

DexMimic: A Generalist Dexterous Hand–Object Tracking Policy via Value-Guided Reset Allocation

Anonymous Authors




Fig. 1: DexMimic learns a generalist dexterous hand–object tracking policy from human demonstrations through Value-Guided Reset Allocation. It consistently outperforms trajectory-specific baselines in tracking success and training efficiency, with zero-shot generalization improving as co-training scales. Real-world experiments validate DexMimic’s ability to perform diverse dexterous manipulation skills.

Abstract—Achieving human-like dexterous manipulation remains a central challenge in robotics. Learning from human hand–object interaction (HOI) datasets offers a promising pathway to diverse dexterous skills, yet existing physics-based approaches train specialist policies for individual trajectories, thereby limiting skill reuse at scale. We introduce **DexMimic**, a generalist dexterous hand–object tracking policy jointly trained across diverse HOI trajectories. Efficient joint training requires allocating each rollout to both an undertrained phase and a starting state that supports continued tracking. DexMimic realizes this insight through Value-Guided Reset Allocation, which adapts the initialization distribution online to prioritize difficult phases and select reset states with higher expected tracking returns. We further introduce a benchmark spanning training, fine-tuning, and generalization. Experiments show improved tracking success and training efficiency, faster fine-tuning, and the potential for zero-shot generalization to unseen trajectories and objects; larger training sets further improve efficiency and generalization. Finally, we validate DexMimic on a real-world dexterous manipulation platform. Supplementary materials and videos are available at <https://dexmimic.github.io/>.

I. INTRODUCTION

Recent advances in physics-based humanoid control show that a single reference-conditioned policy can learn diverse motor skills from large-scale motion datasets and reuse the

learned control prior beyond individual training motions [1], [2], [3], [4]. Extending this paradigm to dexterous hand–object interaction could enable robots to acquire diverse manipulation skills from human demonstrations [5], [6], [7]. These demonstrations capture fine-grained finger coordination, contact transitions [8], and regrasping behaviors that are difficult to specify through task-specific rewards. A generalist hand–object tracking policy could provide a reusable low-level control interface conditioned on human demonstrations. Realizing this vision requires bridging the substantial embodiment gap between human and robot hands in morphology, actuation, and contact dynamics [9], [10], [11].

Recent physics-based reinforcement learning methods have made substantial progress in bridging this gap for individual demonstrations, using general tracking objectives for hand motion, contact states, and object motion [10], [11]. However, these methods typically train a specialist policy for each trajectory or a small trajectory subset. This specialist formulation limits the reuse of learned control priors across trajectories and incurs repeated training costs as the demonstration set grows. To overcome these limitations, we

investigate what prevents existing trajectory-specific training pipelines from scaling to a generalist policy that jointly learns diverse hand–object interaction (HOI) trajectories and transfers the learned interaction prior to unseen references.

We find that Reference State Initialization (RSI) [1], a standard mechanism for initializing training episodes, becomes a key bottleneck when diverse trajectories share a limited rollout budget. RSI initializes each episode with a state sampled from a human motion retargeted to the robot’s morphology. This initialization scheme introduces inefficiencies in both the choice of starting states and the allocation of training across reference phases. First, kinematically retargeted states may contain physically inconsistent contacts, causing *fast-fail resets* before useful interaction experience is collected. Second, uniform timestamp sampling oversamples simple phases while undertraining difficult phases. As the trajectory set grows, these inefficiencies reduce the useful interaction experience available to each trajectory, limiting the scalability of joint policy learning.

We introduce **DexMimic**, a generalist dexterous hand–object tracking policy built on the insight that effective joint training requires adapting the initialization distribution alongside policy learning. Its **Value-Guided Reset Allocation** uses critic estimates to prioritize under-learned reference phases and select starting states that support continued tracking from phase-conditioned reset buffers. Each training iteration collects rollouts from the current initialization distribution and uses the same experience to update the policy and critic, refine reset candidates, and adjust phase priorities, thereby shaping the distribution for subsequent rollouts. On 50 challenging bimanual tasks, DexMimic achieves 88.40% success versus 56.64% for ManipTrans, while reducing average training time per trajectory by a factor of 15.5. Scaling co-training from 20 to 500 trajectories raises zero-shot success on unseen objects from 2.00% to 51.65%, demonstrating an increasingly transferable interaction prior. We further validate DexMimic on a real-world robotic platform.

In summary, our contributions are threefold: (1) We introduce **DexMimic**, a generalist dexterous hand–object tracking policy with **Value-Guided Reset Allocation**, which adapts reset state and phase allocation alongside policy learning. (2) We introduce a benchmark for generalist hand–object tracking that evaluates joint-training performance and efficiency, fine-tuning on unseen references, and zero-shot generalization to unseen trajectories and objects. (3) We demonstrate improved tracking success and training efficiency, faster adaptation, and stronger zero-shot generalization as training scales. We further validate DexMimic on a real-world robotic platform.

II. RELATED WORK

A. Reference tracking in physics simulation

Human hand–object demonstrations provide rich references for learning contact-rich manipulation in physics simulation [12], [13]. ManipTrans [10] uses reinforcement learning to transfer bimanual skills through residual control, while DexMachina learns functional retargeting through a

curriculum over virtual object controllers [11]. Further advances include adaptive trade-offs between object tracking and motion fidelity [14], reference-scoped exploration [15], and generalizable tracking [16]. Physics-based trajectory optimization also enables the conversion of human demonstrations into dynamically feasible robot motions [17].

In humanoid control, shared policies support diverse whole-body skills [18], [19], and SONIC [20] demonstrates the benefits of scaling data, model capacity, and computation for motion tracking. These advances motivate learning reusable control capabilities from diverse references. DexMimic takes a step toward this goal for dexterous hand–object interaction through efficient joint training across diverse HOI trajectories.

B. Reset distributions for reinforcement learning

The initialization distribution shapes the experience available to a policy and thereby influences exploration and learning efficiency [21], [22]. Curricula over starting states and demonstration-guided resets have shown the benefits of organizing where learning begins [23], [24], [25], [26]. In physics-based tracking, reference initialization, reuse of simulated states, and adaptive sampling further improve access to useful training experience [1], [27], [15], [3], [14]. DexMimic treats initialization as an evolving component of learning: a continually updated critic guides the distribution toward under-learned phases and starting states that support continued tracking, allowing reset allocation to adapt alongside policy learning.

III. METHOD

DexMimic learns a shared reference-conditioned tracking policy across diverse HOI trajectories in physics simulation. Our key insight is that achieving effective policy learning at scale hinges on optimizing the initialization distribution. To this end, Value-Guided Reset Allocation uses critic estimates to adapt the initialization distribution alongside policy learning. Reset state allocation selects starting states that support continued tracking, while reset phase allocation directs training toward under-learned reference phases. Fig. 2 illustrates the resulting training loop.

A. Problem Formulation

Given a set of reference trajectories $\mathcal{D} = \{\tau_i\}_{i=1}^N$ from HOI datasets, our goal is to learn a shared tracking policy. Each trajectory $\tau_i = \{\hat{x}_{i,t}\}_{t=0}^{T_i}$ specifies the target hand–object state and contact geometry at each reference frame. We formulate this task as a reference-conditioned Markov decision process (MDP), $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \rho_0, \gamma \rangle$. The task state $(i, t, s_t) \in \mathcal{S}$ includes the reference trajectory index i , reference frame t , and hand–object simulator state s_t . $\mathcal{A}$, $\mathcal{P}$, $\mathcal{R}$, ρ_0 , and γ denote the continuous action space, simulation dynamics, reward function, initialization distribution, and discount factor, respectively.

At the start of each episode, we sample a reference trajectory i , a starting frame t_0 , and an initial simulator state




Fig. 2: Overview of Value-Guided Reset Allocation in DexMimic. The reset map stores candidate states and their critic values for each trajectory and reference phase. As training progresses, reset phase allocation prioritizes under-learned phases, while reset state allocation favors candidates with higher expected tracking returns within each selected phase. Physics rollouts from the sampled initializations update the policy and critic, replenish the reset map, and refine phase sampling probabilities, forming an adaptive training loop.

s_{t_0} jointly from ρ_0 . We optimize the shared policy $\pi_\theta(a_t | o_t)$ to maximize the expected discounted tracking return:

$$J(\theta; \rho_0) = \mathbb{E}_{(i, t_0, s_{t_0}) \sim \rho_0} \left[\sum_{h=0}^{H-1} \gamma^h r_{t_0+h} \right]. \quad (1)$$

Here, o_t is the reference-conditioned observation, r_t is the tracking reward, and H is the episode length until reference completion or early termination.

B. Reference-Conditioned Tracking Policy

Observation. The policy observation o_t includes recent hand-object states, object geometry and mass properties, and future reference targets. The state history includes wrist position and orientation, finger joint angles, object position and orientation, their corresponding velocities, fingertip contact forces, and object translation relative to the wrist. Future reference inputs describe target hand-object states, tracking errors, and desired contact patterns. The observation also includes the kinematic base action and the remaining reference steps ℓ_t .

Action. The policy outputs continuous residual actions for wrist translation, wrist rotation, and finger joint positions. These residuals are combined with a kinematic base action derived from retargeted hand references to produce wrist forces and torques and finger position targets. Object motion follows contact dynamics. Detailed observation features and action parameterization are provided in the supplementary material.

Reward. We use a task-agnostic tracking reward shared across all references,

$$r_t = \lambda_h r_t^{\text{hand}} + \lambda_o r_t^{\text{object}} + \lambda_v r_t^{\text{velocity}} + \lambda_c r_t^{\text{contact}} + \lambda_e r_t^{\text{energy}}, \quad (2)$$

covering hand and object pose tracking, velocity matching, reference-guided contact, and mechanical-power regularization. Episodes terminate upon reference completion or excessive tracking error.

Network and training. The actor π_θ and critic V_ϕ share a multilayer perceptron (MLP) backbone and are jointly trained using standard Proximal Policy Optimization (PPO) [28]. Beyond policy learning, the critic provides value estimates for adapting the initialization distribution.

C. Value-Guided Reset Allocation

At scale, fast-fail resets and repeated sampling of well-learned phases consume a limited rollout budget, leaving insufficient experience for difficult interactions. We address this by adapting two initialization decisions: which reference phase to practice and which physical state to start from. At training iteration m , we factorize the initialization distribution as

$$\rho_{0,m}(i, t, s) = p_{\text{phase},m}(i, t) p_{\text{state},m}(s | i, t). \quad (3)$$

Here, $p_{\text{phase},m}$ selects a reference phase (i, t) , and $p_{\text{state},m}$ selects a starting state conditioned on that phase. By estimating the expected tracking return, the critic identifies promising starting states and under-learned phases. As learning progresses, rollout experience updates the policy and

Algorithm 1 DexMimic training pipeline

```
1: Input: HOI reference trajectories and training iterations
2: Initialize: Policy, critic, and reset buffers
3: for each training iteration do
4:   // Value-Guided Reset Allocation
5:   Sample  $(i, t) \sim p_{\text{phase}}$ 
6:   Sample  $s \sim p_{\text{state}}(\cdot \mid i, t)$  using Eq. (5)
7:   // Physics-based policy learning
8:   Collect physics rollouts from the sampled reset
9:   Update the policy and critic with PPO
10:  // Online reset adaptation
11:  Update reset buffers and phase priorities from rollouts
12:  Form the next initialization distribution
13: end for
14: Output: Trained generalist tracking policy
```

critic while refining reset candidates and phase priorities. The initialization distribution thus adapts alongside the policy and critic, continually directing training toward difficult interactions from states that support continued tracking. We omit the iteration subscript below for clarity.

1) *Reset state allocation:* Kinematically retargeted initializations can contain physically inconsistent contacts, causing fast-fail resets before useful interaction experience is collected. We therefore maintain a reset buffer for each reference phase (i, t) to reuse promising starting states discovered during physical rollouts:

$$\mathcal{B}_{i,t} = \left\{ (s_{i,t}^k, v_{i,t}^k) \right\}_{k=1}^K. \quad (4)$$

Collectively, these phase-conditioned buffers form a reset map over the training trajectories. Each buffer holds K candidate states and their critic scores; each state stores the hand–object configurations, velocities, and controller targets needed to restore the simulation. We initialize every slot with the phase-aligned reference state $s_{i,t}^{\text{ref}}$ and its critic score.

Given a reference phase (i, t) , we sample $s_{\text{reset}} \sim p_{\text{state}}(\cdot \mid i, t)$ using the nonnegative candidate weights $w_k = \max(v_{i,t}^k, 0)$:

$$s_{\text{reset}} = \begin{cases} s_{i,t}^{\text{ref}}, & \text{with probability } \alpha, \\ s_{i,t}^k, & \text{with probability } (1 - \alpha) \frac{w_k}{\sum_{j=1}^K w_j}. \end{cases} \quad (5)$$

Critic-weighted sampling favors candidates with higher expected tracking returns within the selected phase. With probability α , we instead sample the fixed reference anchor $s_{i,t}^{\text{ref}}$, helping limit drift away from the reference demonstrations. The subsequent tracking targets remain those of the original reference trajectory.

As training progresses, states collected during physical rollouts are evaluated as reset candidates for the corresponding reference phases. For a candidate s , we record its critic score $v = V_\phi(o)$ and evaluate its subsequent tracking duration. To retain candidates that support continued tracking, we require their survival fraction $c(s)$ —the fraction of remaining reference steps tracked from s before rollout

termination—to exceed $c_{\min}$. An eligible candidate replaces the lowest-scored slot in $\mathcal{B}_{i,t}$ only when its score exceeds the stored score. During buffer maintenance, stored scores decay as $v \leftarrow (1 - \delta)v$, allowing candidates generated by later policies to enter the buffer. Reference anchors remain unchanged. Together, candidate replacement and score decay adapt p_{state} throughout training, extending reference-based initialization with reusable rollout experience.

2) *Reset phase allocation:* To avoid repeatedly sampling well-learned phases, reset phase allocation uses critic estimates to prioritize under-learned interactions. We partition the reference trajectories into temporal bins, indexed globally by $b = 1, \dots, B$. To account for differences in remaining reference length, we normalize the critic value at each rollout step by the sum of the remaining discount weights:

$$\nu_t = \frac{V_\phi(o_t)}{\sum_{h=0}^{\ell_t-1} \gamma^h}, \quad (6)$$

where $\ell_t > 0$ is the number of remaining reference steps. Averaging normalized values over rollout steps assigned to bin b yields $\bar{\nu}_b$, a critic-based estimate of the policy’s tracking performance in that bin. We compute a difficulty score f_b by subtracting $\bar{\nu}_b$ from the theoretical maximum normalized value determined by the reward definition. Larger scores indicate phases that remain under-learned.

Poor tracking in a difficult phase may also result from insufficient preparation in preceding phases. We therefore propagate each bin’s difficulty score backward over a window of n bins within the same reference trajectory, with geometrically decaying weights:

$$\tilde{f}_b = \sum_{d=0}^{n-1} \lambda^d f_{b+d}, \quad n = 5, \quad 0 < \lambda < 1, \quad (7)$$

Here, $b + d$ denotes the current or a later bin within the same reference trajectory, and terms crossing trajectory boundaries are omitted. This raises reset priorities before difficult interactions, giving the policy opportunities to learn the preparation needed to enter them.

To reduce abrupt changes in reset priorities, we smooth the propagated difficulty scores across training iterations using an exponential moving average (EMA):

$$\bar{f}_b \leftarrow (1 - \beta)\bar{f}_b + \beta\tilde{f}_b, \quad 0 < \beta < 1, \quad (8)$$

where β weights the newly propagated score relative to the previous estimate. We normalize the smoothed scores across all bins to obtain

$$p_{\text{bin}}(b) = \frac{\bar{f}_b}{\sum_{j=1}^B \bar{f}_j}. \quad (9)$$

We sample a bin from p_{bin} , then uniformly sample a reference frame (i, t) within it. Thus, $p_{\text{phase}}(i, t) = p_{\text{bin}}(b(i, t))/n_{b(i, t)}$, where $b(i, t)$ is the bin containing (i, t) and n_b is its number of reference frames. We then sample the reset state from $p_{\text{state}}(\cdot \mid i, t)$ in Eq. (5), including its reference-anchor branch. As new rollout values are collected, phase priorities adapt to the policy’s evolving tracking performance.




Fig. 3: Qualitative simulation rollouts of DexMimic on single-hand (top) and bimanual (middle and bottom) manipulation tasks. Four selected frames are shown for each trajectory, progressing from left to right.

3) *Online training loop*: Algorithm 1 summarizes how policy learning and reset allocation are coupled through shared rollout experience. At each iteration, DexMimic samples reference phases and starting states from the current initialization distribution, then collects physics rollouts to update the policy and critic with PPO. The same rollouts refine reset candidates and phase priorities, directing subsequent training toward under-learned interactions from states that support continued tracking.

IV. EXPERIMENTS

A. Experimental Settings

Benchmark. We comprehensively evaluate DexMimic on 1,944 dexterous manipulation tasks involving 140 objects from the OakInk-V2 [29] and HOT3D [30] human hand-object interaction (HOI) datasets, comprising 1,302 single-hand trajectories and 642 bimanual trajectories involving rigid objects. We construct reference sets for training, fine-tuning, and zero-shot evaluation, with the zero-shot evaluation set organized into three levels of increasing difficulty. All main benchmark experiments use the Wuji dexterous hand, while results on additional hand embodiments are reported separately.

Training and evaluation protocol. We train policies using PPO in Isaac Gym [31] with 65,536 parallel environments across eight NVIDIA A100 GPUs. Value-Guided Reset Allocation is applied only during training to determine the reset phases and initialization states. During evaluation, we use Reference State Initialization (RSI) to initialize the hand and object at the first frame of the original reference and evaluate each task with 50 complete rollouts.

Metrics. We evaluate tracking performance using object rotation error, object translation error, finger keypoint error, fingertip error, and success rate (SR) [10]. We additionally report the average training time per trajectory. Let $\mathbf{T}_t = (\mathbf{R}_t, \mathbf{p}_t)$ denote the executed object pose and $\mathbf{x}_j^t$ the position of finger keypoint j ; barred quantities denote their references.

Methods	SR $\uparrow$ (%)	Time $\downarrow$ (h)	$E_r \downarrow$ (rad)	$E_t \downarrow$ (mm)	$E_j \downarrow$ (mm)	$E_{ft} \downarrow$ (mm)
ManipTrans	56.64	1.86	0.12	6.13	19.39	14.42
SPIDER	39.68	0.78	0.35	12.14	38.97	27.48
DexMimic	88.40	0.12	0.07	3.76	17.25	14.52

TABLE I: Tracking performance and average training or optimization time per trajectory on 50 challenging bimanual tasks. E_r , E_t , E_j , and E_{ft} denote object rotation, object translation, finger keypoint, and fingertip errors, respectively. Best values are highlighted in green.

With J finger keypoints and fingertip indices $\mathcal{F}$, the frame-wise errors are

$$e_{\text{rot}}^t = \arccos \frac{\text{tr}(\bar{\mathbf{R}}_t^\top \mathbf{R}_t) - 1}{2}, \quad e_{\text{trans}}^t = \|\mathbf{p}_t - \bar{\mathbf{p}}_t\|_2, \quad (10)$$

$$e_{\text{finger}}^t = \frac{1}{J} \sum_{j=1}^J \|\mathbf{x}_j^t - \bar{\mathbf{x}}_j^t\|_2, \quad e_{\text{tip}}^t = \frac{1}{|\mathcal{F}|} \sum_{j \in \mathcal{F}} \|\mathbf{x}_j^t - \bar{\mathbf{x}}_j^t\|_2.$$

Rotation error is reported in radians and position errors in millimeters. A rollout succeeds if it completes the reference and its time-averaged rotation, translation, finger keypoint, and fingertip errors are below 30° , 3 cm, 8 cm, and 6 cm, respectively. In bimanual tasks, both hands must satisfy the success criteria. SR is the fraction of successful rollouts. We report tracking errors averaged over time and then over successful rollouts, with hand errors averaged across hands in bimanual tasks.

B. Does DexMimic Improve Tracking Performance and Training Efficiency over Trajectory-Specific Methods?

We evaluate DexMimic on 50 challenging bimanual manipulation tasks from OakInk-V2 and HOT3D, comparing against ManipTrans [10] and SPIDER [17] as representative reinforcement learning and optimization baselines, respectively. ManipTrans learns a trajectory-specific residual policy on top of a pretrained hand-motion imitator to reproduce hand-object interactions. SPIDER directly optimizes a dynamically feasible control sequence for each reference trajectory through contact-guided sampling in parallel physics simulation. Both baselines train or optimize separately for each trajectory, while DexMimic jointly trains a single policy across all 50 trajectories. For a fair efficiency comparison, we report the average training or optimization time per trajectory, dividing DexMimic’s total joint-training time by 50. All methods use identical initial states, rollout counts, and success criteria during evaluation.

Table I summarizes the results on these 50 challenging bimanual tasks. DexMimic achieves an 88.40% success rate, outperforming ManipTrans and SPIDER by 31.76 and 48.72 percentage points, respectively. Its average training time is only 0.12 h per trajectory, compared with 1.86 h for ManipTrans and 0.78 h for SPIDER, corresponding to $15.5\times$ and $6.5\times$ lower time costs. DexMimic also achieves the lowest object rotation and translation errors, at 0.07 rad and 3.76 mm, respectively. These results demonstrate that a single shared tracking policy can efficiently learn diverse dexterous manipulation tasks while achieving higher execution success and more accurate object tracking than trajectory-specific methods.




Fig. 4: Real-world execution of four manipulation motions with DexMimic: pouring, scooping, tilting, and shaking. Frames progress from left to right within each sequence.



Fig. 5: Real-world platform with eight 3D-printed OakInk objects.

Fig. 3 further illustrates DexMimic’s tracking behavior on representative single-hand and bimanual manipulation tasks.

C. Does DexMimic Learn a Transferable Interaction Prior?

We introduce a benchmark to evaluate whether generalist co-training yields a reusable interaction prior that supports both zero-shot generalization and rapid adaptation, and how its transferability scales with training data. The benchmark combines three levels of zero-shot evaluation with fine-tuning on unseen references, covering transfer to new motions and objects as well as the efficiency of acquiring new tracking behaviors. For zero-shot evaluation, we train policies on nested subsets of 20, 50, 200, and 500 trajectories and freeze their parameters at test time. We evaluate three levels: L1, similar trajectories involving seen objects; L2, unseen trajectories involving seen objects; and L3, trajectories involving unseen objects. L1 trajectories are selected heuristically from the motion-capture datasets based on trajectory and hand-posture similarity to the training references. For adaptation, we jointly train a single policy on 20 unseen trajectories, either by fine-tuning the pretrained DexMimic base policy or by training from scratch.

Fig. 6(a) reveals a clear scaling trend: increasing the training set from 20 to 500 trajectories consistently improves zero-shot success rates across all three generalization levels. With 500 training trajectories, DexMimic achieves success rates of 68.50%, 55.89%, and 51.65% on L1, L2, and L3, respectively, up from 8.00%, 10.63%, and 2.00% with 20 trajectories. These substantial gains show that scaling co-training strengthens the learned interaction prior, enabling increasingly effective transfer to unseen trajectories and objects.

Furthermore, the pretrained DexMimic base policy enables rapid adaptation to unseen trajectories. The comparison on an unseen trajectory in Fig. 6(b) shows a $5.3\times$ speedup in reaching 80% training SR through fine-tuning compared with training from scratch. This result highlights the value of pretraining for acquiring new tracking behaviors efficiently. We believe that further scaling the training data can yield more capable generalist policies with stronger generalization and adaptation abilities.

Method	SR $\uparrow$ (%)	Time $\downarrow$ (h)	$E_r \downarrow$ (rad)	$E_t \downarrow$ (mm)	$E_j \downarrow$ (mm)	$E_{ft} \downarrow$ (mm)	FastFail $\downarrow$ (%)	WorstTraj SR $\uparrow$
RSI	66.51	0.150	0.28	5.17	22.64	20.26	96.34	21/50
Uniform	82.59	0.310	0.18	3.81	20.89	18.39	35.65	0/50
DexMimic	84.73	0.096	0.17	3.55	20.50	17.38	41.70	38/50

TABLE II: Ablations validate both reset allocation designs. FastFail and WorstTraj SR show that DexMimic addresses two key bottlenecks in generalist training: fast-fail resets and learning imbalance across trajectories.


Method	SR (% , $\uparrow$) / Time (min, $\downarrow$)			
	20	50	200	500
RSI	74.20 / 30.46	72.40 / 15.68	78.46 / 14.32	66.51 / 9.00
Uniform	76.90 / 24.00	67.64 / 17.43	77.79 / 25.89	82.59 / 18.60
DexMimic	85.80 / 20.39	75.44 / 10.25	88.86 / 8.63	84.73 / 5.76

TABLE III: Reset allocation variants trained on 20–500 trajectories. Each entry reports SR (% , $\uparrow$) / average training time per trajectory (min, $\downarrow$), using the same timing protocol as the main comparison.


D. How Does Value-Guided Reset Allocation Improve Generalist Training?

We separately ablate reset state allocation and reset phase allocation to examine how initialization quality and phase sampling affect generalist training. RSI initializes from kinematically retargeted reference states while retaining reset phase allocation; Uniform retains reset state allocation and samples reference phases uniformly. All three variants are jointly trained on the same 500 reference trajectories. Table II shows that DexMimic achieves the highest SR and the lowest object and hand tracking errors, while reducing training time from 18.60 min to 5.76 min relative to Uniform, a reduction of 69.0%.

Compared with RSI, enabling reset state allocation increases SR from 66.51% to 84.73% and reduces FastFail from 96.34% to 41.70%. FastFail measures the percentage of all training failures that occur within 60 frames after reset. Fig. 7 illustrates how Value-Guided Reset Allocation corrects physically infeasible initial reset states as training progresses. These results show that improving initialization feasibility mitigates fast-fail resets, allowing the policy to



(a)



(b)

Fig. 6: Benchmark for generalization and adaptation. (a) Scaling co-training improves zero-shot SR across three generalization levels. (b) Fine-tuning the pretrained DexMimic base policy accelerates joint adaptation to 20 unseen trajectories.




Fig. 7: Correction of physically infeasible initial reset states during training. Pink and blue show initial reset states and the corresponding states updated by DexMimic as training progresses, respectively. Orange arrows highlight the fingertip corrections.

collect sustained interaction experience.

Compared with Uniform, enabling reset phase allocation increases the minimum per-trajectory success rate (WorstTraj SR) from 0.00% to 76.00%, while improving average SR from 82.59% to 84.73%. The gap in worst-trajectory performance highlights learning imbalance across trajectories that is obscured by average SR. Fig. 8 illustrates phase-wise learning outcomes for four representative trajectories. These results show that reset phase allocation improves the performance of the least successful trajectories, addressing learning imbalance under uniform phase sampling.

We further evaluate both reset allocation mechanisms across co-training scales of 20, 50, 200, and 500 trajectories. Table III shows that DexMimic achieves higher success rates and lower average training time per trajectory than both ablated variants at every evaluated scale. Reset state allocation yields its largest success-rate improvement at 500 trajectories, outperforming RSI by 18.22 percentage points. The efficiency benefit of reset phase allocation grows with training scale: compared with Uniform, the reduction in average training time per trajectory increases from 15.0% at 20 trajectories to 69.0% at 500 trajectories. Together, these results show that feasible initialization supports successful




Fig. 8: Comparison of reset phase allocation and uniform phase sampling on four representative trajectories. Color indicates normalized training iterations to success; light red regions remain unsuccessful at the end of the same training budget.

tracking, while adaptive phase sampling improves training efficiency, enabling generalist tracking to scale to larger trajectory sets.

E. Real-World Deployment

As illustrated in Fig. 5, our real-world platform consists of a 7-DoF xArm equipped with a 20-DoF Wuji right hand and a RealSense D435i RGB-D camera. We execute hand-object interaction trajectories on the physical robot through open-loop replay of successful simulation rollouts generated by DexMimic. Before execution, we estimate the initial object pose using FoundationPose++ [32], [33] and determine a target wrist pose that matches the hand-object relative pose in the first frame of the selected rollout. We use PyRoki [34] to solve inverse kinematics to align the arm flange with the initial wrist pose and follow the wrist trajectory during replay, while the hand executes the corresponding finger joint motions. We do not enforce strict temporal alignment with the original human demonstration during replay, as the real robot cannot always match the speed of human hand motions.

F. Performance across Dexterous Hands

We further evaluate DexMimic on four dexterous hands, Wuji, Inspire, Shadow, and Allegro, as illustrated in Fig. 9.

Hand	Wuji	Inspire	Shadow	Allegro
SR $\uparrow$	872/1000	861/1000	832/1000	759/1000

TABLE IV: DexMimic across dexterous hands.




Fig. 9: Using mouse with Wuji, Inspire, Shadow, and Allegro.

We train one policy for each hand. Each policy is trained on the same set of 20 dexterous manipulation tasks. Evaluation follows the protocol in Sec. IV-A. Table IV reports the overall success rate (SR) for each hand over 1,000 evaluation rollouts.

V. CONCLUSION

We believe that learning dexterous manipulation from large-scale human hand-object interaction datasets offers a promising pathway toward human-level robot dexterity. DexMimic addresses the reset allocation bottleneck in scaling reference-conditioned tracking to a generalist policy through Value-Guided Reset Allocation. Experiments demonstrate improved tracking success and training efficiency, faster adaptation, zero-shot generalization potential, and transfer of learned behaviors to physical robots.

Although DexMimic shows promise in learning diverse dexterous skills from human demonstrations, it remains constrained by the scale of available HOI datasets. Future work will explore tracking methods that are more robust to noise in reconstructed hand-object interactions [35], enabling learning from more abundant in-the-wild videos.

REFERENCES

- [1] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, “DeepMimic: Example-Guided Deep Reinforcement Learning of Physics-Based Character Skills,” *arXiv preprint arXiv:1804.02717*, 2018.
- [2] Z. Luo, J. Cao, A. Winkler, K. Kitani, and W. Xu, “Perpetual Humanoid Control for Real-time Simulated Avatars,” *arXiv preprint arXiv:2305.06456*, 2023.
- [3] Q. Liao, *et al.*, “BeyondMimic: From Motion Tracking to Versatile Humanoid Control via Guided Diffusion,” *arXiv preprint arXiv:2508.08241*, 2025.
- [4] Z. Luo, *et al.*, “Universal Humanoid Motion Representations for Physics-Based Control,” *arXiv preprint arXiv:2310.04582*, 2023.
- [5] K. Shaw, S. Bahl, and D. Pathak, “VideoDex: Learning dexterity from internet videos,” *arXiv preprint arXiv:2212.04498*, 2022.
- [6] 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.
- [7] J. Mu, *et al.*, “DexImit: Learning bimanual dexterous manipulation from monocular human videos,” *arXiv preprint arXiv:2602.10105*, 2026.
- [8] Z. Fan, *et al.*, “ARCTIC: A dataset for dexterous bimanual hand-object manipulation,” *arXiv preprint arXiv:2204.13662*, 2022.
- [9] A. Handa, *et al.*, “DexPilot: Vision Based Teleoperation of Dexterous Robotic Hand-Arm System,” *arXiv preprint arXiv:1910.03135*, 2019.
- [10] K. Li, P. Li, T. Liu, Y. Li, and S. Huang, “ManipTrans: Efficient dexterous bimanual manipulation transfer via residual learning,” *arXiv preprint arXiv:2503.21860*, 2025.
- [11] Z. Mandi, Y. Hou, D. Fox, Y. Narang, A. Mandlekar, and S. Song, “DexMachina: Functional Retargeting for Bimanual Dexterous Manipulation,” *arXiv preprint arXiv:2505.24853*, 2025.
- [12] Y. Qin, *et al.*, “DexMV: Imitation learning for dexterous manipulation from human videos,” *arXiv preprint arXiv:2108.05877*, 2021.
- [13] P. Mandikal and K. Grauman, “DexVIP: Learning dexterous grasping with human hand pose priors from video,” *arXiv preprint arXiv:2202.00164*, 2022.
- [14] Y. Liang, Q. Peng, R.-Z. Qiu, and X. Wang, “ConTrack: Constrained Hand Motion Tracking with Adaptive Trade-off Control,” *arXiv preprint arXiv:2606.03177*, 2026.
- [15] S. Xu, *et al.*, “Dexplore: Scalable Neural Control for Dexterous Manipulation from Reference-Scoped Exploration,” *arXiv preprint arXiv:2509.09671*, 2025.
- [16] X. Liu, J. Adalibieke, Q. Han, Y. Qin, and L. Yi, “DexTrack: Towards Generalizable Neural Tracking Control for Dexterous Manipulation from Human References,” *arXiv preprint arXiv:2502.09614*, 2025.
- [17] C. Pan, *et al.*, “SPIDER: Scalable Physics-Informed Dexterous Retargeting,” *arXiv preprint arXiv:2511.09484*, 2025.
- [18] Y. Ze, *et al.*, “TWIST: Teleoperated Whole-Body Imitation System,” *arXiv preprint arXiv:2505.02833*, 2025.
- [19] X. B. Peng, Y. Guo, L. Halper, S. Levine, and S. Fidler, “ASE: Large-Scale Reusable Adversarial Skill Embeddings for Physically Simulated Characters,” *arXiv preprint arXiv:2205.01906*, 2022.
- [20] Z. Luo, *et al.*, “SONIC: Supersizing Motion Tracking for Natural Humanoid Whole-Body Control,” *arXiv preprint arXiv:2511.07820*, 2025.
- [21] A. Tavakoli, V. Levnik, R. Islam, C. M. Smith, and P. Kormushev, “Exploring Restart Distributions,” *arXiv preprint arXiv:1811.11298*, 2018.
- [22] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune, “First return, then explore,” *arXiv preprint arXiv:2004.12919*, 2020.
- [23] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel, “Reverse Curriculum Generation for Reinforcement Learning,” *arXiv preprint arXiv:1707.05300*, 2017.
- [24] C. Resnick, R. Raileanu, S. Kapoor, A. Peysakhovich, K. Cho, and J. Bruna, “Backplay: “Man muss immer umkehren,”” *arXiv preprint arXiv:1807.06919*, 2018.
- [25] I. Uchendu, *et al.*, “Jump-Start Reinforcement Learning,” *arXiv preprint arXiv:2204.02372*, 2022.
- [26] I. Popov, *et al.*, “Data-efficient Deep Reinforcement Learning for Dexterous Manipulation,” *arXiv preprint arXiv:1704.03073*, 2017.
- [27] S. Xu, H. Y. Ling, Y.-X. Wang, and L.-Y. Gui, “InterMimic: Towards Universal Whole-Body Control for Physics-Based Human-Object Interactions,” *arXiv preprint arXiv:2502.20390*, 2025.
- [28] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [29] X. Zhan, *et al.*, “OAKINK2: A dataset of bimanual hands-object manipulation in complex task completion,” in *Proc. IEEE/CVF CVPR*, 2024, pp. 445–456.
- [30] P. Banerjee, *et al.*, “HOT3D: Hand and object tracking in 3D from egocentric multi-view videos,” in *Proc. IEEE/CVF CVPR*, 2025.
- [31] V. Makovychuk, *et al.*, “Isaac Gym: High performance GPU-based physics simulation for robot learning,” *arXiv preprint arXiv:2108.10470*, 2021.
- [32] W. Yan and J. Chu, “FoundationPose++,” GitHub repository, Mar. 2025. [Online]. Available: <https://github.com/teal024/FoundationPose-plus-plus>
- [33] B. Wen, W. Yang, J. Kautz, and S. Birchfield, “FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects,” *arXiv preprint arXiv:2312.08344*, 2023.
- [34] C. M. Kim, B. Yi, H. Choi, Y. Ma, K. Goldberg, and A. Kanazawa, “PyRoki: A modular toolkit for robot kinematic optimization,” in *Proc. IEEE/RSJ IROS*, 2025.
- [35] Y. Hasson, G. Varol, I. Laptev, and C. Schmid, “Towards unconstrained joint hand-object reconstruction from RGB videos,” *arXiv preprint arXiv:2108.07044*, 2021.

DexMimic: A Generalist Dexterous Hand–Object Tracking Policy via Value-Guided Reset Allocation

Supplementary Material

A. POLICY AND ENVIRONMENT DETAILS

Notation. Let $\tau_i = \{\hat{x}_{i,k}\}_{k=0}^{T_i}$, s_t , o_t , and a_t denote the reference trajectory, simulator state, observation, and residual action. Superscripts w and o index the wrist rigid body and manipulated object, respectively. For $b \in \{w, o\}$, let $(\mathbf{p}_t^b, \mathbf{Q}_t^b, \mathbf{v}_t^b, \boldsymbol{\omega}_t^b)$ denote its simulated position, orientation quaternion, linear velocity, and angular velocity; hatted quantities are reference values from $\hat{x}_{i,k}$. For $x \in \{p, v, \omega\}$, define target-minus-simulation errors as

$$\mathbf{e}_{i,t,k}^{b,x} = \hat{x}_{i,k}^b - \mathbf{x}_t^b. \quad (1)$$

The orientation difference is

$$\mathbf{e}_{i,t,k}^{b,Q} = \mathbf{Q}_t^b \otimes (\hat{\mathbf{Q}}_{i,k}^b)^{-1}. \quad (2)$$

Here $\otimes$ denotes quaternion multiplication; $\mathbf{q}_t$ and $\dot{\mathbf{q}}_t$ are finger joint positions and velocities. Index j pairs a reference hand point with its simulated rigid body.

A.1. Observation Space

The observation contains hand proprioception, hand–object interaction features, and reference targets. History and future offsets in control frames are

$$\mathcal{H} = \{0, -1, -2, -4, -8\}, \quad \mathcal{F} = \{0, 1, 2, 4, 8\}. \quad (3)$$

The proprioceptive vector at frame t is

$$\mathbf{h}_t = [\mathbf{q}_t, \cos \mathbf{q}_t, \sin \mathbf{q}_t, \mathbf{0}_3, \mathbf{Q}_t^w, \mathbf{v}_t^w, \boldsymbol{\omega}_t^w]. \quad (4)$$

Absolute wrist translation is zeroed in this block; relative position features appear below. Missing history frames repeat the reset state. Let $\mathbf{f}_t^j$ denote fingertip contact force, $\mathbf{p}_t^{\text{CoM}}$ the object center of mass, and w_o its weight. Current interaction features are

$$\mathbf{c}_t = [\dot{\mathbf{q}}_t, \mathbf{c}_t^{\text{obj}}, \mathbf{c}_t^{\text{tip}}]. \quad (5)$$

$$\mathbf{c}_t^{\text{obj}} = [\mathbf{p}_t^o - \mathbf{p}_t^w, \mathbf{Q}_t^o, \mathbf{v}_t^o, \boldsymbol{\omega}_t^o, \mathbf{p}_t^{\text{CoM}} - \mathbf{p}_t^w, w_o]. \quad (6)$$

$$\mathbf{c}_t^{\text{tip}} = \{\mathbf{f}_t^j, \|\mathbf{f}_t^j\|_2\}_{j=1}^5. \quad (7)$$

For $k = \min(t + \delta, T_i)$ and $\delta \in \mathcal{F}$, the wrist and object reference blocks are

$$\mathbf{r}_{i,t,k}^b = [\mathbf{r}_{i,t,k}^{b,\text{lin}}, \mathbf{r}_{i,t,k}^{b,\text{rot}}]. \quad (8)$$

$$\mathbf{r}_{i,t,k}^{b,\text{lin}} = [\mathbf{e}_{i,t,k}^{b,p}, \hat{\mathbf{v}}_{i,k}^b, \mathbf{e}_{i,t,k}^{b,v}]. \quad (9)$$

$$\mathbf{r}_{i,t,k}^{b,\text{rot}} = [\hat{\mathbf{Q}}_{i,k}^b, \mathbf{e}_{i,t,k}^{b,Q}, \hat{\boldsymbol{\omega}}_{i,k}^b, \mathbf{e}_{i,t,k}^{b,\omega}]. \quad (10)$$

Reference hand points contribute position errors, target velocities, and velocity errors:

$$\mathbf{r}_{i,t,k}^{\text{link}} = \{\hat{\mathbf{p}}_{i,k}^j - \mathbf{p}_t^j, \hat{\mathbf{v}}_{i,k}^j, \hat{\mathbf{v}}_{i,k}^j - \mathbf{v}_t^j\}_{j=1}^J. \quad (11)$$

Let $\hat{\mathbf{d}}_{i,k}^{\text{tip}}$ contain reference fingertip–object distances. With $\mathbf{h}_t^{\mathcal{H}} = [\mathbf{h}_{t+\delta}]_{\delta \in \mathcal{H}}$ and $\mathbf{r}_{i,t}^{\mathcal{F}} = [\mathbf{r}_{i,t,\min(t+\delta, T_i)}]_{\delta \in \mathcal{F}}$ and

$$\mathbf{r}_{i,t,k} = [\mathbf{r}_{i,t,k}^w, \mathbf{r}_{i,t,k}^o, \mathbf{r}_{i,t,k}^{\text{link}}, \hat{\mathbf{d}}_{i,k}^{\text{tip}}]. \quad (12)$$

the full observation is

$$\mathbf{o}_t = [\mathbf{h}_t^{\mathcal{H}}, \mathbf{c}_t, \mathbf{r}_{i,t}^{\mathcal{F}}, \mathbf{d}_t^o, \mathbf{b}_t^o, \mathbf{a}_t^{\text{base}}, \ell_t]. \quad (13)$$

Here $\mathbf{d}_t^o$ contains object-to-hand rigid-body distances, $\mathbf{b}_t^o$ is the 128-dimensional object geometry encoding, $\mathbf{a}_t^{\text{base}}$ is the kinematic base action, and ℓ_t gives the remaining reference steps. Bimanual observations concatenate hand-specific blocks and share one ℓ_t ; actor and critic receive the same observation.

A.2. Action Space

For n actuated finger joints, the action comprises normalized wrist translation, wrist rotation, and finger target residuals:

$$\mathbf{a}_t = [\mathbf{a}_t^p, \mathbf{a}_t^R, \mathbf{a}_t^q] \in [-1, 1]^{6+n}, \quad \boldsymbol{\delta}_t = 2\mathbf{a}_t. \quad (14)$$

The kinematic base action comes from the retargeted hand reference. Its wrist position error $\mathbf{e}_{i,t,t}^{w,p}$ is clipped to 0.10 m; the axis–angle error of $\hat{\mathbf{Q}}_{i,t}^w \otimes (\mathbf{Q}_t^w)^{-1}$ is clipped to 1 rad. The observation uses the reverse quaternion difference in Eq. (2). The finger base $\mathbf{b}_t^q$ is the normalized retargeted target.

In wrist PID control, residuals augment the base force and torque:

$$\tilde{\mathbf{F}}_t = \text{PID}_p(\bar{\mathbf{e}}_t^p) + 500\Delta t s_p \boldsymbol{\delta}_t^p. \quad (15)$$

$$\tilde{\mathbf{M}}_t = \text{PID}_R(\bar{\mathbf{e}}_t^R) + 200\Delta t s_R \boldsymbol{\delta}_t^R. \quad (16)$$

Here $\bar{\mathbf{e}}_t^p$ and $\bar{\mathbf{e}}_t^R$ denote the clipped errors, $\Delta t = 1/60$ s, $s_p = 1$, and $s_R = 0.1$; PID gains depend on the hand model. Finger residuals are added in normalized joint coordinates and mapped to joint limits:

$$\mathbf{u}_t^q = \text{clip}(\mathbf{b}_t^q + \boldsymbol{\delta}_t^q, -1, 1). \quad (17)$$

$$\mathbf{q}_t^{\text{raw}} = \mathbf{q}_{\min} + \frac{1}{2}(\mathbf{u}_t^q + 1) \odot (\mathbf{q}_{\max} - \mathbf{q}_{\min}). \quad (18)$$

Wrist force, torque, and finger targets use temporal smoothing with $\alpha = 0.4$:

$$\mathbf{y}_t = \alpha \tilde{\mathbf{y}}_t + (1 - \alpha) \mathbf{y}_{t-1}, \quad \mathbf{y} \in \{\mathbf{F}, \mathbf{M}, \mathbf{q}^{\text{cmd}}\}. \quad (19)$$

Term (unit)	Component reward	Weight
Hand pose tracking		
Wrist position (m)	$\exp(-40E_{w,p})$	0.10
Wrist orientation (rad)	$\exp(-E_{w,R})$	0.60
Thumb (m)	$\exp(-100E_T)$	0.90
Index (m)	$\exp(-90E_I)$	0.80
Middle (m)	$\exp(-80E_M)$	0.75
Ring (m)	$\exp(-60E_R)$	0.60
Little (m)	$\exp(-60E_P)$	0.60
Link group L_1 (m)	$\exp(-50E_{L_1})$	0.50
Link group L_2 (m)	$\exp(-40E_{L_2})$	0.30
Object pose tracking		
Object position (m)	$\exp(-80E_{o,p})$	5.00
Object orientation (rad)	$\exp(-6E_{o,R})$	1.00
Velocity matching		
Wrist linear (m/s)	$\exp(-E_{w,v})$	0.10
Wrist angular (rad/s)	$\exp(-E_{w,\omega})$	0.05
Hand links (m/s)	$\exp(-E_{\text{link},v})$	0.10
Object linear (m/s)	$\exp(-E_{o,v})$	0.10
Object angular (rad/s)	$\exp(-E_{o,\omega})$	0.10
Reference-guided contact		
Fingertip contact (N)	$\exp[-(F_t^{\text{tip}}/N + 10^{-5})^{-1}]$	1.00
Mechanical-power regularization		
Joint power (W)	$\exp(-10P_t^q)$	0.50
Wrist power (W)	$\exp(-2P_t^w)$	0.50

TABLE S1: Single-hand reward components. Units refer to the underlying errors or physical quantities.

For finger control, $\tilde{\mathbf{q}}_t^{\text{cmd}} = \mathbf{q}_t^{\text{raw}}$; the smoothed target is clipped to joint limits. Wrist force $\mathbf{F}_t$ (N) and torque $\mathbf{M}_t$ (N·m) act on the wrist rigid body; finger targets are in radians. Object motion follows contact dynamics. Bimanual actions concatenate both hands' residuals.

A.3. Reward Function

The task-agnostic reward comprises hand pose, object pose, velocity, reference-guided contact, and mechanical-power terms. For the set $\mathcal{G}$ of these groups,

$$r_t = \sum_{g \in \mathcal{G}} \lambda_g r_t^g. \quad (20)$$

The hand group includes wrist and finger tracking. Table S1 specifies each component's weight; $\lambda_g = 1$ for all five groups.

At the current reference frame, pose errors are

$$E_{b,p} = \|\mathbf{e}_{i,t,t}^{b,p}\|_2, \quad E_{b,R} = 2 \arccos |\langle \mathbf{Q}_t^b, \hat{\mathbf{Q}}_{i,t}^b \rangle|. \quad (21)$$

For each hand-link group G ,

$$E_G = |\mathcal{I}_G|^{-1} \sum_{j \in \mathcal{I}_G} \|\hat{\mathbf{p}}_{i,t}^j - \mathbf{p}_t^j\|_2. \quad (22)$$

Velocity errors average absolute component differences:

$$E_{b,v} = \frac{1}{3} \|\mathbf{e}_{i,t,t}^{b,v}\|_1, \quad E_{b,\omega} = \frac{1}{3} \|\mathbf{e}_{i,t,t}^{b,\omega}\|_1. \quad (23)$$

For tracked hand links,

$$E_{\text{link},v} = (3J)^{-1} \sum_{j=1}^J \|\hat{\mathbf{v}}_{i,t}^j - \mathbf{v}_t^j\|_1. \quad (24)$$

Reference fingertip–object distance gates the simulated fingertip force. For fingertip j , set $u_t^j = \text{clip}((0.03 - \hat{d}_{i,t}^{\text{tip},j})/0.01, 0, 1)$, with distances in metres, and let

$$F_t^{\text{tip}} = \sum_{j=1}^5 u_t^j \|\mathbf{f}_t^j\|_2. \quad (25)$$

Mechanical-power terms use

$$P_t^q = \sum_{j=1}^n |\tau_{t,j}^q \dot{q}_{t,j}|. \quad (26)$$

$$P_t^w = |\mathbf{F}_t^\top \mathbf{v}_t^w| + |\mathbf{M}_t^\top \boldsymbol{\omega}_t^w|. \quad (27)$$

Component rewards and weights are dimensionless; exponential scales assume SI units. Bimanual rewards sum per-hand terms, including an object term for each hand.

B. TRAINING AND RESET ALLOCATION DETAILS

B.1. PPO and Network Configuration

Each observation block is normalized using running statistics before concatenation. A feature MLP with three 512-unit Swish layers and a 256-dimensional projection feeds a shared ELU trunk of widths 256, 512, 128, and 64. Separate heads output the mean of a diagonal Gaussian policy and the scalar critic $V_\phi(o_t)$; the Gaussian has trainable, observation-independent log standard deviations. Policy samples are bounded to $[-1, 1]$ before residual action synthesis. Clipped PPO jointly updates the shared features and both heads using generalized advantage estimation. Value targets are normalized during optimization, while reset allocation uses critic estimates on the original reward scale.

B.2. Reset Buffer Update

For each reference phase (i, t) , the reset buffer $\mathcal{B}_{i,t}$ holds K scored states, initialized from the phase-aligned reference state $s_{i,t}^{\text{ref}}$ and its critic score. The fixed reference anchor is sampled through a separate α branch. Each state stores hand and object root configurations and velocities, finger joint positions and velocities, and controller targets. Before action execution at relative step u of episode e , the physical state is paired with its critic estimate. At episode end, the pair is assigned to the corresponding reference frame:

$$t = t_e^{\text{start}} + u, \quad v_{e,u} = V_\phi(o_{e,u}). \quad (28)$$

For candidates assigned to the same phase, scores are processed in descending order. Here $c(s)$ is the fraction of remaining reference steps tracked from the candidate before episode termination. A candidate must satisfy $c(s) > c_{\min}$ and strictly exceed the current lowest buffer score to replace that entry; the lowest score is recomputed after each replacement. After a nonempty episode-finalization batch, candidate scores decay once as $v \leftarrow (1 - \delta)v$. The reference anchor remains fixed and is used if all candidate sampling weights vanish.




Fig. S1: A single DexMimic policy tracking different reference trajectories.

B.3. Reset Phase Allocation

Each reference trajectory is partitioned into temporal bins of 15 control frames; its final bin may be shorter. The bins are indexed globally across trajectories. For each rollout step with $\ell_t > 0$, the critic estimate on the original reward scale is normalized by the remaining discounted horizon:

$$\nu_t = \frac{V_\phi(o_t)}{\sum_{h=0}^{\ell_t-1} \gamma^h}. \quad (29)$$

For bins receiving samples, let $\mathcal{D}_b$ be the multiset of rollout steps assigned to bin b , including repeated visits. Their mean is $\bar{\nu}_b = |\mathcal{D}_b|^{-1} \sum_{t \in \mathcal{D}_b} \nu_t$. The difficulty $f_b = \nu_{\max} - \bar{\nu}_b$ uses the theoretical maximum normalized value implied by the reward. It propagates to preceding bins within the same trajectory:

$$\tilde{f}_b = \sum_{d=0}^{n-1} \lambda^d f_{b+d}, \quad (30)$$

with terms crossing the trajectory boundary omitted. The propagated score updates the phase priority by exponential moving average:

$$\bar{f}_b \leftarrow (1 - \beta) \bar{f}_b + \beta \tilde{f}_b. \quad (31)$$

Sampling normalizes $\bar{f}_b$ over all bins, then selects uniformly among the n_b valid frames of the chosen bin. Thus, for a frame (i, t) in bin $b(i, t)$,

$$p_{\text{phase}}(i, t) = \frac{\bar{f}_{b(i, t)}}{n_{b(i, t)} \sum_{j=1}^B \bar{f}_j}. \quad (32)$$

B.4. Hyperparameters

The principal training and reset parameters are summarized below.




Fig. S2: Per-trajectory success rates versus reference trajectory length. Colors distinguish right-hand (RH) and left-hand (LH) trajectories.

Parameter	Value
PPO rollout / mini-epochs / minibatch	32 / 5 / 1024
Discount γ / GAE λ_{GAE}	0.99 / 0.95
Adam learning rate / PPO clip	5×10^{-4} / 0.2
Reset slots K / anchor probability α	3 / 0.1
Survival threshold $c_{\min}$ / score decay δ	0.1 / 10^{-5}
Bin width / lookback n / factor λ	15 / 5 / 0.8




Fig. S3: Evaluation outcomes and estimated failure-category proportions. Failure labels denote tracking errors exceeding their respective thresholds. All percentages are relative to the 1,017 evaluation rollouts. Failure-category proportions are estimated by proportionally allocating unclassified failures across the five categories.




Fig. S4: Representative examples of the five failure categories. Translucent hands and objects indicate the reference.

C. ADDITIONAL RESULTS

C.1. Per-Trajectory Success Rates

We evaluate DexMimic on 1,302 trajectories from OakInk and HOT3D. Figure S2 plots success rate against reference trajectory length, with colors distinguishing right-hand (RH) and left-hand (LH) trajectories. Figure S1 illustrates a single policy tracking different reference trajectories.

C.2. Failure Case Analysis

Across 1,017 evaluation rollouts, DexMimic succeeds in 826 cases (81.219%). Rollouts terminate early when tracking errors in object position, object orientation, wrist position, wrist orientation, or finger joint configuration exceed their respective thresholds. The estimated breakdown in Fig. S3 is dominated by object tracking failures, indicating that maintaining the reference object position and orientation remains the primary challenge in these evaluations. Figure S4 shows representative cases for the five categories.