# AffordSim: A Scalable Data Generator and Benchmark for Affordance-Aware Robotic Manipulation

**Mingyang Li**[1,*], **Haofan Xu**[1,*], **Haowen Sun**[1,*], **Xinzhe Chen**[1], **Sihua Ren**[1], **Liqi Huang**[1], **Xinyang Sui**[1], **Chenyang Miao**[1], **Jiawei Ye**[1], **Qiongjie Cui**[1], **Xingyu Chen**[1], **Zeyang Liu**[1,†], **Xuguang Lan**[1,†]

[1]National Key Laboratory of Human-Machine Hybrid Augmented Intelligence, Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University.

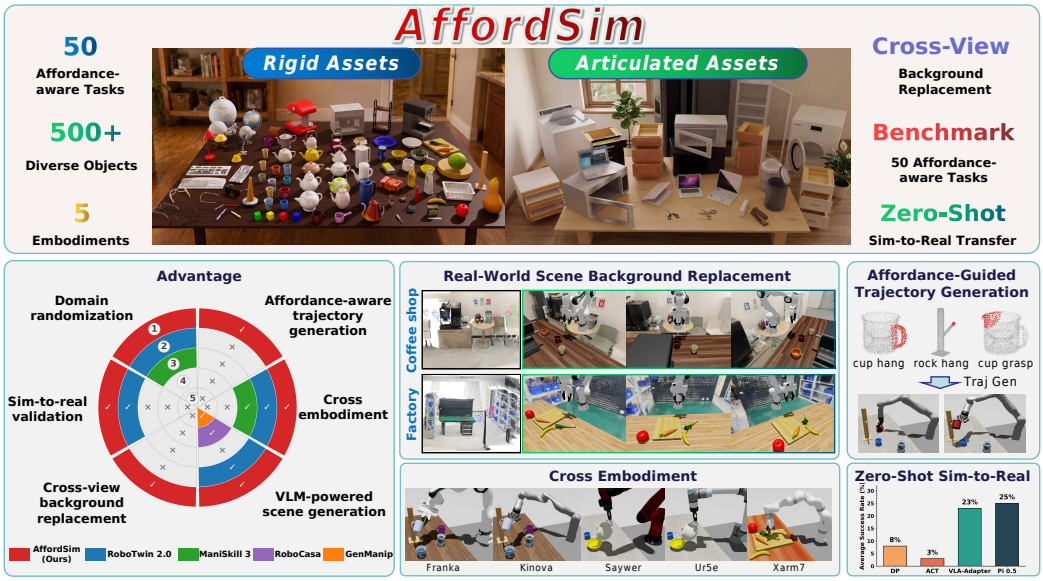



Figure 1: **Overview of AffordSim.** A scalable affordance-aware manipulation benchmark of 50 tasks, 500+ rigid and articulated objects, and 5 robot embodiments, paired with cross-viewpoint background replacement and affordance-guided trajectory generation.

## Abstract

Many everyday robot manipulation skills are affordance-dependent, with success determined by whether the robot contacts the functional object region required by the subsequent action. Current simulation data generators obtain contacts from generic grasp estimators or per-object manual contact annotations, but generic estimators rank stable grasps without task semantics and often select contacts that are misaligned with the downstream action, while manual contact annotations must be rewritten for each new object and task. To solve these challenges, we introduce **AffordSim**, a scalable data generator and benchmark that integrates open-vocabulary 3D affordance prediction into simulation-based trajectory generation.

*Equal contribution.

†Correspondence to: Zeyang Liu (zeyang.liu@xjtu.edu.cn), Xuguang Lan (xglan@xjtu.edu.cn).

40th Conference on Neural Information Processing Systems (NeurIPS 2026). Track on Evaluations and Datasets.

Given a natural-language task description, AffordSim synthesizes a task-relevant scene, emits affordance queries, grounds them on object surfaces, samples region-conditioned grasps, and selects executable candidates with motion planning. It further randomizes object pose, texture, lighting, image noise, and cross-viewpoint backgrounds for sim-to-real transfer. We instantiate AffordSim as a 50-task benchmark across diverse manipulation skills, five robot embodiments, and 500+ rigid and articulated objects. AffordSim achieves $93\%$ of the trajectory collection success rate of manual contact annotations on affordance-critical tasks and $89\%$ on hard composite tasks. Vision-language-action policies trained on AffordSim data transfer zero-shot to a real Franka FR3, reaching $24\%$ average success. The data generator, benchmark, and code are available at project page.

## 1 Introduction

Simulation has become an important source of training trajectories for robotic manipulation [1–5]. A critical but often implicit component of these trajectories is object affordance—the region of an object the robot must engage with for the task to succeed. Its importance becomes most apparent when a stable grasp on the wrong region still forecloses the downstream step. Consider hanging a mug on a rack, the task requires grasping the cup body so that the handle remains free for the hook. Grasping the handle instead, however stable the grasp itself may be, leaves no surface for the hook to engage. Such affordance-dependent tasks are routine in everyday environments, and a simulator that cannot generate data for them excludes a large share of real-world manipulation.

Existing simulation platforms collect trajectories in one of two ways. The first relies on manual contact annotations, where a human authors a task-specific grasp pose for every object [3, 4, 6]. The second relies on generic grasp estimators that score candidates at collection time by grasp quality [7], which removes the per-object annotation cost. Both approaches, however, struggle on affordance-dependent tasks. Manual annotation achieves high success on the annotated set but does not scale, since every new asset and every new task demands a fresh round of annotation. Generic estimators scale, but because they ignore affordance, they rank candidates by grasp quality rather than by task-relevant region, often landing on the wrong region. Patching this with a semantic mask filter on the estimator's output [8] cannot recover candidates the estimator never proposed inside the desired region in the first place. A deeper limitation of the estimator route is that it models only the grasping action, providing no signal for object-to-object interaction [9, 10].

Despite their differences, both routes share the same root cause: affordance enters trajectory collection only implicitly, never computed first as a separate step that constrains grasp generation. Manual annotation freezes the author's affordance reasoning into a static per-object pose at annotation time, demanding fresh annotation for every new object and every new task. Generic estimators ignore affordance entirely in their base form; bolting on a semantic mask after the fact only filters candidates that were already chosen without affordance in mind, which means the mask cannot recover candidates the estimator never proposed inside the desired region.

We instead make affordance explicit, computing it first as a separate step that constrains grasp generation and reshapes the candidate space considered at collection time. Implicit pipelines choose grasps first and apply any affordance signal afterward, so the candidate space is shaped without affordance and the right region can only be recovered by filtering. Our explicit pipeline shapes the candidate space from the start: grasp candidates are sampled only inside the predicted region, and feasibility scoring runs only on those candidates. Stable grasps that land in the wrong region are eliminated at the source rather than filtered out after the fact.

We realize this pipeline in **AffordSim**, a simulation framework built on NVIDIA Isaac Sim [11] that turns a natural-language task description into executable manipulation trajectories. A VLM-based scene generator parses the description into a structured scene graph and emits affordance queries, which an open-vocabulary 3D affordance model [12] resolves into localized interaction regions on the object point cloud. A two-stage grasp selector then samples candidates within each region and ranks them by motion-planning feasibility, yielding executable grasps without per-object annotation. For sim-to-real transfer, a domain randomizer perturbs object pose, lighting, and surface texture, while cross-viewpoint backgrounds are reconstructed from real photographs via 3D Gaussian Splatting [13]. Beyond data collection, AffordSim also serves as a systematic evaluation suite, by curating affordance-

critical tasks, it measures how well a policy actually understands object affordance rather than merely succeeds on grasp-stable scenarios. Such evaluation, however, has been out of reach for existing platforms, which lack benchmarks for affordance-dependent tasks, leaving a gap that AffordSim is designed to fill. Policies trained with AffordSim data further exhibit zero-shot sim-to-real transfer, demonstrating the practical utility of the collected trajectories.

We make three contributions. (1) We present AffordSim, the first simulation framework built around affordance-aware manipulation, driving data generation with open-vocabulary 3D affordance prediction in place of per-object annotation. (2) We release a 50-task affordance-aware benchmark spanning diverse manipulation skills and five robot embodiments, targeting affordance-dependent tasks that existing benchmarks largely omit; on these tasks, AffordSim collects trajectories without any per-object annotation and reaches $93\%$ of the success rate achieved by platforms with manual contact annotations. (3) Vision-language-action policies trained exclusively on AffordSim data reach $30\%$ on placing, $25\%$ on pouring, and $10\%$ on hanging in zero-shot evaluation on a real Franka FR3, averaging $24\%$.

## 2 Related Work

### 2.1 Simulation Platforms for Robotic Manipulation

Existing simulation platforms have scaled data generation along axes such as scene diversity, GPU-parallelized rendering, and unified task suites [1–3, 5, 11, 14–25], but affordance reasoning remains absent from their trajectory collection. RoboTwin 2.0 [6] is the only exception, driving bimanual data generation through manually authored per-object affordance annotations that do not extend to new objects or tasks. All other platforms fall back on generic grasp pipelines that ignore functional constraints.

A complementary line of work [8, 26–28] uses LLMs or VLMs to automate task and scene authoring or to scale demonstrations from a small seed set, but inherits the same affordance gap. GenSim2 and RoboGen still rely on generic grasp pipelines with no per-object affordance signal. GenManip filters AnyGrasp candidates against a task-relevant 2D mask derived from SAM2 [29, 30] segmentation and VLM mark-based selection [31]; affordance therefore enters only as a post-hoc filter on grasps ranked by stability, and when AnyGrasp returns no stable candidate inside the region, no filtering can recover one. Predicting affordance directly on 3D object geometry, by contrast, constrains grasp generation upstream and removes this failure mode. Related VLM manipulation systems on real images [32–34] inherit the same limitation.

### 2.2 3D Affordance Prediction

Affordance prediction has been studied as a standalone perception problem. Where2Act [35], 3D AffordanceNet [36], and GAPartNet [37] predict actionable regions on point clouds over a fixed label set, evaluated on detection metrics rather than manipulation trajectories. A recent line tackles open-vocabulary 3D affordance prediction with free-form language queries [38–40]. In particular, VoxAfford [12] enriches MLLM output tokens with multi-scale geometric features from a frozen 3D VQVAE encoder. Even so, the per-point affordance scores are not directly executable. Turning them into kinematically feasible grasp poses requires a downstream stage that prior affordance work does not provide. AffordSim connects region prediction to trajectory collection through grasp sampling and motion-planning feasibility checks.

### 2.3 Domain Randomization

Domain randomization is a standard approach for sim-to-real transfer [41, 42]. Prior work diversifies visual and physical parameters such as textures, lighting, camera pose, mass, and friction during training, yielding visuomotor policies that transfer to real robots. SIMPLER [43] uses "green screening" to composite simulated robots and interactive objects onto a real background photograph. Its Visual Matching setup is limited to fixed cameras, with camera parameters matched to the reference image. Thus, changing the viewpoint requires a corresponding background image and camera alignment; the photograph itself cannot provide a new view. AffordSim renders reconstructed real scenes from arbitrary camera viewpoints and composites the resulting background images with simulated foreground objects under the same camera configuration (Section 3.3).

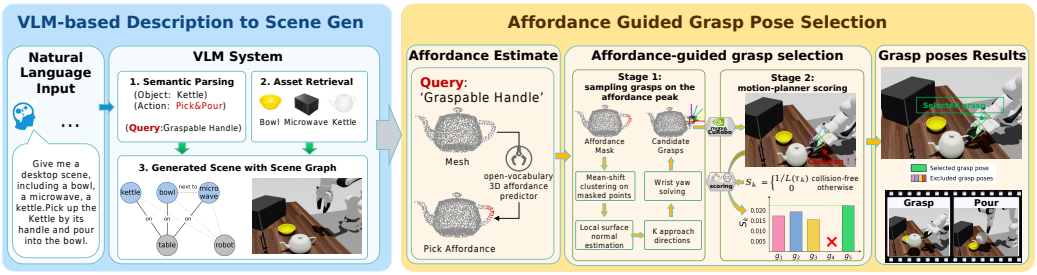



Figure 2: **Affordance-guided data generation pipeline of AffordSim.** The full module breakdown and data flow are described in Section 3. In the Stage 2 bar chart, the x-axis lists the candidate grasps $g_k$ generated in Stage 1 and the y-axis gives their motion-planner score $S_k$, with the red cross marking the colliding candidate ($S_k = 0$) and the green dashed line indicating the maximum-score grasp ultimately selected.

# 3 AffordSim Framework

AffordSim turns a natural-language task description into visually diverse manipulation demonstrations for imitation training. The framework chains three modules. A VLM-based generator (Section 3.1) separately parses scene and task descriptions, assembles the scene from a USD asset library, and converts object–action entries into queries for the downstream affordance model. A two-stage grasp selector (Section 3.2) processes each query. Stage 1 runs an open-vocabulary 3D affordance model on the object point cloud and samples grasp candidates on the predicted affordance peak. Stage 2 scores these candidates with a motion planner and returns the highest-ranked grasp together with its pre-grasp trajectory. A domain randomizer (Section 3.3) perturbs object pose, lighting, surface texture, and cross-viewpoint reconstructed backgrounds to broaden the visual distribution. Figure 2 illustrates the first two modules, and Figure 3 shows the domain randomizer.

## 3.1 VLM-Powered Task and Scene Generation

Authoring a simulation scene by hand for every new manipulation task does not scale. Each task needs compatible assets, a plausible layout, and a description of the functional region on each manipulated object. We use separate VLM calls [44, 45] for scene and task parsing, instantiated with Qwen-VL-Plus [46] in our implementation. The scene parser extracts objects and relations; the task parser emits object–action entries that are converted into functional-region queries. Each node in the scene graph binds a USD asset from an indexed library.

**Scene graph generation.** From the scene description, the VLM extracts objects and relations that the framework expands into a scene graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node $v \in \mathcal{V}$ stores a category, a USD asset path, an initial 6-DoF pose, a color, and a flag indicating whether the object is manipulated or static. Each edge $(u, v) \in \mathcal{E}$ encodes a spatial relation parsed from the prompt. Relations such as "inside", "on top of", and "left of" constrain the placement of the corresponding object pairs. Collision-free placement is enforced by rejection sampling on the proposed poses under the spatial constraints in $\mathcal{E}$. The simulation scene is then assembled by loading each USD asset at its planned pose. The full prompt is given in Appendix R.

**Affordance query extraction.** In a separate VLM call, the task description is parsed into ordered primitives and object–action entries. The system converts these entries into one query per functional contact region, so every region the policy must engage receives its own query. The VLM reasons over the manipulation primitives and the role of each object. For pouring, the entries distinguish the held source object from the receiving container, producing separate queries for each role. Each converted query is then passed to the open-vocabulary affordance model on the corresponding point cloud. The full prompt covers nine manipulation primitives (pick, place, pour, push, pull, open, close, hang, hold) and is given in Appendix S.

**Asset retrieval.** For each node, the VLM specifies a target category and required semantic attributes such as color or articulation type. The framework looks up a matching USD asset in the library and binds its path to the node before instantiation.

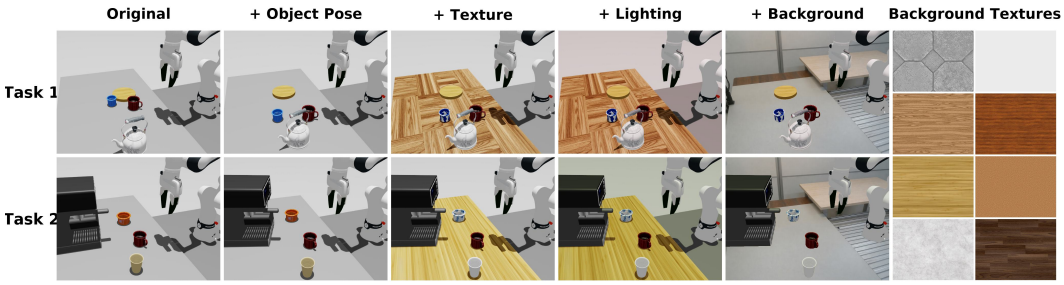



Figure 3: **Domain randomization in AffordSim.** Each column shows the cumulative effect of adding one randomization axis: object pose, surface texture, lighting, and background. The rightmost column shows example background textures used for randomization. Two representative tasks are shown across rows.

### 3.2 Affordance-Aware Trajectory Collection

We propose affordance-guided grasp selection, which turns each affordance query from Section 3.1 into a manipulation trajectory by combining open-vocabulary 3D affordance prediction with motion planning. A trajectory $\boldsymbol{\tau} = \{(o_t, a_t)\}_{t=1}^{T}$ is an observation-action sequence where $o_t$ contains visual, proprioceptive, and scene state, and $a_t$ is the robot action. For each manipulated object, we pass the canonical point cloud $\mathcal{P} \in \mathbb{R}^{N\times 3}$ shipped with its USD asset, together with the corresponding affordance query, to the open-vocabulary 3D affordance model. The model outputs a per-point binary mask $M_a \in \{0,1\}^N$ marking the queried affordance region.

**Affordance-guided grasp selection.** Grasp selection runs in two stages. Stage 1 decides where to grasp by sampling candidates on the predicted affordance peak. Stage 2 decides which candidate the robot can reach by scoring with motion planning and keeping the shortest collision-free trajectory.

*Stage 1: sampling grasps on the affordance peak.* We run mean-shift clustering on the predicted affordance points $\{p \in \mathcal{P} : M_a(p) = 1\}$ and take the cluster centroid $p$ as the representative contact point on the object surface. Appendix N extends this procedure to disconnected masks. At $p$, we estimate the local surface normal $n$ by running PCA on $p$'s $m$ nearest neighbors. We take the eigenvector with the smallest eigenvalue and orient it outward from the object centroid. We then sample $K$ approach directions $\{d_k\}_{k=1}^{K}$ on the hemisphere opposite to $n$, with azimuth $\phi_k = 2\pi k/K$ around $n$ and a fixed polar tilt $\theta_0$ from $-n$. For each direction, the wrist yaw is solved so that the gripper closing axis is perpendicular to $d_k$ and aligned with the local affordance geometry. This produces $K$ candidate grasps $\mathcal{C} = \{\mathbf{g}_k = (p, R_k)\}_{k=1}^{K}$.

*Stage 2: motion-planner scoring.* For each candidate $\mathbf{g}_k$, we use cuRobo [47] to attempt to plan a collision-free pre-grasp trajectory $\boldsymbol{\tau}_k$ from the current joint configuration to $\mathbf{g}_k$. The score is then

$$s_k = \begin{cases} 1/L(\boldsymbol{\tau}_k), & \text{if } \boldsymbol{\tau}_k \text{ is collision-free,} \\ 0, & \text{otherwise,} \end{cases} \tag{1}$$

where $L(\boldsymbol{\tau}_k)$ is the joint-space arc length. The selected grasp $\mathbf{g}^* = \mathbf{g}_{\mathrm{argmax}_k\ s_k}$ reuses its already-planned $\boldsymbol{\tau}_{k^*}$ as the pre-grasp trajectory.

**Motion planning and execution.** The pre-grasp trajectory $\boldsymbol{\tau}_{k^*}$ is followed by gripper closure, lift, and task-completion segments, each planned by cuRobo against the scene's collision meshes. The simulator then executes the full task trajectory $\boldsymbol{\tau}$.

### 3.3 Domain Randomization for Sim-to-Real Transfer

To close the visual gap between simulation and the real world, we render each generated trajectory under domain randomization along five axes (Figure 3). **Background substitution** replaces the workspace background with either random textures from a diverse texture library or cross-viewpoint real-scene renders. For the latter, we reconstruct each real deployment scene with 3D Gaussian Splatting [13] from 10–20 photographs, using Depth Anything 3 [48] for view synthesis. The resulting 3D Gaussian field can be rendered from arbitrary camera viewpoints. The *Real-World Scene Background*

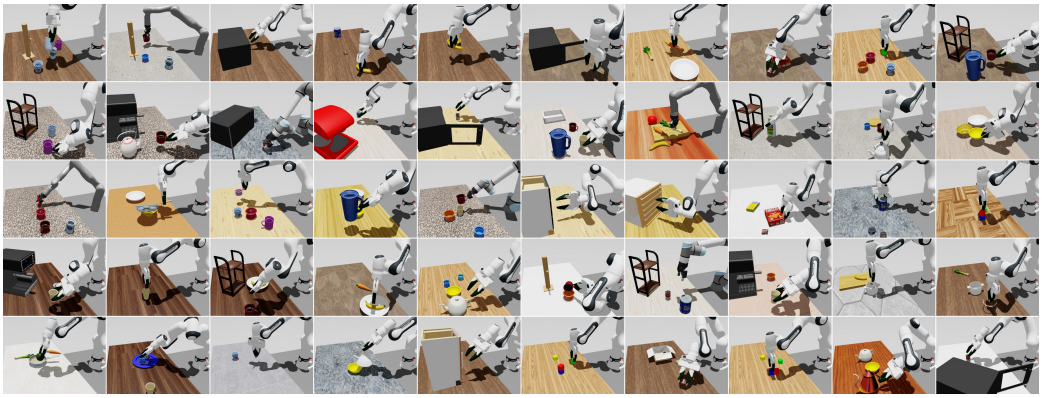


Figure 4: **Task gallery.** Representative frames from all 50 tasks of the AffordSim benchmark, grouped into seven manipulation categories: pick & place (20), open/close (2), pull/push (5), hang (2), pour (7), stack (4), and long-horizon composite (10).

*Replacement* panel of Figure 1 shows an example, and full reconstruction hyperparameters with three example scenes are given in Appendix Q. The remaining four axes are standard, namely **lighting** (type, position, intensity, and color of light sources), **object texture** (PBR albedo, photo asset library), **object pose** (initial xyz position and quaternion orientation), and **image noise** (Gaussian pixel noise on the rendered RGB).

### 3.4 Benchmark Design

We introduce an affordance-aware manipulation benchmark of 50 tasks across seven manipulation skills (Figure 4). The skills are pick & place (20), open/close (2), pull/push (5), hang (2), pour (7), stack (4), and long-horizon composite (10). The benchmark targets tasks whose success hinges on contacting specific functional regions. This allows affordance reasoning to be evaluated directly through task success, instead of through generic grasp performance or reinforcement learning return curves [49, 50]. Each task runs unchanged across five robot embodiments (Franka, Kinova, Sawyer, UR5e, xArm7). Each task is evaluated by its **task success rate (SR)** over 100 simulation rollouts. Success is checked automatically from rigid-body poses and articulation states. Multi-step tasks additionally report per-stage SRs. Per-skill demands are detailed in Appendix J.

**Task difficulty tiers.** We further organize the 50 tasks into three difficulty tiers, primarily graded by their demands on affordance reasoning and manipulation horizon.

- **Easy, Generic-Grasp Tasks** (10 tasks). The manipulated object admits any stable grasp and the manipulation horizon is short, so generic grasp estimators can already collect data successfully.
- **Medium, Affordance-Critical Tasks** (30 tasks). Success requires contact at a specific functional region of the manipulated object while the horizon remains short, and generic grasp estimators fail on most of them.
- **Hard, Composite Affordance Tasks** (10 tasks). These tasks combine functional-region contact constraints with either long-horizon sequential dependencies in multi-step composites, or with unusually tight precision and stability margins in short-horizon settings. Per-task notes for the short-horizon Hard cases are given in Appendix K.

The contact-region demand separates Medium from Easy, and the performance gap between the two reflects the contribution of affordance reasoning to short-horizon trajectory collection. The full task lists for the three tiers are given in Appendix K.

**Rationale for affordance-based difficulty.** Within tabletop organization and manipulation, affordance structure fixes functional contacts and primitive sequence; embodiment controls realizability, and randomization tests visual robustness. We therefore grade difficulty by affordance demand and task horizon to isolate the contribution of affordance understanding. Structurally, small functional regions such as cup rims require more precise grounding than objects graspable from many directions, and longer task chains compound sub-step failures. Empirically, collection and policy SR decrease

from Easy to Hard (Tables 1 and 2), consistent with lower real-world success on precise hanging (Table 3). Including randomization in the tiers would conflate choosing functional contacts with robustness to changing scenes. Keeping tiers fixed allows contact reasoning, arm realizability, and visual robustness to be evaluated separately. Randomization remains important: AffordSim supports it, RoboTwin 2.0 [6] emphasizes that axis, and Section 4.5 quantifies its effect on real-world transfer.

**Asset library.** The benchmark is backed by an indexed library of 500 rigid and articulated USD objects from three sources. 50 objects are scanned in-house with a Raptor Pro 3D scanner to match the physical deployment objects. Another 50 come from YCB [51]. The remaining 400 objects come from PartNet-Mobility [19]. A subset of 50 objects additionally carries human-annotated affordance ground truth, which serves as the oracle reference in Table 1.

# 4 Experiments

We evaluate AffordSim along four axes, covering (1) the contribution of affordance integration to trajectory collection on the full 50-task benchmark (Section 4.1), (2) imitation learning performance on the resulting data (Section 4.2), (3) an ablation of the affordance model (Section 4.3), and (4) zero-shot sim-to-real transfer (Section 4.4).

## 4.1 Affordance Integration on the 50-task Benchmark

The predicted affordance must identify the functional contact region for a feasible grasp to be collected, and in turn for the trajectory to be collected at all. We test this across all 50 tasks of the benchmark by comparing five grasp methods under identical scene and motion-planner settings. Each task is run 100 times under aggressive randomization, with the manipulated object's initial pose drawn from a 30 cm horizontal window and a 180° yaw range. We report the trajectory collection success rate (SR).

**Grasp methods.** **Manual** is a manually tuned heuristic program that reads each object's pose from the simulator and grasps within a per-object xyz range encoding a coarse human prior. The full specification, including the simple-geometry and functional-region regimes, is in Appendix H. **AnyGrasp** [7] is a generic 6-DoF grasp estimator that ranks candidates by stability without task semantics. **AnyGrasp+Filter** adds a 2D SAM2 task-relevant mask on top of AnyGrasp candidates and keeps only those whose contact point falls inside the mask. This is the post-hoc filtering scheme of GenManip [8] discussed in Section 2. **AffordSim** is our affordance-guided grasp selection built on an open-vocabulary 3D affordance model, as described in Section 3.2. **Human Aff.** reads the human-annotated affordance ground truth available on the 50-object oracle subset and serves as an oracle upper bound. Table 1 reports per-task SR, with two representative tasks per tier and the rest abbreviated. Figure 5 summarizes per-tier averages alongside the sim-to-real numbers from Section 4.4.

Table 1: **Affordance integration on the AffordSim benchmark.** Trajectory collection success rate (%) per task across the five grasp methods. AnyGrasp+Filter denotes AnyGrasp candidates filtered against a 2D SAM2 task-relevant mask, mirroring the GenManip-style affordance-filtering pipeline. Two representative tasks per tier are listed; full breakdown in Appendix E. Tier averages bolded.

| Tier | Task | Manual | AnyGrasp | AnyGrasp+Filter | AffordSim | Human Aff. |
|---|---|---|---|---|---|---|
| Easy | *pick banana place bowl* | 100 | 80 | 100 | 100 | 100 |
| | *pick soupcan place table* | 100 | 55 | 75 | 95 | 100 |
| | ··· (8 more tasks) | ··· | ··· | ··· | ··· | ··· |
| | **Average (10 tasks)** | **100** | **67** | **87** | **98** | **100** |
| Medium | *pick mugcup place coffeemachine* | 90 | 25 | 55 | 88 | 95 |
| | *hang cup* | 85 | 0 | 0 | 25 | 90 |
| | ··· (28 more tasks) | ··· | ··· | ··· | ··· | ··· |
| | **Average (30 tasks)** | **85** | **15** | **45** | **79** | **93** |
| Hard | *stack four cube* | 82 | 18 | 33 | 77 | 88 |
| | *composite double pour kettle* | 70 | 0 | 15 | 65 | 85 |
| | ··· (8 more tasks) | ··· | ··· | ··· | ··· | ··· |
| | **Average (10 tasks)** | **72** | **3** | **16** | **64** | **86** |

The five methods trace the two failure modes from Section 1. Manual matches the Human Aff. oracle on Easy and Medium but loses ground on Hard and does not scale (Appendix H), marking the upper bound of manually annotated platforms like RoboTwin 2.0. AnyGrasp collapses beyond Easy, and AnyGrasp+Filter only lifts tier averages by $+20$, $+30$, and $+13$ since the filter cannot recover candidates the stability ranking never sampled. Both fail end to end on hang tasks ($0\%$ in Table 1) because generic grasp estimators emit no signal for the post-grasp object-to-object interaction; substituting the Manual post-grasp routine to isolate the grasp stage, AnyGrasp / AnyGrasp+Filter reach $45\%$ / $60\%$ on *hang cup* and $46\%$ / $63\%$ on *hang mugcup*, but recovering even this much requires per-object Manual tuning. AffordSim itself trails Manual on these hang tasks ($25\%$ / $30\%$ vs $85\%$ / $70\%$) because the configuration is OOD for its affordance backbone, yet across tier averages it closes most of the gap to Manual without per-object tuning and stays close to the oracle. The AnyGrasp-to-AffordSim jump grows monotonically with difficulty while the oracle gap widens, identifying affordance prediction as the missing ingredient and locating the remaining Hard headroom in the affordance model.

AffordSim exhibits an out-of-distribution (OOD) failure mode unique to its affordance-driven pipeline. On every hang task, AffordSim stays below $30\%$, well below its tier average, because grasping the cup body while keeping the handle free for the rack hook is rare in the affordance backbone's training data. AnyGrasp's $0\%$ on the same tasks is a separate, architectural limitation rather than an OOD one as discussed in Section 1, generic grasp estimators only output stable grasp poses and provide no signal for the object-to-object interaction that the downstream hang motion requires, so the hook-engagement constraint never enters their candidate space in the first place.

**Cross-embodiment results.** Across all five embodiments, AffordSim achieves 76–$83\%$ collection SR on the 50-task benchmark, exceeding AnyGrasp+Filter by 32–33 percentage points on every arm (Appendix D).

### 4.2 Imitation Learning on AffordSim Data

We next ask whether modern imitation learning baselines can exploit AffordSim data, covering both classical visuomotor policies and the recent vision-language-action (VLA) lineage [52–59]. We evaluate **BC**, **Diffusion Policy (DP)** [60], **ACT** [61], **VLA-adapter** [62], and **Pi 0.5** [63] on 20 tasks across the seven manipulation skills. Each task uses 300 demonstrations, with RGB and proprioceptive state as input. Architectures and hyperparameters are in Appendix I. We randomize only the object pose, with a $15$ cm window and $30^\circ$ yaw, while keeping background, lighting, and texture fixed. This isolates task structure from visual domain shift.

Two patterns emerge. First, success rate degrades sharply with tier difficulty across all baselines. Easy averages cluster high, Medium drops by roughly a third, and Hard collapses below $30\%$ even for the strongest policies. This identifies affordance precision and long-horizon chaining as the dominant difficulty axes. Second, Pi 0.5 consistently outperforms BC across difficulty tiers. On Hard tasks, their average SRs are $24\%$ and $1\%$, respectively, a gap of 23 percentage points. Among the baselines, the foundation models Pi 0.5 ($55\%$) and VLA-adapter ($54\%$) outperform DP, ACT, and BC. This suggests that the multimodal action distributions of affordance-guided trajectories require expressive policy classes.

**Sub-task analysis.** Per-stage success rates for every multi-stage task are in Appendix O. The pick stage succeeds reliably across the benchmark, with VLA-adapter and Pi 0.5 at 93–$100\%$. AffordSim grasps are therefore executable, and failures concentrate downstream. The post-grasp drop scales with alignment precision and chain length. Tasks with tighter functional alignment achieve lower SR, with pour scoop at 31–$67\%$, pour kettle at 26–$56\%$, and hang red mug at 21–$36\%$. Long-horizon tasks compound per-stage failure multiplicatively. Stack three cubes degrades from $93\%$ at pick 1 to $47\%$ at place 2 for Pi 0.5. The bottleneck is precise post-grasp alignment and long action chains, not the affordance-grounded grasp itself.

### 4.3 Affordance Model Ablation

We isolate the contribution of the affordance backbone by holding every other module of AffordSim fixed and swapping only the affordance model. We compare VoxAfford [12] against four open-vocabulary 3D affordance baselines, namely OpenAD [38], IAGNet [64], LASO [65], and 3D-

Table 2: **Benchmark evaluation results** (Success Rate %) on 20 representative tasks organized by the three difficulty tiers of Section 3.4. Each task is evaluated over 100 rollouts. Best result per task in **bold**; tier averages bolded.

| Tier | Task | BC | DP | ACT | VLA-adapter | Pi 0.5 |
|---|---|---|---|---|---|---|
| Easy | *pick banana place plate* | 40 | 82 | 47 | 87 | **93** |
| | *pick leek place cuttingboard* | 33 | 70 | 68 | 84 | **86** |
| | *pick chocolatebox place sugarbox* | 43 | 76 | 79 | 88 | **89** |
| | **Tier average (3 tasks)** | **39** | **76** | **65** | **86** | **89** |
| Medium | *pick cup place shelf* | 7 | 28 | 26 | 55 | **61** |
| | *pick mug place coffee machine* | 10 | 33 | 27 | 51 | **53** |
| | *pick lid place pot* | 13 | 37 | 17 | **53** | 50 |
| | *pick soupcan place microwave* | 21 | 45 | 41 | 67 | **69** |
| | *pour scoop into bowl* | 13 | 35 | 31 | 62 | **67** |
| | *pour red mug into bowl* | 3 | 37 | 30 | 62 | **68** |
| | *pour kettle into bowl* | 3 | 33 | 26 | **56** | 53 |
| | *stack three cubes* | 4 | 38 | 36 | **57** | 47 |
| | *hang white cup on rack* | 4 | 23 | 10 | **47** | 42 |
| | *push can to target* | 27 | 57 | 53 | 70 | **73** |
| | *pull cabinet open* | 24 | 60 | 57 | **74** | 71 |
| | *open microwave* | 20 | 38 | 40 | **60** | 57 |
| | **Tier average (12 tasks)** | **12** | **39** | **33** | **60** | **59** |
| Hard | *hang red mug on rack* | 3 | 23 | 21 | **36** | 34 |
| | *pour kettle and teapot* | 0 | 4 | 2 | 18 | **21** |
| | *store cups on plates* | 0 | 4 | 0 | 10 | **12** |
| | *pour sugar pick mug into coffeemachine* | 0 | 4 | 4 | 18 | **20** |
| | *organize the clutter into baskets* | 0 | 15 | 9 | 32 | **35** |
| | **Tier average (5 tasks)** | **1** | **10** | **7** | **23** | **24** |
| | **Overall (20 tasks)** | **13** | **37** | **31** | **54** | **55** |

AffordanceLLM [40]. All baselines receive the same affordance queries and simulation settings as in Section 4.1, evaluated on 10 tasks split into 4 Easy, 3 Medium, and 3 Hard with 100 rollouts each. Trajectory SR closely tracks the published mIoU ranking on the OpenAD detection benchmark [12], with VoxAfford leading on every tier and OpenAD trailing. The gap widens with difficulty, since Easy tolerates coarse localization while Medium and Hard make affordance errors unrecoverable downstream. This places affordance prediction, rather than motion planning or visual randomization, as the dominant performance lever on Medium and Hard. Full numbers are in Appendix F.

### 4.4 Zero-Shot Sim-to-Real Transfer

We deploy all five baselines, trained exclusively on AffordSim data, to a real Franka FR3 on four tasks across placing, pouring, and hanging. Each task is run for 10 trials with no real world fine tuning. The workspace is reconstructed via 3D Gaussian Splatting with Depth Anything 3. Table 3 and Figure 6 report the results. Pi 0.5 leads at $25\%$ average success, with VLA-adapter at $23\%$. DP, ACT, and BC lag at $8\%$, $3\%$, and $0\%$ respectively. The baseline ordering and per-task gradient (placing up to $40\%$, pouring $30\%$, hanging $10\%$) mirror Table 2. This confirms that the bottleneck is genuine rather than a simulation artifact.

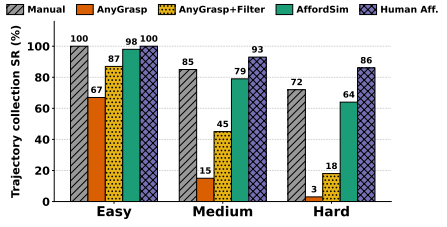
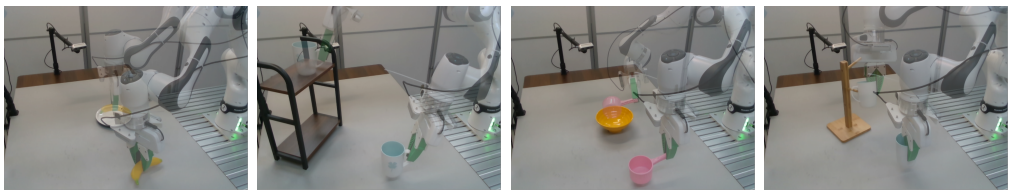



Figure 5: **Per-tier average trajectory collection SR (%).**

Table 3: **Zero-shot sim-to-real transfer.** SR (%) on a real Franka FR3 (10 trials per task, no real world fine tuning). Best per task in **bold**.

| Task | BC | DP | ACT | VLA-A. | Pi 0.5 |
|---|---|---|---|---|---|
| *pick banana place plate* | 0 | 20 | 10 | **40** | **40** |
| *pick cup place shelf* | 0 | 0 | 0 | **20** | **20** |
| *pour scoop into bowl* | 0 | 10 | 0 | 20 | **30** |
| *hang red mug on rack* | 0 | 0 | 0 | **10** | **10** |
| **Average** | **0** | **8** | **3** | **23** | **25** |



*pick banana place plate* *pick cup place shelf* *pour scoop into bowl* *hang red mug on rack*

Figure 6: **Zero-shot sim-to-real deployment.** The four real robot tasks in Table 3, executed by Pi 0.5 on a Franka FR3 with no real world fine tuning.

### 4.5 Domain Randomization Ablation

We ablate the randomization axes of Section 3.3 on a single real-world task, *pick banana place plate* from Section 4.4, using all five baselines zero-shot. Each setting uses 10 trials. Panel (a) follows Section 4.4 training; panel (b) uses B+P. BC remains at $0\%$ throughout and is omitted.

Table 4: **Domain randomization and real-world success.** SR (%) on *pick banana place plate* over 10 trials per setting. B: reconstructed background; T: texture-based background replacement; P: object pose; L: lighting. Panel (b) varies the training position window; deployment uses a fixed 10 cm window and $30^\circ$ yaw.

**(a) Randomization type**

| DR axes | DP | ACT | VLA-A. | Pi 0.5 |
|---|---|---|---|---|
| B+P | 20 | 10 | 40 | 40 |
| P+L | 10 | 0 | 30 | 30 |
| B+L | 0 | 0 | 10 | 10 |
| T+P+L | 20 | 10 | 30 | 40 |
| B+P+L | 20 | 10 | 40 | 50 |

**(b) Object-pose range**

| Range | Demos | DP | ACT | VLA-A. | Pi 0.5 |
|---|---|---|---|---|---|
| None | 300 | 0 | 0 | 10 | 10 |
| 10 cm | 300 | 10 | 10 | 20 | 30 |
| 15 cm | 675 | 20 | 10 | 30 | 40 |
| 20 cm | 1200 | 20 | 10 | 40 | 50 |
| 30 cm | 2700 | 20 | 10 | 40 | 40 |

For Pi 0.5, replacing the reconstructed background with textures lowers SR from $50\%$ to $40\%$, removing background replacement lowers it to $30\%$, and removing pose variation lowers it to $10\%$. Lighting adds 10 points for Pi 0.5, with the other baselines unchanged. The pose sweep holds demonstration density constant for nonzero windows, so data volume grows quadratically with range. A 10 cm training window, matching deployment, yields $30\%$ for Pi 0.5; expanding it to 20 cm reaches $50\%$. At 30 cm, Pi 0.5 falls to $40\%$ while VLA-adapter remains at $40\%$, despite more demonstrations. Thus, training beyond the deployment position range can improve transfer, but further expansion does not continually increase success.

### 4.6 Grasp Source Ablation for Policy Learning

With a fixed demonstration budget, ACT achieves average simulation SRs of $21.3\%$, $24.3\%$, $33.0\%$, and $34.0\%$ using AnyGrasp, AnyGrasp+Filter, AffordSim, and Human Aff. data, respectively, across three tasks (100 rollouts per task). AffordSim improves success over both generic grasp sources on every task, demonstrating the influence of grasp selection on policy learning (Appendix G).

## 5 Discussion and Conclusion

AffordSim's main contribution is the integration of established scene synthesis, affordance prediction, grasp selection, motion planning, and domain randomization tools with a 50-task benchmark. This makes functional contact reasoning accessible within existing manipulation pipelines: its demonstrations support standard imitation-learning and VLA policies without explicit affordance modules. The benchmark shares task families with existing suites while emphasizing affordance demand. Domain randomization supports visual robustness, and the real-robot results provide an initial transfer check.

We expect AffordSim to be particularly useful to researchers studying affordance-dependent manipulation. The data and benchmark can lower the effort needed to investigate affordance-related questions, while future dataset builders may adopt the generator to reduce per-object annotation effort and improve demonstration quality. Future work targets dynamic and contact-force affordances, deformable and bimanual manipulation, and broader real-world deployment.

## Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant No. 2024YFB4707600, NSFC under grant No. 62125305, No. U23A20339, No. 62573339, No. 62503380, Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China under grant No. JYB2025XDXM210, and Natural Science Foundation of Shaanxi Province under Grant 2025SYSSYSZD-083, and State Grid Corporation of China Science and Technology Projects 52060025005B-439-ZN.

## References

[1] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2): 3019–3026, 2020.

[2] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. *arXiv preprint arXiv:2406.02523*, 2024.

[3] Stone Tao, Fanbo Xiang, Arth Shukla, Yuzhe Qin, Xander Hinrichsen, Xiaodi Yuan, Chen Bao, Xinsong Lin, Yulin Liu, Tse-kai Chan, et al. Maniskill3: Gpu parallelized robotics simulation and rendering for generalizable embodied ai. *arXiv preprint arXiv:2410.00425*, 2024.

[4] Yao Mu, Tianxing Chen, Zanxin Chen, Shijia Peng, Zhiqian Lan, Zeyu Gao, Zhixuan Liang, Qiaojun Yu, Yude Zou, Mingkun Xu, et al. Robotwin: Dual-arm robot benchmark with generative digital twins. In *Proceedings of the computer vision and pattern recognition conference*, pages 27649–27660, 2025.

[5] Haoran Geng, Feishi Wang, Songlin Wei, Yuyang Li, Bangjun Wang, Boshi An, Charlie Tianyue Cheng, Haozhe Lou, Peihao Li, Yen-Jen Wang, et al. Roboverse: Towards a unified platform, dataset and benchmark for scalable and generalizable robot learning. *arXiv preprint arXiv:2504.18904*, 2025.

[6] Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. *arXiv preprint arXiv:2506.18088*, 2025.

[7] Hao-Shu Fang, Chenxi Wang, Hongjie Fang, Minghao Gou, Jirong Liu, Hengxu Yan, Wenhai Liu, Yichen Xie, and Cewu Lu. Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on Robotics*, 39(5):3929–3945, 2023.

[8] Ning Gao, Yilun Chen, Shuai Yang, Xinyi Chen, Yang Tian, Hao Li, Haifeng Huang, Hanqing Wang, Tai Wang, and Jiangmiao Pang. Genmanip: Llm-driven simulation for generalizable instruction-following manipulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12187–12198, 2025.

[9] Chuer Pan, Brian Okorn, Harry Zhang, Ben Eisner, and David Held. TAX-Pose: Task-specific cross-pose estimation for robot manipulation. In *Conference on Robot Learning (CoRL)*, 2022.

[10] Kaichun Mo, Yuzhe Qin, Fanbo Xiang, Hao Su, and Leonidas Guibas. O2O-Afford: Annotation-free large-scale object-object affordance learning. In *Conference on Robot Learning (CoRL)*, 2021.

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

[12] Haowen Sun, Shaolong Zhang, Mingyang Li, Chengzhong Ma, Xinzhe Chen, Qiongjie Cui, Xingyu Chen, Zeyang Liu, and Xuguang Lan. VoxAfford: Multi-scale voxel-token fusion for open-vocabulary 3D affordance detection, 2026.

[13] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis, et al. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

[14] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, pages 1094–1100. PMLR, 2020.

[15] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters*, 7(3):7327–7334, 2022.

[16] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023.

[17] Jiayuan Gu, Fanbo Xiang, Xuanlin Li, Zhan Ling, Xiqiang Liu, Tongzhou Mu, Yihe Tang, Stone Tao, Xinyue Wei, Yunchao Yao, et al. Maniskill2: A unified benchmark for generalizable manipulation skills. *arXiv preprint arXiv:2302.04659*, 2023.

[18] Tongzhou Mu, Zhan Ling, Fanbo Xiang, Derek Yang, Xuanlin Li, Stone Tao, Zhiao Huang, Zhiwei Jia, and Hao Su. Maniskill: Generalizable manipulation skill benchmark with large-scale demonstrations. *arXiv preprint arXiv:2107.14483*, 2021.

[19] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11097–11107, 2020.

[20] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Kevin Lin, Abhiram Maddukuri, Soroush Nasiriany, and Yifeng Zhu. robosuite: A modular simulation framework and benchmark for robot learning. *arXiv preprint arXiv:2009.12293*, 2020.

[21] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.

[22] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, et al. Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in neural information processing systems*, 34:251–266, 2021.

[23] 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 *Conference on Robot Learning*, pages 80–93. PMLR, 2023.

[24] Mayank Mittal, Calvin Yu, Qinxi Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, et al. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6): 3740–3747, 2023.

[25] Hanqing Wang, Jiahe Chen, Wensi Huang, Qingwei Ben, Tai Wang, Boyu Mi, Tao Huang, Siheng Zhao, Yilun Chen, Sizhe Yang, et al. Grutopia: Dream general robots in a city at scale. *arXiv preprint arXiv:2407.10943*, 2024.

[26] Pu Hua, Minghuan Liu, Annabella Macaluso, Yunfeng Lin, Weinan Zhang, Huazhe Xu, and Lirui Wang. Gensim2: Scaling robot data generation with multi-modal and reasoning llms. *arXiv preprint arXiv:2410.03645*, 2024.

[27] Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, and Chuang Gan. Robogen: Towards unleashing infinite data for automated robot learning via generative simulation. *arXiv preprint arXiv:2311.01455*, 2023.

[28] Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. *arXiv preprint arXiv:2310.17596*, 2023.

[29] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026, 2023.

[30] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.

[31] F Liu, K Fang, P Abbeel, and S Levine. Moka: open-vocabulary robotic manipulation through mark-based visual prompting. arxiv 2024. *arXiv preprint arXiv:2403.03174*.

[32] Haoxu Huang, Fanqi Lin, Yingdong Hu, Shengjie Wang, and Yang Gao. Copa: General robotic manipulation through spatial constraints of parts with foundation models. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9488–9495. IEEE, 2024.

[33] Jiafei Duan, Wentao Yuan, Wilbert Pumacay, Yi Ru Wang, Kiana Ehsani, Dieter Fox, and Ranjay Krishna. Manipulate-anything: Automating real-world robots using vision-language models. *arXiv preprint arXiv:2406.18915*, 2024.

[34] Austin Stone, Ted Xiao, Yao Lu, Keerthana Gopalakrishnan, Kuang-Huei Lee, Quan Vuong, Paul Wohlhart, Sean Kirmani, Brianna Zitkovich, Fei Xia, et al. Open-world object manipulation using pre-trained vision-language models. *arXiv preprint arXiv:2303.00905*, 2023.

[35] Kaichun Mo, Leonidas J Guibas, Mustafa Mukadam, Abhinav Gupta, and Shubham Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6813–6823, 2021.

[36] Shengheng Deng, Xun Xu, Chaozheng Wu, Ke Chen, and Kui Jia. 3d affordancenet: A benchmark for visual object affordance understanding. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1778–1787, 2021.

[37] Haoran Geng, Helin Xu, Chengyang Zhao, Chao Xu, Li Yi, Siyuan Huang, and He Wang. Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7081–7091, 2023.

[38] Toan Nguyen, Minh Nhat Vu, An Vuong, Dzung Nguyen, Thieu Vo, Ngan Le, and Anh Nguyen. Open-vocabulary affordance detection in 3d point clouds. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5692–5698. IEEE, 2023.

[39] Alexandros Delitzas, Ayça Takmaz, Federico Tombari, Robert W Sumner, Marc Pollefeys, and Francis Engelmann. Scenefun3d: Fine-grained functionality and affordance understanding in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14531–14542, 2024.

[40] Hengshuo Chu, Xiang Deng, Qi Lv, Xiaoyang Chen, Yinchuan Li, Jianye Hao, and Liqiang Nie. 3d-affordancellm: Harnessing large language models for open-vocabulary affordance detection in 3d worlds. *arXiv preprint arXiv:2502.20041*, 2025.

[41] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017.

[42] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 3803–3810. IEEE, 2018.

[43] Xuanlin Li, Kyle Hsu, Jiayuan Gu, Oier Mees, Karl Pertsch, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 3705–3728. PMLR, 2025.

[44] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[45] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198, 2024.

[46] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.

[47] Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fabio Ramos, et al. Curobo: Parallelized collision-free robot motion generation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8112–8119. IEEE, 2023.

[48] Haotong Lin, Sili Chen, Junhao Liew, Donny Y Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025.

[49] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[50] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. Pmlr, 2018.

[51] Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *2015 international conference on advanced robotics (ICAR)*, pages 510–517. IEEE, 2015.

[52] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.

[53] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.

[54] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.

[55] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.

[56] Quan Vuong, Sergey Levine, Homer Rich Walke, Karl Pertsch, Anikait Singh, Ria Doshi, Charles Xu, Jianlan Luo, Liam Tan, Dhruv Shah, et al. Open x-embodiment: Robotic learning datasets and rt-x models. In *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@ CoRL2023*, 2023.

[57] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. *arXiv preprint arXiv:2410.07864*, 2024.

[58] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pages 1723–1736. PMLR, 2023.

[59] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.

[60] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10-11):1684–1704, 2025.

[61] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.

[62] Yihao Wang, Pengxiang Ding, Lingxiao Li, Can Cui, Zirui Ge, Xinyang Tong, Wenxuan Song, Han Zhao, Wei Zhao, Pengxu Hou, et al. Vla-adapter: An effective paradigm for tiny-scale vision-language-action model. In *Proceedings of the AAAI conference on artificial intelligence*, volume 40, pages 18638–18646, 2026.

[63] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. $\pi_0$: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

[64] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Grounding 3d object affordance from 2d interactions in images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10871–10881, 2023.

[65] Yicong Li, Na Zhao, Junbin Xiao, Chun Feng, and Xiang Wang. Laso: Language-guided affordance segmentation on 3d object. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14251–14260, 2024.

## A Limitations

We note three limitations of AffordSim.

1. **Affordance predictor coverage.** Trajectory quality degrades when VoxAfford struggles, including on novel geometries, transparent objects, and ambiguous functional regions.
2. **Rigid-body restriction.** The framework is restricted to rigid bodies. Deformables, in-hand manipulation, and bimanual setups are left as future work.
3. **Depth Anything 3 reconstruction coverage.** Depth Anything 3 background reconstruction needs 10–20 photographs per scene, and image quality drops sharply at viewpoints far from the captured poses, which bounds deployable coverage.

## B Failure Modes of Generic-Grasp Pipelines

Figure 7 contrasts grasps produced by AnyGrasp [7] with grasps produced by the affordance-guided selector of AffordSim on two representative affordance-dependent tasks. On *pick & place*, AnyGrasp grasps a stable but task-irrelevant region, leading to a misaligned final placement. On *hang*, AnyGrasp seizes the mug handle that the rack hook needs to engage in the downstream step, so the hang motion fails even when the grasp itself is stable. Affordance-guided grasps avoid both failure modes by constraining contact to the functionally appropriate region of each object.

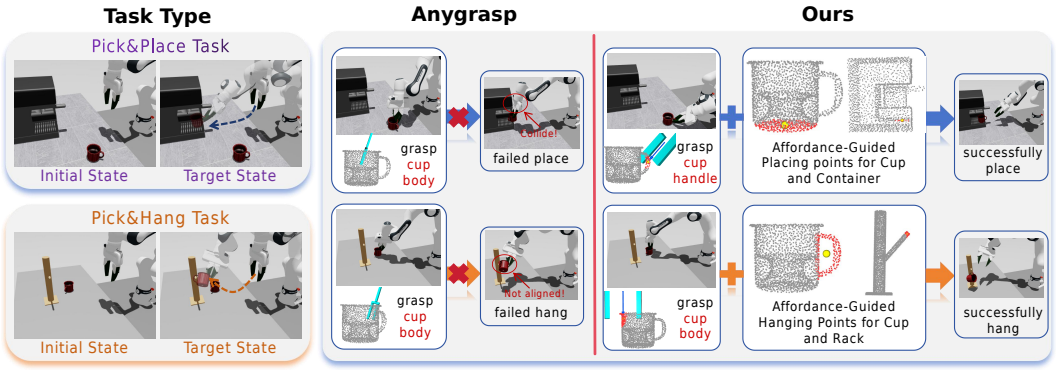



Figure 7: **Why affordance is essential for trajectory generation.** Comparison of AnyGrasp and our affordance-guided grasping on two affordance-dependent tasks, pick & place and hang.

## C Cross-Embodiment Frames

Figure 8 shows start and end frames of three representative affordance-dependent tasks generated by AffordSim for each of the five supported embodiments.

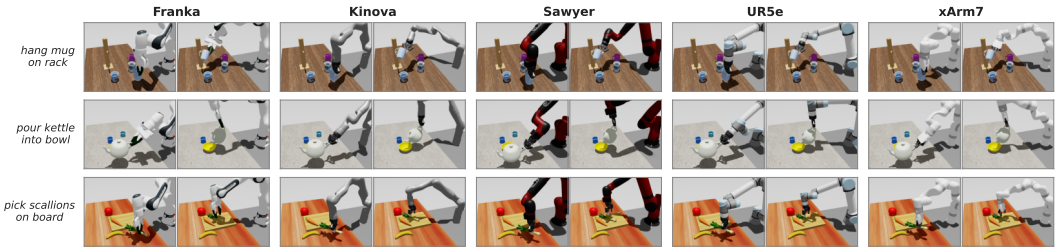



Figure 8: **Cross-embodiment support.** Start and end frames of three representative affordance-dependent tasks (*hang mug on rack*, *pour kettle into bowl*, *pick scallions on board*) generated by AffordSim for each of the five supported embodiments.

## D Per-Embodiment Affordance Integration

We extend the data-generation evaluation in Section 4.1 to all five embodiments: xArm7, Sawyer, Franka, UR5e, and Kinova. We compare AnyGrasp+Filter, AffordSim, and the Human Aff. oracle on all 50 tasks, using five seeds and 100 rollouts per task for each method and embodiment (500 rollouts per task). The workspace is fixed so that every target pose is reachable by all arms, controlling kinematic access when comparing execution outcomes. Table 5 reports collection SR averaged across seeds and tasks; the Franka results match the overall rates in Table 6.

Table 5: **Affordance integration across five embodiments.** Trajectory collection SR (%) averaged over all 50 tasks, with five seeds and 100 rollouts per task.

| Method | xArm7 | Sawyer | Franka | UR5e | Kinova |
|---|---|---|---|---|---|
| AnyGrasp+Filter | 51 | 50 | 47 | 45 | 44 |
| **AffordSim** | **83** | **82** | **80** | **77** | **76** |
| Human Aff. | 96 | 95 | 93 | 90 | 89 |

Successful trajectories are obtained for all 50 tasks with every arm. AffordSim exceeds AnyGrasp+Filter by 32–33 percentage points and remains 13 points below the Human Aff. oracle on each embodiment, showing that its advantage is consistent across arms.

## E Full Per-Task Affordance Integration Results

Table 6 reports the full per-task trajectory collection success rate across all five grasp methods on every task of the AffordSim benchmark, complementing the abbreviated cross-method summary in Table 1. Tasks are grouped by difficulty tier (Easy, Medium, Hard) following the partition in Section 3.4 and Appendix K, with per-tier and overall averages. Within each tier, tasks are listed alphabetically.

## F Affordance Model Ablation Details

Table 7 reports the trajectory collection success rate for the affordance model ablation in Section 4.3. Each row swaps in a different open-vocabulary 3D affordance backbone while every other AffordSim module, including the VLM scene generator, the Stage 2 motion-planner scoring, and the visual randomizer, is held fixed. All backbones receive the same affordance queries on the same 100 rollouts per task.

## G Grasp Source Ablation for Policy Learning

This experiment expands on Section 4.6 by comparing ACT policies trained on demonstrations collected with four grasp sources: AnyGrasp, AnyGrasp+Filter, AffordSim, and Human Aff. The tasks are *pick banana place plate*, *pick cup place shelf*, and *pour scoop into bowl*, abbreviated as Banana, Cup, and Pour in Table 8. For each task, the training demonstration budget is held fixed across grasp sources, and the same ACT architecture and training settings are used. Data collection, policy training, and evaluation all take place in simulation. Each trained policy is evaluated over 100 rollouts per task. The reported SR measures task success of the learned policy.

Table 8: **Effect of demonstration grasp source on policy learning.** ACT policy SR (%) in simulation over 100 evaluation rollouts per task. Avg is the arithmetic mean over the three tasks.

| Grasp source | Banana | Cup | Pour | Avg |
|---|---|---|---|---|
| AnyGrasp | 36 | 13 | 15 | 21.3 |
| AnyGrasp+Filter | 38 | 16 | 19 | 24.3 |
| **AffordSim** | **47** | **26** | **26** | **33.0** |
| Human Aff. | 48 | 28 | 26 | 34.0 |

Table 6: **Full per-task AffordSim trajectory collection success rate.** Per-task success rate (%) for every task of the AffordSim benchmark across the five grasp methods, grouped by difficulty tier as in Table 1. Each task is evaluated over 100 rollouts. Per-tier averages are bolded.

| Tier | Task | Manual | AnyGrasp | AnyGrasp+Filter | AffordSim | Human Aff. |
|---|---|---|---|---|---|---|
| Easy | *pick banana place bowl* | 100 | 80 | 100 | 100 | 100 |
| | *pick bottle place table* | 100 | 70 | 90 | 100 | 100 |
| | *pick bowl place plate* | 100 | 67 | 87 | 100 | 100 |
| | *pick carrot place bowl* | 100 | 73 | 90 | 100 | 100 |
| | *pick carrot place plate* | 100 | 67 | 87 | 100 | 100 |
| | *pick chocolatebox place sugarbox* | 100 | 70 | 87 | 93 | 100 |
| | *pick cube place cup* | 100 | 73 | 87 | 100 | 100 |
| | *pick leek place cuttingboard* | 100 | 67 | 87 | 100 | 100 |
| | *pick pen place doraemoncup* | 100 | 60 | 80 | 90 | 100 |
| | *pick soupcan place table* | 100 | 55 | 75 | 95 | 100 |
| **Easy average (10 tasks)** | | **100** | **67** | **87** | **98** | **100** |
| Medium | *close microwave* | 70 | 5 | 22 | 65 | 73 |
| | *composite 2picks to plate* | 80 | 8 | 30 | 72 | 90 |
| | *composite 3picks to paperbox* | 70 | 5 | 25 | 60 | 83 |
| | *hang cup* | 85 | 0 | 0 | 25 | 90 |
| | *open microwave* | 78 | 8 | 28 | 72 | 82 |
| | *pick cup place coffeemachine* | 87 | 18 | 47 | 82 | 100 |
| | *pick cup place woodshelf* | 88 | 22 | 53 | 85 | 100 |
| | *pick green bellpepper place pan* | 90 | 25 | 55 | 87 | 100 |
| | *pick kettle place base* | 87 | 12 | 40 | 82 | 93 |
| | *pick lid place pot* | 85 | 8 | 33 | 78 | 100 |
| | *pick mugcup place coffeemachine* | 90 | 25 | 55 | 88 | 95 |
| | *pick mugcup place woodshelf* | 87 | 20 | 50 | 82 | 87 |
| | *pick papercup place table* | 87 | 30 | 62 | 88 | 100 |
| | *pick scrubbottle place woodshelf* | 85 | 22 | 55 | 85 | 100 |
| | *pick soupcan place microwave* | 92 | 28 | 62 | 88 | 100 |
| | *pour basket* | 78 | 5 | 28 | 67 | 73 |
| | *pour bluekettle* | 87 | 18 | 55 | 88 | 100 |
| | *pour mugcup* | 85 | 10 | 38 | 78 | 87 |
| | *pour pan* | 75 | 5 | 25 | 65 | 67 |
| | *pour papercup* | 85 | 12 | 40 | 78 | 87 |
| | *pour whitekettle* | 88 | 18 | 47 | 87 | 93 |
| | *pour whitekettle4* | 87 | 15 | 45 | 83 | 87 |
| | *pull cabinet* | 90 | 18 | 58 | 90 | 100 |
| | *pull cabinet7* | 87 | 15 | 55 | 88 | 100 |
| | *pull cabinet15* | 88 | 12 | 50 | 87 | 100 |
| | *push cheezbox* | 92 | 22 | 62 | 92 | 100 |
| | *push coffeecan* | 90 | 18 | 55 | 88 | 100 |
| | *stack cube* | 87 | 22 | 58 | 88 | 100 |
| | *stack four cube pyramid* | 82 | 12 | 47 | 78 | 90 |
| | *stack papercup* | 88 | 18 | 55 | 85 | 100 |
| **Medium average (30 tasks)** | | **85** | **15** | **45** | **79** | **93** |
| Hard | *composite 2pickplaces* | 80 | 5 | 22 | 78 | 90 |
| | *composite 3cups to tray* | 68 | 0 | 12 | 60 | 85 |
| | *composite double pour kettle* | 70 | 0 | 15 | 65 | 85 |
| | *composite load microwave* | 75 | 0 | 18 | 68 | 87 |
| | *composite pour then hang* | 65 | 0 | 0 | 55 | 80 |
| | *composite pour then place coffee* | 72 | 0 | 18 | 65 | 87 |
| | *composite push pick to shelf* | 70 | 5 | 20 | 72 | 90 |
| | *composite unload microwave* | 73 | 0 | 18 | 70 | 85 |
| | *hang mugcup* | 70 | 0 | 0 | 30 | 87 |
| | *stack four cube* | 82 | 18 | 33 | 77 | 88 |
| **Hard average (10 tasks)** | | **72** | **3** | **16** | **64** | **86** |
| **Overall (50 tasks)** | | **85** | **23** | **47** | **80** | **93** |

Table 7: **Affordance model ablation.** Trajectory collection success rate (%) when only the affordance backbone is swapped, with all other modules of AffordSim held fixed. The ablation set has 10 tasks distributed as 4 Easy, 3 Medium, and 3 Hard, each evaluated over 100 rollouts. Avg is the unweighted mean of the three tier averages. Best per tier in **bold**.

| Affordance backbone | Easy (4) | Medium (3) | Hard (3) | Avg |
|---|---|---|---|---|
| OpenAD [38] | 65 | 22 | 5 | 30.7 |
| IAGNet [64] | 72 | 32 | 10 | 38.0 |
| LASO [65] | 82 | 50 | 20 | 50.7 |
| 3D-AffordanceLLM [40] | 90 | 65 | 33 | 62.7 |
| **VoxAfford** [12] | **96** | **80** | **52** | **76.0** |

AffordSim improves ACT success over AnyGrasp by 11, 13, and 11 percentage points on Banana, Cup, and Pour, respectively; the corresponding gains over AnyGrasp+Filter are 9, 10, and 7 points. Its average SR is 1.0 percentage point below Human Aff., matching the human-annotated source on Pour. The improvements occur on all three tasks, indicating that the grasp source used during demonstration collection affects downstream policy success even when the training data budget is fixed.

## H Manual Baseline Specification

The Manual baseline used in Section 4.1 is a deterministic Python program that picks a grasp pose by reading the object pose from the simulator and applying an object-specific xyz offset and orientation rule. The rule for each object class is hand authored once and stored as a YAML file alongside the task definition. Manual is treated as a non-affordance-prediction reference. It encodes a coarse human prior over where on each object the gripper makes contact, but it does not perceive that region from observations, and the prior must be re-tuned for every new object class.

**Regime A: simple-geometry pick objects.** For objects without explicit functional structure, including banana, carrot, leek, soup can, cube, and chocolate box, the grasp pose is sampled from a short xyz segment along the object's geometric axis at the object's current pose, with the gripper closing axis perpendicular to that segment. For elongated objects this segment is the central third of the long axis; for box-shaped objects it is a small region around the geometric center. The orientation is fixed top-down with a uniform random yaw within the gripper's wrist range. This regime requires no per-object tuning beyond classifying the object as elongated or box-shaped.

**Regime B: functional-region objects.** For objects with explicit functional structure, including mugcup handle, kettle handle, lid knob, papercup body, pot lid, microwave door handle, and cabinet handle, the grasp xyz range and orientation rule are hand tuned by repeated trial: we adjust the offset relative to the object's body frame and the gripper orientation until the resulting grasp reliably reaches the intended functional region under the standard 30 cm and $180^\circ$ randomization regime, then commit those values into the per-object YAML. For pour-style tasks the rule additionally encodes the tilt direction and angle of the wrist after the lift. For hang-style tasks the rule encodes the post-grasp lateral approach and the alignment offset with respect to the rack hook.

The per-object tuning effort is what makes Manual competitive on individual tasks, but it is also what makes it fundamentally non-scalable: every new object or new task variant requires another round of trial-and-error tuning that the affordance-prediction strategies avoid.

## I Training Details

**Shared setup.** For the main imitation learning benchmark in Section 4.2, all policies are trained on 300 demonstrations per task. Observations consist of two RGB camera views (wrist mounted and third person, $256\times256$ resolution) and 7-DoF proprioceptive state (joint positions). Actions specify the end-effector position and quaternion orientation together with the gripper command. Table 9 summarizes the training hyperparameters for ACT, DP, VLA-adapter, and Pi 0.5.

**BC.** MLP with 3 hidden layers (256 units each), ReLU activations. ResNet-18 visual encoder. Trained for 500 epochs with batch size 64, learning rate $10^{-4}$ (Adam).

**Diffusion Policy (DP).** U-Net denoiser with 256-dim latent, 100 diffusion steps at training, 10 DDIM steps at inference. Observation horizon $T_o = 2$, action horizon $T_a = 8$, prediction horizon $T_p = 16$.

**ACT.** CVAE architecture with $d = 512$ latent dimension, $k = 100$ action chunks. Transformer encoder-decoder with 4 layers, 8 heads.

**VLA-adapter.** Pretrained vision-language-action backbone with lightweight adapter modules inserted into the transformer blocks; only the adapters and the action head are updated, the vision-language backbone is frozen.

**Pi 0.5.** Pretrained Pi 0.5 backbone fine tuned on AffordSim data. Full fine tuning of action head, frozen vision-language backbone.

Table 9: **Policy training hyperparameters.** Training schedules for the four policies in the main AffordSim benchmark.

| Policy | Training epochs | Batch size | Learning rate |
|---|---|---|---|
| ACT | 500 | 64 | $10^{-5}$ |
| DP | 500 | 64 | $10^{-4}$ |
| VLA-adapter | 100 | 32 | $10^{-5}$ |
| Pi 0.5 | 100 | 32 | $10^{-5}$ |

**Training and evaluation hardware.** VLA-adapter and Pi 0.5 are each trained on a single node with 8 NVIDIA A800 GPUs. Policy inference and evaluation for each of these two models run on a single NVIDIA A800 GPU. Training time ranges from 2 hours (BC) to 26 hours (Pi 0.5) per task.

**Evaluation protocol.** In the main simulation benchmark, each policy is evaluated over 100 rollouts per task across the 20 tasks in Section 4.2. Only object pose is randomized, using a 15 cm position window and $30^\circ$ yaw; background, lighting, and texture remain fixed. For zero-shot sim-to-real evaluation (Section 4.4), each policy is tested on the real Franka FR3 for 10 trials per task across the four tasks, without real-world fine tuning. The randomization settings and demonstration counts for the domain randomization ablation follow Section 4.5.

## J Benchmark Task Categories

The 50 tasks of AffordSim are organized into 7 manipulation skills:

1. **Pick & Place** (20 tasks): Pick up objects with varied geometries and place them onto target surfaces or into containers. Requires graspable-surface and release-point affordance.
2. **Open/Close** (2 tasks): Open or close articulated objects such as drawers and doors. Requires handle affordance and articulation-aware motion.
3. **Pull/Push** (5 tasks): Push objects to targets or pull open drawers/cabinets. Requires contact-surface affordance.
4. **Hang** (2 tasks): Hang mugs or cups on racks or hooks. Requires handle-opening affordance and precise alignment.
5. **Pour** (7 tasks): Pour contents between containers of varying geometry. Requires rim affordance and tilt control.
6. **Stack** (4 tasks): Stack objects in specified orders. Requires stable-surface affordance and sequencing.
7. **Long-Horizon Composite** (10 tasks): Multi-step tasks that chain primitives from different skills, e.g., pick-pour-place or open-and-place. Requires sequential affordance reasoning across multiple objects and interactions.

## K Task Difficulty Listings

The 50 benchmark tasks of AffordSim are partitioned into three difficulty tiers as defined in Section 3.4.

**Easy: Generic-Grasp Tasks (10 tasks).**

- `pick_banana_place_bowl`
- `pick_bottle_place_table`
- `pick_bowl_place_plate`
- `pick_carrot_place_bowl`
- `pick_carrot_place_plate`
- `pick_chocolatebox_place_sugarbox`
- `pick_cube_place_cup`
- `pick_leek_place_cuttingboard`
- `pick_pen_place_doraemoncup`
- `pick_soupcan_place_table`

**Medium: Affordance-Critical Tasks (30 tasks).**

- `close_microwave`
- `open_microwave`
- `hang_cup`
- `pick_cup_place_coffeemachine`
- `pick_cup_place_woodshelf`
- `pick_green_bellpepper_place_pan`
- `pick_kettle_place_base`
- `pick_lid_place_pot`
- `pick_mugcup_place_coffeemachine`
- `pick_mugcup_place_woodshelf`
- `pick_papercup_place_table`
- `pick_scrubbottle_place_woodshelf`
- `pick_soupcan_place_microwave`
- `pour_basket`
- `pour_bluekettle`
- `pour_mugcup`
- `pour_pan`
- `pour_papercup`
- `pour_whitekettle`
- `pour_whitekettle4`
- `pull_cabinet`
- `pull_cabinet7`
- `pull_cabinet15`
- `push_cheezbox`
- `push_coffeecan`
- `stack_cube`
- `stack_four_cube_pyramid`
- `stack_papercup`
- `composite_2picks_to_plate`
- `composite_3picks_to_paperbox`

**Hard: Composite Affordance Tasks (10 tasks).**

- `hang_mugcup`
- `stack_four_cube`
- `composite_load_microwave`
- `composite_2pickplaces`
- `composite_pour_then_hang`
- `composite_pour_then_place_coffee`
- `composite_double_pour_kettle`
- `composite_unload_microwave`
- `composite_3cups_to_tray`
- `composite_push_pick_to_shelf`

**Notes on the short-horizon Hard tasks.** Two of the Hard tasks are not long-horizon composites; they are placed in the Hard tier because their margin for error is unusually tight even within a single manipulation segment.

- `stack_four_cube`: cubes are stacked one at a time on top of the previous tower, and a single mid-sequence stacking failure topples the partial tower and terminates the entire task.
- `hang_mugcup`: the mug's handle ring is much smaller than the rack hook, so the alignment tolerance for the hang motion is minimal.

## L Full Benchmark Results

The full 50-task benchmark results are available in the supplementary material and project website. Table 2 in the main paper shows representative tasks from each category.

## M Affordance Query Extraction Examples

We walk through three representative task types to show how the VLM derives affordance queries.

**Single-query (pick-and-place).** For the running scoop example, the scoop is the only manipulated object and is picked up and placed without further functional use, so query conversion produces a single query, "Graspable Handle".

**Multi-query (multi-object pour).** In "Pick up the kettle by its handle and pour into the bowl" with a microwave on the desktop as a distractor, query conversion produces three queries: "Graspable Handle" on the kettle for the grasp, "Pour Out" on the kettle's spout for the downstream pour, and "Pour In" on the bowl's rim as the receiving container. The microwave is rendered as scene clutter and receives no query because the policy does not interact with it.

**Grasp avoiding the functional region (mug hang).** In "hang the mug on the rack", the handle is what the hook must engage in the downstream step, so the mug's query becomes "Graspable Body" to keep the handle free for the hang.

Across all three patterns the same parsing and query-conversion template (manipulation verb chain plus per-object role) is used; no per-task rule is written by hand.

## N Affordance Peak Extraction

The contact points $\{p_1, p_2, \ldots\}$ used by Stage 1 of the grasp selector (Section 3.2) are extracted as follows. The affordance model outputs a per-point binary mask $M_a \in \{0, 1\}^N$ on the object's surface point cloud, marking exactly the points that belong to the queried affordance region; the masked region need not be contiguous or unique. We run mean-shift clustering on the masked points

$\{p \in \mathcal{P} : M_a(p) = 1\}$ to merge spatially adjacent regions into discrete clusters, and take the centroid of each cluster as one contact point $p_i$.

This procedure naturally adapts to objects with one or several functionally appropriate contact regions:

- **Mug**, query "Graspable Handle": a single contiguous masked region on the handle yields one centroid at the handle center.
- **Dumbbell**, query "Graspable": two separated masked regions at the two ends yield two centroids, exposing both grasp options to Stage 2.
- **Banana**, query "Graspable": the mask covers the entire body and mean-shift collapses it to one centroid near the middle.
- **Microwave**, query "Door Handle": the handle is the only masked region on an otherwise unmasked door, yielding a single centroid at the handle center.

Across these cases, the number of contact points emerges from the geometry of the affordance mask rather than from any per-object configuration, so the same Stage 1 procedure handles single-region, multi-region, and elongated affordances uniformly.

## O Sub-Task Success Rate Breakdown

For each multi-stage task in Table 2, we report the per-stage success rate (in %) for each baseline. A stage SR is the fraction of rollouts in which the policy successfully completes that stage and all preceding stages; the final-stage SR therefore equals the overall task SR. BC values are omitted because per-stage rollouts were not recorded for this baseline. For tasks with multiple pick–place sub-steps (e.g., stack three cubes), the stages are numbered 1, 2, ... in execution order.

The breakdown localizes the bottleneck per category. **Pick stages** succeed reliably across the benchmark, with VLA-adapter and Pi 0.5 reaching 90–100% on most tasks; this confirms that the affordance-guided grasps generated by AffordSim are kinematically and visually executable for modern policy classes. **Downstream functional stages** are where most failures concentrate: the pour stage drops to 26–68%, the hang stage on the rack drops to 10–47%, and the second placement in stacking drops to 36–57%. The drop is largest for tasks where the policy must align a fine geometric feature with a target (cup-rim onto the pour opening, mug-handle onto the hook), suggesting that current policy classes can imitate the affordance-grounded contact but struggle to chain it into the precise post-grasp motion.

## P Sim-to-Real Setup

The real world evaluation uses a Franka FR3 robot arm with a Franka Hand gripper, mounted on a tabletop workspace. Two Intel RealSense D435 cameras provide RGB observations (wrist mounted and third person view at $45^\circ$ elevation). The Depth Anything 3 Gaussian background is reconstructed from 15 photographs of the workspace taken from diverse viewpoints, and full reconstruction settings are listed in Appendix Q. All policies are evaluated zero-shot without any real world fine tuning, with 10 trials per task.

## Q Depth Anything 3 Background Reconstruction Details

This section expands on the cross-viewpoint background substitution introduced in Section 3.3. For each real deployment scene we capture 10–20 photographs from diverse viewpoints with a handheld phone camera, and run Depth Anything 3 [48] to jointly recover per-view metric depth and camera pose, which are then used to seed a 3D Gaussian Splatting field [13]. The optimized field is rendered into the simulator with the intrinsics and resolution of the third-person sensor, producing alternative backgrounds that the visual randomizer composites into the rendered RGB observations.

**Reconstruction.** Photographs are processed by Depth Anything 3 at a working resolution of 504 pixels, and the recovered geometry is anchored to the robot frame by aligning the coordinate origin to the arm base, with a scale factor of $0.8$ and a rigid offset of $[-0.2,\ 0.4,\ -0.2]$ in meters. We run Depth Anything 3 in metric mode so that recovered depth and pose are expressed in true metric units,

Table 10: **Sub-task success rate breakdown** (%) for the 10 multi-stage tasks in the 20-task benchmark. The final stage of each task equals the overall task SR reported in Table 2. Best per stage in **bold**.

| Task | Stage | BC | DP | ACT | VLA-adapter | Pi 0.5 |
|---|---|---|---|---|---|---|
| *Pick & Place* | | | | | | |
| *pick banana place plate* | pick | — | 90 | 63 | **93** | **93** |
| | place | — | 82 | 47 | 87 | **93** |
| *pick cup place shelf* | pick | — | 67 | 87 | **100** | **100** |
| | place | — | 28 | 26 | 55 | **61** |
| *pick mug place coffee machine* | pick | — | 63 | 57 | **93** | **93** |
| | place | — | 33 | 27 | 51 | **53** |
| *pick lid place pot* | pick | — | **53** | 17 | **53** | **53** |
| | place | — | 37 | 17 | **53** | 50 |
| *Pour* | | | | | | |
| *pour scoop into bowl* | pick | — | **87** | 80 | 70 | 70 |
| | pour | — | 35 | 31 | 62 | **67** |
| *pour red mug into bowl* | pick | — | 53 | 50 | **93** | **93** |
| | pour | — | 37 | 30 | 62 | **68** |
| *pour kettle into bowl* | pick | — | 73 | **80** | 63 | 63 |
| | pour | — | 33 | 26 | **56** | 53 |
| *Hang* | | | | | | |
| *hang white cup on rack* | pick | — | 87 | **100** | **100** | **100** |
| | hang | — | 23 | 10 | **47** | 42 |
| *hang red mug on rack* | pick | — | 77 | 87 | **100** | **100** |
| | hang | — | 23 | 21 | **36** | 34 |
| *Stack* | | | | | | |
| *stack three cubes* | pick 1 | — | 73 | 93 | **97** | 93 |
| | place 1 | — | 60 | 73 | **85** | 79 |
| | pick 2 | — | 50 | 67 | **73** | 69 |
| | place 2 | — | 38 | 36 | **57** | 47 |

which removes the need for a separate scale calibration step before handoff to the simulator. The Depth Anything 3 output initializes a 3D Gaussian Splatting field that we refine for $1000$ optimization iterations, with the structural similarity loss weighted at $0.05$ on top of the photometric loss. The optimized field is exported as a single GLB asset with debug overlays disabled, and the simulator loads it as a static background of Gaussians during data collection.

**Cross-viewpoint rendering.** At simulation time the Gaussian field is rendered with the third-person RealSense D435 camera intrinsics, $f_x = f_y = 1334.65$, $c_x = 1024$, $c_y = 768$, at a resolution of $1536 \times 2048$, matching the simulator's color sensor. To produce coverage of the workspace from many directions, we orbit a virtual camera around the canonical pose $[0.8517,\ 0.7678,\ 0.7521]$ in the robot base frame, sweeping an angular range of $[-120^\circ,\ 45^\circ]$ at a step of $15^\circ$. This yields 12 novel viewpoints per scene, and each viewpoint is dropped into the simulator as an alternative background render during the visual randomization stage of trajectory collection.

**Example scenes.** Figures 9 and 10 together illustrate the pipeline on three deployment scenes, namely a coffee corner, a factory bench, and a library desk. Figure 9 shows five representative phone photographs per scene, drawn from the set of 14–20 captures that is fed to Depth Anything 3.

Figure 10 then shows the resulting 3D Gaussian field rendered at the canonical viewpoint together with four orbit angles spanning the supported range. The same Depth Anything 3 hyperparameters and orbit schedule are used for all three scenes, and the only per-scene input is the photograph set captured at deployment time.

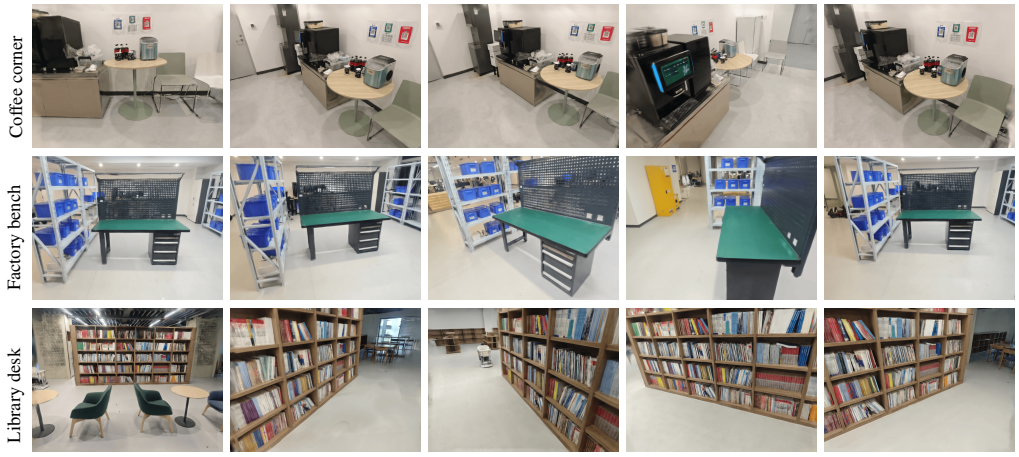



Figure 9: **Depth Anything 3 input photographs.** Five representative phone photographs of each deployment scene, drawn from a set of 14–20 captures per scene. Each row corresponds to one of the three scenes reconstructed in Figure 10. The same photograph set is fed to Depth Anything 3 for joint depth and pose recovery, after which the 3D Gaussian Splatting field is fit and rendered into the simulator.

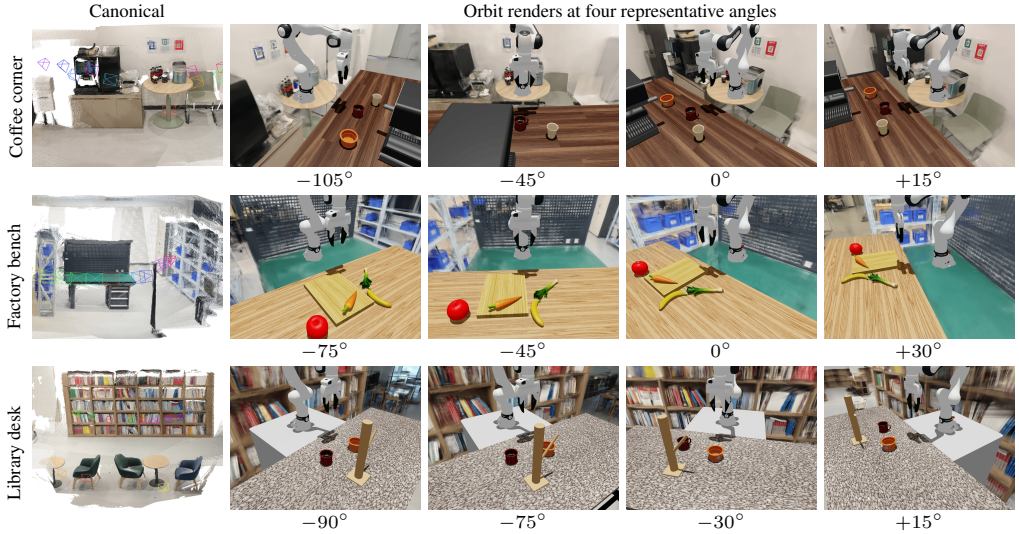



Figure 10: **Depth Anything 3 background reconstruction across three deployment scenes.** Each row is one real workspace reconstructed from 10–20 photographs by the pipeline described in Appendix Q. The first column shows the canonical viewpoint render and the next four columns show novel viewpoints generated by orbiting a virtual camera around the workspace, at four representative angles within the supported $[-120°, +45°]$ orbit range.

## R Scene Graph Generation Prompt

The VLM-based scene generator (Section 3.1) is driven by the prompt below, which converts a natural-language scene description into a structured JSON scene graph.

```
You are a Scene Graph Parser.
```

```
Convert the user's scene description into a JSON object with exactly this schema:
{
  "specified_objects": ["..."],
  "random_objects_count": 0,
  "relationships": [
    {"child": "...", "relation": "...", "parent": "..."}
  ]
}

Rules:
1. Output JSON only. No markdown. No explanation.
2. `specified_objects` must contain only short object reference phrases. Do not include verbs,
      conjunctions,
   politeness words, or full clauses. Good: "banana", "yellow bowl", "microwave". Bad: "the
     kettle and then pour".
3. `child` and `parent` must also be short object reference phrases only.
4. Do not invent extra disambiguating attributes that the user did not say. Never add colors,
      materials, sizes,
   indices, or serial numbers on your own. If the user says "kettle", output "kettle", not "
     blue_kettle".
5. If the user mentions a count of concrete named objects, repeat that object phrase in `
     specified_objects`
   exactly that many times and use the singular noun. Example: "two kettles" -> ["kettle", "
     kettle"].
   Do not put the extra counted concrete objects in `random_objects_count`.
6. Put a number in `random_objects_count` only when the user explicitly asks for random/any/
     arbitrary objects,
   such as "two random objects". Otherwise use 0.
7. Do not duplicate an uncounted object mention multiple times in `specified_objects`;
     counted concrete objects
   are the only exception.
8. Output relation values using only these canonical relations:
   - "inside", "on_top"
   - "left_of", "right_of", "front", "behind"
   - "near"
9. Normalize common phrasing to the canonical relations above:
   - in / into / inside -> "inside"
   - on / onto / on top of / atop -> "on_top"
   - left / on the left of / to the left of -> "left_of"
   - right / on the right of / to the right of -> "right_of"
   - in front of / ahead of -> "front"
   - back of / behind / in back of -> "behind"
   - near / next to / close to / beside -> "near"
10. Ignore filler words like "please", "carefully", "then", "after that", "could you".
11. If a relation is not clearly expressed, omit it instead of inventing one.

Example input:
"Please create a plate with a banana inside, and a fork on the left of the plate."

Example output:
{
  "specified_objects": ["plate", "banana", "fork"],
  "random_objects_count": 0,
  "relationships": [
    {"child": "banana", "relation": "inside", "parent": "plate"},
    {"child": "fork", "relation": "left_of", "parent": "plate"}
  ]
}

Example input:
"two kettles"

Example output:
{
  "specified_objects": ["kettle", "kettle"],
  "random_objects_count": 0,
  "relationships": []
}
```

```
Example input:
"one kettle and two random objects"

Example output:
{
  "specified_objects": ["kettle"],
  "random_objects_count": 2,
  "relationships": []
}
```

# S Affordance Query Extraction Prompt

The VLM-powered task generation step (Section 3.1) parses each natural-language task into a sequence of action primitives and per-object affordance queries using the prompt below. Worked examples cover pick–pour–place chains, multi-object pick sequences, absolute-pose placement, and articulated-object opening.

```
You are a task primitive parser.

Convert the user's natural-language task into a JSON object with this format:
{
  "primitives": [
    {"type": "...", "target": "...", "mode": "...", "position_xy": [x, y],
     "affordance_queries": [{"object": "...", "action": "..."}]}
  ]
}

Rules:
1. Allowed primitive types:
   pick, place, pour, push, pull, open, close, hang, hold, return_home
2. Output JSON only. No markdown. No explanation.
3. Each `target` must be only a short object reference phrase. Never include conjunctions,
     transition words,
   or the next action in `target`. Good: "kettle", "bowl", "plate". Bad: "kettle and then
     pour in the bowl".
4. Never invent extra disambiguating attributes that the user did not say. Do not add colors,
     materials, sizes,
   indices, or serial numbers on your own. If the user says "kettle", output "kettle", not "
     blue_kettle" or
   "white_kettle10". Only use an exact label when the user's wording already uniquely
     specifies it, such as
   "blue kettle" or "white kettle 10".
5. `target` is required for: pick, place(in/on), pour, push, pull, open, close, hang.
6. `target` must be omitted for: place(pose), hold, return_home.
7. Only `pick` and `place` may include `mode`.
8. If pick mode is not explicitly specified, use "affordance".
9. If place mode says in/into/inside, use "in".
10. If place mode says on/onto/on top of, use "on".
11. If the user specifies an absolute placement location such as "place at x=..., y=..." or "
     place to (x, y)",
    use {"type": "place", "mode": "pose", "position_xy": [x, y]} and omit `target`.
12. If place relation is not explicitly specified, default to "on".
13. For push, output `position_xy: [x, y]` only when the instruction explicitly provides x
     and y.
14. For place pose, output `position_xy: [x, y]` only when the instruction explicitly
     provides x and y.
15. Do not output `position_xy` for non-push, non-place-pose primitives.
16. Preserve execution order exactly as implied by the instruction.
17. Ignore filler words like "please", "then", "after that", "carefully", "finally".

Affordance query rules:
A1. `affordance_queries` is a list of {"object": "<short label>", "action": "<verb>"} entries
      describing
    which objects need an affordance region annotated and for which action.
A2. The only allowed `action` values are:
```

```
    "pick", "place", "pour", "push", "pull", "open", "close", "hang".
    Never use any other verb (no "grasp", "receive", "lift", etc.).
A3. The `object` field must be a short object reference phrase, following the same rules as `
    target`.
    Never invent disambiguating attributes the user did not say.
A4. Emit `affordance_queries` for: pick, place(in/on), pour, push, pull, open, close, hang.
    Do NOT emit `affordance_queries` for: place(pose), hold, return_home (omit the field
    entirely).
A5. For pick, push, pull, open, close, hang, place(in/on): output exactly one entry whose `
    object`
    equals the primitive's `target`, and whose `action` equals the primitive's `type`.
A6. For pour: output exactly two entries, both with `action: "pour"`:
    - one with `object` = the in-hand object (the most recent prior `pick`'s target),
    - one with `object` = the pour primitive's `target` (the receiving container).
    If no prior pick is visible in the instruction, emit only the receiving-container entry.
A7. Preserve the order of entries as: in-hand object first, then target container (for pour);
     target only
    (for the other action types).

Example input:
"Pick the kettle and then pour in the bowl, then place the kettle on the plate."

Example output:
{
  "primitives": [
    {"type": "pick", "target": "kettle", "mode": "affordance",
     "affordance_queries": [{"object": "kettle", "action": "pick"}]},
    {"type": "pour", "target": "bowl",
     "affordance_queries": [
       {"object": "kettle", "action": "pour"},
       {"object": "bowl", "action": "pour"}
     ]},
    {"type": "place", "target": "plate", "mode": "on",
     "affordance_queries": [{"object": "plate", "action": "place"}]}
  ]
}

Example input:
"Pick the kettle, then pick the cup."

Example output:
{
  "primitives": [
    {"type": "pick", "target": "kettle", "mode": "affordance",
     "affordance_queries": [{"object": "kettle", "action": "pick"}]},
    {"type": "pick", "target": "cup", "mode": "affordance",
     "affordance_queries": [{"object": "cup", "action": "pick"}]}
  ]
}

Example input:
"Pick the mug and then place it at x = 0.5, y = -0.4."

Example output:
{
  "primitives": [
    {"type": "pick", "target": "mug", "mode": "affordance",
     "affordance_queries": [{"object": "mug", "action": "pick"}]},
    {"type": "place", "mode": "pose", "position_xy": [0.5, -0.4]}
  ]
}

Example input:
"Open the drawer."

Example output:
{
  "primitives": [
```

```
    {"type": "open", "target": "drawer",
     "affordance_queries": [{"object": "drawer", "action": "open"}]}
  ]
}
```

## NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and Section 1 state three contributions (C1–C3): (1) the AffordSim framework, (2) the 50-task affordance-aware benchmark, and (3) zero-shot sim-to-real transfer of VLA policies trained on AffordSim data. Each is supported by experiments in Section 4 (Tables 1, 2, and 3).

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Appendix A discusses three limitations: (i) VoxAfford predictor coverage on novel geometries, transparent objects, and ambiguous functional regions; (ii) restriction to rigid bodies, with deformables, in-hand, and bimanual manipulation deferred to future work; (iii) Depth Anything 3 background reconstruction requires 10–20 photographs per scene and degrades sharply at viewpoints far from the captured poses.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper does not include theoretical results; it is a systems and benchmark contribution.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 4 describes the experimental setup, baselines, and evaluation protocol. Appendix I provides training hyperparameters for all baselines, Appendix H specifies the Manual baseline rules, and Appendix P together with Appendix Q detail the sim-to-real hardware setup and Depth Anything 3 reconstruction.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We will release the AffordSim codebase, benchmark tasks, and generated datasets under a permissive open-source license upon publication. An anonymized version is provided during review.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 4 and Appendix I specify all training details including architectures, hyperparameters, optimizers, learning rates, batch sizes, and training epochs for each baseline (BC, Diffusion Policy, ACT, VLA-adapter, Pi 0.5).

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Each task is evaluated over 100 simulation rollouts for trajectory collection (100 for benchmark policy evaluation, 10 for real-robot experiments). Success rates are reported as percentages. We report per-task results in addition to aggregated tier and category averages, allowing readers to assess variance across tasks.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Appendix I specifies that all imitation-learning experiments use a single node with 8 NVIDIA A800 80GB GPUs and reports per-task training times (2–26 hours depending on the baseline).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research involves simulation-based and tabletop robotic manipulation with no human subjects, no sensitive data, and no IRB-relevant procedures.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Beyond data collection, AffordSim also serves as a systematic evaluation suite: by curating affordance-critical tasks, it measures how well a policy actually understands object affordance rather than merely succeeds on grasp-stable scenarios, an evaluation that has been out of reach for existing platforms which lack benchmarks for affordance-dependent tasks. Positive impacts therefore include accelerating affordance-aware manipulation research, exposing failure modes that grasp-stability-only benchmarks hide, and enabling more reliable deployment in manufacturing, logistics, and assistive robotics; policies trained with AffordSim data further exhibit zero-shot sim-to-real transfer, demonstrating the practical utility of the collected trajectories. Potential negative impacts include enabling autonomous manipulation systems that may be deployed without sufficient human oversight or that displace manual labor; we view these as inherent to robotic-manipulation research rather than uniquely raised by this work, and the framework imposes no application-specific deployment.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The released assets (simulation tasks, benchmark data, evaluation code) do not pose risks for misuse beyond standard robotic manipulation research.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All external assets (VoxAfford, AnyGrasp, Isaac Sim, Depth Anything 3, Qwen-VL-Plus, baseline policies BC/DP/ACT/VLA-adapter/Pi 0.5) are cited in the main paper and bibliography. License information will be provided in the code release.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The AffordSim framework, benchmark tasks (Appendix J, Appendix K), and generated datasets are released at `https://anonymous.4open.science/r/AffordSim-8D94/` with documentation including installation instructions, task descriptions, and evaluation scripts.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or research with human subjects.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve research with human subjects.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [Yes]

    Justification: Section 3.1 describes the use of Qwen-VL-Plus as the VLM backbone for separate scene and task parsing followed by affordance-query conversion, which is a core component of the AffordSim pipeline. The full prompts are provided in Appendix R and Appendix S.