

Efficiently Reconstructing Dynamic Scenes One D4RT at a Time

Chuhan Zhang^{*} Guillaume Le Moing^{*} Skanda Koppula^{*◊} Ignacio Rocco^{*}
Liliane Momeni^{*} Junyu Xie^{◊1} Shuyang Sun^{*} Rahul Sukthankar^{*} Joëlle K. Barral^{*}
Raia Hadsell^{*} Zoubin Ghahramani^{*} Andrew Zisserman^{*◊} Junlin Zhang^{*}
Mehdi S. M. Sajjadi^{*2}

^{*}Google DeepMind [◊]University College London [◊]University of Oxford

Abstract

Understanding and reconstructing the complex geometry and motion of dynamic scenes from video remains a formidable challenge in computer vision. This paper introduces D4RT, a simple yet powerful feedforward model designed to efficiently solve this task. D4RT utilizes a unified transformer architecture to jointly infer depth, spatio-temporal correspondence, and full camera parameters from a single video. Its core innovation is a novel querying mechanism that sidesteps the heavy computation of dense, per-frame decoding and the complexity of managing multiple, task-specific decoders. Our decoding interface allows the model to independently and flexibly probe the 3D position of any point in space and time. The result is a lightweight and highly scalable method that enables remarkably efficient training and inference. We demonstrate that our approach sets a new state of the art, outperforming previous methods across a wide spectrum of 4D reconstruction tasks. We refer to the project webpage for animated results.³

1. Introduction

Traditional 3D reconstruction asks: ‘What is the geometry of everything, everywhere, all at once?’ We argue this exhaustive, rigid approach is fundamentally ill-equipped for a dynamic world. Despite the clear need for unified 4D understanding, leading approaches often tackle the problem by dividing it into discrete, task-specific components.

For instance, MegaSaM [29] relies on a complex mosaic of off-the-shelf models to separately estimate mono-depth, metric depth, and motion segmentation. Fusing these disparate signals requires computationally expensive test-time optimization to enforce geometric consistency. Recent feedforward approaches such as VGGT [48] employ

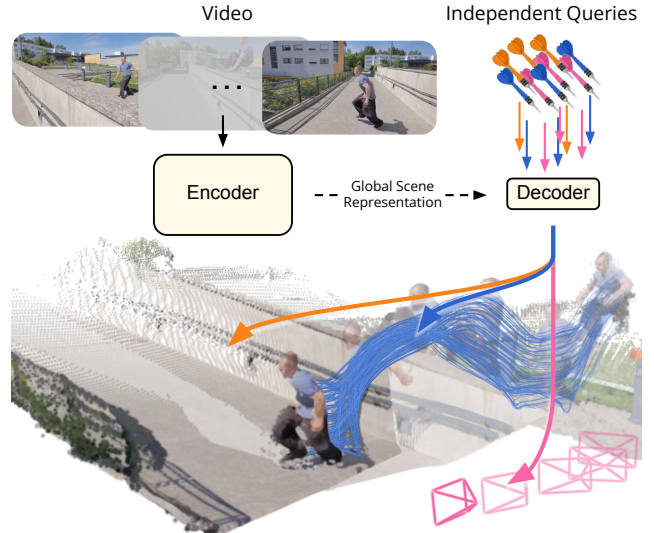
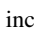
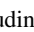
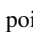


Figure 1. **D4RT** is a unified, efficient, feedforward method for *Dynamic 4D Reconstruction and Tracking*, unlocking a variety of outputs including point cloud () , point tracks () , camera parameters () through a single interface.

separate, specialized decoders for distinct modalities. Crucially, neither of these methods is capable of establishing correspondences for *dynamic* portions of the scene. While SpatialTrackerV2 [56] incorporates dynamics, it still lacks a unified, single-stage formulation, instead relying on costly iterative refinement.

We propose shifting the paradigm from fragmented, frame-level decoding to efficient, on-demand querying. We introduce D4RT, a feedforward method leveraging a flexible and scalable architecture to achieve full 4D reconstruction. As shown in Fig. 1, our model first encodes the input video, generating a latent scene representation which is then used to independently decode any number of spatio-temporal point queries. This simple, novel design unifies all 4D reconstruction tasks into a single interface and unlocks efficient training and inference.

¹Work was done during an internship at Google DeepMind.

²Correspondence: d4rt@msajjadi.com

³Project website: <https://d4rt-paper.github.io/>

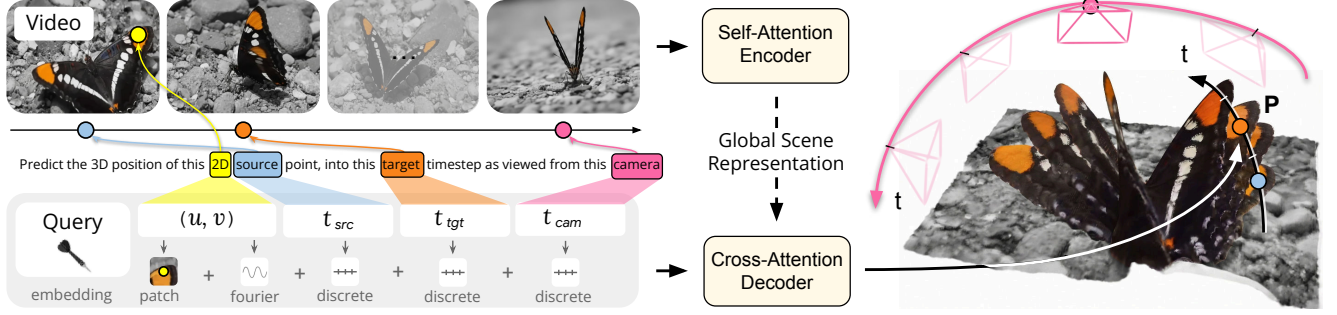


Figure 2. **D4RT model overview** – A global self-attention encoder first transforms the input video into the latent *Global Scene Representation* F , which is passed to a lightweight decoder. The decoder can be independently queried for the *3D position* \mathbf{P} of any given 2D point (u, v) from the source timestep t_{src} at target timestep t_{tgt} in camera coordinate t_{cam} , unlocking full decoding at any point in space and time. The query also contains an embedding of the local video patch centered around (u, v) , providing additional spatial context.

Our key contributions are as follows:

- We propose D4RT, a novel method for efficient feedforward querying of point-level 4D scene information captured in a video.
- We demonstrate how our unified approach unlocks 4D correspondence, point clouds, depth maps, and camera parameters for both static and dynamic scenes through a single interface.
- In an extensive set of experiments, we show that D4RT sets a new state of the art in *Dynamic 4D Reconstruction and Tracking* while outperforming existing approaches in both speed and accuracy.
- Finally, we demonstrate how our flexible decoder unlocks an efficient algorithm to track all pixels in a video, enabling dense, holistic scene reconstruction.

2. Method

D4RT is based on a simple encoder-decoder architecture inspired by the Scene Representation Transformer [39, 40]. As shown in Fig. 2, the video is first processed by a powerful encoder producing the *Global Scene Representation* F . The role of the encoder is to capture information about the full environment, identifying dense correspondence across all video frames, as well as understanding the flow of time and its effect on the scene. In a second stage, a lightweight decoder queries F through a simple low-level interface. Specifically, given a 2D point in a *source frame*, the decoder predicts the *3D position* of the point at a given *target timestep* (defining the temporal state) and expressed relative to a given *camera* reference (defined by the frame timestep where the camera viewpoint corresponds to this reference).

We draw attention to three desirable properties of this formulation: first, the indices need not coincide, allowing a full disentanglement of space and time; second, each query is decoded independently, allowing for both efficient training and inference as well as flexible decoding (both sparse and dense); and third, this interface unlocks a suite of down-

stream applications in a unified, consistent manner (Tab. 1).

2.1. D4RT Framework

Given a video $V \in \mathbb{R}^{T \times H \times W \times 3}$, the encoder \mathcal{E} extracts the latent *Global Scene Representation*

$$F = \mathcal{E}(V) \in \mathbb{R}^{N \times C}.$$

Once F is calculated, it remains fixed throughout the second stage, where the decoder \mathcal{D} cross-attends from any number of queries into F . We define a query $\mathbf{q} = (u, v, t_{src}, t_{tgt}, t_{cam})$, where (u, v, t_{src}) correspond to *source* parameters and (t_{tgt}, t_{cam}) correspond to *target* parameters. Here, $(u, v) \in [0, 1]^2$ represent the normalized 2D coordinates of a point of interest in the source frame t_{src} , while $(t_{tgt}, t_{cam}) \in [1, \dots, T]$ denote the temporal indices of the target timestep and the reference camera coordinate system (illustrated in Fig. 2). Each query \mathbf{q} is processed *fully independently* with the video features F to produce its corresponding 3D point position \mathbf{P} :

$$\mathbf{P} = \mathcal{D}(\mathbf{q}, F) \in \mathbb{R}^3.$$

From query to 4D reconstruction. Through a simple variation of queries, our framework allows us to address a broad range of 4D tasks as shown in Tab. 1. Choosing any fixed point (u, v) from a source frame t_{src} in the video while varying $t_{tgt} = t_{cam} = \{1 \dots T\}$ produces its *point track*, the 3D trajectory of the corresponding point throughout the video. For full *point cloud* reconstruction, the 3D position of all pixels in the video can be directly predicted in a shared reference frame t_{cam} by the model. This alleviates the need for coordinate transformations to map pixels from different video frames into a unified coordinate system using explicit, potentially noisy camera estimates. *Depth maps* can be recovered by simply querying any pixel in the video with $t_{src} = t_{tgt} = t_{cam}$ and only keeping the Z-dimension of the output \mathbf{P} .






Task	Query				
	u	v	t_{src}	t_{tgt}	t_{cam}
 Point Track	Fixed	Fixed	Fixed		$1 \dots T$
 Point Cloud	$1 \dots W$	$1 \dots H$		$1 \dots T$	Fixed
 Depth Map	$1 \dots W$	$1 \dots H$		$1 \dots T$	
 Extrinsics	$1 \dots h$	$1 \dots w$		Fixed	$1 \dots T$
 Intrinsics	$1 \dots h$	$1 \dots w$		$1 \dots T$	

Table 1. **Unified decoding** – A diverse set of geometry-related tasks can be inferred by querying the Cartesian product of the respective entries. Note that for intrinsics and extrinsics, we only query a coarse (h, w) grid for faster inference.

We next detail how *camera extrinsics* and *intrinsics* predictions are obtained. To derive the relative camera pose between any pair of video frames $i, j \in [1 \dots T]$, we create queries for an ensemble of source points $\{(u_k, v_k)\}_k$ sampled on a (h, w) grid in both reference frames:

$$\mathbf{q}_{i,k} = (u_k, v_k, i, i, i), \quad \mathbf{q}_{j,k} = (u_k, v_k, i, i, j).$$

The resulting sets $\{\mathcal{D}(\mathbf{q}_{i,k}, F)\}_k$ and $\{\mathcal{D}(\mathbf{q}_{j,k}, F)\}_k$ describe the same 3D points in different reference frames. We therefore only need to find the rigid transformation between them, which can be efficiently derived through Umeyama’s algorithm [46] that solves a 3×3 SVD decomposition.

To recover intrinsics for video frame $i \in [1 \dots T]$, we construct a set of queries for different source points, again sampled on a (h, w) grid. We decode all corresponding 3D positions $\mathbf{P} = (p_x, p_y, p_z)$. Assuming a pinhole camera model with a principal point at $(0.5, 0.5)$, we get focal length parameters as follows:

$$f_x = p_z(u - 0.5)/p_x, \quad f_y = p_z(v - 0.5)/p_y.$$

We take the median over the k estimates for robustness. Camera models with distortion (e.g., fisheye) can also be seamlessly incorporated by adding a non-linear refinement step on top of the initial estimation [61].

2.2. Model Architecture

Our encoder \mathcal{E} is based on the Vision Transformer [9] with interleaved *local frame-wise*, and *global* self-attention layers [48]. For simplicity, and to support arbitrary aspect ratios, we resize the input video to a fixed square resolution before tokenizing it. To incorporate the original aspect ratio, we embed it into a separate token and pass it to the transformer along with the main video tokens.

Pointwise decoder. The decoder \mathcal{D} is a small cross-attention transformer. A query token is first constructed by adding the Fourier feature embedding [47] of the 2D coordinates (u, v) to the learned discrete timestep embeddings for t_{src} , t_{tgt} , and t_{cam} . We empirically observe that

Algorithm 1 Efficient Dense Tracking of All Pixels.

Require: Input video V , model encoder \mathcal{E} and decoder \mathcal{D} .

```

1:  $F \leftarrow \mathcal{E}(V)$  ▷ Compute Global Scene Representation
2:  $G \leftarrow \{\text{false}\}^{T \times H \times W}$  ▷ Initialize Occupancy Grid
3:  $\mathcal{T} \leftarrow \emptyset$  ▷ Initialize Set of Dense Tracks
4: while any( $G = \text{false}$ ) do
5:   Sample a batch  $B$  of unvisited source points from  $G$ 
6:   for each  $(u, v, t_{src}) \in B$  do ▷ Process batch in parallel
7:      $Q \leftarrow \{u, v, t_{src}, t_{tgt}=t_{cam}\}_{k=1}^T$  ▷ Get Track Queries
8:      $P \leftarrow \{\mathcal{D}(\mathbf{q}_k, F)\}_{k=1}^T$  ▷ Run the decoder
9:      $G \leftarrow \text{Visible}(P)$  ▷ Set visible track pixels as visited
10:     $\mathcal{T} \leftarrow \mathcal{T} \cup P$  ▷ Add new track to the output
11:   end for
12: end while
13: return  $\mathcal{T}$ 

```

augmenting the query with an embedding of the local 9×9 pixel RGB patch centered at (u, v) dramatically improves performance, see Sec. 4.4.

Each query is decoded independently through cross-attention into the *Global Scene Representation* F , ensuring that queries do not interact. The resulting output feature is then mapped to a 3D point position \mathbf{P} via a simple learned projection. Decoding queries independently is a deliberate design decision with major advantages. It allows efficient training, as only a small number of queries need to be decoded to provide a supervision signal to the model. Equally importantly, at inference time, the queries can be chosen freely across all video frames as described in Sec. 2.3, and need not be correlated to each other to avoid out-of-distribution effects – indeed, we have empirically observed major performance drops when enabling self-attention between queries in early experiments. Finally, this design allows highly efficient inference thanks to its trivial parallelism. As shown in Fig. 3 and Tab. 3, we obtain the best trade-off between performance and latency across multiple tasks.

2.3. Training and Inference

The model is implemented in Kauldron [15] and trained end-to-end by minimizing the weighted sum of losses computed over a batch of N sampled queries. The primary supervision signal is derived from an L_1 loss applied to the normalized 3D point position \mathbf{P} . Specifically, both the target and the estimated point sets are normalized by their respective mean depths [51] and then passed through the transform $\text{sign}(x) \cdot \log(1+|x|)$ to dampen the influence of far-away points on the loss. We also supervise a set of auxiliary predictions from additional linear projections on the decoder output: An L_1 loss on *2D coordinates* of the point positions in image space; cosine similarity for *3D surface normals* [54]; binary cross-entropy for *target point visibility*; and L_1 on the vector of *point motion*. All loss terms are

Model	Task		Functionality		Architecture	
	3D Recon- struction	Dynamic Cor- respondence	Flexible Ref. Frames	Sparse Decoding	Global Context	Single Decoder
MegaSaM [29]	✓	✗	✗	✗	✗	✗
DUST3R[51]	✓	✗	✗	✗	✗	✗
VGGT [48]	✓	✗	✗	✗	✓	✗
π^3 [54]	✓	✗	✓	✗	✓	✗
St4RTrack [11]	✓	✓	✗	✗	✗	✗
STv2 [56]	✓	✓	✗	✓	✓	✗
D4RT (Ours)	✓	✓	✓	✓	✓	✓

Table 2. **Model capabilities** – We highlight both the tasks our model executes but also its comprehensive functionality and simple model architecture.

applied only where ground truth supervision is available. We finally incorporate a confidence penalty $-\log(c)$ where c additionally weights the 3D point error [25].

Efficient dense dynamic correspondence. A key capability of our query-based model is that it can efficiently compute dense correspondences for *all* pixels in a video, both static and dynamic. As shown in Fig. 4, this capability is crucial for building a complete, holistic scene reconstruction, which in turn is key to eliminating the occlusion-induced discontinuities and sparse artifacts common in prior works. However, a naive approach to reconstruct tracks for all pixels across the video would involve $O(T^2HW)$ queries, the majority of which are not required.

We introduce Alg. 1 which exploits spatio-temporal redundancy using an occupancy grid $G \in \{0, 1\}^{T \times H \times W}$ to speed this procedure up significantly. The algorithm only initiates new tracks from unvisited pixels. Each *full-video track* marks all spatio-temporal pixels it *visibly* passes through as visited. Empirically, we find that this yields an adaptive speedup between 5–15 \times depending on the motion complexity in the video. This dense, flexible strategy is feasible for our method precisely because our decoder is *both* sparse, and lightweight. It is notable that prior works are ill-suited for this task for various reasons. Most methods fail to provide any correspondences for dynamic portions of the scenes [29, 48, 54], dense frame-level decoding remains locked in the costly naive approach [11], and models with heavy sparse decoders face a large per-query cost [23, 56].

3. Related Work

Classical approaches to 3D reconstruction are fundamentally anchored in multi-view geometry, specifically Structure-from-Motion (SfM) [19, 33] and Multi-View Stereo (MVS) [12, 43]. The standard pipeline exemplified by COLMAP [42] incrementally estimates and optimizes sparse geometry and camera poses. The advantage of these

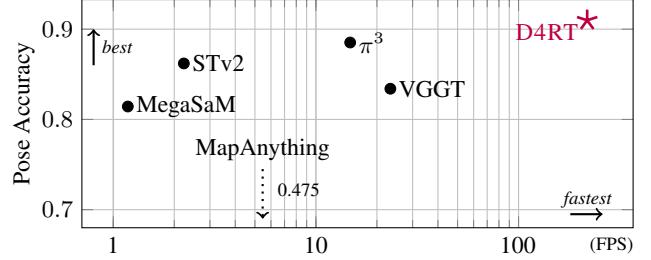


Figure 3. **Pose accuracy vs. speed** – We compare pose accuracy vs. throughput against recent state-of-the-art methods. Pose accuracy is $1 - \text{error}$, averaged over ATE/RTE/RPE on Sintel and ScanNet. Throughput is measured in FPS on an A100 GPU. D4RT achieves 200+ FPS pose estimation, 9 \times faster than VGGT, and 100 \times faster than MegaSaM, while delivering superior accuracy.

methods is their explicit enforcement of geometric consistency and mathematically interpretable reconstructions. However, they are computationally intensive and brittle.

3.1. Feedforward 3D Reconstruction

The field of 3D reconstruction has recently shifted towards end-to-end, feedforward models that directly infer geometry from images. The seminal work DUST3R [51] demonstrated that Transformer-based networks [47] can perform 3D reconstruction from unposed and uncalibrated image pairs in an end-to-end approach. VGGT [48] later scaled this approach beyond pairs by using a Vision Transformer [9] with global attention. Building on this feedforward paradigm, several works have extended the methodology to dynamic videos [3, 50, 52, 60] and more efficient inference [54, 57]. However, these models share significant limitations. For instance, many do not support changing camera intrinsics within a video [29, 48, 54], and several are restricted to using only the first frame as the camera reference [29, 50, 51]. More importantly, architectures inspired by VGGT [48] either use separate decoder heads for each task [50, 54, 60] (*e.g.*, depth, pose, point cloud), or they incorporate separate models for sub-tasks of 4D reconstruction [31, 44, 52], making the full pipeline cumbersome and computationally expensive to run. Finally, these methods are not directly capable of providing correspondences for dynamic regions of the scene.

3.2. From 2D to 3D Point Tracking

The task of point tracking aims to establish long-term 2D correspondences through challenges like occlusions and non-rigid motion, evolving from early methods such as Particle Video [41] to newer deep-learning approaches [7, 8, 17, 63]. This field has further progressed from tracking a sparse set of points to the dense tracking of every pixel in a video [18, 28, 32]. A parallel line of work has explored lifting these 2D tracks into 3D, which is achieved by project-

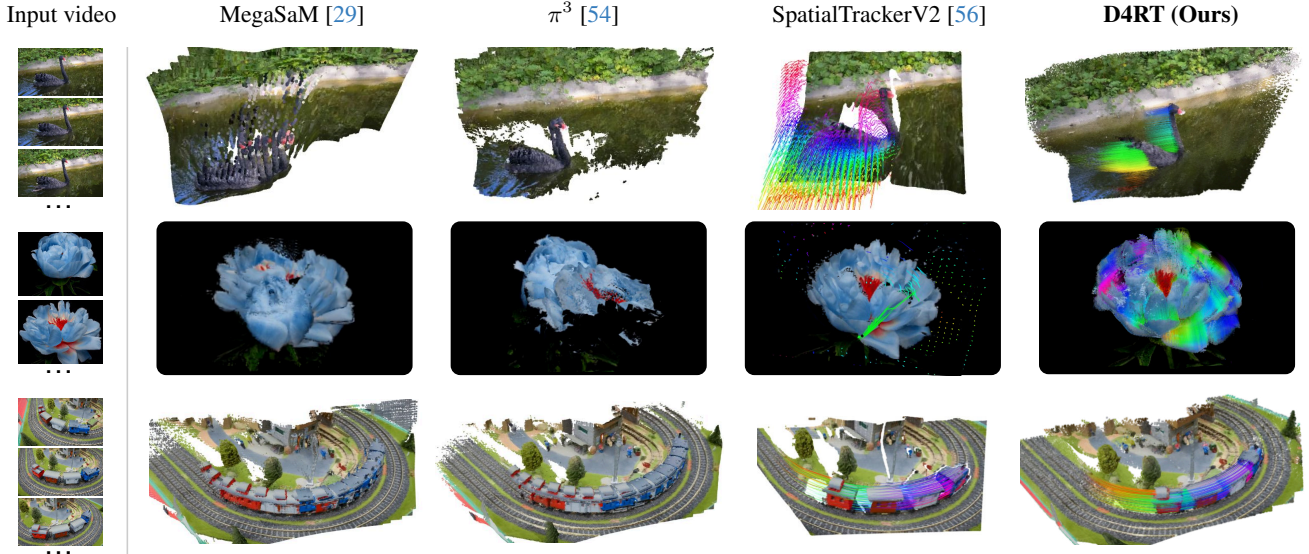


Figure 4. **Reconstruction results across methods** – Pure reconstruction methods (MegaSaM and π^3) are only able to accumulate point clouds of all pixels; exhibiting clear failure cases in dynamic scenes. For example, the swan is repeated in MegaSaM’s reconstruction, and π^3 is failing entirely to reconstruct the flower. SpatialTrackerV2, a state-of-the-art tracking method, successfully captures dynamics, however its design only allows tracking points from *one* frame, leaving gaps in the reconstruction (behind the swan and train). D4RT is the only method that successfully reconstructs a full 4D representation of the scene including *all* pixels of the video.

ing tracks from image space to camera coordinates using ground truth or predicted depth and camera intrinsics [27].

Most recently, models like SpatialTracker [55, 56], L4P [1], DPM [44] and St4RTracker [11] have moved to directly predicting 3D tracks without priors. Yet, limitations remain: St4RTracker [11] and DPM [44] follow the pairwise paradigm of DUST3R [51], preventing holistic video processing. SpatialTrackerV2 [56] is a multi-stage approach and exhibits slow inference speeds as it relies on iterative refinements. Meanwhile, L4P [1] requires different heads for sparse and dense outputs of the same nature (*e.g.*, flow vs. 2D tracking or depth vs. 3D tracking). In contrast, our approach unifies these tasks within a single architecture. Tab. 2 summarizes capabilities across a recent set of state-of-the-art models.

4. Experiments

We benchmark D4RT against recent state-of-the-art methods. After inspecting qualitative differences in capability in Sec. 4.1, we proceed with evaluating 4D Reconstruction and Tracking performance in Sec. 4.2, focusing the comparison on the best methods from Sec. 3.2. We then continue with pure reconstruction tasks in Sec. 4.3, comparing to the recent SOTA models mentioned in Sec. 3.1. We conclude with a set of key ablations in Sec. 4.4.

Training setup. For the encoder, we use the ViT-g model variant with 40 layers on a spatio-temporal patch size of $2 \times 16 \times 16$. This encoder and our 8-layer cross-

Method	Max. Track Count @ Target FPS			
	60 FPS	24 FPS	10 FPS	1 FPS
DELTA [32]	0	5	408	5,770
SpatialTrackerV2 [56]	29	84	219	2,290
D4RT (Ours)	550	1,570	3,890	40,180

Table 3. **3D tracking throughput** – We measure the maximum number of full-video 3D point tracks that different model can produce while maintaining a given FPS target on a single A100 GPU. Note that for D4RT, each track consists of T independent queries processed by the decoder. D4RT is 18–300× faster than others.

attention decoder contain 1 B and 144 M parameters, respectively. Our training mixture contains public and internal datasets including BlendedMVS [58], Co3Dv2 [38], Dynamic Replica [22], Kubric [16], MVS-Synth [20], PointOdyssey [62], ScanNet++ [59], ScanNet [4], Tantanair [53], VirtualKitti [13], and Waymo Open [45]. We train on 48-frame clips at 256×256 resolution, decoding 2048 random queries that are oversampled on specific regions for more efficient training, see the appendix for details. The model is trained for 500 k steps using the AdamW optimizer [26, 30] with a local batch size of 1 across 64 TPU chips, taking a total of just over 2 days to complete.

4.1. Qualitative Analysis

We begin our evaluations by inspecting video reconstruction quality in Fig. 4. Dynamic objects (*e.g.*, swan and



Figure 5. **Visualizations on in-the-wild videos** – D4RT demonstrates accurate reconstructions on static (top row) and dynamic scenes (bottom row). In the presence of motion, D4RT additionally produces robust 3D point trajectories.

		TAPVid-3D												
w/ GT intrin.	Method	Camera Coordinate 3D tracking									World Coordinate 3D tracking			
		DriveTrack			ADT			PStudio			DriveTrack		ADT	
		AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	AJ \uparrow	APD _{3D} \uparrow	OA \uparrow	APD _{3D} \uparrow	L1 \downarrow	APD _{3D} \uparrow	L1 \downarrow
✗	St4RTrack [11]	-	-	-	-	-	-	-	-	-	0.020	0.499	0.062	0.839
	CoTracker3 [23] + UniDepthV2 [36]	0.038	0.062	0.856	0.102	0.146	0.937	0.119	0.176	0.862	-	-	-	-
	CoTracker3 [23] + VGGT [48]	0.129	0.189	0.856	0.132	0.190	0.937	0.045	0.070	0.862	0.205	0.170	0.192	0.736
	SpatialTrackerV2 [56]	0.064	0.100	0.865	0.260	0.342	0.936	0.097	0.152	0.805	0.101	0.139	0.336	0.238
	D4RT (Ours)	0.257	0.345	0.875	0.240	0.319	0.926	0.138	0.186	0.897	0.373	0.020	0.319	0.096
✓	CoTracker3 [23] + UniDepthV2 [36]	0.147	0.225	0.856	0.173	0.254	0.937	0.205	0.311	0.862	-	-	-	-
	CoTracker3 [23] + VGGT [48]	0.245	0.342	0.856	0.175	0.250	0.937	0.215	0.318	0.862	0.292	0.160	0.234	0.755
	DELTA [32]	0.170	0.238	0.874	0.199	0.258	0.879	0.175	0.261	0.760	-	-	-	-
	SpatialTrackerV2 [56]	0.195	0.275	0.865	0.303	0.404	0.936	0.175	0.270	0.805	0.201	0.117	0.378	0.263
	D4RT (Ours)	0.304	0.410	0.875	0.307	0.408	0.926	0.372	0.498	0.897	0.470	0.017	0.398	0.093

Table 4. **4D reconstruction and tracking** – We evaluate 3D tracking capability on dynamic videos, with tracks predicted in both local camera coordinates (left) and world coordinates (right). Our model achieves superior performance compared to the prior state-of-the-art.

train) reveal qualitative differences, highlighting performance gaps between different methods. Pure reconstruction methods (see Sec. 3.1), exemplified by the state-of-the-art models MegaSaM and π^3 , fail to understand dynamics in the scenes, producing either repeated images of the entities as they move through the scene, or failing to reconstruct them altogether. Conversely, tracking-based models successfully reconstruct and track dynamics. However, they commonly only track points from one frame, leaving gaps in the reconstruction in all areas that are occluded in the first frame as seen for SpatialTrackerV2 [56] here. Meanwhile, D4RT is able to track all dynamic pixels of the video in a unified reference frame. We further evaluate our method on diverse in-the-wild videos in Fig. 5, demonstrating its performance to both static and dynamic scenes.

4.2. 4D Reconstruction and Tracking

We now turn to the quantitative evaluation of 4D Reconstruction and Tracking, comparing D4RT to other methods

that can estimate dynamic correspondences. We evaluate on TAPVid-3D [27] which is comprised of three subsets of real-world videos in challenging settings.

We first evaluate 3D tracking in the local camera frame using the standard protocol from Koppula et al. [27], which measures a model’s ability to predict and track the 3D position of every observed point in the *local* camera reference frames. We report standard 3D tracking metrics *Average percent of points within delta error* APD_{3D}, *Occlusion Accuracy* (OA) and *3D Average Jaccard* (AJ). As shown in Tab. 4 (left), D4RT achieves state-of-the-art results against competing methods, both with and without known ground-truth intrinsics.

Tab. 4 (right) shows results for *world coordinate 3D tracking*. This task measures a model’s ability to predict tracks within a single, consistent world coordinate system, thereby also measuring the models’ ability to implicitly change reference frames. PStudio is omitted here as it has no camera motion, rendering this task effectively identical

Method	Point Cloud		Video Depth							
	Sintel	Scannet	Sintel		ScanNet		KITTI		Bonn	
	L1 ↓	L1 ↓	AbsRel (S) ↓	AbsRel (SS) ↓	AbsRel (S) ↓	AbsRel (SS) ↓	AbsRel (S) ↓	AbsRel (SS) ↓	AbsRel (S) ↓	AbsRel (SS) ↓
MegaSaM [29]	1.531	0.072	0.342	0.249	0.050	0.047	0.109	0.101	0.056	0.056
VGGT [48]	1.582	0.063	0.318	0.247	0.044	0.040	0.094	0.067	0.055	0.051
MapAnything [24]	1.718	0.064	0.397	0.306	0.043	0.035	0.096	0.090	0.076	0.049
SpatialTrackerV2 [56]	1.375	0.036	0.209	0.175	0.027	0.025	0.075	0.064	0.042	0.978
π^3 [54]	1.139	0.030	0.241	0.163	0.021	0.019	0.055	0.053	0.033	0.032
D4RT (Ours)	0.768	0.028	0.171	0.148	0.020	0.018	0.055	0.051	0.036	0.036

Table 5. **Video depth and point map estimation** – Quantitative results for both video depth and point map estimation across four benchmarks. D4RT achieves top-tier performance on the depth estimation task under both scale-only (S) and scale-and-shift (SS) alignments.

Method	Sintel			ScanNet			Re10K Pose AUC ↑
	ATE ↓	RPE-T ↓	RPE-R ↓	ATE ↓	RPE-T ↓	RPE-R ↓	
MegaSaM [29]	0.074	0.030	0.126	0.029	0.016	0.839	71.0
VGGT [48]	0.168	0.056	0.428	0.016	0.012	0.316	70.2
MapAnything [24]	0.202	0.089	2.383	0.023	0.016	0.438	68.7
SpatialTrackerV2 [56]	0.126	0.053	1.052	0.018	0.012	0.324	75.7
π^3 [54]	0.086	0.039	0.248	0.015	0.010	0.291	78.7
D4RT (Ours)	0.065	0.024	0.126	0.014	0.010	0.302	83.5

Table 6. **Camera pose estimation** – We evaluate D4RT against state-of-the-art methods on static indoor scenes (ScanNet, Re10K) and dynamic outdoor scenes (Sintel).

to the preceding one. Since there is no concept of occlusion for this task, we report APD_{3D} as well as the L1 distance between the predicted and ground-truth tracks which provides a comprehensive, global measure of deviation. We observe that our model excels at this task as well, exhibiting strong improvements across metrics.

We finally compare efficiency in Tab. 3, by counting the number of tracks each model can produce for a given target frame-per-second rate (FPS). The measurements show that our model is 18-300 \times faster than prior methods.

4.3. 3D Reconstruction

We now evaluate the 3D reconstruction capabilities of our model. The most holistic task here is 3D point cloud reconstruction, which requires predicting all observed pixels in the same world coordinate system.

We evaluate on two standard benchmarks: MPI Sintel [2] and ScanNet [5], which consist of dynamic and static scenes, respectively. For evaluation, we first align the predicted and ground-truth point clouds via mean-shifting following Li et al. [29] before reporting the mean L1 distance. As shown in Tab. 5 (left), our model outperforms all recent state-of-the-art models across datasets.

Depth estimation. We evaluate depth estimation performance across four datasets: Sintel [2], ScanNet [5], KITTI [14], and Bonn [34]. Following prior works [29, 50], we align the predicted video depth with the ground-truth depth before computing metrics. Specifically, we adopt two alignment settings: *scale-only* (S) and *scale-and-shift* (SS)

alignment. Both settings determine a single global scale factor between the predicted and ground-truth depths across the sequence, while the latter introduces an additional 3D translation term, following an affine-invariant setup. After alignment, we compute the *Absolute Relative Error* (AbsRel) [10] between the predicted and ground-truth depths.

We show results in Tab. 5 (right). D4RT achieves top-tier performance across all benchmarks, with particularly strong results on Sintel which is the most challenging dataset due to its highly dynamic scenes.

Camera pose estimation. We evaluate camera pose estimation performance on the same datasets. Following prior work [21, 50], we evaluate on the 14-sequence subset of Sintel which provides full rendering with motion blur and atmospheric effects to measure in-the-wild performance. We report the *Absolute Translation Error* (ATE), *Relative Translation Error* (RPE trans), and *Relative Rotation Error* (RPE rot) after Sim(3) alignment with the ground truth as in [50, 54, 60], as well as the *Pose AUC@30* as in [48]. Tab. 6 presents the results, where our model strongly outperforms competing baselines on both indoor scenes (ScanNet and Re10K) and simulated outdoor scenes (Sintel).

Fig. 3 visualizes efficiency and accuracy for camera pose estimation across several recent state-of-the-art methods. To ensure a fair comparison, we remove the decoding heads of the baselines which are unrelated to camera estimation, making them faster. Despite this, we find that D4RT surpasses all existing methods in both accuracy and efficiency, notably outperforming MegaSaM in throughput by two orders of magnitude.

4.4. Ablation Studies

We present a series of ablation studies to validate our design choices and hyperparameters. Unless otherwise noted, we adopt a ViT-L backbone with 6 decoder layers as the default setting. The model is trained for 300k iterations with a batch size 16. All results are reported on the challenging Sintel dataset for both video depth and camera pose.

Local RGB patch. As described in Sec. 2.2, we introduce an additional embedding that encodes the local appearance

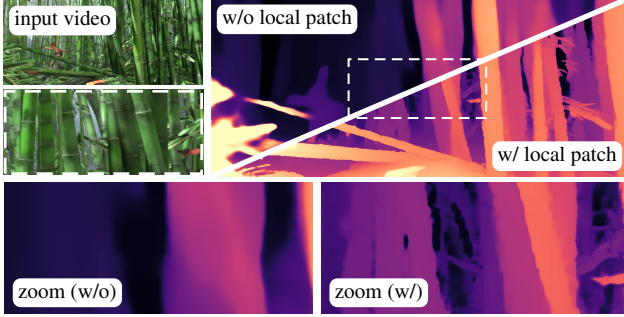


Figure 6. **Preservation of low-level details** – We ablate the effect of including local patch information into the model’s queries, visualizing the resulting depth maps on an example from the Sintel dataset. We find that the local RGB patches help preserve fine-grained details and produce sharper object boundaries.

w/ local appear. patch	Video Depth		Camera Pose		
	AbsRel (S) ↓	AbsRel (SS) ↓	ATE ↓	RPE-T ↓	RPE-R ↓
✗	0.366	0.306	0.173	0.031	0.262
✓	0.302	0.257	0.091	0.028	0.245

Table 7. **Local RGB patch** – We evaluate a ViT-L model, trained with and without the appearance patch as input to the decoder, on Sintel video depth and camera pose estimation.

around the query’s source location. This additional embedding yields dramatic performance improvements as shown in Tab. 7 and Fig. 6. These improvements can be attributed to two factors: (i) the local appearance information helps the query establish more reliable correspondences with the encoded spatiotemporal features; and (ii) these patches provide low-level cues that help to segment objects from their surroundings, leading to fine-grained predictions, as exemplified by the sharper depth estimates. We note that while most related works have specialized modules to preserve low level details, such as DPT [37] with its skip connections between encoder and decoder layers, our approach is much simpler. In Sec. C of the appendix, we further highlight how the patches unlock subpixel precision.

Encoder size. We now examine how performance scales with the size of our encoder. Tab. 9 summarizes our findings for ViT-B (90 M) to ViT-g (1 B). As the number of trainable parameters increases, we observe a significant performance improvement, particularly in depth estimation and RPE-R in camera pose estimation.

Auxiliary losses. We finally investigate the contribution of each auxiliary loss in Tab. 8. We observe a range of behaviors across the auxiliary losses: some significantly boost depth performance (e.g., 2D position, normal), others are critical for better camera pose estimation (e.g., confidence), and displacement has small improvements across the board.

Losses	Video Depth		Camera Pose		
	AbsRel(S) ↓	AbsRel(SS) ↓	ATE ↓	RPE-T ↓	RPE-R ↓
D4RT	0.302	0.257	0.091	0.028	0.245
w/o 2D position	+0.071	+0.063	+0.002	+0.002	-0.028
w/o normal	+0.043	+0.026	+0.003	+0.003	-0.022
w/o displacement	+0.011	+0.007	+0.011	+0.003	+0.007
w/o visibility	-0.003	-0.025	+0.012	+0.011	+0.022
w/o confidence	+0.002	-0.025	+0.126	+0.061	+0.115

Table 8. **Auxiliary losses** – We remove auxiliary losses individually to evaluate their impact on model performance. A slight trade-off between depth and camera estimation pose is observed, though we find that all auxiliary losses improve overall performance.

Backbone size	Video Depth		Camera Pose		
	AbsRel (S) ↓	AbsRel (SS) ↓	ATE ↓	RPE-T ↓	RPE-R ↓
ViT-B	0.319	0.232	0.145	0.034	0.266
ViT-L	0.256	0.214	0.073	0.027	0.191
ViT-H	0.226	0.173	0.070	0.028	0.186
ViT-g	0.191	0.168	0.078	0.026	0.160

Table 9. **Backbone size** – We examine how D4RT’s performance scales with the size of the pretrained ViT encoder backbone, evaluating video depth and camera pose estimation performance on Sintel. We observe a clear improvement as the backbone size increases from ViT-B to ViT-g.

5. Conclusion

In this work, we introduce D4RT, a simple yet highly scalable feedforward network that reconstructs dynamic 4D scenes with temporal correspondence. Our core innovation is an efficient, query-based decoder that allows for the independent prediction of any point’s 3D position in space and time. This flexible parametrization avoids the computational bottlenecks of dense per-frame decoders, enabling inference that scales linearly with the number of points to be reconstructed. Crucially, we demonstrate that D4RT achieves state-of-the-art results across a wide range of 4D tasks including depth, point cloud estimation, and 3D point tracking. D4RT demonstrates that scaling to complex, dynamic environments does not require sacrificing precision, offering a unified framework for the next generation of 4D perception.

Contribution and Acknowledgements. MS led the project, with management support from JZ. MS proposed the SRT-style decoder, and GL proposed local RGB patches and tracking-all-pixels. CZ, GL, IR, and MS contributed to the core model design. LM, SS, IR, and SK designed and created training datasets. CZ, GL, and MS carried out the major implementation, with significant contributions from SK, IR, LM, and JX. CZ drove the model experimentation. Comprehensive evaluations and data pipelines were set up by CZ, GL, SK, IR, LM, JX, SS, and MS. Visualizations were produced by GL, IR and JX. RS, JB, RH, ZG, and AZ

provided project support, advising, and guidance. We thank a number of colleagues and advisors for making this work possible. We thank Saurabh Saxena, Kaiming He, Carl Doersch, Leonidas Guibas, Noah Snavely, Ben Poole, Joao Carreira, Pauline Luc, Yi Yang, Howard Huang, Huizhong Chen, Cordelia Schmid for providing advice during the project; Gabriel Brostow, for advising SK during the course of the project and for providing feedback on the manuscript; Relja Arandelović and Maks Ovsjanikov for providing feedback on the manuscript; Ross Goroshin, Tengda Han and Dilara Gokay for their help during early-stage development; Aravindh Mahendran for helping with code reviews; Daniel Duckworth for helping with visualizations and comparisons against baselines; and Alberto García and Jesús Pérez for the invaluable contributions to data generation and collection.

References

- [1] Abhishek Badki, Hang Su, Bowen Wen, and Orazio Gallo. L4P: Towards unified low-level 4D vision perception. In *3DV*, 2026. 5
- [2] Daniel J. Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 7
- [3] Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Easi3r: Estimating disentangled motion from dust3r without training. *ICCV*, 2025. 4
- [4] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 5
- [5] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR*, 2017. 7
- [6] Kai Deng, Zexin Ti, Jiawei Xu, Jian Yang, and Jin Xie. VGGT-Long: Chunk it, loop it, align it – pushing VGGT’s limits on kilometer-scale long RGB sequences. *arXiv preprint*, 2025. 13
- [7] 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. 4
- [8] 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. 4
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3, 4
- [10] David Eigen, Christian Puhersch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, 2014. 7
- [11] 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. *ICCV*, 2025. 4, 5, 6
- [12] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and trends® in Computer Graphics and Vision*, 2015. 4
- [13] A Gaidon, Q Wang, Y Cabon, and E Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 2016. 5
- [14] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *IJRR*, 2013. 7, 13
- [15] Klaus Greff, Etienne Pot, and Mehdi S. M. Sajjadi. Kauldron: A neural network training framework, optimized for research velocity and modularity. <https://github.com/google-research/kauldron>. 3
- [16] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanapragasam, Florian Golemo, Charles Herrmann, et al. Kubric: A scalable dataset generator. In *CVPR*, 2022. 5
- [17] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *ECCV*, 2022. 4
- [18] Adam W. Harley, Yang You, Xinglong Sun, Yang Zheng, Nikhil Raghuraman, Yunqi Gu, Sheldon Liang, Wen-Hsuan Chu, Achal Dave, Pavel Tokmakov, Suyu You, Rares Ambrus, Katerina Fragkiadaki, and Leonidas J. Guibas. All-Tracker: Efficient dense point tracking at high resolution. In *ICCV*, 2025. 4
- [19] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 4
- [20] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *CVPR*, 2018. 5
- [21] Zeren Jiang, Chuanxia Zheng, Iro Laina, Diane Larlus, and Andrea Vedaldi. Geo4d: Leveraging video generators for geometric 4d scene reconstruction. In *ICCV*, 2025. 7
- [22] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Dynamicstereo: Consistent dynamic depth from stereo videos. In *CVPR*, 2023. 5
- [23] Nikita Karaev, Yuri Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker3: Simpler and better point tracking by pseudo-labelling real videos. In *ICCV*, 2025. 4, 6
- [24] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zauss, Ethan Weber, Nelson Antunes, et al. Mapanything: Universal feed-forward metric 3d reconstruction. *arXiv preprint*, 2025. 7
- [25] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017. 4
- [26] Diederik P Kingma. Adam: A method for stochastic optimization. *ICLR*, 2015. 5

- [27] Skanda Koppula, Ignacio Rocco, Yi Yang, Joe Heyward, Joao Carreira, Andrew Zisserman, Gabriel Brostow, and Carl Doersch. Tapvid-3d: A benchmark for tracking any point in 3d. In *NeurIPS*, 2024. 5, 6
- [28] Guillaume Le Moing, Jean Ponce, and Cordelia Schmid. Dense optical tracking: Connecting the dots. In *CVPR*, 2024. 4
- [29] Zhengqi Li, Richard Tucker, Forrester Cole, Qianqian Wang, Linyi Jin, Vickie Ye, Angjoo Kanazawa, Aleksander Holynski, and Noah Snavely. MegaSaM: Accurate, fast and robust structure and motion from casual dynamic videos. In *CVPR*, 2025. 1, 4, 5, 7, 17
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *ICLR*, 2019. 5
- [31] Jiahao Lu, Tianyu Huang, Peng Li, Zhiyang Dou, Cheng Lin, Zhiming Cui, Zhen Dong, Sai-Kit Yeung, Wenping Wang, and Yuan Liu. Align3r: Aligned monocular depth estimation for dynamic videos. In *CVPR*, 2025. 4
- [32] 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. *ICLR*, 2025. 4, 5, 6
- [33] John Oliensis. A critique of structure-from-motion algorithms. *CVIU*, 2000. 4
- [34] Emanuele Palazzolo, Jens Behley, Philipp Lottes, Philippe Giguère, and Cyrill Stachniss. Refusion: 3d reconstruction in dynamic environments for rgb-d cameras exploiting residuals. In *IROS*, 2019. 7
- [35] Duc-Hai Pham, Tung Do, Phong Nguyen, Binh-Son Hua, Khoi Nguyen, and Rang Nguyen. SharpDepth: Sharpening metric depth predictions using diffusion distillation. In *CVPR*, 2025. 14
- [36] Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Matia Segu, Siyuan Li, Wim Abbeloos, and Luc Van Gool. UniDepthV2: Universal Monocular Metric Depth Estimation Made Simpler. *PAMI*, 2025. 6
- [37] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ICCV*, 2021. 8
- [38] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *ICCV*, 2021. 5
- [39] Mehdi SM Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lučić, Daniel Duckworth, Alexey Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *CVPR*, 2022. 2
- [40] Mehdi S. M. Sajjadi, Aravindh Mahendran, Thomas Kipf, Etienne Pot, Daniel Duckworth, Mario Lučić, and Klaus Greff. RUST: Latent Neural Scene Representations from Unposed Imagery. *CVPR*, 2023. 2
- [41] Peter Sand and Seth Teller. Particle video: Long-range motion estimation using point trajectories. *IJCV*, 2008. 4
- [42] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 4
- [43] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, 2016. 4
- [44] 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. 4, 5
- [45] Pei Sun, Henrik Kretzschmar, 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 *CVPR*, 2020. 5
- [46] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *TPAMI*, 1991. 3, 13
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 3, 4
- [48] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *CVPR*, 2025. 1, 3, 4, 6, 7
- [49] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. VideoMAE v2: Scaling video masked autoencoders with dual masking. In *CVPR*, 2023. 15
- [50] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *CVPR*, 2025. 4, 7
- [51] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *CVPR*, 2024. 3, 4, 5
- [52] Shizun Wang, Zhenxiang Jiang, Xingyi Yang, and Xinchao Wang. C4d: 4d made from 3d through dual correspondences. In *ICCV*, 2025. 4
- [53] Wenshan Wang, Delong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. TartanAir: A dataset to push the limits of visual slam. In *IROS*, 2020. 5
- [54] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. π^3 : Permutation-equivariant visual geometry learning. *arXiv preprint*, 2025. 3, 4, 5, 7, 17
- [55] 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. 5
- [56] Yuxi Xiao, Jianyuan Wang, Nan Xue, Nikita Karaev, Yuri Makarov, Bingyi Kang, Xing Zhu, Hujun Bao, Yujun Shen, and Xiaowei Zhou. SpatialTrackerV2: Advancing 3d point tracking with explicit camera motion. In *ICCV*, 2025. 1, 4, 5, 6, 7, 17
- [57] Jianing Yang, Alexander Sax, Kevin J Liang, Mikael Henaff, Hao Tang, Ang Cao, Joyce Chai, Franziska Meier, and Matt Feiszli. Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. In *CVPR*, 2025. 4

- [58] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *CVPR*, 2020. [5](#)
- [59] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *ICCV*, 2023. [5](#)
- [60] 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. *ICLR*, 2025. [4](#), [7](#)
- [61] Zhengyou Zhang. A flexible new technique for camera calibration. *TPAMI*, 2002. [3](#)
- [62] Yang Zheng, Adam W. Harley, Bokui Shen, Gordon Wetstein, and Leonidas J. Guibas. Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In *ICCV*, 2023. [5](#)
- [63] Artem Zholus, Carl Doersch, Yi Yang, Skanda Koppula, Viorica Patraucean, Xu Owen He, Ignacio Rocco, Mehdi SM Sajjadi, Sarath Chandar, and Ross Goroshin. TAPNext: Tracking any point as next token prediction. In *ICCV*, 2025. [4](#)

Efficiently Reconstructing Dynamic Scenes One D4RT at a Time

Appendix

A. Model Overview & Training Details

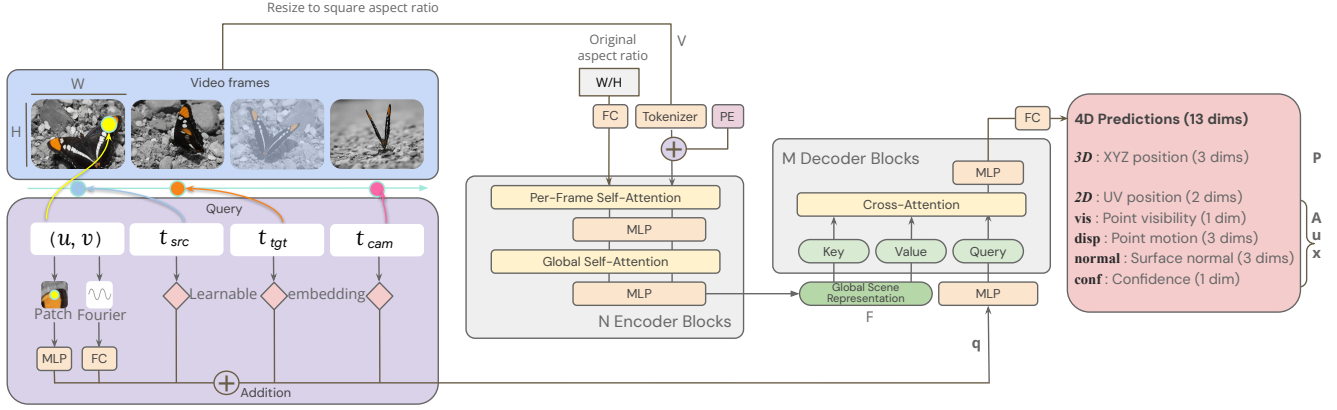


Figure 7. **Full D4RT model overview** – We provide a holistic overview of the model together with its inputs and outputs. FC corresponds to a fully connected layer, and PE for positional encoding. See Sec. 2 of the main paper for reference.

The model is trained end-to-end by minimizing a composite loss \mathcal{L} , which is a weighted sum of task-specific losses computed over a batch of N sampled queries:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left(c \lambda_{3D} \mathcal{L}_{3D} - \lambda_{\text{conf}} \log c + \lambda_{2D} \mathcal{L}_{2D} + \lambda_{\text{vis}} \mathcal{L}_{\text{vis}} + \lambda_{\text{disp}} \mathcal{L}_{\text{disp}} + \lambda_{\text{conf}} \mathcal{L}_{\text{conf}} + \lambda_{\text{normal}} \mathcal{L}_{\text{normal}} \right)_i$$

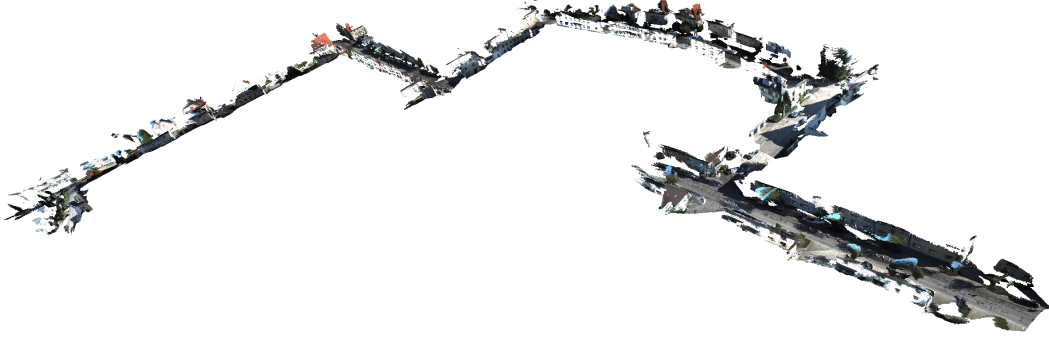
where c is the confidence score predicted by the model. We use the following loss weights: $\lambda_{3D}=1.0$, $\lambda_{2D}=0.1$, $\lambda_{\text{vis}}=0.1$, $\lambda_{\text{disp}}=0.1$, $\lambda_{\text{conf}}=0.2$, $\lambda_{\text{normal}}=0.5$, $\lambda_{\text{conf}}=0.2$. We train using the AdamW optimizer with a weight decay of 0.03. The learning rate is warmed up for 2,500 steps until it reaches a peak value of 10^{-4} . Subsequently, it follows a cosine annealing schedule, decaying to a final value of 10^{-6} . Gradients are clipped to a maximum L^2 -norm of 10.

Data augmentation. Extensive data augmentation techniques are applied to the video during training to improve model generalization. We apply temporally consistent color jittering by performing random brightness, saturation, contrast, and hue adjustments. We also apply random color drop with a probability of 0.2, and Gaussian blur augmentation with a probability of 0.4. For spatial augmentation, we use random crop augmentations with a scale ratio between 0.3 and 1.0 of the original size. After determining the crop size, a random aspect ratio is selected by sampling from a uniform distribution in the logarithmic domain; this ensures equal probability for wide and tall crops while respecting image boundaries. Additionally, during the random crop augmentation, the image is randomly zoomed in with a probability of 0.05. On the temporal dimension, frames are subsampled from the full video with a random stride.

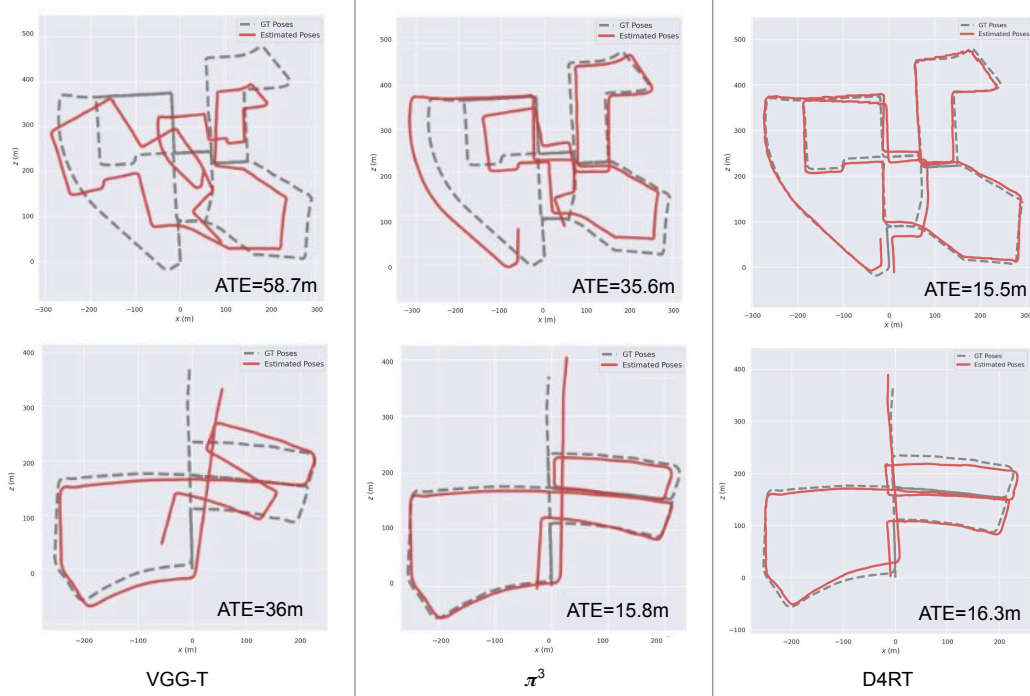
Training queries. Queries are sampled from available ground truth point trajectories. To focus the model on challenging regions, 30% of the queries (along dimensions u, v) are sampled near depth discontinuities or motion boundaries, which are pre-computed using a Sobel filter on depth maps. Timesteps t_{src} , t_{tgt} , and t_{cam} are sampled uniformly at random, except that we enforce $t_{\text{tgt}} = t_{\text{cam}}$ with probability 0.4 to improve downstream performance.

B. Generalization to Long Videos

We implement a long-sequence processing algorithm by partitioning videos into overlapping segments. We then align these segments by estimating Sim(3) transformations using the Umeyama algorithm [46], based on the top 85% of points with the highest confidence in the overlapping regions. The process is similar to the first stage proposed in VGGT-Long [6], we omit loop detection and optimization stage in [6] to directly evaluate the reconstruction model’s raw precision. As shown in Fig. 8 on KITTI [14], our model yields the best ATE in the first example, significantly outperforming VGG-T and π^3 by a large margin. On the second example, we significantly outperform VGG-T and produce a result on par with π^3 .



(a) Visualization of reconstruction of 1000 frames from KITTI sequence 00.



(b) Comparison of raw chunked prediction alignment results against VGG-T and π^3 baselines (no loop closure).

Figure 8. Long sequence results on KITTI sequences – We align and stitch predictions from overlapped chunks by Umeyama algorithm, without the loop detection and global optimization proposed in [6]. Our model obtains consistently better results than VGG-T, and significantly better results than π^3 on sequence “00” (top row).

C. High-Resolution Decoding with Subpixel Precision

A key advantage of D4RT’s architecture is the decoupling of the global scene encoding from the point-wise decoding. As the query coordinates (u, v) are defined in a continuous normalized space $[0, 1]^2$, our decoder is able to probe the scene at arbitrary resolutions, independent of the resolution of the Global Scene Representation F .

We explore this capability in Tab. 10. To quantify the preservation of high-frequency details, we report the Pseudo Depth Boundary Error accuracy ($\epsilon_{\text{PDBE}}^{\text{acc}}$) proposed by Pham et al. [35], in addition to standard depth metrics. We compare four configurations using a ViT-g encoder fixed at a resolution of 256×256 . The baseline output at the encoder’s native resolution (**Config ①**) is naturally coarse. Incorporating the local appearance patch mechanism described in the main paper (**Config ②**) allows for the recovery of finer low-level structures, although pixelation artifacts persist. We further leverage the continuous nature of our query mechanism to predict at the *original* resolution. While naive dense querying (**Config ③**) recovers smoother edges, it still fails at recovering high frequencies.

In **Config ④**, we extract the local RGB patches from the source frames at their *original* resolution for decoding.

The significant drop in $\epsilon_{\text{PDBE}}^{\text{acc}}$ demonstrates how this allows D4RT to recover finer details. We show qualitative results in Fig. 9 where we observe that hair strands and object boundaries are resolved more accurately.

Config.	Encoder Resolution	RGB Patch	Output Resolution	RGB patch Resolution	Scale		Scale and Shift	
					AbsRel ↓	$\epsilon_{\text{PDBE}}^{\text{acc}}$ ↓	AbsRel ↓	$\epsilon_{\text{PDBE}}^{\text{acc}}$ ↓
①	256×256	✗	256×256	256×256	0.254	3.323	0.219	3.307
②	256×256	✓	256×256	256×256	0.218	2.254	0.179	2.243
③	256×256	✓	<i>Original</i>	256×256	0.217	2.266	0.178	2.258
④	256×256	✓	<i>Original</i>	<i>Original</i>	0.220	2.193	0.176	2.185

Table 10. **Quantitative impact of query density and patch fidelity** – Feeding RGB patches from the high-resolution video into the decoder (**Config ④**) yields significantly sharper edges in depth maps as measured by $\epsilon_{\text{PDBE}}^{\text{acc}}$.

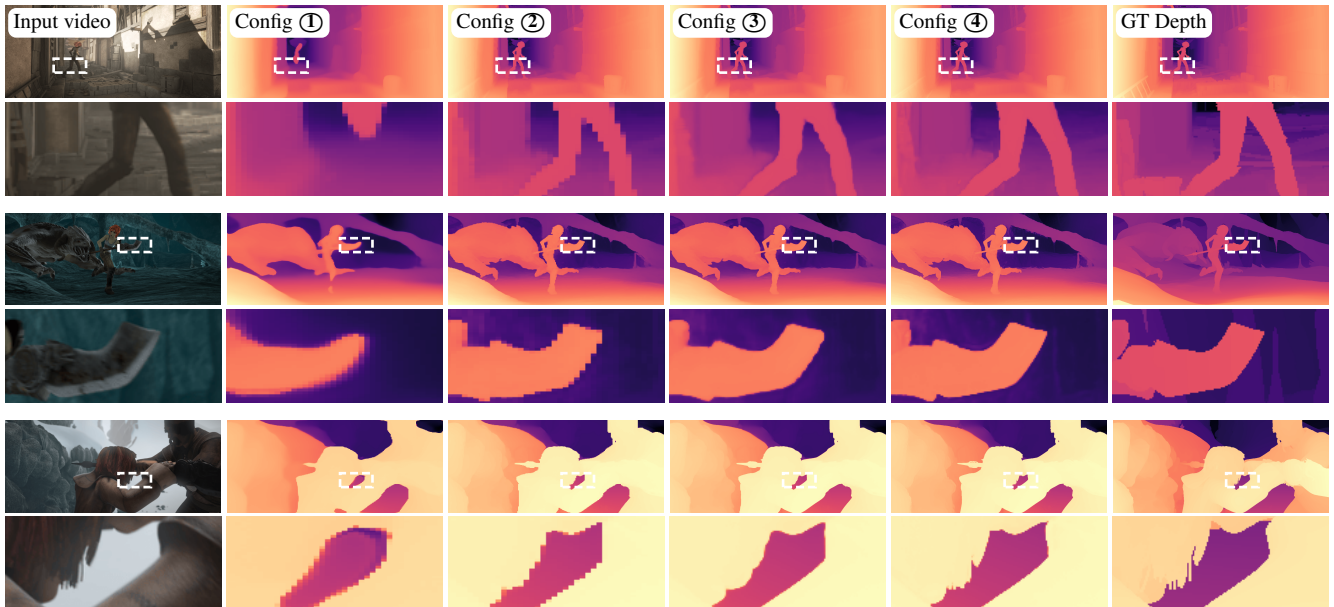


Figure 9. **Visualizing sub-pixel detail recovery** – We propose a visual comparison of the different high-res configurations. Config ④ achieves the highest fidelity, it preserves sharp edges and recovers fine details—such as the hair in the bottom row—without increasing the computational cost or memory requirements of the overall model.

D. Further Ablations

Pretrained Encoder. In Tab. 11, we show our model performance when the video encoder is initialized with random weights compared to using pre-trained VideoMAE [49] weights. We observe significant improvements across the board.

Model weight initialization	Video Depth Estimation		Camera Pose Estimation		
	AbsRel (S) ↓	AbsRel (SS) ↓	ATE ↓	RPE-T ↓	RPE-R ↓
None	0.738	0.520	0.334	0.139	1.126
VideoMAE [49]	0.302	0.257	0.091	0.028	0.245

Table 11. **Model initialization** – Initializing the model with VideoMAE [49] weights leads to significant improvements.

Local RGB patch size. We perform an ablation study on the size of the local RGB patch. As shown in Fig. 10, our results indicate that the patch size 9×9 yields the best overall performance across camera pose and depth estimation tasks.

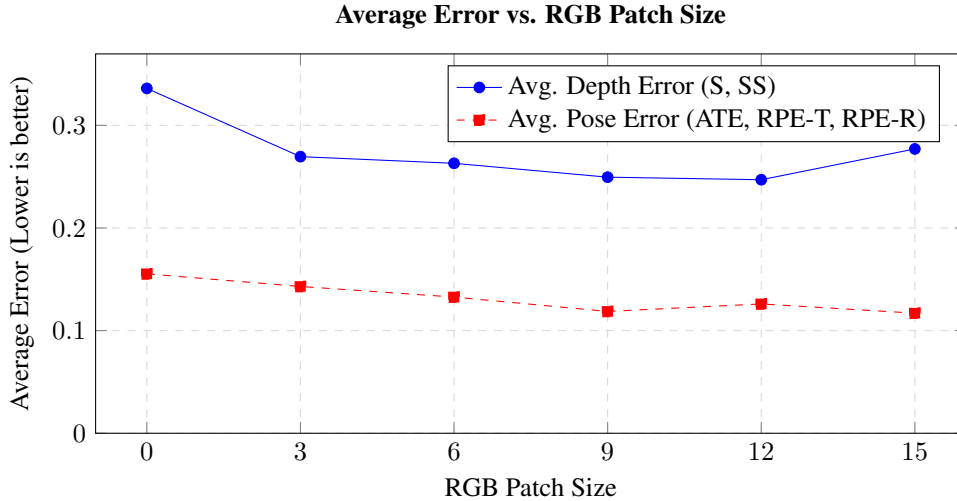


Figure 10. **Ablation study on RGB patch size** – The plot shows the average error for depth and pose estimation on Sintel. A patch size between 9 and 12 yields the best performance for both tasks.

E. Further Qualitative Results

Complementing the qualitative visualizations in the main text, we provide additional examples of our reconstruction results in Fig. 11, along with further baseline comparisons in Fig. 12 and Fig. 13.

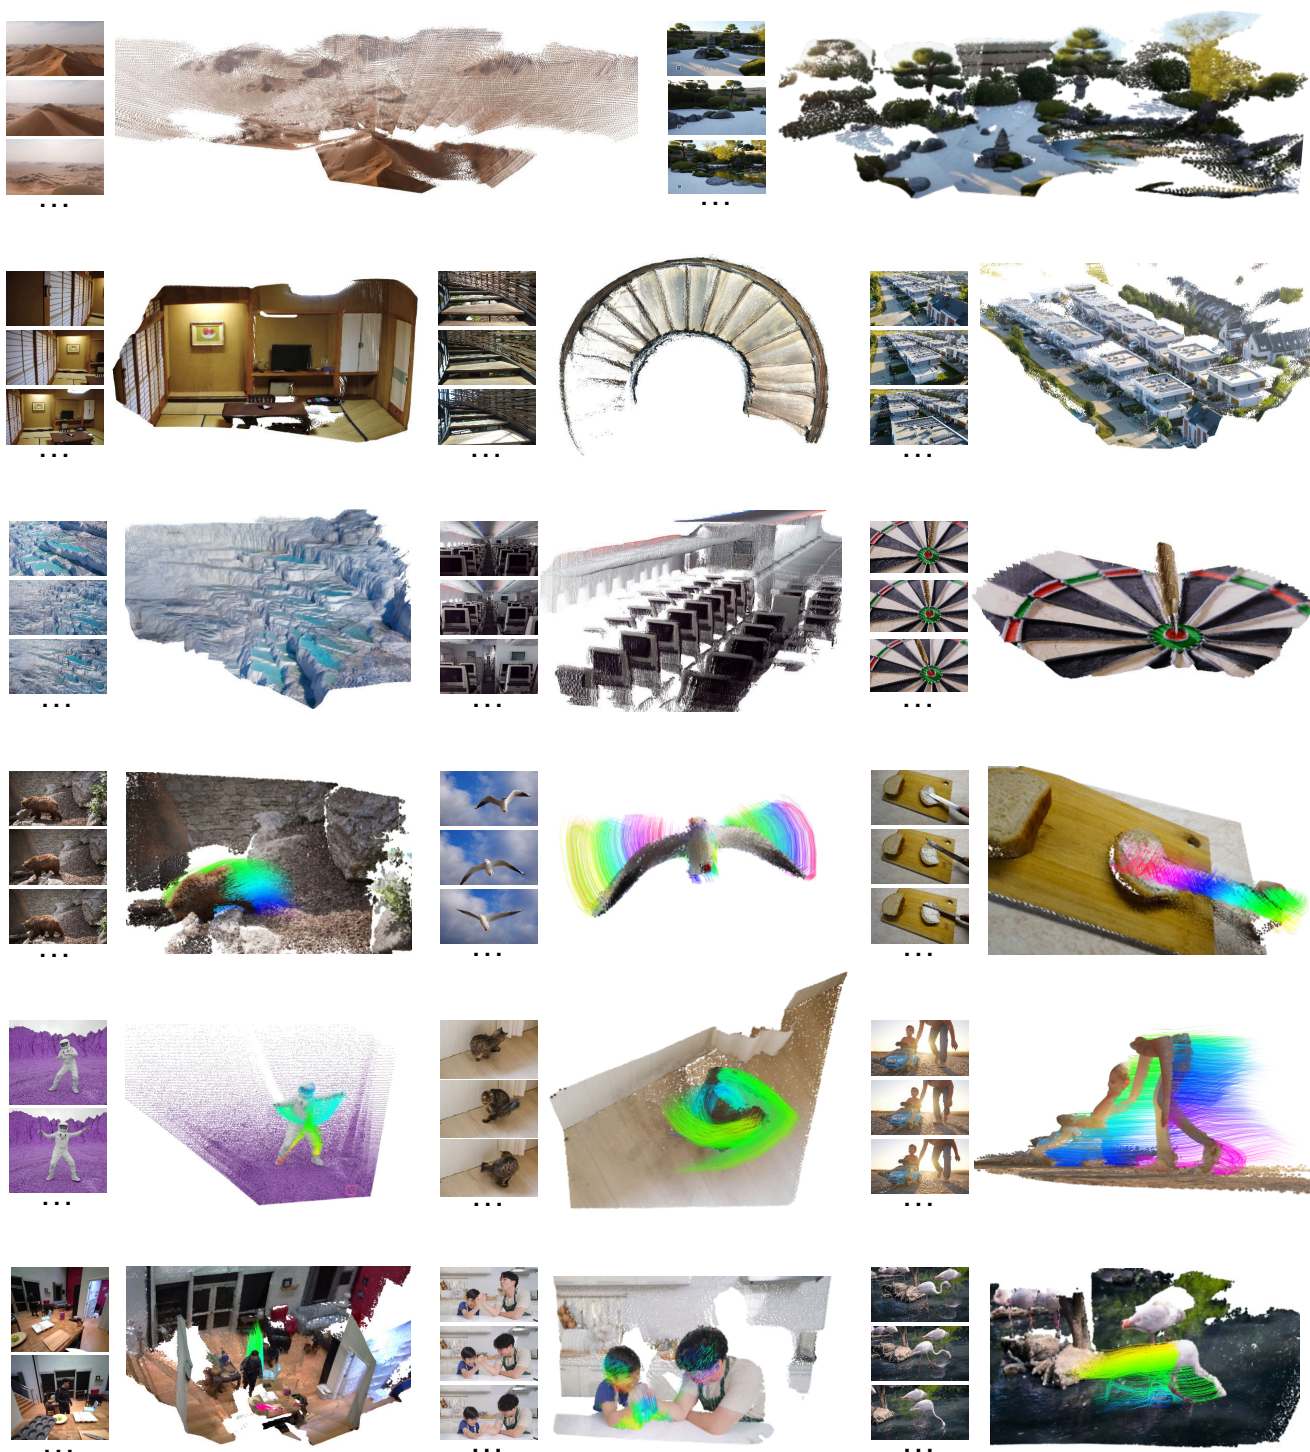


Figure 11. **Additional visualizations of D4RT** – D4RT produces accurate reconstructions for both static environments (top three rows) and dynamic sequences (bottom three rows).

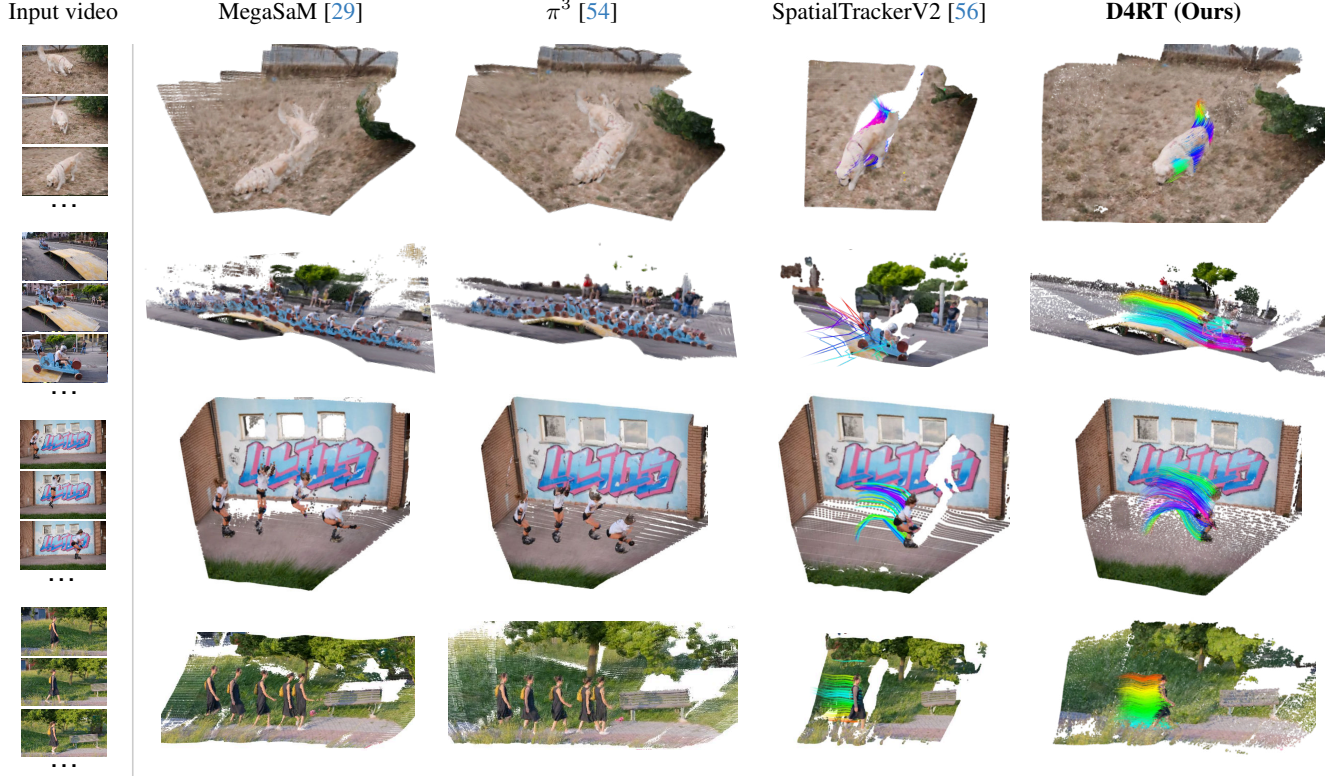


Figure 12. **Additional reconstruction results across methods** – Pure reconstruction methods (MegaSaM and π^3) are visualized as accumulated point clouds. For SpatialTrackerV2, we visualize sparse tracks on a representative frame. In contrast, D4RT reconstructs a complete 4D scene representation, tracking *all* pixels across the entire video.

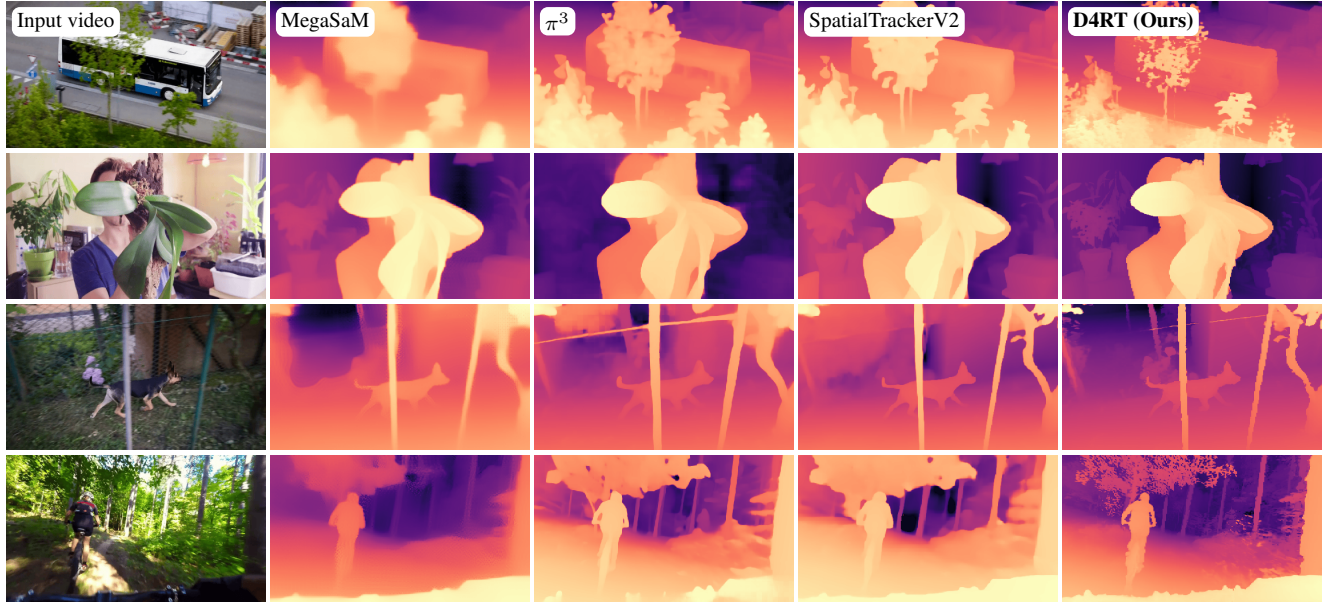


Figure 13. **Qualitative depth comparison across methods** – D4RT is able to perform dense depth estimation with finer details than current state-of-the-art methods, preserving geometric accuracy even in scenarios with large motion blur.