

Worlds in One Demo: A Synthetic Data Engine for Learning Open-World Mobile Manipulation

Lingxiao Guo* Huanyu Li* Guanya Shi

Carnegie Mellon University

*Equal contribution; order decided by a coin flip.




Figure 1: WANDA enables mobile manipulation policies achieve extreme spatial generalization, long horizon robustness, cross-environment generalization and cross-embodiment generalization from **only one human demonstration.** This is achieved by the proposed data engine that synthesizes diverse trajectories across reconstructed worlds and generated worlds on different robot embodiments.

Correspondence: {lingxiag, huanyuli, guanyas}@andrew.cmu.edu

Website: <https://wanda.lecar-lab.org>

Abstract

Learning open-world mobile manipulation policies requires vast data to achieve spatial generalization, long-horizon robustness, and scene generalization. Current prevailing data collection paradigms, teleoperation and UMI, demand prohibitive human effort and cost at scale. To

scale beyond the limits of manual data collection, we seek to maximize the value of each human demonstration by scalable data generation. To this end, we introduce **WANDA**: learning open-World mobile mANipulation from one demonstration via a synthetic **DA**ta engine. **WANDA** first reconstructs background Gaussian splats and robot-object interaction trajectories from source RGBD observations, as a world substrate for later planning and rendering. It then rearranges contact-rich robot-object interaction segments into extensive spatial configurations, utilizing whole-body motion planning to chain them into new trajectories. To enhance long-horizon robustness, it applies Corrective State Expansion to increase the robot and object state diversity at different stages of mobile manipulation. To unlock cross-environment generalization, trajectories are synthesized on diverse generated 3D worlds from everyday photos. Furthermore, we synthesize photo-realistic observations by compositing rendered robot and object meshes with Gaussian splatting backgrounds. We evaluate our approach on extensive simulation and real-world tasks in various scenes. Experiments show that policies trained with **WANDA** achieve long-horizon robustness, broad spatial generalization and cross-environment generalization from one real demonstration. Moreover, **WANDA** naturally supports cross-embodiment data generation, validated by zero-shot deployment on another mobile manipulator with a distinct morphology.

1 Introduction

Achieving mobile manipulation in unstructured, open-ended environments, like homes, cafes and offices, has been a long-standing goal in the robotics community. Learning foundation models end-to-end from data has shown great promise in robot manipulation [Physical Intelligence et al., 2025, 2026, Fang et al., 2026, Kim et al., 2024, Ye et al., 2026, Pertsch et al., 2025]. However, learning open-world mobile manipulation requires substantial data [Sunday Team, 2025] due to its challenging nature: a mobile manipulator has more degrees of freedom and a much larger operation space than tabletop settings. It requires the policy to have strong spatial generalization across diverse configurations [Physical Intelligence et al., 2025, Yang et al., 2025], be robust to long-horizon compounding errors [BEHAVIOR Team, 2025], and generalize across diverse environments. This necessitates thousands of demonstrations to learn a single task.

Current prevailing data collection paradigms, teleoperation [Cheng et al., 2024, He et al., 2024, Fu et al., 2024a] and UMI [Chi et al., 2024, Zhaxizhuoma et al., 2025, Xu et al., 2026, Liu et al., 2026], require great human effort to collect data at scale. Collecting mobile manipulation data is even more costly. It requires humans to bring physical hardware into diverse environments and repeat the same tasks across different spaces. UMI and human data [Punamiya et al., 2026, Kareer et al., 2025, Zheng et al., 2026] also require precise odometry and localization techniques to capture accurate poses, which is challenging in the noisy open world.

Another line of work uses synthetic data to train robot policies [Mandlekar et al., 2023, Jang et al., 2025, Xue et al., 2025]. MimicGen [Mandlekar et al., 2023] and its subsequent extensions [Xue et al., 2025, Jiang et al., 2025b, Lin et al., 2026, Garrett et al., 2024] have proposed to replace the tedious relocate-and-recollect procedure with automatic demonstration generation. MoMaGen [Li et al., 2025a] extends it to the mobile manipulation setting via whole-body motion planning, but is still limited to simulation. Due to the sim-to-real gap, they still need real data for sim-to-real co-training [Mandlekar et al., 2023, Lin et al., 2026, Li et al., 2025a] to improve real-world performance. Some work extends the MimicGen-style approach directly to the real-world setting [Xue et al., 2025, Xu et al., 2025]. But most of them operate in a sparse point-cloud observation space, which is unsuitable for training large robot foundation models [Physical Intelligence et al., 2025, Zhen et al., 2024]. Moreover, none of these methods can synthesize diverse demonstrations for open-world mobile manipulation.

In this work, we argue that the potential of a single demonstration is not fully leveraged for policy learning. For mobile manipulation, a single demonstration already contains rich information

about the world it acts in. The moving robot’s changing views enable consistent 3D scene reconstruction. The robot observations can capture 4D object motions by fusing information across cameras and time. A single demonstration is sufficient to specify the task and to capture the contact-rich interactions essential for success. Learning pixel-to-action policies alone loses this useful information. To maximize the data efficiency of one demonstration, we push the goal to the limit: ***to learn an end-to-end open-world mobile manipulation policy from a single human demonstration.***

To this end, we propose **WANDA**, learning open-World mobile mANipulation from one demonstration via a synthetic **DA**ta engine. **WANDA** starts from world reconstruction and generation. For backgrounds, it leverages a sparse-view 2D Gaussian splatting [Huang et al., 2024] method, **MA**tCha [Guédon et al., 2025], to reconstruct both splats and mesh of the scene. The inputs to **MA**tCha are the images captured by the head camera during navigation. For foreground objects, it uses **B**undleSDF [Wen et al., 2023] to track the object’s 6D pose and reconstruct the mesh from RGBD observations. Inspired by **Real2Render2Real** [Yu et al., 2025], we reconstruct a render-only world without dynamics, mainly for planning and rendering. The reconstructed world serves as a workspace for subsequent trajectory planning and rendering.

To handle the challenges of open-world mobile manipulation, **WANDA** adopts several strategies to synthesize novel trajectories. For **spatial generalization**, the contact-rich robot-object interaction segments are automatically relocated into diverse spatial configurations. Whole-body motion planning is used to chain those segments into complete trajectories. For **long-horizon robustness**, we propose **C**orrective State Expansion to simulate policy drifts and increase the robot and object state diversity at different mobile manipulation stages. For **scene generalization**, we adopt **MA**rble [World Labs, 2026a] to scale 3D world generation from everyday photos for novel trajectory synthesis in diverse environments. This makes new scene data generation as simple as taking one picture. We utilize factorized rendering to produce the visual observations. Furthermore, **WANDA** naturally supports cross-embodiment data generation, and to our knowledge, demonstrates the first zero-shot **cross-embodiment generalization** capability for mobile manipulation.

Empirically, we validate **WANDA** across simulation and real-world tasks on both VLAs and visuo-motor policies. In simulation, **WANDA** achieves $\sim 50\times$ sample efficiency over teleoperated baselines in the single scene setting. In the real world, **WANDA** uses one demonstration to achieve 54.8% average progress score on five long-horizon mobile-manipulation tasks shown in Figure 4 across 16 different environments, unlocking broad spatial and cross-environment generalization. Moreover, we achieve zero-shot deployment on Linearbot using policies trained solely from synthetic data generated from a single Agibot G1 demonstration, demonstrating cross-embodiment transfer.

2 Related Work

Data generation for robot learning. Collecting teleoperation data is costly, requiring many operators and robot setups [Fu et al., 2024b, Jiang et al., 2025a]. Simulation-based generation methods [Mandlekar et al., 2023, Garrett et al., 2024, Jiang et al., 2025b, Li et al., 2025a] synthesize diverse demonstrations from few demos, but real deployment often requires sim-to-real effort or additional real fine-tuning. **D**emoGen [Xue et al., 2025] avoids dynamics simulation through point-cloud editing and action replanning, but focuses on tabletop spatial generalization rather than long-horizon mobile manipulation. **Real2Render2Real** [Yu et al., 2025] relies on manually scanned objects, and **1001Demos** [Pan et al., 2025] requires pre-scanned scenes for action-view augmentation; both depend on dedicated capture passes that limit scalability, and target tabletop tasks. **WANDA** reconstructs scenes and objects directly from the robot demonstration, scaling to long-horizon mobile manipulation in the wild.

Imitation learning for mobile manipulation. Mobile manipulation is harder than fixed-base manipulation, requiring coordinated base-arm control; teleoperation is correspondingly more ex-




Figure 2: Pipeline overview of WANDA. From a single real demonstration, **WANDA** reconstructs and generates 3D worlds, rearranges contact-rich robot-object interaction segments into diverse spatial configurations, chains them into new trajectories with whole-body motion planning, applies Corrective State Expansion to increase robot and object state coverage, and produces photo-realistic observations by factorized rendering.

pensive [Fu et al., 2024b, Jiang et al., 2025a]. Mobi- π [Yang et al., 2025] addresses base-pose selection for pretrained fixed-base policies rather than generating training trajectories. DynaMem [Liu et al., 2025] focuses on modular memory/planning with off-the-shelf grasping rather than end-to-end policy learning. Large vision-language-action policies [Physical Intelligence et al., 2025, 2026] provide policy architectures for open-world robotics, but their scaling remains tied to large heterogeneous datasets. **WANDA** targets this data bottleneck by turning one mobile-manipulation demonstration into synthetic trajectories across reconstructed and generated scenes.

3 Method

In this section, we propose **WANDA**, a data engine that turns one demonstration into diverse mobile manipulation data. As shown in Figure 2, we first reconstruct the 4D world from the source RGBD observations and diversify the backgrounds with 3D scenes generated from photos. Then we rearrange contact-rich robot-object interaction segments into extensive spatial configurations in both reconstructed and generated worlds, using whole-body motion planning to chain them into new trajectories. Corrective State Expansion is adopted to increase the robot and object state coverage at different mobile manipulation stages. Finally, we use factorized rendering to achieve fast, photo-realistic rendering of the visual observations for the generated trajectories.

3.1 Worlds Reconstruction and Generation

3D background scene reconstruction. To fully leverage the information from one demonstration, **WANDA** first performs world reconstruction for both backgrounds and foreground objects from the



Figure 3: Sampled initial configuration results on reconstructed and generated 3D scenes. Scene and robot meshes are shown in Blender. Red regions indicate the object initial pose sampling regions.

source observation. This acts as a world substrate for later planning and rendering. For background reconstruction, we first subsample a set of diverse mobile camera views from the source demonstrations. Then we inpaint the foreground from the background using Nano Banana Pro automatically. The background images are fed into a pose-free sparse-view 2D Gaussian Splatting method, MATCha, to reconstruct both splats and mesh of the scene. To align the reconstructed scene with real scale, we register the scene by aligning the predicted camera height with the corresponding camera height calculated from forward kinematics. This only requires easily accessible robot proprioception and avoids the demand for high-precision SLAM to get accurate camera poses.

Robot-object interaction reconstruction. We utilize BundleSDF [Wen et al., 2023] to simultaneously reconstruct the object mesh and track the 6D pose from RGBD observation streams. Compared with image-to-3D generation methods like SAM3D [Chen et al., 2026], this achieves high-fidelity object texture and geometry reconstruction, while eliminating the need to scan each object manually. The reconstruction is also compatible with articulated objects, where it can reconstruct and track the pose for different object parts separately. We provide a full reconstruction running example in Appendix A.

Scalable 3D scene generation. To enable in-the-wild scene generalization for the policy, we use Marble [World Labs, 2026b] to generate 3D Gaussian Splats and extract the mesh. We capture only one photo each scene across diverse environments, which are fed into Marble to generate 3D worlds. By augmenting backgrounds with generated 3D worlds, **WANDA** scales trajectory generation across diverse scenes to achieve navigation and manipulation generalization in different environments.

VLM-guided spatial configuration. To achieve large spatial coverage, we annotate the contact-rich End Effector-object interaction trajectories and rearrange them into different positions in the scenes. We utilize VLM’s commonsense to automatically output the initial randomization regions for replaying the contact-rich manipulation segments. Given one reconstructed or generated scene, we render orbit views of the scene. Then we input the rendered images as well as the manipulation task instruction to prompt Gemini 3.1 Pro to identify entities suitable for initializing the manipulated objects. Then it calls SAM3 to predict the entity masks, which are back-projected to 3D to get the regions. Since Marble-generated scenes often lack metric scale, the VLM leverages commonsense knowledge to infer plausible entity heights, which are then used to recover the scene’s metric scale. We showcase some initial configuration results on reconstructed and generated scenes in Figure 3.

3.2 Novel Trajectory Generation for Mobile Manipulation

Whole-body motion planning. After rearranging the contact-rich segments, we use whole-body motion planning to chain them into complete new trajectories. It includes navigation phase and manipulation phase planning. For navigation phases, we first sample a feasible base position and solve whole-body IK that can reach the start EE pose of the contact-rich segments. Then we utilize RRT-Connect [Kuffner and LaValle, 2000] to plan a collision-free path for the robot to reach the target base position. For manipulation phases, we utilize RRT-Connect to plan a collision-free arm path to reach the target EE pose from the navigation pose. Then during contact-rich segments, we jointly replay the object tracked pose and the EE motion, with IK applied to solve the arm joint angles. Inspired by MoMaGen, we also apply a visibility cost to keep objects in view during navigation. We detail the full planning procedure in Appendix B.

Corrective State Expansion. Long-horizon mobile manipulation can easily drift into out-of-distribution situations due to compounding error during navigation and manipulation. Dataset state coverage is crucial for the policy to handle long-horizon drifts. However, plain planning can cause narrow state coverage at different mobile manipulation stages. For example, the navigation planning may always go to the same position relative to the object. But during policy inference, the robot may navigate to a different position relative to the object. To simulate those drifts and increase policy robustness, we propose Corrective State Expansion, including object and robot state expansion.

First we apply *Object State Expansion* to diversify the object state during manipulation. During planning, after the robot reaches the target base position, we perturb the object poses, so that the object position distribution in robot base frame during manipulation has a large spatial coverage. Then we perform IK and arm motion planning again to adapt to the new object poses. This simulates the navigation drift relative to objects and enhances the spatial generalization of manipulation.

Second, we perform *Robot State Expansion* to enlarge the robot state coverage during mobile manipulation. Robot proprioceptive state can drift in the long run during manipulation and navigation. Therefore, in the navigation phase, we sample new navigation starts inside a stratified cone behind the original target position, and let the robot turn to the target object and go straight towards it. We also randomly perturb the arm joints during navigation and ease them back to the default carried pose to counter arm drift. In the manipulation phase, we replace the approaching-object segments with a cone-cycloid EE trajectory [Wang et al., 2025] whose start is sampled in a cone behind the target gripper pose; per-frame IK solves the active arm joint poses. This broadens the state configurations during approaching objects. Together these expansions teach the policy not just on-nominal execution but how to recover from realistic drift accumulated across long horizons.

Factorized rendering. After novel trajectory generation, we use factorized rendering to get the visual observation. Specifically, we render the background with Gaussian splatting and render the foreground robot and objects with Isaac Sim as a render engine, applying domain randomization on lighting. We then merge the background and foreground to get final frames. This avoids time-consuming large-scene rendering in Isaac Sim and achieves photo-realistic rendering fast.

4 Experiments

4.1 Simulation Experiments

We evaluate **WANDA** across different policy architectures and simulation benchmarks to validate the general effectiveness. For all the simulation experiments, we only utilize the RGBD observation and joint state information from the demonstrations. No privileged information is leveraged to align with the real-world setting.

Bigym. First, we choose three mobile manipulation tasks from Bigym [Chernyadev et al., 2024]



Figure 4: Real-world tasks. Five long-horizon mobile-manipulation tasks, each learned from a single real demonstration: LUNCHBOX, UTENSIL, DROPTRASH, POUR, and FRIDGE.




Figure 5: Demonstration of simulation tasks and subtasks from Bigym [Chernyadev et al., 2024] and B1K [BEHAVIOR Team, 2025].

Table 1: Simulation results.

Bigym	Saucepan to Hob	Move Plate	Drawer Close	Avg
WANDA (Ours)	86.7	40.0	100.0	75.6
full source demo	86.0	48.0	100.0	78.0

BEHAVIOR	A	B	C	Q-score
WANDA	20.0	16.7	13.3	16.67
20 demos	6.7	3.3	0.0	3.33
50 demos	16.7	13.3	3.3	11.11
200 demos	80.0	76.7	30.0	62.22

shown in Figure 5 and generate 100 demonstrations from one source demonstration. Then we choose ACT for policy learning, to validate the effectiveness of **WANDA** on small policies. Baselines are ACTs trained on the full source dataset, around 40–60 demos per task. We report success rate over 50 rollouts.

BEHAVIOR Challenge. We choose a three-minute-long task from BEHAVIOR Challenge [BEHAVIOR Team, 2025], which contains four navigation stages and four manipulation subtasks shown in Figure 5 to stress-test our method on long-horizon mobile manipulation. We perform full-parameter fine-tuning of $\pi_{0.5}$ [Physical Intelligence et al., 2025] on the datasets. To investigate the data efficiency, we generate 1360 demonstrations from one official demonstration and compare with $\pi_{0.5}$ trained on 20, 50, and 200 official demonstrations separately. We report the Q-score (progress score defined by official implementation) over 30 evaluation episodes.

Results. The simulation evaluation results are presented in Table 1. For Bigym, using data generated from one source demonstration, **WANDA** achieves comparable results to baselines trained on ~50 demonstrations. This indicates the effectiveness of our method for localized mobile manipula-

Table 2: Real-world mobile manipulation results on Agibot G1. Progress scores are mean completed-subtask percentages over 10 trials; Average is the unweighted task mean. The Utensil ablation retains its legacy two-subtask score.

Method	Lunch Box	Utensil	Drop Trash	Fridge	Pour	Average
WANDA (Ours)	55.0%	52.5%	75.0%	46.7%	45.0%	54.8%
WANDA w/o CSE	10.0%	15.0%	35.0%	13.3%	5.0%	15.7%




Figure 6: Per-subtask success rate on the real-world tasks. Conditional success with raw counts; 0/0 means never reached and N/A denotes stages unrecorded in the legacy Utensil ablation.

tion. Even with limited viewpoint changes during movement, our pipeline still reliably reconstructs the world and generates high-fidelity demonstrations. For BEHAVIOR Challenge, our approach even outperforms those trained on 20 and 50 source demonstrations. This validates our pipeline’s effectiveness for generating high-quality trajectories for extremely long-horizon mobile manipulation. Together, the simulation results show that **WANDA** achieves $\sim 50\times$ data efficiency, reducing human effort for mobile manipulation data collection.

4.2 Real-World Experiments

We design five mobile manipulation tasks on Agibot G1, as illustrated in Figure 4. These tasks emphasize long horizon robustness and spatial generalization over different regions in the scenes:

- *Lunch Box*: The robot needs to first navigate to the lunch box on different tables, fold it, then take it to the kitchen sink and place it in the sink. We generate 396 trajectories for this task.
- *Utensil*: The robot is required to first search the knife on the plate, navigate to it, pick up the knife precisely and place it in the storage basket. 510 generated trajectories are used for training.
- *Drop Trash*: The robot needs to first find the cola trash over different tables in the rooms, grasp it and navigate to the trash can and throw the trash. We generate 407 trajectories for this task.




Figure 7: Scalability results. (a) Performance versus generated-data proportion. (b) Per-task progress scores in three novel scenes; each task–scene pair compares **WANDA** with the 0% source-only baseline.

- *Pour*: The robot needs to navigate to the teapot randomized in the large kitchen table, take the teapot to the cup and pour into it. 740 trajectories are generated for the task.
- *Fridge*: The robot is required to take the grocery bag randomized on the round table, store it in the fridge and close the fridge door. We utilize 263 generated demonstrations to train the policy.

We use PICO VR to teleoperate Agibot G1 to collect one source demonstration for each task. We perform full-parameter fine-tuning on $\pi_{0.5}$ on the generated data. We compare the full **WANDA** with a variant that removes Corrective State Expansion (CSE) to assess the necessity of this component.

Metric. We conduct 10 trials for each task and report the progress score (the percentage of completed subtasks). The progress score for each task is defined in Appendix C.

Results. The results are presented in Table 2, with per-subtask success rates in Figure 6 and subtask definitions in Appendix C. **WANDA** produces meaningful progress on all five long-horizon real-world mobile manipulation tasks from a single human demonstration. For *Lunch Box*, the robot must complete a multi-stage task lasting around three minutes; **WANDA** remains robust to long-horizon drift and completes the full sequence. For *Drop Trash*, the robot steadily navigates to the trash placed on different tables, grasps it, and carries it to the trash can. This indicates the effectiveness of our method in achieving spatial generalization. For *Utensil*, the robot navigates a long distance to reach the target base position and grasps the knife precisely. Together, these tasks show that our method is compatible with both rigid and articulated objects.

CSE Ablation. The performance of the non-CSE variant degrades significantly compared with full **WANDA**. **WANDA** without CSE is vulnerable to compounding state deviations. Before manipulation, the robot may reach a pose with noticeable error relative to the source demonstration, leading to unseen robot-object configurations that are not covered by non-CSE data. During manipulation, small drifts in the arm joint state can further push the policy out of distribution when approaching the object. In contrast, **WANDA** remains robust under such deviations, demonstrating the importance of Corrective State Expansion. Figure 10 visualizes how CSE mitigates these deviations.

4.3 Scalability Study

To assess the scalability of the proposed data generation pipeline, we evaluate our methods on two scaling dimensions: scaling in-domain demonstrations in the source scene, and scaling generated scene diversity in the generated demonstrations.

Scaling generated demonstrations for in-domain improvement. We perform this experiment both in BEHAVIOR Challenge and in real-world *Drop Trash* task. We sample subsets of the full




Figure 8: Per-subtask success rate across novel scenes. WANDA is trained on Marble-generated worlds and evaluated on Drop Trash, Pour, Lunch Box, and Utensil in three real scenes; we use the same conditional metric as Figure 6. Subtask definitions are given in Appendix C.

Table 3: Data-generation time cost. Wall-clock hours to generate one hour of 30 fps demonstration data.

	Human teleop	Ours		MoMaGen	
	1 teleoperator	1 GPU	8 GPUs	1 GPU	8 GPUs
Time to generate 1 h of data (h)	1.5	3.0	0.38	19.4	2.4

generated dataset and perform $\pi_{0.5}$ fine-tuning under the same training configurations. As shown in Figure 7(a), scaling the generated data achieves steady improvement in both simulation and real world. This implies the scalability of WANDA for in-domain performance improvement.

Scaling generated scenes for open-world generalization. To validate the scene generalization ability of our approach, we test WANDA on different environments. For each scene, we take a single photo and feed it into Marble to generate a corresponding 3D world. We train the policy on data synthesized from both source scenes and these photo-generated worlds, and evaluate it in the same real-world scenes where the photos were taken. Although the generated worlds are derived from these environments, they still contain substantial visual and geometric discrepancies from reality. Consequently, successful task execution requires strong scene generalization to bridge this gap.

During evaluation, we roll out the policy trained on only the source scene and the policy trained on generated scenes in each environment for 10 episodes and report the progress score. As shown in Figure 7(b), WANDA achieves generalization in the open world, while the variant without generated scenes fails to generalize. The source-only baseline stays near a zero progress score on every novel scene, whereas WANDA sustains substantial progress across all four tasks and three novel scenes. Figure 8 reports the per-subtask breakdown across the three scenes, where the gains persist into the later manipulation subtasks. This shows WANDA scales scene diversity, making open-world data generation as easy as taking one photo.

Data Generation Cost. Our method is scalable with the compute to generate data. Table 3 compares the wall-clock time to generate one hour of 30 fps demonstration data. Through factorized rendering, **WANDA** generates one data-hour with about 3 GPU-hours, $6.4\times$ less than MoMaGen [Li et al., 2025a] at equal GPU count; on a single 8×L40S node this is only 0.38 wall-clock hours, about $4\times$ faster than a human teleoperator. Crucially, **WANDA** needs only one human demonstration per task upfront, whose labor amortizes toward zero as the generated dataset scales, whereas teleoperation incurs roughly 1.5 human-hours for every data-hour, estimating the average environment reset time to be half the demonstration time. A per-task cost breakdown is provided in Appendix E.

4.4 Zero-shot Cross-Embodiment Generalization

WANDA naturally supports cross-embodiment data generation, since the reconstructed world is embodiment-agnostic for planning and rendering. Also, a single demonstration is sufficient to specify the task and capture the contact-rich interactions essential for success, which is also embodiment-agnostic for most non-prehensile manipulation tasks.

Here we perform cross-embodiment data generation from Agibot G1 to Linearbot, shown in Figure 9. These two robots have large morphology gaps and different camera placements. We choose *Drop Trash* for validation. 400 Linearbot demonstrations are generated from one Agibot G1 real demonstration and fed into $\pi_{0.5}$ training. The policy is then zero-shot deployed on a physical Linearbot. Over ten rollouts it reaches a 40% progress score (per-subtask success rates in Table 4); the generated observations are shown in Appendix F. To our knowledge, this demonstrates the first zero-shot cross-embodiment generalization capability in mobile manipulation.

4.5 Visualization

Generated demonstration visualization. We show the visualization of head camera views from the synthetic demonstrations in Figure 1. For all generated camera streams including wrist cameras, please refer to Appendix D. Our approach generates photo-realistic observations for the source scenes and diverse generated scenes, providing high-quality data for downstream policy learning.




Figure 10: Corrective State Expansion.

Corrective State Expansion visualization. To understand the role of Corrective State Expansion




Figure 9: Cross-embodiment data generation from Agibot G1 to Linearbot.

Table 4: Linearbot subtask success rate.

Subtask	Success rate
Grasp the cola can	5/10
Throw into the trash bin	3/5

sion, we visualize the object position distribution in robot base frame during manipulation and the robot state distribution during navigation and manipulation from a subset of the dataset. As shown in Figure 10(a), CSE achieves a much larger spatial distribution of object positions during manipulation than the non-CSE version, which enhances the manipulation spatial robustness. Figure 10(b) shows the robot pose distributions over the target. CSE greatly enhances the density and diversity of robot poses for the target, making it able to recover from pose drift.

5 Limitations and Future Work

WANDA has two main limitations. First, it handles only rigid and articulated objects. Reconstruction tracks the 6D poses of rigid parts and replays them, so soft and deformable objects, such as cloth and fluids, cannot be reconstructed or generated; in such cases **WANDA** replays the demonstrated motion rather than modeling the underlying deformation. Second, as a multi-stage data engine, errors can compound across stages: a reconstruction or planning failure in one stage propagates to the subsequent navigation and manipulation stages, so generating long-horizon trajectories occasionally requires human effort to intervene and correct failed stages.

These limitations point to several directions for future work. To relax the rigid-object assumption, the reconstruction stage could be extended with 4D Gaussian Splatting [Wu et al., 2024] to capture complex, deformable object interactions. Because evaluating open-world mobile manipulation is time-consuming and stochastic—it requires physically relocating the robot across many sites—incorporating physics into **WANDA** and extending it toward real-to-sim evaluation is a promising route to benchmarking mobile manipulation more efficiently and reproducibly. Finally, co-training the synthesized data with UMI data and human videos is an interesting avenue for further broadening the diversity and coverage of the learned policies.

6 Acknowledgments

This work was partially funded by NSF Award #2512805. Research supported by the NVIDIA Academic Grant Program using NVIDIA RTX PRO 6000 Blackwell Max-Q GPUs and the NVIDIA Jetson AGX Thor platform. Guanya Shi holds concurrent appointments as an Assistant Professor at Carnegie Mellon University and as an Amazon Scholar. This paper describes work performed at Carnegie Mellon University and is not associated with Amazon. We greatly thank AppLovin for providing us with the compute resources for this project. We thank Tairan He, Chaoyi Pan and Junfeng Ni for their insightful discussions and feedback. We also thank Zeji Yi, Yi Yang and Angchen Xie for their support with the experiments.

References

- BEHAVIOR Team. BEHAVIOR: Benchmark for everyday household activities in virtual, interactive, and realistic environments. <https://behavior.stanford.edu/>, 2025.
- Xingyu Chen, Fu-Jen Chu, Pierre Gleize, Kevin J Liang, Alexander Sax, Hao Tang, Weiyao Wang, Michelle Guo, Thibaut Hardin, Xiang Li, et al. Sam 3d: 3dfy anything in images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7220–7232, 2026.
- Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European conference on computer vision*, pages 640–658. Springer, 2022.
- Xuxin Cheng, Jialong Li, Shiqi Yang, Ge Yang, and Xiaolong Wang. Open-television: Teleoperation with immersive active visual feedback. *arXiv preprint arXiv:2407.01512*, 2024.
- Nikita Chernyadev et al. Bigym: A demo-driven mobile bi-manual manipulation benchmark. *arXiv preprint arXiv:2407.07788*, 2024.
- Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. *arXiv preprint arXiv:2402.10329*, 2024.
- Haoquan Fang, Jiafei Duan, Donovan Clay, Sam Wang, Shuo Liu, Weikai Huang, Xiang Fan, Wei-Chuan Tsai, Shirui Chen, Yi Ru Wang, et al. Molmoact2: Action reasoning models for real-world deployment. *arXiv preprint arXiv:2605.02881*, 2026.
- Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. *arXiv preprint arXiv:2406.10454*, 2024a.
- Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*, 2024b.
- Caelan Reed Garrett, Ajay Mandlekar, Bowen Wen, and Dieter Fox. Skillmimicgen: Automated demonstration generation for efficient skill learning and deployment. In *8th Annual Conference on Robot Learning*, 2024.
- Antoine Guédon, Tomoki Ichikawa, Kohei Yamashita, and Ko Nishino. Matcha gaussians: Atlas of charts for high-quality geometry and photorealism from sparse views. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6001–6011, 2025.
- Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8944–8951. IEEE, 2024.
- Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024.
- Arhan Jain, Mingtong Zhang, Kanav Arora, William Chen, Marcel Torne, Muhammad Zubair Irshad, Sergey Zakharov, Yue Wang, Sergey Levine, Chelsea Finn, et al. Polaris: Scalable real-to-sim evaluations for generalist robot policies. *arXiv preprint arXiv:2512.16881*, 2025.

- Joel Jang, Seonghyeon Ye, Zongyu Lin, Jiannan Xiang, Johan Bjorck, Yu Fang, Fengyuan Hu, Spencer Huang, Kaushil Kundalia, Yen-Chen Lin, et al. Dreamgen: Unlocking generalization in robot learning through video world models. *arXiv preprint arXiv:2505.12705*, 2025.
- Yunfan Jiang, Ruohan Zhang, Josiah Wong, Chen Wang, Yanjie Ze, Hang Yin, Cem Gokmen, Shuran Song, Jiajun Wu, and Li Fei-Fei. Behavior robot suite: Streamlining real-world whole-body manipulation for everyday household activities. *arXiv preprint arXiv:2503.05652*, 2025a.
- Zhenyu Jiang, Yuqi Xie, Kevin Lin, Zhenjia Xu, Weikang Wan, Ajay Mandlekar, Linxi Jim Fan, and Yuke Zhu. Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 16923–16930. IEEE, 2025b.
- Simar Kareer, Dhruv Patel, Ryan Punamiya, Pranay Mathur, Shuo Cheng, Chen Wang, Judy Hoffman, and Danfei Xu. Egomimic: Scaling imitation learning via egocentric video. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13226–13233. IEEE, 2025.
- Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- James J Kuffner and Steven M LaValle. 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, pages 995–1001. IEEE, 2000.
- Chengshu Li, Mengdi Xu, Arpit Bahety, Hang Yin, Yunfan Jiang, Huang Huang, Josiah Wong, Sujay Garlanka, Cem Gokmen, Ruohan Zhang, Weiyu Liu, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Momagen: Generating demonstrations under soft and hard constraints for multi-step bimanual mobile manipulation, 2025a. URL <https://arxiv.org/abs/2510.18316>.
- Ming-Feng Li, Xin Yang, Fu-En Wang, Hritam Basak, Yuyin Sun, Shreekanth Gayaka, Min Sun, and Cheng-Hao Kuo. Ua-pose: Uncertainty-aware 6d object pose estimation and online object completion with partial references. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1180–1189, 2025b.
- Kevin Lin, Ajay Mandlekar, Caelan Reed Garrett, Nikita Chernyadev, Yu Fang, Runyu Ding, Yuqi Xie, Linxi Fan, and Yuke Zhu. Humanoidmimicgen: Data generation for loco-manipulation via whole-body planning and adaptation. In *ICRA 2026 Workshop on Synthetic Data for Robot Learning*, 2026.
- Peiqi Liu, Zhanqiu Guo, Mohit Warke, Soumith Chintala, Chris Paxton, Nur Muhammad Mahi Shafillah, and Lerrel Pinto. Dynamem: Online dynamic spatio-semantic memory for open world mobile manipulation. In *ICRA 2025 Workshop: Human-Centered Robot Learning in the Era of Big Data and Large Models*, 2025.
- Songming Liu, Bangguo Li, Kai Ma, Lingxuan Wu, Hengkai Tan, Xiao Ouyang, Hang Su, and Jun Zhu. Rdt2: Exploring the scaling limit of umi data towards zero-shot cross-embodiment generalization. *arXiv preprint arXiv:2602.03310*, 2026.
- Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Ireteyio Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. In *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 1820–1864, 2023.

- Jiageng Mao, Sicheng He, Hao-Ning Wu, Yang You, Shuyang Sun, Zhicheng Wang, Yanan Bao, Huizhong Chen, Leonidas Guibas, Vitor Guizilini, et al. Robot learning from a physical world model. *arXiv preprint arXiv:2511.07416*, 2025.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- Chuer Pan, Litian Liang, Dominik Bauer, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, and Shuran Song. One demo is worth a thousand trajectories: Action-view augmentation for visuomotor policies. In *Conference on Robot Learning (CoRL)*, 2025.
- Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.
- Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- Physical Intelligence, Bo Ai, Ali Amin, Raichelle Aniceto, Ashwin Balakrishna, Greg Balke, Kevin Black, George Bokinsky, Shihao Cao, Thomas Charbonnier, et al. $\pi_{0.7}$: a steerable generalist robotic foundation model with emergent capabilities. *arXiv preprint arXiv:2604.15483*, 2026.
- Ryan Punamiya, Simar Kareer, Zeyi Liu, Josh Citron, Ri-Zhao Qiu, Xiongyi Cai, Alexey Gavryushin, Jiaqi Chen, Davide Liconti, Lawrence Y Zhu, et al. Egoverse: An egocentric human dataset for robot learning from around the world. *arXiv preprint arXiv:2604.07607*, 2026.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- Sunday Team. Act-1: A robot foundation model trained on zero robot data. <https://www.sunday.ai/journal/no-robot-data>, 2025.
- Wenhao Wang, Kehe Ye, Xinyu Zhou, Tianxing Chen, Cao Min, Qiaoming Zhu, Xiaokang Yang, Yongjian Shen, Yang Yang, Maoqing Yao, et al. Fieldgen: From teleoperated pre-manipulation trajectories to field-guided data generation. *arXiv preprint arXiv:2510.20774*, 2025.
- Bowen Wen, Jonathan Tremblay, Valts Blukis, Stephen Tyree, Thomas Muller, Alex Evans, Dieter Fox, Jan Kautz, and Stan Birchfield. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- World Labs. Marble. <https://marble.worldlabs.ai/>, 2026a. Accessed: 2026-05-29.
- World Labs. Marble documentation. <https://docs.worldlabs.ai/>, 2026b.
- Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20310–20320, 2024.
- Xiaomeng Xu, Jisang Park, Han Zhang, Eric Cousineau, Aditya Bhat, Jose Barreiros, Dian Wang, and Shuran Song. Hommi: Learning whole-body mobile manipulation from human demonstrations. *arXiv preprint arXiv:2603.03243*, 2026.

- Xiuwei Xu, Angyuan Ma, Hankun Li, Bingyao Yu, Zheng Zhu, Jie Zhou, and Jiwen Lu. R2rgen: Real-to-real 3d data generation for spatially generalized manipulation. *arXiv preprint arXiv:2510.08547*, 2025.
- Zhengrong Xue, Shuying Deng, Zhenyang Chen, Yixuan Wang, Zhecheng Yuan, and Huazhe Xu. Demogen: Synthetic demonstration generation for data-efficient visuomotor policy learning. *arXiv preprint arXiv:2502.16932*, 2025.
- Jingyun Yang, Isabella Huang, Brandon Vu, Max Bajracharya, Rika Antonova, and Jeannette Bohg. Mobi- π : Mobilizing your robot learning policy. *arXiv preprint arXiv:2505.23692*, 2025.
- Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, et al. World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026.
- Justin Yu, Letian Fu, Huang Huang, Karim El-Refai, Rares Ambrus, Richard Cheng, Muhammad Zubair Irshad, and Ken Goldberg. Real2render2real: Scaling robot data without dynamics simulation or robot hardware. *arXiv preprint arXiv:2505.09601*, 2025.
- Zhaxizhuom Zhaxizhuoma, Kehui Liu, Chuyue Guan, Zhongjie Jia, Ziniu Wu, Xin Liu, Tianyu Wang, Shuai Liang, Pengan Chen, Pingrui Zhang, et al. Fastumi: A scalable and hardware-independent universal manipulation interface with dataset. In *Conference on Robot Learning*, pages 3069–3093. PMLR, 2025.
- Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and Chuang Gan. 3d-vla: A 3d vision-language-action generative world model. *arXiv preprint arXiv:2403.09631*, 2024.
- Ruijie Zheng, Dantong Niu, Yuqi Xie, Jing Wang, Mengda Xu, Yunfan Jiang, Fernando Castañeda, Fengyuan Hu, You Liang Tan, Letian Fu, et al. Egoscale: Scaling dexterous manipulation with diverse egocentric human data. *arXiv preprint arXiv:2602.16710*, 2026.

Contents

1	Introduction	2
2	Related Work	3
3	Method	4
3.1	Worlds Reconstruction and Generation	4
3.2	Novel Trajectory Generation for Mobile Manipulation	6
4	Experiments	6
4.1	Simulation Experiments	6
4.2	Real-World Experiments	8
4.3	Scalability Study	9
4.4	Zero-shot Cross-Embodiment Generalization	11
4.5	Visualization	11
5	Limitations and Future Work	12
6	Acknowledgments	12
A	Reconstruction Running Example	18
A.1	Demonstration Post-process Interface	18
A.2	Background Reconstruction	19
A.3	Foreground Reconstruction	19
A.4	VLM-guided Initial Configuration	21
B	Novel Trajectory Generation	21
B.1	Whole-Body Motion Planning	22
B.2	Corrective State Expansion	23
C	Progress Score and Subtask Evaluation	24
C.1	Progress Score Definition	24
C.2	Subtask Success Rate	24
C.3	Scene Generalization Subtask Success Rate	25
D	Multi-View Videos of Synthesized Data	25
E	Per-Task Data-Generation Cost	25
F	Cross-embodiment Data Generation	29
G	Policy Training Details	29
H	Prompt	29
H.1	Inpaint Prompt	29
H.2	Initial Configuration Prompt	30

A Reconstruction Running Example

A.1 Demonstration Post-process Interface

Given one collected real demonstration, we first segment the foreground object masks, annotate the contact-rich segments and navigation segments via a user-friendly interface (Figure 11). We utilize XMem [Cheng and Schwing, 2022], a long-term video segmentation approach, to get the foreground object masks along the video. For background reconstruction, we annotate the navigation segments, and fill the object mask region with green-screen color. Then we prompt Nano Banana Pro (Gemini 3 Pro Image) to remove the foreground robot arms and object masks and inpaint the background, which are then fed into background reconstruction pipeline. The prompt is shown in Appendix H.1. For foreground reconstruction, we annotate the contact-rich manipulation segments to get corresponding object masks. Then the masks with corresponding RGB and Depth images are then fed into BundleSDF [Wen et al., 2023] for foreground reconstruction. The annotation process only requires the user to click the segmented object a few times. Together with the segmentation model’s propagation time, annotation costs 5 to 30 minutes, varying among tasks. This process could even be replaced by an automatic agent workflow in the future to remove the human annotation effort.




Figure 11: Demonstration post-process interface used to annotate contact-rich and navigation segments.



Figure 12: Background reconstruction results (textured mesh and geometry mesh) visualized in Blender. From left to right: *Pour*, *Drop Trash*, *Fridge*, *Utensil*.

A.2 Background Reconstruction

Given inpainted background images, we utilize a pose-free sparse-view 2D Gaussian splatting method, MAtCha [Guédon et al., 2025] to reconstruct both splats and geometry meshes, and predict camera poses in the world frame. To align the scale and orientation of the reconstructed splats and meshes, we perform registration to minimize the predicted camera pose z and the actual camera height calculated via forward kinematics from recorded joint states. This avoids the demand on high-precision odometry or SLAM, making the pipeline robust to noisy real-world observation. More background reconstruction results are shown in Figure 12, where we show the textured mesh and geometry mesh in Blender.

A.3 Foreground Reconstruction



Figure 13: Foreground reconstruction: BundleSDF vs. BundleSDF with generative completion.

We adopt BundleSDF [Wen et al., 2023], a SDF [Sitzmann et al., 2020] and NeRF [Mildenhall et al., 2021] based approach to jointly reconstruct the object geometry and texture and track the object 6D poses. It accepts RGB, object segmentation mask and depth image streams from the observation streams of the contact-rich segments. We adjust the parameters of BundleSDF to make it suitable for fine-grained texture reconstruction. Besides, sometimes the source demonstration cannot have a complete viewpoint coverage over the objects, therefore the reconstruction can be incomplete. For such cases, we adopt a generative completion method inspired by UA-Pose [Li et al., 2025b]. Specifically, we utilize SAM3D [Chen et al., 2026] to generate 3D mesh from single image of the object from the demonstration. Although SAM3D-generated mesh fails to reconstruct the photo-realistic texture, it provides a good geometry shape prior for the non-seen part of objects. Therefore, we first reconstruct the raw object mesh from incomplete real RGBD observation and compute the visibility of each reconstructed face. Then we perform registration to align the SAM3D generated mesh scale and 6D pose with the reconstructed mesh, then render the images from SAM3D generated mesh on those viewpoints unseen in real observation. Then we fuse the real RGBD images and synthetic images rendered by SAM3D mesh, and train the object SDF and NeRF to get final textured meshes.

To validate the effectiveness of the generative completion method, we show comparisons of the object mesh generated by **BundleSDF**, **SAM3D** and **BundleSDF w/ Gen** method in Figure 13. Note that we only apply this method to tasks that do not contain the complete observation over the object. As shown in Figure 13, our approach achieves both complete geometry shape and precise texture reconstruction. Due to incomplete observation, **BundleSDF** fails to get reasonable shape at the unseen faces. **SAM3D** fails to generate photo-realistic textures, while our approach can even recover the text on the objects, as shown in the first row of Figure 13 and the first column of Figure 14.

For demonstrations that contains complete observation over the object, we directly adopt **BundleSDF** and provide a comparison with **SAM3D** in Figure 14. Compared with the image-to-3D methods that are commonly used in previous real2sim work [Mao et al., 2025, Jain et al., 2025], our method achieves photo-realistic texture reconstruction and precise geometry reconstruction with minimal sim-and-real gap.

For robot, we utilize the robot URDF or USD to directly replay the joint state for contact-rich segments recorded during data collection. We utilize Isaac Sim to render both robots and objects, align the background light color and apply with minimal light brightness domain randomization to simulate the light changes in the real environment.



Figure 14: Additional foreground reconstruction results.




Figure 15: Initial configuration results on reconstructed and generated 3D scenes.

A.4 VLM-guided Initial Configuration

We provide the detailed prompt used in Appendix H.2. Here we show the results of initial configuration on both reconstructed 3D scenes and generated 3D scenes. The VLM successfully recognize the reasonable region for randomizing the objects, and place the objects in extensive spaces to enhance the spatial coverage of the generated data.

B Novel Trajectory Generation

Notation. Homogeneous transforms $\mathbf{T}_A^B \in \text{SE}(3)$ map frame B into frame A ; W is the world frame and o the target-object frame, with object position $\mathbf{p}_o$. The mobile base is the planar state $\mathbf{b} = [x_b, y_b, \psi_b]^\top \in \text{SE}(2)$. End-effector (EE) targets are $\mathbf{T}_W^{E_k}$ for the left/right arm $k \in \{\text{L}, \text{R}\}$, with rotation and translation blocks $\mathbf{R}_W^{E_k}, \mathbf{p}_W^{E_k}$ and constant object-relative grasp $\mathbf{T}_o^{E_k} = (\mathbf{T}_W^o)^{-1} \mathbf{T}_W^{E_k}$. $\mathbf{R}_z(\cdot) \in$

SO(3) is a world-z yaw rotation, $\hat{\mathbf{d}}$ a unit direction, $(\cdot)^*$ an IK-achieved quantity, and the superscript src marks source-demonstration quantities. All weights and tolerances ($w_\bullet, \theta_{\max}, c_{\min}, \epsilon_p, \epsilon_R, \mathcal{R}_t, \mathcal{R}_\psi, \alpha, R$) are configuration constants, while Ψ_{yaw} below is a path quantity. The navigation endpoint is p_e with yaw ψ_e and unit approach direction $\hat{\mathbf{d}}_{\text{app}}$; $\hat{\mathbf{d}}_{\text{head}}$ is the head forward axis, $\hat{\mathbf{d}}_{H \rightarrow o}$ the head-to-object direction, $\mathbf{p}_{\text{piv}}$ the centroid of the perturbed object group, and $\mathcal{K} \subseteq \{\text{L}, \text{R}\}$ the set of tracked arms.

B.1 Whole-Body Motion Planning

For each contact-rich segment, the planner regenerates a complete whole-body trajectory in the target world through three stages—navigation, arm approach, and replay—while preserving the segment itself verbatim in object-relative coordinates.

Navigation. We first sample a set of base candidates on SE(2) around the target object (concentric rings and rays radiating from the source base pose). Each candidate is collision-checked (base footprint against the scene mesh of the world substrate) and has its visibility evaluated; visibility hard-rejects a candidate only when a strict visibility gate is enabled, and otherwise enters the ranking cost below. For surviving candidates we solve dual-arm whole-body IK to the segment’s start EE poses with a multi-seed solver (seeded from the current and source postures), and plan a collision-free base path with RRT-Connect over SE(2), which is then retimed into a smooth turn–drive–turn motion. If line of sight is lost mid-route, the planner inserts a yaw-only spin to re-face the object and replans the remainder from the new heading, and for rotationally symmetric objects several object-yaw variants are tried. A candidate is feasible ($\mathbf{b} \in \mathcal{B}_{\text{feas}}$) when its retimed base path $P_{\mathbf{b}}$ passes the collision and IK gates (and the visibility gate, when configured). Each feasible path is scored by a weighted cost, and—by the configured policy—either the first feasible candidate or the lowest-cost one is committed:

$$C_{\text{vis}}(P) = w_\rho(1 - \rho_{\text{vis}}) + \frac{w_{\bar{c}}}{2}(1 - \bar{c}) + w_f \frac{|\theta_{\text{err}}|}{\pi} + w_{xy} L_{xy} + w_\psi \frac{\Psi_{\text{yaw}}}{\pi},$$

$$\mathbf{b}^* = \begin{cases} \text{first } \mathbf{b} \in \mathcal{B}_{\text{feas}} & (\text{first-feasible policy}), \\ \arg \min_{\mathbf{b} \in \mathcal{B}_{\text{feas}}} C_{\text{vis}}(P_{\mathbf{b}}) & (\text{min-cost policy}). \end{cases} \quad (\text{B.1})$$

Over the N retimed steps of the base path, $\hat{\mathbf{h}}_i$ is the head forward direction (the head camera’s forward axis, or the planar base heading $[\cos \psi_i, \sin \psi_i]^\top$ as a fallback), the direction to the object is $\hat{\mathbf{d}}_i = \text{normalize}(\mathbf{p}_o - \mathbf{p}_i)$, and the facing cosine is $c_i = \hat{\mathbf{h}}_i^\top \hat{\mathbf{d}}_i$; a step is visible when it clears the field-of-view, facing, and occlusion tests,

$$\text{vis}_i = \mathbb{1}[\arccos(c_i) \leq \theta_{\max}] \mathbb{1}[c_i \geq c_{\min}] \mathbb{1}[\text{ray}_i \text{ clear}]. \quad (\text{B.2})$$

Here $\rho_{\text{vis}} = \frac{1}{N} \sum_i \text{vis}_i$ is the visible ratio, $\bar{c} = \frac{1}{N} \sum_i c_i$ the mean facing cosine, $\theta_{\text{err}} = |\text{atan2}(p_{o,y} - y_b, p_{o,x} - x_b) - \psi_b|_{\text{wrap}}$ the terminal facing error, $L_{xy} = \sum_i \|\mathbf{p}_i^{xy} - \mathbf{p}_{i-1}^{xy}\|$ the travelled distance, and $\Psi_{\text{yaw}} = \sum_i |\Delta \psi_i|_{\text{wrap}}$ the cumulative yaw change. In contrast to a hard visibility constraint, this cost lets the planner trade off visibility against path effort when ranking feasible candidates.

Arm approach. From the committed navigation pose, we plan a collision-free arm path with joint-space RRT-Connect to the segment’s start EE pose.

Replay. During the contact-rich segment, the per-frame dual-arm targets follow the standard object-relative (contact-relative) transfer of the source EE poses [Mandlekar et al., 2023, Li et al., 2025a]:

$$\mathbf{T}_W^{E_k}(i) = \mathbf{T}_W^o(\mathbf{T}_W^{o, \text{src}}(i))^{-1} \mathbf{T}_W^{E_k, \text{src}}(i), \quad k \in \{\text{L}, \text{R}\}, \quad (\text{B.3})$$

which we realize by per-frame whole-body IK while tracking the recorded object motion. Replaying in object-relative coordinates preserves the contact interaction exactly, while the surrounding trajectory is regenerated for the new spatial configuration.

B.2 Corrective State Expansion

A planner-feasible trajectory follows a narrow, on-manifold distribution, whereas a real mobile robot deviates from it over long horizons: navigation overshoots, the arms drift while carrying, and the gripper approaches the target from unfamiliar angles. Corrective State Expansion (CSE) broadens state coverage at different stages, and every generated sample is re-validated for kinematic feasibility and collision before it is accepted.

Object State Expansion. Applied inside the planner, this form simulates navigation drift relative to the object. With the base parked at $\mathbf{b}$ and frozen, the (grouped) target object is perturbed rigidly about the group centroid $\mathbf{p}_{\text{piv}}$. A translation and yaw are drawn uniformly from configured (possibly asymmetric) ranges, applied in the world frame by default (optionally the base frame):

$$(\delta x, \delta y, \delta z) \sim \mathcal{U}(\mathcal{R}_t), \quad \delta \psi \sim \mathcal{U}(\mathcal{R}_\psi), \quad (\text{B.4})$$

$$\Delta = \begin{bmatrix} \mathbf{R}_z(\delta \psi) & \mathbf{p}_{\text{piv}} + [\delta x, \delta y, \delta z]^\top - \mathbf{R}_z(\delta \psi) \mathbf{p}_{\text{piv}} \\ \mathbf{0}^\top & 1 \end{bmatrix}. \quad (\text{B.5})$$

The rigid transform $\Delta \in \text{SE}(3)$ rotates about $\mathbf{p}_{\text{piv}}$; applying it to the object pose and propagating through the fixed grasp yields the updated targets

$$\tilde{\mathbf{T}}_W^o = \Delta \mathbf{T}_W^o, \quad \mathbf{T}_W^{E_k} = \tilde{\mathbf{T}}_W^o \mathbf{T}_o^{E_k}, \quad k \in \{\text{L}, \text{R}\}, \quad (\text{B.6})$$

which are re-solved by whole-body IK over an enlarged seed set (the perturbed target may leave the basin of the original seed). A sample is accepted only if every tracked arm ($k \in \mathcal{K}$) is within tolerance and the configuration is collision-free,

$$\mathbb{1}_{\text{feas}} = \prod_{k \in \mathcal{K}} \mathbb{1}[\|\mathbf{p}_W^{E_k, \star} - \mathbf{p}_W^{E_k}\| \leq \epsilon_p] \mathbb{1}\left[\arccos\left(\frac{\text{tr}(\mathbf{R}_W^{E_k, \star \top} \mathbf{R}_W^{E_k}) - 1}{2}\right) \leq \epsilon_R\right] \mathbb{1}[\text{collision-free}], \quad (\text{B.7})$$

optionally re-checking a few evenly spaced replay frames under the fixed base; among accepted candidates we keep the one minimizing the worst-arm position error. This widens the distribution of object positions seen during manipulation and improves spatial generalization.

Robot State Expansion. This form perturbs the robot’s own state at two stages. In the *navigation* phase we apply two perturbations. First, a new start p_s is sampled in an area-uniform cone opening backward from the approach direction $\hat{\mathbf{d}}_{\text{app}}$ (cone axis $-\hat{\mathbf{d}}_{\text{app}}$):

$$\begin{aligned} p_s &= p_e + R\sqrt{u}[\cos \phi, \sin \phi]^\top, & u &\sim \mathcal{U}(0, 1), \\ \phi &= \text{atan2}(-\hat{d}_{\text{app}, y}, -\hat{d}_{\text{app}, x}) + \Delta\theta, & \Delta\theta &\sim \mathcal{U}(-\alpha, \alpha). \end{aligned} \quad (\text{B.8})$$

To preserve the object-relative endpoint pose, the object and endpoint are co-rotated about $\mathbf{p}_o$ by a closed-form yaw $\Delta\psi$ chosen so that, on arrival, the base heads from p_s toward the rotated endpoint. With $D = \|p_s - p_o\|$, $r = \|p_e - p_o\|$, $\alpha_0 = \text{atan2}(p_{e,y} - p_{o,y}, p_{e,x} - p_{o,x})$, $C = \psi_e - \alpha_0$, and $\varphi = \text{atan2}(p_{o,y} - p_{s,y}, p_{s,x} - p_{o,x})$, the orbit angle α^* of the rotated endpoint satisfies

$$\sin(\alpha^* + C + \varphi) = \frac{r \sin C}{D}, \quad \Delta\psi = (\alpha^* - \alpha_0)_{\text{wrap}}, \quad (\text{B.9})$$

selecting the forward-facing root; the object, endpoint, and both end-effectors are then advanced by $\mathbf{R}_z(\Delta\psi)$ about $\mathbf{p}_o$, and the robot drives a straight, retimed path from p_s . The start yaw is chosen so the target is visible at the start, i.e. $\mathbb{1}[\arccos(\hat{\mathbf{d}}_{\text{head}}^\top \hat{\mathbf{d}}_{H \rightarrow o}) \leq \theta_{\max}] \mathbb{1}[\text{ray clear}] = 1$ (with a fallback if none qualifies), and the trajectory is then accepted on its densified path’s collision check alone:

$$\mathbb{1}_{\text{nav}} = \mathbb{1}[\text{path collision-free}]. \quad (\text{B.10})$$

Second, at a random navigation frame we inject bounded Gaussian noise into the arm joints, graduated proximal-to-distal (small at the shoulder, larger at the wrist), followed by an ease-out recovery back to the carried pose, with any held object moved rigidly by the corresponding forward-kinematics delta; this teaches the policy to correct arm drift that accumulates while the robot is moving. In the *manipulation* phase, the approach-to-object segment is replaced by a cone-cycloid EE trajectory [Wang et al., 2025] whose start is sampled inside a cone behind the target gripper pose and whose end matches the original target; per-frame inverse kinematics solves the active arm, and the trajectory is accepted only if every frame meets tight position and rotation tolerances and is collision-free. Together these expansions teach the policy not only to execute the nominal trajectory but to recover to it from the realistic drift that accumulates over long horizons.

C Progress Score and Subtask Evaluation

C.1 Progress Score Definition

We provide the detailed definition of the progress score for each task here.

- *Lunch Box*: This task is scored out of 4: **1.** the robot needs to successfully navigate to the lunch box and grasp the lunch box edge with left gripper; **2.** successfully close the lunch box using right gripper; **3.** open left gripper, retract left arm and succeed to grasp the closed lunch box with right gripper; **4.** take the lunch box, navigate to the kitchen sink and place it in the sink.
- *Utensil*: The task is scored out of 4: **1.** the robot searches for the knife and navigates to it; **2.** grasps the knife precisely with its right hand; **3.** carries the knife and navigates to the storage basket; **4.** places the knife in the basket.
- *Drop Trash*: The task is scored out of 2. **1.** the robot needs to first search and navigate to the empty cola can; **2.** take the can, find the trash can, navigate to the trash can and throw it in.
- *Pour*: The task is scored out of 2. **1.** it needs to first navigate to the teapot and grasp it; **2.** then take the teapot to the cup and pour into it.
- *Fridge*: The task is scored out of 3. **1.** the robot needs to first navigate to the grocery bag and grasp it; **2.** navigate to the fridge, and store the bag in it; **3.** close the fridge door.

C.2 Subtask Success Rate

The per-subtask success rate for each real-world task is reported in Figure 6, comparing **WANDA** with the ablation without Corrective State Expansion. As tasks run without resets, for subtask k we report the fraction of rollouts that complete it among those that reach it (i.e. that completed subtask $k-1$), annotating each bar with the raw count; a method that never reaches a subtask has an undefined rate (0/0), drawn as a short stub on the axis. **WANDA** completes each individual subtask with high reliability, including the contact-rich late stages, whereas without Corrective State Expansion the policy is markedly weaker at every stage and often fails to reach the later ones. This shows Corrective State Expansion curbs the error accumulation of long-horizon mobile manipulation.

C.3 Scene Generalization Subtask Success Rate

Figure 8 breaks the scene-generalization evaluation in Figure 7(b) down by subtask. We train **WANDA** with Marble-generated worlds and evaluate it on four tasks in three real scenes (10 rollouts per task and scene), reporting the same per-subtask success rate as above. **WANDA** completes subtasks across all three novel scenes, confirming that synthesizing data on generated worlds transfers to the corresponding real environments.

D Multi-View Videos of Synthesized Data

We show the visualization of head and wrist camera videos from our synthesized data in Figures 16–18. Our approach achieves realistic rendering of the visual observations at high resolution (head camera: 1280×800 , wrist cameras: 800×480). The gap between rendered and real observations is negligible at the lower 224×224 resolution fed into the policy. For each trajectory we show eight frames sampled along the trajectory for the head and the two wrist cameras. Figure 16 shows the reconstructed real scenes for all five tasks, and Figures 17 and 18 show the world-lab scenes for *Pour*, *Lunch Box*, and *Drop Trash*.

E Per-Task Data-Generation Cost

Table 5 breaks down the compute to generate each real-world dataset, for the reconstructed source scene and the Marble-generated worlds. *Data* is the duration of the generated 30 fps multi-view video; *Plan* is wall-clock planning time on a single 192-core AMD EPYC 9R14 node; *FG* and *BG* are the foreground (Isaac Sim) and background (Gaussian-splat) render time, and *Render* their sum, in L40S GPU-hours. In total, generating 251.6 data-hours costs 24.6 planning hours and 765 GPU-hours.


Table 5: Per-task data-generation cost.

Task	Scene	Data (h)	Plan (h)	FG (GPU-h)	BG (GPU-h)	Render (GPU-h)	GPU-h / data-h
Lunch Box	Source	15.3	0.2	28	34	62	4.05
	Worlds	85.0	4.3	153	69	222	2.61
Drop Trash	Source	13.0	0.4	23	26	50	3.81
	Worlds	24.4	9.7	44	20	64	2.61
Pour	Source	17.8	0.6	32	36	68	3.84
	Worlds	49.7	4.2	89	40	130	2.61
Utensil	Source	37.0	3.4	67	68	134	3.63
Fridge	Source	9.4	1.7	17	18	35	3.72
Total		251.6	24.6	453	311	765	3.04

Foreground rendering is robot-only and thus task-independent. Background rendering is cheaper for generated worlds, whose lighter splats render about $3\times$ faster, lowering their cost to 2.61 GPU-hours per data-hour against 3.6–4.1 for source scenes. Planning accounts for only a small fraction of the total cost.

Real Scenes

**Drop
Trash**



**Lunch
Box**



Utensil



Fridge



Pour




Figure 16: Representative frames from the real-scene trajectories.

World-Lab Scenes


**Pour
Scene 1**



**Pour
Scene 2**



**Pour
Scene 3**



**Pour
Scene 4**



**Lunch
Box
Scene 1**



**Lunch
Box
Scene 2**




Figure 17: Representative frames from the world-lab trajectories.

World-Lab Scenes (continued)

**Lunch
Box
Scene 3**




**Lunch
Box
Scene 4**



**Drop
Trash
Scene 1**



**Drop
Trash
Scene 2**



**Drop
Trash
Scene 3**



Figure 18: Representative frames from the world-lab trajectories (continued).

F Cross-embodiment Data Generation

Figure 19 shows the generated visual observations for the cross-embodiment data generation from Agibot G1 to Linearbot on *Drop Trash*, described in Section 4.4. Despite a large kinematic gap (different DoFs) and a large visual gap (different camera extrinsics and intrinsics) between the two robots, **WANDA** generates reasonable visual observations.




Figure 19: Cross-embodiment data generation visual observations from Agibot to Linearbot.

G Policy Training Details

We train our policies by full-parameter fine-tuning of $\pi_{0.5}$ [Physical Intelligence et al., 2025] from the public pi05_base checkpoint on data generated by **WANDA**. $\pi_{0.5}$ is a PaliGemma-based vision-language-action model—a Gemma-2B vision-language backbone with a Gemma-300M action expert. The policy consumes three 224×224 RGB views (the head camera and the two wrist cameras), the robot proprioceptive state, and a tokenized language instruction (up to 200 tokens), and predicts an action chunk of horizon 32. The proprioceptive-state and action dimensionality is 20 for the dual-arm Agibot G1, padded to the model’s 32-dimensional action space.

When fine-tuning $\pi_{0.5}$, we use AdamW ($\beta_1=0.9$, $\beta_2=0.95$, $\epsilon=10^{-8}$, weight decay 10^{-10} , global gradient-norm clip 1.0) with a cosine-decay learning-rate schedule (peak 5×10^{-5} , 1,000-step warmup, decaying to 2.5×10^{-6}). We train for 30,000 steps at batch size 256 on 4 nodes of $8 \times$ NVIDIA H100 GPUs, taking roughly 12 hours. Training uses mixed precision: the trainable weights and AdamW state are kept in float32, while the forward and backward passes compute in bfloat16. Table 6 summarizes the full set of hyperparameters.

Each task is specified by a fixed natural-language instruction that conditions the policy throughout the episode. Table 7 lists the instructions for the five real-world tasks.

H Prompt

H.1 Inpaint Prompt

Inpainting Prompt. The inpainting model is instructed to remove all bright green masked regions and any visible robot parts, including arms, grippers, torso, mobile base, links, or other robot body components. It must reconstruct the missing background so that it appears natural, seamless, and consistent with the surrounding pixels. The prompt explicitly requires that pixels outside the removed masked or robot regions remain unchanged.

Table 6: $\pi_{0.5}$ fine-tuning hyperparameters.

Hyperparameter	Value
Initialization	pi05_base checkpoint
Fine-tuning	full-parameter
Training steps	30,000
Batch size	256
Optimizer	AdamW ($\beta_1=0.9$, $\beta_2=0.95$, $\epsilon=10^{-8}$, wd 10^{-10})
Gradient-norm clip	1.0
LR schedule	cosine decay
Peak / final LR	5×10^{-5} / 2.5×10^{-6}
Warmup steps	1,000
Action chunk horizon	32
Precision	mixed (bf16 compute, fp32 weights)

Table 7: Language instructions for the real-world tasks.

Task	Language instruction
Lunch Box	“Close the lunch box on the table, pick it up, and throw it into the sink.”
Utensil	“Pick up the knife from the plate and place it in the basket.”
Drop Trash	“Pick up the cola can on the table and drop it into the recycling bin.”
Fridge	“Pick up the bag on the table, put it in the refrigerator, and close the refrigerator.”
Pour	“Pick up the teapot and pour water into the cups.”

Judging Prompt. A separate reviewer model evaluates the generated result by comparing it with the original image. The reviewer checks whether:

1. all bright green masked regions were fully removed;
2. all visible robot parts were fully removed;
3. the filled background looks natural and seamless;
4. no unintended changes occurred outside the original masked or robot regions.

The reviewer must output either KEEP or RETRY, followed by a concise reason.

Corrective Prompt. If the result fails review, a corrective prompt is generated using the previous inpainting prompt, the original image, the failed result, and the failure reason. The corrected prompt focuses on fixing the observed failure while preserving the core constraints: removing green masks and robot parts, completing the background naturally, and avoiding edits outside the masked or robot regions.

H.2 Initial Configuration Prompt

Task Randomization Prompt. The pipeline takes a user-provided natural-language task description, such as pick up the cola can, and uses it as the main task prompt for selecting scene regions relevant to randomized manipulation. This task description is passed to the task-driven reasoning stage, which identifies candidate objects, surfaces, or spatial regions that are important for the requested task.

Task-Driven Region Selection. Given rendered RGB-D views of the scene and the task description, the pipeline runs task-driven reasoning, segmentation, backprojection, and top- k region selection. The goal is to find the most task-relevant randomization regions in the 3D scene. The selected regions are saved in a manifest file containing the top-ranked candidate regions.

Scene Rendering Context. Before task reasoning, the pipeline renders multiple RGB-D views of the input mesh. For WorldLab scenes, orbit views are rendered using configurable camera parameters such as azimuth count, field of view, orbit origin, and up-axis. For real scenes, camera poses are loaded from a corresponding `cameras.json` file.

Scale and Packaging Prompt Context. For non-real scenes, the pipeline estimates scene scale using a vision-language model. The support surface is described by the configurable label `PIPELINE_SUPPORT_LABEL`, whose default value is `tabletop surface`. This label guides the scale-estimation stage when interpreting the selected scene regions.

Final Output. The pipeline packages the rendered images, depth maps, selected and scaled regions, visualization files, and debug candidates into the final output directory. The retained outputs include:

- `images/`, `depth/`, and `depth_mm/`;
- `orbit_views.json`;
- `scaled_assets/` and `scaled_regions/`;
- `visualization/` and `debug_candidates/`.