

ART: Articulated Reconstruction Transformer

Zizhang Li^{1,2*} Cheng Zhang¹ Zhengqin Li¹ Henry Howard-Jenkins¹ Zhaoyang Lv¹ Chen Geng²
Jiajun Wu² Richard Newcombe¹ Jakob Engel¹ Zhao Dong¹

¹Reality Labs Research, Meta ²Stanford University

<https://kyleleey.github.io/ART/>



Figure 1. We propose ART, a feed-forward model that reconstructs articulated objects from images by decomposing them into rigid parts. ART predicts each part’s geometry, texture, and articulation structure. Here we showcase its results on diverse categories, with each image pair displaying the part-based decomposition (left, each movable part in a unique color) and the final textured mesh (right).

Abstract

We introduce ART, Articulated Reconstruction Transformer—a category-agnostic, feed-forward model that reconstructs complete 3D articulated objects from only sparse, multi-state RGB images. Previous methods for articulated object reconstruction either rely on slow optimization with fragile cross-state correspondences or use feed-forward models limited to specific object categories. In contrast, ART treats articulated objects as assemblies of rigid parts, formulating reconstruction as part-based prediction. Our newly designed transformer architecture maps sparse image inputs to a set of learnable part slots, from which ART jointly decodes unified representations for individual parts, including their 3D geometry, texture, and explicit articulation parameters. The resulting reconstructions are physically interpretable and readily exportable for simulation. Trained on a large-scale, diverse dataset

*Work done during internship at Meta

with per-part supervision, and evaluated across diverse benchmarks, ART achieves significant improvements over existing baselines and establishes a new state of the art for articulated object reconstruction from image inputs.

1. Introduction

Articulated objects are ubiquitous in daily lives and central to human–scene interactions [38]. Accurately constructing their digital replicas is important for VR/AR, robotics, and embodied AI [9, 10, 25, 27, 28, 53, 68]. While recent 3D generation and reconstruction methods have significantly advanced the automatic creation of static assets [16, 34, 43, 48, 54, 61, 65], articulated objects remain challenging as they require recovering both geometry and underlying kinematic structure. Today, building such models still demands extensive expert effort, making the process labor-intensive and hard to scale, and ultimately limiting the accessibility and realism of articulated content for

large-scale interactive environments.

To automate this process, we tackle the challenging problem of *image-based articulated object reconstruction*: recovering a complete 3D representation for an articulated object, including geometry, texture, and its underlying articulation structure. We specifically focus on a practical, yet difficult setting: reconstructing articulated objects in a feed-forward manner from only a sparse set of multi-state RGB images. This setup is important for scalability, as dense multi-view, multi-state capture is often infeasible in real-world scenarios. However, the sparsity of inputs poses a significant challenge, requiring the inference of complex 3D shape, material, and articulation structure from limited visual cues, a task where existing articulated object reconstruction methods typically fall short.

Existing approaches to image-based articulated object reconstruction can be categorized into *per-object optimization* and *feed-forward learning*, both of which are ill-suited to our target setting. Per-object optimization methods [15, 21, 36, 42, 62] achieve high-fidelity reconstructions but are impractically slow due to the lengthy test-time optimization; they also depend on dense observations (often ~ 100 views) and fragile cross-state matching, making them unsuitable for sparse inputs. In contrast, while feed-forward models [4, 13, 20, 37] offer fast inference, they are typically trained on limited datasets (e.g., PartNet-Mobility [63], restricting them to a few categories and limiting generalization to diverse, unseen objects).

To address this gap, we introduce the Articulated Reconstruction Transformer (ART), a category-agnostic, feed-forward model that reconstructs complete articulated 3D objects from sparse, multi-state RGB images. Our key insight is that articulated objects can be effectively represented as a collection of rigid parts, with articulation defining their kinematic relationships. Accordingly, ART formulates the reconstruction of articulated objects as a part-based prediction task.

Inspired by the success of large-scale static reconstruction models [16], ART adopts a transformer architecture that maps sparse image inputs to a set of learnable part slots, each trained to capture one object part. From each slot, ART jointly decodes a unified part representation—3D geometry, texture, and explicit articulation parameters (e.g., motion type, axis, and pivot). Training with per-part supervision on a large-scale, diverse dataset yields a transferable prior, allowing ART to function as a single, unified model across categories. The part-based output is physically interpretable and directly exportable to standard simulation formats (e.g., URDF), producing simulation-ready assets.

Through comprehensive experiments, we show that ART significantly outperforms both optimization-based and feed-forward baselines. Our method sets a new state-of-the-art for articulated object reconstruction from sparse im-

age inputs, demonstrating the potential of large-scale, part-based feed-forward models for this challenging task.

In summary, our main contributions are as follows:

- We tackle image-based articulated object reconstruction from sparse-view, multi-state inputs by formulating it as a part-level prediction of geometry, texture, and articulation properties.
- We propose ART, a category-agnostic feed-forward transformer trained on large-scale articulated object datasets, capable of inferring not only per-part geometry/texture, but also kinematically consistent articulation structures.
- We demonstrate that ART significantly outperforms both optimization-based and feed-forward baselines, establishing a new state-of-the-art for holistic articulated object reconstruction.

2. Related Work

2.1. Articulation structure understanding

A large body of work studies articulation understanding across diverse input modalities, including RGB [59] and RGB-D [1, 18, 39] images, point clouds [12, 40, 60], videos [41, 50] and 3D meshes [51]. Within this space, one line of work identifies movable parts to reveal the potential degrees of freedom of the object [20, 50, 56] from the given input. Another complementary line directly estimates articulation parameters—from high-level kinematic graphs to low-level joint directions, pivots, and motion angles [21, 31, 64]—with recent advances adopting generative or diffusion-based formulations to improve robustness and accuracy [26, 37]. However, most prior methods treat articulation in isolation from geometry and appearance: they recover motion structure but do not reconstruct photorealistic shape and texture suitable for simulation or rendering. In contrast, our method jointly predicts articulation structure and part-level geometry/texture, producing a unified, simulation-ready representation.

2.2. Articulated object reconstruction

Beyond understanding alone, another line of research targets *reconstruction* of articulated objects—i.e., recovering geometry, texture, and articulation structure from observations. Methods typically fall into two categories. The first is the *per-object optimization* methods. Many approaches formulate reconstruction as inverse rendering with neural radiance fields [45] or 3D gaussian splatting [24]. These approaches iteratively optimize geometry and appearance over multi-view, multi-state sequences while inferring articulation [9, 36, 42, 46, 57, 62]. Although often yielding high fidelity, these pipelines require very dense viewpoints or carefully staged sequences, entail lengthy per-instance optimization, and rely on fragile cross-state correspondences, making them sensitive to occlusion and initialization.

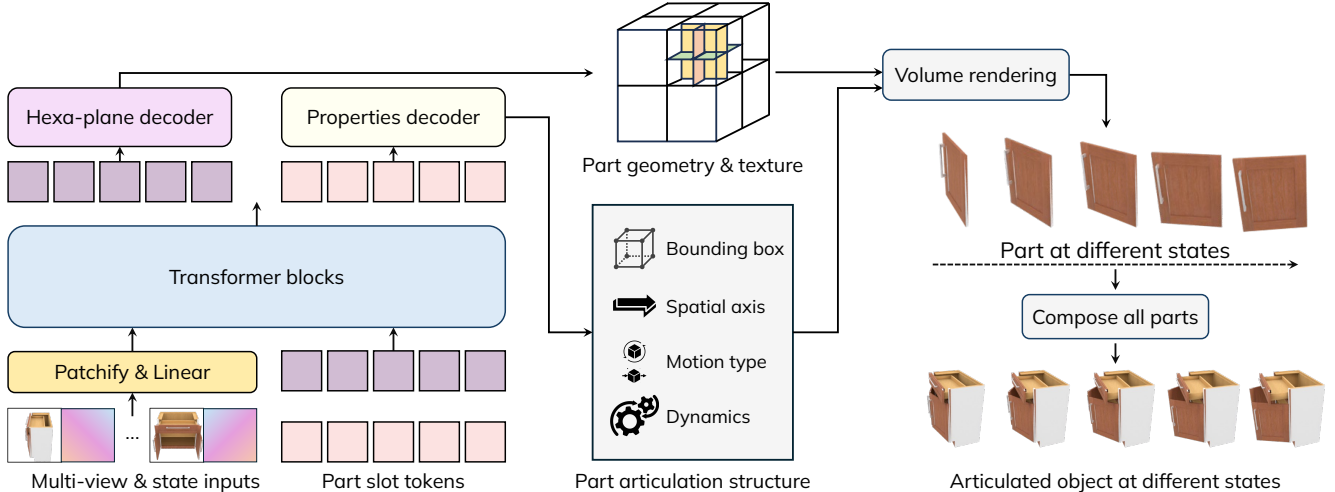


Figure 2. **Model Architecture** of ART. Multi-view, multi-state image inputs with known camera poses are tokenized and processed by a transformer alongside learnable part slot tokens. Two separate decoders then predict each part’s geometry/texture and the articulation structure, and SDF volume rendering composes these components to render and reconstruct the articulated object at different states.

The alternative is the *feed-forward prediction* models. To improve scalability, these models infer articulation without test-time optimization. For instance, SINGAPO [37] predicts a kinematic graph and retrieves parts to assemble full assets, while other approaches [4, 6, 15, 44] reconstruct articulated objects directly from single- or multi-stage image inputs. Nonetheless, data scarcity often restricts such models to a small set of categories. More recent work [14] leverages strong pre-trained generative prior to expand generalization, but is designed for multi-stage image generation and is typically combined with an optimization-based refinement for final reconstruction. By contrast, our approach differs remains purely feed-forward at inference from sparse RGB inputs, directly predicting a part-based 3D geometry/texture and explicit articulation parameters, and is trained on a substantially larger and more diverse dataset.

2.3. Feed-forward 3D reconstruction

The availability of large-scale 3D datasets [7, 8] has enabled powerful feed-forward reconstruction systems that combine scalable transformers architectures with differentiable rendering supervision [16, 22, 29, 34, 65]. Recent work extends this approach to part-based reconstruction/generation [3, 35], indicating that structured outputs can be produced in a single forward pass. Building on these insights, ART exploits a transformer backbone for articulated 3D reconstruction: it decomposes an object into consistent components and predicts part-level geometry and texture along with explicit articulation parameters (motion type, axis, pivot/limits) for each dynamic part, yielding a unified representation that is both photorealistic and kinematically interpretable.

3. Articulated Reconstruction Transformer

3.1. Problem Formulation

The input is a multi-view and multi-state image set $\mathcal{I} = \{\mathcal{I}_{v,t} \in \mathbb{R}^{H \times W \times 3} \mid v \in [1, V], t \in [1, T]\}$, where V is the number of camera views and T stands for the number of articulation stages (states). During training, each view has known intrinsics and extrinsics to cast sampling rays, and the object is normalized to a bounding sphere with radius r .

We aim to reconstruct an articulated object as a set of *parts*. Let P denote the number of parts (including a static base part). For each part $p \in [1, P]$, we predict a unified representation \mathcal{X}_p :

$$\mathcal{X}_p = \{\mathcal{T}_p, \mathcal{A}_p\}, \quad \mathcal{A}_p = (\mathbf{B}_p, \mathbf{C}_p, \mathbf{D}_p, \mathbf{O}_p, \mathbf{S}_p). \quad (1)$$

where \mathcal{T}_p encodes geometry/texture, and \mathcal{A}_p denotes articulation parameters. The detailed definitions of these articulation parameters are given below (omitting the part index for simplicity):

- $\mathbf{B} = (\mathbf{B}_{\text{center}}, \mathbf{B}_{\text{size}}) \in \mathbb{R}^6$ denotes an axis-aligned bounding box in the canonical object frame, defined by its center position and the side lengths along each axis.
- $\mathbf{C} \in \{\text{static, prismatic, revolute}\}$ is the motion type. The base part is static by definition; all other parts are either prismatic (translational) or revolute (rotational).
- $\mathbf{D} \in \mathbb{S}^2$ is the joint axis direction in the canonical object frame, represented as a unit-length 3D vector.
- $\mathbf{O} \in \mathbb{R}^3$ is a point on the joint axis in canonical frame. For a revolute joint, it is essentially the hinge pivot. For a prismatic joint, \mathbf{O} is defined and predicted but is not required in inference since the direction of the axis suffices to interpret the part motion.
- $\mathbf{S} \in \mathbb{R}^T$ represents the normalized motion value (“dynamics”) for each input stage: angles in radians for revo-

lute motion, or translations in object-scale units for prismatic motion. During training, \mathcal{S} aligns rendered images to the observed state t ; at inference, it can be used to control the articulated configuration.

For \mathcal{T}_p , following transformer-based reconstruction models [16, 34, 61], we represent each part’s geometry and texture with a hexa-plane [2, 34] parameterization (details in Section 3.2 and Section 3.3). Concretely, we denote:

$$\mathcal{T}_p = \{\mathbf{T}_{p,k}^{(+,-)} \mid k \in \{xy, yz, xz\}\}, \quad (2)$$

where each predicted plane \mathbf{T} stores the features that will be queried during volume rendering process. All parts in the model are predicted in a shared canonical object frame (the *rest state*). Given an articulation configuration q (rotational angle or translation), part p can be posed by a rigid transform $T_p(q; \mathbf{C}_p, \mathbf{D}_p, \mathbf{O}_p)$ that (i) rotates about axis $(\mathbf{D}_p, \mathbf{O}_p)$ when $\mathbf{C}_p = \text{revolute}$, or (ii) translates along direction \mathbf{D}_p when $\mathbf{C}_p = \text{prismatic}$. The base part uses the identity transform. During training, we enforce consistency between the posed prediction and the observed stage t using the predicted per-state dynamics.

We allocate P_0 learnable part slots in the network and predict $P \leq P_0$ active parts. At inference time, similar to previous methods [36, 42, 62], we assume the part count is known. In practice, existing available VLMs [5, 17, 67] can also provide accurate part-count estimates.

An important design choice is to use a canonical *rest state* frame for articulation parameters. This rest state is a predefined pose configuration for each object instance (e.g., all drawers closed, microwave shut), set during the data construction. In contrast to parameterizing motion relative to the first observed frame—which is sequence-dependent and thus inconsistent—this canonicalization ensures identical ground truth across different sequences of the same object for both part bounding boxes and underlying geometry/texture. This results in more stable training and substantially faster convergence, an important benefit given the limited availability of articulated 3D data.

3.2. Model Architecture

Given the multi-view, multi-stage inputs described above, ART maps images into a shared token space, routes these token features to a fixed set of learnable part slots, and decodes for each part both the geometry/texture plane representations and articulation structure parameters.

Encoding image tokens. Each image $I_{v,t}$ is first divided into non-overlapping 8×8 patches and projected by a small MLP [49] into a sequence of image tokens. To disambiguate tokens across views and articulation stages, we augment them with three types of side information.

Stage embeddings. We add a learnable embedding e_t to every token originating from stage t , enabling the network to separate information from different articulation states.



Procedural dataset

StorageFurniture dataset

Figure 3. **Articulated object data samples** from two of our data collections: Procedural (left two columns) and StorageFurniture (right two columns). For each object, we show two states from one dynamic sequence under the same viewpoint.

Viewpoint information. For each input image we compute the Plücker ray representation [19, 55] using known camera intrinsics and extrinsics, denoted as $(\mathbf{v}, \mathbf{v} \times \mathbf{o})$, where \mathbf{v} is the unit ray direction and \mathbf{o} is the camera origin.

High-level semantics. We concatenate features from a pre-trained DINOv2 encoder [47]; the image inputs to DINOv2 are resized to account for different patch sizes. These features provide rich semantics [11, 33, 66, 70], which is especially helpful under sparse views and limited training data.

After the concatenation of the above features with patch embeddings, we obtain the final token sequence from the input images, serving as the conditional input to the transformer layers.

Learnable part slots. We introduce P_0 learnable part slots in the model. Each slot has a set of tokens trained to predict part information given the articulated object observation inputs. One slot is reserved for the base part, and the remaining slots model movable parts. During training, slot tokens are updated via attention with image tokens [58]. With part number P given, we simply keep the first P slots to obtain the final prediction in both training and inference.

Transformer layers. Inputs to the transformer are part tokens and image tokens (Figure 2). Stacked layers update

part tokens while continuously progressing on information in image tokens.

We employ two complementary layer types. In the *self-attention* layers, we concatenate image and part tokens and apply a single attention operation over the entire set. This promotes global context sharing across views, stages, and parts, which is beneficial for long-range correspondence and enforcing inter-part consistency. In the *cross-attention* layers, image tokens act as queries and part tokens provide keys and values, explicitly routing visual information to a compact set of part slots and mitigating inter-part interference—an important factor to further improve performance.

Most LRM-style models [16, 34, 61, 65] adopt only self-attention layer. In our setting this led to slower training and unstable specialization of part slots, so we replace $\sim 75\%$ of layers with cross-attention, for two reasons: (1) *Token efficiency*. Multi-stage, multi-part inputs introduce many more tokens than single-object settings; cross-attention uses a smaller effective attention window and is more efficient. (2) *Convergence and accuracy*. Interleaving cross-attention accelerates convergence and improves final quality by encouraging distinct roles for image vs. part tokens, focusing the model on their interactions, and enabling it to learn stronger reconstruction priors.

Decoding part properties. As shown in Fig. 2, we split the final part tokens into two branches and use separate MLP heads to predict (i) the hexa-plane representation \mathcal{T}_p for geometry/texture and (ii) the articulation vector $\hat{\mathcal{A}}_p \in \mathbb{R}^{14+T}$. We partition $\hat{\mathcal{A}}_p$ along the channel axis into $(\hat{\mathbf{B}}_p, \hat{\mathbf{C}}_p, \hat{\mathbf{D}}_p, \hat{\mathbf{O}}_p, \hat{\mathbf{S}}_p)$ and remap these raw outputs to the final properties:

$$\mathbf{B}_p = (2r \cdot \psi(\hat{\mathbf{B}}_{p,\text{center}}) - r, 2r \cdot \psi(\hat{\mathbf{B}}_{p,\text{size}})) \quad (3)$$

$$\mathbf{D}_p = \hat{\mathbf{D}}_p / \|\hat{\mathbf{D}}_p\|_2 \quad (4)$$

$$\mathbf{O}_p = 2r \cdot \psi(\hat{\mathbf{O}}_p) - r \quad (5)$$

$$\mathbf{S}_p = 2 \cdot \psi(\hat{\mathbf{S}}_p) - 1, \quad (6)$$

where $\psi(\cdot)$ is the sigmoid and r is the radius of the normalized bounding sphere. For the motion type, $\hat{\mathbf{C}}_p$ produces two logits (i.e., prismatic vs. revolute), which is determined by `softmax` during training and `argmax` at inference. By convention, the first part slot is reserved for the static base with motion type fixed to static. The two-way classification is only applied to the remaining $P_0 - 1$ movable slots.

3.3. Rendering Articulated Object

As shown in Fig. 2, during each training iteration we render the articulated object both per part and as a composite of all parts to generate the final supervision image. Our renderer follows signed-distance-function (SDF) volume rendering [69] to allow the model to learn both underlying geometry and appearance. Specifically for each dynamic part,

we transform sampling rays into the object coordinate space of the corresponding state. Implementation details for rendering static and dynamic parts and for composing *all* parts via volumetric rendering are provided in Sec B in the supplementary.

3.4. Training Scheme

Training objectives. Our loss combines *rendering* objectives with *direct* supervision on articulation parameters. — All rendering losses are computed on *per-part* renderings rather than the final composite. Empirically, supervising only the composite image hinders learning in occluded regions and biases geometry/texture near part boundaries. For each part p , view v , and stage t , we apply mean-squared error \mathcal{L}_2 on RGB and masks, and a perceptual loss $\mathcal{L}_{\text{LPIPS}}$ [71] on RGB. For articulation parameters, we use cross-entropy \mathcal{L}_{CE} for motion-type classification C_p and MSE for the remaining parameters B_p, D_p, O_p, S_p .

Pre-training stage. Articulated-object datasets are inherently less diverse and scalable than static 3D corpora [7, 8]. To learn a strong prior over geometry, texture, and part decomposition, we introduce an optional pre-training stage. For pre-training, we curate 130k static 3D objects with part decomposition from a collection of 3D-artist generated assets that we licensed for AI training from a commercial source. And we further filter the assets whose native glTF/GLB hierarchy contains at most P_0 parts (as defined by their mesh-based composition). During pre-training, the model is optimized only with rendering losses ($\mathcal{L}_2, \mathcal{L}_{\text{LPIPS}}$) and MSE on part bounding-box centers/sizes; articulation parameters (motion type, axis, dynamics) are not applicable. Empirically, this stage consistently boosts downstream performance across all metrics.

Coarse-to-fine articulation training stage. Following static pre-training, we fine-tune ART on our articulated dataset using a coarse-to-fine curriculum designed to gradually increase the network ability. First, we gradually sharpen the rendered surfaces by linearly increasing the reciprocal of the standard derivation of the SDF [34]. Second, we employ resolution annealing: training begins at a 128×128 resolution for rendering-based objectives, and the supervision resolution is later increased to 256×256 to encourage finer geometric predictions.

3.5. Articulated Object Datasets

Our model is trained with multi-view, multi-state RGB images from dynamic articulated sequences generated from a large and diverse collection of articulated 3D object assets. Complete details of the data pipeline (asset collection/construction and per-asset sequence generation) and statistics are given in the supplementary material. Below we briefly introduce the names and characteristics of each collection used. In total, we aggregate articulated objects

from three primary sources:

PartNet-mobility. We utilize several common indoor categories from the PartNet-Mobility benchmark [63], including buckets, microwaves, and a range of furniture classes.

Procedural dataset. To enhance diversity and realism, we generated a new articulated-object dataset using a procedural generation method motivated by [23]. It includes 2,000 high-quality articulated models across six categories, with rich variation in shape and texture.

StorageFurniture dataset. We specifically focus on this category because it is ubiquitous in real-world environments and widely used as a training domain in prior work [4, 20, 37]. Our collection is procedurally generated from the PartNet-Mobility storage-furniture category: using a compositional assembly system, we replace parts of source objects with alternative geometries from other assets, yielding a large and diverse set of realistic storage-furniture models.

Figure 3 shows random samples from the latter two collections. For every articulated asset described above, we automatically synthesize multiple articulated sequences, which are used for training and evaluation.

4. Experiments

Evaluation datasets. We evaluate on two distinct test sets.

StorageFurniture Test Set: For comparisons to feed-forward methods, we use the held-out split of our large-scale StorageFurniture dataset, comprising 631 objects, enabling robust large-scale evaluation.

PartNet-mobility Test Set: For comparisons to optimization-based methods, we use a held-out set of 10 articulated sequences from PartNet-mobility [63], aligning with the evaluation protocol of previous optimization methods on small sets due to their long inference times.

Baseline methods. We compare against two classes of methods. For feed-forward models, we select URDFormer [4] and SINGAPO [37], two recent state-of-the-art methods for articulated reconstruction with a particular focus on part-level articulation prediction for storage-furniture objects. For optimization-based methods, we include PARIS [36], DTA [62] and ArtGS [42], which recover geometry, texture, and articulation structure from two-state, multi-view inputs.

Evaluation metrics. To provide a holistic evaluation, we assess geometry, texture, and part-level accuracy.

- Image-level metrics: We report PSNR and LPIPS [71] on renderings from 8 novel test viewpoints.
- Geometry-level metrics: We compute Chamfer Distance (CD) and F-Score between the predicted and ground-truth meshes.
- Part-level metrics: Following the evaluation protocol in Liu et al. [37], we match predicted parts and ground-truth

	$d_{gIoU} \downarrow$	$d_{cDist} \downarrow$	CD \downarrow
URDFormer [4]	1.0710	0.1622	0.0536
SINGAPO [37]	<u>0.8306</u>	<u>0.0947</u>	<u>0.0059</u>
ART (Ours)	0.4717	0.0538	0.0019

Table 1. **Quantitative comparison** on the StorageFurniture test set (all metrics: lower is better). Our method surpasses baselines in both part-level prediction and holistic geometry reconstruction.

	PSNR \uparrow	LPIPS \downarrow	CD \downarrow	F-Score \uparrow
PARIS [36]	<u>22.851</u>	0.183	0.023	0.486
DTA* [62]	21.587	<u>0.165</u>	0.008	0.821
ArtGS [42]	22.352	0.176	0.016	0.520
ART (Ours)	27.059	0.049	<u>0.009</u>	<u>0.762</u>

Table 2. **Quantitative comparison** of the reconstruction quality on the PartNet-mobility [63] test set. * indicates requiring additional depth inputs.

parts via Hungarian algorithm, then report (i) part-level prediction accuracy using distance of generalized Intersection over Union (1-gIoU) [52] and (ii) Euclidean distance between part centroids. For both, lower is better.

Implementation details. We train two versions of ART: a multi-view model ($V=4$) for comparison with optimization methods and a monocular model ($V=1$) for fair comparison with feed-forward baselines. We fix $T=2$ (“start & end”) states. All training and inference images are resized to $H=W=128$ to balance compute.

4.1. Results

Comparison with feed-forward baselines. We first compare the monocular version of ART to URDFormer and SINGAPO. Table 1 reports quantitative results: ART outperforms both baselines by a large margin across all metrics. In particular, it achieves substantially better part-level prediction, as evidenced by lower d_{gIoU} and d_{cDist} —and more accurate overall geometry, as indicated by a reduced Chamfer Distance. These gains are also evident qualitatively in Figure 5, which contrasts our predicted part bounding boxes with SINGAPO and highlights ART’s enhanced ability to jointly reconstruct part structure and geometry.

Comparison with optimization-based baselines. Next, we evaluate the multi-view ART against per-object optimization baselines on the PartNet-mobility test set, using sparse inputs (4 views across 2 states) for all methods; DTA additionally receives depth maps, as required. As shown in Table 2, ART achieves state-of-the-art results on image-level metrics, outperforming all optimization-based competitors by a clear margin. While DTA reports comparable geometry metrics, this is expected given its depth supervision; nonetheless, its lower PSNR and higher LPIPS scores indicate poor appearance recovery. This disparity is rooted in sparsity: optimization approaches rely on dense cross-

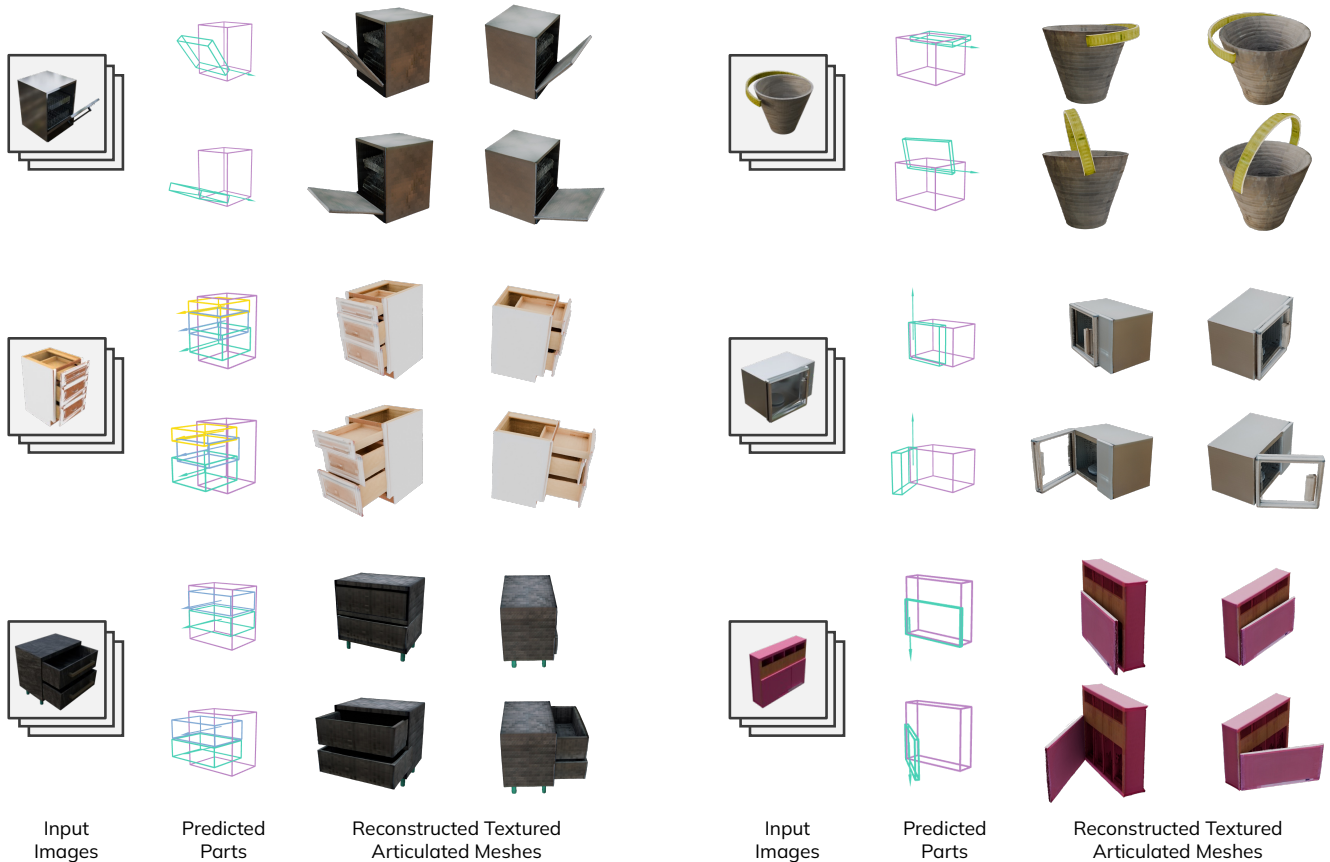


Figure 4. **Qualitative results.** “Input Images” shows the multi-state inputs provided to the model (visualized by one image). “Predicted Parts” visualizes the predicted per-part bounding boxes in the start and end states; “Reconstructed Textured Articulated Meshes” displays the textured 3D reconstructions for the same two articulated states.




Figure 5. **Qualitative comparison with feed-forward baseline** of predicted part bounding box and articulation structure.

state correspondences, which are fragile and difficult to establish from few views. As shown by the failure example in Figure 6, due to the lack of robust correspondences, ArtGS yields fragmented, noisy geometry, whereas ART —bene-



Figure 6. **Qualitative comparison with optimization baseline.** With the same sparse inputs, our method reconstructs complete, high-fidelity textured meshes in both states. ArtGS [42] lacks reliable correspondences under sparsity, leading to fragmented and inaccurate results.

fitting from a strong learned prior—reconstructs coherent, high-fidelity textured meshes.

Qualitative results. We present qualitative results in Figure 1 and Figure 4. Our core idea is to cast articulated object reconstruction as a part-based prediction problem, decomposing objects into a set of rigid parts. As shown in Figure 1, ART successfully reconstructs diverse object cat-

	PSNR \uparrow	LPIPS \downarrow	CD \downarrow	F-Score \uparrow	d_{gIoU} \downarrow	d_{cDist} \downarrow
monocular view	25.961	0.037	0.017	0.591	0.731	0.086
w/o rest-state	23.587	0.081	0.015	<u>0.672</u>	0.878	0.113
w/o defined part order	22.588	0.060	0.015	0.564	1.118	0.208
w/o part rendering loss	24.465	0.047	0.022	0.512	0.681	<u>0.075</u>
w/o pre-train	<u>27.495</u>	<u>0.034</u>	<u>0.013</u>	0.641	<u>0.665</u>	0.082
ART (Ours)	28.678	0.030	0.008	0.716	0.629	0.062

Table 3. **Ablation study** of the components in proposed ART.




Figure 7. **Real-world images results.**

egories and yields clear part decompositions, visualized by rendering each predicted part with a consistent color.

Figure 4 provides a closer look at our method’s outputs, showing reconstructions for the start/end states provided in the image inputs. The *Predicted Parts* column visualizes the inferred structure by overlaying per-part bounding boxes for both states, along with recovered articulation structures (e.g., axes for movable parts). Together, these visualizations demonstrate that ART reconstructs high-fidelity geometry and texture while accurately recovering the articulation structure that drives object motion.

Real-world images results. We further present a real-world example in Fig. 7. The capture uses only approximate camera poses and no background masking. Despite no real-image training, ART recovers the correct articulation structure with plausible geometry and texture.

4.2. Ablation Study

We conduct a detailed ablation study (Table 3) to validate our key design choices. All model variants are evaluated on the overall held-out test set of PartNet-mobility [63], which includes more than 130 multi-view, multi-state sequences. Our reference is the *base model w/o pre-train*, which is the multi-view ($V=4$) ART model trained from scratch on only the articulated object dataset. We then measure the impact of adding the pre-training stage (last row) and of removing key components (first four rows) relative to this base model.

Pre-training stage helps. By adding the pre-training stage, the full ART shows a substantial improvement over the version without pre-training, with gains across geometry, texture, and articulation metrics. This confirms that a robust prior on geometry and part decomposition learned from

large-scale 3D data is highly beneficial for part-based articulated reconstruction.

Monocular vs. multi-view inputs. We compare the base multi-view model with a variant trained using only a single view as input. This change produces a clear drop in quality across all metrics, as expected, confirming that multiple views are crucial for resolving ambiguities and achieving high-fidelity, geometrically accurate reconstruction.

Rest-state formulation. Removing the rest-state formulation and predicting parts relative to the first observed frame leads to a severe quality drop (PSNR $27.495 \rightarrow 23.587$) and higher part-level errors (d_{cDist} $0.082 \rightarrow 0.113$). The decline is unsurprising: because the test set contains multiple sequences of the same object, the first-frame scheme induces a pose-dependent “canonical” per sequence (e.g., one starting closed, another open), effectively treating the same object as different identities. This ambiguity hampers learning and underscores the benefit of a consistent, predefined rest state.

Defined Part Order. Without forcing a pre-defined part ordering in dataset construction, performance degrades most severely across the board (PSNR $27.495 \rightarrow 22.588$, d_{cDist} $0.082 \rightarrow 0.208$). This is expected: Lacking a consistent ordering forces the network to learn both reconstruction and part matching, which is especially difficult for objects with many similar dynamic parts (e.g., cabinets with many drawers). We observe slot collapse—multiple slots predicting parts in the same location—directly reflected in the part-level metrics. This confirms the benefit of enforcing a consistent part ordering.

Per-part vs. composite rendering losses. Finally, eliminating per-part rendering losses and supervising only the composited image consistently harms performance. Per-part supervision is essential for learning occluded content; otherwise, the result is smeared geometry and inferior texture quality.

5. Discussions and Conclusion

Limitations. Our method assumes a known part count for the target object and relies on pre-calibrated camera poses. Future work should include learning a pose-free variant (self-calibrated cameras) with larger datasets and integrating part-count estimation directly into the model.

Conclusion. In this work, we propose ART, a feed-forward model to reconstruct complete 3D articulated objects from sparse, multi-state images. By casting reconstruction as a part-based prediction problem, ART jointly decodes geometry, texture, and articulation structure for each part. Experiments across a broad range of articulated objects demonstrate that ART consistently outperforms strong feed-forward and optimization-based baselines.

Acknowledgments. We thank Yawar Siddiqui, Ruocheng Wang, and Qiao Gu for the comments and fruitful discussions, Ka Chen, David Clabaugh, and Samir Aroudj for the help on dataset collecting and constructions, and Yanjie Ze and Pengyu Mo for the robot simulation results. This work is in part supported by NSF RI #2338203 and ONR MURI N00014-22-1-2740.

References

- [1] Ben Abbatematteo, Stefanie Tellex, and George Konidaris. Learning to generalize kinematic models to novel objects. In *Proceedings of the 3rd Conference on Robot Learning*, 2019. 2
- [2] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 130–141, 2023. 4
- [3] Minghao Chen, Jianyuan Wang, Roman Shapovalov, Tom Monnier, Hyunyoung Jung, Dilin Wang, Rakesh Ranjan, Iro Laina, and Andrea Vedaldi. Autopartgen: Autogressive 3d part generation and discovery. *arXiv preprint arXiv:2507.13346*, 2025. 3
- [4] Zoey Chen, Aaron Walsman, Marius Memmel, Kaichun Mo, Alex Fang, Karthikeya Vemuri, Alan Wu, Dieter Fox, and Abhishek Gupta. Urdformer: A pipeline for constructing articulated simulation environments from real-world images. *arXiv preprint arXiv:2405.11656*, 2024. 2, 3, 6, S2
- [5] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasapat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blisstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 4
- [6] Tianyuan Dai, Josiah Wong, Yunfan Jiang, Chen Wang, Cem Gokmen, Ruohan Zhang, Jiajun Wu, and Li Fei-Fei. Automated creation of digital cousins for robust policy learning. *arXiv preprint arXiv:2410.07408*, 2024. 3
- [7] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36:35799–35813, 2023. 3, 5
- [8] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13142–13153, 2023. 3, 5
- [9] Jianning Deng, Kartic Subr, and Hakan Bilen. Articulate your nerf: Unsupervised articulated object modeling via conditional view synthesis. *Advances in Neural Information Processing Systems*, 37:119717–119741, 2024. 1, 2
- [10] Yufan Deng, Yuhao Zhang, Chen Geng, Shangzhe Wu, and Jiajun Wu. Anymate: A dataset and baselines for learning 3d object rigging. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–10, 2025. 1
- [11] Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3d awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806, 2024. 4
- [12] Lian Fu, Ryoichi Ishikawa, Yoshihiro Sato, and Takeshi Oishi. Capt: Category-level articulation estimation from a single point cloud using transformer. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 751–757. IEEE, 2024. 2
- [13] Daoyi Gao, Yawar Siddiqui, Lei Li, and Angela Dai. Meshart: Generating articulated meshes with structure-guided transformers. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 618–627, 2025. 2
- [14] Mingju Gao, Yike Pan, Huan-ang Gao, Zongzheng Zhang, Wenyi Li, Hao Dong, Hao Tang, Li Yi, and Hao Zhao. Partrm: Modeling part-level dynamics with large cross-state reconstruction model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7004–7014, 2025. 3
- [15] Nick Heppert, Muhammad Zubair Irshad, Sergey Zakharov, Katherine Liu, Rares Andrei Ambrus, Jeannette Bohg, Abhinav Valada, and Thomas Kollar. Carto: Category and joint agnostic reconstruction of articulated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21201–21210, 2023. 2, 3
- [16] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400*, 2023. 1, 2, 3, 4, 5
- [17] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 4
- [18] 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. PMLR, 2022. 2
- [19] Yan-Bin Jia. Plücker coordinates for lines in the space. *Problem Solver Techniques for Applied Computer Science, COMS-477/577 Course Handout*, 3, 2020. 4
- [20] Hanxiao Jiang, Yongsan Mao, Manolis Savva, and Angel X Chang. Opd: Single-view 3d openable part detection. In *European Conference on Computer Vision*, pages 410–426. Springer, 2022. 2, 6, S2
- [21] Zhenyu Jiang, Cheng-Chun Hsu, and Yuke Zhu. Ditto: Building digital twins of articulated objects from interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5616–5626, 2022. 2
- [22] Haian Jin, Hanwen Jiang, Hao Tan, Kai Zhang, Sai Bi, Tianyuan Zhang, Fujun Luan, Noah Snavely, and Zexiang Xu. Lvsrn: A large view synthesis model with minimal 3d inductive bias. *arXiv preprint arXiv:2410.17242*, 2024. 3
- [23] Abhishek Joshi, Beining Han, Jack Nugent, Max Gonzalez Saez-Diez, Yiming Zuo, Jonathan Liu, Hongyu Wen, Stamatios Alexandropoulos, Karhan Kayan, Anna Calveri, Tao Sun,

- Gaowen Liu, Yi Shao, Alexander Raistrick, and Jia Deng. Procedural generation of articulated simulation-ready assets, 2025. 6, S2
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2
- [25] Jeonghwan Kim, Jisoo Kim, Jeonghyeon Na, and Hanbyul Joo. Parahome: Parameterizing everyday home activities towards 3d generative modeling of human-object interactions. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1816–1828, 2025. 1
- [26] Jiahui Lei, Congyue Deng, William B Shen, Leonidas J Guibas, and Kostas Daniilidis. Nap: Neural 3d articulated object prior. *Advances in Neural Information Processing Systems*, 36:31878–31894, 2023. 2
- [27] Chengshu Li, Fei Xia, Roberto Martín-Martín, Michael Lingelbach, Sanjana Srivastava, Bokui Shen, Kent Vainio, Cem Gokmen, Gokul Dharan, Tanish Jain, et al. igibson 2.0: Object-centric simulation for robot learning of everyday household tasks. *arXiv preprint arXiv:2108.03272*, 2021. 1
- [28] Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Wensi Ai, Benjamin Martinez, et al. Behavior-1k: A human-centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation. *arXiv preprint arXiv:2403.09227*, 2024. 1
- [29] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. *arXiv preprint arXiv:2311.06214*, 2023. 3
- [30] Ruilong Li, Matthew Tancik, and Angjoo Kanazawa. Nerfacc: A general nerf acceleration toolbox. *arXiv preprint arXiv:2210.04847*, 2022. S1
- [31] Xiaolong Li, He Wang, Li Yi, Leonidas J Guibas, A Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3706–3715, 2020. 2
- [32] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023. S1
- [33] Zizhang Li, Dor Litvak, Ruining Li, Yunzhi Zhang, Tomas Jakab, Christian Rupprecht, Shangzhe Wu, Andrea Vedaldi, and Jiajun Wu. Learning the 3d fauna of the web. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9752–9762, 2024. 4
- [34] Zhengqin Li, Dilin Wang, Ka Chen, Zhaoyang Lv, Thu Nguyen-Phuoc, Milim Lee, Jia-Bin Huang, Lei Xiao, Yufeng Zhu, Carl S Marshall, et al. Lirm: Large inverse rendering model for progressive reconstruction of shape, materials and view-dependent radiance fields. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 505–517, 2025. 1, 3, 4, 5, S1
- [35] Yuchen Lin, Chenguo Lin, Panwang Pan, Honglei Yan, Yiqiang Feng, Yadong Mu, and Katerina Fragkiadaki. Partcrafter: Structured 3d mesh generation via compositional latent diffusion transformers. *arXiv preprint arXiv:2506.05573*, 2025. 3
- [36] Jiayi Liu, Ali Mahdavi-Amiri, and Manolis Savva. Paris: Part-level reconstruction and motion analysis for articulated objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 352–363, 2023. 2, 4, 6
- [37] Jiayi Liu, Denys Iliash, Angel X Chang, Manolis Savva, and Ali Mahdavi-Amiri. Singapo: Single image controlled generation of articulated parts in objects. *arXiv preprint arXiv:2410.16499*, 2024. 2, 3, 6, S2
- [38] Jiayi Liu, Manolis Savva, and Ali Mahdavi-Amiri. Survey on modeling of human-made articulated objects. In *Computer Graphics Forum*, page e70092. Wiley Online Library, 2025. 1
- [39] Liu Liu, Han Xue, Wenqiang Xu, Haoyuan Fu, and Cewu Lu. Toward real-world category-level articulation pose estimation. *IEEE Transactions on Image Processing*, 31:1072–1083, 2022. 2
- [40] Liu Liu, Jianming Du, Hao Wu, Xun Yang, Zhenguang Liu, Richang Hong, and Meng Wang. Category-level articulated object 9d pose estimation via reinforcement learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 728–736, 2023. 2
- [41] Qihao Liu, Weichao Qiu, Weiyao Wang, Gregory D Hager, and Alan L Yuille. Nothing but geometric constraints: A model-free method for articulated object pose estimation. *arXiv preprint arXiv:2012.00088*, 2020. 2
- [42] 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 *The Thirteenth International Conference on Learning Representations*, 2025. 2, 4, 6, 7
- [43] Ying-Tian Liu, Yuan-Chen Guo, Vikram Voleti, Ruizhi Shao, Chia-Hao Chen, Guan Luo, Zixin Zou, Chen Wang, Christian Laforte, Yan-Pei Cao, et al. Threestudio: A modular framework for diffusion-guided 3d generation. *cg.cs.tsinghua.edu.cn*, 2023. 1
- [44] Zhao Mandi, Yijia Weng, Dominik Bauer, and Shuran Song. Real2code: Reconstruct articulated objects via code generation. *arXiv preprint arXiv:2406.08474*, 2024. 3
- [45] 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. *Communications of the ACM*, 65(1):99–106, 2021. 2, S1
- [46] Jiteng Mu, Weichao Qiu, Adam Kortylewski, Alan Yuille, Nuno Vasconcelos, and Xiaolong Wang. A-sdf: Learning disentangled signed distance functions for articulated shape representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13001–13011, 2021. 2
- [47] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al.

- Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 4
- [48] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 1
- [49] Marius-Constantin Popescu, Valentina E Balas, Liliana Perescu-Popescu, and Nikos Mastorakis. Multilayer perceptron and neural networks. *WSEAS Transactions on Circuits and Systems*, 8(7):579–588, 2009. 4
- [50] Shengyi Qian, Linyi Jin, Chris Rockwell, Siyi Chen, and David F Fouhey. Understanding 3d object articulation in internet videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1599–1609, 2022. 2
- [51] Xiaowen Qiu, Jincheng Yang, Yian Wang, Zhehuan Chen, Yufei Wang, Tsun-Hsuan Wang, Zhou Xian, and Chuang Gan. Articulate anymesh: Open-vocabulary 3d articulated objects modeling. *arXiv preprint arXiv:2502.02590*, 2025. 2
- [52] Hamid Rezaatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. 6
- [53] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Claudia Pérez-D’Arpino, Shyamal Buch, Sanjana Srivastava, Lyne Tchapmi, et al. igibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7520–7527. IEEE, 2021. 1
- [54] Yawar Siddiqui, Tom Monnier, Filippos Kokkinos, Mahendra Kariya, Yanir Kleiman, Emilien Garreau, Oran Gafni, Natalia Neverova, Andrea Vedaldi, Roman Shapovalov, et al. Meta 3d assetgen: Text-to-mesh generation with high-quality geometry, texture, and pbr materials. *Advances in Neural Information Processing Systems*, 37:9532–9564, 2024. 1
- [55] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34:19313–19325, 2021. 4
- [56] Xiaohao Sun, Hanxiao Jiang, Manolis Savva, and Angel Chang. Opdmulti: Openable part detection for multiple objects. In *2024 International Conference on 3D Vision (3DV)*, pages 169–178. IEEE, 2024. 2
- [57] Archana Swaminathan, Anubhav Gupta, Kamal Gupta, Shishira R Maiya, Vatsal Agarwal, and Abhinav Shrivastava. Leia: Latent view-invariant embeddings for implicit 3d articulation. In *European Conference on Computer Vision*, pages 210–227. Springer, 2024. 2
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [59] Ruiqi Wang, Akshay Gadi Patil, Fenggen Yu, and Hao Zhang. Active coarse-to-fine segmentation of moveable parts from real images. In *European Conference on Computer Vision*, pages 111–127. Springer, 2024. 2
- [60] Xiaogang Wang, Bin Zhou, Yahao Shi, Xiaowu Chen, Qingping Zhao, and Kai Xu. Shape2motion: Joint analysis of motion parts and attributes from 3d shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8876–8884, 2019. 2
- [61] Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Meshlrn: Large reconstruction model for high-quality meshes. *arXiv preprint arXiv:2404.12385*, 2024. 1, 4, 5, S1
- [62] 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. 2, 4, 6
- [63] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11097–11107, 2020. 2, 6, 8, S2
- [64] Xianghao Xu, Yifan Ruan, Srinath Sridhar, and Daniel Ritchie. Unsupervised kinematic motion detection for part-segmented 3d shape collections. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. 2
- [65] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. In *European Conference on Computer Vision*, pages 1–20. Springer, 2024. 1, 3, 5
- [66] Zhen Xu, Zhengqin Li, Zhao Dong, Xiaowei Zhou, Richard Newcombe, and Zhaoyang Lv. 4dgt: Learning a 4d gaussian transformer using real-world monocular videos. *arXiv preprint arXiv:2506.08015*, 2025. 4
- [67] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 4
- [68] Gengshan Yang, Chaoyang Wang, N Dinesh Reddy, and Deva Ramanan. Reconstructing animatable categories from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16995–17005, 2023. 1
- [69] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in neural information processing systems*, 34:4805–4815, 2021. 5, S1
- [70] Junyi Zhang, Charles Herrmann, Junhwa Hur, Luisa Polania Cabrera, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. *Advances in Neural Information Processing Systems*, 36:45533–45547, 2023. 4

- [71] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [5](#), [6](#)

ART: Articulated Reconstruction Transformer

Supplementary Material

A. More results

We provide additional results in the project [webpage](#).

Video results. We provide both fixed-view and rotational-view renderings of the reconstructed articulated objects, with the dynamic parts moving according to the predicted articulation structure.

Export into simulator. As discussed in the main text, our part-based output can be directly converted to URDF. Combining the exported URDF with each part’s textured mesh yields simulation-ready assets. We showcase several interaction scenes in the MuJoCo simulator—each featuring a humanoid robot and an articulated object—and include the corresponding videos in the project [webpage](#).

Detailed runtime analysis. We provide a quantitative runtime comparison on a single A100 GPU in Table S1. ART completes reconstruction in 0.85s (multi-view) and 0.42s (monocular). This is orders of magnitude faster than optimization baselines with lengthy per-instance processing. This efficiency stems in part from ART’s use of cross-attention layers, which reduces token processing time.

Real-world results Here we add four additional real-world results (Figure S1). While texture reconstruction is inherently limited by the synthetic training domain, ART demonstrates robust articulation-structure understanding across diverse opening configurations.

B. Rendering Articulated Object

At each training iteration, we render the articulated object both per part and in a compositional manner to obtain the final image for supervision. The rendering process queries hexa-plane feature outputs from the model to compute the per-pixel RGB and uses the predicted articulation to correctly place each dynamic part.

Our rendering pipeline follows the signed distance function (SDF) volume rendering [69]. For a camera ray defined by origin \mathbf{o} and unit direction \mathbf{v} , we first intersect the ray with the part’s axis-aligned bounding box \mathbf{B}_p (in canonical space) and sample 3D points along the valid ray segment. Each sampled world-space point $\mathbf{x}_p(\mathbf{o}, \mathbf{v}, \mathbf{B}_p)$ is mapped to the part’s normalized local coordinates $\hat{\mathbf{x}}_p = (x, y, z) \in [-1, 1]^3$, which are then used to query the hexa-plane features:

$$\mathbf{f}_{p,xy} = \begin{cases} \text{Bilinear}(\mathbf{T}_{p,xy+}; x, y), & z \geq 0, \\ \text{Bilinear}(\mathbf{T}_{p,xy-}; x, y), & z < 0. \end{cases} \quad (\text{S1})$$

Here $\mathbf{T}_{p,xy+}$, $\mathbf{T}_{p,xy-}$ are the model’s hexa-plane repre-

Optimization-based Runtime (min)	PARIS	DTA	ArtGS
Feed-forward Runtime (s)	URDFormer	SINGAPO	Ours
	2.37	0.54	0.85 / 0.42

Table S1. **Runtime comparison** between baselines and ART.

sentation output from Eq. 2; $\mathbf{f}_{p,yz}$, $\mathbf{f}_{p,xz}$ can be obtained analogously. These features are concatenated to form the feature vector \mathbf{f}_p for the spatial point $\mathbf{x} + p$. Following the architecture in [34, 61], we use two small MLPs to predict SDF value s_p and RGB color \mathbf{c}_p , respectively. In particular, the SDF value is computed as:

$$s_p = \text{MLP}(\mathbf{f}_p) + \mathbf{s}_{\text{bias}}(\mathbf{x}_p). \quad (\text{S2})$$

where the bias term $\mathbf{s}_{\text{bias}}(\mathbf{x}_p) = \|\mathbf{x}_p\| - 0.1r$ is a prior defined in the part’s local space, to initialize the shape as a sphere to stabilize the training [34, 69]. The SDF value s_p can be converted further to volume density using the Laplace CDF [69]:

$$\sigma_p = \begin{cases} \frac{1}{2} \exp(-\frac{s_p}{\beta}), & s_p \geq 0, \\ 1 - \frac{1}{2} \exp(-\frac{s_p}{\beta}), & s_p < 0. \end{cases} \quad (\text{S3})$$

where β is the standard deviation that controls the sharpness of the underlying surface. As noted in the main text, $\frac{1}{\beta}$ is linearly annealed over training to progressively sharpen surfaces. With σ_p and \mathbf{c}_p evaluated at ray samples, we compute *per-part* color, opacity (mask), and depth via the usual transmittance accumulation [45, 69]. Normals are obtained by numerically differentiating the SDF at each sample [32].

To obtain the composited image over *all parts*, we merge the sampled points from all parts along each ray, sort their orderings by the sampled ray distance, and apply standard alpha compositing to compute the rendered images. We employ Nerfacc [30] to accelerate point sampling, ray distance sorting, and compositing.

Dynamic parts rendering. The equations above describe static per-part rendering in each part’s canonical space (from the model-predicted bounding-box region). For dynamic parts at stage t , instead of physically moving the part’s volume and rebuilding an oriented bounding box, we transform the camera rays into the part’s instantaneous local frame (i.e., move the rays inversely), keeping the axis-aligned bounding box \mathbf{B}_p unchanged.

Recall that \mathbf{C}_p denotes the motion type, \mathbf{D}_p the unit joint axis, \mathbf{O}_p a point on that axis, and $\mathbf{S}_{p,t} \in [-1, 1]$ the normalized motion value at stage t . For a point \mathbf{x} in the part’s canonical space, the stage- t rigid transform $T_{p,t}$ for different motion types can be defined as:



Figure S1. Real-world images results.

Prismatic:

$$T_{p,t}(\mathbf{x}) = \mathbf{x} + (2r) \mathbf{S}_{p,t} \mathbf{D}_p. \quad (\text{S4})$$

Revolute:

$$T_{p,t}(\mathbf{x}) = \mathbf{O}_p + R(\mathbf{D}_p, 2\pi \mathbf{S}_{p,t}) (\mathbf{x} - \mathbf{O}_p), \quad (\text{S5})$$

where $R(\mathbf{d}, \theta)$ denotes the rotation by angle θ (radians) about axis \mathbf{d} .

Now, instead of actually transforming the bounding box (point), we’ll inversely transform each ray (\mathbf{o}, \mathbf{v}) to reach the equivalent rendering results. It’s straightforward that for prismatic parts, after transformation:

$$\hat{\mathbf{o}} = \mathbf{o} - 2r \cdot \mathbf{S}_{p,t} \cdot \mathbf{D}_p, \quad \hat{\mathbf{v}} = \mathbf{v}, \quad (\text{S6})$$

and for revolute parts, similarly:

$$\hat{\mathbf{o}} = \mathbf{o} + \text{Rotate}(\mathbf{D}_p, -2\pi \cdot \mathbf{S}_{p,t})(\mathbf{o} - \mathbf{O}_p) + \mathbf{O}_p, \quad (\text{S7})$$

$$\hat{\mathbf{v}} = \text{Rotate}(\mathbf{D}_p, -2\pi \cdot \mathbf{S}_{p,t})\mathbf{v}. \quad (\text{S8})$$

We then proceed exactly as in the static case: intersect the transformed ray $(\hat{\mathbf{o}}, \hat{\mathbf{v}})$ with bounding box B_p , sample points, query hexa-planes, and composite to obtain the desired renderings. The computed values from the queried features are further assigned to the original ray sampled points, which are in the current world coordinate space, for volume rendering composition.

C. Dataset Construction

A key advantage of ART is its training data: we substantially increase both the quantity and diversity of articulated assets. Prior feed-forward models [20, 37] typically rely on subsets of PartNet-Mobility [63], which offer limited geometric diversity and often unrealistic textures.

In this work, we combine three articulated-object data sources to construct our training dataset, increasing diversity and fidelity in geometry, texture, and articulation complexity. Beyond the basic information in the main text, further details on these sources are provided below.

PartNet-mobility. As previously mentioned, PartNet-Mobility provides common indoor articulated categories—bucket, dishwasher, door, laptop, microwave, oven, refrigerator, and storage furniture. In total, we collect over

300 objects from this dataset. However, many assets have unrealistic textures and low-quality surface geometry, motivating our exploration of more diverse and realistic articulated object sources.

Procedural dataset. To scale both quantity and diversity, and inspired by Infinigen-Sim [23], we adopt procedural generation to author articulated assets in Blender (with TA support). The resulting dataset includes 2,000 high-quality articulated models across six categories: laptop, dishwasher, beverage refrigerator, cabinet with drawers, bucket, and microwave oven. 3D geometry is provided in GLB/OBJ and articulation in URDF. Each category is governed by procedural rules over shape, appearance, and articulation; for each, we generate several hundred variants with randomized shapes, sizes, and materials. In principle, this pipeline can produce an unlimited number of articulated objects.

StorageFurniture dataset. We recognize that the Storage-Furniture category in PartNet-Mobility spans many commonly used articulated assets, a property also noted by prior methods [4, 37] that primarily train on this category. Based on this, we construct a *StorageFurniture* dataset by recombining parts from the PartNet-Mobility storage-furniture class.

For a given object, we use its articulation tree (kinematic structure) to procedurally create new instances via compositional part assembly: original parts are replaced with geometries from other objects (rescaled as needed), followed by UV-map correction and material randomization. This yields a large, realistic, and varied set of assets—over 10,000 articulated models for training.

These three data sources collectively provide a large pool of articulated objects. Using each asset’s articulation definition, we further generate random per-part trajectories, resulting in a large dataset of articulated object sequences.

D. More Implementation Details

During sequence data construction, we explicitly set the order of the parts following a certain rule: the static base part is the first, and the remaining dynamic parts are ordered from low to high, front to back and left to right. This consistent ground-truth ordering greatly improves training stability and convergence speed. We also define the *rest state* at this stage as the configuration where all dynamic parts are “closed.”

The transformer in ART has 32 attention blocks with a 3 : 1 cross-/self-attention ratio, 16 heads, and a 256-dimensional embedding. It outputs 24×24 hexa-plane features that are later upsampled to the spatial resolution of 192×192 . We set the maximum part count to $P_0=8$. Training uses AdamW with $(\beta_1, \beta_2)=(0.9, 0.95)$. The multi-view model trains for 5 days on 64 H100 GPUs; the monoc-

ular model trains for 3 days.

ART is trained with known camera poses as part of the input to distinguish tokens across views. This design follows common LRM-style practice, where camera information serves as implicit geometric cues for multi-view 3D reconstruction. Empirically, providing calibrated poses significantly accelerates training convergence and improves accuracy compared to using fully optimizable view embeddings. While we currently assume calibrated poses, we observe that with results remain reasonable under noisy real-world pose estimates. As discussed in our limitations, developing a pose-free variant by integrating optimizable view embeddings is a meaningful direction for future work to make the system more practical for arbitrary user-captured images.

Inference details Similar to other LRM-style models, ART targets an object-centric reconstruction scenario. Consequently, it requires input images containing a single articulated object accompanied by a corresponding segmentation mask produced by off-the-shelf segmentation models. To prepare the input, we identify the longer side of the object’s bounding box and crop a square region that preserves the full foreground, which is then resized to the input resolution. These standard pre-processing steps are used in both training and inference.