP’s original paper (Wang et al., 2022a) and confirmed by our experiments, NTP struggles to converge beyond 25 frames, so we only compare against it in a 20 frame settings. As is typical in the scene flow literature (Chodosh et al., 2023), we perform ego compensation and ground point removal on both Argoverse 2 and Waymo Open using the dataset provided map and ego pose.

5.1 HOW DOES EULERFLOW COMPARE TO PRIOR ART ON REAL DATA?

EulerFlow achieves **state-of-the-art** performance on the *Argoverse 2 2024 Scene Flow Challenge* leaderboard. Despite being unsupervised, EulerFlow **surpasses all prior art, supervised or unsupervised**, including Flow4D (Kim et al., 2024)⁵ and ICP Flow (Lin & Caesar, 2024)⁶. Notably, EulerFlow achieves $< 2.5 \times$ lower error mean Dynamic EPE than ICP Flow and beats Flow4D by over 10%.

³Hyperparameters (e.g. filter width of 128) of NSFP’s prior are kept fixed, except for depth (Section 5.2.3).

⁴<https://github.com/kylevedder/SceneFlowZoo>, from Vedder et al. (2024).

⁵Flow4D is the winner of the 2024 Argoverse 2 Scene Flow Challenge supervised track.

⁶ICP Flow is the winner of the 2024 Argoverse 2 Scene Flow Challenge unsupervised track.

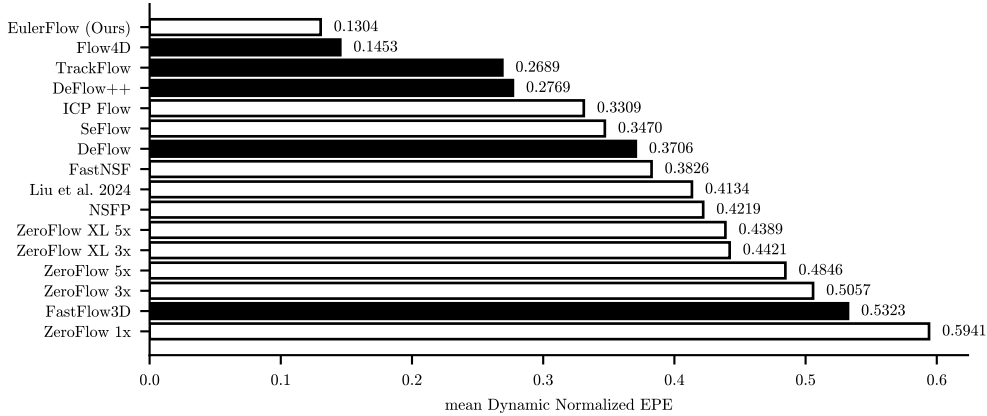


Figure 4: Mean Dynamic Normalized EPE of EulerFlow compared to prior art on the Argoverse 2 2024 Scene Flow Challenge test set. EulerFlow is state-of-the-art, beating all supervised (shown in black) and unsupervised (shown in white) methods. Lower is better.




Figure 5: Mean Dynamic Normalized EPE of EulerFlow compared to prior art on the Waymo Open validation set. EulerFlow is state-of-the-art, beating all supervised (shown in black) and unsupervised (shown in white) methods. Lower is better.

EulerFlow’s dominant performance also holds on Waymo Open (Sun et al., 2020); we compare against several popular methods (Figure 5), and EulerFlow again out-performs the baselines by a wide margin, more than halving the error over the next best method.

5.2 WHAT CONTRIBUTES TO EULERFLOW’S STATE-OF-THE-ART PERFORMANCE?

We find that EulerFlow’s lower mean Dynamic EPE can be attributed to better performance on smaller objects. On Argoverse 2, compared to Flow4D, EulerFlow’s improves on WHEELED VRU (Figure 6d), a small, rare, fast moving class. Compared to ICP Flow, EulerFlow’s improves on all classes (at least halving the error on every class!), with the largest improvements coming from the smaller and harder to detect objects PEDESTRAIN and WHEELED VRU (Figures 6c–6d). On Waymo Open, the same story holds; the most dramatic performance improvements come from the small object classes of CYCLIST and PEDESTRIAN (Figure 7).

These results are consistent with our qualitative visualizations. Figure 12 shows EulerFlow is able to cleanly extract the motion of a bird flying past the ego vehicle. Euler integration using EulerFlow’s PDE, starting at the bird’s takeoff position and ending when it loses lidar returns, produces emergent 3D point tracking behavior on the bird through its trajectory (Figure 8), further demonstrating the quality of EulerFlow’s model of the scene’s motion.

5.2.1 HOW DOES OBSERVATION SEQUENCE LENGTH IMPACT EULERFLOW?

As we discuss in Section 3, EulerFlow benefits from constructive interference from PDE estimation over many observations. Does this sufficiently explain EulerFlow’s performance? Figure 9 shows the performance of EulerFlow at length 5, 20, 50, and full sequence (roughly 160 frames) compared to NSFP and NTP at length 20. EulerFlow sees clear continual improvements as the number of frames increases without signs of saturation. However, sequence length alone does not explain EulerFlow’s performance; even at the same sequence length of 20, EulerFlow demonstrates significantly better performance than NTP.

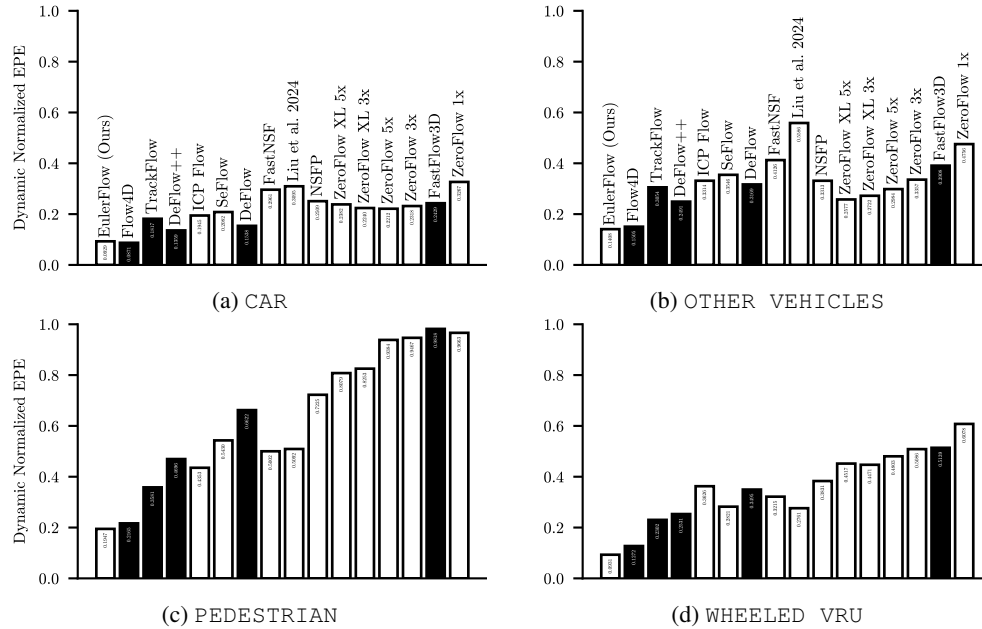


Figure 6: Per class Dynamic Normalized EPE of EulerFlow compared to prior art on the Argoverse 2 2024 Scene Flow Challenge test set. Supervised methods shown in black, unsupervised methods shown in white. Methods are ordered left to right by increasing mean Dynamic Normalized EPE. Lower is better.

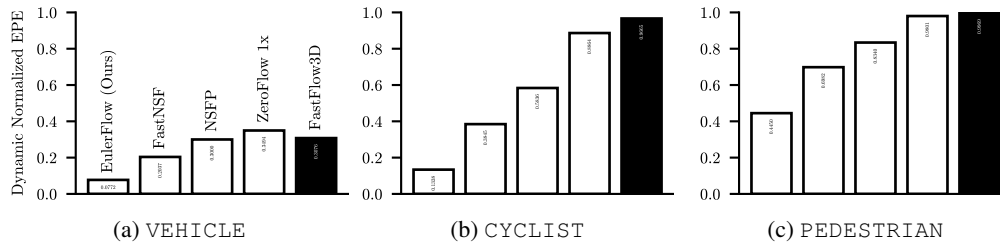


Figure 7: Per class Dynamic Normalized EPE of EulerFlow compared to prior art on the Waymo Open validation set. Supervised methods shown in black, unsupervised methods shown in white. Methods are ordered left to right by increasing mean Dynamic Normalized EPE. Lower is better.

5.2.2 HOW DO MULTI-FRAME OPTIMIZATION OBJECTIVES IMPACT EULERFLOW?

Equation 3 outlines two major components of EulerFlow’s loss: multi-frame Euler integration for Chamfer Distance reconstruction, and cycle consistency. Figure 10 compares EulerFlow without more than one integration step (No $k > 1$) and without cycle consistency rollouts (No Cycle) to better




Figure 8: EulerFlow is able to track the bird over 20 frames by performing Euler integration starting from takeoff until it loses all point cloud lidar returns.




Figure 9: Mean Dynamic Normalized EPE of EulerFlow for various sequence lengths on the Argoverse 2 val split, compared against representative baselines. These results demonstrate that EulerFlow improves with sequence length; however, at a sequence length of 20, our method significantly outperforms NTP, suggesting that EulerFlow’s performance cannot solely be attributed to longer sequence modeling.

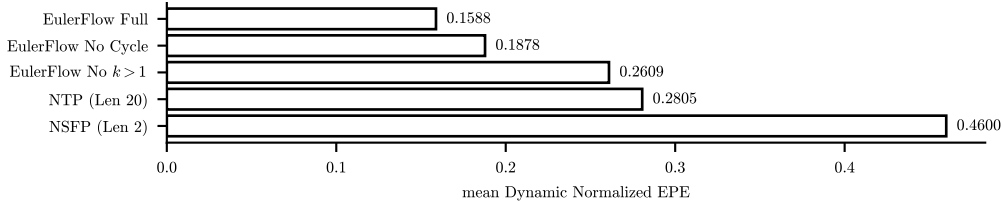


Figure 10: Mean Dynamic Normalized EPE of EulerFlow for various losses on the Argoverse 2 val split, compared against representative baselines. These results demonstrate that EulerFlow’s multi-observation optimization objectives significantly improve overall performance.

understand the impact of these components. Ablating multi-step Euler integrated rollouts results in significant degradation, as they are a strong forcing function to have consistent, smooth flow volumes; indeed, despite consuming the entire sequence, EulerFlow (No $k > 1$) is only slightly better than NTP with a sequence length of 20. These results highlight the power of multi-step rollouts and their potential as a objective for other test-time optimization methods and feedforward methods.

5.2.3 HOW DOES THE CAPACITY OF THE NEURAL PRIOR IMPACT EULERFLOW?

Li et al. ablate the capacity of NSFP’s neural prior to characterize underfitting and overfitting to optimization objective noise, ultimately selecting a depth of 8. EulerFlow’s neural prior is structured similarly; however, NSFP is fitting a single snapshot in time, while EulerFlow is fitting an entire PDE over significant time intervals. Intuitively, one would expect that full sequence modeling would benefit from greater network capacity.

To evaluate this, we perform a sweep of EulerFlow’s network depth on the Argoverse 2 validation split (Figure 11). While EulerFlow with NSFP’s default of depth 8 performs well on our Argoverse 2 evaluations (0.1% worse than the supervised state-of-the-art Flow4D), we see that performance improves as the neural prior’s depth increases until depth 18 (indicating underfitting), where we then start to see degradation (indicating overfitting to surrogate objective noise). Based on these results, as we discuss in Section 5.1, our submission to the Argoverse 2 2024 Scene Flow Challenge leaderboard uses a depth 18 neural prior (Figure 4).

5.3 BEYOND AUTONOMOUS VEHICLES

Due to a dearth of real-world, labeled scene flow data, prior scene flow work on real data overwhelmingly evaluates on autonomous vehicle datasets (Dewan et al., 2016; Li et al., 2021b; Jund et al., 2021; Li et al., 2023; Chodosh et al., 2023; Liu et al., 2024; Vedder et al., 2024; Khatri et al., 2024); consequently, motion understanding in other important domains like tabletop manipulation has been neglected. To showcase EulerFlow’s out-of-the-box flexibility and generalizability, we visualize EulerFlow on several dynamic tabletop scenes we collected using the ORBBEC Astra, a low cost depth camera commonly used in robotics (Figure 13). For viewing ease, we paint our point clouds with color; however, RGB information is not provided to EulerFlow during optimization. Furthermore, while EulerFlow only reasons about point clouds, we show that it can leverage video

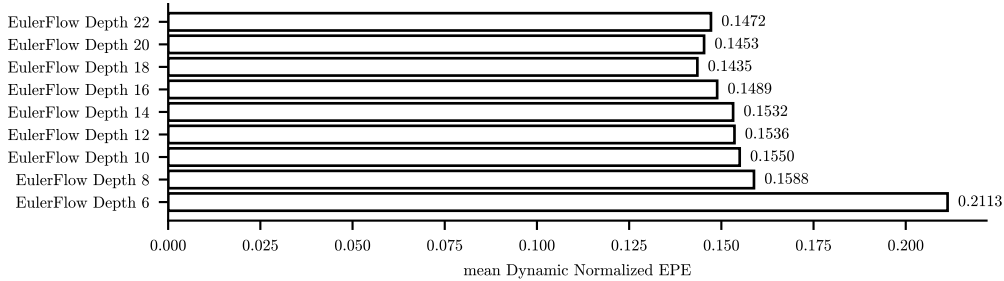


Figure 11: Mean Dynamic Normalized EPE of EulerFlow on the Argoverse 2 val split for different neural prior capacities. Shallow networks underfit the PDE, while deeper networks overfit to noise in the optimization objectives.

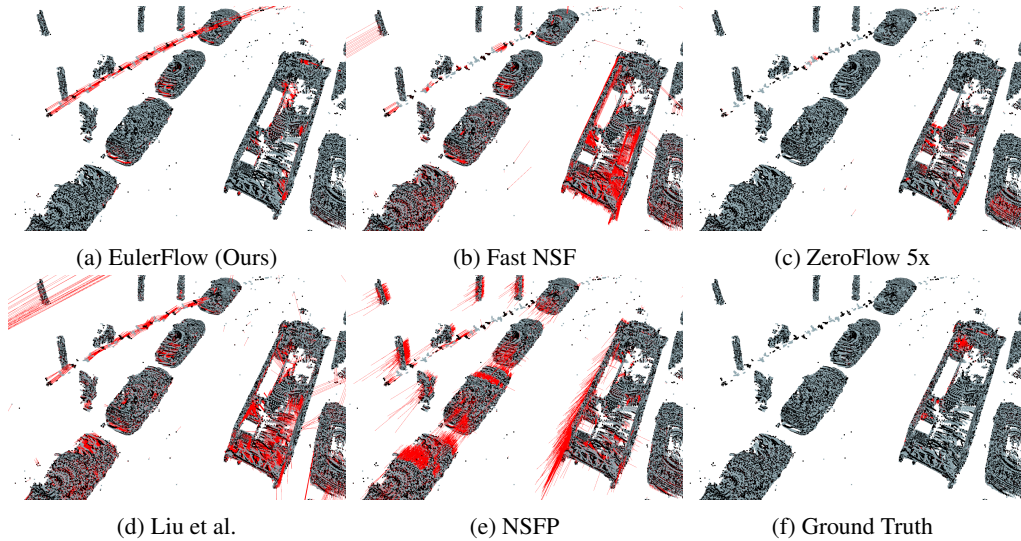


Figure 12: Visualization of EulerFlow compared to prior art for the same scene as Figure 1a and Figure 8a. EulerFlow is able to extract the bird’s trajectory; however, all other methods except Liu et al. fail to recognize this motion, and Liu et al.’s flow is marred by severe scene artifacts. The bird is outside the labeled object taxonomy, and so its motion is unlabeled in the ground truth (Figure 12f).

monocular depth estimators to describe RGB-only scene flow in Supp. A.2. Interactive visuals are available at vedder.io/eulerflow.

6 CONCLUSION

By reframing scene flow as fitting a PDE over positions for a full sequence of observations, we are able to construct EulerFlow, a simple unsupervised scene flow method that achieves state-of-the-art performance on the Argoverse 2 2024 Scene Flow Challenge and Waymo Scene Flow benchmark, where it beats all prior art, supervised or unsupervised. EulerFlow is able to describe motion on small, fast moving, out of distribution objects unable to be captured by prior art, suggesting that it makes good on the promises of scene flow as a powerful primitive for understanding the dynamic world. It also exhibits other emergent capabilities, like basic 3D point tracking behavior.

We believe that this PDE formulation has implications for scene flow at large, including beyond test-time optimization methods; the power of multi-step Euler integration may translate to feedforward network training. Future work should explore feedforward models that perform autoregressive rollouts or directly learn to estimate multiple steps into the future.

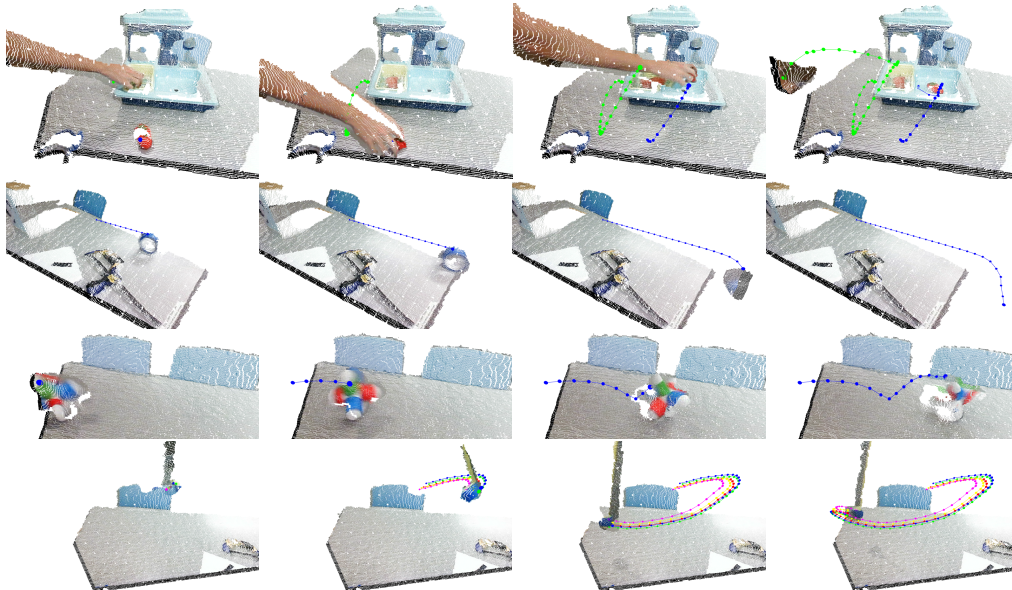


Figure 13: Visualizations of EulerFlow’s emergent 3D point tracking behavior that demonstrate the quality of its PDE estimate. Row 1 depicts tracking a tomato placed in the sink by a human hand; note the point does not move despite the hand grasping the tomato. Row 2 depicts tracking of painters tape rolling off a table; EulerFlow is able to estimate its trajectory even after it disappears out of frame. Row 3 depicts tracking of the motion of a jack commonly used in tabletop manipulation experiments (Venkatesh et al., 2023). Row 4 depicts tracking of a tennis ball taped to a flexible rod. All tracks are produced by Euler integration through the estimated PDE from the initial conditions shown in the left column. Note that point clouds are shown in color for visualization purposes only; RGB is not used during optimization.

6.1 LIMITATIONS AND FUTURE WORK

EulerFlow’s strong performance opens the book on an exciting new line of work; however, we feel that it’s important to be candid about EulerFlow’s current limitations in order to make future progress.

EulerFlow is point cloud only. Point cloud sparsity bottlenecks performance; for instance, in Figure 8 and Figure 12 we were only able to track the bird for 20 frames because we lost lidar observations of the bird, while it remained visible in the car’s RGB cameras. Future works should explore multi-modal fusion for better long-term motion descriptions.

EulerFlow is expensive to optimize. With our implementation, optimizing EulerFlow for a single Argoverse 2 sequence takes 24 hours on one NVIDIA V100 16GB GPU, putting it on par with the original NeRF paper’s computation expense (Mildenhall et al., 2021). However, like with NeRF, we believe further optimization, algorithmic, and engineering innovation can significantly improve runtime.

EulerFlow does not understand ray casting geometry. During ego-motion, a static foreground occluding object casts a moving shadow on the background; this causes Chamfer Distance to estimate this as a leading edge of moving structure, encouraging false motion artifacts (Li et al., 2021b). This can be addressed with optimization losses that model point clouds as originating from a time of flight sensor with limited visibility, as has been successfully demonstrated in the reconstruction (Chodosh et al., 2024) and forecasting literature (Khurana et al., 2023; Agro et al., 2024), rather than an unstructured set of points to be associated via local point distance.

REFERENCES

Ben Agro, Quin Sykora, Sergio Casas, Thomas Gilles, and Raquel Urtasun. UnO: Unsupervised Occupancy Fields for Perception and Forecasting. In *CVPR*, 2024.

- Ramy Batraway, René Schuster, Mohammad-Ali Nikouei Mahani, and Didier Stricker. RMS-FlowNet: Efficient and Robust Multi-Scale Scene Flow Estimation for Large-Scale Point Clouds. In *Int. Conf. Rob. Aut.*, pp. 883–889. IEEE, 2022.
- Aseem Behl, Despoina Paschalidou, Simon Donné, and Andreas Geiger. Pointflownet: Learning representations for rigid motion estimation from point clouds. In *Int. Conf. Comput. Vis.*, pp. 7962–7971, 2019.
- D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.) (ed.), *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pp. 611–625. Springer-Verlag, October 2012.
- Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NeurIPS’18, pp. 6572–6583, Red Hook, NY, USA, 2018.
- Shin-Fang Chng, Sameera Ramasinghe, Jamie Sherrah, and Simon Lucey. Gaussian Activated Neural Radiance Fields for High Fidelity Reconstruction and Pose Estimation. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pp. 264–280, Berlin, Heidelberg, 2022. Springer-Verlag. ISBN 978-3-031-19826-7.
- Nathaniel Chodosh, Deva Ramanan, and Simon Lucey. Re-Evaluating LiDAR Scene Flow for Autonomous Driving. *arXiv preprint*, 2023.
- Nathaniel Chodosh, Anish Madan, Deva Ramanan, and Simon Lucey. Simultaneous Map and Object Reconstruction, 2024. URL <https://arxiv.org/abs/2406.13896>.
- Ayush Dewan, Tim Caselitz, Gian Diego Tipaldi, and Wolfram Burgard. Rigid scene flow for 3d lidar scans. In *Int. Conf. Intel. Rob. Sys.*, pp. 1765–1770. IEEE, 2016.
- Emeç Erçelik, Ekim Yurtsever, Mingyu Liu, Zhijie Yang, Hanzhen Zhang, Pınar Topçam, Maximilian Listl, Yılmaz Kaan Çaylı, and Alois Knoll. 3D Object Detection with a Self-supervised Lidar Scene Flow Backbone. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision – ECCV 2022*, pp. 247–265, Cham, 2022. Springer Nature Switzerland.
- 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.*, pp. 3254–3263, 2019.
- Ramin Hasani, Mathias Lechner, Alexander Amini, Daniela Rus, and Radu Grosu. Liquid Time-constant Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35:7657–7666, May 2021.
- Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. DepthCrafter: Generating Consistent Long Depth Sequences for Open-world Videos. *arXiv preprint arXiv:2409.02095*, 2024.
- 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*, 12 2021.
- Ishan Khatri, Kyle Vedder, Neehar Peri, Deva Ramanan, and James Hays. I Can’t Believe It’s Not Scene Flow! In *European Conference on Computer Vision (ECCV)*, 2024.
- Tarasha Khurana, Peiyun Hu, David Held, and Deva Ramanan. Point Cloud Forecasting as a Proxy for 4D Occupancy Forecasting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- Jaeyeul Kim, Jungwan Woo, Ukcheol Shin, Jean Oh, and Sunghoon Im. Flow4D: Leveraging 4D Voxel Network for LiDAR Scene Flow Estimation, 2024. URL <https://arxiv.org/abs/2407.07995>.
- Yair Kittenplon, Yonina C Eldar, and Dan Raviv. Flowstep3d: Model unrolling for self-supervised scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 4114–4123, 2021.

- Ruibao Li, Guosheng Lin, Tong He, Fayao Liu, and Chunhua Shen. HCRF-Flow: Scene flow from point clouds with continuous high-order CRFs and position-aware flow embedding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 364–373, 2021a.
- Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural Scene Flow Prior. *Advances in Neural Information Processing Systems*, 34, 2021b.
- 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)*, pp. 9878–9890, October 2023.
- Yancong Lin and Holger Caesar. ICP-Flow: LiDAR Scene Flow Estimation with ICP. 2024.
- Dongrui Liu, Daqi Liu, Xueqian Li, Sihao Lin, Hongwei xie, Bing Wang, Xiaojun Chang, and Lei Chu. Self-supervised multi-frame neural scene flow, 2024. URL <https://arxiv.org/abs/2403.16116>.
- Xingyu Liu, Charles R Qi, and Leonidas J Guibas. FlowNet3D: Learning Scene Flow in 3D Point Clouds. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM*, 65(1):99–106, dec 2021. ISSN 0001-0782.
- Mahyar Najibi, Jingwei Ji, Yin Zhou, Charles R. Qi, Xinchen Yan, Scott Ettinger, and Dragomir Anguelov. Motion Inspired Unsupervised Perception and Prediction in Autonomous Driving. *European Conference on Computer Vision (ECCV)*, 2022.
- Chrisopher Nalty, Neehar Peri, Joshua Gleason, Carlos Castillo, Shuowen Hu, Thirimachos Bourlai, and Rama Chellappa. A Brief Survey on Person Recognition at a Distance. 12 2022. doi: 10.48550/arXiv.2212.08969.
- Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 343–352, 2015. doi: 10.1109/CVPR.2015.7298631.
- Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable Neural Radiance Fields. *ICCV*, 2021.
- Neehar Peri, Achal Dave, Deva Ramanan, and Shu Kong. Towards Long Tailed 3D Detection. *CoRL*, 2022a.
- Neehar Peri, Jonathon Luiten, Mengtian Li, Aljosa Osep, Laura Leal-Taixe, and Deva Ramanan. Forecasting from LiDAR via Future Object Detection. *arXiv:2203.16297*, 2022b.
- Neehar Peri, Mengtian Li, Benjamin Wilson, Yu-Xiong Wang, James Hays, and Deva Ramanan. An empirical analysis of range for 3d object detection. *arXiv preprint arXiv:2308.04054*, 2023.
- Gilles Puy, Alexandre Boulch, and Renaud Marlet. Flot: Scene flow on point clouds guided by optimal transport. In *Eur. Conf. Comput. Vis.*, pp. 527–544. Springer, 2020.
- Sameera Ramasinghe, Hemanth Saratchandran, Violetta Shevchenko, Alexander Long, and Simon Lucey. On the Optimality of Activations in Implicit Neural Representations, 2024. URL <https://openreview.net/forum?id=0Lqyut1y7M>.

- Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- 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.*, pp. 150–159. IEEE, 2020.
- Kyle Vedder and Eric Eaton. Sparse PointPillars: Maintaining and Exploiting Input Sparsity to Improve Runtime on Embedded Systems. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- Kyle Vedder, Neehar Peri, Nathaniel Chodosh, Ishan Khatri, Eric Eaton, Dinesh Jayaraman, Yang Liu, Deva Ramanan, and James Hays. ZeroFlow: Scalable Scene Flow via Distillation. In *Twelfth International Conference on Learning Representations (ICLR)*, 2024.
- Sharanya Venkatesh, Bibit Bianchini, Alp Aydinoglu, and Michael Posa. Sampling-Based Model Predictive Control for Contact-Rich Manipulation. In *IROS 2023 Workshop on Leveraging Models for Contact-Rich Manipulation*, 2023.
- Kavisha Vidanapathirana, Shin-Fang Chng, Xueqian Li, and Simon Lucey. Multi-body neural scene flow. In *2024 International Conference on 3D Vision (3DV)*. IEEE, 2024.
- Chaoyang Wang, Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural Prior for Trajectory Estimation. In *CVPR*, pp. 6522–6532, 2022a. doi: 10.1109/CVPR52688.2022.00642.
- Jun Wang, Xiaolong Li, Alan Sullivan, Lynn Abbott, and Siheng Chen. PointMotionNet: Point-Wise Motion Learning for Large-Scale LiDAR Point Clouds Sequences. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 4418–4427, 2022b.
- Thomas Weng, Sujay Man Bajracharya, Yufei Wang, Khush Agrawal, and David Held. Fabricflownet: Bimanual cloth manipulation with a flow-based policy. In *Conference on Robot Learning*, pp. 192–202. PMLR, 2022.
- 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.
- 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.*, pp. 88–107. Springer, 2020.
- Guangyao Zhai, Xin Kong, Jinhao Cui, Yong Liu, and Zhen Yang. FlowMOT: 3D Multi-Object Tracking by Scene Flow Association. *ArXiv*, abs/2012.07541, 2020.
- Qingwen Zhang, Yi Yang, Heng Fang, Ruoyu Geng, and Patric Jensfelt. DeFlow: Decoder of Scene Flow Network in Autonomous Driving. *ICRA*, 2024a.
- Qingwen Zhang, Yi Yang, Peizheng Li, Olov Andersson, and Patric Jensfelt. Seflow: A self-supervised scene flow method in autonomous driving. *arXiv preprint arXiv:2407.01702*, 2024b.
- Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced Grouping and Sampling for Point Cloud 3D Object Detection. *arXiv preprint arXiv:1908.09492*, 2019.

A ADDITIONAL RESULTS

A.1 HOW DOES THE CHOICE OF LEARNABLE FUNCTION CLASS AND DESIGN OF ENCODINGS IMPACT EULERFLOW?

EulerFlow at its core is an optimization loop over a simple, feedforward ReLU-based multi-layer perception inherited from Neural Scene Flow Prior (Li et al., 2021b). How does this choice of learnable function class impact the performance of EulerFlow? To better understand these design choices we examine the choice of non-linearity and time feature encoding.

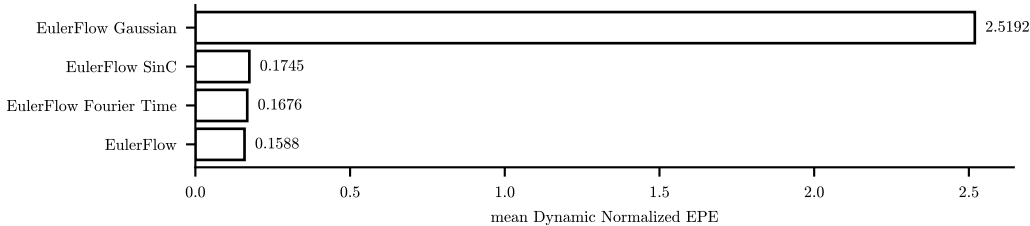


Figure 14: Mean Dynamic Normalized EPE of EulerFlow on the Argoverse 2 val split for less-smooth configurations of its learnable function class. These results indicate that the smoothness of the ReLU non-linearity proposed by Li et al. transfers well to EulerFlow.

One of Li et al.’s core theoretical contributions demonstrates that NSFP’s ReLU MLP is a good prior for scene flow because it represents a smooth learnable function class, and scene flow is often locally smooth with respect to input position. However, unlike NSFP, EulerFlow is fitting flow over a full PDE; while it seems reasonable to assume that this PDE is typically also locally smooth, cases like adjacent cars moving rapidly in opposite directions may benefit from the ability to model higher frequency, less locally smooth functions. To test this hypothesis, we ablate EulerFlow by replacing its normalized time with higher frequency sinusoidal time embeddings (mirroring Wang et al.’s proposed time embedding for NTP), as well as try other popular non-linearities like SinC (Ramasinghe et al., 2024) and Gaussian (Chng et al., 2022) from the coordinate network literature. Figure 14 features negative results on these ablations across the board; Gaussians were unable to converge due the extremely high frequency representation triggering early stopping, while the use of SinC and higher frequency time embeddings both resulted in worse overall performance, indicating that Li et al.’s smooth function prior does indeed seem appropriate for EulerFlow’s neural prior.

A.2 EULERFLOW WITH MONOCULAR DEPTH ESTIMATES

While EulerFlow only consumes point clouds, we can leverage RGB-based video monocular depth estimators to fit scene flow. In Figure 15, we use DepthCrafter (Hu et al., 2024) to generate a point cloud from the raw RGB of the tabletop video from Figure 13, Row 4.




Figure 15: Visualizations of EulerFlow’s emergent 3D point tracking behavior on monocular depth estimates from DepthCrafter (Hu et al., 2024). Interactive visualizations available at vedder.io/eulerflow.

A.3 HOW DOES EULERFLOW FAIL?

As we discuss in Section 6.1, EulerFlow does not understand projective geometry — its optimization losses use Chamfer Distance which directly associates points, sometimes resulting in moving shadows




Figure 16: Visualizations of one of the failure modes of EulerFlow where flow is predicted on the edges of the moving "shadow" in the point cloud. Interactive visualizations available at vedder.io/eulerflow.

on background objects. To demonstrate this, we select a particularly egregious example in Figure 16, featuring a frame from the jack being thrown across the table. Due to the moving shadow cast by the jack onto the table, EulerFlow incorrectly assigns flow to the table surface nearby the jack, particularly on the leading edge, even though the table surface is stationary.

B FAQ

B.1 WHAT DATASETS DID YOU PRETRAIN ON?

EulerFlow is not pretrained on any datasets. It is a test-time optimization method (akin to NeRFs), and as we show with our tabletop data, this means it runs out-of-the-box on arbitrary point cloud data.

B.2 WHY DIDN'T YOU USE A NEURAL ODE OR A LIQUID NEURAL NETWORK?

Neural ODEs (Chen et al., 2018) take variable size and number of steps in latent space to do inference; imagine a ResNet that can use an ODE solver to dynamically scale the impact of the residual block, as well as decide the number of residual blocks. They are not a function class specially designed to fit a PDE well. Similar to Neural ODEs, Liquid Neural Networks (Hasani et al., 2021) focus on the same class of problems and are similarly not applicable.

B.3 WHY DIDN'T YOU DO EXPERIMENTS ON FLYINGTHINGS3D / <SIMULATED DATASET>?

Most popular synthetic datasets do not contain long observation sequences (Mayer et al., 2016; Butler et al., 2012), but instead include standalone frame pairs. Our method leverages the long sequence of observations to refine our neural estimate of the true PDE. Indeed, on two frames, EulerFlow collapses to NSFP.

More importantly, these datasets are also not representative of real world environments. To quote Chodosh et al.: "[FlyingThings3D has] 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." Khatri et al. also make this point by highlighting the importance of meaningfully breaking down the object distribution during evaluation identify performance on rare safety-critical categories. FlyingThings3D does not have meaningful semantics; it's not obvious what things even matter or how to appropriately break down the scene.

Instead, we want to turn our attention to the sort of workloads that *do* clearly matter — describing motion in domains like manipulation or autonomous vehicles, where it seems clear that scene flow, if solved, will serve as powerful primitive for downstream systems. This is why we performed qualitative experiments on the tabletop data we collected ourselves; to our knowledge, no real-world dynamic datasets of this nature exist with ground truth annotations, but we want to emphasize that EulerFlow works in such domains, and consequently EulerFlow and other Scene Flow via PDE-based methods can be used as a primitive in these real world domains.