

RE³SIM: Generating High-Fidelity Simulation Data via 3D-Photorealistic Real-to-Sim for Robotic Manipulation

Xiaoshen Han^{1,2} Minghuan Liu¹ Yilun Chen² Junqiu Yu² Xiaoyang Lyu³ Yang Tian² Bolun Wang²
Weinan Zhang¹ Jiangmiao Pang²

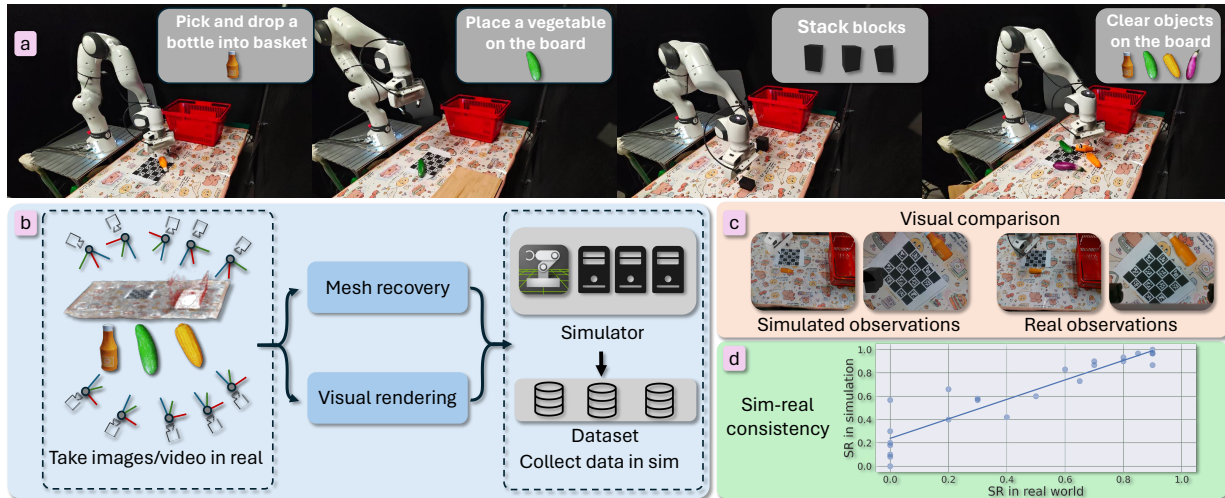


Figure 1: **Illustration of RE³SIM.** a) RE³SIM allows zero-shot policy transfer on various tasks. b) The system pipeline to generate high-quality data. c) High-fidelity rendering results. d) Consistency in success rates between real and simulated environments.

Abstract

Real-world data collection for robotics is costly and resource-intensive, requiring skilled operators and expensive hardware. Simulations offer a scalable alternative but often fail to achieve sim-to-real generalization due to geometric and visual gaps. To address these challenges, we propose a 3D-photorealistic real-to-sim system, namely, RE³SIM, addressing geometric and visual sim-to-real gaps. RE³SIM employs advanced 3D reconstruction and neural rendering techniques to faithfully recreate real-world scenarios, enabling real-time rendering of simulated cross-view cameras within a physics-based simulator. By utilizing privileged information to collect expert demonstrations efficiently in simulation, and train robot policies with imitation learning, we vali-

date the effectiveness of the real-to-sim-to-real pipeline across various manipulation task scenarios. Notably, with only simulated data, we can achieve zero-shot sim-to-real transfer with an average success rate exceeding 58%. To push the limit of real-to-sim, we further generate a large-scale simulation dataset, demonstrating how a robust policy can be built from simulation data that generalizes across various objects. Codes and demos are available at: <http://xshenhan.github.io/Re3Sim/>.

1. Introduction

We have witnessed impressive generalization capabilities of robotic models (Brohan et al., 2023; Cheang et al., 2024; Kim et al., 2024; Li et al., 2023; 2024b; Tian et al., 2024) in certain tabletop or kitchen scenarios, achieved through high-quality teleoperation data. However, collecting real-world expert data (Black et al., 2024; Brohan et al., 2022) remains a time-consuming and costly process. Despite advancements in teleoperation systems (Cheng et al., 2024;

¹Shanghai Jiao Tong University ²Shanghai AI Lab ³The University of Hong Kong. Correspondence to: Yilun Chen <chenyilun@pjlab.org.cn>, Jiangmiao Pang <pangjiangmiao@pjlab.org.cn>.

Aldaco et al., 2024; Yang et al., 2024; Fu et al., 2024), the labor involved is still intensive, making this a key challenge in robotics research. Recent efforts (O’Neill et al., 2023; Khazatsky et al., 2024) have explored inter-institutional collaboration to address this issue. In contrast, simulation data offers the advantage of exponential scalability with computational resources, making it an appealing and renewable alternative for training robotic policies. Research on sim-to-real transfer (Mandlekar et al., 2023; Jiang et al., 2024; Nasiriany et al., 2024) has increasingly focused on generating high-quality simulated or synthetic data to train real-world robotic policies. Unfortunately, simulation data often suffers from large sim-to-real gaps, necessitating the collection of target domain data for policy fine-tuning.

Among these gaps, two notable ones are the geometric gap and the visual gap. The geometric gap arises from shape mismatches between objects, often due to the reliance on a limited set of pre-defined CAD models in simulation, which assume perfect structures. The visual gap, on the other hand, comes with noiseless and imperfect rendering, posing challenges for tasks that rely heavily on appearance information, such as picking fresh fruits, pressing colorful buttons, sorting products based on packaging, and so on.

To bridge these gaps between real and simulated environments, we propose a 3D-photorealistic real-to-sim-to-real system, **RE³SIM**, namely, **reconstruction-rendering-based real-to-sim**. RE³SIM faithfully replicates real-world scenarios by combining realistic 3D geometry reconstruction and visual RGB rendering. The 3D geometry is reconstructed using multi-view stereo (MVS) techniques (Schönberger & Frahm, 2016; Cernea, 2020), while the visual rendering is achieved through Gaussian rasterization (Kerbl et al., 2023). This dual approach enables efficient and high-fidelity reconstruction, with simulated dynamics computed through physics engine backends (Makovychuk et al., 2021; Xiang et al., 2020; Todorov et al., 2012) and real-time rendering handled by a dedicated hybrid rendering engine.

Specifically, RE³SIM adopts a sequential real-to-sim pipeline comprising three key steps: (a) *mesh recovery* for reconstructing scene and object geometry; (b) *hybrid visual rendering* for compositing foreground and background elements, and (c) *real-world alignment* to synchronize world coordinates between the real and simulated environments. Human involvement is minimal in the real-to-sim process and limited to: (a) Placing ArUco markers in the target scenario for real-world alignment. (b) Capturing photos or videos of the scene and objects, including an additional capture step to align the robot base.

Once all assets and robots are imported into the simulator, RE³SIM utilizes a privileged policy to generate high-fidelity expert data. It incorporates the following key features to enhance sim-to-real transfer:

- **High-fidelity geometry and vision:** Achieves superior reconstruction and rendering quality, reducing the sim-to-real gap in both geometry and vision.
- **Rapid scene reconstruction:** Enables novel scene reconstruction in under three minutes of manual setup.
- **Efficient rendering:** Provides 24 FPS rendering for 480p images across two independent camera views.

To validate the effectiveness of the pipeline, we design several tabletop tasks demonstrating the generality and applicability of the real-to-sim-to-real process, as shown in Fig. 1. The zero-shot sim-to-real policies achieve an average success rate of 58% with only about 10 minutes of data collection in simulation, showcasing the effectiveness of our approach. We expect these experiments to highlight the potential of the real-to-sim-to-real pipeline in reducing sim-to-real gaps, paving the way for future advancements in scalable, efficient, and generalizable robotic policy.

2. Background

2.1. 3D Reconstruction

The goal of 3D reconstruction is to recover the geometry of 3D scenes from various signals, such as images, point clouds, and more. To this end, the first common step is taking structure-from-motion approaches (e.g., Colmap) to estimate camera poses and sparse point clouds (Schönberger & Frahm, 2016). Once the camera poses are determined, dense point clouds can be obtained through multi-view stereo (MVS) techniques. These include traditional methods (Schönberger et al., 2016), as well as network-based approaches (Yao et al., 2018). This process effectively enables the reconstruction of geometry from images. Besides, commercial software solutions, such as ARCode and PolyCam, leverage RGB-D SLAM methods to simultaneously estimate camera poses and depth information, facilitating the recovery of 3D models. These software solutions often include some post-processing techniques, such as smoothing and inpainting, together with user-friendly GUIs, enhancing the usability and quality of the reconstructed 3D models.

2.2. Rendering

Rendering involves generating photorealistic images from specified camera views based on either explicit or implicit 3D representations. Implicit methods, like Neural Radiance Fields (NeRF) (Mildenhall et al., 2021), represent the 3D space using a density field encoded in a neural network and render images through volumetric rendering techniques. Explicit representations, such as meshes and point clouds, use texture mapping (Blinn & Newell, 1976) for rendering and synthesizing novel views. Another recent progress in explicit representation is Gaussian primitives,

which adopt Gaussian rasterization for high-fidelity rendering, also known as 3D Gaussian splatting (3DGS) (Kerbl et al., 2023). Compared with other methods, 3DGS achieves more realistic rendering results and offers higher rendering efficiency. In detail, each Gaussian primitive is modeled as an ellipse with a full 3D covariance matrix Σ defined in world space, centered at point μ :

$$G(x) = e^{-\frac{1}{2}x^T\Sigma^{-1}x},$$

where the covariance matrix Σ of a Gaussian primitive is derived from its rotation R and scaling factor S as:

$$\Sigma = RSS^T R^T.$$

Additionally, each Gaussian is assigned a color c_i and depth-ordered using the Z-buffer. Using the alpha-blending formula, the rendered color C for a view is calculated as:

$$C = \sum_{i=1}^N \prod_{j<i} \alpha_j (1 - \alpha_i) c_i.$$

3. RE³SIM

Physic-based robot simulators, such as Isaac Sim (NVIDIA, 2021a) (PhysX (NVIDIA, 2021b)), PyBullet (Bullet) (Coumans & Bai, 2016–2021), Mujoco (Todorov et al., 2012), and Gazebo (Koenig & Howard, 2004) (Bullet, ODE (Smith., 2001), DART (Lee et al., 2018), and Simbody (Sherman et al., 2011)), offer high-accurate simulation results based on underlying physics engines. However, it’s hard to directly obtain an accurate 3D model of the real-world scene, and the simulators are not good at rendering high-fidelity images, leaving a large gap from real-world images. Recent advancements in 3D reconstruction (Mildenhall et al., 2021; Kerbl et al., 2023; Cernea, 2020; Yu et al., 2024; Guédon & Lepetit, 2024) provide high-quality rendering results and detailed 3D geometric models, promising to bridge these gaps. We introduce RE³SIM, a 3D reconstruction-based system, to bridge the visual gap and recover their realistic 3D meshes for sim-to-real transfer.

A straightforward approach is to reconstruct objects and the background together. However, using Multi-View Stereo (MVS) methods for background reconstruction often leads to suboptimal visual quality, while techniques like SuGaR (Guédon & Lepetit, 2024) struggle with accurately reconstructing flat planes, causing protrusions and indentations that reduce realism in object movement. Additionally, segmenting foreground objects for manipulation is labor-intensive. To address these issues, we propose separately reconstructing background meshes for collision estimation and leveraging 3DGS to improve rendering realism, aligning the two. Since object rendering constitutes a small portion of the robot’s visual observations, we opt not to use 3DGS to render objects.

3.1. System Overview

RE³SIM is a real-to-sim-real pipeline that reconstructs meshes for collision estimation, utilizes 3D Gaussian scene representations for rendering, and generates simulation data within a physics-based simulator¹. It is worth noting that we only consider reconstructing the background and foreground objects, robots are not involved in our system since their accurate 3D assets are usually easily accessible. Further, we take this work as an initial study of RE³SIM, validated on a table-top robot arm for rigid body tasks. As shown in Fig. 2, RE³SIM are composed of four components: mesh recovery for object collision computation, hybrid visual rendering combining 3DGS for the background and mesh-based rendering for the foreground, real-world alignment of the reconstructed scene and robot positions, and large-scale data generation for manipulation tasks.

3.2. Mesh Recovery

This step reconstructs realistic objects and background geometries for accurate collision detection and simulation.

Mesh reconstruction. The meshes of backgrounds and objects are reconstructed separately for flexibility and are represented in USD for extensibility. For the background, we first employ SfM techniques (COLMAP) to estimate the pose for each image and get the sparse point cloud based on a set of images, and then utilize OpenMVS (Cernea, 2020) for detailed mesh reconstruction. PolyCam (Polycam, 2020) is also able to reconstruct the mesh of the background, but empirically, the geometric reconstruction results of OpenMVS exhibit fewer minor protrusions and depressions compared to those of PolyCam. Subsequently, we simplify this initial mesh to create collision bodies. This approach effectively captures intricate details while maintaining accurate planar reconstruction. For foreground objects, we employ AR-Code (Code, 2022), which can automatically segment the object, balancing usability and quality.

Post-processing. Reconstructed collision bodies often have voids or sharp edges due to imperfect image quality, causing erratic behavior in simulation. For example, objects may fail to rest stably on surfaces. Void filling and surface smoothing help improve stability and realism in simulation. As acquiring physical parameters like mass and friction is challenging with only initial static visual data, we use default values to simplify the process without sacrificing performance and find it works well for many manipulation tasks.

¹Isaac Sim (NVIDIA, 2021a) is used by default, but other ray tracing options are also feasible.

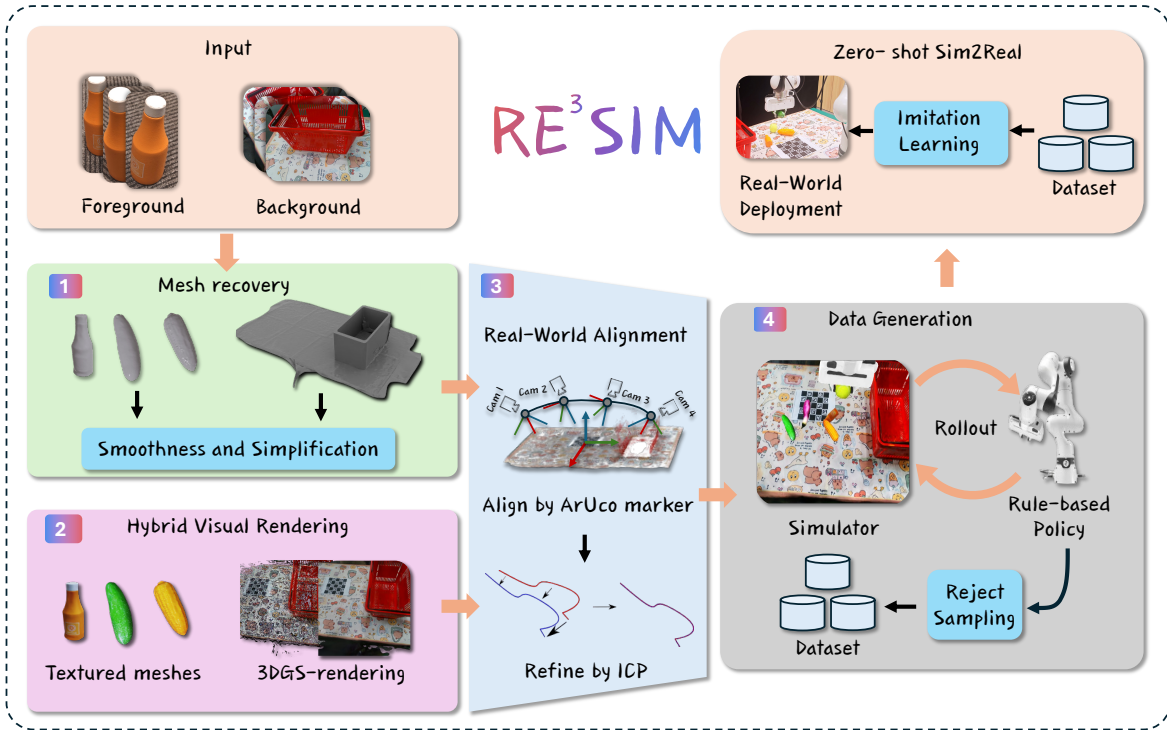


Figure 2: **Illustration of the proposed real-to-sim-to-real system, RE³SIM.** It leverages 3D reconstruction and a physics-based simulator, providing small 3D gaps that enable large-scale simulation data generation for learning manipulation skills via sim-to-real transfer.

3.3. Hybrid Visual Rendering

Color images, as a key type of perception signal, often yield a large visual gap between simulation and the real world. To close such a gap, we render the foreground objects and backgrounds with hybrid real-to-sim rendering techniques. For the foreground objects, object renderings are applied from the mesh with ray tracing supports. Background reconstruction contains most parts of the scenes and applies more photorealistic rendering using 3DGS. Z-buffer rendering is applied to mix objects according to ground-truth depth. Notably, hybrid visual rendering supports multi-view rendering according to the perspective views. The rendering time consumption of each component is shown in Tab. 5.

3.4. Real-World Alignment

Despite the prior alignments between rendering and meshes, another key step is to align with real-world scenarios, specifically, the positions of the background/foreground relative to the robots. For foreground objects, we define a range of possible placements in the simulation. Regarding the background, we leverage the ArUco marker on the table to align the 3DGS coordinates with the marker coordinates. Then the ICP (Schenker, 1992; Censi, 2008; Low, 2004) post-processing is applied to optimize the coordinate alignment gap. Specifically, ICP optimizes the relative transformation

between the partial point cloud obtained from the depth camera and the complete point cloud sampled from the reconstructed mesh.

3.5. Expert Data Collection

To verify the effectiveness of the real-to-sim-to-real system, following the designed task scenarios such as pick-and-place, we collect large-scale demonstrations using the script policy. During each rollout, we introduced domain randomizations including source/target object randomization, and robot arm’s base position randomization. Define an expert policy $\pi_{\text{priv}}(a_t | o_{\text{priv},t})$ based on privileged information $o_{\text{priv},t}$, such as the exact pose of target objects. We can then make use of π_{priv} to interact with the simulator and generate observation-action pairs (o_t, a_t) . Here, o_t represents observations that are accessible from the real world, including images and proprioception data. Specifically, we compute the 6D pose of the gripper at some key steps and use the motion planner based on the RRTConnect (Kuffner & LaValle, 2000) algorithm to plan the path between the key steps in the joint space. During data generation, reject sampling is applied to filter out failed rollouts, improving the dataset quality and ensuring reliability. We can continue to generate such data with the amount according to our needs, obtaining the simulated datasets \mathcal{D} .

4. Experiment

We mainly investigate the following four research questions based on the proposed RE³SIM system:

- Q1** Does RE³SIM produce high-quality and well-aligned reconstruction results?
- Q2** Can RE³SIM generate high-fidelity simulation data with small 3D sim-to-real gaps to benefit real-world manipulation problems?
- Q3** How does large-scale simulation data help in more challenging real-world tasks?
- Q4** Can RE³SIM build the whole scene in simulation in a time-effective way and synthesize data with low cost?

4.1. Experimental Setups

Hardware platform. In this paper, we evaluate RE³SIM with a Franka Research 3 robot equipped with a parallel gripper. Two Realsense D435i depth cameras capture visual observations—one mounted on the end-effector and another positioned beside the robot for a third-person view.

Policy model and training details. During our sim-to-real experiments, we adopt an ACT-structured policy (Zhao et al., 2023) with DINOv2 (Oquab et al., 2023). The images in the dataset are compressed using JPEG format. During training, we apply visual augmentations and for evaluation, we use the temporal ensemble (Zhao et al., 2023) technique. Appendix B provides further details.

4.2. Rendering in Robot Task

Table 1: **Background rendering methods comparison.** Quantitative results of the reconstruction result in terms of PSNR and SSIM. Results are computed from 1029 robot arm positions in the wrist view.

Background Rendering	PSNR	SSIM
Polycam	11.52 ± 1.40	0.34 ± 0.04
OpenMVS	13.40 ± 0.96	0.27 ± 0.03
3DGS	13.29 ± 1.11	0.37 ± 0.04

High-quality visual results not only require high-quality reconstruction but also precise alignment. To address **Q1**, we evaluate the quality and alignment of reconstruction results by selecting a trajectory from the simulation dataset of tasks in Sec. 4.3. This trajectory is then replayed in the real world to obtain ground truth (GT) images captured by calibrated cameras. We compare the visual similarity from both the wrist view and a third-person perspective, showcasing qualitative rendering results in Fig. 3. More rendering results in the simulation are in Sub. E. The high similarity

to real-world images further demonstrates the accuracy of our alignment. Furthermore, Tab. 1 shows the PSNR (peak signal-to-noise ratio) and SSIM (structural similarity index measure) evaluating the quantitative vision gap. We compare RE³SIM with the reconstruction methods OpenMVS and PolyCam (the method used in Torne et al. (2024)). Both methods reconstruct textured meshes, which can be rendered to generate visual observations. The objects are manually aligned with those in the simulation, and some pixel-level discrepancies remain. The background alignment also has some pixel-level deviations. These factors collectively lead to the relatively low PSNR and SSIM values, especially in the texture-rich scene. 3DGS outperforms Polycam in both RSNR and SSIM. Its PSNR is comparable to OpenMVS, but SSIM is notably higher. Qualitative visual comparisons and more analysis are shown in Appendix B. OpenMVS’s reconstruction has cracks, causing an obvious sim-to-real gap. These results demonstrate the advantage of the design choice of the rendering component in RE³SIM, which benefits small vision gaps.

4.3. Zero-Shot Sim-to-Real

Q2 is the main question that motivates us to build RE³SIM, thus we design three table-top manipulation tasks to investigate the usage of RE³SIM:

1. **Pick and drop a bottle into the basket.** The robot must pick an orange bottle that is randomly placed within a 25cm × 35cm area on a table and place it into a basket.
2. **Place a vegetable on the board.** A cucumber model is placed on the table in an area of 35cm × 50cm, and the cutting board’s position is varied within a 35cm × 80cm area. The robot must grasp the vegetable and accurately place it on the cutting board.
3. **Stack blocks.** The robot is tasked with stacking three black cubes on the table. Each cube is placed randomly within a 30cm × 10cm area separately, and the robot must grasp and stack them successfully.

We rebuild these tasks in simulation via RE³SIM, collect 100 trajectories, and train independent policies. More details of the implementation of the rule-based policy and the evaluation methods in simulators are in Sub. A. We evaluate the policy performances in terms of success rate evaluated in the real world, comparing the one trained on simulation data collected by RE³SIM, by a real-to-sim baseline method called RialTo (Torne et al., 2024) and the policy trained on 50 real-world trajectories. Using double simulation data leads to similar performance compared with using real data, probably both because of the small simulation-to-real gap and because the different data collection methods for simulation and real-world data lead to varying data quality, as detailed in Appendix C. We also realized a state-of-the-art grasping




Figure 3: **Visual comparison between real and simulation.** Rendering results from our hybrid rendering method compared with photos captured by real-world cameras, highlighting the high fidelity and realism achieved by our approach.

Table 2: **Quantitative results of real-to-sim-to-real evaluation for real-world manipulation tasks.** Each task is tested 20 times without allowing retries in the event of failure. The policies are trained separately with 100 episodes of simulation data (RialTo+IL and RE³SIM +IL) and 50 episodes of real data (Real+IL). AnyGrasp combined with script primitives (AnyGrasp+Prim) is also tested for comparison in the picking-only scenario. RE³SIM reduces the sim-to-real gap, enabling successful real-world deployment, achieving the best zero-shot sim-to-real performance for an average success rate of 58%.

Real-World Task	RialTo+IL	AnyGrasp+Prim.	Real+IL	RE ³ SIM +IL
Pick and drop a bottle into the basket	0.4	0.9	0.8	0.75
Stack cubes	0	-	0.15	0.25
Place a vegetable on board	0.45	-	0.6	0.75

method, AnyGrasp (Fang et al., 2023), combined with script primitives on the `pick` and `drop` a bottle into the basket task for comparison. This method only predicts the grasp pose instead of the whole grasping trajectory and thus cannot be applied to other tasks.

The numerical results listed in Tab. 2 demonstrate that the data generated by RE³SIM help the policy achieve zero-shot sim-to-real transfer, showing strong performance even compared with the existing real-to-sim and state-of-the-art grasping method. Furthermore, policies trained on our extensive synthetic dataset can achieve comparable or even slightly better performance than those only trained on real-world data, indicating the effectiveness and huge potential of high-fidelity simulation data.

4.4. Large-Scale Sim-to-Real

To push the limit of utilizing synthetic data for real-world manipulation problems by answering **Q3**, we choose a clear objects on the table task: the robot has to clean up the table by picking all items that are randomly placed on the table in a 40cm × 50cm area, and put them into the basket. This is a long-horizon task that requires the

policy to sequentially pick both seen and unseen objects one by one. To provide a sufficient evaluation, we design four different setups:

Seen: four seen objects (the bottle, cucumber, corn, and eggplant) included in the data for training.

Unseen: four objects (the green pepper, banana, red bowl, and momordica charantia) that are unseen during training.

Cluttered: all eight objects from *seen* and *unseen*.

Darkness: the testing brightness is significantly lower than that of the simulation data. The model is evaluated on four *seen* objects.

The results are presented in Fig. 4. Each setup was tested 10 times, with two additional automatic grasp attempts per object in case of failure.

Result analysis. Though trained on only five items, the policy generalizes well to grasping new objects of similar size despite differences in shape and color. We hypothesize that the robot leverages the fixed scene and identifies objects based on color differences in the background. Additionally, objects with varying shapes often have similar grasping positions, enabling the robot to execute successful grasps. A larger dataset with multiple objects allows the policy to

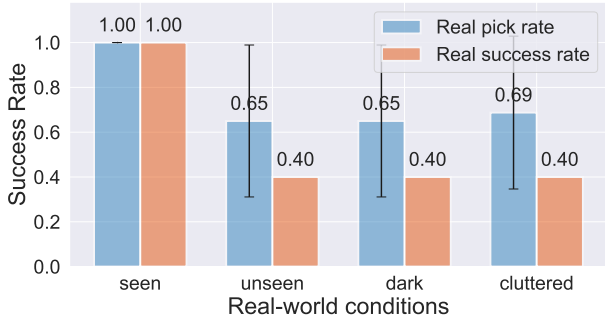


Figure 4: **Real-world evaluation and robustness test for large-scale sim-to-real.** The *success rate* reflects the proportion of trials in which all objects were successfully grasped, while the *grasp rate* indicates the proportion of objects grasped relative to the total number on the table. See qualitative results in the website.




Figure 5: **Data scaling effects**, tested on seen objects in the real world.

maintain a relatively higher success rate, while a smaller dataset limits the policy’s ability to exhibit this capability. Furthermore, the policy demonstrates strong robustness to changes in lighting conditions.

We further explored how the success rate in simulation and real-world settings changes with varying amounts of data. As shown in Fig. 5, increasing the size of the synthetic dataset generated by RE³SIM leads to a significant improvement in imitation learning performance. When the amount of simulation data is limited, the real-world performance is almost negligible. Doubling the data size often results in a significant improvement in success rate until convergence at a high performance.

4.5. Sim-Real Consistency

We evaluated the consistency of policy performance between simulated and real-world environments using policy ACT. In this experiment, we evaluate policies trained at different stages (from 0 to 120000 training steps) and with

Table 3: **Human effort in reconstruction.** The table presents estimated reconstruction times at the table level. Additionally, we show the human effort for reconstructing an object with ARCode.

Input Types	Video	Images	ARCode
Human Efforts (s)	51.5	84.5	60.5

3 different seeds on the pick and drop a bottle into a basket task.

As illustrated in Fig. 1 (d), the Pearson correlation coefficient is 0.924, indicating the clear correlation between performance in simulation and the real world. Specifically, models that achieve higher success rates in simulation tend to exhibit similar success rates in real-world testing. When the success rate in simulation is low, it is often difficult for the model to perform well in real-world scenarios. This could be due to alignment errors, leading to a small sim-to-real gap, or the fact that the number of real-world test instances is much lower than in simulation, potentially introducing bias. This consistency suggests that policies trained with RE³SIM are well-suited for real-world deployment, reflecting a minimal sim-to-real gap.

4.6. Pipeline Efficiency

To validate Q4, we record and compute the time of RE³SIM to reconstruct both the background and objects, along with the time of generating demonstrations.

Human effort during reconstruction. In RE³SIM, background reconstruction requires only a set of images or a short video of the background, along with an additional image for alignment, which contains depth information. Given the substantial human involvement in this process, we provide estimates for the time required for both foreground and background reconstruction. Tab. 3 summarizes the time required for different input types: video and images. Using video as input may result in slight quality degradation due to motion blur while using images provides better results but requires slightly more time. Additionally, we present the time required to reconstruct an object using ARCode. The total reconstruction time for the scene is the sum of the background and object reconstruction times.

Data collection cost. Tab. 4 shows the time required to collect 100 simulation episodes for each task in Sec. 4.3, using 8 RTX 4090 GPUs. The time required for data collection is much lower than the time needed for teleoperation in real-world scenarios—especially considering RE³SIM only involves machine runtime, and the latter mostly requires the time of data collection experts. This demonstrates our capability to efficiently generate large-scale simulation data

Table 4: **Time cost for simulation data collection.** Time needed to collect 100 episodes of simulation data for each task, using a machine equipped with 8 RTX 4090 GPUs. Each GPU can run two Isaac Sim processes.

Tasks	Time Cost (minutes)
Pick and drop a bottle into the basket	12.35
Place a vegetable on the board	13.78
Stack blocks	6.45

Table 5: **Average time consumption per step.** Breakdown of time consumed by each process.

Process	Time (ms)
Physics Simulating	26.64
Gaussian Splatting Rendering	12.93
Motion Planning	0.36
Others	1.53
Total Time	41.46

at minimal cost.

Time consumption breakdown. We measured the average time spent on each phase during the data collection of the *pick and drop a bottle* task in the simulator over 60 steps, as shown in Tab. 5. The reconstructed mesh, with numerous vertices and faces, demands more time for physics simulation. Rendering 480p images from two camera views using 3DGS adds 12.93 milliseconds, representing one-third of the total time per step. RE³SIM operates at approximately 24 frames per second (FPS) with two cameras, ensuring real-time performance and visual fidelity.

5. Related Work

Sim-to-real. Sim-to-real transfer requires techniques to enable policies to successfully adapt from simulation to the real world. The most direct approach is improving simulators (Todorov et al., 2012; Makoviychuk et al., 2021; Mittal et al., 2023; Xiang et al., 2020), which reduces the sim-to-real gap. Other methods, such as domain randomization (Tobin et al., 2017; Mehta et al., 2020; Chen et al., 2021; Tremblay et al., 2018; Loquercio et al., 2019; Tobin et al., 2018) and system identification (Ramaswamy & Tangirala, 2024; Allevato et al., 2020b; Song et al., 2024; Allevato et al., 2020a), also aim to bridge this gap.

Real-to-sim-to-real. Many recent works leverage real-world data to enhance simulation models. Reconstruction methods integrated with grasping techniques, such as EvoNeRF (Kerr et al., 2023) and LERF-TOGO (Rashid et al., 2023), enable the grasping of objects using only RGB images. Meanwhile, GaussianGrasper (Zheng et al., 2024),

SplatMover (Shorinwa et al., 2024), and GraspSplats (Ji et al., 2024), use 3D Gaussian splatting (Kerbl et al., 2023) for fast rendering and explicit representation in robotics tasks. Additionally, methods like URDFFormer (Chen et al., 2024), Digital Cousins (Dai et al., 2024), and Articulate Anything (Le et al., 2024) use a single image to reconstruct the environment directly, allowing for collecting large amounts of data for imitation learning or reinforcement learning. These approaches enhance data by varying articulations and leveraging various articulations from simulation datasets to train models deployable in the real world.

With the help of 3DGS works like RoboStudio (Lou et al., 2024), SplatSim (Qureshi et al., 2024) and RoboGSim (Li et al., 2024a) use multi-view images or video for world reconstruction, improving multi-view rendering quality with minimal cost. These methods are especially effective in manipulation tasks involving multiple objects or occlusions. 1) RoboStudio focuses on reconstructing the URDF of a robot, offering a photorealistic rendering result and accurate collision mesh. 2) SplatSim utilizes pre-obtained 3D models of objects and backgrounds to collect trajectories in the physics simulator and then re-render them with 3DGS to reduce the visual gap between simulated and real-world environments, but acquiring the 3D models is hard for many tasks. 3) RoboGSim compares rendering quality, validates high-quality rendering results for novel pose synthesis, and shows the potential for evaluating various manipulation algorithms. However, its sim-to-real validation remains limited. Different from them, RE³SIM reconstructs both geometric and visual aspects with small gaps, and validates robot policies trained on simulated data through extensive experiments in the real world.

6. Conclusion

We introduced RE³SIM, a novel Real-to-Sim-to-Real pipeline that integrates Gaussian splatting with NVIDIA Isaac Sim’s PhysX engine, improving scene reconstruction and sim-to-real transfer for robotic manipulation tasks. Extensive experiments demonstrate its efficacy and scalability, showing that RE³SIM significantly reduces the sim-to-real gap and enables reliable real-world task performance. Models trained on large simulated datasets achieved competitive success rates compared with those trained on real-world data, highlighting the potential RE³SIM to generate diverse, high-quality data for pre-training large-scale robot models.

Limitations. RE³SIM is currently limited to rigid objects and has not yet been extended to reconstruct articulated, deformable objects, or liquids. Additionally, it relies on manually defined physics parameters rather than system identification methods. Rule-based policies for data collection become challenging as task complexity increases. We leave addressing these limitations to future work.

Impact Statement

The proposed RE³SIM system provides an intriguing way of generating simulated robot demonstrations for real-world tasks. This work is meant to release the human burden of collecting large-scale real-world data, improve efficiency, and push the limit of generalist robot policies with data scaling law. RE³SIM can be applied across various robot applications and thus may cause potential unemployment issues.

References

- Aldaco, J., Armstrong, T., Baruch, R., Bingham, J., Chan, S., Draper, K., Dwibedi, D., Finn, C., Florence, P., Goodrich, S., et al. Aloha 2: An enhanced low-cost hardware for bimanual teleoperation. *arXiv preprint arXiv:2405.02292*, 2024.
- Allevato, A., Short, E. S., Pryor, M., and Thomaz, A. Tunenet: One-shot residual tuning for system identification and sim-to-real robot task transfer. In Kaelbling, L. P., Kragic, D., and Sugiura, K. (eds.), *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pp. 445–455. PMLR, 30 Oct–01 Nov 2020a. URL <https://proceedings.mlr.press/v100/allevato20a.html>.
- Allevato, A. D., Schaertl Short, E., Pryor, M., and Thomaz, A. L. Iterative residual tuning for system identification and sim-to-real robot learning. *Autonomous Robots*, 44 (7):1167–1182, 2020b.
- Black, K., Brown, N., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Groom, L., Hausman, K., Ichter, B., et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- Blinn, J. F. and Newell, M. E. Texture and reflection in computer generated images. *Commun. ACM*, 19(10): 542–547, October 1976. ISSN 0001-0782. doi: 10.1145/360349.360353. URL <https://doi.org/10.1145/360349.360353>.
- Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Chen, X., Choromanski, K., Ding, T., Driess, D., Dubey, A., Finn, C., et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- Buslaev, A., Iglovikov, V. I., Khvedchenya, E., Parinov, A., Druzhinin, M., and Kalinin, A. A. AlbuMentations: Fast and flexible image augmentations. *Information*, 11 (2), 2020. ISSN 2078-2489. doi: 10.3390/info11020125. URL <https://www.mdpi.com/2078-2489/11/2/125>.
- Censi, A. An icp variant using a point-to-line metric. In *2008 IEEE international conference on robotics and automation*, pp. 19–25. Ieee, 2008.
- Cernea, D. OpenMVS: Multi-view stereo reconstruction library. 2020. URL <https://cdcseacave.github.io/openMVS>.
- Cheang, C.-L., Chen, G., Jing, Y., Kong, T., Li, H., Li, Y., Liu, Y., Wu, H., Xu, J., Yang, Y., Zhang, H., and Zhu, M. Gr-2: A generative video-language-action model with web-scale knowledge for robot manipulation. *arXiv preprint arXiv:2410.06158*, 2024.
- Chen, X., Hu, J., Jin, C., Li, L., and Wang, L. Understanding domain randomization for sim-to-real transfer. *arXiv preprint arXiv:2110.03239*, 2021.
- Chen, Z., Walsman, A., Memmel, M., Mo, K., Fang, A., Vemuri, K., Wu, A., Fox, D., and Gupta, A. Urdformer: A pipeline for constructing articulated simulation environments from real-world images. *arXiv preprint arXiv:2405.11656*, 2024.
- Cheng, X., Li, J., Yang, S., Yang, G., and Wang, X. Open-television: Teleoperation with immersive active visual feedback. *arXiv preprint arXiv:2407.01512*, 2024.
- Code, A. Ar code. <https://ar-code.com/>, 2022.
- Coumans, E. and Bai, Y. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.
- Dai, T., Wong, J., Jiang, Y., Wang, C., Gokmen, C., Zhang, R., Wu, J., and Fei-Fei, L. Acdc: Automated creation of digital cousins for robust policy learning. *arXiv e-prints*, pp. arXiv–2410, 2024.
- Fang, H.-S., Wang, C., Fang, H., Gou, M., Liu, J., Yan, H., Liu, W., Xie, Y., and Lu, C. Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on Robotics*, 2023.
- Fu, Z., Zhao, T. Z., and Finn, C. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *Conference on Robot Learning (CoRL)*, 2024.
- Guédon, A. and Lepetit, V. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, pp. 5354–5363, 2024.
- Ji, M., Qiu, R.-Z., Zou, X., and Wang, X. Graspplats: Efficient manipulation with 3d feature splatting. *arXiv preprint arXiv:2409.02084*, 2024.
- Jiang, Z., Xie, Y., Lin, K., Xu, Z., Wan, W., Mandlekar, A., Fan, L., and Zhu, Y. Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning. 2024.
- Kerbl, B., Kopanas, G., Leimkühler, T., and Drettakis, G. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023. URL <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>.
- Kerr, J., Fu, L., Huang, H., Avigal, Y., Tancik, M., Ichnowski, J., Kanazawa, A., and Goldberg, K. Evo-nerf: Evolving nerf for sequential robot grasping of transparent objects. In *Conference on Robot Learning*, pp. 353–367. PMLR, 2023.
- Khazatsky, A., Pertsch, K., Nair, S., Balakrishna, A., Dasari, S., Karamcheti, S., Nasiriany, S., Srirama, M. K., Chen, L. Y., Ellis, K., Fagan, P. D., Hejna, J., Itkina, M., Lepert, M., Ma, Y. J., Miller, P. T., Wu, J., Belkhale, S., Dass, S., Ha, H., Jain, A., Lee, A., Lee, Y., Memmel, M., Park, S., Radosavovic, I., Wang, K., Zhan, A., Black, K., Chi, C., Hatch, K. B., Lin, S., Lu, J., Mercat, J., Rehman, A., Sankeketi, P. R., Sharma, A., Simpson, C., Vuong, Q., Walke, H. R., Wulfe, B., Xiao, T., Yang, J. H., Yavary, A., Zhao, T. Z., Agia, C., Bajjal, R., Castro, M. G., Chen, D., Chen, Q., Chung, T., Drake, J., Foster, E. P., Gao, J., Herrera, D. A., Heo, M., Hsu, K., Hu, J., Jackson, D., Le, C., Li, Y., Lin, K., Lin, R., Ma, Z., Maddukuri, A., Mirchandani, S., Morton, D., Nguyen, T., O’Neill, A., Scalise, R., Seale, D., Son, V., Tian, S., Tran, E., Wang, A. E., Wu, Y., Xie, A., Yang, J., Yin, P., Zhang, Y., Bastani, O., Berseth, G., Bohg, J., Goldberg, K., Gupta, A., Gupta, A., Jayaraman, D., Lim, J. J., Malik, J., Martín-Martín, R., Ramamoorthy, S., Sadigh, D., Song, S., Wu, J., Yip, M. C., Zhu, Y., Kollar, T., Levine, S., and Finn, C. Droid: A large-scale in-the-wild robot manipulation dataset, 2024. URL <https://arxiv.org/abs/2403.12945>.
- Kim, M. J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E., Lam, G., Sankeketi, P., et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- Koenig, N. and Howard, A. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, volume 3, pp. 2149–2154 vol.3, 2004. doi: 10.1109/IROS.2004.1389727.
- Kuffner, J. and LaValle, S. Rrt-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, volume 2, pp. 995–1001 vol.2, 2000. doi: 10.1109/ROBOT.2000.844730.
- Le, L., Xie, J., Liang, W., Wang, H.-J., Yang, Y., Ma, Y. J., Vedder, K., Krishna, A., Jayaraman, D., and Eaton, E. Articulate-anything: Automatic modeling of articulated objects via a vision-language foundation model. *arXiv preprint arXiv:2410.13882*, 2024.
- Lee, J., X. Grey, M., Ha, S., Kunz, T., Jain, S., Ye, Y., S. Srinivasa, S., Stilman, M., and Karen Liu, C. Dart: Dynamic animation and robotics toolkit. *The Journal of Open Source Software*, 3(22):500, 2018.
- Li, X., Liu, M., Zhang, H., Yu, C., Xu, J., Wu, H., Cheang, C., Jing, Y., Zhang, W., Liu, H., et al. Vision-language foundation models as effective robot imitators. *arXiv preprint arXiv:2311.01378*, 2023.
- Li, X., Li, J., Zhang, Z., Zhang, R., Jia, F., Wang, T., Fan, H., Tseng, K.-K., and Wang, R. Robogsim: A real2sim2real robotic gaussian splatting simulator. *arXiv preprint arXiv:2411.11839*, 2024a.
- Li, X., Li, P., Liu, M., Wang, D., Liu, J., Kang, B., Ma, X., Kong, T., Zhang, H., and Liu, H. Towards generalist robot policies: What matters in building vision-language-action models. *arXiv preprint arXiv:2412.14058*, 2024b.
- Loquercio, A., Kaufmann, E., Ranftl, R., Dosovitskiy, A., Koltun, V., and Scaramuzza, D. Deep drone racing: From simulation to reality with domain randomization. *IEEE Transactions on Robotics*, 36(1):1–14, 2019.
- Lou, H., Liu, Y., Pan, Y., Geng, Y., Chen, J., Ma, W., Li, C., Wang, L., Feng, H., Shi, L., et al. Robo-gs: A physics consistent spatial-temporal model for robotic arm with hybrid representation. *arXiv preprint arXiv:2408.14873*, 2024.
- Low, K.-L. Linear least-squares optimization for point-to-plane icp surface registration. *Chapel Hill, University of North Carolina*, 4(10):1–3, 2004.
- Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., and State, G. Isaac gym: High performance gpu-based physics simulation for robot learning, 2021.
- Mandlekar, A., Nasiriany, S., Wen, B., Akinola, I., Narang, Y., Fan, L., Zhu, Y., and Fox, D. Mimicgen: A data

- generation system for scalable robot learning using human demonstrations. In *7th Annual Conference on Robot Learning*, 2023.
- Mehta, B., Diaz, M., Golemo, F., Pal, C. J., and Paull, L. Active domain randomization. In *Conference on Robot Learning*, pp. 1162–1176. PMLR, 2020.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- Mittal, M., Yu, C., Yu, Q., Liu, J., Rudin, N., Hoeller, D., Yuan, J. L., Singh, R., Guo, Y., Mazhar, H., Mandlekar, A., Babich, B., State, G., Hutter, M., and Garg, A. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023. doi: 10.1109/LRA.2023.3270034.
- Nasiriany, S., Maddukuri, A., Zhang, L., Parikh, A., Lo, A., Joshi, A., Mandlekar, A., and Zhu, Y. Robocasa: Large-scale simulation of everyday tasks for generalist robots. In *Robotics: Science and Systems (RSS)*, 2024.
- NVIDIA. Nvidia isaac sim, 2021a. URL <https://developer.nvidia.com/isaac-sim>.
- NVIDIA. Nvidia physx, 2021b. URL <https://nvidia-omniverse.github.io/PhysX/>.
- O’Neill, A., Rehman, A., Gupta, A., Maddukuri, A., Gupta, A., Padalkar, A., Lee, A., Pooley, A., Gupta, A., Mandlekar, A., et al. Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864*, 2023.
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- Polycam. Polycam. <https://polycam.com>, 2020.
- Qureshi, M. N., Garg, S., Yandun, F., Held, D., Kantor, G., and Silwal, A. SplatSim: Zero-shot sim2real transfer of rgb manipulation policies using gaussian splatting. *arXiv preprint arXiv:2409.10161*, 2024.
- Ramaswamy, J. and Tangirala, A. K. Enhancing system identification through transfer learning in gaussian process models: Bridging sim-to-real and cross-environment applications. In *2024 SICE Festival with Annual Conference (SICE FES)*, pp. 806–811, 2024.
- Rashid, A., Sharma, S., Kim, C. M., Kerr, J., Chen, L. Y., Kanazawa, A., and Goldberg, K. Language embedded radiance fields for zero-shot task-oriented grasping. In *7th Annual Conference on Robot Learning*, 2023. URL <https://openreview.net/forum?id=k-Fg8JDQmc>.
- Schenker, P. S. *Sensor fusion iv: Control paradigms and data structures; proceedings of the meeting, boston, ma, nov. 12-15, 1991*. Number SPIE-1611. Society of Photo-Optical Instrumentation Engineers (SPIE Proceedings. Vol. 1611), 1992.
- Schönberger, J. L. and Frahm, J.-M. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Schönberger, J. L., Zheng, E., Pollefeys, M., and Frahm, J.-M. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016.
- Sherman, M. A., Seth, A., and Delp, S. L. Simbody: multi-body dynamics for biomedical research. *Procedia Iutam*, 2:241–261, 2011.
- Shorinwa, O., Tucker, J., Smith, A., Swann, A., Chen, T., Firoozi, R., Kennedy, M. D., and Schwager, M. Splat-mover: Multi-stage, open-vocabulary robotic manipulation via editable gaussian splatting. 2024.
- Smith, R. L. Open dynamics engine (ode). <https://ode.org>, 2001.
- Song, R., Gao, S., and Li, Y. Sim-to-real in unmanned surface vehicle control: A system identification-based approach for enhanced training environments. In *2024 9th International Conference on Electronic Technology and Information Science (ICETIS)*, pp. 563–570, 2024. doi: 10.1109/ICETIS61828.2024.10593794.
- Tian, Y., Yang, S., Zeng, J., Wang, P., Lin, D., Dong, H., and Pang, J. Predictive inverse dynamics models are scalable learners for robotic manipulation. *arXiv preprint arXiv:2412.15109*, 2024.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 23–30. IEEE, 2017.
- Tobin, J., Biewald, L., Duan, R., Andrychowicz, M., Handa, A., Kumar, V., McGrew, B., Ray, A., Schneider, J., Welinder, P., et al. Domain randomization and generative models for robotic grasping. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3482–3489. IEEE, 2018.

- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
- Torne, M., Simeonov, A., Li, Z., Chan, A., Chen, T., Gupta, A., and Agrawal, P. Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation. *Arxiv*, 2024.
- Tremblay, J., Prakash, A., Acuna, D., Brophy, M., Jampani, V., Anil, C., To, T., Cameracci, E., Boochoon, S., and Birchfield, S. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 969–977, 2018.
- Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., Yi, L., Chang, A. X., Guibas, L. J., and Su, H. SAPIEN: A simulated part-based interactive environment. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Yang, S., Liu, M., Qin, Y., Runyu, D., Jialong, L., Cheng, X., Yang, R., Yi, S., and Wang, X. Ace: A cross-platform visual-exoskeletons for low-cost dexterous teleoperation. *arXiv preprint arXiv:240*, 2024.
- Yao, Y., Luo, Z., Li, S., Fang, T., and Quan, L. Mvsnet: Depth inference for unstructured multi-view stereo. *European Conference on Computer Vision (ECCV)*, 2018.
- Yu, M., Lu, T., Xu, L., Jiang, L., Xiangli, Y., and Dai, B. Gsdg: 3dgs meets sdf for improved rendering and reconstruction. *arXiv preprint arXiv:2403.16964*, 2024.
- Zhao, T. Z., Kumar, V., Levine, S., and Finn, C. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- Zheng, Y., Chen, X., Zheng, Y., Gu, S., Yang, R., Jin, B., Li, P., Zhong, C., Wang, Z., Liu, L., et al. Gaussiangrasper: 3d language gaussian splatting for open-vocabulary robotic grasping. *arXiv preprint arXiv:2403.09637*, 2024.

The appendix is organized as follows:

- Sec. A illustrates the detailed design of the rule-based policy and the evaluation method of the tasks in Sec. 4.3 and Sec. 4.4.
- Sec. B illustrates the details of the policy training and image augmentation.
- Sec. C explains the differences in properties between simulated and real data, providing insights into why policies trained on these datasets exhibit different performance characteristics.
- Sec. E presents additional rendering results of various tasks from different camera angles.

A. Task Details

This section describes the design of the rule-based policy and the evaluation method for each task.

Pick and drop a bottle into the basket. This task consists of two steps: 1) determining the 6D pose of the gripper before picking, and 2) identifying the 6D pose of the placement position. The pre-pick pose is derived from the object’s point cloud, while the placement position is fixed. Motion planning in joint space guides the robotic arm from the starting position to the two key positions for picking or releasing the object. Picking success depends on the positional relationship between the robotic arm, the gripper, and the object’s height relative to the table. Upon release, the object is considered successfully placed if its horizontal position lies within the basket’s area.

Placing a vegetable on the board The board’s center serves as the placement position. The rule-based policy follows a similar approach to the previous tasks. Success is achieved only if the object’s horizontal position during placement lies within a small region near the board’s center.

Stack blocks. This task involves sequentially picking and placing two objects. Using the objects’ 3D models, the pre-pick pose and placement position are calculated. The implementation resembles the `pick and drop a bottle into the basket` task. The second block is picked only after the first is successfully placed; otherwise, the trajectory is terminated early and discarded. Stacking is considered successful if the horizontal positions of the two objects align and their height difference equals the height of one block.

Clear objects on the table. To minimize action variability in identical states and simplify model training, the rule-based policy prioritizes picking objects closer to the robot base. Due to the large number of objects, randomizing their initial positions often results in overlapping point clouds, leading to unpredictable physics simulation outcomes. Therefore, during domain randomization, configurations with intersecting objects are filtered out before simulation. The task is considered complete only when all objects are successfully picked and placed in their designated positions.

B. More Experiment Details

Qualitative comparison of background rendering methods. We collected a trajectory using teleoperation in the simulator and replayed it in the real world. Visual comparisons of reconstruction results from various methods are shown in Figure 6. PolyCam tends to blur details in certain areas, OpenMVS’s reconstructed mesh texture shows cracks, while 3DGS produces high-quality renderings. Table 1 shows that OpenMVS’s PSNR is slightly higher than 3DGS’s, likely because the

Table 6: Hyper-parameters for Data Augmentation.

Params	Value
Batch size	8
Epochs	100
Learning rate for warmup steps	1000
Kl weight	10
Hidden dim	512
Dim feedforward	3200



Figure 6: **Visual comparison of rendering approaches.** Rendering results of reconstruction outputs from PolyCam, OpenMVS, and 3DGS, compared with real-world photos.

Table 7: **Ablation on data augmentation methods.** Success rates (SR) of the single-item picking task in both simulated and real-world environments for different image augmentation techniques.

Augmentation	Sim SR	Real SR
-	0.77	0.25
+ Gaussian Blur	0.45	0.4
+ Defocus	0.97	0.6
+ ColorJitter	0.37	0.6
+ Gaussian Noise	0.94	0.8

background and crack colors are similar, making PSNR less sensitive. However, the cracks cause increased variance within patches, leading to a lower SSIM compared to 3DGS.

Policy Training. The policy predicts a sequence of future k -step actions from single-step observations, including two-view images and proprioception data. To better capture information from the images, we replace the original ResNet architecture with DINOv2 (Oquab et al., 2023) small as the model’s encoder. During training, we apply visual augmentations, including *Gaussian Blur*, *Defocus*, *Color Jittering*, and *Gaussian Noise*, to enhance robustness. Additionally, we apply cosine decay to the learning rate (lr) to facilitate easier convergence during model training. The hyperparameters related to model training for tasks in Sec. 4.3 are listed in Table 6, and because we use a much larger dataset for the `Clear` objects on the `table` task, we train the policy only for 8 epochs. During the evaluation, the temporal ensemble (Zhao et al., 2023) technique is used to smooth the final action for rollouts.

Image augmentation. Recognizing that Gaussian splatting inherently cannot replicate certain visual effects such as motion blur and focus-dependent blur, however, due to camera motion, photos taken during the actual deployment of the device may appear blurred. Additionally, the brightness of the room may change over time. To simulate this phenomenon during model training, we apply random transformations to the images using the implementation from (Buslaev et al., 2020), including *Gaussian Blur*, *Decocus*, *Color Jitter* and *Gaussian Noise*.

As presented in Tab. 7, applying appropriate image augmentation techniques enhances model performance in both simulated and real-world environments. Specifically, augmentations such as blur and defocus simulate the blurring caused by camera motion. Adding noise and blur together may help the model learn key information for completing tasks, reducing sensitivity to changes in irrelevant information and enabling more robust performance in real-world scenarios. Using these image

augmentation methods can achieve the highest success rate in real-world environments, indicating their effectiveness in bridging the sim-to-real gap between 3D Gaussian splatting rendering results and real-world camera outputs.

Evaluation in simulation. To employ rejection sampling during data collection and evaluate the training policy, a mechanism is required to determine whether the task has been successfully completed. The relationship between the gripper and the object’s point cloud, as well as the distance between the object and the tabletop, can help verify if the object has been successfully grasped. Furthermore, the relative position between the object to be grasped and the target can be used to confirm whether the item has been placed in the correct position.




Figure 7: **Compression analysis.** Qualitative and quantitative analysis using third-person view images to demonstrate the impact of different JPEG quality parameters on image quality, task success rate, and compression ratio in the `place a vegetable on the board` task. The policies are trained on compressed data, but uncompressed data is used during deployment.

Data saving format. For complex tasks, we require a large amount of simulation data to achieve satisfactory results. Storing image data without any compression is costly in terms of space. We compress the images using the JPEG format. During deployment, uncompressed data is used to retain more information. Fig. 7 shows the image quality at different quality settings, the success rate of models trained on the `place a vegetable on the board` task, and the corresponding compression rates. In our experiments, we chose a quality setting of 40, as it results in almost no perceptible loss in quality while saving a significant amount of storage space.

C. Comparison over Simulated and Real Data

We train the policy with 100 episodes of simulation data and 50 episodes of real-world data, achieving comparable performance in the real-world setting. In this section, we analyze the differences and similarities between simulated and real data, exploring the underlying causes. Real-world and simulation data often exhibit variations in both distribution and quality, because of differences in scene initialization methods and trajectory preferences between human operators and the rule-based policy.

C.1. Object Location

In simulation, we define a rectangular region and use the standard library’s random implementation to achieve an almost uniform distribution. In the real world, random initialization can only rely on operator estimation. In Fig. 8, we visualize the initial object position distributions from simulation data (100 episodes) and real-world data (50 episodes) used for training. The visualization reveals that despite operators consciously attempting to randomize object positions within the corresponding area, the data distributions show slight differences. Specifically, we find that in the `pick and drop a bottle into the basket` task, the real-world data shows a wider distribution, while in the other two tasks, the simulation data has a slightly broader spread. As seen in Table 2, tasks with a broader distribution tend to yield better performance.




Figure 8: **Data distribution.** We show the initial locations of objects in each task described in Sec. 4.3. There are 100 episodes of simulation data(Sim) and 50 episodes of real-world data(Real). We show each cube’s location in `stack blocks` tasks separately.

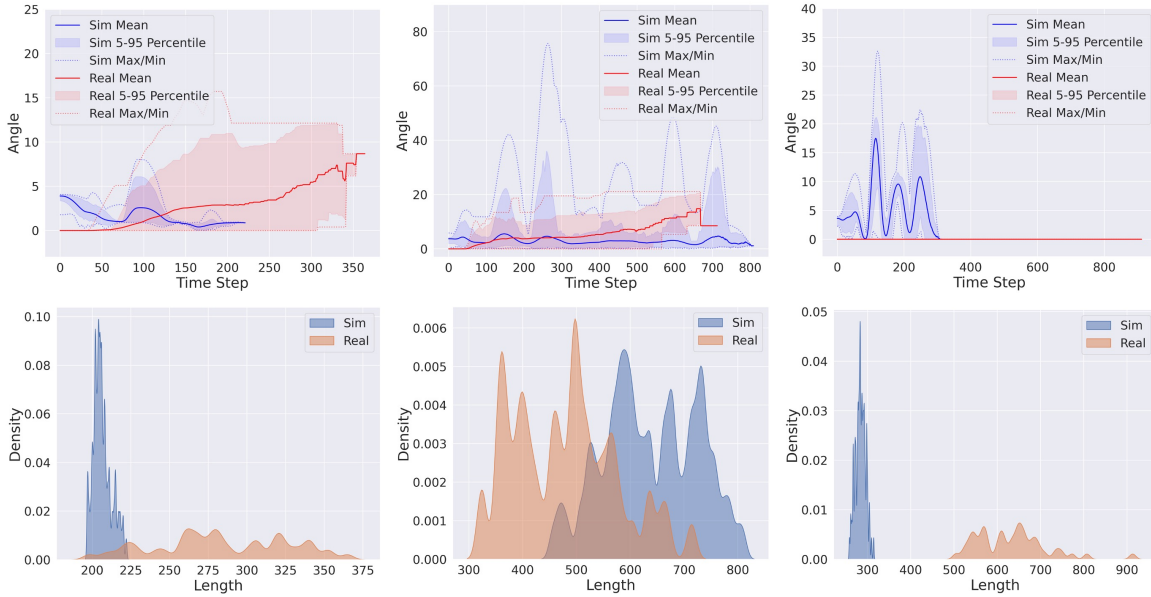


Figure 9: **Data quality.** Analysis of differences between real-world and simulation data properties. The first row presents statistical data on the distribution of gripper angles per time step across episodes. The second row displays the Kernel Density Estimate (KDE) of episode lengths.

C.2. Data Quality

In Fig. 9, we visualize the distribution of episode lengths and gripper angles relative to the gravity direction. Due to different data collection methods, there are significant differences in trajectory lengths. In simulation, the motion planner tends to take the shortest path, resulting in shorter trajectories but with larger angular variations. Additionally, we empirically observe that the motion planner simultaneously adjusts both the gripper’s rotation and position, which rarely occurs in real-world teleoperation using a space mouse. When reconstructing the scene, we did not capture images from positions with large angles, which led to reduced rendering realism in high-angle simulations. This could lead to the need for more simulated data to compensate for these discrepancies. On the other hand, longer trajectories may include more pauses, which can negatively impact model training due to reduced action continuity.

D. Co-training and Fine-tuning

We tested the effectiveness of co-training and pretrain-finetune on the `Stack blocks` task. Specifically, we report the speed characteristics during model execution. As shown in Figure 10(a), the policy trained on simulation data exhibits




Figure 10: **Features of co-training and fine-tuning models.** (a) KDE of the Euclidean distance traveled by the robotic arm’s end effector between adjacent time steps. (b) The number of time steps taken by the robotic arm from the start of movement to the first closure of the gripper. ”Sim” and ”Real” indicate models trained on simulated and real data, respectively, while ”Co-train” and ”Fine-tune” refer to models trained on a mix of data and pre-trained with real data, respectively.

relatively fast speeds for many time steps, while the model trained on real data mostly demonstrates slower speeds. The co-training model lies between these two extremes. After fine-tuning with real data, the model’s speed decreases a lot compared to the simulation model but remains slightly faster than the real data-trained model. Since only real data was used for fine-tuning, the final model is closer to the real data-trained model than the co-training model. Figure 10(b) illustrates the number of time steps required for the robotic arm to move from the start to the first closure of the gripper, reflecting the average speed of the arm’s motion. The models trained with real data or fine-tuning with real data are relatively slow, while models trained with simulation data or jointly trained with simulation data exhibit a higher speed.

These suggest that the distribution of simulation and real data is generally similar and that the data generated by our method can be integrated into real data through pretraining or co-training, introducing new features without causing the training process to collapse.

E. Result Visualization

OpenMVS is also capable of providing a textured mesh for acquiring visual observations. A qualitative comparison of the rendering quality between RE³SIM and OpenMVS is illustrated in Figure 11. It is evident that, powered by 3DGS, RE³SIM achieves superior visual rendering outcomes. This rationale explains why we do not directly utilize the results from OpenMVS for obtaining visual rendering results.

We already show the rendering results of `pick and drop a bottle into the basket` task in Fig. 3. Here we present multi-view renderings of the data collection process across the other three `sim-to-real` tasks: `place a vegetable on the board`, `stack blocks`, and `clear objects on the table` in Fig. 12, Fig. 13 and Fig. 14 separately. These visualizations provide a detailed demonstration of our data collection procedure.




Figure 11: Visual comparison between RE³SIM and OpenMVS.




Figure 12: Rendering results of place a vegetable on the board task.




Figure 13: Rendering results of stack blocks task.




Figure 14: Rendering results of clear objects on the table task.