

Tracking by 3D Model Estimation of Unknown Objects in Videos

Denys Rozumnyi^{1,4} Jiří Matas⁴ Marc Pollefeys¹ Vittorio Ferrari³ Martin R. Oswald^{1,2}
¹Department of Computer Science, ETH Zurich ²University of Amsterdam
³Google Research ⁴Czech Technical University in Prague

Abstract

Most model-free visual object tracking methods formulate the tracking task as object location estimation given by a 2D segmentation or a bounding box in each video frame. We argue that this representation is limited and instead propose to guide and improve 2D tracking with an explicit object representation, namely the textured 3D shape and 6DoF pose in each video frame. Our representation tackles a complex long-term dense correspondence problem between all 3D points on the object for all video frames, including frames where some points are invisible. To achieve that, the estimation is driven by re-rendering the input video frames as well as possible through differentiable rendering, which has not been used for tracking before. The proposed optimization minimizes a novel loss function to estimate the best 3D shape, texture, and 6DoF pose. We improve the state-of-the-art in 2D segmentation tracking on three different datasets with mostly rigid objects.

1. Introduction

Visual object tracking is an important task in computer vision. A wide range of challenges [20, 21], benchmarks [22], and methods [34, 33, 49] have been proposed in recent years. Given an input video and an annotation of the object in the first frame, tracking methods localize the object with a 2D segmentation (or a 2D bounding box) in all following frames. However, they do not attempt to reconstruct the 3D shape and its 6DoF motion and settle for 2D segmentation as a low-level representation of the object state. On the one hand, this makes those methods general in terms of prior assumptions on object shape and motion. On the other hand, the information extracted by the segmentation is limited, e.g. to whole object editing, and does not support many applications.

In this paper, we propose a generic approach for tracking that is fundamentally different from what standard trackers perform [22, 21]. Instead of a mere 2D segmentation, we estimate the latent object 3D shape, texture, and 6DoF pose in each video frame to improve 2D tracking. This formula-

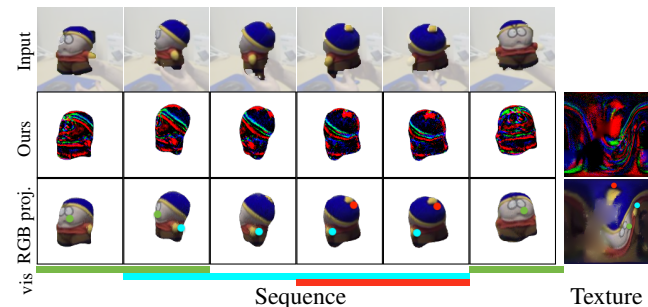


Figure 1. **Improving 2D tracking with explicit 3D modeling.** Starting from initial segmentations given by a tracker [49] (top row), we jointly optimize all object parameters to re-render the input video as close as possible. Long-term correspondences are visualized by colored points and their visibility by colored lines underneath. In contrast, the standard tracker reports essentially no motion (just slightly changing 2D segmentation). We use S2DNet features [8] for reprojection error (middle row) and back-project RGB values only for visualization (bottom row).

tion of tracking solves a more complex long-term dense correspondence problem between all 3D points on the object, even of points that are not visible (Fig. 1, colored lines). This is useful in many applications, e.g. augmented reality, motion interpretation and analysis, point matching, per-pixel editing, robotics, object manipulation and grasping.

The core of the approach is to find the object shape, pose, and texture – the parameters that define the object – whose reprojection is most consistent with input video frames. The parameter estimation is driven by re-rendering the input video via differentiable rendering instead of directly predicting the 2D segmentation in the next frame based on features extracted in the previous frames. This modern differentiable rendering way of thinking was demonstrated to be successful in various applications, such as 3D reconstruction [28, 3], structure-from-motion [26], reconstruction of clothed humans [16], and deblurring of motion-blurred objects [39, 40], but never in tracking. To initialize our optimization process, we use segmentations given by a standard 2D tracker. We demonstrate experimentally that the proposed tracking method improves segmentation accuracy of this initial tracker.

Another related line of research is multi-view 3D object reconstruction [5, 35, 52]. Those methods solve the task with given additional information, usually with camera poses (or alternatively object poses), camera calibration matrix, depth information, object segmentation, *etc.* Moreover, these methods do not process the sequences causally and reconstruct the whole scene at once. Most methods assume either a perfect segmentation or objects in front of a blank background [5, 38]. They usually require distinctive texture and are trained on large datasets, in some cases even for a particular object class [38].

A third group, methods for 6DoF pose estimation [13, 14, 23], require another type of additional information as input – the 3D shape of the object appearing in the video [13] (or a small set of 3D shapes [27, 23]). Some methods reconstruct higher than 6DoF poses [51] or poses of articulated objects [30]. In contrast, our method requires no additional information and only uses RGB information as input. To summarize, we make the following contributions:

- (1) we propose the first model-free object tracking method that explicitly estimates the latent object 3D shape, texture, and 6DoF motion from an RGB sequence. The tracker is aware of which part of its surface is visible,
- (2) we introduce differentiable rendering to tracking to optimize object parameters to re-render the input video,
- (3) we introduce deep surface textures to improve tracking robustness under view-dependent illumination changes,
- (4) experiments demonstrate improved tracking accuracy compared to the initial tracker, more accurate 2D segmentation, and dense long-term correspondences on the object surface, even after self-occlusion. It sets new state-of-the-art results on tracking datasets with rigid objects such as CDTB (part of VOT challenge), coin tracking, and In-hand Object Manipulation.

2. Related work

Visual object tracking in 2D. Visual object tracking is one of the oldest tasks in computer vision dating back to the Lucas-Kanade tracker [31, 42]. Discriminative Correlation Filters have been successfully used for the tracking task [11, 34]. More recently, methods based on deep learning have been developed [7, 33, 49]. Most trackers localize the object only in 2D by a 2D bounding box or a 2D segmentation. In contrast, we aim to improve tracking by explicitly recovering rough 3D shape, texture, and 6DoF pose.

3D reconstruction. In recent years, 3D reconstruction has gained significant attention. Structure-from-Motion pipelines, such as COLMAP [41], have become widely used. Other methods are based on Neural Radiance Fields [35, 52], Transformers [44], Recurrent Neural Networks [5]. Methods designed specifically for 3D object

reconstruction have also been proposed [44, 38, 5] However, they always assume a perfect segmentation or trivial backgrounds and impose strong priors on object shape [38]. Similarly, BackFlow [45] reconstructs simple objects that are manipulated by human hands in RGB-D videos (we use only RGB). Moreover, they assume that a user provides scribbles to segment out the background in the first frame.

3D and 6DoF object pose estimation or tracking. An approach to enhance 2D tracking with 3D reconstruction was proposed by [17], but it requires depth information as input and estimates latent 3D representation as a point cloud by concatenating RGB-D segmentations, without any optimization, viewpoint consistency, and explicit surface modeling. Thus, it loses many desirable properties, *e.g.* self-occlusions awareness and long-term point-to-point correspondence. The field of 6DoF object pose estimation is vast. A number of challenges and benchmarks have been proposed [13, 14], comprising of various methods [23, 12, 46]. However, the main limiting factor of these methods is that they need the known 3D shape of the object as input. Some methods were proposed for 6D tracking of unknown objects [37, 24, 9, 47, 36]. In contrast to our approach, these methods do not estimate the 3D shape but only try to predict 6DoF object pose given initial object segmentation.

Differentiable rendering. The classical computer graphics rendering pipeline was recently made differentiable in various ways. Most methods are based on soft rasterization [28, 3, 4], differentiable sphere tracing [29], neural radiance fields (NeRF) [35], *etc.* NeRF was also combined with pose estimation, *e.g.* [50, 25]. Differentiable rendering techniques were also used in various applications: 3D reconstruction [28, 3], monocular 6DoF pose estimation [46], reconstruction of clothed humans [16], and deblurring of motion-blurred objects [39, 40]. Differentiable rendering was used in [10] with a primary focus on 3D object scanning. It is intended for offline use as it requires all video frames at the same time. Hence, it is non-causal and cannot be used for a general tracking case. These methods are related to our work in their use of differentiable rendering, but we apply it to a different task.

As a consequence, the proposed method is the first to use differentiable rendering in the object tracking task in videos with joint 3D reconstruction and 6DoF pose estimation of unknown objects in the wild.

3. Method

We assume that the input is a video stream $\{I_1, \dots, I_N\}$ of N frames and the position of the object in the first frame, given either by a segmentation or a bounding box. As the first step of our method, we approximately segment the object with a state-of-the-art 2D tracker, *e.g.* [49, 48], resulting in a binary segmentation per frame $\{M_1, \dots, M_N\}$. The desired output of our method is improved binary segmenta-

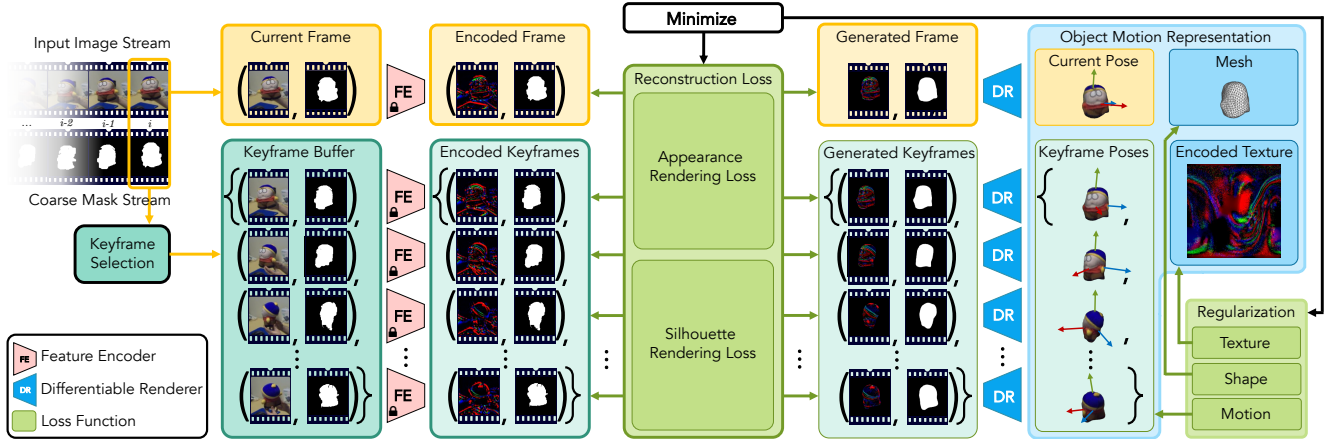


Figure 2. **Method Overview.** **Left-to-middle:** The proposed method processes a stream of images and coarse masks and jointly estimates 3D object shape, texture, and 6DoF poses. As a first step, the input images are transferred using a feature encoder (FE) to better handle view-dependent brightness changes. The feature encoder is pre-trained and remains unchanged in our optimization (as indicated by the \mathbb{L} -symbol). The feature encoder is optional and can be replaced with an identity map in which case our method computes the reconstruction loss in RGB space rather than the more general feature space, *e.g.* S2DNet [8]. Concurrently, we collect keyframes that serve as reference information for object motion estimation. **Right-to-middle:** The method can be seen as a generative approach: Starting from an object shape (mesh), texture, and 6DoF object poses, we generate images and corresponding masks with a differentiable renderer (DR). At test time, the inverse problem is solved by refining the object motion representation to minimize the reconstruction loss. Further regularization on the object shape, texture, and motion further stabilizes the optimization and helps avoiding undesired local minima.

tion and additional latent object representation, which is a textured 3D shape of the tracked object and its 6DoF pose in each frame. We obtain improved segmentations by rendering the estimated 3D object at the estimated 6DoF pose in each frame. The improvement comes from seeking for a 3D shape, texture, and 6DoF poses that re-render the input video frames as accurately as possible. We optimize over the space of 3D shapes, textures, and 6DoF poses with differentiable rendering.

The whole pipeline (Fig. 2) is designed in such a way that the initial segmentations only guide the object’s 6DoF tracking and 3D reconstruction without being overly dependent on their accuracy. We run the optimization sequentially for each frame, making the method causal, *i.e.* depending only on the past. Hence, we perform online tracking.

In principle, at each new frame we could optimize the object 3D shape, texture, and 6DoF poses over all past video frames. However, for efficiency we maintain a subset of so-called *keyframes* indexed by $K_n \subseteq [1, \dots, n-1]$, which act as sparse anchors on which the optimization is performed. The main idea is that each keyframe provides a significantly different view of the object. The subset of keyframes is sequentially updated every time the process moves to a new frame, based on criteria described in Sec. 3.3.

3.1. Representations

Object shape representation. We estimate a single 3D object shape that is fixed throughout the video. The object is represented in its canonical space by a textured triangulation

with a set of vertices, faces, and texture mapping. The shape parameters \mathcal{P} are vertex offsets from a prototype mesh, which deform it into the actual shape of the object in the video. The mesh topology and faces are fixed.

Object appearance representation. The appearance of the object is represented by a texture map of features \mathcal{T} . The function that maps RGB values to a feature representation is denoted by $F(\cdot)$. In the simplest case, F can be an identity mapping. Then, the features are simply RGB values. However, in most cases, when the object moves, the RGB values change due to lighting effects, shadows, illumination changes, *etc.* To overcome this problem, invariant features have been studied in image matching and structure-from-motion, *e.g.* [26]. We use a similar approach and extract 128-dimensional S2DNet [8] features, leading to a deep surface texture. Texture mapping is a function that denotes the correspondences between pixels in the texture map and mesh faces, and it is fixed throughout the optimization.

Object motion representation. The object motion \mathcal{N} is a temporal sequence of 6DoF poses. The 6DoF pose $\mathcal{N}_n = (T_n, Q_n)$ in each frame n is represented by 3D translation and 3D rotation. Translations are represented by 3D vectors T_n and rotations by normalized 4D quaternions Q_n .

3.2. Loss terms

Appearance and silhouette rendering. Let us define a rendering function \mathcal{R} that takes three inputs, object shape \mathcal{P} , object appearance \mathcal{T} , and 6DoF pose (T_n, Q_n) , and ren-




Figure 3. **Evaluation on a vacuum cleaner sequence.** This example highlights the need for more robust features like S2DNet [8] (middle row) that are robust to various illumination changes. RGB color values vary significantly within the input sequence (top row), and directly optimizing RGB texture leads to failures. In contrast, we back-project RGB features (bottom row) after optimization has finished. Long-term correspondences are marked with colored points and their visibility with colored lines. Texture visualization is a UV map [6].

ders the object at the given pose. This function provides two outputs, namely object appearance $\mathcal{R}_F(\mathcal{P}, \mathcal{T}, \mathcal{N}_n)$ and object silhouette $\mathcal{R}_S(\mathcal{P}, \mathcal{T}, \mathcal{N}_n)$ (subscript F for ‘Features’, and S for ‘Silhouette’). We assume that the camera is static. The camera projection matrix is fixed to 45° field-of-view angle. Given the rendered appearance and silhouette, we define below loss terms for every frame n .

Appearance rendering loss. This loss measures how well the estimated 3D shape and texture re-renders the input video frames. This re-rendering is the main driving force of the proposed approach. Concretely, we compare the object appearance features rendered at the given pose to the observed video features $F(I_i)$. We use the Cauchy loss γ with a scale of 0.25 as in [26]:

$$\mathcal{L}_F = \mu_n \sum_{i \in \{K_n \cup n\}} \left\| \mathcal{R}_S(\mathcal{P}, \mathcal{T}, \mathcal{N}_i) \cdot \left(\mathcal{R}_F(\mathcal{P}, \mathcal{T}, \mathcal{N}_i) - F(I_i) \right) \right\|_\gamma \quad (1)$$

Here, the summation runs over the current set of frames K_n and the current frame n . The factor $\mu_n = (|K_n| + 1)^{-1}$ normalizes the loss.

Silhouette rendering loss. This loss is based on the initial segmentations of the 2D tracker we start from. This loss keeps the method anchored to the correct object since there might be multiple objects in the scene. We define it as a

sum of two terms:

$$\mathcal{L}_S = \mu_n \sum_{i \in \{K_n \cup n\}} \left(\left(1 - \text{IoU}(M_i, \mathcal{R}_S(\mathcal{P}, \mathcal{T}, \mathcal{N}_i)) \right) + \|\text{DT}(M_i) \cdot \mathcal{R}_S(\mathcal{P}, \mathcal{T}, \mathcal{N}_i)\| \right). \quad (2)$$

The first term is the intersection-over-union (IoU) between the rendered silhouette of the estimated 3D shape and the initial segmentation M_i . When the rendered silhouette and the initial segmentation are near, this loss behaves well. However, when they are far away, this term is zero, leading to zero gradient flow during optimization (Sec. 3.3). To address this case, the second term in (2) attracts the object silhouette to move closer to the initial segmentation. It measures the average distance of each pixel of the silhouette to the closest point in the initial segmentation M_i . It is efficiently implemented by computing a distance transform (DT) image of the initial segmentation and multiplying it by the rendered silhouette.

Motion regularization. To encourage motion smoothness, we penalize the difference between the translation T_n in the current frame and the latest previous keyframe. Similarly, we penalize the angular difference in rotations Q_n :

$$\mathcal{L}_M = \max \left(0, \frac{1}{K_n[-1] - n} \|T_{K_n[-1]} - T_n\|_2 - \nu_T \right) + \max \left(0, \frac{1}{K_n[-1] - n} \angle(Q_{K_n[-1]}, Q_n) - \nu_Q \right), \quad (3)$$

where $K_n[-1]$ denotes the index of the most recent keyframe before the current frame. We do not penalize differences of rotations and translations to the last keyframe smaller than thresholds $\nu_T = 0.1$ and $\nu_T = 30^\circ$, as we consider these motions to be in a normal range. Effectively, the loss penalizes drastic pose changes w.r.t. the last keyframe.

Shape regularization. To promote 3D shape smoothness, we use the Laplacian loss \mathcal{L}_L as defined in [28].

Texture regularization. For a smooth texture map, we use a total variation loss \mathcal{L}_T that computes the average appearance difference between neighboring pixels in the texture map (using \mathcal{L}_F to measure the difference).

Total loss. The total loss is a weighted sum of all the above loss terms: $\mathcal{L} = \alpha_F \mathcal{L}_F + \alpha_S \mathcal{L}_S + \alpha_M \mathcal{L}_M + \alpha_L \mathcal{L}_L + \alpha_T \mathcal{L}_T$. In our experiments, we set $\alpha_F = 1$, $\alpha_S = 1$, $\alpha_M = 1$, $\alpha_L = 1000$, and $\alpha_T = 0.001$.

3.3. Optimization

To minimize the total loss \mathcal{L} in each frame, we use the ADAM [18] optimizer with a learning rate of 0.1. The optimization is stopped early when the appearance rendering loss reaches $\mathcal{L}_F < \tau_F$, where τ_F is the convergence threshold. For RGB features we use $\tau_F = 0.2$, while for S2DNet [8] features we use $\tau_F = 0.05$. If this stopping criterion is not reached within 500 iterations, we declare the current frame a failure. This automatic self-assessment declares that the initial segmentation cannot be improved by the proposed method. Failure can happen due to various reasons, *e.g.* bad initial segmentation, difficult object motion, non-rigid object deformation (we assume rigid object), extreme lighting or texture changes on the object itself in the real world (we assume constant texture throughout the sequence). The self-assessment is especially helpful when the method is evaluated on large datasets, where a rigid object in the scene cannot be guaranteed. When a failure is detected, our method automatically degenerates to returning the initial segmentation for that frame. This allows our method to operate in more general conditions, improving over the initial tracker when it can and simply returning its output when it cannot. In our experiments, self-assessment detects a failure in 1-5% of frames depending on the dataset.

Keyframe selection criteria. Since the optimization is anchored to the subset of keyframes, they have to be carefully selected. The subset of keyframes is dynamic, as it changes over time. Selection criteria are re-evaluated after the optimization has finished for each frame. Thus, keyframes can fall out of the subset when one of the selection criteria do not hold anymore. We have four selection criteria for the current frame to become a keyframe:

- (1) The rendered object silhouette aligns well with the 2D segmentations, *i.e.* $\mathcal{L}_S < 0.3$.
- (2) The rendered object appearance is consistent with the

Method	IoU	Δ
CSR-DCF [34]	0.498	-
CSR-DCF [34] + Ours	0.512	2.8 %
D3S [33]	0.732	-
D3S [33] + Ours	0.765	4.5 %
OTrack [49]	0.758	-
OTrack [49] + Ours	0.788	4.0 %

Table 1. **Segmentation accuracy on the CDTB dataset [32].** Our method improves initial segmentations from trackers in all cases.

input frame, *i.e.* $\mathcal{L}_F < \tau_F$.

- (3) The current frame and the latest keyframe provide significantly different views of the object, *i.e.* relative translation is greater than half the object size or relative rotation is greater than 45° .
- (4) Finally, the number of keyframes is restricted to a maximum α_K frames; thus, only the most recent α_K are kept (discussed in Sec. 4.4).

Differentiable rendering. As part of the forward model, we render object appearance and silhouette. To make this step differentiable, we apply differentiable rendering proposed in [3].

Initialization. The object 3D shape is initialized to a prototype mesh, which is a sphere with 1212 vertices. Vertex offsets are initialized to zero. The texture map has dimension 300×300 and is initialized to zero values. The 6DoF pose in the first frame is initialized to 2D location in the center of the image, depth such that the object covers the whole image, and zero 3D rotation. Every next 6DoF pose is initialized to the pose estimated in the previous frame.

Runtime. Currently, the method runs at 2 sec. per frame on average. However, the runtime can be improved by simpler shape representations, better motion prediction, and hierarchical optimization. Moreover, the method is highly parallelizable, *e.g.* all keyframes can be rendered in parallel.

4. Evaluation

The first step of the proposed method is running a state-of-the-art 2D tracker to obtain initial segmentation masks (by default OTrack [49] with AlphaRefine [48], the winner of VOT-ST 2022 challenge [19]). Hence, we evaluate the improvement in 2D segmentation brought by our method by rendering the silhouette of the reconstructed object on the video frames.

Our method assumes the object is rigid. Most tracking datasets, *e.g.* [22, 15], contain non-rigid objects such as humans and animals. Therefore, we selected datasets with mostly rigid objects. First, we evaluate on the large CDTB [32] dataset that contains mostly rigid objects (part of the VOT dataset collection [22]). Second, we evaluate



Figure 4. **Results with slight non-rigid deformations.** Our method successfully reconstructed and improved the segmentation accuracy (by 11.1% IoU on average) on a human body sequence from the CDTB dataset [32].

Method	Input	IoU
OSVOS [1]	RGB	87.97 %
OTrack [49]	RGB	93.26 %
Ours (OTrack init)	RGB	94.05 %
BackFlow [45]	RGB-D	93.52 %

Table 2. **In-hand object manipulation dataset [45].** Our method achieves the highest segmentation accuracy.

on two smaller datasets, the Coin Tracking dataset [43] and the in-hand object manipulation dataset [45], both of which contain only rigid objects.

4.1. CDTB dataset

The “Color-and-Depth general visual object Tracking Benchmark” (CDTB) [32] is part of the VOT-RGBD challenge suite [22]. It is the only standard tracking dataset that contains mostly rigid objects. Since we do not require depth information, we drop the depth channel. The dataset contains 80 sequences with an average of 1274 frames each.

To test how much the initial 2D tracker influences the final results of our method, we use three modern 2D trackers: D3S [33], CSR-DCF [34], and OTrack [49] (the most recent state-of-the-art from ECCV 2022). When the tracker output is a bounding box (as most of the trackers, *e.g.* CSR-DCF, OTrack), we apply AlphaRefine [48] to extract foreground segmentations (as suggested by OTrack [49]). As in standard tracking tasks, we are given the ground truth segmentation (or bounding box) in the first frame of the video. Then, the 2D tracker runs for the whole sequence without any re-initialization, producing a segmentation in each frame, which we feed as initialization to our method.

The proposed method processes every sequence in an on-

Method	IoU	Δ
Coin tracking [43]	0.704	-
Ours (RGB feat.)	0.718	2.0 %
Ours (RGB feat., flat prior)	0.731	3.8 %
Ours (S2DNet [8] feat.)	0.740	5.1 %
Ours (S2DNet [8] feat., flat prior)	0.764	8.5 %

Table 3. **Coin tracking dataset [43].** We achieve the highest segmentation accuracy with S2DNet [8] features and flat-object prior.

line manner. For evaluation, we compute the Intersection-over-Union (IoU) between the ground truth segmentation and our output (Sec. 3.3).

As Table 1 shows, our approach improves segmentation performance when built on top of any of the evaluated trackers. Average improvement is above 4% for recent methods [33, 49]. For an older discriminative correlation filter based tracker with bounding box output, the improvement is only 2.8%, probably because the initial segmentations are poor. Many objects in this dataset are close to rigid but not perfectly rigid. Nevertheless, our method still successfully reconstructs such objects in motion, *e.g.* a human body in Fig. 4 and a backpack in Fig. 7 (top).

4.2. Coin Tracking dataset

The Coin Tracking dataset [43] was proposed to address a new problem, namely tracking of both sides of flat, potentially fast-rotating rigid objects as cards, books, and coins. The dataset creators also introduced a method to tackle the flat-object tracking problem. Therefore, we also integrate a flat-object prior to our method by projecting all vertices to a single plane after each iteration of optimization.

When using simple RGB values as appearance features,

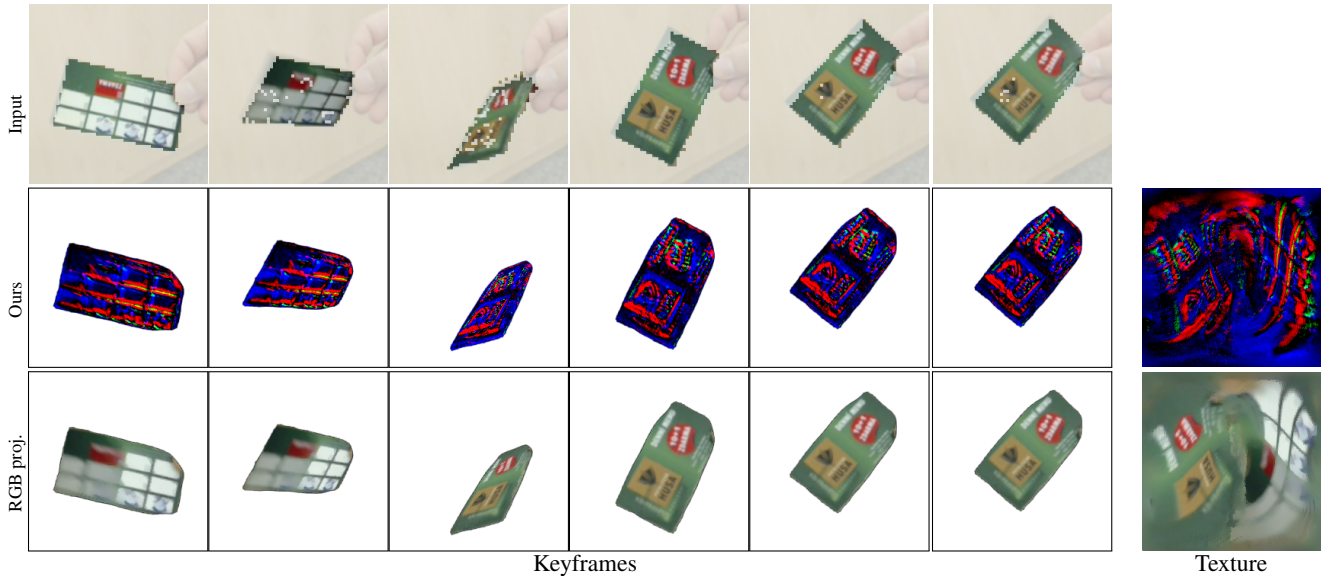


Figure 5. **Coin tracking dataset.** The proposed method successfully tracks with 6DoF and estimates 3D shape of flat objects. Since the object in motion has different lighting effects, e.g. first and second columns, the RGB back-projection contains noisy estimates.

Ablated version	IoU	Δ
OTrack [49]	0.703	-
Ours (RGB feat.)	0.717	+2.0 %
Ours (S2DNet [8] feat.)	0.740	+5.3 %
No \mathcal{L}_F (1)	0.727	+3.4 %
No \mathcal{L}_M (3)	0.689	-2.0 %
No \mathcal{L}_L (Sec. 3.2)	0.692	-1.6 %
No \mathcal{L}_T (Sec. 3.2)	0.734	+4.4 %

Table 4. **Ablation study.** Each row in the bottom part represents segmentation accuracy change w.r.t. OTrack [49] when each of the introduced loss terms (in Sec. 3.2) is disabled.

the segmentation accuracy improves by 2%, and by 3.8% with the flat-object prior (Table 3). In contrast, when using S2DNet [8] features, the improvement is 5.1%, and 8.5% with the flat-object prior. The final reconstruction with the flat-object prior is shown in Fig. 5 and Fig. 7 (third and fourth rows).

4.3. In-hand object manipulation dataset

This dataset proposed by [45] contains 13 sequences of in-hand manipulation of objects from the YCB dataset [2]. Each sequence ranges from 300 to 700 frames in length. The ground truth is provided as manually labeled segmentation masks for every tenth frame. We again evaluate accuracy of our reprojected silhouettes (with OTrack initial segmentations and S2DNet features). As shown in Table 2, our method brings a small improvement from 93.26% to 94.05%. We believe that the performance on this dataset is already saturated and larger improvements are hardly achievable due to inaccurate ground truth annotations. Final




Figure 6. **Influence of the number of keyframes.** When the number of keyframes is increased, the average IoU increases as well thanks to more constraints from different views.

reconstructions are shown in Fig. 7 (bottom row).

4.4. Ablation study

We ablate the main components of the proposed method on 16 random sequences from the CDTB dataset [32]. As shown in Table 4, disabling any part of the total loss function \mathcal{L} (Sec. 3.2) decreases performance, indicating that all components of our method contribute to the overall accuracy. The silhouette rendering loss (2) cannot be disabled since in that case, the method does not know where the object of interest is and never converges. Another interesting observation is that using RGB features ('RGB feat.' row) is worse than completely discarding the feature rendering

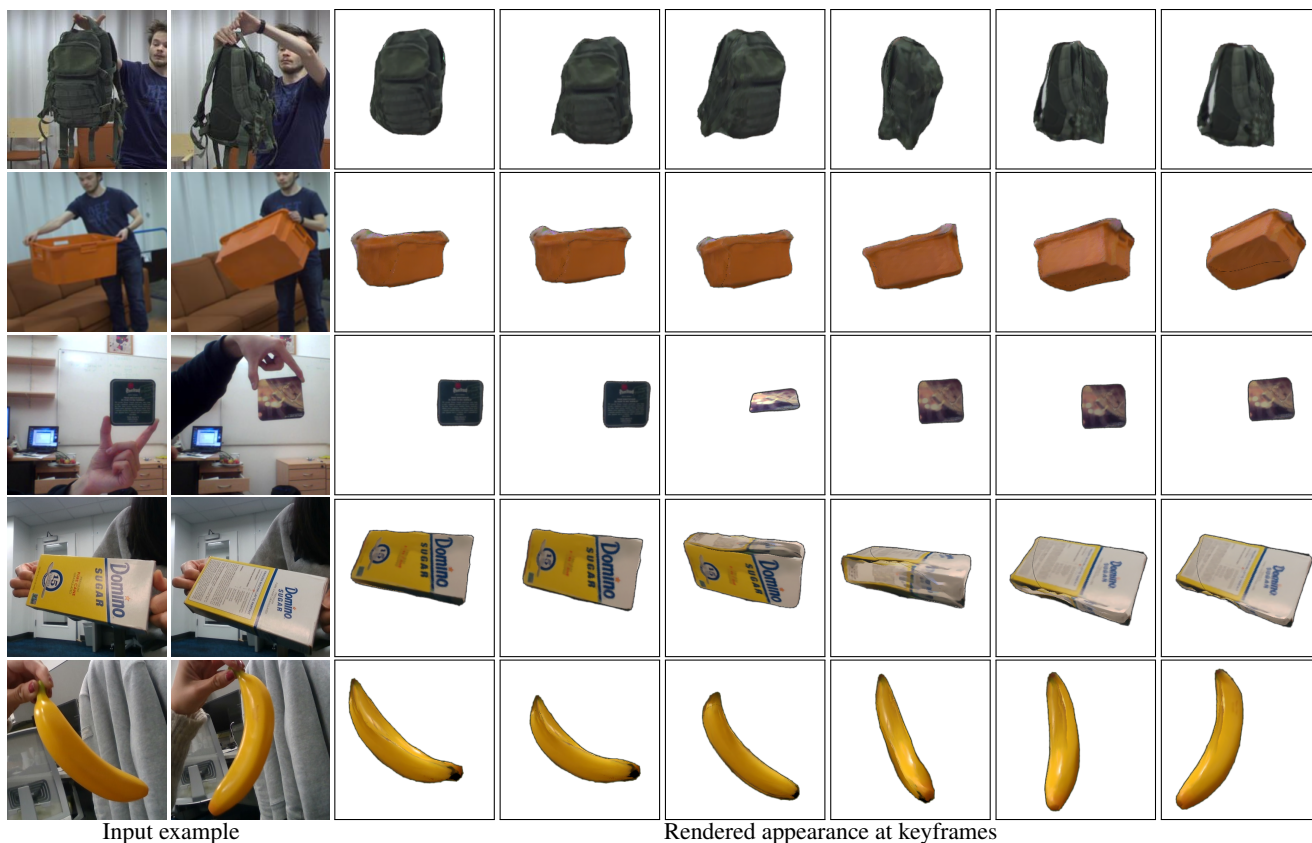


Figure 7. **Final 3D textured shape reconstructions.** We visualize RGB back-projections (after optimizing with S2DNet [8] features) on a backpack and container sequences from the CDTB dataset [32] (top two rows), a beer mat sequence from the coin tracking dataset [43] (third and fourth rows), and sugar and banana sequences from [45] (bottom row).

loss ('No \mathcal{L}_F row'; +2% vs +3.4% improvement) since the lighting and illumination effects make the RGB features rather unreliable. For instance, this can be observed in a vacuum cleaner example where the direct light does not move with the object (Fig. 3, top row), and in a card tracking example where shadows change colors drastically in the first and second columns (Fig. 5). The influence of the number of keyframes α_K is shown in Fig. 6. As expected, the more keyframes are used, the better are results, but it comes at the cost of the increased runtime. In all experiments, we used $\alpha_K = 6$, which is a reasonable trade-off value.

4.5. Benchmark for 6DoF Object Pose

As a by-product of the proposed tracker, we additionally produce an approximate 3D shape and 6DoF pose of the object. To evaluate these estimates, we use the TUD-L dataset from the Benchmark for 6D Object Pose (BOP) [14]. We chose this dataset because it contains videos of one moving object with limited occlusions, matching the common setting in standard 2D tracking datasets, which we target. As shown in Table 5, our method produces results competitive with several recent methods, *while being the only one that*

does not need the 3D object model as input, neither at the training nor the test time. Instead, we *estimate it* along with the 6DoF pose. Hence, this is a strong result. Moreover, we do not train on the TUD-L dataset at all: our optimization-based method works out of the box, showing strong generalization across datasets. In fact, all other compared methods are trained on large synthetic and real datasets and use various refinement steps (ICP).

4.6. Limitations

Runtime. Currently, our method runs on average at 5 seconds per video frame. This prevents it from being used in real-time tracking applications. However, the runtime can be improved by simpler shape representations, better motion prediction, and hierarchical optimization. Moreover, the method is highly parallelizable, *e.g.* all keyframes can be rendered in parallel.

Fixed topology. The topology of the object of interest is one of the hyperparameters of our method. While the spherical prototype we currently use is sufficient for many objects, in future work, we could have multiple prototypes as done by [39] in a different domain.

Method	AR \uparrow	given 3D model	requires training
CosyPose-ECCV20-Synt+Real	82.3	yes	yes
CDPNv2 BOP20-RGB	77.2	yes	yes
CDPN BOP19-RGB	76.9	yes	yes
Zhigang-CDPN-ICCV19	75.7	yes	yes
Leaping from 2D to 6D-ECCVW20	75.1	yes	yes
CosyPose-ECCV20-PBR	68.5	yes	yes
CDPNv2 BOP20-PBR	58.8	yes	yes
EPOS-BOP20-PBR	55.8	yes	yes
Ours	54.5	no	no
Pix2Pose-BOP20-ICCV19	42.0	yes	yes
Sundermeyer-IJCV19	40.1	yes	yes
Pix2Pose-BOP19-ICCV19	34.9	yes	yes
SingleMultiPathEncoder-CVPR20	33.4	yes	yes
DPOD (synthetic)	24.2	yes	yes

Table 5. **Evaluation on TUD-L dataset from the BOP benchmark [14].** We compare to methods that operate on RGB images at test time. Only our method does not require a given 3D model and requires no training. AR: average recall as in [14].

Camera motion. While we assumed a static camera, our method does not strictly require it. If the camera is moving, our method still works, but the estimated “object pose” combines the camera and the object pose. They could potentially be disentangled in a post-processing stage by estimating camera motion, *e.g.* with SfM, and then subtracting it from the estimated object pose.

5. Conclusion

We proposed a novel model-free tracking method that jointly estimates 3D shape, texture, and 6DoF pose of unknown rigid objects in videos, as opposed to just a 2D segmentation, which is commonly used in tracking. As a result, the proposed tracker provides dense long-term correspondences on the object surface. The core of our method uses differentiable rendering, which is used in visual object tracking for the first time. The experiments demonstrated improved tracking accuracy on several tracking benchmarks. The proposed method produces an object representation that supports diverse applications, such as augmented reality, object manipulation, and grasping.

Acknowledgements. This research was supported by a Google Focused Research Award, Innosuisse grant No. 34475.1 IP-ICT, Research Center for Informatics (project CZ.02.1.01/0.0/0.0/16.019/0000765 funded by OP VVV), and a research grant by FIFA.

References

- [1] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *CVPR*, 2017. 6
- [2] Berk Calli, Arjun Singh, James Bruce, Aaron Walsman, Kurt Konolige, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M

- Dollar. Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research*, 36(3):261–268, 2017. 7
- [3] Wenzheng Chen, Jun Gao, Huan Ling, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *NeurIPS*, 2019. 1, 2, 5
- [4] Wenzheng Chen, Joey Litalien, Jun Gao, Zian Wang, Clement Fuji Tsang, Sameh Khalis, Or Litany, and Sanja Fidler. DIB-R++: Learning to predict lighting and material with a hybrid differentiable renderer. In *NeurIPS*, 2021. 2
- [5] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 2
- [6] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. 4
- [7] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ATOM: accurate tracking by overlap maximization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4660–4669. Computer Vision Foundation / IEEE, 2019. 2
- [8] Hugo Germain, Guillaume Bourmaud, and Vincent Lepetit. S2dnet: Learning image features for accurate sparse-to-dense matching. In *ECCV*, 2020. 1, 3, 4, 5, 6, 7, 8
- [9] Mariia Gladkova, Nikita Korobov, Nikolaus Demmel, Aljoša Ošep, Laura Leal-Taixé, and Daniel Cremers. Directtracker: 3d multi-object tracking using direct image alignment and photometric bundle adjustment. *IROS*, 2022. 2
- [10] Shreyas Hampali, Tomas Hodan, Luan Tran, Lingni Ma, Cem Keskin, and Vincent Lepetit. In-hand 3d object scanning from an rgb sequence. In *CVPR*, pages 17079–17088, June 2023. 2
- [11] João F. Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):583–596, 2015. 2
- [12] Tomáš Hodaň, Dániel Baráth, and Jiří Matas. EPOS: Estimating 6D pose of objects with symmetries. *CVPR*, 2020. 2
- [13] Tomáš Hodaň, Frank Michel, Eric Brachmann, Wadim Kehl, Anders Glent Buch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, Caner Sahin, Fabian Manhardt, Federico Tombari, Tae-Kyun Kim, Jiří Matas, and Carsten Rother. BOP: Benchmark for 6D object pose estimation. *ECCV*, 2018. 2
- [14] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. BOP challenge 2020 on 6D object localization. *ECCVW*, 2020. 2, 8, 9
- [15] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE TPAMI*, 43(5):1562–1577, 2021. 5

- [16] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 2
- [17] Ugur Kart, Alan Lukezic, Matej Kristan, Joni-Kristian Kamarainen, and Jiri Matas. Object tracking by reconstruction with view-specific discriminative correlation filters. In *CVPR*, June 2019. 2
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 5
- [19] Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kamarainen, Hyung Jin Chang, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, Johanna Bjorklund, Yushan Zhang, Zhongqun Zhang, Song Yan, Wenyan Yang, Dingding Cai, Christoph Mayer, and Gustavo Fernandez. The tenth visual object tracking vot2022 challenge results, 2022. 5
- [20] Matej Kristan, Ales Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kamarainen, Luka Čehovin Zajc, Martin Danelljan, Alan Lukezic, Ondrej Drbohlav, Linbo He, Yushan Zhang, Song Yan, Jinyu Yang, Gustavo Fernandez, and et al. The eighth visual object tracking vot2020 challenge results, 2020. 1
- [21] Matej Kristan, Jiří Matas, Aleš Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kamarainen, Hyung Jin Chang, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, Jani Kapyla, Gustav Hager, Song Yan, Jinyu Yang, Zhongqun Zhang, Gustavo Fernandez, and et. al. The ninth visual object tracking vot2021 challenge results, 2021. 1
- [22] Matej Kristan, Jiří Matas, Aleš Leonardis, Tomas Vojiř, Roman Pflugfelder, Gustavo Fernandez, Georg Nebehay, Fatih Porikli, and Luka Čehovin. A novel performance evaluation methodology for single-target trackers. *IEEE TPAMI*, 38(11):2137–2155, Nov 2016. 1, 5, 6
- [23] Y. Labbe, J. Carpentier, M. Aubry, and J. Sivic. Cosypose: Consistent multi-view multi-object 6d pose estimation. In *ECCV*, 2020. 2
- [24] Felix Leeb, Arunkumar Byravan, and Dieter Fox. Motionnets: 6d tracking of unknown objects in unseen environments using rgb. *IROS Workshop on The Importance of Uncertainty in Deep Learning for Robotics*, 2019. 2
- [25] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *ICCV*, 2021. 2
- [26] Philipp Lindenberger, Paul-Edouard Sarlin, Viktor Larsson, and Marc Pollefeys. Pixel-Perfect Structure-from-Motion with Featuremetric Refinement. In *ICCV*, 2021. 1, 3, 4
- [27] Lahav Lipson, Zachary Teed, Ankit Goyal, and Jia Deng. Coupled iterative refinement for 6d multi-object pose estimation. In *CVPR*, 2022. 2
- [28] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. *ICCV*, Oct 2019. 1, 2, 5
- [29] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In *CVPR*, 2020. 2
- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 2
- [31] Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'81*, page 674–679, San Francisco, CA, USA, 1981. Morgan Kaufmann Publishers Inc. 2
- [32] Alan Lukezic, Ugur Kart, Jani Käpylä, Ahmed Durmush, Joni-Kristian Kamarainen, Jiří Matas, and Matej Kristan. Cdtb: A color and depth visual object tracking dataset and benchmark. In *ICCV*, pages 10012–10021, 2019. 5, 6, 7, 8
- [33] Alan Lukezic, Jiří Matas, and Matej Kristan. D3s - a discriminative single shot segmentation tracker. In *CVPR*, 2020. 1, 2, 5, 6
- [34] Alan Lukezic, Tomas Vojíš, Luka Čehovin Zajc, Jiří Matas, and Matej Kristan. Discriminative correlation filter with channel and spatial reliability. In *CVPR*, July 2017. 1, 2, 5, 6
- [35] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2
- [36] Van Nguyen Nguyen, Yuming Du, Yang Xiao, Michael Ramamonjisoa, and Vincent Lepetit. Pizza: A powerful image-only zero-shot zero-cad approach to 6 dof tracking. 2022. 2
- [37] Alessandro Pieropan, Niklas Bergström, Masatoshi Ishikawa, and Hedvig Kjellström. Robust 3d tracking of unknown objects. In *ICRA*, pages 2410–2417, 2015. 2
- [38] Konstantinos Rematas, Ricardo Martin-Brualla, and Vittorio Ferrari. Sharf: Shape-conditioned radiance fields from a single view. In *ICML*, 2021. 2
- [39] Denys Rozumnyi, Martin R. Oswald, Vittorio Ferrari, and Marc Pollefeys. Shape from blur: Recovering textured 3d shape and motion of fast moving objects. In *NeurIPS*, 2021. 1, 2, 8
- [40] Denys Rozumnyi, Martin R. Oswald, Vittorio Ferrari, and Marc Pollefeys. Motion-from-blur: 3d shape and motion estimation of motion-blurred objects in videos. In *CVPR*, Jun 2022. 1, 2
- [41] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [42] Jianbo Shi and Tomasi. Good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, 1994. 2
- [43] Jonáš Šerých and Jiří Matas. Visual coin-tracking: Tracking of planar double-sided objects. In *GCPR*, page 317–330, Berlin, Heidelberg, 2019. Springer-Verlag. 6, 8
- [44] Dan Wang, Xinrui Cui, Xun Chen, Zhengxia Zou, Tianyang Shi, Septimiu Salcudean, Z. Jane Wang, and Rabab Ward. Multi-view 3d reconstruction with transformers. In *ICCV*, pages 5722–5731, October 2021. 2

- [45] Fan Wang and Kris Hauser. In-hand object scanning via rgb-d video segmentation. *ICRA*, 2019. [2](#), [6](#), [7](#), [8](#)
- [46] Gu Wang, Fabian Manhardt, Jianzhun Shao, Xiangyang Ji, Nassir Navab, and Federico Tombari. Self6d: Self-supervised monocular 6d object pose estimation. In *The European Conference on Computer Vision (ECCV)*, August 2020. [2](#)
- [47] B Wen and Kostas E Bekris. Bundletrack: 6d pose tracking for novel objects without instance or category-level 3d models. In *IROS*, 2021. [2](#)
- [48] Bin Yan, Xinyu Zhang, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Alpha-refine: Boosting tracking performance by precise bounding box estimation. In *CVPR*, pages 5285–5294, 2021. [2](#), [5](#), [6](#)
- [49] Botao Ye, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Joint feature learning and relation modeling for tracking: A one-stream framework. In *ECCV*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [50] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. iNeRF: Inverting neural radiance fields for pose estimation. *IROS*, 2020. [2](#)
- [51] Yang You, Ruoxi Shi, Weiming Wang, and Cewu Lu. Cppf: Towards robust category-level 9d pose estimation in the wild. In *CVPR*, 2022. [2](#)
- [52] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. [2](#)