

Grounding Sim-to-Real Generalization in Robotic Manipulation: An Empirical Study with Vision-Language-Action Models

Ruixing Jin¹, Zicheng Zhu¹, Ruixiang Ouyang¹, Sheng Xu¹, Bo Yue¹, Zhizheng Wu¹, and Guiliang Liu^{*1,2}

¹School of Data Science, The Chinese University of Hong Kong, Shenzhen

²Shenzhen Loop Area Institute

{ruixingjin, zichengzhu, ruixiangouyang, shengxu1,
boyue}@link.cuhk.edu.cn

{wuzhizheng, liuguiliang}@cuhk.edu.cn

Abstract. Learning a generalist control policy for robotic manipulation typically relies on large-scale datasets. Given the high cost of real-world data collection, a practical alternative is to generate synthetic data through simulation. However, the resulting synthetic data often exhibits a significant gap from real-world distributions. While many prior studies have proposed algorithms to bridge the Sim-to-Real discrepancy, there remains a lack of principled research that grounds these methods in real-world manipulation tasks, particularly their performance on generalist policies such as Vision-Language-Action (VLA) models. In this study, we empirically examine the primary determinants of Sim-to-Real generalization across four dimensions: multi-level domain randomization, photorealistic rendering, physics-realistic modeling, and reinforcement learning updates. To support this study, we design a comprehensive evaluation protocol to quantify the real-world performance of manipulation tasks. The protocol accounts for key variations in background, lighting, distractors, object types, and spatial features. Through experiments involving over 10k real-world trials, we derive critical insights into Sim-to-Real transfer. To inform and advance future studies, we release both the robotic platforms and the evaluation protocol for public access to facilitate independent verification, thereby establishing a realistic and standardized benchmark for robotic manipulation policies.

1 Introduction

Developing precise and scalable robotic manipulation policy represents a key milestone toward realizing artificial general intelligence (AGI) [42]. To achieve this goal, Vision-Language-Action (VLA) models have emerged as a principled design for building general-purpose robotic agents, enabling an end-to-end mapping from visual observations and language instructions to continuous motor

* Corresponding Author.

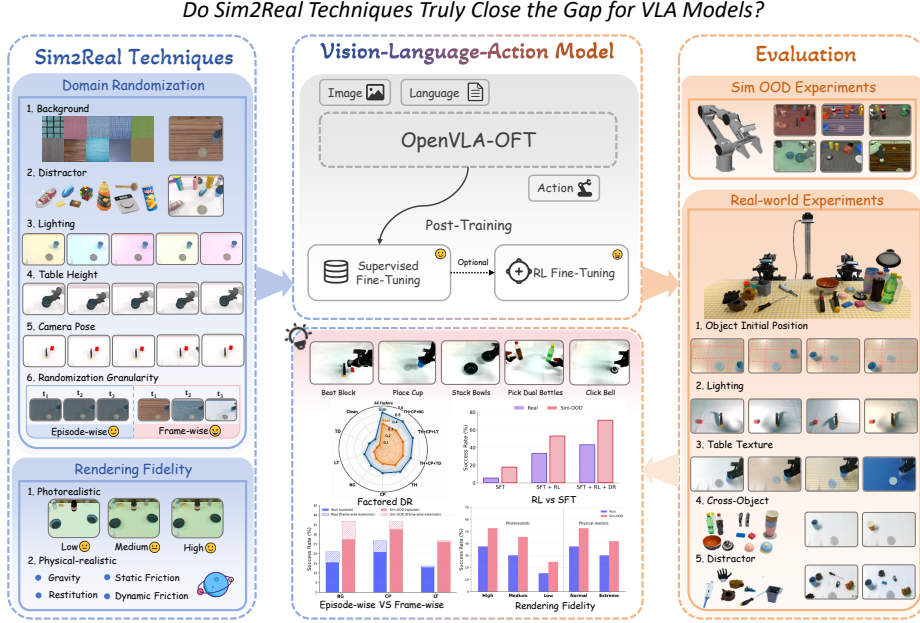


Fig. 1: Overview of our framework for analyzing Sim2Real generalization in Vision-Language-Action (VLA) models. We study how different Sim2Real techniques, including domain randomization, rendering fidelity, and reinforcement learning fine-tuning, influence generalization across Vision, Semantics, and Execution under both simulation OOD and real-world evaluations.

actions [26, 61]. VLA methods typically adapt multi-modal foundation models [5, 21, 49] to embodied control and fine-tune them on diverse robotic demonstrations [16, 39], thereby exhibiting compositional reasoning capabilities from foundation models [1, 18, 37, 45, 63]. This formulation allows a single policy to perform multi-task manipulation across heterogeneous environments, demonstrating promising language grounding and cross-task generalization [53, 60].

As a learning-based framework for robotic manipulation, VLA training typically follows a data-driven pipeline. However, this process requires a substantial amount of robotic operation trajectories, which are fundamentally scarce, costly to acquire, and inherently hardware-specific. These characteristics make it difficult to scale across diverse tasks, robotic platforms, and real-world scenarios [6]. To address these challenges, recent studies [9, 22, 23, 28, 59] have explored training VLA models using large-scale simulated data generated by diverse simulation platforms. These platforms include rule-based simulators, such as MuJoCo [48] and Isaac Lab [31], as well as learning-based environments built upon world models [44]. Compared to data collection in real-world environments, synthetic data can be produced efficiently at scale. However, significant discrepancies in dynamics, kinematics, sensor characteristics, or other factors between simulated and real settings result in a pronounced simulation-to-reality (Sim2Real) gap. As

a result, VLA models trained in simulation often exhibit degraded performance when deployed in real-world scenarios [4, 10, 35].

The ability to deploy policies trained in simulation to the physical world is therefore essential for scalable and practical robotics. Existing Sim2Real approaches primarily fall into four categories: 1) Domain randomization [7] improves robustness to out-of-distribution (OOD) real-world conditions by introducing diverse random perturbations during simulation training. 2) Domain adaptation [2, 13, 58] reduces the discrepancy between simulated and real domains by aligning them within a shared feature space. 3) Photo- or physics-realistic rendering mitigates the Sim2Real gap by increasing the visual fidelity of observations and the physical accuracy of environmental dynamics. 4) Reinforcement fine-tuning (RFT) [19] further improves VLA models by optimizing them through interaction with simulated environments using reinforcement learning (RL). However, these approaches are often evaluated independently on general robotics tasks or control models. It remains unclear *which specific factors influence, or to what extent they contribute to, the Sim2Real transferability of VLA models in robotic manipulation tasks*. Such a lack of mechanistic understanding limits the principled optimization of Sim2Real strategies and hinders systematic diagnosis of VLA model failures under real-world deployment.

In this work, we address this question through a factorized empirical analysis of the zero-shot Sim2Real performance of VLA models, where models are trained entirely in simulation and directly deployed to real-world tasks without any real-world data fine-tuning [59]. To evaluate **zero-shot Sim2Real transfer**, we build a unified benchmark for VLA systems based on the RoboTwin 2.0 simulation framework [4]. The benchmark includes a set of representative dual-arm manipulation tasks covering both short-horizon precision interactions and multi-stage behaviors. In addition to simulation OOD evaluation, we conduct real-world tests under controlled variations, including background textures, lighting perturbations, object instances, distractors, and spatial position changes across predefined grid layouts.

Within this unified framework, we conduct a factorized comparison of Sim2Real methods on zero-shot transfer performance, examining the effects of domain randomization factors, randomization granularity, photorealistic rendering fidelity, and RL fine-tuning. Through a massive real-world evaluation including more than 10,000 real-world trials, our study offers the following insights:

- **Spatial features speak louder than appearance.** Augmenting spatial features, such as table height and camera pose, yields larger improvements than purely visual perturbations like background textures or lighting. This suggests that spatial variation enhances the policy’s ability to adapt to different spatial configurations and strengthens the coupling between visual perception and motor execution.
- **Granularity matters.** Frame-wise domain randomization yields better zero-shot transfer than episode-wise strategies by improving the model’s attention to task-relevant objects rather than background variations.

- **Simulation fidelity enhances Sim2Real transfer.** Our results reinforce the conclusion that higher levels of photorealistic and physics-based fidelity consistently improve Sim2Real performance; however, the magnitude of improvement diminishes as the simulation fidelity reaches a threshold.
- **RL enhances robustness to distribution shifts.** Compared to SFT-only policies, RFT demonstrates greater resilience to object variations and environmental perturbations. When combined with structured domain randomization, the performance gains are further amplified, highlighting the complementary benefits of policy optimization and environmental diversity.

Inspired by RoboChallenge [55], we develop an online system that offers open access to standardized protocols, experimental settings, and real-world deployment interfaces of bimanual robots. This platform ensures result reproducibility and enables practitioners to gain deeper insights into Sim2Real transfer. It supports continued investigation and systematic benchmarking of VLA models in physical environments.

2 Related Work

2.1 Vision-Language-Action (VLA) Models

Recently, with the advancement of large language models (LLMs) and vision-language models (VLMs) [5, 21, 49], training Vision-Language-Action (VLA) models on large-scale robotic datasets [16, 39] has become a significant research direction in the robotics field. Early works such as RT-1 [3] and RT-2 [63] demonstrated the feasibility of training transformer architectures on large-scale robotic datasets by adding action heads to VLMs. Similarly, open-source models such as OpenVLA [18], and CogACT [20] have also proven that VLA architectures can achieve competitive performance across diverse hardware platforms. Meanwhile, recent studies have continuously optimized model architectures; for instance, π_0 [1] adopted flow matching, and OpenVLA-OFT [17] altered the decoding method for action tokens. However, these works primarily focus on simulated environments or deployment fine-tuned on real-world data. Our research will focus on empirical studies of Sim2Real transfer for VLA models.

2.2 Techniques for Sim-to-Real Transfer in Robotics

To bridge the gap between simulation and real-world deployment, several mainstream Sim2Real techniques have been widely studied in robotics. Representative approaches include domain randomization, domain adaptation, improving rendering realism, and RL-based post-training.

Domain Randomization (DR) remains the foundational approach for zero-shot transfer [11, 47, 62], which expands the simulated distribution by randomizing visual and physical parameters during training to encompass real-world variations. Tobin et al. [47] proposed that randomizing textures, lighting, and camera parameters enables zero-shot transfer for visual grasping tasks by deep

neural networks. Subsequently, DR has been widely adopted in object detection [11], robotics [43], autonomous driving [36, 40], and other fields. However, DR is typically applied monolithically [4, 59, 62], where multiple factors are randomized simultaneously, making it difficult to unravel their individual contributions. Recent studies also have determined relevant parameters by automatic learning [38], active exploration [29], Bayesian update [33, 34], offline inference [46] and continual learning [12]. In this work, we address this gap through a factorized empirical study that disentangles and ranks key domain randomization factors to identify their impact on zero-shot Sim2Real transfer for VLA models.

Photo/Physics Realistic Rendering provides a complementary strategy to reduce visual Sim-to-Real discrepancies by improving the realism of simulated observations. Instead of expanding the training distribution as in domain randomization, photorealistic simulation aims to narrow the visual reality gap at its source by generating high-fidelity synthetic data. Recent advances in neural scene reconstruction and rendering further enable realistic simulation environments derived from real-world captures [15, 50]. In robotics, modern simulation platforms support high-fidelity GPU-accelerated rendering and physics simulation [27, 51], making photorealism increasingly practical for large-scale training. In this work, we treat rendering fidelity as a controlled experimental factor and systematically evaluate its effect on zero-shot Sim2Real transfer.

Reinforcement Fine-Tuning (RFT) has emerged as a powerful paradigm for enhancing the capabilities of foundation models [25]. In large language models (LLMs) and vision-language models (VLMs), RL has been shown to significantly improve reasoning abilities [41] and align outputs with human preferences [14]. Beyond standard supervised fine-tuning (SFT), RL fine-tuning can unlock out-of-distribution generalization and stronger reasoning performance [8, 24]. These benefits have recently extended to vision-language-action (VLA) models, where RL fine-tuning has been shown to improve policy performance in both simulation and real-world settings [19, 24]. In this work, we further examine how RL fine-tuning influences zero-shot Sim2Real generalization in VLA models.

2.3 Real-Robot Evaluation

While simulation platforms are gradually maturing, real-robot evaluation remains indispensable. RoboChallenge [54] identifies this as a significant challenge in robotics and has constructed a large-scale online real-robot evaluation infrastructure, introducing multiple tabletop manipulation tasks. In establishing real-world test sets, Xie et al. [52] discovered that different environmental factors have varying impacts on model performance in the real world, and these factors are largely independent of each other, which means most pairs of factors exhibit no compounding effects. This finding also serves as an important basis for our design of real-world test sets. In this work, we design a controlled benchmark to comprehensively evaluate the zero-shot Sim-to-Real capability of VLA models.

3 Preliminaries

3.1 VLA Models for Robotic Manipulation

Robotic manipulation in the Sim2Real setting can be formulated as a Partially Observable Markov Decision Process (POMDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, P_T, R, \mu_0, \gamma)$ where: 1) The state $s_t \in \mathcal{S}$ captures the semantic information of a scene, encompassing the configuration (e.g., layouts, appearance, and physical characteristics) of various types of objects and the robots. 2) The action $a_t \in \mathcal{A}$ corresponds to low-level robot control commands. In our work, actions are parameterized as continuous target joint angles. 3) The observation $o_t \in \mathcal{O}$ represents the perceptual signals captured by sensors. These observations are typically non-Markovian and provide only partial information about the current state. 4) The transition function P_T describes the physical dynamics that map the current state s_t and executed action a_t to the next state s_{t+1} . 5) The reward function $R(s, a)$ measures task performance and, in many robotic manipulation settings, is defined by task goals (e.g., language instructions) and often simplified as a binary success signal indicating task completion [19, 57]. 6) μ_0 denotes the initial state distribution, and $\gamma \in (0, 1]$ is the discount factor that balances immediate and future rewards.

The VLA policy is conditioned on a language instruction l that specifies the task goal and learns a mapping $\pi(a_{t:t+M} \mid o_t, l)$ to predict a short-horizon sequence of future actions from the current observation and instruction [17]. The objective is to maximize expected task success under instruction-conditioned trajectories, allowing a unified policy to map high-level semantic commands to low-level motor control. This language interface enables generalization across diverse manipulation goals, object variations, and spatial configurations [1, 17, 18].

3.2 Zero-Shot Sim2Real Paradigm for VLA Models

Let $\hat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, P, R, \mu_0, \gamma)$ denote the real-world environment, and let $\mathcal{M}$ denote the corresponding simulated environment. In simulation, we collect trajectories $\tau = (o_0, a_0, \dots, o_T, a_T, l)$ generated by interacting with $\mathcal{M}$, where l specifies the manipulation task.

The policy is optimized in simulation as:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\tau \sim \mathcal{M}} [J(\tau, \pi; l)]. \quad (1)$$

We present a unified framework for analyzing zero-shot Sim2Real transfer in VLA systems. Our method organizes the problem into two components: policy learning and simulation design.

Supervised Fine-Tuning (SFT) We first train the VLA policy in simulation using supervised behavioral cloning. Given a demonstration dataset $\mathcal{D} =$

$\{(o_t^{(i)}, a_{t:t+M}^{(i)}, l^{(i)})\}$, where i denotes the data index, o_t denotes the current observation, and l the language instruction, the policy predicts a short-horizon sequence of future continuous actions conditioned on multimodal inputs.

Actions are parameterized as continuous low-level control commands (e.g., joint targets and gripper states), forming an action vector $a_{t:t+M} \in \mathbb{R}^{(M+1)d_a}$ over a prediction horizon of $M + 1$ steps. The policy directly outputs continuous action values rather than discretized tokens.

The supervised objective minimizes the L1 regression loss [17]:

$$\mathcal{L}_{\text{SFT}}(\theta) = \mathbb{E}_{(o_t, a_{t:t+M}, l) \sim \mathcal{D}} [\|\pi_\theta(o_t, l) - a_{t:t+M}\|_1], \quad (2)$$

where θ denotes the learnable model parameters.

Reinforcement Learning Fine-Tuning To further enhance robustness, we optionally perform reinforcement learning (RL) fine-tuning in simulation. In our study, we adopt Group Relative Policy Optimization (GRPO) [41] as used in SimpleVLA-RL [19]. GRPO optimizes the policy by comparing relative returns within a group of sampled trajectories, encouraging actions that outperform group-level baselines while stabilizing training for large VLA models.

Given a task-conditioned reward function $R(s_t, l)$, we optimize the policy using a group-relative objective. For each update, we sample G trajectories $\{\tau_i\}_{i=1}^G$ from the current policy and compute trajectory-level returns R_i . Instead of using a fixed reference model with KL regularization, we follow the DAPO-style modification and remove the KL penalty to encourage exploration [56].

The policy is optimized via the following clipped objective:

$$\mathcal{J}(\theta) = \mathbb{E}_{\{a_t\} \sim \pi_{\theta_{\text{old}}}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|\tau_i|} \sum_t \min \left(r_{i,t}(\theta) \hat{A}_i, \text{clip}(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) \right], \quad (3)$$

where

$$r_{i,t}(\theta) = \frac{\pi_\theta(a_{i,t} \mid o_{i,t})}{\pi_{\theta_{\text{old}}}(a_{i,t} \mid o_{i,t})}, \quad \hat{A}_i = \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G)}. \quad (4)$$

Here $r_{i,t}(\theta)$ denotes the policy likelihood ratio, $\hat{A}_i$ is the normalized advantage computed from trajectory returns, and ϵ is a clipping hyperparameter that constrains $r_{i,t}(\theta)$ within $[1 - \epsilon, 1 + \epsilon]$ to stabilize policy updates. Only trajectory groups containing both successful and unsuccessful rollouts are used for updates, ensuring meaningful relative comparisons within each group [19].

Notably, Sim2Real transfer in our setting is conducted in a zero-shot manner, where no real-world demonstrations are used during training and the learned policy must directly generalize to the real environment $\mathcal{M}$. However, discrepancies between $\mathcal{M}$ and $\hat{\mathcal{M}}$ introduce distribution shifts in visual appearance and geometry. To mitigate this gap, we focus on simulation-side techniques that improve robustness without requiring real-world adaptation data. While domain adaptation methods align simulated and real observations using real-world data, they fall outside the zero-shot setting considered in this work. Instead, we adopt

two complementary strategies for constructing simulation training data: **Domain Randomization**, which exposes the policy to diverse environment configurations, and **Rendering Fidelity**, which improves the realism of simulated observations, including photorealism and physical realism.

Domain Randomization (DR). We train policies over a distribution of simulator parameters rather than a fixed environment [6]. Let x denote an observation generated by the simulator, and let ξ represent the simulator parameters controlling the scene configuration. In a fixed simulator, observations follow $x \sim p_\xi(x)$.

Under DR, simulator parameters are sampled from a distribution $\xi \sim p(\xi)$, leading to a training distribution $x \sim \int p(x | \xi)p(\xi)d\xi$.

The learning objective for a policy f_ϕ becomes

$$\mathcal{L} = \mathbb{E}_{\xi \sim p(\xi)} [\mathbb{E}_{x \sim p(x|\xi)} \ell(f_\phi(x), y)] . \quad (5)$$

The objective of DR is not to exactly match the real distribution $p_{\text{real}}(x)$, but to expand the synthetic distribution such that $\text{supp}(p_{\text{real}}) \subseteq \text{supp}(p_{\text{DR}})$, so that real-world observations appear as a special case within the randomized synthetic distribution.

Following structured domain randomization principles from prior manufacturing DR literature [11, 62], we factorize domain randomization into structured components as $\xi = \{\xi_{\text{bg}}, \xi_{\text{dist}}, \xi_{\text{cam}}, \xi_{\text{light}}, \xi_{\text{table}}\}$.

Domain Randomization Factors. We decompose domain randomization into five structured components to analyze their individual contributions: **Background** (ξ_{BG}), which randomizes wall and table textures from a predefined image set; **Table Distractor** (ξ_{TD}), which samples the number, mesh, pose, and texture of irrelevant objects while avoiding overlap with task-relevant ones; **Camera Pose** (ξ_{CP}), which applies random translational offsets to the head camera position at initialization; **Lighting** (ξ_{LT}), which randomizes light color and position within scene bounds; and **Table Height** (ξ_{TH}), which samples table height from a predefined range.

Randomization Granularity. We further study the temporal scope of randomization by comparing two strategies: **episode-wise**, where factors are sampled once per episode and fixed during rollout, and **frame-wise**, where factors are resampled at every simulation step.

Rendering Fidelity. Beyond domain randomization, the fidelity of both visual rendering and physics simulation plays a critical role in bridging the Sim2Real gap. **Photorealism** influences the statistical properties of simulated observations by improving light transport, shadow consistency, and global illumination, producing images that more closely resemble real-world sensor data [15]. In practice, we vary rendering configurations—including ray tracing (RT), samples per pixel, path depth, and denoising—to control the level of visual realism. **Physical**

realism affects the accuracy of object dynamics and contact interactions during manipulation. We vary key physics parameters such as gravitational acceleration, static/dynamic friction coefficients, and restitution, which govern object motion and contact behavior in simulation. Together, these factors determine how closely the simulated environment approximates real-world conditions.

3.3 Problem Formulation: Empirical Study for Sim2Real VLAs

Real-world robotic data is scarce, expensive to collect, and highly environment-specific [6], making large-scale training of Vision-Language-Action (VLA) models using real demonstrations impractical. As a result, simulation becomes the primary scalable source of training data.

However, discrepancies between simulated and real environments often cause policies trained in simulation to degrade when deployed in the physical world [4, 10, 35]. Although some sim-to-real techniques, such as domain randomization, have shown promising results in robotics [30, 32, 59], its role in large-scale VLA models remains insufficiently understood. Compared with conventional control policies, VLA models jointly integrate visual perception, language grounding, and long-horizon decision making, which may introduce new sensitivities to simulation–reality discrepancies.

This motivates a systematic empirical study to disentangle how individual simulation design factors affect Sim2Real transfer in VLA models.

4 Experiments

4.1 Evaluation Protocol.

We conduct a systematic study of zero-shot simulation-to-real generalization under controlled distribution shifts. All experiments use the OpenVLA-OFT policy [17] trained with supervised fine-tuning (SFT) in simulation. Within this framework, we evaluate how structured domain randomization, randomization granularity, and rendering fidelity affect Sim2Real transfer in both simulation and real-world environments.

Simulation Dataset. All policies are trained in RobotWin 2.0 [4], a physics-based manipulation simulator that supports configurable lighting, camera pose, table height, background textures, and distractor objects.

We evaluate five manipulation tasks with varying horizon lengths and spatial complexity. For each task and each training setting, we collect 100 demonstration trajectories in simulation.

The simulated robot embodiment uses the Cobot Magic platform to ensure consistency with real-world deployment [4]. Visual input is provided by a single forward-facing Intel RealSense D435 RGB camera at 640×480 resolution, adopted for simplicity and clearer interpretation of experimental results. Camera intrinsics remain fixed across experiments, while camera pose may be randomized depending on the experimental condition.




Fig. 2: Real-world manipulation setup and randomized factors used in our experiments. Left: the physical platform with robot arms, camera, lighting, objects, and distractors. Right: examples of variations including object positions, lighting, backgrounds, object instances, and distractor configurations.

Policy Architecture We base our study on OpenVLA-OFT [17], a state-of-the-art Vision-Language-Action model that enhances the original OpenVLA [18] with optimized designs for efficient and robust control. Built on a fused visual encoder and a Llama2 7B language backbone [49], OpenVLA-OFT incorporates two key improvements: (1) Parallel decoding with action chunking, which replaces autoregressive generation to enable single-pass prediction of multiple future actions. (2) Continuous action representation with L1 regression, which avoids information loss from discretization and enables finer-grained control.

Simulation OOD Evaluation. To measure robustness within simulation, we construct Out-of-Distribution (Sim-OOD) environments by introducing unseen variations along individual factors. These include unseen background and table textures, unseen distractor objects, table height perturbations, camera pose shifts, and lighting disturbances.

Each factor is varied independently to isolate its effect on performance. Sim OOD evaluation serves as a controlled proxy for real-world generalization.

Real-World Evaluation Setup. Zero-shot real-world evaluation is conducted on a physical Piper robot platform equipped with a single RealSense D435 RGB camera (640×480 input resolution). An overview of the real-world setup is shown in Fig. 2. The workspace consists of an adjustable-height table with replaceable tablecloth textures and configurable lighting.

We evaluate robustness across multiple real-world factors:

- **Background variation:** three table textures (wood, blue fabric, yellow grid).
- **Lighting variation:** three distinct light positions with varying illumination colors.

Table 1: Factorized domain randomization results under Simulation OOD and zero-shot real-world evaluation across five tasks.

Data setting	Click Bell		Place Empty Cup		Beat Block Hammer		Stack Bowls Two		Place Dual Bottles	
	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real
Clean	14%	2.7%	11%	5.4%	17%	0.0%	36%	26.2%	19%	1.5%
BG	23%	11.5%	15%	10.2%	30%	2.7%	42%	40.4%	27%	12.3%
LT	22%	12.3%	12%	8.6%	29%	4.6%	43%	31.5%	24%	7.3%
TD	15%	3.1%	12%	6.9%	23%	0.0%	38%	27.3%	21%	4.6%
CP	34%	23.5%	20%	17.5%	32%	4.6%	47%	42.3%	30%	16.2%
TH	40%	36.9%	26%	24.8%	37%	6.5%	58%	49.6%	32%	15.4%
TH + CP + BG	50%	47.7%	34%	33.5%	43%	8.5%	62%	60.0%	42%	21.2%
TH + CP + LT	48%	44.2%	31%	27.9%	46%	7.3%	60%	54.6%	39%	20.0%
TH + CP + TD	43%	40.0%	29%	26.3%	38%	6.5%	56%	52.7%	32%	16.5%
All Factors	54%	49.7%	44%	41.0%	49%	11.5%	65%	63.1%	52%	23.8%

- **Object variation:** both *seen* and *unseen* object instances within the same task category, evaluating generalization to novel appearances.
- **Distractor presence:** task-irrelevant objects placed in the scene at three difficulty levels (2, 4, and 8 distractors), sampled from a separate object set distinct from task objects.
- **Spatial generalization:** evaluation on unseen positions within a predefined 3×3 grid.

For each task, we first evaluate the policy under a canonical base configuration. We then systematically vary one factor at a time while keeping all other factors fixed. For each configuration, we evaluate the task under multiple object positions by applying small random shifts, and report the resulting success rate across these variations.

4.2 Domain Randomization for Sim2Real Generalization

Effects of Domain Randomization To analyze the role of domain randomization, we decompose the five randomization factors into two categories: appearance factors, including Background (BG), Lighting (LT), and Table Distractor (TD); and spatial factors, including Table Height (TH) and Camera Position (CP). In this experiment, all randomization is applied in an episode-wise manner, where parameters are sampled once at scene initialization and kept fixed throughout each rollout.

Tab. 1 reports the success rates under Simulation OOD and zero-shot real-world evaluation across different domain randomization factors. Training with clean data leads to poor generalization, resulting in low success rates in both Sim-OOD and real environments. Introducing individual randomization factors consistently improves performance across tasks.

Lesson 1: Spatial randomization is the primary driver of Sim2Real generalization. Among all factors, spatial perturbations such as camera pose (CP) and table height (TH) provide the largest performance gains, significantly improving success rates across both Sim-OOD and real-world evaluations. In

Table 2: Results of randomization granularity. Δ denotes the gain of frame-wise over episode-wise.

Data Setting	Click Bell		Place Empty Cup		Beat Block Hammer		Stack Bowls Two		Pick Dual Bottles	
	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real
Episode-wise CP	34%	23.5%	20%	17.5%	32%	4.6%	47%	42.3%	30%	16.2%
Frame-wise CP	37%	31.9%	27%	25.6%	34%	7.3%	51%	49.2%	34%	19.6%
Δ (Frame – Episode)	3.0%	8.4%	7.0%	8.1%	2.0%	2.7%	4.0%	6.9%	4.0%	3.4%
Episode-wise BG	23%	11.5%	15%	10.2%	30%	2.7%	42%	40.4%	27%	12.3%
Frame-wise BG	31%	15.0%	27%	23.1%	41%	7.7%	53%	45.0%	32%	14.2%
Δ (Frame – Episode)	8.0%	3.5%	12.0%	12.9%	11.0%	5.0%	11.0%	4.6%	5.0%	1.9%
Episode-wise LT	22%	12.3%	12%	8.6%	29%	4.6%	43%	31.5%	24%	7.3%
Frame-wise LT	23%	14.6%	13%	9.0%	30%	5.0%	43%	32.6%	26%	7.3%
Δ (Frame – Episode)	1.0%	2.3%	1.0%	0.4%	1.0%	0.4%	0.0%	1.1%	2.0%	0.0%

contrast, appearance-level changes such as background (BG) and lighting (LT) yield smaller improvements when applied individually. This suggests that variations affecting geometric relationships and viewpoint are more critical for robust policy learning than purely visual changes.

Lesson 2: Appearance randomization provides complementary benefits when combined with spatial perturbations. Although background and lighting perturbations alone produce relatively modest gains, combining them with spatial randomization further improves performance. In particular, configurations such as TH + CP + BG and TH + CP + LT consistently outperform single-factor settings across most tasks. The best results are achieved when all factors are applied simultaneously, yielding the highest success rates in both Sim-OOD and real-world evaluations.

These results indicate that while spatial variations play a dominant role in improving policy robustness, appearance-level randomization helps cover additional visual discrepancies between simulation and reality when combined with spatial perturbations.

Effects of Randomization Granularity Building upon the episode-wise setting used in Experiment 1, we next investigate whether increasing the temporal frequency of randomization further improves zero-shot Sim2Real transfer. Specifically, we consider episode-wise and frame-wise strategies as two representative boundary cases of temporal stochasticity. In this study, we limit frame-wise randomization to background textures, lighting configuration, and camera position. Factors that directly affect task feasibility, including distractor placement and table height, remain episode-wise to preserve task stability, maintain consistent task planning, and avoid disrupting motion execution during rollout.

Lesson 3: Granularity matters. Applying randomization at the frame level consistently outperforms episode-wise randomization across tasks (Tab. 2). Frame-wise perturbations introduce continuous variation within each rollout, preventing the policy from overfitting to static scene configurations. This effect is particularly evident for background (BG) and camera pose (CP), where frame-wise randomization improves real-world success rates by roughly 3–8%

for CP and up to around 13% for BG, with similarly consistent gains observed under Sim-OOD evaluation. In contrast, lighting perturbations provide comparatively smaller improvements, typically around 0–2%, suggesting that illumination changes alone contribute less to generalization when applied at a finer temporal granularity.

4.3 Rendering Fidelity for Sim2Real Generalization

We next investigate whether increasing rendering fidelity improves zero-shot Sim2Real transfer (Fig. 3b). We compare three rendering presets that differ in ray-tracing activation and sampling configuration: The **Low** setting disables RT and uses low sampling and path depth. The **Medium** setting enables RT with the same sampling and path depth as the Low setting, while introducing additional rendering effects. The **High** setting further increases the sampling rate and path depth to achieve higher visual fidelity.

In addition to visual fidelity, we also consider physical realism. To examine the effect of physics fidelity, we construct a setting with intentionally distorted dynamics by modifying key physical parameters, including gravitational acceleration, static and dynamic friction coefficients, and restitution. These parameters are set to relatively extreme values to reduce the physical realism of the simulator, allowing us to evaluate how violations of realistic dynamics affect VLA training and Sim2Real transfer.

Lesson 4: Simulation fidelity enhances Sim2Real transfer. Fig. 3a reports the performance comparison across five tasks under different rendering fidelity settings. Increasing photorealism consistently improves real-world success rates, indicating that more realistic visual appearance helps the policy learn representations that transfer better to the physical environment. However, the improvement becomes marginal beyond the medium setting, suggesting that moderate photorealism already captures most transferable visual cues.

The figure also shows that improving physical realism leads to certain gains in Sim2Real performance for VLA models. However, the improvement is smaller compared with that obtained from increasing photorealism, indicating that visual fidelity plays a more dominant role in facilitating Sim2Real transfer.

4.4 Reinforcement Learning for Sim2Real Generalization

Finally, we evaluate the effect of reinforcement learning (RL) fine-tuning on zero-shot Sim2Real transfer. Following the design in SimpleVLA [19], we modify the action prediction head of OpenVLA-OFT by replacing the original MLP-based continuous action regression with the LLaMA2 output head that generates discretized action tokens optimized with a cross-entropy loss.

Based on this modification, we first re-train an SFT policy using 100 demonstrations per task, which serves as the initialization for all RL experiments. Starting from this SFT checkpoint, we further perform RL fine-tuning in simulation under different rollout settings. We compare three training variants: (1)




Fig. 3: Rendering fidelity analysis. (a) Quantitative results showing the impact of photorealism and physical realism on Sim-OOD and real-world success rates. (b) Example renderings under different photorealism levels (Low, Medium, High).

Table 3: Effect of reinforcement learning (RL) and domain randomization (DR) on zero-shot Sim2Real performance across five tasks.

Training Variant	Click Bell		Place Empty Cup		Beat Block Hammer		Stack Bowls Two		Pick Dual Bottles	
	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real	Sim-OOD	Real
SFT	11%	0%	23%	6.8%	15%	0%	33%	21.2%	10%	0%
SFT + RL	42%	36.9%	65%	34.8%	52%	9.6%	60%	59.2%	48%	26.5%
SFT + RL + DR	60%	50.8%	87%	51.4%	72%	16.2%	68%	64.6%	67%	30.8%

SFT trained in a clean simulation environment, (2) SFT followed by RL fine-tuning using rollouts collected in a clean simulation environment, and (3) SFT followed by RL fine-tuning using rollouts collected under domain randomization.

Lesson 5: RL enhances robustness to distribution shifts. As shown in Tab. 3, even when RL is trained under clean simulation conditions, it substantially improves performance over the SFT baseline, increasing the average real-world success rate from 5.6% to 33.4%. Notably, this clean RL setting already approaches the performance achieved by SFT combined with domain randomization as shown in Tab. 1, suggesting that RL itself introduces stronger policy generalization beyond supervised imitation. Furthermore, when domain randomization is incorporated during RL rollouts, the performance further improves to 42.8% in real-world evaluation and 70.8% in Sim-OOD, indicating that RL can effectively exploit diverse training experiences to further enhance Sim2Real robustness.

5 Conclusion

In this work, we systematically study how different Sim2Real techniques affect the generalization of Vision-Language-Action (VLA) models. Our experiments show that spatial domain randomization plays a dominant role in improving robustness, while appearance randomization provides complementary benefits. We further find that frame-wise randomization is more effective than episode-wise

perturbations, higher simulation fidelity enhances Sim2Real transfer, and reinforcement learning fine-tuning significantly improves robustness to distribution shifts.

In addition, we develop an online benchmarking platform that provides standardized protocols, experimental settings, and real-world deployment interfaces for bimanual robots, enabling reproducible evaluation and continued benchmarking of Sim2Real performance for VLA models.

Acknowledgments

This work is supported in part by Shenzhen Science and Technology Program under grant KJZD20240903104008012, Shenzhen Science and Technology Program under grant ZDCY202509011130000001, CUHK-CUHK(SZ)-GDSTC Joint Collaboration Fund No. 2025A0505000053, and GuangDong Key Laboratory of Big Data Computing (2021B1212040002).

References

1. Black, K., Brown, N., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Shi, L.X., Tanner, J., Vuong, Q., Walling, A., Wang, H., Zhilinsky, U.: $\pi 0$: A vision-language-action flow model for general robot control. ArXiv **abs/2410.24164** (2024)
2. Bousmalis, K., Irpan, A., Wohlhart, P., Bai, Y., Kelcey, M., Kalakrishnan, M., Downs, L., Ibarz, J., Pastor, P., Konolige, K., Levine, S., Vanhoucke, V.: Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In: International Conference on Robotics and Automation, ICRA. pp. 4243–4250 (2018)
3. Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., Ibarz, J., Ichter, B., Irpan, A., Jackson, T., Jesmonth, S., Joshi, N.J., et al.: Rt-1: Robotics transformer for real-world control at scale. ArXiv **abs/2212.06817** (2022)
4. Chen, T., Chen, Z., Chen, B., Cai, Z., Liu, Y., Liang, Q., Li, Z., Lin, X., Ge, Y., Gu, Z., Deng, W., Guo, Y., Nian, T., Xie, X., Chen, Q., Su, K., Xu, T., Liu, G., Hu, M., ang Gao, H., Wang, K., Liang, Z., Qin, Y., Yang, X., Luo, P., Mu, Y.: Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. ArXiv **abs/2506.18088** (2025)
5. Chen, X., Djolonga, J., Padlewski, P., Mustafa, B., Changpinyo, S., Wu, J., Ruiz, C.R., Goodman, S., Wang, X., Tay, Y., Shakeri, S., Dehghani, M., others: Pali-x: On scaling up a multilingual vision and language model. ArXiv **abs/2305.18565** (2023)
6. Chen, X., Hu, J., Jin, C., Li, L., Wang, L.: Understanding domain randomization for sim-to-real transfer. In: International Conference on Learning Representations, ICLR (2022)
7. Chen, X., Hu, J., Jin, C., Li, L., Wang, L.: Understanding domain randomization for sim-to-real transfer. In: International Conference on Learning Representations, ICLR (2022)

8. Chu, T., Zhai, Y., Yang, J., Tong, S., Xie, S., Schuurmans, D., Le, Q.V., Levine, S., Ma, Y.: SFT memorizes, RL generalizes: A comparative study of foundation model post-training. In: International Conference on Machine Learning, ICML. vol. 267 (2025)
9. Deng, S., Yan, M., Wei, S., Ma, H., Yang, Y., Chen, J., Zhang, Z., Yang, T., Zhang, X., Zhang, W., et al.: Graspvla: a grasping foundation model pre-trained on billion-scale synthetic action data. arXiv preprint arXiv:2505.03233 (2025)
10. ud Din, M., Akram, W., Saoud, L.S., Rosell, J., Hussain, I.: Vision language action models in robotic manipulation: A systematic review. ArXiv **abs/2507.10672** (2025)
11. Fan, J., Xie, Y., Li, Z.: Domain randomization for object detection: A survey. In: International Conference on Advanced Electronic Materials, Computers and Software Engineering , AEMCSE (2025)
12. Josifovski, J., Auddy, S., Malmir, M., Piater, J.H., Knoll, A., Navarro-Guerrero, N.: Continual domain randomization. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS. pp. 4965–4972 (2024)
13. Josifovski, J., Gu, S., Malmir, M., Huang, H., Auddy, S., Navarro-Guerrero, N., Spanos, C., Knoll, A.: Safe continual domain adaptation after sim2real transfer of reinforcement learning policies in robotics. arXiv preprint arXiv:2503.10949 (2025)
14. Kaufmann, T., Weng, P., Bengs, V., Hüllermeier, E.: A survey of reinforcement learning from human feedback. ArXiv **abs/2312.14925** (2023)
15. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., et al.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. (2023)
16. Khazatsky, A., et al.: Droid: A large-scale in-the-wild robot manipulation dataset. ArXiv **abs/2403.12945** (2024), <https://api.semanticscholar.org/CorpusID:268531351>
17. Kim, M.J., Finn, C., Liang, P.: Fine-tuning vision-language-action models: Optimizing speed and success. ArXiv **abs/2502.19645** (2025)
18. Kim, M.J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E.P., Lam, G., Sanketi, P.R., Vuong, Q., Kollar, T., Burchfiel, B., Tedrake, R., Sadigh, D., Levine, S., Liang, P., Finn, C.: Openvla: An open-source vision-language-action model. ArXiv **abs/2406.09246** (2024)
19. Li, H., Zuo, Y., Yu, J., Zhang, Y., Yang, Z., Zhang, K., Zhu, X., Zhang, Y., Chen, T., Cui, G., Wang, D., Luo, D., Fan, Y., Sun, Y., Zeng, J., Pang, J., Zhang, S., Wang, Y., Mu, Y., Zhou, B., Ding, N.: Simplevla-rl: Scaling vla training via reinforcement learning. In: International Conference on Learning Representations, ICLR (2026)
20. Li, Q., Liang, Y., Wang, Z., Luo, L., Chen, X., Liao, M., Wei, F., Deng, Y., Xu, S., Zhang, Y., et al.: Cogact: A foundational vision-language-action model for synergizing cognition and action in robotic manipulation. arXiv preprint arXiv:2411.19650 (2024)
21. Li, Z., Chen, G., Liu, S., Wang, S., Vibashan, V., Ji, Y., Lan, S., Zhang, H., Zhao, Y., Radhakrishnan, S., Chang, N., Sapra, K., Deshmukh, A.S., Rintamaki, T., Le, M., Karmanov, I., Voegtli, L., Fischer, P., Huang, D.A., Roman, T., Lu, T., Álvarez, J.M., Catanzaro, B., Kautz, J., Tao, A., Liu, G., Yu, Z.: Eagle 2: Building post-training data strategies from scratch for frontier vision-language models. ArXiv **abs/2501.14818** (2025), <https://api.semanticscholar.org/CorpusID:275921957>
22. Lin, S., Qing, Y., Liu, L., Zhou, M., Jin, R., Fan, X., Liu, G.: Dygro-vla: Cross-task scaling of vision-language-action models via dynamic grouped residual optimization (2026), <https://api.semanticscholar.org/CorpusID:288654089>

23. Liu, G., Deng, Y., Zhao, R., Zhou, H., Chen, J., Chen, J., Xu, R., Tai, Y., Jia, K.: Dexscale: Automating data scaling for sim2real generalizable robot control. In: International Conference on Machine Learning, ICML (2025)
24. Liu, J., Gao, F., Wei, B., Chen, X., Liao, Q., Wu, Y., Yu, C., Wang, Y.: What can RL bring to VLA generalization? an empirical study. In: Conference and Workshop on Neural Information Processing Systems, NeurIPS (2025)
25. Liu, K., Yang, D., Qian, Z., Yin, W., Wang, Y., Li, H., Liu, J., Zhai, P., Liu, Y., Zhang, L.: Reinforcement learning meets large language models: A survey of advancements and applications across the llm lifecycle. ArXiv **abs/2509.16679** (2025), <https://api.semanticscholar.org/CorpusID:281421479>
26. Ma, Y., Song, Z., Zhuang, Y., Hao, J., King, I.: A survey on vision-language-action models for embodied ai. arXiv preprint arXiv:2405.14093 (2024)
27. Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., State, G.: Isaac gym: High performance GPU based physics simulation for robot learning. In: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2) (2021)
28. Mandlekar, A., Nasiriany, S., Wen, B., Akinola, I., Narang, Y., Fan, L., Zhu, Y., Fox, D.: Mimicgen: A data generation system for scalable robot learning using human demonstrations. In: Annual Conference on Robot Learning, CoRL (2023)
29. Mehta, B., Diaz, M., Golemo, F., Pal, C.J., Paull, L.: Active domain randomization. In: Conference on Robot Learning, CoRL. vol. 100, pp. 1162–1176 (2019)
30. Mehta, B., Diaz, M., Golemo, F., Pal, C.J., Paull, L.: Active domain randomization. ArXiv **abs/1904.04762** (2019), <https://api.semanticscholar.org/CorpusID:104291994>
31. Mittal, M., Roth, P., Tigue, J., Richard, A., Zhang, O., Du, P., Serrano-Muñoz, A., Yao, X., Zurbügg, R., Rudin, N., Wawrzyniak, L., Rakhsha, M., Denzler, A., Heiden, E., Borovicka, A., Ahmed, O., Akinola, I., Anwar, A., Carlson, M.T., Feng, J.Y., Garg, A., Gasoto, R., Gulich, L., Guo, Y., Gussert, M., Hansen, A., Kulkarni, M., Li, C., Liu, W., Makoviychuk, V., Malczyk, G., Mazhar, H., Moghani, M., Murali, A., Noseworthy, M., Poddubny, A., Ratliff, N., Rehberg, W., Schwarke, C., Singh, R., Smith, J.L., Tang, B., Thaker, R., Treppe, M., Wyk, K.V., Yu, F., Millane, A., Ramasamy, V., Steiner, R., Subramanian, S., Volk, C., Chen, C., Jawale, N., Kuruttukulam, A.V., Lin, M.A., Mandlekar, A., Patzwaldt, K., Welsh, J., Zhao, H., Anes, F., Laffleche, J.F., Moënné-Loccoz, N., Park, S., Stepinski, R., Gelder, D.V., Ameyor, C., Carius, J., Chang, J., Chen, A.H., de Heras Ciechowski, P., Daviet, G., Mohajerani, M., von Muralt, J., Reutsky, V., Sauter, M., Schirm, S., Shi, E.L., Terdiman, P., Vilella, K., Widmer, T., Yeoman, G., Chen, T., Grizan, S., Li, C., Li, L., Smith, C., Wiltz, R., Alexis, K., Chang, Y., Chu, D., Fan, L.J., Farshidian, F., Handa, A., Huang, S., Hutter, M., Narang, Y., Pouya, S., Sheng, S., Zhu, Y., Macklin, M., Moravanszky, A., Reist, P., Guo, Y., Hoeller, D., State, G.: Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning. arXiv preprint arXiv:2511.04831 (2025), <https://arxiv.org/abs/2511.04831>
32. Muratore, F., Eilers, C., Gienger, M., Peters, J.: Data-efficient domain randomization with bayesian optimization. IEEE Robotics and Automation Letters **6**, 911–918 (2020), <https://api.semanticscholar.org/CorpusID:230524079>
33. Muratore, F., Eilers, C., Gienger, M., Peters, J.: Data-efficient domain randomization with bayesian optimization. IEEE Robotics and Automation Letters **6**(2), 911–918 (2021)

34. Muratore, F., Gruner, T., Wiese, F., Belousov, B., Gienger, M., Peters, J.: Neural posterior domain randomization. In: Conference on Robot Learning, CoRL. vol. 164, pp. 1532–1542 (2021)
35. Nasiriany, S., Maddukuri, A., Zhang, L., Parikh, A., Lo, A., Joshi, A., Mandlekar, A., Zhu, Y.: Robocasa: Large-scale simulation of everyday tasks for generalist robots. ArXiv **abs/2406.02523** (2024), <https://api.semanticscholar.org/CorpusID:270226600>
36. Niu, H., Hu, J., Cui, Z., Zhang, Y.: Dr2l: Surfacing corner cases to robustify autonomous driving via domain randomization reinforcement learning. In: International Conference on Computer Science and Application Engineering (2021)
37. Nvidia, Bjorck, J., Castañeda, F., Cherniadev, N., Da, X., Ding, R., LinxiJimFan, Fang, Y., Fox, D., Hu, F., Huang, S., Jang, J., Jiang, Z., Kautz, J., Kundalia, K., Lao, L., Li, Z., Lin, Z., Lin, K., Liu, G., Llontop, E., Magne, L., Mandlekar, A., Narayan, A., Nasiriany, S., Reed, S., Tan, Y.L., Wang, G., Wang, Z., Wang, J., Wang, Q., Xiang, J., Xie, Y., Xu, Y., Xu, Z.T., Ye, S., Yu, Z., Zhang, A., Zhang, H., Zhao, Y., Zheng, R., Zhu, Y.: Gr00t n1: An open foundation model for generalist humanoid robots. ArXiv **abs/2503.14734** (2025), <https://api.semanticscholar.org/CorpusID:277113335>
38. OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., Zhang, L.: Solving rubik’s cube with a robot hand. CoRR **abs/1910.07113** (2019)
39. Padalkar, A., et al.: Open x-embodiment: Robotic learning datasets and rt-x models : Open x-embodiment collaboration0. International Conference on Robotics and Automation ,ICRA pp. 6892–6903 (2023), <https://api.semanticscholar.org/CorpusID:263626099>
40. Pouyanfar, S., Saleem, M., George, N., Chen, S.C.: Roads: Randomization for obstacle avoidance and driving in simulation. In: Conference on Computer Vision and Pattern Recognition Workshops, CVPRW (2019)
41. Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J.M., Zhang, M., Li, Y.K., Wu, Y., Guo, D.: Deepseekmath: Pushing the limits of mathematical reasoning in open language models. ArXiv **abs/2402.03300** (2024)
42. Smith, C., Karayiannidis, Y., Nalpantidis, L., Gratal, X., Qi, P., Dimarogonas, D.V., Kragic, D.: Dual arm manipulation - A survey. Robotics and Autonomous Systems **60**(10), 1340–1353 (2012)
43. Tan, J., Zhang, T., Coumans, E., Iscen, A., Bai, Y., Hafner, D., Bohez, S., Vanhoucke, V.: Sim-to-real: Learning agile locomotion for quadruped robots. ArXiv **abs/1804.10332** (2018)
44. Team, G., Ye, A., Wang, B., Ni, C., Huang, G., Zhao, G., Li, H., Zhu, J., Li, K., Xu, M., et al.: Gigaworld-0: World models as data engine to empower embodied ai. arXiv preprint arXiv:2511.19861 (2025)
45. Team, O.M., Ghosh, D., Walke, H.R., Pertsch, K., Black, K., Mees, O., Dasari, S., Hejna, J., Kreiman, T., Xu, C., Luo, J., Tan, Y.L., Sanketi, P.R., Vuong, Q., Xiao, T., Sadigh, D., Finn, C., Levine, S.: Octo: An open-source generalist robot policy. ArXiv **abs/2405.12213** (2024), <https://api.semanticscholar.org/CorpusID:266379116>
46. Tiboni, G., Arndt, K., Kyrki, V.: DROPO: sim-to-real transfer with offline domain randomization. Robotics and Autonomous Systems **166**, 104432 (2023)
47. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: International Conference on Intelligent Robots and Systems, IROS (2017)

48. Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS. pp. 5026–5033 (2012)
49. Touvron, H., et al.: Llama 2: Open foundation and fine-tuned chat models. ArXiv **abs/2307.09288** (2023), <https://api.semanticscholar.org/CorpusID:259950998>
50. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: conference on computer vision and pattern recognition, CVPR (2024)
51. Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., et al.: Sapien: A simulated part-based interactive environment. In: conference on computer vision and pattern recognition, CVPR (2020)
52. Xie, A., Lee, L., Xiao, T., Finn, C.: Decomposing the generalization gap in imitation learning for visual robotic manipulation. Conference on Robotics and Automation, ICRA (2023)
53. Xu, S., Jin, R., Zhou, H., Yue, B., Qiao, G., Tai, Y., Deng, Y., Jia, K., Liu, G.: From reaction to anticipation: Proactive failure recovery through agentic task graph for robotic manipulation (2026), <https://api.semanticscholar.org/CorpusID:288259953>
54. Yakefu, A., Xie, B., Xu, C., Zhang, E., Zhou, E., Jia, F., Yang, H., Fan, H., Zhang, H., Peng, H., Tan, J., Huang, J., et al.: Robochallenge: Large-scale real-robot evaluation of embodied policies. ArXiv **abs/2510.17950** (2025)
55. Yakefu, A., Xie, B., Xu, C., Zhang, E., Zhou, E., Jia, F., Yang, H., Fan, H., Zhang, H., Peng, H., et al.: Robochallenge: Large-scale real-robot evaluation of embodied policies. arXiv preprint arXiv:2510.17950 (2025)
56. Yu, Q., Zhang, Z., et al.: Dapo: An open-source llm reinforcement learning system at scale. ArXiv **abs/2503.14476** (2025), <https://api.semanticscholar.org/CorpusID:277104124>
57. Yue, B., Guo, S., Hu, K., Wang, C., Wang, B., Jia, K., Liu, G.: Real-time verification of embodied reasoning for generative skill acquisition. arXiv preprint arXiv:2505.11175 (2025)
58. Zhang, J., Tai, L., Yun, P., Xiong, Y., Liu, M., Boedecker, J., Burgard, W.: Vr-goggles for robots: Real-to-sim domain adaptation for visual control. IEEE Robotics and Automation Letters **4**(2), 1148–1155 (2019)
59. Zhao, R., Xu, S., Jin, R., Deng, Y., Tai, Y., Jia, K., Liu, G.: Sim2real VLA: Zero-shot generalization of synthesized skills to realistic manipulation. In: International Conference on Learning Representations, ICLR (2026)
60. Zhao, W.X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., Du, Y., Yang, C., Chen, Y., Chen, Z., Jiang, J., Ren, R., Li, Y., Tang, X., Liu, Z., Liu, P., Nie, J., rong Wen, J.: A survey of large language models. ArXiv **abs/2303.18223** (2023), <https://api.semanticscholar.org/CorpusID:257900969>
61. Zheng, Y., Yao, L., Su, Y., Zhang, Y., Wang, Y., Zhao, S., Zhang, Y., Chau, L.P.: A survey of embodied learning for object-centric robotic manipulation. Machine Intelligence Research pp. 1–39 (2025)
62. Zhu, X., Henningson, J., Li, D., Mårtensson, P., Hanson, L., Björkman, M., Maki, A.: Domain randomization for object detection in manufacturing applications using synthetic data: A comprehensive study. In: International Conference on Robotics and Automation (ICRA). pp. 16715–16721. IEEE (2025)

63. Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F., Wu, J., Wohlhart, P., Welker, S., Wahid, A., et al.: Rt-2: Vision-language-action models transfer web knowledge to robotic control. In: Conference on Robot Learning, CoRL. PMLR (2023)