



SimLauncher: Launching Sample-Efficient Real-world Robotic Reinforcement Learning via Simulation Pre-training

Mingdong Wu*, Lehong Wu*, Yizhuo Wu*, Weiyao Huang, Hongwei Fan, Zheyuan Hu, Haoran Geng, Jinzhou Li, Jiahe Ying, Long Yang, Yuanpei Chen, Hao Dong

Abstract—Autonomous learning of dexterous, long-horizon robotic skills has been a longstanding pursuit of embodied AI. Recent advances in robotic reinforcement learning (RL) have demonstrated remarkable performance and robustness in real-world visuomotor control tasks. However, applying RL in the real world faces challenges such as low sample efficiency, slow exploration, and significant reliance on human intervention. In contrast, simulators offer a safe and efficient environment for extensive exploration and data collection, while the visual sim-to-real gap, often a limiting factor, can be mitigated using real-to-sim techniques. Building on these, we propose SimLauncher, a novel framework that combines the strengths of real-world RL and real-to-sim-to-real approaches to overcome these challenges. Specifically, we first pre-train a visuomotor policy in the digital twin simulation environment, which then benefits real-world RL in two ways: (1) bootstrapping target values using extensive simulated demonstrations and real-world demonstrations derived from pre-trained policy rollouts, and (2) incorporating action proposals from the pre-trained policy for better exploration. We conduct comprehensive experiments across multi-stage, contact-rich, and dexterous hand manipulation tasks. Compared to prior real-world RL approaches, SimLauncher significantly improves sample efficiency and achieves near-perfect success rates. We hope this work serves as a proof of concept and inspires further research on leveraging large-scale simulation pre-training to benefit real-world robotic RL.

I. INTRODUCTION

For decades, researchers have been captivated by the pursuit of embodied AI that could seamlessly generalize across diverse tasks, demonstrate dexterous manipulation, and achieve flawless performance upon deployment. While significant strides have been made in generalizable learning [1], [2], [3], [4] and dexterous manipulation [5], [6], [7], the challenge of robust manipulation—ensuring consistent success across real-world conditions—remains an underexplored frontier, yet one that is indispensable for the industrialization of embodied intelligence.

While directly applying reinforcement learning to robotic tasks in real-world settings is a conceptually promising approach for acquiring robust robot policies, training policies in physical environments remains both unsafe and costly.

Mingdong Wu, Lehong Wu, Yizhuo Wu, Weiyao Huang, Hongwei Fan, Jinzhou Li, Jiahe Ying, Long Yang and Hao Dong are with the Center on Frontiers of Computing Studies, School of Computer Science, Peking University, also with PKUAgibot Lab. Haoran Geng is with the University of California, Berkeley. Zheyuan Hu is with the Robotics Institute at Carnegie Mellon University. Yuanpei Chen is with the PKU-Psibot Joint Lab.

*indicates equal contribution.

Corresponding to hao.dong@pku.edu.cn.

Even with carefully designed, compliant controllers to address safety concerns [8], the challenges of exploration still hinder sample efficiency. Recent advancements [9], [10] have demonstrated that integrating human-collected demonstrations to bootstrap the critic or using human interventions to guide exploration can significantly enhance learning efficiency, enabling near-optimal visuomotor policies for a diverse range of precise and dexterous skills. However, these methods rely on labor-intensive data collection and manual interventions, making them costly and difficult to scale.

How can we enhance data coverage for real-world robotic reinforcement learning bootstrapping and improve exploration while minimizing human labor costs? One potential approach is to leverage simulator. On the one hand, significant advancements have been made in automatically generating large-scale robotic task trajectories in simulation [11], [12], [13], [14], [15], [16], [17], [18], [19], making it increasingly feasible to develop a simulation-pretrained generalist policy. On the other hand, Real-to-Sim-to-Real techniques have been extensively studied in recent works [20], [21], [22], [23], [24], [22], [25] and have shown promising results in mitigating the Sim-to-Real gap [26], [27] and building a digital twin more efficiently [28], [29], [30]. These insights lead us to the central research question of this study:

How can simulation-pretrained policies, together with a digital twin, improve the sample efficiency of real-world robotic reinforcement learning?

To this end, we conduct a proof-of-concept study that restricts our task scenarios to fixed objects and backgrounds and pre-trains a specialized policy within the digital twin.

We propose **SimLauncher**, a simple yet effective vision-based reinforcement learning framework that leverages a simulated pre-trained policy along with a digital twin to improve real-world RL’s sample efficiency. We simulated rollouts generated by the pre-trained policy as demonstrations for critic bootstrapping—providing hundreds of trajectories that significantly expand state coverage. To prevent the critic from over-exploiting the task-irrelevant features arising from differences between simulated demonstrations and the real-world replay buffer, we further incorporate a limited number of real-world rollouts from the pre-trained policy to regularize training. Moreover, following IBRL [31], we use the pre-trained policy to accelerate exploration by selecting the action with the higher critic score between those proposed by the pre-trained and the RL policy. Unlike

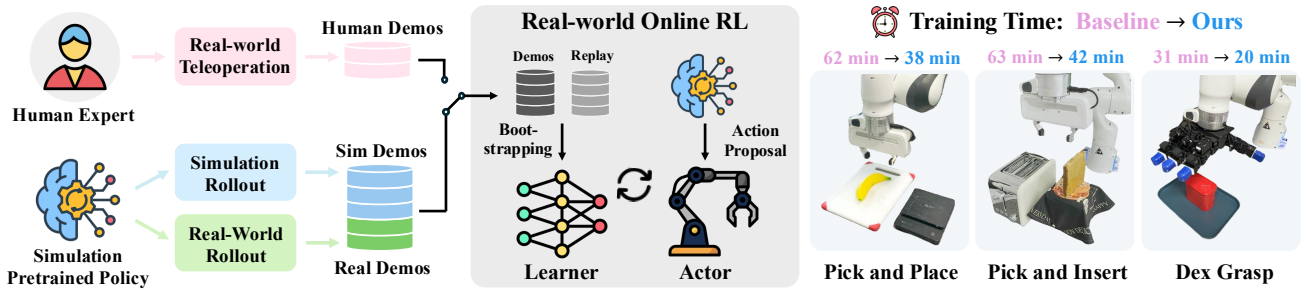


Fig. 1: **Illustration of our motivation.** Given a simulation-pretrained policy, SimLauncher leverages simulated and real-world rollouts as demonstrations for critic bootstrapping and incorporates the policy for action proposal. SimLauncher significantly improves the sample efficiency of real-world RL compared with conventional RL methods using human-collected data.

existing methods that integrate state-based simulation digital twins with real-world RL, SimLauncher adopts a vision-based setting, offering enhanced robustness and adaptability.

We conduct extensive experiments on three challenging real-world tasks: a multi-stage task, a precise manipulation task, and a dexterous hand manipulation task with high-dimensional action space. SimLauncher outperforms conventional hybrid RL baselines that rely on human-collected demonstrations. Our analysis suggests that simulated data alone can effectively support bootstrapping.

II. RELATED WORKS

A. Real-World Robotic Reinforcement Learning

Real-world robotic reinforcement learning (RL) requires sample-efficient algorithms for high-dimensional inputs like onboard perception, with easy reward and reset specification. Several methods have shown efficient learning from scratch in the real world [10], [9], [31], [32] or fine-tuning generalist policies [33], [34]. Recent advances in reset-free learning [35], [36], [37], [38] aim to minimize human intervention. While prior work has focused on improving sample efficiency in off-policy RL via offline pretraining [39], [40], [41], hybrid RL [32], [31], or leveraging foundation model priors [42], our approach accelerates real-world RL by incorporating simulation. In this line, previous studies have used simulation-trained value functions for exploration [43], simulation-pretrained actors for exploration efficiency [44], or fine-tuning pre-trained actors with significant unlearning periods [45]. In contrast, we propose leveraging a simulation-pretrained policy for both exploration and generating demonstration rollouts to bootstrap critic training. To our knowledge, this is the first bootstrapping-based approach that uses simulation to enhance real-world RL sample efficiency. Furthermore, unlike prior state-based methods, we consider a vision-based setting which improves robustness and generalizability across diverse real-world scenarios.

B. Real-to-Sim-to-Real Approaches for Policy Learning

Sim2Real transfer is a promising and widely explored approach for robotic policy learning [46] [47] [48]. However, the visual and physical discrepancies between simulation and the real world present significant challenges [46] [49].

To address the visual gap, [20] propose the first Real2Sim2Real framework, utilizing augmented demonstrations synthesized from a digital twin to enhance the robustness of behavioral cloning. To reduce the burden of recon-

structing a digital twin, ACDC [50] introduces a framework for retrieving digital assets that exhibit similar geometric and semantic affordances to the target task scenario.

Inspired by the promising results of 3D Gaussian Splatting (3DGS), recent studies have incorporated 3DGS for Real2Sim reconstruction [21], [22], [23], enabling training of RGB-based visuomotor policies via reinforcement learning [24], reconstructing articulated robot arms [22], and augmenting real-world data. Additionally, another line of research explores leveraging Real2Sim for scalable, high-quality data generation [51], [25], [52], facilitating future research on training generalist policies in simulation.

Our work integrates 3DGS-based Real2Sim to train a visuomotor policy and generate simulated demonstrations for the target task. Unlike previous approaches, we utilize the trained policy for action proposal in real-world RL, while the generated demonstrations are used to bootstrap the critic.

III. METHOD

A. Overview

Preliminaries. We model our robotic tasks as Markov Decision Processes $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \rho, \mathcal{P}, r, \gamma\}$, where \mathcal{S} is the observation space, \mathcal{A} is the action space, $\rho(s_0)$ is the distribution over initial states, \mathcal{P} is the transition probability function, $r : \mathcal{S} \times \mathcal{A} \rightarrow \{0, 1\}$ is the sparse reward function, and γ is the discount factor. Reinforcement learning (RL) aims to find the optimal policy π that maximizes the cumulative reward $E \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) \right]$. The core algorithm of our method is RLPD [32], a hybrid RL algorithm that bootstraps on prior data. For each batch, RLPD performs 50/50 sampling from the replay buffer and the demo buffer. The following objectives are optimized for the parametric critic $Q_\varphi(s, a)$ and actor $\pi_\theta(a|s)$:

$$y = r + \gamma E_{a' \sim \pi_\theta} [Q_{\tilde{\varphi}}(s', a')], \quad (1)$$

$$\mathcal{L}_Q(\varphi) = E_{s, a, s'} \left[(Q_\varphi(s, a) - y)^2 \right], \quad (2)$$

$$\mathcal{L}_\pi(\theta) = -E_s \left[E_{a \sim \pi_\theta} [Q_\varphi(s, a) - \alpha \log \pi_\theta(a|s)] \right], \quad (3)$$

where $Q_{\tilde{\varphi}}(s, a)$ is the target Q network, and α is the entropy temperature controlling policy randomness.

Overview. We present a proof-of-concept study on improving the sample efficiency of real-world reinforcement learning (RL) by leveraging a simulation-pretrained policy and a digital twin. Specifically, we focus on constrained environments with fixed objects and backgrounds, where we

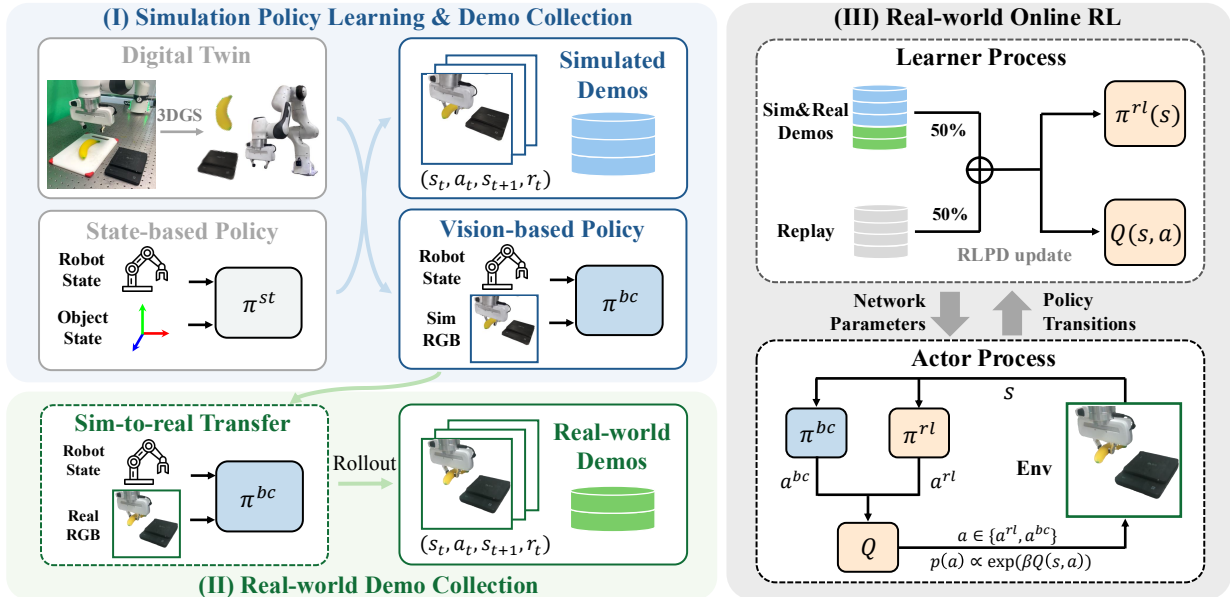


Fig. 2: **Overview of SimLauncher.** In simulation, we collect simulated demos and train a vision-based policy for each task. We then rollout the pre-trained policy in the real world to collect real demos. The simulated and real-world demos are used for critic bootstrapping in real-world RL. The pre-trained policy also provides action and bootstrap proposal for the actor.

pretrain a task-specific policy within the corresponding digital twin. Building on this foundation, we propose a hybrid RL framework based on two key designs: (1) utilizing rollouts from the pretrained policy in both simulation and the real world to bootstrap the critic’s learning, and (2) employing the pretrained policy to propose alternative actions for online exploration and critic bootstrapping, analogous to IBRL [31]. Section III-B details the digital twin reconstruction process and visuomotor policy training. In Section III-C, we outline our demonstration collection strategies and real-world RL pipeline, which integrates the pre-trained policy and digital twin. We evaluate our approach on a diverse set of manipulation tasks, including **Pick and Place**, **Pick and Insert**, and **Dex Grasp**, as illustrated in Fig. 3.

B. Simulation Policy Learning and Demo Collection

Digital Twin Environments. We initially construct digital twin environments for real-world tasks to enhance the sim-to-real transferability of subsequently trained visuomotor policies. For **Pick and Place** and **Pick and Insert**, we employ 3D Gaussian Splatting (3DGS) to reconstruct the geometry and texture of robots and task-related objects while using Isaac Gym [28] for physics simulation. To render visual observations under different robot and object states, we follow [21] and transform the objects’ Gaussian kernels based on their poses. For **Dex Grasp**, we use Mujoco [53] for both physics simulation and visual rendering.

Pre-training Visuomotor Policy in Simulation. Our visuomotor policy is pre-trained using a distillation-based approach, following PartManip [54]. For each task, we first train a state-based policy with privileged information (proprioceptive data and object states) using RL. This policy generates successful trajectories in simulation, which serve as visual demonstrations. We then apply behavior cloning (BC) to these demonstrations to train a visuomotor policy that

takes RGB images and robot proprioceptive data as input.

Mitigating the Sim-to-Real Gap. For the *physical* sim-to-real gap, we first calibrate physical parameters, camera pose, and robot controller by rolling out the same sequence of actions in both the simulator and the real world and then comparing the trajectories. For the *visual* sim-to-real gap, we apply the following strategies: (1) During simulated demonstration collection, we randomize the camera pose within a small range to simulate camera extrinsic calibration errors. (2) When training the behavior cloning policy, we use random cropping and color jittering as data augmentation. (3) We mask out the background in both simulated and real-world observations. For real-world policy rollout and RL training, we annotate the first frame’s observations to generate initial segmentation masks using SAM2 [55]. Leveraging SAM2’s efficient mask-tracking capability, subsequent masks can be obtained within 0.05s, making it compatible with our 10Hz control frequency.

C. Real-World Online RL with Simulation Bootstrapping

SimLauncher is a simple and effective approach that leverages a simulation-pretrained policy to perform rollouts in both simulation and the real world, enabling the bootstrapping of the critic. Additionally, it utilizes the pre-trained policy to propose actions for online exploration. The corresponding pseudo-code is provided in Appendix B. The key design choices are detailed below:

Simulated Demo Collection. Simulation provides a safe and efficient environment for data collection. We leverage simulation rollouts for critic bootstrapping, denoted as \mathcal{D}_{sim} . Compared to limited real-world demos from policies or humans, simulated demos offer broader coverage of initial conditions and intermediate states. Our base approach relies on *success-only* demos, but we also explore an extension:




Fig. 3: **Task illustrations, initialization ranges, and common failure modes.** (A) Pick and Place. This task involves relocating a banana to an electronic scale. Common failure includes the banana slipping off the scale. (B) Pick and Insert. This task involves grasping a toast and inserting it into the correct slot of a toaster. Common failure includes the toast getting stuck on the toaster edge. (C) Dex Grasp. This task involves coordinating multiple fingers to achieve a force closure on a can and lift it by 5 cm. Common failure includes the can slipping off during lifting or being knocked over by the hand.

collecting rollouts uniformly during state-based policy training with post-rendered image observations, referred to as *hybrid* demos. With better state coverage, hybrid demos improve performance, as discussed in IV-D.

Real-world Demo Collection. Relying solely on simulated demos can lead to the value underestimation of real-world transitions, negatively impacting online RL performance, as noted by [41]. To mitigate this, we deploy the BC policy in the real world to collect successful trajectories, denoted as \mathcal{D}_{real} . To balance the disparity between simulated and real-world demos, we sample equally from \mathcal{D}_{sim} and \mathcal{D}_{real} . Additionally, we incorporate the latest successful online trajectories into the real demo buffer to minimize the distribution gap between the replay buffer and the demo buffer. Each training batch is composed of 25% data from \mathcal{D}_{sim} , 25% from \mathcal{D}_{real} , and 50% from the replay buffer \mathcal{R} .

Action and Bootstrap Proposal. Following IBRL [31], we incorporate the behavior cloning policy during online interaction for better exploration. Specifically, we select either the BC action π^{bc} or the RL action π^{rl} by sampling from a Boltzmann distribution over Q-values, i.e., $p_Q(a) \propto \exp(\beta Q(s, a))$ for $a \in \{a^{bc}, a^{rl}\}$, where β is the inverse temperature that controls the sharpness of the distribution. The proposed action is also used for critic bootstrapping. This alters the value target in Eq. 1 to:

$$y = r + \gamma \max_{a' \in \{a^{rl}, a^{bc}\}} Q_{\tilde{\varphi}}(s', a'). \quad (4)$$

IV. EXPERIMENTS

Our experiments aim to evaluate SimLauncher’s effectiveness in improving the sample efficiency of real-world RL for learning a visuomotor policy. We use a diverse test suite of tasks that challenge online exploration. Specifically, we focus on three key questions: (1) Can SimLauncher, which leverages simulated demonstrations for bootstrapping and a pre-trained policy for exploration, achieve better sample efficiency than the current state-of-the-art hybrid RL methods that rely on human demonstrations? (Sec. IV-B) (2) How

crucial are SimLauncher’s key design choices, including simulated demos, a limited number of sim-to-real demos, and the action proposal (AP) module? (Sec. IV-C) (3) Why can simulation pre-trained policy benefit real-world RL despite the presence of the sim-to-real gap? (Sec. IV-D)

A. Experimental Setup

Tasks and Hardware Setup. We evaluate online exploration using a diverse suite of tasks. While autonomous reset and reward assignment are key goals for real-world RL, they are beyond this paper’s scope. To reduce confounding factors, we manually reset the object and assign a binary reward upon success. We set the target control frequency for all tasks to 10Hz. Figure 3 illustrates the task setups, randomization ranges, success criteria, and common failure modes. Details for each task are as follows:

Pick and Place: A daily multi-stage robotic manipulation task [56]. The setup features a Franka arm with a Franka Hand parallel gripper. Observations include two third-person RGB images and the gripper state. The action space comprises a 3-DoF delta TCP translation and a binary gripper action. The object is randomly initialized in each episode.

Pick and Insert: Another multi-stage task involves contact-rich manipulation, which adds complexity. The hardware setup, observation, and action space match those of Pick and Place. The object is randomly initialized in each episode.

Dexterous Hand Grasping: Grasping a columnar can requires leveraging both support and friction forces for stability. We use a Franka arm with a Leap Hand, freezing 5 of its 16 DoFs to prevent excessive finger collisions. Observations include a third-person camera image and hand joint state, while the action space consists of a 3-DoF delta TCP translation and 11-DoF delta hand joint. The object and hand wrist pose are randomly initialized per episode.

Baselines. We do not have a strict baseline that perfectly matches our setting. Nevertheless, we compare our method with state-of-the-art RL approaches [31], [32] that integrate human-collected offline data with online RL. **RLPD** is our core online RL algorithm, as detailed in III-A, demonstrating

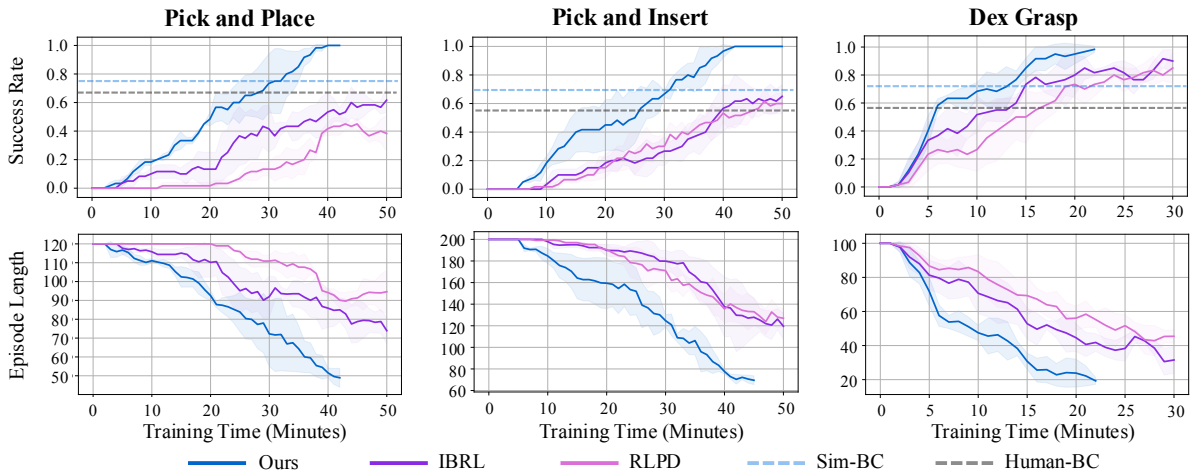


Fig. 4: **Comparison with baselines.** SimLauncher significantly outperforms state-of-the-art RL approaches that leverage human-collected demos and behavior-cloning methods. We report the mean and standard deviation over 3 seeds.

Task	Training Time (min)	Success Rate (%)		
		Ours	IBRL	RLPD
Pick and Place	37.5 ± 5.3	100.0	58.3 ± 10.3	20.0 ± 8.2
Pick and Insert	41.8 ± 2.4	100.0	53.3 ± 14.3	46.7 ± 8.5
Dex Grasp	20.0 ± 2.7	100.0	81.7 ± 2.4	65.0 ± 10.8

TABLE I: **Evaluation of RL-based methods.** The success rates are reported over 20 trials. For baselines, we evaluate the checkpoints when Ours achieve 100% success rate. We report the mean and standard deviation over 3 seeds.

superior results compared to offline-to-online RL methods. **IBRL** further incorporates a policy obtained via behavior cloning from the human demos and utilizes the BC policy for action and bootstrap proposal, as detailed in III-C. IBRL achieves state-of-the-art performance among hybrid RL approaches. Therefore, we do not compare with other hybrid RL algorithms in this work. Following [10], we provide 20 human demonstrations for IBRL and RLPD, which is a commonly used and user-affordable setting.

B. Comparison with Baselines

To answer Question 1, we compare our method with baselines across three tasks illustrated in Fig. 3. While our setup does not perfectly match the baselines, we ensure a fair comparison by setting the number of real demonstrations in SimLauncher to 20, identical to the baselines using human demos. Additionally, we align the usage of offline data in IBRL with our approach for fair comparison, rather than using buffer initialization as in the original paper. We present learning curves for all tasks in Fig. 4, where success rate and episode length are computed as running averages over the latest 20 episodes. Each experiment is **repeated 3 times** with different seeds, and we report the mean and variance. We also compare the earliest training checkpoint where our method reaches 100% success with baseline checkpoints at the same timestep, as shown in Tab. I.

As illustrated in Fig. 4, our method consistently outperforms all baselines across all tasks. When our method converges to near-perfect performance (100%), the strongest

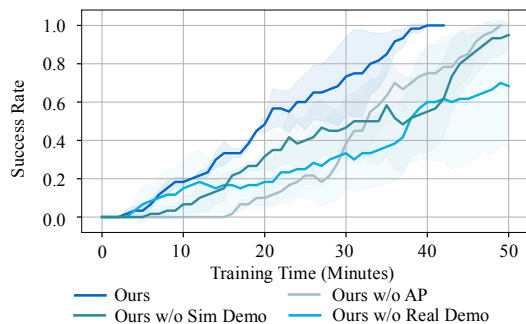


Fig. 5: Ablation study on our key design choices, 3 seeds.

baseline, IBRL, lags behind by 20–40% (Tab. I). These results demonstrate that our approach significantly improves sample efficiency compared to conventional hybrid RL methods that rely on human-collected data. Notably, on **Pick and Place** and **Pick and Insert**, which involve more sequential stages than **Dex Grasp**, our method exhibits a greater advantage, as shown in Fig. 4 and Tab. I. This aligns with our intuition that the state coverage provided by simulation pre-training is particularly beneficial for tasks with more complex exploration challenges. Surprisingly, both our method and baselines converge within 30 minutes on **Dex Grasp**, even faster than on **Pick and Place**. This may be due to a higher FPS of the actor and learner nodes during training, reaching 10 Hz and 12 Hz, respectively. However, our method requires 16000 learner steps to converge on **Dex Grasp**, which exceeds the 10000 steps needed on **Pick and Place**.

C. Ablation Studies

To answer Question 2, we compare our method with three ablated versions: **Ours w/o Sim Demo**, **Ours w/o Real Demo**, and **Ours w/o AP**, which respectively remove simulated demos, sim-to-real rollout demos, and action proposals by the simulation-pretrained policy, on **Pick and Place**. The evaluation follows the same procedure as in Sec. IV-B.

As shown in Fig. 5, all three ablated versions underperform the full method, underscoring the importance of each design choice. **Ours w/o Real Demo** suffers the most significant drop, likely due to critic overfitting, emphasizing




Fig. 6: Simulated demos alone enable effective bootstrapping, 3 seeds.

the necessity of real-world demos for regularization. Its critic assigns lower values to real-world interactions, possibly because bootstrapping from both successful simulated demos and low-success-rate replay exacerbates overfitting. **Ours w/o AP** struggles early on, as other methods leverage a pre-trained policy with a 73.3% success rate. However, after the “cold start,” it rapidly improves and converges faster than **Ours w/o Real Demo** and **Ours w/o Sim Demo**, likely due to the combined benefits of simulated and real-world demos enhancing stage coverage and stabilizing training.

D. Analysis

To answer Question 3, we study the robustness of our proposed action proposal and simulated demonstration bootstrapping approach to the sim-to-real gap. Below all the experiments are conducted on **Pick and Place**.

Takeaway 1: Scaling simulation pre-training improves sim-to-real policy transfer. We compare a BC policy trained on different numbers of simulated demos with one trained on human-collected demonstrations. As shown in Tab. II, the performance of the BC policy on simulated data consistently improves as the training data increases. Although Sim-BC underperforms Human-BC when trained on the same amount of data (e.g., 20 trajectories), likely due to the sim-to-real gap, simulation allows for scalable data collection. With extensive training (e.g., 1000 trajectories), Sim-BC achieves high real-world success despite the sim-to-real gap.

	Human	Simulation				
# of Demos	20	10	20	50	100	1000
Success Rate (%)	65	25	45	55	70	75

TABLE II: Scaling simulation data for behavior cloning.

Takeaway 2: Simulated demonstrations alone enable effective bootstrapping. A key concern when using simulated demos for bootstrapping is that the critic might overfit to simulation-specific features, allowing it to distinguish between simulated demos and the real-world replay buffer. To examine this, we compare the RLPD method using different sources of demos: 100 simulated demos, 20 human-collected real demos, and 20 simulated demos. As shown in Fig. 6, simulated demos alone can effectively bootstrap real-world RL. While RLPD with 20 simulated demos performs slightly worse than using 20 human demos, scaling up the number of simulated demos to 100 significantly improves efficiency, surpassing RLPD with 20 human demos.




Fig. 7: Enlarging state-coverage of simulated demonstrations can further improve sample-efficiency of **Ours**, 3 seeds.




Fig. 8: Visualization of state coverage in hybrid and success-only simulated demonstrations on the Pick and Place task.

Takeaway 3: Expanding state coverage in simulated demonstrations enhances bootstrapping. A key advantage of simulation is its scalability in data generation. This raises the question of whether increasing state coverage in simulated demonstrations can improve sample efficiency. To investigate this, we construct a “hybrid demo” by collecting rollouts uniformly throughout the state-based policy training process, and post-render the image observations. As illustrated in Fig. 8, this approach provides better state coverage. Furthermore, as shown in Fig. 7, using the hybrid demo buffer leads to slightly improved sample efficiency, suggesting that expanding state coverage in simulated demonstrations can further enhance bootstrapping.

V. CONCLUSION

We present SimLauncher, a vision-based real-world reinforcement learning approach that integrates digital twins to bridge simulation pretraining and real-world policy optimization. Simlauncher leverages simulated and real-world rollouts from the simulation pre-trained policy for critic bootstrapping and combines action proposals from the pre-trained policy for better exploration. Our method achieves superior sample efficiency compared to conventional hybrid RL approaches with an affordable scale of human demonstration across multi-stage tasks, precision manipulation challenges, and high-DoF dexterous manipulation.

Limitations and Future Works. SimLauncher’s adaptability is fundamentally constrained by the simulation environment, which struggles to accurately replicate highly dynamic scenarios, high-precision tasks, or interactions with deformable objects. Our current implementation relies on real-time segmentation for real-world observations. In future work, we may leverage large-scale training or domain randomization to reduce this dependency. Additionally, this proof of concept does not include a fully autonomous system for self-resetting or reward control; future works might incorporate advances in autonomous reinforcement learning, such as training a reward classifier [10] and a reset policy [35].

VI. ACKNOWLEDGEMENT

We thank Yankai Fu and Tianhao Wu for their insightful discussions. This project was supported by the National Youth Talent Support Program (8200800081) and the National Natural Science Foundation of China (62376006).

APPENDIX

A. Implementation Details

TABLE III: Implementation details across tasks

Parameter	Pick & Place	Pick & Insert	Dex Grasp
Randomization	$\pm 10\text{cm}$ (xy)	$\pm 8\text{cm}$ (xy)	$\pm 10\text{cm}$ (xy) $\pm 4\text{cm}$ (z)
Initial α	0.02	0.02	0.1
Critic ensemble	10	10	10
Sub-samples	2	2	10
Discount (γ)		0.97	
Agent		DrQv2	

B. PSEUDO CODE of SimLauncher

REFERENCES

- [1] A. Xie, L. Lee, T. Xiao, and C. Finn, "Decomposing the generalization gap in imitation learning for visual robotic manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3153–3160.
- [2] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, "Bc-z: Zero-shot task generalization with robotic imitation learning," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 991–1002. [Online]. Available: <https://proceedings.mlr.press/v164/jang22a.html>
- [3] Y. Ju, K. Hu, G. Zhang, M. Jiang, and H. Xu, "Robo-abc: Affordance generalization beyond categories via semantic correspondence for robot manipulation," in *European Conference on Computer Vision*. Springer, 2024, pp. 222–239.
- [4] W. Pumacay, I. Singh, J. Duan, R. Krishna, J. Thomason, and D. Fox, "The colosseum: A benchmark for evaluating generalization for robotic manipulation," *arXiv preprint arXiv:2402.08191*, 2024.
- [5] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.
- [6] Y. Chen, T. Wu, S. Wang, X. Feng, J. Jiang, Z. Lu, S. McAleer, H. Dong, S.-C. Zhu, and Y. Yang, "Towards human-level bimanual dexterous manipulation with reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5150–5163, 2022.
- [7] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, K. Ghasemipour, C. Finn, and A. Wahid, "Aloha unleashed: A simple recipe for robot dexterity," *arXiv preprint arXiv:2410.13126*, 2024.
- [8] J. Luo, E. Solowjow, C. Wen, J. A. Ojea, A. M. Agogino, A. Tamar, and P. Abbeel, "Reinforcement learning on variable impedance controller for high-precision robotic assembly," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3080–3087.
- [9] J. Luo, C. Xu, J. Wu, and S. Levine, "Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning," *arXiv preprint arXiv:2410.21845*, 2024.
- [10] J. Luo, Z. Hu, C. Xu, Y. L. Tan, J. Berg, A. Sharma, S. Schaal, C. Finn, A. Gupta, and S. Levine, "Serl: A software suite for sample-efficient robotic reinforcement learning," *arXiv preprint arXiv:2401.16013*, 2024.
- [11] J. Duan, W. Yuan, W. Pumacay, Y. R. Wang, K. Ehsani, D. Fox, and R. Krishna, "Manipulate-anything: Automating real-world robots using vision-language models," *arXiv preprint arXiv:2406.18915*, 2024.

Algorithm 1 SimLauncher

Hyperparameters: Entropy Temperature α , Action Proposal Inverse Temperature β , Gradient Steps G
Randomly initialize Critic φ_i (set targets $\varphi' = \varphi$) and Actor θ parameters.

Initialize empty replay buffer \mathcal{R}

Initialize real demo buffer \mathcal{D}_{real} with real-world demos and sim demo buffer \mathcal{D}_{sim} with simulated demos

while True do

Receive initial observation state s_0

for $t = 0, T$ **do**

Compute action $a_t^{rl} \sim \pi_\theta(\cdot|s_t)$ and $a_t^{bc} \sim \pi_{bc}(\cdot|s_t)$

Take action $a_t \sim \text{softmax}_{a \in \{a^{bc}, a^{rl}\}}(\beta Q(s, a))$

Store transition (s_t, a_t, r_t, s_{t+1}) in \mathcal{R}

for $g = 1, G$ **do**

Sample minibatch b_R of $\frac{N}{2}$ from \mathcal{R}

Sample minibatch $b_{\mathcal{D}_{real}}$ of $\frac{N}{4}$ from \mathcal{D}_{real}

Sample minibatch $b_{\mathcal{D}_{sim}}$ of $\frac{N}{4}$ from \mathcal{D}_{sim}

Concatenate $b_R, b_{\mathcal{D}_{real}}$ and $b_{\mathcal{D}_{sim}}$ to form batch b of size N set

For element j in b

$$y^j = r^j + \gamma \max_{a' \in \{a^{rl}, a^{bc}\}} Q_{\varphi}(s', a'). \quad (5)$$

Update φ minimizing loss:

$$\mathcal{L}_Q(\varphi) = \frac{1}{N} \sum_j (Q_\varphi(s, a) - y) \quad (6)$$

Update target networks $\varphi'_i \leftarrow \rho \varphi'_i + (1 - \rho) \varphi_i$

end for

Update θ maximizing objective:

$$\frac{1}{N} \sum_j Q_\varphi(s^j, \pi_\theta(s^j)) - \alpha \log \pi_\theta(\pi_\theta(s^j)|s^j), \quad (7)$$

end for

if trajectory succeed: **then**

Append trajectory into \mathcal{D}_{real}

end if

end while

- [12] Y. Wang, Z. Xian, F. Chen, T.-H. Wang, Y. Wang, K. Fragkiadaki, Z. Erickson, D. Held, and C. Gan, "Robogen: Towards unleashing infinite data for automated robot learning via generative simulation," *arXiv preprint arXiv:2311.01455*, 2023.
- [13] L. Wang, Y. Ling, Z. Yuan, M. Shridhar, C. Bao, Y. Qin, B. Wang, H. Xu, and X. Wang, "Gensim: Generating robotic simulation tasks via large language models," *arXiv preprint arXiv:2310.01361*, 2023.
- [14] P. Hua, M. Liu, A. Macaluso, Y. Lin, W. Zhang, H. Xu, and L. Wang, "Gensim2: Scaling robot data generation with multi-modal and reasoning llms," *arXiv preprint arXiv:2410.03645*, 2024.
- [15] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox, "Mimicgen: A data generation system for scalable robot learning using human demonstrations," *arXiv preprint arXiv:2310.17596*, 2023.
- [16] Z. Jiang, Y. Xie, K. Lin, Z. Xu, W. Wan, A. Mandlekar, L. Fan, and Y. Zhu, "Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning," *arXiv preprint arXiv:2410.24185*, 2024.
- [17] C. R. Garrett, A. Mandlekar, B. Wen, and D. Fox, "Skillgen: Auto-

- mated demonstration generation for efficient skill learning and deployment,” in *2nd CoRL Workshop on Learning Effective Abstractions for Planning*.
- [18] C. Garrett, A. Mandlekar, B. Wen, and D. Fox, “Skillmimicgen: Automated demonstration generation for efficient skill learning and deployment,” *arXiv preprint arXiv:2410.18907*, 2024.
 - [19] S. Nasiriany, A. Maddukuri, L. Zhang, A. Parikh, A. Lo, A. Joshi, A. Mandlekar, and Y. Zhu, “Robocasa: Large-scale simulation of everyday tasks for generalist robots,” *arXiv preprint arXiv:2406.02523*, 2024.
 - [20] M. Torne, A. Simeonov, Z. Li, A. Chan, T. Chen, A. Gupta, and P. Agrawal, “Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation,” *arXiv preprint arXiv:2403.03949*, 2024.
 - [21] M. N. Qureshi, S. Garg, F. Yandun, D. Held, G. Kantor, and A. Silwal, “Splatsim: Zero-shot sim2real transfer of rgb manipulation policies using gaussian splatting,” *arXiv preprint arXiv:2409.10161*, 2024.
 - [22] H. Lou, Y. Liu, Y. Pan, Y. Geng, J. Chen, W. Ma, C. Li, L. Wang, H. Feng, L. Shi, *et al.*, “Robo-gs: A physics consistent spatial-temporal model for robotic arm with hybrid representation,” *arXiv preprint arXiv:2408.14873*, 2024.
 - [23] X. Han, M. Liu, Y. Chen, J. Yu, X. Lyu, Y. Tian, B. Wang, W. Zhang, and J. Pang, “Resim: Generating high-fidelity simulation data via 3d-photorealistic real-to-sim for robotic manipulation,” *arXiv preprint arXiv:2502.08645*, 2025.
 - [24] Y. Wu, L. Pan, W. Wu, G. Wang, Y. Miao, and H. Wang, “Rl-gsbridge: 3d gaussian splatting based real2sim2real method for robotic manipulation learning,” *arXiv preprint arXiv:2409.20291*, 2024.
 - [25] S. Patel, X. Yin, W. Huang, S. Garg, H. Nayyeri, L. Fei-Fei, S. Lazebnik, and Y. Li, “A real-to-sim-to-real approach to robotic manipulation with vlm-generated iterative keypoint rewards,” *arXiv preprint arXiv:2502.08643*, 2025.
 - [26] P. Huang, X. Zhang, Z. Cao, S. Liu, M. Xu, W. Ding, J. Francis, B. Chen, and D. Zhao, “What went wrong? closing the sim-to-real gap via differentiable causal discovery,” in *Conference on Robot Learning*. PMLR, 2023, pp. 734–760.
 - [27] H. He, P. Wu, C. Bai, H. Lai, L. Wang, L. Pan, X. Hu, and W. Zhang, “Bridging the sim-to-real gap from the information bottleneck perspective,” *arXiv preprint arXiv:2305.18464*, 2023.
 - [28] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, *et al.*, “Isaac gym: High performance gpu-based physics simulation for robot learning,” *arXiv preprint arXiv:2108.10470*, 2021.
 - [29] Z. Zhou, J. Song, X. Xie, Z. Shu, L. Ma, D. Liu, J. Yin, and S. See, “Towards building ai-cps with nvidia isaac sim: An industrial benchmark and case study for robotics manipulation,” in *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice*, ser. ICSE-SEIP ’24. ACM, Apr. 2024, p. 263–274. [Online]. Available: <http://dx.doi.org/10.1145/3639477.3639740>
 - [30] K. Zakka, B. Tabanpour, Q. Liao, M. Haiderbhai, S. Holt, J. Y. Luo, A. Allshire, E. Frey, K. Sreenath, L. A. Kahrs, C. Sferazza, Y. Tassa, and P. Abbeel, “Mujoco playground,” 2025. [Online]. Available: <https://arxiv.org/abs/2502.08844>
 - [31] H. Hu, S. Mirchandani, and D. Sadigh, “Imitation bootstrapped reinforcement learning,” *arXiv preprint arXiv:2311.02198*, 2023.
 - [32] P. J. Ball, L. Smith, I. Kostrikov, and S. Levine, “Efficient online reinforcement learning with offline data,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 1577–1594.
 - [33] Y. Guo, J. Zhang, X. Chen, X. Ji, Y.-J. Wang, Y. Hu, and J. Chen, “Improving vision-language-action model with online reinforcement learning,” *arXiv preprint arXiv:2501.16664*, 2025.
 - [34] Y. Chen, S. Tian, S. Liu, Y. Zhou, H. Li, and D. Zhao, “Confrft: A reinforced fine-tuning method for vla models via consistency policy,” *arXiv preprint arXiv:2502.05450*, 2025.
 - [35] Z. Hu, A. Rovinsky, J. Luo, V. Kumar, A. Gupta, and S. Levine, “Reboot: Reuse data for bootstrapping efficient real-world dexterous manipulation,” *arXiv preprint arXiv:2309.03322*, 2023.
 - [36] A. Gupta, J. Yu, T. Z. Zhao, V. Kumar, A. Rovinsky, K. Xu, T. Devlin, and S. Levine, “Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6664–6671.
 - [37] K. Xu, Z. Hu, R. Doshi, A. Rovinsky, V. Kumar, A. Gupta, and S. Levine, “Dexterous manipulation from images: Autonomous real-world rl via substep guidance,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5938–5945.
 - [38] H. Zhu, J. Yu, A. Gupta, D. Shah, K. Hartikainen, A. Singh, V. Kumar, and S. Levine, “The ingredients of real-world robotic reinforcement learning,” *arXiv preprint arXiv:2004.12570*, 2020.
 - [39] A. Nair, A. Gupta, M. Dalal, and S. Levine, “Awac: Accelerating online reinforcement learning with offline datasets,” *arXiv preprint arXiv:2006.09359*, 2020.
 - [40] J. Yang, M. S. Mark, B. Vu, A. Sharma, J. Bohg, and C. Finn, “Robot fine-tuning made easy: Pre-training rewards and policies for autonomous real-world reinforcement learning,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 4804–4811.
 - [41] Z. Zhou, A. Peng, Q. Li, S. Levine, and A. Kumar, “Efficient online reinforcement learning fine-tuning need not retain offline data,” *arXiv preprint arXiv:2412.07762*, 2024.
 - [42] W. Ye, Y. Zhang, H. Weng, X. Gu, S. Wang, T. Zhang, M. Wang, P. Abbeel, and Y. Gao, “Reinforcement learning with foundation priors: Let the embodied agent efficiently learn on its own,” *arXiv preprint arXiv:2310.02635*, 2023.
 - [43] P. Yin, T. Westenbroek, S. Bagaria, K. Huang, C.-a. Cheng, A. Kobolov, and A. Gupta, “Rapidly adapting policies to the real world via simulation-guided fine-tuning,” *arXiv preprint arXiv:2502.02705*, 2025.
 - [44] A. Wagenmaker, K. Huang, L. Ke, B. Boots, K. Jamieson, and A. Gupta, “Overcoming the sim-to-real gap: Leveraging simulation to learn to explore for real-world rl,” *arXiv preprint arXiv:2410.20254*, 2024.
 - [45] Y. Zhang, L. Ke, A. Deshpande, A. Gupta, and S. S. Srinivasa, “Cherry-picking with reinforcement learning,” in *Robotics: Science and Systems*, 2023.
 - [46] W. Zhao, J. P. Queralt, and T. Westerlund, “Sim-to-real transfer in deep reinforcement learning for robotics: a survey,” in *2020 IEEE symposium series on computational intelligence (SSCI)*. IEEE, 2020, pp. 737–744.
 - [47] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik, “In-hand object rotation via rapid motor adaptation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1722–1732.
 - [48] J. Wang, Y. Yuan, H. Che, H. Qi, Y. Ma, J. Malik, and X. Wang, “Lessons from learning to spin” pens,”” *arXiv preprint arXiv:2407.18902*, 2024.
 - [49] R. Singh, A. Allshire, A. Handa, N. Ratliff, and K. Van Wyk, “Dextrah-rgb: Visuomotor policies to grasp anything with dexterous hands,” *arXiv preprint arXiv:2412.01791*, 2024.
 - [50] T. Dai, J. Wong, Y. Jiang, C. Wang, C. Gokmen, R. Zhang, J. Wu, and L. Fei-Fei, “Automated creation of digital cousins for robust policy learning,” *arXiv preprint arXiv:2410.07408*, 2024.
 - [51] Y. Mu, T. Chen, S. Peng, Z. Chen, Z. Gao, Y. Zou, L. Lin, Z. Xie, and P. Luo, “Robotwin: Dual-arm robot benchmark with generative digital twins (early version),” *arXiv preprint arXiv:2409.02920*, 2024.
 - [52] W. Ye, F. Liu, Z. Ding, Y. Gao, O. Rybkin, and P. Abbeel, “Video2policy: Scaling up manipulation tasks in simulation through internet videos,” *arXiv preprint arXiv:2502.09886*, 2025.
 - [53] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
 - [54] H. Geng, Z. Li, Y. Geng, J. Chen, H. Dong, and H. Wang, “Partmanip: Learning cross-category generalizable part manipulation policy from point cloud observations,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2978–2988.
 - [55] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryal, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, and C. Feichtenhofer, “Sam 2: Segment anything in images and videos,” *arXiv preprint arXiv:2408.00714*, 2024. [Online]. Available: <https://arxiv.org/abs/2408.00714>
 - [56] Q. Vuong, S. Levine, H. R. Walke, K. Pertsch, A. Singh, R. Doshi, C. Xu, J. Luo, L. Tan, D. Shah, *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models,” in *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@ CoRL2023*, 2023.