arXiv:2610.42136v1 [cs.CV] 30 Oct 2026

RELIFT: DENSE 3D TRACKING AS LIFTED TRAJECTORY REPAIR

Assaf Singer[1*] **Maroun Khriesh**[1*] **Noam Rotstein**[1,2] **Yoni Kasten**[1,2] **Or Litany**[1,2]
[1] Technion – Israel Institute of Technology [2] NVIDIA
[*] Equal contribution

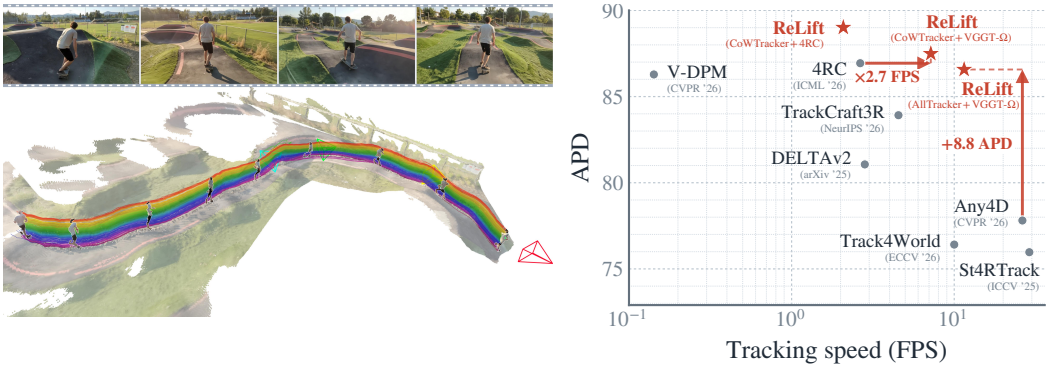


Figure 1: ReLift combines pretrained 2D tracking and 3D geometry with learned trajectory repair and completion, achieving state-of-the-art dense 3D tracking accuracy on WorldTrack and a strong speed–accuracy trade-off. **Left:** Long-range 3D trajectories recovered from a real-world video. **Right:** WorldTrack accuracy versus full-pipeline FPS (see Appendix D.4 for details).

ABSTRACT

Dense 3D tracking from monocular video requires establishing point correspondences over time, estimating their depth, and inferring their trajectories during occlusion. Existing methods typically learn these tasks jointly, a demanding objective given the scarcity of 4D supervision. We show that strong pretrained 2D trackers and 3D geometry models provide an effective foundation: lifted trajectories, with occluded segments simply interpolated, already rival dedicated 3D trackers using the same geometry. However, lifted trajectories inherit correspondence and depth errors, and under occlusion, their depth comes from the occluder rather than the tracked point. We introduce ReLift, an efficient feed-forward framework that recasts dense 3D tracking as geometric trajectory repair and completion. At its core, a trajectory corrector alternates attention along and across trajectories to preserve reliable observations, repair corrupted positions, and complete missing motion. Using only geometric inputs allows us to train directly on 3D trajectories with simulated lifting errors, decoupling training from images and from any specific tracker or geometry model. Concentrating 4D supervision on trajectory refinement yields a reusable model that consistently improves lifted trajectories across trackers and geometry providers without retraining, so it benefits directly from future advances in either component. On WorldTrack, ReLift achieves state-of-the-art mean accuracy, improving on the strongest baselines by 2.1 APD on all points and 3.0 on dynamic points, with faster configurations remaining competitive with leading 3D trackers while running 2.7–4.4× faster than 4RC. Visit our project page for more examples and code.

1 INTRODUCTION

Dense 3D point trajectories describe how scene content moves through space over time, and are key in downstream tasks such as manipulation policy learning, dynamic scene reconstruction, sensor simulation, and motion-controlled video generation (Chen et al., 2025; Bharadhwaj et al., 2024; Lu et al., 2024; Geng et al., 2025; Borreda et al., 2026; Singer et al., 2026). Recovering them from

monocular video is usually posed as a single learning problem: one feed-forward network is trained end to end to output 3D tracks.

Such a network has to learn three things at once: finding correspondence across frames, recovering the geometry of the scene, and inferring motion behind occlusions. However, pretrained models already excel at two of these tasks. Dedicated 2D trackers establish strong correspondences (Harley et al., 2025; Lai et al., 2026) and can exceed the image-plane accuracy of projected 3D predictions (Appendix C.4), while geometry models provide depth and cameras learned from scalable geometric supervision without requiring annotated dynamic 4D trajectories (Lin et al., 2025; Wang et al., 2026). This motivates focusing the scarce 4D supervision on what remains: repairing errors in the combined estimates and completing occluded motion.

We therefore introduce ReLift, which combines strong 2D correspondences and pretrained scene geometry for dense 3D tracking. We take the viewing ray of each point from a 2D tracker and its range along that ray from a depth model, and lift the two into a 3D trajectory we call the lift. Since occluded tracks sample the occluder's depth, we discard predictions flagged as occluded and initialize missing positions by interpolation. Even before learned correction, these initialized lifts can match or exceed the mean accuracy of learned 3D trackers using the same geometry (Tables 1 and 2). However, these lifts still lack observations during occlusion and suffer from jumps at depth boundaries and jitter from per-frame depth fluctuations. These failures can produce severe trajectory discontinuities and leave hidden motion unresolved. We recast dense 3D tracking as geometric trajectory repair and completion. A compact learned corrector preserves reliable observations while repairing corrupted positions and recovering missing motion. Alternating attention along and across tracks shares temporal and spatial evidence, while processing track groups independently keeps dense feed-forward correction efficient.

A natural training strategy is to run pretrained tracking and geometry models on synthetic videos with ground-truth 3D trajectories, then train the corrector to recover those trajectories from the resulting lifts. However, this ties the training inputs to the errors of the chosen estimators. We avoid this by training without ever running a tracker, a depth model, or a renderer. We model the geometric effects of lifting failures directly as randomized corruptions of clean 3D trajectories from PointOdyssey (Zheng et al., 2023), producing *simulated lifts* from which we train a network to recover the clean trajectories. Training uses geometric annotations without images or provider predictions, concentrating supervision on learning motion priors for trajectory repair and completion (Sec. 3.4). Our ablation supports this choice: training on simulated lifts improves accuracy over training on lifts generated by AllTracker and Depth Anything 3 (Sec. 4.4).

Training on simulated lifts supports transfer across providers. We evaluate the same ReLift corrector without retraining with AllTracker (Harley et al., 2025) or CoWTracker (Lai et al., 2026), each paired with depth and camera estimates from Depth Anything 3 (Lin et al., 2025), VGGT-Ω (Wang et al., 2026), or 4RC (Luo et al., 2026). Correction improves all six lifts (Fig. 3). Moreover, using the same depth and camera estimates as 4RC, our 2D tracking and correction pipeline achieves higher mean tracking accuracy than 4RC's trajectory predictions (Table 1).

Our contributions can be summarized as follows:

- **A strong lifting baseline.** We show that pretrained 2D tracks lifted with estimated geometry and simple gap initialization can match or exceed dedicated 3D trackers using the same geometry.
- **ReLift,** an efficient feed-forward framework for dense monocular 3D tracking that combines pretrained 2D motion and 3D geometry through learned trajectory repair and completion.
- A provider-agnostic training strategy based on randomized geometric corruptions of synthetic trajectories, without running tracking or geometry models during training. The same weights transfer across evaluated providers without retraining, allowing the framework to incorporate future advances in either component.
- Our framework achieves state-of-the-art mean dense 3D tracking accuracy on WorldTrack, with faster configurations retaining competitive accuracy at substantially lower runtime than comparably accurate 3D trackers.

2 Related Work

Dynamic 3D reconstruction from monocular video. Classical non-rigid structure-from-motion methods recover time-varying 3D structure from image correspondences using local rigidity and low-rank constraints to reduce monocular ambiguity (Bregler et al., 2000; Akhter et al., 2008; Taylor et al., 2010; Dai et al., 2014). Modern per-video optimization methods combine image evidence with depth, tracking, and deformation priors (Zhang et al., 2022; Liu et al., 2023). Shape of Motion (Wang et al., 2025b) and MoSca (Lei et al., 2024) combine pretrained depth and tracking cues with structured motion representations for reconstruction and rendering. These cues provide the initialization that rendering-based optimization needs, since photometric gradients vanish when rendered and observed content do not overlap (Cohen Rimon et al., 2026). Uni4D (Yao et al., 2025) and TrackingWorld (Lu et al., 2025) optimize camera poses and dense 3D trajectories from pretrained depth, tracking and segmentation models, or from densified 2D tracks. These methods demonstrate the strength of combining 2D motion with 3D geometry, but require lengthy per-scene optimization; we pursue this combination with a feed-forward model.

Feed-forward 3D geometry models. DUSt3R and MASt3R predict geometry and correspondences from image pairs (Wang et al., 2024; Leroy et al., 2024), while VGGT, VGGT-Ω, and Depth Anything 3 extend geometry estimation to multiple views (Wang et al., 2025a; 2026; Lin et al., 2025). MonST3R adapts point-map prediction to dynamic scenes and aligns predictions across a video (Zhang et al., 2025b); St4RTrack and Dynamic Point Maps also represent motion and correspondences (Feng et al., 2025; Sucar et al., 2025), with V-DPM extending pairwise dynamic point maps to video (Sucar et al., 2026). Geometry and tracking capabilities thus overlap. In our pipeline, frozen geometry models supply depth and cameras, while 2D trackers establish point identities across frames.

2D tracking. 2D tracking follows the image projections of scene points over time. CoTracker jointly tracks multiple points (Karaev et al., 2024), while AllTracker and CoWTracker extend tracking to every reference-frame pixel (Harley et al., 2025; Lai et al., 2026). Per-video self-supervision also yields strong correspondences without ground-truth annotations for the input video: DINO-Tracker adapts pretrained image features (Tumanyan et al., 2024), while OmniMotion optimizes globally consistent tracks (Wang et al., 2023). These approaches show that correspondence estimation can exploit real-video evidence without annotated 3D motion, motivating its use as a separate component. Our comparison supports reusing strong 2D correspondences: native CoWTracker surpasses the evaluated 3D methods in image-space accuracy (Appendix C.4).

Established tracks also provide a useful representation for recovering dynamic 3D structure. TracksTo4D reconstructs geometry and motion from 2D tracks using a feed-forward, low-rank formulation without 3D supervision (Kasten et al., 2024). Seurat learns depth changes from 2D trajectories and combines them with pretrained depth estimates (Cho et al., 2025). Building on these precedents, ReLift starts from tracks already lifted with pretrained depth and cameras, and learns 3D residuals to repair corrupted observations and complete missing motion.

Dense feed-forward 3D tracking. Feed-forward 3D trackers estimate trajectories from video, either jointly predicting geometry and motion or conditioning on externally supplied depth. DELTA and DELTAv2 predict dense long-range 3D tracks using video and depth inputs (Ngo et al., 2025b;a), while SpatialTrackerV2 jointly estimates geometry, camera motion, and point motion (Xiao et al., 2025). Trace Anything, Track4World, and 4RC predict dense trajectories or queryable geometry and motion (Liu et al., 2026; Lu et al., 2026; Luo et al., 2026). TrackCraft3R repurposes pretrained video-diffusion models for dense 3D tracking (Nam et al., 2026); MotionCrafter uses video generative priors for geometry and scene flow (Zhu et al., 2026). Available 3D motion supervision includes complete synthetic trajectories from PointOdyssey (Zheng et al., 2023) and trajectories reconstructed from real stereo videos in Stereo4D (Jin et al., 2025). We use PointOdyssey to learn repair and completion of lifted trajectories. Under matched geometry, these lifts already rival dedicated 3D trackers, and correction further improves them (Sec. 4.2).

3 Method

Given a monocular video, we recover dense 3D trajectories in the camera coordinates of the reference frame I_0. Frozen pretrained models supply 2D tracks, depth and camera estimates, which we combine into lifted trajectories (Sec. 3.1). The lift is accurate where tracking and depth are reliable.

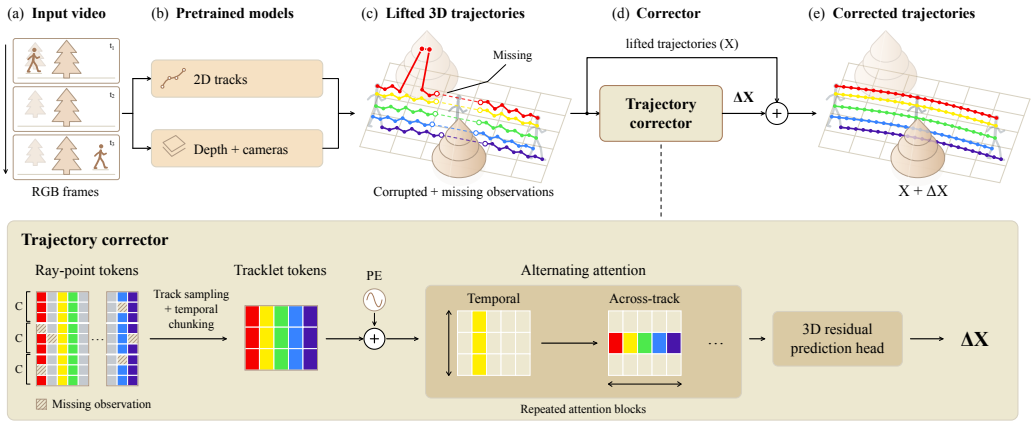


Figure 2: **Inference pipeline of ReLift.** Frozen pretrained models lift video observations (a–b) into corrupted and incomplete 3D trajectories (c). The corrector adds residuals to initialized XYZ coordinates to produce the final trajectories (d–e). Below: ray–point tokens $\boldsymbol{\phi}_{ti}$ are grouped into C-frame tracklets and embedded as tokens $\mathbf{h}_{ni}$, then processed by alternating temporal and across-track attention. The output head predicts frame-level XYZ residuals, which are added to the initialized trajectories. We train the corrector on simulated lifts from clean trajectory annotations (Sec. 3.4).

We use predicted visibility to mask observations, leaving gaps to complete. Retained observations may still be corrupted by missed occlusions, depth errors or tracking drift, and therefore require repair.

Our corrector model uses positions, viewing rays and availability masks (Sec. 3.2). We hypothesize that temporal inconsistencies and disagreements between tracks provide cues for repair, while observations before and after gaps, together with related motion, can guide completion. We therefore use alternating attention along and across tracks to aggregate these cues and predict 3D residual corrections (Sec. 3.3). Temporal chunking and batched track groups enable efficient feed-forward correction over all reference pixels (Fig. 2).

Because the corrector uses only geometry, we train it on simulated lifts constructed directly from clean 3D trajectory annotations, without rendering images or running providers (Sec. 3.4).

3.1 Lifting Tracks into 3D

For a video $\{I_t\}_{t=0}^{T-1}$, a dense 2D tracker provides image trajectories $\mathbf{u}_{ti}$ and visibility scores $v_{ti} \in [0,1]$ for pixels i of the reference frame I_0. A geometry model provides optical-axis depth maps Z_t, intrinsics K_t and camera poses. We express geometry in the reference camera coordinate system: $R_t \in \mathrm{SO}(3)$ rotates from camera t to camera 0, and $\mathbf{o}_t$ is camera t's center in that coordinate system, with $R_0 = I$ and $\mathbf{o}_0 = \mathbf{0}$. For homogeneous image coordinates $\tilde{\mathbf{u}}_{ti} = (\mathbf{u}_{ti}, 1)$, let $\mathbf{q}_{ti} = K_t^{-1}\tilde{\mathbf{u}}_{ti}$. The unit viewing direction, range and lifted position are

$$\mathbf{d}_{ti} = \frac{R_t\,\mathbf{q}_{ti}}{\|\mathbf{q}_{ti}\|}, \qquad r_{ti} = Z_t(\mathbf{u}_{ti})\,\|\mathbf{q}_{ti}\|, \qquad \mathbf{y}_{ti} = \mathbf{o}_t + r_{ti}\,\mathbf{d}_{ti}. \tag{1}$$

Since $\mathbf{q}_{ti}$ has third coordinate 1, its norm converts optical-axis depth into distance along the ray.

We distinguish four common failure patterns in lifted trajectories. *Correspondence errors* introduced by the 2D tracker perturb the viewing direction. *Depth jitter and drift* displace points along the viewing ray. Jitter produces rapid frame-to-frame fluctuations in depth error. Drift is a gradual change in the difference between estimated and true depth over successive frames. *Wrong-surface observations* arise when the sampled depth belongs to another surface, for example near a depth boundary or when a track follows an occluder. *Missing observations* can arise when a point is occluded or the 2D tracker loses track of it.

We use the tracker's predicted visibility to determine observation availability. We retain a lifted position and its ray direction when its range is positive and the tracker's visibility score exceeds a tracker-specific threshold γ; otherwise, we mask both. The binary mask m_{ti} records this choice,

with $m_{0i} = 1$ for valid reference points. Correctly detected occlusions therefore produce gaps, while missed occlusions can leave observations on the occluder's surface. Observations before and after a gap retain the same track index. Preprocessing details are given in Appendix A.1, and Appendix D.2 evaluates the choice to mask ray directions.

3.2 Trajectory Representation

Normalization and initialization. We normalize positions and camera centers by the median reference depth s and initialize each trajectory with the normalized lift, $\bar{\mathbf{x}}_{ti} = \mathbf{y}_{ti}/s$, at available observations. Missing segments are filled by linear interpolation between observations; after the last observation, we hold its position.

Ray–point tokenization. For each track and frame, we construct a token combining initialized displacement, retained observations and viewing geometry. Let $\mathbf{a}_i = \mathbf{y}_{0i}/s$ be the track's normalized initial 3D position (the reference anchor), and g_{ti} the distance in frames to the nearest available observation. The token is

$$\boldsymbol{\phi}_{ti} = \big(\bar{\mathbf{x}}_{ti} - \mathbf{a}_i,\ m_{ti}\,(\bar{\mathbf{x}}_{ti} - \mathbf{a}_i),\ m_{ti},\ g_{ti},\ m_{ti}\,\mathbf{d}_{ti},\ \mathbf{o}_t/s\big) \in \mathbb{R}^{14}. \tag{2}$$

The first displacement provides an estimate at every frame, including gaps. The second exposes only retained observations: it equals $\bar{\mathbf{x}}_{ti} - \mathbf{a}_i$ when $m_{ti} = 1$ and is zero otherwise. The mask indicates whether an observation is available, while g_{ti} measures the temporal distance to the nearest available observation. The ray identifies the direction of range errors and is masked inside gaps; the camera center supplies its origin.

For efficiency, we concatenate non-overlapping chunks of C consecutive ray–point tokens from each track and project each chunk into a single *tracklet token* for attention. This gives $N = \lceil T/C \rceil$ tokens per track, padding the final chunk when necessary. For the frame indices $\mathcal{T}_n$ of chunk n,

$$\mathbf{h}_{ni} = W_{\text{in}}\big[\boldsymbol{\phi}_{ti}\big]_{t\in\mathcal{T}_n} + \mathbf{p}_n + \boldsymbol{\psi}(\mathbf{a}_i) \in \mathbb{R}^D. \tag{3}$$

Here $W_{\text{in}} \in \mathbb{R}^{D\times 14C}$ projects the concatenated features to width D, and $\mathbf{p}_n$ encodes the chunk's temporal position. The anchor embedding $\boldsymbol{\psi}(\mathbf{a}_i)$ encodes the track's initial 3D location, providing spatial context alongside the anchor-relative displacements. Additional details are in Appendix A.2.

3.3 Trajectory Corrector

To limit attention cost, the corrector processes groups of P tracks spanning all T frames. Their tracklet tokens form a grid $H \in \mathbb{R}^{P\times N\times D}$, with one row per track and one column per chunk.

We process H with alternating temporal and across-track self-attention, following the factorization of TracksTo4D (Kasten et al., 2024). Temporal attention connects all N chunks of each track, aggregating evidence across the full sequence along the correspondences supplied by the 2D tracker. Across-track attention connects the P tracks within each temporal chunk, sharing motion information across the group. For fixed feature width, attention costs $O(PN^2 + NP^2)$ per block, compared with $O(P^2N^2)$ for joint attention over all tokens.

A prediction head maps each tracklet token to C XYZ residuals $\boldsymbol{\Delta}_{ti}$, one per frame. Viewing rays condition these predictions, but the residuals are unrestricted in 3D, allowing correction of range and correspondence errors. We add them to the initialized trajectories and restore the original scale:

$$\hat{\mathbf{x}}_{ti} = s\,\big(\bar{\mathbf{x}}_{ti} + \boldsymbol{\Delta}_{ti}\big). \tag{4}$$

We retain the tracker's visibility scores v_{ti} unchanged.

Efficient dense inference. At inference, we randomly partition the reference pixels into groups of P tracks and process the groups independently in batches. All groups use the clip's shared scale s, and their predictions are assembled into the dense trajectory field at the corresponding reference pixels. For fixed sequence length and group size, the corrector's cost grows linearly with the number of pixels. Implementation details and scaling measurements are given in Appendices A.4 and D.4.

3.4 Training from Simulated Lifts

We train the corrector on 3D trajectories from the PointOdyssey training split (Zheng et al., 2023). Each sample is a group of trajectories $\mathbf{x}_{ti}$ over T frames, ground-truth visibility $v^\star_{ti} \in \{0, 1\}$, and

camera centers $\mathbf{o}_t$ in reference camera coordinates. We generate simulated lifts directly from these annotations, without images or provider inference, so lifting errors can be varied independently of any tracker or geometry model.

Simulated lifting errors. We simulate lifting failures caused by imperfect 2D trackers and 3D geometry models (Sec. 3.1) by perturbing viewing ray directions, ranges and observation availability. For each clean point $\mathbf{x}_{ti}$, its range $r^\star_{ti} = \|\mathbf{x}_{ti} - \mathbf{o}_t\|$ is its distance from the camera, and $\mathbf{d}^\star_{ti} = (\mathbf{x}_{ti} - \mathbf{o}_t)/r^\star_{ti}$ is its unit viewing direction. We construct

$$\mathbf{d}_{ti} = Q_{ti}\mathbf{d}^\star_{ti}, \qquad r_{ti} = r^\star_{ti}\exp(\varepsilon_{ti}), \qquad \mathbf{y}_{ti} = \mathbf{o}_t + r_{ti}\,\mathbf{d}_{ti}. \tag{5}$$

The rotation $Q_{ti} \in \mathrm{SO}(3)$ applies small angular perturbations to approximate *correspondence errors* introduced by the 2D tracker. The log-range perturbation ε_{ti} uses independent noise to model *depth jitter* and temporally correlated offsets to approximate *depth drift*. Occasional large offsets additionally approximate *wrong-surface observations*. The exponential converts these log-range offsets to multiplicative range changes. Camera centers remain unchanged; we do not explicitly simulate camera-pose errors or persistent switches to another object's trajectory.

To generate *missing observations*, we initialize the input mask m_{ti} from ground-truth visibility $v^\star_{ti}$ and remove observations using moving geometric occluders and contiguous temporal dropouts. To model wrong-surface observations caused by missed occlusions, we retain some observations near the start or end of a synthetic occlusion ($m_{ti} = 1$) at ranges corresponding to the occluder's depth. Clean trajectories remain the targets at every frame, including those with $m_{ti} = 0$. The simulated observations undergo the same normalization, initialization and tokenization as at inference (Sec. 3.2), with scale and anchors computed from the corrupted input. Appendix A.3 specifies the corruption distributions, and Sec. 4.4 evaluates their contributions.

Training for transfer. To encourage transfer across trackers and geometry models, we vary the parameters of these corruption processes across training examples, following domain randomization in robotics (Tobin et al., 2017; Peng et al., 2018). We sample quantities such as noise strengths, temporal correlation coefficients and outlier rates from fixed distributions (Appendix A.3). This exposes the corrector to diverse patterns of lifting errors without using provider-generated training inputs. The same trained corrector improves the initialized lifts of all six evaluated tracker–geometry combinations (Fig. 3).

Reconstruction objective. We supervise 3D positions and frame-to-frame displacements, including where input observations are unavailable, with a weighted squared-error objective,

$$\mathcal{L} = \frac{\sum_{t,i} w_{ti}\|\mathbf{e}_{ti}\|^2}{\sum_{t,i} w_{ti}} + \lambda\,\frac{\sum_{t\geq 1,i} w_{ti}\|\mathbf{e}_{ti} - \mathbf{e}_{t-1,i}\|^2}{\sum_{t\geq 1,i} w_{ti}}. \tag{6}$$

Here $\mathbf{e}_{ti} = (\hat{\mathbf{x}}_{ti} - \mathbf{x}_{ti})/(s\rho_i)$ is the relative position error, with $\rho_i = \max(\|\mathbf{a}_i\|, \rho_{\min})$ the floored reference range in normalized coordinates. This scaling prevents distant points from dominating the objective. The first term supervises positions; the second, weighted by λ, matches predicted and true displacements between consecutive sampled frames, discouraging jitter while preserving the target motion. The weights w_{ti} count moving tracks four times as much as static ones, and unavailable observations twice as much as available ones (Appendix A.3).

4 Experiments

4.1 Experimental Setup

Evaluation Datasets. We evaluate dense 3D point tracking in world coordinates on the four WorldTrack (Feng et al., 2025) datasets: the real-world Aria Digital Twin (ADT) (Pan et al., 2023) and Panoptic Studio (PS) (Joo et al., 2019) from TAPVid-3D (Koppula et al., 2024), and the synthetic PointOdyssey (PO) (Zheng et al., 2023) and Dynamic Replica (DR) (Karaev et al., 2023). We score dense predictions at the sparse ground-truth trajectories of points visible in the reference frame, over the first 84 frames. We also evaluate on TrackCraft3R's (Nam et al., 2026) 50 24-frame Kubric (Greff et al., 2022) clips, with ground-truth trajectories for every reference pixel (Appendix C.1).

Metrics. We report the Average Percentage of 3D Points within Delta (APD3D↑) (Koppula et al., 2024), averaged over distance thresholds $\delta \in \{0.1, 0.3, 0.5, 1.0\}$ m. Following 4RC (Luo et al.,

Table 1: **Dense 3D tracking on WorldTrack.** APD3D (%) on all (All) and dynamic (Dyn) points after RANSAC Sim(3) alignment; bold best, underline second best. Time includes the complete pipeline. †Timing at 256^2 instead of 512^2. ‡Restricted ADT coverage for MotionCrafter, also affecting its means (Appendix B.2).

	ADT		PS		PO		DR		Mean		
Method	All	Dyn	All	Dyn	All	Dyn	All	Dyn	All	Dyn	Time (s)
4RC	87.07	78.57	86.25	86.10	**85.92**	81.17	88.51	80.56	86.94	81.60	31.8
Any4D	76.17	78.32	77.57	77.61	73.96	73.56	83.54	72.83	77.81	75.58	3.2
MotionCrafter	67.04‡	83.34‡	62.60	62.63	46.29	49.09	70.83	62.39	61.69‡	64.36‡	20.4
OpenD4RT	71.71	79.68	84.36	85.27	68.39	74.21	77.55	79.69	75.50	79.71	284†
St4RTrack	81.39	79.39	72.84	72.24	71.36	71.44	78.29	67.23	75.97	72.58	2.9
Track4World	77.48	91.80	80.10	79.39	69.13	69.43	78.94	71.98	76.41	78.15	8.4
TraceAnything	81.76	77.01	71.88	71.66	49.77	53.14	68.91	61.80	68.08	65.90	5.2
V-DPM	88.42	79.40	89.26	89.49	84.46	**85.95**	82.99	80.80	86.28	83.91	591
DELTAv2 + DA3	83.63	92.19	85.38	85.37	66.32	67.35	76.88	64.50	78.06	77.35	33.4
DELTAv2 + VGGT-Ω	84.57	91.67	89.41	89.22	67.49	68.60	82.79	72.23	81.06	80.43	29.8
TrackCraft3R + DA3	91.75	90.82	84.39	83.95	74.04	71.42	78.44	65.84	82.16	78.01	22.1
TrackCraft3R + VGGT-Ω	92.56	89.54	86.56	86.65	72.66	71.48	83.89	73.74	83.92	80.36	18.5
Ours, one trained corrector for all rows (2D tracker + geometry provider)											
AllTracker + DA3	93.69	**94.98**	85.46	85.59	76.04	77.44	79.50	68.94	83.68	81.74	11.4
AllTracker + VGGT-Ω	94.04	92.17	90.40	90.46	76.10	78.96	85.74	78.10	86.57	84.92	7.3
AllTracker + 4RC	91.07	94.13	88.73	88.65	83.63	80.26	88.81	81.08	88.06	86.03	36.4
CoWTracker + DA3	93.78	94.85	86.99	86.64	77.63	79.19	79.72	68.19	84.53	82.22	15.3
CoWTracker + VGGT-Ω	**94.14**	92.09	**92.20**	**92.06**	77.63	80.49	86.00	78.70	87.49	85.84	11.7
CoWTracker + 4RC	91.41	94.31	89.97	89.83	85.60	82.14	**89.05**	**81.49**	**89.01**	**86.95**	40.4

2026), we evaluate all points (*All*) and dynamic points (*Dyn*) after RANSAC Sim(3) alignment, fitted separately to each scoring set. We also evaluate image-plane accuracy using APD2D (Appendix C.4). Appendices B and C provide scoring details, complementary metrics (AJ, OA and EPE), and global-scale alignment results.

Baselines. We pair AllTracker (Harley et al., 2025) or CoWTracker (Lai et al., 2026) with depth and camera estimates from Depth Anything 3 (DA3) (Lin et al., 2025) or VGGT-Ω (Wang et al., 2026). TrackCraft3R (Nam et al., 2026) and DELTAv2 (Ngo et al., 2025a) receive the same estimated depth, cameras and intrinsics as ours. We also pair both 2D trackers with 4RC's (Luo et al., 2026) own depth and camera estimates, comparing our tracking and correction pipeline with 4RC's trajectory predictions under matched geometry. Additional baselines include V-DPM (Sucar et al., 2026), St4RTrack (Feng et al., 2025), Any4D (Karhade et al., 2025), Track4World (Lu et al., 2026), TraceAnything (Liu et al., 2026), MotionCrafter (Zhu et al., 2026), and OpenD4RT (RHOS Team, 2026), an open reimplementation of D4RT (Zhang et al., 2025a). No method receives ground-truth intrinsics (Appendices A and B.2).

4.2 Quantitative Results

Accuracy. Our method achieves state-of-the-art mean APD3D on all and dynamic points with CoWTracker and 4RC geometry (Table 1). It is particularly strong on the real-world datasets ADT and PS, where our variants lead both metrics. All six configurations share one trained corrector: for both trackers, mean accuracy increases from DA3 to VGGT-Ω to 4RC geometry, demonstrating gains from swapping geometry providers without retraining.

Combining 2D motion and 3D geometry. To test the decomposition introduced in Sec. 1, we compare methods with the same geometry inputs. Even before correction, CoWTracker lifts with initialized gaps and VGGT-Ω geometry reach 84.9 mean all-point APD3D (Table 2, first row), compared with 83.9 for TrackCraft3R and 81.1 for DELTAv2 using the same geometry (Table 1). After correction, our method outperforms both in every APD3D column of Table 1. Given 4RC's own geometry, our CoWTracker variant improves on 4RC's mean accuracy by 2.1 (All) and 5.4 (Dyn) points, demonstrating the advantage of our pipeline with geometry held fixed. On Kubric, our DA3 and VGGT-Ω variants trail 4RC but lead methods sharing their geometry on all-point APD3D; with 4RC geometry, ours outperform 4RC on dynamic points (Appendix C.1).

Gains from trajectory repair and completion. Figure 3 measures what correction adds to the initialized lift (Sec. 3.2). Accuracy improves across datasets and all six provider combinations, with

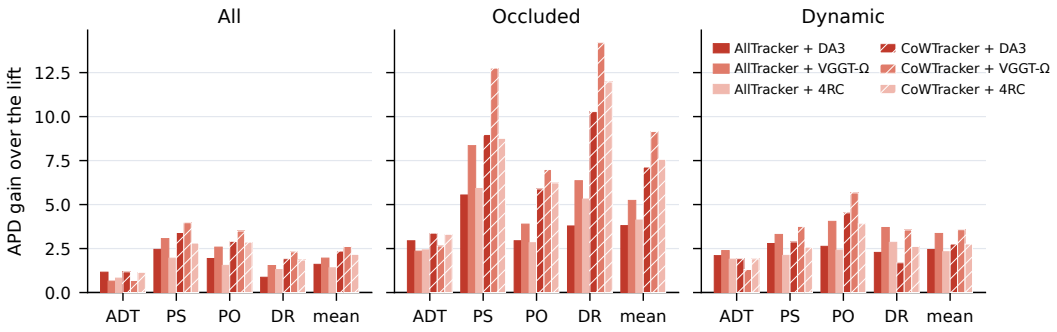


Figure 3: **Gains from trajectory repair and completion.** APD3D gains (%) over the initialized lift (Sec. 3.2), per dataset and averaged, for all six tracker–geometry combinations; hatching denotes CoWTracker. *Occluded* selects ground-truth-occluded point–frames from *All*, retaining the All-point alignment. This subset may differ from the corrector's masked inputs.

the largest gains on ground-truth-occluded points. Correction adds 1.6–3.6 points to mean APD3D in Table 2. With CoWTracker and VGGT-Ω, all-point APD3D rises from 84.9 to 87.5, and the corrected model surpasses 4RC and V-DPM on all and dynamic points.

Image-plane accuracy. Native CoWTracker outperforms the projected trajectories of every evaluated 3D method in APD2D, on each dataset for both all and dynamic points (Table 11), supporting our use of strong pretrained 2D correspondences. Among 3D methods, our CoWTracker with 4RC geometry leads mean all-point APD2D and ranks second on dynamic points, alongside the best mean APD3D on both sets. Our projections remain below native CoWTracker; Appendix C.4 analyzes this gap and the distinction between image-plane and depth accuracy.

Efficiency. Our VGGT-Ω configurations offer competitive accuracy while running 2.7–$4.4\times$ faster than 4RC, including all provider inference (Table 1). Using 4RC geometry requires more time but delivers our highest mean accuracy. Appendix D.4 analyzes runtime and memory across resolutions.

4.3 Qualitative Results

Figure 4 illustrates tracking under rotational motion, large displacements and occlusion. In *spin*, our trajectories form coherent arcs, while TrackCraft3R produces scattered paths. Our tracks follow the car through shadow-induced appearance changes in *car-shadow* and recover the dog's head after occlusion in *libby*, where competing predictions stop short. Additional qualitative comparisons in Appendix E include the same initialized lift before and after correction, showing reduced scatter and more coherent motion with tracks and geometry held fixed. These comparisons complement the quantitative gains in Fig. 3. Videos are available on our project page.

4.4 Ablation Study

Simulated versus provider-generated lifts. We test whether training on provider-generated lifting errors transfers better than training on simulated corruptions. Both approaches use PointOdyssey with the same network and training budget (Fig. 5). A single simulated-lift model improves over the initialized lift on every dataset with both test geometry providers, outperforming models trained on AllTracker+DA3 or AllTracker+VGGT-Ω lifts. Both provider-trained models improve on their training scenes ($+14.1$ and $+8.3$ APD3D points for DA3 and VGGT-Ω, respectively), but these gains transfer inconsistently. DA3-lift training improves accuracy on held-out PointOdyssey scenes, but simulated training yields larger gains. The model trained on VGGT-Ω lifts reduces accuracy on every evaluated dataset.

Provider-generated training requires aligning ground-truth targets, and discarding windows with unreliable alignment (Appendix D.1). To control for this selection, we repeat simulated-lift training on the same windows retained for DA3-lift training. Its advantage persists with both test geometry providers (Fig. 6), showing that differences in the selected windows alone do not explain the result. Simulation also avoids provider inference and target alignment during data preparation.

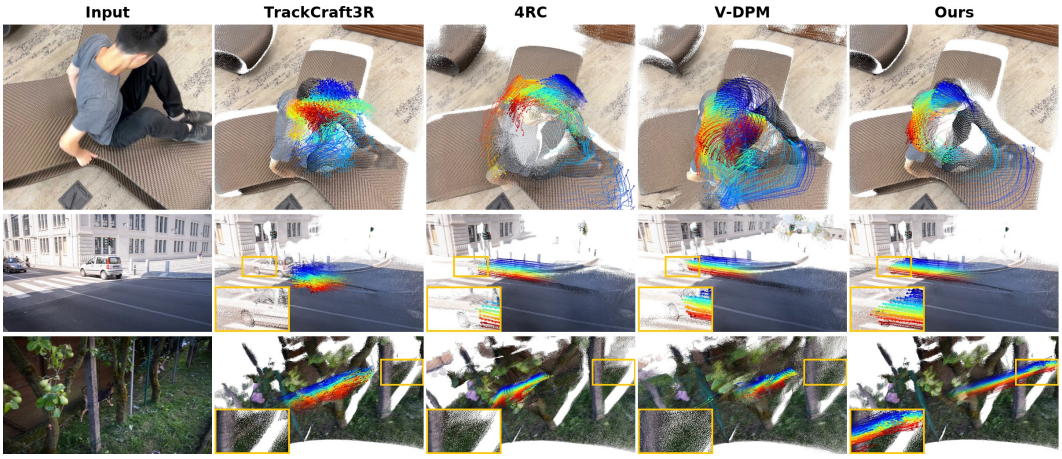


Figure 4: **Qualitative comparison.** Top to bottom: DyCheck *spin* (Gao et al., 2022) and DAVIS *car-shadow* and *libby* (Perazzi et al., 2016). Trajectories of the same reference pixels are shown from matched novel viewpoints. Ours uses AllTracker and VGGT-Ω; TrackCraft3R shares this geometry, while 4RC and V-DPM use their own reconstructions. Yellow boxes mark magnified regions.

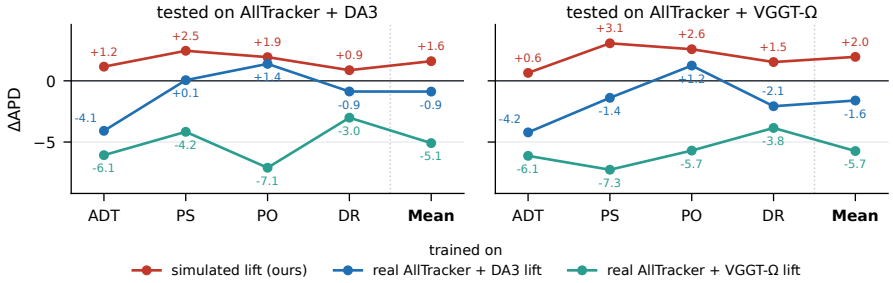


Figure 5: **Simulated versus provider-generated lifts.** APD3D gains on all points, in percentage points over the initialized lift, with AllTracker and DA3 (left) or VGGT-Ω (right). "Real" denotes non-simulated, provider-generated training lifts. All runs use the same architecture and 150,000 training steps. Appendix D.1 reports a control with matched training windows.

Contribution of corruption families. We test the simulation of Sec. 3.4 by removing *correlated noise* (flicker, tilt, bias and ray noise), *wrong-surface observations* (range outliers and observations retained at an occluder's depth), or *synthetic gaps* (occluders and tracker dropouts). Table 2 shows that removing wrong-surface corruptions causes the largest loss of accuracy across both trackers and geometry models, supporting their role in learning to repair incorrect observations. Correlated noise provides smaller mean improvements. Natural PointOdyssey occlusions remain in every training variant. Additional synthetic gaps have little effect on overall mean accuracy but yield small mean gains on occluded points across all four tracker–geometry combinations (Appendix Table 12), supporting their inclusion during training.

Chunking and track-group size. With $C = 4$ frames per token, correction is more accurate and about three times faster than with $C = 1$. Increasing to $C = 8$ further reduces runtime but slightly lowers accuracy. We use $P = 1024$ tracks per group because smaller groups reduce accuracy and larger groups offer little improvement at higher attention cost (Appendix D.3).

Table 2: **Contribution of training corruptions.** Removing wrong-surface corruptions causes the largest drop. APD3D (%), averaged over four datasets; gains over the initialized lift are shown in parentheses, in percentage points. Each variant is retrained with the same data, network, loss and budget (Appendix D.1).

	AllTracker + DA3		AllTracker + VGGT-Ω		CoWTracker + DA3		CoWTracker + VGGT-Ω	
Training corruptions	All	Dyn	All	Dyn	All	Dyn	All	Dyn
Lift (the corrector's input)	82.1	79.3	84.6	81.6	82.2	79.5	84.9	82.3
Full simulation (ours)	83.7 (+1.6)	81.7 (+2.4)	86.6 (+2.0)	84.9 (+3.4)	84.5 (+2.4)	82.2 (+2.8)	87.5 (+2.6)	85.8 (+3.6)
without correlated noise	83.5 (+1.5)	81.3 (+2.0)	86.4 (+1.8)	84.5 (+2.9)	84.2 (+2.0)	82.1 (+2.7)	87.3 (+2.5)	85.7 (+3.5)
without wrong surface	82.3 (+0.2)	79.6 (+0.3)	84.9 (+0.3)	82.1 (+0.5)	82.8 (+0.6)	80.2 (+0.7)	85.5 (+0.6)	83.1 (+0.9)
without synthetic gaps	83.7 (+1.6)	81.5 (+2.2)	86.6 (+2.0)	84.9 (+3.4)	84.5 (+2.4)	82.3 (+2.8)	87.5 (+2.6)	85.8 (+3.6)
white noise only	82.4 (+0.3)	79.8 (+0.5)	85.0 (+0.4)	82.6 (+1.0)	82.8 (+0.6)	80.5 (+1.0)	85.5 (+0.6)	83.2 (+1.0)

5 Conclusion

We showed that pretrained 2D tracks lifted with estimated geometry and simple gap initialization already rival dedicated 3D trackers using the same geometry. Building on this observation, ReLift combines these providers with efficient feed-forward trajectory repair and completion for dense monocular 3D tracking. Training on simulated lifting errors enables one corrector to improve all evaluated tracker–geometry combinations without retraining, with the largest gains on occluded points. Our framework achieves state-of-the-art mean accuracy on WorldTrack, while faster configurations retain competitive accuracy at substantially lower runtime than comparably accurate 3D trackers. Performance remains sensitive to input geometry (Appendix C.1), motivating future work on robustness to geometric and persistent correspondence errors.

Acknowledgments

Yoni Kasten is a Dr. Natalie Shaffer Faculty Fellow and acknowledges support from the Israel Science Foundation (grant 1425/26). Or Litany acknowledges support from the Israel Science Foundation (grant 624/25) and the Azrieli Foundation Early Career Faculty Fellowship.

References

Ijaz Akhter, Yaser Sheikh, Sohaib Khan, and Takeo Kanade. Nonrigid structure from motion in trajectory space. *Advances in Neural Information Processing Systems*, 21, 2008.

Homanga Bharadhwaj, Roozbeh Mottaghi, Abhinav Gupta, and Shubham Tulsiani. Track2Act: Predicting point tracks from internet videos enables generalizable robot manipulation. In *European Conference on Computer Vision*, pp. 306–324, 2024. URL https://arxiv.org/abs/2405.01527.

Tomer Borreda, Fangqiang Ding, Sanja Fidler, Shengyu Huang, and Or Litany. RadarGen: Automotive radar point cloud generation from cameras. In *ECCV*, 2026.

Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3D shape from image streams. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000*, volume 2, pp. 690–696. IEEE, 2000.

Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo de Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gojcic, Sanja Fidler, Marco Pavone, Li Song, and Yue Wang. OmniRe: Omni urban scene reconstruction. In *The Thirteenth International Conference on Learning Representations*, 2025.

Seokju Cho, Jiahui Huang, Seungryong Kim, and Joon-Young Lee. Seurat: From moving points to depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7211–7221, 2025.

Avigail Cohen Rimon, Amir Mann, Mirela Ben-Chen, and Or Litany. SpectralSplats: Robust differentiable tracking via spectral moment supervision. In *European Conference on Computer Vision*. Springer, 2026.

Yuchao Dai, Hongdong Li, and Mingyi He. A simple prior-free method for non-rigid structure-from-motion factorization. *International Journal of Computer Vision*, 107(2):101–122, 2014.

Carl Doersch, Ankush Gupta, Larisa Markeeva, Adrià Recasens, Lucas Smaira, Yusuf Aytar, João Carreira, Andrew Zisserman, and Yi Yang. TAP-vid: A benchmark for tracking any point in a video. In *Advances in Neural Information Processing Systems*, volume 35, pp. 13610–13626, 2022.

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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. URL https://st4rtrack.github.io/.

Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. *Advances in Neural Information Processing Systems*, 2022.

Daniel Geng, Charles Herrmann, Junhwa Hur, Forrester Cole, Serena Zhang, Tobias Pfaff, Tatiana Lopez-Guevara, Carl Doersch, Yusuf Aytar, Michael Rubinstein, Chen Sun, Oliver Wang, Andrew Owens, and Deqing Sun. Motion prompting: Controlling video generation with motion trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. URL https://motion-prompting.github.io/.

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 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3739–3751, 2022. doi: 10.1109/CVPR52688.2022.00373.

Adam W Harley, Yang You, Xinglong Sun, Yang Zheng, Nikhil Raghuraman, Yunqi Gu, Sheldon Liang, Wen-Hsuan Chu, Achal Dave, Suya You, et al. Alltracker: Efficient dense point tracking at high resolution. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5253–5262. IEEE, 2025.

Linyi Jin, Richard Tucker, Zhengqi Li, David Fouhey, Noah Snavely, and Aleksander Holynski. Stereo4D: Learning how things move in 3D from internet stereo videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. URL https://arxiv.org/abs/2412.09621.

Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic Studio: A massively multiview system for social interaction capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(1):190–204, 2019. doi: 10.1109/TPAMI.2017.2782743.

Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. DynamicStereo: Consistent dynamic depth from stereo videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13229–13239, 2023. doi: 10.1109/CVPR52729.2023.01271.

Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker: It is better to track together. In *European Conference on Computer Vision*, pp. 18–35. Springer, 2024.

Jay Karhade, Nikhil Keetha, Yuchen Zhang, Tanisha Gupta, Akash Sharma, Sebastian Scherer, and Deva Ramanan. Any4D: Unified feed-forward metric 4D reconstruction. *arXiv preprint arXiv:2512.10935*, 2025. URL https://arxiv.org/abs/2512.10935.

Yoni Kasten, Wuyue Lu, and Haggai Maron. Fast encoder-based 3D from casual videos via point track processing. *Advances in Neural Information Processing Systems*, 37:96150–96180, 2024.

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. In *Advances in Neural Information Processing Systems*, 2024. URL https://arxiv.org/abs/2407.05921.

Zihang Lai, Eldar Insafutdinov, Edgar Sucar, and Andrea Vedaldi. CoWTracker: Tracking by warping instead of correlation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 42571–42580, 2026. arXiv:2602.04877.

Jiahui Lei, Yijia Weng, Adam W. Harley, Leonidas Guibas, and Kostas Daniilidis. MoSca: Dynamic gaussian fusion from casual videos via 4D motion scaffolds. *arXiv preprint arXiv:2405.17421*, 2024. URL https://arxiv.org/abs/2405.17421v2.

Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3D with MASt3R. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 71–91, 2024.

Haotong Lin, Sili Chen, Junhao Liew, Donny Y Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025.

Xinhang Liu, Yuxi Xiao, Donny Y Chen, Jiashi Feng, Yu-Wing Tai, Chi-Keung Tang, and Bingyi Kang. Trace anything: Representing any video in 4d via trajectory fields. In *International Conference on Learning Representations*, volume 2026, pp. 37738–37767, 2026.

Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13–23, 2023.

Guanxing Lu, Shiyi Zhang, Ziwei Wang, Changliu Liu, Jiwen Lu, and Yansong Tang. Manigaussian: Dynamic gaussian splatting for multi-task robotic manipulation. In *European Conference on Computer Vision*, pp. 349–366. Springer, 2024.

Jiahao Lu, Weitao Xiong, Jiacheng Deng, Peng Li, Tianyu Huang, Zhiyang Dou, Cheng Lin, Sai-Kit Yeung, and Yuan Liu. TrackingWorld: World-centric monocular 3D tracking of almost all pixels. In *Advances in Neural Information Processing Systems*, 2025. URL https://arxiv.org/abs/2512.08358.

Jiahao Lu, Jiayi Xu, Wenbo Hu, Ruijie Zhu, Chengfeng Zhao, Sai-Kit Yeung, Ying Shan, and Yuan Liu. Track4World: Feedforward world-centric dense 3D tracking of all pixels. *arXiv preprint arXiv:2603.02573*, 2026.

Yihang Luo, Shangchen Zhou, Yushi Lan, Xingang Pan, and Chen Change Loy. 4RC: 4D reconstruction via conditional querying anytime and anywhere. In *International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2602.10094.

Jisu Nam, Jahyeok Koo, Soowon Son, Jaewoo Jung, Honggyu An, Junhwa Hur, and Seungryong Kim. TrackCraft3R: Repurposing video diffusion transformers for dense 3D tracking. *arXiv preprint arXiv:2605.12587*, 2026. URL https://arxiv.org/abs/2605.12587.

Tuan Duc Ngo, Ashkan Mirzaei, Guocheng Qian, Hanwen Liang, Chuang Gan, Evangelos Kalogerakis, Peter Wonka, and Chaoyang Wang. DELTAv2: Accelerating dense 3D tracking. *arXiv preprint arXiv:2508.01170*, 2025a. URL https://arxiv.org/abs/2508.01170.

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. In *International Conference on Learning Representations*, 2025b. URL https://snap-research.github.io/DELTA/.

Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Carl Yuheng Ren. Aria Digital Twin: A new benchmark dataset for egocentric 3D machine perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 20076–20086, 2023. doi: 10.1109/ICCV51070.2023.01842.

Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *IEEE International Conference on Robotics and Automation*, 2018. doi: 10.1109/ICRA.2018.8460528. URL https://arxiv.org/abs/1710.06537.

Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.

RHOS Team. OpenD4RT: An unofficial PyTorch/GPU implementation of D4RT for 4D reconstruction and tracking. https://github.com/Lijiaxin0111/Open-d4rt, 2026.

Assaf Singer, Noam Rotstein, Amir Mann, Ron Kimmel, and Or Litany. Time-to-Move: Training-free motion-controlled video generation via dual-clock denoising. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/1eaa5146756be028ad6fff1efcc8e6bd-Abstract-Conference.html.

Edgar Sucar, Zihang Lai, Eldar Insafutdinov, and Andrea Vedaldi. Dynamic point maps: A versatile representation for dynamic 3D reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7295–7305, 2025.

Edgar Sucar, Eldar Insafutdinov, Zihang Lai, and Andrea Vedaldi. V-DPM: 4D video reconstruction with dynamic point maps. *arXiv preprint arXiv:2601.09499*, 2026. URL https://arxiv.org/abs/2601.09499.

Jonathan Taylor, Allan D. Jepson, and Kiriakos N. Kutulakos. Non-rigid structure from locally-rigid motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2761–2768, 2010.

Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *arXiv preprint arXiv:1703.06907*, 2017. URL https://arxiv.org/abs/1703.06907.

Narek Tumanyan, Assaf Singer, Shai Bagon, and Tali Dekel. DINO-Tracker: Taming DINO for self-supervised point tracking in a single video. *arXiv preprint arXiv:2403.14548*, 2024.

Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. VGGT: Visual geometry grounded transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5294–5306, 2025a.

Jianyuan Wang, Minghao Chen, Shangzhan Zhang, Nikita Karaev, Johannes Schönberger, Patrick Labatut, Piotr Bojanowski, David Novotny, Andrea Vedaldi, and Christian Rupprecht. VGGT-Ω. *arXiv preprint arXiv:2605.15195*, 2026.

Qianqian Wang, Yen-Yu Chang, Ruojin Cai, Zhengqi Li, Bharath Hariharan, Aleksander Holynski, and Noah Snavely. Tracking everything everywhere all at once. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. URL https://omnimotion.github.io/.

Qianqian Wang, Vickie Ye, Hang Gao, Weijia Zeng, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4D reconstruction from a single video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025b. URL https://arxiv.org/abs/2407.13764.

Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jérôme Revaud. DUSt3R: Geometric 3D vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20697–20709, June 2024.

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 *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6726–6737. IEEE, 2025.

David Yifan Yao, Albert J. Zhai, and Shenlong Wang. Uni4D: Unifying visual foundation models for 4D modeling from a single video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1116–1126, 2025. URL https://arxiv.org/abs/2503.21761.

Chuhan Zhang, Guillaume Le Moing, Skanda Koppula, Ignacio Rocco, Liliane Momeni, Junyu Xie, Shuyang Sun, Rahul Sukthankar, Joëlle K. Barral, Raia Hadsell, Zoubin Ghahramani, Andrew Zisserman, Junlin Zhang, and Mehdi S. M. Sajjadi. Efficiently reconstructing dynamic scenes one D4RT at a time. *arXiv preprint arXiv:2512.08924*, 2025a. URL https://arxiv.org/abs/2512.08924.

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. In *International Conference on Learning Representations (ICLR)*, 2025b.

Zhoutong Zhang, Forrester Cole, Zhengqi Li, Michael Rubinstein, Noah Snavely, and William T. Freeman. Structure and motion from casual videos. In *European Conference on Computer Vision*, pp. 20–37. Springer, 2022.

Yang Zheng, Adam W. Harley, Bokui Shen, Gordon Wetzstein, and Leonidas J. Guibas. PointOdyssey: A large-scale synthetic dataset for long-term point tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. URL https://arxiv.org/abs/2307.15055.

Ruijie Zhu, Jiahao Lu, Wenbo Hu, Xiaoguang Han, Jianfei Cai, Ying Shan, and Chuanxia Zheng. MotionCrafter: Dense geometry and motion reconstruction with a 4D VAE. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026. URL https://arxiv.org/abs/2602.08961.

A METHOD DETAILS

A.1 OBSERVATION PREPROCESSING

Depth lookup uses the nearest pixel, clamped to the image border, while the viewing ray uses the sub-pixel track position. The observation mask in Sec. 3.1 uses $\gamma = 0.85$ for AllTracker and $\gamma = 0.65$ for CoWTracker; the latter is calibrated to match AllTracker's false-visible rate on 24 PointOdyssey training windows. These thresholds apply to visibility scores, not tracker confidence outputs.

A.2 CORRECTOR IMPLEMENTATION

Input and output scaling. After scene normalization (Sec. 3.2), we multiply the initialized displacement, its masked counterpart and the normalized camera centre in Eq. (2) by $S = 10$. The gap feature is $\min(g_{ti}, 40)/40$. For raw prediction-head output $\mathbf{z}_{ti}$, the residual in Eq. (4) is $\boldsymbol{\Delta}_{ti} = \mathbf{z}_{ti}/S$. The loss uses scaled errors $S\mathbf{e}_{ti}$ with unchanged range normalization, giving $S^2\mathcal{L}$.

Position embeddings. For an even embedding width d, define

$$\mathrm{PE}(x; d, \omega) = \left[\sin(x\,\omega^{-k/(d/2)}),\ \cos(x\,\omega^{-k/(d/2)})\right]_{k=0}^{d/2-1}. \tag{7}$$

The anchor embedding $\boldsymbol{\psi}(\mathbf{a}_i)$ concatenates $\mathrm{PE}(8a_{i,c}; 32, 10^4)$ over the three coordinates c, giving 96 features, then applies Linear($96 \to 384$), GELU and Linear($384 \to 384$). The temporal embedding is $\mathbf{p}_n = \mathrm{PE}(n; 384, 10^3)$, using the chunk index n.

Architecture. Each block applies temporal self-attention, across-track self-attention and an MLP in sequence, each with pre-layer normalization and a residual connection. There is no track-index or pixel embedding, so reordering tracks reorders the outputs. Predictions can nevertheless depend on which other tracks are included in the group. Table 3 specifies the architecture. The zero-initialized prediction head makes the untrained corrector reproduce the initialized trajectories.

Table 3: **Corrector architecture settings.**

Setting	Value
Frames per tracklet C	4
Feature width D	384
Factorized attention blocks	8
Attention heads	6, with width 64 each
Feed-forward hidden width / activation	1536 / GELU
Output normalization	Final layer norm
Prediction head	Zero-initialized linear map to $3C$ coordinates
Parameters	19.1 M

A.3 TRAINING DETAILS

Training data and sampling. We use 100 scenes (6,137 windows) for training and 8 for validation from the PointOdyssey training split (Sec. 3.4). We exclude `cab_e_ego2`, an alternate view of WorldTrack's `cab_e_3rd` evaluation recording. Cached windows contain 84 frames at temporal strides 1 and 2, with starts separated by 42 and 84 source frames, respectively. Eligible tracks have valid trajectories and are visible, inside the image and in front of the camera in the reference frame; windows with fewer than 1,024 are discarded. We cache 2,048 tracks per window, at most half moving unless static tracks are insufficient. A track is marked moving, $\mu_i = 1$, if its displacement from its reference position exceeds 5 cm anywhere in the cached window; otherwise $\mu_i = 0$.

Each training sample draws 512 cached tracks. Each batch retains the first T frames, with T sampled uniformly from $\{48, 68, 84\}$, and mirrors trajectory and camera x coordinates with probability one half. Motion labels remain those of the full cached window.

Corruption processes. The simulated lift in Eq. (5) uses the log-range perturbation

$$\varepsilon_{ti} = f_t + \mathbf{b}_t^\top \boldsymbol{\delta}_{ti} + \beta_{ti} + \eta_{ti} + j_{ti}. \tag{8}$$

Here f_t is frame-shared flicker, $\mathbf{b}_t \in \mathbb{R}^2$ a tilt acting on the first two coordinates $\boldsymbol{\delta}_{ti}$ of the true ray in reference camera coordinates, β_{ti} a per-track bias, η_{ti} white noise and j_{ti} sparse outliers. The

processes f_t, $\mathbf{b}_t$ and β_{ti} use unit-variance autoregressive sequences $z_t = \kappa z_{t-1} + \sqrt{1-\kappa^2}\,\xi_t$, with $z_0, \xi_t \sim \mathcal{N}(0,1)$, scaled as in Table 4. We clip ε_{ti} to $[-1,1]$ before exponentiation. To perturb a ray, we project Gaussian noise $\mathcal{N}(\mathbf{0}, \sigma^2 I_3)$ onto its orthogonal plane, add it to the ray and renormalize.

Synthetic occluders are rectangles moving in the reference camera's normalized image plane $(x/z, y/z)$; they hide points projected inside and lying behind them. We remove these observations and contiguous dropouts from ground-truth visibility. Ground-truth-visible cells hidden by an occluder and adjacent to an available observation can instead be retained as false-visible. For these cells, we replace the perturbed range by the clean range times the occluder-to-point depth ratio; angular noise still applies. We enforce $m_{0i} = 1$ and disable occluder-depth replacement in the reference frame. Table 4 gives the sampling distributions, fixed across training and not fitted to evaluation sequences.

Table 4: **Training corruption distributions.** $\mathcal{U}$ denotes uniform sampling. Noise levels and outlier rates are drawn once per training example.

Term	Distribution
Flicker f_t	std $\mathcal{U}(0, 0.03)$, autoregressive with $\kappa \sim \mathcal{U}(0.3, 0.95)$
Tilt $\mathbf{b}_t$	std $\mathcal{U}(0, 0.03)$, autoregressive with the flicker's κ
Bias β_{ti}	std $\mathcal{U}(0, 0.02)$, autoregressive with $\kappa = 0.9$
White noise η_{ti}	std $\mathcal{U}(0.002, 0.02)$
Outliers j_{ti}	rate $\mathcal{U}(0, 0.05)$ per cell, next-frame persistence probability 0.5, jump $\pm\mathcal{U}(0.1, 0.5)$
Angular ray noise	std $\mathcal{U}(0.0005, 0.005)$ rad
Occluders	up to two, with independent probabilities 0.6 and 0.2; target track drawn in proportion to $\mu_i + 10^{-3}$
placement	centred on the target at frame $t_s \sim \mathcal{U}\{\lfloor T/4 \rfloor, \dots, T-1\}$, half-size $\mathcal{U}(0.03, 0.25)$ per axis
motion, depth	constant velocity $\mathcal{N}(0, 0.002^2)$ per axis in image-plane units per frame; depth $\mathcal{U}(0.5, 0.9)$ times the target's at t_s
Dropouts	10% of tracks, start $\mathcal{U}\{1, \dots, T-1\}$, length $\mathcal{U}\{2, \dots, 20\}$
False-visible observations	probability 0.3 for ground-truth-visible cells hidden by an occluder and adjacent to an available observation

Loss settings. In Eq. (6), we set $\rho_{\min} = 0.3$, $\lambda = 4$ and $w_{ti} = (1 + 3\mu_i)(2 - m_{ti})$, where μ_i is the moving-track indicator and m_{ti} the input availability mask. Displacements use consecutive sampled frames, including at temporal stride 2.

Training optimization. All runs, including ablations, use 150,000 AdamW steps with learning rate $3 \cdot 10^{-4}$, weight decay 0.01 and $\beta = (0.9, 0.95)$. A 1,000-step linear warm-up is followed by step-indexed cosine decay to 3% of the peak. Training uses batches of 16, bfloat16 forward passes and gradient clipping at norm 1 on one RTX PRO 6000 GPU. We evaluate the final exponential-moving-average checkpoint, with decay $\min(0.999, (1+k)/(10+k))$ at step k.

A.4 Dense Inference

We correct tracks with positive reference depth; other pixels retain their uncorrected lift. Eligible tracks are shuffled with seed 0 and processed in groups of $P = 1024$, up to 8 groups per batch (Appendix D.3 evaluates P). The shared scale s is the median reference depth over these eligible pixels; during training, it uses the sampled tracks.

A short final group is filled by cyclically repeating pixels from its batch. Repeated tracks participate in attention without masking, and their predictions are written to their original pixel indices. For temporal padding, sequences are extended to a multiple of C with at most $C-1$ unavailable frames, holding the camera at its last position. Padded-frame predictions are discarded.

B Evaluation Details

B.1 Evaluation Protocol

We evaluate ground-truth points visible and inside the image in the first frame. One ADT clip has no eligible queries and is excluded for all methods, leaving 199 WorldTrack clips. Dyn selects complete

tracks whose cumulative ground-truth motion over the evaluated frames, $\sum_{t=0}^{T-2} \|\mathbf{x}_{t+1,i} - \mathbf{x}_{ti}\|$, exceeds 0.01: metres on WorldTrack and scene units on Kubric.

Following 4RC (Luo et al., 2026), RANSAC Sim(3) alignment fits scale, rotation and translation, whereas global-scale alignment fits only scale. Each alignment is estimated per clip, separately for All and Dyn. RANSAC uses the same seeds across methods for each clip and scoring set.

Average Jaccard (AJ) and Occlusion Accuracy (OA) follow TAP-Vid (Doersch et al., 2022), with predicted visibility thresholded at 0.5. AJ averages the 3D Jaccard score over the APD3D thresholds of Sec. 4.1, using the all-point alignment; OA measures the fraction of point–frames with correct visibility predictions.

Metrics are averaged over clips within each dataset. The WorldTrack mean weights the four dataset scores equally; Kubric results are averaged over its 50 clips.

B.2 Baseline Settings

Table 5 lists the released-code configurations used for the WorldTrack comparisons in the main paper and appendix. Kubric uses the same settings with fewer temporal windows. Native frame sizes are 512^2 for ADT, 360×640 for PS, 720×1280 for DR, and 540×960 for PO; the latter three have aspect ratio 16:9. Predictions are read at query pixels scaled to each method's output grid. Queries outside a crop use the nearest prediction, except for MotionCrafter on ADT.

Table 5: **Baseline configurations.** Input grids are $H \times W$; settings apply to the accuracy comparisons. Runtime measurements use the settings in Appendix D.4.

Method	PS / DR / PO	ADT	Inference settings
4RC (Luo et al., 2026)	280×504	504^2	
Any4D (Karhade et al., 2025)	294×518	518^2	
V-DPM (Sucar et al., 2026)			
St4RTrack (Feng et al., 2025)	288×512	512^2	
TraceAnything (Liu et al., 2026)			
MotionCrafter (Zhu et al., 2026)	320×640	320×640	Centre crop.
OpenD4RT (RHOS Team, 2026)	256^2	256^2	Resize, no crop; stride-2, 128^2 grid.
TrackCraft3R (Nam et al., 2026)	480×832	480×832	Interleaved 12-frame clips, stride 8.
Track4World (Lu et al., 2026)	280×504	512^2	
DELTAv2 (Ngo et al., 2025a)	280×504	512^2	Dense mode.

Evaluation exceptions. MotionCrafter is scored on the 63.6% of ADT queries inside its crop.[1] Affected ADT scores and benchmark means are marked ‡. Sim(3) handles the non-metric scale of OpenD4RT and MotionCrafter, and MotionCrafter's coordinate-frame mismatch with the first camera. Scale-only alignment cannot correct that mismatch.

C Additional Quantitative Results

C.1 Dense Tracking on Kubric

Table 6 tests dense tracking and sensitivity to geometry providers, complementing WorldTrack's sparse queries. We use the alignments and WorldTrack baseline settings in Appendices B.1 and B.2.

Under RANSAC Sim(3), our DA3 and VGGT-Ω configurations trail 4RC but outperform methods using the same geometry on all-point APD3D. With 4RC geometry, both tracker variants approach 4RC's all-point accuracy and exceed its dynamic-point accuracy. Under global-scale alignment, our CoWTracker variant with 4RC geometry leads all three reported metrics. These results show that transfer across providers remains sensitive to the quality of the input geometry.

C.2 Visibility Metrics and End-Point Error

Tables 7 and 8 complement APD3D with visibility-aware accuracy and mean position error. Our CoWTracker variant with 4RC geometry achieves the highest mean AJ and lowest mean EPE, while

[1] With all ADT queries: APD3D (All/Dyn) is 60.98/78.51 under Sim(3) and 2.83/5.81 under global scale; EPE is 0.610 m under Sim(3); APD2D is 12.33/5.98.

Table 6: **Dense tracking on Kubric.** TrackCraft3R's (Nam et al., 2026) 50 24-frame Kubric (Greff et al., 2022) clips, with ground-truth trajectories for every reference pixel. APD3D (%) on all and dynamic tracks and EPE (scene units) under RANSAC Sim(3) and global-scale alignment; bold best, underline second best.

	RANSAC Sim(3)			Global scale		
Method	All	Dyn	EPE	All	Dyn	EPE
4RC	**85.33**	78.81	**0.299**	<u>82.40</u>	73.68	<u>0.342</u>
Any4D	65.31	55.51	0.854	42.69	39.77	1.107
MotionCrafter	52.98	47.89	1.765	–	–	–
OpenD4RT	38.40	37.92	1.432	24.32	26.92	1.801
St4RTrack	59.45	54.24	3.640	56.56	51.77	3.585
Track4World	53.55	55.53	2.074	37.37	38.82	2.552
TraceAnything	59.29	50.40	1.023	30.46	28.02	1.252
V-DPM	75.19	**81.03**	3.369	63.70	75.18	3.405
DELTAv2 + DA3	63.27	62.05	1.659	44.73	43.37	2.028
DELTAv2 + VGGT-Ω	62.35	59.19	5.048	41.19	38.52	5.888
TrackCraft3R + DA3	62.76	55.72	1.539	41.42	38.08	2.088
TrackCraft3R + VGGT-Ω	59.27	50.48	4.059	34.81	31.85	6.352
Ours: the corrector on the lift of						
AllTracker + DA3	67.10	61.32	1.794	46.72	44.01	2.144
AllTracker + VGGT-Ω	65.89	59.66	4.999	43.45	39.37	5.813
AllTracker + 4RC	84.28	79.79	0.336	81.99	<u>76.23</u>	0.347
CoWTracker + DA3	67.29	61.40	1.773	47.07	44.33	2.128
CoWTracker + VGGT-Ω	66.57	59.76	4.863	43.54	39.82	5.659
CoWTracker + 4RC	<u>84.88</u>	<u>80.66</u>	<u>0.324</u>	**83.04**	**77.20**	**0.335**

Table 7: **AJ and OA on WorldTrack (%).** Predicted visibility is thresholded at 0.5. AJ uses all-point RANSAC Sim(3) alignment; OA requires no geometric alignment. Higher is better.

	ADT		PS		PO		DR		Mean	
Method	AJ	OA	AJ	OA	AJ	OA	AJ	OA	AJ	OA
OpenD4RT	58.82	91.04	69.01	87.34	45.52	**78.63**	63.04	90.03	59.10	86.76
Track4World	72.88	<u>98.06</u>	64.76	86.51	47.53	76.25	67.46	92.79	63.16	<u>88.40</u>
DELTAv2 + DA3	84.18	96.94	55.18	71.57	44.90	74.35	62.74	87.36	61.75	82.56
DELTAv2 + VGGT-Ω	83.83	95.09	57.91	71.32	45.61	74.21	68.54	87.94	63.97	82.14
TrackCraft3R + DA3	82.95	94.26	70.18	88.13	50.29	75.59	66.79	92.69	67.55	87.67
TrackCraft3R + VGGT-Ω	83.48	93.48	72.40	<u>88.29</u>	49.72	75.64	72.68	92.92	69.57	87.58
Ours: the corrector on the lift of										
AllTracker + DA3	89.18	97.82	70.88	86.56	51.64	74.49	67.33	<u>93.02</u>	69.76	87.97
AllTracker + VGGT-Ω	89.75	97.82	76.17	86.56	52.21	74.49	74.27	<u>93.02</u>	73.10	87.97
AllTracker + 4RC	85.76	97.82	74.60	86.56	<u>58.79</u>	74.49	<u>78.17</u>	<u>93.02</u>	74.33	87.97
CoWTracker + DA3	<u>89.99</u>	**98.51**	74.34	**90.03**	54.10	<u>77.06</u>	67.42	**93.03**	71.46	**89.66**
CoWTracker + VGGT-Ω	**90.54**	**98.51**	**80.13**	**90.03**	54.53	<u>77.06</u>	74.42	**93.03**	<u>74.91</u>	**89.66**
CoWTracker + 4RC	86.73	**98.51**	<u>77.93</u>	**90.03**	**61.58**	<u>77.06</u>	**78.34**	**93.03**	**76.14**	**89.66**

CoWTracker with VGGT-Ω gives the lowest EPE on both real-world datasets, ADT and PS. Our CoWTracker variants also lead mean OA. Visibility accuracy comes from the 2D tracker, whose scores we return unchanged (Sec. 3.3); the corrector refines the positions used by AJ and EPE.

C.3 Global-Scale Alignment

Table 9 evaluates the predictions of Table 1 under global-scale alignment, fitting only scale without adjusting rotation or translation. This tests whether our conclusions depend on the additional freedom of Sim(3) alignment. Our configurations using 4RC geometry still lead mean APD3D on both all and dynamic points. Correction improves all six lifts, adding 1.4–2.6 points on all points and 2.6–3.8 on dynamic points. Mean accuracy decreases across methods, with substantial drops on PS, where V-DPM now leads. The mean-accuracy advantage and correction gains therefore persist, while dataset rankings depend on the alignment used.

Table 8: **End-point error on WorldTrack.** Mean 3D position error over evaluated point–frames, in metres after all-point RANSAC Sim(3) alignment. Lower is better. ‡Restricted ADT coverage for MotionCrafter, also affecting its mean (Appendix B.2).

Method	ADT	PS	PO	DR	Mean
4RC	0.153	0.138	0.249	0.149	0.172
Any4D	0.286	0.216	0.335	0.191	0.257
MotionCrafter	0.474‡	0.358	0.820	0.278	0.482‡
OpenD4RT	0.317	0.161	0.390	0.252	0.280
St4RTrack	0.242	0.258	0.311	0.236	0.262
Track4World	0.269	0.184	0.366	0.219	0.260
TraceAnything	0.198	0.284	0.735	0.378	0.399
V-DPM	0.141	0.116	**0.179**	0.189	0.156
DELTAv2 + DA3	0.330	0.160	0.516	0.270	0.319
DELTAv2 + VGGT-Ω	0.333	0.125	0.595	0.203	0.314
TrackCraft3R + DA3	0.106	0.161	0.348	0.245	0.215
TrackCraft3R + VGGT-Ω	0.100	0.139	0.412	0.181	0.208
Ours: the corrector on the lift of					
AllTracker + DA3	0.089	0.151	0.344	0.242	0.207
AllTracker + VGGT-Ω	0.087	0.109	0.372	0.169	0.184
AllTracker + 4RC	0.114	0.119	0.266	**0.145**	0.161
CoWTracker + DA3	0.087	0.139	0.311	0.241	0.195
CoWTracker + VGGT-Ω	**0.085**	**0.095**	0.347	0.168	0.174
CoWTracker + 4RC	0.110	0.110	0.231	**0.145**	**0.149**

Table 9: **WorldTrack under global-scale alignment.** APD3D (%) on all (All) and dynamic (Dyn) points, fitting only scale (Appendix B.1). MotionCrafter's coordinate-frame mismatch makes its scores non-comparable under this alignment. ‡Restricted ADT coverage for MotionCrafter, also affecting its means (Appendix B.2).

	ADT		PS		PO		DR		Mean	
Method	All	Dyn	All	Dyn	All	Dyn	All	Dyn	All	Dyn
4RC	83.09	68.24	69.15	69.15	80.23	76.44	82.95	73.13	78.86	71.74
Any4D	57.60	65.80	63.39	63.39	66.16	66.89	78.15	68.16	66.32	66.06
MotionCrafter	2.12‡	6.91‡	0.00	0.00	0.11	0.02	1.59	1.39	0.96‡	2.08‡
OpenD4RT	69.28	67.96	76.87	76.87	65.19	70.39	74.80	76.25	71.53	72.87
St4RTrack	74.33	71.13	69.12	69.12	67.86	68.34	74.10	67.31	71.35	68.97
Track4World	75.63	85.43	71.39	71.39	63.93	65.61	70.74	68.12	70.42	72.64
TraceAnything	73.64	67.34	67.60	67.60	38.44	44.02	60.85	58.59	60.13	59.39
V-DPM	84.41	69.61	**78.40**	**78.40**	81.43	**82.19**	75.56	**78.13**	79.95	77.08
DELTAv2 + DA3	78.63	85.85	65.05	65.05	54.83	56.57	69.41	57.51	66.98	66.25
DELTAv2 + VGGT-Ω	79.43	85.12	70.50	70.50	58.61	60.48	76.19	68.11	71.18	71.05
TrackCraft3R + DA3	89.12	84.04	63.91	63.91	65.51	62.25	71.26	60.14	72.45	67.59
TrackCraft3R + VGGT-Ω	90.46	83.55	68.13	68.13	67.72	64.21	77.69	70.43	76.00	71.58
Ours: the corrector on the lift of										
AllTracker + DA3	90.71	91.40	64.80	64.80	70.22	70.71	72.49	61.92	74.55	72.21
AllTracker + VGGT-Ω	91.93	88.14	71.38	71.38	72.65	73.10	81.15	74.68	79.28	76.82
AllTracker + 4RC	89.22	90.71	71.31	71.31	80.42	76.78	**84.63**	75.18	81.40	78.49
CoWTracker + DA3	90.73	**91.90**	65.93	65.93	72.12	72.93	72.33	62.17	75.28	73.23
CoWTracker + VGGT-Ω	**92.32**	88.82	72.49	72.49	74.15	74.62	80.82	74.10	79.94	77.51
CoWTracker + 4RC	89.08	91.30	72.96	72.96	**82.27**	78.82	84.12	75.80	**82.11**	**79.72**

C.4 Image-Plane Accuracy

Table 11 complements the 3D evaluation with the Average Percentage of 2D Points within Delta (APD2D), averaging accuracy over $\{1, 2, 4, 8, 16\}$-pixel thresholds after rescaling coordinates to 256×256. We use the 84-frame clips, initial queries, All/Dyn definitions and aggregation of Appendix B.1.

Table 10: **Camera sources for APD2D.** Cameras are predicted, supplied by a geometry provider, or recovered from model outputs; no ground-truth camera parameters are used.

Methods	Projection cameras
4RC, Any4D, Track4World, V-DPM	Predicted by the method
TrackCraft3R, DELTAv2	Input geometry cameras
Our configurations	Predicted by the geometry provider
MotionCrafter	Robust joint pose/focal fit to predicted point maps
St4RTrack, OpenD4RT, TraceAnything	Recovered from each method's own outputs

Table 11: **Image-plane accuracy on WorldTrack.** APD2D (%) on all (All) and dynamic (Dyn) points. 3D trajectories are projected through the cameras in Table 10; native 2D references are scored without projection. Bold and underline rank 3D methods only. Scoring details are in Appendix C.4. ‡Restricted ADT coverage for MotionCrafter, also affecting its means (Appendix B.2).

	ADT		PS		PO		DR		Mean	
Method	All	Dyn	All	Dyn	All	Dyn	All	Dyn	All	Dyn
4RC	62.06	27.36	51.79	51.79	71.50	60.91	79.15	65.26	66.13	51.33
Any4D	40.90	23.60	31.46	31.46	56.99	44.39	76.93	50.10	51.57	37.39
MotionCrafter	17.47‡	10.34‡	15.11	15.11	21.88	13.70	40.69	30.58	23.79‡	17.43‡
OpenD4RT	45.94	21.77	46.90	46.90	56.13	49.43	72.58	57.41	55.39	43.88
St4RTrack	59.53	28.47	39.12	39.12	63.95	47.47	79.77	50.73	60.59	41.45
Track4World	76.31	**66.42**	63.67	63.67	<u>76.82</u>	**68.95**	**89.73**	**74.30**	76.63	**68.34**
TraceAnything	23.13	11.51	32.30	32.30	35.29	28.14	59.93	41.65	37.66	28.40
V-DPM	50.15	22.88	55.04	55.04	**77.00**	64.66	85.10	64.42	66.82	51.75
DELTAv2 + DA3	76.27	<u>65.59</u>	66.20	66.20	72.94	61.94	87.28	66.65	75.68	65.09
DELTAv2 + VGGT-Ω	74.75	62.12	66.15	66.15	72.52	62.32	87.52	67.33	75.23	64.48
TrackCraft3R + DA3	48.89	20.04	42.44	42.44	63.29	45.66	75.85	54.15	57.62	40.57
TrackCraft3R + VGGT-Ω	52.14	22.21	41.59	41.59	61.10	44.52	77.27	55.37	58.02	40.92
Ours (2D tracker + geometry provider)										
AllTracker + DA3	76.59	60.36	63.41	63.41	74.51	66.05	86.35	70.14	75.22	64.99
AllTracker + VGGT-Ω	76.47	59.19	63.56	63.56	74.33	65.92	86.67	70.88	75.26	64.89
AllTracker + 4RC	76.45	60.84	63.66	63.66	74.68	66.29	86.69	71.00	75.37	65.45
CoWTracker + DA3	**76.98**	62.64	67.18	67.18	76.31	67.41	88.08	72.15	<u>77.14</u>	67.35
CoWTracker + VGGT-Ω	76.78	61.43	<u>67.23</u>	<u>67.23</u>	76.11	67.28	88.35	72.88	77.12	67.21
CoWTracker + 4RC	<u>76.91</u>	63.23	**67.35**	**67.35**	76.41	<u>67.62</u>	<u>88.39</u>	<u>72.93</u>	**77.27**	<u>67.78</u>
Native 2D references (no 3D prediction)										
AllTracker (2D tracker alone)	77.18	67.90	64.29	64.29	75.00	68.68	88.31	73.55	76.19	68.60
CoWTracker (2D tracker alone)	78.32	71.22	68.59	68.59	77.44	70.58	90.67	76.00	78.75	71.60

Evaluation protocol. We project 3D predictions through their associated cameras without ground-truth alignment. Projections are mapped to original image coordinates, accounting for resizing and crop offsets, before the common rescaling. Point–frames behind the ground-truth camera are excluded; predictions behind their camera or with non-finite projections count as misses at every threshold. Native AllTracker and CoWTracker predictions are scored directly in 2D at the same queries, with the same scaling and ground-truth masks. Track4World and DELTAv2 use inputs capped to our grid, with ADT uncapped (Appendix B.2). PS has identical All and Dyn scores because all its queries meet the dynamic-track criterion.

Cameras for APD2D. Table 10 lists the projection cameras. Our method, TrackCraft3R and DELTAv2 use their geometry provider's predicted intrinsics and extrinsics, following the shared estimated-geometry setup in Sec. 4.1. For methods without camera outputs, cameras are recovered from their own predictions without ground-truth camera parameters.

Results. Native CoWTracker achieves the highest APD2D in every reported column, outperforming projected 3D trajectories and motivating our use of strong pretrained 2D correspondences (Sec. 4.2). Our CoWTracker variant with 4RC geometry leads mean All APD2D among 3D methods and ranks second on Dyn, behind Track4World. Switching from AllTracker to CoWTracker improves every dataset and geometry pairing, adding 1.86–1.92 points to mean All and 2.32–2.36 to mean Dyn. For either tracker, the means vary across geometry providers by only 0.15 (All) and at most 0.57 (Dyn) points. Tracker choice therefore has a larger effect on APD2D here.

Relation to 3D accuracy. All six of our configurations remain below their native tracker in both means. With CoWTracker and 4RC geometry, the gaps are 1.48 points on All and 3.82 on Dyn. This comparison evaluates the complete pipeline; it neither isolates the corrector's effect nor separates visible and occluded points.

APD2D evaluates trajectories together with their projection cameras, including any camera recovery. It cannot measure range errors along a valid viewing ray: lifting and reprojecting through the same camera returns the input pixel. The small variation across geometry providers therefore does not imply similar 3D accuracy. The APD3D gains over initialized lifts in Fig. 3 provide complementary evidence of improved 3D estimation.

D Ablations and Efficiency

D.1 Training Ablation Settings

Simulated versus provider-generated lifts. Provider-generated lifts use AllTracker with DA3 or VGGT-Ω. To account for differences in coordinate frame and scale, we align ground-truth frame-0 positions $\mathbf{x}_{0i}$ to lifted positions $\mathbf{y}_{0i}$ using a similarity transform $\alpha R\mathbf{x}_{0i} + \mathbf{b}$. Valid points have positive frame-0 range. We first fit the transform by least squares on all valid points, then perform three refits to reduce the influence of outliers. After each fit, we compute the relative residual

$$e_i = \frac{\|\alpha R\mathbf{x}_{0i} + \mathbf{b} - \mathbf{y}_{0i}\|}{\|\mathbf{y}_{0i}\|}.$$

The next fitting set comprises valid points satisfying $e_i \le \max(3\,\mathrm{median}_{j\in K} e_j,\, 0.01)$, where K is the set used for the current fit. The 1% floor limits trimming when residuals are already low. After the final fit, we recompute the median residual over all valid points, including those excluded from fitting. We retain windows with at least 256 valid points, positive fitted scale, and a finite median residual below 0.2.

We apply this single transform to the ground truth at every frame, leaving the input lift unchanged. There is no per-frame refitting or additional alignment inside the loss, so errors in the estimated camera motion remain part of the correction task. After filtering, the runs trained on AllTracker+DA3 and AllTracker+VGGT-Ω lifts use 3,890 and 3,837 training windows, respectively. The latter must pass both providers' filters.

These models improve training-scene APD3D on all points by 14.14 and 8.33 percentage points over their corresponding lifts, respectively.

We train an additional simulated-lift model on the same 3,890 filtered windows as the model trained on AllTracker+DA3 lifts, with the same architecture, optimization settings and 150,000 training steps. Figure 6 shows that simulated training still improves APD3D on every dataset and outperforms training on AllTracker+DA3 lifts with both test geometry providers.

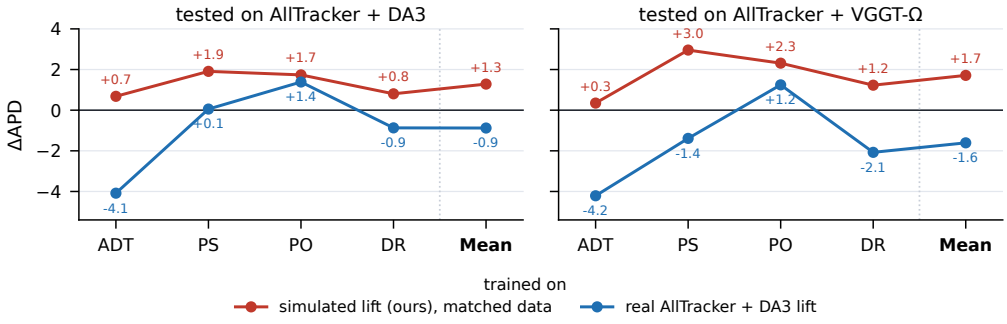


Figure 6: **Simulated versus AllTracker+DA3 training lifts with matched data.** Both models train on the same 3,890 filtered windows for 150,000 steps. APD3D gains on all points are reported in percentage points over the initialized lift, with AllTracker and DA3 (left) or VGGT-Ω (right). "Mean" averages the four datasets.

Corruption families. The variants in Table 2 are retrained with the same data, network, loss and optimization settings, removing one corruption family or retaining only white noise among the added

Table 12: **Contribution of training corruptions on occluded points.** APD3D (%) on ground-truth-occluded point–frames, averaged over the four datasets. Scoring retains the all-point RANSAC Sim(3) alignment. Parentheses report percentage-point gains over the initialized lift.

	AllTracker		CoWTracker	
Training corruptions	DA3	VGGT-Ω	DA3	VGGT-Ω
Lift (the corrector's input)	71.9	73.7	70.0	71.6
Full simulation (ours)	**75.7 (+3.8)**	**79.0 (+5.2)**	**77.1 (+7.1)**	**80.8 (+9.2)**
without correlated noise	75.3 (+3.4)	78.7 (+5.0)	76.7 (+6.7)	80.2 (+8.6)
without wrong surface	72.7 (+0.8)	75.3 (+1.5)	71.9 (+1.9)	74.0 (+2.4)
without synthetic gaps	75.6 (+3.7)	78.9 (+5.1)	77.0 (+7.0)	80.6 (+8.9)
white noise only	72.3 (+0.4)	74.8 (+1.1)	71.4 (+1.4)	73.3 (+1.7)

perturbations. Natural PointOdyssey occlusions remain in all variants, so completion supervision is present even without synthetic gaps.

Table 12 compares these variants on ground-truth-occluded point–frames. Full simulation achieves slightly higher mean APD3D than the variant without synthetic gaps across all four tracker–geometry combinations. This comparison supports a modest benefit from adding synthetic gaps to the occlusions already present in the training data. Removing synthetic occluders also removes their false-visible observations at the occluder's depth, so this ablation does not isolate the effect of missing observations alone.

D.2 Ray Masking

Table 13 evaluates the ray-masking choice in Sec. 3.1. Masking improves ADT APD3D by about 2.8 points overall and 8.1 on occluded points with either geometry provider. Changes elsewhere are below one point and mixed.

Table 13: **Ray masking.** APD3D (%) for separately trained correctors that mask or retain rays where the input mask is zero, using AllTracker and $C = 4$. All and Occ share the all-point RANSAC Sim(3) alignment; Occ scores ground-truth-occluded point–frames.

		ADT		PS		PO		DR	
Geometry	Subset	With masking	Without masking	With masking	Without masking	With masking	Without masking	With masking	Without masking
DA3	All	93.69	90.93	85.46	85.44	76.04	76.03	79.50	79.52
DA3	Occ	90.27	82.22	76.08	76.30	73.19	73.03	63.17	64.05
VGGT-Ω	All	94.04	91.23	90.40	90.36	76.10	76.10	85.74	85.75
VGGT-Ω	Occ	90.12	82.03	82.04	82.54	72.49	72.23	71.26	71.85

D.3 Chunk Size and Group Size

Table 14 compares separately trained correctors with $C \in \{1, 4, 8\}$, each trained for 150,000 steps. Our setting, $C = 4$, gives the highest accuracy; $C = 8$ reduces correction time by about a third at a cost of at most 0.3 APD3D points. Without chunking ($C = 1$), correction is slower and less accurate under this training setup.

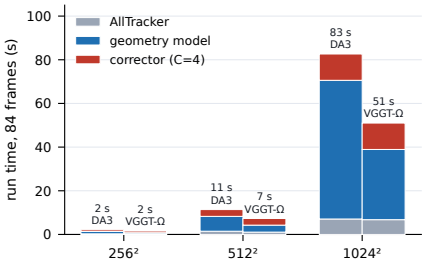
Unlike chunk size, group size varies only at inference, with the checkpoint fixed (Fig. 7). Smaller groups reduce accuracy, particularly on dynamic points; increasing P from 1024 to 4096 changes every score by at most 0.2 points. We therefore use $P = 1024$. For M total tracks and N temporal chunks, the across-track attention cost in Sec. 3.3 sums to $O(MNP)$ over all groups, increasing linearly with P. Complete pipeline timings are reported in Appendix D.4.

D.4 Run Time and Memory

Table 15 reports complete pipeline costs for 84 frames from one Dynamic Replica clip at three nominal resolutions on a 96 GB RTX PRO 6000 GPU, using each method's released inference code. Timing includes all models (the 2D tracker, geometry model and corrector for ours) and runs from

Table 14: **Temporal chunking and pipeline runtime.** AllTracker with DA3 or VGGT-Ω geometry. **Left:** correctors trained separately for each chunk size C (frames per token), reporting mean WorldTrack APD3D (%) under RANSAC Sim(3) alignment and corrector-only runtime at 512^2. **Right:** pipeline runtime by stage across resolutions, using our default $C = 4$. All timings use 84 frames on one RTX PRO 6000; profiling details appear in Appendix D.4.

	DA3 input		VGGT-Ω input		Corrector time (s)	
C	All	Dyn	All	Dyn	DA3	VGGT-Ω
1	82.93	81.08	85.83	83.79	10.0	9.7
4	83.68	81.74	86.57	84.92	3.1	3.0
8	83.59	81.48	86.48	84.75	2.0	1.9



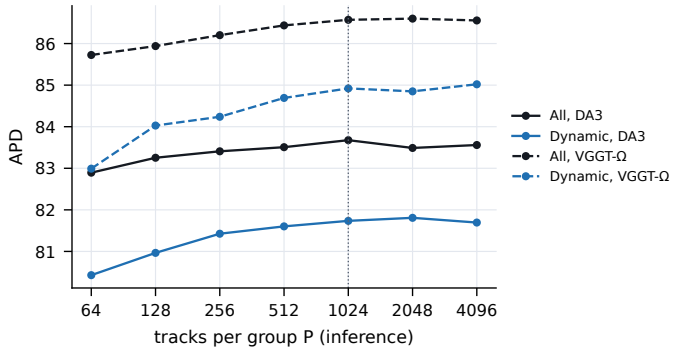


Figure 7: **Group size at inference.** Mean WorldTrack APD3D (%) on all and dynamic points under RANSAC Sim(3) alignment, varying tracks per group P with AllTracker and fixed corrector weights ($C = 4$). The dotted line marks our setting, $P = 1024$.

decoded frames in host memory to dense trajectories returned to host memory; video decoding and model loading are excluded. At 256^2 and 512^2, we report the median of three runs after one warm-up; 1024^2 uses a single run without warm-up. For separately profiled stages, we sum their median times. Memory is the maximum GPU allocation across stages and timed runs. Input sizes follow each method's patch-size rounding, with resolution overrides marked in the table.

Our VGGT-Ω configurations offer competitive WorldTrack accuracy at substantially lower runtime than 4RC and V-DPM, while faster baselines are less accurate (Fig. 1). The stage breakdown in Table 14 (right) connects these measurements to grouped inference (Sec. 3.3). With AllTracker, correction takes roughly 3 s at 512^2 and 12 s at 1024^2, consistent with linear scaling in pixel count. Geometry estimation grows faster, so the complete pipeline does not share this scaling. Our AllTracker configurations with DA3 or VGGT-Ω remain feasible at 1024^2, where 4RC, V-DPM and DELTAv2 run out of memory.

Table 15: **Run time and peak allocated GPU memory.** Complete pipelines at three nominal resolutions. OOM: out of memory; FAIL: St4RTrack's upsampling exceeds an implementation tensor-size limit; –: not run. †Hard-coded input resolution overridden for this experiment. OpenD4RT runs only at 256^2; ours uses $C = 4$.

	Run time (s)			Peak memory (GiB)		
Method	256^2	512^2	1024^2	256^2	512^2	1024^2
Ours: the corrector on the lift of						
AllTracker + DA3	2.1	11.4	82.6	10.1	21.3	63.7
AllTracker + VGGT-Ω	1.5	7.3	51.0	7.9	12.5	43.9
CoWTracker + DA3	2.8	15.3	OOM	10.6	36.5	OOM
CoWTracker + VGGT-Ω	2.2	11.7	OOM	10.6	36.5	OOM
4RC	4.3	31.8	OOM	11.0	26.8	OOM
V-DPM†	77.3	590.8	OOM	26.9	85.7	OOM
TrackCraft3R + DA3 / VGGT-Ω	5.0 / 4.4	22.1 / 18.5	162.7 / 132.3	9.8 / 7.6	21.0 / 12.2	63.4 / 30.9
DELTAv2 + DA3 / VGGT-Ω	7.8 / 7.2	33.4 / 29.8	OOM / OOM	9.8 / 8.3	32.5 / 32.5	OOM / OOM
Any4D†	0.6	3.2	OOM	13.7	45.2	OOM
TraceAnything†	1.9	5.2	20.4	3.0	4.8	11.8
Track4World	1.5	8.4	85.7	10.1	21.1	70.2
St4RTrack	0.6	2.9	FAIL	6.8	20.8	FAIL
MotionCrafter†	4.8	20.4	OOM	10.9	29.1	OOM
OpenD4RT (256^2 only)	284.2	–	–	7.8	–	–

E Additional Qualitative Results

Additional method comparisons. Figure 8 adds two DAVIS (Perazzi et al., 2016) examples. In *soccerball*, which includes occlusion by a tree trunk, our trajectories form cleaner arcs ending on the ball. In *lady-running*, ours follow the body in an ordered band, while TrackCraft3R's scatter below it; 4RC and V-DPM also produce compact trajectories.

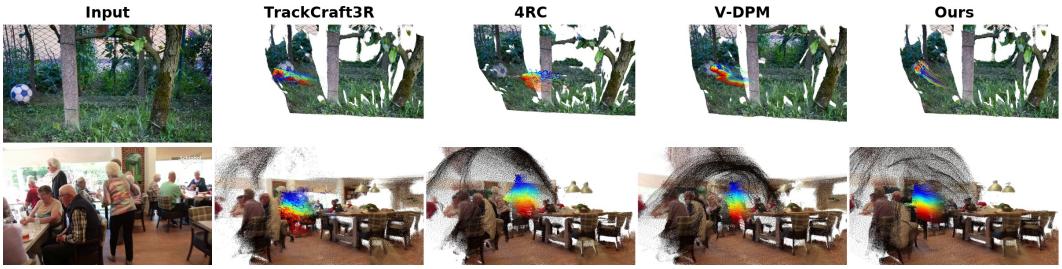


Figure 8: **Additional qualitative comparisons.** Top: DAVIS *soccerball*; bottom: *lady-running*. Ours uses CoWTracker with VGGT-Ω (top) and AllTracker with DA3 (bottom), with the same corrector weights ($C = 4$). TrackCraft3R uses VGGT-Ω in both rows (unlike ours in *lady-running*); 4RC and V-DPM use their own reconstructions. Methods track the same reference pixels, shown on their respective reconstructions from matched novel viewpoints.

Isolating trajectory repair. Figure 9 isolates the effect of correction on the initialized lift (Sec. 3.2). In *camel*, correction reduces trajectory outliers near the ground and produces more coherent paths along the body. In *skate-park*, scattered trajectories become a more coherent bundle. These comparisons complement the quantitative gains over the initialized lift in Fig. 3.

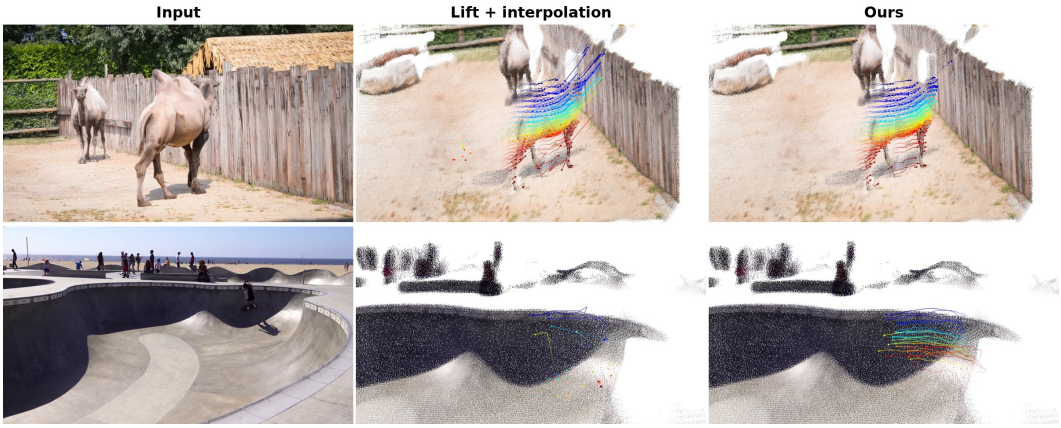


Figure 9: **Effect of trajectory correction.** Left to right: input frame, initialized lift, and corrected trajectories ($C = 4$). Top: DAVIS *camel*; bottom: *skate-park*. CoWTracker's 2D tracks and DA3 geometry are held fixed; both trajectory panels use the same point cloud and novel viewpoint.

Visualization details. Figures 4, 8 and 9 use queries spaced 16 pixels apart at native resolution inside the eroded first-frame object mask. Colors identify queries in raster order; trails show the current and preceding 20 frames, with dots marking current positions. Backdrops combine confident background points from all frames, outside the dilated object mask, with object points from the displayed frame. The method comparisons show each clip's third-to-last frame. 4RC and V-DPM outputs are read at the nearest pixels of their grids; TrackCraft3R uses its interleaved benchmark mode with queries lifted using VGGT-Ω depth. To match visualization viewpoints across methods, a novel camera in VGGT-Ω coordinates is mapped into each reconstruction by a similarity fitted to first-frame point maps, trimming 20% outliers over two rounds.