

TOWARDS PHYSICALLY GROUNDED ARTICULATED 3D ASSETS: A MULTI-AGENT PIPELINE AND BENCHMARK

Xudong Cai^{1#*} Ruijie Zhu^{1*} Haoliang Shang^{1,2*} Ruining Li^{3*}
 Marc Pollefeys^{2,4} Shangzhe Wu⁵ Andrea Vedaldi³ Chuanxia Zheng¹

¹PVG, Nanyang Technological University ²ETH Zurich
³VGG, University of Oxford ⁴Microsoft ⁵University of Cambridge
[physicalvision.github.io/artisan](https://github.com/physicalvision/artisan)




Figure 1: **Physically grounded articulated 3D objects generated by ARTISAN** (assembled into scenes). **Orange** and **blue** annotations indicate revolute and prismatic axes, respectively. Insets list the physical attributes composition of individual parts for examples objects. **Zoom in for more details.**

ABSTRACT

In this paper, we study *physically grounded articulated 3D asset generation*—moving the goal from reproducing how objects look to capturing how they move and function. Inspired by the literature on procedural and agentic generation, we present ARTISAN, a system that extends the latest method Articraft by introducing a *multi-agent* pipeline with *physical simulation feedback*, enabling comprehensive attributes beyond geometry, appearance, and physically plausible articulation, such as collision-free motion, stability under gravity, and accurate mass and material properties. Furthermore, we construct a unified benchmark spanning four key aspects: spatial fidelity, physical plausibility, downstream task performance, and overall quality. The benchmark encompasses 100 reference assets from ArtVIP and LightWheel, 50 stability prompts, and 50 ABO objects with reference masses, with protocols evaluating various metrics. Through extensive evaluation of more than 10 representative models, ARTISAN achieves competitive spatial fidelity while attaining the lowest collision rates and highest stability rates among the evaluated methods. Together, our system and benchmark establish insights for advancing 3D generation towards physically plausible, functionally usable assets.

Project lead. * Core technical contribution.

1 INTRODUCTION

We consider the problem of generating physically grounded articulated 3D assets from natural language and, optionally, reference images. Specifically, we aim to reconstruct the objects’ appearance and geometry, as well as structural, functional, and physical properties such as part hierarchies, joints, and mass. The intent is to capture the configuration space of the object and to support simulation of contact and gravity. By solving this problem, we can reduce the cost of creating assets for physics simulation (Genesis Authors, 2024; Genesis AI Team, 2026), product design (Willis et al., 2021), and robotics (Li et al., 2023; Nasiriany et al., 2026).

Appearance, shape, articulation, and physical attributes are interdependent and should be modeled *holistically*. However, most existing methods address these properties separately, focusing on shape and appearance (Xiang et al., 2025; Hunyuan3D et al., 2025; Wu et al., 2025c), articulation (Wu et al., 2026; Li et al., 2026a;b), or specific physical properties, such as materials (Wu et al., 2025b; Youwang et al., 2026), mass (Zhai et al., 2024; Cao et al., 2026a;b; Li et al., 2026c), stability (Chen et al., 2024b; Li et al., 2025), and collision-free motion (Luo et al., 2025; Kreber & Stueckler, 2025). Addressing each aspect in isolation makes it difficult to satisfy all of them, resulting in artifacts and defects in the reconstructions and ultimately limiting the usability of the generated assets.

As alternatives to these learned models, procedural generators (Lian et al., 2025; Joshi et al., 2025) assemble assets programmatically, based on predefined rules that enforce desirable constraints and make at least some physical properties explicit and controllable. However, such rule-based systems limit the types of objects that can be generated and are difficult to use in tasks like image-based reconstruction. This limitation is addressed by several recently proposed agentic methods. For instance, Artcraft (Zhou et al., 2026) uses a Large Language Model (LLM) to write, validate, and revise programs that construct 3D objects, which is much more flexible than using canned rules of traditional procedural generators. However, Artcraft and other similar approaches use a single agent to both construct and verify the ‘3D program’, which can lead to *self-evaluation bias*, where the agent may satisfy its own checks without genuinely ensuring the quality of the generated assets. Moreover, while their checks are meant to increase the likelihood that the generated assets would work in a physics-simulated environment, the latter is not directly tested and may still fail.

In this paper, we address these gaps by introducing ARTISAN, a new pipeline that improves on the blueprint of Artcraft (Zhou et al., 2026) by switching to a *multi-agent* design and introducing *physical simulator feedback*. The new pipeline (Figure 2) consists of: **(i)** the *Design Agent*, which produces a blueprint of the object specifying its parts, joints, and articulations; **(ii)** the *Coding Agent*, which translates the blueprint into a program that outputs the geometry, articulations, and physical parameters of the object; and **(iii)** the *Review Agent*, which validates the built asset and provides feedback to the Coding Agent to revise the program. The agents work *independently*, exchanging only the blueprint, 3D asset, and feedback as explained. Their independence reduces the self-evaluation bias mentioned above. Furthermore, the validation step is augmented with feedback obtained from a physical simulator, which verifies physical properties like collisions and stability within the target simulator environment.

Existing benchmarks offer limited assessment of whether visual fidelity translates into physical plausibility and even functional usability to evaluate the assets produced by such pipelines. Our second key contribution is therefore to introduce a new comprehensive benchmark that jointly assesses all of these properties. It evaluates: *spatial fidelity*, comparing the object’s appearance, geometry, and kinematics to the reference inputs; *physical attributes*, measuring mass, material, stability under gravity, and collision under motion; *downstream tasks*, assessing the usability of the generated assets in tasks like robotic manipulation; and *overall quality*, based on human judgments after manipulating objects in the simulator. Each aspect is evaluated with multiple metrics.

Using the new benchmark, we conduct extensive experiments to compare our method with state-of-the-art models. Empirically, we find that the assets constructed by ARTISAN: **(i)** are faithful to the input references; **(ii)** are physically valid and articulate correctly; and **(iii)** work well in downstream tasks such as simulated robotic manipulation. We also show that ARTISAN outperforms the recent and strong models in most metrics and in human evaluation, especially improving physical properties like collision and stability. In most cases, our general-purpose pipeline is more accurate than task-specific feed-forward 3D reconstruction models. We also show that the constructed objects are functional and support downstream robotic manipulation tasks.

2 RELATED WORK

Creating Functional 3D Assets. Early methods build articulated objects through per-instance optimization (Mu et al., 2021; Wu et al., 2022; Tseng et al., 2022; Liu et al., 2023; Weng et al., 2024; Wu et al., 2025a), but are limited to specific instances. To improve generalizability, learning-based methods create articulated objects from pure noise (Lei et al., 2023; Gao et al., 2025), text (Liu et al., 2024; Su et al., 2025), images (Liu et al., 2025; Jiang et al., 2022; Chen et al., 2024a; He et al., 2026; Chen et al., 2025; Huang et al., 2026b; Wu et al., 2026), videos (Delitzas et al., 2026) or 3D models (Li et al., 2026a; Zhang et al., 2026; Li et al., 2026b), learning from large-scale datasets (Xiang et al., 2020; Li et al., 2026b). Recent methods focus on inferring independent physical properties from visual representations (Zhai et al., 2024; Shuai et al., 2025), jointly predicting these alongside joint motion characteristics (Cao et al., 2026a;b; Li et al., 2026c; Huang et al., 2026a), and leveraging simulation feedback to refine aspects such as support, clearances, or mobility (Chen et al., 2024b; Li et al., 2025; Luo et al., 2025; Kreber & Stueckler, 2025; Xu et al., 2026). In contrast, ARTISAN generates articulated objects that take diverse physical properties into account.

Procedural and Agentic 3D Asset Creation. Procedural systems construct editable and simulation-ready variations within the authored rules (Lian et al., 2025; Joshi et al., 2025). The latest agent-based pipelines extend this coverage through executable articulation, retrieval with visual critique, or semantic search over kinematic structures (Mandi et al., 2025; Le et al., 2025; Qiu et al., 2025; Wang et al., 2025). Closer to our approach, Articraft (Zhou et al., 2026) builds articulated objects through an agent with Software Development Kit (SDK)s. Recent systems improve reliability through geometric and visual checks, and revisions to earlier construction stages when checks fail (Wang et al., 2026; Shui et al., 2026; Lin et al., 2026). In contrast, ARTISAN introduces a multi-agent pipeline with physical simulator feedback, enabling physically grounded articulated 3D asset generation.

Datasets and Benchmarks for Physically Grounded 3D Assets. Articulated datasets provide part geometry and motion annotations (Xiang et al., 2020; Geng et al., 2023; Liu et al., 2022; Delitzas et al., 2024; Jin et al., 2026; Engelbracht et al., 2026). Program-generated collections further expand the available assets (Zhou et al., 2026; Cao et al., 2026b; Li et al., 2026c). Existing benchmarks separately assess geometric and semantic fidelity (Li et al., 2026a; Duggal et al., 2025), programmatic modelling quality (Yang et al., 2026; Gao et al., 2026), material properties (Cao et al., 2026b; Li et al., 2026c), and motion behavior (Luo et al., 2025; Pham, 2026). Our benchmark comprehensively evaluates the spatial fidelity, physical plausibility, and downstream task utility of generated articulated assets using a set of distinct metrics.

3 METHOD: A MULTI-AGENTIC PIPELINE WITH PHYSICAL VALIDATION

Our goal is to construct a system $\mathcal{P}$ that, given as input a natural language prompt p and, optionally, one or more reference images $\mathcal{I}$, generates as output an articulated 3D asset $\mathcal{O}$:

$$\mathcal{P} : (p, \mathcal{I}) \mapsto \mathcal{O} = \{\{\mathcal{G}_i, \rho_i, \mathcal{M}_i\}_{i=1}^P, (\mathcal{K}, \mathcal{J})\}. \quad (1)$$

The object $\mathcal{O}$ consists of P rigid parts defined by their geometry $\mathcal{G}_i$, mass ρ_i , and material $\mathcal{M}_i = \{c_i, m_i, r_i, \alpha_i\}$ specified by its base color c_i , metalness m_i , roughness r_i and opacity α_i (i.e., standard Physically Based Rendering (PBR) parameters).¹ $\mathcal{K}$ specifies the *kinematic structure*, defining the parent-child relationships among the P parts, while $\mathcal{J}$ specifies the joint types, axes, motion limits, and any coupling constraints between dependent joints. The pipeline aims to output an object $\mathcal{O}$ that faithfully reproduces the inputs p and $\mathcal{I}$ and is physically sound: parts should not collide, freestanding objects should remain stable under gravity, and it must be possible to grasp and manipulate it. The pipeline consists of three independent LLM agents equipped with tools and execution harnesses, $\mathcal{P} = \{\mathcal{A}_i\}_{i=1}^3$, described below, with implementation details given in Section A.

3.1 DESIGN AGENT

The *Design Agent* $\mathcal{A}_{\text{design}}$ takes the prompt and image as input, and outputs a structured textual blueprint $\mathcal{B}$ and a design board $\mathcal{D}$:

$$\mathcal{A}_{\text{design}}(p, \mathcal{I}) \rightarrow (\mathcal{B}, \mathcal{D}), \text{ where } \mathcal{B} = \{\{\tilde{\mathcal{G}}_i^{\text{text}}, \tilde{\mathcal{M}}_i^{\text{text}}\}_{i=1}^P, (\tilde{\mathcal{K}}^{\text{text}}, \tilde{\mathcal{J}}^{\text{text}}), \mathcal{Q}^{\text{text}}\}. \quad (2)$$

¹The SDK supports distinct PBR materials with optional textures for components within each part.




Figure 2: The staged pipeline of the three agents: Design Agent, Coding Agent, and Review Agent.

The superscript $*^{\text{text}}$ indicates that *these are all textual descriptions* rather than the final actual representations. Text $\tilde{G}_i^{\text{text}}$ and $\tilde{M}_i^{\text{text}}$ describe the geometry and material of part i and $\tilde{K}^{\text{text}}$ and $\tilde{J}^{\text{text}}$ describe the kinematic structure and joints between these parts of the asset.

Beyond geometry, materials, and articulation considered in prior work (Zhou et al., 2026), Q^{text} describes novel *functional* properties. First, it specifies coupled motions, such as a bin lid opening when its pedal is pressed. It also defines interaction tasks in plain language, such as “*raise the mixing head and lift out the mixing bowl*”. These functional constraints guide asset verification in Section 3.3 and the generation of downstream task annotations in Section 5.2, helping ensure that the asset is both physically plausible and functionally usable in downstream applications.

Blueprint Verification and Visual Reference. $\mathcal{A}_{\text{design}}$ checks the blueprint against the input requirements and corrects inconsistencies in parts, motions, and functions. An image-generation model (OpenAI, 2026b) converts the blueprint into a *design board* $\mathcal{D}$ (Figure 2). It lays out orthographic views of the articulated object, an exploded view of the parts, a kinematic diagram of the moving parts in their closed and open states, and cutaway sections of the parts whose interiors matter.

3.2 CODING AGENT

The *Coding Agent* $\mathcal{A}_{\text{coding}}$ takes as input the user prompt and the optional reference images, the blueprint $\mathcal{B}$ and design board $\mathcal{D}$ from the Design Agent, the program Π_{t-1} and the corresponding 3D asset $\mathcal{O}_{t-1}$ from its previous iteration, the physics report $\mathcal{E}_{t-1}$, and the review scores $\mathcal{S}_{t-1}$ and list of issues $\mathcal{R}_{t-1}$ from the Review Agent, and produces the next program Π_t and asset $\mathcal{O}_t$:

$$\mathcal{A}_{\text{coding}} \left(\underbrace{(\mathcal{P}, \mathcal{I})}_{\text{Input}}, \underbrace{(\mathcal{B}, \mathcal{D})}_{\mathcal{A}_{\text{design}}}, \underbrace{(\Pi_{t-1}, \mathcal{O}_{t-1})}_{\mathcal{A}_{\text{coding}}}, \underbrace{(\mathcal{E}_{t-1}, (\mathcal{S}_{t-1}, \mathcal{R}_{t-1}))}_{\mathcal{A}_{\text{review}}} \right) \rightarrow (\Pi_t, \mathcal{O}_t), \quad (3)$$

Here, t indexes a coding–review iteration. The previous program Π_0 , asset $\mathcal{O}_0$, physics report $\mathcal{E}_0$ (from the Physics Simulator in Section 3.4), and review feedback $\mathcal{S}_0$ and $\mathcal{R}_0$ (returned by the Review Agent $\mathcal{A}_{\text{review}}$, described in Section 3.3) are absent on the first iteration.

Building on Artcraft’s SDK and execution harness (Zhou et al., 2026), the agent iteratively *edits* the program Π_t (written in PYTHON), *executes* it to obtain a corresponding asset $\mathcal{O}_t$, and then *estimates* its quality for iterative improvement. We refer readers to Zhou et al. (2026) for more details. We extend the workflow for physical plausibility with three components: (i) *tools* for geometric and visual inspection (detailed in Section A.2), (ii) reusable *standard connectors* for joint connections, and (iii) *physical simulation feedback* (Section 3.4) that guides repairs to physical attributes.

Standard Connectors for Joint Connections. Inspired by standard connectors in mechanical design (Poli, 2001), we extend the Artcraft SDK with reusable, parameterized connectors, such as hinges, slide rails, and rotary connectors (Figure 3). $\mathcal{A}_{\text{coding}}$ only needs to select a connector




Figure 3: Illustrations of the six families of parameterized standard parts.

and adjust its dimensions and mounting positions, instead of designing a geometrically accurate connector from scratch. This reduces floating connections and unintended interpenetration, making the generated asset more collision-free as measured in Section 3.4.

Action Annotation. Let τ denote a task described in natural language in $\mathcal{Q}^{\text{text}}$. For each task τ , the Coding Agent first interprets the task to determine the required actions, and then the Coding Agent records the target rigid part ℓ_τ , the geometric components forming its grasp or press region V_τ , the joint to actuate j_τ , and the fraction of joint travel required for success σ_τ . These are written to a manifest beside the exported meshes so that a downstream pipeline needs no manual labelling.

3.3 REVIEW AGENT

The *Review Agent* $\mathcal{A}_{\text{review}}$ evaluates the asset against the intended design and physical constraints. It takes as input the user prompt and optional images, the blueprint $\mathcal{B}$ and design board $\mathcal{D}$, the current program Π_t and asset $\mathcal{O}_t$, and the corresponding physics report $\mathcal{E}_t$ from Section 3.4 and outputs a set of scores $\mathcal{S}_t$, one for each evaluation criterion, as well as a structured list of issues $\mathcal{R}_t$:

$$\mathcal{A}_{\text{review}} \left(\underbrace{(\mathcal{p}, \mathcal{I})}_{\text{Input}}, \underbrace{(\mathcal{B}, \mathcal{D})}_{\mathcal{A}_{\text{design}}}, \underbrace{(\Pi_t, \mathcal{O}_t)}_{\mathcal{A}_{\text{coding}}}, \mathcal{E}_t \right) \rightarrow (\mathcal{S}_t, \mathcal{R}_t). \quad (4)$$

The Review Agent critically reviews repairs against the blueprint and visual references to check that intended functions are preserved. For example, $\mathcal{A}_{\text{coding}}$ may shorten a door’s travel to pass collision checks, while $\mathcal{A}_{\text{review}}$ flags the reduced opening range as inconsistent with the blueprint.

Each issue identifies the affected component, supporting evidence, severity, and what the Coding Agent needs to fix. The pipeline uses the scores and issue severity to request repairs when necessary; critical structural or functional failures cannot be offset by a high overall score. Note that the feedback $(\mathcal{S}_t, \mathcal{R}_t)$ guides the next coding iteration, but the Review Agent itself does not alter the asset program.

3.4 PHYSICAL VALIDATION

We validate each candidate asset $\mathcal{O}_t$ through collision and stability checks in MuJoCo (Todorov et al., 2012). The resulting physics report $\mathcal{E}_t$ records colliding part pairs and penetration depths at the tested joint configurations, together with tilt measurements under gravity, to guide coding and review.

Collision under Motion. The collisions are determined based on whether different parts interpenetrate during the simulated motion sequences, across sampled configurations of single-joint, double-joint, and multi-joint motions (details in Appendix A.3). Specifically, we first perform convex decomposition (Wei et al., 2022) to construct mesh collision geometry for simulation. Then, we simulate the motion of each joint in isolation and record any collisions between parts. Collisions that persist during individual joint operation require remediation. Collision findings across motions include the affected part pairs, poses, and penetration depths, enabling the Review Agent to assess whether these collisions occur during normal operation (Section 3.3).

Stability under Gravity. We simulate the assembly under gravity and measure its tilt angle relative to the initial pose to assess whether it remains upright. We run the MuJoCo simulation for up to 5 seconds under gravity. The 3D asset is considered unstable if it tilts beyond a certain threshold during the simulation. The tilt measurements are recorded in $\mathcal{E}_t$ for review against the intended use.

4 BENCHMARK

A key contribution of our paper is to build a comprehensive benchmark for physically grounded articulated 3D assets. To this end, we design a set of evaluation protocols and metrics, and curate benchmark data with abundant annotations for these metrics.

4.1 EVALUATION PROTOCOL AND METRICS

Spatial Fidelity. The spatial fidelity metrics evaluate how well the 3D shape of the generated asset matches the reference asset in both the rest-pose and the articulated states. We first perform point sampling and align the meshes using ICP registration (Besl & McKay, 1992), then establish correspondences between parts based on their centers using the Hungarian matching algorithm (Kuhn, 1955). Due to the difficulty of clearly distinguishing between states, outputs from different states may sometimes be confused. To avoid erroneous penalties, we test both state pairings and retain the one with the lower alignment error across the two states as the final aligned result.

Appearance: After aligning the generated 3D asset to the reference asset, we render both using Blender at 1024^2 resolution from three viewpoints per state. We report the Intersection over Union (IoU) of the rendered masks, pixel-level PSNR, patch-level SSIM (Wang et al., 2004), feature-level LPIPS (Zhang et al., 2018), and CLIP score (Radford et al., 2021) over all six images.

Rest-Pose and Articulated Geometry: We calculate 3D metrics in both states. We report the object Chamfer Distance (oCD) for the object as a whole. We also report the part-aware CD (pCD) and mean part IoU (pIoU) by averaging the metrics of all matched parts, following (Li et al., 2026a).

Physical Properties. Except for spatial fidelity on appearance, geometric and kinematics, we also assess the physical properties of the generated assets, including the reconstructed *materials* and *mass*, the absence of *collisions* during articulated motion, and *stability* under gravity.

Materials: Following PBR-NeRF (Wu et al., 2025b) and LMRM (Youwang et al., 2026), we evaluate the materials of the generated assets by rendering the colour, metallic, and roughness buffers from the same viewpoints used for appearance evaluation, and then calculating PSNR against corresponding renders of the reference asset. To mitigate the impact of geometric inconsistencies, we also develop a protocol to transfer predicted materials to the ground-truth geometry so that the material attributes can be rendered and compared on the same geometry (results in Appendix Table 10).

Mass: We compare the mass $\hat{m}$ of the object against the corresponding ground-truth mass m . Here, we report three scores following (Standley et al., 2017; Zhai et al., 2024): $ADE = |\hat{m} - m|$ in kilograms, $APE = |\hat{m} - m|/m$ and $MnRE = \min(r, 1/r)$, where $r = \hat{m}/m$.

Collision: We evaluate the collision of the generated asset at rest and during articulated motion. We sample the motion in three levels: single-joint, double-joint, and multi-joint. Specifically, for single-joint motion, we uniformly sample 9 states ranging from the rest state to its (predicted) upper-limit state. For double-joint motion, we then combine two joints to form a 9×9 grid of motion states. To efficiently evaluate multi-joint motion, we randomly sample a state for each joint and perform a total of 200 trials. To avoid penalizing collisions from joint combinations that violate motion dependencies, we establish joint correspondences using Hungarian matching (Kuhn, 1955), reproduce each sampled configuration on the ground-truth asset, and exclude colliding part pairs that also contact in that asset. At each sampling level, we report the percentage of collisions over all sampled states.

Stability: Prior work (Guo et al., 2024; Li et al., 2025) evaluates stability under gravity primarily considering the whole object as a single rigid body, with uniform mass distribution. We extend this evaluation to a challenging benchmark of 50 articulated objects with high centers of mass. We simulate each object in MuJoCo (Todorov et al., 2012) for 10 seconds in both its rest and fully articulated states, and report the percentage of stable assets separately for each state. An object is considered stable if its maximum tilt angle relative to its initial pose remains below 30° , its positional change is below 1 mm, and its rotational change does not exceed 0.01 rad.

Robotic Task Performance. Beyond spatial fidelity and physical plausibility, we evaluate whether the generated assets are functionally usable and improve downstream task performance.

Table 1: **Spatial fidelity scores** on ArtVIP (Jin et al., 2026), reported in the rest and the open state. For simplicity, we only report the up-to-date results of Articraft and Codex, while the full results are available in Table 7. **Bold** marks the best value and underline marks the second-best value per metric.

Method	Appearance					Rest-Pose Geometry			Articulated Geometry		
	IoU $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$
Instruct-Particulate	0.6843	0.8581	17.12	0.2705	0.9167	0.0027	0.1143	0.2436	0.0010	0.1132	0.2732
PhysX-Anything	0.5259	0.8457	15.26	0.3899	0.8474	0.0098	0.2703	0.0965	0.0123	0.2989	0.0920
URDF-Anything+	0.6294	0.8598	16.11	0.3819	0.8319	0.0062	0.2913	0.1262	0.0068	0.3034	0.1236
Articulate-Anything	0.6054	0.8197	14.44	0.3893	0.8444	0.0087	0.1923	0.1223	0.0097	0.2076	0.1254
Articraft (GPT-5.5)	0.7195	0.8580	17.48	0.2781	0.9033	0.0026	0.0211	0.2795	0.0040	0.0235	0.2731
Articraft (GPT-6 Astra)	<u>0.7602</u>	0.8754	<u>18.62</u>	<u>0.2347</u>	<u>0.9197</u>	0.0022	0.0135	0.3244	0.0034	<u>0.0229</u>	<u>0.3224</u>
Codex (GPT-6 Astra)	0.7504	0.8700	18.16	0.2486	0.9169	0.0032	0.0206	0.3331	0.0039	0.0265	0.3095
Claude (Fable 5.1)	0.7316	0.8767	18.21	0.2619	0.9048	0.0015	0.0283	0.3309	0.0050	0.0407	0.3024
Agy (Gemini-Flash-3.8)	0.7542	0.8771	18.01	0.2562	0.9130	0.0010	<u>0.0129</u>	<u>0.3553</u>	0.0050	0.0235	0.3113
Ours (GPT-5.5 w/o phys.)	0.7196	0.8669	17.64	0.2784	0.9100	0.0021	0.0162	0.2974	0.0052	0.0260	0.2825
Ours (GPT-5.5)	0.7374	0.8552	17.57	0.2826	0.9086	0.0016	0.0199	0.2956	0.0024	0.0248	0.2937
Ours (GPT-5.6 Sol)	0.7283	0.8701	17.63	0.2721	0.9150	0.0017	0.0179	0.2995	0.0026	0.0249	0.2991
Ours (GPT-6 Astra)	0.7675	<u>0.8769</u>	18.65	0.2319	0.9239	<u>0.0011</u>	0.0101	0.3587	<u>0.0021</u>	0.0160	0.3417

Success Rate: We evaluate whether training with generated articulated assets improves the performance of the vision-language-action policy $\pi_{0.5}$ (Black et al., 2025) on unseen instances. For each held-out asset, we perform policy rollouts to evaluate the performance of $\pi_{0.5}$, which is fine-tuned on robotic demonstrations generated with predefined interactions on articulated assets. We report the mean success rate for each task and the aggregate success rate across all tasks.

Asset Qualification: To further assess the simulator-readiness of the assets, we additionally report how many generated assets can be utilized for robotic tasks. The qualification consists of two stages: **(i) physical stability and task feasibility (Round 1)**, which assesses whether an asset remains in its intended initial configuration after being placed in the simulator and whether it can physically reach the benchmark-defined success state under valid interactions; **(ii) demonstration generation success (Round 2)**, which assesses whether the asset supports the generation of successful demonstrations using the benchmark’s synthetic data generation pipeline within a predefined trial budget.

Overall Quality. We also add a *human* evaluation to assess the overall quality. Participants are presented with the input text prompt and reference images, rendered images of the generated assets in rest and open states, and interactive *MuJoCo* scenes in which each asset can be operated to inspect how its parts move, with six methods’ outputs anonymized and randomly ordered. Each participant selects and ranks the top three of six results according to overall quality, considering visual fidelity, articulation, and physical interaction (detailed in Section B.4).

4.2 DATASETS

We evaluate spatial fidelity on ArtVIP (Jin et al., 2026) and LightWheel (Lightwheel, 2025), each containing 50 articulated objects with diverse categories, shapes, and kinematics. We use our system to reconstruct each of the assets from a text description and a reference image. For physical evaluation, we use ArtVIP for material and collision, 50 articulated objects from ABO-500 (Collins et al., 2022) for mass, and a curated set of 50 objects for stability. The downstream evaluation covers three tasks of RoboTwin 2.0 (Chen et al., 2026) and eight from RoboCasa365 (Nasiriany et al., 2026), with disjoint native training and test assets, generated assets that correspond to the training distribution, and 50 demonstrations per asset. More details of the dataset are provided in Section B.3 and Section E.

4.3 BASELINES

We compare against three groups of baselines: **(i) task-specific learned articulation models**, including the latest feed-forward method Instruct-Particulate (Li et al., 2026b), PhysX-Anything (Cao et al., 2026a) and URDF-Anything+ (Wu et al., 2026); **(ii) specialized agentic methods** for articulated asset construction, Articraft (Zhou et al., 2026) and Articulate-Anything (Le et al., 2025); and **(iii) general-purpose coding agents**, Codex (OpenAI, 2026a), Claude (Fable 5.1 high), and Antigravity (Agy) with Gemini-Flash-3.8. For Articraft, Codex, and ARTISAN, we report GPT-5.5, GPT-5.6 Sol, and GPT-6 Astra variants where available, as indicated in the tables.



Figure 4: **Qualitative comparison** of articulated assets generated by different models. **Orange** rods indicate revolute axes; **blue** arrows indicate prismatic directions.

Table 2: **Physical attribute scores** on ArtVIP (Jin et al., 2026) (Collision, Material), our own dataset (Stability), and ABO-500 (Collins et al., 2022) (Mass). The full results are available in Table 9. **Bold** marks the best value and underline marks the second-best value per metric. — indicates an unavailable result due to the problem scope of each method.

Method	Collision			Stability		Material			Mass		
	Single ↓	Double ↓	Multiple ↓	Rest ↑	Open ↑	Base color ↑	Metallic ↑	Roughness ↑	MnRE ↑	ADE (kg) ↓	APE ↓
NeRF2Physics	—	—	—	—	—	—	—	—	0.601	10.78	0.509
PUGS	—	—	—	—	—	—	—	—	0.412	44.50	3.016
Instruct-Particulate	70.4	90.4	98.7	68%	42%	12.00	18.02	10.40	—	—	—
PhysX-Anything	58.4	77.6	81.3	74%	22%	10.06	—	—	0.302	198.07	19.292
URDF-Anything+	84.6	99.3	100.0	64%	48%	—	—	—	—	—	—
Articulate-Anything	31.2	39.2	32.8	64%	44%	7.92	—	—	—	—	—
Artcraft (GPT-5.5)	44.5	57.5	79.4	82%	78%	10.75	—	—	—	—	—
Artcraft (GPT-6 Astra)	45.5	60.7	78.8	82%	82%	11.75	—	—	—	—	—
Codex (GPT-6 Astra)	32.4	47.0	63.7	86%	82%	9.98	—	—	0.763	6.63	0.303
Claude (Fable 5.1)	39.9	51.7	64.7	84%	86%	10.50	—	—	0.743	6.51	0.388
Agy (Gemini-Flash-3.8)	46.3	65.6	81.5	80%	78%	10.15	—	—	0.775	5.49	0.296
Ours (GPT-5.5 w/o phys.)	37.5	47.3	77.6	84%	82%	10.69	16.32	12.97	0.742	6.07	0.340
Ours (GPT-5.5)	3.5	6.4	12.3	96%	94%	11.48	16.28	<u>13.05</u>	0.720	7.22	0.385
Ours (GPT-5.6 Sol)	<u>1.9</u>	<u>4.8</u>	16.4	98%	96%	10.57	15.98	14.32	0.670	8.25	0.566
Ours (GPT-6 Astra)	1.7	3.7	<u>13.0</u>	98%	96%	<u>11.90</u>	<u>16.38</u>	12.36	<u>0.773</u>	<u>5.86</u>	0.260

5 RESULTS

5.1 ANALYSIS OF BENCHMARK RESULTS

Main Results. We compare ARTISAN to the state-of-the-art baselines described in Section 4.3. We show spatial fidelity results in Table 1 and Figure 4 and physical attribute results in Table 2 and Figure 5. Without bells and whistles, ours (GPT-6 Astra) outperforms the other methods in 8 of 11 metrics for spatial fidelity in Table 1. More importantly, it dramatically improves physical plausibility, reducing collision occurrences to 1.7% on single-joint and 13.0% on multi-joint evaluations, and achieving 98% and 96% stability in the rest and articulated states, respectively. This highlights the significant advantage of our method in generating physically grounded articulated assets. In the following, we summarize several key findings. More details are provided in Section C.

Agentic generation methods can rival or even surpass state-of-the-art task-specific methods.

Compared to very recent state-of-the-art task-specific model Instruct-Particulate, whose input mesh is generated by Hunyuan-3.1 (Tencent Hunyuan3D Team, 2026) from the reference image, our method consistently achieves superior results in most metrics in Tables 1 and 2. Figure 4 further illustrates the visual results in two additional reference-guided examples: a motorcycle and a miniature steam traction engine. Ours provides better appearance and articulated structure. This reveals the potential impact and versatility of agentic generation methods.

Physical validation is the key to generating a ready-to-use asset. As shown in Tables 1 and 2, we ablate ours (GPT-5.5) by disabling the physical validation (GPT-5.5 w/o phys.). Experiments demonstrate that, without physical validation, models can still achieve comparable performance in terms of appearance, geometry, and even kinematics in Table 1. However, their physical behavior degrades drastically with a high rate of collision detection in Table 2, indicating that such assets




Figure 5: **Qualitative collision and stability comparisons.** (a) Visualization of Collision. Red and blue denote the major intersected part pair. ARTISAN is clean, whereas the other baselines involve collisions between components. (b) Visualization of stability under gravity. ARTISAN remains upright, whereas all five baselines topple after gravity is applied.

Table 3: **Downstream robotic evaluation of articulated assets.** (a) Task success rate (%) of $\pi_{0.5}$ (Black et al., 2025) fine-tuned with different training-asset compositions, averaged over 3 RoboTwin 2.0 (Chen et al., 2026) tasks and 8 RoboCasa365 (Nasiriany et al., 2026) tasks. We evaluate all policies on held-out native assets not seen during fine-tuning. (b) Percentage of all assets that pass the qualification test (Round 1) and support successful demonstration generation (Round 2).

(a) Task Success Rate (%)			(b) Asset Qualification Rate (%)				
Assets	RoboTwin	RoboCasa	Assets	RoboTwin		RoboCasa	
				Round 1	Round 2	Round 1	Round 2
Native only	80%	39%	Native	100%	100%	98%	76%
Articraft only	60%	5%	Articraft	62%	62%	39%	8%
ARTISAN only	63%	27%	ARTISAN	100%	100%	95%	44%
Native + ARTISAN	84%	43%					

cannot execute their intended articulated motion within the simulator. As shown in Figure 5, our model with physical validation is the only method that avoids component intersections and maintains stability within the simulator. Therefore, compared with existing agentic or feed-forward methods, our method brings the generated assets much closer to being truly ready to use. This suggests that physical validation in the generation loop is crucial for producing physically grounded assets.

Multi-agent design and independent review help eliminate the self-evaluation bias. In Tables 1 and 2, a fair comparison between “Ours (GPT-5.5 w/o phys.)” and “Articraft (GPT-5.5)” shows average improvements of 5.4% in pCD, 4.9% in pIoU, and 11.9% in collision rates. This demonstrates the effectiveness of the multi-agent design. This further suggests that designing effective and efficient harnessing of agentic systems for physical 3D asset generation is useful, but it is a new and emerging area of physical 3D asset generation research.

5.2 ASSESSING ROBOTIC TASKS AND HUMAN JUDGEMENTS

Robotic Tasks. We report task success and asset qualification rates in Table 3. Based on the experimental results, we have the following key findings: (i) *Our assets are more suitable for robotic tasks:* When trained on robotic demonstrations generated with ARTISAN, the $\pi_{0.5}$ model achieves a higher success rate than that trained on Articraft demonstrations on both benchmarks (rows 2 & 3 in Table 3(a)); (ii) *Our assets can be a good extension to the current robotic training sets:* After incorporating our assets into the original training set, the newly trained model achieves a higher success rate on the same test set with unseen instances (rows 1 & 4 in Table 3(a)); (iii) *Our assets have better availability:* Compared to Articraft assets, our assets achieve higher pass rates (37 percentage points higher on average) across two rounds of qualification reviews on two benchmarks (Table 3(b)), suggesting that our assets are more ready for robotic tasks.




Figure 6: **Human evaluation results.** Each bar stacks the rates at which a method is ranked first, second, and third; labels give each segment’s rate.

Human Judgements. Details of the human evaluation are provided in Section B.4. The results in Figure 6 show that ARTISAN achieves 43.3% of the Top-1 votes and 85.1% of the Top-3 votes, far surpassing the runner-up Codex (21.1% in Top-1 and 68.7% in Top-3). These results are based on user feedback following object manipulation in the MuJoCo simulation environment, further demonstrating that the assets produced by our method are closer to being simulation-ready.

6 CONCLUSION

We presented ARTISAN, a pipeline that achieves physically grounded, articulated 3D asset generation by combining multiple agents and physical simulation feedback. We also introduced a new comprehensive benchmark to evaluate various properties of the generated assets, including appearance, geometry, kinematics, materials, collision, mass, stability, and even robotic task performance. Empirically, we found that ARTISAN achieves the best overall performance across a wide range of evaluation metrics, and even outperforms existing task-specific models in most areas. We also found that multi-agent harnessing with physical simulation feedback significantly improves the spatial and physical fidelity of the generated assets.

Our findings motivate further research for our community: revisiting the role of feed-forward models, developing efficient agent harnesses across tasks, providing richer and more actionable simulator feedback, and extending simulators to capture a broader range of physical phenomena. We hope this work will help the community measure and advance progress toward physically grounded 3D worlds.

Contribution and Acknowledgements. XC led the project with management support from CZ, implemented the core algorithms, and conducted the experiments. XC initiated the multi-agent pipeline, and SW and AV proposed the benchmark concept. RZ contributed to the spatial and physical evaluation and analysis, with additional support from XC on the physical evaluation. HS supported the downstream robotic evaluation by generating training data and conducting experiments on RoboTwin 2.0, while RL assisted with the RoboCasa365 experiments. XC, RZ, RL, SW, AV, and CZ participated in weekly project meetings. XC and RZ led the manuscript writing, with input from all other authors. XC prepared the video demo and other visualizations. RZ made the project page, with feedback from SW, XC and CZ.

We thank Matt Zhou for clarifying the technical details of Articraft; Brandon Smart for proofreading the manuscript; Tat-Jen Cham, Mengdi Wang, Stanislaw Szymanowicz, Edgar Sucar, Yuedong Chen, Qianyi Wu, and Jiawang Bian for helpful discussions on agentic 3D asset generation and related topics; Mengdi Wang for suggesting the virtual laboratory concept; and Jianfei Yang for providing access to the kitchen of their real-world robotic experiments.

This work is supported by NTU SUG-NAP and National Research Foundation, Singapore, under its NRF Fellowship Award NRF-NRFF17-2025-0009. Ruining Li is supported by a Toshiba Research Studentship. Andrea Vedaldi is supported by a Royal Academy Faraday Discovery Fellowship FDF/S2\251003.

AI USE STATEMENT

We used generative artificial intelligence tools to assist in manuscript editing, including correcting grammar, refining sentence structure, and improving clarity and conciseness. The authors reviewed and revised the AI-assisted text and verified its consistency with the methods and experimental results. The authors assume full responsibility for the final manuscript, including its scientific claims and reported results.

ETHICS STATEMENT

Our human evaluation involves inspecting rendered assets and interacting with virtual objects in a physics simulator, as described in Appendix B.4. Participants assess visual quality, articulation, and physical behavior through simulated interactions. We report aggregate preference statistics, and the supplementary offline demo contains no real participant data. The generated assets are intended for simulation and evaluation; their use in real-world robotic applications should undergo additional validation appropriate to the deployment setting.

REPRODUCIBILITY STATEMENT

We describe the generation pipeline in Section 3 and the datasets and evaluation metrics in Section 4. Appendices A and B provide implementation details, dataset construction procedures, and evaluation protocols, including the human-study setup. Appendices C–F document complete results, component ablations, generation budgets and computational costs, robotic experiment settings, and representative agent traces. The supplementary material also includes an offline demonstration of the human-evaluation interface, allowing readers to inspect the rendered candidates and interact with the simulated assets.

REFERENCES

- Paul J. Besl and Neil D. McKay. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992. doi: 10.1109/34.121791. URL <https://doi.org/10.1109/34.121791>.
- Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 17–40. PMLR, 2025. URL <https://proceedings.mlr.press/v305/black25a.html>.
- Ziang Cao, Fangzhou Hong, Zhaoxi Chen, Liang Pan, and Ziwei Liu. PhysX-Anything: Simulation-ready physical 3D assets from single image. In *CVPR*, pp. 5839–5848, 2026a.
- Ziang Cao, Yinghao Liu, Haitian Li, Runmao Yao, Fangzhou Hong, Zhaoxi Chen, Liang Pan, and Ziwei Liu. PhysX-Omni: Unified simulation-ready physical 3D generation for rigid, deformable, and articulated objects. *arXiv preprint arXiv:2605.21572*, 2026b.
- Honghua Chen, Yushi Lan, Yongwei Chen, and Xingang Pan. ArtiLatent: Realistic articulated 3D object generation via structured latents. In *SIGGRAPH Asia*, 2025.
- Qiuyu Chen, Aaron Walsman, Marius Memmel, Kaichun Mo, Alex Fang, Dieter Fox, and Abhishek Gupta. URDFormer: A pipeline for constructing articulated simulation environments from real-world images. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024a. doi: 10.15607/RSS.2024.XX.124.

- Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie, Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan-ang Gao, Kaixuan Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. RoboTwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. In *ICML*, 2026. URL <https://openreview.net/forum?id=itonej9GIV>.
- Yunuo Chen, Tianyi Xie, Zeshun Zong, Xuan Li, Feng Gao, Yin Yang, Ying Nian Wu, and Chenfanfu Jiang. Atlas3D: Physically constrained self-supporting text-to-3D for simulation and fabrication. In *NeurIPS*, pp. 102501–102530, 2024b. doi: 10.52202/079017-3255.
- Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F. Yago Vicente, Thomas Dideriksen, Himanshu Arora, Matthieu Guillaumin, and Jitendra Malik. ABO: Dataset and benchmarks for real-world 3D object understanding. In *CVPR*, 2022.
- Alexandros Delitzas, Ayca Takmaz, Federico Tombari, Robert Sumner, Marc Pollefeys, and Francis Engelmann. SceneFun3D: Fine-Grained Functionality and Affordance Understanding in 3D Scenes. In *CVPR*, 2024.
- Alexandros Delitzas, Chenyangguang Zhang, Alexey Gavryushin, Tommaso Di Mario, Boyang Sun, Rishabh Dabral, Leonidas Guibas, Christian Theobalt, Marc Pollefeys, Francis Engelmann, and Daniel Barath. Reconstructing Functional 3D Scenes from Egocentric Interaction Videos. In *CVPR*, 2026.
- Shivam Duggal, Yushi Hu, Oscar Michel, Aniruddha Kembhavi, William T. Freeman, Noah A. Smith, Ranjay Krishna, Antonio Torralba, Ali Farhadi, and Wei-Chiu Ma. Eval3D: Interpretable and fine-grained evaluation for 3D generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025*, pp. 13326–13336. IEEE, 2025. doi: 10.1109/CVPR52734.2025.01244.
- Tim Engelbracht, René Zurbrügg, Matteo Wohlrapp, Martin Büchner, Abhinav Valada, Marc Pollefeys, Hermann Blum, and Zuria Bauer. Hoi! - a multimodal dataset for force-grounded, cross-view articulated manipulation. In *CVPR*, pp. 8880–8890, June 2026.
- 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*, pp. 618–627, 2025.
- Yipeng Gao, Lei Shu, Genzhi Ye, Xi Xiong, Ameesh Makadia, Meiqi Guo, Laurent Itti, and Jindong Chen. 3DCodeBench: Benchmarking agentic procedural 3D modeling via code. *arXiv preprint arXiv:2606.01057*, 2026.
- Genesis AI Team. The role of simulation in scalable robotics, Genesis World 1.0, and the path forward. Genesis AI Blog, May 2026. URL <https://www.genesis.ai/blog>.
- Genesis Authors. Genesis: A generative and universal physics engine for robotics and beyond, December 2024. URL <https://github.com/Genesis-Embodied-AI/genesis-world>.
- Haoran Geng, Helin Xu, Chengyang Zhao, Chao Xu, Li Yi, Siyuan Huang, and He Wang. GPartNet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts. In *CVPR*, pp. 7081–7091, June 2023.
- Minghao Guo, Bohan Wang, Pingchuan Ma, Tianyuan Zhang, Crystal E. Owens, Chuang Gan, Joshua B. Tenenbaum, Kaiming He, and Wojciech Matusik. Physically compatible 3D object modeling from a single image. *NeurIPS*, 37:119260–119282, 2024.
- Yumeng He, Ying Jiang, Jiayin Lu, Yin Yang, and Chenfanfu Jiang. SPARK: Sim-ready part-level articulated reconstruction with VLM knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7233–7243, 2026.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.

- Jie Huang, Xiaohe Li, Jiahao Li, Fangli Mou, Chen Qian, Yuqiang Fang, Junhao Fan, Kaixin Zhang, and Zide Fan. PhysX-CoT: Structured physical reasoning from a single image to simulation-ready 3D assets. *arXiv preprint arXiv:2608.08053*, 2026a.
- Zhao Huang, Boyang Sun, Alexandros Delitzas, Jiaqi Chen, and Marc Pollefeys. React3d: Recovering articulations for interactive physical 3d scenes. *IEEE Robotics and Automation Letters*, 11(5): 5954–5961, 2026b. doi: 10.1109/LRA.2026.3674028.
- Team Hunyuan3D, Shuhui Yang, Mingxin Yang, Yifei Feng, Xin Huang, Sheng Zhang, Zebin He, Di Luo, Haolin Liu, Yunfei Zhao, Qingxiang Lin, Zeqiang Lai, Xianghui Yang, Huiwen Shi, Zibo Zhao, Bowen Zhang, Hongyu Yan, Lifu Wang, Sicong Liu, Jihong Zhang, Meng Chen, Liang Dong, Yiwen Jia, Yulin Cai, Jiaao Yu, Yixuan Tang, Dongyuan Guo, Junlin Yu, Hao Zhang, Zheng Ye, Peng He, Runzhou Wu, Shida Wei, Chao Zhang, Yonghao Tan, Yifu Sun, Lin Niu, Shirui Huang, Bojian Zheng, Shu Liu, Shilin Chen, Xiang Yuan, Xiaofeng Yang, Kai Liu, Jianchen Zhu, Peng Chen, Tian Liu, Di Wang, Yuhong Liu, Linus, Jie Jiang, Jingwei Huang, and Chunchao Guo. Hunyuan3D 2.1: From images to high-fidelity 3D assets with production-ready PBR material. *arXiv*, 2506.15442, 2025.
- Zhenyu Jiang, Cheng-Chun Hsu, and Yuke Zhu. Ditto: Building digital twins of articulated objects from interaction. In *CVPR*, pp. 5616–5626, 2022.
- Zhao Jin, Zhengping Che, Tao Li, Zhen Zhao, Kun Wu, Yuheng Zhang, Yinuo Zhao, Zehui Liu, Qiang Zhang, Xiaozhu Ju, Jing Tian, Yousong Xue, and Jian Tang. ArtVIP: Articulated digital assets of visual realism, modular interaction, and physical fidelity for robot learning. In *International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=SqPLEZ66BO>.
- Abhishek Joshi, Beining Han, Jack Nugent, Max Gonzalez Saez-Diez, Yiming Zuo, Jonathan Liu, Hongyu Wen, Stamatis Alexandropoulos, Karhan Kayan, Anna Calveri, Tao Sun, Gaowen Liu, Yi Shao, Alexander Raistrick, and Jia Deng. Procedural generation of articulated simulation-ready assets, 2025. URL <https://arxiv.org/abs/2505.10755>.
- Jens U. Kreber and Joerg Stueckler. Guiding diffusion-based articulated object generation by partial point cloud alignment and physical plausibility constraints. In *IEEE/CVF International Conference on Computer Vision, ICCV 2025*, pp. 3206–3214. IEEE, 2025. doi: 10.1109/ICCV51701.2025.00307.
- Harold W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2):83–97, 1955.
- Long Le, Jason Xie, William Liang, Hung-Ju Wang, Yue Yang, Yecheng Jason Ma, Kyle Vedder, Arjun Krishna, Dinesh Jayaraman, and Eric Eaton. Articulate-Anything: Automatic modeling of articulated objects via a vision-language foundation model. In *ICLR*, 2025.
- Jiahui Lei, Congyue Deng, William B. Shen, Leonidas J. Guibas, and Kostas Daniilidis. NAP: Neural 3D articulated object prior. *NeurIPS*, 36:31878–31894, 2023.
- Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Michael Lingelbach, Jiankai Sun, et al. BEHAVIOR-1K: A benchmark for embodied AI with 1,000 everyday activities and realistic simulation. In *Conf. Robot Learn.*, pp. 80–93. PMLR, 2023.
- Ruining Li, Chuanxia Zheng, Christian Rupprecht, and Andrea Vedaldi. Dso: Aligning 3d generators with simulation feedback for physical soundness. In *ICCV*, 2025.
- Ruining Li, Yuxin Yao, Chuanxia Zheng, Christian Rupprecht, Joan Lasenby, Shangzhe Wu, and Andrea Vedaldi. Particulate: Feed-forward 3d object articulation. In *CVPR*, 2026a.
- Ruining Li, Yuxin Yao, Matt Zhou, Chuanxia Zheng, Christian Rupprecht, Joan Lasenby, Shangzhe Wu, and Andrea Vedaldi. Instruct-particulate: Scaling feed-forward 3d object articulation with kinematic control. In *NeurIPS*, 2026b.

- Xian Li, Rong Wei, Lujie Yang, Haolin Huang, Junyuan Fang, Siliang Tang, Jun Xiao, Rui Tang, and Juncheng Li. UniPhysGen: Unified physical grounding for simulation-ready 3D assets. *arXiv preprint arXiv:2607.13586*, 2026c.
- Xinyu Lian, Zichao Yu, Ruiming Liang, Yitong Wang, Li Ray Luo, Kaixu Chen, Yuanzhen Zhou, Qihong Tang, Xudong Xu, Zhaoyang Lyu, et al. Infinite Mobility: Scalable high-fidelity synthesis of articulated objects via procedural generation. *arXiv preprint arXiv:2503.13424*, 2025.
- Lightwheel. SimReady: Simulation-ready 3D assets. <https://simready.com/>, 2025. Accessed: 2026.
- Youtian Lin, Yikang Yang, Zhanpeng Hu, Mengqi Zhou, Feihu Zhang, Xun Cao, Jiaheng Liu, and Yao Yao. Procedura: Agentic 3D modeling with procedural control. *arXiv preprint arXiv:2608.26238*, 2026.
- Jiayi Liu, Ali Mahdavi-Amiri, and Manolis Savva. PARIS: Part-level reconstruction and motion analysis for articulated objects. In *ICCV*, pp. 352–363, 2023.
- Jiayi Liu, Hou In Ivan Tam, Ali Mahdavi-Amiri, and Manolis Savva. CAGE: Controllable articulation generation. In *CVPR*, pp. 17880–17889, 2024.
- Jiayi Liu, Denys Iliash, Angel X. Chang, Manolis Savva, and Ali Mahdavi-Amiri. SINGAPO: Single image controlled generation of articulated parts in objects. In *ICLR*, 2025.
- Liu Liu, Wenqiang Xu, Haoyuan Fu, Sucheng Qian, Qiaojun Yu, Yang Han, and Cewu Lu. AKB-48: A real-world articulated object knowledge base. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 14789–14798. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01439.
- Rundong Luo, Haoran Geng, Congyue Deng, Puhao Li, Zan Wang, Baoxiong Jia, Leonidas J. Guibas, and Siyuan Huang. PhysPart: Physically plausible part completion for interactable objects. In *IEEE International Conference on Robotics and Automation, ICRA 2025*, pp. 12386–12393. IEEE, 2025. doi: 10.1109/ICRA55743.2025.11127496.
- Zhao Mandi, Yijia Weng, Dominik Bauer, and Shuran Song. Real2Code: Reconstruct articulated objects via code generation. In *ICLR*, 2025.
- Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Ireteyio Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. MimicGen: A data generation system for scalable robot learning using human demonstrations. In Jie Tan, Marc Toussaint, and Kourosh Darvish (eds.), *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pp. 1820–1864. PMLR, 2023. URL <https://proceedings.mlr.press/v229/mandlekar23a.html>.
- 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 *ICCV*, pp. 12981–12991. IEEE, 2021.
- Soroush Nasiriany, Sepehr Nasiriany, Abhiram Maddukuri, and Yuke Zhu. RoboCasa365: A large-scale simulation framework for training and benchmarking generalist robots. In *Proc. ICLR*, 2026.
- OpenAI. Codex CLI. <https://developers.openai.com/codex/cli>, 2026a. [Online; accessed 06-May-2026].
- OpenAI. GPT-Image-2. <https://developers.openai.com/api/docs/models/gpt-image-2>, 2026b. [Online; accessed 03-Sep-2026].
- Anh-Quan Pham. Automatically improving simulation physics for articulated objects. *arXiv preprint arXiv:2605.19136*, 2026.
- Corrado Poli. *Design for Manufacturing: A Structured Approach*. Butterworth-Heinemann, 2001. ISBN 9780750673419.

- 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. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 4585–4606. PMLR, 2025. URL <https://proceedings.mlr.press/v305/qiu25b.html>.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021. URL <http://proceedings.mlr.press/v139/radford21a.html>.
- Yinghao Shuai, Ran Yu, Yuantao Chen, Zijian Jiang, Xiaowei Song, Nan Wang, Jv Zheng, Jianzhu Ma, Meng Yang, Zhicheng Wang, Wenbo Ding, and Hao Zhao. PUGS: zero-shot physical understanding with Gaussian splatting. In *IEEE International Conference on Robotics and Automation, ICRA 2025, Atlanta, GA, USA, May 19-23, 2025*, pp. 4478–4485. IEEE, 2025. doi: 10.1109/ICRA55743.2025.11128683. URL <https://doi.org/10.1109/ICRA55743.2025.11128683>.
- Yuan Shui, Yandong Guan, Zhanwei Zhang, Juncheng Hu, Jing Zhang, Dong Xu, and Qian Yu. ArtiCAD: Articulated CAD assembly design via multi-agent code generation. *arXiv preprint arXiv:2604.10992*, 2026.
- Trevor Standley, Ozan Sener, Dawn Chen, and Silvio Savarese. image2mass: Estimating the mass of an object from its image. In *Conf. Robot Learn.*, pp. 324–333. PMLR, 2017.
- Jiayi Su, Youhe Feng, Zheng Li, Jinhua Song, Yangfan He, Botao Ren, and Botian Xu. ArtFormer: Controllable generation of diverse 3D articulated objects. In *CVPR*, pp. 1894–1904, June 2025.
- Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Reed Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fábio Ramos, Nathan Ratliff, and Dieter Fox. CuRobo: Parallelized collision-free robot motion generation. In *IEEE International Conference on Robotics and Automation, ICRA 2023*, pp. 8112–8119. IEEE, 2023. doi: 10.1109/ICRA48891.2023.10160765.
- Tencent Hunyuan3D Team. Hunyuan3D 3.1. <https://3d.hunyuan.tencent.com/>, 2026.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*, pp. 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109. URL <https://doi.org/10.1109/IROS.2012.6386109>.
- Wei-Cheng Tseng, Hung-Ju Liao, Lin Yen-Chen, and Min Sun. CLA-NeRF: Category-level articulated neural radiance field. In *IEEE Int. Conf. Robot. Autom.*, pp. 8454–8460. IEEE, 2022.
- Jiawei Wang, Dingyou Wang, Jiaming Hu, Qixuan Zhang, Jingyi Yu, and Lan Xu. Kinematify: Open-vocabulary synthesis of high-dof articulated objects. *arXiv preprint arXiv:2511.01294*, 2025.
- Rui-Huan Wang, Si-Tong Wei, Jia-Qi He, Heng-Yi Wei, Baoquan Chen, and Peng-Shuai Wang. aDSL: Agentic 3D creation via joint agent-program design. *arXiv preprint arXiv:2608.17975*, 2026.
- Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.
- Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. Approximate convex decomposition for 3D meshes with collision-aware concavity and tree search. *ACM Trans. Graph.*, 41(4):42:1–42:18, 2022. doi: 10.1145/3528223.3530103. URL <https://doi.org/10.1145/3528223.3530103>.

- 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 *CVPR*, pp. 3141–3150, 2024.
- Karl D. D. Willis, Yewen Pu, Jieliang Luo, Hang Chu, Tao Du, Joseph G. Lambourne, Armando Solar-Lezama, and Wojciech Matusik. Fusion 360 Gallery: A dataset and environment for programmatic CAD construction from human design sequences. *ACM Transactions on Graphics (TOG)*, 40(4): 1–24, 2021.
- Di Wu, Liu Liu, Zhou Linli, Anran Huang, Liangtu Song, Qiaojun Yu, Qi Wu, and Cewu Lu. REArtGS: Reconstructing and generating articulated objects via 3D Gaussian splatting with geometric and motion constraints. In *NeurIPS*, 2025a.
- Sean Wu, Shamik Basu, Tim Brödermann, Luc Van Gool, and Christos Sakaridis. PBR-NeRF: Inverse rendering with physics-based neural fields. In *CVPR*, pp. 10974–10984. IEEE, 2025b.
- Tianhao Wu, Fangcheng Zhong, Andrea Tagliasacchi, Forrester Cole, and Cengiz Oztireli. D²NeRF: Self-supervised decoupling of dynamic and static objects from a monocular video. *NeurIPS*, 35: 32653–32666, 2022.
- Tianhao Wu, Chuanxia Zheng, Frank Guan, Andrea Vedaldi, and Tat-Jen Cham. Amodal3r: Amodal 3d reconstruction from occluded 2d images. In *ICCV*, 2025c.
- Zhuangzhe Wu, Yue Xin, Chengkai Hou, Minghao Chen, Yaoxu Lyu, Jieyu Zhang, and Shanghang Zhang. URDF-Anything+: End-to-end generation for simulation-ready articulated assets. *CoRR*, abs/2603.14010, 2026. doi: 10.48550/ARXIV.2603.14010. URL <https://doi.org/10.48550/arXiv.2603.14010>.
- 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 *CVPR*, pp. 11097–11107, 2020.
- Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3D latents for scalable and versatile 3D generation. In *Proc. CVPR*, 2025.
- Wenbo Xu, Liu Liu, Li Zhang, Dan Guo, and Ruonan Liu. MotionAnymesh: Physics-grounded articulation for simulation-ready digital twins. *arXiv preprint arXiv:2603.12936*, 2026.
- Yikang Yang, Zhanpeng Hu, Youtian Lin, Mengqi Zhou, Jingxi Xu, Feihu Zhang, Jiaheng Liu, and Yao Yao. P3D-Bench: Benchmarking MLLMs for parametric 3D generation and structural reasoning. *arXiv preprint arXiv:2606.11152*, 2026.
- Kim Youwang, Jon Hasselgren, Peter Kocsis, Andrea Weidlich, Tae-Hyun Oh, and Jacob Munkberg. Extracting neural materials from multi-view images. *arXiv preprint arXiv:2606.26715*, 2026.
- Albert J. Zhai, Yuan Shen, Emily Y. Chen, Gloria X. Wang, Xinlei Wang, Sheng Wang, Kaiyu Guan, and Shenlong Wang. Physical property understanding from language-embedded feature fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pp. 28296–28305. IEEE, 2024. doi: 10.1109/CVPR52733.2024.02673. URL <https://doi.org/10.1109/CVPR52733.2024.02673>.
- Chuanrui Zhang, Minghan Qin, Yuang Wang, Baifeng Xie, Hang Li, and Ziwei Wang. SimArt: Decomposing monolithic meshes into sim-ready articulated assets via MLLM. In *ACM SIGGRAPH 2026 Conference Papers*, pp. 1–12. ACM, 2026. doi: 10.1145/3799902.3811173.
- Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 586–595. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00068. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html.

Matt Zhou, Ruining Li, Xiaoyang Lyu, Zhaomou Song, Zhening Huang, Chuanxia Zheng, Christian Rupprecht, Andrea Vedaldi, and Shangzhe Wu. Articraft: An agentic system for scalable articulated 3d asset generation. In *NeurIPS*, 2026.

APPENDIX CONTENTS

A	Implementation Details	19
A.1	SDK Extensions	19
A.2	Agent Tools	19
A.3	Physical Simulation	20
B	Benchmark Details	20
B.1	Spatial Fidelity Protocol	20
B.2	Material Transfer Protocol	21
B.3	Evaluation Datasets	21
B.4	Human Evaluation	23
C	Additional Results	25
C.1	Full Metric Tables	25
C.2	Ablation	27
C.3	Additional Qualitative Examples	28
D	API and Time Cost	32
E	Robotic Experiments	32
E.1	RoboTwin 2.0	33
E.2	RoboCasa365	34
F	Agent Trace Examples	37
F.1	Stand Mixer: Control Accessibility	37
F.2	Shoe Cabinet: Drawer Connections	38

A IMPLEMENTATION DETAILS

A.1 SDK EXTENSIONS

Physical Property Abstraction. The SDK exposes three groups of physical attributes: part materials, part masses, and joint damping and friction. For materials, the Coding Agent selects a named material preset and can directly adjust its PBR values, including albedo, metallicity, roughness, and opacity. For mass, the agent reasons from the blueprint to estimate each part’s mass, considering its dimensions, materials, and construction. We deliberately avoid computing mass as material density times geometric volume, since real objects often combine multiple materials and contain hollow regions, making this approximation prone to large errors. The SDK computes each part’s default inertia from its geometry and estimated mass. The agent can also adjust each part’s inertial properties to reflect its mass distribution. Joint damping and friction are configured according to the motions and mechanisms specified in the blueprint.

Standard Parts for Joint Connections. We provide six commonly used families of parameterized *standard parts*: hinges, slide rails, pinned joints, rotary connectors, telescoping slides, and threaded sleeves (Figure 3). Each connector is built from matching components with shared dimensions, alignment, and clearance parameters. For example, hinge halves share a pin axis, while a slide rail pairs a guide with a matching runner. The Coding Agent adapts their dimensions and mounting positions to the asset, reducing repeated geometric design and mismatches between mating components.

Anchor-based Placement. Authoring every pose in explicit coordinates requires the agent to calculate distances and orientations, and to update them whenever geometry changes. We instead let the agent specify the intended spatial relationship and let the SDK solve the corresponding pose from the built geometry. An *anchor* names a geometric feature on an existing element: an oriented plane from a face, an axis from an edge or a round feature’s center line, or a point from a corner. Given a subject anchor a_{sub} , reference anchors $\mathcal{A}_{\text{ref}}$, and placement parameters θ , the interface computes $T = \omega(a_{\text{sub}}, \mathcal{A}_{\text{ref}}; \theta)$, where T is the subject’s rigid pose and θ specifies offsets, clearances, or a translation direction. Common operations are $\omega \in \{\text{ATTACH}, \text{CENTRE}, \text{SNAP}\}$: ATTACH aligns mating anchors with the requested orientation and clearance; CENTER centers the subject between two references along a chosen axis; and SNAP translates it along a specified direction until it reaches a support at the prescribed gap. The SDK resolves the anchors in a shared coordinate frame and computes the rotation and translation required by the selected relation, preserving pose components that the operation does not constrain. Operations can be composed, for example, to center a shelf between side panels and then seat it against the back panel. Changes to the underlying geometry automatically propagate through these relationships, keeping dependent placements consistent without manually recomputing coordinates.

A.2 AGENT TOOLS

Following Articraft’s tool interface (Zhou et al., 2026), we expose actions for editing, execution, inspection, and reference retrieval (Table 4). The Coding Agent can modify the asset program, while the Review Agent can inspect the asset and edit its own tests without changing the asset itself.

Editing and execution. The Coding Agent writes or patches the construction program and can execute Blender scripts to create the corresponding geometry. Compilation builds the asset and returns diagnostics for invalid structure, disconnected geometry, unintended overlaps, and failed tests. The Review Agent separately writes and revises tests against the blueprint, then executes them on the asset to obtain failures and supporting details.

Inspection. Both agents can query the part hierarchy and joint axes and limits, or run read-only probes to measure dimensions, clearances, overlaps, and connectivity. These measurements identify the parts and spatial relationships involved in a suspected defect. They can also retrieve the physics evidence associated with the current asset to examine failures under articulation or gravity.

Visual Inspection. `request_render` uses Blender to produce RGB images or semantic maps that assign a distinct flat colour to each part, making part boundaries easier to distinguish. Both agents can request front, side, top, or perspective views in open or closed configurations and focus on named parts or joints. The Review Agent can additionally specify joint positions and camera angles,

Family	Tool(s)	C	R	Purpose
Edit	write_file, replace	✓		Create or revise the asset program and geometry scripts.
	write_test, edit_test		✓	Create or revise independent tests against the blueprint.
Execute	run_blender	✓		Execute a geometry script and return the mesh and its diagnostics.
	compile_model	✓		Build the asset and return structural, geometric, and test feedback.
	run_tests		✓	Run the Review Agent’s tests and return failures or errors.
Inspect	probe_model	✓	✓	Query part and joint structure or geometric measurements.
	read_gate_evidence	✓	✓	Retrieve motion and physics evidence for the current asset.
	request_render	✓	✓	Return RGB views or semantic maps with a distinct colour per part.
	crop_region		✓	Enlarge a selected region of an existing render.
Reference	read_file	✓	✓	Read source code, task specifications, or SDK documentation.
	find_examples	✓		Retrieve curated examples of geometry and mechanisms.
	lookup_cadquery_api,	✓		Look up signatures and documentation in the installed APIs.
	lookup_blender_api			

Table 4: Tools available to the Coding (C) and Review (R) Agents. The Coding Agent edits the asset program; the Review Agent edits only its own tests.

or crop an existing image to inspect a local detail. RGB views reveal appearance and proportions, while semantic maps help locate unexpected part visibility or apparent interpenetration for further geometric inspection.

A.3 PHYSICAL SIMULATION

Collision under Motion. We use MuJoCo to check collisions at sampled joint configurations, with mesh collision geometry constructed by CoACD using a decomposition threshold of 0.05, at most 32 convex hulls per mesh, and random seed 0. The initial collision screen includes the rest pose, seven evenly spaced values along each joint’s travel, pairwise 7×7 grids, and 200 random multi-joint samples drawn from the same joint-value grids. We cap this initial set at 1,500 configurations and use sampling seed 0. The functional constraints in Q^{text} guide checks of dependent motions: for example, an internal drawer is tested with its access door opened, so an invalid operation sequence is not mistaken for a geometric defect. Collisions that persist during individual joint operation are sent to the Coding Agent for repair. For collisions found only when multiple joints move together, the Review Agent determines whether they occur during normal use and need repair. For each affected part pair, the report records a representative configuration and penetration depth, allowing the agents to reproduce and inspect the finding.

Stability under Gravity. For free-standing assets, we test stability under gravity in MuJoCo at rest and with each joint opened in turn. In each test, the assembly is placed on a ground plane with a free root, allowing it to settle or topple. We simulate for up to 5 s, stopping early when the root’s linear and angular speeds remain at most 0.02 m/s and 0.05 rad/s, respectively, for 0.4 s. Tilt is measured as the angle between the root’s initial and current up directions; a final tilt above 3° is flagged for review. The physics report $\mathcal{E}_t$ records, for each pose, the final and maximum tilt, whether the assembly settled, and the simulated duration, so the Review Agent can assess the result against the asset’s intended use.

B BENCHMARK DETAILS

B.1 SPATIAL FIDELITY PROTOCOL

Alignment and Correspondence. We sample 100,000 points uniformly over each asset’s surface and record their positions in both endpoint poses. To resolve ambiguity in joint-axis motion direction,

we test both ways of pairing the predicted poses with the ground-truth closed and open states. For each pairing, we align both states with ICP from several initial orientations. We select the pairing with the lower alignment error across both states, retaining the original ordering when the difference is negligible. After alignment, we normalize both assets so that the longest side of the ground-truth asset’s closed-state bounding box has unit length. We compute each part’s centroid as the mean of its sampled surface points. Hungarian matching minimizes the total Euclidean distance between predicted and reference centroids in the aligned closed state, yielding one-to-one correspondences that are reused in the open state.

B.2 MATERIAL TRANSFER PROTOCOL

To separate material errors from geometric incoherence, we transfer predicted materials onto the reference geometry. We use Hungarian assignment to match parts by their relative positions, shapes, and motion between the rest and open states. Within each matched part pair, component matching minimizes a weighted cost combining differences in normalized centers, bounding-box dimensions and aspect ratios, and relative surface area and bounding-box volume. We transfer materials between the resulting component matches. Numeric material parameters are copied directly, while texture inputs are reduced to surface-area-weighted means because the meshes do not share UV layouts.

Keeping the ground-truth geometry and joint configurations fixed, we render base-color, metallic, and roughness buffers in Blender at 512^2 resolution from three viewpoints in each state. We compute PSNR over visible object interiors, excluding boundary pixels. Scores are averaged over the six images and then over objects. Table 10 reports the transferred-material scores.

B.3 EVALUATION DATASETS

Spatial Fidelity Evaluation Datasets. LightWheel (Lightwheel, 2025) is a library of simulation-ready OpenUSD assets modelled by artists and released under a CC-BY-NC license. Recent articulation work adopts 243 of these assets, spanning 14 categories, as an evaluation set (Li et al., 2026a;b). It contains fourteen kitchen-appliance categories, including blenders, electric kettles, range hoods, and stovetops. We sample 50 objects evenly across the fourteen categories (Figure 7a).

ArtVIP (Jin et al., 2026) comprises 992 digital-twin articulated objects over 9 categories and 37 subcategories. Each asset is built by professional 3D modellers to a unified standard and ships with tuned dynamic parameters, embedded modular interaction behaviours, and pixel-level affordance annotations. Its categories extend beyond kitchen appliances to furniture, fixtures, cleaning tools, stationery, storage, and mechanical equipment. This diversity challenges methods to construct a wider range of geometries and mechanisms, motivating our use of ArtVIP as the main benchmark. We select 50 objects through category-stratified sampling to span diverse object types and joint counts, and refine the selection by inspecting the models and renders. Figure 7b–c shows the subset’s category and joint-count distributions. Table 5 lists all 50 objects using the same category grouping.

Stability Evaluation Dataset. We construct 50 object prompts around common destabilization mechanisms in everyday articulated objects, such as cantilevered loads, telescoping extensions, and tilting panels, as shown in Figure 8. By varying both object categories and mechanisms, the set tests whether generated assets remain stable across diverse support geometries and joint configurations.

For example, one prompt used in the benchmark reads:

A narrow hallway shoe cabinet with a slim upright carcass and three stacked tilt-out shoe bins. Each bin rotates outward and downward on its own horizontal revolute hinge at the bottom front edge of its opening, carrying the whole bin body clear of the carcass front. Hall-furniture scale, tall and shallow with a footprint barely deeper than a shoe.

Figure 8 shows the distributions by the destabilization mechanism (left) and the category (right).

Mass Evaluation Dataset. Following NeRF2Physics (Zhai et al., 2024), we use ABO-500, a collection of 500 objects from Amazon Berkeley Objects (ABO) (Collins et al., 2022) with calibrated multi-view images and ground-truth masses from product listings. We select 50 articulated objects from distinct product families to cover diverse articulation mechanisms and a ground-truth mass range of 0.086–205 kg. For each object, we provide a brief prompt specifying its category and overall




Figure 7: Composition of the spatial-fidelity evaluation subsets. **(a)** The LightWheel subset contains 50 objects across 14 categories, but all for kitchen-appliance categories. **(b)** The ArtVIP subset contains 50 objects across 9 large categories with 37 subcategories. **(c)** The ArtVIP objects have 1–13 movable joints each, with fixed joints excluded.

Table 5: The 50 ArtVIP objects used for evaluation. Names follow the source asset identifiers.

Category	Count	Object names
Furniture	18	GULLABERG_Shoe_cabinet_with_2_compartments_grey_88x30x122cm, KILBERGET_chest_of_6_drawers_white_anchor_unlock_function_144x78cm, MUSKEN_bedside_table_brown_45x58cm, NORDKISA_Dressing_table_bamboo_76x47cm(2), RIDSPO_Desk_oak_140X70cm, STOCKHOLM_2025_Sideboard_oak_veneer_161x42x83cm, bookcase006, cabinet_1, chest_of_drawers_1, clock007, double_door_4, lab_bench001, office_chair_3, poison_cabinet001, table_3, wardrobe_1, wardrobe_21, wardrobe_22
Kitchen appliances	9	capsule_coffee_maker_3, dishwasher_2, fridge_1, l_dishwasher_1, microwave_2, oven_2, refrigerator_6, rice_cooker_7, steamer
Mechanical equipment	7	dolly, elevator, incubator, medical_cart001, medical_monitor, microscope001, oscillator
Appliances	5	big_electric_fan_9, desk_lamp_008, electric_fan_10, television003, washing_machine_2
Cleaning tools	4	mop_10, mop_9, steping_dustbin_20, stepping_dustbin_6
Storage	3	basket_21, carton_4, clamshell_tool_box_2
Fixtures	2	closestool_3, faucet_7
Stationery	1	folder_25
Kitchenware	1	tongs_10




Figure 8: Composition of the stability dataset. **Left:** The 50 objects cover 9 destabilization mechanisms. **Right:** The objects span 10 categories defined by their primary use.

dimensions, together with a reference image. Figure 9 shows the ground-truth mass distribution (left) and object categories (right), and Table 6 lists the selected objects and their ground-truth masses.




Figure 9: Composition of the mass-evaluation dataset. **Left:** The ground-truth masses range from 0.086 to 205 kg. **Right:** The 50 objects span 12 categories grouped by primary use.

Downstream Tasks Datasets. We evaluate three articulated operation tasks from RoboTwin 2.0 (Chen et al., 2026) and eight from RoboCasa365 (Nasiriany et al., 2026), using disjoint native training and test assets and generating additional assets matched to the training distribution. For RoboTwin 2.0, each task contains 2–11 native assets; we hold out 1–2 for testing and generate one asset per remaining training asset. For RoboCasa365, we generate demonstrations on 10 native training assets and 10 generated assets, and hold out the remaining native training assets as test assets for each task. We collect 50 demonstrations per asset using RoboTwin 2.0’s task scripts with cuRobo (Sundaralingam et al., 2023) or RoboCasa365’s MimicGen pipeline (Mandlekar et al., 2023). Further details on the train-test split and demonstration collection can be found in Section E.

B.4 HUMAN EVALUATION

We introduce an interactive user-study protocol in which participants assess generated assets through direct manipulation in MuJoCo alongside rendered views. Making physical interaction part of human evaluation is essential for physically grounded assets: an object may look convincing yet fail to operate as expected. Participants can move articulated parts and observe their responses to applied forces and contact, allowing functional failures to inform their judgments.

Assets and Comparisons. The study covers 100 ground-truth assets, comprising the 50-object evaluation subsets from ArtVIP and LightWheel. For each object, we compare ARTISAN (GPT-6-Astra) with five representative baselines: Instruct-Particulate, Codex (GPT-6-Astra), Articraft (GPT-6-Astra), Claude (Fable-5.1), and agy (Gemini-3.8-Flash). Participants receive the text prompt and

ASIN	Category	Mass (kg)	ASIN	Category	Mass (kg)
B006ZT4VA0	tablet stand	0.086	B07TFQM3JZ	utility cart	10.478
B082L6NQZ7	speaker wall mount	0.617	B084W2S628	barstool	10.977
B075HXHKVF	jewelry box	0.717	B075X463J2	ottoman	12.247
B00WRDS8H0	laptop stand	1.102	B073GGGF7H	storage bench	17.236
B0742D9X4G	table lamp	1.361	B082L28P18	file cabinet	18.144
B07MSYK43S	projector ceiling bracket	1.383	B007IHF36E	sideboard	19.999
B07MF1S1WL	wall sconce lamp	1.397	B07DMGM9T6	kitchen cart	21.319
B07BL7H8N5	frying pan with a lid	1.896	B07L48L3BH	kitchen cart	23.501
B07B8PXMRT	wall clock	2.336	B07H8JNCF3	bedside table	24
B07WFW48N	step stool	3.452	B07N6RT9SM	portable air conditioner	25.301
B00MIBN16O	monitor arm	3.538	B07TLDMSQ2	desk	26
B07XMV34PX	floor fan	3.992	B075X4WBNH	storage ottoman	29.483
B07DHNHZ2Y	dual monitor arm	4.286	B07H8P3K2R	TV stand	32
B07QVC3SLK	shop stool	4.5	B07TR659GF	TV stand	32
B07KR4BT31	ceiling TV mount	4.64	B07GFFPKBG	TV stand	33
B07Q16JDH5	keyboard stand	4.74	B07GFRP9RY	chest of drawers	40
B07T9NM5QR	DJ laptop stand	4.931	B07K6S2G5Z	corner desk	47
B07QWC7NYS	garage creeper seat	5.198	B07QBMQ737	recliner	48.081
B07TX2D4WW	storage cube	5.418	B07H8T5XPT	TV stand	50
B07M9KPR1P	TV wall mount	6.21	B07K6RH9HS	desk	51
B07RKJCGXR	piano bench	6.532	B07Y2B738L	window air conditioner	51.891
B000S6N026	paper shredder	6.68	B07H8SSY4F	sideboard	56
B07WBYQQ8F	garden utility wagon	8.8	B075YZPG5Y	dresser	60.328
B079VKDKC1	bedside cabinet	9.7	B07GFRWJS7	double bed frame with storage	60.5
B07TF7YLJC	desk chair	10.4	B07H8VCM68	wardrobe	205


Table 6: The 50 ABO objects used for mass evaluation, ordered by ground-truth mass.

reference images, inspect closed- and open-state renders from matched viewpoints, and interact with all six candidates in the simulator. Method identities are hidden and candidate order is randomized.


Evaluation Criteria. Participants select and rank the three best candidates by prompt alignment and overall quality, considering three complementary aspects: *visual quality*, whether appearance matches the reference image and description; *articulation*, whether movable parts are correctly separated and move about the intended axes and in the intended directions; and *physical interaction*, whether the asset behaves as its real-world counterpart would during manipulation. These criteria guide a single overall ranking for each object. For each method, we report top three selection rates, averaged equally across objects.

Study Interface. Figure 10 shows the web interface, which presents each object on one page with a 2D-render tab and a 3D-interaction tab that runs all candidates in MuJoCo in the browser. Before submitting a ranking, participants must drag each candidate at least twice and spend at least three seconds in its 3D viewer.


Participants and Results. We recruit 30 non-expert college students from diverse academic backgrounds, each evaluating 10 randomly assigned objects. The study yields 275 valid responses. Figure 6 reports the selection rates for all six methods.



(a) 2D renders tab



(b) 3D interaction tab



(c) 3D interaction tab with Color by part

Figure 10: User study interface. (a) Closed- and open-state renders of the six anonymized candidates from a shared viewpoint, next to the reference images and text prompt; the right panel lists the three criteria and collects the top-three ranking. (b) The same page on the 3D interaction tab, cropped to the interaction area. Each candidate is loaded into an in-browser MuJoCo viewer, where the participant drags a part (red arrow: applied force); buttons below the viewer locate the movable parts, and the bottom bar tracks each candidate’s interactions. (c) The same moment with *Color by part* enabled, which draws each rigid body in its own color to reveal the part decomposition.

C ADDITIONAL RESULTS

C.1 FULL METRIC TABLES

Tables 7 and 8 report complete spatial fidelity results on ArtVIP and LightWheel, including all evaluated model variants. Table 9 reports collision rates on both datasets, together with stability,

Table 7: **Complete spatial fidelity results on ArtVIP.** **Bold** and underline indicate the best and second-best values per metric, including ties at the reported precision.

Method	Appearance					Rest-Pose Geometry			Articulated Geometry		
	IoU $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$
Instruct-Particulate	0.6843	0.8581	17.12	0.2705	0.9167	0.0027	0.1143	0.2436	0.0010	0.1132	0.2732
PhysX-Anything	0.5259	0.8457	15.26	0.3899	0.8474	0.0098	0.2703	0.0965	0.0123	0.2989	0.0920
URDF-Anything+	0.6294	0.8598	16.11	0.3819	0.8319	0.0062	0.2913	0.1262	0.0068	0.3034	0.1236
Articulate-Anything	0.6054	0.8197	14.44	0.3893	0.8444	0.0087	0.1923	0.1223	0.0097	0.2076	0.1254
Articraft (GPT-5.5)	0.7195	0.8580	17.48	0.2781	0.9033	0.0026	0.0211	0.2795	0.0040	0.0235	0.2731
Articraft (GPT-5.6-Sol)	0.7297	0.8669	17.79	0.2658	0.9127	0.0024	0.0231	0.2755	0.0041	0.0282	0.2737
Articraft (GPT-6-Astra)	<u>0.7602</u>	0.8754	<u>18.62</u>	<u>0.2347</u>	<u>0.9197</u>	0.0022	0.0135	0.3244	0.0034	<u>0.0229</u>	<u>0.3224</u>
Codex (GPT-5.5)	0.6732	0.8543	16.92	0.3010	0.9003	0.0037	0.0259	0.2399	0.0049	0.0372	0.2421
Codex (GPT-5.6-Sol)	0.7126	0.8649	17.49	0.2826	0.9028	0.0040	0.0380	0.2617	0.0044	0.0413	0.2431
Codex (GPT-6-Astra)	0.7504	0.8700	18.16	0.2486	0.9169	0.0032	0.0206	0.3331	0.0039	0.0265	0.3095
Claude (Fable 5.1)	0.7316	0.8767	18.21	0.2619	0.9048	0.0015	0.0283	0.3309	0.0050	0.0407	0.3024
Agy (Gemini-Flash-3.8)	0.7542	0.8771	18.01	0.2562	0.9130	0.0010	<u>0.0129</u>	<u>0.3553</u>	0.0050	0.0235	0.3113
Ours (GPT-5.5 w/o phys.)	0.7196	0.8669	17.64	0.2784	0.9100	0.0021	0.0162	0.2974	0.0052	0.0260	0.2825
Ours (GPT-5.5)	0.7374	0.8552	17.57	0.2826	0.9086	0.0016	0.0199	0.2956	0.0024	0.0248	0.2937
Ours (GPT-5.6-Sol)	0.7283	0.8701	17.63	0.2721	0.9150	0.0017	0.0179	0.2995	0.0026	0.0249	0.2991
Ours (GPT-6-Astra)	0.7675	<u>0.8769</u>	18.65	0.2319	0.9239	<u>0.0011</u>	0.0101	0.3587	<u>0.0021</u>	0.0160	0.3417

Table 8: **Complete spatial fidelity results on LightWheel.** **Bold** and underline indicate the best and second-best values per metric.

Method	Appearance					Rest-Pose Geometry			Articulated Geometry		
	IoU $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$
Instruct-Particulate	0.8517	<u>0.9250</u>	20.80	0.1085	0.9347	0.0039	0.0845	<u>0.2631</u>	0.0033	0.0938	0.2890
PhysX-Anything	0.6225	0.9021	16.69	0.2158	0.8948	0.0082	0.1743	0.0986	0.0163	0.1967	0.0762
URDF-Anything+	0.7126	0.9145	16.17	0.2455	0.8368	0.0063	0.3620	0.1200	0.0087	0.3728	0.1144
Articulate-Anything	0.6120	0.9066	16.45	0.2306	0.8596	0.0103	0.1627	0.1098	0.0149	0.1928	0.0924
Articraft (GPT-5.5)	0.7891	0.9110	18.49	0.1534	0.9117	0.0019	0.0516	0.2050	0.0041	0.0635	0.1726
Articraft (GPT-6-Astra)	0.8301	0.9240	<u>20.46</u>	<u>0.1154</u>	0.9399	0.0018	0.0264	0.2511	<u>0.0032</u>	<u>0.0357</u>	0.2269
Codex (GPT-5.5)	0.7547	0.9054	18.01	0.1668	0.9052	0.0027	0.0657	0.1859	0.0047	0.0767	0.1648
Codex (GPT-6-Astra)	0.8124	0.9188	19.92	0.1248	0.9322	<u>0.0015</u>	0.0329	0.2603	0.0045	0.0460	0.2282
Claude (Fable 5.1)	0.7771	0.9215	19.26	0.1525	0.9186	0.0023	0.0671	0.2097	0.0061	0.0815	0.1715
Agy (Gemini-Flash-3.8)	0.7826	0.9070	18.38	0.1504	0.9151	0.0022	0.0689	0.2290	0.0055	0.0813	0.1969
Ours (GPT-5.5 w/o phys.)	0.7875	0.9153	18.55	0.1599	0.9095	0.0017	0.0437	0.2186	0.0045	0.0591	0.1974
Ours (GPT-5.5)	0.7933	0.9153	18.50	0.1562	0.9107	0.0018	0.0433	0.2127	0.0040	0.0583	0.1923
Ours (GPT-6-Astra)	<u>0.8320</u>	0.9254	20.32	0.1174	<u>0.9355</u>	0.0011	<u>0.0272</u>	0.2901	0.0031	0.0320	<u>0.2512</u>

direct-material, and mass results. Table 10 separately reports material transfer onto the ground-truth geometry on ArtVIP. All results follow the protocols in Sections B.1, B.2 and 4.1.

Spatial Fidelity. With GPT-6-Astra, ARTISAN leads on 8 of 11 ArtVIP metrics. On LightWheel, it achieves competitive appearance scores and ranks first on 4 of the 6 geometry metrics, with the lowest object Chamfer distance in both states. Instruct-Particulate, whose training mixture includes household assets from PartNet-Mobility and GRScenes (Li et al., 2026b), leads on 3 of the 5 LightWheel appearance metrics and open-state part IoU. These results show that agentic construction can match or surpass learned articulation pipelines in spatial fidelity; its stronger performance on the more diverse ArtVIP benchmark further suggests broader generalization across object categories.

Physical Attributes. As shown in Table 9, ARTISAN (GPT-6-Astra) achieves lower collision rates than all baselines at all three sweep levels on both datasets, together with stability rates of 98% at rest and 96% when open. It also achieves the lowest mass APE, although Agy has lower ADE and higher MnRE. We also estimate mass by multiplying the densities of the agent-assigned materials by their geometric volumes. This density-times-volume estimate yields an MnRE of 0.525, an ADE of 23.70 kg, and an APE of 2.379, substantially underperforming our direct mass estimates. This

Table 9: **Complete physical attribute results.** Collision rates are reported on ArtVIP and LightWheel; direct material PSNR uses ArtVIP, stability our stability dataset, and mass uses the ABO subset. Bold and underline mark the best and second-best values per metric; — denotes unavailable results.

Method	Collision (%)						Stability			Material (dB)			Mass		
	ArtVIP			LightWheel			Rest↑	Open↑	Base color↑	Metallic↑	Roughness↑	MnRE↑	ADE (kg)↓	APE↓	
	Single↓	Double↓	Multi↓	Single↓	Double↓	Multi↓									
NeRF2Physics	—	—	—	—	—	—	—	—	—	—	—	0.601	10.78	0.509	
PUGS	—	—	—	—	—	—	—	—	—	—	—	0.412	44.50	3.016	
Instruct-Particulate	70.4	90.4	98.7	45.1	66.2	80.6	68%	42%	12.00	18.02	10.40	—	—	—	
PhysX-Anything	58.4	77.6	81.3	58.6	86.5	96.3	74%	22%	10.06	—	—	0.302	198.07	19.292	
URDF-Anything+	84.6	99.3	100.0	92.0	100.0	100.0	64%	48%	—	—	—	—	—	—	
Articulate-Anything	31.2	39.2	32.8	24.1	38.7	62.5	64%	44%	7.92	—	—	—	—	—	
Articraft (GPT-5.5)	44.5	57.5	79.4	43.4	66.0	83.1	82%	78%	10.75	—	—	—	—	—	
Articraft (GPT-5.6-Sol)	49.0	67.6	81.7	—	—	—	82%	76%	11.06	—	—	—	—	—	
Articraft (GPT-6-Astra)	45.5	60.7	78.8	42.8	58.6	75.6	82%	82%	11.75	—	—	—	—	—	
Codex (GPT-5.5)	61.6	73.4	80.0	49.7	68.6	82.5	—	—	9.64	—	—	—	—	—	
Codex (GPT-5.6-Sol)	37.7	54.6	65.9	—	—	—	80%	86%	10.66	—	—	—	—	—	
Codex (GPT-6-Astra)	32.4	47.0	63.7	25.2	38.0	55.6	86%	82%	9.98	—	—	0.763	6.63	0.303	
Claude (Fable 5.1)	39.9	51.7	64.7	27.2	39.2	62.6	84%	86%	10.50	—	—	0.743	6.51	0.388	
Agy (Gemini-Flash-3.8)	46.3	65.6	81.5	33.9	52.2	63.9	80%	78%	10.15	—	—	0.775	5.49	0.296	
Ours (GPT-5.5 w/o phys.)	37.5	47.3	77.6	48.9	67.2	74.8	84%	82%	10.69	16.32	12.97	0.742	6.07	0.340	
Ours (GPT-5.5)	3.5	6.4	12.3	2.1	3.9	7.8	96%	94%	11.48	16.28	13.05	0.720	7.22	0.385	
Ours (GPT-5.6-Sol)	1.9	4.8	16.4	—	—	—	98%	96%	10.57	15.98	14.32	0.670	8.25	0.566	
Ours (GPT-6-Astra)	1.7	3.7	13.0	3.2	5.4	9.7	98%	96%	11.90	16.38	12.36	0.773	5.86	0.260	

Table 10: **Material transfer evaluation on ArtVIP.** PSNR (dB) uses predicted materials transferred to ground-truth geometry. **Bold** and underline mark the best and second-best scores; — denotes unavailable results due to the problem scope of methods.

Method	Base color $\uparrow$	Metallic $\uparrow$	Roughness $\uparrow$
Instruct-Particulate	10.07	18.14	10.54
PhysX-Anything	8.84	—	—
Articulate-Anything	7.55	—	—
Articraft (GPT-5.5)	11.15	—	—
Articraft (GPT-5.6-Sol)	11.36	—	—
Articraft (GPT-6-Astra)	13.24	—	—
Codex (GPT-5.5)	10.32	—	—
Codex (GPT-5.6-Sol)	11.97	—	—
Codex (GPT-6-Astra)	10.72	—	—
Claude (Fable 5.1)	11.16	—	—
Agy (Gemini-Flash-3.8)	10.66	—	—
Ours (GPT-5.5 w/o phys.)	10.75	<u>16.93</u>	<u>12.80</u>
Ours (GPT-5.5)	11.37	16.10	12.48
Ours (GPT-5.6-Sol)	11.29	16.26	14.39
Ours (GPT-6-Astra)	<u>12.08</u>	16.40	12.12

supports our design of reasoning directly about part mass to account for hollow and composite construction. Compared with Instruct-Particulate, it achieves comparable base-color PSNR and higher roughness PSNR.

Model Capacity. For Articraft, Codex, and ARTISAN, moving from GPT-5.5 to GPT-6-Astra improves every reported spatial fidelity metric on both datasets. However, the GPT-6-Astra baselines still exhibit substantially higher collision rates than ARTISAN. Stronger models improve reconstruction quality, but do not remove the need for explicit physical validation.

C.2 ABLATION

To assess the contributions of each component in ARTISAN, we examine three components by removing the Design Agent, the Review Agent, or physical validation from ARTISAN with GPT-5.6-Sol (Table 11). The geometry scores here are the average of the rest and open states.

Table 11: **Component ablations with GPT-5.6-Sol.** Each variant removes the Design Agent, Review Agent, or physical validation. Appearance, geometry, collision, and mean part count are measured on ArtVIP; stability and mass use their respective datasets. Geometry scores average the rest and open states, and mass ADE is in kilograms. **Bold** and underline mark the best and second-best scores per metric; part count is descriptive.

Variant	Appearance		Geometry			Parts	Collision (%)			Stability		Mass		
	PSNR $\uparrow$	CLIP $\uparrow$	oCD $\downarrow$	pCD $\downarrow$	pIoU $\uparrow$	N_p	Single $\downarrow$	Double $\downarrow$	Multi $\downarrow$	Rest $\uparrow$	Open $\uparrow$	MnRE $\uparrow$	ADE $\downarrow$	APE $\downarrow$
w/o Design	17.84	0.9130	0.0035	0.0195	0.2788	5.82	1.0	3.5	<u>15.1</u>	96%	96%	0.630	9.57	0.710
w/o Review	17.95	0.9128	<u>0.0030</u>	0.0211	0.2951	6.50	3.1	5.8	11.8	<u>96%</u>	<u>94%</u>	0.640	8.44	0.736
w/o Physics	<u>17.90</u>	0.9091	<u>0.0030</u>	<u>0.0199</u>	<u>0.2952</u>	6.66	35.7	54.1	73.2	90%	90%	<u>0.649</u>	7.68	<u>0.628</u>
Full	17.63	0.9150	0.0022	0.0214	0.2993	6.78	<u>1.9</u>	<u>4.8</u>	16.4	98%	96%	0.670	<u>8.25</u>	0.566

Design Agent. Without the Design Agent, part IoU falls from 0.2993 to 0.2788 and mass APE rises from 0.566 to 0.710. The mean part count also falls from 6.78 to 5.82, alongside lower single- and double-joint collision rates. This combination shows why collision rates must be interpreted together with structural fidelity: simpler assemblies can exhibit fewer collisions while reproducing the target less accurately. These results suggest that explicit design encourages richer part structures while preserving the intended functionality.

Review Agent. Removing the Review Agent reduces CLIP from 0.9150 to 0.9128 and part IoU from 0.2993 to 0.2951, while increasing mass APE from 0.566 to 0.736. Its collision rates remain close to those of the full system, with a lower multi-joint rate. The independent review thus helps assess whether the generated asset preserves the intended appearance, structure, and physical properties.

Physical Validation. Without physical validation, single-joint collisions rise from 1.9% to 35.7% and multi-joint collisions from 16.4% to 73.2%, while open-state stability falls from 96% to 90%. These findings reinforce the need for simulation feedback when constructing physically grounded assets.

C.3 ADDITIONAL QUALITATIVE EXAMPLES

In this section, we provide more qualitative comparisons for the ArtVIP and LightWheel benchmarks, as well as for stability and collision. For Artcraft, Codex, and our method, we use GPT-6-Astra to generate the assets. For Agy, we use Gemini-3.8-Flash, and for Claude we use Claude-Fable-5.1. For all LLMs, the thinking level is set to high.

Figure 11 provides more qualitative comparisons on the ArtVIP and LightWheel benchmarks, showing the reference image, the ground-truth asset, and the assets generated by the baselines and our method. The results show that our method reproduces the reference appearance more faithfully, recovers part decompositions that follow the ground-truth structure, and matches the ground-truth part motion better, whereas the baselines lose the reference colour, merge or fragment parts, and deviate from the ground-truth asset in the open state.

Figure 12 visualizes more stability results under gravity in MuJoCo. The results show that our method can generate stable assets under gravity, while the baselines often produce unstable assets that fail the stability check.

Figure 13 visualizes more collision results. Our assets are collision-free while most baselines exhibit interpenetration, such as drawers, doors, and racks passing into the surrounding body.




Figure 11: Additional qualitative comparisons on the ArtVIP and LightWheel benchmarks.




Figure 12: Additional visualizations of stability under gravity, each shown at rest and after ten seconds in MuJoCo. A green frame marks an asset that passes the stability check and a red frame one that does not.




Figure 13: Additional collision comparisons. Red and blue highlight the most severely interpenetrating pair of parts, enlarged in the lower-right inset.

D API AND TIME COST

The Design Agent is allocated up to 12 turns for blueprint generation, and the Coding and Review Agents up to 150 turns each. We set an API cost budget for each complete asset-generation run of \$10 for GPT-5.5 and GPT-5.6-Sol, and \$15 for GPT-6-Astra. The repair loop between the Coding Agent and the Review Agent is bounded at five rounds and stops early after two consecutive rounds in which no rubric score rises. To curb the rapid increase in API costs as conversation contexts grow, we trigger context compression at 110,000 tokens by capping and deduplicating tool outputs, dropping superseded renders, and summarizing earlier turns. For each method-model configuration, we report the mean API cost and generation time per asset over all evaluated samples pooled from the available benchmarks.

Table 12: Average API cost and generation time per asset.

Method	Model	Avg. cost (USD)	Avg. time (min)
Ours (ARTISAN)	GPT-5.6-Sol	6.03	34.7
Ours (ARTISAN)	GPT-6-Astra	5.79	34.3
Ours (ARTISAN)	GPT-5.5	7.35	33.9
Vanilla Codex	GPT-5.6-Sol	0.33	9.9
Vanilla Codex	GPT-6-Astra	1.07	6.3
Vanilla Codex	GPT-5.5	0.51	1.9
Articraft	GPT-5.6-Sol	0.94	8.2
Articraft	GPT-6-Astra	0.95	5.4
Articraft	GPT-5.5	1.32	6.0
Articulate-Anything	GPT-5.6-Sol	0.46	6.8
Claude Code CLI	claude-fable-5-1	0.54	1.5
agy CLI	gemini-3.8-flash-high	0.92	7.3

Although our method incurs higher API costs, it produces assets with fewer interpenetrations and greater stability, making them more usable in physical simulation (Table 2). Figure 14 shows the API cost and generation time breakdown of ARTISAN, with simulation time reported separately.




Figure 14: **API cost and generation time breakdown of ARTISAN.** The left three pies show API cost shares of the Design, Coding, and Review Agents; the right three show generation time shares, including the simulator.

E ROBOTIC EXPERIMENTS

We evaluate whether our generated articulated assets provide usable training data for robotic manipulation. Specifically, we assess three properties: (1) physical validity and task feasibility, (2) compatibility with benchmark demonstration-generation pipelines, and (3) downstream policy generalization to unseen native assets.

Experimental Protocol. For each task, we partition native benchmark assets into disjoint training and test sets. We generate the same number of additional assets as the native training set using both

Table 13: **Asset split in each RoboTwin 2.0 task.** Native assets are split into disjoint training and held-out test sets. ARTISAN and Articraft denote generated assets that are physically stable and can be used for downstream evaluation and policy training.

Task	Train	Test	ARTISAN	Train + ARTISAN	Articraft
OpenMicrowave	1	1	6	7	3
OpenLaptop	9	2	9	18	4
TurnSwitch	6	2	6	12	6
Total	16	5	21	37	13

Table 14: **Round 2 qualification test results on RoboTwin 2.0.** Number of assets that pass the qualification test and are successful in generating demonstrations on the RoboTwin 2.0 tasks, for the native assets, the assets generated by ARTISAN, and the assets generated by Articraft.

Task	Train	ARTISAN	Articraft
OpenMicrowave	1/1	6/6	3/6
OpenLaptop	9/9	9/9	4/9
TurnSwitch	6/6	6/6	6/6

Entries denote # successful assets / # total assets.

ARTISAN and Articraft, where applicable, with a minimum number of 6. All generated assets are then checked for readiness through the 2 round quality tests and used exclusively for demonstration collection and policy training, while policy evaluation is performed on held-out native assets.

As shown in Table 3 (b) and discussed in Section 4.1, asset qualification consists of two rounds. **Round 1** verifies that an asset remains physically stable in its intended initial configuration and can reach the benchmark-defined success state under valid interactions. **Round 2** evaluates whether successful demonstrations can be collected using the benchmark’s data-generation pipeline. We initially perform 50 collection trials per asset and retain assets yielding at least one successful demonstration. For retained assets, we allow up to 3,000 trials to collect 50 successful demonstrations. Assets failing qualification are excluded from policy training.

To isolate the effect of asset diversity, we use a clean tabletop configuration for RoboTwin and a fixed kitchen scene for RoboCasa. We compare policies trained on native assets (**Train**), generated assets (**ARTISAN**), and their combination (**Train + ARTISAN**), with Articraft included as a baseline where feasible.

Policy Training. We fine-tune pretrained $\pi_{0.5}$ (Black et al., 2025) policies using LoRA (Hu et al., 2022), with rank 16 for the PaliGemma 2B backbone and rank 32 for the 300M action expert. All policies are trained until convergence with batch size 64 using AdamW ($\beta_1 = 0.9$, $\beta_2 = 0.95$, $\epsilon = 10^{-8}$, weight decay 10^{-10} , gradient clipping 1.0). The learning rate follows cosine decay with 1,000 warmup steps, a peak of 5×10^{-5} , and a configured 30,000-step decay horizon to 5×10^{-6} . EMA is disabled, and the vision encoder remains trainable.

E.1 ROBOTWIN 2.0

Setup. We evaluate three articulated-object manipulation tasks: *Open Microwave*, *Open Laptop*, and *Turn Switch*. Demonstrations are collected using the benchmark-provided scripts and cuRobo (Sundaralingam et al., 2023). Contact points are annotated on assets for compatibility with the scripted manipulation pipeline. The asset splits are provided in Table 13.

Results. Table 15 reports policy success rates on held-out native objects. Augmenting native training assets with ARTISAN improves average success rates over training with native assets alone, indicating that our generated assets introduce useful object-level variation rather than merely increasing performance on generated instances.

Table 15: **Task execution success rate (%) of $\pi_{0.5}$ on RoboTwin 2.0.** All policies are evaluated on held-out native test assets unseen during training. Native, **Ours**, Native + **Ours**, and Articraft denote policies fine-tuned on demonstrations collected from the corresponding asset sets. Higher is better.

Task	Train	Train + ARTISAN	ARTISAN	Articraft
Open Microwave	64%	78%	40%	46%
Open Laptop	98%	97%	97%	86%
Turn Switch	78%	76%	53%	49%
Average	80%	84%	63%	60%

Table 16: **Asset split in each RoboCasa365 task.** Native assets are split into disjoint training and held-out test sets. ARTISAN and Articraft denote generated assets physically able to be used for downstream evaluation and policy training. Each training asset contributes 50 demonstration episodes.

Task	Train	Test	ARTISAN	Train + ARTISAN	Articraft
OpenElectricKettleLid	8	9	10	18	0
OpenFridgeDrawer	5	11	1	6	0
OpenStandMixerHead	9	12	8	17	0
OpenToasterOvenDoor	6	5	3	9	1
SlideDishwasherRack	10	4	1	11	0
SlideOvenRack	3	9	2	5	0
TurnOnStove	10	1	5	15	0
TurnSinkSpout	10	12	5	15	5
Total	61	63	35	96	6

Entries denote the number of assets.

Failure Analysis. As shown in Table 14, Articraft exhibits collision-geometry failures that prevent successful manipulation despite the assets being loadable in simulation. In particular, 5 of 9 generated laptop assets and 3 of 6 microwave assets contain incorrect collision meshes that cause grasp failures and are excluded from the demonstration collection. These cases show that visually plausible articulated assets are not necessarily usable for robotic interaction: collision geometry must also permit physically valid contact and manipulation.

E.2 ROBOCASA365

Setup. We evaluate eight articulated-object manipulation tasks in RoboCasa365, using its MimicGen-based demonstration-generation workflow (Mandlekar et al., 2023). Unlike RoboTwin’s scripted motion-planning pipeline, MimicGen transfers manipulation trajectories from reference demonstrations to new object configurations. The evaluated asset splits are summarized in Table 16.

Asset Qualification and Demonstration Generation. RoboCasa reveals a substantial difference between physical feasibility and compatibility with demonstration generation. Among 80 assets per group, 78 native assets, 76 assets generated by ARTISAN, and 31 Articraft assets pass Round 1. After Round 2, these counts drop to 61, 35, and 6, respectively.

These results reveal two distinct bottlenecks. First, Articraft frequently fails physical qualification, with only 31 of 80 assets satisfying the stability and task-feasibility requirements. Second, even physically qualified assets do not necessarily support successful demonstration collection: 35 of our 76 Round-1-qualified assets and 61 of the 78 qualified native assets pass Round 2. Thus, demonstration-generation compatibility introduces additional constraints beyond physical plausibility.

Table 17 further reports per-task demonstration-generation success rates. Articraft achieves zero successful demonstrations in six of eight task groups, with observed failures including physical instability, incorrect collision geometry, and mesh interpenetration. The results indicate that physically grounded geometry and articulation are essential for generating usable robotic training data.

Table 17: **Round 2 qualification test results on RoboCasa365.** Number of assets that pass the qualification test and are successful in generating demonstrations on the RoboCasa365 tasks, for the native assets, the assets generated by ARTISAN, and the assets generated by Articraft.

Task	Train	ARTISAN	Articraft
OpenElectricKettleLid	8/10	10/10	0/10
OpenFridgeDrawer	5/10	1/10	0/10
OpenStandMixerHead	9/10	8/10	0/10
OpenToasterOvenDoor	6/10	3/10	1/10
SlideDishwasherRack	10/10	1/10	0/10
SlideOvenRack	3/10	2/10	0/10
TurnOnStove	10/10	5/10	0/10
TurnSinkSpout	10/10	5/10	5/10

Entries denote # successful assets / # total assets.

Table 18: **Task execution success rate (%) of $\pi_{0.5}$ on RoboCasa365.** All policies are evaluated on the held-out native test assets listed in Table 16, with 50 evaluation rollouts per test asset. Native, Native + ARTISAN, ARTISAN, and Articraft denote policies fine-tuned on demonstrations collected from the corresponding asset sets. For tasks with no qualified training assets, we assign a success rate of 0% when computing the benchmark average. Higher is better.

Task	Train	Train + ARTISAN	ARTISAN	Articraft
OpenElectricKettleLid	11%	12%	14%	0%
OpenFridgeDrawer	18%	23%	4%	0%
OpenStandMixerHead	17%	29%	12%	0%
OpenToasterOvenDoor	30%	35%	21%	23%
SlideDishwasherRack	6%	13%	72%	0%
SlideOvenRack	87%	84%	59%	0%
TurnOnStove	92%	98%	8%	0%
TurnSinkSpout	51%	49%	28%	19%
Average	39%	43%	27%	5%

Downstream Policy Evaluation. Table 18 reports policy success rates on held-out native assets under different training-asset compositions. These evaluations measure whether demonstrations collected from generated objects provide transferable training signals rather than performance gains restricted to generated instances.

Failure Analysis. We identify two distinct sources of demonstration-generation failure.

Physical and collision failures. Some generated assets cannot be manipulated successfully because of unstable configurations, incorrect collision geometry, or mesh interpenetration. Such failures violate the requirements for physically executable manipulation and prevent reliable demonstration collection.

Trajectory-transfer failures. Other assets are physically feasible but incompatible with the reference interactions used by MimicGen. For example, in *OpenElectricKettleLid*, a generated kettle has a hinged button rather than the sliding button present in the source demonstrations. Although the generated mechanism is physically plausible, the transferred manipulation trajectory does not successfully actuate it (Figure 15).




Figure 15: Close-up comparison of electric kettle buttons. **Top:** The native asset uses a sliding button in the human demonstration. **Bottom:** The generated asset uses a hinged button, which is plausible but prevents successful demonstration transfer with MimicGen.

F AGENT TRACE EXAMPLES

We show outputs, tool calls, and images from two real agent traces; [. . .] marks omissions. Turns count recorded responses within each agent across rounds, with separate intake and blueprint counters for Design.

F.1 STAND MIXER: CONTROL ACCESSIBILITY

Prompt. “A slate blue stand mixer with a heavy rounded head, a matching pedestal and a steel mixing bowl. The head tips back a quarter turn, the bowl lifts off its seat, the whisk spins on the head, the side speed dial turns, and a release button presses in.”

Design Agent Blueprint and design board		design specified
Intake Turn 1. Identifies the appliance, visible motions, and user interactions from the input image. Blueprint Turns 1–4. Authors the blueprint (1), checks its consistency (2), corrects the accessory-cap location (3), and verifies the revision (4). A design board is generated from the resulting blueprint. Selected outputs (abbreviated). Parts: base, mixing head, whisk, speed selector, mixing bowl. Geometry/materials: a rounded $0.29 \times 0.46 \times 0.31$ m slate-blue base; a hollow stainless-steel bowl with an integral handle; an open wire whisk; [...]. Structure/joints: base $\rightarrow$ head, revolute about $+X$, travel 1.48 rad; head $\rightarrow$ whisk, continuous rotation; selector rotation 2.36 rad and bowl lift 0.22 m. Dependencies/tasks: raise the head before lifting out the bowl; turn the selector to choose a setting. Electrical speed selection imposes no mechanical angle coupling to the whisk. Functional constraint: the selector tab must be reachable with the bowl seated. Design board: orthographic, articulated, exploded, and mechanism views.		
Coding Agent First build		construction and checks
Turns 1–13 Construction. uses <code>read_file</code> and <code>lookup_cadquery_api</code>, then <code>write_file</code>, <code>replace</code>, and <code>compile_model</code> to build the five parts and their joints. <code>probe_model</code> and <code>request_render</code> check dimensions and the head connection. Turns 14–33 Motion repair. examines motion feedback, probes intermediate poses, and revises the head housing and bearing geometry. Further compilation and rendering precede handoff.		
   Early build: head raised, bowl removed After motion-clearance revisions Candidate sent for review		
Physical validation First candidate		checks passed
MuJoCo reports passing collision and gravity-stability checks for the submitted candidate. The report supplies the tested poses and measurements to the Review Agent; it does not assess finger access to the selector.		
Review Agent First review		functional repair requested
Turns 1–3 Evidence and tests. reads the program, task bindings, and physics report; uses <code>write_test</code> and <code>run_tests</code> to check articulation, clearances, and bowl removal. All 24 checks pass. Turns 4–9 Focused inspection. uses <code>crop_region</code>, <code>request_render</code>, and <code>probe_model</code> to inspect the neck and selector. The tab is trapped behind the seated bowl, with only 4.499 mm clearance in the rest pose. Turns 10 Feedback. requests a selector location that satisfies the access requirement, together with a more enclosed rear hinge housing. “Revise the stationary socket/bezel location and matching selector joint together so the control is beside the bowl rather than trapped behind its wall. [...] do not solve access by requiring bowl removal.”		
   Original review crop: exposed hinge Original review crop: selector Matched detail: 4.499 mm gap		
Coding Agent Second build		control and housing revised
Turns 34–36 Revision. reads the review, probes the affected parts, renders the hinge, and uses <code>replace</code> to move the selector/socket assembly to an outer pedestal shoulder and add a hinge cover. Turns 37–41 Verification. probes the motion and selector clearance, recompiles, and inspects full and focused renders before resubmission. Preserved: bowl size, all part masses, and all joint ranges. The selector’s placement is changed to meet the blueprint’s higher-level access requirement.		





Figure 16: **Stand-mixer construction and review trace.** Inspection identifies restricted control access despite passing physical checks; the Coding Agent relocates the selector while preserving the bowl and joint ranges. Small panels include original intermediate renders/crops and matched detail views of the archived candidates.

F.2 SHOE CABINET: DRAWER CONNECTIONS

Prompt. “Reconstruct this articulated object. A slim grey shoe cabinet with two small drawers across the top and two tall shoe compartments below, all fitted with round knobs. The drawers pull straight out, while each shoe compartment tilts forward from its lower edge like a hopper. [...]”






Figure 17: **Shoe-cabinet construction and review trace.** Longer runners and revised front rails restore engagement without shortening travel. Under-drawer views use matched cameras and supplementary lighting on the unmodified archived models.