


SimToolReal: An Object-Centric Policy for Zero-Shot Dexterous Tool Manipulation

Kushal Kedia^{*1} Tyler Ga Wei Lum^{*2} Jeannette Bohg^{†2} C. Karen Liu^{†2}

¹Cornell University ²Stanford University ^{*}Equal contribution [†]Equal advising

simtoolreal.github.io

Single Policy: Zero-Shot on Novel Tools & Tasks



Grasp → In-Hand Rotation → Tool-Use



Fig. 1: *SimToolReal* is a framework for training a single general-purpose, object-centric RL policy in simulation and transferring it to real-world tool use. (Top) Zero-shot deployment on novel real tools and tasks, spanning thin markers to thick hammers. (Bottom) Tool use typically involves grasping objects from flat surfaces, reorienting them in-hand, and performing the task.

Abstract—The ability to manipulate tools significantly expands the set of tasks a robot can perform. Yet, tool manipulation represents a challenging class of dexterity, requiring grasping thin objects, in-hand object rotations, and stable object control during environmental contact. Since collecting teleoperation data for these behaviors is challenging, sim-to-real reinforcement learning (RL) is a promising alternative. However, prior approaches typically require substantial engineering effort to model objects and tune reward functions for each task. In this work, we propose *SimToolReal*, taking a step towards generalizing sim-to-real RL policies for tool manipulation. Instead of focusing on a single object and task, we procedurally generate a large variety of tool-like object primitives in simulation and train a single RL policy with the universal goal of manipulating each object to random goal poses. This approach enables *SimToolReal* to perform general dexterous tool manipulation at test-time without any object or task-specific training. We demonstrate that *SimToolReal* outperforms prior retargeting and fixed-grasp methods by 37% in average *Task Progress* while matching the performance of specialist RL policies trained on specific target objects and tasks. Finally, we show that *SimToolReal* generalizes across a diverse set of everyday tools, achieving strong zero-shot performance over 120 real-world rollouts spanning 24 tasks, 12 object instances, and 6 tool categories. Website: simtoolreal.github.io.

I. INTRODUCTION

From household chores to industrial work, tool use expands a robot’s capacity to act on its environment. However, tool use represents a challenging class of dexterous manipulation tasks. Success requires grasping tools lying flat on a surface, reorienting them into functional poses, and maintaining stable control while executing object motions that involve environmental contact (see Fig. 1). Consider hammering a nail: the robot must grasp the hammer by its thin handle, rotate it in-hand into a striking configuration, and deliver impact without losing the grasp. Such interactions expose a fundamental limitation of parallel-jaw grippers: with only two opposing contacts along a single grasp axis, they provide limited resistance against externally induced torques, making grasp stability primarily rely on friction and grip force. This motivates multi-fingered dexterous hands for stable tool use.

Imitation learning from teleoperated demonstrations is a common paradigm for learning manipulation policies [21, 92], but teleoperation is a poor fit for collecting high-quality dexterous tool-use data [13]. Human and robot hands differ in kinematics and actuation, creating a human-to-robot corre-




Fig. 2: **Overview of *SimToolReal*.** (Top) Training in Simulation: We train a goal-conditioned RL policy in simulation that manipulates a wide variety of procedurally-generated objects to randomly sampled goal poses. (Bottom) Inference in Real: We deploy this policy zero-shot on real-world tools from *DexToolBench*, following tool trajectories from human videos.

spondence gap [70, 7] that makes precise dexterous control unintuitive for the operator [83]. At the same time, operators receive limited to no force and tactile feedback, preventing reliable regulation of contacts during the task [54, 56, 52, 75].

Sim-to-real reinforcement learning (RL) is a promising alternative that has successfully demonstrated agile, dynamic, and dexterous behaviors in the real world [3, 16]. Although effective for specific skills, extending its success to diverse tools and tasks is bottlenecked by substantial per-object simulation setup and task-specific reward engineering [58]. Consequently, existing approaches are limited to sub-problems of dexterous tool-use, such as grasping [46, 1, 19], in-hand object reorientation [15, 55, 28], or spinning objects [80, 41].

Our goal is to learn a dexterous manipulation policy that can be applied zero-shot to novel tools and tasks at test-time. Our key insight is to view dexterous tool use from an *object-centric* lens, framing a task as moving a tool through a sequence of goal poses. Thus, we focus on training a single goal-conditioned policy that manipulates diverse objects to random goal poses. In simulation, we procedurally generate a large set of primitive objects and train one unified goal-conditioned policy. At test time, we compose this policy for zero-shot tool use by sequentially conditioning it on a goal pose trajectory extracted from a human video [22, 47]. This provides dense guidance from the initial grasp, through in-hand reorientation into a functional configuration, and into the ensuing tool-use motion. Our approach avoids per-task reward shaping while allowing generalization to novel tools and trajectories at test-time. Intuitively, mastering the ability of manipulating objects to any random pose induces the core skills required for tool use: establishing an initial grasp, in-hand object reorientation, and maintaining stable contact.

We instantiate this approach in *SimToolReal*, a framework for training a single general-purpose RL policy in simulation and deploying it for real-world tool use (Fig. 2). To enable generalization to novel objects, the policy is conditioned on

the tool’s current 6D pose and a coarse 3D bounding box over its graspable region (e.g., the handle of a hammer). This abstraction enables zero-shot transfer to the real-world by effectively bypassing the sim-to-real visual gap. At deployment, we recover this representation on real tools using a perception pipeline based on vision foundation models, combining SAM 3D [17] for segmentation and mesh extraction with FoundationPose [82] for 6D pose tracking.

To evaluate generalization, we introduce *DexToolBench*, a benchmark of daily tool-use behaviors, both in simulation and in the real-world. Each task is paired with a human video demonstration, and the robot is evaluated on its ability to follow the demonstrated tool trajectory. Notably, *SimToolReal* is trained exclusively on procedurally-generated primitives with randomly sampled goal pose trajectories. Yet, it transfers zero-shot to the trajectory following tasks in *DexToolBench* and matches or outperforms specialist policies trained on a single object instance and task. Our contributions:

- 1) We propose *SimToolReal*, a unified RL framework that frames dexterous tool use as manipulating a tool through a sequence of desired poses. Training a single policy under this objective induces dexterous skills, such as stable grasping and in-hand object reorientation. This policy is used zero-shot for tool-use without task-specific engineering.
- 2) We develop an object-centric perception pipeline that estimates the tool’s 6D pose and a coarse 3D grasp bounding box using vision foundation models, enabling zero-shot sim-to-real deployment of the RL policy.
- 3) We introduce *DexToolBench*, a benchmark of dexterous tool-use tasks. We show that *SimToolReal* achieves strong zero-shot performance on this benchmark over 120 real-world rollouts spanning 24 tasks, 12 object instances, and 6 tool categories. We outperform fixed-grasp and retargeting baselines by 37% in average *Task Progress* while matching the performance of specialist policies.

II. RELATED WORK

Imitation Learning from Teleoperated Data. Imitation learning (IL) is the predominant approach to train robot manipulation policies [90, 42, 20]. However, IL relies on high-quality demonstrations, typically collected via human teleoperation. Many systems have been proposed to teleoperate multi-fingered dexterous hands via motion capture gloves [69, 79], VR devices [23, 31, 43, 8], or even direct camera input [65, 27, 74]. However, directly mapping hand motion to robot actions can lead to imprecise, unintuitive control [13, 70, 7]. Further, operators receive limited haptic and tactile feedback during teleoperation, which limits the collected demonstrations to simple manipulation tasks [54, 56, 52, 75].

Recent works propose wearable exoskeleton systems that physically couple the robot hand to the operator during data collection [76, 84, 25]. Still, these bulky exoskeleton systems carry a considerable burden on the operator, constrain the range of human motion, and fail to demonstrate complex skills, such as in-hand reorientations. Another line of work treats the problem of teleoperation as shared autonomy. These approaches learn structured action-spaces from human hand motion data [53] or RL in simulation [29, 45], which are used by the human operator to control the robot. Recently, DexterityGen [88] trains policies conditioned on target grasps. In contrast, our policy conditions on object goal poses, avoiding an additional grasp-generation pipeline for training and desired grasp poses from a human operator at inference. Largely, shared-autonomy methods are *object-blind* at execution time and rely on skilled operators to bridge the gap between perception and control. In contrast to all these approaches, we focus on autonomous manipulation policies without any human teleoperation effort.

Learning Dexterity from Human Video Demonstrations.

Learning dexterous manipulation directly from human video demonstrations is a promising alternative to teleoperation. Kinematic motion retargeting pipelines can translate human hand motions into robot actions by reconstructing hand-object interaction [39, 36]. More recently, these pipelines have been extended to convert demonstrations synthesized by video-generation models into robot-executable trajectories [59, 38, 12]. However, because these methods primarily rely on kinematic hand-motion references, they struggle to produce actions that reliably grasp and maintain contact with tools.

To address this limitation, another line of work performs *functional retargeting*, focusing on matching the demonstrated task outcome rather than precisely matching human hand motion. Recent methods reproduce object motion from 3D hand-object motion-capture data in simulation while discovering robot-specific contact strategies using RL [50, 18, 40] or sampling-based optimization [57]. More recently, *functional retargeting* has been extended to RGB-D human video demonstrations [47, 22] and even AI-generated videos [51]. Despite their success, RL-based retargeting requires time-intensive per-demonstration training and yields policies tied to the demonstrated task and object, while sampling-based

optimization methods produce open-loop plans that cannot recover from errors and struggle on complex manipulation. In contrast, our approach can imitate complex object motion directly from a single RGB-D human demonstration without any additional demonstration-specific training. We can use the same policy across diverse tasks.

Sim-to-Real Learning. Policies learned in simulation can acquire dexterous manipulation behaviors without relying on human demonstrations, but transferring these skills to the real world often requires substantial *environment shaping* [58]. This includes closing the sim-to-real gap, building accurate task-specific object models, and carefully designing and tuning rewards for each behavior. Domain randomization over physical and visual parameters [28, 3, 6] and adaptation methods that infer real-world dynamics online [35, 64, 77] can reduce the sim-to-real gap significantly. Still, sim-to-real dexterous manipulation requires per-task engineering, which has led many systems to focus on narrower skill subsets such as grasping [1, 86, 46, 72, 71], in-hand object reorientation [14, 15, 45], or dynamic object spinning [80, 41]. While prior RL reorientation works only focus on rotating objects in-hand [14, 15], *SimToolReal* uses a single arm-hand policy for tabletop grasping, in-hand reorientation, and 6-DoF tool-motion tracking from human-video goal sequences. Recent research in locomotion has begun to unify RL policies across embodiments by leveraging large-scale simulation with procedurally generated robots and unified reward functions [44, 2]. Inspired by this direction, we train a single RL policy across procedurally generated object primitives and the universal reward function of reaching random goal poses. We show that this formulation induces highly dexterous skills (e.g., in-hand reorientation and dynamic spinning) and transfers to diverse real-world tools without task-specific engineering effort.

III. SIMTOOLREAL

We propose *SimToolReal*, a sim-to-real RL framework for zero-shot, dexterous tool manipulation that generalizes to unseen tools and manipulation sequences. Our key insight is to view dexterous tool-use tasks through an *object-centric lens*: a wide range of tasks can be specified as manipulating a tool through a sequence of goal poses. This reduces the tool-use problem to learning a single goal-reaching RL policy in simulation that manipulates procedurally generated tool primitives toward random goal poses. This training objective induces dexterous behaviors, such as stable grasping and in-hand reorientation, that are essential for tool use. At test time, the policy takes as input a sequence of object goal poses (e.g., from a human video) and executes actions to track this sequence, enabling diverse zero-shot tool-use without any real-world object modeling in simulation or task-specific training.

A. Problem Formulation

We consider the problem of training dexterous policies that can manipulate any tool to arbitrary goal poses. Let $\mathbf{o}_t \in SE(3)$ denote the current object pose at time t , $\mathbf{s}_t$ denote robot proprioception, and ϕ denote a coarse object descriptor




Fig. 3: **Real-World Deployment.** (Left) Human Video Processing: We collect an RGB-D human video and process it using vision foundation models. We use SAM 3D [17] to generate a metric-scale object mesh and segment a 3D grasp bounding box. Then, we use FoundationPose [82] to extract a sequence of 6D goal poses. (Right) Inference-Time Pipeline: Our LSTM policy takes in proprioception, object pose, grasp bounding box, and goal pose, and it outputs joint position targets for the 29-DoF dexterous robot (arm + hand).

(e.g., approximate geometry or physical properties). Given a goal pose $\mathbf{g} \in SE(3)$, we learn a policy $\mathbf{a}_t = \pi_\theta(\mathbf{s}_t, \mathbf{o}_t, \phi, \mathbf{g})$ that outputs joint position targets for the robot’s arm and hand. Task execution then reduces to repeatedly reaching the current goal pose and advancing to the next goal. This abstraction lets a single goal-reaching controller support diverse tool-use tasks, without task-specific policy training or reward design.

To perform a tool-use task, we assume access to a sequence of goal object poses $\{\mathbf{g}^k\}_{k=1}^K$ with $\mathbf{g}^k \in SE(3)$, extracted from a human demonstration video. For example, in Fig. 2, we visualize the goal sequence of reorienting the hammer, aligning it above the nail, and swinging. During execution, we condition the policy on the current goal pose $\mathbf{g}^k$ and advance to the next goal $\mathbf{g}^{k+1}$ when the object pose $\mathbf{o}_t$ is sufficiently close, i.e., when $d(\mathbf{o}_t, \mathbf{g}^k) < \epsilon$. Here, d is a pose distance function and ϵ is the success threshold.

In the following, we describe how we train our goal-conditioned RL policy in simulation, and the sim-to-real pipeline used for real-world deployment.

B. Training a General-Purpose Goal-Reaching Policy

Environment Setup. We design a general environment to model tool-use tasks. At the start of each episode, we place a randomly selected object on the table in a random pose and initialize the robot in a randomized joint configuration. We then sample a sequence of goal object poses $\{\mathbf{g}^k\}_{k=1}^K$. The robot must first learn to grasp the object from the flat table surface and subsequently manipulate it to reach each goal. If the object is dropped, the episode terminates and is reset. We sample the first goal randomly in the robot’s reachable workspace to expose the robot to a wide range of object poses and large reorientations. Subsequent goals are sampled close to

the previous goal to encourage smooth, trajectory-like motion.

Reward Function. Our primary objective is to train a goal-pose reaching RL policy. Practically, we optimize [61]:

$$r = r_{\text{smooth}} + r_{\text{grasp}} + \mathbb{I}_{\text{grasped}} r_{\text{goal}}. \quad (1)$$

r_{smooth} is a regularization term that encourages smooth actions, r_{grasp} is a shaping term that encourages grasping the object at the start of the episode, and $\mathbb{I}_{\text{grasped}}$ indicates whether the object has been grasped. After the object is grasped, r_{goal} is the dominant reward term, driving goal-pose reaching. We define:

$$r_{\text{goal}} = \max(d^* - d(\mathbf{o}_t, \mathbf{g}), 0) + B_{\text{succ}} \mathbb{I}[d(\mathbf{o}_t, \mathbf{g}) < \epsilon], \quad (2)$$

which consists of (i) a dense term that minimizes the distance $d(\mathbf{o}_t, \mathbf{g})$ to the goal, and (ii) a large, sparse success bonus B_{succ} . d^* is a stateful variable that tracks the minimum distance achieved so far, which is reset when a new goal is sampled. Positive reward is only provided for progress towards the goal, preventing the agent from exploiting the reward function by maintaining static proximity. Whenever the goal pose is reached (i.e., $d(\mathbf{o}_t, \mathbf{g}) < \epsilon$), the agent receives the success bonus and a new goal is sampled. Following [4, 61], we define the distance to the goal pose as $d(\mathbf{o}_t, \mathbf{g}) = \max_i \|\mathbf{o}_{t,i} - \mathbf{g}_i\|$, where we represent each pose using $D = 4$ keypoints in its local frame. $\mathbf{o}_{t,i}$ and $\mathbf{g}_i$ denote the positions of the i -th keypoint of the current and goal poses, respectively. See Appendix A for more details about the reward terms.

Procedural Generation of Tools in Simulation. We generate a large diversity of tool-like primitive objects in simulation such that they span the variation seen in real-world tools. Each tool primitive is generated as a combination of a *handle* and a *head*: we sample their geometries as cylinders and cuboids with varying dimensions. While simple, this design

captures the key structure of many everyday tools, such as brushes, spatulas, markers, and hammers. A random subset of our generated primitives is visualized in Fig 2. In addition to geometric variation, we randomize mass distribution by assigning different densities to the handle and head, representative of many real-world tools (e.g., hammer heads are typically much denser than handles). Additional generation details are provided in Appendix B.

Object-Centric Policy Inputs. To deploy a single policy across diverse real-world tools, we need an object representation that makes learning efficient while also being feasible to extract in the real-world. Real-world tools are diverse in their geometry and physical properties, and accurately estimating detailed geometric or physical parameters from visual sensing alone is impractical. Instead, we provide the policy only the object inputs that are reliably available at deployment: the current 6D tool pose and a coarse 3D grasp-region bounding box. The bounding box encodes the intended graspable region (center + extents in the object frame; Fig. 3) and is held fixed during an episode. We use an LSTM backbone so the policy can integrate interaction history and implicitly infer latent physical and geometric properties that are not directly observed. This design follows the spirit of prior work such as DexFunc [1] and RMA [64], which demonstrate dexterous manipulation behaviors with limited object perception to reduce the sim-to-real observation gap.

RL Training Details. We highlight key RL design choices essential for policy learning and sim-to-real transfer.

a) *SAPG* [73]: We train our policy using SAPG, a variant of PPO [68]. We find that standard PPO encounters exploration bottlenecks in massively parallel simulations. SAPG mitigates this by maintaining a population of policies to promote exploration diversity and updating the leader policy using their collective experience.

b) *Domain Randomization*: We apply targeted domain randomization during training to aid sim-to-real transfer. This includes observation delays, action-execution latency, noise and delay in object-pose estimates, and perturbations to the grasp-region bounding box. We additionally apply random force and torque perturbations to the object, encouraging the policy to learn strong, stable grasps. Randomization ranges and other details are provided in Appendix D.

c) *Asymmetric Critic* [62]: While the actor receives only the minimal object representation available at test time, we train with an asymmetric actor-critic setup in which the critic has access to privileged states in simulation. We also provide the critic with clean (noise-free, undelayed) observations to improve value estimation and stabilize training. Additional details are included in Appendix D.

C. Real-World Deployment

In this section, we describe how we deploy our object-centric policy in the real world to perform a wide range of tool-use tasks. Fig. 3 provides an overview. Given a third-person RGB-D view of a human demonstration, we extract (i) a 3D metric-scale object mesh and grasp region bounding

box, and (ii) a sequence of target object poses that serves as the goal trajectory. During execution, we run the policy in a closed loop to reach the goal poses sequentially.

In contrast to Lum et al. [47], our approach does not require access to the real-world object or target trajectory during training, enabling zero-shot generalization to novel objects using demonstrations collected entirely after training.

Human Video Processing. From the first RGB-D frame, we reconstruct a metric-scale 3D mesh of the object using SAM 3D [17]. We then segment the intended grasp region using SAM 2 [67], and convert this region into a coarse 3D bounding box that is provided as part of the policy input. This semi-automated pipeline renders the mesh from multiple views, uses sparse user point prompts to initialize a SAM 2 mask for the graspable region, propagates the mask across the rendered-view sequence, and aggregates the masked geometry into the 3D grasp-region bounding box. Finally, we run FoundationPose [82] on the RGB-D video conditioned on the extracted mesh to obtain the 6D object-pose trajectory. We process this trajectory by downsampling to 3 Hz. We use this processed trajectory as the goal sequence at test time. See Appendix E for implementation details.

Inference-time Object Tracking. During inference, we estimate the current 6D object pose with an update frequency of 30 Hz. For live pose tracking, we use FoundationPose [82] with RGB-D observations from a third-person camera view combined with the 3D mesh extracted from the human video. At each control step, the policy is conditioned on proprioception, the current object pose, the (fixed) grasp-region bounding box, and the current goal pose from the demonstration trajectory. We initialize the goal as the first pose in the sequence, switching it to the next goal whenever the pose error $d(\mathbf{o}_t, \mathbf{g})$ falls within a tolerance ϵ . The episode ends when the final goal is reached or upon failure (e.g., loss of grasp or tracking).

D. DexToolBench

Existing robot manipulation benchmarks primarily focus on robots with parallel-jaw grippers [48, 32, 9, 63, 33, 24, 37] or dexterous grasping [86, 81, 91, 85, 78]. We introduce *DexToolBench*: a real-world, dexterous manipulation benchmark of challenging tool-use tasks paired with digital-twin simulation environments. These tasks require grasping tools from flat surfaces, in-hand rotations to functional configurations, and environment interactions. This benchmark consists of 24 daily tool-use tasks, 12 unique object instances across 6 categories.

Each task is defined by an RGB-D human video, captured via a ZED 1 stereo camera. We process each video with our perception pipeline (Fig. 3) to extract a 3D object mesh and a 6D object pose trajectory. To facilitate reproducibility, we provide the raw RGB-D videos, processed data, simulation environments, and purchase links for the real-world tools.

Each tool instance is visualized in Fig 4. We briefly describe the skills required in each tool-use category in the benchmark:

- a) *Hammer*: Grasp, 90° in-hand rotation, swing down or side.
- b) *Marker*: Grasp a thin object, write on a whiteboard.
- c) *Eraser*: Grasp, wipe marker ink on whiteboard.




Fig. 4: **Generalization to Unseen *DexToolBench* Tools and Tasks in the Real World.** We evaluate our policy in the real world on unseen tool-use tasks in *DexToolBench*. Our evaluations span 24 unique task trajectories across 6 different object categories and 12 object instances. Each bar corresponds to 1 task trajectory on 1 object instance. We report the average *Task Progress* across 5 rollouts. Despite not being trained on these objects or trajectories, our policy demonstrates strong generalization to diverse tools of varying masses and geometries.

d) *Brush*: Grasp, 90° in-hand rotation, sweep forward or right.

e) *Spatula*: Grasp, serve or 180° in-hand rotation and flip.

f) *Screwdriver*: Grasp, 90° in-hand rotation, spin in free space along vertical or horizontal axis.

See Appendix F for more details about *DexToolBench*.

IV. EXPERIMENTS

Experimental Setup: We evaluate *SimToolReal* on a robot consisting of a 22-DoF Sharpa five-fingered left hand mounted on a 7-DoF KUKA iiwa 14 arm. We evaluate our policies on *DexToolBench* (Sec. III-D). For each task, we evaluate *Task Progress*, measuring the percentage of demonstrated goal poses tracked successfully. We focus on trajectory following of a fixed goal sequence rather than evaluating functional task completion. We consider a goal pose to be reached if the distance between the object pose and goal pose $d(o_t, g)$ is below our defined success tolerance $\epsilon = 2cm$. We design experiments to answer the following questions:

- What is the zero-shot transfer performance of *SimToolReal* on unseen objects and trajectories in the real world?
- How does *SimToolReal* compare to prior methods that use kinematic motion retargeting or fixed-grasps?
- How does *SimToolReal* compare to task-specific specialists trained on a single object and trajectory?
- To what extent does our training objective correlate with downstream performance on unseen tools and tasks?
- Which RL design decisions are important for maximizing training performance?

A. Zero-Shot Real-World Tool-Use

Setup. We evaluate whether our single policy trained in simulation transfers zero-shot to real-world tool use, without any object- or task-specific fine-tuning. As shown in Fig. 4, our

evaluation spans 24 task trajectories across 6 tool categories and 12 object instances in *DexToolBench* (Sec. III-D), all of which were unseen during training. We report average *Task Progress* over 5 trials, totaling 120 real-world rollouts.

Results. Overall, the policy demonstrates strong zero-shot generalization across tools with diverse masses and geometries. We observe the highest *Task Progress* on eraser trajectories, which rely primarily on translation motion rather than in-hand object rotation. Similarly, marker trajectories don’t require in-hand rotations, but because the tool is much thinner, the grasps are less reliable. Further, the marker’s small size makes it prone to pose tracking loss during occlusion.

The remaining four categories (hammer, brush, spatula, screwdriver) require substantial in-hand rotation to reach functional tool configurations. While the policy achieves strong *Task Progress* across these tasks, performance degrades on thinner tools (e.g., performs better on the ~3cm thick spoon spatula than the ~1cm thick flat spatula) and heavier tools (e.g., performs better on the 36g claw hammer than the 331g mallet hammer). We also note that tool geometry impacts task difficulty. For instance, the red brush’s sideways head makes sweep forward easier, but the blue brush’s frontal head makes sweep right easier. The screwdriver is the most challenging category, as the task requires both functional reorientation and continuous spinning.

Failure Analysis. The most common failure mode was pose tracking loss (43.7% of failures), followed by object drops (34.5%), failure to reach the goal pose due to incomplete in-hand rotation (18.2%), and grasp failure (3.6%). See Appendix G for additional quantitative and qualitative analysis of our results. Notably, we observe strong recovery behavior when the policy makes mistakes. For example, when the object




Fig. 5: **Comparison against Baselines in the Real World.** We compare *SimToolReal* against baselines on two variations of sweeping a table with a brush: with and without requiring tool rotation based on the initial states shown on the left. Average *Task Progress* is indicated in parentheses. *SimToolReal* succeeds on both variations, performing dexterous in-hand tool rotations in the harder variation. *Fixed Grasp* succeeds on the simpler variation of this task without tool rotation. However, when rotation is required, enforcing a fixed grasp causes the arm to collide with the table while tracking the target trajectory. *Kinematic Retargeting* fails to reason about contact forces, and is unable to grasp the brush in both variations.

is dropped, the policy consistently attempts to re-grasp the object, provided the object remains within the workspace and pose tracking is not lost.

B. Comparisons to Retargeting and Fixed Grasp Baselines

Setup. In this section, we compare *SimToolReal* with:

(i) *Kinematic Retargeting*: Following prior works [39, 47, 26, 66, 60], we retarget hand motion from a human video into dexterous robot actions. From an RGB-D video, we estimate 3D hand finger positions and then solve for arm and hand joint positions using an Inverse Kinematics (IK) solver. Although this method bypasses the need for RL, the resulting motion is purely kinematic and fails to account for contact interactions. See Appendix H for implementation details.

(ii) *Fixed Grasp*: Recent works [1, 59, 30, 93, 5] frame manipulation tasks as establishing a grasp, then subsequently fixing the grasp while moving the object. To perform any object rotations, the robot must strictly rely on arm motion. This baseline isolates the necessity of in-hand manipulation; failure here demonstrates that the arm lacks the kinematic workspace to rotate tools arbitrarily without the assistance of in-hand rotation. See Appendix I for implementation details.

To systematically evaluate the baselines, we set up two variations of sweep forward with the red brush in the real-world (see Fig. 5): (i) *Variation #1* starts with the brush top-down and does not require tool rotation, and (ii) *Variation #2* starts with the brush sideways, which necessitates a 90° tool rotation. We evaluate average *Task Progress* across 5 rollouts.

Results. *Kinematic Retargeting* fails to grasp the brush in both variations as kinematic motion alone does not establish stable contacts. *Fixed Grasp* succeeds on *Variation #1* as tool rotation is not required and arm motion alone is sufficient to finish the task. Still, it has lower *Task Progress* than *SimToolReal* since it is an open-loop method and cannot react

to small errors. On *Variation #2*, enforcing a fixed grasp causes the robot arm to collide with the table while tracking the target trajectory, as the optimizer is unable to find an arm motion trajectory that satisfies the tool rotation while avoiding table collision. In contrast, *SimToolReal* performs efficient in-hand object rotations to complete the task.




Fig. 6: **Comparison against Specialists.** We compare *SimToolReal* in simulation against specialist policies trained on a single object (Obj A) and trajectory (Traj A). We train one specialist policy for each of the 6 object categories in *DexToolBench* and report the average *Task Progress* across these categories. While the specialists succeed on their training setup (Obj A / Traj A), performance degrades under deviation in the trajectory (Obj A / Traj B) or the object (Obj B / Traj A). *SimToolReal* has high zero-shot performance across all variants, despite not being trained on these objects or trajectories.

C. Comparisons to Specialist Baseline

Setup. In simulation, we compare *SimToolReal* against task-specific specialist RL baselines trained on a single tool and task. Following [47], we train 6 specialist policies, one for each

of the 6 tool-use categories in *DexToolBench*. Each specialist policy trains on a single object instance (Obj A) and task trajectory (Traj A). We evaluate both methods in simulation on (i) the training setup for the specialist: (Obj A / Traj A), (ii) the training object with a novel trajectory (Obj A / Traj B), and (iii) a novel object instance with the same trajectory (Obj B / Traj A). For each variation, we perform 10 rollouts for each of the 6 specialist policies, reporting the average *Task Progress* across all trials.

Results. Fig. 6 summarizes these results. On the setup used for training the specialist (Obj A / Traj A), *SimToolReal* matches the specialist without training on that object or trajectory. However, the specialist’s performance drops significantly under any variation from the training setup. When the object stays the same, with only a change in trajectory (Obj A / Traj B), the specialist can only track the first few goals that lift the object. The performance is lowest when the object instance is changed (e.g., from red brush to blue brush) even if the trajectory remains the same (Obj B / Traj A). This indicates that the specialist overfits to the training conditions. In contrast, *SimToolReal* shows strong zero-shot *Task Progress* across both object and trajectory variations.

D. Training Objective Predicts Generalization Performance

Setup. We study how improving on our training objective: reaching random goal poses across procedurally generated primitive objects, relates to performance on our test objective: executing human-demonstrated tool-use trajectories in *DexToolBench*. Fig. 7 reports (i) training rewards on the goal-pose reaching objective on primitive objects and (ii) average *Task Progress* on *DexToolBench* trajectory following tasks in simulation, plotted against the number of environment steps.




Fig. 7: **Training Objective Drives Generalization.** In simulation, we evaluate (Left) the episode reward on procedurally-generated objects and (Right) the zero-shot *Task Progress* on unseen *DexToolBench* tools on different policy checkpoints throughout training. The strong correlation between the two curves validates our core hypothesis: improving random goal-pose reaching performance on diverse object primitives drives corresponding gains in generalization to unseen tool-use tasks.

Results. As training proceeds, the policy’s reward on the goal-pose reaching objective increases steadily. In parallel, the policy’s average *Task Progress* on *DexToolBench* trajectories

also improves consistently. The correlation between these curves indicates that optimizing our training objective is a reliable predictor of downstream generalization to tool-use behaviors, despite the mismatch between random goal-pose reaching on primitive objects during training and trajectory following on real tool objects at test time. This strong correlation suggests that our randomized primitive training effectively covers the space of skills required for real-world tool use.

E. Ablations on RL Training

Setup. We investigate which RL design decisions are critical for maximizing training performance. We focus on two key components of our RL pipeline: the choice of optimization algorithm (SAPG [73] vs. PPO [68]) and the use of privileged information (Asymmetric Critic [62] vs. standard critic). We train these variants with identical hyperparameters and environment settings, averaging results across 5 random seeds.




Fig. 8: **Ablation of RL Training Components.** We compare the training reward across environment steps of *SimToolReal* against ablations averaged across 5 seeds. Replacing SAPG [73] with PPO [68] or not using Asymmetric Critic [62] results in a significant performance drop, highlighting their importance.

Results. Fig. 8 compares the training reward curves of our full method against these ablations. *Importance of SAPG.* We observe that replacing SAPG [73] with PPO [68] leads to a significant drop in performance. While PPO suffers from exploration saturation at scale, SAPG mitigates this by training separate policies across environment chunks to increase data diversity, and then fusing gradients via importance sampling. Our results confirm that this is essential for learning complex dexterous tool-use. *Importance of Asymmetric Critic.* We observe that removing the Asymmetric Critic [62] and forcing the critic to rely solely on the same partial observations as the actor severely hinders learning. By accessing privileged simulation states, the critic learns a more accurate value function to guide policy learning. This guidance is essential for overcoming the task’s partial observability.

V. DISCUSSION AND LIMITATIONS

We introduced *SimToolReal*, an RL framework for zero-shot dexterous tool manipulation that generalizes to unseen tool-use tasks. Our key insight is an object-centric reduction: tool-use can be specified as manipulating a tool through a sequence

of goal poses. This shifts the problem from per-task reward design to a universal objective of training goal-pose reaching policies. Training with this objective over procedurally generated primitives induces dexterous skills, such as stable grasping and in-hand reorientation, critical for tool-use.

Limitations. While our approach can track tool-use goal sequences, it does not guarantee functional task completion, especially for high-force interactions. Conditioning on object pose goals alone is environment-blind, which can lead to collisions in cluttered scenes. Our grasp-region abstraction also assumes a distinct, localized graspable region that can be approximated by a coarse 3D bounding box, which is less natural for objects whose grasp depends on an edge or broad surface, such as plates or cutting boards. We also currently assume tools are rigid. Pose alone can be insufficient to describe the state of non-rigid tools (e.g., scissors). Finally, our high-level goal sequence is fixed and is not replanned dynamically.

ACKNOWLEDGEMENTS

This work is supported by Stanford Human-Centered Artificial Intelligence (HAI), ONR Young Investigator Award, the National Science Foundation (NSF) under Grant Numbers 2153854, 2327974, 2312956, 2327973, and 2342246, and the Natural Sciences and Engineering Research Council of Canada (NSERC) under Award Number 526541680. We thank Sharpa for their research collaboration and for the technical support provided by their team, specifically Kaifeng Zhang, Wenjie Mei, Yi Zhou, Yunfang Yang, Jie Yin, and Jason Lee.

REFERENCES

- [1] Ananye Agarwal, Shagun Uppal, Kenneth Shaw, and Deepak Pathak. Dexterous functional grasping, 2023. URL <https://arxiv.org/abs/2312.02975>.
- [2] Bo Ai, Liu Dai, Nico Bohlinger, Dichen Li, Tongzhou Mu, Zhanxin Wu, K Fay, Henrik I Christensen, Jan Peters, and Hao Su. Towards embodiment scaling laws in robot locomotion. *arXiv preprint arXiv:2505.05753*, 2025.
- [3] Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- [4] Arthur Allshire, Mayank Mittal, Varun Lodaya, Viktor Makoviychuk, Denys Makoviichuk, Felix Widmaier, Manuel Wüthrich, Stefan Bauer, Ankur Handa, and Animesh Garg. Transferring dexterous manipulation from gpu simulation to a remote real-world trifinger. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11802–11809. IEEE, 2022.
- [5] Sai Haneesh Allu, Jishnu Jaykumar P, Ninad Khargonkar, Tyler Summers, Jian Yao, and Yu Xiang. Hrt1: One-shot human-to-robot trajectory transfer for mobile manipulation, 2025. URL <https://arxiv.org/abs/2510.21026>.
- [6] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39 (1):3–20, 2020.
- [7] Miguel Arduengo, Ana Arduengo, Adrià Colomé, Joan Lobo-Prat, and Carme Torras. Human to robot whole-body motion transfer. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 299–305, 2021. doi: 10.1109/HUMANOIDS47582.2021.9555769.
- [8] Sridhar Pandian Arunachalam, Irmak Güzey, Soumith Chintala, and Lerrel Pinto. Holo-dex: Teaching dexterity with immersive mixed reality. *arXiv preprint arXiv:2210.06463*, 2022.
- [9] Pranav Atreya, Karl Pertsch, Tony Lee, Moo Jin Kim, Arhan Jain, Artur Kuramshin, Clemens Eppner, Cyrus Neary, Edward Hu, Fabio Ramos, et al. Roboarena: Distributed real-world evaluation of generalist robot policies. In *Proceedings of the Conference on Robot Learning (CoRL 2025)*, 2025.
- [10] Genesis Authors. Genesis: A generative and universal physics engine for robotics and beyond, December 2024. URL <https://github.com/Genesis-Embodied-AI/Genesis>.
- [11] Samuel R Buss. Introduction to inverse kinematics with jacobian transpose, pseudoinverse and damped least squares methods. Unpublished manuscript, University of California, San Diego, 2004. Available at <https://www.math.ucsd.edu/~sbuss/ResearchWeb/ikmethods/iksurvey.pdf>.
- [12] Boyuan Chen, Tianyuan Zhang, Haoran Geng, Kiwhan Song, Caiyi Zhang, Peihao Li, William T Freeman, Jitendra Malik, Pieter Abbeel, Russ Tedrake, et al. Large video planner enables generalizable robot control. *arXiv preprint arXiv:2512.15840*, 2025.
- [13] Claire Chen, Zhongchun Yu, Hojung Choi, Mark Cutkosky, and Jeannette Bohg. Dexforce: Extracting force-informed actions from kinesthetic demonstrations for dexterous manipulation. *IEEE Robotics and Automation Letters*, 10(6):6416–6423, 2025.
- [14] Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. *Conference on Robot Learning*, 2021.
- [15] Tao Chen, Megha Tippur, Siyang Wu, Vikash Kumar, Edward Adelson, and Pulkit Agrawal. Visual dexterity: In-hand reorientation of novel and complex object shapes. *Science Robotics*, 8(84), 2023. ISSN 2470-9476. doi: 10.1126/scirobotics.adc9244. URL <http://dx.doi.org/10.1126/scirobotics.adc9244>.
- [16] Tao Chen, Eric Cousineau, Naveen Kuppaswamy, and Pulkit Agrawal. Vegetable peeling: A case study in constrained dexterous manipulation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4542–4550, 2025. doi: 10.1109/ICRA55743.2025.11127224.

- [17] Xingyu Chen, Fu-Jen Chu, Pierre Gleize, Kevin J Liang, Alexander Sax, Hao Tang, Weiyao Wang, Michelle Guo, Thibaut Hardin, Xiang Li, et al. Sam 3d: 3dfy anything in images. *arXiv preprint arXiv:2511.16624*, 2025.
- [18] Yuanpei Chen, Chen Wang, Yaodong Yang, and Karen Liu. Object-centric dexterous manipulation from human motion data. In *8th Annual Conference on Robot Learning*, 2024.
- [19] Zeyuan Chen, Qiyang Yan, Yuanpei Chen, Tianhao Wu, Jiayao Zhang, Zihan Ding, Jinzhou Li, Yaodong Yang, and Hao Dong. Clutterdexgrasp: A sim-to-real system for general dexterous grasping in cluttered scenes. *arXiv preprint arXiv:2506.14317*, 2025.
- [20] Baiye Cheng, Tianhai Liang, Suning Huang, Maanping Shao, Feihong Zhang, Botian Xu, Zhengrong Xue, and Huazhe Xu. Moe-dp: An moe-enhanced diffusion policy for robust long-horizon robotic manipulation with skill decomposition and failure recovery. *arXiv preprint arXiv:2511.05007*, 2025.
- [21] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Robotics: Science and Systems*, 2023.
- [22] Prithwish Dan, Kushal Kedia, Angela Chao, Edward Weiye Duan, Maximus Adrian Pace, Wei-Chiu Ma, and Sanjiban Choudhury. X-sim: Cross-embodiment learning via real-to-sim-to-real. *arXiv preprint arXiv:2505.07096*, 2025.
- [23] Runyu Ding, Yuzhe Qin, Jiyue Zhu, Chengzhe Jia, Shiqi Yang, Ruihan Yang, Xiaojuan Qi, and Xiaolong Wang. Bunny-visionpro: Real-time bimanual dexterous teleoperation for imitation learning. 2024.
- [24] Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11441–11450, 2020. doi: 10.1109/CVPR42600.2020.01146.
- [25] Hao-Shu Fang, Branden Romero, Yichen Xie, Arthur Hu, Bo-Ruei Huang, Juan Alvarez, Matthew Kim, Gabriel Margolis, Kavya Anbarasu, Masayoshi Tomizuka, Edward Adelson, and Pulkit Agrawal. Dexop: A device for robotic transfer of dexterous human manipulation. *arXiv preprint arXiv:2509.04441*, 2025.
- [26] Irmak Guzey, Yinlong Dai, Georgy Savva, Raunaq Bhiringi, and Lerrel Pinto. Bridging the human to robot dexterity gap through object-oriented rewards, 2024. URL <https://arxiv.org/abs/2410.23289>.
- [27] Ankur Handa, Karl Van Wyk, Wei Yang, Jacky Liang, Yu-Wei Chao, Qian Wan, Stan Birchfield, Nathan Ratliff, and Dieter Fox. Dexpivot: Vision based teleoperation of dexterous robotic hand-arm system, 2019. URL <https://arxiv.org/abs/1910.03135>.
- [28] Ankur Handa, Arthur Allshire, Viktor Makoviychuk, Aleksei Petrenko, Ritvik Singh, Jingzhou Liu, Denys Makoviichuk, Karl Van Wyk, Alexander Zhurkevich, Balakumar Sundaralingam, and Yashraj Narang. Dextreme: Transfer of agile in-hand manipulation from simulation to reality. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5977–5984, 2023. doi: 10.1109/ICRA48891.2023.10160216.
- [29] Elvis Hsieh, Wen-Han Hsieh, Yen-Jen Wang, Toru Lin, Jitendra Malik, Koushil Sreenath, and Haozhi Qi. Learning dexterous manipulation skills from imperfect simulations. *arXiv:2512.02011*, 2025.
- [30] Cheng-Chun Hsu, Bowen Wen, Jie Xu, Yashraj Narang, Xiaolong Wang, Yuke Zhu, Joydeep Biswas, and Stan Birchfield. Spot: Se(3) pose trajectory diffusion for object-centric manipulation, 2025. URL <https://arxiv.org/abs/2411.00965>.
- [31] Aadithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Haldar, Soumith Chintala, and Lerrel Pinto. Open teach: A versatile teleoperation system for robotic manipulation. *arXiv preprint arXiv:2403.07870*, 2024.
- [32] Arhan Jain, Mingtong Zhang, Kanav Arora, William Chen, Marcel Torne, Muhammad Zubair Irshad, Sergey Zakharov, Yue Wang, Sergey Levine, Chelsea Finn, Wei-Chiu Ma, Dhruv Shah, Abhishek Gupta, and Karl Pertsch. Polaris: Scalable real-to-sim evaluations for generalist robot policies, 2025. URL <https://arxiv.org/abs/2512.16881>.
- [33] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 2020.
- [34] Chung Min Kim*, Brent Yi*, Hongsuk Choi, Yi Ma, Ken Goldberg, and Angjoo Kanazawa. Pyroki: A modular toolkit for robot kinematic optimization. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025. URL <https://arxiv.org/abs/2505.03728>.
- [35] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021.
- [36] Arjun S Lakshmipathy, Jessica K Hodgins, and Nancy S Pollard. Kinematic motion retargeting for contact-rich anthropomorphic manipulations. *ACM Transactions on Graphics*, 44(2):1–20, 2025.
- [37] Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabriel Levine, Wensi Ai, Benjamin Martinez, Hang Yin, Michael Lingelbach, Minjune Hwang, Ayano Hiranaka, Sujay Garlanka, Arman Aydin, Sharon Lee, Jiankai Sun, Mona Anvari, Manasi Sharma, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Yunzhu Li, Silvio Savarese, Hyowon Gweon, C. Karen Liu, Jiajun Wu, and Li Fei-Fei. Behavior-1k: A human-centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation. *arXiv preprint arXiv:2403.09227*, 2024.

- [38] Hongyu Li, Lingfeng Sun, Yafei Hu, Duy Ta, Jennifer Barry, George Konidaris, and Jiahui Fu. Novaflow: Zero-shot manipulation via actionable flow from generated videos. *arXiv preprint arXiv:2510.08568*, 2025.
- [39] Jinhan Li, Yifeng Zhu, Yuqi Xie, Zhenyu Jiang, Mingyo Seo, Georgios Pavlakos, and Yuke Zhu. Okami: Teaching humanoid robots manipulation skills through single video imitation. In *8th Annual Conference on Robot Learning (CoRL)*, 2024.
- [40] Kailin Li, Puhao Li, Tengyu Liu, Yuyang Li, and Siyuan Huang. Maniptrans: Efficient dexterous bimanual manipulation transfer via residual learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6991–7003, 2025.
- [41] Toru Lin, Zhao-Heng Yin, Haozhi Qi, Pieter Abbeel, and Jitendra Malik. Twisting lids off with two hands. *arXiv:2403.02338*, 2024.
- [42] Toru Lin, Yu Zhang, Qiyang Li, Haozhi Qi, Brent Yi, Sergey Levine, and Jitendra Malik. Learning visuotactile skills with two multifingered hands. *arXiv:2404.16823*, 2024.
- [43] Toru Lin, Yu Zhang, Qiyang Li, Haozhi Qi, Brent Yi, Sergey Levine, and Jitendra Malik. Learning visuotactile skills with two multifingered hands. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5637–5643. IEEE, 2025.
- [44] Min Liu, Deepak Pathak, and Ananye Agarwal. Locomformer: Generalist locomotion via long-context adaptation. *arXiv preprint arXiv:2509.23745*, 2025.
- [45] Xueyi Liu, He Wang, and Li Yi. Dexndm: Closing the reality gap for dexterous in-hand rotation via joint-wise neural dynamics model, 2025. URL <https://arxiv.org/abs/2510.08556>.
- [46] Tyler Ga Wei Lum, Martin Matak, Viktor Makoviychuk, Ankur Handa, Arthur Allshire, Tucker Hermans, Nathan D. Ratliff, and Karl Van Wyk. DextrAH-g: Pixels-to-action dexterous arm-hand grasping with geometric fabrics. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=S2Jwb0i7HN>.
- [47] Tyler Ga Wei Lum, Olivia Y. Lee, C. Karen Liu, and Jeannette Bohg. Crossing the human-robot embodiment gap with sim-to-real rl using one human demonstration, 2025. URL <https://arxiv.org/abs/2504.12609>.
- [48] Jianlan Luo, Charles Xu, Fangchen Liu, Liam Tan, Zipeng Lin, Jeffrey Wu, Pieter Abbeel, and Sergey Levine. Fmb: a functional manipulation benchmark for generalizable robotic learning, 2024. URL <https://arxiv.org/abs/2401.08553>.
- [49] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning, 2021. URL <https://arxiv.org/abs/2108.10470>.
- [50] Zhao Mandi, Yifan Hou, Dieter Fox, Yashraj Narang, Ajay Mandlekar, and Shuran Song. Dexmachina: Functional retargeting for bimanual dexterous manipulation, 2025. URL <https://arxiv.org/abs/2505.24853>.
- [51] Jiageng Mao, Sicheng He, Hao-Ning Wu, Yang You, Shuyang Sun, Zhicheng Wang, Yanan Bao, Huizhong Chen, Leonidas Guibas, Vitor Guizilini, et al. Robot learning from a physical world model. *arXiv preprint arXiv:2511.07416*, 2025.
- [52] Dan Morris, Hong Tan, Federico Barbagli, Timothy Chang, and Kenneth Salisbury. Haptic feedback enhances force skill learning. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC’07)*, pages 21–26, 2007. doi: 10.1109/WHC.2007.65.
- [53] Patrick Naughton, Jinda Cui, Karankumar Patel, and Soshi Iba. Respilot: Teleoperated finger gaiting via gaussian process residual learning. In *Conference on Robot Learning (CoRL)*, 2024.
- [54] Allison M Okamura. Haptic feedback in robot-assisted minimally invasive surgery. *Current opinion in urology*, 19(1):102–107, 2009.
- [55] OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation, 2019. URL <https://arxiv.org/abs/1808.00177>.
- [56] Claudio Pacchierotti and Domenico Prattichizzo. Cutaneous/tactile haptic feedback in robotic teleoperation: Motivation, survey, and perspectives. *IEEE Transactions on Robotics*, 40:978–998, 2023.
- [57] Chaoyi Pan, Changhao Wang, Haozhi Qi, Zixi Liu, Homanga Bharadhwaj, Akash Sharma, Tingfan Wu, Guanya Shi, Jitendra Malik, and Francois Hogan. Spider: Scalable physics-informed dexterous retargeting, 2025. URL <https://arxiv.org/abs/2511.09484>.
- [58] Younghyo Park, Gabriel B Margolis, and Pulkit Agrawal. Position: Automatic environment shaping is the next frontier in rl. In *Forty-first International Conference on Machine Learning*, 2024.
- [59] Shivansh Patel, Shraddhaa Mohan, Hanlin Mai, Unnat Jain, Svetlana Lazebnik, and Yunzhu Li. Robotic manipulation by imitating generated videos without physical demonstrations. *arXiv preprint arXiv:2507.00990*, 2025.
- [60] Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. Reconstructing hands in 3D with transformers. In *CVPR*, 2024.
- [61] Aleksei Petrenko, Arthur Allshire, Gavriel State, Ankur Handa, and Viktor Makoviychuk. Dexpbt: Scaling up dexterous manipulation for hand-arm systems with population based training. *CoRR*, abs/2305.12127, 2023. URL <https://doi.org/10.48550/arXiv.2305.12127>.
- [62] Lerrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Asymmetric actor

- critic for image-based robot learning. In Hadas Kress-Gazit, Siddhartha S. Srinivasa, Tom Howard, and Nikolay Atanasov, editors, *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018. doi: 10.15607/RSS.2018.XIV.008. URL <http://www.roboticsproceedings.org/rss14/p08.html>.
- [63] Wilbert Pumacay, Ishika Singh, Jiafei Duan, Ranjay Krishna, Jesse Thomason, and Dieter Fox. The colosseum: A benchmark for evaluating generalization for robotic manipulation. 2024.
- [64] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. In-Hand Object Rotation via Rapid Motor Adaptation. In *Conference on Robot Learning (CoRL)*, 2022.
- [65] Yuzhe Qin, Wei Yang, Binghao Huang, Karl Van Wyk, Hao Su, Xiaolong Wang, Yu-Wei Chao, and Dieter Fox. Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system, 2024. URL <https://arxiv.org/abs/2307.04577>.
- [66] Ri-Zhao Qiu, Shiqi Yang, Xuxin Cheng, Chaitanya Chawla, Jialong Li, Tairan He, Ge Yan, David J. Yoon, Ryan Hoque, Lars Paulsen, Ge Yang, Jian Zhang, Sha Yi, Guanya Shi, and Xiaolong Wang. Humanoid policy ~ human policy. *arXiv preprint arXiv:2503.13441*, 2025.
- [67] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. URL <https://arxiv.org/abs/2408.00714>.
- [68] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [69] Kenneth Shaw, Yulong Li, Jiahui Yang, Mohan Kumar Srirama, Ray Liu, Haoyu Xiong, Russell Mendonca, and Deepak Pathak. Bimanual dexterity for complex tasks. In *8th Annual Conference on Robot Learning*, 2024.
- [70] Zilin Si, Kevin Lee Zhang, Zeynep Temel, and Oliver Kroemer. Tilde: Teleoperation for Dexterous In-Hand Manipulation Learning with a DeltaHand. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. doi: 10.15607/RSS.2024.XX.128.
- [71] Ritvik Singh, Arthur Allshire, Ankur Handa, Nathan Ratliff, and Karl Van Wyk. Dextrah-rgb: Visuomotor policies to grasp anything with dexterous hands. *arXiv preprint arXiv:2412.01791*, 2024.
- [72] Ritvik Singh, Karl Van Wyk, Pieter Abbeel, Jitendra Malik, Nathan Ratliff, and Ankur Handa. End-to-end rl improves dexterous grasping policies. *arXiv preprint arXiv:2509.16434*, 2025.
- [73] Jayesh Singla, Ananye Agarwal, and Deepak Pathak. Sapg: Split and aggregate policy gradients. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, Proceedings of Machine Learning Research, Vienna, Austria, July 2024. PMLR.
- [74] Aravind Sivakumar, Kenneth Shaw, and Deepak Pathak. Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube. *arXiv preprint arXiv:2202.10448*, 2022.
- [75] P Ström, Leif Hedman, L Särnå, Ann Kjellin, T Wredmark, and Li Felländer-Tsai. Early exposure to haptic feedback enhances performance in surgical simulator training: A prospective randomized crossover study in surgical residents. *Surgical endoscopy*, 20:1383–8, 10 2006. doi: 10.1007/s00464-005-0545-3.
- [76] Tony Tao, Mohan Kumar Srirama, Jason Jingzhou Liu, Kenneth Shaw, and Deepak Pathak. Dexwild: Dexterous human interactions for in-the-wild robot policies. *arXiv preprint arXiv:2505.07813*, 2025.
- [77] Shagun Uppal, Ananye Agarwal, Haoyu Xiong, Kenneth Shaw, and Deepak Pathak. Spin: Simultaneous perception interaction and navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18133–18142, 2024.
- [78] Weikang Wan, Haoran Geng, Yun Liu, Zikang Shan, Yaodong Yang, Li Yi, and He Wang. Unidexgrasp++: Improving dexterous grasping policy learning via geometry-aware curriculum and iterative generalist-specialist learning. *arXiv preprint arXiv:2304.00464*, 2023.
- [79] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *arXiv preprint arXiv:2403.07788*, 2024.
- [80] Jun Wang, Ying Yuan, Haichuan Che, Haozhi Qi, Yi Ma, Jitendra Malik, and Xiaolong Wang. Lessons from learning to spin “pens”. In *CoRL*, 2024.
- [81] Youzhuo Wang, Jiayi Ye, Chuyang Xiao, Yiming Zhong, Heng Tao, Hang Yu, Yumeng Liu, Jingyi Yu, and Yuexin Ma. Dexh2r: A benchmark for dynamic dexterous grasping in human-to-robot handover, 2025. URL <https://arxiv.org/abs/2506.23152>.
- [82] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17868–17879, June 2024.
- [83] Philipp Wu, Yide Shentu, Zhongke Yi, Xingyu Lin, and Pieter 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, 2024. doi: 10.1109/IROS58592.2024.10801581.
- [84] Mengda Xu, Han Zhang, Yifan Hou, Zhenjia Xu, Linxi Fan, Manuela Veloso, and Shuran Song. Dexumi: Using human hand as the universal manipulation interface for dexterous manipulation. *arXiv preprint arXiv:2505.21864*, 2025.
- [85] Yinzheng Xu, Weikang Wan, Jialiang Zhang, Haoran Liu, Zikang Shan, Hao Shen, Ruicheng Wang, Haoran Geng,

- Yijia Weng, Jiayi Chen, et al. Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy. *arXiv preprint arXiv:2303.00938*, 2023.
- [86] Jianglong Ye, Keyi Wang, Chengjing Yuan, Ruihan Yang, Yiquan Li, Jiyue Zhu, Yuzhe Qin, Xueyan Zou, and Xiaolong Wang. Dex1b: Learning with 1b demonstrations for dexterous manipulation. In *Robotics: Science and Systems (RSS)*, 2025.
- [87] Brent Yi, Chung Min Kim, Justin Kerr, Gina Wu, Rebecca Feng, Anthony Zhang, Jonas Kulhanek, Hong Suk Choi, Yi Ma, Matthew Tancik, and Angjoo Kanazawa. Viser: Imperative, web-based 3d visualization in python. *arXiv preprint arXiv:2507.22885*, 2025.
- [88] Zhao-Heng Yin, Changhao Wang, Luis Pineda, Francois Hogan, Krishna Bodduluri, Akash Sharma, Patrick Lancaster, Ishita Prasad, Mrinal Kalakrishnan, Jitendra Malik, Mike Lambeta, Tingfan Wu, Pieter Abbeel, and Mustafa Mukadam. Dexteritygen: Foundation controller for unprecedented dexterity, 2025. URL <https://arxiv.org/abs/2502.04307>.
- [89] Kevin Zakka, Baruch Tabanpour, Qiayuan Liao, Mustafa Haiderbhai, Samuel Holt, Jing Yuan Luo, Arthur Allshire, Erik Frey, Koushil Sreenath, Lueder A. Kahrs, Carlo Sferrazza, Yuval Tassa, and Pieter Abbeel. Mujoco playground: An open-source framework for gpu-accelerated robot learning and sim-to-real transfer., 2025. URL https://github.com/google-deepmind/mujoco_playground.
- [90] Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. In *Robotics: Science and Systems*, 2024.
- [91] Jialiang Zhang, Haoran Liu, Danshi Li, XinQiang Yu, Haoran Geng, Yufei Ding, Jiayi Chen, and He Wang. Dexgraspnet 2.0: Learning generative dexterous grasping in large-scale synthetic cluttered scenes. In *8th Annual Conference on Robot Learning*.
- [92] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [93] Yifeng Zhu, Arisrei Lim, Peter Stone, and Yuke Zhu. Vision-based manipulation from single human video with open-world object graphs. *arXiv preprint arXiv:2405.20321*, 2024.