

DELTAv2: Accelerating Dense 3D Tracking

Tuan Duc Ngo*
UMass Amherst

Ashkan Mirzaei
Snap Inc.

Guocheng Qian
Snap Inc.

Hanwen Liang
University of Toronto

Chuang Gan
UMass Amherst

Evangelos Kalogerakis
UMass Amherst & TU Crete

Peter Wonka
Snap Inc. & KAUST

Chaoyang Wang
Snap Inc.

<https://snap-research.github.io/DELTAv2/>

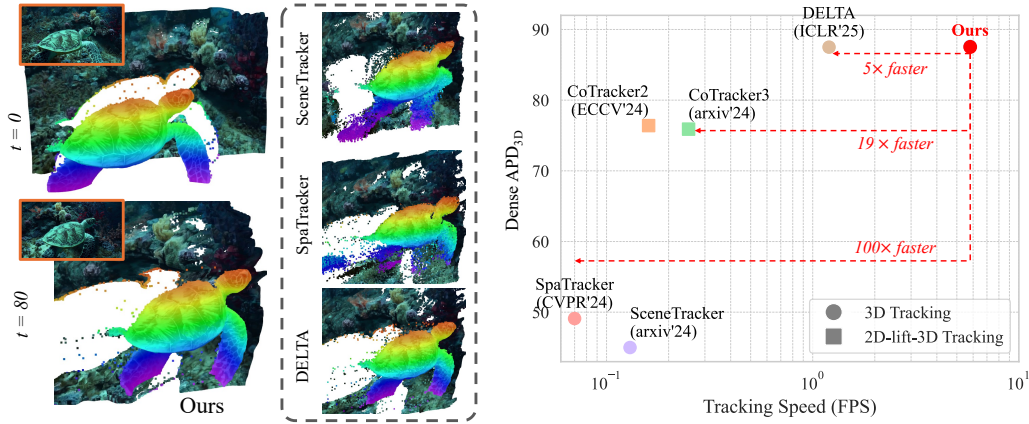


Figure 1: **DELTAv2** achieves state-of-the-art for dense 3D tracking — matching the accuracy of DELTA while being $5\times$ faster, and outperforming all prior methods in speed–accuracy tradeoff. (Left) Long-range 3D trajectories on real-world videos. (Right) Performance vs. FPS comparison.

Abstract

We propose a novel algorithm for accelerating dense long-term 3D point tracking in videos. Through analysis of existing state-of-the-art methods, we identify two major computational bottlenecks. First, transformer-based iterative tracking becomes expensive when handling a large number of trajectories. To address this, we introduce a coarse-to-fine strategy that begins tracking with a small subset of points and progressively expands the set of tracked trajectories. The newly added trajectories are initialized using a learnable interpolation module, which is trained end-to-end alongside the tracking network. Second, we propose an optimization that significantly reduces the cost of correlation feature computation, another key bottleneck in prior methods. Together, these improvements lead to a $5\text{--}100\times$ speedup over existing approaches while maintaining state-of-the-art tracking accuracy.

1 Introduction

We focus on *dense 3D tracking*: given an RGB video, we predict the 3D trajectory of every pixel in the first frame across all subsequent frames in the local camera coordinate system. This task is more challenging than related problems such as optical flow [41, 20, 5, 27, 51, 71, 57, 58, 24, 70] and scene flow [17, 21, 50, 59, 39, 67, 45], which estimate motion only between adjacent frames. In

*Work done during internship at Snap Inc.




Figure 2: Runtime breakdown of DELTA [44] with a single iteration and one sliding window.

Table 1: Comparison of cost reduction strategies on Kubric3D *val* set [44]. All methods use 1 iteration unless noted; full-resolution outputs are obtained via bilinear interpolation.

Strategy	APD _{3D} ↑	Runtime (ms) ↓
DELTA (4 iterations)	87.3	8404
DELTA (1 iteration)	74.3	2275
Downsample video reso. (4×)	35.8	383
Subsample video frames (4×)	65.4	833
Subsample trajectories (16×)	71.2	585

contrast, dense 3D tracking requires maintaining accurate long-term correspondences under larger motions and appearance changes.

To achieve long-term point tracking, recent methods leverage powerful transformer-based networks to estimate trajectories from videos [18, 12, 37, 36, 32, 31]. Some works [69, 62] further extend tracking from 2D to 3D by incorporating accurate monocular depth estimation [4, 48, 22]. However, these approaches remain limited to sparse trajectories, short temporal windows, or both.

DELTA [44] addressed these limitations by introducing the first framework for dense, long-range 3D tracking. While DELTA achieves strong performance and is relatively fast compared to prior work, it remains computationally expensive: tracking every pixel in a 100-frame video takes around 2 minutes, too slow for real-time or latency-sensitive applications.

In this paper, we present a faster and more scalable successor to DELTA. We begin by identifying two key computational bottlenecks in the DELTA pipeline.

First, despite the efficient transformer design, its iterative refinement must process a large number of trajectory tokens repeatedly for every tracking iteration, leading to high computational cost. We observe that in a dense motion field, many nearby pixels exhibit similar motion patterns. As such, tracking all of them through every iteration is often redundant. Motivated by this, we analyze three possible coarse-to-fine strategies and propose an algorithm that reduces computation by subsampling a sparse set of trajectories in early iterations. The tracking density is then progressively increased in later iterations, culminating in fully dense tracking in the final iteration to ensure accuracy is preserved. To recover the motion of untracked pixels during intermediate iterations, we introduce a learnable interpolation module. This module dynamically predicts blending weights to infer the motion of untracked pixels from their nearby tracked neighbors, leveraging both spatial and feature similarity.

Second, through a runtime breakdown for a single iteration of DELTA (see Fig. 2), we observe that 4D correlation feature computation becomes increasingly expensive as the number of tracking points grows, contributing to significant runtime overhead. We identify an inefficiency in the implementation used by prior work that leads to suboptimal GPU utilization. To address this, we introduce an effective optimization that significantly reduces runtime without sacrificing performance.

Together, these optimizations achieve a $5\times$ speed-up over DELTA on dense 3D tracking of 100-frame videos, while maintaining state-of-the-art accuracy (Fig. 1).

In summary, our main contributions are: 1) a coarse-to-fine tracking algorithm that starts with sparsely sampled trajectories and densifies them over iterations; 2) a learnable interpolator that produces dense trajectories in each iteration and supports adaptive resampling; 3) an accelerated 4D correlation computation, bringing dense 3D tracking closer to real-time.

2 Related Work

Optical Flow estimates motion by establishing dense correspondences between consecutive frames. Classical variational methods [41, 20, 5] faced limitations in handling complex dynamics such as fast motion, occlusions, and large displacements. The advent of CNN-based techniques [27, 51, 71, 57] brought substantial improvements to short-term motion estimation. RAFT [58] introduced a key

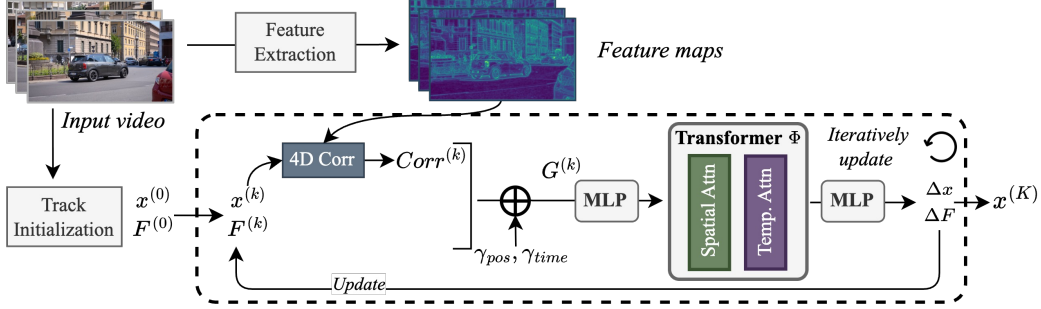


Figure 3: A modern long-range tracking pipeline (we omit the 3D tracking parts here for simplicity)

innovation by computing dense correlation volumes across all pixel pairs. Later approaches expanded on this architecture by introducing transformer-based tokenization of correlation volumes [24], employing global feature aggregation to better address occlusions [28], and modeling optical flow as a softmax-based matching task [70]. Attempts to extend optical flow to long-term sequences using multi-frame methods [58, 14, 53] or point-tracking integration [34, 6] often face challenges with drift and occlusion, limiting long-term tracking performance.

Scene Flow extends optical flow into 3D by estimating dense motion in 3D space. Some methods rely on RGB-D inputs [17, 21, 50, 59, 72], while others operate on monocular video alone [25, 26, 3, 29]. A separate line of work estimates 3D motion directly from point clouds [39, 67, 16, 45]. Several approaches [19, 40, 23] leverage temporal context to predict scene flow from videos. More recently, methods such as [56, 38] build on top of foundation 3D reconstruction models [66, 74, 35] to estimate 3D scene flow in a more general and scalable manner.

Point Tracking estimates long-range motion trajectories in videos. Particle Video [52] introduced particle trajectories for long-range video motion. TAP-Vid [10] provided a comprehensive benchmark to evaluate point tracking and TAPNet, a baseline that predicts tracking locations using global cost volume and soft-argmax operation. PIPs [18] revisited the concept of particle video and proposed a network that updates trajectories iteratively over fixed temporal windows, but ignored spatial context with independent point tracking and struggle with occlusion. TAPIR [12] combines the global estimation of TAPNet [10] and the iterative refinement of [18] with variable-length window. Subsequent efforts further improve tracking performance by jointly process multiple points with temporal attention and spatial attention [32]; using 4D correlation feature [8]. CoTracker3 [31] proposes a lightweight architecture and a semi-supervised training strategy to further improve the 2D tracking performance. SceneTracker [62] and SpaTracker [69] extend point tracking to 3D by incorporating depth information, but remain inefficient for dense tracking due to computationally expensive cross-track attention. DELTA [44] is the first approach tackling the dense 3D tracking problem with an efficient transformer and upsample layer to obtain high-resolution dense tracking. **Concurrently**, several other efforts also explore 3D tracking. Seurat [7] estimates depth changes on top of a 2D tracker [32, 8] to recover 3D motion. TAPI3D [73] proposes a 3D-space correlation mechanism for directly tracking points in 3D space. St4RTrack [13] leverages recent 3D reconstruction models [66, 74] and extends them for dense 3D tracking via joint optimization.

Tracking by Reconstructing tackles long-range motion estimation by explicitly reconstructing a deformable scene representation. OmniMotion [63] jointly optimizes a NeRF [42] with a bijective deformation field [9] and extracts 2D point trajectories through this mapping. More recent efforts leverage DINOv2 [46] to compute long-range correspondences, either through enhanced invertible deformation fields [55] or self-supervised techniques [60]. Shape-of-Motion [65] uses 3D Gaussian Splatting to jointly learn geometry and motion, enabling point tracking by tracing Gaussian positions across frames. While these reconstruction-based methods are capable of producing dense trajectories, they rely on computationally intensive per-video optimization. As a result, they remain significantly slower and less accurate than feedforward, data-driven tracking models on standard benchmarks.

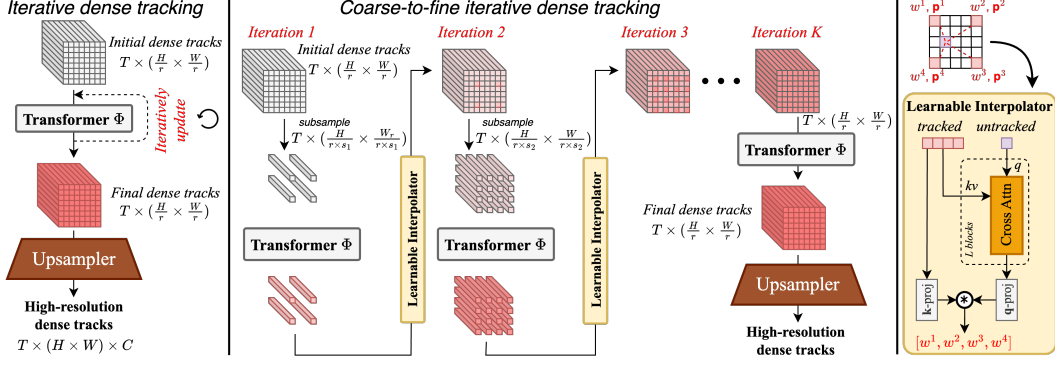


Figure 4: **Overview of our proposed framework.** (Left) Traditional iterative dense tracking refines all trajectories at every iteration, leading to high computational cost. (Middle) Our **coarse-to-fine iterative dense tracking** reduces computation by subsampling trajectory points in early iterations and progressively increasing the density across iterations. (Right) A **learnable interpolation module** leverages attention to infer untracked motions from nearby tracked pixels, enabling efficient and adaptive trajectory propagation.

3 Method

Problem setup: The input to our method is an RGB-D video, where the RGB frames are represented as $V \in \mathbb{R}^{T \times H \times W \times 3}$, with T , H , and W denoting the temporal and spatial dimensions. The corresponding depth maps $D \in \mathbb{R}^{T \times H \times W}$ are generated using a monocular depth estimation model. Our model outputs dense, occlusion-aware 3D trajectories $P \in \mathbb{R}^{T \times H \times W \times 4}$, where each vector $\mathbf{p}_{(t,u,v)} = (u_t, v_t, d_t, o_t)$ captures the motion of a pixel located at (u, v) in the first frame as it propagates to frame t . Here, (u_t, v_t) denote the projected 2D pixel coordinates in frame t , d_t is the predicted depth, and $o_t \in \{0, 1\}$ indicates the visibility status.

3.1 Preliminary: Modern Tracking Frameworks

Our work builds upon a modern tracking framework. We select DELTA [44], a recent approach for dense 3D tracking from RGB-D videos, but other frameworks have a similar structure (See Fig. 3 for an overview). First, a feature extraction network F (similar to [58]) creates image feature map for each frame. Then, the core of the algorithm iteratively evolves a data structure of size of dense tracks $\mathcal{D} \in \mathbb{R}^{T \times \frac{H}{r} \times \frac{W}{r}}$, denoted as $\{P_i\}$, where i is the trajectory index, and $P_i = [\mathbf{p}_1^i, \mathbf{p}_2^i, \dots, \mathbf{p}_T^i]$, with $\mathbf{p}_t^i = (u_t^i, v_t^i, d_t^i, o_t^i)$ being the 3D location and visibility of the point associated with the i -th trajectory at the t -th frame. The downsampling factor r is beneficial for faster processing.

At each iteration, a trajectory is represented by a list of tokens $G^i = [G_1^i, G_2^i, \dots, G_T^i]$, each token G_t^i encodes position, visibility, appearance and correlation of the trajectory at t -th frame:

$$G_t^i = [\mathcal{F}_t^i, \text{Corr}_t^i, \text{DCorr}_t^i, o_t^i, \gamma(\mathbf{x}_t^i - \mathbf{x}_1^i)] + \gamma_{\text{pos}}(\mathbf{x}_t^i) + \gamma_{\text{time}}(t), \quad (1)$$

where \mathcal{F}_t^i are image features extracted by F , Corr_t^i are correlation features [58, 18, 32, 8, 44], DCorr_t^i is depth correlation, and γ_{pos} and γ_{time} are the positional embedding of the input position $\mathbf{x}_t^i = (u_t^i, v_t^i, d_t^i)$ and time t , respectively. For more details about the individual components, refer to [44].

The tokens G_t^i are processed by a transformer network Φ . From the output tokens of Φ we can compute new dense tracks using a pointwise MLP and use the computed point tracks and the output tokens to compute the input tokens for the next iteration. The framework uses K iterations, e.g., $K = 4$. The network Φ itself employs a joint global-local spatial attention mechanism: global attention over a sparse set of anchor tracks captures long-range motion, while local attention focuses on fine-grained motion within small spatial neighborhoods.

Finally, a transformer-based upsampling module is used to upsample the dense tracks to the full resolution $T \times H \times W \times 4$. This upsampler models each pixel’s trajectory as a soft combination of neighboring trajectories (see [44] for details of this architecture).

While DELTA supports end-to-end dense tracking and has already been optimized for runtime, we undertake the challenging task of finding further significant improvements. We will discuss this next.

3.2 Analysis of Strategies for Accelerating Tracking

To accelerate the tracking framework, we first profile the time usage of each stage in the tracking algorithm, as shown in Fig. 2. This analysis reveals two primary bottlenecks: (1) the point tracking transformer Φ , and (2) the computation of the correlation features $Corr$. In this section, we focus on analyzing of the transformer Φ .

The computational cost of Φ scales linearly with the number of trajectories [32], and quadratically with the number of frames. Therefore, reducing either the number of trajectories or the number of frames can lead to substantial speedups. We consider three factors that influence these quantities: the spatial resolution of the input RGB images, the temporal resolution of the video, and the number of trajectory points tracked. To understand the impact of these factors on both performance and efficiency, we design a series of tests to evaluate one iteration of tracking. For each setting, we evaluate two aspects: (1) the time required for computation, and (2) the degradation in tracking accuracy, measured by the APD_{3D} metric on a synthetic dataset. The specific tests are described below.

1) Spatial downsampling of input RGB images: We reduce the resolution of the input frames by downsampling both the width and height by a factor of 4. Point tracking is then performed on the low-resolution frames using the pretrained transformer Φ . We then upsample the resulting trajectories back to the original resolution for evaluation. This analysis reduces both the number of trajectories and the spatial cost of attention, providing insights into how resolution affects accuracy and speed.

2) Temporal subsampling of input frames: We subsample the input video along the temporal dimension by a factor of 4, reducing the number of frames used in trajectory computation. The trajectories are still computed at full spatial resolution. Afterward, we apply linear interpolation to upsample the trajectories back to match the original temporal length of the video. Although more advanced interpolation methods such as B-splines or non-linear motion kernels could be used to model more complex dynamics, linear interpolation is sufficient for analyzing the effects of temporal subsampling on performance.

3) Subsampling the number of trajectories: We keep the input frames at full spatial and temporal resolution, but subsample the initial set of trajectory points by a factor of 16. This reduces the number of tokens passed through the transformer while preserving the original feature resolution.

4) Baseline: no downsampling is used as reference. We denote this as DELTA (1 iteration) as reference and DELTA (4 iterations) for the complete algorithm.

We present the results of this analysis in Tab. 1 (see the *supplementary* for more details). Among the evaluated strategies, *subsampling the number of trajectories* emerges as the most effective, as other approaches incur significantly higher tracking errors with clear drawbacks. *Spatial downsampling* fails due to information loss caused by reduced feature map resolution. *Temporal subsampling* assumes linear motion over time, but DELTA tracks in UVD space, where 3D linear motion does not correspond to linear 2D and depth trajectories. As a result, interpolating skipped frames introduces distortion and lowers accuracy. Based on this analysis, we propose a new algorithm that leverages trajectory subsampling as a core component (see Sec. 3.3.1). To complement this, we introduce a novel learnable interpolation architecture to recover dense trajectories, forming a key component of our method (see Sec. 3.3.2). Finally, we present an accelerated design for computing correlation features, further improving overall efficiency (see Sec. 3.3.3).

In addition to the three major contributions discussed above, we explored several alternative design choices to further accelerate the tracking framework. These include reducing the number of transformer layers (e.g., from 6 to 3, see Tab. 6), lowering the number of iterative refinements (see Fig. 5), and experimenting with various trajectory subsampling strategies (see Tab. 5).

3.3 Accelerating Dense Point Tracking

3.3.1 Coarse to Fine Tracking with Trajectory Subsampling

As shown in Tab. 1, tracking on a subsampled set of points followed by interpolation significantly reduces runtime while maintaining accuracy for the interpolated points. This observation naturally leads to a coarse-to-fine tracking algorithm. We begin by sampling a sparse set of points and perform one iteration of point tracking. Next, we interpolate the positions of the remaining untracked points. These interpolated positions provide reasonable estimates that serve as initialization for a denser set of points in the next iteration. The process is repeated, using the interpolated estimates to guide the position updates for increasingly dense sets of points, progressively refining the tracking results.

More concretely, the algorithm proceeds as follows.

1) Subsample points. We begin by subsampling a small set of points to track from a $\frac{H}{r} \times \frac{W}{r}$ pixel grid, where r is the upsampling ratio of the upsampler used in DELTA. In our implementation (illustrated in Fig. 4), we uniformly sample points on a sparse grid of size $\frac{H}{r \times s_1} \times \frac{W}{r \times s_1}$, where s_1 denotes the initial spatial subsampling scale for the first iteration.

Each sampled point is initialized with its original pixel position. After each coarse-to-fine tracking iteration, we update the positions of the points that have already been tracked and introduce additional points for tracking in the next iteration. At the k -th iteration, the number of points tracked corresponds to a grid of size $\frac{H}{r \times s_k} \times \frac{W}{r \times s_k}$, where s_k is the subsampling scale at step k . In our implementation, we decrease s_k by a factor of 2 at each step, progressively refining the tracking resolution.

In our ablation study, we compare different subsampling schedules and explore alternative strategies, such as random sampling and importance sampling focused on distinctive features (e.g., object boundaries). However, we find that the simple uniform grid sampling performs on par with these more complex strategies. Therefore, we adopt the grid-based approach for its simplicity.

2) Single-step point tracking. We compute the correlation features for each sampled point based on its current estimated or initialized position, and feed these features into the transformer Φ to produce updated point positions.

3) Interpolate positions for new points. Given the current estimated positions of the sampled points, we use an interpolation method to estimate the initial positions of the additional points introduced at the next scale, which have not yet been tracked. We experimented with various interpolation strategies, including bilinear and nearest-neighbor interpolation, and found that nearest-neighbor consistently outperforms bilinear (see Fig. 5). We hypothesize that bilinear interpolation tends to oversmooth positional estimates, especially in regions with large motion differences, which is a problem that becomes more pronounced in long-term tracking scenarios. To further improve the accuracy of interpolated positions, we introduce a learnable interpolation module (see Sec. 3.3.2), which is trained end-to-end alongside the coarse-to-fine tracking pipeline.

The algorithm then returns to step 1) and repeats the process until the full resolution is reached.

3.3.2 Learnable Interpolation Module

To propagate motion from tracked to untracked pixels during coarse-to-fine iterations, we introduce a *learnable interpolation module* that estimates the 3D motion of each untracked pixel as an adaptive blend of nearby tracked motions. Unlike fixed interpolation methods (e.g., bilinear or nearest-neighbor), our approach dynamically predicts interpolation weights using attention over spatial features. Let $\mathcal{P}_{\text{track}}$ be the set of tracked pixels and $\mathcal{P}_{\text{query}}$ the untracked ones. For each query $(u, v) \in \mathcal{P}_{\text{query}}$, we predefine its four nearest neighbors $\{(u_j, v_j)\}_{j=1}^4 \in \mathcal{P}_{\text{track}}$. The interpolated 3D motion $\mathbf{p}_{(u,v)}$ (for clarity, we omit the frame subscript t) is computed as:

$$\mathbf{p}_{(u,v)} = \sum_{j=1}^4 w_{(u,v)}^j \cdot \mathbf{p}_{(u_j, v_j)} \quad (2)$$

where weights $w_{(u,v)}^j \in [0, 1]$ sum to 1 and are predicted by a lightweight attention module.

Let $\mathcal{F} \in \mathbb{R}^{\frac{H}{r} \times \frac{W}{r} \times C_{\mathcal{F}}}$ denote the dense feature map of the query frame of the dense tracking from the feature extractor F . The initial query feature $\mathcal{F}_{(u,v)}$ is first refined via L multi-head cross-attention blocks over the support set. We follow the Alibi attention scheme [49] in these blocks, where a

Table 2: Long-range optical flow results on CVO [68, 34].

Methods	CVO-Clean (7 frames)		CVO-Final (7 frames)		CVO-Extend (48 frames)	
	EPE↓ (all/vis/occ)	IoU↑	EPE↓ (all/vis/occ)	IoU↑	EPE↓ (all/vis/occ)	IoU↑
RAFT	2.48 / 1.40 / 7.42	57.6	2.63 / 1.57 / 7.50	56.7	21.80 / 15.4 / 33.4	65.0
MFT	2.91 / 1.39 / 9.93	19.4	3.16 / 1.56 / 10.3	19.5	21.40 / 9.20 / 41.8	37.6
TAPIR	3.80 / 1.49 / 14.7	73.5	4.19 / 1.86 / 15.3	72.4	19.8 / 4.74 / 42.5	68.4
CoTracker2	1.51 / 0.88 / 4.57	75.5	1.52 / 0.93 / 4.38	75.3	5.20 / 3.84 / 7.70	70.4
DOT	1.29 / 0.72 / 4.03	80.4	1.34 / 0.80 / 3.99	80.4	4.98 / 3.59 / 7.17	71.1
SceneTracker	4.40 / 3.44 / 9.47	-	4.61 / 3.70 / 9.62	-	11.5 / 8.49 / 17.0	-
SpatialTracker	1.84 / 1.32 / 4.72	68.5	1.88 / 1.37 / 4.68	68.1	5.53 / 4.18 / 8.68	66.6
DOT-3D	1.33 / 0.75 / 4.16	79.0	1.38 / 0.83 / 4.10	78.8	5.20 / 3.58 / 7.95	70.9
DELTA	0.94 / 0.51 / 2.97	78.7	1.03 / 0.61 / 3.03	78.3	3.67 / 2.64 / 5.30	70.1
Ours	1.04 / 0.61 / 3.25	77.6	1.12 / 0.69 / 3.27	77.3	3.53 / 2.57 / 5.10	70.6

relative positional bias is added to the attention logits, computed as the L1 distance between the query and support pixel locations. The refined query and the final weights are predicted as:

$$\tilde{\mathcal{F}}_{(u,v)} = \text{CrossAttn}^{(L)}(\mathcal{F}_{(u,v)}, \{\mathcal{F}_{(u_j, v_j)}\}), \mathbf{w}_{(u,v)} = \text{softmax}\left(\mathbf{q}(\tilde{\mathcal{F}}_{(u,v)}) \cdot \mathbf{k}(\{\mathcal{F}_{(u_j, v_j)}\})\right) \quad (3)$$

where $\mathbf{q}(\cdot)$ and $\mathbf{k}(\cdot)$ are linear projections for the refined query and support features. This mechanism allows sparse trajectories to efficiently guide the evolution of the full dense motion field over time, maintaining temporal/spatial consistency without the cost of processing full dense trajectories.

3.3.3 Accelerating Computation of Correlation Features

A major bottleneck in DELTA is the computation of 4D correlation features, introduced in [8], which measure pairwise similarities between local neighborhoods around query and predicted positions. These features form a 4D tensor processed by a dual-convolutional module. As shown in our profiling (Fig. 2), this becomes increasingly expensive with a large number of trajectories.

The main inefficiency arises from the input channel size (typically $7 \times 7 = 49$), which is not divisible by 8, leading to poor GPU utilization in dense settings. To address this, we add a lightweight MLP projection that reduces the dimension to 32, followed by LayerNorm and ReLU. This adds minimal overhead for sparse tracking and yields a significant speedup in the dense setting. An alternative from [31] replaces the dual-conv module with pure MLPs on flattened 4D correlation, but we find it performs worse on dense 2D/3D tracking (See Tab. 6).

4 Experiments

Implementation Details: Following prior work [44, 32, 69], we use the Kubric engine [15] to generate 5,631 RGB-D training videos with both sparse and dense tracking annotations. During training, we apply online data augmentation, including color jittering, random scaling, and random cropping. As in [44], we supervise the model on both sparse and dense tracks. The total loss combines the 2D coordinate loss \mathcal{L}_{2D} , inverse depth loss \mathcal{L}_{depth} , and visibility loss \mathcal{L}_{visib} . The loss weights $\lambda_{2d}, \lambda_{depth}, \lambda_{visib}$ follow those in [44]. We initialize the model from the pretrained DELTA checkpoint and add parameters for the interpolation module. The model is trained for 100,000 iterations on 8 A100 GPUs using AdamW with a one-cycle learning rate schedule, starting from an initial learning rate of 2×10^{-4} .

4.1 Comparison with prior work

Baselines. We compare our method with previous optical flow and point tracking approaches. In particular, we closely compare it with DELTA [44], the current SOTA for dense 3D tracking. We also include DOT [34], a method for dense 2D tracking, its 3D variant DOT-3D [44], and other recent point tracking methods, including CoTracker2 [32], CoTracker3 [31], SpaTracker [69], and SceneTracker [62].

Benchmark Datasets. We evaluate our method on a diverse set of tracking benchmarks covering 2D optical flow, dense 3D pixel tracking, and sparse 3D point tracking. (1) Long-range 2D optical

Table 3: **Dense 3D tracking results** on the Kubric3D dataset.

Methods	Kubric-3D (24 frames)		Runtime↓	
	AJ↑ (<i>all/vis</i>)	APD _{3D} ↑ (<i>all/vis</i>)	OA↑	(second)
CoTracker2	70.0 / 80.7	76.4 / 85.1	96.7	145
CoTracker3	68.9 / 79.3	75.9 / 84.3	95.8	94
SpatialTracker	35.3 / 42.7	49.1 / 51.6	96.5	350
SceneTracker	-	45.0 / 65.5	-	179
DOT-3D	68.1 / 72.3	75.3 / 77.5	88.7	11.8
DELTA	81.7 / 85.1	87.5 / 90.9	97.5	18.2
Ours	81.7 / 84.8	87.6 / 91.0	97.0	3.5

flow: We use the **CVO** dataset [68, 34], which includes three splits: *CVO-Clean*, *CVO-Final*, and *CVO-Extended*. Each split contains about 500 videos with 7 frames (48 frames for *CVO-Extended*), captured at 60 FPS, with dense long-range optical flow annotations and occlusion masks. (2) Dense 3D pixel tracking: Following [44], we evaluate on a held-out test split of the **Kubric** dataset [15], which includes 143 RGB-D videos, each 24 frames long, with dense ground-truth 3D trajectories for every pixel. (3) 3D point tracking: We benchmark on the large-scale **TAP-Vid3D** dataset [33], which contains 4,569 videos from DriveTrack [2], PStudio [30], and Aria [47], covering a range of real-world and simulated environments, with video lengths from 25 to 300 frames.

Metrics. For the *long-range optical flow*, we adopt the evaluation protocol from [34, 44], reporting the end-point error (EPE) between predicted and ground-truth flows and the occlusion prediction accuracy using the intersection-over-union (IoU) between the predicted and ground-truth visibility masks. For the *dense 3D tracking* and *sparse 3D point tracking* benchmarks, we follow the metrics proposed in [33]. Specifically, we report APD_{3D} (Average Percent of Points within a threshold to measure spatial accuracy, OA (Occlusion Accuracy) for evaluating visibility prediction, and AJ (Average Jaccard), which jointly captures both spatial and occlusion correctness.

Long-range 2D optical flow. Our method achieves comparable performance to DELTA [44], a state-of-the-art approach for dense tracking (see Tab. 2). While we observe slightly lower accuracy on the *Clean* and *Final* subsets, our model outperforms DELTA on the more challenging *Extended* subset, which contains significantly longer video sequences. This demonstrates our model’s robustness in long-range tracking scenarios.

Dense 3D Tracking. Table 3 reports results on the Kubric3D test set. Our method matches the performance of DELTA while significantly outperforming other baselines, including 3D tracking methods such as SpaTracker[69] and SceneTracker[62], as well as 2D-to-3D lifted approaches like CoTracker2[32] and CoTracker3[31]. Importantly, our approach is substantially more efficient: it achieves up to a $5\times$ speed-up over DELTA and is approximately $100\times$ faster than SPATRACKER, enabling practical deployment in real-time or large-scale applications. Qualitative results of dense 3D tracking on in-the-wild videos are provided in the *supplementary materials*.

3D Tracking. We evaluate our approach on sparse 3D point tracking in Table 4. Although the sparse setting does not directly benefit from our coarse-to-fine strategy, the results demonstrate that training with this strategy—and incorporating our redesigned 4D correlation module—does not degrade sparse tracking performance. Our method achieves slightly better results on the *Aria* and *DriveTrack* subsets, with a minor drop on *PStudio*, showing overall competitive and robust performance.

4.2 Ablation Study

We conduct a series of ablation studies on the *CVO Extended* split [34]. Runtime is measured as the average time required to densely track all 384×512 pixels across a 48-frame video using a machine with a single A100 GPU. The number of iterations is set to 4 by default, unless otherwise specified.

Analysis of the Coarse-to-fine strategy. Figure 5a compares our coarse-to-fine strategy (schedule: 8,4,2,1) with the baseline that tracks all trajectories in every iteration. Our method achieves similar accuracy while being $2.75\times$ faster. In Fig.5b, we ablate different interpolators, showing that our attention-based design outperforms fixed alternatives like nearest-neighbor and bilinear. Fig. 5c compares different subsampling schedules, where (s_1, s_2, s_3, s_4) denotes the subsampling factors across 4 iterations. A finer schedule, such as (2,2,2,1), improves accuracy but with increased runtime.

Table 4: **3D tracking results** on the TAP-Vid3D Benchmark. We use UniDepth for depth estimation.[†] denotes using depth to lift 2D tracks to 3D tracks.

Methods	Aria			DriveTrack			PStudio			Average		
	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow
TAPIR [†] + COLMAP	7.1	11.9	72.6	8.9	14.7	80.4	6.1	10.7	75.2	7.4	12.4	76.1
CoTracker2 [†] + COLMAP	8.0	12.3	78.6	11.7	19.1	81.7	8.1	13.5	77.2	9.3	15.0	79.1
BootsTAPIR [†] + COLMAP	9.1	14.5	78.6	11.8	18.6	83.8	6.9	11.6	81.8	9.3	14.9	81.4
CoTracker2 [†] + UniDepth	13.0	20.9	84.9	12.5	19.9	80.1	6.2	13.5	67.8	10.6	18.1	77.6
TAPTR [†] + UniDepth	15.7	24.2	87.8	12.4	19.1	84.8	7.3	13.5	84.3	11.8	18.9	85.6
LocoTrack [†] + UniDepth	15.1	24.0	83.5	13.0	19.8	82.8	7.2	13.1	80.1	11.8	19.0	82.3
SpatialTracker + UniDepth	13.6	20.9	90.5	8.3	14.5	82.8	8.0	15.0	75.8	10.0	16.8	83.0
SceneTracker + UniDepth	-	23.1	-	-	6.8	-	-	12.7	-	-	14.2	-
DOT-3D + UniDepth	13.8	22.1	85.5	11.8	17.9	82.3	3.2	5.3	52.5	9.6	15.1	73.4
DELTA + UniDepth	16.6	24.4	86.8	14.6	22.5	85.8	8.2	15.0	76.4	13.1	20.6	83.0
Ours + UniDepth	17.0	24.7	87.2	15.6	23.8	84.6	7.7	14.4	74.1	13.4	21.0	81.2

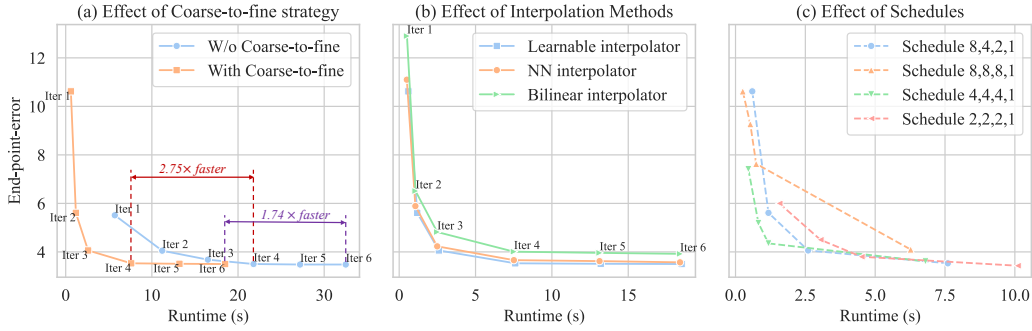


Figure 5: **Analysis of the coarse-to-fine strategy.** We visualize how accuracy evolves with runtime across different methods. (a) We evaluate runtime per iteration and observe that the coarse-to-fine strategy consistently reduces runtime compared to the baseline while achieving similar accuracy. (b) Nearest neighbor interpolation outperforms bilinear, and our learnable interpolation further improves accuracy. (c) We compare different coarse-to-fine scheduling strategies.

Table 5 presents different pixel sampling strategies. In addition to the default uniform grid sampling, we evaluate two alternatives: random sampling and SIFT-based keypoint selection. Both alternatives require an additional step to identify neighbors for the interpolation module. While the random strategy yields slightly better results, the SIFT-based approach performs worse, likely due to non-uniform coverage. We retain the uniform grid sampling as our default due to its simplicity, efficiency, and competitive performance.

Runtime Analysis. Table 6 summarizes our runtime improvements. Replacing the original 4D correlation in DELTA [44] reduces runtime by 33%. Adding our coarse-to-fine strategy yields a further $3\times$ speed-up. A 3-layer transformer brings runtime to 5.8s, though with some loss in accuracy.

5 Conclusion

In this work, we presented an efficient framework for dense 3D video tracking that significantly improves the runtime of DELTA [44] while preserving its strong performance. We proposed a coarse-to-fine tracking algorithm that progressively increases spatial coverage across iterations, combined with a learnable interpolation module for dense supervision and a faster 4D correlation computation. These improvements yield $5 - 100\times$ speedup over previous approaches, making our approach more suitable for real-time applications. Nonetheless, our method inherits common **limitations** of data-driven tracking: it is trained on synthetic data and may struggle under fast motion, severe occlusions, or poor depth estimation, which affect 2D/3D tracking accuracy in complex real-world scenes.

Table 5: Ablation of coarse-to-fine sampling strategies.

Sampling	EPE↓	Runtime (s)
None	3.50 / 2.55 / 5.09	21.8
Uniform grid (default)	3.53 / 2.57 / 5.10	7.6
Random	3.51 / 2.61 / 5.04	7.8
SIFT	3.79 / 2.88 / 5.72	7.8

Table 6: Impact of runtime optimization components.

Model Variant	EPE↓	Runtime (s)
DELTA [44]	3.67 / 2.64 / 5.30	32.1
+ new 4D Corr	3.50 / 2.55 / 5.09	21.8
+ Coarse-to-fine	3.53 / 2.57 / 5.10	7.6
+ 3-layer Trans	3.91 / 2.85 / 5.82	5.8
+ MLP 4D Corr	4.76 / 2.81 / 7.75	5.4

References

- [1] Nasir Ahmed, T. Natarajan, and Kamisetty R Rao. Discrete cosine transform. *IEEE transactions on Computers*, 100(1):90–93, 2006.
- [2] Arjun Balasingam, Joseph Chandler, Chenning Li, Zhoutong Zhang, and Hari Balakrishnan. Drivetrack: A benchmark for long-range point tracking in real-world videos. In *CVPR*, 2024.
- [3] Bayram Bayramli, Junhwa Hur, and Hongtao Lu. Raft-msf: Self-supervised monocular scene flow using recurrent optimizer. *International Journal of Computer Vision*, 131(11):2757–2769, 2023.
- [4] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth, 2023.
- [5] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, 2004.
- [6] Seokju Cho, Jiahui Huang, Seungryong Kim, and Joon-Young Lee. Flowtrack: Revisiting optical flow for long-range dense tracking. In *CVPR*, 2024.
- [7] Seokju Cho, Jiahui Huang, Seungryong Kim, and Joon-Young Lee. Seurat: From moving points to depth. *arXiv preprint arXiv:2504.14687*, 2025.
- [8] Seokju Cho, Jiahui Huang, Jisu Nam, Honggyu An, Seungryong Kim, and Joon-Young Lee. Local all-pair correspondence for point tracking. *arXiv preprint arXiv:2407.15420*, 2024.
- [9] Laurent Dinh, Jascha Narain Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *ArXiv*, abs/1605.08803, 2016.
- [10] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. TAP-vid: A benchmark for tracking any point in a video. *NeurIPS*, 2022.
- [11] Carl Doersch, Pauline Luc, Yi Yang, Dilara Gokay, Skanda Koppula, Ankush Gupta, Joseph Heyward, Ignacio Rocco, Ross Goroshin, João Carreira, and Andrew Zisserman. BootsTAP: Bootstrapped training for tracking-any-point. *Asian Conference on Computer Vision*, 2024.
- [12] Carl Doersch, Yi Yang, Mel Vecerik, Dilara Gokay, Ankush Gupta, Yusuf Aytar, Joao Carreira, and Andrew Zisserman. TAPIR: Tracking any point with per-frame initialization and temporal refinement. In *ICCV*, 2023.
- [13] Haiwen Feng, Junyi Zhang, Qianqian Wang, Yufei Ye, Pengcheng Yu, Michael J Black, Trevor Darrell, and Angjoo Kanazawa. St4rtrack: Simultaneous 4d reconstruction and tracking in the world. *arXiv preprint arXiv:2504.13152*, 2025.
- [14] Pierre Godet, Alexandre Boulch, Aurélien Plyer, and Guy Le Besnerais. Starflow: A spatiotemporal recurrent cell for lightweight multi-frame optical flow estimation. In *ICPR*, 2021.
- [15] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanapragasam, Florian Golemo, Charles Herrmann, Thomas Kipf, Abhijit Kundu, Dmitry Lagun, Issam Laradji, Hsueh-Ti (Derek) Liu, Henning Meyer, Yishu Miao, Derek Nowrouzezahrai, Cengiz Oztireli, Etienne Pot, Noha Radwan, Daniel Rebain, Sara Sabour, Mehdi S. M. Sajjadi, Matan Sela, Vincent Sitzmann, Austin Stone, Deqing Sun, Suhani Vora, Ziyu Wang, Tianhao Wu, Kwang Moo Yi, Fangcheng Zhong, and Andrea Tagliasacchi. Kubric: a scalable dataset generator. In *CVPR*, 2022.
- [16] 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 *CVPR*, 2019.
- [17] Simon Hadfield and Richard Bowden. Kinecting the dots: Particle based scene flow from depth sensors. In *ICCV*, 2011.

- [18] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *European Conference on Computer Vision*, pages 59–75. Springer, 2022.
- [19] Pan He, Patrick Emami, Sanjay Ranka, and Anand Rangarajan. Learning scene dynamics from point cloud sequences. *International Journal of Computer Vision*, 130(3):669–695, 2022.
- [20] BK Horn and B Schunck. Determining optical flow (artificial intelligence laboratory). *Massachusetts Institute of Technology, Cambridge, MA*, 1980.
- [21] Michael Hornacek, Andrew Fitzgibbon, and Carsten Rother. Sphreflow: 6 dof scene flow from rgb-d pairs. In *CVPR*, 2014.
- [22] 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, 2024.
- [23] Shengyu Huang, Zan Gojcic, Jiahui Huang, Andreas Wieser, and Konrad Schindler. Dynamic 3d scene analysis by point cloud accumulation. In *European Conference on Computer Vision*, pages 674–690. Springer, 2022.
- [24] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *ECCV*, 2022.
- [25] Junhwa Hur and Stefan Roth. Self-supervised monocular scene flow estimation. In *CVPR*, 2020.
- [26] Junhwa Hur and Stefan Roth. Self-supervised multi-frame monocular scene flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2684–2694, 2021.
- [27] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017.
- [28] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. In *ICCV*, 2021.
- [29] Zijie Jiang and Masatoshi Okutomi. Emr-msf: Self-supervised recurrent monocular scene flow exploiting ego-motion rigidity. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 69–78, 2023.
- [30] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, et al. Panoptic studio: A massively multiview system for social interaction. *TPAMI*, 2017.
- [31] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024.
- [32] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker: It is better to track together. *arXiv:2307.07635*, 2023.
- [33] Skanda Koppula, Ignacio Rocco, Yi Yang, Joe Heyward, João Carreira, Andrew Zisserman, Gabriel Brostow, and Carl Doersch. Tapvid-3d: A benchmark for tracking any point in 3d, 2024.
- [34] Guillaume Le Moing, Jean Ponce, and Cordelia Schmid. Dense optical tracking: Connecting the dots. In *CVPR*, 2024.
- [35] Vincent Leroy, Yohann Cabon, and Jerome Revaud. Grounding image matching in 3d with mast3r, 2024.
- [36] Hongyang Li, Hao Zhang, Shilong Liu, Zhaoyang Zeng, Feng Li, Tianhe Ren, Bohan Li, and Lei Zhang. Taptrv2: Attention-based position update improves tracking any point. *arXiv preprint arXiv:2407.16291*, 2024.

- [37] Hongyang Li, Hao Zhang, Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, and Lei Zhang. Taptr: Tracking any point with transformers as detection. *arXiv preprint arXiv:2403.13042*, 2024.
- [38] Yiqing Liang, Abhishek Badki, Hang Su, James Tompkin, and Orazio Gallo. Zero-shot monocular scene flow estimation in the wild. *arXiv preprint arXiv:2501.10357*, 2025.
- [39] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. Flownet3d: Learning scene flow in 3d point clouds. In *CVPR*, 2019.
- [40] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. Meteornet: Deep learning on dynamic 3d point cloud sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9246–9255, 2019.
- [41] Etienne Mémín and Patrick Pérez. Dense estimation and object-based segmentation of the optical flow with robust techniques. *TIP*, 1998.
- [42] 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. In *ECCV*, 2020.
- [43] Michal Neoral, Jonávs vSerých, and Jivří Matas. Mft: Long-term tracking of every pixel. In *WACV*, 2024.
- [44] Tuan Duc Ngo, Peiye Zhuang, Chuang Gan, Evangelos Kalogerakis, Sergey Tulyakov, Hsin-Ying Lee, and Chaoyang Wang. Delta: Dense efficient long-range 3d tracking for any video. *arXiv preprint arXiv:2410.24211*, 2024.
- [45] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *ICCV*, 2019.
- [46] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023.
- [47] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Yuheng Carl Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. In *ICCV*, 2023.
- [48] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. UniDepth: Universal monocular metric depth estimation. In *CVPR*, 2024.
- [49] Ofir Press, Noah Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. In *ICLR*, 2022.
- [50] Julian Quiroga, Thomas Brox, Frédéric Devernay, and James Crowley. Dense semi-rigid scene flow estimation from rgbd images. In *ECCV*, 2014.
- [51] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, 2017.
- [52] Peter Sand and Seth Teller. Particle video: Long-range motion estimation using point trajectories. *IJCV*, 2008.
- [53] Xiaoyu Shi, Zhaoyang Huang, Weikang Bian, Dasong Li, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Videoflow: Exploiting temporal cues for multi-frame optical flow estimation. In *ICCV*, 2023.
- [54] Ivan Skorokhodov, Sharath Girish, Benran Hu, Willi Menapace, Yanyu Li, Rameen Abdal, Sergey Tulyakov, and Aliaksandr Siarohin. Improving the diffusability of autoencoders. *arXiv preprint arXiv:2502.14831*, 2025.

- [55] Yunzhou Song, Jiahui Lei, Ziyun Wang, Lingjie Liu, and Kostas Daniilidis. Track everything everywhere fast and robustly. *ArXiv*, abs/2403.17931, 2024.
- [56] Edgar Sucar, Zihang Lai, Eldar Insafutdinov, and Andrea Vedaldi. Dynamic point maps: A versatile representation for dynamic 3d reconstruction. *arXiv preprint arXiv:2503.16318*, 2025.
- [57] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018.
- [58] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020.
- [59] Zachary Teed and Jia Deng. Raft-3d: Scene flow using rigid-motion embeddings. In *CVPR*, 2021.
- [60] Narek Tumanyan, Assaf Singer, Shai Bagon, and Tali Dekel. Dino-tracker: Taming dino for self-supervised point tracking in a single video, 2024.
- [61] Gregory K Wallace. The jpeg still picture compression standard. *Communications of the ACM*, 34(4):30–44, 1991.
- [62] Bo Wang, Jian Li, Yang Yu, Li Liu, Zhenping Sun, and Dewen Hu. Scenetraacker: Long-term scene flow estimation network. *arXiv preprint arXiv:2403.19924*, 2024.
- [63] Qianqian Wang, Yen-Yu Chang, Ruojin Cai, Zhengqi Li, Bharath Hariharan, Aleksander Holynski, and Noah Snavely. Tracking everything everywhere all at once. In *ICCV*, 2023.
- [64] Qianqian Wang, Yen-Yu Chang, Ruojin Cai, Zhengqi Li, Bharath Hariharan, Aleksander Holynski, and Noah Snavely. Tracking everything everywhere all at once. *ICCV*, 2023.
- [65] Qianqian Wang, Vickie Ye, Hang Gao, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4d reconstruction from a single video. *arXiv preprint arXiv:2407.13764*, 2024.
- [66] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024.
- [67] Zirui Wang, Shuda Li, Henry Howard-Jenkins, Victor Prisacariu, and Min Chen. Flownet3d++: Geometric losses for deep scene flow estimation. In *WACV*, 2020.
- [68] Guangyang Wu, Xiaohong Liu, Kunming Luo, Xi Liu, Qingqing Zheng, Shuaicheng Liu, Xinyang Jiang, Guangtao Zhai, and Wenyi Wang. Accflow: Backward accumulation for long-range optical flow. In *ICCV*, 2023.
- [69] Yuxi Xiao, Qianqian Wang, Shangzhan Zhang, Nan Xue, Sida Peng, Yujun Shen, and Xiaowei Zhou. Spatialtracker: Tracking any 2d pixels in 3d space. In *CVPR*, 2024.
- [70] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *CVPR*, 2022.
- [71] Jia Xu, René Ranftl, and Vladlen Koltun. Accurate optical flow via direct cost volume processing. In *CVPR*, 2017.
- [72] Gengshan Yang and Deva Ramanan. Learning to segment rigid motions from two frames. In *CVPR*, 2021.
- [73] Bowei Zhang, Lei Ke, Adam W Harley, and Katerina Fragkiadaki. Tapip3d: Tracking any point in persistent 3d geometry. *arXiv preprint arXiv:2504.14717*, 2025.
- [74] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. *arXiv preprint arXiv:2410.03825*, 2024.

Table 7: **3D tracking results** on the TAP-Vid3D Benchmark. We use ZoeDepth for depth estimation.[†] denotes using depth to lift 2D tracks to 3D tracks.

Methods	Aria			DriveTrack			PStudio			Average		
	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow
TAPIR [†] + ZoeDepth	9.0	14.3	79.7	5.2	8.8	81.6	10.7	18.2	78.7	8.3	13.8	80.0
CoTracker2 [†] + ZoeDepth	10.0	15.9	87.8	5.0	9.1	82.6	11.2	19.4	80.0	8.7	14.8	83.4
BootsTAPIR [†] + ZoeDepth	9.9	16.3	86.5	5.4	9.2	85.3	11.3	19.0	82.7	8.8	14.8	84.8
TAPTR [†] + ZoeDepth	9.1	15.3	87.8	7.4	12.4	84.8	10.0	17.8	84.3	8.8	15.2	85.6
LocoTrack [†] + ZoeDepth	8.9	15.1	83.5	7.5	12.3	82.8	9.7	17.1	80.1	8.7	14.8	82.1
SpatialTracker + ZoeDepth	9.2	15.1	89.9	5.8	10.2	82.0	9.8	17.7	78.0	8.3	14.3	83.3
SceneTracker + ZoeDepth	-	15.1	-	-	5.6	-	-	16.3	-	-	12.3	-
DELTA + ZoeDepth	10.1	16.2	84.7	7.8	12.8	87.2	10.2	17.8	74.5	9.4	15.6	82.1
Ours + ZoeDepth	10.0	15.9	86.5	8.0	13.5	87.3	9.7	18.0	73.6	9.2	15.8	81.5

Table 8: **2D tracking results** on the TAP-Vid2D Benchmark.

Methods	Kinetics			DAVIS			RGB-Stacking		
	AJ \uparrow	APD _{2D} \uparrow	OA \uparrow	AJ \uparrow	APD _{2D} \uparrow	OA \uparrow	AJ \uparrow	APD _{2D} \uparrow	OA \uparrow
TAP-Net [10]	38.5	54.4	80.6	33.0	48.6	78.8	54.6	68.3	87.7
MFT [43]	39.6	60.4	72.7	47.3	66.8	77.8	-	-	-
PIPs [18]	31.7	53.7	72.9	42.2	64.8	77.7	15.7	28.4	77.1
OmniMotion [64]	-	-	-	46.4	62.7	85.3	69.5	82.5	90.3
TAPIR [12]	49.6	64.2	85.0	56.2	70.0	86.5	54.2	69.8	84.4
CoTracker2 [32]	48.7	64.3	86.5	60.6	75.4	89.3	63.1	77.0	87.8
DOT [34]	48.4	63.8	85.2	60.1	74.5	89.0	77.1	87.7	93.3
BootsTAPIR [11]	54.6	68.4	86.5	61.4	73.6	88.7	70.8	83.0	89.9
TAPTR [37]	49.0	64.4	85.2	63.0	76.1	91.1	-	-	-
TAPTRv2 [36]	49.7	64.2	85.7	63.5	75.9	91.4	-	-	-
LocoTrack [8]	52.9	66.8	85.3	63.0	75.3	87.2	69.0	83.2	89.5
SpatialTracker [69]	50.1	65.9	86.9	61.1	76.3	89.5	63.5	77.6	88.2
SceneTracker [62]	-	66.5	-	-	71.8	-	-	73.3	-
DOT-3D	48.1	63.7	85.9	61.2	75.3	88.1	76.3	86.6	92.1
DELTA	49.5	63.3	82.2	62.7	76.7	88.2	74.2	83.5	90.0
Ours	47.8	61.3	82.0	62.5	76.2	87.9	74.0	82.9	89.3

A More Results

A.1 3D Tracking

We further evaluate our method on the TAP-Vid3D benchmark using depth maps predicted by ZoeDepth [4]. Results are summarized in Tab. 7.

A.2 Dense 3D Tracking

Figure 6 shows visual comparisons of dense 3D tracking between our method and baselines. Our model produces highly stable and accurate trajectories, clearly outperforming SceneTracker [62] and SpaTracker [69], and achieving results on par with the strong DELTA [44] baseline.

B Analysis of Coarse-to-fine Dense Tracking

B.1 Further Analysis of Strategies for Accelerating Tracking

Fig. 7 illustrates the conceptual overview of each acceleration strategy described in Sec. 3.2 of the main paper. To better understand their practical impact, we further evaluate these strategies across a broader range of hyperparameters:

1) **Spatial downsampling of RGB input:** applied with spatial reduction factors of $2\times$, $4\times$, and $8\times$.

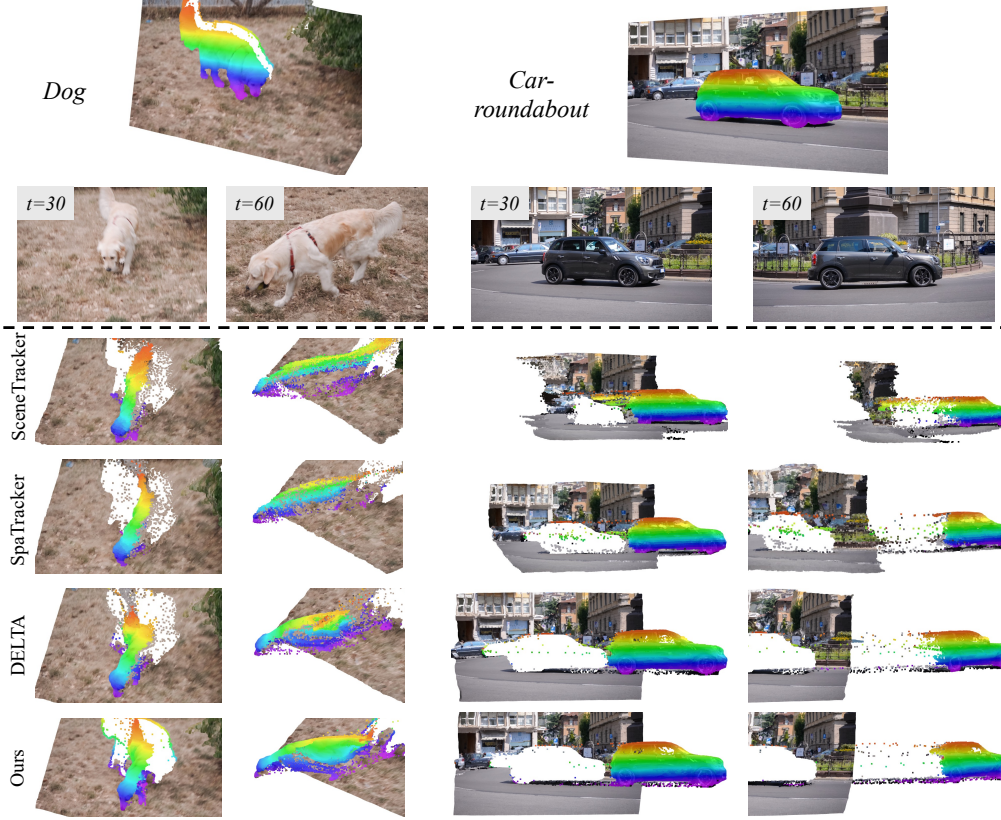


Figure 6: **Qualitative comparison of dense 3D tracking** between SceneTracker [62], SpaTracker [69], DELTA [44], and our method. We track every pixel from the first frame through 3D space, with moving objects highlighted in rainbow colors. Our approach accurately captures foreground motion, preserves stable backgrounds, and operates significantly faster with a $5\times$ to $100\times$ speedup compared to prior methods.

2) Temporal subsampling of input frames: applied with temporal reduction factors of $2\times$, $4\times$, and $8\times$.

3) Trajectory point subsampling: applied with trajectory subsampling ratios of $4\times$, $16\times$, and $64\times$.

Each variant is followed by upsampling to the original spatial and temporal resolution using either (1) bilinear interpolation or (2) nearest-neighbor interpolation. The corresponding quantitative results are summarized in Tab. 9.

B.2 Spectral Analysis of the optical flow predictions

To analyze the spatial frequency characteristics of optical flow, we follow prior work [54] on spectral analysis using the two-dimensional discrete cosine transform (2D DCT) [1]. Given a flow field $\mathcal{F}_{flow} \in \mathbb{R}^{H \times W \times 2}$, we process the horizontal and vertical components separately by dividing each into non-overlapping blocks of size $B \times B$. We then apply the type-II 2D DCT to each block, which projects the flow values onto a set of cosine basis functions at different spatial frequencies. Within each block, we arrange the coefficients into a one-dimensional sequence using the standard JPEG zigzag ordering [61], which traverses from low to high spatial frequencies. To summarize the overall spectral content of the flow, we average these zigzag-ordered coefficients across all blocks and across both flow components (horizontal and vertical). This produces a compact frequency profile that captures the distribution of motion energy across spatial scales.

Table 9: Comparison of cost reduction strategies on the Kubric3D *val* set [44]. All methods use 1 iteration unless noted; full-resolution outputs are obtained using either bilinear or nearest-neighbor interpolation.

Baseline	APD _{3D} ↑				Runtime↓	
DELTA (4 iterations)	87.3				8404	
DELTA (1 iteration)	74.3				2275	
Strategy	2×		4×		8×	
	APD _{3D} ↑	Runtime↓	APD _{3D} ↑	Runtime↓	APD _{3D} ↑	Runtime↓
(a) Bilinear interpolation						
Downsample video reso.	54.7	759	35.8	383	17.8	214
Subsample video frames	70.8	1119	65.4	833	42.3	491
Subsample trajectories	72.9	917	71.2	585	68.5	406
(b) Nearest-neighbor interpolation						
Downsample video reso.	53.9	748	34.4	382	17.6	212
Subsample video frames	68.8	1187	60.0	821	49.3	487
Subsample trajectories	73.7	902	72.1	580	67.6	405

We apply this process independently to the predicted optical flow at each iteration of the model, as well as to the ground-truth flow. This allows us to track how the spectral properties of the predicted flow evolve during refinement and how they compare to the underlying ground-truth signal.

The DCT spectrum is shown in Fig. 8, revealing a trade-off introduced by the coarse-to-fine strategy. While this approach helps accelerate inference and preserve low-frequency motion structures, it tends to suppress high-frequency components, which are essential for capturing fine-grained motion details and sharp flow boundaries. In contrast, the model without coarse-to-fine better preserves high-frequency content and more closely matches the spectral distribution of the ground-truth flow across all frequency bands. Notably, this limitation of the coarse-to-fine model is largely mitigated in the final iteration, where full-resolution trajectories are used. As shown in the left plot, the spectrum of the last iteration closely aligns with the ground truth, indicating that fine details can still be recovered at the final refinement stage.

We visualize the predicted optical flow and corresponding error maps across different transformer iterations in Fig. 9,10,11,12, comparing models with and without the coarse-to-fine strategy and using different interpolation methods. The results show that nearest-neighbor interpolation introduces noticeable grid artifacts, while bilinear interpolation tends to oversmooth motion boundaries. In contrast, our proposed learnable interpolator yields more accurate and spatially coherent predictions, particularly in early iterations where only a sparse grid of points is tracked and interpolation plays a critical role in producing dense motion estimates.

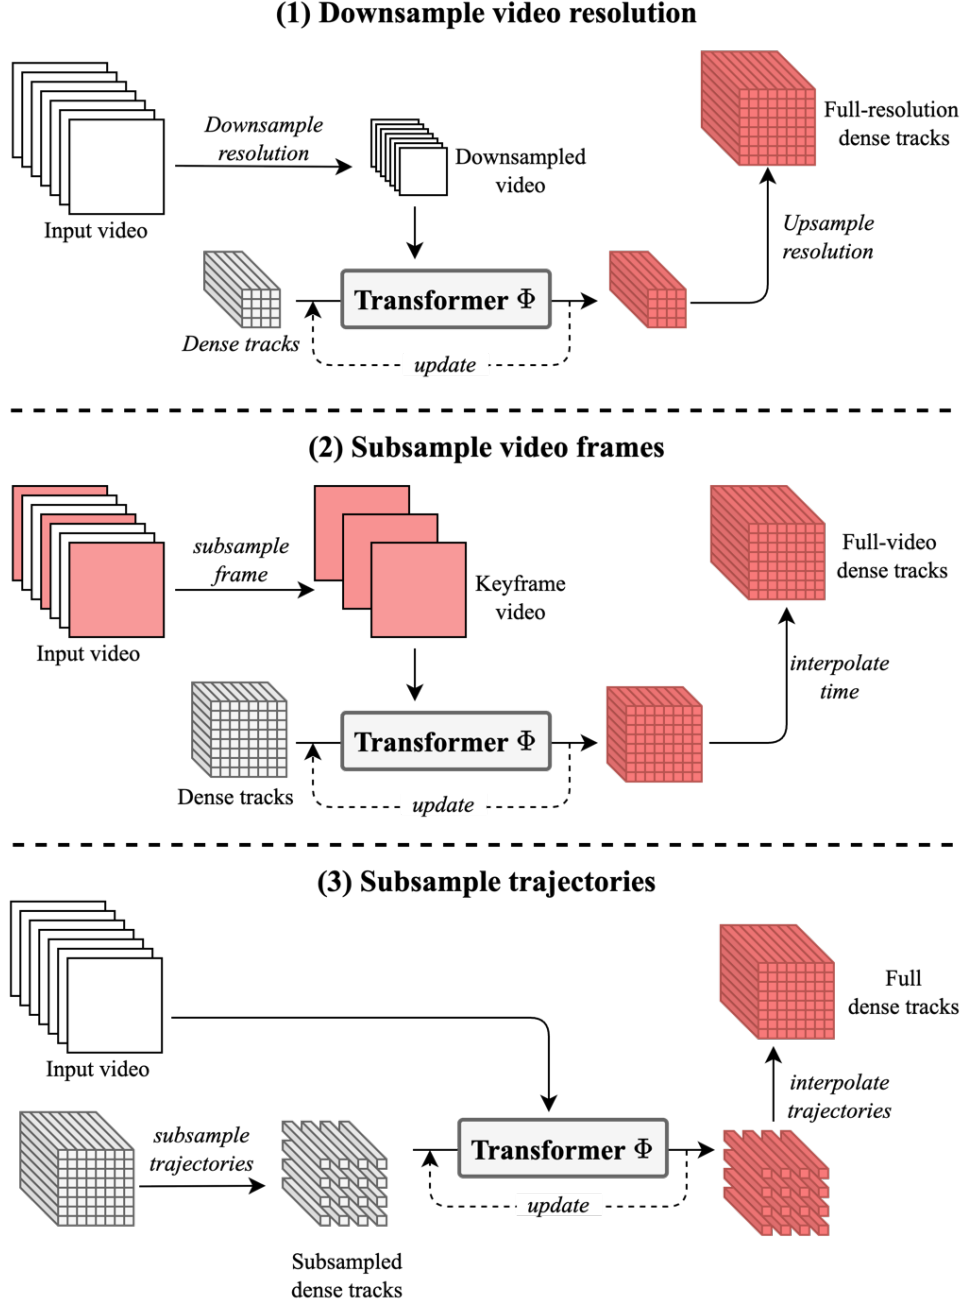




Figure 7: Illustration of three acceleration strategies. (1): Spatial downsampling of RGB input; (2) Temporal subsampling of input ; (3) Trajectory point subsampling;



(a) W/o Coarse-to-fine



(b) With Coarse-to-fine

Figure 8: DCT Spectrum of the predicted optical flows at every iteration and groundtruth between (a) w/o coarse-to-fine and (b) with coarse-to-fine.

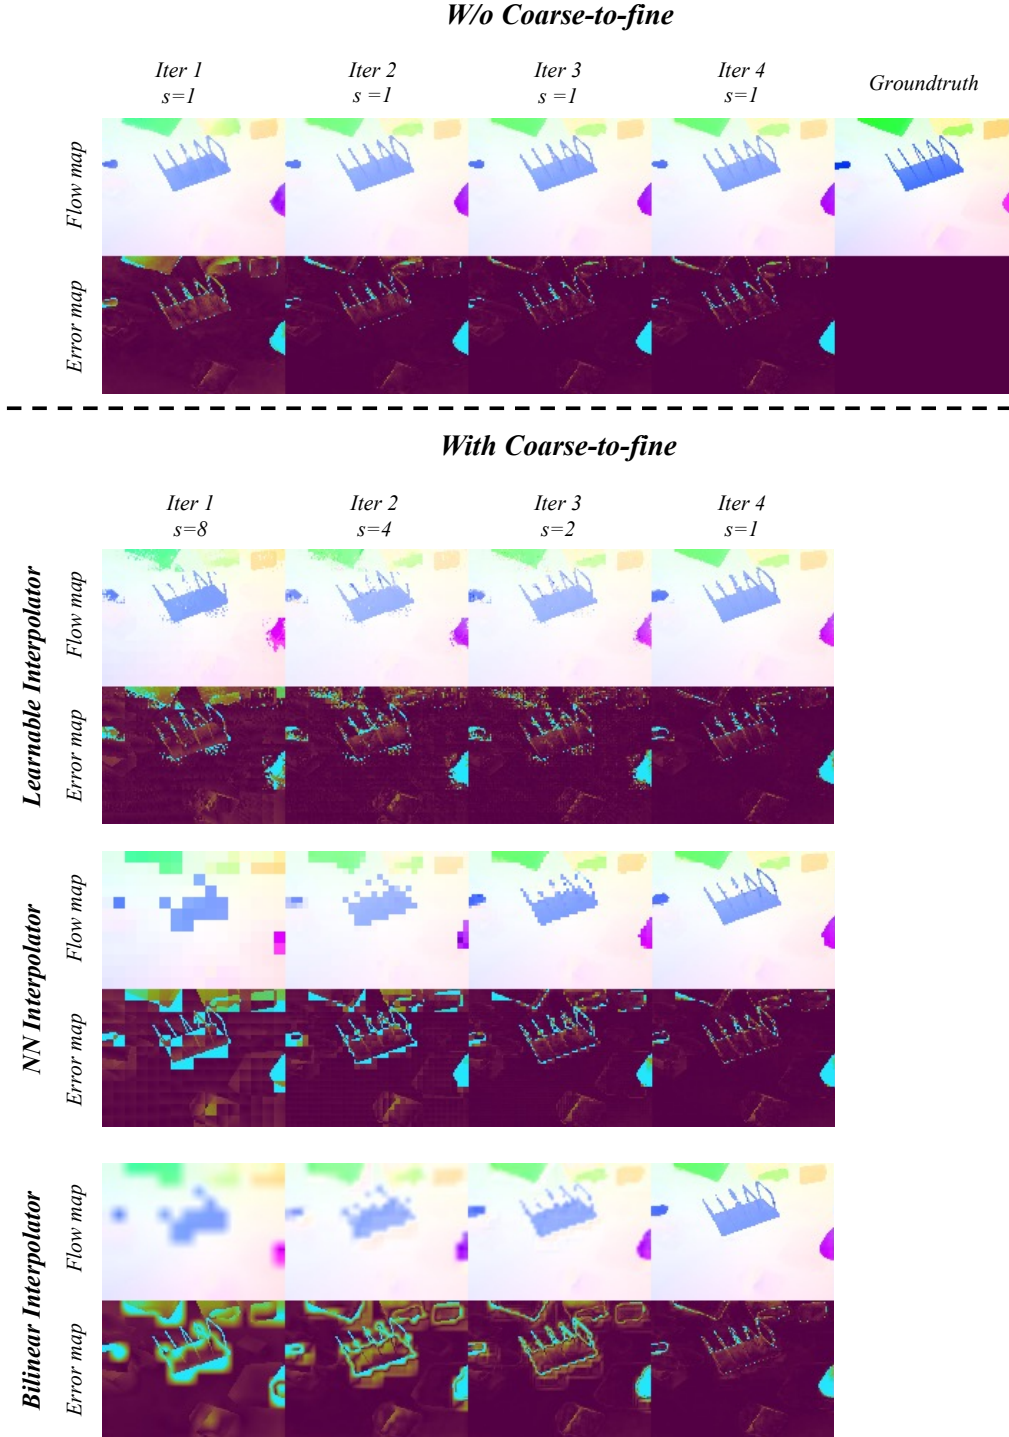


Figure 9: Visualization of predicted optical flow and corresponding error maps across different iterations. We compare models with and without the coarse-to-fine strategy, using various interpolation methods.

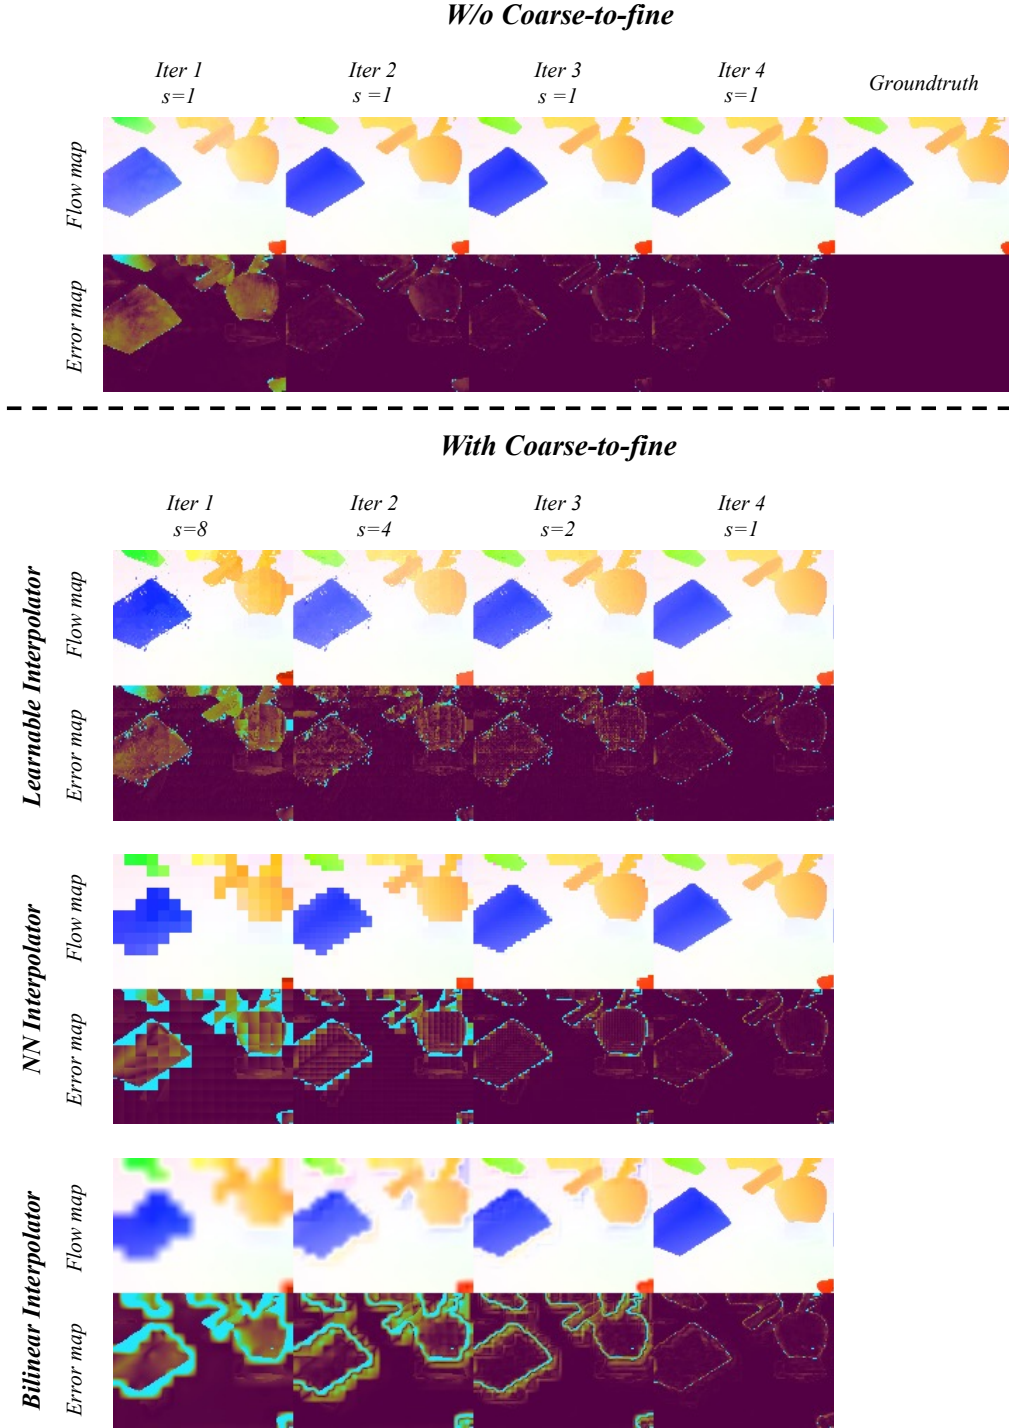


Figure 10: Visualization of predicted optical flow and corresponding error maps across different iterations. We compare models with and without the coarse-to-fine strategy, using various interpolation methods.

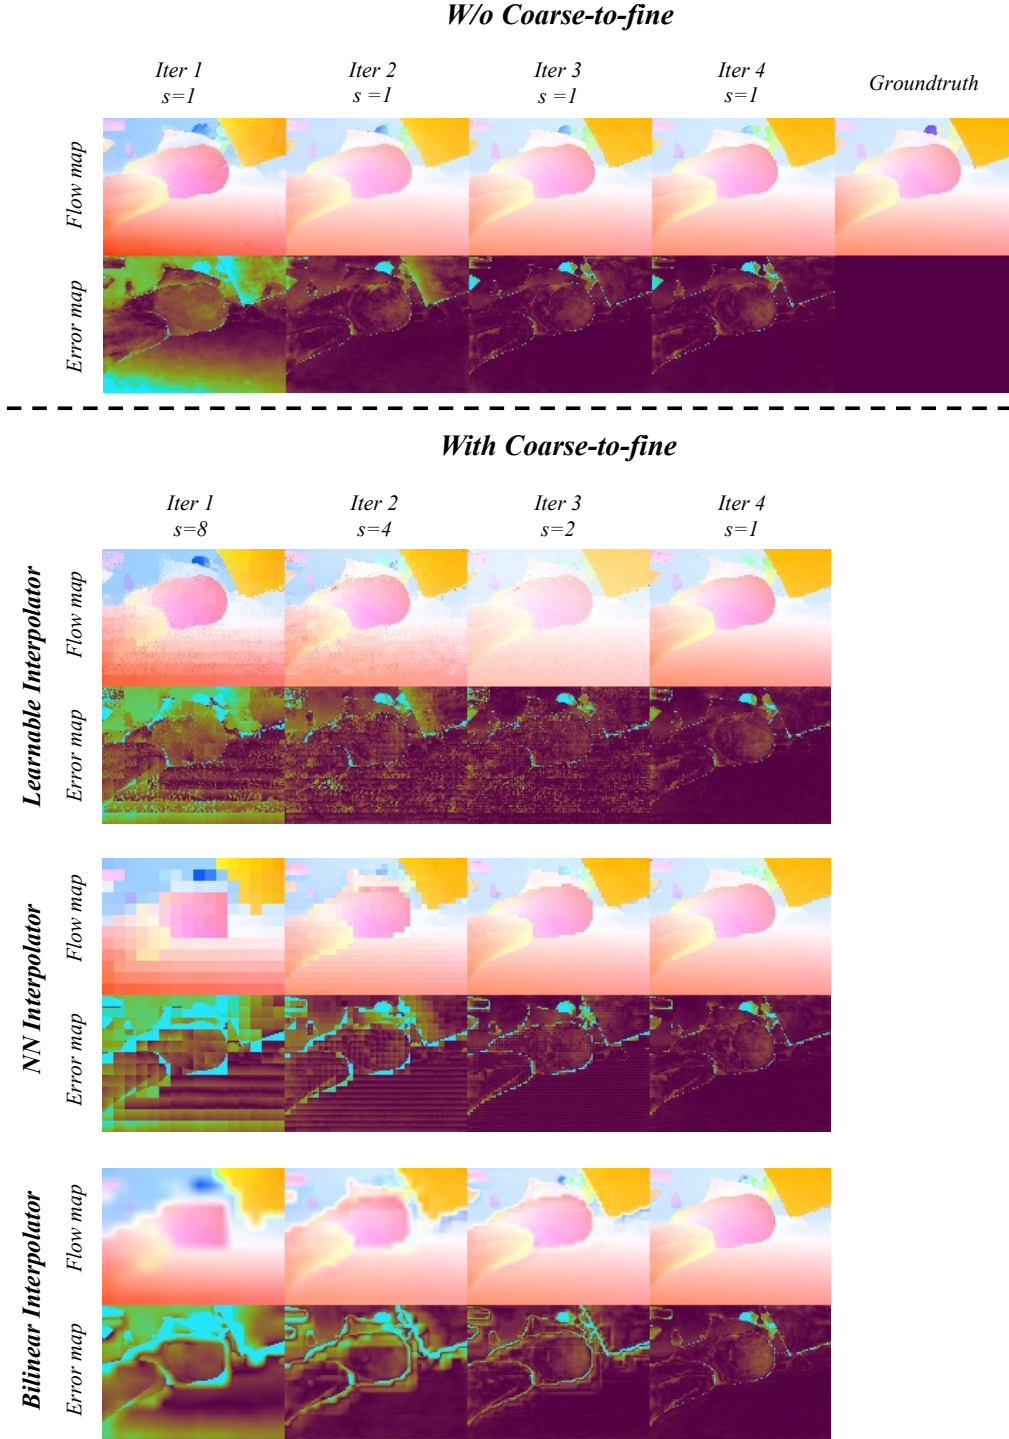


Figure 11: Visualization of predicted optical flow and corresponding error maps across different iterations. We compare models with and without the coarse-to-fine strategy, using various interpolation methods.

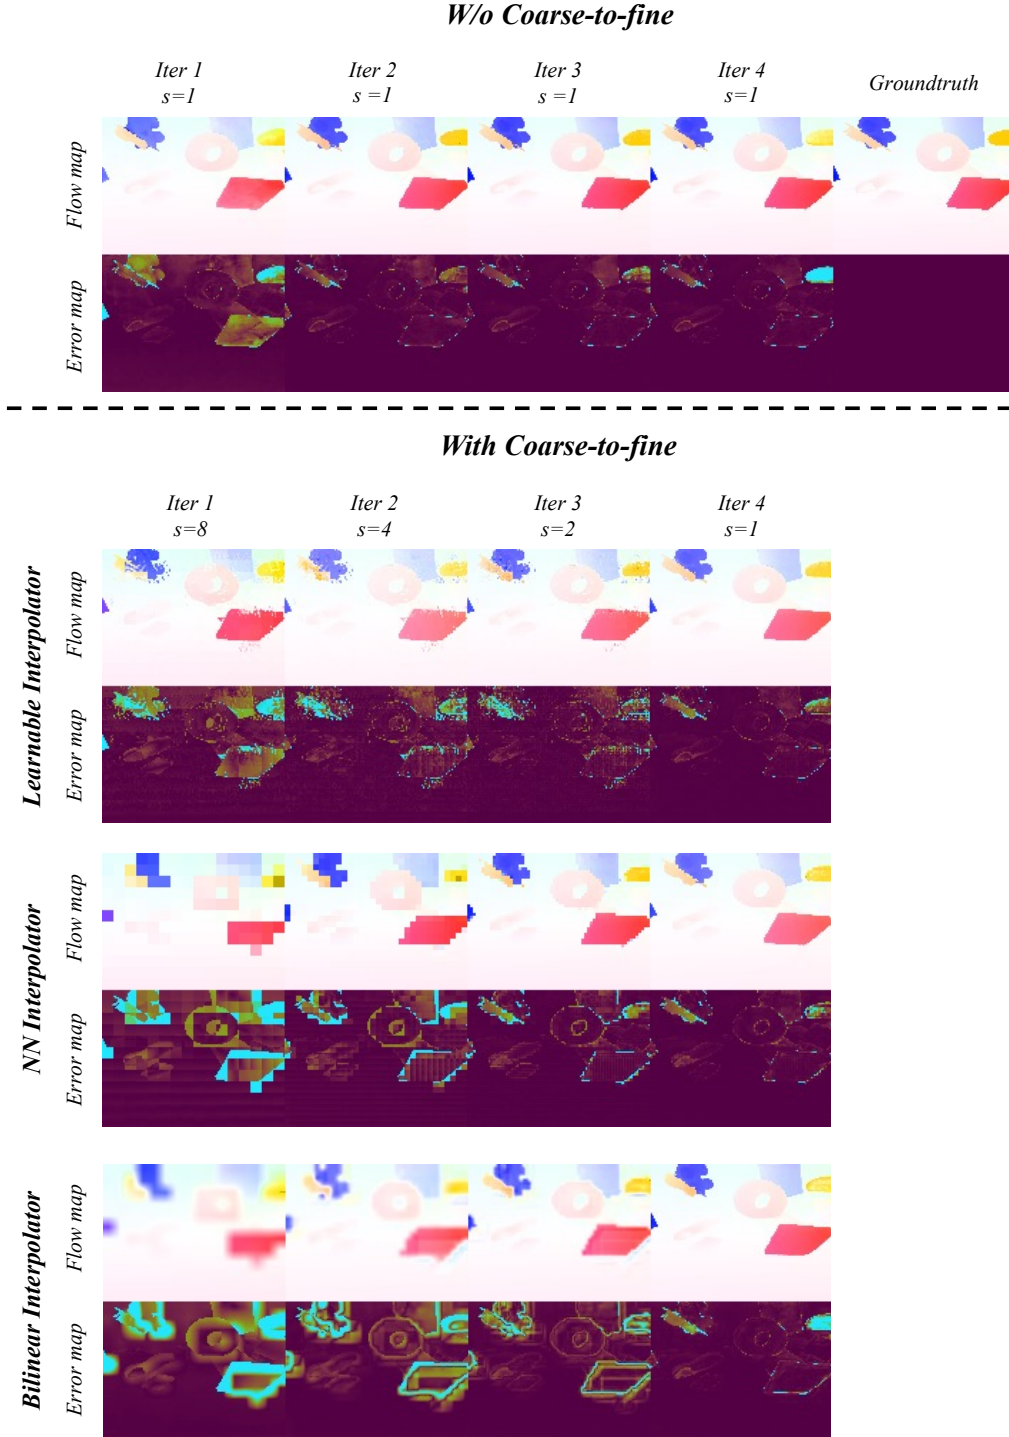


Figure 12: Visualization of predicted optical flow and corresponding error maps across different iterations. We compare models with and without the coarse-to-fine strategy, using various interpolation methods.