

000 ARTVIP: ARTICULATED DIGITAL ASSETS OF 001 VISUAL REALISM, MODULAR INTERACTION, AND 002 PHYSICAL FIDELITY FOR ROBOT LEARNING 003

004 **Anonymous authors**
005
006

007 Paper under double-blind review
008
009

010 ABSTRACT 011

012 Robot learning increasingly relies on simulation to advance complex ability such
013 as dexterous manipulations and precise interactions, necessitating high-quality
014 digital assets to bridge the sim-to-real gap. However, existing open-source
015 articulated-object datasets for simulation are limited by insufficient visual realism
016 and low physical fidelity, which hinder their utility for training models mastering
017 robotic tasks in real world. To address these challenges, we introduce ArtVIP,
018 a comprehensive open-source dataset comprising high-quality digital-twin articu-
019 lated objects, accompanied by indoor-scene assets. Crafted by professional 3D
020 modelers adhering to unified standards, ArtVIP ensures visual realism through
021 precise geometric meshes and high-resolution textures, while physical fidelity is
022 achieved via fine-tuned dynamic parameters. Meanwhile, the dataset pioneers
023 embedded modular interaction behaviors within assets and pixel-level affordance
024 annotations. Feature-map visualization and optical motion capture are employed
025 to quantitatively demonstrate ArtVIP’s visual and physical fidelity, with its ap-
026 plicability validated across imitation learning and reinforcement learning ex-
027 periments. Provided in USD format with detailed production guidelines, ArtVIP is
028 fully open-source, benefiting the research community and advancing robot learn-
029 ing research.



049 Figure 1: ArtVIP (Articulated Digital Assets of Visual Realism, Modular Interaction, and Physical
050 Fidelity for Robot Learning). The collection spans 9 categories, 37 subcategories, and 992 digital-
051 twin articulated objects. (a) Representative assets across categories and articulation types. (b)
052 High-fidelity physics enables realistic interactions; for example, when pushing an ergonomic
053 chair, its casters rotate accordingly. (c) Six sim-ready scenes in which all objects support real-
world-consistent interactions. (d) Pixel-level annotations and sim2real evaluations.

054
055
056
057
1 INTRODUCTION058
059
060
061
062
063
064
Embodied AI is catalyzing the transformation of robotic systems from constrained laboratory settings (Billard & Kragic, 2019; Spong et al., 2020) to complex, unstructured real-world environments (Brohan et al., 2023b; Zhao et al., 2023; Brohan et al., 2023a). The emergence of large-scale pretrained models (Zhang & Yan, 2023; Kim et al., 2024; Wang et al., 2024b) and novel learning paradigms (Team, 2025; Intelligence, 2025) has ushered in a data-centric era. In this new era, the availability of high-quality data is a critical bottleneck for developing scalable and generalizable embodied intelligence.065
066
067
068
069
070
071
While collecting data and deploy robots in real-world is resource-intensive and challenging to scale, simulation provides an efficient alternative to enhance robot learning. Simulation supports imitation learning by collecting unlimited and low-cost training data (Wu et al., 2024) and reinforcement learning by providing virtual environments (Makoviychuk et al., 2021; Torne et al., 2024). Meanwhile, simulations enable rapid deployment and standardized test (Ramasubramanian et al., 2022; Do et al., 2025) of algorithms without concerns about hardware damage or safety issues. Overall, simulation facilitates the exploration of innovative strategies for robot learning.072
073
074
075
076
077
078
079
080
081
082
083
084
High-quality digital assets are vital to simulation for robot learning. Simulation platforms (Koenig & Howard, 2004; Todorov et al., 2012; Kolve et al., 2017; Makoviychuk et al., 2021; Puig et al., 2023) depend on digital assets to accurately represent the real world digitally and to simulate its physical characteristics (Choi et al., 2021). High-quality digital assets can effectively reduce the sim-to-real gap, thereby enhancing the performance of robot learning algorithms. For instance, digital-twin assets, which are virtual replicas created via reverse-modeling techniques, can benefit pre-deployment validation and optimization of robotic systems (Straub et al., 2019; Ramakrishnan et al., 2021). Moreover, high-quality digital assets can serve as training data or seed models for synthetic-asset methods such as 3D reconstruction (Liu et al., 2023a; Li et al., 2020; Sun et al., 2023; Liu et al., 2023b) and domain-randomization (Dai et al., 2024; Ge et al., 2024; Torne et al., 2024) techniques, enhancing the data distribution and providing limitless diversity of objects and environments. Conversely, utilizing poor-quality data for synthetic-data generation exacerbates the sim-to-real gap and impair robot learning models (Schraml & Notni, 2024; Osvaldová et al., 2024; Kim et al., 2022).085
086
087
088
089
090
091
092
093
094
095
096
097
098
As robot learning turn form mastering simple tasks such as pick and grasp to dexterous manipulation and interaction tasks, high quality articulated-object assets is of great demand. Current open-source articulated-object datasets fail to meet the needs of robot learning. For instance, PartNet-Mobility (Xiang et al., 2020) suffers from limited visual realism and insufficient physical fidelity of dynamic joints. BEHAVIOR-1K (Li et al., 2024a) offers better visual fidelity, but it is locked into the OmniGibson simulator (Li et al., 2024a) and its physical parameters have not been fine-tuned. Moreover, both datasets are largely sourced from internet-searchable 3D model repositories (Inc., 2024a;b) without adhering to consistent modeling standards, leading to inconsistency in quality. Apart from using existing datasets, people attempts to obtain simulation assets in other ways, facing further challenges. Retrieval-based methods (Liu et al., 2024b;a) and reconstruction techniques (Chen et al., 2024; Eppner et al., 2024) often inherit stylistic biases from their training data and have limited geometric generalization. More recent pipelines (Qiu et al., 2025; Mandi et al., 2024; Le et al., 2025) introduce promising directions, yet face challenges such as mesh quality variance, segmentation noise, and lack of robust joint parameter tuning.099
100
The main bottleneck for articulated object datasets lies in asset quality rather than quantity; to this end, we identify four key aspects that require careful consideration.101
102
103
104
105
106
107

- **Visual Realism.** Assets should be constructed with precise geometric meshes and high-resolution textures to ensure a photorealistic appearance. The amount of triangular faces should be optimized to guarantee real-time simulation performance.
- **Modular Interaction.** Assets should support interactivity (e.g., toggling a switch to turn on a light). These interactions should be modular to enable reuse across scenarios.
- **Physical Fidelity.** Accurate collision geometry and joint dynamics (stiffness, damping, friction) of articulated assets are essential for simulated motion to faithfully reproduce real-world kinematics and dynamics.

108 • **Simulation Friendliness.** Information expanding simulation usages such as pixel-level affordance annotations and accompanied scenes are encouraged. Meanwhile, open-source assets compatible to various simulation platforms and replicable asset creation process should be provided.

111 To meet the mentioned requirements, we introduce ArtVIP, a high-quality and readily deployable
 112 suite of *Articulated-object digital assets with Visual realism, modular Interaction, and Physical*
 113 *fidelity*, designed to facilitate the learning and evaluation of diverse manipulation skills such as ro-
 114 tating, clicking, pulling, and pressing. As illustrated in Fig. 1, ArtVIP encompasses both articulated
 115 object models and complementary indoor-scene assets, all meticulously authored by professional 3D
 116 modelers under a unified asset specification to ensure consistent visual quality and realism. Physi-
 117 cal properties are precisely tuned to reproduce real-world dynamics, thereby enhancing the physical
 118 fidelity. Furthermore, ArtVIP provides pixel-level affordance annotations and uniquely embeds in-
 119 teraction semantics directly into the assets, enabling modular reuse and scalable behavior modeling.

120 In conclusion, ArtVIP offers the following contributions:

121 • We release a collection comprising 9 categories, 37 subcategories, and 992 high-quality digital-
 122 twin articulated objects. All assets exhibit both visual realism and physical fidelity, supported by
 123 quantitative evaluations.

124 • We provide digital-twin scene assets and configured scenarios integrating articulated objects
 125 within scene for immediate use. Extensive experiments on imitation learning, reinforcement
 126 learning, and 3D construction algorithms demonstrate the broader applicability of the assets.

127 • All assets are provided in the modern USD format and remain compatible with established
 128 robotics workflows via conversion to legacy formats such as URDF or MJCF. The detailed pro-
 129 duction process offer comprehensive guidance to facilitate community adoption and replication.

131 2 RELATED WORKS

132 **Simulation Platforms.** A typical simulation platform integrates a physics engine (Smith et al.,
 133 2005; Todorov et al., 2012; Coumans & Bai, 2016; Corporation, 2025; Tasora et al., 2016) and a
 134 rendering engine (Matl, 2019; Chociej et al., 2019; Rojtberg, Pavel and Rogers, David and Streeting,
 135 Steve and others, 2001 – 2024). Game engines (Technologies, 2025.05.14; Games, 2025) offer simi-
 136 lar features but do not natively support ROS (Quigley et al., 2009; Macenski et al., 2022) for robotics.
 137 MuJoCo (Todorov et al., 2012) and Webots (Webots, 2018) excel in simulating rigid body and multi-
 138 joint dynamics but prioritize computational efficiency over high-fidelity rendering. Gazebo (Koenig
 139 & Howard, 2004), despite its large community and robust integration with ROS, provides out-
 140 dated rendering performance and exhibits lower accuracy in physical simulation. Frameworks like
 141 AI2THOR (Kolve et al., 2017), Habitat (Savva et al., 2019; Szot et al., 2021; Puig et al., 2023) and
 142 ALFRED (Shridhar et al., 2020) are designed for mobile manipulation and instruction-following,
 143 fail to deliver precise physical interactions. In contrast, Isaac Sim (Nvidia, 2025.05.14) offers the
 144 highest-fidelity visual rendering and leverages powerful GPU-parallel physics computation, making
 145 it well-suited for robot learning. Other platforms, such as RoboCasa (Nasiriany et al., 2024) (built
 146 upon MuJoCo) and OmniGibson (Li et al., 2024a) (built upon Isaac Sim), have become challeng-
 147 ing to maintain. Consequently, we developed ArtVIP specifically for Isaac Sim to capitalize on its
 148 superior rendering and physics capabilities.

149 **Datasets for Robot Simulation.** Many datasets provide digital assets suitable for robot simula-
 150 tion. Indoor-scene assets (Straub et al., 2019; Shen et al., 2021; Ramakrishnan et al., 2021; Li
 151 et al., 2022) contribute significantly to robot navigation tasks but lacking support for graphical
 152 user interface (GUI)-based editing. Object digital assets includes ShapeNet (Chang et al., 2015),
 153 Objaverse (Deitke et al., 2023) and other digital-twin datasets (Kuang et al., 2023; Dong et al.,
 154 2025). However, these assets can only function as rigid bodies in simulations, preventing robots
 155 from performing articulated manipulation tasks with them. Limited studies addressed articulated
 156 object assets. PartNet-Mobility (Xiang et al., 2020) provides 2,346 articulated-object assets across
 157 46 categories, with many assets suffering from unsmoothed geometric surfaces, low rendering
 158 quality, and imprecise dynamic joint. RoboCasa (Nasiriany et al., 2024) offers 2,508 digital assets, but
 159 only 24 are articulated objects. BEHAVIOR-1K (Li et al., 2024a) includes 543 articulated-object
 160 assets with improved visual fidelity, yet all assets are encrypted and accessible only through Om-
 161 niGibson. These limitations underscore the need for a high-quality, open-source articulated-object
 dataset.

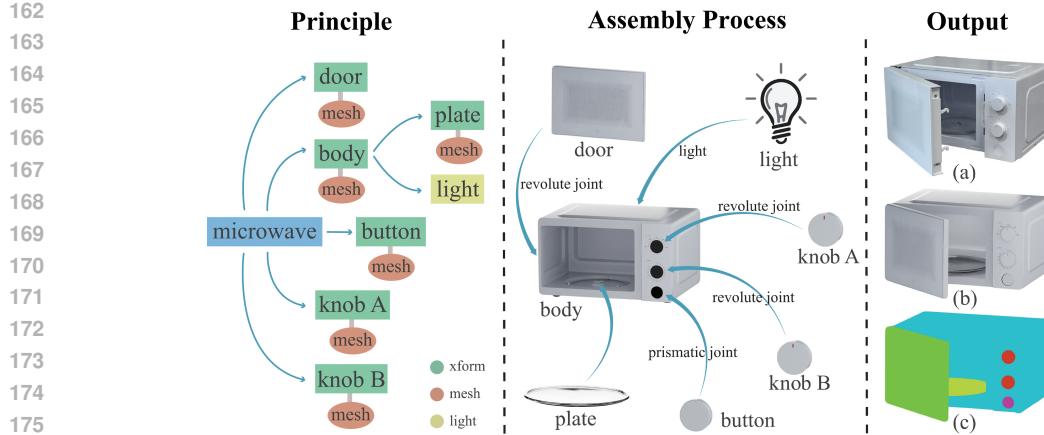


Figure 2: An asset in ArtVIP. **Left:** Top-down assembly principle. **Middle:** Assembly process. **Right:** Comparison between the real object (a) with its digital-twin (b), and annotations (c).

Articulated Objects Construction and Generation Methods. Construction methods (Liu et al., 2024a; Chen et al., 2024; Su et al., 2024; Xue et al., 2021; Wang et al., 2024a) can generate articulated objects from images and reduce the labor cost. However, these methods perform reliably only on objects with simple joints, such as cabinets and desks, and produce assets with compromised visual realism. Generative methods (Yang et al., 2022; Liu et al., 2023c; Long et al., 2023; Xu et al., 2023; Koo et al., 2024), are currently limited to static rigid-body objects. These assets often exhibit distorted and unreasonable meshes, coupled with poor rendering quality. The absence of support for articulated objects in generative methods further limits their applicability to robot learning tasks.

3 ARTVIP COLLECTION AND METHODOLOGY

Existing datasets are largely sourced from pre-made models from public repositories. This leads to inconsistent modeling quality, disorganized part hierarchies, and non-standardized coordinate systems, all of which typically require manual preprocessing for simulation use. While current generative and reconstruction methods can easily scale up, they are still not mature enough to ensure quality. Given these constraints, we opted to prioritize fidelity over scale at this stage. ArtVIP emphasizes both visual realism and physical fidelity across a comprehensive collection of articulated objects. It covers 9 categories and 37 subcategories, encompassing 992 articulated assets (see Appendix Sec. A). Complementary sim-ready scenes (see Appendix Sec. B) and pixel-level annotations (see Appendix Sec. C) are also provided.

3.1 VISUAL REALISM

To ensure visual realism, professional 3D modelers follow unified modeling and assembly guidelines when manually crafting articulated objects. As shown in Fig. 2, we adopt a top-down mechanical modeling approach that decomposes each articulated object into three hierarchical levels: assembly, module, and mesh. An assembly constitutes the complete functional unit, encompassing multiple modules and meshes. Modelers first establish the assembly’s base coordinate frame at the geometric center of the object’s bottom surface. Guided by the assembly’s affordances, functionality, and joint locations, they partition it into rigid-body modules of the Xform type, which expose dynamic information such as transforms, velocities, and world coordinates. Each rigid-body module contains mesh parts that provide geometric detail, visual appearance, and static physical properties, including collision shapes and mass. Modelers follow strict rules regarding meshes, textures, and materials (see Appendix Sec. D) to ensure visual realism. After modeling individual meshes, they assemble them bottom-up—mesh, module, assembly—and integrate dynamic motion by connecting modules with joints (middle panel of Fig. 2), ensuring the asset preserves intended affordances and appearance. Finally, for the finished asset (right panel of Fig. 2), each module is annotated with pixel-level labels to enable precise identification of interaction affordances.

216
217

3.2 PHYSICAL FIDELITY

218 In addition to visual realism, physical fidelity plays a critical role in reducing the sim-to-real gap.
 219 Optimized collision modeling ensures accurate rigid-body contact, improving precision in tasks
 220 such as grasping handles and other force-mediated contact scenarios. Similarly, joint optimization
 221 yields precise joint dynamics, improving the fidelity of articulated components' motion trajectories
 222 during fine-grained operations (e.g., opening cabinet doors or pressing switches). ArtVIP adopts the
 223 following processes.

224 **Collision.** To strike a balance between physical fidelity, interaction consistency, and computational
 225 efficiency, ArtVIP represents each mesh's collision shape using a mix of convex hulls, convex de-
 226 composition, and fine-tuned collision meshes. For relatively regular or simple geometry, ArtVIP
 227 relies on Isaac Sim's default convex hull generation. When a complex mesh can be decomposed
 228 without sacrificing its affordance, 3D modelers split its collision volume into multiple primitive
 229 meshes (e.g., cubes, cylinders). If neither a convex hull nor fine-tuned collision suffices, ArtVIP
 230 employs Isaac Sim's built-in convex decomposition tool, which leverages mesh normals and related
 231 methods to produce accurate collision geometry.

232 **Joints.** To achieve physical fidelity of dynamic joints and simulate variable joint motions in the real
 233 world, we enhance the joint drive equation (NVIDIA, 2025) originally provided by Isaac Sim:

$$234 \quad \tau = K(q) \cdot (q - q_{\text{target}}(q)) + D \cdot (\dot{q} - \dot{q}_{\text{target}}(q)) \quad (1)$$

236 where τ denotes the generalized force or torque applied to drive the joint; q and \dot{q} are the joint
 237 position and velocity, respectively; D denotes damping; and K denotes stiffness. While this equa-
 238 tion models basic joint motions, it does not fully capture complex joint dynamics observed in the
 239 real world. For complex joints such as door closers and light switches, τ may vary with q and \dot{q} .
 240 To accommodate these cases, we parameterize K and the target terms as functions of q , and allow
 241 dependence on \dot{q} when needed. The details are described in the Appendix Sec. E.

242

243
244

3.3 MODULAR INTERACTION

245 A key innovation of this work is embedding customizable behaviors directly within each asset to
 246 enable interactive functionality without writing additional code.

247
248
249

Reproducing complex interactions. We abstract five canonical behavior primitives (Fig. 3) for
 250 articulated objects and instantiate them across ArtVIP, covering 394 assets and more than 900 joints.

- 251 • *Latching/magnetic closure*: Simulates automatic self-closing when the articulation enters a cap-
 252 ture angle range, driven by magnetic attraction or mechanical spring/closer assemblies; once cap-
 253 tured, a closing torque is applied until fully latched. Examples include refrigerator doors (self-
 254 closing hinge with magnetic gasket) and doors equipped with overhead closers.
- 255 • *Damping*: Simulates sliding components and rotational hinges whose effective damping peaks
 256 near the closed position and varies smoothly along the motion, enabling gentle starts and stops.
 257 Examples include nightstand drawers, dishwashers, and cabinets.

258
259
260
261
262
263
264

265 Figure 3: Green arrows denote applied force, yellow dashed lines indicate object motion, and blue
 266 arrows show damping. From left to right: (i) latching/magnetic closure — the door automatically
 267 closes when near shut; (ii) damping — the damping magnitude increases as the drawer is pushed in;
 268 (iii) cross-asset effects — triggering the switch opens the door; (iv) within-asset effects — pressing
 269 the microwave button opens the door and turns on the interior light; (v) hover/hold position — the
 oven door can hold at any angle.

- *Cross-asset effects*: Simulates trigger-based coupling between distinct objects, allowing one object’s state or event to drive another’s behavior. Examples include button-triggered door opening and light switching.
- *Within-asset effects*: Simulates instantaneous, mechanism-internal triggers. For example, pressing a microwave button pops the door open; similar behaviors occur in foot-pedal trash bins and height-adjustable desks.
- *Hover/hold position*: Simulates static-friction-mediated holding in sliding or rotational joints so that, once external forces are removed, the mechanism can remain at any intermediate pose. Examples include oven doors and drawers.

Improving asset reusability. Enhancing simulation development efficiency hinges on modularizing digital assets and maximizing their reusability. Our approach binds behaviors to assets at design time: researchers and artists can simply import the USD file and instantly obtain interaction affordances. This modular, reusable design reduces development overhead and accelerates algorithm iteration, allowing researchers to focus on advancing embodied AI rather than asset programming.

4 EVALUATION

We evaluate ArtVIP along two axes: visual realism and physical fidelity, using quantitative comparisons in simulation and the real world.

4.1 VISUAL REALISM EVALUATIONS

A comparative analysis is conducted among ArtVIP, BEHAVIOR-1K, and PartNet-Mobility (detailed chart in the Appendix Sec. F). As shown on the right of Fig. 4, both BEHAVIOR-1K and PartNet-Mobility exhibit distorted geometry and implausible appearance. In addition, we quantify geometric detail via triangle count, evaluate reconstruction performance, and visualize feature distributions to assess visual realism.

Geometric Detail. Meshes built from densely triangular faces preserve the core geometric detail. A high count of triangular faces improves surface smoothness and minimizes faceting. The left of Fig. 4 illustrates the comparison results on object categories that appear in all three datasets, demonstrating the rich geometric detail in ArtVIP. More analysis and relative profiling are in the Appendix Sec. G.

Reconstruction Performance Evaluation. To assess differences in reconstruction quality across data assets, we conducted experiments using VGGT (Wang et al., 2025), a widely adopted method that has demonstrated strong generalization in real-world reconstruction tasks. Using identical multi-view sampling strategies on the OmniGibson and ArtVIP assets, we generated reconstruction inputs, with results shown on the left portion of Fig. 5. Reconstructions from ArtVIP assets exhibit higher structural fidelity and finer detail preservation compared to those from OmniGibson. This suggests that ArtVIP’s more realistic geometry and material representation enhance the quality and compatibility of sampled images for reconstruction tasks. The results underscore the role of high-fidelity assets in supporting viewpoint diversity and accurate structure recovery.

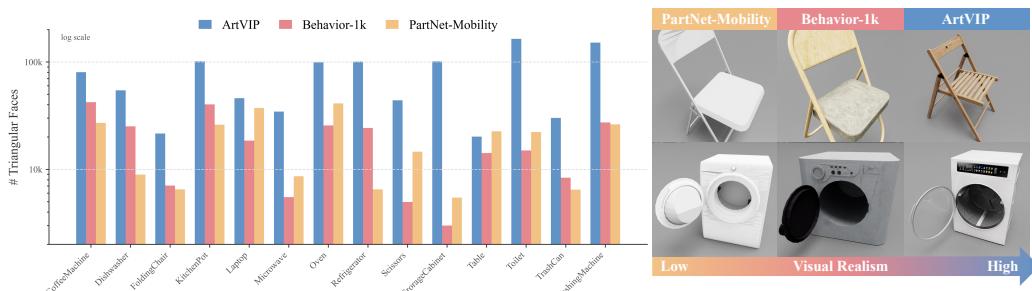


Figure 4: **Left:** Comparison of triangle count. **Right:** Rendering comparison.

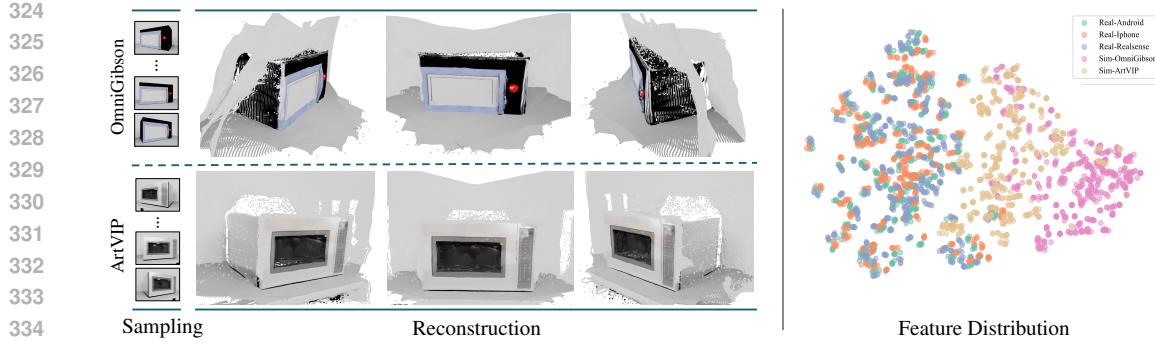


Figure 5: **Left:** Reconstruction of a microwave. OmniGibson yields poor results due to weak visual appearance, while ArtVIP enables better reconstruction via more realistic details. **Right:** CLIP-based (Radford et al., 2021) feature distribution. Each color denotes a data source and ArtVIP features align more closely with real-world data.

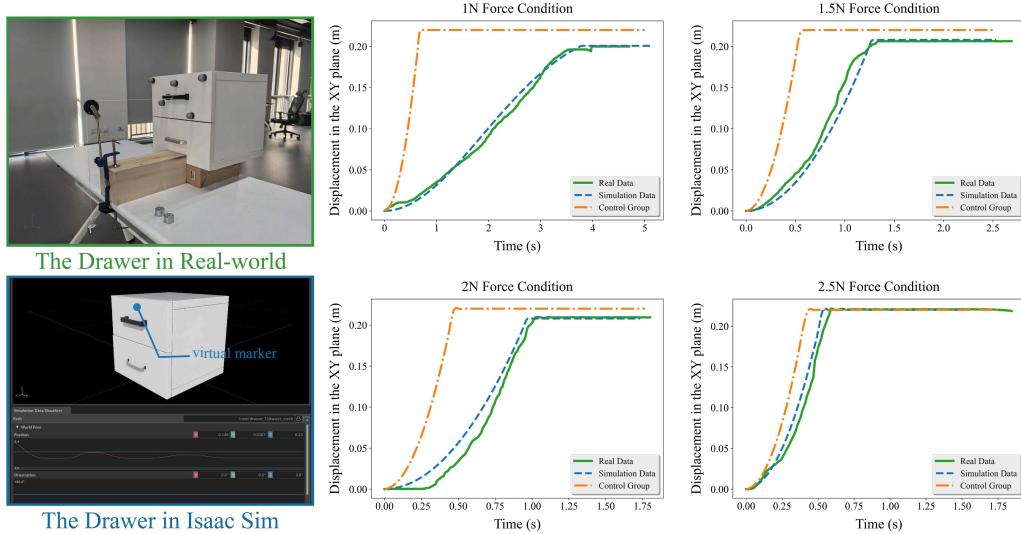


Figure 6: **Left:** Digital-twin asset examples in real-world and simulation. **Right:** Analysis of the drawer’s displacement driven by different forces.

Feature Distribution Visualization Analysis. To verify the visual realism of ArtVIP assets, we randomly sampled 100 3D models and selected corresponding or semantically similar objects from OmniGibson and the real world for comparison. Real-world images were captured using three devices (an Android phone, an iPhone, and an Intel RealSense D435) under multi-view settings. In Isaac Sim, we rendered samples of the ArtVIP and OmniGibson assets using matched camera viewpoints to ensure consistency across domains. We applied t-SNE (Van der Maaten & Hinton, 2008) to visualize the extracted CLIP (Radford et al., 2021) features. As shown on the right portion of Fig. 5, ArtVIP features align more closely with real-world data, indicating higher consistency in visual semantics, texture, and material. This fidelity enhances the value of ArtVIP for simulation-to-reality transfer in downstream tasks.

4.2 PHYSICAL FIDELITY AND INTERACTION EVALUATIONS

To demonstrate the physical fidelity of joint motion within articulated objects, we employed an optical tracking system (0.1 mm spatial resolution and 90 Hz sampling rate) to record motion trajectories of joints on real-world objects. These recordings were compared with the joint motions of their corresponding digital-twin articulated objects in simulation to evaluate the discrepancy between simulated and real-world joint behavior. We test in a common scenario where joint motion

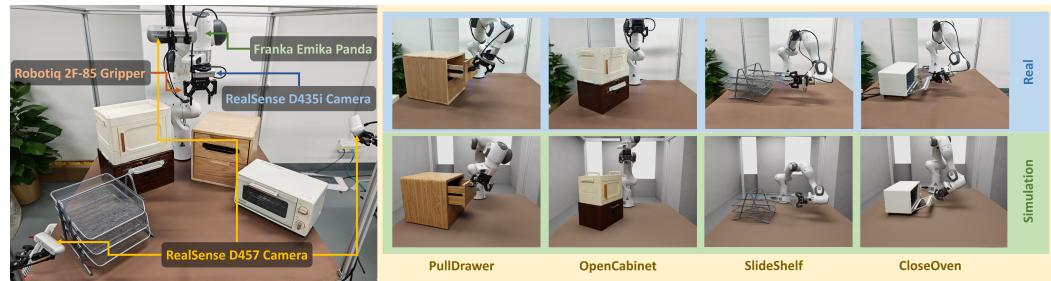


Figure 7: Experimental Setup. We conducted 4 real-world tasks for imitation learning.

triggered by external force. More setting descriptions and evaluation results are described in the Appendix Sec. H.

As shown in Fig. 6, in the real-world experiment, horizontal pulling forces of 1 N, 1.5 N, 2 N, and 2.5 N were applied to the drawer by suspending calibrated weights from the end of the fixed pulley system, ensuring consistent force direction. The drawer’s displacement in the XY plane was recorded in real time. In the simulation environment, two configurations were evaluated: one with default joint parameters and the other with optimized parameters. Both were subjected to the same force configuration as the real-world setup, and the spatial trajectories of the drawer’s keypoints were tracked. The close agreement between the displacement obtained from simulation and real-world experiments, as shown in the right of Fig. 6, demonstrates the physical fidelity of the joints in ArtVIP.

5 APPLICATIONS

To further verify the capability of ArtVIP in supporting downstream robotic learning tasks, we conducted extensive experiments in both the real-world and simulated environments following two primary paradigms in robotic learning: Imitation Learning and Reinforcement Learning.

5.1 IMITATION LEARNING IN REAL WORLD ENVIRONMENTS

Experimental Setup. As illustrated in Fig. 7, we used a Franka robotic arm equipped with a Robotiq 2F-85 gripper and four RealSense cameras to create the real-world experimental environment. These cameras include three external RealSense D457 cameras (placed on the left, right, and top of the table) and one hand-eye RealSense D435i camera mounted at the wrist of the robotic arm. For simulation, we used Isaac Sim and replicated this real-world setup, including the Franka robotic arm, the operating table, camera settings, and the manipulated objects from ArtVIP. We constructed the simulated scene to match the real-world experiment environment as closely as possible.

Task Design and Data Collection. As shown in Fig. 7, we design four challenging articulated-object manipulation tasks: (1) **PullDrawer**, (2) **OpenCabinet**, (3) **SlideShelf**, and (4) **CloseOven**. These tasks demand precise and flexible motions, including rotation, angled pushing, and horizontal translation (see Appendix Sec. I). Data was collected via teleoperation in both real and simulated environments, where articulated objects were randomly placed within a predefined workspace and human operators completed each task. For each task, we gathered 100 successful trajectories in the real world and 100 in simulation. Each trajectory includes RGB streams from four camera viewpoints and full proprioceptive robot states (e.g., joint positions) throughout execution.

Imitation Learning Algorithm. We used two canonical imitation learning baselines, Action Chunking Transformer (ACT) (Zhao et al., 2023) and Diffusion Policy (DP) (Chi et al., 2023), to train the robotic policies for the articulated object manipulation task (more details in Appendix Sec. I).

Experimental Results on Imitation Learning. For each of the four articulated-object manipulation tasks, we trained ACT and DP under the following dataset settings: (1) **Real-Only (RO)**: 100 real-world trajectories; (2) **Sim-Only (SO)**: 100 simulated trajectories; (3) **Real-Sim-Mixed (RSM100+10/20/50/100)**: 100 real-world + 10, 20, 50, 100 simulated trajectories. For each experi-

432 Table 1: Success rates of ACT and DP across dataset settings: RO (real-only), SO (sim-only), and
 433 RSM variants for all tasks.

434

435

Method	Dataset	PullDrawer	OpenCabinet	SlideShelf	CloseOven
ACT (Zhao et al., 2023)	RO	64%	34%	27%	58%
	SO	39%	12%	13%	23%
	RSM100+10	64%	36%	26%	59%
	RSM100+20	68%	38%	27%	60%
	RSM100+50	78%	44%	32%	66%
DP (Chi et al., 2023)	RSM100+100	81%	46%	36%	68%
	RO	66%	49%	44%	66%
	SO	20%	10%	18%	28%
	RSM100+10	65%	53%	47%	67%
	RSM100+20	69%	58%	53%	70%
	RSM100+50	73%	62%	56%	73%
	RSM100+100	79%	66%	59%	78%

441

442

443

444

445

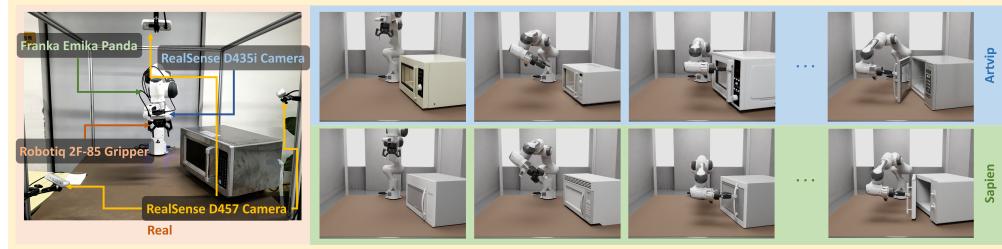
446

447

448 we trained ACT and DP for 50k gradient descent iterations with three different random seeds,
 449 and evaluated the final checkpoint from each run with 60 rollouts to compute per-task success rates.

450

451 Tab. 1 summarizes success rates for ACT and DP under three dataset settings (RO, SO, RSM). We
 452 highlight three findings: (1) **Simulation-trained models achieve zero-shot success in the real**
 453 **world** (e.g., ACT 39% on PullDrawer), reflecting ArtVIP’s high-fidelity visuals and physics that
 454 reduce the sim-to-real gap. (2) With equal data volume, **real-world training outperforms simulation**
 455 (e.g., DP 49% vs. 10% on OpenCabinet), underscoring persistent sim-to-real challenges. (3)
 456 **Mixing real and simulated data boosts performance** (e.g., SlideShelf: DP from 44% to 59%),
 457 indicating that articulated assets in ArtVIP align well with real-world data distributions.



458 Figure 8: We collect data on five microwaves each from ArtVIP and PartNet-Mobility.

460

461

462

463

464

465

466

471 5.2 COMPARISON WITH OTHER ASSETS VIA IMITATION LEARNING.

472

473 To validate the quality of assets in ArtVIP with other datasets, we conduct a digital-cousin compari-
 474 son against PartNet-Mobility. We select five microwave ovens from PartNet-Mobility and five from
 475 ArtVIP. We select microwaves with pull-to-open doors and deliberately exclude button-triggered
 476 opening, as it is operationally trivial. We collect data via teleoperation following the same proce-
 477 dure as in the digital-twin experiments mentioned before, obtaining 100 simulated trajectories per
 478 microwave (500 in total). For the real-world task, we purchase an unseen microwave oven for which
 479 neither PartNet-Mobility nor ArtVIP provides a corresponding digital-twin model. We train ACT
 480 and DP under the following dataset settings: (1) **Real-Only (RO)**: 100 real-world trajectories; (2)
 481 **Sim-Only (SO)**: 500 simulated trajectories; (3) **Real-Sim-Mixed (RSM100+500)**: 100 real-world
 482 + 500 simulated trajectories. All runs use the same training hyperparameters.

483

484

485

486 Tab. 2 summarizes success rates across three dataset settings. We highlight: **higher-quality ArtVIP**
 487 **assets yield stronger zero-shot sim-to-real transfer under SO and higher success under RSM**,
 488 supporting the conclusion that higher-quality assets reduce the sim-to-real gap and lead to higher
 489 success rates.

486
487 Table 2: Success rates of ACT and DP across dataset settings: RO (real-only), SO (sim-only), and
488 RSM100+500, comparing ArtVIP and PartNet-Mobility on the microwave door-pull task.
489

Method	Dataset	ArtVIP	PartNet-Mobility
ACT	RO		56%
	SO	41%	32%
	RSM100+500	79%	68%
DP	RO		62%
	SO	45%	35%
	RSM100+500	83%	70%

498 5.3 REINFORCEMENT LEARNING IN HIGH-FIDELITY SIMULATORS

500 Reinforcement learning (RL) requires training environments that mirror real-world physical and
501 perceptual complexity. To validate the quality of articulated assets in ArtVIP, we designed a Clos-
502 eTrashcan task with a Franka robotic arm and trained a two-stage agent with the state-of-the-art
503 visual RL framework EAGLE (Zhao et al., 2025) in Isaac Sim (see Appendix Sec. J).

504
505 Table 3: Sim vs. real performance across training checkpoints (k iterations).

Checkpoint (k)	500	450	400	350	300
Sim	0.98	0.91	0.85	0.81	0.73
Real	0.93	0.86	0.80	0.70	0.60

511 EAGLE enables efficient training of visuomotor policies. In Stage 1, we train a PPO expert (Schul-
512 man et al., 2017) with low-level state inputs. In Stage 2, we distill this expert into a visuomotor
513 policy, applying EAGLE’s self-supervised attention masks and control-aware augmentation. Ran-
514 domConv (Lee et al., 2019) is used to diversify control-irrelevant backgrounds. To bridge the back-
515 ground visual gap between simulation and the real world, we adopt the Visual Matching technique
516 introduced in (Li et al., 2024b).

517 We train the RL policy in simulation and then deploy it in the real world on the same task, ensuring
518 sim-to-real visual consistency. Tab. 3 reports success rates at five checkpoints between 300k and
519 500k training iterations, each evaluated with 100 simulation trials and 30 real-world trials under di-
520 verse initial object poses. The RL policy trained in ArtVIP exhibits an absolute sim-to-real success
521 rate gap of 0.05, indicating that ArtVIP provides high physical fidelity and visual realism. Follow-
522 ing (Li et al., 2024b), we compute the Pearson correlation coefficient from the success rates in Tab. 3
523 as

$$524 \quad r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}},$$

525 where x_i and y_i denote the corresponding success rates in simulation and the real world at the i -th
526 checkpoint. A high Pearson correlation indicates a strong linear relationship between simulated and
527 real-world performance. The value $r = 0.9886$ using data from Tab. 3 shows that ArtVIP provides
528 a reliable simulated training and evaluation pipeline for RL.

532 6 LIMITATION AND CONCLUSION

535 We introduced ArtVIP, a high-quality dataset of articulated objects for robotic manipulation, fea-
536 turing visual realism, accurate physical properties, and modular interaction capabilities. We validated
537 its quality via diverse evaluations and demonstrated effectiveness in both imitation learning and
538 reinforcement learning. Scaling remains bottlenecked by intensive human labor for asset model-
539 ing; future work will explore generative methods to automate synthesis, reduce manual effort, and
broaden object diversity.

540 REFERENCES
541

542 Aude Billard and Danica Kragic. Trends and challenges in robot manipulation. *Science*, 364:
543 eaat8414, 2019.

544 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski,
545 Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action
546 models transfer web knowledge to robotic control, 2023a.

547 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn,
548 Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics
549 transformer for real-world control at scale, 2023b.

550 Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li,
551 Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d
552 model repository. *arXiv preprint arXiv:1512.03012*, 2015.

553 Zoey Chen, Aaron Walsman, Marius Memmel, Kaichun Mo, Alex Fang, Karthikeya Vemuri, Alan
554 Wu, Dieter Fox, and Abhishek Gupta. Urdformer: A pipeline for constructing articulated simula-
555 tion environments from real-world images. *arXiv preprint arXiv:2405.11656*, 2024.

556 Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake,
557 and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The
558 International Journal of Robotics Research*, pp. 02783649241273668, 2023.

559 Maciek Chociej, Peter Welinder, and Lilian Weng. Orrb – openai remote rendering backend, 2019.
560 URL <https://arxiv.org/abs/1906.11633>.

561 HeeSun Choi, Cindy Crump, Christian Duriez, Asher Elmquist, Gregory Hager, David Han, Frank
562 Hearl, Jessica Hodgins, Abhinandan Jain, Frederick Leve, et al. On the use of simulation in
563 robotics: Opportunities, challenges, and suggestions for moving forward. *Proceedings of the
564 National Academy of Sciences*, 118:e1907856118, 2021.

565 NVIDIA Corporation. Nvidia physx sdk, 2025. URL <https://developer.nvidia.com/physx-sdk>.

566 Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games,
567 robotics and machine learning, 2016.

568 Tianyuan Dai, Josiah Wong, Yunfan Jiang, Chen Wang, Cem Gokmen, Ruohan Zhang, Jiajun Wu,
569 and Li Fei-Fei. Automated creation of digital cousins for robust policy learning. In *8th Annual
570 Conference on Robot Learning*, 2024.

571 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig
572 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-
573 tated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
574 Recognition*, pp. 13142–13153, 2023.

575 Tan-Dzung Do, Nandiraju Gireesh, Jilong Wang, and He Wang. Watch less, feel more: Sim-to-real
576 rl for generalizable articulated object manipulation via motion adaptation and impedance control.
577 *arXiv preprint arXiv:2502.14457*, 2025.

578 Zhao Dong, Ka Chen, Zhaoyang Lv, Hong-Xing Yu, Yunzhi Zhang, et al. Digital twin catalog: A
579 large-scale photorealistic 3d object digital twin dataset, 2025.

580 Jonathan Embley-Riches, Jianwei Liu, Simon Julier, and Dimitrios Kanoulas. Unreal robotics lab:
581 A high-fidelity robotics simulator with advanced physics and rendering, 2025.

582 Clemens Eppner, Adithyavairavan Murali, Caelan Garrett, Rowland O’Flaherty, Tucker Hermans,
583 Wei Yang, and Dieter Fox. scene_synthesizer: A python library for procedural scene generation
584 in robot manipulation. *Journal of Open Source Software*, 2024.

585 Epic Games. Unreal engine, 2025. URL <https://www.unrealengine.com>.

594 Yunhao Ge, Yihe Tang, Jiashu Xu, Cem Gokmen, Chengshu Li, Wensi Ai, Benjamin Jose Martinez,
 595 Arman Aydin, Mona Anvari, Ayush K Chakravarthy, et al. Behavior vision suite: Customizable
 596 dataset generation via simulation. In *Proceedings of the IEEE/CVF Conference on Computer*
 597 *Vision and Pattern Recognition*, pp. 22401–22412, 2024.

598
 599 Xiaoshen Han, Minghuan Liu, Yilun Chen, Junqiu Yu, Xiaoyang Lyu, Yang Tian, Bolun Wang,
 600 Weinan Zhang, and Jiangmiao Pang. Re³sim: Generating high-fidelity simulation data via 3d-
 601 photorealistic real-to-sim for robotic manipulation, 2025.

602 Trimble Inc. 3d warehouse, 2024a. URL <https://3dwarehouse.sketchup.com/>. Online
 603 3D model repository for SketchUp.

604
 605 TurboSquid Inc. Turbosquid, 2024b. URL <https://www.turbosquid.com/>. Professional
 606 3D model marketplace.

607 Physical Intelligence. $\pi_{0.5}$: a vision-language-action model with open-world generalization, 2025.

608
 609 Alexander Kim, Kyuhyup Lee, Sejoon Lee, Jinwoo Song, Soonwook Kwon, and Suwan Chung.
 610 Synthetic data and computer-vision-based automated quality inspection system for reused scaf-
 611 folding. *Applied Sciences*, 12(19), 2022. ISSN 2076-3417. doi: 10.3390/app121910097. URL
 612 <https://www.mdpi.com/2076-3417/12/19/10097>.

613 Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, et al. Openvla:
 614 An open-source vision-language-action model, 2024.

615
 616 Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-
 617 robot simulator. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol-
 618 ume 3, pp. 2149–2154, 2004.

619
 620 Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel
 621 Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An Interactive 3D Environment
 622 for Visual AI. *arXiv*, 2017.

623
 624 Juil Koo, Seungwoo Yoo, Minh Hieu Nguyen, and Minhyuk Sung. Salad: Part-level latent diffusion
 625 for 3d shape generation and manipulation, 2024. URL <https://arxiv.org/abs/2303.12236>.

626
 627 Zhengfei Kuang, Yunzhi Zhang, Hong-Xing Yu, Samir Agarwala, Elliott Wu, Jiajun Wu, et al.
 628 Stanford-orb: a real-world 3d object inverse rendering benchmark. *Advances in Neural*
 629 *Information Processing Systems*, 36:46938–46957, 2023.

630
 631 Long Le, Jason Xie, William Liang, Hung-Ju Wang, Yue Yang, Yecheng Jason Ma, Kyle Vedder,
 632 Arjun Krishna, Dinesh Jayaraman, and Eric Eaton. Articulate-anything: Automatic modeling
 633 of articulated objects via a vision-language foundation model, 2025. URL <https://arxiv.org/abs/2410.13882>.

634
 635 Kimin Lee, Kibok Lee, Jinwoo Shin, and Honglak Lee. Network randomization: A simple tech-
 636 nique for generalization in deep reinforcement learning. *International Conference on Learning*
 637 *Representations*, 2019.

638
 639 Chengshu Li, Fei Xia, Roberto Martín-Martín, Michael Lingelbach, Sanjana Srivastava, Bokui Shen,
 640 Kent Elliott Vainio, Cem Gokmen, Gokul Dharan, Tanish Jain, et al. igibson 2.0: Object-centric
 641 simulation for robot learning of everyday household tasks. In *5th Annual Conference on Robot*
642 Learning, 2022.

643
 644 Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-
 645 Martín, Chen Wang, Gabrael Levine, Michael Lingelbach, Jiankai Sun, et al. Behavior-1k: A
 646 human-centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation,
 647 2024a. URL <https://arxiv.org/abs/2403.09227>.

648
 649 Xiaolong Li, He Wang, Li Yi, Leonidas Guibas, A. Lynn Abbott, and Shuran Song. Category-level
 650 articulated object pose estimation, 2020. URL <https://arxiv.org/abs/1912.11913>.

648 Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu,
649 Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su,
650 Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation,
651 2024b. URL <https://arxiv.org/abs/2405.05941>.

652 Shengjie Lin, Jiading Fang, Muhammad Zubair Irshad, Vitor Campagnolo Guizilini, Rares Andrei
653 Ambrus, Greg Shakhnarovich, and Matthew R Walter. Splart: Articulation estimation and part-
654 level reconstruction with 3d gaussian splatting. *arXiv preprint arXiv:2506.03594*, 2025.

655

656 Jiayi Liu, Denys Iliash, Angel X Chang, Manolis Savva, and Ali Mahdavi-Amiri. Singapo: Single
657 image controlled generation of articulated parts in objects. *arXiv preprint arXiv:2410.16499*,
658 2024a.

659

660 Jiayi Liu, Hou In Ivan Tam, Ali Mahdavi-Amiri, and Manolis Savva. Cage: Controllable articulation
661 generation, 2024b. URL <https://arxiv.org/abs/2312.09570>.

662

663 Xueyi Liu, Bin Wang, He Wang, and Li Yi. Few-shot physically-aware articulated mesh generation
664 via hierarchical deformation, 2023a.

665

666 Xueyi Liu, Ji Zhang, Ruizhen Hu, Haibin Huang, He Wang, and Li Yi. Self-supervised category-
667 level articulated object pose estimation with part-level se(3) equivariance, 2023b. URL <https://arxiv.org/abs/2302.14268>.

668

669 Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu.
670 Meshdiffusion: Score-based generative 3d mesh modeling. *arXiv preprint arXiv:2303.08133*,
2023c.

671

672 Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma,
673 Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d
674 using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023.

675

676 Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot op-
677 erating system 2: Design, architecture, and uses in the wild. *Science Robotics*, 7:eabm6074,
2022.

678

679 Viktor Makovychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin,
680 David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance
681 gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.

682

683 Zhao Mandi, Yijia Weng, Dominik Bauer, and Shuran Song. Real2code: Reconstruct articulated
684 objects via code generation, 2024. URL <https://arxiv.org/abs/2406.08474>.

685

686 Matthew Matl. Pyrender. <https://github.com/mmatl/pyrender>, 2019.

687

688 Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi,
Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for gener-
689 alist robots. In *Robotics: Science and Systems*, 2024.

690

691 Federico Nesti, Gianluca D’Amico, Mauro Marinoni, and Giorgio Buttazzo. Simpriv: a simulation
692 framework for physical robot interaction with virtual environments, 2025.

693

NVIDIA. Joint tuning — Isaac Sim documentation, 2025. URL https://docs.isaacsim.omniverse.nvidia.com/latest/robot_setup/joint_tuning.html#gain-tuning.

694

Nvidia. Understanding physically-based rendering, 2025a. URL <https://docs.omniverse.nvidia.com/simready/latest/simready-asset-creation/material-best-practices.html>.

695

Nvidia. Omniverse rtx renderer, 2025b. URL <https://docs.omniverse.nvidia.com/materials-and-rendering/latest/rtx-renderer.html>.

696

697

Nvidia. Nvidia isaac sim, 2025.05.14. URL <https://developer.nvidia.com/isaac-sim>. Isaac Sim.

698

699

700

701

702 Katarína Osvaldová, Lukáš Gajdošech, Viktor Kocur, and Martin Madaras. Enhancement of 3d
 703 camera synthetic training data with noise models. [arXiv preprint arXiv:2402.16514](https://arxiv.org/abs/2402.16514), 2024.

704

705 Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dallaire Cote, Tsung-Yen Yang, Ruslan Part-
 706 sey, Ruta Desai, Alexander William Clegg, Michal Hlavac, So Yeon Min, et al. Habitat 3.0: A
 707 co-habitat for humans, avatars and robots, 2023.

708

709 Xiaowen Qiu, Jincheng Yang, Yian Wang, Zhehuan Chen, Yufei Wang, Tsun-Hsuan Wang, Zhou
 710 Xian, and Chuang Gan. Articulate anymesh: Open-vocabulary 3d articulated objects modeling,
 711 2025. URL <https://arxiv.org/abs/2502.02590>.

712

713 Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler,
 714 Andrew Y Ng, et al. Ros: an open-source robot operating system. In [ICRA workshop on open](#)
 715 [source software](#), volume 3, pp. 5, 2009.

716

717 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 718 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 719 models from natural language supervision. In [International Conference on Machine Learning](#), pp.
 720 8748–8763, 2021.

721

722 Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexan-
 723 der Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X
 724 Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3d): 1000
 725 large-scale 3d environments for embodied AI. In [Thirty-fifth Conference on Neural Information](#)
 726 [Processing Systems Datasets and Benchmarks Track](#), 2021. URL <https://arxiv.org/abs/2109.08238>.

727

728 Aswin K Ramasubramanian, Robins Mathew, Matthew Kelly, Vincent Hargaden, and Nikolaos Pa-
 729 pakostas. Digital twin for human–robot collaboration in manufacturing: Review and outlook.
 730 [Applied Sciences](#), 12:4811, 2022.

731

732 Rojberg, Pavel and Rogers, David and Streeting, Steve and others. Ogre scene-oriented, flexible 3d
 733 engine. <https://www.ogre3d.org/>, 2001 – 2024.

734

735 Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain,
 736 Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied
 737 ai research. In [Proceedings of the IEEE/CVF international conference on computer vision](#), pp.
 738 9339–9347, 2019.

739

740 Dominik Schraml and Gunther Notni. Synthetic training data in ai-driven quality inspection: The
 741 significance of camera, lighting, and noise parameters. [Sensors](#), 24(2), 2024. ISSN 1424-8220.
 742 doi: 10.3390/s24020649. URL <https://www.mdpi.com/1424-8220/24/2/649>.

743

744 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 745 optimization algorithms. [arXiv preprint arXiv:1707.06347](https://arxiv.org/abs/1707.06347), 2017.

746

747 Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Claudia
 748 Pérez-D’Arpino, Shyamal Buch, Sanjana Srivastava, Lyne Tchapmi, et al. igibson 1.0: A sim-
 749 ulation environment for interactive tasks in large realistic scenes. In [IEEE/RSJ International](#)
 750 [Conference on Intelligent Robots and Systems](#), pp. 7520–7527, 2021.

751

752 Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi,
 753 Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions
 754 for everyday tasks. In [Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern](#)
 755 [Recognition](#), pp. 10740–10749, 2020.

756

757 Russell Smith et al. Open dynamics engine, 2005.

758

759 Mark W Spong, Seth Hutchinson, and M Vidyasagar. Robot modeling and control. [John Wiley](#)
 760 & 2020.

761

762 Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel,
 763 Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor
 764 spaces. [arXiv preprint arXiv:1906.05797](https://arxiv.org/abs/1906.05797), 2019.

756 Jiayi Su, Youhe Feng, Zheng Li, Jinhua Song, Yangfan He, Botao Ren, and Botian Xu. Artformer
757 Controllable generation of diverse 3d articulated objects. [arXiv preprint arXiv:2412.07237](https://arxiv.org/abs/2412.07237), 2024.

758

759 Xiaohao Sun, Hanxiao Jiang, Manolis Savva, and Angel Xuan Chang. Opdmulti: Openable part
760 detection for multiple objects, 2023. URL <https://arxiv.org/abs/2303.14087>.

761

762 Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah
763 Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, et al. Habitat 2.0:
764 Training home assistants to rearrange their habitat. [Advances in Neural Information Processing
765 Systems](https://paperswithcode.com/paper/habitat-2-0-training-home-assistants-to-rearrange-their-habitat), 34:251–266, 2021.

766

767 Alessandro Tasora, Radu Serban, Hammad Mazhar, Arman Pazouki, Daniel Melanz, Jonathan Fleis-
768 chmann, Michael Taylor, Hiroyuki Sugiyama, and Dan Negru. Chrono: An open source multi-
769 physics dynamics engine. In [High Performance Computing in Science and Engineering: Second
770 International Conference, HPCSE 2015, Soláň, Czech Republic, May 25–28, 2015, Revised
771 Selected Papers 2](https://hpcse2015.solaen.cz/paper/selected_papers_2.pdf), pp. 19–49, 2016.

772

773 Gemini Robotics Team. Gemini robotics: Bringing ai into the physical world, 2025.

774

775 Unity Technologies. Unity, 2025.05.14. URL <https://unity.com/>. Game development
776 platform.

777

778 Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control.
779 In [IEEE/RSJ International Conference on Intelligent Robots and Systems](https://ieeexplore.ieee.org/abstract/document/6285771), pp. 5026–5033, 2012.

780

781 Marcel Torne, Anthony Simeonov, Zechu Li, April Chan, Tao Chen, Abhishek Gupta, and Pulkit
782 Agrawal. Reconciling reality through simulation: A real-to-sim-to-real approach for robust ma-
783 nipulation. [Arxiv](https://arxiv.org/abs/2403.07001), 2024.

784

785 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. [Journal of machine
786 learning research](https://www.jmlr.org/papers/v09/vandermaaten08a/vandermaaten08a.pdf), 9(11), 2008.

787

788 Haowen Wang, Zhen Zhao, Zhao Jin, Zhengping Che, Liang Qiao, Yakun Huang, Zhipeng Fan,
789 Xiuquan Qiao, and Jian Tang. Sm 3: Self-supervised multi-task modeling with multi-view 2d
790 images for articulated objects. In [International Conference on Robotics and Automation](https://arxiv.org/abs/2403.12498), pp.
791 12492–12498, 2024a.

792

793 Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David
794 Novotny. Vggt: Visual geometry grounded transformer. In [Proceedings of the IEEE/CVF
795 Conference on Computer Vision and Pattern Recognition](https://arxiv.org/abs/2403.12498), 2025.

796

797 Lirui Wang, Xinlei Chen, Jialiang Zhao, and Kaiming He. Scaling proprioceptive-visual learning
798 with heterogeneous pre-trained transformers, 2024b.

799

800 Webots. <http://www.cyberbotics.com>, 2018. URL <http://www.cyberbotics.com>. Open-
801 source Mobile Robot Simulation Software.

802

803 Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo
804 Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence
805 normative data for robot manipulation. [arXiv preprint arXiv:2412.13877](https://arxiv.org/abs/2412.13877), 2024.

806

807 Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao
808 Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN:
809 A simulated part-based interactive environment. In [Proceedings of the IEEE/CVF Conference on
810 Computer Vision and Pattern Recognition](https://arxiv.org/abs/2003.11107), pp. 11097–11107, 2020.

811

812 Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli,
813 Gordon Wetzstein, Zexiang Xu, and Kai Zhang. Dmv3d: Denoising multi-view diffusion using
814 3d large reconstruction model, 2023. URL <https://arxiv.org/abs/2311.09217>.

815

816 Han Xue, Liu Liu, Wenqiang Xu, Haoyuan Fu, and Cewu Lu. Omad: Object model with articulated
817 deformations for pose estimation and retrieval, 2021.

810 Jie Yang, Kaichun Mo, Yu-Kun Lai, Leonidas J. Guibas, and Lin Gao. Dsg-net: Learning disen-
811 tangled structure and geometry for 3d shape generation, 2022. URL [https://arxiv.org/](https://arxiv.org/abs/2008.05440)
812 abs/2008.05440.

813
814 Yunhao Zhang and Junchi Yan. Crossformer: Transformer utilizing cross-dimension dependency
815 for multivariate time series forecasting. In International Conference on Learning Representations,
816 2023.

817 Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual
818 manipulation with low-cost hardware, 2023.

819
820 Yinuo Zhao, Kun Wu, Tianjiao Yi, Zhiyuan Xu, Zhengping Che, Chi Harold Liu, and Jian
821 Tang. Efficient training of generalizable visuomotor policies via control-aware augmentation.
822 In Proceedings of the 24th International Conference on Autonomous Agents and Multiagent
823 Systems, 2025.

824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863



(a) Kitchen

(b) Small livingroom

Figure 9: Scenes: all articulated joints in the open state.

A ARTICULATED OBJECTS

ArtVIP comprises 992 articulated objects, encompassing 9 categories and 37 subcategories, with a total of 2156 prismatic joints and 1809 revolute joints. The detailed breakdown, including approximate human labor time, is presented in Tab. 4.

B SCENES

We provide sim-ready complex, dynamic environments—six in total: childrenroom, diningroom, kitchen, kitchen with parlor, large livingroom, and small livingroom (see Fig. 9 for two example scenes). Every object in these environments, including fixed furniture, supports physical interaction. This includes switches, small appliances, plush toys, laptops, books, spice jars, and more. For example, the kitchen environment contains a total of 65 joints, and all objects can be used just like their real-world counterparts. Robots can operate the light switch on the wall, open the refrigerator door, place items on shelves, or challenge their motion capabilities by crouching to open drawers beneath the stove top. Additionally, users can freely place the 992 articulated objects provided in ArtVIP into any of these environments via the Isaac Sim GUI, enabling the creation of rich robot interaction scenarios such as grasping, pulling, pressing, and placing. Users can also utilize open-source tools like mjcf2usd and urdf2usd to convert assets from other datasets into the USD format, allowing seamless integration with ArtVIP assets. This kind of sim-ready, complex environment is currently unique to ArtVIP. Moreover, the ability to edit and save assets directly through a GUI reflects an open-source spirit that is not yet common in other datasets.

C ANNOTATIONS

Annotations in ArtVIP provide objective descriptions of object parts, thereby supporting robots' ability to infer task-appropriate interaction behaviors. We further argue that annotations are most meaningful when aligned with consistent modeling standards. For example, for a desk, modelers often merge the legs and tabletop into a single mesh, which limits part-level annotation based on distinct interaction functions. To address this, we highlight functional components in Tab. 5 that frequently participate in interactions yet are commonly overlooked during mesh segmentation. An example segmentation is shown in Fig. 10.

D MODELING STANDARDS

In simulation systems, the use of high-quality meshes, textures, and materials confers several advantages. High-fidelity visuals reduce the disparity between simulation and reality (Nesti et al., 2025), thereby narrowing the sim-to-real gap and enabling robotic policies to be deployed in real-world environments with minimal or even zero-shot adaptation (Han et al., 2025; Embley-Riches et al., 2025). Photorealistic simulation data can be employed to train and validate visual perception algorithms, such as object detection, semantic segmentation, and SLAM. Moreover, realistic models not

918
 919 Table 4: Detailed breakdown of object categories, modeling time (each), physics tuning time (each),
 920 and count.

category	subcategories	modeling time	physics tuning time	count
furniture	chair	2h	0.3h	23
	table	1.5h	0.2h	131
	cabinet	3.1h	0.4h	183
	cupboard	15h	2h	11
	bed	3h	0.3h	28
kitchenware	home decor	2.5h	0.3h	30
	cookware	1.9h	0.3h	81
kitchen appliances	coffee machine	3h	0.3h	14
	built-in oven	5h	0.3h	14
	microwave	3h	0.3h	8
	oven	4h	0.3h	11
	dishwasher	5h	0.4h	19
	water dispenser	3h	0.3h	6
	rice cooker	3h	0.3h	14
fixtures	fridge	6h	0.5h	22
	juicer	4h	0.3h	6
	faucet	2h	0.2h	14
appliances	toilet	4h	0.4h	14
	door	2h	0.3h	10
	computer	2.5h	0.3h	13
	fan	1.8h	0.3h	34
	air conditioner	4h	0.3h	3
	washing machine	5.7h	0.5h	30
	speaker	1.5h	0.3h	14
cleaning tools	floor lamp	1h	0.3h	28
	mop	2h	0.3h	8
	pump bottle	2h	0.3h	14
stationery	trash can	2h	0.3h	18
	scissors	1h	0.2h	28
	stapler	1.5h	0.2h	11
storage	utility knife	1h	0.2h	19
	folder	0.5h	0.2h	8
	storage box	2h	0.3h	25
Mechanical equipment	toolbox	2.5h	0.3h	22
	cardboard box	1.5h	0.2h	28
	electrical equipment	3h	0.3h	17
	non-electrical equipment	3h	0.3h	33

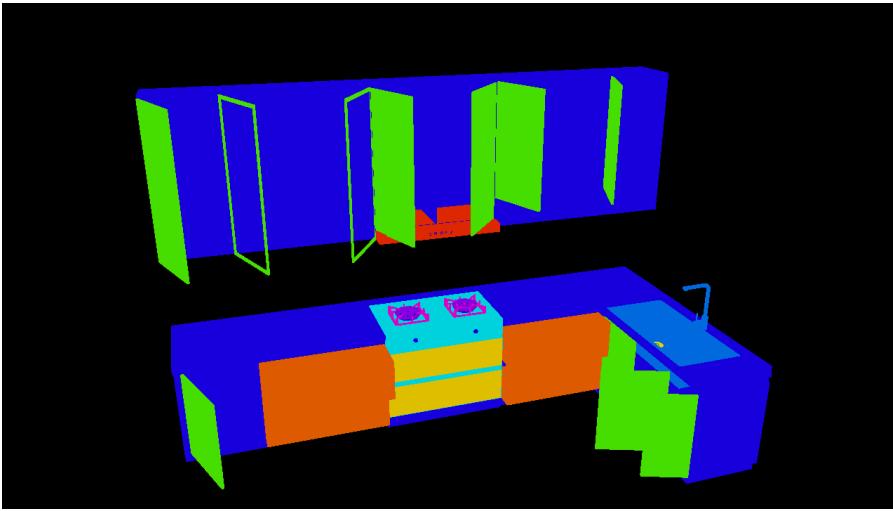


Figure 10: Segmentation result of the kitchen scene.

only enhance visual fidelity but also improve interaction effects within simulations. When robots perform actions such as grasping, collision, or force-based interactions, accurate geometry ensures stable and reliable feedback. To achieve photorealistic appearance and minimize the sim-to-real visual gap, we addressed the following standards:

Mesh. Manifold meshes form the core geometric foundation of each asset, defining the object’s overall contour and spatial occupancy. These meshes are critical for generating collision bodies that maintain accuracy in physical interactions. ArtVIP ensures that mesh details produce smooth surfaces and lifelike contours, avoiding jagged or blocky appearances. Additionally, through normal vector optimization algorithms, redundant vertices are merged, reducing geometric data volume and thereby alleviating computational burdens in simulation.

Texture. Textures are mapped onto mesh surfaces via UV coordinates to provide visual details. ArtVIP employs high-resolution textures to capture fine surface characteristics, such as the metallic sheen of a refrigerator or the subtle grain of wood on a chair. Furthermore, textures are meticulously aligned with the UV map to prevent stretching, distortion, or visible seams.

Material. A material is a collection of rendering parameters, including references to textures, that defines how an object’s surface responds to light. ArtVIP leverages RTX Renderer (Nvidia, 2025b) in Isaac Sim and adopts Physically Based Rendering (PBR) (Nvidia, 2025a) to accurately simulate diffuse and specular reflections, enabling rendering effects such as roughness and emissive properties. This approach allows for the realistic representation of diverse materials, achieving true-to-life visual fidelity.

E PHYSICAL FIDELITY OF JOINTS

To achieve physical fidelity of dynamic joint and simulate variable joints motions in the real world, we enhance the joint drive equation originally provided by Isaac Sim:

$$\tau = K(q) \cdot (q - q_{\text{target}}(q)) + D \cdot (\dot{q} - \dot{q}_{\text{target}}(q)) \quad (2)$$

where τ represents the force(F) and torque(T) applied to drive the joint, q and \dot{q} are the joint position and velocity, respectively, D donates damping, and K donates stiffness. While this equation can model basic joint motions, it fails to fully replicate complex dynamic joint motions in the real world. For complex joints such as door closers and light switches, τ may vary with q and \dot{q} . To accommodate the above situations, we design functions of q and \dot{q} .

Impact from \dot{q} . Friction must be accounted for in simulation and cannot be modeled as a constant. It imposes resistance to the force generated by the joint drive τ , and we propose the following equation

Table 5: Annotation labels and descriptions in ArtVIP.

Label	Description
armrest	Chair armrest
backrest	Chair backrest
ball_handle	Handle for lifting the main body, such as the handle of a toolbox
blade	Blade of a utility knife, scissors, or fan blades
body	Parts that need labeling excluding base and lid
button	Applies to all push-button switch components of models
door	Door of cabinets, refrigerators, ovens, etc.
drawer	Drawer of cabinets, refrigerators, toolboxes, etc.
front_cover	Cover of a folder
fun_guard	Fan protective cover
handle	Any handles
headrest	Chair headrest
jaw	Head of pliers, the part that contacts the gripped item
keyboard	Computer keyboard
knob	Applies to all rotary switch components of models
lid	Such as cardboard box lid, electric steamer lid, trash can lid
light	All types of lights
mop_head	Mop head
pedal	Foot pedal, such as on a step-on trash can
pipe	Water pipe part of faucet
plate	All types of plates
pole	Rod-shaped component
portafilter	A handle holds the coffee grounds
pot	Inner pot of rice cookers, steamers, etc.
rack	Rack in an oven, refrigerator door shelf
roller	Washing machine drum
screen	Electronic product screen
seat	Chair seat
shelf	Shelf part of cabinets, refrigerators, etc.
spout	Spout of a pump bottle, water dispenser, etc.
stapler_magazine	Staple compartment of a stapler
tabletop	Top surface of a table
toilet_seat	Toilet seat
touch_pad	Computer touchpad
wheel	Chair wheels

with three different conditions:

$$F_{\text{friction}}(\dot{q}) = \begin{cases} -F_{\text{ext}} & \dot{q} = 0 \text{ and } |F_{\text{ext}}| \leq \mu_s \cdot (|F| + |T|) \\ -\mu_s \cdot (|F| + |T|) \cdot \text{sign}(F_{\text{ext}}) & \dot{q} = 0 \text{ and } |F_{\text{ext}}| > \mu_s \cdot (|F| + |T|) \\ -D \cdot \dot{q} \cdot \text{sign}(\dot{q}) & \dot{q} \neq 0 \end{cases} \quad (3a)$$

$$(3b) \quad (3c)$$

We illustrate the friction from static friction, to maximum static friction, and finally to dynamic friction, corresponding to conditions from Eqn. equation 3a through Eqn. equation 3c. F_{ext} denotes the static friction. The coefficient μ_s denotes the static friction coefficient, which can be configured in Isaac Sim via the `Joint Friction` parameter. The `sign` function ensures that the frictional force is applied in the correct direction.

Impact from q . The latch release mechanism exemplifies the position-dependent joint drive, we analyze a button-actuated trash bin lid mechanism. When the button is depressed, it triggers a linkage to retract the spring-loaded latch, enabling the lid to freely rotate under torsional spring torque to $q_{\text{upper_bound}}$.

$$q_{\text{target}}(q) = \begin{cases} q_{\text{upper_bound}} & \text{if } q > q_{\text{threshold}} \text{ and } S_{\text{open}} = 1 \\ q_{\text{lower_bound}} & \text{if } q < q_{\text{threshold}} \text{ and } S_{\text{open}} = 0 \end{cases} \quad (4a)$$

$$(4b)$$

1080

1081

1082 We further investigate joint motion with abrupt stiffness variations, exemplified by refrigerator door
 1083 closers and magnetic latching mechanisms. To maintain static equilibrium in the stationary state,
 1084 a high stiffness value k_{high} is employed. When $S_{\text{open}} = 1$ (door opening phase), the stiffness
 1085 progressively decreases with increasing q . Upon exceeding the critical position $q_{\text{threshold}}$, the stiffness
 1086 reaches its minimum k_{low} , and the joint target position switches to $q_{\text{upper_bound}}$. During door closure,
 1087 as q approaches $q_{\text{threshold}}$ from above, the target position abruptly transitions to $q_{\text{lower_bound}}$, accom-
 1088 panied by an exponential stiffness surge to rapidly complete closure, emulating commercial door
 1089 closer dynamics. This behavior is formalized as:

1090

$$K(q) = \begin{cases} k_{\text{high}}, & q = q_{\text{lower_bound}} \\ k_{\text{high}} - \alpha q, & \text{if } q_{\text{lower_bound}} < q \leq q_{\text{threshold}} \text{ and } S_{\text{open}} = 1 \\ k_{\text{low}} + k_{\text{max}} e^{-\lambda q}, & \text{if } q_{\text{lower_bound}} < q \leq q_{\text{threshold}} \text{ and } S_{\text{open}} = 0 \\ k_{\text{low}}, & q_{\text{threshold}} < q < q_{\text{upper_bound}} \end{cases} \quad (5a)$$

$$(5b) \quad (5c) \quad (5d)$$

1091

1092

1093

1094

1095

1096

1097

F CHART COMPARISON WITH EXISTING DATASETS

1098

1099

We present a detailed comparison of ArtVIP with existing articulated-object datasets in Tab. 6

1100

Table 6: Detailed comparison with existing articulated-object datasets.

1101

1102

1103

1104

1105

1106

1107

G VISUAL REALISM COMPARISON

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

We present further comparative analysis in Fig. 11. PartNet-Mobility employs the URDF format, with meshes stored in OBJ format and material information defined in MTL files. Although the OBJ files are manually crafted, they frequently exhibit distorted meshes, significantly compromising visual quality. The MTL material format inherently lacks the capability to model physically accurate light reflection, resulting in a lack of environmental realism across all PartNet-Mobility assets. Our analysis reveals that many materials in PartNet-Mobility rely solely on base color for rendering, and the absence of textures substantially degrades the overall rendering quality. Although BEHAVIOR-1K adopts the USD format, which supports physically based rendering (PBR), it still suffers from issues related to distorted meshes and poor texture quality.

1118

1119

1120

1121

1122

1123

1124

1125

1126

To mitigate issues such as distorted meshes and angular surfaces, we employed a high number of triangular faces to ensure smooth surfaces and enhanced geometric detail. For categories such as toilets and refrigerators, ArtVIP significantly surpasses BEHAVIOR-1K and PartNet-Mobility in the number of triangular faces utilized. However, this approach entails a trade-off, as it reduces the simulation frame rate. To address this, we conducted profiling analysis to optimize the simulation frame rate for each object. In our experiments, we selected the kitchen, which contains the highest number of articulated objects, and the living room, which features the most extensive texture rendering, as testing environments. Each asset from ArtVIP was individually placed within these scenes, ensuring that the overall rendering frame rate consistently exceeds 60 Hz (i7-13700, Nvidia 4090, 64 GB).

1127

1128

1129

1130

1131

1132

1133

To study the effect of triangle count, we report comprehensive statistics in Tab. 7 for ArtVIP, PartNet-Mobility, and BEHAVIOR-1K: the average triangle count, the average number of active joints, the average FPS with a single asset, and the average FPS in the kitchen scene. The kitchen scene is the most complex environment, containing 65 actuated joints. ArtVIP and PartNet-Mobility are evaluated in Isaac Sim 5.1. BEHAVIOR-1K assets are encrypted and accessible only through OmniGibson (Isaac Sim 4.5). We attribute the large FPS fluctuations observed for BEHAVIOR-1K to overhead introduced by the derivative framework. Based on the FPS results for ArtVIP and PartNet-Mobility, we conclude: 1) For a single object, under Isaac Sim’s iterative optimizations, triangle

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
Table 7: Category-wise averages for triangle count, active joints, and FPS across datasets.

Item Category	Avg triangle count			Avg active joints			Avg FPS (single item)			Avg FPS (in kitchen)		
	ArtVIP	PartNet-Mobility	BEHAVIOR-1K	ArtVIP	PartNet-Mobility	BEHAVIOR-1K	ArtVIP	PartNet-Mobility	BEHAVIOR-1K	ArtVIP	PartNet-Mobility	BEHAVIOR-1K
Coffee Machine	80484.8	27104.7	42256	2.2	5.759	5.5	91.97	91.95	114.77	72.01	73.83	48.33
Microwave	34494.6	8620.5	5521	4	4.313	1.857	91.93	91.88	87.51	69.38	72.95	50.04
Oven	99049	41206.2	25638	4.5	6.133	1	91.98	91.95	109.97	72.66	74.64	49.4
Dishwasher	54427.1	8932.6	25162.2	1.429	1.333	2.5	91.94	91.95	115.2	65.75	69.95	49.73
Rice Cooker	101573.3	26068.7	40245.3	3.333	1.12	1	91.96	91.96	115.58	70.97	74.91	48.29
Laptop	46053.6	37378.7	18546.3	1	1	1	91.97	91.93	112	74.59	74.89	46.99
Washing Machine	151705.4	26269.8	27380.8	2.57	7.471	1.538	91.95	91.94	107.17	70.97	70.74	48.27
Toilet	164271.6	22276.49	15011.1	3.6	2.319	2.611	91.95	91.95	120.58	74.94	74.93	47.18
Refrigerator	100909.1	6317	24273.4	6.25	1.682	1.538	91.94	91.96	99.06	68.38	73.39	49.82
Table	20184.7	22307.6	14210.6	5.85	3.158	2.633	91.96	91.96	116.37	68.13	71.68	47.92
Folding Chair	21567.5	6519.2	7064.6	2	1.231	2	91.95	91.97	125.93	74.5	74.92	51.41
Scissors	43953	14601	4972	2	1.963	1	91.96	91.96	129.71	72.5	74.92	52.45
Trash Can	30139.6	6468.33	8370.17	1.77	1.971	1	91.94	91.93	121.28	71.66	74.93	52.24

counts up to approximately 100k and up to 20 active joints have negligible impact on FPS. 2) In complex scenes, both triangle count and the number of active joints reduce FPS.

H PHYSICAL FIDELITY AND INTERACTION EVALUATIONS

Motion Triggered by Latch Release. To validate the modular interaction within assets, we compared the triggered joint in both real-world and virtual microwave. We conducted button-press experiments in each environment to initiate the door-opening action and recorded the resulting door motion trajectories. In the real-world tests we tracked a marker on the door using the optical tracking system to capture its spatial motion after the button pressed. In the simulation we set a virtual marker at the same position as the real-world marker on the door, and we triggered the door opening via pressing the button as well (for which the activation configured in modular interaction) and logged the virtual marker’s trajectories. We performed ten trials in each environment and computed the average spatial trajectory as Fig. 12 shown.

Motion Triggered by Joint Position Threshold. Appliances equipped with door closers typically exhibit a dynamic change in motion once the door reaches a certain angle during closing. After arriving at a certain angle, the door closer causes the door to accelerate and snap shut against the appliance body. To evaluate how well the simulation captures this physical transition, we focus on analyzing the door’s linear and angular velocities during the transition from the threshold state to full closure. In both the simulation and real-world experiments, a force of no more than 1.0 N is applied when the door is within the threshold range to trigger the door closer mechanism. We then record the kinematic behavior following the activation of the door closer. In the real-world setup, the optical motion capture system is used to track the spatial displacement of markers on the door. Both the simulation and real-world experiments are repeated ten times, and we compute the average spatial trajectories and changes in velocity along the X-axis for quantitative comparison (Fig. 13).

I IMITATION LEARNING APPLICATION

Task Summary. As shown in Fig. 14, we design four challenging articulated-object manipulation tasks: (1) **PullDrawer**, (2) **OpenCabinet**, (3) **SlideShelf**, and (4) **CloseOven**. These tasks demand precise and flexible motions, including rotation, angled pushing, and horizontal translation. We define these tasks as follows:

- **PullDrawer.** This task requires the robot to insert the gripper into the handle of the drawer, securely press the handle, and gradually pull the drawer out along a linear trajectory using a smooth and consistent motion.
- **OpenCabinet.** For this task, the robotic arm needs to precisely locate the thin vertical handle of the cabinet door. The gripper has to align vertically, firmly grip the handle, and pull the door outward along a curved path while maintaining a stable trajectory.
- **SlideShelf.** This task involves horizontal manipulation of the shelf. First, the gripper needs to rotate around 90 degrees to align parallel to the shelf’s direction. It then grips the base of the shelf and moves horizontally, pulling the shelf out along its guide rails in a stable and controlled manner.
- **CloseOven.** To complete this task, the robotic arm needs to close its gripper to push against the bottom edge of the oven door. The arm then rotates and lifts under the door, applying a curved upward force to close the door.

1188

1189

1190

1191 **PartNet-Mobility**

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

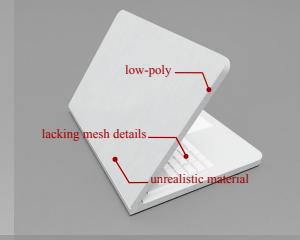
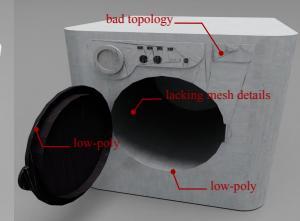
Behavior-1k**ArtVIP****Low****Visual Realism****High**

Figure 11: Comparisons of ArtVIP, BEHAVIOR-1K, and PartNet-Mobility.

1242

1243

1244

1245

1246

1247

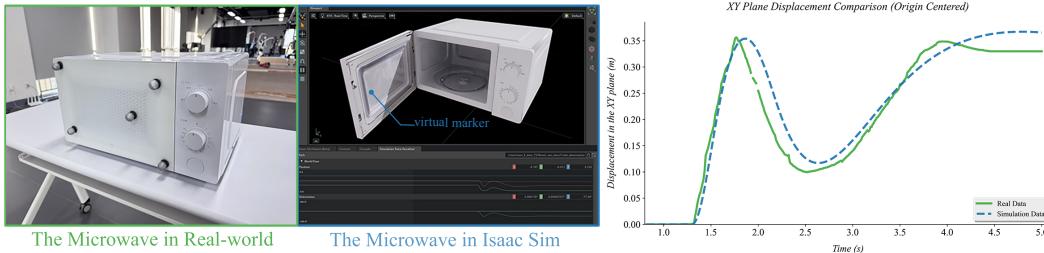


Figure 12: **Left and Middle:** Digital-twin asset examples in real-world and simulation. **Right:** Analysis of the Microwave’s displacement.

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

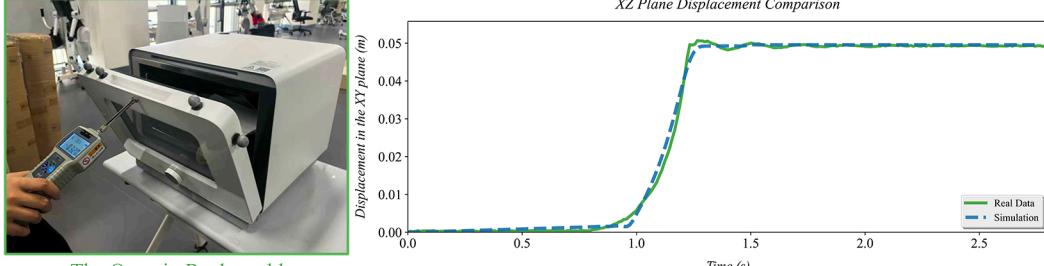
1266

1267

1268

1269

1270



The Oven in Real-world

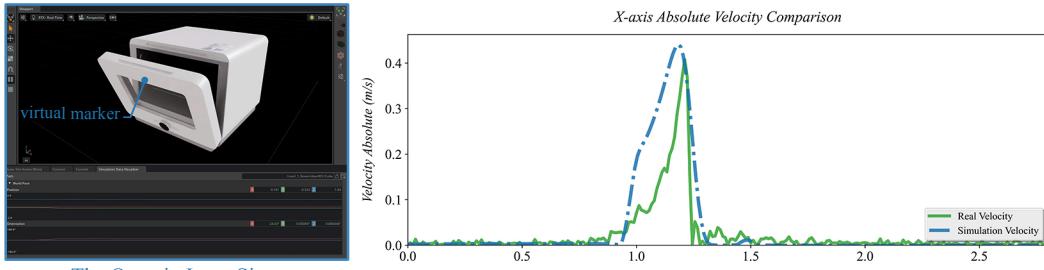


Figure 13: **Left:** Digital-twin asset examples in real-world and simulation. **Right:** Analysis of the oven’s displacement.

1291

1292

1293

1294

1295

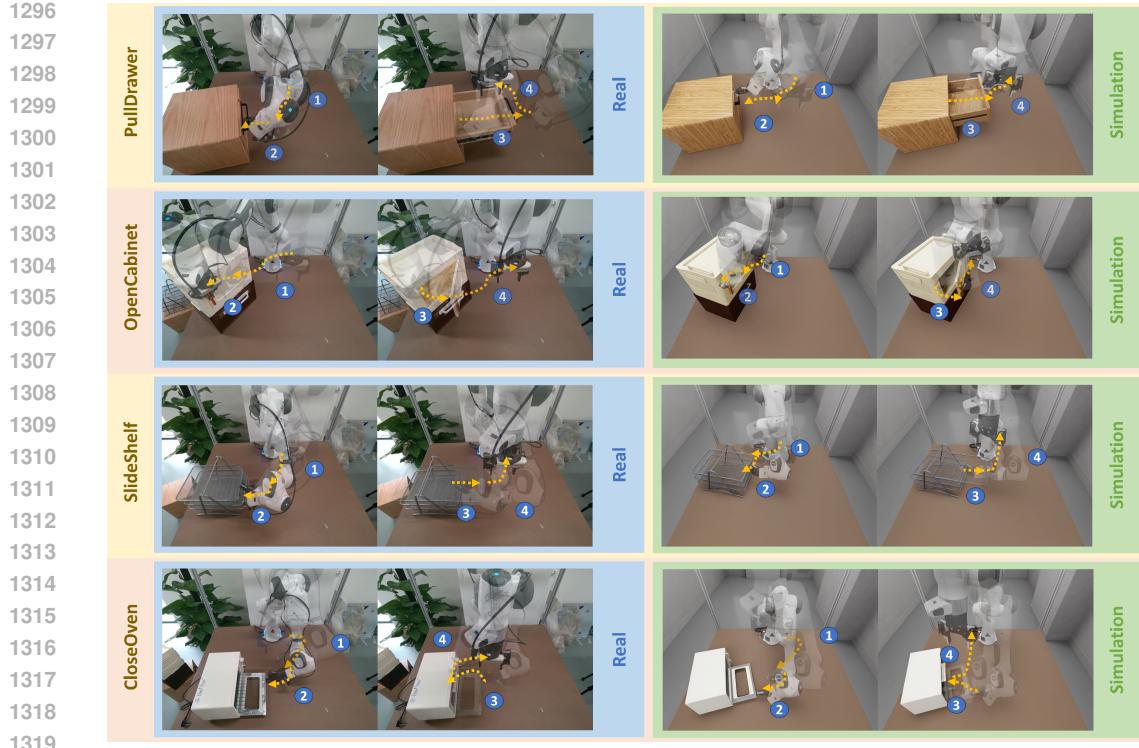


Figure 14: The four articulated-object manipulation tasks conducted for imitation learning.

	Hyperparameter	Value		Hyperparameter	Value
Training	Batch size	48	Network Architectures	Encoder layer	4
	Learning rate	1e-4		Decoder layer	7
	Optimizer	AdamW		Forward dim	3200
	KL weight	10		Heads num	8
	Action sequence	50		Transformer hidden dim	512
	Training step	50k		Backbone	ResNet50

Table 8: Implementation details of Action Chunking Transformer (ACT).

1338 Imitation Learning Algorithm. The input to the imitation learning models consists of RGB image
 1339 data from multiple camera views and the robot’s proprioceptive states. The output is the robot control
 1340 signals, such as joint positions, enabling end-to-end task execution. We used two state-of-the-art
 1341 imitation learning methods, Action Chunking Transformer (ACT) (Zhao et al., 2023) and Diffusion
 1342 Policy (DP) (Chi et al., 2023), to train the robotic policies for the articulated object manipulation
 1343 task. Hyperparameters of both methods are demonstrated in Tab. 8 and Tab. 9.

- **Action Chunking Transformer (ACT)** (Zhao et al., 2023): ACT is built on the transformer network architecture and leverages temporal ensemble techniques to produce fluid and precise action sequences.
- **Diffusion Policy (DP)** (Chi et al., 2023): DP employs a diffusion-based generative model that captures multi-modal action distributions, offering robustness and high success rates for complex robotic tasks.

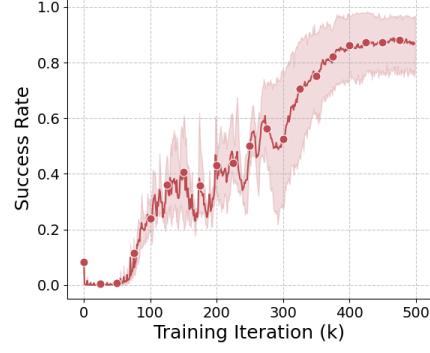
Detailed Experiment Results. The Full experiment results are presented in Tab. 10.

	Hyperparameter	Value		Hyperparameter	Value
Training	Batch size	48	Network Architectures	Diffusion Network	Unet1D
	Learning rate	1e-4		Pooling	SpatialSoftmax
	Optimizer	AdamW		Noise scheduler	DDIM
	EMA power	0.75		EMA model	True
	Action sequence	16		Noise schedule	SquaredcosCap
	Training step	50k		Backbone	ResNet50

Table 9: Implementation details of Diffusion Policy (DP).



(a) Training task.



(b) Training curve over five random seeds.

Figure 15: RL-based training of visuomotor policy with ArtVIP.

J REINFORCEMENT LEARNING APPLICATION

Training Details. We extend the visual RL framework EAGLE (Zhao et al., 2025) to articulated-object tasks in ArtVIP. Fig. 15a shows the CloseTrashcan task, where the robot arm is required to close the trashcan within a given time limit. EAGLE is a two-stage visual RL framework designed for efficiency and generalization. In Stage 1, the teacher policy receives low-level states, including the robot arm’s proprioceptive input, the lid’s joint value, and the 3D relative position between the trashbin and the gripper. In Stage 2, the student policy is provided only with the wrist camera image and the robot’s proprioceptive state—no object-related states are available. Fig. 15b presents the training curves in Stage 2.

For implementation details, in Stage 1, we replace EAGLE’s original RL agent with PPO; In Stage 2, a privileged-state teacher is distilled into a visuomotor student while a self-supervised attention mask learned as follows:

$$\mathcal{L}_{att} = \mathcal{L}_{rec} + \mathcal{L}_{ae} + \beta \mathcal{L}_{ctl} + \lambda \mathcal{L}_{sps}, \quad (6)$$

where \mathcal{L}_{rec} and \mathcal{L}_{ae} are reconstruction losses, \mathcal{L}_{ctl} predicts dynamics, and \mathcal{L}_{sps} enforces mask sparsity. Hyper-parameters β and λ weight auxiliary losses.

The student policy is trained with the distillation loss:

$$\hat{\mathcal{L}}(\pi_\theta) = \mathbb{E}_{(\mathbf{o}, \mathbf{s}) \sim \mathcal{D}} [\|\pi_\theta(\mathbf{o}_{aug}) - \pi_e(\mathbf{s})\|_2^2], \quad (7)$$

where \mathbf{s} contains privileged states and \mathbf{o}_{aug} are images augmented by the learned mask with Eqn. equation 6. Hyper-parameters used in EAGLE are listed in Tab. 11.

Reward Functions. The **CloseTrashcan** task is a long-horizon challenge requiring the robot to first approach the trashcan lid and then close it smoothly. To facilitate efficient RL training, we design a multi-objective reward function as follows:

$$r_t(\mathbf{s}_t, \mathbf{a}_t) = \lambda_1 r_{dst}(\mathbf{s}_t) + \lambda_2 r_{dir}(\mathbf{s}_t) + \lambda_3 r_{cls}(\mathbf{s}_t) + \lambda_4 r_{smth}(\mathbf{a}_t), \quad (8)$$

where r_{dst} rewards proximity between the gripper and the lid, r_{dir} encourages alignment toward the lid, r_{cls} measures lid closure progress, and r_{smth} promotes smooth actions. The reward weights are set as: $\lambda_1 = 0.5$, $\lambda_2 = 0.125$, $\lambda_3 = 10$, $\lambda_4 = -0.01$.

Table 10: Performance results (scheme A): per-seed scores and mean \pm 90% CI

1404	Task	Method	Strategy	seed1	seed2	seed3	Mean \pm CI90
1405	PullDrawer	ACT	RO	0.567	0.767	0.600	0.644 \pm 0.059
1406			SO	0.433	0.433	0.300	0.389 \pm 0.060
1407			RSM100+10	0.500	0.667	0.767	0.640 \pm 0.059
1408			RSM100+20	0.667	0.600	0.767	0.678 \pm 0.057
1409			RSM100+50	0.833	0.767	0.733	0.778 \pm 0.051
1410			RSM100+100	0.767	0.867	0.800	0.811 \pm 0.048
1411	DP	DP	RO	0.600	0.733	0.633	0.656 \pm 0.058
1412			SO	0.133	0.233	0.233	0.200 \pm 0.049
1413			RSM100+10	0.600	0.650	0.700	0.650 \pm 0.057
1414			RSM100+20	0.650	0.700	0.733	0.694 \pm 0.056
1415			RSM100+50	0.700	0.733	0.750	0.728 \pm 0.055
1416			RSM100+100	0.733	0.767	0.867	0.789 \pm 0.050
1417	OpenCabinet	ACT	RO	0.300	0.400	0.333	0.344 \pm 0.058
1418			SO	0.167	0.100	0.100	0.122 \pm 0.040
1419			RSM100+10	0.333	0.367	0.367	0.356 \pm 0.059
1420			RSM100+20	0.367	0.400	0.367	0.378 \pm 0.059
1421			RSM100+50	0.433	0.500	0.400	0.444 \pm 0.061
1422			RSM100+100	0.567	0.367	0.433	0.456 \pm 0.061
1423	DP	DP	RO	0.467	0.500	0.500	0.489 \pm 0.061
1424			SO	0.133	0.033	0.133	0.100 \pm 0.037
1425			RSM100+10	0.500	0.533	0.567	0.533 \pm 0.058
1426			RSM100+20	0.550	0.583	0.600	0.578 \pm 0.057
1427			RSM100+50	0.600	0.617	0.633	0.617 \pm 0.057
1428			RSM100+100	0.667	0.700	0.600	0.656 \pm 0.058
1429	SlideShelf	ACT	RO	0.233	0.233	0.333	0.267 \pm 0.054
1430			SO	0.100	0.167	0.133	0.133 \pm 0.042
1431			RSM100+10	0.200	0.267	0.300	0.256 \pm 0.053
1432			RSM100+20	0.233	0.300	0.267	0.267 \pm 0.054
1433			RSM100+50	0.300	0.367	0.300	0.322 \pm 0.057
1434			RSM100+100	0.333	0.333	0.400	0.356 \pm 0.059
1435	DP	DP	RO	0.467	0.433	0.433	0.444 \pm 0.061
1436			SO	0.167	0.167	0.200	0.178 \pm 0.047
1437			RSM100+10	0.433	0.467	0.500	0.467 \pm 0.058
1438			RSM100+20	0.500	0.533	0.550	0.528 \pm 0.057
1439			RSM100+50	0.533	0.567	0.583	0.561 \pm 0.056
1440			RSM100+100	0.567	0.600	0.600	0.589 \pm 0.060
1441	CloseOven	ACT	RO	0.500	0.633	0.600	0.578 \pm 0.061
1442			SO	0.267	0.267	0.167	0.233 \pm 0.052
1443			RSM100+10	0.500	0.600	0.667	0.589 \pm 0.060
1444			RSM100+20	0.533	0.633	0.633	0.600 \pm 0.060
1445			RSM100+50	0.733	0.533	0.700	0.656 \pm 0.058
1446			RSM100+100	0.667	0.800	0.567	0.678 \pm 0.057
1447	DP	DP	RO	0.600	0.700	0.667	0.656 \pm 0.058
1448			SO	0.267	0.233	0.333	0.278 \pm 0.055
1449			RSM100+10	0.633	0.667	0.700	0.667 \pm 0.057
1450			RSM100+20	0.667	0.700	0.733	0.700 \pm 0.056
1451			RSM100+50	0.700	0.733	0.750	0.728 \pm 0.055
1452			RSM100+100	0.767	0.733	0.833	0.778 \pm 0.051
1453							
1454							
1455	Baseline Comparison. To put EAGLE’s performance in context, we compare it with a vision-based PPO method. As shown in Tab. 12, due to the high computation complexity and low data diversity, the baseline performs poorly on the CloseTrashcan task, while EAGLE achieves a 98% success rate after 500k training iterations.						

	Hyperparameter	Value
1460	Learning rate for all net	5e-4
1461	Optimizer	Adam
1462	Teacher	Batch size
1463	(Stage 1)	12 × 4096
1464		Discount factor
1465		0.99
1466		Clip ratio
1467		0.2
1468		Rollout size
1469		96 × 4096
1470		Observation
1471		128 × 128
1472		Learning rate for all net
1473		1e-4
1474		Optimizer
1475		Adam
1476		Batch size
1477		256
1478		Frame stack
1479		1
1480		Replay buffer size
1481		100k
1482		λ
1483		0.01
1484		β
1485		0.5
1486		α in <i>random overlay</i>
1487		linear schedule from 0.4 to 0.9

Table 11: Hyperparamters for EAGLE.

Table 12: EAGLE vs. vision-based PPO: success rate across training checkpoints (k).

Checkpoint (k)	100	200	300	400	500
EAGLE	0.23	0.28	0.73	0.85	0.98
Vision-based PPO	0.16	0.19	0.21	0.22	0.24

K GENERATIVE PIPELINES CAN'T BE INTEGRATED INTO SIMULATION WORKFLOWS

To evaluate the quality of generated assets, we reproduced SplArt (Lin et al., 2025), a state-of-the-art generative articulated-object method, and generated assets for a two-drawer cabinet and a side-by-side refrigerator. As shown in Fig. 16, the outputs are low quality compared with our digital-twin assets. The generated assets exhibit the following issues:

- Self-collisions between parts
- Severe mesh distortions and breakage
- Incorrect joint limits, positions, and axes
- Materials and colors deviating markedly from reality
- Severe lack of interior details

These issues trigger runtime errors in Isaac Sim. Consequently, current generative baselines fail to produce simulation-ready articulated assets for robotics.

L THE USE OF LARGE LANGUAGE MODELS.

A large language model (LLM) was used strictly as a writing aid for language polishing (grammar, clarity, and style). All ideas, methodological designs, datasets, code, analyses, and results are original and solely produced by the authors.

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533



(a) real cabinet



(b) cabinet reconstructed by SplArt



(c) cabinet from ArtVIP



(d) real refrigerator



(e) fridge reconstructed by SplArt



(f) fridge from ArtVIP

1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

Figure 16: Comparison of real-world objects, generated outputs, and digital-twin assets.