

GSWorld: Closed-Loop Photorealistic Simulation Suite for Robot Manipulation

Guangqi Jiang^{*1} Haoran Chang^{*1,2} Ri-Zhao Qiu¹ Yutong Liang¹ Mazeyu Ji¹ Jiyue Zhu¹
 Zhao Dong³ Xueyan Zou¹ Xiaolong Wang¹
¹UC San Diego, ²UC Los Angeles, ³Meta
 *Equal Contributions <https://3dgsworld.github.io>

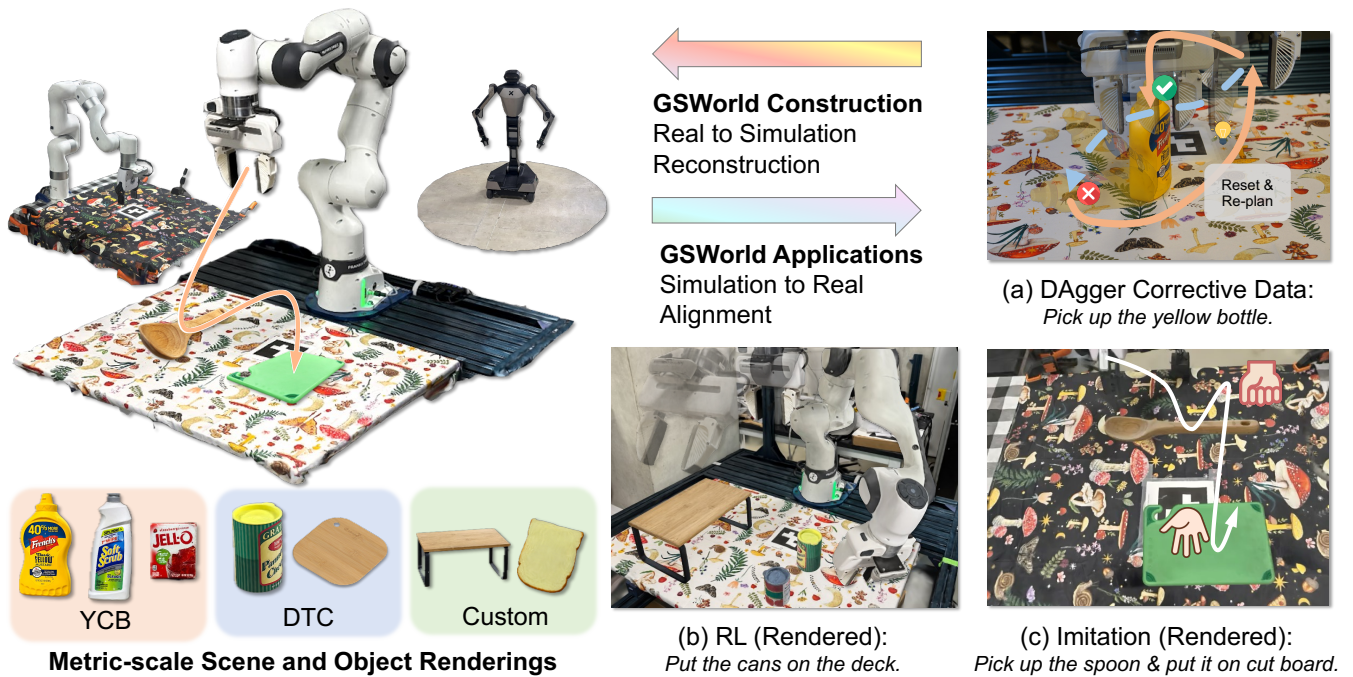


Fig. 1: GSWorld leverages 3DGS reconstruction to render photo-realistic robot scenes, supports multiple policy learning recipes in the simulation, and realizes zero-shot sim2real transfer. GSWorld also applies to policy visual benchmarking and virtual teleoperation for data collection.

Abstract—This paper presents GSWorld, a robust, photo-realistic simulator for robotics manipulation that combines 3D Gaussian Splatting with physics engines. Our framework advocates ‘closing the loop’ of developing manipulation policies with reproducible evaluation of policies learned from real-robot data and sim2real policy training without using real robots. To enable photo-realistic rendering of diverse scenes, we propose a new asset format, which we term GSDF (Gaussian Scene Description File), that infuses Gaussian-on-Mesh representation with robot URDF and other objects. With a streamlined reconstruction pipeline, we curate a database of GSDF that contains 3 robot embodiments for single-arm and bimanual manipulation, as well as more than 40 objects. Combining GSDF with physics engines, we demonstrate several immediate interesting applications: (1) learning zero-shot sim2real pixel-to-action manipulation policy with photo-realistic rendering, (2) automated high-quality DAGger data collection for adapting policies to deployment environments, (3) reproducible benchmarking of real-robot manipulation policies in simulation, (4) simulation data collection by virtual teleoperation, and (5) zero-shot sim2real visual reinforcement learning.

I. INTRODUCTION

Training manipulation policies typically relies on three data sources—simulation, human videos, and real teleoperation—each presenting a distinct trade-off. While simulation provides a perfectly aligned action space for the robot, it often suffers from many sim-to-real gaps. Human videos offer the benefit of photo-realistic scenes and real physics, but lack temporally aligned robot actions and operate in a mismatched action space. Teleoperation successfully aligns both perception and actions, yet its high cost and difficulty to scale are significant limitations.

To resolve these trade-offs, we introduce GSWorld, a closed-loop, photo-realistic simulation suite that couples 3D Gaussian Splatting (3DGS) with physics to narrow both the visual and action-space gaps for manipulation. “Closed-loop” here means the same environment can be used to train, evaluate, diagnose failures, and relabel, enabling rapid iteration: policies perceive photo-realistic renderings while issuing controls in the robot’s native space, all inside a

simulator that mirrors real scenes closely enough to support zero-shot transfer and efficient adaptation. We demonstrate that GSWorld enables a range of downstream applications and, in particular, effective sim-to-real transfer for imitation learning, reinforcement learning, and DAGger-style data collection.

This closed-loop capability is powered by a bidirectional pipeline that ensures tight alignment between the physical world and its digital twin. In the *real-to-sim* direction, our pipeline reconstructs a metric-accurate digital twin from short multi-view captures, sets an absolute scale with ArUco markers, and aligns the robot URDF to the scene via surface fitting (e.g., ICP). We then attach collision meshes and material properties to produce a versatile GSDF asset. The *sim-to-real* direction is the reverse: policies trained in GSWorld deploy on hardware without interface translation because their control and observation spaces match the robot’s native APIs. Policies trained with both sim and real data can be deployed in the sim to evaluate, detect failures, and gather DAGger corrections, thus closing the iteration loop. Achieving this sim-to-real alignment requires that the scene’s geometry, camera properties, and action semantics remain consistent across the real-to-sim divide. We measure the degree of this correspondence through a suite of metrics evaluating the visual, geometric, and functional similarity between the two worlds.

With photo-realistic perception and native action-space control in one loop, GSWorld supports zero-shot sim-to-real for both visual imitation learning and visual RL, while exploiting scalable parallelism to accelerate data generation and training. Its closed-loop DAGger workflow lets practitioners reproduce on-robot failures inside the digital twin, step through them frame-by-frame, and collect targeted corrective labels with far less teleoperation overhead. Finally, GSWorld provides reproducible visual benchmarking: shared GSDF assets, fixed camera intrinsics/extrinsics, consistent lighting/materials, and standardized action semantics enable apples-to-apples comparisons across robots, scenes, and tasks—so improvements reflect algorithmic progress rather than environmental variance.

Existing GS-based simulators either target single-setup photo-realistic rendering, provide engine-tied pipelines without a portable asset standard, or limit the reproducible cross-embodiment benchmarking and deployment-oriented on-policy data collection [5, 26, 37, 44]. While GSWorld delivers an effective real-to-sim-to-real workflow that unifies photo-realistic 3DGS with contact-accurate physics, enabling scalable cross-embodiment benchmarking, zero-shot imitation and reinforcement learning, and automated high-quality DAGger data collection for continual deployment-time improvement.

In summary, our contributions are:

- A solid real-to-sim-to-real pipeline. Our robust real-to-sim-to-real pipeline accurately aligns the simulation with the real environment, enabling a wide range of subsequent applications.
- Simulation Data Collection and Visual Imitation Learn-

ing (IL). GSWorld supports multiple sim data collection methods, e.g. motion planning, teleoperation. IL policies trained with GSWorld data can be directly deployed to the reconstructed real-world scenes.

- Visual RL. GSWorld is designed to utilize parallel environments in the simulation to train RL policies. We provide an analysis to show that GSWorld reduces RL sim2real visual gaps.
- Closed-loop DAGger Learning with Visual Benchmarking. GSWorld shows reliable policy evaluation results in correlation with real-world deployments, contributing to using DAGger to iteratively improve real-world policies.

II. RELATED WORK

Robotics Simulation. With developments in computer graphics for rendering [53] and object material simulation [49], the robotics community has designed various physics engines [9, 12] and simulators [46, 55, 58] to support various robotics tasks. Recently, more and more simulators have started to improve efficiency and fidelity. For example, Mujoco-Jax [58] exploits just-in-time compilers [17] to achieve impressive simulation efficiency in Python. To improve rendering fidelity and reduce **sim2real visual gaps**, recent papers have turned to sophisticated ray-tracing techniques [53] and generative AI tools [38] to reduce the visual observation gap between simulation renderings and the real world. A recent ‘meta-simulator’, Roboverse [14], attempts to provide a unified simulation interface to make use of the advantages in individual simulators. This paper focuses on improving the fidelity of simulation renderings by combining recent advances in 3D Gaussian Splatting and simulators.

Real2Sim (Real-to-Simulation). As an alternative to modeling basic physics and elements from the bottom-up in simulation, real2sim approaches take a different approach to build simulation assets by virtualizing real-world assets. Recent real2sim methods can be roughly divided into two categories: photo-realistic 3D reconstruction [1, 6, 16, 26, 33, 34, 43, 44, 52] and part-level (articulated) object understanding [7, 10, 35, 51].

Real2Sim (Real-to-Simulation) - Photo-realistic Rendering. Early method in Real2Sim reconstruction [43] use NeRFs [36], Diffusion [27] or Mesh [51] for photo-realistic modeling. Due to the inherent implicit representation of NeRFs, they often rely on surrogates such as deformation fields [20] to deform the visual renderings to accommodate object motion, which is unnatural and inefficient. On the other hand, mesh-based representations contain many artifacts [51] that lead to visual gap. Simpler [27] proposes to use generative modeling to advance reproducible benchmarking [21, 23] by supporting photorealistic rendering. However, Simpler requires expensive manual labor to match green screen and textures, hindering its scalability. With developments in 3DGS [25], a rasterization-based method, SplatSim [44] successfully combined 3DGS with PyBullet [9] to build a photo-realistic simulator and demonstrated zero-shot sim2real policy deployment. As 3DGS can be



Fig. 2: GSWorld provides an interface on top of existing simulators to render photorealistic assets. Our GSDF assets are compatible with existing simulators to use standard formats for rendering visuals (e.g., depth, segmentation) and computing physics collisions. GSWorld provides a rendering wrapper on top of simulators to make the RGB rendering photo-realistic to support various domain randomization and applications.

explicitly represented as ‘Gaussian Blobs’, the photo-realistic appearance can be displaced consistently with object physics. However, SplatSim relies on manual 3D segmentation of both the robot arm and objects, which is overfitted to a single scene. In parallel to the development of SplatSim, Embodied-GS [1] learns arm-object interaction without using physics engine in a single scene; Robo-GS [33] focuses on identifying physical parameters of individual objects from rendering; and ManiGaussian [34] investigates optimizing Gaussian representations from simulators with ideally synchronized multi-view information. Most recently, Re3Sim [16] extends SplatSim [44] and found that the capability to perform photo-realistic simulation leads to more robust manipulation policies with domain randomization and mixed simulation (e.g., non-photorealistic simulation assets with photo-realistic scenes). We continue to advance the progress in photo-realistic simulation by combining the latest advancements in 3DGS and introducing more 3DGS assets with a unified asset format.

Real2Sim (Real-to-Simulation) - Others. Here we provide a concise review of other progress in real2sim that is orthogonal to our method. Part-level (articulated) object understanding methods [7, 10, 35, 51] apply internet-scale pre-trained visual and language models [31, 45] to create physics and articulation of simulation assets from real-world observations. While these methods focus on understanding articulation of objects, recent methods have also worked on more challenging tasks such as estimating physics for deformable objects [24] and rigid contracts [41]. PhysTwin [24] optimizes physics models for elastic objects by assuming a Spring-Mass model and estimating physics parameters from video observations. Scalable Real2Sim [41] is an advancement from previous physics estimation method [33], in which the authors built a pipeline that uses robot arm and camera setups to automate estimation of rigid physics parameters including mass, center of mass, and inertia tensor [41]. Liu et al. [30] proposes to optimize robot kinematics from differentiable rendering.

III. METHOD

A. Problem Formulation

We consider the problem of learning robot policies from visual observations. Let \mathcal{S}_{sim} denote simulated scene and $\mathcal{G}_{\text{real}}$ denote the real scene $\mathcal{S}_{\text{real}}$ reconstructed by 3DGS from multiple RGB viewpoints $\mathcal{V} = \Sigma v_i$. $\mathcal{G}_{\text{real}}$ can be used to render novel-view RGB images $I^{\text{gs}} = \mathcal{G}_{\text{real}}(p, s)$, which enables photorealistic rendering of the scene under arbitrary camera poses p and environment states s .

Our goal is to replace raw real-world RGB observations I^{real} with 3DGS-rendered images I^{gs} for downstream robot learning tasks, including IL, RL, and DAgger. Formally, at each time step t , the underlying state of the system is represented as

$$s_t = (q_t, x_t^1, \dots, x_t^n), \quad (1)$$

where $q_t \in \mathbb{R}^m$ denotes the robot’s joint position, and x_t^k represents the 6D pose of the k -th object in $\mathcal{S}_{\text{real}}$. We mainly use joint position control for the robots. The robot receives an observation

$$o_t = I_t^{\text{gs}} = \mathcal{G}_{\text{real}}(p_t, s_t), \quad (2)$$

rendered from $\mathcal{G}_{\text{real}}$ given the current camera pose p_t and environment state s_t .

In RL, the policy π_θ is trained by interacting with the environment and receiving rewards. In IL, the expert \mathcal{E} provides demonstrations $\tau_{\mathcal{E}} = \{(q_1, o_1, a_1), \dots, (q_T, o_T, a_T)\}$, which are used to supervise π_θ . In DAgger, π_θ is iteratively refined by collecting expert rollouts from previous failure cases, getting $\tau_{\mathcal{D}}$. In all cases, the policy takes as input I_t^{gs} instead of I_t^{real} , and robot proprioceptions q_t , i.e.,

$$a_t \sim \pi_\theta(I_t^{\text{gs}}, q_t). \quad (3)$$

At test time, the trained policy π_θ must generalize to real-world observations I^{real} , ensuring that GSWorld bridges sim2real visual gap.

B. Recipe for Real2Sim Reconstruction

This section describes how GSWorld creates GSDF assets to construct photo-realistic scenes. Compared with existing work [16, 26, 33, 44, 57] that focus on constructing robots in a single scene, GSWorld is designed (1) with an easy-to-use streamlined procedure to reduce manual alignment efforts and (2) to incorporate recent advancements in 3DGS such as geometric accuracy [18].

a) Collecting Training Views: To construct a scene with a robot, we use both robot sensors (*i.e.*, wrist cameras and third-person cameras) and mobile phone cameras while saving the current joint pose of robot during the scene capture.

b) Aligning Scale for Metric Representation: Existing real2sim2real methods [42, 44] relies on COLMAP [48], which introduces scale ambiguity. While such ambiguity can be dealt with manually for a single scene, it affects the scalability for multiple robot embodiments and scenes. To avoid manual scale alignment, we use a simple solution to include a printed ArUco marker [13, 22] on the tabletop during the data collection (qualitative examples in Fig. 1). The detected keypoints of ArUco markers are projected onto the point cloud formed by 3DGS. We then scale the point cloud using the known scales of the ArUco marker. In addition, the ArUco marker helps identify the support surface in collision and estimate gravity direction.

c) Aligning Robots and Table: Given a metric-scale \mathcal{G}_{real} of a static robot \mathcal{R} and a metric-scale robot URDF in \mathcal{S}_{sim} , we align the simulation joint position with the real world. Then, we sample and densify surface point clouds from the visual mesh of the robot URDF. We then perform an ICP to compute the rigid transform $\mathcal{T}_{\mathcal{R},sim}^{gs}$: $\mathcal{G}_{real} = \mathcal{T}_{\mathcal{R},sim}^{gs} \cdot \mathcal{S}_{sim}$. Compared to previous methods [44], our ICP has fewer degrees of freedom since the scale is fixed. Given the alignment, we use K-NN to segment robot links in \mathcal{G}_{real} .

d) Object Assets: Our prior reconstruction stage focuses on background and robot scans. For moveable objects \mathcal{O} , we consider integrating existing large-scale datasets and supporting custom objects for generalizability. Specifically, we use DTC [11] for its photo-realistic visual quality, and YCB [4]. For custom objects, we use 2DGS [18] to get the reconstruction \mathcal{O}^{gs} and mesh reconstruction. Mass is estimated by weighing. The unobserved bottom regions of the object can be optionally inpainted using amodal reconstruction [2, 54] or 3D object generation [29] method. Similarly, we use ICP to get the transform $\mathcal{T}_{k,sim}^{gs}$ for the k -th object: $\mathcal{O}_k^{gs} = \mathcal{T}_{k,sim}^{gs} \cdot \mathcal{O}_k$.

C. Applications - Closing the Loop for Developing Visual Manipulation Policies

a) Closed-loop DAgger Training: In the traditional imitation learning settings, policy weights are not updated once they are trained and deployed. During deployment, policies often run into failures. DAgger [47] is a solution to this case where corrective data is used to train the model to adapt to failure cases. DAgger data has been shown to have much better data efficiency than plain data by previous

works [28, 59]. However, collecting DAgger data is hard as it requires reproducing scene setups for the model to ‘re-experience’ the failure case.

GSWorld provides an interface for automatic DAgger data collection in simulation. Given the GSDF of the target deployment environment and tasks with scripted policies, we roll out the policies with GSWorld. For failure recordings $\mathcal{D}_f = (s_1, \dots, s_T)$, we randomly sample recovery states $s_r \sim \mathcal{D}_f$ with a uniform sampler, where the task is still achievable in s_r , and run the motion planner to obtain corrective data, as illustrated in Fig. 3.

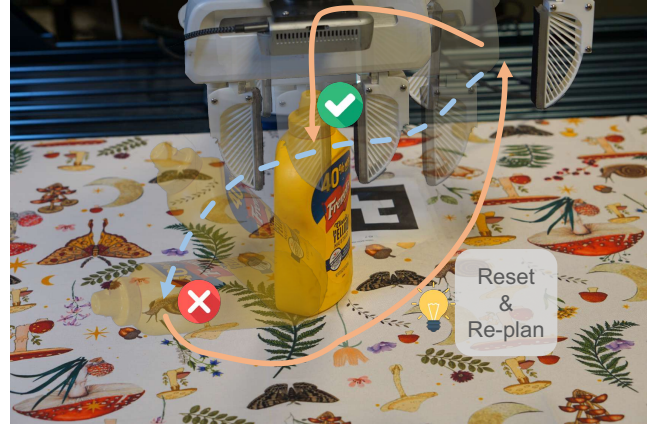


Fig. 3: **DAgger Data Collection in Simulation.** With privileged information provided by simulation, we can record and relive the failure cases during rollouts and generate corrective data for policy adaptation.

To iteratively improve zero-shot sim2real policies, we collect data, evaluate policies, record failures, and recover, resulting in an aggregated dataset τ_S :

$$\tau_S = \Sigma_i (Q_s, O_s, A_s)_i \quad (4)$$

, where Q_s, A_s denote simulation robot joint positions, and action labels. $O_s = I^{gs}$ is the GSWorld-rendered observations and i is the DAgger iteration.

To improve policies trained with real-world data, we first evaluate the polices and repeat the same loop, leading to dataset $\tau_{\mathcal{R}}$:

$$\tau_{\mathcal{R}} = (Q_r, O_r, A_r) \cup \tau_S \quad (5)$$

, where Q_r, O_r, A_r are collected in the real world by teleoperation.

b) Visual Benchmarking.: Recent VLAs [3, 32, 39] train generalist policies that work on many different robot embodiments. Learning from extensive cross-embodiment training data, these base models adapt to unseen robot hardware and language instructions in a few-shot or even zero-shot manner [3, 32]. However, due to their reliance on real robot data during training, there does not exist a standardized visual manipulation benchmark that studies the quality of the base models and their data sampling efficiency for novel embodiments.

GSWorld provides a photo-realistic rendering interface to address such an issue. With our GSDF assets, we provide photo-realistic simulation of various scenes with different robot embodiments, and a wide selection of interactable

objects. We define a range of manipulation tasks using ManiSkill [50] as the simulator backend. Moreover, we use image augmentation to further reduce the visual gap between sim and real.

c) *Reinforcement Learning*.: RL requires massive interaction between agents and the environments, usually with parallelism [50]. To better support RL, GSWorld optimizes its implementation by only parallelizing 3DGS points that are linked with moving parts of the scene, *i.e.*, robot \mathcal{R} and objects \mathcal{O} , and keeps other points cached with just one copy. This enables us to use a single GPU to run large parallelism to accelerate RL convergence.

IV. EXPERIMENTS

This section presents an empirical evaluation of GSWorld to demonstrate its effectiveness. Our experiments are designed to address the following research questions:

- **Zero-shot Sim2real Imitation Learning.** Can GSWorld effectively bridge the sim-to-real gap to enable zero-shot policy transfer?
- **Closed-loop Policy Improvement.** Does access to a digital twin of the target deployment environment increase sampling efficiency and enable continual improvement after policy deployment via DAgger?
- **Visual Benchmarking.** Does performance in GSWorld correlate with real-world performance?
- **Virtual Teleoperation.** Does GSWorld enable simulated data collection through human teleoperation?
- **Reinforcement Learning (RL).** Can GSWorld narrow down the sim2real visual gap for visual RL?

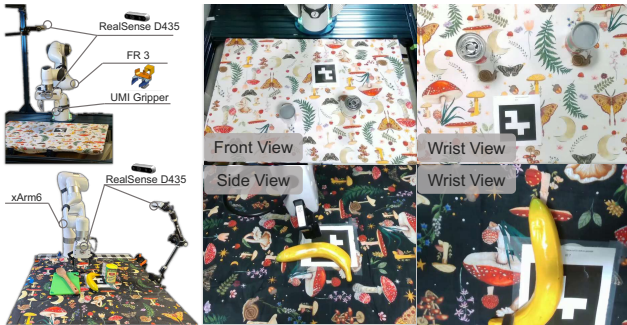


Fig. 4: **Real World Hardware Platforms.** For FR3, we set up two cameras: a third-person camera positioned in front of the robot (front view), and a wrist-mounted camera attached to the robot’s end effector (wrist view). For xArm6, we set up a third-person camera besides the robot (side view) and a wrist view.

a) *Hardware Platforms*: Though GSWorld scales to many robot embodiments, we consider three robot platforms for evaluation: a Franka Research 3 (FR3) robot with a UMI gripper [8]; a UF xArm6 with a parallel gripper; and a bimanual Galaxea R1 robot equipped with two 6-DoF arms.

b) *Experiment Protocol*: In total, we evaluate GSWorld on 4 manipulation tasks on FR3 and 3 tasks on xArm. We also use R1 to demonstrate that GSWorld supports virtual simulation teleoperation to collect data. For policy implementations, we use ACT [60] and Pi0 [3] to show GSWorld to be policy-agnostic. For visual benchmarking, we intentionally train policies with varying training settings



Fig. 5: **FR3 Task Visualizations.** We design 4 real-world robotic tasks that involve distinct manipulation skills and diverse objects on FR3.



Fig. 6: **xArm Task Visualizations.** We design 3 manipulation tasks on xArm6.

and data sizes to evaluate the performance of ‘good’ and ‘bad’ policies. Please kindly refer to our released codebase for full implementation details due to limited space.

c) *Manipulation Tasks Design*: We designed the following table-top manipulation tasks to evaluate performances, as shown in Fig. 5 and Fig. 6:

- **Place Box.** The bottle and the box are initialized randomly within a $45\text{ cm} \times 45\text{ cm}$ area. FR3 must pick up the bottle and place it onto the box.
- **Pour Sauce.** The mustard bottle and the bread slice are put within a $45\text{ cm} \times 45\text{ cm}$ area. FR3 is required to pour the sauce onto the bread slice.
- **Stack Cans.** Two cans are randomly placed within a $45\text{ cm} \times 45\text{ cm}$ area. FR3 must stack them.
- **Arrange Cans.** Two cans are randomly placed within a $20\text{ cm} \times 15\text{ cm}$ area. A rack is randomly placed beside them. FR3 is tasked with placing one object on the rack, followed by placing the second object next to it.
- **Align Cans.** xArm needs to grasp one can and put it next to the other can.
- **Grasp Banana.** xArm should grasp the banana and rotate it by 30 to 60 degrees along the z-axis.
- **Tidy Table.** xArm is required to pick up the kitchen spoon and place it onto the cutting board to clean the table.

A. Zero-shot Sim2real Imitation Learning

Fig. 7 shows the performances of policies trained with only simulation data. We leverage the MPLib motion planner [50], utilizing predefined motions and poses. For each task, we collect 100 trajectories per iteration and train on the sum of all generated data. We can safely conclude that GSWorld allows zero-shot sim2real policy transfer and shows promising success rates. More details are included in section IV-B.

B. Closed-loop DAgger for Continuous Policy Improvement

Data collection follows the zero-shot setup. We start DAgger with 100 expert trajectories, then in each of four additional iterations identify failed rollouts, reset to a solvable pre-failure state, and collect corrective demonstrations (100 trajectories/iteration). As shown in Fig. 7, DAgger significantly improves all four tasks over training from scratch (TfS).

DAgger also improves deployed real-world policies via a real2sim2real pipeline. We first train ACT with a few real demonstrations, then continue DAgger in GSWorld from that checkpoint. Fig. 8 shows clear post-deployment gains.

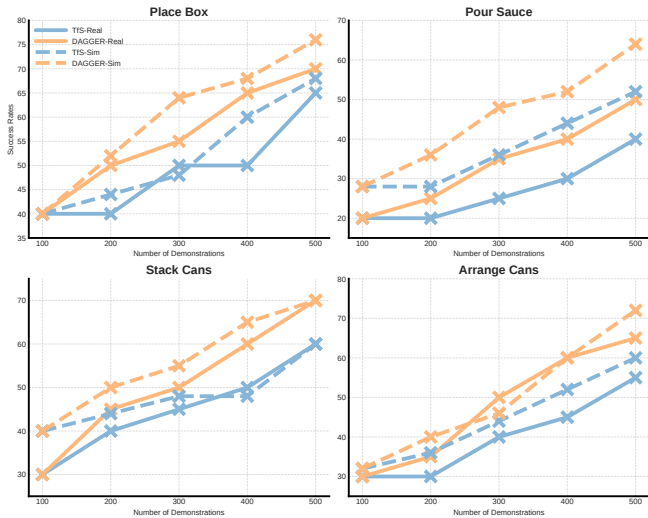


Fig. 7: **Closed-loop Sim2real DAGger Training on FR3.** Policies are trained with sim data and deployed in both sim and real environments. DAGger consistently improves policy performances and outperforms training from scratch.

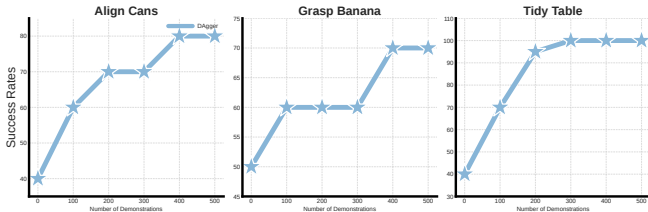


Fig. 8: **Closed-loop Real2sim2real DAGger on xArm.** DAGger can also be used to improve real-world policies.

These results highlight the value of closed-loop DAGger: our photo-realistic digital twin enables corrective data collection that is hard to obtain in the real world due to precise reset requirements.

C. Benchmarking

In Fig. 9, we observe a strong correlation between simulation and real-world performance across all evaluated tasks and different policy architectures. This correlation indicates that GSWorld can reliably predict real-world outcomes without requiring physical deployment of policies in the real world scene. As shown in Table I, higher simulated performance consistently corresponds to higher success rates in real-world experiments. By leveraging the photorealistic rendering capabilities of 3DGS [25], we establish a benchmarking framework that reflects real-world behaviors.

D. Virtual Teleoperation

Simulation data can be used to scale up robot policy learning [19, 56]. We show that through mouse and keyboard, we can collect teleoperation data in the simulation with real renderings, as illustrated in Fig. 10.

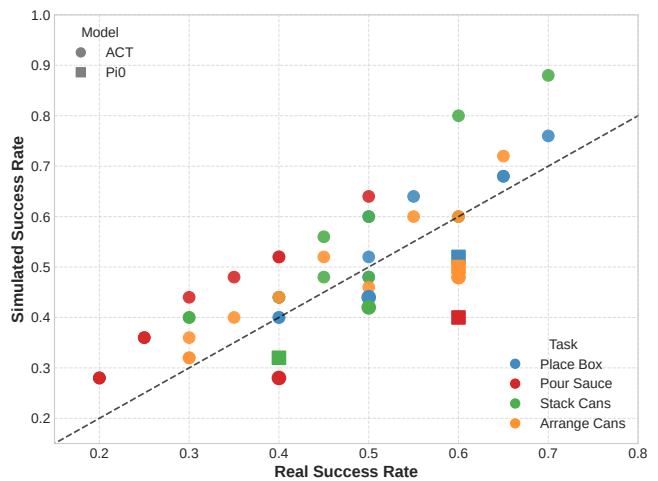


Fig. 9: **Visual Benchmarking with FR3.** We roll out various policies in both real and sim. We intentionally use different policies and data sizes to show positive correlations of sim-real performance regardless of the quality of the policies.

	ACT [60]		Pi0 [3]	
	real	sim	real	sim
Place Box	50.0%	44.0%	60.0%	52.0%
Pour Sauce	40.0%	28.0%	60.0%	40.0%
Stack Cans	50.0%	42.0%	40.0%	32.0%
Arrange Cans	60.0%	48.0%	60.0%	50.0%
Avg.	50.0%	41.0%	55.0%	43.5%

TABLE I: **Visual Benchmarking of Real-world Policies on FR3.** We show the evaluation performance in both sim and real of policies trained with only real-world data.

E. Visual Reinforcement Learning

GSWorld is designed to support parallel environments so that we can use this to learn visual RL policies, which has great potential in robot learning [19, 40]. We trained asymmetric SAC [15] with GSWorld, where the critic sees simulation-privileged information and the actor only uses robot joint position. As we aim to show that GSWorld can reduce the visual gap for RL instead of acquiring good sim2real VRL policies, we train with no domain randomization except for color jittering for efficiency. We only use the third-person view since wrist camera shows significant gaps during RL exploration. Our real-world success rates for *Grasp Banana* and *Tidy Table* are 30% and 20%, while baseline ManiSkill reaches 0% and 5%, respectively. Training results are shown in Fig. 11.

V. CONCLUSION

In this work, we present a pipeline for constructing a photorealistic digital twin that delivers highly correlated performance metrics between simulation and real-world deployments across different policy architectures and embodiments. Furthermore, we demonstrate that our environment can efficiently collect corrective data at scale, enabling more effective policy training.

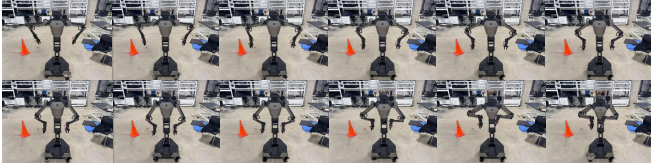


Fig. 10: **Galaxea R1 Virtual Teleoperation.** We use a keyboard to teleop R1 and render photo-realistic videos.

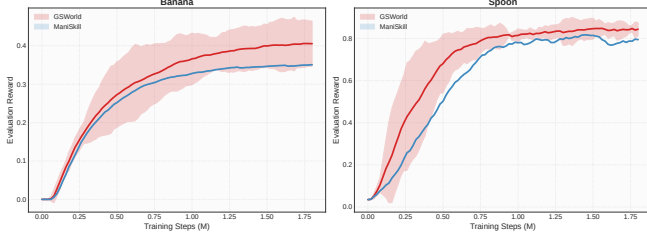


Fig. 11: **SAC.** Results are aggregated over 3 runs. Directly training from ManiSkill is also plotted for comparison.

REFERENCES

- [1] Jad Abou-Chakra et al. “Physically Embodied Gaussian Splatting: A Visually Learnt and Physically Grounded 3D Representation for Robotics”. In: *CoRL*. 2024.
- [2] William Agnew et al. “Amodal 3d reconstruction for robotic manipulation via stability and connectivity”. In: *CoRL*. 2021.
- [3] Kevin Black et al. “ π_0 : A Vision-Language-Action Flow Model for General Robot Control”. In: *arXiv preprint arXiv:2410.24164* (2024).
- [4] Berk Calli et al. “Yale-CMU-Berkeley dataset for robotic manipulation research”. In: *The International Journal of Robotics Research* 36.3 (2017), pp. 261–268.
- [5] Tianxing Chen et al. “RoboTwin 2.0: A Scalable Data Generator and Benchmark with Strong Domain Randomization for Robust Bimanual Robotic Manipulation”. In: *arXiv preprint arXiv:2506.18088* (2025).
- [6] Ziyu Chen et al. “Omnire: Omni urban scene reconstruction”. In: *arXiv preprint arXiv:2408.16760* (2024).
- [7] Zoey Chen et al. “Urdformer: A pipeline for constructing articulated simulation environments from real-world images”. In: *RSS*. 2024.
- [8] Cheng Chi et al. “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots”. In: *arXiv preprint arXiv:2402.10329* (2024).
- [9] Erwin Coumans and Yunfei Bai. *Pybullet, a python module for physics simulation for games, robotics and machine learning*. 2016.
- [10] Tianyuan Dai et al. “Automated creation of digital cousins for robust policy learning”. In: *CoRL*. 2024.
- [11] Zhao Dong et al. “Digital Twin Catalog: A Large-Scale Photorealistic 3D Object Digital Twin Dataset”. In: *arXiv preprint arXiv:2504.08541* (2025).
- [12] Tom Erez, Yuval Tassa, and Emanuel Todorov. “Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ode and physx”. In: *ICRA*. 2015.
- [13] Sergio Garrido-Jurado et al. “Automatic generation and detection of highly reliable fiducial markers under occlusion”. In: *Pattern Recognition* (2014).
- [14] Haoran Geng et al. “RoboVerse: Towards a unified platform, dataset and benchmark for scalable and generalizable robot learning”. In: *arXiv preprint arXiv:2504.18904* (2025).
- [15] Tuomas Haarnoja et al. “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor”. In: *International conference on machine learning*. Pmlr. 2018, pp. 1861–1870.
- [16] Xiaoshen Han et al. “Re³Sim: Generating High-Fidelity Simulation Data via 3D-Photorealistic Real-to-Sim for Robotic Manipulation”. In: *arXiv preprint arXiv:2502.08645* (2025).
- [17] Yuanming Hu et al. “Taichi: a language for high-performance computation on spatially sparse data structures”. In: *ACM Transactions on Graphics (TOG)* (2019).
- [18] Binbin Huang et al. “2d gaussian splatting for geometrically accurate radiance fields”. In: *ACM SIGGRAPH*. 2024.
- [19] Tao Huang et al. “Diffusion reward: Learning rewards via conditional video diffusion”. In: *European Conference on Computer Vision*. Springer. 2024, pp. 478–495.
- [20] Clément Jambon et al. “NeRFshop: Interactive Editing of Neural Radiance Fields”. In: *Proceedings of the ACM on Computer Graphics and Interactive Techniques* (2023).
- [21] Stephen James et al. “Rlbench: The robot learning benchmark & learning environment”. In: *IEEE Robotics and Automation Letters* (2020).
- [22] Mazeyu Ji et al. “Graspsplats: Efficient manipulation with 3d feature splatting”. In: *CoRL*. 2024.
- [23] Guangqi Jiang et al. “Robots Pre-Train Robots: Manipulation-Centric Robotic Representation from Large-Scale Robot Datasets”. In: *arXiv preprint arXiv:2410.22325* (2024).
- [24] Hanxiao Jiang et al. “Phystwin: Physics-informed reconstruction and simulation of deformable objects from videos”. In: *arXiv preprint arXiv:2503.17973* (2025).
- [25] Bernhard Kerbl et al. “3d gaussian splatting for real-time radiance field rendering.” In: *ACM Trans. Graph.* 42.4 (2023), pp. 139–1.
- [26] Xinhai Li et al. “RoboGSim: A Real2Sim2Real Robotic Gaussian Splatting Simulator”. In: *arXiv preprint arXiv:2411.11839* (2024).
- [27] Xuanlin Li et al. “Evaluating Real-World Robot Manipulation Policies in Simulation”. In: *arXiv preprint arXiv:2405.05941* (2024).
- [28] Huihan Liu et al. “Robot learning on the job: Human-in-the-loop autonomy and learning during deploy-

- ment”. In: *The International Journal of Robotics Research* (2022).
- [29] Ruoshi Liu et al. “Zero-1-to-3: Zero-shot one image to 3d object”. In: *ICCV*. 2023.
- [30] Ruoshi Liu et al. “Differentiable robot rendering”. In: *arXiv preprint arXiv:2410.13851* (2024).
- [31] Shilong Liu et al. “Grounding dino: Marrying dino with grounded pre-training for open-set object detection”. In: *ECCV*. 2024.
- [32] Songming Liu et al. “RDT-1B: a Diffusion Foundation Model for Bimanual Manipulation”. In: *arXiv preprint arXiv:2410.07864* (2024).
- [33] Haozhe Lou et al. *Robo-GS: A Physics Consistent Spatial-Temporal Model for Robotic Arm with Hybrid Representation*. 2024.
- [34] Guanxing Lu et al. “Manigaussian: Dynamic gaussian splatting for multi-task robotic manipulation”. In: *ECCV*. 2024.
- [35] Zhao Mandi et al. “Real2code: Reconstruct articulated objects via code generation”. In: *arXiv preprint arXiv:2406.08474* (2024).
- [36] Ben Mildenhall et al. “Nerf: Representing scenes as neural radiance fields for view synthesis”. In: *Communications of the ACM* 65.1 (2021), pp. 99–106.
- [37] Yao Mu et al. “RoboTwin: Dual-Arm Robot Benchmark with Generative Digital Twins”. In: *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*. 2025, pp. 27649–27660.
- [38] Sorous Nasiriany et al. “Robocasa: Large-scale simulation of everyday tasks for generalist robots”. In: *arXiv preprint arXiv:2406.02523* (2024).
- [39] Octo Model Team et al. “Octo: An Open-Source Generalist Robot Policy”. In: *Proceedings of Robotics: Science and Systems*. 2024.
- [40] Guoping Pan et al. “RoboDuet: Learning a Cooperative Policy for Whole-body Legged Locomanipulation”. In: *IEEE Robotics and Automation Letters* (2025).
- [41] Nicholas Pfaff et al. “Scalable real2sim: Physics-aware asset generation via robotic pick-and-place setups”. In: *arXiv preprint arXiv:2503.00370* (2025).
- [42] Ri-Zhao Qiu et al. “Language-driven physics-based scene synthesis and editing via feature splatting”. In: *ECCV*. 2024.
- [43] Delin Qu et al. “Livescene: Language embedding interactive radiance fields for physical scene rendering and control”. In: *arXiv preprint arXiv:2406.16038* (2024).
- [44] M Nomaan Qureshi et al. “SplatSim: Zero-shot sim2real transfer of rgb manipulation policies using gaussian splatting”. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2025, pp. 6502–6509.
- [45] Alec Radford et al. “Learning transferable visual models from natural language supervision”. In: *ICML*. 2021.
- [46] Eric Rohmer, Surya PN Singh, and Marc Freese. “V-REP: A versatile and scalable robot simulation framework”. In: *IROS*. 2013.
- [47] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. “A reduction of imitation learning and structured prediction to no-regret online learning”. In: *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings. 2011.
- [48] Johannes L Schonberger and Jan-Michael Frahm. “Structure-from-motion revisited”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 4104–4113.
- [49] Deborah Sulsky, Zhen Chen, and Howard L Schreyer. “A particle method for history-dependent materials”. In: *Computer methods in applied mechanics and engineering* (1994).
- [50] Stone Tao et al. “ManiSkill3: GPU Parallelized Robotics Simulation and Rendering for Generalizable Embodied AI”. In: *arXiv preprint arXiv:2410.00425* (2024).
- [51] Marcel Torne et al. “Reconciling Reality through Simulation: A Real-to-Sim-to-Real Approach for Robust Manipulation”. In: *arXiv preprint arXiv:2403.03949* (2024).
- [52] Yijia Weng et al. “Neural Implicit Representation for Building Digital Twins of Unknown Articulated Objects”. In: *CVPR*. 2024.
- [53] Turner Whitted. “An improved illumination model for shaded display”. In: *Proceedings of the 6th annual conference on Computer graphics and interactive techniques*. 1979, p. 14.
- [54] Tianhao Wu et al. “Amodal3R: Amodal 3D Reconstruction from Occluded 2D Images”. In: *arXiv preprint arXiv:2503.13439* (2025).
- [55] Fanbo Xiang et al. “Sapien: A simulated part-based interactive environment”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020, pp. 11097–11107.
- [56] Ruihan Yang et al. *EgoVLA: Learning Vision-Language-Action Models from Egocentric Human Videos*. 2025.
- [57] Sizhe Yang et al. “Novel Demonstration Generation with Gaussian Splatting Enables Robust One-Shot Manipulation”. In: *arXiv preprint arXiv:2504.13175* (2025).
- [58] Kevin Zakka et al. “MuJoCo Playground”. In: *arXiv preprint arXiv:2502.08844* (2025).
- [59] Xiaoyu Zhang et al. “Diffusion meets dagger: Supercharging eye-in-hand imitation learning”. In: *RSS*. 2024.
- [60] Tony Z Zhao et al. “Learning fine-grained bimanual manipulation with low-cost hardware”. In: *arXiv preprint arXiv:2304.13705* (2023).