

Articulated Object Estimation in the Wild

Abdelrhman Werby^{1*}, Martin Büchner^{1*}, Adrian Röfer^{1*}, Chenguang Huang²,
Wolfram Burgard², Abhinav Valada¹

¹University of Freiburg

²University of Technology Nuremberg

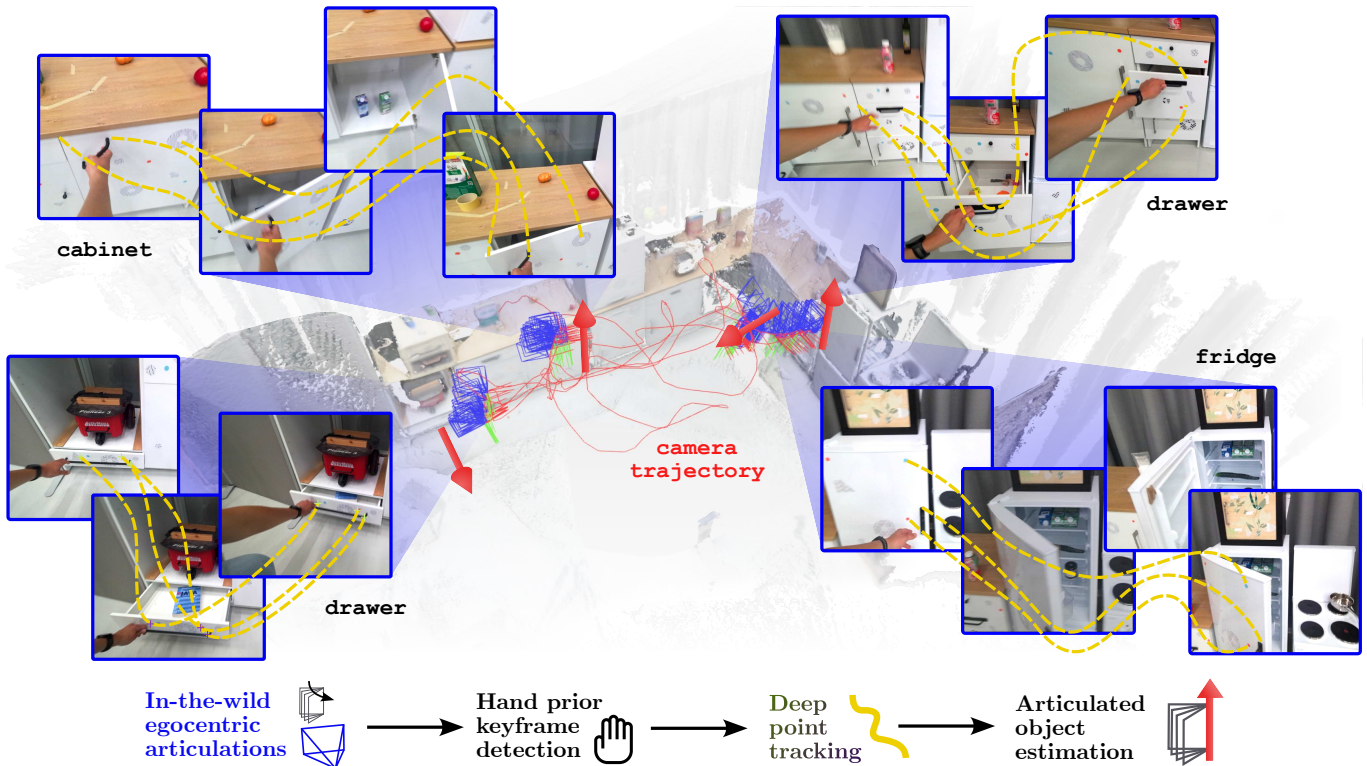


Fig. 1: We present *ArtiPoint*, our novel approach to articulated object estimation in the wild. *ArtiPoint* makes use of deep point tracking and factor graph optimization. It does not rely on costly scene-wise optimization or deep models that are prone to overfit. We evaluate our approach on the *Arti4D* dataset, a novel scene-level articulated object dataset that is recorded with a moving camera, compared to previous benchmarks that are limited to isolated objects and static camera poses.

Abstract—Understanding the 3D motion of articulated objects is essential in robotic scene understanding, mobile manipulation, and motion planning. Prior methods for articulation estimation have primarily focused on controlled settings, assuming either fixed camera viewpoints or direct observations of various object states, which tend to fail in more realistic unconstrained environments. In contrast, humans effortlessly infer articulation by watching others manipulate objects. Inspired by this, we introduce *ArtiPoint*, a novel estimation framework that can infer articulated object models under dynamic camera motion and partial observability. By combining deep point tracking with a factor graph optimization framework, *ArtiPoint* robustly estimates articulated part trajectories and articulation axes directly from raw RGB-D videos. To foster future research in this domain, we introduce *Arti4D*, the first ego-centric in-the-wild dataset that captures articulated

object interactions at a scene level, accompanied by articulation labels and ground-truth camera poses. We benchmark *ArtiPoint* against a range of classical and modern deep learning baselines, demonstrating its superior performance on *Arti4D*. We make code and *Arti4D* publicly available at <https://artipoint.cs.uni-freiburg.de>.

I. INTRODUCTION

Robotic manipulation is concerned with moving external objects to achieve a task-relevant goal. While recent approaches have made tremendous progress in manipulating unconstrained objects [25, 6], manipulating constrained *articulated* objects remains challenging in model-free settings that rely on affordances [13]. On the other hand, model-based approaches have thoroughly investigated manipulating articulated objects, demonstrating both manipulation success and transferability to novel objects [35, 43, 16, 4]. However, the aforementioned

* Equal contribution.

This work was funded by the BrainLinks-BrainTools Center of the University of Freiburg and an academic grant from NVIDIA.

approaches focus on offline learning of predictive models in curated settings, where only the articulated object itself moves [27, 28]. If we consider deploying a robot in a novel environment, these requirements are not met [13]. As a consequence, most approaches for reconstructing articulated objects and estimating motion models fail when operating *in-the-wild*, i.e., when faced with dynamic camera poses, occlusions, and significant clutter, where objects are not isolated and interactions are less constrained. While humans acquire manipulation skills through observing others interacting with various entities such as articulated objects [1], endowing robotic systems with those capabilities remains an open problem. Ultimately, this limits the transfer of insights from articulated object estimation [35, 4] to the problem of human-to-robot imitation in everyday robotics.

In this work, we propose to address this limitation by exploiting interaction priors. As humans manipulate environments with their hands, we aim to observe those interactions. First, we extract interaction segments containing object interactions. Secondly, we obtain point trajectories within regions close to the detected hand, which describe both the static and dynamic parts of the observed scene. We leverage recent state-of-the-art models in any-point tracking [11, 23], which yields point trajectories throughout whole interaction segments, including an estimated visibility score. Third, we lift those point trajectories into 3D using depth measurements and compensate for camera motion using accurate camera odometry. We separate all static 3D trajectories with a negligible length from dynamic trajectories, which potentially represent the motion of dynamic objects. This enables us to estimate the motion of the moving part triggered through hand interaction. Finally, given the acquired point trajectories, we estimate the underlying articulation model, classify the joint type (prismatic or revolute), and globally register the model in the observed scene.

As we address the novel task of estimating articulated object motion *in-the-wild*, we introduce Arti4D, the first egocentric in-the-wild dataset for scene-level articulated object manipulation from human demonstrations. It comprises 45 RGB-D sequences across four diverse scenes, capturing 411 human-object interactions under dynamic camera motion, occlusions, and uncurated environments ranging from kitchen scenes to robot lab rooms. The dataset was recorded by a human operator holding an RGB-D camera while actively exploring scenes and interacting with articulated objects. Unlike prior datasets, Arti4D includes $\mathfrak{se}(3)$ -labeled ground-truth articulation axes, difficulty annotations, reconstructed scenes, and accurate camera pose ground truth. It offers a challenging benchmark not only for articulated object estimation but also for visual odometry and SLAM research in complex, real-world settings.

Concretely, we make the following contributions:

- 1) We present a novel articulated object estimation framework called *ArtiPoint* that operates on deep point trajectories, enabling robust deployment in uncontrolled settings.
- 2) We introduce *Arti4D*, the first real-world *in-the-wild* dataset of articulated object interactions providing human

exploration and articulation demonstrations on a scene-level. The dataset includes articulation axes labels and 3D ground truth camera poses.

- 3) We compare our method against a set of classical and deep articulation model estimation pipelines and present extensive ablations.
- 4) We publish code, data, and model predictions of our method at <https://artipoint.cs.uni-freiburg.de>.

II. RELATED WORK

In the following, we review previous work in articulated object estimation and review recent developments in any-point tracking.

Probabilistic Methods: Early work in articulation model estimation often employed probabilistic formulations to infer the relationships among articulated object parts. Sturm *et al.* [37, 35] proposed a probabilistic framework that infers the articulation model from 6D trajectories of object parts. While the initial work used marker-based tracking to obtain these trajectories, subsequent work extended these frameworks to handle sparse (markerless) objects by extracting visual features, such as SURF [3] or by detecting rectangles in dense-depth images [36].

Deep-Learning-based Methods: Recent advancements in deep learning diversify the articulation estimation methods. By consuming a sequence of depth images with a fixed viewing pose, ScrewNet [19] enables category-agnostic articulation motion prediction in an end-to-end manner. A follow-up work, DUST-net [20] further estimates the uncertainty of the motion to introduce interpretability. While these works take a depth image sequence as input, ANCSH [26] predicts articulation joint parameters and state for unseen instances using a depth image, assuming prior knowledge of the instance category. Other works utilize neural networks as a part of their pipelines for articulation motion prediction. While Heppert *et al.* [15] formulate a two-stage pipeline consisting of a learning-based part tracking stage and a factor graph optimization stage, Buchanan *et al.* [4] combine the learning-based method with interactive perception, leveraging a neural network to propose an initial estimation of the articulation motion and interactive perception for prediction refinement. FormNet [40] predicted the motion field residual and the part connectedness before a post-processing step for the final prediction.

3D Reconstruction of Articulated Objects: Reconstructing articulated objects in 3D is more challenging than reconstructing rigid objects, as it requires modeling the articulation between their constituent parts. Early work from Huang *et al.* [18] leverages structure from motion to reconstruct the articulated object at different configurations, followed by post-processing techniques to estimate joint axes and types. More recent works inspired by Gaussian splatting offer a smooth and continuous representation for complex shapes. For instance, ArtGS [28] utilizes 3D Gaussians to jointly optimize for canonical Gaussians and the articulation model. Ditto [21] is a neural network-based method for estimating the articulation model and the 3D geometry of an articulated object, given point

cloud observations before and after interaction. PARIS [27] takes multiview images at different articulation states and jointly estimates the articulation parameters and the 3D geometry of the object.

Tracking Any Point (TAP): Tracking points is a key problem in various computer vision and robotics tasks, such as 3D reconstruction, video analysis, and object tracking. Early works relied on sparse feature matching, such as SIFT [29], SURF [3], and GFTT [33], as well as optical flow. Tracking any point as a task was introduced in PIP [14], where the goal was to track pixel locations across entire sequences. TAP-Vid [10] enhances the problem formulation by providing a benchmark and introducing TAP-Net. Subsequent works [11, 44, 7, 12] improved performance while Karaev *et al.* [23, 22] introduce joint tracking of points to account for their dependencies.

In this work, we address the novel problem of estimating articulations from mobile observations *in-the-wild*. We connect deep any-point tracking [22] to the estimation approach of Buchanan *et al.* [4] to form a holistic approach to estimating the parameters of multiple articulated objects in large environments.

III. APPROACH

The goal of *ArtiPoint* is to estimate the underlying motion models of articulated objects given posed RGB-D observations of human interactions with such objects. To support this goal, we exploit as a prior that humans typically use their hands to interact with objects. Our approach consists of four stages, which we visualize in Fig. 1: First, we extract keyframe segments that contain hand-object interactions. We then sample 2D tracking points near the hand and track them throughout the segment to capture the object in motion. To obtain 3D motion, we lift the 2D tracks to 3D, compensate for camera motion, filter out static points in 3D, and apply trajectory smoothing in our third stage to mitigate noise. Finally, we input refined 3D tracks into a factor graph framework to jointly estimate the trajectory of the object parts and the parameters of the articulation motion model.

A. Extraction of Interaction Intervals

Given a scene level ego-centric RGB-D sequence $\{I_t, D_t\}_{t=1}^T$ of human operating and interacting with articulated object in the scene, we extract segments $S = \{(t_{start}^{(n)}, t_{end}^{(n)})\}_{n=1}^N$, every segment contains keyframes where hand-object interactions occurred. We use a preexisting hand segmentation model [5] to extract hand mask H_t from each RGB frame I_t , and assign a binary label $d_t \in \{0, 1\}$ to the frame based on the hand visibility. To avoid irregular false positives/negatives, we compute the moving average \bar{d}_t for the last w_h frames as $\bar{d}_t = \frac{1}{w_h} \sum_{i=t-w_h}^t d_i$. A keyframe is identified if the smoothed detection \bar{d}_t exceeds τ_h , initiating a segment s_n at time $t_{start} = t$. The segment continues as long as $\bar{d}_t \geq \tau_h$, and terminates at $t_{end} = t$ when \bar{d}_t drops below this threshold. Finally, the extracted segments S are filtered based on their duration, ensuring that $T_{min} \leq t_{end}^{(i)} - t_{start}^{(i)} \leq T_{max}$,

where T_{min} and T_{max} represent the minimum and maximum allowed segment lengths, respectively.

B. Deep Point Tracking

For each obtained segment s_n and hand masks $\{H_t\}_{t=t_{start}^{(n)}}^{t_{end}^{(n)}}$ associate with every keyframe within the segment, we uniformly sample points around the hand. We use the sampled points as prompts for the class-agnostic instance segmentation model [41] to identify masks for potential articulated objects near the hand within each keyframe. Given the resulting K objects' 2D masks $\{O_k^t\}_{k=1}^K$ in frame t , our goal is to select F query points $q_t \in \mathbb{R}^{F \times 2}$ that lie on the aforementioned masks and allow tracking throughout the entire segment. We achieve this by extracting GFTT keypoints via the lightweight Shi-Tomasi method [33] within each object mask in frame t . By concatenating all selected query points for a segment s_n , we obtain a set of query points $Q_i \in \mathbb{R}^{K \times F \times 2}$ where K represents the number of keyframes selected from the segment s_n . Subsequently, we employ Co-tracker3 [22] to track both the trajectory $\mathcal{X}_n \in \mathbb{R}^{T_n \times F \times 2}$ and the visibility $\mathcal{V}_i \in \{0, 1\}^{T_n \times F}$ of the query points across all T_n frames of each segment s_n .

C. 3D Track Estimation and Filtering

Given the resulting 2D trajectory of query points X_n for segment s_n , first, we lift the trajectory to 3D using the corresponding depth frame D_t and camera intrinsics \mathbf{K} , and we filter out points with invalid depth values. This results in a set of 3D tracks $\mathcal{P}_n \in \mathbb{R}^{T_n \times F \times 3}$ and a corresponding visibility mask $\mathcal{V}_n \in \{0, 1\}^{T_n \times F}$ indicating the validity of each 3D point at each time step. Second, to compensate for camera motion and obtain 3D trajectories in global coordinates, we transform all the 3D point tracks within each segment to the global frame assuming accurate camera odometry. Third, since some points may lie on static objects, we filter out these static points by computing the variance of each 3D track within the segment. Tracks with a variance below a certain percentile threshold σ_{static} are considered static and are discarded. Fourth, points may undergo occlusion during the interaction, potentially leading to unreliable point tracks. Therefore, we filter out tracks that are occluded for more than a specified percentage $\sigma_{reliable}$ of the frames within each segment. Lifting the 2D trajectory to 3D introduces high-frequency temporal jitter in the resulting 3D trajectory. We employ trajectory smoothing via minimizing a cost function (Eq. (1)) penalizing spikes in acceleration and high velocities while preserving fidelity

$$E(\mathbf{p}) = \sum_{t=1}^T v_t \|\mathbf{p}_t - \hat{\mathbf{p}}_t\|^2 + \lambda_{vel} \|\mathbf{p}_t - \mathbf{p}_{t-1}\|^2 + \lambda_{jerk} \|\mathbf{p}_t - 3\mathbf{p}_{t-1} + 3\mathbf{p}_{t-2} - \mathbf{p}_{t-3}\|^2, \quad (1)$$

where the first term promotes closeness between smoothed points \hat{p}_t and observed points p_t weighted by the point visibility v_t . The second term regularizes significant velocity changes weighted by λ_{vel} whereas the third term minimizes sudden changes in acceleration weighted by λ_{jerk} .

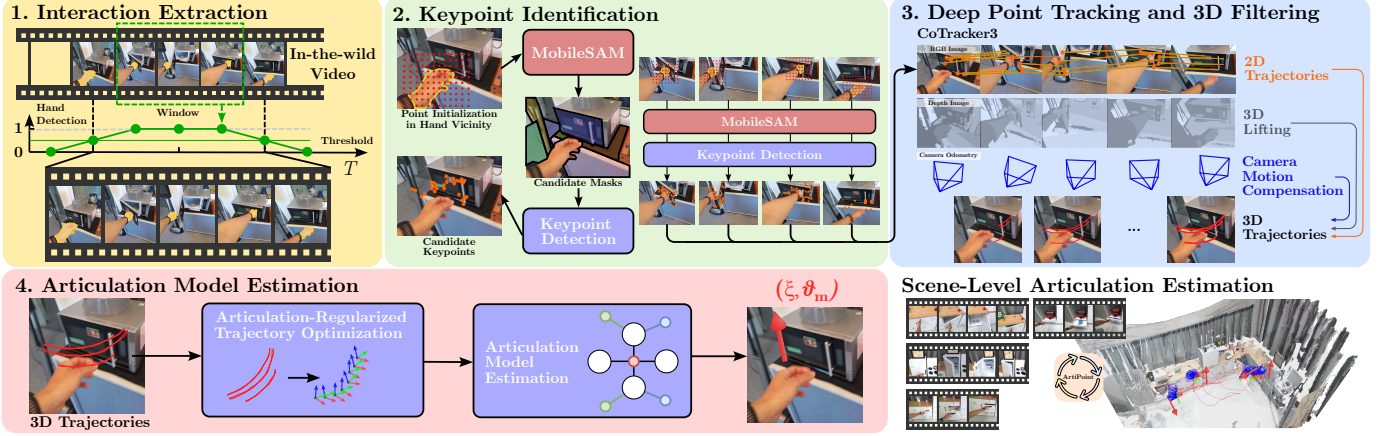


Fig. 1: Overview of our method: We take an ego-centric RGB-D video as input and employ hand tracking as a trigger signal to identify interaction segments (top left). We uniformly sample points around the hand masks and prompt a class-agnostic instance segmentation model (MobileSAM [41]), which yields object masks in the immediate vicinity that may be undergoing articulation. Given those masks, we detect stable keypoints (bottom left) that are fed into an any-point tracking model (CoTracker3 [22]) in order to obtain point trajectories throughout each entire articulation segment (top right). Finally, we estimate the underlying articulation model of the object through a factor graph formulation that operates directly on the obtained point trajectories (bottom right).

D. Exploiting the Articulation Prior

The most recent works examining articulation model estimation [19, 15, 4] opt for representing 1-DoF articulations as an articulation parameter $\xi = \langle \omega, v \rangle \in \mathfrak{se}(3) \subseteq \mathbb{R}^6$. Scaling an element $\xi \in \mathfrak{se}(3)$ with a configurational position θ is sufficient to represent any rigid-body transformation which can be decomposed into a proportional linear translation and rotation. Such a transformation $\xi\theta$ can be converted to the equivalent, non-linear $SE(3)$ rigid-body transform using the *exponential map* $\exp : \mathfrak{se}(3) \rightarrow SE(3)$, while the inverse of the exponential map is referred to as the *log map* $\log : SE(3) \rightarrow \mathfrak{se}(3)$ [2]. The advantage of this representation is its ability to represent diverse articulations, i.e. prismatic, revolute, and screw joints, while also being differentiable which makes them suitable for optimization.

In this work, we employ the factor graph formulation of Buchanan *et al.* [4] to estimate the articulation model which best fits an extracted point trajectory. However, this estimator operates on a sequence of pose observations between two moving parts, not directly on points of these parts. Thus, we extract poses as follows: Given our extracted points \mathcal{P}_n and visibility masks \mathcal{V}_n , we interpret these as observation \mathcal{Z} as

$$\mathcal{Z} = \left\{ P_m = \left\{ \langle p_f^t, p_f^{t+\epsilon} \rangle \mid \mathcal{V}_{n,f}^t \wedge \mathcal{V}_{n,f}^{t+\epsilon} \right\}_{m=1}^M \right\}, \quad (2)$$

where $p_f^t, p_f^{t+\epsilon}$ are two observations of point f at times t and $t + \epsilon$ taken from \mathcal{P}_n . Extracting trajectories from these observations is equivalent to identifying the transformations ${}^t\mathbf{T}_{t+\epsilon,m}$ minimizing

$${}^t\mathbf{T}_{t+\epsilon,m}^* = \min_{{}^t\mathbf{T}_{t+\epsilon,m}} \sum_{j=1}^{|P_m|} \|p_j^{t+\epsilon} - {}^t\mathbf{T}_{t+\epsilon,m} \cdot p_j^t\|, \quad (3)$$

where $\langle p_j^t, p_j^{t+\epsilon} \rangle \in P_m$. These local transformations can be integrated to yield poses ${}^W\mathbf{T}_m$ in the global frame of reference as by recursion: We define ${}^W\mathbf{T}_0 = {}^W\mathbf{T}_A$ and build on this root with ${}^W\mathbf{T}_m = {}^W\mathbf{T}_{m-1}$ where ${}^W\mathbf{T}_A$ describes a global

offset of the local trajectory. This local offset can reasonably be obtained as the mean of \mathcal{P}_n^1 . In our approach, we propose exploiting the context of our problem, an extend the estimation in Eq. (3) to be *articulation-regularized* across the entire motion. We do so, by introducing a shared base twist $\hat{\xi}$ for the entire path as

$$\hat{\xi}^*, \theta_1^*, \dots, \theta_M^* = \min_{\hat{\xi}, \theta_1, \dots, \theta_M} \sum_{m=1}^M \sum_{j=1}^{|P_m|} \|p_j^{t+\epsilon} - \exp(\hat{\xi}\theta_m) \cdot p_j^t\|, \quad (4)$$

where ${}^t\mathbf{T}_{t+\epsilon}$ can finally be extracted as ${}^t\mathbf{T}_{t+\epsilon,m} = \exp(\hat{\xi}^*\theta_m^*)$. As this formulation can be interpreted as a factor graph, we use GTSAM [9] to solve for it.

Finally, we extract the relative poses required by the articulation estimator [4], as ${}^A\mathbf{T}_m = {}^W\mathbf{T}_A^{-1} \cdot {}^W\mathbf{T}_m$ which is similar to the technique described in [4], where the authors use the initial frame of the articulation as the static frame. From this data the estimator then determines ξ_n^* for segment s_n .

IV. ARTI4D DATASET

We present the *Arti4D* dataset, which, to the best of our knowledge, is the first ego-centric *in-the-wild* human demonstration dataset capturing scene-level articulated object manipulation. With this dataset, we aim to create a foundation for studying the problem of *circumstantial articulated object estimation from human demonstrations*. In this problem, a human collects videos of interactions with multiple objects. The aim is to detect all interactions and estimate the articulation parameters of all manipulated objects. *Arti4D* consists of 45 egocentric RGB-D sequences across four distinct scenes (RR080, DR080, RH078, and RH201), featuring 411 human-object interactions recorded *in-the-wild*.

Unlike previous datasets such as *PARIS* [27], *DTA-Multi* [38], or *ArtGS-Multi* [28], the *in-the-wild* character presents novel challenges: the camera poses are dynamically changing as the scene is explored, the articulated objects are partially occluded

throughout the interaction, and the objects are not isolated, i. e., are an integrated part of the larger environment. We present exemplary interactions and the associated camera trajectory in Fig. 1. In this manner, our dataset markedly differentiates itself from existing datasets in the domains of articulated object estimation [27, 39, 30] and 3D scene understanding [42, 24, 8].

In tandem with the sequences, our dataset provides ground-truth axis labels for all articulations in the scenes as $\mathfrak{se}(3)$ points, as well as difficulty ratings of all interactions ranked as either EASY or HARD. This is based on the level of hand visibility throughout each interaction, whether an object exhibits a reasonable number of depth measurements, whether the articulating hand is fully retrieved in between articulations, and whether large extents of the object are occluded. Furthermore, we provide cm-accurate ground-truth camera poses, obtained via external tracking, to ease the task of estimating articulations from dynamic camera observations. In addition, we believe that Arti4D constitutes a challenging benchmark for the tasks of visual odometry and simultaneous localization and mapping. Given that the articulations cover significant extents of the field of view, finding stable correspondences becomes difficult [31], ultimately complicating odometry estimation.

V. EXPERIMENTAL RESULTS

In the following, we present our experimental findings by comparing against two sets of capable baselines, demonstrating the performance of ArtiPoint on various splits of Arti4D, and lastly, ablating key components of our proposed pipeline. As introduced in Sec. IV, we employ the Arti4D dataset to evaluate the performance of estimating articulation models *in-the-wild*. The reported results are averaged across all object interactions recorded across 44 RGB-D sequences stemming from four diverse environments: RR080, DR080, RH078, and RH201.

Baselines: We consider two sets of baselines. The first set consists of the deep learning-based approach Ditto [17] and the Gaussian splatting-based method ArtGS [28]. Both methods are tailored towards semi-static variations of articulation states of isolated objects. As this is somewhat different from the in-the-wild paradigm introduced as part of the Arti4D dataset, we adapted these methods by masking out hands and selecting a representative number of frames from before, during, and after the interaction (Sec. III-A) from which to reconstruct the motion model. However, both Ditto and ArtGS are prone to occlusions and variations in camera poses as they aim to identify correspondences among various articulation states. Our second set of baselines consists of two factor graph-based articulation estimation pipelines. As Sturm *et al.* [37, 35, 36] only provides a back-end for estimating the articulation model, we employ the ArtiPoint frontend but fit oriented bounding boxes to all detected objects based on the masks extracted by MobileSAM [41]. Using the 3D bounding boxes of the objects, we derive part poses across all frames and feed these into the estimation framework of Sturm *et al.* [37]. Lastly, we also employ our front-end pipeline to provide object poses ${}^W\mathbf{T}_A$ from any-point tracking for Sturm *et al.* [37].

Metrics: In order to quantify the accuracy of the estimated articulation model, we employ the following metrics: 1. The angular error $\theta_{err} = \arccos(|\hat{\mathbf{a}} \cdot \mathbf{a}|)$ evaluated between ground truth and predicted axes, 2. the Euclidean distance d_{L2} between predicted and ground truth axes (revolute joints only), and the accuracy of predicted joint types.

Implementation Details: For our interaction extraction module (Sec. III-A) we use $w_h = 6$, $T_{min} = 30$, and $T_{max} = 120$. For our 3D track estimation and filtering (Sec. III-C) we choose a track reliability threshold $\sigma_{reliable} = 0.5$, and for the optimization-based smoothing we choose $\lambda_{vel} = 0.5$, $\lambda_{jerk} = 5$. We employ an NVIDIA 3090 GPU, 128GB RAM, and an AMD Ryzen 9 7950X CPU.

A. Quantitative Results on Arti4D

We report an overall comparison on the Arti4D in Tab. I. We note that ArtGS [28] and Ditto [21] perform poorly in both estimating and classifying joints, which is primarily due to partial observability throughout interactions. In addition, we observe that the non-isolatedness of objects complicates estimation accuracy in the case of ArtGS and Ditto. Using Sturm [37] with the bounding box-based front-end leads to better results compared to ArtGS [28] and Ditto [21]. In addition, we compare the impact of our articulation-regularized trajectory estimation against a bounding-box based approach similar to [34]. We note that this approach underperforms greatly on estimating the direction of the articulation axes, incurring over twice the angular error. Also the trajectories obtained using the RANSAC estimator, incur more error than the ones estimated using our estimator (Eq. (4)). We note the increased positional error on revolute joints. articulation-regularized trajectory pose estimation from point observations outperforms a naive RANSAC model estimating transforms on a per frame-pair basis. We note that our trajectory estimator also benefits the estimation approach by Sturm *et al.* We also study the performance of the different approaches across easy and difficult estimation problems in Tab. II. We find that ArtiPoint performs best on both in parameter estimation.

B. Qualitative Results on Arti4D

In addition to the quantitative comparisons, we depict predictions on two drawers (prismatic) and a revolute joint of a storage case in Fig. 2. The two drawers are observed under a small number of keypoints whereas the case exhibits frequently missing depth measurements. Nonetheless, ArtiPoint is able to estimate the underlying motion model robustly under adverse conditions.

C. Ablation Study

To further understand the contribution of different components of the ArtiPoint pipeline, we conducted an ablation study on the Arti4D dataset. From Tab. III, we can conclude that using ORB keypoints [32] leads to lower results across all metrics. Increasing the keyframe stride to 4 in Sec. III-B reduces the computation overhead but also reduces performance. Furthermore, deactivating our trajectory smoothing impacts the accuracy of estimating the position of revolute joints, as well

TABLE I: Overall comparison of classical and deep articulation estimation methods on the *Arti4D* dataset. We report the axis-angle errors θ_{err} , the positional errors d_{L2} of revolute joints, and the accuracy of joint type detection. We note that *ArtiPoint* achieves the lowest error in estimating articulation parameters.

Method	Prismatic joints		Revolute joints		Type accuracy [%]	
	θ_{err} [deg]	d_{L2} [m]	θ_{err} [deg]	d_{L2} [m]	Prismatic	Revolute
ArtGS	28.86	–	35.09	0.90	1.00	0.00
Ditto	54.00	–	56.09	0.54	0.17	0.80
Sturm <i>et al.</i> [37]	36.59	–	34.26	1.04	0.67	0.66
Sturm <i>et al.</i> [37] w/ bbox	42.59	–	45.27	1.15	0.65	0.56
Buchanan [4] w/ bbox	51.39	–	54.30	0.29	0.01	0.97
ArtiPoint w/ RANSAC	20.13	–	28.70	0.34	0.55	0.86
ArtiPoint	17.12	–	26.50	0.28	0.69	0.89

TABLE II: Difficulty-level comparison of estimation methods on *Arti4D*.

Diff.	Method	Prismatic joints		Revolute joints		Type accuracy [%]	
		θ_{err} [deg]	d_{L2} [m]	θ_{err} [deg]	d_{L2} [m]	Prismatic	Revolute
EASY	ArtGS	23.80	-	26.40	0.71	1.00	0.00
	Ditto	51.99	-	51.22	0.44	0.19	0.85
	ArtiPoint w/ [37]-AE	32.82	-	24.47	0.82	0.72	0.83
	ArtiPoint	15.57	-	24.21	0.36	0.74	0.82
HARD	ArtGS	37.26	-	40.82	1.03	1.00	0.00
	Ditto	57.88	-	59.04	0.61	0.11	0.76
	ArtiPoint w/ [37]-AE	44.38	-	40.78	1.17	0.55	0.56
	ArtiPoint	20.26	-	27.99	0.23	0.60	0.93

TABLE III: Ablations of key components of our approach on the *Arti4D* dataset. As before, we report the same set of metrics as in Tab. I.

	Prismatic joints		Revolute joints		Type Accuracy [%]	
	θ_{err} [deg]	d_{L2} [cm]	θ_{err} [deg]	d_{L2} [cm]	Prismatic	Revolute
w/ ORB keypoints	20.70	-	33.06	0.34	0.69	0.83
w/ keyframe stride 4	17.57	-	28.55	0.30	0.66	0.87
w/o trajectory smoothing	17.63	-	26.79	0.36	0.61	0.81
w/ unreliable tracks	17.45	-	28.17	0.36	0.67	0.85
ArtiPoint (ours)	17.12	–	26.50	0.28	0.69	0.89




Fig. 2: Qualitative results on *Arti4D*: Estimated joint axis and pose trajectory of a drawer in a kitchen scene (RH201), a revolute joint of a storage case (RH078), and another drawer in RH201.

as reducing the overall joint type prediction accuracy. Finally, using unreliable tracks or highly occluded tracks has the same impact as disabling smoothing.

VI. LIMITATIONS

Our proposed *ArtiPoint* framework relies on key point tracking, which inherently requires the presence of somewhat distinct features that allow tracking. However, state-of-the-art any-point tracking methods as discussed in Sec. II show strong robustness even under feature-sparse conditions. However, when objects do not yield dependable depth estimates, due to properties such as color, surface paint, or reflections, tracking their motion becomes particularly challenging. To address this limitation, future work will explore the integration of inpainting methods informed by monocular depth estimation. Additionally, our system is currently limited to relatively simple articulation models that involve only two-body kinematics,

making it infeasible to estimate more complex articulations. Another limitation of our approach is the requirement for a trigger signal to segment the RGB-D sequence. In our implementation, this trigger takes the form of human hand detection, which is chosen due to its reliability in indicating articulations. Alternative methods in action recognition, whether in egocentric or third-person views, often lack robustness or require extensive inference times, similar to more sophisticated models.

VII. CONCLUSION

We presented ArtiPoint, a novel framework for estimating articulation motion models from human object interactions *in-the-wild*. Unlike prior works operating in controlled settings involving fixed camera poses, isolated objects, and full observability, we employ deep any-point tracking and factor graph optimization to infer articulation models from ego-centric human demonstrations. As part of our work, we introduced Arti4D, the first real-world, ego-centric articulation demonstration dataset that includes odometry and axis ground truth. ArtiPoint demonstrates robust performance while outperforming both classical and previous deep and object-wise rendering methods in articulated object estimation. To foster further research in articulation understanding in realistic, unconstrained settings, we make code, model predictions, and dataset publicly available.

REFERENCES

- [1] Albert Bandura and Richard H Walters. *Social learning theory*, volume 1. Prentice hall Englewood Cliffs, NJ, 1977.
- [2] Timothy D Barfoot. *State estimation for robotics*. Cambridge University Press, 2024.
- [3] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European Conference on Computer Vision*, pages 404–417, 2006.
- [4] Russell Buchanan, Adrian Röfer, João Moura, Abhinav Valada, and Sethu Vijayakumar. Online estimation of articulated objects with factor graphs using vision and proprioceptive sensing. In *IEEE International Conference on Robotics and Automation*, pages 16111–16117, 2024.
- [5] Guglielmo Camporese. Hands segmentation is all you need. <https://github.com/guglielmocamporese>, 2021. URL <https://github.com/guglielmocamporese/hands-segmentation-pytorch>.
- [6] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *International Journal of Robotics Research*, 2024.
- [7] Seokju Cho, Jiahui Huang, Jisu Nam, Honggyu An, Seungryong Kim, and Joon-Young Lee. Local all-pair correspondence for point tracking. In *European Conference on Computer Vision*, pages 306–325, 2024.
- [8] Alexandros Delitzas, Ayca Takmaz, Federico Tombari, Robert Sumner, Marc Pollefeys, and Francis Engelmann. SceneFun3D: Fine-Grained Functionality and Affordance Understanding in 3D Scenes. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2024.
- [9] Frank Dellaert and GTSAM Contributors. borglab/gtsam, May 2022. URL <https://github.com/borglab/gtsam>.
- [10] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. Tap-vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems*, 35:13610–13626, 2022.
- [11] Carl Doersch, Yi Yang, Mel Vecerik, Dilara Gokay, Ankush Gupta, Yusuf Aytar, Joao Carreira, and Andrew Zisserman. Tapir: Tracking any point with per-frame initialization and temporal refinement. In *International Conference on Computer Vision*, pages 10061–10072, 2023.
- [12] Carl Doersch, Pauline Luc, Yi Yang, Dilara Gokay, Skanda Koppula, Ankush Gupta, Joseph Heyward, Ignacio Rocco, Ross Goroshin, João Carreira, et al. Bootstrap: Bootstrapped training for tracking-any-point. In *Proceedings of the Asian Conference on Computer Vision*, pages 3257–3274, 2024.
- [13] Arjun Gupta, Michelle Zhang, Rishik Sathua, and Saurabh Gupta. Opening cabinets and drawers in the real world using a commodity mobile manipulator. *Robotics: Science and Systems*, 2025.
- [14] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *European Conference on Computer Vision*, pages 59–75, 2022.
- [15] Nick Heppert, Toki Migimatsu, Brent Yi, Claire Chen, and Jeannette Bohg. Category-independent articulated object tracking with factor graphs. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3800–3807, 2022.
- [16] Daniel Honerkamp, Tim Welschehold, and Abhinav Valada. N2m2: Learning navigation for arbitrary mobile manipulation motions in unseen and dynamic environments. *IEEE Transactions on Robotics*, 39(5):3601–3619, 2023.
- [17] Cheng-Chun Hsu, Zhenyu Jiang, and Yuke Zhu. Ditto in the house: Building articulation models of indoor scenes through interactive perception. In *IEEE International Conference on Robotics and Automation*, pages 3933–3939, 2023.
- [18] Xiaoxia Huang, Ian Walker, and Stan Birchfield. Occlusion-aware reconstruction and manipulation of 3d articulated objects. In *IEEE International Conference on Robotics and Automation*, pages 1365–1371, 2012.
- [19] Ajinkya Jain, Rudolf Lioutikov, Caleb Chuck, and Scott Niekum. Screwnet: Category-independent articulation model estimation from depth images using screw theory. In *IEEE International Conference on Robotics and Automation*, pages 13670–13677, 2021.
- [20] Ajinkya Jain, Stephen Giguere, Rudolf Lioutikov, and Scott Niekum. Distributional depth-based estimation of object articulation models. In *Conference on Robot*

- Learning*, pages 1611–1621, 2022.
- [21] Zhenyu Jiang, Cheng-Chun Hsu, and Yuke Zhu. Ditto: Building digital twins of articulated objects from interaction. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, pages 5616–5626, 2022.
 - [22] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024.
 - [23] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker: It is better to track together. In *European Conference on Computer Vision*, pages 18–35, 2024.
 - [24] Christina Kassab, Sacha Morin, Martin Büchner, Matías Mattamala, Kumaraditya Gupta, Abhinav Valada, Liam Paull, and Maurice Fallon. Openlex3d: A new evaluation benchmark for open-vocabulary 3d scene representations. *arXiv preprint arXiv:2503.19764*, 2025.
 - [25] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model. *Conference on Robot Learning*, 2024.
 - [26] Xiaolong Li, He Wang, Li Yi, Leonidas J Guibas, A Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, pages 3706–3715, 2020.
 - [27] Jiayi Liu, Ali Mahdavi-Amiri, and Manolis Savva. Paris: Part-level reconstruction and motion analysis for articulated objects. In *International Conference on Computer Vision*, pages 352–363, 2023.
 - [28] Yu Liu, Baoxiong Jia, Ruijie Lu, Junfeng Ni, Song-Chun Zhu, and Siyuan Huang. Building interactable replicas of complex articulated objects via gaussian splatting. In *International Conf. on Learning Representations*, 2025.
 - [29] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60:91–110, 2004.
 - [30] Roberto Martín-Martín, Clemens Eppner, and Oliver Brock. The rbo dataset of articulated objects and interactions. *International Journal of Robotics Research*, 38(9):1013–1019, 2019.
 - [31] E. Palazzolo, J. Behley, P. Lottes, P. Giguère, and C. Stachniss. ReFusion: 3D Reconstruction in Dynamic Environments for RGB-D Cameras Exploiting Residuals. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
 - [32] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *International Conference on Computer Vision*, pages 2564–2571, 2011.
 - [33] Jianbo Shi et al. Good features to track. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, pages 593–600, 1994.
 - [34] J Sturm, Vijay Pradeep, Cyrill Stachniss, Christian Plagemann, Kurt Konolige, and Wolfram Burgard. Learning kinematic models for articulated objects. 2009.
 - [35] Jurgen Sturm, Christian Plagemann, and Wolfram Burgard. Unsupervised body scheme learning through self-perception. In *IEEE International Conference on Robotics and Automation*, pages 3328–3333, 2008.
 - [36] Jürgen Sturm, Kurt Konolige, Cyrill Stachniss, and Wolfram Burgard. 3d pose estimation, tracking and model learning of articulated objects from dense depth video using projected texture stereo. In *RGB-D: Advanced Reasoning with Depth Cameras Workshop, RSS*, 2010.
 - [37] Jürgen Sturm, Cyrill Stachniss, and Wolfram Burgard. A probabilistic framework for learning kinematic models of articulated objects. *Journal of Artificial Intelligence Research*, 41:477–526, 2011.
 - [38] Yijia Weng, Bowen Wen, Jonathan Tremblay, Valts Blukis, Dieter Fox, Leonidas Guibas, and Stan Birchfield. Neural implicit representation for building digital twins of unknown articulated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3141–3150, 2024.
 - [39] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*.
 - [40] Vicky Zeng, Tabitha Edith Lee, Jacky Liang, and Oliver Kroemer. Visual identification of articulated object parts. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2443–2450, 2021.
 - [41] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023.
 - [42] Chenyangguang Zhang, Alexandros Delitzas, Fangjinhua Wang, Ruida Zhang, Xiangyang Ji, Marc Pollefeys, and Francis Engelmann. Open-vocabulary functional 3d scene graphs for real-world indoor spaces. *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2025.
 - [43] Harry Zhang, Ben Eisner, and David Held. Flowbot++: Learning generalized articulated objects manipulation via articulation projection. *Conference on Robot Learning*, 2023.
 - [44] Yang Zheng, Adam W Harley, Bokui Shen, Gordon Wetzstein, and Leonidas J Guibas. Pointodysey: A large-scale synthetic dataset for long-term point tracking. In *International Conference on Computer Vision*, pages 19855–19865, 2023.

Supplementary Material

In this supplementary to the main manuscript, we provide additional insights on the utilized metrics in Sec. S.1.1, provide additional quantitative and qualitative results in Sec. S.1.2, and provide in-depth explanations regarding the introduced Arti4D dataset in Sec. S.2.

S.1. EXPERIMENTAL RESULTS

In this section, we present additional details on our experiments, the metrics we employ, and the additional ablation study.

A. Metrics

In Sec. V, we report quantitative results of our approach and the baselines. In this section, we detail the definition of the metrics that we employ to quantify performance. As stated in the main manuscript, our metrics consist of positional and angular error. We compute the positional error d_i for the placement of an articulation’s rotation axis $\hat{\mathbf{a}}$ and the ground truth axis \mathbf{a}_{gt} using the help of supporting points $\hat{\mathbf{p}}, \mathbf{p}_{gt}$ as

$$d_i = \begin{cases} \frac{(\hat{\mathbf{p}}_i - \mathbf{p}_{gt})^\top (\hat{\mathbf{a}}_i \times \mathbf{a}_{gt})}{\|\hat{\mathbf{a}}_i \times \mathbf{a}_{gt}\|} & \text{if } \|\hat{\mathbf{a}}_i \times \mathbf{a}_{gt}\| > \epsilon \\ \|(\hat{\mathbf{p}}_i - \mathbf{p}_{gt}) \times \mathbf{a}_{gt}\| & \text{else} \end{cases}, \quad (1)$$

where the first case covers the case in which the axes are not parallel with $\epsilon = 10^{-4}$. If the inference does not provide a point on the axis directly, but a twist, we compute $\hat{\mathbf{p}}_i = \frac{\omega_i \times \mathbf{v}_i}{\|\omega_i\|^2}$. The angular error of a prediction, we compute simply as using the dot-product of the axes

$$\Theta_{err,i} = \cos^{-1} |\mathbf{a}_{gt}^\top \hat{\mathbf{a}}_i|. \quad (2)$$

We report only angular errors for prismatic joints, as the location of the axis does not have any effect on the motion of the parts of the articulated object.

B. Additional Results

In the following sections, we share additional ablations on hyperparameters of ArtiPoint and scene-respective results.

1) *Ablation on Hand Detection*: Extracting the interaction intervals, as described in Sec. III-A and illustrated in Fig. S.1, is a critical component of the ArtiPoint pipeline as it directly affects the number of articulated objects detected. As such, it requires careful parameter tuning, as small values of w_h or T_{min} increase the number of false positives. To better understand the impact of these parameters on the final results, we conduct a detailed ablation study with its findings presented in Tab. S.1. Decreasing T_{min} leads to a noticeable degradation in both angular error (θ_{err}) and joint type classification accuracy for both joint types. Increasing w_h leads to over-smoothing of the raw hand detection signal, causing the segments to contain elongated idle phases at the start and the end, thus inducing unfavorable noise, and resulting in a noticeable degradation in both angular error (θ_{err}) and joint type classification accuracy for both joint types.

2) *Ablation on Point Tracking*: In this section, we present ablation results on the any-point tracking component as illustrated in Fig. S.2. In particular, we evaluate to which degree different keyframe strides used as input to the point tracking component affect the downstream axis prediction performance (see Sec. III-B). Choosing the keyframe stride is vital hyperparameter of the point tracking stage. While a smaller stride leads to an increase in the number of detected points to be tracked by Cotracker3 [22], larger strides reduce the number of detected points to be tracked, thereby lowering the computational load. Retaining a sufficient number of points is necessary for estimating an object’s 3D trajectory over time, therefore, careful tuning of this parameter is important. Overall, we observe the lowest prediction errors for a keyframe stride of 2. However, except for the revolute angular error, we do not observe a clear trends among the other two metrics.

C. Scene-Respective Results

In addition, to the EASY/HARD differentiation evaluated in Tab. II, we show scene-respective results in Tab. S.2. This involves averaging the predictions of all object interactions contained in a scene split. We list the number of sequences per scene as well as the number of labeled objects per sequence in Sec. S.2. As reported in Tab. S.2, RH078 constitutes the most difficult split of the Arti4D dataset. In comparison, DR080 and RH201 seem to represent simpler environments.

We attribute worse results on RH078 to a number of non-separable interactions as the hand is occasionally not fully retrieved between interactions. As a consequence, the proposed interaction extraction baseline potentially fails at differentiating two different interactions. We have mentioned the hand trigger limitation in Sec. VI and leave improvements on that front to future work. In addition to that, we observe that there are comparably more revolute joints in scene RH078 whose associated objects are rather textureless and of metallic character, thus hindering consistent depth observations.

D. Qualitative Results

In the following, we provide additional qualitative results. In Fig. S.3, we visualize the output of our proposed interaction extraction, any-point tracking and track filtering components. We observe a sufficient number of point trajectories even under partially missing depth measurements or feature-sparse textures. In addition to that, we visualize a full scene-level output of ArtiPoint on DR080 scene in Fig. S.4. Overall, ArtiPoint detects the majority of interactions and produces reliable estimates considering the in-the-wild character of the recorded Arti4D sequences.

S.2. ARTI4D DATASET

In the following, we provide additional insight on the in-the-wild object articulation dataset Arti4D. We provide 45 sequences across four distinct environments as listed in Sec. S.2 and visualized in Fig. S.7. In addition to the sequence IDs and the recording names of the produced sequences, we report the number of labeled objects for each sequence, the ratio between prismatic and revolute joints as well as the ratio between easy

TABLE S.1: Ablations of key parameters for extracting the interaction intervals component Sec. III-A on the Arti4D dataset. As before, we report the same set of metrics as in Tab. I

Method	Prismatic joints		Revolute joints		Type accuracy [%]	
	$\theta_{err}[\text{deg}]$	$d_{L2}[\text{m}]$	$\theta_{err}[\text{deg}]$	$d_{L2}[\text{m}]$	Prismatic	Revolute
$w_h = 12$	19.35	–	31.95	0.36	0.65	0.83
$T_{max} = 90$	18.45	–	27.46	0.24	0.64	0.89
$T_{min} = 15$	24.78	–	32.75	0.35	0.59	0.85
ArtiPoint	17.12	–	26.50	0.28	0.69	0.89




Fig. S.1: Hand detection and interaction extraction: We visualize a live frame of an interaction, including a hand mask marked violet (top). In addition, we visualize the frequency of raw hand detections over time up to the live frame as well as its moving average (bottom). The horizontal red line indicates the threshold at which an interaction segment is created, given the moving average signal. The vertical dashed red line indicates the current frame.

and hard objects. First, note that the number of objects is not equal to the number of object interactions per sequence, as several sequences contain repeated interactions with the same single object instance. This constitutes a corner case in terms of articulated object interaction as it requires prediction methods to fuse, e.g., two predictions belonging to a single object. While most interactions are separable by detecting the absence of a hand mask, especially the RH078 split which contains a number of hard-to-separate interactions. This is due to the fact that the interacting hand was not always fully retrieved in-between interactions of two distinct objects. Similar to the repeated interactions case mentioned before, this represents another corner case requiring advanced action recognition.

As part of the dataset, we make both rosbags and processed

raw data public. While the rosbags include TF data at a higher frequency, the raw data includes aligned RGB, depth, and camera poses at 15 Hz. We employed an Azure Kinect RGB-D camera that was handheld throughout all interactions. In terms of ground truth camera pose retrieval, we employed external tracking using HTC Vive trackers, which provide cm-level accuracy. In case of sudden odometry glitches induced by considerable occlusions or reflections on glass or metal, we have removed those sequences from the dataset. We found that running classical structure-from-motion approaches to reconstruct the underlying sequence fails as significant parts of the camera field of view cover articulations. In turn, the contained articulations break assumptions towards mostly *static* visual correspondence made in structure-from-motion methods.

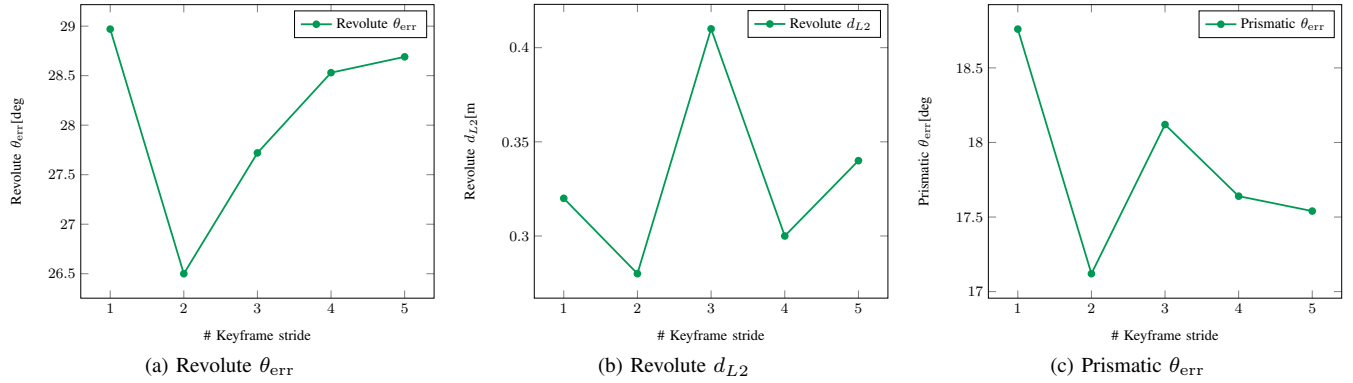


Fig. S.2: Impact of keyframe stride on tracking accuracy and error. (a) The revolute joint angular error θ_{err} [deg] exhibits a minimum at stride 2. (b) The revolute joint positional error d_{L2} [m] is lowest at a stride of 2, with higher errors observed for both smaller and larger strides. (c) Similarly, the prismatic joint angular error θ_{err} exhibits its minimum at stride of 2. We conclude that a stride of 2 is optimal in terms of point density and computational efficiency.

TABLE S.2: Scene-respective results: We report the scene-wise results using the established set of metrics. We find that RH078 constitutes the most difficult split of the Arti4D dataset.

Method	Prismatic joints		Revolute joints		Type accuracy [%]	
	θ_{err} [deg]	d_{L2} [m]	θ_{err} [deg]	d_{L2} [m]	Prismatic	Revolute
RH078	30.36	–	27.16	0.395	0.622	0.778
RR080	16.17	–	21.84	0.332	0.598	0.889
DR080	14.62	–	27.50	0.279	0.824	0.926
RH201	14.91	–	25.44	0.246	0.723	0.886
Overall	17.12	–	26.50	0.28	0.69	0.89

The ground truth object axes were labeled based on the reconstructed sequences using Blender, exported as JSONs, and verified by a second reviewer. We also make the sequence reconstructions in the form of meshes and point clouds available. Furthermore, we provide metadata on difficulty levels and interaction keyframe ground truth labels. We provide four exemplary depth-masked RGB frames covering an interaction of a microwave featuring semi-transparent glass and metal surfaces in Fig. S.6. As depicted, a considerable part of the object does not produce depth estimates. Objects of that kind are labeled using the HARD category.

In addition, we depict four TSDF-reconstructed sequences stemming from each of the splits in Fig. S.5. The reconstructions reflect minimal ground truth odometry drift and enable precise anchoring object axes.

We depict several RGB frames drawn from the four distinct scene splits in Fig. S.7, underlining the variety of objects and challenging conditions of the proposed dataset.





Fig. S.3: Smoothed point trajectories: We visualize a cabinet (revolute) and its tracked keypoints (left) as well as a linear slider shelf (prismatic) on the right. Each keypoint trajectory is represented with a unique color. Both sets of point trajectories visualized constitute the output of our track filtering introduced in Sec. III-C.



Fig. S.4: Scene-level prediction: We depict a full scene-level output of the ArtiPoint framework on sequence `scene_2025-04-11-11-44-32` of the DR080 scene. Yellow arrows denote axes of motion of predicted object interactions while coordinate frames represent the estimate part poses throughout articulation based on the proposed estimation framework (see Sec. III-D).




DR080



RH078



RH201



RR080

Fig. S.5: We visualize the reconstructed scenes of the four Arti4D environments: DR080, RH078, RH201, and RR080.




Fig. S.6: We visualize an exemplary frame with the depth projected onto the RGB observation in red whenever a depth reading was available for the particular image coordinate. As depicted, metallic, semi-transparent, or dark-colored objects do not produce reliable depth estimates, ultimately complicating the lifting of 3D point trajectories used to estimate the underlying articulation.

TABLE S.3: Overview of all Arti4D demonstration sequences: We assign sequence identifiers to each recording and list the number of objects interacted with per sequence as well as the distribution of prismatic vs. revolute joints (PRISM REV) and objects classified as either EASY or HARD. The number of objects corresponds to the number of annotated axes per sequence. As such, it does not account for repeated interactions of the same object.

<i>Scene</i>	<i>Sequence ID</i>	<i>Recording</i>	<i># Objects</i>	<i># PRISM / REV</i>	<i># EASY / HARD</i>
RH078	RH078-00	scene_2025-04-04-19-14-38	7	3 / 4	7 / 0
	RH078-01	scene_2025-04-04-19-18-54	6	2 / 4	4 / 2
	RH078-02	scene_2025-04-07-11-39-17	7	3 / 4	7 / 0
	RH078-03	scene_2025-04-07-11-41-52	8	3 / 5	6 / 2
	RH078-04	scene_2025-04-07-11-48-40	7	3 / 4	7 / 0
	RH078-05	scene_2025-04-09-10-30-11	8	5 / 3	3 / 5
	RH078-06	scene_2025-04-09-10-32-52	6	5 / 1	3 / 3
	RH078-07	scene_2025-04-09-10-35-47	7	6 / 1	2 / 5
	RH078-08	scene_2025-04-09-10-38-38	8	5 / 3	4 / 4
	RH078-09	scene_2025-04-09-10-46-48	7	7 / 0	4 / 3
	RH078-10	scene_2025-04-09-10-49-20	7	6 / 1	2 / 5
RR080	RR080-00	scene_2025-04-10-13-11-16	18	15 / 3	10 / 8
	RR080-01	scene_2025-04-10-16-05-09	15	13 / 2	9 / 6
	RR080-02	scene_2025-04-17-15-25-14	11	11 / 0	7 / 4
	RR080-03	scene_2025-04-17-15-33-44	9	9 / 0	7 / 2
	RR080-04	scene_2025-04-22-09-53-49	8	8 / 0	5 / 3
	RR080-05	scene_2025-04-22-09-56-24	10	9 / 1	9 / 1
	RR080-06	scene_2025-04-22-09-58-49	7	7 / 0	4 / 3
	RR080-07	scene_2025-04-22-11-45-15	9	8 / 1	7 / 2
	RR080-08	scene_2025-04-22-11-48-01	9	8 / 1	7 / 2
	RR080-09	scene_2025-04-22-11-50-40	8	7 / 1	6 / 2
DR080	DR080-00	scene_2025-04-11-11-44-32	11	7 / 4	5 / 6
	DR080-01	scene_2025-04-11-12-58-58	9	5 / 4	4 / 5
	DR080-02	scene_2025-04-11-13-01-59	9	5 / 4	4 / 5
	DR080-03	scene_2025-04-11-13-18-00	9	4 / 5	3 / 6
	DR080-04	scene_2025-04-11-13-43-03	11	7 / 4	5 / 6
	DR080-05	scene_2025-04-11-14-01-06	11	7 / 4	5 / 6
	DR080-06	scene_2025-04-11-15-43-24	11	7 / 4	5 / 6
	DR080-07	scene_2025-04-11-15-46-48	11	6 / 5	4 / 7
RH201	RH201-00	scene_2025-04-24-17-52-21	11	5 / 6	6 / 5
	RH201-01	scene_2025-04-24-17-54-13	9	5 / 4	5 / 4
	RH201-02	scene_2025-04-24-19-18-42	11	4 / 7	7 / 4
	RH201-03	scene_2025-04-24-19-21-50	8	2 / 6	5 / 3
	RH201-04	scene_2025-04-24-19-24-09	9	5 / 4	6 / 3
	RH201-05	scene_2025-04-25-10-36-37	9	5 / 4	7 / 2
	RH201-06	scene_2025-04-25-10-53-40	8	4 / 4	3 / 5
	RH201-07	scene_2025-04-25-10-56-33	9	4 / 5	4 / 5
	RH201-08	scene_2025-04-25-11-11-47	16	7 / 9	7 / 9
	RH201-09	scene_2025-04-25-11-15-47	7	5 / 2	4 / 3
	RH201-10	scene_2025-04-25-14-58-42	9	6 / 3	4 / 5
	RH201-11	scene_2025-04-25-15-02-14	7	3 / 4	4 / 3
	RH201-12	scene_2025-04-25-15-04-48	7	4 / 3	3 / 4
	RH201-13	scene_2025-04-25-15-16-29	9	5 / 4	3 / 6
	RH201-14	scene_2025-04-25-15-19-22	10	5 / 5	5 / 5
	RH201-15	scene_2025-04-25-15-22-54	8	4 / 4	3 / 5

DR080



RH078



RH201



RR080



Fig. S.7: We visualize several object interactions across the four different environments (DR080, RH078, RH201, RR080) captured as part of the Arti4D dataset.