RobotArena ∞ : Unlimited Robot Benchmarking via Real-to-Sim Translation

Anonymous authorsPaper under double-blind review

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

025

026

027

028

029

031

032

033 034

035

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

The pursuit of robot generalists—instructable agents capable of performing diverse tasks across diverse environments—demands rigorous and scalable evaluation. Yet real-world testing of robot policies remains fundamentally constrained: it is labor-intensive, slow, unsafe at scale, and difficult to reproduce. Existing simulation benchmarks are similarly limited, as they train and test policies within the same synthetic domains and cannot assess models trained primarily on real-world demonstrations, which is the dominant paradigm for today's vision-language-action (VLA) models. As policies expand in scope and complexity, these barriers only intensify, since defining "success" in robotics often hinges on nuanced human judgments of execution quality. In this paper, we introduce a new benchmarking framework that overcomes these challenges by shifting VLA evaluation into large-scale simulated environments augmented with online human feedback. Leveraging advances in vision-language models, 2D-to-3D generative modeling, and differentiable rendering, our approach automatically converts video demonstrations from widely used robot datasets into simulated counterparts. Within these digital twins, we assess VLA policies using both automated VLM-guided scoring and scalable human preference judgments collected from crowdworkers—transforming human involvement from tedious scene setup, resetting, and safety supervision into lightweight preference comparisons. To measure robustness, we systematically perturb simulated environments along multiple axes—textures and object placements — stress-testing policy generalization under controlled variation. The result is a continuously evolving, reproducible, and scalable benchmark for realworld-trained robot manipulation policies, addressing a critical missing capability in today's robotics landscape.

1 Introduction

While recent years have witnessed substantial progress in developing more capable and general robot policies, their evaluation remains a persistent challenge and lacks standardization. This problem becomes especially acute as policies grow more generalist, requiring broader and more diverse evaluation scenarios. **Real-world evaluation is inherently unscalable**: it is limited by logistics, safety concerns, and reproducibility issues, and requires significant human involvement for setup, execution, and scoring. Human operators must supervise trials and manually reset scenes, which restricts the scale and frequency of evaluations Vincent et al. (2024); Abou-Chakra et al. (2025); Li et al. (2024). Such manual oversight also raises concerns about consistency and fairness, particularly when baselines and new models are compared under slightly different conditions.

Centralized physical evaluation provides a gold standard, where policies are tested under identical conditions, typically by submitting containers or shipping robots to a shared testing site. Notable examples include the Amazon Picking Challenge Correll et al. (2016), the Open-Vocabulary Mobile Manipulation Challenge Yenamandra et al. (2023), and RoboCup@Home Matamoros et al. (2019). However, the high cost for both organizers and participants means such events occur infrequently, often no more than once a year. In contrast, fields such as computer vision and natural language processing have advanced rapidly thanks to standardized benchmarks that provide consistent metrics, clear performance targets, and a foundation for fair comparison across methods Deng et al. (2009); Schuhmann et al. (2022). Our work is particularly inspired by LMarena Chiang et al. (2024), a large-scale, crowdsourced evaluation framework that benchmarks LLMs and VLMs through direct pairwise



Figure 1: RobotArena ∞ provides a scalable and extensible robot benchmarking framework by automating environment construction and evaluation. It automatically generates simulated environment seeded from real videos, deploys robot policies, and evaluates them using VLMs and crowdsourced workers that cast preferences between pairs of execution videos. The simulated environments are derived from both in-distribution and out-of-distribution videos, enabling rigorous tests of generalization in contemporary VLAs.

comparisons of responses to the same prompt by human annotators. By aggregating thousands of such head-to-head matchups across diverse prompts, LMarena produces an Elo-style ranking that reflects collective judgments of model quality. Motivated by this success, we ask: **What would be the analogue of LMArena for Robotics?**

RobotArena ∞ : We introduce RobotArena ∞ , a new benchmarking framework that scales robot evaluation by deploying policies in automatically constructed simulated environments and assessing them through automatic VLM score and online human preference feedback. Our framework first automatically translates real videos into corresponding simulation environments. This reality-to-simulation translation method builds on recent advances in vision-language models for scene understanding, 2D-to-3D generative models for 3D asset creation from image crops and videos, and our differentiable rendering techniques for object pose estimation and robot–camera calibration.

Once constructed, we deploy VLAs in these environments and evaluate their execution trajectories using two complementary strategies: (1) absolute evaluation, in which prompted VLMs estimate task progress scores for each video frame, and (2) relative evaluation, in which human annotators express pairwise preferences between execution videos from different VLAs performing the same task. A key strength of RobotArena ∞ is its ability to measure both **in-distribution** performance—by testing on simulation environments seeded from training videos in established datasets such as Bridge Walke et al. (2023b)—and **out-of-distribution** performance, by testing on environments generated from videos outside the training set. To assess robustness and generalization, we also introduce systematic perturbations, such as variations in lighting, object placement, and background appearance. Our initial benchmark aggregates more than 7000 preference pairs across one hundred nominal environments and hundreds of perturbations, comparing four VLAs from independent labs worldwide. To the best of our knowledge, this constitutes the largest-scale robot evaluation effort to date.

The diversity of our reproducible evaluations yields several insights. First, VLAs perform poorly when tested on datasets they were not trained on, indicating that current models are not true generalists but instead specialists in recognizing the environments underlying their training data, as also found in Xing et al. (2025). Second, within the same environment, performance degrades under controlled perturbations—especially color changes—highlighting their limited robustness to distribution shifts. Yet, despite these limitations, consistent performance rankings emerge across models, demonstrating that architectural and data choices do produce measurable differences.

RobotArena ∞ is inspired by prior efforts to design scalable robot benchmarks, particularly the seminal contributions of BEHAVIOR Li et al. (2024) and SIMPLER Li et al. (2024). Unlike prior work, RobotArena ∞ provides a smooth scaling evaluation of VLA abilities. BEHAVIOR boasts an impressive manual effort of asset and environment creation, while SIMPLER reconstructs four

real-world Bridge scenes and includes hand-designed reward functions. Compared to these efforts, RobotArena ∞ offers a far more scalable and extensible framework by automating environment construction and evaluation. In summary, our contributions are as follows:

- 1. We present a scalable and extensible benchmarking protocol for robotics, by coupling physics engines, real-to-sim translation and human preference feedback.
- 2. We introduce a fully automated reality-to-simulation translation pipeline built upon VLMs, 2D-to-3D generative models and differentiable rendering.
- 3. We evaluate VLAs from labs worldwide across hundreds of environments and thousands of human preferences, the most extensive robot evaluation to date.
- 4. We present **key evaluation results** that reveal how current generalist policies generalize—or fail to—under distribution shifts.

Our benchmark is not without limitations. We outline these and discuss future directions and extensions. Importantly, RobotArena ∞ will continue to benefit from advances in physics engines and real-to-sim research. Both our benchmark environments and evaluation code will be publicly released and centrally maintained for continual support.

2 RELATED WORK

Generalist Robot Manipulation Policies Recent advances in robot foundation models have led to significant progress in generalist manipulation policies (Brohan et al., 2023; Kim et al., 2024; Octo Model Team et al., 2023; Kalashnikov et al., 2021; Ehsani et al., 2023; Bharadhwaj et al., 2024; Liu et al., 2024; Anil et al., 2023; Sridhar et al., 2023; Ye et al., 2024; Black et al., 2024; Pertsch et al., 2025), driven by the availability of large-scale robot datasets (O'Neill et al., 2024; Walke et al., 2023a; Khazatsky et al., 2024b; Shah et al., 2023). These models are trained to perform a wide range of tasks –including pick-and-place, cloth folding, and tool use (Walke et al., 2023a; Kim et al., 2024; Black et al., 2024; Pertsch et al., 2025)—and to generalize across diverse environments with varying backgrounds and distractors (Zhou et al., 2024; Fu et al., 2024), as well as across different robot embodiments such as arms, quadrupeds, and drones (Yang et al., 2024; Doshi et al., 2024). As these capabilities grow, evaluating such policies becomes increasingly labor intensive, requiring performance assessments across many different tasks, scenes, and embodiments.

Evaluating Robot Policies in The Real World Evaluating robot policies in the real world in a fair, comprehensive, and reproducible way remains a major challenge. Most methods are evaluated in custom lab-specific settings using proprietary hardware, task definitions, and success metrics, which makes cross-institution comparisons difficult. While standardized benchmarks exist for focused areas such as grasp prediction (Fang et al., 2020) and motion planning (Moll et al., 2014; Chamzas et al., 2022), extending these to cover generalist skills remains complex. Standardization efforts using shared object sets (e.g., YCB (Calli et al., 2015), IKEA (Heo et al., 2023), NIST (Kimble et al., 2020), ACRV (Leitner et al., 2016)) help reduce some variability, but differences in hardware, camera placement, lighting, and workspace setup still hinder consistent evaluation across labs. Real-world evaluation is often time-consuming and labor intensive. Human operators are typically required to supervise trials and manually reset scenes, limiting the scale and frequency of evaluations. For example, Cheng et al.(Chi et al., 2024) report manually aligning a T-shaped object into 20 predefined start configurations for each rollout, a process repeated across all baselines. This reset step is nonparallelizable and presents a significant bottleneck when evaluating policies across different tasks, agents, and environments(Vincent et al., 2024; Abou-Chakra et al., 2025). While recent systems such as AutoEval (Zhou et al., 2025) aim to automate evaluation, they are often limited in scope—e.g., supporting only five tasks in three static real-world scenes. RoboArena Atreya et al. (2025) builds a system of distribution real-world evaluation where human users reset the scene, run the robot policies and evaluate the resulting robot executions. It targets scenes from the DROID environment and evaluatea a set of policies finetuned on DROID. As real-world robot datasets and generalist policies continue to expand in scale and complexity, the need for a more scalable, general, and continuous evaluation framework becomes increasingly urgent.

Evaluating Robot Policies in Simulation Simulation offers a scalable and safe alternative for evaluating robot policies, with numerous benchmarks like: RLBench James et al. (2020), Colosseum (Pumacay et al., 2024), CALVIN (Mees et al., 2022), LIBERO (Liu et al., 2024), Per-ACT2 (Grotz et al., 2024), Meta-World (Yu et al., 2020), IKEA Simulation (Lee et al., 2019),

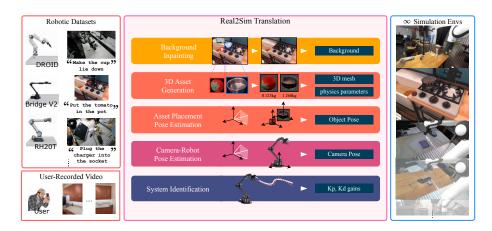


Figure 2: **Automated video-to-simulation translation** in RobotArena ∞ . Given a scene from a robotic dataset, we automatically create a simulated environment for the scene.

and Behavior-1K (Li et al., 2024). However, these typically assume policies are trained and tested in the same simulated environments, potentially favoring specialist policies that exploit simulation-specific advantages over generalist models trained on real-world data. Our approach, in contrast, uses simulation strictly as an evaluation environment, independent of the policy's training origin. We contend that simulation is increasingly viable for policy evaluation due to improving physics engines and advancements in generative models and VLMs, which can automate scene asset creation and task success detection. RobotArena ∞ leverages these by automatically generating diverse simulation-based evaluation environments. The closest work, SIMPLER Li et al. (2024), creates high-fidelity replicas of real scenes with human efforts, showing strong sim-to-real correlation. Unlike SIMPLER (Li et al., 2024), RobotArena ∞ automates both scene generation and task evaluation across varied tasks and datasets, and evaluates policy robustness and generalization by systematic perturbations.

3 Translating Videos to Simulation for Policy Evaluation

RobotArena ∞ leverages VLMs for scene understanding and 2D-to-3D generative models to automatically create scalable simulation environments from video. We mainly consider robot demonstration datasets, where cameras are typically static sources of video. Our goal is to generate environments and tasks that include both in-distribution and out-of-distribution cases for current VLAs—avoiding inflated performance while revealing their relative strengths and limitations.

3.1 Mapping Demonstration Videos to Simulation

We develop an automated method for mapping robot demonstration videos to simulation environments in physics engines, shown in Figure 2. Each robot demonstration is typically annotated with a language task description and per frame robot joint angle trajectories. We also know the robot used and assume access to its URDF file. Our method extracts five key elements from the demonstration video: (1) the camera's 6-DoF pose relative to the robot body frame, (2) 3D mesh reconstructions of task-relevant objects, their orientations, sizes, and material properties, (3) a scene depth map, (4) a clean background image, (5) proportional—derivative control gains. Together, these components enable realistic, physics-consistent simulations derived directly from video, without requiring manual calibration or curated annotations beyond standard robot joint trajectories.

Automated Robot-Camera Calibration through Differentiable Robot Rendering Robot demonstration videos are typically uncalibrated. We therefore estimate the camera-to-robot transformation using an analysis-by-synthesis approach (Figure 3). Specifically, we construct a joint angle–conditioned 3D Gaussian model of the robot via differentiable rendering in simulation, based on its URDF file following DR-Robot (Liu et al., 2024), shown in Figure 3 Step 1.

Given an RGB video with per-frame joint angles, we render the Gaussian robot model and optimize the camera's 3D translation and orientation to minimize a composite alignment loss with three terms:

(i) an RGB loss penalizing pixel-level appearance differences, (ii) a flow loss enforcing consistency between rendered motion fields and optical flow from the video Karaev et al. (2024), and (iii) a feature loss aligning DI-NOv2 embeddings between rendered and observed frames, following differentiable rendering practices (Chu et al., 2024) (Figure 3 Step 2). When calibration metadata is available (e.g., from DROID (Khazatsky et al., 2024a)), it is used for initialization; otherwise, we perform a coarse flow-based grid search to provide a robust starting point, as in BridgeV2 (Walke et al., 2023a) and RoboMind (Wu et al., 2024). More details can be found in Appendix B.

Object and Scene 3D Reconstruction, Completion, and Physics Estimation We begin by prompting Gemini Team et al. (2023) to segment the robot and all task-relevant foreground

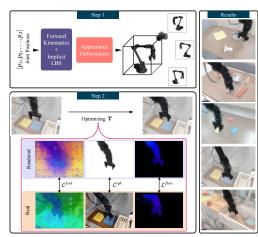


Figure 3: Automated robot-camera calibration through differentiable rendering of poseconditioned 3D robot Gaussians.

objects (Appendix C.1). Each segmented crop is super-resolved with InvSR Yue et al. (2024) and converted into a textured 3D asset using Hunyuan-3D (Team, 2025), which, like most image-to-3D mesh generation models, reconstructs objects in a canonical frame. To recover each object's correct 3D pose, we render synthetic views of the reconstructed mesh and match them to 2D object crops from the original image using the correspondence model MINIMA (Ren et al., 2024). The view with the most feature matches is selected, and these correspondences are lifted into 3D using monocular depth estimates for the real image (Wang et al., 2024) and simulated depth for the rendered view. The final pose is then solved via singular value decomposition (SVD) on the resulting 3D–3D correspondences (Appendix C.2). Physical and material properties are inferred by prompting Gemini and incorporated into the simulation to ensure realistic interactions. To complete the scene, we generate a static background by inpainting the robot and object regions in the first video frame with LaMa Suvorov et al. (2021), producing a clean backdrop for the reconstructed assets (Appendix D). Finally, to accurately reproduce robot dynamics, we perform system identification to tune the proportional derivative (PD) controller gains, aligning simulated end-effector trajectories with those observed in real demonstrations (Appendix E). Apart from standard joint trajectory annotations, our method requires no additional human supervision. More details on our translation method can be found in the Appendix.

Controllable Domain Perturbations We introduce controlled perturbations to the generated environments in order to stress-test policy generalization under changes in background, scene arrangement, viewpoint, and color. Specifically, we consider:

- Background Change (ΔBG): Replaces the original scene background with different inpainted textures drawn from a diverse background dataset, isolating the policy's dependence on contextual appearance cues (Appendix G.1).
- Color Shift (ΔColor): Alters the RGB channel configuration of the scene (e.g., converting RGB to BGR), applied at intensities from 0% to 100% in increments of approximately 33%, to test robustness against low-level color variation (Appendix G.2).
- Object Pose Change (ΔObjPose): Randomly permutes the location of objects in the scene (Appendix G.3).

4 EVALUATING ROBOT TRAJECTORIES WITH HUMANS AND VLMS

In some environments, robot policies may fail entirely, making pairwise preference comparisons noisy or even meaningless. To address this, we first use vision-language models (VLMs) to assign per-frame task progress scores in each execution video. Only videos with non-zero progress scores are then forwarded for human preference evaluation.

4.1 EVALUATING ROBOT TASK PROGRESS SCORING WITH VLMS

Our goal is to automate success detection and progress scoring, allowing evaluation to scale across many environments without humancrafted success detectors. We thus assess task progress in execution videos using prompting techniques for vision-language models (VLMs) (Ma et al., 2024), without relying on privileged simulator state. Specifically, a VLM is prompted with a shuffled sequence of video frames, augmented with the initial frame as a zero-progress reference, and asked to assign progress scores (shown in Figure 4). Shuffling prevents the model from exploiting temporal order, forcing evaluation based purely on visual cues. This approach is effective for both successful and failed trajectories, and we observed

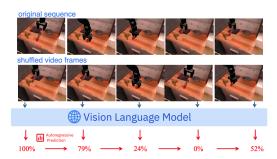


Figure 4: We obtain task progress scores for execution videos automatically by prompting Gemini with a shuffled frame sequence and the language instruction, following Ma et al. (2024).

no improvement from few-shot prompting compared to zero-shot. From the per-frame scores, we considered the mean score over the full execution trajectory, the mean score over the final 30% of frames, and the mean score over the highest scored 30% of frames. Among these, we found *the mean score averaged over the final 30% of frames to align best with human-annotated task progress scores*, and retain it for subsequent analysis.

4.2 EVALUATING ROBOT PERFORMANCE WITH HUMAN PREFERENCE FEEDBACK

While automated scoring provides scalability, human preference feedback remains essential for capturing nuanced aspects of robot behavior that numerical metrics may overlook. Our human evaluations are conducted through pairwise, double-blind comparisons of two policy execution videos drawn from the same simulation environment, under identical initial conditions and task instructions, following the protocol of Chiang et al. (2024). For each comparison, evaluators provide two forms of feedback: (i) a preference label specifying which policy performed better overall or whether they are tied, and (ii) a free-form natural language explanation describing the rationale behind their choice. We use free-form explanations as a way to increase evaluator engagement and attention. In sanity-check experiments, we found that requiring written justification led to more accurate human annotations. We have included the web interface for pairwise comparisons in the Appendix J.

4.2.1 GLOBAL RANKING FROM PAIRWISE PREFERENCES

In this section, we describe our algorithm for computing a global policy ranking from the pairwise feedback provided by evaluators. Formally, let the set of N policies be $\Pi = \{\pi_1, \dots, \pi_N\}$ and let the dataset of pairwise preferences be $\mathcal{D}_p = \{(P_{\pi_A,\pi_B},t)\}$, where $P_{\pi_A,\pi_B} \in \{-1,0,1\}$ indicates a preference for π_A over π_B or a tie, and t identifies the task on which the A/B evaluation was run (e.g., a specific scene and language instruction). Our goal is to compute a global ranking $\mathcal R$ over policies, e.g., $\pi_i \succ \pi_j \succ \cdots \succ \pi_k$.

We will use the Bradley–Terry (BT) model Bradley & Terry (1952) which is a widely used probabilistic framework for ranking items based on pairwise comparisons, modeling the probability that item i is preferred over item j as $P(i \succ j) = \frac{\theta_i}{\theta_i + \theta_j}$, where θ_i and θ_j are latent ability scores. In our benchmark, comparisons may result in ties, which the classical BT formulation cannot handle. Extensions such as the Davidson model Davidson (1970) incorporate ties by introducing an additional tie parameter ν that accounts for the likelihood of indifference between two items. In this formulation, the probabilities are defined as:

$$P(i \succ j) = \frac{\theta_i}{\theta_i + \theta_j + \nu \sqrt{\theta_i \theta_j}}, \quad P(i \equiv j) = \frac{\nu \sqrt{\theta_i \theta_j}}{\theta_i + \theta_j + \nu \sqrt{\theta_i \theta_j}}.$$

This modification enables the model to capture not only strict preferences but also situations where two items are judged equally good, making it more robust for human preference aggregation in settings such as robot evaluation or language model comparison.

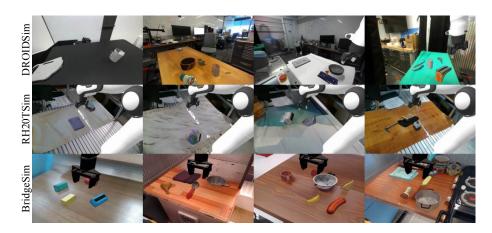


Figure 5: We show simulation environments in RobotArena ∞ seeded from videos demonstrations in the datasets of Bridge, RH20T and DROID.

Estimation. The parameters $\theta = \{\theta_1, \dots, \theta_N\}$ (and the tie parameter ν) are estimated by maximizing the likelihood of the observed preferences. For the Davidson extension, the log-likelihood is:

$$\mathcal{L}(\theta, \nu) = \sum_{(i,j) \in \mathcal{D}_p} \Big[\mathbf{1}[P_{i,j} = 1] \log P(i \succ j) + \mathbf{1}[P_{i,j} = -1] \log P(j \succ i) + \mathbf{1}[P_{i,j} = 0] \log P(i \equiv j) \Big].$$

This objective is concave in the natural parameterization $\log \theta_i$, and can be efficiently optimized using iterative methods such as minorization–maximization or gradient ascent. The resulting θ_i values provide an interpretable ability score for each policy, from which a global ranking \mathcal{R} can be obtained by simple ordering.

5 Benchmarking Robot Policies in RobotArena ∞

Generated Simulation Arenas Our benchmark spans a diverse set of environments and tasks derived from both existing datasets and user-provided scenes (shown in Figure 5):

- 1. *BridgeSim* contains simulated environments and tasks generated from robot demonstrations in the BridgeV2 dataset Walke et al. (2023b), a widely used subset of the Open X-Embodiment (OXE) dataset Open X-Embodiment Collaboration et al. (2023), frequently employed to pretrain generalist robot policies.
- 2. *DROIDSim* consists of environments and tasks created from demonstrations in the DROID dataset Khazatsky et al. (2024b). Unlike BridgeV2, DROID is often excluded from pretraining pipelines for generalist policies due to its higher noise levels Ma et al. (2024).
- 3. *Rh20TSim* includes environments and tasks derived from the RH20T dataset Fang et al. (2023). Notably, among the candidate policies we evaluate, only *SpatialVLA* has been trained on this dataset.

Candidate Policies to Evaluate We benchmark the following open-source generalist robot policies. All policies operate with a fixed egocentric camera and do not make use of wrist-mounted cameras:

- 1. *Octo* Octo Model Team et al. (2024) is a transformer-based manipulation policy pre-trained on 800k demonstrations from the OXE dataset Open X-Embodiment Collaboration et al. (2023). For our experiments, we evaluate the 93M-parameter Octo-Base model.
- 2. *RoboVLM* Li et al. (2024) extends vision-language models (VLMs) into vision-language-action (VLA) policies by adding a continuous action prediction head. We evaluate the variant built on the KosMos backbone Peng et al. (2023).
- SpatialVLA Qu et al. (2025) augments standard VLAs with 3D spatial reasoning through Ego3D Position Encodings and Adaptive Action Grids. It is trained on 1.1 million real-world robot demo demonstrations.

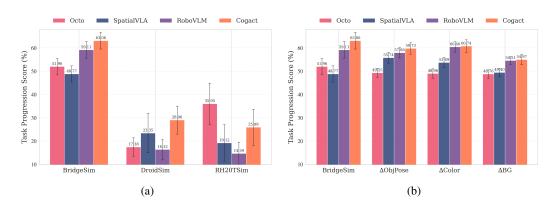


Figure 6: Task completion policy evaluation results obtained automatically from VLMs, (a) in RobotArena ∞ environments and (b) in perturbations of *BridgeSim* environments.

4. CogAct Li et al. (2024) combines a 7B-parameter VLM backbone with a diffusion transformer for action prediction. It is pre-trained on the large-scale OXE dataset (22.5M frames, 60 datasets, 22 robot embodiments) and further fine-tuned on smaller real-world datasets for embodiment-specific adaptation.

We show task VLM derived progress evaluation results of robot policies in Figure 6 *left* with Standard Error of the Mean (SEM). We show the performance in perturbed environments in Figure 6 *right* with SEM for the *BridgeSim* for which the polices perform the best. We draw the following conclusions: (1) **Cross-dataset generalization is weak:** Policies perform substantially worse on environments derived from datasets they were not trained on (e.g., DROID and RH20T). (2) **Model choice matters:** RoboVLM and CogACT consistently outperform Octo-Base and SpatialVLA. (3) **3D structure helps:** SpatialVLA shows improved robustness under object position perturbations, suggesting that explicit modeling of 3D spatial relationships enhances generalization and reduces reliance on fixed asset layouts.

(4) **Backbone strength drives robustness:** Policies with stronger VLM backbones (CogACT, RoboVLM) are more resilient to color perturbations, relying on structural rather than superficial appearance cues. (5) **Background matters most:** Performance across all policies degrades sharply when backgrounds are replaced (e.g., sampling alternatives from BridgeV2), revealing that VLAs often overfit to scene-level correlations between background cues and spatial priors.

We show BT rankings of the models from human preferences in Figure 7 with the chi-squared interval. We note that human preferences align with the VLM task success score; RoboVLM and CogACT are much more preferred than Octo and SpatialVLA.

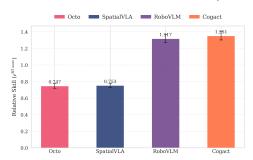


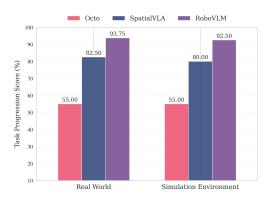
Figure 7: BT scores for human preferences of the models.

5.1 Policy ranking in RobotArena ∞ versus the real world

In Figure 8, we further test the sim-real evaluation alignment. Specifically, we tested the three VLAs, i.e., RoboVLMs, Octo, and SpatialVLA, on the real-world task "Put the carrot in the plate" as presented by Qu et al. (2025). We recreated the same scenes in simulation and deployed the same policies. RoboVLMs and SpatialVLA succeeded in both real and simulated settings, while Octo failed in both, consistently attempting but failing to grasp the carrot.

5.2 ROBOTARENA ∞ VERSUS SIMPLER OF LI ET AL. (2024)

In Figure 9, we compare policy performance on our reproduction of the four SIMPLER environments with our BridgeSim benchmark, which includes 70 environments derived from the same Bridge V2 dataset. *All evaluated VLA models achieve much higher scores on SIMPLER than BridgeSim*,



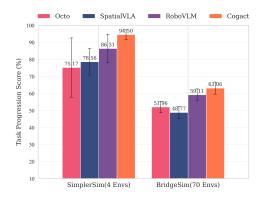


Figure 8: Validation of Simulation-Based Robot Evaluation Against Real-World Robot Evaluations.

Figure 9: Task completion policy evaluation results in RobotArena ∞ versus SIMPLER benchmark of Li et al. (2024).

indicating that the benchmark is nearing saturation—the strongest model nearly solves all tasks. While relative model rankings remain consistent across the two benchmarks, SIMPLER's limited difficulty makes it less suitable for distinguishing future, more advanced models.

Please find additional video results in our project's website at https://submitsquirtel.github.io/.

6 LIMITATIONS / FUTURE DIRECTIONS

By leveraging recent advances in reality-to-simulation translation and crowdsourced evaluation, RobotArena ∞ provides the most scalable and extensible robot benchmark to date. Our evaluators are not domain experts or robotics researchers, but everyday end-users—the very audience that robots are ultimately intended to serve.

That said, the current benchmark has limitations. First, the policies evaluated do not yet incorporate wrist-camera inputs, which restricts the fidelity of certain manipulations. We are actively extending our reality-to-simulation pipeline to generate complete 3D interactive environments that will support richer observations. Second, current simulators still struggle to model fine-grained contact dynamics, such as inserting a charger into a socket. Despite progress in physics engines and automated asset generation Wang et al. (2025); Narang et al. (2022), these tasks remain difficult to reproduce faithfully, highlighting a crucial direction for future research.

Looking forward, RobotArena ∞ is well positioned to benefit from advances in simulation, physics engines, and environment generation, and to serve as a continually improving platform for evaluating the next generation of robotic foundation models.

7 Conclusion

In this work, we introduced RobotArena ∞ , a new benchmarking framework that scales robot evaluation through automated reality-to-simulation translation and online human preference feedback. By coupling advances in vision–language models, 2D-to-3D generative models, and differentiable rendering, our framework enables the automatic construction of rich simulated environments directly from real-world videos. We demonstrated its utility by evaluating multiple VLAs across hundreds of environments and over 7000 human preference judgments, yielding the most extensive study of policy robustness and generalization to date. Our findings highlight significant cross-dataset failures, sensitivity to perturbations, and systematic differences across architectures—revealing both the limitations of current VLAs and the promise of scalable evaluation protocols.

Looking ahead, we envision RobotArena ∞ as a continually evolving benchmark that grows alongside advances in robot learning. Future directions include expanding the range of tasks, incorporating more diverse real-world data sources, and leveraging improvements in physics engines and real-to-sim translation. By releasing both the benchmark environments and evaluation code, we aim to provide the community with an open, extensible platform for rigorous policy evaluation.

REFERENCES

- Jad Abou-Chakra, Lingfeng Sun, Krishan Rana, Brandon May, Karl Schmeckpeper, Maria Vittoria Minniti, and Laura Herlant. Real-is-sim: Bridging the sim-to-real gap with a dynamic digital twin for real-world robot policy evaluation, 2025. URL https://arxiv.org/abs/2504.03597.
- Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv* preprint arXiv:2305.10403, 2023.
- Pranav Atreya, Karl Pertsch, Tony Lee, Moo Jin Kim, Arhan Jain, Artur Kuramshin, Clemens Eppner, Cyrus Neary, Edward Hu, Fabio Ramos, Jonathan Tremblay, Kanav Arora, Kirsty Ellis, Luca Macesanu, Matthew Leonard, Meedeum Cho, Ozgur Aslan, Shivin Dass, Jie Wang, Xingfang Yuan, Xuning Yang, Abhishek Gupta, Dinesh Jayaraman, Glen Berseth, Kostas Daniilidis, Roberto Martin-Martin, Youngwoon Lee, Percy Liang, Chelsea Finn, and Sergey Levine. Roboarena: Distributed real-world evaluation of generalist robot policies, 2025. URL https://arxiv.org/abs/2506.18123.
- Genesis Authors. Genesis: A generative and universal physics engine for robotics and beyond, December 2024. URL https://github.com/Genesis-Embodied-AI/Genesis.
- Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. In 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 4788–4795. IEEE, 2024.
- Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π 0: A vision-language-action flow model for general robot control, 2024. *arXiv preprint arXiv:2410.24164*, 2024.
- Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39:324, 1952. URL https://api.semanticscholar.org/CorpusID:125209808.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M. Dollar. The yeb object and model set: Towards common benchmarks for manipulation research. In *2015 International Conference on Advanced Robotics (ICAR)*, pp. 510–517, 2015. doi: 10.1109/ICAR. 2015.7251504.
- Constantinos Chamzas, Carlos Quintero-Pena, Zachary Kingston, Andreas Orthey, Daniel Rakita, Michael Gleicher, Marc Toussaint, and Lydia E. Kavraki. Motionbenchmaker: A tool to generate and benchmark motion planning datasets. *IEEE Robotics and Automation Letters*, 7(2):882–889, April 2022. ISSN 2377-3774. doi: 10.1109/lra.2021.3133603. URL http://dx.doi.org/10.1109/LRA.2021.3133603.
- 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, 2024. URL https://arxiv.org/abs/2303.04137.
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating llms by human preference, 2024. URL https://arxiv.org/abs/2403.04132.
- Wen-Hsuan Chu, Lei Ke, and Katerina Fragkiadaki. Dreamscene4d: Dynamic multi-object scene generation from monocular videos. *arXiv preprint arXiv:2405.02280*, 2024.

Nikolaus Correll, Kostas E Bekris, Dmitry Berenson, Oliver Brock, Albert Causo, Kris Hauser, Kei Okada, Alberto Rodriguez, Joseph M Romano, and Peter R Wurman. Analysis and observations from the first amazon picking challenge. *IEEE Transactions on Automation Science and Engineering*, 15(1):172–188, 2016.

- Roger R. Davidson. On extending the bradley-terry model to accommodate ties in paired comparison experiments. *Journal of the American Statistical Association*, 65:317–328, 1970. URL https://api.semanticscholar.org/CorpusID:121759206.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Ria Doshi, Homer Walke, Oier Mees, Sudeep Dasari, and Sergey Levine. Scaling cross-embodied learning: One policy for manipulation, navigation, locomotion and aviation. *arXiv* preprint *arXiv*:2408.11812, 2024.
- Kiana Ehsani, Tanmay Gupta, Rose Hendrix, Jordi Salvador, Luca Weihs, Kuo-Hao Zeng, Kunal Pratap Singh, Yejin Kim, Winson Han, Alvaro Herrasti, et al. Spoc: Imitating shortest paths in simulation enables effective navigation and manipulation in the real world. *arXiv preprint arXiv:2312.02976*, 2023.
- Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11441–11450, 2020. doi: 10.1109/CVPR42600.2020.01146.
- Hao-Shu Fang, Hongjie Fang, Zhenyu Tang, Jirong Liu, Chenxi Wang, Junbo Wang, Haoyi Zhu, and Cewu Lu. Rh20t: A comprehensive robotic dataset for learning diverse skills in one-shot, 2023. URL https://arxiv.org/abs/2307.00595.
- Letian Fu, Huang Huang, Gaurav Datta, Lawrence Yunliang Chen, William Chung-Ho Panitch, Fangchen Liu, Hui Li, and Ken Goldberg. In-context imitation learning via next-token prediction. *arXiv* preprint arXiv:2408.15980, 2024.
- Markus Grotz, Mohit Shridhar, Tamim Asfour, and Dieter Fox. Peract2: A perceiver actor framework for bimanual manipulation tasks. *arXiv preprint arXiv:2407.00278*, 2024.
- Minho Heo, Youngwoon Lee, Doohyun Lee, and Joseph J. Lim. Furniturebench: Reproducible real-world benchmark for long-horizon complex manipulation, 2023. URL https://arxiv.org/abs/2305.12821.
- 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.
- Dmitry Kalashnikov, Jacob Varley, Yevgen Chebotar, Benjamin Swanson, Rico Jonschkowski, Chelsea Finn, Sergey Levine, and Karol Hausman. Mt-opt: Continuous multi-task robotic reinforcement learning at scale. *arXiv* preprint arXiv:2104.08212, 2021.
- Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker3: Simpler and Better Point Tracking by Pseudo-Labelling Real Videos. *arXiv e-prints*, art. arXiv:2410.11831, October 2024. doi: 10.48550/arXiv.2410.11831.
- 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*, 2024a.
- 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*, 2024b.

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. OpenVLA: An Open-Source Vision-Language-Action Model. *arXiv e-prints*, art. arXiv:2406.09246, June 2024. doi: 10.48550/arXiv.2406.09246.

- Kenneth Kimble, Karl Van Wyk, Joe Falco, Elena Messina, Yu Sun, Mizuho Shibata, Wataru Uemura, and Yasuyoshi Yokokohji. Benchmarking protocols for evaluating small parts robotic assembly systems. *IEEE Robotics and Automation Letters*, 5(2):883–889, 2020. doi: 10.1109/LRA.2020. 2965869.
- Youngwoon Lee, Edward S. Hu, Zhengyu Yang, Alex Yin, and Joseph J. Lim. Ikea furniture assembly environment for long-horizon complex manipulation tasks, 2019. URL https://arxiv.org/abs/1911.07246.
- Jürgen Leitner, Adam W. Tow, Jake E. Dean, Niko Suenderhauf, Joseph W. Durham, Matthew Cooper, Markus Eich, Christopher Lehnert, Ruben Mangels, Christopher McCool, Peter Kujala, Lachlan Nicholson, Trung Pham, James Sergeant, Liao Wu, Fangyi Zhang, Ben Upcroft, and Peter Corke. The acrv picking benchmark (apb): A robotic shelf picking benchmark to foster reproducible research, 2016. URL https://arxiv.org/abs/1609.05258.
- Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Wensi Ai, Benjamin Martinez, et al. Behavior-1k: A human-centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation. arXiv preprint arXiv:2403.09227, 2024.
- Xinghang Li, Peiyan Li, Minghuan Liu, Dong Wang, Jirong Liu, Bingyi Kang, Xiao Ma, Tao Kong, Hanbo Zhang, and Huaping Liu. Towards Generalist Robot Policies: What Matters in Building Vision-Language-Action Models. *arXiv e-prints*, art. arXiv:2412.14058, December 2024. doi: 10.48550/arXiv.2412.14058.
- Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, et al. Evaluating real-world robot manipulation policies in simulation. *arXiv preprint arXiv:2405.05941*, 2024.
- 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, 2024.
- Ruoshi Liu, Alper Canberk, Shuran Song, and Carl Vondrick. Differentiable Robot Rendering. *arXiv e-prints*, art. arXiv:2410.13851, October 2024. doi: 10.48550/arXiv.2410.13851.
- 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.
- Yecheng Jason Ma, Joey Hejna, Ayzaan Wahid, Chuyuan Fu, Dhruv Shah, Jacky Liang, Zhuo Xu, Sean Kirmani, Peng Xu, Danny Driess, Ted Xiao, Jonathan Tompson, Osbert Bastani, Dinesh Jayaraman, Wenhao Yu, Tingnan Zhang, Dorsa Sadigh, and Fei Xia. Vision language models are in-context value learners, 2024. URL https://arxiv.org/abs/2411.04549.
- Mauricio Matamoros, SEIB Viktor, and Dietrich Paulus. Trends, challenges and adopted strategies in robocup@ home. In 2019 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC), pp. 1–6. IEEE, 2019.
- 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* (RA-L), 7(3):7327–7334, 2022.
- Mark Moll, Ioan Alexandru Sucan, and Lydia E. Kavraki. Benchmarking motion planning algorithms: An extensible infrastructure for analysis and visualization. *IEEE Robotics & Automation Magazine*, 22:96–102, 2014. URL https://api.semanticscholar.org/CorpusID:1791284.

Yashraj Narang, Kier Storey, Iretiayo Akinola, Miles Macklin, Philipp Reist, Lukasz Wawrzyniak, Yunrong Guo, Adam Moravanszky, Gavriel State, Michelle Lu, Ankur Handa, and Dieter Fox. Factory: Fast contact for robotic assembly, 2022. URL https://arxiv.org/abs/2205.03532.

648

649

650

651

652

653

654

655

656 657

658

659

660

661 662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

683

684

685

686

687

688

689

690

691

692

693

694

696

697

699

700

701

Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. https://octo-models.github.io, 2023.

Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An Open-Source Generalist Robot Policy. *arXiv e-prints*, art. arXiv:2405.12213, May 2024. doi: 10.48550/arXiv.2405.12213.

Open X-Embodiment Collaboration, Abby O'Neill, Abdul Rehman, Abhinav Gupta, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, Albert Tung, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anchit Gupta, Andrew Wang, Andrey Kolobov, Anikait Singh, Animesh Garg, Aniruddha Kembhavi, Annie Xie, Anthony Brohan, Antonin Raffin, Archit Sharma, Arefeh Yavary, Arhan Jain, Ashwin Balakrishna, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Blake Wulfe, Brian Ichter, Cewu Lu, Charles Xu, Charlotte Le, Chelsea Finn, Chen Wang, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Christopher Agia, Chuer Pan, Chuyuan Fu, Coline Devin, Danfei Xu, Daniel Morton, Danny Driess, Daphne Chen, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dinesh Jayaraman, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Ethan Foster, Fangchen Liu, Federico Ceola, Fei Xia, Feiyu Zhao, Felipe Vieira Frujeri, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Gilbert Feng, Giulio Schiavi, Glen Berseth, Gregory Kahn, Guangwen Yang, Guanzhi Wang, Hao Su, Hao-Shu Fang, Haochen Shi, Henghui Bao, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homanga Bharadhwaj, Homer Walke, Hongjie Fang, Huy Ha, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jad Abou-Chakra, Jaehyung Kim, Jaimyn Drake, Jan Peters, Jan Schneider, Jasmine Hsu, Jay Vakil, Jeannette Bohg, Jeffrey Bingham, Jeffrey Wu, Jensen Gao, Jiaheng Hu, Jiajun Wu, Jialin Wu, Jiankai Sun, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jimmy Wu, Jingpei Lu, Jingyun Yang, Jitendra Malik, João Silvério, Joey Hejna, Jonathan Booher, Jonathan Tompson, Jonathan Yang, Jordi Salvador, Joseph J. Lim, Junhyek Han, Kaiyuan Wang, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Black, Kevin Lin, Kevin Zhang, Kiana Ehsani, Kiran Lekkala, Kirsty Ellis, Krishan Rana, Krishnan Srinivasan, Kuan Fang, Kunal Pratap Singh, Kuo-Hao Zeng, Kyle Hatch, Kyle Hsu, Laurent Itti, Lawrence Yunliang Chen, Lerrel Pinto, Li Fei-Fei, Liam Tan, Linxi "Jim" Fan, Lionel Ott, Lisa Lee, Luca Weihs, Magnum Chen, Marion Lepert, Marius Memmel, Masayoshi Tomizuka, Masha Itkina, Mateo Guaman Castro, Max Spero, Maximilian Du, Michael Ahn, Michael C. Yip, Mingtong Zhang, Mingyu Ding, Minho Heo, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Suenderhauf, Ning Liu, Norman Di Palo, Nur Muhammad Mahi Shafiullah, Oier Mees, Oliver Kroemer, Osbert Bastani, Pannag R Sanketi, Patrick "Tree" Miller, Patrick Yin, Paul Wohlhart, Peng Xu, Peter David Fagan, Peter Mitrano, Pierre Sermanet, Pieter Abbeel, Priya Sundaresan, Qiuyu Chen, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Mart'in-Mart'in, and Rohan Baijal. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. arXiv e-prints, art. arXiv:2310.08864, October 2023. doi: 10.48550/arXiv.2310.08864.

Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 6892–6903. IEEE, 2024.

Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding Multimodal Large Language Models to the World. *arXiv e-prints*, art. arXiv:2306.14824, June 2023. doi: 10.48550/arXiv.2306.14824.

Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.

- Wilbert Pumacay, Ishika Singh, Jiafei Duan, Ranjay Krishna, Jesse Thomason, and Dieter Fox. THE COLOSSEUM: A Benchmark for Evaluating Generalization for Robotic Manipulation. *arXiv e-prints*, art. arXiv:2402.08191, February 2024. doi: 10.48550/arXiv.2402.08191.
- Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, and Xuelong Li. SpatialVLA: Exploring Spatial Representations for Visual-Language-Action Model. *arXiv e-prints*, art. arXiv:2501.15830, January 2025. doi: 10.48550/arXiv.2501.15830.
- Jiangwei Ren, Xingyu Jiang, Zizhuo Li, Dingkang Liang, Xin Zhou, and Xiang Bai. MINIMA: Modality Invariant Image Matching. *arXiv e-prints*, art. arXiv:2412.19412, December 2024. doi: 10.48550/arXiv.2412.19412.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022.
- Dhruv Shah, Ajay Sridhar, Arjun Bhorkar, Noriaki Hirose, and Sergey Levine. Gnm: A general navigation model to drive any robot. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pp. 7226–7233. IEEE, 2023.
- Ajay Sridhar, Dhruv Shah, Catherine Glossop, and Sergey Levine. NoMaD: Goal Masked Diffusion Policies for Navigation and Exploration. *arXiv pre-print*, 2023. URL https://arxiv.org/abs/2310.07896.
- Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. *arXiv preprint arXiv:2109.07161*, 2021.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Tencent Hunyuan3D Team. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation, 2025.
- Joseph A. Vincent, Haruki Nishimura, Masha Itkina, Paarth Shah, Mac Schwager, and Thomas Kollar. How generalizable is my behavior cloning policy? a statistical approach to trustworthy performance evaluation, 2024. URL https://arxiv.org/abs/2405.05439.
- Homer Walke, Kevin Black, Abraham Lee, Moo Jin Kim, Max Du, Chongyi Zheng, Tony Zhao, Philippe Hansen-Estruch, Quan Vuong, Andre He, Vivek Myers, Kuan Fang, Chelsea Finn, and Sergey Levine. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning (CoRL)*, 2023a.
- Homer Walke, Kevin Black, Abraham Lee, Moo Jin Kim, Max Du, Chongyi Zheng, Tony Zhao, Philippe Hansen-Estruch, Quan Vuong, Andre He, Vivek Myers, Kuan Fang, Chelsea Finn, and Sergey Levine. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning (CoRL)*, 2023b.
- Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. MoGe: Unlocking Accurate Monocular Geometry Estimation for Open-Domain Images with Optimal Training Supervision. *arXiv e-prints*, art. arXiv:2410.19115, October 2024. doi: 10. 48550/arXiv.2410.19115.
- Yian Wang, Bingjie Tang, Chuang Gan, Dieter Fox, Kaichun Mo, Yashraj Narang, and Iretiayo Akinola. Matchmaker: Automated asset generation for robotic assembly, 2025. URL https://arxiv.org/abs/2503.05887.

Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. *arXiv preprint arXiv:2412.13877*, 2024.

- Youguang Xing, Xu Luo, Junlin Xie, Lianli Gao, Hengtao Shen, and Jingkuan Song. Shortcut learning in generalist robot policies: The role of dataset diversity and fragmentation, 2025. URL https://arxiv.org/abs/2508.06426.
- Jonathan Yang, Catherine Glossop, Arjun Bhorkar, Dhruv Shah, Quan Vuong, Chelsea Finn, Dorsa Sadigh, and Sergey Levine. Pushing the limits of cross-embodiment learning for manipulation and navigation. *arXiv preprint arXiv:2402.19432*, 2024.
- Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Sejune Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, et al. Latent action pretraining from videos. *arXiv* preprint arXiv:2410.11758, 2024.
- Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, et al. Homerobot: Open-vocabulary mobile manipulation. *arXiv preprint arXiv:2306.11565*, 2023.
- 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*, pp. 1094–1100. PMLR, 2020.
- Zongsheng Yue, Kang Liao, and Chen Change Loy. Arbitrary-steps image super-resolution via diffusion inversion. *arXiv preprint arXiv:2412.09013*, 2024.
- Zhiyuan Zhou, Pranav Atreya, Abraham Lee, Homer Walke, Oier Mees, and Sergey Levine. Autonomous improvement of instruction following skills via foundation models. *arXiv preprint arXiv:2407.20635*, 2024.
- Zhiyuan Zhou, Pranav Atreya, You Liang Tan, Karl Pertsch, and Sergey Levine. AutoEval: Autonomous Evaluation of Generalist Robot Manipulation Policies in the Real World. *arXiv e-prints*, art. arXiv:2503.24278, March 2025. doi: 10.48550/arXiv.2503.24278.

A WEBSITE AND CODE

The project website is available at https://submitsquirtel.github.io/, and the source code can be accessed at https://github.com/submitsquirtel/ICLR-paper-13938.

B INFERRING CAMERA-ROBOT POSE TRANSFORMATION VIA ANALYSIS-BY-SYNTHESIS

In this optimization problem, we work with a sequence of video frames capturing the robot's motion, denoted I_t^{gt} , where t ranges from 1 to T. We also have the robot's joint angles q^t at each time t, provided by the dataset. Our task is to estimate the transformation $\mathbf{T} \in SE(3)$, which aligns the robot's coordinate system with the camera's. Using \mathbf{T} and q^t , we produce rendered a image I_t , a synthetic depiction of the robot at time t. Motion is captured through optical flow F_t^{gt} , estimated between consecutive video frames I_t^{gt} and I_{t+1}^{gt} , and rendered flow F_t . Additionally, we use the DINOv2 model to extract feature maps $\phi_t^{gt} = \mathrm{DINOv2}(I_t^{gt})$ from the video frames and $\phi_t = \mathrm{DINOv2}(I_t)$ from the rendered images, enhancing alignment with high-level image features.

The optimization aligns the rendered and observed data using three loss terms:

 RGB Loss: Measures the color difference between the rendered and video images at each time t:

$$\mathcal{L}_{rgb}(t) = \|I_t - I_t^{gt}\|_2^2 \tag{1}$$

2. Flow Loss: Compares the rendered flow to the optical flow for each pair of frames from t=1 to T-1:

$$\mathcal{L}_{flow}(t) = \|F_t - F_t^{gt}\|_2^2 \tag{2}$$

3. **Feature Loss**: Aligns the feature maps using cosine loss at each time t:

$$\mathcal{L}_{feat}(t) = 1 - \cos\left(\phi_t, \phi_t^{gt}\right) \tag{3}$$

The total loss combines these terms over the sequence:

$$\mathcal{L} = \sum_{t=1}^{T} \lambda_{rgb} \mathcal{L}_{rgb}(t) + \sum_{t=1}^{T} \lambda_{feat} \mathcal{L}_{feat}(t) + \sum_{t=1}^{T-1} \lambda_{flow} \mathcal{L}_{flow}(t)$$
(4)

where λ_{rgb} , λ_{feat} , and λ_{flow} are weights balancing each term. The optimal **T** is found by minimizing:

$$\mathbf{T}^* = \underset{\mathbf{T} \in SE(3)}{\operatorname{argmin}} \mathcal{L} \tag{5}$$

C RELEVANT OBJECT SEGMENTATION, 3D RECONSTRUCTION, AND MATERIAL PROPERTY ESTIMATION

The creation of simulation-ready 3D assets for movable objects is accomplished through a multistage pipeline. This process begins with semantic understanding of the scene using a vision language model (VLM), specifically Gemini Team et al. (2023). We prompt Gemini to segment the robot and all foreground objects it interacts with. The resulting segmentation masks provide fine-grained object localization. To enhance visual details critical for reconstruction, the original image is first super-resolved, then resized to its original dimensions. The segmentation masks are applied to isolate and crop each object of interest. These object-specific image patches are further refined using a second stage of super-resolution with InvSR Yue et al. (2024), producing high-fidelity inputs for 3D reconstruction. The enhanced patches are then processed by an image-to-3D model, Hunyuan-3D Team (2025), which outputs detailed meshes capturing both the geometry and texture of each object. To ensure physical plausibility within simulation, Gemini Team et al. (2023) is prompted to infer object-specific physical parameters such as mass and friction coefficients. Final asset preparation includes appropriate scaling and placement within the simulated environment to ensure consistent alignment with the original scene configuration.

This section describes the method of recovering the 3D position and orientation of objects from a single RGB image. The approach leverages camera parameters recovered by our previously described method, detailed in Section B, and an initial depth map estimated using (MoGE Wang et al. (2024)).

C.1 OBJECT SEGMENTATION

To obtain accurate localization of task-relevant entities, we employ Gemini Team et al. (2023) with three structured prompts tailored for different goals.

Robot Segmentation Prompt We first segment the robot, as jointly detecting all entities may result in missed or inaccurate detection of the robot. The prompt below is used to obtain the robot's segmentation mask:

Give segmentation masks for the robot in the scene. Output a JSON list of segmentation masks where each entry contains the 2D bounding box in "box_2d", a descriptive text label in "label", and the mask in "mask".

Foreground Object Segmentation Prompt After isolating the robot, we proceed to segment the foreground objects it interacts with, along with estimating their physical properties. The following prompt is used to obtain segmentation masks and associated physical attributes:

Give segmentation masks for all important complete foreground objects on the plane which the robot interacts with. You must ignore any background objects, irrelevant surfaces, or any objects that are occluded, covered, or severely blocked by other objects.

Output a JSON list of segmentation masks where each entry contains:

- "box_2d": the 2D bounding box
- "mask": the segmentation mask (in image format)
- "label": a descriptive text label
- "mass": estimated object mass in kilograms (float)
- "friction": estimated friction coefficient (float)
- "surface_type": one of [Glass, Water, Emission, Plastic, Rough, Smooth, Reflective, Metal, Iron, Aluminium, Copper, Gold]

Task-Relevant Object Selection Prompt (conditional) If a task description is provided, we use the following prompt to analyze the previously segmented objects, identifying those with semantic relevance to the task and determining their functional roles (e.g., as a manipulation *target* or *destination*):

Consider this specific task: {task}, given the object names: {object_names}, please select the object names that are relevant to the task and identify their role in the task from "target" and "destination".

Output in this format exactly as: target: <target_object>, destination: <destination_object>. Please try to identify the target object (If you can't find an exact match, use the closest one from given object names) and if there is no destination object, only output the target object.

For example, given the task: put the pot on top of the yellow cloth and the segmented object names: sauce pan, cloth, the output would be: target: sauce pan, destination: cloth

C.2 OBJECT POSE ESTIMATION

Initial 3D Point Cloud Reconstruction With the 2D mask of the target object generated using Gemini Team et al. (2023), we unproject each masked pixel into a 3D point cloud in world coordinates. This process uses the recovered camera intrinsic and extrinsic parameters \mathbf{K} , \mathbf{R} , \mathbf{t} and a metric-scale depth map D(u,v). To recover metric depth maps, we compute the scale factor by comparing MoGE'sWang et al. (2024) relative depth estimates of the robot arm against the simulated robot arm's ground-truth depth; applying this factor converts the relative predictions into accurately scaled metric depth maps.

For each pixel (u, v) within the mask, its world-space coordinate $\mathbf{P}_{\text{world}}$ is then given by

$$\mathbf{P}_{\text{world}} = \mathbf{R}^{\top} D(u, v) \, \mathbf{K}^{-1} \left[u, v, 1 \right]^{\top} - \mathbf{t}.$$

Simulated Viewpoint Generation and 2D Correspondence Establishment To robustly estimate the object's pose, we generate multiple synthetic views. This process begins with the generated 3D mesh of the object. To ensure dimensional consistency for the simulation, the scale of this generated mesh is aligned by utilizing its bounding box and the bounding box of the original point cloud. The scale-aligned mesh is then imported into our simulation environment and initialized at a reference position. In this simulated environment, multiple views are rendered by a camera programmed to navigate around this central object position. The camera's trajectory is systematically defined by varying its elevation and azimuth angles relative to the reference position.

Le

$$z_{\text{levels}} = \{z_1, z_2, \dots, z_L\}, \quad \theta_{\text{values}} = \{\theta_1, \theta_2, \dots, \theta_M\}, \quad N_\theta = M,$$

For each view index $i \in \{0, ..., N-1\}$, we compute:

$$\theta_i = \theta_{\text{values}} [i \mod N_\theta], \tag{6}$$

$$z_i = z_{\text{levels}}[\lfloor i/N_{\theta} \rfloor],$$
 (7)

$$\hat{\mathbf{v}}_i = \frac{1}{\sqrt{\cos^2 \theta_i + \sin^2 \theta_i + z_i^2}} \begin{bmatrix} \cos \theta_i \\ \sin \theta_i \\ z_i \end{bmatrix}, \tag{8}$$

$$\mathbf{p}_i = \mathbf{p}_{\text{ref}} + d\,\hat{\mathbf{v}}_i. \tag{9}$$

where $\mathbf{p}_{\mathrm{ref}}$ is the target's centroid in world coordinates and d>0 is the fixed radial distance. The simulated camera is then oriented such that its viewing direction is aimed at $\mathbf{p}_{\mathrm{ref}}$.

For each camera pose p_i , we render a synthetic image:

$$I_{\text{sim}}^i = \text{Render}(\mathbf{p}_i)$$
,

We then apply the MINIMA algorithm Ren et al. (2024) to establish 2D keypoint correspondences between I_{sim}^i and the original masked image. This produces, for each view i, a set of M_i matched keypoint pairs:

$$C_i = \left\{ \left(\mathbf{k}_{\text{sim}}^{i,j}, \, \mathbf{k}_{\text{orig}}^{i,j} \right) \right\}_{j=1}^{M_i}.$$

where $\mathbf{k}_{\text{sim}}^{i,j}$ is the j-th keypoint in I_{sim}^i and $\mathbf{k}_{\text{orig}}^{i,j}$ is its corresponding keypoint in the original image.

Optimal View Selection and 3D Keypoint Pair Generation We select the optimal view index

$$i^* = \underset{i \in \{0, \dots, N-1\}}{\operatorname{argmax}} |\mathcal{C}_i|,$$

and denote the corresponding rendered image by I_{sim}^* . The set of matched keypoints for this view is

$$C_{i^*} = \left\{ \left(\mathbf{k}_{\text{sim}}^{*,j}, \, \mathbf{k}_{\text{orig}}^{*,j} \right) \right\}_{i=1}^{M},$$

The optimal 2D keypoint correspondences $\{(\mathbf{k}_{\text{sim}}^{*,j},\mathbf{k}_{\text{orig}}^{*,j})\}_{j=1}^{M}$ are then lifted into 3D by unprojection.

Specifically, for each simulated keypoint $\mathbf{k}_{\text{sim}}^{*,j}$ we compute its 3D coordinate via

$$\mathbf{P}_{\text{sim},j} = \text{Unproject}(\mathbf{k}_{\text{sim}}^{*,j}; \mathbf{K}_{\text{sim}}^{*}, \mathbf{R}_{\text{sim}}^{*}, \mathbf{t}_{\text{sim}}^{*}, D_{\text{sim}}^{j}), \tag{10}$$

where $\mathbf{K}_{\text{sim}}^*, \mathbf{R}_{\text{sim}}^*, \mathbf{t}_{\text{sim}}^*$ are the intrinsic and extrinsic parameters of the known simulated camera and D_{sim}^j is the depth rendered in the simulation.

Similarly, for each original image keypoint $\mathbf{k}_{\mathrm{orig}}^{*,j}$ we obtain

$$\mathbf{P}_{\text{orig},j} = \text{Unproject}(\mathbf{k}_{\text{orig}}^{*,j}; \mathbf{K}, \mathbf{R}, \mathbf{t}, D(u_{\text{orig}}^{j}, v_{\text{orig}}^{j})), \tag{11}$$



Figure 10: Background Impainting Results. The top row shows the original RGB images with task-relevant objects present; the bottom row shows the corresponding inpainted images after foreground object removal, where only backgrounds remain.

where \mathbf{K} , \mathbf{R} and \mathbf{t} are the recovered camera intrinsic and extrinsic parameters, and $D(u_{\mathrm{orig}}^j, v_{\mathrm{orig}}^j)$ is the aligned MoGEWang et al. (2024) depth at pixel $(u_{\mathrm{orig}}^j, v_{\mathrm{orig}}^j)$.

This yields a set of M 3D–3D correspondences

$$\mathcal{S} = \left\{ (\mathbf{P}_{\text{sim},j}, \, \mathbf{P}_{\text{orig},j}) \right\}_{j=1}^{M}.$$

Rigid Transformation Estimation using SVD Given the set of M 3D–3D keypoint correspondences $\{(\mathbf{P}_{\text{sim},j},\mathbf{P}_{\text{orig},j})\}_{j=1}^{M}$, our goal is to estimate the rigid transformation (\mathbf{R},\mathbf{t}) that best aligns these two point sets. This transformation consists of a rotation matrix $\mathbf{R} \in SO(3)$ and a translation vector $\mathbf{t} \in \mathbb{R}^{3}$, such that:

$$\mathbf{P}_{\text{orig},j} \approx \mathbf{R} \mathbf{P}_{\text{sim},j} + \mathbf{t}$$
.

We employ the Singular Value Decomposition (SVD) method for this estimation, which yields the rotation ${\bf R}$ and translation ${\bf t}$ that map the 3D keypoints of the simulated object to the original point cloud and thus recover the position and orientation of the generated object.

D BACKGROUND INPAINTING

We generate a static background by inpainting the robot and object regions in the first video frame with LaMa Suvorov et al. (2021). We visualize the results of background inpainting for 4 different scenes in Bridge Dataset Walke et al. (2023b) in Figure 10.

E SYSTEM IDENTIFICATION

To improve the simulation fidelity of the robot's open-loop motion, we apply system identification (SysID) to tune the proportional—derivative (PD) controller gains, \mathbf{kp} and \mathbf{kd} . The goal is to minimize the discrepancy between simulated and real end-effector trajectories, using only the joint position commands as input. For each trajectory in a dataset, we extract the end-effector pose sequence and run a physics-based simulation using the same joint commands. The simulated robot is controlled with a PD controller, and the gains are optimized to minimize the difference in end-effector positions over time.

$$kp^*, kd^* = \underset{kp,kd}{\operatorname{argmin}} \sum_{t=1}^{T} \left(\|\mathbf{x}_t^{gt} - \mathbf{x}_t\|_2 + \arcsin\left(\frac{1}{2\sqrt{2}} \|\mathbf{R}_t^{gt} - \mathbf{R}_t\|_F \right) \right)$$
(12)

where \mathbf{x}_t^{gt} and \mathbf{R}_t^{gt} denote the ground-truth end-effector position and orientation from the real world dataset, \mathbf{x}_t and \mathbf{R}_t denote those of the simulated robot, and $\|\cdot\|_F$ denotes the Frobenius norm.

Note that we exclude the gripper from the system identification process, as its state in the dataset is typically represented by a binary open/close signal, which does not support continuous control tuning.

To search for optimal parameters, we use a gradient-free Simulated Annealing (SA) strategy. In each iteration, a candidate set of gains is evaluated in parallel using the Genesis simulator Authors (2024). All environments are initialized simultaneously, and the same candidate gains are applied to each. At each step:

- Inverse kinematics is solved in batch to produce joint targets for each trajectory.
- PD control is applied in each environment using the current candidate (kp, kd).
- The mean Euclidean distance between the simulated and dataset end-effector positions is computed.
- The SA policy perturbs the parameters and accepts new candidates based on the change in average error.

Gains are optimized over 5000 steps of SA. Each iteration involves simulating all trajectories for a fixed number 5000 of steps (typically), making the process feasible in parallel with GPU acceleration. We use the following parameter ranges for the search: $\mathbf{kp} \in [2000, 15000]$, $\mathbf{kd} \in [10, 2000]$.

A visual comparison for Bridge Dataset Walke et al. (2023b) before and after system identification is shown in Figure 11, where the dataset (red) and simulated (blue) end-effector trajectories are overlaid on the XY plane.

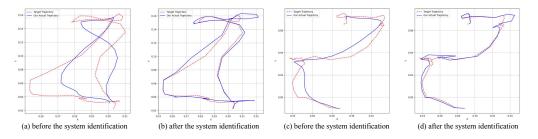


Figure 11: **Inferring robot control gains for matching robot trajectories in reality and simulation.** The end-effector trajectory, projected onto the XY plane, is shown for both the pre-identification (a,c) and post-identification (b,d) stages. After system identification, the simulated trajectory (blue) aligns closely with the recorded dataset trajectory (red), whereas significant deviations are observed prior to identification.

F BridgeSim Environment Visualizations

We present additional visualizations of results from our Real2Sim pipeline in Figure 12.

G ENVIRONMENT PERTURBATIONS

To quantify policy robustness under realistic visual and spatial variability, we systematically introduce four controlled environment perturbations.

G.1 BACKGROUND CHANGE (Δ BG)

For each scene, we replace the original background with five different inpainted backgrounds generated by our Reality-to-Simulation pipeline and keep all foreground objects exactly the same, so that for every scene we have five 'background flipped' variants that differ only in their contextual appearance, as illustrated in Figure 14.

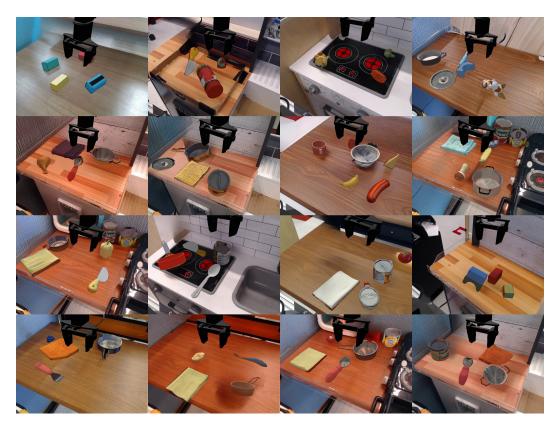


Figure 12: Visualizations of a subset of BridgeSim Environments

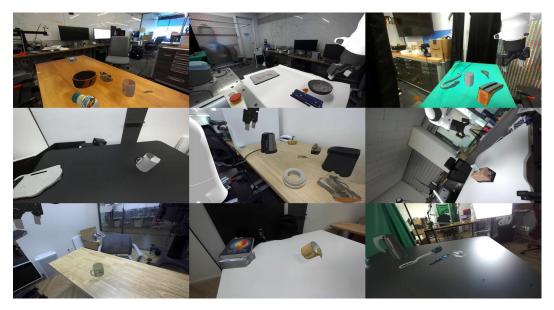


Figure 13: Visualizations of a subset of *DROIDSim* Environments

G.2 Color Shift (Δ Color)

For each scene, we leave the foreground (objects, lighting, dynamics) completely unchanged and apply only a low-level color perturbation to the background by remapping its RGB channels to BGR and blending back at four intensity levels. Concretely, for every background pixel with the original



Figure 14: Background Change Example. The top-left image shows the original image without background perturbations.



Figure 15: Color Shift Example. The leftmost image shows the original image without color perturbations.

color vector [R, G, B], we compute:

$$C'(\alpha) = (1 - \alpha) \begin{bmatrix} R \\ G \\ B \end{bmatrix} + \alpha \begin{bmatrix} B \\ G \\ R \end{bmatrix},$$

where

$$\alpha \in \{0.00, 0.33, 0.66, 1.00\}$$

corresponds respectively to 0%, 33%, 66% and 100% BGR swap intensity. This yields four "color-swapped" background variants per scene, as illustrated in Figure 15.

G.3 Object Pose Change (Δ ObjPose)

For each scene containing N distinct assets with original world-space positions $\{x_i\}_{i=1}^N$, we perform N independent random permutations $\{\pi^{(k)}\}_{k=1}^N$ (including the identity permutation for the original layout). For permutation k, we reassign asset i to the position of asset $\pi^{(k)}(i)$, yielding

$$x_i^{\prime(k)} = x_{\pi^{(k)}(i)}, \quad i = 1, \dots, N.$$

This procedure generates N scene variants that differ only in object arrangement, as illustrated in Figure 16.

H AUTOMATED TASK PROGRESS SCORING WITH VLMS

We present the results for generative value learning based progressive score prediction in Figure 17 18 19.

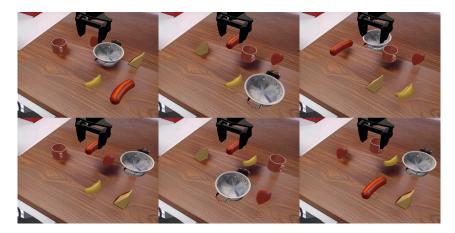


Figure 16: Object Position Perturbation Example. The top-left image shows the original image with base permutation.

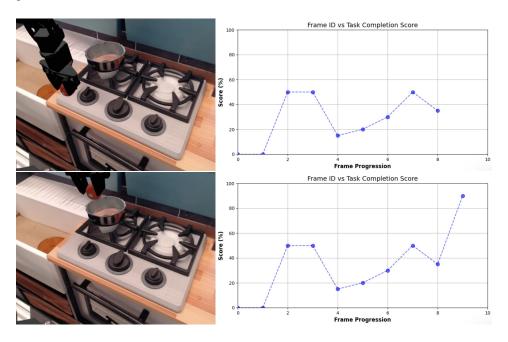


Figure 17: Example VLM-generated task evaluation curves on base environment. **Left panels:** Representative frames sampled at low- and high-progress points. **Right panels:** VLM-assigned completion score (y-axis) across frame ID (x-axis).

I VLM EVALUATIONS

In the *Rh20TSim* environment, Octo achieves a significantly higher VLM score than the other models. Qualitative analysis reveals this is due to a minimalistic motion pattern; Octo just slowly lowers its arm directly onto the workbench. This simple trajectory succeeds by coincidence, as the target object happens to be positioned in the arm's path. This shortcut solution scores higher than the baseline policies, which engage in more extensive but ultimately unsuccessful exploration.

J HUMAN EVALUATIONS

To assess the reliability of our automated task success metrics, we conducted a validation study by comparing them against human judgments collected through a user study. We developed a web-based

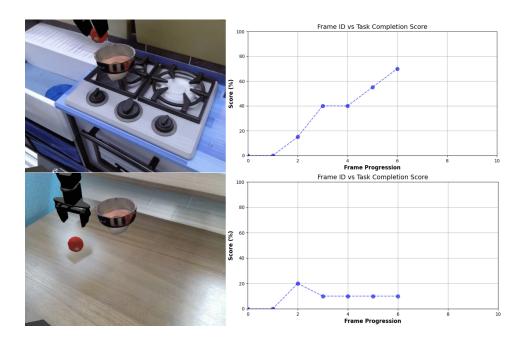


Figure 18: Example VLM-generated task evaluation curves on perturbed environments. **Top:** A high-progress moment immediately after the object lift, for which the VLM predicts a completion score of approximately 70%. **Bottom:** An inconsequential action with no strong effect on task progress, for which the VLM predicts a low completion score of approximately 10%.

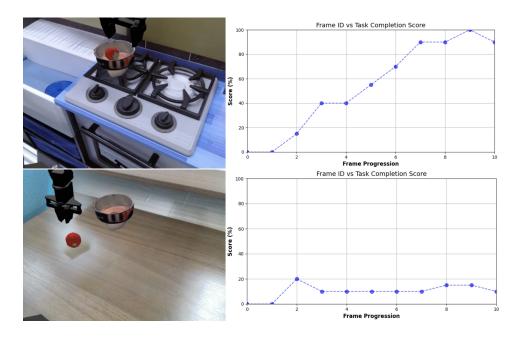


Figure 19: Example VLM-generated task evaluation curves on perturbed environments. **Top:** A successful pick-and-place execution—after the object lift the VLM score climbs steadily and correctly shows task completion. **Bottom:** An unsuccessful trajectory with completion score remaining below 20%, demonstrating the VLM's capacity to detect failure to complete the task.

interface that presented annotators with two side-by-side videos, each showing a different policy attempting the same task described by a natural language instruction. For each pair, participants are first asked to provide descriptions of each robot's attempt and then judge which policy performed better or if it was a tie. An example of this interface is shown in Figure 20. To ensure high data quality, we require participants to pass a qualification quiz by correctly evaluating at least 8 out of 10 video pairs. These pairs were intentionally chosen with clear performance differences, which makes them straightforward to judge.

From a total pool of 70 Environments, we collected 7,104 pairwise comparisons on the Amazon Mechanical Turk (AMT) platform, with each participant assigned a random subset to evaluate. All participants were compensated according to the platform's guidelines, ensuring a diverse pool of evaluators and promoting reliable, unbiased annotations.

Study Overview Thank you for participating in our study! Your feedback will help us understand how different robot policies perform on simulated manipulation tasks. In this study, you will watch short video clips showing two robots (slide-by-side) attempting to complete the same task (e.g., stacking, moving objects). Each robot follows a different control policy. Your Task For each pair of videos, you will first briefly describe what each robot did. After submitting your descriptions, you will choose the robot that you think performed better. If their performance was similar, you may indicate that they performed equally well. Qualification Outz The first 10 video pairs will serve as a qualification quiz. These trials have a known correct answer, and you must answer at least 8 out of 10 correctly to proceed to the main part of the study. This is to ensure everyone understands the task and is providing consistent, high-quality ratings. If you do not pass, the study will end, for can review the example below to familiatize yourself with the process. To begin, please enter your Participant ID below.



Figure 20: Human evaluation interface. Participants viewed two policies side-by-side and selected the one that performed better or if it was a tie.

K ETHICS STATEMENT

This work complies with the ICLR Code of Ethics and established research integrity principles. No human subjects, personal, or sensitive data were used for training or evaluating policies; all human involvement in the benchmarking framework is limited to online preference annotation of video content from robot simulation, with participants recruited and compensated through recognized platforms according to their guidelines. Annotator tasks were designed to be non-intrusive, requiring only high-level performance comparisons in a double-blind interface, and a qualification quiz was used to promote reliable, unbiased evaluations. Research artifacts, algorithms, and datasets are used in compliance with their licenses, and no intellectual property or privacy concerns are raised by the experimental protocol. The study sought to anticipate and mitigate unfair outcomes and measured bias and robustness across diverse input sources, including systematic stress-tests for failure modes. Limitations concerning simulation realism and generalization beyond the studied benchmarks are acknowledged, as discussed in the main text and the limitations section. The full evaluation framework and results will be released for broad, equitable access by the research community.

L REPRODUCIBILITY STATEMENT

Significant effort has been devoted to ensuring that all results in this paper are reproducible. Details of the reality-to-simulation translation pipeline, environment and object generation, evaluation methodology (including VLM progress scoring, human preference infrastructure, and global policy ranking algorithm), as well as hyperparameter settings, are provided in the main text and in Appendices B–G.

The codebase for environment construction, simulation, and scoring, along with evaluation logs and benchmark scenarios, will be anonymously released as supplementary materials. Source datasets, simulation assets, and parameter optimization procedures are fully documented in the Appendix and will be indexed for direct replication. Further, systematic perturbation experiments and robustness analyses can be reproduced using the public release. The aim is to enable independent verification and extension of all primary claims and experimental findings using the open-source pipeline described in the manuscript.

M LLM USAGE STATEMENT

In the development of this manuscript, Large Language Models (LLMs) were used as general-purpose assistive tools for grammar checking, rewording, preparation of summary paragraphs, and technical editing. No sections of the manuscript, experimental results, or research claims were generated solely by an LLM. All original research content, analyses, and experiment design were authored by the stated researchers. When LLMs contributed to polishing language or format, the outputs were reviewed, edited, and verified for accuracy to ensure compliance with the ICLR Code of Ethics, integrity, and originality. The authors take full responsibility for all content presented here and confirm that no LLM-generated text constitutes plagiarism, fabrication, or any form of scientific misconduct. LLMs are not considered authors and have not contributed to research ideation, methodology, or core scientific findings