

4D Primitive-Mâché: Glueing Primitives for Persistent 4D Scene Reconstruction

Kirill Mazur, Marwan Taher, Andrew J. Davison
Dyson Robotics Lab, Imperial College London
{k.mazur21, m.taher, a.davison}@imperial.ac.uk




Figure 1. Our method (**4DPM**) takes in casual monocular videos (captured by an iPhone) and outputs complete 3D scene reconstructions at every observed timestamp, using all scene observations. The method takes in the outputs of a feedforward reconstruction model (**top row**) and glues dynamic geometry observations across time (**middle row**). This results in a complete and accurate geometric reconstruction, which re-uses observations from all timestamps (**bottom row**).

Abstract

We present **4DPM**, a dynamic reconstruction system that receives a casual monocular RGB video as input, and outputs a complete and persistent reconstruction of the scene. In other words, we reconstruct not only the currently visi-

ble parts of the scene, but also all previously viewed parts, which enables replaying the complete reconstruction across all timesteps.

Our method decomposes the scene into a set of rigid 3D primitives, which are assumed to be moving through the scene. Using estimated dense 2D correspondences, we

jointly infer the rigid motion of these primitives through an optimisation pipeline, yielding a 4D reconstruction of the scene, i.e. providing 3D geometry dynamically moving through time. To achieve this, we also introduce a mechanism to extrapolate motion for objects that become invisible, employing motion-grouping techniques to maintain continuity.

The resulting system enables 4D spatio-temporal awareness, offering capabilities such as replayable 3D reconstructions of articulated objects through time, multi-object scanning, and object permanence. On object scanning and multi-object datasets, our system significantly outperforms existing methods both quantitatively and qualitatively. Project page: <https://makezur.github.io/4DPM/>

1. Introduction

Building accurate and complete geometric reconstructions of dynamic scenes is a critical task in computer vision with broad applications in robotics, embodied intelligence, and augmented reality. Spatial AI systems typically operate in dynamic environments, requiring maps that are both complete and precise. Whilst SLAM [4, 8, 15, 24] and SfM [29, 36] methods excel in building maps of static environments, they abstain from mapping dynamic scene parts. Agents typically interact with and modify their surroundings, and many environments are inherently dynamic, with objects and actors constantly changing position and configuration.

The main goal of mapping is persistence, i.e. to retain as much relevant observed information as possible. For dynamic scenes, only the latest frame depth represents up-to-date geometric observations. Most dynamic mapping systems lack persistence and do not aggregate observed geometry across time. Thus, a lot of the information about previously observed scene elements is lost. Our goal in this paper is to build the most complete and persistent reconstruction possible, by not discarding previously observed information about moving objects.

The most general case of persistent, non-rigid dynamic reconstruction remains very challenging even for RGB-D systems and can only be performed on short sequences in controlled scenarios. Here, we make a piecewise-rigid motion assumption but show that this still covers a wide range of scenes and allows for replayable moving reconstructions consisting of many parts.

We present a new scene reconstruction approach, 4D Primitive-Mâché (4DPM), which enables corresponding reconstruction of every observed keyframe *at all given timestamps*. This contrasts with previous methods [13, 20, 41] that reconstruct geometry (possibly coupled with point tracks) only at the time it was observed. This allows us to replay the 3D reconstruction in 4D.

The key to our method is to decompose scenes into object-like chunks, as in SuperPrimitive [23] (SP), a primitive-based approach to SLAM and SfM. We tackle dynamic scene reconstruction using scene primitives as the underlying representation. This enables compact and optimisable representation of dense scene geometry. Whilst SP assumes pixel-aligned primitives and estimates only unknown depth scales, we extend this framework by endowing each primitive with rigid motion parameters.

This formulation allows us to efficiently reduce the complex all-to-all temporal dense mapping problem to estimating a single $\mathbb{SE}(3)$ pose per primitive. The pose encodes the transformation to the last observed timestamp of that object. We demonstrate that this compact representation suffices to encode the complete temporal trajectory of each primitive, enabling tractable motion inference across a diverse range of dynamic scenes.

In order to estimate the motion of these primitives (“glue” them across time), we employ an off-the-shelf dense 2D correspondence estimation network [12]. Given the estimated correspondence between 3D primitives, we estimate their motion by direct 3D alignment in their respective coordinate systems.

To briefly summarise our contributions:

- We introduce a primitive-based motion parameterisation that represents per-pixel motion fields through sparse per-primitive $\mathbb{SE}(3)$ poses, dramatically reducing the dimensionality of dynamic scene reconstruction whilst maintaining expressiveness for piecewise-rigid motion.
- We demonstrate geometric accuracy significantly higher than the existing monocular methods on challenging scenarios including multi-object interactions, validated against multi-view ground truth.
- Our primitive persistence enables spatial memory capabilities, maintaining representations of temporarily occluded objects. To the best of the authors’ knowledge, our method is the first monocular reconstruction system demonstrating such capabilities.

2. Related Work

Compact Scene Representations Rather than optimising dense geometric (such as depth) values for every observed pixel, compact representations parameterise scenes through learned models or structured primitives. Code-based methods [2, 6, 22] learn depth prediction networks conditioned on optimisable latent codes, enabling joint optimisation during reconstruction. Similarly, COMO [9] employs explicit 3D control points as an optimisable scene representation. Primitive-based approaches take this further by representing scenes as compositions of higher-level structures: SLAM++ [28] represents scenes as collections of CAD models and optimises their poses, while SuperPrimitive [23] decomposes frames into 2.5D primitives obtained

through surface normal integration within image segments, estimating their depth scales through optimisation.

Compact Motion Representations Early sparse reconstruction systems of moving bodies segmented motion into multiple rigid components [5], which was later extended to articulated objects [33]. To move beyond rigid-body assumptions, subsequent work explored compact motion representations through decomposition: [3] proposed decomposing moving objects as linear combinations of shape bases, whilst [1] instead decomposed trajectories using a linear basis set.

With the advent of commercially available depth sensors, fusion-based methods enabled dense dynamic reconstruction. DynamicFusion [25] pioneered RGB-D non-rigid reconstruction by representing motion via a sparse set of deformation nodes and fusing observations into a canonical, voxel-based TSDF model. Object-centric approaches such as Co-fusion [27] and MID-fusion [40] spawned individual dense maps — surfel-based and octree-based respectively — for different objects. In contrast with [25], these methods enabled dynamic scene level reconstruction. Whilst these methods share our object-based motion representation, they require continuous depth streams for map fusion and model-based tracking, making them inapplicable to monocular scenarios.

For monocular dense reconstruction, SuperPixel Soup [18, 19] represents scene motion by decomposing the scene into rigidly moving superpixels; however, these methods are largely limited to two-view, small-baseline configurations.

4D reconstruction methods In recent years, joint depth and camera pose estimation from casual videos containing dynamic objects has gained significant attention. Initial approaches combined learned priors with scene-specific optimisation [13, 17, 20, 43], which were later followed by feed-forward reconstruction models [39, 46]. Our method is complementary to these approaches: whilst they estimate geometry only at the time of observation, we focus on remapping observed reconstructions to all timestamps. Consequently, the outputs of these methods can be used as the input to ours.

Inspired by the recent success of DUS_t3R [37], a series of recent methods [10, 32, 42] has extended DUS_t3R’s shared coordinate frame mapping paradigm to dynamic scenes by establishing temporal correspondences, thereby warping point maps across different time frames. While these methods are theoretically more expressive, real data for training supervision remains extremely sparse, leading to limited performance even on scenes with predominantly rigid motion. Additionally, pair-wise methods suffer from quadratic complexity when the number of frames increases.

3. Method

3.1. Formulation

Our geometry representation is based on point maps $\mathbf{X}_k^t \in \mathbb{R}^{H \times W \times 3}$, which represent 3D geometry captured from camera k and warped to time t in world coordinates. For a video sequence $\mathcal{I} = \{I_0, \dots, I_n \mid I_i \in \mathbb{R}^{H \times W \times 3}\}$, the scene reconstruction at time t is the collection $\mathcal{X}^t = \{\mathbf{X}_0^t, \dots, \mathbf{X}_n^t\}$.

Prior work reconstructs image-aligned point maps only at their respective observation times, yielding $\{\mathbf{X}_0^0, \mathbf{X}_1^1, \dots, \mathbf{X}_n^n\}$. In contrast, our goal is to reconstruct the complete scene geometry at every observed timestamp, obtaining $\{\mathcal{X}^0, \dots, \mathcal{X}^n\}$. We achieve this by taking observation-time point maps from a feed-forward model [39] as input to our 3D primitive glueing method.

Primitives Every keyframe I_i is partitioned into a set of non-overlapping image regions $S_p \in I_i$. These regions cut 3D primitives out of \mathbf{X}_i , which we denote as $S_p \odot \mathbf{X}_i$. To represent scene motion, we allow these primitives to move as rigid bodies over time, which is parametrised with $\mathbb{SE}(3)$ poses $T(S_p)$ of each individual 3D primitive. In Sec. 3.4 we describe how to infer this motion from estimated 2D correspondences between images.

Pose representation Object and camera poses $T \in \mathbb{SE}(3)$ are stored as 4×4 matrices. For optimisation, we employ a Lie group parametrisation following the notation in [30]. Pose updates are represented as Lie algebra elements $\tau \in \mathfrak{se}(3) \simeq \mathbb{R}^6$, and transformations are updated via:

$$T \leftarrow T \oplus \tau \tag{1}$$

3.2. Objects

Our 3D primitives are matched across time to form object-like structures for two key reasons. First, this enables filtering of spurious correspondences, which frequently arise at object boundaries. Second, and more importantly, this matching facilitates a compact motion representation: each 3D primitive requires only a single $\mathbb{SE}(3)$ to represent its motion at all visible timestamps, see Sec. 3.5.

The scene primitives are clustered into *object* groups \mathcal{O} . Intuitively, our method assumes each object-like entity to be rigidly moving in the scene. Each object \mathcal{O} is a set of primitives, grouped by time $\mathcal{O} = \{S^{t_{start}}, \dots, S^{t_{end}}\}$.

To represent an object \mathcal{O} ’s motion, we couple each segment $S \in \mathcal{O}$ with an $\mathbb{SE}(3)$ pose $T(S)$, which maps it in the coordinate frame of the last observed segment, $S^{t_{end}}$. The pose $T(S^{t_{end}})$ corresponding to the last observed segment is set to identity, since pose estimation for each object entity has a natural $\mathbb{SE}(3)$ gauge freedom [31, 34] (since the scale between objects is fixed).

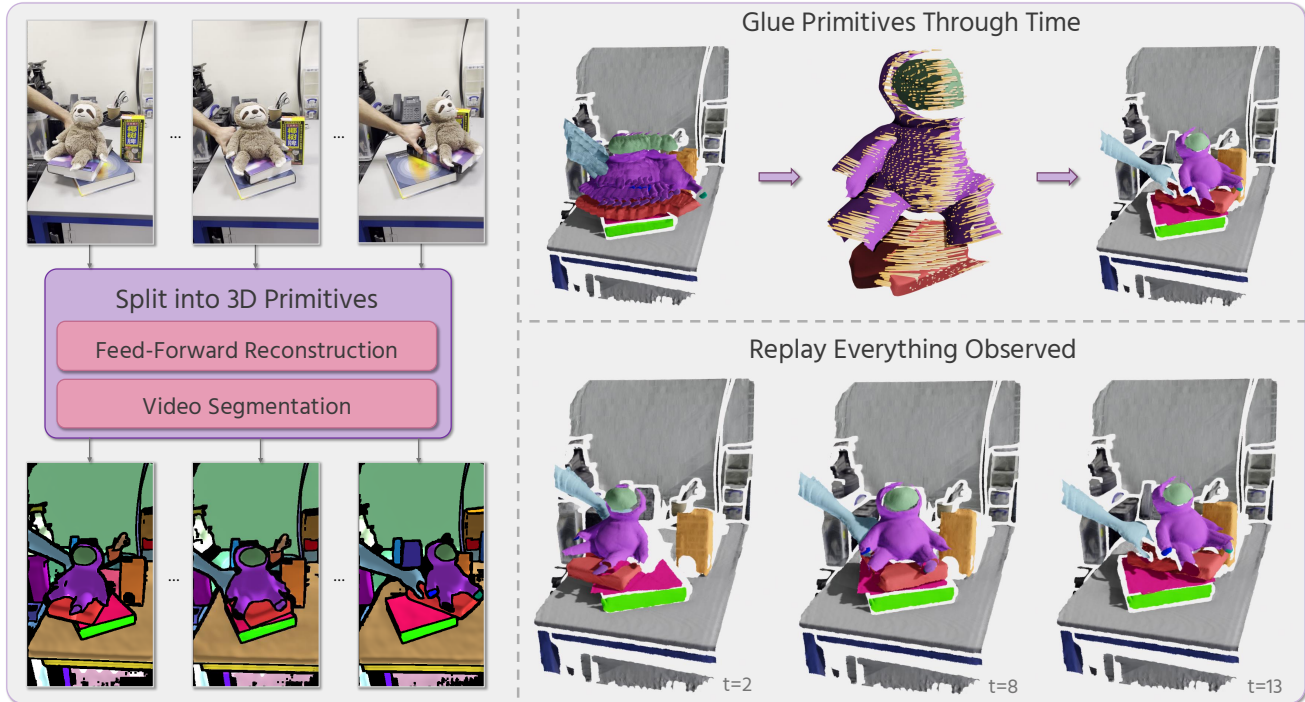


Figure 2. **4D reconstruction with 4DPM.** (left) Our frontend takes in a monocular RGB video and splits it into a set of 3D primitives. Each primitive is represented as a 3D point map in the world coordinate space, cut out by a segmentation mask. These primitives are matched across time (visualised with consistent colours) to form consistent entities across time, to which we refer as *objects*. (top right) Given geometric observations positioned at their respective timestamps, we “glue” primitives belonging to the same object across time according to their estimated dense 2D correspondences. (bottom right) The resulting complete reconstruction can be replayed across all observed timestamps.

Static-Dynamic classification Monocular videos present a natural geometric ambiguity: from 2D correspondences alone, motion induced by camera movement cannot be easily disambiguated from actual scene motion. To resolve this gauge freedom in the backend optimisation, we freeze objects based on their initial correspondence residuals. These frozen objects are assumed to be static in the world frame, such that their observed motion arises solely from camera movement.

3.3. Frontend

Geometry Given a set of keyframes \mathcal{I} , we run a feed-forward scene reconstruction model π^3 [39], which estimates point map at observation times $\pi^3(\mathcal{I}) = \{\mathbf{X}_0^0, \mathbf{X}_1^1, \dots, \mathbf{X}_n^n\}$. This serves as the basis for our Primitive-Mâché glueing algorithm.

Segmentation Given a video sequence \mathcal{I} , our frontend should decompose this video into a set of non-overlapping objects, covering most of the video sequence.

To this end, we segment the first frame of the video as in SuperPrimitive [23]. These masks are propagated to all other keyframes using SAMv2 [26]. Our algorithm sequentially processes subsequent frames, instantiating new objects in uncovered image regions via active sampling and

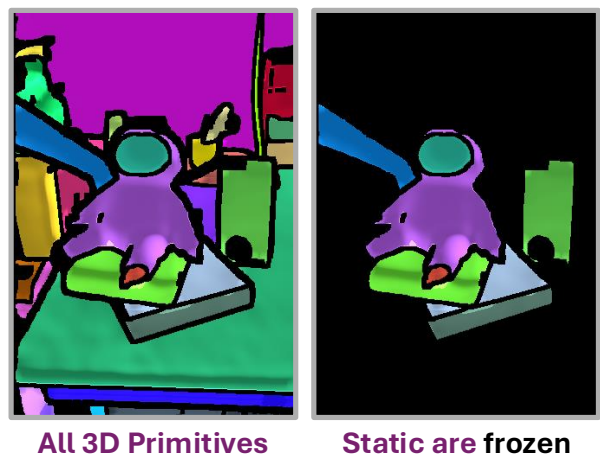


Figure 3. **Static vs dynamic segmentation.** We visualise all estimated primitives on the left. Before motion estimation, we freeze primitives with insufficiently high correspondence residuals, assuming they are static. On the right, only dynamic primitives are shown. Our system produces motion segmentation as a by-product.

propagating them forwards, akin to [45].

This process partitions the imagery into the set of objects introduced in Sec. 3.2.

Correspondence To extract dense correspondence between subsequent keyframes, we run a dense point tracking model [12] on the videos between the keyframes. These correspondences are then filtered to only align points that belong to the same object.

3.4. Backend

To estimate 3D primitives’ poses, we employ dense 2D correspondences between consecutive keyframes as constraints in a 3D optimisation. We perform joint optimisation across all objects and keyframes to find poses that best align corresponding 3D points. Intuitively, this process “glues” the objects out of their constituent primitives.

We use a dense correspondence network to estimate pixel-wise flow between temporally adjacent keyframes I_k and I_{k+1} . While non-adjacent keyframes can also be used, we find adjacent pairs provide the best efficiency-robustness trade-off. The network[39] also provides per-pixel correspondence confidence weights which we use in the cost calculation.

Given a pointmap \mathbf{X}_k^k (which we further denote as \mathbf{X}_k for brevity) and flow field, we obtain corresponding 3D points $\widehat{\mathbf{X}}_{k+1}$ by warping \mathbf{X}_{k+1} according to the flow. Similarly, segment masks S_k are warped to produce \widehat{S}_{k+1} .

For a single object $\mathcal{O} = \{S_n, \dots, S_m\}$, our dense optimisation aims to estimate the poses $T(S_n), \dots, T(S_m)$ of its constituent primitives. Our cost function directly minimises the distance between the 3D points of primitives belonging to the same object, weighted by correspondence confidence w_{ij} :

$$E(\mathcal{O}) = \sum_{(i,j) \in \mathcal{T}(\mathcal{O})} \left\| w_{ij} \cdot S_i \cdot \widehat{S}_j \left(T_j^{-1} T_i \mathbf{X}_i - \widehat{\mathbf{X}}_j \right) \right\|_{\rho} \quad (2)$$

where $\mathcal{T}(\mathcal{O})$ is the set of temporally-adjacent primitives which belong to object \mathcal{O} and $\|\cdot\|_{\rho}$ is the Huber cost function.

The final cost is combined across all objects present in the scene that are not classified as static:

$$E_{\text{final}} = \sum_i E(\mathcal{O}_i) \quad (3)$$

Jacobians \mathbf{J} for all poses are derived analytically and efficiently inferred in parallel for all objects in the scene. We solve this optimisation problem for all objects jointly via iteratively reweighted least squares using Gauss-Newton optimisation [34]:

$$\mathbf{J}^T \mathbf{W} \mathbf{J} \tau = -\mathbf{J}^T \mathbf{W} \mathbf{r}, \quad T_i = T_i \oplus \tau \quad (4)$$

3.5. Time Remapping

Having estimated poses for all primitives, we now aim to produce reconstructions of every frame at each times-

tamp, namely $\mathcal{X}^0, \dots, \mathcal{X}^n$. In contrast to the non-rigid scenario, having mapped rigid objects onto their latest observed frame allows us to infer these objects’ positions at all observed timestamps.

More precisely, given an object \mathcal{O} represented with primitives $\{S^{t_{start}}, \dots, S^{t_{end}}\}$ together with their estimated poses $T(S^{t_{start}}), \dots, T(S^{t_{end}})$ we can express the warping as follows. As covered in Sec. 3.2, each pose maps a primitive onto the coordinate system of the last observed primitive, $S^{t_{end}}$. Then, the warping of S^p to time q can be naturally expressed as:

$$T^{p \rightarrow q} = [T(S^q)]^{-1} T(S^p), \quad (5)$$

which corresponds to first transforming S^p onto the coordinate frame of $S^{t_{end}}$ and then pulling it back to time q .

3.6. Motion Segmentation

While some objects might become unobservable, humans can still reason about their positions in the 3D space. Our representation enables such reasoning.

When an object becomes unobservable (e.g., an item placed in a closing drawer), we infer its continued motion (when possible) by linking it to a parent object that remains visible.

We identify potential parent objects using two criteria: spatial contact and velocity similarity. Parents are only assigned across objects in spatial contact, up to transitive closure with the preference to objects with clustered velocities.

Spatial Contact To determine whether or not two objects are in contact, we fit Oriented Bounding Boxes (OBBs) to every object at each timestamp. We consider the objects to be *in contact* if their bounding boxes extended by α have non-zero intersection. In our experiments, we set $\alpha = 1.1$.

Velocity Clustering Every object has its own coordinate system for poses, as discussed in Sec. 3.2. Hence object poses $T'_{\text{obj}} = \mathcal{F} T_{\text{obj}}$ are defined up to an unknown gauge freedom $\mathcal{F} \in \mathbb{SE}(3)$. Thus comparing their poses directly is ill-posed. However, comparing velocities is gauge-invariant:

$$\begin{aligned} T'(t)^{-1} T'(t-1) &= T(t)^{-1} \mathcal{F}^{-1} \mathcal{F} T(t-1) \\ &= T(t)^{-1} T(t-1) \end{aligned} \quad (6)$$

Velocities \mathcal{V} and \mathcal{W} of two co-observed objects are then compared $\log(\mathcal{V}^{-1} \mathcal{W})$ under the Mahalanobis distance with diagonal covariance parameters σ_{τ} (translation) and σ_{ψ} (rotation). Objects with distance below threshold are considered to have similar motion.

4. Experiments

We evaluate our dynamic geometry reconstruction quality, both in terms of accuracy and completeness. This is done by

warping all observations to the final timestamp and comparing against the GT geometry of the scene. In this setup, we assume a scene is captured by a set of synchronised RGB-D cameras, providing multi-view GT. Note that only one RGB camera stream is used as input to all methods, to which we further refer as the main camera.

Ideally, after processing the whole sequence, a method should reuse all observations to produce as complete and accurate reconstruction as possible. Specifically, for all methods, we time-warp each pointmap onto the latest keyframe n , yielding the final reconstruction $\mathcal{X}^n = \{\mathbf{X}_0^n, \mathbf{X}_1^n, \dots, \mathbf{X}_n^n\}$.

Due to gauge freedom, the estimated geometry has a $\text{Sim}(3)$ transformation ambiguity. For every method, we align its prediction to the ground truth coordinate system using Umeyama alignment [35], applying the transformation computed from the last observed keyframe’s pointmap.

For DUST3R-based baselines (St4track [10] and POMATO [42]), we first extract all pointmaps temporally aligned to the final observation X_n^n , and align them to a shared coordinate frame.

Metrics We evaluate reconstruction quality using two complementary metrics against pseudo-ground-truth pointmaps: (1) *accuracy* — the percentage of predicted points within a 1cm threshold of the ground truth, measuring precision; and (2) *recall* — the percentage of ground truth points covered by at least one predicted point within 1cm, measuring completeness. Following the established practice [16], we report the F-score, the harmonic mean of accuracy and recall, which reflects both metrics in a single value.

In all our experiments, we evaluate reconstruction quality solely on dynamic scene parts. Including static parts in the evaluation would dominate the completeness (i.e., recall) metric. Specifically, auxiliary cameras in multi-view datasets typically capture static scene regions that are never visible from the main camera. Evaluating completeness on these regions would penalise a method for failing to reconstruct geometry that was never observed in the input video, making the recall measure uninformative for assessing the quality of dynamic scene reconstruction.

Long sequences are split into chunks of 150 consecutive frames and all methods are tested on each chunk separately. In this case, we report average result per sequence.

Baselines We propose two natural baselines for our method. First, (π^3 last view in the tables) only uses the latest frame pointmap as estimated geometry. In the last frame the positioning of dynamic geometry is up to date, but it lacks completeness.

Another natural baseline (π^3 in the tables) is simply the untouched estimate of a feed-forward reconstruction model run on the input keyframes. While the geometrical information in this case is complete, it is not properly positioned for

early keyframes.

Besides these baselines, we also compare our system to state-of-the-art methods, that are capable of dense time-warping of observed scene geometry: St4track [10], POMATO [42] and TraceAnything [21].

All our experiments are conducted with a single NVIDIA GeForce RTX 4090 GPU.

4.1. Object Scanning

We evaluate our system on HO3D [11], an object scanning dataset in which stationary cameras capture humans manipulating objects of interest. We select sequences recorded with four calibrated depth cameras, which provide ground truth depth for evaluation. Although human hands are non-rigid, hand poses remain largely static in this dataset. We therefore report geometric accuracy for both hands and scanned objects (i.e., dynamic scene parts) using the ground truth segmentation masks provided.

In Tab. 1, we report per-sequence F-scores for all methods, alongside average F-score, recall, and precision. While (π^3 last frame) achieves higher precision, its recall (i.e., scene coverage) is poor because information from all keyframes except the last is discarded. Other baselines perform well in terms of recall, indicating decent object coverage. However, most fall short in precision. Our method provides the best balance between precision and recall, achieving the highest F-score across all compared methods.

4.2. Multi-Object Dataset

Due to the scarcity of real-world datasets capturing complex multi-object rigid motion, we collected our own dataset using four time-synchronised Azure Kinect cameras. Following our evaluation protocol for HO3D, only one camera stream serves as input to all methods, while all four synchronised views are used to generate more reliable pseudo-ground-truth (pseudo-GT) geometry. Although these cameras include depth sensors, we observed unsatisfactory depth quality for most objects in practice. We therefore employed the metric-aligned feed-forward reconstruction from [39] on all four synchronised views to obtain pseudo-GT geometry estimates. We provide ground-truth masks for all dynamic scene parts and evaluate using the same geometric metrics as on HO3D, namely accuracy, recall and F-score all evaluated at 1cm threshold.

Quantitative results are presented in Tab. 2: our method substantially outperforms all other baselines in terms of our main evaluation metric, F-score. Qualitatively, as seen in Fig. 4 our method produces significantly more accurate and complete reconstructions than others.

4.3. Object Permanence

Our primitive-based representation naturally allows us to reason about spatial relationships between moving scene

	ABF1	BB1	GPMF1	GSF1	MDF1	ShSu1	SiBF1	avg. F-score	Precision	Recall
π^3 last view	0.3185	0.3399	0.2553	0.4330	0.2604	0.2794	0.3581	0.3206	0.9255	0.2018
π^3	0.4405	0.4913	0.5257	0.5394	0.4855	0.6744	0.4968	0.5219	0.4735	0.6296
St4Track	0.4717	0.3015	0.5699	0.6414	0.4913	0.7892	0.5095	0.5392	0.4549	0.7293
POMATO	0.4572	0.4304	0.5254	0.6747	0.5581	0.5271	0.4766	0.5214	0.4065	<u>0.7650</u>
TraceAny	<u>0.5656</u>	<u>0.6323</u>	<u>0.6199</u>	<u>0.7158</u>	<u>0.5772</u>	0.7375	<u>0.6069</u>	<u>0.6365</u>	0.5748	0.7398
Ours	0.7259	0.7306	0.7645	0.7636	0.7586	<u>0.7758</u>	0.7824	0.7573	<u>0.7630</u>	0.7774

Table 1. **Quantitative evaluation on HO3D dataset.** We report F-score (threshold at 1 cm) per sequence for all methods. Average F-score, precision, and recall across all sequences are also reported. Our method outperforms all baselines by a substantial margin in terms of F-score for dynamic object scanning, providing the best balance between completeness and accuracy. Best is highlighted as **bold**, while second-best is underscored.

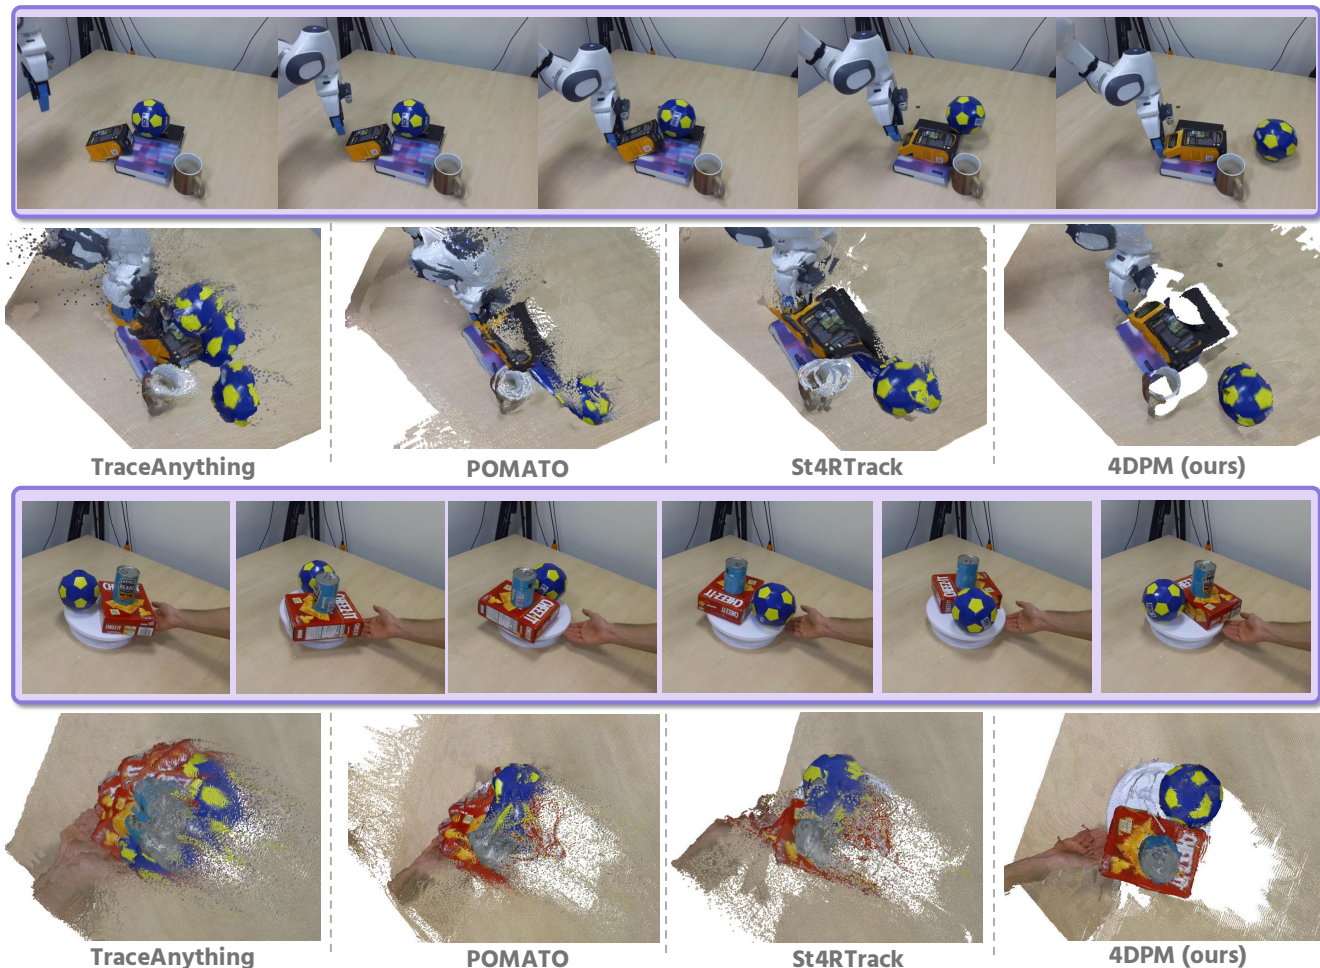


Figure 4. **Qualitative comparison on Multi-Object dataset.** Input video frames are shown in purple. Below each video, we visualise all observed point-maps time-warped to the latest timestamp. Our system successfully handles multi-object motion and performs well on particularly challenging objects such as the spinning ball and robot gripper (**top row**). We provide a top-down view of multiple objects spinning on a rotating base (**bottom row**). Our method correctly aggregates all observations, resulting in complete and accurate object scans.

parts. This capability proves especially useful when objects become occluded, due to motion. In Figure 5, we demonstrate our system operating on a closing drawer sequence. We show the top-down view of our 4D reconstruction. Ini-

tially, objects inside the drawer are visible and can be ordinarily reconstructed. As the drawer closes, these objects gradually become occluded, yet their motion and position can be inferred from the drawer front’s movement. Our

	BallPush	BoxSpin	MultiObj1	MultiObj2	PanStir	Spin1	Spin2	Mean	Precision	Recall
π^3 last view	0.3627	0.5296	0.6470	<u>0.8318</u>	0.1898	0.4711	0.5179	0.5071	0.8837	0.3707
π^3	0.2176	0.6128	<u>0.6828</u>	0.6900	0.1640	0.3987	<u>0.5934</u>	0.4799	0.3637	0.7382
St4track	0.3968	<u>0.7955</u>	0.4514	0.6608	0.0401	0.4778	0.4189	0.4630	0.3585	0.6792
POMATO	0.4311	0.7175	0.6541	0.6373	0.4597	<u>0.6225</u>	0.5827	<u>0.5864</u>	0.4668	<u>0.8071</u>
TraceAny	<u>0.4581</u>	0.6876	0.3191	0.6175	<u>0.4616</u>	0.3706	0.4578	0.4817	0.3773	0.6946
Ours	0.7683	0.9179	0.8226	0.8903	0.6359	0.8544	0.6746	0.7948	<u>0.7195</u>	0.9000

Table 2. **Quantitative evaluation on Multi-Object dataset.** We report F-score (threshold at 1 cm) per sequence for all methods. Average F-score, precision, and recall across all sequences are also reported. Our method significantly outperforms all baselines in terms of F-score on multi-object dynamic reconstruction, providing the best balance between completeness and accuracy. Best is highlighted as **bold**, while second-best is underscored.




Figure 5. **Object permanence capabilities.** In (top row) we show input frames of a drawer closing sequence. The resulting reconstruction estimated with 4DPM from the top-down view in the (bottom row). When the drawer is fully closed (rightmost column), our method still reconstructs objects inside the drawer and the drawer body, despite it being completely occluded. This showcases object permanence capabilities of 4DPM. The top of the drawer is removed from reconstruction for better viewing.

primitive-based 3D representation, coupled with the motion segmentation technique described in Sec. 3.6, enables such capability.

Based on motion segmentation, our system groups the drawer body with its front. Importantly, the objects inside are not in direct contact with the drawer front; hence their motion must be transitively inferred through the drawer body. Consequently, these occluded objects remain spatially associated and motion-grouped with the drawer, preserving them in the reconstruction.

To the best of the authors’ knowledge, this is the first system to demonstrate such capabilities from casual monocular videos. We strongly encourage readers to consult the supplementary materials for video demonstrations of this capability.

5. Limitations

Our system assumes that each primitive is rigid and thus cannot represent more intricate non-rigid deformations. Extending the method to handle such deformations whilst maintaining computational efficiency remains an important

direction for future work. Additionally, incremental mapping capabilities, where the scene representation is built and updated over extended sequences, have yet to be explored.

6. Conclusion

We have presented 4D Primitive-Mâché (4DPM), a novel method for monocular 4D dense scene reconstruction from casual RGB videos. Our primitive-based motion parametrisation achieves more accurate and complete reconstructions than existing monocular methods across object scanning and multi-object interaction scenarios. A persistence mechanism further enables spatial memory of occluded objects — a capability not previously demonstrated in monocular reconstruction. By dramatically reducing the dimensionality of the dynamic reconstruction problem without sacrificing expressiveness, 4DPM opens new possibilities for robust scene understanding from monocular video.

7. Acknowledgments

Research presented in this paper was supported by Dyson Technology Ltd.

4D Primitive-Mâché: Glueing Primitives for Persistent 4D Scene Reconstruction

Supplementary Material

8. System Analysis

We analyse the behaviour of our system across several dimensions.

Robustness to geometry frontend Since our method relies on the outputs of feed-forward geometry estimation networks, we first examine its sensitivity to the choice of model. Table Tab. 3 shows that 4DPM is robust to the choice of frontend, yielding a consistent improvement in F-score across all configurations (by at least +0.235).

Component ablation Table Sec. 8 presents a detailed ablation of our system’s components. Our correspondence filtering mechanism is critical to performance: removing it causes a drop in F-score of 0.15 (−19%). Second-order optimisation substantially outperforms first-order Adam. Replacing Alltracker with SEA-RAFT [38] results in a modest performance degradation. Finally, our method is robust to segmentation quality: using unfiltered SAMv2 masks or the lighter SAMv2-Small variant yields performance comparable to the full pipeline.

	F-score	delta F-score
π^3	0.5219	-
VGGT	0.5330	-
DAv3	0.4587	-
π^3 + Ours	0.7573	+0.235
VGGT + Ours	0.7700	+0.237
DAv3 + Ours	0.7103	+0.256

Table 3. **Robustness to different geometry estimators on HO3D.**

	F-score
Unfiltered correspondence	0.6401
FO optimiser	0.7228
SEA-RAFT for correspondence	0.7727
Unfiltered masks from SAMv2	0.7765
SAMv2-small	0.7947
Ours	0.7948

Table 4. **Ablation study on Multi-obj dataset.**

9. Analytical Jacobian Derivation

We provide analytical Jacobians for a single pairwise residual of an object. In practice, the Hessian is block-diagonal with respect to objects.

Given relative residuals $r = T_j^{-1}T_iX_i - \widehat{X}_j$, we derive its analytical right Jacobians $\mathbf{J} = [\mathbf{J}_{T_i} | \mathbf{J}_{T_j}]$ with respect to object poses T_i and T_j .

Let $Z = T_j^{-1}T_i$, then:

$$\frac{dZ}{dT_i} = \text{Id} \quad (7)$$

$$\frac{dZ}{dT_j} = -\text{Ad}(Z^{-1}) \quad (8)$$

From [30], if $Z = \begin{bmatrix} R_Z & t_Z \\ 0 & 1 \end{bmatrix}$, $R_Z \in \mathbb{SO}(3)$, $t_Z \in \mathbb{R}^3$ and Act is the action operator of $\mathbb{SE}(3)$ group on \mathbb{R}^3 , i.e $\text{Act}(Z, X) = R_Z X + t_Z$, then:

$$\frac{d \text{Act}(Z, X)}{dZ} = \begin{bmatrix} R_Z & -R_Z[X]_{\times} \end{bmatrix} \quad (9)$$

Then, using chainrule,

$$\begin{aligned} \frac{dr}{dT_i} &= \frac{d(ZX_i - \widehat{X}_j)}{dT_i} = \\ &= \frac{d \text{Act}(Z, X)}{dZ} \frac{dZ}{dT_i} = \begin{bmatrix} R_Z & -R_Z[X_i]_{\times} \end{bmatrix} \end{aligned} \quad (10)$$

$$\begin{aligned} \frac{dr}{dT_j} &= \frac{d(ZX_i - \widehat{X}_j)}{dT_j} = \frac{d \text{Act}(Z, X_i)}{dZ} \frac{dZ}{dT_j} = \\ &= \begin{bmatrix} R_Z & -R_Z[X_i]_{\times} \end{bmatrix} (-\text{Ad}(Z^{-1})) \end{aligned} \quad (11)$$

The last equation can be simplified further. Recall that $Z^{-1} = \begin{bmatrix} R_Z^T & -R_Z^T t_Z \\ 0 & 1 \end{bmatrix}$ and, therefore, its adjoint is:

$$\text{Ad}(Z^{-1}) = \begin{bmatrix} R_Z^T & [-R_Z t_Z]_{\times} R_Z^T \\ 0 & R_Z^T \end{bmatrix} \quad (12)$$

Then the second-block column of $\frac{dr}{dT_j}$ corresponding to rotation is:

$$\begin{aligned} & -R_Z[-R_Z t_Z]_{\times} R_Z^T + R_Z[X_i]_{\times} R_Z^T = \\ & = [t_Z]_{\times} + [R_Z X_i]_{\times} = \\ & = [R_Z X_i + t_Z]_{\times} = [\text{Act}(Z, X_i)]_{\times} \end{aligned} \quad (13)$$

10. First Order vs Second Order Study

We compare the performance of our system using Gauss-Newton optimisation against Adam. In Tab. 5, we report the resulting F-score on our Multi-Object dataset for our




Figure 6. **Reconstruction on ego-centric data.** Our system runs on different data distributions, such as ego-centric data, including TasteRob [44] and EpicKitchens [7] datasets.

system compared to a variant using the first-order Adam optimiser [14]. Interestingly, even with a generous time budget, the first-order optimiser never converged to the same level of quality, both quantitatively and qualitatively. In our experiments, the translational component of the transformation was often estimated correctly, while the rotational component remained problematic.

11. Run-Time Analysis

In Tab. 6, we report a performance breakdown of our system. The frontend dominates overall runtime, with video segmentation being the most expensive component as we

	F-Score	Time Spent
Adam 500 steps	0.6342	20s
Adam 1k steps	0.6474	40s
Adam 10k steps	0.7228	400s
Ours (10 steps)	0.7843	2s
Ours (50 steps)	0.7948	10s

Table 5. **Reconstruction quality of our method with different optimisers and computational budgets.**

propagate masks for all potential objects in the scene (typically around 50-100 objects). Backend performance scales with image resolution: our dense alignment at 512×512




Figure 7. **Failure Cases.** Our system struggles with rapid motion (e.g. shaking the milk carton) and object re-identification when hands disappear from and re-enter the camera view.

Component	Time Spent	Peak Memory
Frontend (correspondence)	2.3s	6.9GB
Frontend (π^3)	5.6s	8.7GB
Frontend (segmentation)	42s	13.1GB
Frontend (combined)	50s	-
Backend	9.6s	14.5GB
Motion Segmentation	2s	7.6GB

Table 6. **Performance and Memory benchmarking.**

takes approximately 10 seconds, reducing to 2.5 seconds at 256×256 .

Further optimisation of the backend is possible, though the frontend now represents the primary opportunity for improvement. Additionally, our current implementation does not include early termination of the optimisation; incorporating this could yield further performance gains.

References

- [1] Ijaz Akhter, Yaser Sheikh, Sohaib Khan, and Takeo Kanade. Nonrigid structure from motion in trajectory space. In *Neural Information Processing Systems (NeurIPS)*, 2008. 3
- [2] M. Bloesch, J. Czarnowski, R. Clark, S. Leutenegger, and A. J. Davison. CodeSLAM — learning a compact, optimisable representation for dense visual SLAM. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [3] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3d shape from image streams. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2000. 3
- [4] Carlos Campos, Richard Elvira, Juan J. Gomez, Jose M. M. Montiel, and Juan D. Tardos. ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM. *IEEE Transactions on Robotics*, 2021. 2
- [5] Joao Costeira and Takeo Kanade. A multi-body factorization method for motion analysis. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 1995. 3
- [6] J. Czarnowski, T. Laidlow, R. Clark, and A. J. Davison. Deepfactors: Real-time probabilistic dense monocular SLAM. In *IEEE Robotics and Automation Letters*, 2020. 2
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [8] A. J. Davison, N. D. Molton, I. Reid, and O. Stasse. MonoSLAM: Real-Time Single Camera SLAM. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2007. 2
- [9] Eric Dexheimer and Andrew J Davison. Como: Compact

- mapping and odometry. In *European Conference on Computer Vision*, pages 349–365. Springer, 2024. 2
- [10] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3, 6
- [11] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *CVPR*, 2020. 6
- [12] Adam W Harley, Yang You, Xinglong Sun, Yang Zheng, Nikhil Raghuraman, Yunqi Gu, Sheldon Liang, Wen-Hsuan Chu, Achal Dave, Suyu You, et al. Alltracker: Efficient dense point tracking at high resolution. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2025. 2, 5
- [13] Jiahui Huang, Qunjie Zhou, Hesam Rabeti, Aleksandr Kovko, Huan Ling, Xuanchi Ren, Tianchang Shen, Jun Gao, Dmitry Slepichev, Chen-Hsuan Lin, et al. Vipe: Video pose engine for 3d geometric perception. *arXiv preprint arXiv:2508.10934*, 2025. 2, 3
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 3
- [15] G. Klein and D. W. Murray. Parallel Tracking and Mapping for Small AR Workspaces. In *Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR)*, 2007. 2
- [16] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017. 6
- [17] 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. 3
- [18] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames. In *Proceedings of the IEEE international conference on computer vision*, pages 4649–4657, 2017. 3
- [19] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Superpixel soup: Monocular dense 3d reconstruction of a complex dynamic scene. *IEEE transactions on pattern analysis and machine intelligence*, 43(5):1705–1717, 2019. 3
- [20] Zhengqi Li, Richard Tucker, Forrester Cole, Qianqian Wang, Linyi Jin, Vickie Ye, Angjoo Kanazawa, Aleksander Holynski, and Noah Snavely. Megasam: Accurate, fast and robust structure and motion from casual dynamic videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10486–10496, 2025. 2, 3
- [21] 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. *arXiv preprint arXiv:2510.13802*, 2025. 6
- [22] H. Matsuki, R. Scona, J. Czarnowski, and A. J. Davison. CodeMapping: Real-time dense mapping for sparse SLAM using compact scene representations. *IEEE Robotics and Automation Letters*, 6(4):7105–7112, 2021. 2
- [23] Kirill Mazur, Gwangbin Bae, and Andrew Davison. SuperPrimitive: Scene reconstruction at a primitive level. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 4
- [24] Riku Murai, Eric Dexheimer, and Andrew J Davison. Mast3r-slam: Real-time dense slam with 3d reconstruction priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [25] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [26] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 4
- [27] Martin Rünz and Lourdes Agapito. Co-fusion: Real-time segmentation, tracking and fusion of multiple objects. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017. 3
- [28] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. J. Kelly, and A. J. Davison. SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. 2
- [29] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [30] Joan Sola, Jeremie Deray, and Dinesh Atchuthan. A micro lie theory for state estimation in robotics. *arXiv preprint arXiv:1812.01537*, 2018. 3, 2
- [31] H. Strasdat, J. M. M. Montiel, and A. J. Davison. Scale Drift-Aware Large Scale Monocular SLAM. In *Proceedings of Robotics: Science and Systems (RSS)*, 2010. 3
- [32] Edgar Sucar, Zihang Lai, Eldar Insafutdinov, and Andrea Vedaldi. Dynamic point maps: A versatile representation for dynamic 3d reconstruction. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2025. 3
- [33] Phil Tresadern and Ian Reid. Articulated structure from motion by factorization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 3
- [34] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle Adjustment — A Modern Synthesis. In *Proceedings of the International Workshop on Vision Algorithms, in association with ICCV*, 1999. 3, 5
- [35] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13:376–380, 1991. 6

- [36] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [37] 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. 3
- [38] Yihan Wang, Lahav Lipson, and Jia Deng. Sea-raft: Simple, efficient, accurate raft for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. 2
- [39] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. π^3 : Scalable permutation-equivariant visual geometry learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026. 3, 4, 5, 6
- [40] Binbin Xu, Wenbin Li, Dimos Tzoumanikas, Michael Bloesch, Andrew Davison, and Stefan Leutenegger. MID-Fusion: Octree-based object-level multi-instance dynamic slam. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2019. 3
- [41] 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 *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. 2
- [42] Songyan Zhang, Yongtao Ge, Jinyuan Tian, Guangkai Xu, Hao Chen, Chen Lv, and Chunhua Shen. Pomato: Marrying pointmap matching with temporal motions for dynamic 3d reconstruction. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2025. 3, 6
- [43] 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. 3
- [44] Hongxiang Zhao, Xingchen Liu, Mutian Xu, Yiming Hao, Weikai Chen, and Xiaoguang Han. Taste-rob: Advancing video generation of task-oriented hand-object interaction for generalizable robotic manipulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [45] Chenhao Zheng, Jieyu Zhang, Mohammadreza Salehi, Ziqi Gao, Vishnu Iyengar, Norimasa Kobori, Quan Kong, and Ranjay Krishna. One trajectory, one token: Grounded video tokenization via panoptic sub-object trajectory. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2025. 4
- [46] Kaichen Zhou, Yuhan Wang, Grace Chen, Xinhai Chang, Gaspard Beaudouin, Fangneng Zhan, Paul Pu Liang, and Mengyu Wang. Page-4d: Disentangled pose and geometry estimation for 4d perception. *arXiv preprint arXiv:2510.17568*, 2025. 3