

DexPIE: Stable Dexterous Policy Improvement from Real-World Experience

Ruizhe Liao¹ Wenrui Chen^{1,†} Liangji Zeng¹ Haoran Lin¹ Fan Yang¹
 Kailun Yang¹ Yaonan Wang¹

¹Hunan University [†]Corresponding author

<https://siiuuuuuu.github.io/DexPIE>

Abstract: Dexterous manipulation presents substantial challenges for imitation learning due to its high-dimensional action space and complex contact-rich dynamics. Policies trained purely from demonstrations often suffer from compounding errors during deployment and require large amounts of expert data to achieve reliable performance. To move beyond the limitations of demonstration data, in this work, we propose DexPIE, a post-training framework for dexterous policy improvement from experience collected through real-world deployment. First, DexPIE enables effective exploration coverage through a dexterous-hand-adapted intervention system and multi-stage DAgger-style data collection across initial and intermediate task stages, providing reliable supervision for accurate policy evaluation. To reduce temporal noise between post-training rollouts and demonstration data, we introduce asynchronous inference in the relative action space, which better aligns rollout data with demonstrated behavior and allows the critic to learn a value function induced by a more consistent underlying policy. Finally, DexPIE improves the policy through conditioning on a continuous optimality indicator, allowing the policy to leverage the quality of data in a more fine-grained manner. Across three challenging real-world dexterous manipulation tasks, DexPIE achieves a 37% improvement in success rate over the demonstration-based reference policy, outperforming all baseline methods and demonstrating stronger robustness. The source code and dataset will be made publicly available.

Keywords: Real-World Reinforcement Learning, Dexterous Manipulation

1 Introduction

The rapid development of Vision-Language-Action (VLA) models has significantly advanced general-purpose robotic manipulation [1, 2, 3, 4]. However, most existing VLA methods still heavily rely on Imitation Learning (IL) [5, 6]. Beyond the requirement for large-scale demonstration data, IL suffers from compounding errors during deployment [7], which often leads to performance saturation. Moreover, simply incorporating DAgger-style post-training remains within the imitation-learning paradigm. Reinforcement Learning (RL) [8], in contrast, provides a promising alternative for enabling policies to improve autonomously through real-world deployment beyond expert demonstrations. This suggests that effective robot learning should not only leverage demonstration data, but also exploit experience collected from real-world rollouts, allowing the policy to correct deployment-time errors and complete tasks more robustly [9, 10, 11, 12].

Extending this paradigm to dexterous manipulation, however, remains challenging. Recent works [10, 9, 13, 12, 14] have shown that Human-In-the-Loop (HIL) RL post-training can effectively address exploration bottlenecks of policies trained from demonstrations, substantially alleviate Out-Of-Distribution (OOD) issues, and improve sample efficiency. Nevertheless, most of these methods are limited to parallel grippers and cannot be directly adapted to dexterous hands. Furthermore, effective RL for long-horizon manipulation requires informative and fine-grained reward signals. Sparse rewards, such as binary success indicators, are widely used in real-world reinforcement learning due to their simplicity and generality, but they introduce severe credit-assignment challenges in

long-horizon tasks. Existing works [15, 16, 17] mainly address this issue by manually or automatically annotating task stages and progress labels, *e.g.*, with human or LLM supervision. Recent work [9, 18] achieves more stable and scalable policy improvement through advantage-conditioned policy extraction. Nevertheless, they use binary optimality labels, which fail to preserve the relative ordering of action quality. In addition, previous research works [18, 19, 20, 21] show that improving train-inference consistency through asynchronous inference can enhance policy performance. We further observe that the resulting inconsistent deployment behaviors introduce temporal noise and distribution shift between demonstrations and deployment rollouts. We refer to this mismatch as the demonstration-deployment gap. Such a gap not only provides low-quality data for imitation learning but also leads to a heterogeneous mixture of mismatched behaviors in the post-training dataset, making critic learning difficult.

To address these challenges, we propose DexPIE, a stable dexterous policy improvement framework from experience collected through real-world deployment. First, to correct deployment-time errors and overcome exploration bottlenecks, we develop a human-following intervention system that enables intuitive human intervention from arbitrary robot states for dexterous manipulation. Second, under sparse rewards in long-horizon tasks, we introduce staged DAgger, which performs DAgger-style data collection from both the initial task stage and selected intermediate task stages. It can be viewed as a practical relaxation of the exploring-starts assumption [8], improving exploration coverage in long-horizon manipulation, providing intermediate anchors for critic learning. Furthermore, to bridge the demonstration-deployment gap, we extend training-time RTC [19] to the relative action space through a future-state-referenced relative action padding scheme. The improved temporal consistency allows the critic to learn a value function induced by a more consistent underlying policy. Finally, we introduce a continuous optimality function that enables stable policy improvement by exploiting fine-grained behavioral differences, instead of collapsing diverse behaviors into binary labels. Across three real-world dexterous manipulation tasks, after only one iteration, our method improves the success rate by more than 30% compared with the demonstration-based policy.

2 Related Work

Human-in-the-Loop Robot Learning. Imitation learning collects expert demonstrations through leader–follower robotic arms, VR devices, exoskeletons, and other teleoperation interfaces [5, 22, 23, 24, 25]. However, policies trained purely from expert demonstrations often fail to recover from errors accumulated during online execution. To address this issue, classical approaches build upon DAgger [26]. HG-DAgger [27] further allows a teleoperator to intervene when the policy enters undesirable states. Human-in-the-loop mechanisms have also demonstrated their effectiveness in real-world RL. Under sparse reward settings, human intervention can accelerate early-stage policy exploration and alleviate the exploration challenges faced by demonstration-based policies [13, 10, 14, 9]. Recent systems [28, 29, 30] improve teleoperation and intervention interfaces to enable intuitive intervention for gripper–arm systems. For dexterous manipulation systems, DexGrasp-VLA [31] and DexHiL [32] enable intervention through shared autonomy assisted by an autonomous hand policy and hand retargeting, respectively. However, these systems are based on incremental end-effector control, which is less intuitive than the leader-follower intervention mode, where the operator starts intervention from a pose aligned with the policy output. Our human-following intervention system is designed to provide intuitive corrective control from arbitrary robot states.

Algorithms and Systems for Real-World RL. Recent works have achieved promising policy improvement in simulation environments using PPO and its variants [33, 34, 35, 36]. However, due to the low sample efficiency and high noise of real-world interaction, such methods remain difficult to transfer successfully to real-world scenarios. Meanwhile, some works focus on reward specification or progress modeling in the real world through manually designed reward functions, VLM-based reward annotation, or stage-wise labeling [37, 34, 16, 17, 15]. Complementary to these approaches, we collect multi-stage data with staged DAgger, which provides intermediate anchors for critic learning and alleviates the credit assignment problem in long-horizon tasks. Since real-world datasets are often composed of multi-source heterogeneous data, another line of work studies offline and

off-policy reinforcement learning in real-world settings [12, 10, 14, 13, 11, 38], which can efficiently reuse historical trajectory data and improve sample efficiency, leading to significant progress in real-world RL. Some recent works further leverage world models as substitutes for real-world interaction [39, 40, 41, 42]. RECAP [9] improves the policy through binary advantage-label conditioning, achieving significant performance gains on real-world long-horizon tasks. Compared with binary labels, our continuous optimality function provides fine-grained supervision by exploiting the relative quality of trajectories.

3 Preliminaries

Reinforcement Learning. We consider the standard Reinforcement Learning (RL) setting [43], where an agent, represented by a policy $\pi(a_t | o_t)$, selects an action $a_t \in \mathcal{A}$ given an observation $o_t \in \mathcal{O}$. A trajectory is denoted by $\tau = (o_0, a_0, \dots, o_T)$, and the induced trajectory distribution is $\rho_\pi(\tau) = p(o_0) \prod_{t=0}^{T-1} \pi(a_t | o_t) p(o_{t+1} | o_t, a_t)$. The reward function is given by $r(o_t, a_t)$, abbreviated as r_t , and $\gamma \in [0, 1]$ denotes the discount factor. The discounted cumulative reward, or return, is defined as $R(\tau) = \sum_{t=0}^T \gamma^t r_t$. The goal of RL is to learn a policy that maximizes the expected return, *i.e.*, $J(\pi) = \mathbb{E}_{\tau \sim \rho_\pi(\tau)}[R(\tau)] = \mathbb{E}_{\tau \sim \rho_\pi(\tau)}[\sum_{t=0}^T \gamma^t r_t]$. The value function of policy π is then defined as $V^\pi(o_t) = \mathbb{E}_{\tau_t:T} [\sum_{l=t}^T \gamma^{l-t} r_l]$. We estimate the advantage of action a_t at observation o_t using an N -step return, written as $A^\pi(o_t, a_t) = \mathbb{E}_{\rho_\pi(\tau)} [\sum_{l=t}^{t+N-1} \gamma^{l-t} r_l + \gamma^N V^\pi(o_{t+N})] - V^\pi(o_t)$.

Product-Policy View of Optimality Guidance. To achieve stable improvement over a reference policy π_{ref} , we adopt the product-policy view, which performs policy improvement by constructing an improved target distribution around the reference policy. Specifically, rather than directly optimizing a regularized RL objective, we define the target policy $\hat{\pi}$ by reweighting π_{ref} with an optimality indicator I_t :

$$\hat{\pi}(a_t | o_t) \propto \pi_{\text{ref}}(a_t | o_t) \cdot p(I | o_t, a_t)^\beta, \quad (1)$$

where β denotes the guidance strength, and $p(I | o_t, a_t) = \frac{f(A^{\pi_{\text{ref}}}(o_t, a_t))}{Z(o_t)}$, where $Z(o_t) > 0$ is an action-independent normalizing constant. Following the product-policy improvement theorem in cfgRL [44], improvement over π_{ref} is guaranteed when f is chosen as a non-negative, monotonically increasing function of $A^{\pi_{\text{ref}}}(o_t, a_t)$. Taking the logarithm of Eq. (1) and differentiating with respect to the action yields an additive decomposition of the product-policy score:

$$\nabla_{a_t} \log \hat{\pi}(a_t | o_t) = \nabla_{a_t} \log \pi_{\text{ref}}(a_t | o_t) + \beta \nabla_{a_t} \log p(I | o_t, a_t). \quad (2)$$

By applying Bayes' rule to the optimality indicator $p(I | o_t, a_t)$, its action-gradient can be written as the difference between the optimality-conditioned score and the unconditional score. Substituting this result into Eq. (2), we obtain the guided score:

$$\nabla_{a_t} \log \hat{\pi}(a_t | o_t) = (1 - \beta) \nabla_{a_t} \log \pi_{\text{ref}}(a_t | o_t) + \beta \nabla_{a_t} \log \pi_{\text{ref}}(a_t | o_t, I). \quad (3)$$

This is analogous to Classifier-Free Guidance (CFG) [45]. In practice, we implement this guidance in a diffusion policy by conditioning the denoising network on the optimality indicator I .

4 Method

As shown in Fig. 1, our method aims to provide a complete post-training pipeline for dexterous manipulation. In Sec. 4.1, we integrate a human-in-the-loop data collection pipeline for dexterous manipulation. In Sec. 4.2, we introduce asynchronous inference in the relative action space to mitigate the demonstration-deployment gap. In Sec. 4.3, based on the offline dataset, we first train the critic using Monte Carlo return estimates. After obtaining stable advantage estimates, we perform policy improvement conditioned on the continuous optimality indicator.

4.1 Human-in-the-Loop Data Collection for Dexterous Manipulation

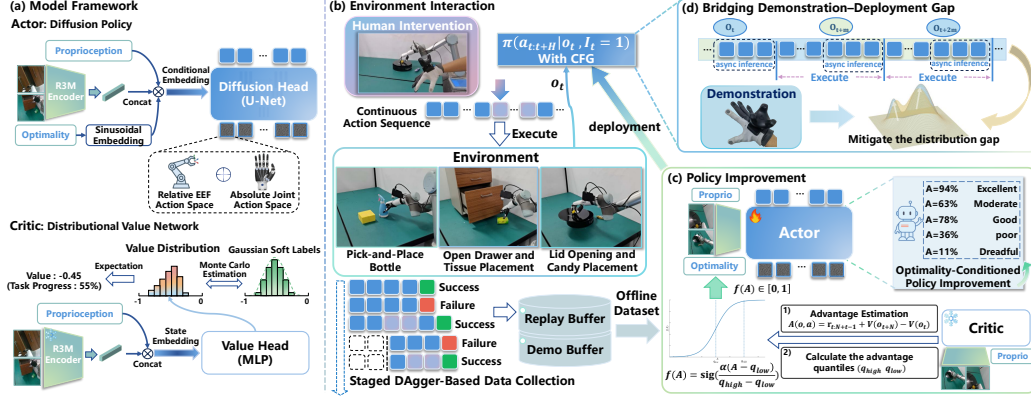


Figure 1: **Overview of DexPIE framework.** (a) The model architecture consists of an actor and a critic. The actor is an optimality-conditioned diffusion policy, with an action space defined as relative EEf actions [46] concatenated with absolute dexterous-hand joint actions, while the critic is a distributional value network. (b) The policy is warm-started with demonstration data, and CFG is applied to the optimality indicator. During deployment, failures or exploration bottlenecks trigger human intervention. We further perform DAGger-style data collection across different task stages. (c) After obtaining stable advantage estimates, we perform policy improvement conditioned on a continuous optimality indicator. (d) We reduce the distribution shift between post-training data and teleoperation demonstrations through asynchronous inference.

Human-Following Intervention System. To provide simple and intuitive corrections during policy deployment for error recovery and exploration bottlenecks, we implement a system integration that enables seamless switching between arm-hand teleoperation and policy execution. The system adopts a multi-threaded architecture, where the teleoperation system and the autonomous policy run in parallel, and execution is switched via keyboard commands. For the teleoperation interface, we use a Vive tracker to provide the wrist pose and Manus gloves to capture the full hand posture, which is then retargeted to the absolute joint space of the Inspire hand. To enable intuitive corrective interventions from arbitrary robot states, we introduce a simple human-as-follower strategy, inspired by the intervention mode transition in leader-follower arm control [29, 30]. As shown in Fig. 2, the operator first aligns the wrist orientation and hand gesture with the end-effector orientation and dexterous-hand posture predicted by the policy before taking over. Human intervention is then triggered via a keyboard command. Once triggered, the system records the current tracker pose $\mathbf{T}_{t_0}^v \in SE(3)$ and the current robot end-effector pose $\mathbf{T}_{t_0}^{ee}$ as the teleoperation reference poses. At each subsequent timestep t , the relative tracker motion $\Delta \mathbf{T}_t^v = (\mathbf{T}_{t_0}^v)^{-1} \mathbf{T}_t^v$ is computed and applied to the robot end-effector reference pose, resulting in $\mathbf{T}_t^{ee} = \mathbf{T}_{t_0}^{ee} \Delta \mathbf{T}_t^v$. Meanwhile, the policy-predicted hand action is replaced by the retargeted glove action for dexterous-hand control, with smoothing applied during the transition to reduce abrupt action changes at the switching moment. This human-as-follower design allows the operator to follow the robot state before intervention and take over from an aligned pose, thereby enabling a smooth transition into the intervention state and supporting intuitive corrective control.

Staged DAGger-based Data Collection. In long-horizon manipulation tasks with sparse rewards, the critic must infer the contribution of each action from delayed terminal outcomes, which often results in inaccurate estimates for intermediate and later-stage states under accumulated execution errors. To better approximate the exploring-starts assumption [8] and achieve sufficient exploration coverage, we introduce staged DAGger, where rollouts are initialized from both the initial state and selected intermediate stages for DAGger-style data collection. Specifically, after an autonomous rollout fails, we restore the environment to the same stage and intervene before the failure occurs,




Figure 2: **Human-as-follower strategy.** The operator aligns the wrist orientation and hand gesture with the robot end-effector and dexterous-hand posture before takeover, thereby enabling intuitive human corrective control to be initiated from arbitrary robot states.

collecting paired failed and corrected trajectories. Such paired trajectories provide sufficient positive and negative samples for critic learning. These later-stage trajectories shorten the effective horizon for value supervision and provide intermediate anchors for critic learning. Consequently, the long-horizon value estimation problem is decomposed into shorter stage-wise subproblems, allowing the critic to learn reliable progress-aware values that can be propagated backward to earlier task stages.

4.2 Bridging Demonstration-Deployment Gap via Asynchronous Inference

Synchronous inference can introduce latency-induced pauses or action stalls, disrupting the smooth action streams observed in training demonstrations. These temporal noise not only degrade the quality of data used for imitation learning but also lead to a heterogeneous mixture of mismatched behaviors in the post-training dataset, thereby complicating critic learning. To mitigate these issues, we extend training-time RTC [19] to a future-state-referenced relative action padding scheme, enabling smooth asynchronous inference that is temporally matched to the training action streams. Let $A_t = [a_t, a_{t+1}, \dots, a_{t+H-1}]$ denote an action chunk of horizon H , and let n denote the maximum inference delay. During training, we enable the masking mechanism with a certain probability p_m^{pre} , where the first i actions in A_t are randomly masked and $0 \leq i \leq n$. This encourages the policy to remain flexible under varying amounts of historical action context. During deployment, once the current action chunk A_t reaches the maximum-delay horizon, asynchronous inference is initiated with the current observation o_{t+m} . The remaining actions of A_t , originally relative to o_t , are transformed into the reference frame of o_{t+m} and used as the relative-action prefix of the next chunk A_{t+1} . While these remaining actions are executed, the policy predicts A_{t+1} asynchronously. Once A_t is completed, the prefix is removed from A_{t+1} , and the remaining actions are executed, producing a continuous action stream that is temporally aligned with the teleoperation.




Figure 3: **Future-State-Referenced relative action padding.** After asynchronous inference is triggered, the remaining actions in A_t are transformed from the reference frame of o_t to the current observation o_{t+m} and used as the relative-action prefix of the next chunk A_{t+1} .

4.3 Policy Improvement from Deployment Experience

Following the RECAP process, we next introduce how to perform policy evaluation on an offline dataset composed of autonomous rollouts and human interventions, and how to construct an optimality indicator from advantage signals for policy improvement.

Policy Evaluation. We adopt the progress-based reward used in $\pi_{0.6}^*$ [9] as our reward design. Its sparsity introduces a credit assignment problem. To mitigate this issue, staged DAgger provides empirical returns at different levels of task progress, which alleviates the difficulty of directly using Monte Carlo returns for supervision in long-horizon tasks. Meanwhile, it also introduces heterogeneous data by collecting trajectories from diverse initial states, behavior policies, and outcome qualities, which naturally induces multimodal value distributions. To address both properties, we employ a distributional value network to model the full value distribution [47, 9]. The value is obtained by computing the expectation of the output distribution. Using $R_t(\tau) = \sum_{t'=t}^T r_{t'}$ to denote the empirical return of a trajectory τ from time step t until termination, Considering the continuous and multimodal nature of value data, we model each empirical return as a Gaussian distribution and discretize it into a soft target over B value bins. Specifically, for each empirical return $R_t(\tau)$, we construct a Gaussian soft label over the discretized bins as

$$q_b(\tau, t) = \frac{\exp\left(-\frac{(v_b - R_t(\tau))^2}{2\sigma^2}\right)}{\sum_{j=1}^B \exp\left(-\frac{(v_j - R_t(\tau))^2}{2\sigma^2}\right)}, \quad b = 1, \dots, B, \quad (4)$$

where σ controls the smoothness of the soft target distribution. The resulting vector

$$\mathbf{q}(\tau, t) = [q_1(\tau, t), q_2(\tau, t), \dots, q_B(\tau, t)] \quad (5)$$

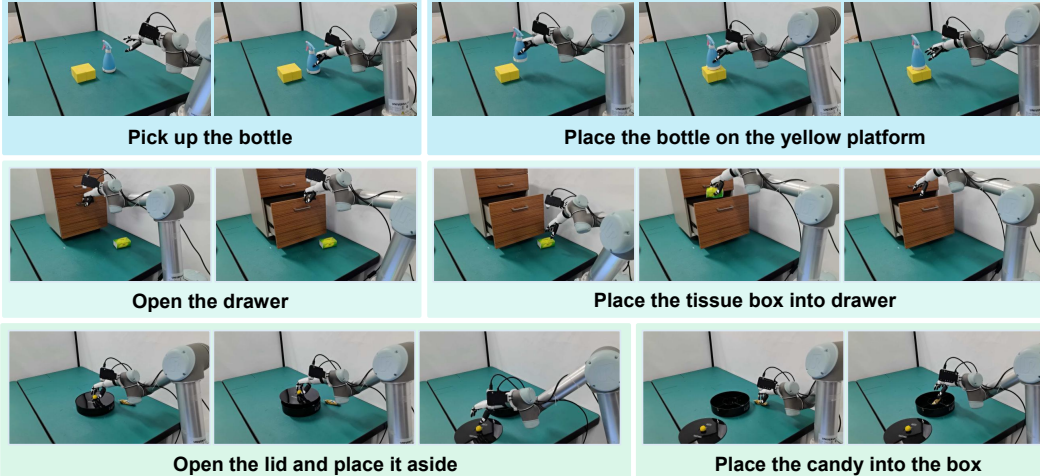


Figure 4: **Illustrations of the dexterous manipulation tasks.** **Task A (top)** requires the robot to grasp a tapered bottle and finely adjust its pose for stable placement. **Task B (middle)** requires the robot to insert a finger into the narrow drawer-handle gap, pull the drawer open, and place the tissue box inside. **Task C (bottom)** requires the robot to manipulate the spherical handle to open the lid, place the lid aside, and then grasp and place the candy into the box.

serves as the distributional supervision target for value learning. The distributional critic $p_\phi(V | o_t)$ is then trained to match this Gaussian soft target by minimizing the cross-entropy loss:

$$\mathcal{L}_{critic}(\phi) = \mathbb{E}_{(\tau, t) \sim \mathcal{D}} [H(\mathbf{q}(\tau, t), p_\phi(V | o_t))], \quad (6)$$

where $H(\cdot, \cdot)$ denotes the cross-entropy between the Gaussian soft label and the predicted value distribution.

Optimality-Conditioned Policy Improvement. Based on the formulation and conclusions in Sec 3, policy improvement is guaranteed as long as the constructed optimality function f is non-negative and monotonically increasing. Unlike RECAP, which relies on binary optimality labels $f_{bin}(A^{\pi_{ref}}(o_t, a_t)) = \mathbf{1}[A^{\pi_{ref}}(o_t, a_t) > q_{low}]$ [9], we construct f as a continuous optimality function,

$$f(A^{\pi_{ref}}(o_t, a_t)) = sig\left(\frac{\alpha(A^{\pi_{ref}}(o_t, a_t) - q_{low})}{q_{high} - q_{low}}\right), \quad (7)$$

where $sig(\cdot)$ denotes the sigmoid function, α is a temperature coefficient controlling the sharpness of the mapping, and q_{low} and q_{high} denote dataset-level quantiles of the advantage estimates, which mitigates the effect of the absolute advantage scale. Unlike exponential mappings that concentrate the learning signal on a few high-advantage samples, the sigmoid function provides a more robust mapping due to its smooth and bounded nature. Continuous optimality conditioning preserves the relative ordering of action quality instead of collapsing diverse behaviors into binary labels, allowing the policy to exploit fine-grained behavioral differences. We condition the diffusion policy on the continuous optimality value and train the actor with the following objective:

$$\mathcal{L}_{actor} = \mathbb{E}_{\mathcal{D}, \eta} \left[\|\epsilon - \epsilon_\theta(\tilde{\mathbf{a}}_{t:t+h}, o_t, I_t, \eta)\|_2^2 \right], \quad (8)$$

where $I_t = f(A^{\pi_{ref}}(o_t, a_t))$.

where $(o_t, \mathbf{a}_{t:t+h}) \sim \mathcal{D}$, $\eta \sim \mathcal{U}\{1, \dots, T\}$ denotes a sampled diffusion timestep, and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is Gaussian noise. The noised action $\tilde{\mathbf{a}}_{t:t+h}$ is obtained by applying the DDPM [48] forward noising process at timestep η . During training, we randomly mask the optimality indicator I_t with probability p_m , replacing it with a null condition. This design supports both direct policy sampling conditioned on $I_t = 1$ and CFG through the optimality indicator during inference.




Figure 5: **Main results.** Comparison with the BC reference policy [6] and post-training baselines on three real-world dexterous manipulation tasks. All post-training methods improve over BC, while our method consistently outperforms HG-Dagger [27] and RECAP [9].

5 Experiments

Setup. To verify our proposed human-in-the-loop data collection pipeline and policy improvement method, we evaluate them on three real-world robotic manipulation tasks: **Task A** (Pick-and-Place Bottle), **Task B** (Open Drawer and Tissue Placement), and **Task C** (Lid Opening and Candy Placement), as shown in Fig. 4. These tasks require both long-horizon operational capability and dexterous manipulation skills. For our method, we evaluate the policy after one iteration of post-training data collection and policy update. The success rate is computed over 50 evaluation trials. The robot setup, data collection details, and DexPIE implementation details are provided in Appendices A.1, A.2, and A.3, respectively.

Comparisons and Baselines. We compare against two baselines, RECAP [9] and HG-Dagger [27]. We use the same BC policy as the reference policy to warm-start our method and all baselines. This policy is trained only on demonstration data using the diffusion policy objective [6]. RECAP uses binary optimality labels as embedding conditions for the diffusion policy, whereas HG-Dagger trains the policy using only successful rollouts from the post-training data. For a fair comparison, we keep the amount of post-training data approximately consistent across methods. The implementation details are provided in the Appendix A.4.

Quantitative Results. As shown in Fig. 5, all three post-training methods achieve improvements over the reference policy. Overall, our method achieves the largest improvement of 37%. The performance of the HG-Dagger baseline confirms the effectiveness of our proposed human-following intervention system, which provides efficient corrective guidance and thereby alleviates covariate shift. However, since HG-Dagger does not leverage value estimation for long-horizon credit assignment, it may imitate low-quality segments collected during post-training, leading to inferior performance compared with our method. Compared with the discrete binary labels used in RECAP, our continuous optimality function yields consistent improvements. This demonstrates that continuous advantage conditioning can represent the degree of action optimality in a more fine-grained manner than discrete binary labels. Meanwhile, we observe that our method exhibits stronger robustness to positional variations. This improvement mainly stems from learning both failure cases induced by shifted grasping targets and the corresponding corrective interventions. Additional qualitative results and discussions are provided in Appendix B.1.

Ablation on Temporal Consistency. We conduct an ablation study on Task B to examine the effect of the temporal consistency between post-training rollouts and demonstration data. Given the BC reference policy trained with training-time RTC, deployment can be performed using either synchronous or asynchronous inference. We compare one iteration of our method using rollouts collected with synchronous inference against those collected with asynchronous inference from the same reference policy. For a fair comparison, both resulting policies are evaluated under the same asynchronous inference setting. As shown in Fig. 6, collecting post-training data with asynchronous inference improves performance by 14% compared with synchronous inference. This result suggests that, by reducing the demonstration-




Figure 6: **Ablation on temporal consistency.**

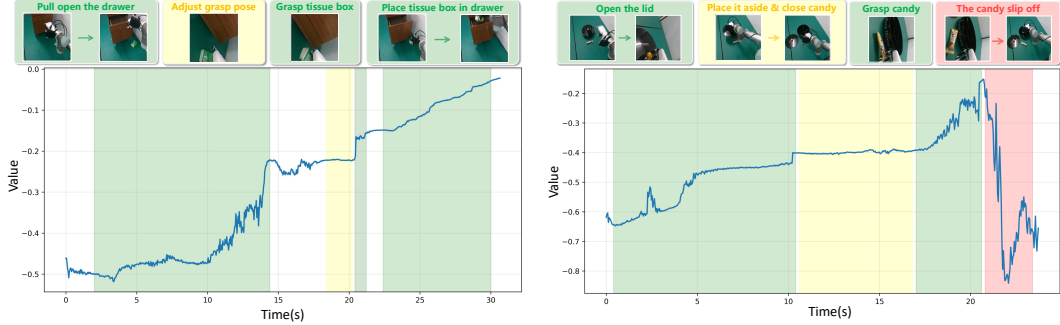


Figure 7: **Visualization of the value functions.** We visualize the value function output on two episodes. The green regions highlight successful progress toward task completion, the yellow regions indicate transitional or adjustment stages, and the red regions highlight failure-related segments.

deployment gap, asynchronous inference makes post-training rollouts with human corrections better align with the demonstrated behavior, allowing the critic to learn a value function induced by a more consistent underlying policy rather than by a heterogeneous mixture of mismatched behaviors. Additional qualitative examples are provided in Appendix B.2.

Ablation on Staged DAgger. To evaluate whether staged DAgger facilitates value-function learning and thereby improves policy performance, we conduct an ablation study on Task C. Starting from the same reference policy, we collect post-training data using two different strategies: staged DAgger and standard DAgger that initializes rollouts only from the initial state. As shown in Fig. 8, staged DAgger yields an 8% improvement in success rate. This demonstrates that by introducing later-stage trajectories, staged DAgger provides intermediate anchors for learning progress-aware values, effectively decomposing long-horizon tasks into shorter task segments where the critic can better capture progress regression and failure modes, thereby enabling more effective long-horizon credit assignment. As shown in Fig. 7, we visualize the value curves of two episodes. Benefiting from exploration coverage across different task stages, the learned value function accurately captures task progress and identifies failure modes, providing reliable credit assignment. Additional visualization results and a discussion of a special credit-assignment failure case are provided in Appendix B.3.




Figure 8: **Ablation on Staged DAgger.**

6 Conclusion

We presented DexPIE, a post-training framework for dexterous policy improvement from experience collected through real-world deployment. Experiments on three real-world long-horizon dexterous manipulation tasks show that DexPIE achieves consistent improvements over the reference policy and competitive post-training baselines. These results suggest that DexPIE provides an effective recipe for post-training in real-world dexterous manipulation. Meanwhile, its effective policy evaluation provides quality annotations for autonomous experience. As validated in DexPIE, even failure data can be effectively leveraged for policy training, thereby unlocking the potential of incorporating these real-world post-training data into large-scale pre-training.

7 Limitations

Due to hardware and computational constraints, our evaluation is limited to single-arm manipulation tasks. Future work should explore policy improvement in bimanual manipulation settings, as well as its effectiveness on minute-level long-horizon tasks. For higher-DoF dexterous hands, more human-hand-aligned retargeting algorithms are required to ensure consistency during takeover. Incorporating tactile sensing could further refine coarse hand gestures from either the policy or teleoperation, thereby enabling smoother transitions. Finally, the policy exploration method relies on human in-

tervention and manually selected stages. More diverse exploration strategies remain an important direction for future work.

References

- [1] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [2] Physical Intelligence, K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, et al. $\pi_{0.5}$: A Vision-Language-Action Model with Open-World Generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- [3] H. Luo, Y. Wang, W. Zhang, S. Zheng, Z. Xi, C. Xu, H. Xu, H. Yuan, C. Zhang, Y. Wang, et al. Being-h0. 5: Scaling human-centric robot learning for cross-embodiment generalization. *arXiv preprint arXiv:2601.12993*, 2026.
- [4] J. Bjorck, F. Castañeda, N. Cherniadev, X. Da, R. Ding, L. Fan, Y. Fang, D. Fox, F. Hu, S. Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- [5] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [6] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10-11):1684–1704, 2025.
- [7] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.
- [8] R. S. Sutton, A. G. Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [9] A. Amin, R. Aniceto, A. Balakrishna, K. Black, K. Conley, G. Connors, J. Darpinian, K. Dhabalia, J. DiCarlo, et al. $\pi_{0.6}^*$: A vla that learns from experience. *arXiv preprint arXiv:2511.14759*, 2025.
- [10] Y. Chen, S. Tian, S. Liu, Y. Zhou, H. Li, and D. Zhao. Conrft: A reinforced fine-tuning method for vla models via consistency policy. *arXiv preprint arXiv:2502.05450*, 2025.
- [11] Y. Li, X. Ma, J. Xu, Y. Cui, Z. Cui, Z. Han, L. Huang, T. Kong, Y. Liu, H. Niu, et al. Gr-rl: Going dexterous and precise for long-horizon robotic manipulation. *arXiv preprint arXiv:2512.01801*, 2025.
- [12] R. Yang, H. Wang, C. Liu, X. Yan, Y. Wang, X. Du, S. Yue, Y. Liu, C. Zhang, L. Qi, et al. Aloe: Action-level off-policy evaluation for vision-language-action model post-training. *arXiv preprint arXiv:2602.12691*, 2026.
- [13] J. Luo, C. Xu, J. Wu, and S. Levine. Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *Science Robotics*, 10(105):eads5033, 2025.
- [14] C. Xu, Q. Li, J. Luo, and S. Levine. Rldg: Robotic generalist policy distillation via reinforcement learning. *arXiv preprint arXiv:2412.09858*, 2024.
- [15] Q. Chen, J. Yu, M. Schwager, P. Abbeel, Y. Shentu, and P. Wu. Sarm: Stage-aware reward modeling for long horizon robot manipulation. *arXiv preprint arXiv:2509.25358*, 2025.

- [16] J. Zhang, Y. Luo, A. Anwar, S. A. Sontakke, J. J. Lim, J. Thomason, E. Biyik, and J. Zhang. Rewind: Language-guided rewards teach robot policies without new demonstrations. *arXiv preprint arXiv:2505.10911*, 2025.
- [17] Y. Mao, Z. Yu, W. Mao, Y. Li, Q. Hu, Z. Lan, M. Zhu, and H. Chen. Arm: Advantage reward modeling for long-horizon manipulation. *arXiv preprint arXiv:2604.03037*, 2026.
- [18] C. Yu, C. Sima, G. Jiang, H. Zhang, H. Mai, H. Li, H. Wang, J. Chen, K. Wu, L. Chen, et al. χ_0 : Resource-aware robust manipulation via taming distributional inconsistencies. *arXiv preprint arXiv:2602.09021*, 2026.
- [19] K. Black, A. Z. Ren, M. Equi, and S. Levine. Training-time action conditioning for efficient real-time chunking. *arXiv preprint arXiv:2512.05964*, 2025.
- [20] J. Tang, Y. Sun, Y. Zhao, S. Yang, Y. Lin, Z. Zhang, J. Hou, Y. Lu, Z. Liu, and S. Han. Vlash: Real-time vlas via future-state-aware asynchronous inference. *arXiv preprint arXiv:2512.01031*, 2025.
- [21] K. Black, M. Galliker, and S. Levine. Real-time execution of action chunking flow policies. *Advances in Neural Information Processing Systems*, 38:33383–33407, 2026.
- [22] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel. Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12156–12163. IEEE, 2024.
- [23] S. Yang, M. Liu, Y. Qin, R. Ding, J. Li, X. Cheng, R. Yang, S. Yi, and X. Wang. Ace: A cross-platform visual-exoskeletons system for low-cost dexterous teleoperation. *arXiv preprint arXiv:2408.11805*, 2024.
- [24] C. Wang, H. Shi, W. Wang, R. Zhang, L. Fei-Fei, and C. K. Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *arXiv preprint arXiv:2403.07788*, 2024.
- [25] Y. Qin, W. Yang, B. Huang, K. Van Wyk, H. Su, X. Wang, Y.-W. Chao, and D. Fox. Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system. *arXiv preprint arXiv:2307.04577*, 2023.
- [26] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [27] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8077–8083. IEEE, 2019.
- [28] Z. Hu, R. Wu, N. Enock, J. Li, R. Kadakia, Z. Erickson, and A. Kumar. Rac: Robot learning for long-horizon tasks by scaling recovery and correction. *arXiv preprint arXiv:2509.07953*, 2025.
- [29] P. Wu, Y. Shentu, Q. Liao, D. Jin, M. Guo, K. Sreenath, X. Lin, and P. Abbeel. Robocopilot: Human-in-the-loop interactive imitation learning for robot manipulation. *arXiv preprint arXiv:2503.07771*, 2025.
- [30] E.-R. Contributors. Evo-rl: Towards iterative policy improvement in real-world offline rl. <https://github.com/MINT-SJTU/Evo-RL>, 2026.
- [31] Y. Cui, Y. Zhang, L. Tao, Y. Li, X. Yi, and Z. Li. End-to-end dexterous arm-hand vla policies via shared autonomy: Vr teleoperation augmented by autonomous hand vla policy for efficient data collection. *arXiv preprint arXiv:2511.00139*, 2025.

- [32] Y. Han, Z. Chen, Y. Zhao, C. Xu, Y. Shao, Y. Peng, Y. Mu, and W. Lian. Dexhil: A human-in-the-loop framework for vision-language-action model post-training in dexterous manipulation. *arXiv preprint arXiv:2603.09121*, 2026.
- [33] H. Li, Y. Zuo, J. Yu, Y. Zhang, Z. Yang, K. Zhang, X. Zhu, Y. Zhang, T. Chen, G. Cui, et al. Simplevla-rl: Scaling vla training via reinforcement learning. *arXiv preprint arXiv:2509.09674*, 2025.
- [34] S. Zhai, Q. Zhang, T. Zhang, F. Huang, H. Zhang, M. Zhou, S. Zhang, L. Liu, S. Lin, and J. Pang. A vision-language-action-critic model for robotic real-world reinforcement learning. *arXiv preprint arXiv:2509.15937*, 2025.
- [35] T. Zhang, C. Yu, S. Su, and Y. Wang. Reinfo: Fine-tuning flow matching policy with online reinforcement learning. *Advances in Neural Information Processing Systems*, 38:106282–106319, 2026.
- [36] K. Chen, Z. Liu, T. Zhang, Z. Guo, S. Xu, H. Lin, H. Zang, Q. Zhang, Z. Yu, G. Fan, et al. π rl: Online rl fine-tuning for flow-based vision-language-action models. *arXiv preprint arXiv:2510.25889*, 2025.
- [37] Y. J. Ma, J. Hejna, C. Fu, D. Shah, J. Liang, Z. Xu, S. Kirmani, P. Xu, D. Driess, T. Xiao, et al. Vision language models are in-context value learners. In *International Conference on Learning Representations*, volume 2025, pages 33984–34009, 2025.
- [38] D. Huang, Z. Fang, T. Zhang, Y. Li, L. Zhao, and C. Xia. Co-rft: Efficient fine-tuning of vision-language-action models through chunked offline reinforcement learning. *arXiv preprint arXiv:2508.02219*, 2025.
- [39] F. Zhu, Z. Yan, Z. Hong, Q. Shou, X. Ma, and S. Guo. Wmpo: World model-based policy optimization for vision-language-action models. *arXiv preprint arXiv:2511.09515*, 2025.
- [40] J. Yang, K. Lin, J. Li, W. Zhang, T. Lin, L. Wu, Z. Su, H. Zhao, Y.-Q. Zhang, L. Chen, et al. Rise: Self-improving robot policy with compositional world model. *arXiv preprint arXiv:2602.11075*, 2026.
- [41] Z. Jiang, S. Zhou, Y. Jiang, Z. Huang, M. Wei, Y. Chen, T. Zhou, Z. Guo, H. Lin, Q. Zhang, et al. Wovr: World models as reliable simulators for post-training vla policies with rl. *arXiv preprint arXiv:2602.13977*, 2026.
- [42] Y. Guo, T. Lee, L. X. Shi, J. Chen, P. Liang, and C. Finn. Vlaw: Iterative co-improvement of vision-language-action policy and world model. *arXiv preprint arXiv:2602.12063*, 2026.
- [43] S. Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- [44] K. Frans, S. Park, P. Abbeel, and S. Levine. Diffusion guidance is a controllable policy improvement operator. *arXiv preprint arXiv:2505.23458*, 2025.
- [45] J. Ho and T. Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [46] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. *arXiv preprint arXiv:2402.10329*, 2024.
- [47] M. G. Bellemare, W. Dabney, and R. Munos. A distributional perspective on reinforcement learning. In *International conference on machine learning*, pages 449–458. Pmlr, 2017.
- [48] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

- [49] Y. Ze, Z. Chen, W. Wang, T. Chen, X. He, Y. Yuan, X. B. Peng, and J. Wu. Generalizable humanoid manipulation with 3d diffusion policies. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2873–2880. IEEE, 2025.
- [50] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022.

A Implementation Details

A.1 Robot Platform

This section describes the robotic platform and perception setup used in all experiments. As shown in Fig. 9, our experimental platform is built upon an arm-hand system consisting of a 6-DoF UR5 robot arm and a 6-DoF RH56DFX Inspire dexterous hand. The UR5 arm enables smooth Cartesian-space control, while the Inspire hand supports dexterous manipulation. The control frequency is set to 25 Hz. For hand retargeting, we retarget the operator’s hand gesture to the Inspire hand using selected manus ergonomics data, including the spread and flexion values of the thumb CMC joint and the MCP flexion values of the remaining fingers. The perception system adopts a multi-camera configuration, including one RealSense D415 camera that provides global front-view observations and an additional RealSense D435 camera mounted on the end-effector for close-range manipulation views. Both cameras run at 60 Hz and provide 640×480 RGB observations, which are resized to 224×224 using bilinear interpolation. All data, including camera observations, robot proprioceptive states, and glove action data, are synchronized and recorded at 25 Hz.




Figure 9: **Robot setup.**

A.2 Data Collection Process

For each task, we first collect multiple human demonstration trajectories and warm-start the policy through imitation learning. In particular, during each data collection episode, we randomize the environment and may initialize the rollout from later task stages, rather than always starting from the initial state. During policy deployment, we perform staged DAgger. Each rollout is executed until success or failure, and the corresponding trajectory is recorded. After a successful rollout, the environment is randomized before the next deployment. If the policy fails or encounters an exploration bottleneck, the failed trajectory is first recorded, and the environment is then reset to the same initial state of that rollout for redeployment. During redeployment, when the policy is about to enter a failure state or exhibits hesitation, the human operator intervenes and guides the robot to complete the task. This procedure improves exploration over task states and provides coverage of both failure cases and their corresponding corrective trajectories, thereby offering sufficient supervision for critic learning. We empirically observe that diverse failure trajectories are essential for the critic to recognize failure modes. Without sufficient coverage of such erroneous behaviors, the critic may overestimate the value of states or actions that are likely to fail, making it difficult to learn reliable distinctions between successful progress and failure-prone behaviors.

All trajectories, including both autonomous execution segments and human-intervention segments, are stored in the experience replay buffer. The experience replay buffer, together with the original demonstration data, forms the offline dataset. We first train the value function on this dataset to obtain advantage estimates and dataset-level advantage quantiles, which are then used for policy improvement. In the next iteration, the improved policy replaces the previous deployed policy for further data collection. For the initial demonstration data, we collect 41 trajectories for Task A, 30

trajectories for Task B, and 43 trajectories for Task C. All tasks are trained with one post-training iteration, where the numbers of collected post-training trajectories are 61 for Task A, 37 for Task B, and 59 for Task C. The dataset will be released publicly as example data. Each trajectory consists of standard camera observations, robot proprioceptive states and actions, as well as intervention annotations and binary success/failure labels.

A.3 DexPIE Implementation Detail

Actor: We adopt a diffusion policy with a U-Net architecture as the actor [49], using R3M encoder [50] as the visual encoder. The optimality value is first encoded with sinusoidal embedding, and then concatenated with the proprioceptive state and visual features as the conditioning input to the diffusion head, which predicts continuous action chunks. We set the action chunk size to 24 steps and the maximum delay for asynchronous inference to 4 steps. The policy is deployed on an RTX 4060 Ti GPU. **Critic:** we use a frozen R3M encoder to extract image features. The visual features are concatenated with the proprioceptive state and fed into a four-layer MLP head, which maps the input to B discretized value bins, where $B = 201$. We use sparse terminal rewards with per-step penalties. For an episode terminated at timestep T , the reward is defined as

$$r_t = \begin{cases} 0, & t = T \text{ and success,} \\ -C_{\text{fail}}, & t = T \text{ and failure,} \\ -1, & \text{otherwise,} \end{cases} \quad (9)$$

where C_{fail} is set to the maximum episode length of each task, and the discount factor is set to $\gamma = 1$. We normalize the Monte Carlo return by the maximum episode length and then clip it to the range $[-1, 0]$. The hyperparameters of the optimality function are set to $q_{\text{low}} = 0.6$, $q_{\text{high}} = 0.8$, and $\alpha = 5$. The masking probabilities of the optimality indicator and the action prefix are set to $p_m = 0.3$ and $p_m^{\text{pre}} = 0.9$, respectively. The number of denoising steps during policy inference is set to 10 and guidance strength $\beta = 1.5$. The advantage of action a_t at observation o_t is estimated using an N -step return: $A^\pi(o_t, a_t) = \sum_{l=t}^{t+N-1} \gamma^{l-t} r_l + \gamma^N V^\pi(o_{t+N}) - V^\pi(o_t)$, where $N = 24$. During training, we assign an optimality value of 1 to demonstration data and human-intervention segments. We use AdamW for optimization with $\beta_1 = 0.95$, $\beta_2 = 0.999$, and a learning rate of 1×10^{-4} . The batch sizes for the actor and critic are 256 and 512, respectively. Both models are trained on a single RTX 3090 GPU, with the actor trained for 300 epochs and the critic trained for 250 epochs.

A.4 Baselines Implementation Detail

For behavior cloning (BC), we use the same network architecture as our method, except that the optimality-conditioning injection is removed. We then minimize the diffusion-policy objective on the demonstration dataset:

$$\mathcal{L}_{\text{BC}} = \mathbb{E}_{\mathcal{D}_{\text{demo}}, \eta} \left[\|\epsilon - \epsilon_\theta(\tilde{\mathbf{a}}_{t:t+h}, o_t, \eta)\|_2^2 \right]. \quad (10)$$

For HG-DAGger, we retain only the successful trajectories from the post-training data and discard failed trajectories. The policy is trained with the same objective as BC. For RECAP, we use the same complete post-training dataset as in our method. The only difference is that RECAP uses binary optimality labels instead of continuous ones. In implementation, we replace the sinusoidal embedding with a binary embedding to inject the binary labels into the policy network. Its advantage threshold is set to $q_{\text{low}} = 0.6$. Apart from these differences, all training settings of the baselines are kept the same as those of our method.

B Qualitative Results

B.1 Robustness Qualitative Results

We mainly evaluate the robustness to positional variations. As shown in Fig. 10, our policy can robustly handle such positional variations across all three tasks. In practice, we find that a major




Figure 10: **Qualitative examples of robustness to positional variations across three tasks.** (Top) Robust adaptation to positional variations of the tapered bottle and the yellow platform. (Middle) Robust handling of diverse poses and positions of the tissue box. (Bottom) Generalization to different positional variations of the box and the candy.




Figure 11: **Example of the demonstration-deployment gap caused by temporal noise.** (Top) With asynchronous inference, the policy maintains temporally continuous execution and successfully completes the task. (Bottom) With synchronous inference, despite a nearly identical environment setup, latency-induced temporal noise causes repeated grasp failures and eventually results in a failed rollout.

failure mode in these tasks is misalignment with the grasping target. Therefore, during post-training data collection, we collect diverse trajectories where the policy deviates from the target and label them as failures, together with the corresponding corrective trajectories. In fact, with the additional optimality guidance, the policy learns to extract informative representations that distinguish undesirable trajectories from corrective ones, thereby moving away from failure patterns and guiding the generation of successful trajectories. thereby mitigating failure modes caused by positional variations.

B.2 Example of Demonstration-Deployment Gap

As shown in Fig. 11, we present a representative example to illustrate the impact of temporal inconsistency. Under nearly identical environment settings, the same policy succeeds when the action stream remains temporally continuous, but fails under synchronous inference due to latency-induced temporal noise. Although incorporating failure data is desirable in our post-training setting, we expect such failures to arise from insufficient coverage of the model’s intrinsic behavior distribution, rather than from distribution shifts caused by temporal noise. In fact, failure data can still be effectively leveraged for policy improvement as long as the optimality of the failure-inducing segments is correctly identified. However, as illustrated by this example, different inference settings can lead to substantially different outcome qualities under nearly identical states. This mismatch introduces a demonstration-deployment gap, forcing the critic to fit a value function induced by a heterogeneous mixture of inconsistent behaviors, making value estimates unreliable and weakening the credit assignment signals used to identify truly suboptimal trajectory segments. Moreover, when many low-

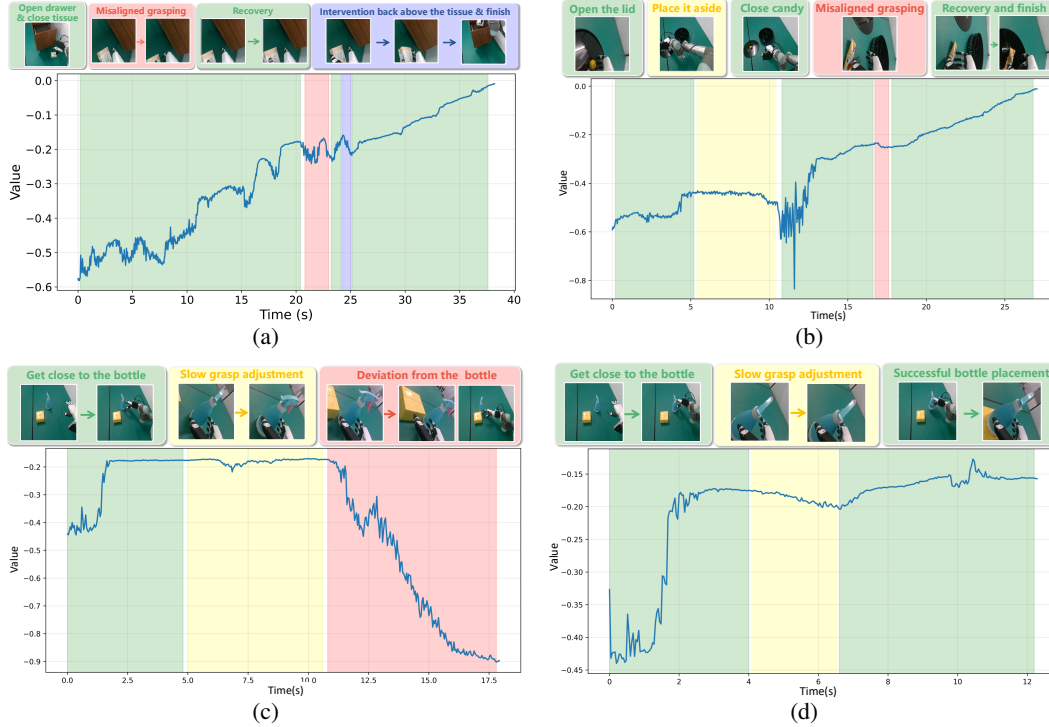


Figure 12: **Additional value visualization.** (a) Value visualization on a Task B trajectory with human intervention, where the blue marker indicates the intervention moment. The learned value function captures the temporary progress regression after intervention, the subsequent recovery toward the tissue box, and the final task completion. (b–d) Value visualization on fully autonomous rollout trajectories. Benefiting from sufficient exploration coverage, the learned value function can distinguish task progress, hesitation, and failure patterns.

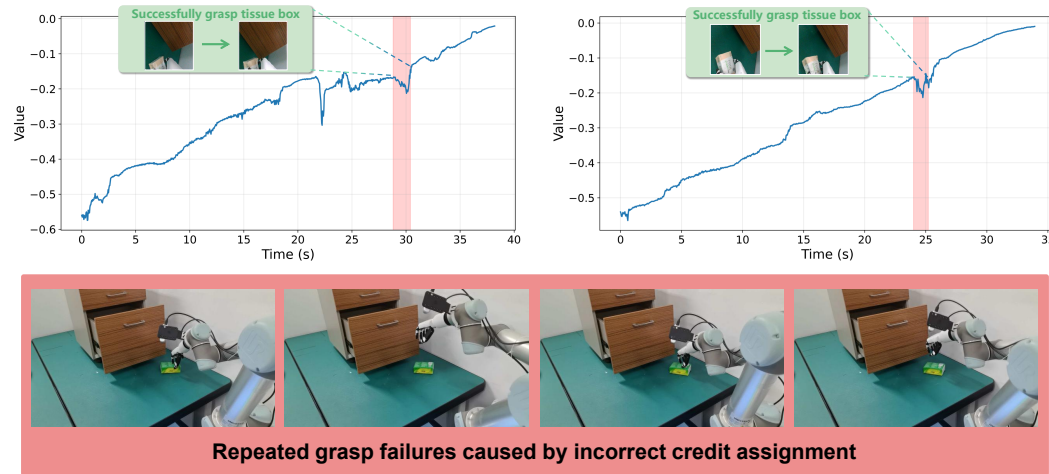


Figure 13: **Example of credit assignment failure.** (Top) The top row shows the learned value curves for two successful trajectories. Influenced by a special failure trajectory, the value function incorrectly predicts a value drop during an otherwise correct grasping process. (Bottom) The bottom row shows a deployment example where incorrect credit assignment causes the policy to avoid approaching the tissue box for grasping, leading to repeated grasp failures.

quality trajectories are introduced without being correctly distinguished, the policy may incorrectly imitate these temporally corrupted behaviors, leading to the learning of undesirable failure patterns.

B.3 Additional Value Visualization Results

As shown in Fig. 12, we further visualize the learned value curves of several trajectories. In particular, Fig. 12a includes a trajectory with human intervention, where the value function is able to capture both progress regression and subsequent recovery. Benefiting from multi-stage exploration coverage, the learned value function can generally capture task progress, hesitation, and failure patterns.

Special Credit-Assignment Failure Case. During our experiments, we have observed a special case of incorrect credit assignment. In one data-collection process, the robot repeatedly collides with the table, triggering collision detection and terminating the rollout. We label these trajectories as failure trajectories. As shown in Fig. 14, we visualize the image observation at the terminal state of one such failure trajectory. However, from the image observation alone, the value function cannot accurately attribute the failure to the grasp position being too low, which causes a collision with the table. Instead, as shown in the top row of Fig. 13, it assigns lower values to states where the robot approaches the tissue box for grasping, resulting in incorrect credit assignment. This misattribution further causes the policy to avoid approaching the tissue box during deployment, leading to repeated grasp failures, as shown in the bottom row of Fig. 13. To mitigate this issue, we filter out such failure trajectories to avoid misleading credit assignment and use human intervention to correct the collision behavior. This observation suggests that such incorrect credit assignment can be mitigated in two ways: either by incorporating richer information into the critic to better identify failure modes, or by filtering failure trajectories whose causes are difficult to infer from visual observations alone, making the remaining failure data more suitable for visual critic learning. In practice, it is not sufficient to simply introduce failure data; the critic must also be able to correctly recognize and attribute the underlying erroneous behaviors.



Figure 14: **Special failure trajectory labeled due to robot–table collision.**