

C4D: 4D Made from 3D through Dual Correspondences

Shizun Wang¹ Zhenxiang Jiang¹ Xingyi Yang² Xinchao Wang^{1*}

¹National University of Singapore ²The Hong Kong Polytechnic University

{shizun.wang, zhenxiang.jiang}@nus.u.edu, xingyi.yang@polyu.edu.hk, xinchao@nus.edu.sg




Figure 1. Given a monocular video that contains both camera movement and object movement, **C4D** can recover the dynamic scene in 4D, including per-frame dense point cloud, camera poses and intrinsic parameters. Video depth, motion masks, and point tracking in both 2D and 3D space are also available in the outputs.

Abstract

Recovering 4D from monocular video, which jointly estimates dynamic geometry and camera poses, is an inevitably challenging problem. While recent pointmap-based 3D reconstruction methods (e.g., DUST3R) have made great progress in reconstructing static scenes, directly applying them to dynamic scenes leads to inaccurate results. This discrepancy arises because moving objects violate multi-view geometric constraints, disrupting the reconstruction. To address this, we introduce **C4D**, a framework that leverages temporal **C**orrespondences to extend existing 3D reconstruction formulation to **4D**. Specifically, apart from predicting pointmaps, C4D captures two types of correspondences: short-term optical flow and long-term point tracking. We train a dynamic-aware point tracker that provides additional mobility information, facilitating the estimation of motion masks to separate moving elements from the static background, thus offering more reliable guidance for dynamic scenes. Furthermore, we introduce a set of dynamic scene optimization objectives to recover per-frame 3D geometry and camera parameters. Simultaneously, the correspondences lift 2D trajectories into smooth 3D trajec-

ories, enabling fully integrated 4D reconstruction. Experiments show that our framework achieves complete 4D recovery and demonstrates strong performance across multiple downstream tasks, including depth estimation, camera pose estimation, and point tracking. Project Page: <https://littlepure2333.github.io/C4D>

1. Introduction

Recovering complete 4D representations from monocular videos, which involves estimating dynamic scene geometry, camera poses, and 3D point tracking, is a highly challenging task. While extending 3D reconstruction methods over the time dimension might seem straightforward, achieving accurate and smooth time-varying geometries and consistent camera pose trajectories is far from simple.

Recent paradigm shifts in 3D reconstruction, such as DUST3R [58], have shown significant success in reconstructing static scenes from unordered images. It directly predicts dense 3D pointmaps from images and makes many 3D downstream tasks, like recovering camera parameters and global 3D reconstruction, become easy by just applying global alignment optimization on 3D pointmaps.

However, when applied to dynamic scenes, these for-

*Corresponding author.

mulations often produce substantial inaccuracies. This is because their reliance on multi-view geometric constraints breaks down as moving objects violate the assumptions of global alignment. As a result, they struggle to achieve accurate 4D reconstructions in dynamic scenes.

Our key insight is that the interplay between temporal correspondences and 3D reconstruction naturally leads to 4D. By capturing 2D correspondences over time, we can effectively separate moving regions from static ones. By calibrating the camera in the static region only, we improve the quality of the 3D reconstruction. In turn, the improved 3D model helps connect these correspondences, creating a consistent 4D representation that integrates temporal details into the 3D structure.

This motivation drives our framework, **C4D**, a framework designed to upgrade the current 3D reconstruction formulation by using temporal *Correspondences* to achieve **4D** reconstruction. Apart from 3D pointmap prediction, C4D captures *short-term* optical flow and *long-term* point tracking. These temporal correspondences are essential: they generate motion masks that guide the 3D reconstruction process, while also contributing to optimizing the smoothness of the 4D representation.

To achieve this, we introduce the Dynamic-aware Point Tracker (DynPT), which not only tracks points but also predicts whether they are moving in the world coordinates. Using this information, we create a correspondence-guided strategy that combines static points and optical flow to generate motion masks. These motion masks guide the 3D reconstruction by focusing on static regions, enabling more accurate estimation of camera parameters from the point maps and further enhancing geometric consistency.

To further improve the 4D reconstruction, we introduce a set of correspondence-aided optimization techniques. These include ensuring the camera movements are consistent, keeping the camera path smooth, and maintaining smooth trajectories for the 3D points. Together, these improvements result in a refined and stable 4D reconstruction that is both accurate and smooth over time. Extensive experiments show that C4D delivers strong performance in dynamic scene reconstruction. When applied to various downstream tasks, such as depth estimation, camera pose estimation, and point tracking, C4D performs competitively, even compared to specialized methods.

In summary, our key contributions are as follows:

- We introduce C4D, a framework that upgrades the current 3D reconstruction formulation to 4D reconstruction by incorporating two temporal correspondences.
- We propose a dynamic-aware point tracker (DynPT) that not only tracks points but also predicts whether a point is dynamic in world coordinates.
- We present a motion mask prediction mechanism guided by optical flow and our DynPT.

- We introduce correspondence-aided optimization techniques to improve the consistency and smoothness of 4D reconstruction.
- We conduct experiments on depth estimation, camera pose estimation, and point tracking, demonstrating that C4D achieves strong performance, even compared to specialized methods.

2. Related Work

2.1. Temporal Correspondences

Optical flow represents dense pixel-level motion displacement between consecutive frames, capturing short-term dense correspondences. Modern deep learning methods have transformed optical flow estimation, leveraging large datasets [3, 36], CNNs [12, 50], ViTs [64], and iterative refinement [52, 61], resulting in significant improvements in accuracy and robustness. We leverage the motion information contained in optical flow to generate motion masks in this work. **Point tracking** aims to track a set of query points and predict their position and occlusion in a video [10], providing long-term sparse pixel correspondences. Tracking Any Point (TAP) methods [6, 11, 18, 23] extract correlation maps between frames and use a neural network to predict tracking positions and occlusions, achieving strong performance on causal videos. While these methods are effective, they all lack the ability to predict the mobility of points in world coordinates, which we achieve in this work.

2.2. 3D Reconstruction

Recovering 3D structures and camera poses from image collections has been studied for decades [19]. Classic methods such as Structure-from-motion (SfM) [44] and visual SLAM [9, 39] operate in sequential pipelines, often involving keypoint detection [2, 34, 35, 42], matching [43, 63], triangulation, and bundle adjustment [1, 55]. However, the sequential pipeline is complex and vulnerable to errors in each sub-task. To address these, DUST3R [58] introduces a significant paradigm shift by directly predicting pointmaps from image pairs, and dense 3D reconstruction can be obtained by a global alignment optimization.

2.3. 4D Reconstruction

Since the world is dynamic, 4D reconstruction naturally extends 3D reconstruction. Recent works [5, 7, 17, 27–29, 31, 33, 46, 49, 57, 59] explore 4D reconstruction from monocular video. Building on either 3DGS [26] or pointmap [58] representation, most of these methods are optimization-based and rely on off-the-shelf priors for supervision, such as depth, optical flow, and tracking trajectories. Concurrent work MonST3R [67] explores pointmap-based 4D reconstruction by fine-tuning DUST3R on dynamic scene data, whereas we directly use pretrained pointmap-based model

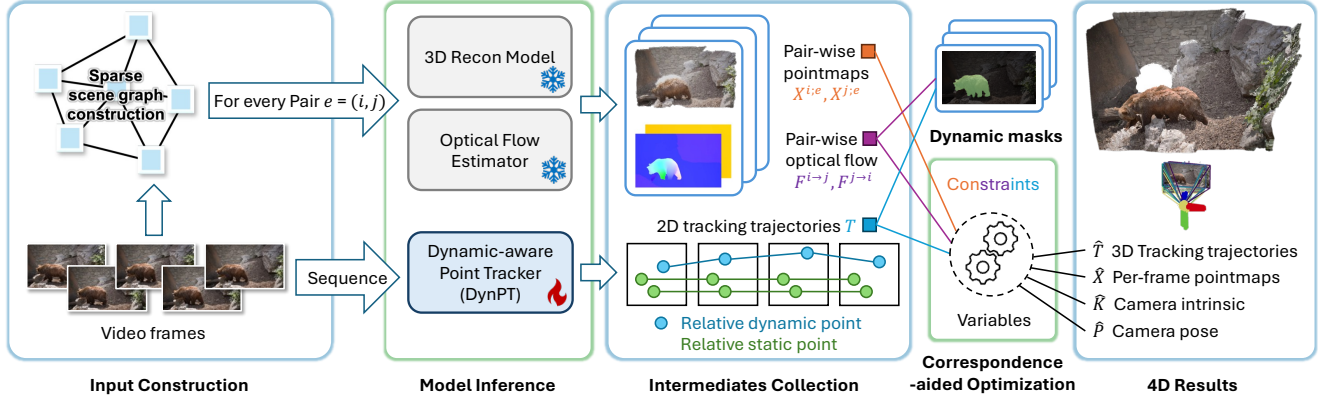


Figure 2. **Overview of C4D.** C4D takes monocular video as input and jointly predicts dense 3D pointmaps (Sec. 3.1) and temporal correspondences (Sec. 3.2), including short-term optical flow and long-term point tracking (Sec. 3.2.1). These correspondences are utilized to predict motion masks (Sec. 3.2.2) and participate in the optimization process (Sec. 3.3) with 3D pointmaps to obtain 4D outputs.

weights and complement them with correspondence-guided optimization for 4D reconstruction.

3. Method

The core idea of our method is to jointly predict dense 3D pointmaps and temporal correspondences from an input video, leveraging these correspondences to improve 4D reconstruction in dynamic scenes. These correspondences are obtained from both *short-term* optical flow and *long-term* point tracking. The whole pipeline is shown in Figure 2.

We begin by reviewing the 3D reconstruction formulation in Sec. 3.1, which provides dense 3D pointmaps. Next, we introduce our dynamic-aware point tracker (DynPT) in Sec. 3.2.1, designed to track points while also identifying whether they are dynamic in world coordinates. In Sec. 3.2.2, we describe how DynPT is combined with optical flow to estimate reliable motion masks. Finally, Sec. 3.3 details our correspondence-aided optimization, which utilizes pointmaps, optical flow, point tracks, and motion masks to refine the 4D reconstruction.

3.1. 3D Reconstruction Formulation

Our method complements the recent feed-forward 3D reconstruction paradigm, DUS3R [58], and can be applied to any DUS3R-based model weights [32, 67]. Given a video with T frames $\{I^1, I^2, \dots, I^T\}$, a scene graph \mathcal{G} is constructed, where an edge represents a pair of images $e = (I^n, I^m) \in \mathcal{G}$. Then DUS3R operates in two steps:

(1) A ViT-based network Φ that takes a pair of two images $I^n, I^m \in \mathbb{R}^{W \times H \times 3}$ as input and directly outputs two dense pointmaps $X^n, X^m \in \mathbb{R}^{W \times H \times 3}$ with associated confidence maps $C^n, C^m \in \mathbb{R}^{W \times H}$.

$$X^n, C^n, X^m, C^m = \Phi(I^n, I^m) \quad (1)$$

(2) Since these pointmaps are represented in the local

coordinate of each pair, DUS3R employs global optimization to all pairs of pointmaps, to recover *global aligned* pointmaps $\{\mathcal{X}^t \in \mathbb{R}^{W \times H \times 3}\}$ for all frames $t = 1, \dots, T$.

$$\mathcal{L}_{\text{GA}}(\mathcal{X}, P, \sigma) = \sum_{e \in \mathcal{G}} \sum_{t \in e} C^{t;e} \|\mathcal{X}^t - \sigma_e P_e X^{t;e}\| \quad (2)$$

Where $P_e \in \mathbb{R}^{3 \times 4}$ and $\sigma_e > 0$ are pairwise pose and scaling. To reduce computational cost, we use a sparse scene graph based on a strided sliding window, as in [13, 58, 67], where only pairs within a local temporal window are used for optimization.

While this 3D formulation performs well on static scenes, its performance drops in dynamic scenes, as discussed in Sec. 4.2. This is primarily due to moving objects violating multi-view geometric constraints, which motivates us to extend the current 3D formulation to a 4D one.

3.2. Capturing Dual Correspondences

We capture two correspondences to help 4D recovery: *long-term* point tracking and *short-term* optical flow.

3.2.1. Dynamic-aware Point Tracker

Current 2D point tracking methods like Tracking Any Point (TAP) [10, 11, 23, 24] can robustly track query points in videos. However, they cannot distinguish whether the movement of the tracking point is caused by camera movement or object movement. To segment moving objects in the *world coordinate* system, we enhance these trackers by enabling them to predict the mobility of tracking points. We introduce the **Dynamic-aware Point Tracker (DynPT)**, which differentiates between motion caused by the camera and true object dynamics. This helps us identify and segment moving objects even when both the camera and the objects are in motion.

Tracker Architecture We adopt a similar design of Co-Tracker [23, 24] to design our DynPT, which is illustrated




Figure 3. **Architecture of Dynamic-aware Point Tracker (DynPT).** For given video input and sampled initial query points, DynPT uses Transformer to iteratively update the tracks with features obtained from both 3D-aware ViT encoder and CNN.

in Figure 3. Original CoTracker only uses one CNN [20] to extract features. While to better capture the spatial dynamic relationships, we additionally employ a 3D-aware ViT encoder, which comes from DUST3R’s encoder, to enhance the 3D spatial information [62]. And different from all other TAP methods, DynPT directly predicts one more attribute, mobility, along with other attributes of tracks.

Specifically, for an input video of length T , DynPT first extracts each frame’s multi-scale features from a 3D-aware encoder and CNN, which are used to construct 4D correlation features $Corr$ to provide richer information for tracking [6]. Given a query point $P_0 \in \mathbb{R}^2$ at the first frame, we initialize the track positions P_t with the same position of P_0 for all remaining times $t = 1, \dots, T$, and initialize the confidence C_t , visibility V_t and mobility M_t with zeros for all times. Then we iteratively update these attributes with a transformer for M times. At each iteration, the transformer takes a grid of input tokens spanning time T : $G_t^i = (\eta_{t-1 \rightarrow t}^i, \eta_{t \rightarrow t+1}^i, C_t^i, V_t^i, M_t^i, Corr_t^i)$ for every query point $i = 1, \dots, N$, where $\eta_{t \rightarrow t+1}^i = \eta(P_{t+1} - P_t)$ represents Fourier Encoded embedding of per-frame displacements. Inside the Transformer, the attention operation is applied across time and track dimensions.

Training and Inference We train DynPT on Kubric [16], a synthetic dataset from which ground-truth mobility labels can be obtained. We use Huber loss to supervise position. And we employ cross-entropy loss to supervise confidence, visibility and mobility. When performing inference on a video, DynPT predicts tracks in a sliding window manner. More details about the DynPT can be found in the supplementary materials.

3.2.2. Correspondence-Guided Motion Mask Estimation

The most important part of 4D reconstruction in dynamic scenes is to separate dynamic areas from static areas in world coordinates. To achieve this, we utilize two temporal correspondences: short-term optical flow F_{est} estimated by off-the-shelf models [52, 61, 64], and long-term point tracking trajectory T predicted by DynPT. Figure 4 shows this strategy of correspondence-guided motion mask prediction.

Since DynPT provides mobility predictions of tracks, at

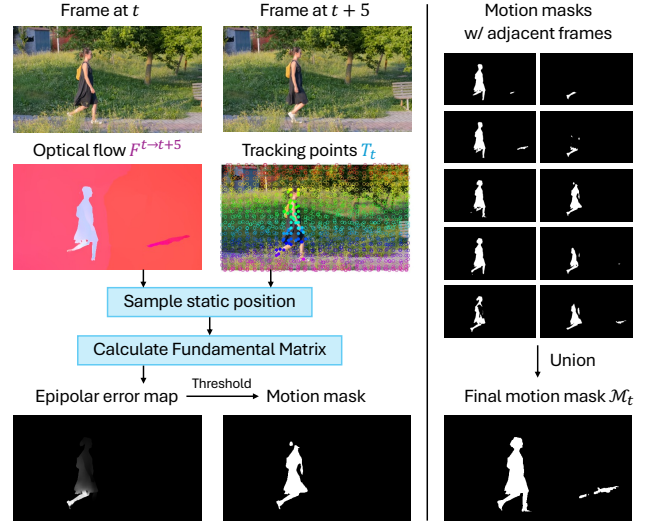


Figure 4. **Correspondence-guided motion mask prediction.** The solid circle indicates predicted dynamic points, the hollow circle indicates predicted static points. Adjacent frames are from constructed image pairs containing the current frame.

time t , we can retrieve the positions of static points $\{P_t^j\}$ where $M_t^j = 0$. And given an optical flow $F^{t \rightarrow t'}$ from time t to adjacent time t' , we can sample the pixel correspondences of these static points $\{(P_t^j, P_{t'}^j)\}$. With these correspondences, we then estimate the fundamental matrix \mathcal{F} between the two frames via the Least Median of Squares (LMedS) method [41], which does not require known camera parameters and is robust to outliers. Since the fundamental matrix is estimated solely based on static points, it reflects only the underlying camera motion, unaffected by dynamic objects in the scene. So using this \mathcal{F} to calculate the epipolar error map on all correspondences in $F^{t \rightarrow t'}$, the area with large error indicates there violates the epipolar constraints, that is, dynamic area. In practice, we compute the error map using the Sampson error [19], which provides a more robust approximation of the epipolar error by accounting for scale and orientation. Then a threshold is applied to obtain the motion mask.

While considering a longer temporal range, calculating the motion mask based on only two frames is not sufficient. For example, a person’s standing foot may remain still for several frames before lifting off to step, as shown in Figure 4. To address this, we calculate the motion mask of the current frame using adjacent frames from the constructed image pairs that include the current frame t , then take the union of these masks to produce the final motion mask \mathcal{M}_t .

3.3. Correspondence-aided Optimization for 4D

Based on the **Global Alignment (GA)** objective described in Sec 3.1, we introduce additional optimization objectives to improve the accuracy and smoothness in dynamic

scenes: *camera movement alignment*, *camera trajectory smoothness*, and *point trajectory smoothness*. The optimizable variables are per-frame depthmap D^t , camera intrinsic K^t and camera pose $P^t = [R^t|T^t]$. Then we re-parameterize the global pointmaps \mathcal{X}^t as $\mathcal{X}_{i,j}^t := P^{t-1}h(K^{t-1}[iD_{i,j}^t; jD_{i,j}^t; D_{i,j}^t])$, where (i, j) is the pixel coordinate and $h(\cdot)$ is the homogeneous mapping. So that, optimizing \mathcal{X}^t is equivalent to optimizing P^t, K^t, D^t .

Since global alignment tends to align moving objects to the same position, it can negatively impact camera pose estimation. To address this, and leveraging the fact that optical flow provides a prior on camera motion, we introduce the **Camera Movement Alignment (CMA)** objective [22, 51, 59, 67, 68]. CMA encourages the estimated ego motion to be consistent with optical flow in static regions. Specifically, for two frames I^t and $I^{t'}$, we compute the ego-motion field $F_{ego}^{t \rightarrow t'}$ as the 2D displacement of \mathcal{X}^t by moving camera from t to t' . Then we can encourage this field to be close to the optical flow field $\mathbf{F}^{t \rightarrow t'}$, in the static regions $\mathcal{S}_t \approx \mathcal{M}_t$, which is the opposite region of the motion mask \mathcal{M}_t computed in Sec 3.2.2:

$$\mathcal{L}_{\text{CMA}}(\mathcal{X}) = \sum_{e \in \mathcal{G}} \sum_{(t, t') \in e} \left\| \mathcal{S} \cdot \left(\mathbf{F}_{\text{ego}}^{t \rightarrow t'} - \mathbf{F}^{t \rightarrow t'} \right) \right\|_1 \quad (3)$$

The **Camera Trajectory Smoothness (CTS)** objective is commonly used in visual odometry [37, 47, 67] to enforce smooth camera motion by penalizing abrupt changes in camera rotation and translation between consecutive frames:

$$\mathcal{L}_{\text{CTS}}(\mathbf{X}) = \sum_{t=0}^N \left\| \mathbf{R}^t \mathbf{R}^{t+1} - \mathbf{I} \right\|_{\text{F}} + \left\| (\mathbf{T}^{t+1} - \mathbf{T}^t) \right\|_2 \quad (4)$$

where $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm, and \mathbf{I} is the identity matrix.

Lastly, we propose the **Point Trajectory Smoothness (PTS)** objective to smooth world coordinate pointmaps over time. Within a local temporal window, we first select 2D tracking trajectories T that remain visible throughout the window and lift them to 3D trajectories. We then smooth these 3D trajectories using a 1D convolution with adaptive weights, where weights are reduced for outlier points based on their temporal deviations. For each frame within the window, we treat the smoothed points as control points and apply a linear blend of control point displacements to transform all other points, weighting each control point’s influence based on proximity, resulting in dense smoothed pointmaps $\tilde{\mathcal{X}}^t$. (More details in supplementary.)

We then minimize per-frame distance between global pointmaps and smoothed pointmaps using L1 loss:

$$\mathcal{L}_{\text{PTS}}(\mathcal{X}) = \sum_{t=0}^N \left(\|\mathcal{X}^t - \tilde{\mathcal{X}}^t\|_1 \right) \quad (5)$$

The complete optimization objective for recovering the 4D scene is:

$$\hat{\mathcal{X}} = \arg \min_{\mathcal{X}, P, \sigma} \left[w_{\text{GA}} \mathcal{L}_{\text{GA}}(\mathcal{X}, \sigma, P) + w_{\text{CMA}} \mathcal{L}_{\text{CMA}}(\mathcal{X}) + w_{\text{CTS}} \mathcal{L}_{\text{CTS}}(\mathcal{X}) + w_{\text{PTS}} \mathcal{L}_{\text{PTS}}(\mathcal{X}) \right] \quad (6)$$

where $w_{\text{GA}}, w_{\text{CMA}}, w_{\text{CTS}}, w_{\text{PTS}}$ are the loss weights. The completed outputs of C4D contain world-coordinate pointmaps $\hat{\mathcal{X}}$, depthmaps \hat{D} , camera poses \hat{P} , camera intrinsics \hat{K} , motion masks \mathcal{M} , 2D tracking trajectories T , and lifted 3D tracking trajectories \hat{T} .

4. Experiments

We evaluate C4D on multiple downstream tasks, comparing it with specialized methods (Sec. 4.3), and 3D formulations (Sec. 4.2). The ablation study in Sec. 4.4 justifies our design choices, and implementation details are provided in supplementary materials.

4.1. Datasets and Metrics

We evaluate camera pose estimation on Sintel [3], TUM-dynamics [48] and ScanNet [8] following [4, 69, 70]. Sintel is a synthetic dataset featuring challenging motion blur and large camera movements. TUM-Dynamics and ScanNet are real-world datasets for dynamic scenes and static scenes, respectively. We report the metrics of Absolute Translation Error (ATE), Relative Translation Error (RPE trans), and Relative Rotation Error (RPE rot).

For depth estimation, we evaluate on Sintel, Bonn [40], and KITTI [14], following [21, 67]. Bonn [40] and KITTI [14] are real-world indoor dynamic scene and outdoor datasets. The evaluation metrics for depth estimation are Absolute Relative Error (Abs Rel), Root Mean Squared Error (RMSE), and the percentage of inlier points $\delta < 1.25$, as used in prior works [21, 65].

For point tracking, we evaluate our method on the TAP-Vid benchmark [10] and Kubric [16]. TAP-Vid contains videos with annotations of tracking point positions and occlusion. We use the metrics of occlusion accuracy (OA), position accuracy (δ_{avg}^x), and average Jaccard (AJ) to evaluate this benchmark, following [11, 18, 23, 56]. Kubric is a generator that synthesizes semi-realistic multi-object falling videos with rich annotations, including the moving status of tracking points in world coordinates. To fully evaluate the diverse dynamic patterns in the real world, we use three datasets from Kubric to assess dynamic accuracy (D-ACC): 1) MOVi-E, which introduces simple (linear) camera movement while always “looking at” the center point in world coordinates; 2) Panning MOVi-E, which modifies MOVi-E with panning camera movement; 3) MOVi-F, similar to MOVi-E but adds some random motion blur.

3D Model Weights	Optimization Formulation	Sintel			TUM-dynamics			ScanNet (static)		
		ATE ↓	RPE trans ↓	RPE rot ↓	ATE ↓	RPE trans ↓	RPE rot ↓	ATE ↓	RPE trans ↓	RPE rot ↓
DUS _t 3R	Global Alignment	0.416	0.216	18.038	0.127	0.058	2.033	0.060	0.024	0.751
	C4D	0.334	0.154	0.948	0.093	0.018	0.906	0.064	0.018	0.570
MAS _t 3R	Global Alignment	0.437	0.329	12.760	0.084	0.052	1.245	0.073	0.027	0.706
	C4D	0.448	0.199	1.730	0.048	0.012	0.671	0.067	0.018	0.467
MonSt3R	Global Alignment	0.158	0.099	1.924	0.099	0.041	1.912	0.075	0.026	0.707
	C4D (Ours)	0.103	0.040	0.705	0.071	0.019	0.897	0.061	0.017	0.538

Table 1. **Camera pose estimation results across 3D/4D formulations.** Evaluation on the Sintel, TUM-dynamic, and ScanNet datasets. The best results are highlighted in **bold**. Our 4D formulation, C4D, consistently improves the performance based on 3D models.

3D Model Weights	Optimization Formulation	Sintel			Bonn			KITTI		
		Abs Rel ↓	RMSE ↓	$\delta < 1.25$ ↑	Abs Rel ↓	RMSE ↓	$\delta < 1.25$ ↑	Abs Rel ↓	RMSE ↓	$\delta < 1.25$ ↑
DUS _t 3R	Global Alignment	0.502	5.141	54.9	0.149	0.422	84.4	0.129	5.162	84.2
	C4D (Ours)	0.478	5.052	57.9	0.143	0.411	84.7	0.126	5.140	85.0
MAS _t 3R	Global Alignment	0.370	4.669	57.8	0.174	0.503	78.4	0.092	4.000	89.8
	C4D (Ours)	0.379	4.756	58.3	0.168	0.485	78.6	0.092	4.000	89.7
MonSt3R	Global Alignment	0.335	4.467	57.5	0.065	0.254	96.2	0.090	4.128	90.6
	C4D (Ours)	0.327	4.465	60.7	0.061	0.249	96.5	0.089	4.128	90.6

Table 2. **Video depth estimation results across 3D/4D formulations.** We evaluate scale-and-shift-invariant depth on Sintel, Bonn, and KITTI. The best results are highlighted in **bold**. Our 4D formulation, C4D, consistently improve the performance based on 3D models.

4.2. Comparison across 3D/4D Formulations

3D Baselines We choose the currently available DUS_t3R-based models as our 3D baseline models: 1) DUS_t3R [58], trained on millions of image pairs in static scenes, demonstrating impressive performance and generalization across various real-world static scenarios with different camera parameters. 2) MAS_t3R [32], the follow-up work to DUS_t3R, which initializes its weights from DUS_t3R and is fine-tuned on the matching task, also using large-scale data from static scenes. 3) MonSt3R [67], which fine-tunes the decoder and head of DUS_t3R on selected dynamic scene datasets. The global alignment is the default optimization strategy in the 3D formulation, as described in Sec. 3.3.

Results We evaluate the results of camera pose estimation and video depth estimation, as shown in Table 1 and Table 2. Our C4D achieves consistent performance improvements compared to 3D formulation across different 3D model weights. For camera pose estimation, C4D significantly improves performance (e.g., reducing RPE_r from 18.038 to 0.948) even on the challenging Sintel dataset, demonstrating the effectiveness of our method. The results on the ScanNet dataset, which consists of static scenes, also show that our method further enhances performance in static environments. C4D also outperforms 3D formulations in terms of video depth accuracy. Moreover, these

results highlight a comparison among 3D model weights: DUS_t3R and MAS_t3R perform comparably overall, while MonSt3R achieves better results as it is fine-tuned on dynamic scene datasets.

4.3. Comparison with Other Methods

Since C4D produces multiple outputs, we compare our method with others specifically designed for individual tasks, including camera pose estimation, video depth estimation, and point tracking.

Evaluation on camera pose estimation We compare with methods that can predict camera pose and video depth jointly: Robust-CVD [30], CasualSAM [69], and the concurrent work MonSt3R [67]. We re-evaluated MonSt3R using their publicly available codes and checkpoints for a fair comparison. For a broader evaluation, we also compare with learning-based visual odometry methods: DROID-SLAM [53], DPVO [54], ParticleSfM [70], and LEAP-VO [4]. Note that DROID-SLAM, DPVO, and LEAP-VO require ground-truth camera intrinsics as input, while our C4D can estimate camera intrinsics and camera poses using only a monocular video as input. The results are presented in Table 3, showing that C4D achieves highly competitive performance even compared to specialized visual odometry methods and generalizes well on static scenes, such as the

Category	Method	Sintel			TUM-dynamics			ScanNet (static)		
		ATE ↓	RPE trans ↓	RPE rot ↓	ATE ↓	RPE trans ↓	RPE rot ↓	ATE ↓	RPE trans ↓	RPE rot ↓
Pose only	DROID-SLAM [†]	0.175	0.084	1.912	-	-	-	-	-	-
	DPVO [†]	0.115	0.072	1.975	-	-	-	-	-	-
	ParticleSfM	0.129	0.031	0.535	-	-	-	0.136	0.023	0.836
	LEAP-VO [†]	0.089	0.066	1.250	<u>0.068</u>	0.008	1.686	<u>0.070</u>	0.018	0.535
Joint depth & pose	Robust-CVD	0.360	0.154	3.443	0.153	0.026	3.528	0.227	0.064	7.374
	CasualSAM	0.141	<u>0.035</u>	<u>0.615</u>	0.071	<u>0.010</u>	1.712	0.158	0.034	1.618
	MonST3R	0.109	0.043	0.737	0.104	0.223	<u>1.037</u>	<u>0.068</u>	<u>0.017</u>	0.545
	C4D-M (Ours)	<u>0.103</u>	0.040	0.705	<u>0.071</u>	0.019	0.897	0.061	0.017	<u>0.538</u>

Table 3. **Camera pose evaluation** on Sintel, TUM-dynamic, and ScanNet. The best and second best results are highlighted in **bold** and underlined, respectively. [†] means the method requires ground truth camera intrinsics as input. “C4D-M” denotes C4D with MonST3R’s model weights.

Alignment	Category	Method	Sintel		Bonn		KITTI	
			Abs Rel ↓	$\delta < 1.25$ ↑	Abs Rel ↓	$\delta < 1.25$ ↑	Abs Rel ↓	$\delta < 1.25$ ↑
	Single-frame depth	Marigold	0.532	51.5	0.091	93.1	0.149	79.6
		DepthAnything-V2	0.367	55.4	0.106	92.1	0.140	80.4
Per-sequence scale & shift	Video depth	NVDS	0.408	48.3	0.167	76.6	0.253	58.8
		ChronoDepth	0.687	48.6	0.100	91.1	0.167	75.9
		DepthCrafter	0.292	69.7	<u>0.075</u>	97.1	<u>0.110</u>	<u>88.1</u>
Joint video depth & pose		Robust-CVD	0.703	47.8	-	-	-	-
		CasualSAM	0.387	54.7	0.169	73.7	0.246	62.2
		MonST3R	0.335	58.5	0.063	96.2	0.157	73.8
		C4D-M (Ours)	<u>0.327</u>	<u>60.7</u>	0.061	<u>96.5</u>	0.089	90.6
Per-sequence scale	Video depth	DepthCrafter	0.692	53.5	0.217	57.6	<u>0.141</u>	<u>81.8</u>
	Joint video depth & pose	MonST3R	<u>0.345</u>	<u>55.8</u>	<u>0.065</u>	<u>96.2</u>	0.159	73.5
	Joint video depth & pose	C4D-M (Ours)	0.338	58.1	0.063	96.4	0.091	90.6

Table 4. **Video depth evaluation** on Sintel, Bonn, and KITTI. Two types of depth range alignment are evaluated: scale & shift and scale-only. “C4D-M” denotes C4D with MonST3R’s model weights.

ScanNet dataset.

Evaluation on video depth estimation Table 4 shows the evaluation results on video depth estimation. We compare with various kinds of depth estimation methods: single-frame depth methods such as Marigold [25] and DepthAnything-V2 [65], and video depth methods such as NVDS [60], ChronoDepth [45], and DepthCrafter [21]. Note that these methods predict relative depth, which leads to inconsistencies across multiple views when projecting to world coordinates [15]. We also compare with methods that can predict video depth and camera pose jointly: Robust-CVD [30], CasualSAM [69], and MonST3R [67]. The evaluation is conducted using two kinds of depth range alignment: scale & shift, and scale-only. C4D achieves highly competitive results in scale & shift alignment. However, as demonstrated in [66], a shift in depth will affect the x, y, and

z coordinates non-uniformly when recovering the 3D geometry of a scene, resulting in shape distortions. Therefore, a more important evaluation is under scale-only alignment, where C4D achieves the best performance.

Evaluation on point tracking As part of the C4D outputs, we evaluate point tracking results in Table 5 and compare them with other TAP methods: RAFT [52], TAP-Net [10], PIPs [18], MFT [38], TAPIR [11], and Co-tracker [23]. Note that all previous TAP methods can only predict the position and occlusion of tracking points, whereas our method can additionally predict mobility, contributing to a robust motion mask prediction as described in Sec. 3.2.2. Despite this more challenging learning objective, our method still achieves comparable performance with SOTA methods and demonstrates high accuracy in predicting mobility.

Method	MOVi-E	Pan. MOVi-E	MOVi-F	TAP-Vid DAVIS			TAP-Vid Kinetics		
	D-ACC \uparrow	D-ACC \uparrow	D-ACC \uparrow	AJ \uparrow	$\delta_{\text{avg}}^x \uparrow$	OA \uparrow	AJ \uparrow	$\delta_{\text{avg}}^x \uparrow$	OA \uparrow
<i>Predict Position & Occlusion</i>									
RAFT	-	-	-	30.0	46.3	79.6	34.5	52.5	79.7
TAP-Net	-	-	-	38.4	53.1	82.3	46.6	60.9	85.0
PIPs	-	-	-	39.9	56.0	81.3	39.1	55.3	82.9
MFT	-	-	-	47.3	66.8	77.8	39.6	60.4	72.7
TAPIR	-	-	-	56.2	70.0	86.5	49.6	<u>64.2</u>	85.0
CoTracker	-	-	-	61.8	76.1	88.3	49.6	64.3	<u>83.3</u>
<i>Predict Position & Occlusion & Mobility</i>									
DynPT (Ours)	87.9	94.1	91.5	<u>61.6</u>	<u>75.4</u>	<u>87.4</u>	<u>47.8</u>	62.6	82.3

Table 5. **Point tracking evaluation results** on the TAP-Vid and Kubric (MOVi-E, Panning MOVi-E, and MOVi-F) Datasets. Apart from achieving competitive results with SOTA TAP methods, DynPT offers a unique capability: predicting the mobility of tracking points, which is crucial for determining whether a point is dynamic in world coordinates.

Method	Camera pose			Video depth		
	ATE \downarrow	RPE.t \downarrow	RPE.r \downarrow	Abs Rel \downarrow	RMSE \downarrow	$\delta < 1.25 \uparrow$
w/o CMA	0.140	0.051	0.905	0.335	4.501	0.582
w/o CTS	0.131	0.058	1.348	0.322	4.442	<u>0.608</u>
w/o PTS	<u>0.103</u>	<u>0.040</u>	<u>0.705</u>	0.327	4.465	0.607
C4D (Ours)	0.103	0.040	0.705	<u>0.327</u>	<u>4.459</u>	0.609

Table 6. **Ablation study** on the Sintel dataset.

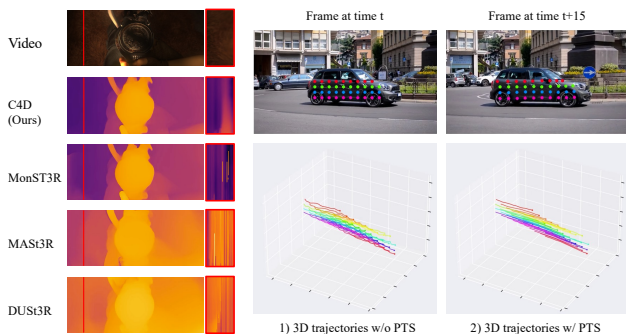


Figure 5. **Ablation illustration of Point Trajectory Smoothness (PTS) objective.** The temporal depth and 3D trajectories become more smooth after applying PTS objective.

4.4. Ablation Study

Ablation results in Table 6 indicate that all loss functions are crucial. The proposed loss suite achieves the best pose estimation with minimal impact on video depth accuracy. Since the temporal smoothness of depth cannot be reflected by the quantitative metrics in Table 6, we show the temporal depth slice changes in Figure 5, following [21, 65], which demonstrates that our PTS objective is effective in producing more temporally smooth depth and 3D point trajectories. Note that while MonST3R also employs the CMA objective, the motion mask used in this objective is crucial, and our mo-

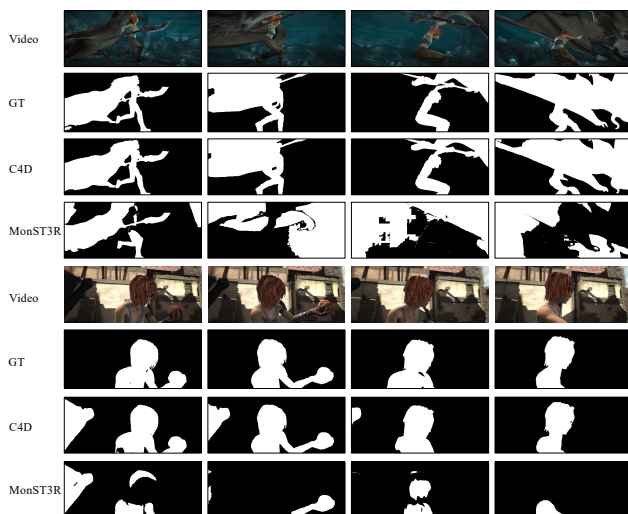


Figure 6. **Qualitative comparison of motion mask on Sintel.** Our motion mask is more accurate than MonST3R's.

tion mask is more accurate than MonST3R's, as shown in Figure 6. Due to page limitations, the ablation of DynPT is provided in the supplement.

5. Conclusion

In this paper, we introduce C4D, a framework for recovering 4D representations from monocular videos through joint prediction of dense pointmaps and temporal correspondences. Within this framework, a Dynamic-aware Point Tracker (DynPT), a correspondence-guided motion mask prediction, and correspondence-aided optimization are proposed to achieve accurate and smooth 4D reconstruction and camera pose estimation. Experiments demonstrate that C4D effectively reconstructs dynamic scenes, delivering competitive performance in depth estimation, camera pose estimation, and point tracking.

Acknowledgement

This project is supported by the National Research Foundation, Singapore, under its Medium Sized Center for Advanced Robotics Technology Innovation.

References

- [1] Sameer Agarwal, Noah Snavely, Steven M Seitz, and Richard Szeliski. Bundle adjustment in the large. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part II 11*, pages 29–42. Springer, 2010. 2
- [2] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110(3):346–359, 2008. 2
- [3] Daniel J. Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, pages 611–625, 2012. 2, 5
- [4] Weirong Chen, Le Chen, Rui Wang, and Marc Pollefeys. LEAP-VO: Long-term effective any point tracking for visual odometry. In *CVPR*, pages 19844–19853, 2024. 5, 6
- [5] Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Easi3r: Estimating disentangled motion from dust3r without training. *arXiv preprint arXiv:2503.24391*, 2025. 2
- [6] Seokju Cho, Jiahui Huang, Jisu Nam, Honggyu An, Seungryong Kim, and Joon-Young Lee. Local all-pair correspondence for point tracking. In *European Conference on Computer Vision*, pages 306–325. Springer, 2025. 2, 4
- [7] Wen-Hsuan Chu, Lei Ke, and Katerina Fragkiadaki. Dreamscene4d: Dynamic multi-object scene generation from monocular videos. *arXiv preprint arXiv:2405.02280*, 2024. 2
- [8] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 5
- [9] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelligence*, 29(6):1052–1067, 2007. 2
- [10] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. Tap-vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems*, 35:13610–13626, 2022. 2, 3, 5, 7
- [11] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10061–10072, 2023. 2, 3, 5, 7
- [12] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015. 2
- [13] Bardenus Duisterhof, Lojze Zust, Philippe Weinzaepfel, Vincent Leroy, Yohann Cabon, and Jerome Revaud. Mast3r-sfm: a fully-integrated solution for unconstrained structure-from-motion. *arXiv preprint arXiv:2409.19152*, 2024. 3
- [14] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. 32(11): 1231–1237, 2013. 5
- [15] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. 7
- [16] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanaprasgam, Florian Golemo, Charles Herrmann, et al. Kubric: A scalable dataset generator. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3749–3761, 2022. 4, 5
- [17] Jisang Han, Honggyu An, Jaewoo Jung, Takuya Narihira, Junyoung Seo, Kazumi Fukuda, Chaehyun Kim, Sunghwan Hong, Yuki Mitsufuji, and Seungryong Kim. D²ust3r: Enhancing 3d reconstruction with 4d pointmaps for dynamic scenes. *arXiv preprint arXiv:2504.06264*, 2025. 2
- [18] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *European Conference on Computer Vision*, pages 59–75. Springer, 2022. 2, 5, 7
- [19] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 2, 4
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [21] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. DepthCrafter: Generating consistent long depth sequences for open-world videos. *arXiv preprint arXiv:2409.02095*, 2024. 5, 7, 8
- [22] Moritz Kappel, Florian Hahlbohm, Timon Scholz, Susana Castillo, Christian Theobalt, Martin Eisemann, Vladislav Golyanik, and Marcus Magnor. D-npc: Dynamic neural point clouds for non-rigid view synthesis from monocular video. *arXiv preprint arXiv:2406.10078*, 2024. 5
- [23] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker: It is better to track together. *arXiv preprint arXiv:2307.07635*, 2023. 2, 3, 5, 7
- [24] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024. 3

- [25] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. 7
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2
- [27] Hanyang Kong, Xingyi Yang, and Xinchao Wang. Efficient gaussian splatting for monocular dynamic scene rendering via sparse time-variant attribute modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4374–4382, 2025. 2
- [28] Hanyang Kong, Xingyi Yang, and Xinchao Wang. Generative sparse-view gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26745–26755, 2025.
- [29] Hanyang Kong, Xingyi Yang, and Xinchao Wang. Rogsplat: Robust gaussian splatting via generative priors. In *Proceedings of the IEEE International Conference on Computer Vision*, 2025. 2
- [30] Johannes Kopf, Xuejian Rong, and Jia-Bin Huang. Robust consistent video depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1611–1621, 2021. 6, 7
- [31] Jiahui Lei, Yijia Weng, Adam Harley, Leonidas Guibas, and Kostas Daniilidis. Mosca: Dynamic gaussian fusion from casual videos via 4d motion scaffolds. *arXiv preprint arXiv:2405.17421*, 2024. 2
- [32] Vincent Leroy, Johann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. *arXiv preprint arXiv:2406.09756*, 2024. 3, 6
- [33] Qingming Liu, Yuan Liu, Jiepeng Wang, Xianqiang Lv, Peng Wang, Wenping Wang, and Junhui Hou. Modgs: Dynamic gaussian splatting from causally-captured monocular videos. *arXiv preprint arXiv:2406.00434*, 2024. 2
- [34] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, pages 1150–1157. Ieee, 1999. 2
- [35] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 2
- [36] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. 2
- [37] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: A versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. 5
- [38] Michal Neoral, Jonáš Šerých, and Jiří Matas. Mft: Long-term tracking of every pixel. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6837–6847, 2024. 7
- [39] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In *2011 international conference on computer vision*, pages 2320–2327. IEEE, 2011. 2
- [40] Emanuele Palazzolo, Jens Behley, Philipp Lottes, Philippe Giguere, and Cyrill Stachniss. Refusion: 3d reconstruction in dynamic environments for RGB-D cameras exploiting residuals. pages 7855–7862, 2019. 5
- [41] Peter J Rousseeuw. Least median of squares regression. *Journal of the American statistical association*, 79(388): 871–880, 1984. 4
- [42] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011. 2
- [43] Agarwal Sameer. Building rome in a day. *Proc. ICCV, 2009*, 2009. 2
- [44] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 2
- [45] Jiahao Shao, Yuanbo Yang, Hongyu Zhou, Youmin Zhang, Yujun Shen, Matteo Poggi, and Yiyi Liao. Learning temporally consistent video depth from video diffusion priors. *arXiv preprint arXiv:2406.01493*, 2024. 7
- [46] Colton Stearns, Adam Harley, Mikaela Uy, Florian Dubost, Federico Tombari, Gordon Wetzstein, and Leonidas Guibas. Dynamic gaussian marbles for novel view synthesis of casual monocular videos. *arXiv preprint arXiv:2406.18717*, 2024. 2
- [47] Frank Steinbrücker, Jürgen Sturm, and Daniel Cremers. Real-time visual odometry from dense rgb-d images. In *2011 IEEE international conference on computer vision workshops (ICCV Workshops)*, pages 719–722. IEEE, 2011. 5
- [48] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of RGB-D SLAM systems. pages 573–580, 2012. 5
- [49] 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. 2
- [50] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. 2
- [51] Yang-Tian Sun, Yihua Huang, Lin Ma, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Splatter a video: Video gaussian representation for versatile processing. *Advances in Neural Information Processing Systems*, 37:50401–50425, 2024. 5
- [52] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. 2, 4, 7

- [53] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021. 6
- [54] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. *Advances in Neural Information Processing Systems*, 36, 2024. 6
- [55] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *Vision Algorithms: Theory and Practice: International Workshop on Vision Algorithms Corfu, Greece, September 21–22, 1999 Proceedings*, pages 298–372. Springer, 2000. 2
- [56] 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*, pages 19795–19806, 2023. 5
- [57] Qianqian Wang, Vickie Ye, Hang Gao, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4d reconstruction from a single video. *arXiv preprint arXiv:2407.13764*, 2024. 2
- [58] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 1, 2, 3, 6
- [59] Shizun Wang, Xingyi Yang, Qihong Shen, Zhenxiang Jiang, and Xinchao Wang. Gflow: Recovering 4d world from monocular video. *arXiv preprint arXiv:2405.18426*, 2024. 2, 5
- [60] Yiran Wang, Min Shi, Jiaqi Li, Zihao Huang, Zhiguo Cao, Jianming Zhang, Ke Xian, and Guosheng Lin. Neural video depth stabilizer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9466–9476, 2023. 7
- [61] Yihan Wang, Lahav Lipson, and Jia Deng. Sea-raft: Simple, efficient, accurate raft for optical flow. In *European Conference on Computer Vision*, pages 36–54. Springer, 2025. 2, 4
- [62] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yohann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jérôme Revaud. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17969–17980, 2023. 4
- [63] Changchang Wu. Towards linear-time incremental structure from motion. In *2013 International Conference on 3D Vision-3DV 2013*, pages 127–134. IEEE, 2013. 2
- [64] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8121–8130, 2022. 2, 4
- [65] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything V2. *arXiv preprint arXiv:2406.09414*, 2024. 5, 7, 8
- [66] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 204–213, 2021. 7
- [67] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. *arXiv preprint arXiv:2410.03825*, 2024. 2, 3, 5, 6, 7
- [68] Zhoutong Zhang, Forrester Cole, Richard Tucker, William T Freeman, and Tali Dekel. Consistent depth of moving objects in video. *ACM Transactions on Graphics (ToG)*, 40(4):1–12, 2021. 5
- [69] 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*, pages 20–37. Springer, 2022. 5, 6, 7
- [70] Wang Zhao, Shaohui Liu, Hengkai Guo, Wenping Wang, and Yong-Jin Liu. Particlesfm: Exploiting dense point trajectories for localizing moving cameras in the wild. In *European Conference on Computer Vision*, pages 523–542. Springer, 2022. 5, 6