

KV-Tracker: Real-Time Pose Tracking with Transformers

Marwan Taher Ignacio Alzugaray Kirill Mazur Xin Kong Andrew J. Davison
Dyson Robotics Lab, Imperial College London

{m.taher, i.alzugaray, k.mazur21, x.kong21, a.davison}@imperial.ac.uk

Abstract

Multi-view 3D geometry networks offer a powerful prior but are prohibitively slow for real-time applications. We propose a novel way to adapt them for online use, enabling real-time 6-DoF pose tracking and online reconstruction of objects and scenes from monocular RGB videos.

Our method rapidly selects and manages a set of images as keyframes to map a scene or object via π^3 [33] with full bidirectional attention. We then cache the global self-attention block’s key-value (KV) pairs and use them as the sole scene representation for online tracking. This allows for up to $15\times$ speedup during inference without the fear of drift or catastrophic forgetting. Our caching strategy is model-agnostic and can be applied to other off-the-shelf multi-view networks without retraining.

We demonstrate KV-Tracker on both scene-level tracking and the more challenging task of on-the-fly object tracking and reconstruction without depth measurements or object priors. Experiments on the TUM RGB-D, 7-Scenes, Arctic and OnePose datasets show the strong performance of our system while maintaining high frame-rates up to ~ 27 FPS.

Project Page: https://marwan99.github.io/kv_tracker/

1. Introduction

Scene reconstruction and tracking, especially at the object level, are important for applications such as AR/VR, self-driving, and robotics. While it is a well-studied topic, accurate and robust real-time 3D tracking remains challenging even if precise object shape models are available in advance. It is even more difficult when a representation of the object must also be acquired as part of the real-time process.

Multi-view 3D reconstruction models based on transformers and trained on large datasets are currently having a huge impact on 3D vision. The rapid development of these models began with DUST3R [32] which can regress point-maps from a pair of views. This work led to MAST3R [11], which provided features that can be used for explicit 2-view matching. Models with multi-view inputs emerged




Figure 1. Real-time object tracking and online reconstruction using an RGB images from a camera scanning a telephone booth. Red frustums indicate keyframes used for mapping; blue frustums indicate tracking frames processed at ~ 27 FPS. The geometry shows the global reconstruction obtained by fusing point maps from keyframes.

next, such as VGGT [29], π^3 [33], and MapAnything [8]. Since these models are attention-based, they have $O(N^2)$ computational cost in the number input frames. Pair-wise networks ($N = 2$) have successfully been used as the front-end for real-time 3D vision systems (e.g. MAST3R-SLAM [15]), where they provide strong calibration-free priors in a single pass, complementing explicit geometry-based optimisation.

However, it is more challenging to make real-time use of multi-view models ($N \geq 2$), which restricts their application in real-time systems [13]. These models are powerful but monolithic, making it unclear how to apply them to a continuous image stream. For instance, after running a model with 50 frames to generate a 3D reconstruction, what should one do when a new frame arrives? Run the model again from scratch with all 51 frames? This is not a scalable approach to online tracking and reconstruction.

Interesting architectures designed and trained to consume a stream of images rather than an off-line set have recently emerged. These build an implicit memory of previous observations [28, 31]. While locally the output might be consistent, globally these methods suffer from drift and

lack the ability to “close the loop” when trajectories get long and camera positions are revisited. This is particularly problematic in object tracking, where a rotating object quickly returns to previous relative camera/object poses.

In this paper we propose *KV-Tracker*, which uses multi-view networks to track in real-time while maintaining the rich multi-view observation from previous frames. We achieve this by extracting a set of key-value (KV) pairs from the global attention blocks and saving them as a representation of the observations during mapping inferences. During tracking, we can perform inference with a single query frame, re-localising the current observation against the saved KV representation without altering it, thus not poisoning the stored representation.

This significantly improves the runtime inference speed, enabling 27 FPS tracking, a $15\times$ speed-up over recomputing the keys and values for frames whose geometry we already know. This makes it possible to use these multi-view models in an online fashion, since the redundant, repetitive compute is significantly reduced from quadratic to linear complexity in the number of keyframes.

We showcase applications of such a method for camera tracking against a scene and zero-shot object tracking with no prior information about the object or depth, with the latter being more challenging since it has fewer visual cues than a typical full scene observation. Typically, object tracking methods have required access to object CAD models or depth measurements, which are restrictive requirements that we relax with our proposed method.

Our key contributions are the following:

- A novel application of multi-view feed-forward geometry networks to real-time object and scene tracking.
- A new scene representation based on cached Key-Value pairs.
- A method for zero-shot object tracking and reconstruction on the fly with no object-specific prior information.

2. Related Work

2.1. Classical Reconstruction and Representations

Traditional SLAM and reconstruction methods can be broadly categorized by their tracking strategy. Dense methods like DTAM [17] and KinectFusion [16] track against explicit dense models of the scene, enabling robust localization. In contrast, keyframe-based methods like PTAM [9] and ORB-SLAM [14] track against sparse features from selected keyframes, achieving real-time performance through reduced representation complexity.

Our approach fits within the tracking against a model paradigm. The cached key-value pairs from multi-view attention serve as a learned scene representation that encodes geometric and appearance information from keyframes.

2.2. Multi-view Reconstruction Models

DUST3R [32] introduced a 2-view reconstruction model that, without requiring any calibration information, predicted point maps end-to-end from pixels, with one of the frames acting as a reference frame. To enhance the estimated geometry, MAST3R [11] added a feature prediction head along with the point map prediction head for dense correspondences, which was later used by MAST3R-SfM [4] to solve an offline structure-from-motion problem. MAST3R-SLAM [15] introduced a real-time SLAM system built with MAST3R as its front-end.

To overcome the pairwise limitation, Fast3R [27] operated on multi-view information with global all-to-all self-attention between frames, inferring geometry from multi-view inputs in a single forward pass. VGGT [29] added local self-attention layers to alternate with global self-attention and trained a network with multi-task learning on other outputs, providing strong generalisation.

π^3 [33] adopted VGGT’s network architecture but reformulated the scene geometry output to local point maps instead of global point maps, dropped the special camera register tokens, and trained a model with a permutation-invariant loss, so the model is less sensitive to the reference view choice. Following the same architecture, MapAnything [8] added optional input modalities as conditioning, such as depth and camera intrinsics, along with more decoding heads for depth maps and other modalities.

In all these multi-view networks, information sharing happens between the multi-views via global all-to-all self-attention, with the computation growing quadratically with the number of input views, causing the inference speed to drop. Addressing this limitation for streaming settings is the core motivation behind our work.

Streaming Models To address online operation on image streams, Spann3R [28] processed a sliding window of images incrementally online, predicting global point maps for the input views in a shared reference frame. Unlike our approach, their memory was generated from each frame independently, rather than jointly.

CUT3R [30] processed input images online via a recurrent neural network, incorporating the current view’s information into the hidden state representation and updating it. Since it was trained on a maximum of 64 views, it struggled to handle long sequences, leading to drift and catastrophic forgetting. To increase robustness, TTT3R [2] added a confidence-guided state update rule via the attention map between state queries and the observation keys; this allowed for better handling of longer sequences. However, since this state is updated with every observation, TTT3R is still prone to drifting and requires state resets to operate on long sequences. Our method does not have this state corruption issue, since our KV-cache-based state representation is an-

chored in a set of keyframes and not updated on every new input frame.

Long3R [3] maintains an implicit memory that is updated via an attention-based memory gating mechanism. Kinaema [20] introduced an RNN-based network that can re-localise query views against previously observed parts of the scene that are encoded into a latent memory. MUST3R [1] uses a latent memory that gets updated with sufficient viewpoint change in the online setting. This operating mode is similar to ours, though ours leverages off-the-shelf trained models and does not require extra training or fine-tuning.

StreamingVGGT [38] trains a network with causal attention to model streaming inputs. Point3R [37] trains a network to run in a streaming fashion with spatial features encoded at explicit 3D locations in the scene serving as memory that gets updated and fused with new information. Similarly, our KV-cache memory is anchored in the keyframes; however, we do not require any training or fine-tuning.

2.3. Object Tracking and Pose Estimation

Object-level 6-DOF tracking is challenging due to limited visual context compared to full scenes, leading most methods to rely on priors such as CAD models or depth sensors.

CAD-Based Methods Many object pose estimation methods rely on having CAD models of the objects a priori. DeepIM [12] was an RGB-based method that predicted the relative 6-DOF pose between an input image and a rendering of the object’s CAD model via an iterative matching network. Having a CAD model is a strong assumption that can be limiting in different applications. Similar to ours, FoundPose [18] used off-the-shelf pretrained models for object pose estimation; however, they relied on having CAD models. Several lines of work, such as MegaPose [10] and FoundationPose [36], followed this recipe of render-and-compare, scaling up network training by training on large synthetic datasets. These methods, however, required depth information at estimation time.

RGB-D Reconstruction Like our method, BundleTrack [34] reconstructed objects on the fly. An object pose graph was built and solved online, recovering the poses of keyframes and using them for tracking. BundleSDF [35] built on top of this and learned a neural signed distance field as an object representation. These methods relied on depth measurements.

RGB Reconstruction OnePose [24] and OnePose++ [7] scan objects and build offline sparse and semi-dense object reconstructions, respectively. They also trained 2D-3D matching networks through which they solve PnP and recover the object pose. These two methods are the closest

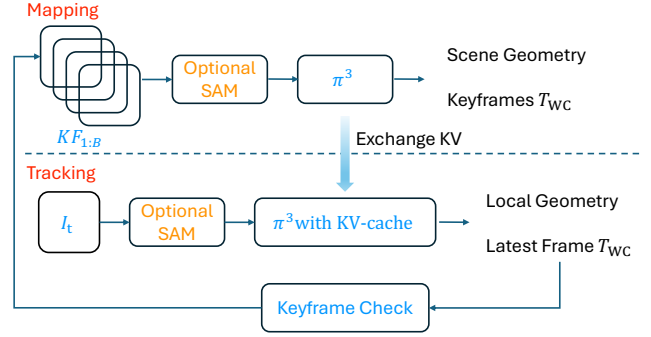


Figure 2. This is an overview of our real-time tracking and online reconstruction method. Our system is decomposed into 2 interleaved steps mapping and tracking. During the mapping stage a set of keyframes $KF_{1:B}$ are used to generate a KV-cache. During the tracking the latest frame I_t tracked using the latest available cached key-value pairs. The KV-cache is refreshed with new keyframe insertions. During object level mode, segmentation masks can be used, along with our keyframing system.

to ours in the object-level setting, as they operate on RGB inputs; however, unlike our approach, their reconstruction is performed offline rather than online.

3. Method

3.1. Overview

Our objective is to use π^3 [33] in an online fashion, building a map via a set of keyframes and tracking the latest frames in real-time as shown in Figure 2. We provide an analysis of the model’s architecture and present a novel way to enable online operation on streaming images. We demonstrate this via 2 tasks: real-time camera tracking and real-time object tracking. Our method does not require any further training or fine-tuning.

Our method consists of 2 interleaved processes, that can be parallelised similar to PTAM [9]:

- **Mapping**, where a set of input images is automatically selected as keyframes, from which we obtain a KV cache as an implicit scene representation.
- **Tracking**, where the latest frame’s I_t state can be estimated in real-time against the latest KV-cache scene representation, achieving up to **27 FPS**.

3.2. Details of π^3

π^3 is feed-forward decoder only transformer based model. Given a set of input images $I_n \in \mathbb{R}^{H \times W \times 3}$, $n = 1, \dots, N$, the model will predict a set camera poses in a world frame $T_n \in \mathbb{SE}(3)$, point maps in the local camera frames $P_n^c \in \mathbb{R}^{H \times W \times 3}$ and a confidence score per 3D point $C_n \in \mathbb{R}^{H \times W}$.

Initially the input images are patchified into M patches and encoded via a ViT backbone to produce a set of tokens

$\text{Enc}(I_{1:N}) = X_{1:N}$, $X \in \mathbb{R}^{M \times d_k}$, where d_k is the feature vector dimension.

These set of tokens are then passed through L decoding layers, alternating between frame-wise self-attention layers and global self-attention layers. These feature aggregating blocks enable multi-view information exchange. In the global self-attention blocks, every input image patch can attend to all the other input patches. This is the key to producing globally consistent 3D output.

Finally, decoding heads are used to predict the final outputs, $\text{decoder.heads}(X_{1:N}) = T_{1:N}, P_{1:N}, C_{1:N}$. A key insight is that each frame's tokens X_n are decoded independently from the others, which is a crucial feature needed for adapting this model for online use.

3.3. Processing Bottleneck

The model uses Scaled Dot-Product Attention [26] for the frame-wise and global self-attention blocks. Recall that attention is computed via:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (1)$$

where queries Q , keys K and values V are computed via linear projection of the tokens X , parametrised by θ_l , where l is the layer index, as below:

$$\text{Proj}(X; \theta_l) = Q, K, V. \quad (2)$$

The complexity of the attention computation is a function of the size of Q and K . In the frame-wise self-attention blocks, a batch of per-frame tokens $X_{1:N}$ gets projected into $Q_{1:N} \in \mathbb{R}^{M \times d_k}$ and $V_{1:N} \in \mathbb{R}^{M \times d_k}$ which yields attention maps which are quadratic with the number of patches per frame, giving $\mathcal{O}(M^2)$ computational complexity.

However in global self-attention block, all tokens from all frames attend to each other. More formally, all the tokens get concatenated together into $X \in \mathbb{R}^{NM \times d_k}$ and projected to $Q \in \mathbb{R}^{NM \times d_k}$ and $V \in \mathbb{R}^{NM \times d_k}$ which yields attention maps which are quadratic in the number of input images, giving $\mathcal{O}((NM)^2)$ computational complexity. As a result of this all-to-all bidirectional attention, the computation becomes prohibitively slow when the number of frames grows, which limits the application of such model in online settings.

3.4. Tracking via Caching

To circumvent the computational bottleneck caused by the global self-attention layer, we propose running the full model on a set of keyframes $KF_{1:B} \subseteq I_{1:N}$, where B is the number of keyframes in the keyframe buffer, with full all-to-all bidirectional attention in the global self-attention blocks, while caching the computed $\tilde{K}_{1:B}^l, \tilde{V}_{1:B}^l, l \in \{1, \dots, L\}$ from each global self-attention block.

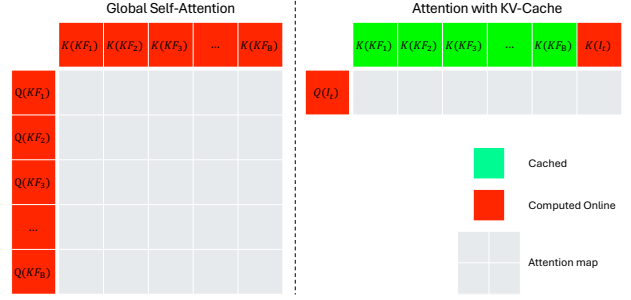


Figure 3. Visualisation of the full bidirectional global self-attention used during mapping to generate the KV-cache (left) vs. cross-attention with KV-cache + self-attention used for tracking of live frames (right).

Then, when a new live frame I_t arrives, we obtain state estimates online via the following method. We encode only I_t into $X_t = \text{Enc}(I_t)$, and during the feature aggregation step we run frame-wise self-attention on X_t only. Then, during global self-attention we project X_t obtaining $\text{proj}(X_t, \theta) = Q_t, K_t, V_t$, and attend the cached keys \tilde{K} and values \tilde{K} via $\text{Attention}(Q_t, [\tilde{K}, K_t], [\tilde{V}, V_t])$. The attention map is now obtained via:

$$QK^T = \underbrace{[Q_t]}_{M \times d_k} \underbrace{\begin{bmatrix} \tilde{K}_{kf_1}^T & \dots & \tilde{K}_{kf_M}^T & \tilde{K}_t^T \end{bmatrix}}_{d_k \times M(N+1)} \quad (3)$$

This brings down the complexity from $\mathcal{O}((NM)^2)$ to $\mathcal{O}(M^2(N+1))$ which is much cheaper as we demonstrate in our experiments. Finally, the tokens X_t processed through the feature aggregator are decoded normally via $\text{Decoder}(X_t)$, obtaining the camera pose and scene geometry. Figure 3 shows the difference between the Bidirectional attention used in the mapping process and the the unidirectional attention used for tracking.

3.5. KV-cache as a Scene Representation

In scene reconstruction and SLAM settings, a scene gets represented by a set of primitives which can be explicit like sparse 3D points [14] or implicit like with an MLP [23]. These primitives are used to encode the scene from which the scene geometry and camera poses can be recovered during mapping and tracking.

In our formulation the KV-cache obtained via the keyframes can be interpreted as a scene representation since it was generated by exchanging multi-view information and it allows for the recover of scene geometry and camera for a given query frame.

3.6. Object Tracking

Along with using this system for scene level camera tracking, we propose applying it to the task of online object re-

construction and tracking, which is more difficult both because an object usually only fills a small fraction of the pixels in an image, and because manipulated objects often move and rotate rapidly relative to the camera. The prior information embedded in our approach can be powerful in dealing with the challenges of object tracking, which also fits nicely with what tracking via caching can offer, since less than 50-60 keyframes are usually sufficient to represent a comprehensive view of an object, enabling full reconstruction of the object while maintaining high frame rate tracking.

Real-time segmentation is now a standard computer vision component, and we make use of this in object tracking. Given segmentation masks of the object of interest, the background pixels can be masked out as black pixels. We use a simple keyframing criterion to trigger adding new keyframes. When a new keyframe is added, the KV-cache gets recomputed fully.

Keyframing We employ an angular threshold-based keyframe selection strategy to ensure sufficient baseline and viewpoint diversity for multi-view reconstruction. Specifically, a new frame I_t is designated as a keyframe when the minimum azimuth or elevation difference between the current frame and all existing keyframes exceeds a threshold:

$$\min_{kf \in KF} |\phi_t - \phi_{kf}| > \tau \quad \text{or} \quad \min_{kf \in KF} |\theta_t - \theta_{kf}| > \tau,$$

ϕ and θ represent the camera’s azimuth and elevation angles, while τ is the threshold angle, balancing reconstruction coverage with memory efficiency.

This view-based criterion naturally adapts to camera motion, selecting more keyframes during viewpoint changes while avoiding redundancy during small motions or stationary periods. By comparing against all existing keyframes rather than only the most recent, our strategy prevents adding redundant views even when the camera revisits previously observed regions.

To improve robustness of the tracking and the quality of the mapping, we employ a keyframe rejection policy, where keyframes with a predicted low confidence are pruned and the system reverts back to the previous KV-cache.

4. Experiments

In this section, we provide results on camera pose estimation (4.1), object pose tracking (4.2), as well as run time analysis of our system and an ablation over the effectiveness of our KV-cache self-attention scheme (4.3).

Our method adapts to any model with an architecture similar to π^3 , which contains a feature aggregation blocks with global self-attention blocks [8, 29]. We choose π^3 since it drops the camera register token, making it less sensitive to reference views compared with VGGT.

Since π^3 decodes the estimated camera, poses local point maps and confidence maps with 3 independent decoders, we leverage this feature and turn off the the point map and confidence decoding head, to provide speed up during tracking. Having them switched on or off only impacts the runtime performance and not the quality of tracking. To recover the scene geometry, we run with all the decoding heads enabled.

During mapping, not only do we cache the the keys and values associated with the image patch tokens, we also cache the associated per frame register tokens key value pairs. Our memory usage grows linearly with the number of keyframes used in mapping. From that it follows that our memory is also a function of the number of input images resolution, making the cache size linear with the number of patches per image.

The reported results were performed on a machine with an NVIDIA GeForce RTX 4090 graphics card, unless otherwise stated.

4.1. Camera Tracking Evaluation

Baselines We focus evaluation on other baselines that represent the scene implicitly via hidden states or memory mechanism. For that, we choose Point3R [37], CUT3R [30] and TTT3R [2] as our baselines, since they can run in a streaming fashion similar to ours. For a fair evaluation, we run CUT3R and TTT3R with state reset every 100 Frames, to avoid catastrophic drifting. Since LONG3R [3] and Ki-naema [20] haven’t open sourced, we couldn’t run their methods and compare against them.

For this task, we run our method with automatic keyframing every 50 frames. We evaluate on all the frames in the scenes.

We run our camera tracking evaluation on 7-scene and TUM RGB-D, we report standard RMSE Absolute Translation Error (ATE) metrics for the camera trajectory over the full sequences. Since our method and the baselines predict geometry up to a scale, we align all the estimated trajectories to the ground truth trajectories with a Sim(3) alignment using Umeyama Algorithm [6].

We present our trajectory evaluation results for the TUM-RGBD dataset [22] in Table 1 and the 7-scenes dataset [21] on seq-01 in Table 2. Our method achieves the lowest average ATE score in both datasets, outperforming our strongest baseline TTT3R by 18% on TUM-RGBD and by 44% on 7-Scenes. Notably, our approach wins on 6 out of 8 scenes on TUM RGB-D and 6 out of 7 scenes on 7-Scenes, demonstrating consistent performance across diverse indoor environments.

The improvement is particularly pronounced on challenging sequences: on TUM RGB-D’s “teddy” scene, we achieve 0.071m error compared to TTT3R’s 0.214m (67% reduction), and on 7-Scenes’ “fire” sequence, we obtain 0.042m versus 0.124m (66% reduction).




Figure 4. Tracking and reconstruction demo on a moving car, captured using a phone in the wild.

Table 1. Absolute Trajectory Error (ATE) RMSE in meters on TUM-RGBD dataset.

Scene	Point3R	CUT3R	TTT3R	DPVO	Ours
360	0.200	0.176	0.110	0.135	0.166
desk	0.366	0.196	0.104	0.038	0.060
desk2	0.321	0.437	0.147	0.048	0.083
plant	0.423	0.383	0.092	0.036	0.048
room	0.558	0.423	0.253	0.394	0.366
rpy	0.062	0.054	0.054	0.034	0.045
teddy	0.580	0.399	0.214	0.064	0.071
xyz	0.134	0.109	0.083	0.012	0.021
Average	0.331	0.272	0.132	0.095	0.108

Table 2. Absolute Trajectory Error (ATE) RMSE in meters on 7-Scenes dataset.

Scene	Point3R	CUT3R	TTT3R	Ours
chess	0.427	0.297	0.154	0.091
fire	0.280	0.218	0.124	0.042
heads	0.389	0.115	0.097	0.054
office	0.436	0.356	0.196	0.065
pumpkin	0.644	0.249	0.228	0.142
redkitchen	0.502	0.118	0.136	0.038
stairs	0.398	0.079	0.063	0.128
Average	0.439	0.205	0.143	0.080

These results were obtained by running Point3r on 224x224, since it was running out of memory on higher resolutions. CUT3R and TTT3R were run with 512x384, to evaluate the ceiling of their performance and ours was run with 350x266. Not only can our method run at a lower resolution and outperforms all our baselines, it also does this while maintaining a higher frame rate as shown in Table 3. In the supplementary we provide extra results running with different resolutions. we also provide more results in the supplementary with the CUT3R and TTT3R run without state resets for completeness.

We also included DPVO [25] in our results on the TUM-

Table 3. The estimated average FPS given the time consumed to run a single frame, for the resolutions reported.

Method	Point3R	CUT3R	TTT3R	Ours
FPS	5	17	17	27

RGBD dataset in Table 1. Since it is a sparse-patch odometry system we do not consider it as a baseline since it lacks dense geometry prediction. But we include it as reference since it is a strong system that can be used for tracking. Our method performs very competitively against it, with only 0.013m difference on the average ATE RMSE.

4.2. Object Tracking Evaluation

Object tracking provides a challenging test case for evaluating how well off-the-shelf reconstruction models, trained on general-purpose datasets, generalize to object-centric scenarios. We evaluate our method in two settings: (1) comparison with other streaming reconstruction methods (CUT3R, TTT3R) on the ARCTIC dataset, and (2) comparison with methods specifically designed for object tracking (OnePose, OnePose++) on their respective benchmarks.

Tracking protocol. Given an initial segmentation mask of the target object, we use SAM 2 [19] to propagate masks throughout the sequence for every input frame. Image acquisition and segmentation run in a separate process from tracking and mapping. All methods process images with backgrounds masked to black pixels. We use our object-level keyframing strategy with azimuth and elevation thresholds of 10 for keyframe selection.

4.2.1. Evaluation on ARCTIC Dataset

The ARCTIC dataset [5] contains sequences of various objects being manipulated by a person in tabletop settings. We use the egocentric camera view from sequences from S01 for evaluation.

We compare against CUT3R and TTT3R as baselines, running all methods on the same set of segmented frames. Table 4 shows that KV-Tracker achieves the lowest average ATE of 0.228m, outperforming both baselines (CUT3R:

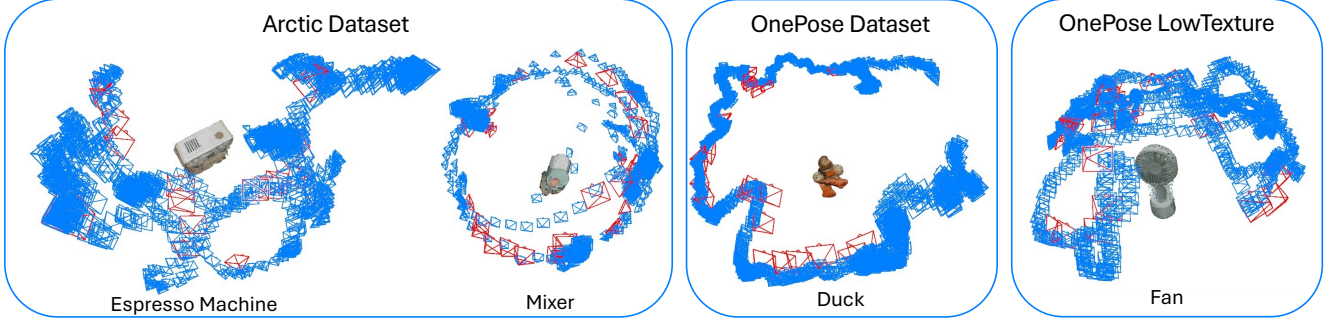


Figure 5. Sample trajectory and reconstruction results from the object tracking evaluation datasets. The Arctic dataset samples are run with 308×308 images. While the OnePose dataset and OnePose Low Texture dataset samples are run on 518×518 .

Table 4. Absolute Trajectory Error (ATE) RMSE in meters on the Arctic dataset [5].

Scene	CUT3R	TTT3R	Ours@308
espressomachine	0.253	0.175	0.151
ketchup	0.369	0.319	0.249
microwave	0.181	0.156	0.135
box	0.176	0.207	0.200
laptop	0.319	0.305	0.248
waffleiron	0.336	0.342	0.204
scissors	0.237	0.264	0.188
capsulemachine	0.458	0.552	0.300
phone	0.408	0.399	0.402
mixer	0.313	0.309	0.198
Average	0.305	0.303	0.228

0.305m, TTT3R: 0.303m).

Notably, none of the evaluated methods, including ours, were explicitly trained on masked images. The strong performance of our approach, which is based on π^3 , suggests that the learned priors generalize well to object-centric tracking with background masking. CUT3R and TTT3R showed similar average performance to each other but exhibited greater variance across individual scenes.

4.2.2. Eval on OnePose Datasets

OnePose [24] introduced a dataset of everyday objects, while its successor OnePose++ [7] contributed a dataset focused on low-texture objects. Ground-truth 3D bounding boxes are manually annotated, and camera poses are obtained via ARKit tracking.

Both baseline methods perform an offline reconstruction phase before tracking, providing them with complete 3D object models—information not available to our online method. Additionally, the baselines do not require object segmentation as input; instead, they leverage annotated 3D bounding boxes from the offline mapping phase, which they project into subsequent frames to estimate object locations.

These bounding boxes are sometimes oversized relative to the actual objects.

We evaluate our system with segmentation masks from SAM 2 [19]. For fair comparison, we also report results using 2D bounding boxes extracted from these masks with 50-pixel dilation to approximate the spatial extent used by the baselines.

We follow the standard evaluation protocol, reporting recall at error tolerances of 1cm–1°, 3cm–3°, and 5cm–5°. We additionally ablate our method across two input resolutions: 308×308 and 518×518 pixels.

OnePose Dataset Despite building the object map online without any offline reconstruction phase, KV-Tracker achieves competitive performance compared to the baselines. At the most permissive 5cm, 5° threshold, our high-resolution variant (Ours@518 with 2D bounding box) achieves 92.85% recall, outperforming both OnePose (84.1%) and OnePose++ (87.7%). At the intermediate 3cm, 3° threshold, our method achieves 69.29%-75.51% recall compared to OnePose++’s 80.8%, a modest gap considering that the baseline methods leverage pre-built maps with known 3D structure.

The performance gap widens at the strictest 1cm, 1° threshold, where our method achieves 10.71% recall (segmentation mask) compared to OnePose++’s 51.1%. OnePose++ benefits from offline multi-view reconstruction that provides more complete and accurate 3D models, while our online reconstruction may have incomplete coverage during initial frames. Nevertheless, our method’s ability to achieve competitive coarse-grained tracking while operating fully online at higher frame rates (16-27 FPS vs. 11-15 FPS) demonstrates a practical trade-off between accuracy and the flexibility of online operation. The reported FPS values of OnePose and OnePose++ were obtained on a V100 GPU.

Table 5. Performance comparison on OnePose and OnePose Low Texture datasets. Accuracy (%) at different error thresholds. Baseline methods use offline 3D reconstruction.

Method	Input Type	OnePose			OnePose Low Texture			FPS
		1cm, 1°	3cm, 3°	5cm, 5°	1cm, 1°	3cm, 3°	5cm, 5°	
OnePose	3D Bbox	49.7	77.5	84.1	12.4	35.7	45.4	15
OnePose++	3D Bbox	51.1	80.8	87.7	16.8	57.7	72.1	11
Ours @ 308	Seg. Mask	2.25	47.0	75.7	1.21	32.2	62.2	27
Ours @ 308	2D Bbox	2.9	52.8	83.2	4.15	57.3	83.3	27
Ours @ 518	Seg. Mask	10.7	75.5	92.1	6.85	62.0	85.7	16
Ours @ 518	2D Bbox	5.3	69.3	92.9	12.1	80.0	94.4	16

OnePose LowTexture Dataset Our method outperforms the baselines on the low-texture dataset at coarser thresholds. Our high-resolution variant with 2D bounding boxes achieves 94.42% recall at 5cm, 5° (vs. OnePose++ 72.1%) and 80.04% at 3cm, 3° (vs. 57.7%), while maintaining comparable performance at the strictest 1cm, 1° threshold (12.1% vs. 16.8%) and operating at 16 FPS compared to OnePose++’s 11 FPS.

Notably, even our lower-resolution variant (308px with 2D bounding box) outperforms both baselines at 5cm, 5° and 3cm, 3° thresholds while running at nearly 3× the frame rate (27 FPS vs. 11 FPS). This demonstrates that our approach generalizes better to challenging low-texture scenarios than traditional feature-based methods, even without offline reconstruction.

Across both datasets, on average our variants using dilated 2D bounding boxes outperform those using segmentation masks. We attribute this to the inclusion of background context around object boundaries, which provides additional discriminative features. This effect is particularly pronounced, since the objects in these datasets are stationary while the camera is scanning them, where boundary regions contain informative visual cues.

4.3. Runtime Analysis

We conduct an experiment to demonstrate the effectiveness of our adaptation to π^3 for accelerating inference through KV-cache reuse from the mapping phase, compared to performing full bidirectional attention with fresh key-value computation for all keyframes.

We run a synthetic workload using images of size 308 × 308 pixels, measuring frames-per-second (FPS) throughput under two conditions: (1) full inference on N frames with all-to-all attention, and (2) inference on a single query frame using cached keys and values from N keyframes. As shown in Figure 6, the full bidirectional attention approach exhibits quadratic scaling with $O(N^2)$ complexity, resulting in rapidly degrading speed as the number of frames increases. In contrast, our KV-cache approach maintains 30




Figure 6. Frames per second (FPS) throughput comparison: processing N frames with full bidirectional attention vs. processing a single query frame with KV-cache from N frames.

FPS for up to 50 frames, then degrades more gradually, sustaining 25 FPS with cache sizes up to 70 frames and remaining above 20 FPS for cache sizes up to 110 frames. We report results up to the point of running out of memory (24GB) during inference.

These results demonstrate that our system can efficiently map small workspaces or objects and relocalize with respect to them in real-time, maintaining practical frame rates even with substantial mapping history.

5. Conclusion

We presented KV-Tracker, a training-free approach for adapting 3D reconstruction models to real-time camera tracking and online reconstruction. By caching key-value pairs from keyframes as memory, our method enables efficient tracking and online reconstruction of objects and small workspaces.

Our system’s memory requirements currently limit its application to spatially confined environments or objects. Future work includes exploring cache pruning, compression, and efficient incremental KV computation to scale toward full SLAM systems capable of handling larger scenes.

References

- [1] Johann Cabon, Lucas Stofl, Leonid Antsfeld, Gabriela Csurka, Boris Chidlovskii, Jerome Revaud, and Vincent Leroy. Must3r: Multi-view network for stereo 3d reconstruction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1050–1060, 2025. 3
- [2] Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Ttt3r: 3d reconstruction as test-time training. *arXiv preprint arXiv:2509.26645*, 2025. 2, 5
- [3] Zhuoguang Chen, Minghui Qin, Tianyuan Yuan, Zhe Liu, and Hang Zhao. Long3r: Long sequence streaming 3d reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5273–5284, 2025. 3, 5
- [4] Bardienus Pieter Duisterhof, Lojze Zust, Philippe Weinzaepfel, Vincent Leroy, Johann Cabon, and Jerome Revaud. Mast3r-sfm: a fully-integrated solution for unconstrained structure-from-motion. In *2025 International Conference on 3D Vision (3DV)*, pages 1–10. IEEE, 2025. 2
- [5] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12943–12954, 2023. 6, 7
- [6] Michael Grupp. evo: Python package for the evaluation of odometry and slam. <https://github.com/MichaelGrupp/evo>, 2017. 5
- [7] Xingyi He, Jiaming Sun, Yuang Wang, Di Huang, Hujun Bao, and Xiaowei Zhou. Onepose++: Keypoint-free one-shot object pose estimation without cad models. *Advances in Neural Information Processing Systems*, 35:35103–35115, 2022. 3, 7
- [8] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zauss, Ethan Weber, Nelson Antunes, et al. Mapanything: Universal feed-forward metric 3d reconstruction. *arXiv preprint arXiv:2509.13414*, 2025. 1, 2, 5
- [9] Georg Klein and David Murray. Parallel tracking and mapping for small ar workspaces. In *2007 6th IEEE and ACM international symposium on mixed and augmented reality*, pages 225–234. IEEE, 2007. 2, 3
- [10] Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. Megapose: 6d pose estimation of novel objects via render & compare. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022. 3
- [11] Vincent Leroy, Johann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer, 2024. 1, 2
- [12] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 683–698, 2018. 3
- [13] Dominic Maggio, Hyungtae Lim, and Luca Carlone. Vggt-slam: Dense rgb slam optimized on the sl (4) manifold. *arXiv preprint arXiv:2505.12549*, 2025. 1
- [14] 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. 2, 4
- [15] Riku Murai, Eric Dexheimer, and Andrew J. Davison. Mast3r-slam: Real-time dense slam with 3d reconstruction priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16695–16705, 2025. 1, 2
- [16] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. Ieee, 2011. 2
- [17] 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
- [18] Evin Pinar Örnek, Yann Labbé, Bugra Tekin, Lingni Ma, Cem Keskin, Christian Forster, and Tomas Hodan. Foundpose: Unseen object pose estimation with foundation features. In *European Conference on Computer Vision*, pages 163–182. Springer, 2024. 3
- [19] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 6, 7
- [20] Mert Bulent Sariyildiz, Philippe Weinzaepfel, Guillaume Bono, Gianluca Monaci, and Christian Wolf. Kinaema: a recurrent sequence model for memory and pose in motion. *Advances in Neural Information Processing Systems*, 2025. 3, 5
- [21] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2930–2937, 2013. 5
- [22] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of rgb-d slam systems. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 573–580. IEEE, 2012. 5
- [23] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6229–6238, 2021. 4
- [24] Jiaming Sun, Zihao Wang, Siyu Zhang, Xingyi He, Hongcheng Zhao, Guofeng Zhang, and Xiaowei Zhou. Onepose: One-shot object pose estimation without cad models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6825–6834, 2022. 3, 7

- [25] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. *Advances in Neural Information Processing Systems*, 36:39033–39051, 2023. [6](#)
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [4](#)
- [27] Chung-Shien Brian Wang, Christian Schmidt, Jens Piekenbrinck, and Bastian Leibe. Faster vggv with block-sparse global attention. *arXiv preprint arXiv:2509.07120*, 2025. [2](#)
- [28] Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. In *2025 International Conference on 3D Vision (3DV)*, pages 78–89. IEEE, 2025. [1](#), [2](#)
- [29] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggv: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025. [1](#), [2](#), [5](#)
- [30] Qianqian Wang*, Yifei Zhang*, Aleksander Holynski, Alexei A. Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *CVPR*, 2025. [2](#), [5](#)
- [31] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10510–10522, 2025. [1](#)
- [32] 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](#)
- [33] 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, 2025. [1](#), [2](#), [3](#)
- [34] Bowen Wen and Kostas Bekris. Bundletrack: 6d pose tracking for novel objects without instance or category-level 3d models. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8067–8074. IEEE, 2021. [3](#)
- [35] Bowen Wen, Jonathan Tremblay, Valts Blukis, Stephen Tyree, Thomas Müller, Alex Evans, Dieter Fox, Jan Kautz, and Stan Birchfield. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *CVPR*, pages 606–617, 2023. [3](#)
- [36] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17868–17879, 2024. [3](#)
- [37] Yuqi Wu, Wenzhao Zheng, Jie Zhou, and Jiwen Lu. Point3r: Streaming 3d reconstruction with explicit spatial pointer memory. *Advances in Neural Information Processing Systems*, 2025. [3](#), [5](#)
- [38] Dong Zhuo, Wenzhao Zheng, Jiahe Guo, Yuqi Wu, Jie Zhou, and Jiwen Lu. Streaming 4d visual geometry transformer. *arXiv preprint arXiv:2507.11539*, 2025. [3](#)