

mimic-one: a Scalable Model Recipe for General Purpose Robot Dexterity

Anonymous Author(s)

Affiliation

Address

email

Abstract: We present a diffusion-based model recipe for real-world control of a highly dexterous humanoid robotic hand, designed for sample-efficient learning and smooth fine-motor action inference. Our system features a newly designed 16-DoF tendon-driven hand, equipped with wide angle wrist cameras and mounted on a Franka Emika Panda arm. We develop a versatile teleoperation pipeline and data collection protocol using both glove-based and VR interfaces, enabling high-quality data collection across diverse tasks such as pick and place, item sorting and assembly insertion. Leveraging high-frequency generative control, we train end-to-end policies from raw sensory inputs, enabling smooth, self-correcting motions in complex manipulation scenarios. Real-world evaluations demonstrate up to 93.3% out of distribution success rates, with up to a +33.3% performance boost due to emergent self-correcting behaviors, while also revealing scaling trends in policy performance. Our results advance the state-of-the-art in dexterous robotic manipulation through a fully integrated, practical approach to hardware, learning, and real-world deployment.

Keywords: Dexterity, Manipulation, Self-Correction

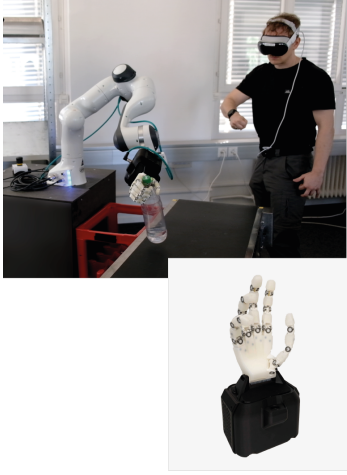
1 Introduction

Robotic dexterity—the ability to manipulate a wide range of objects with precision, speed, and adaptability—remains one of the grand challenges in robotics. Although recent advances in learning-based control, model architectures, and robotic hardware have enabled significant progress in domains like locomotion and grasping, achieving general-purpose, real-world dexterous manipulation still demands a holistic integration of perception, data, hardware, and control. Human hands, through evolution, have become masterful tools for contact-rich manipulation. Can we equip robots with similarly versatile capabilities?

In this work, we introduce **mimic-one**, a scalable recipe for real-world, general-purpose robotic dexterity based on high-frequency generative models trained end-to-end with imitation learning. Our system centers around a newly developed 16 degree-of-freedom (DoF) humanoid robotic hand, capable of fine-grained, tendon-driven control. The hand is mounted on a 7-DoF Franka Emika Panda arm and equipped with wide-FOV wrist cameras to provide rich, low-latency feedback from diverse viewpoints. The entire system is optimized for smooth inference and sample efficiency and is powered by a diffusion-based control policy trained from raw image and proprioceptive observations.

To collect diverse high-quality demonstrations, we design an extensible teleoperation system that supports both glove-based and Apple Vision Pro-based control interfaces. These enable rapid data acquisition across a variety of real-world tasks—including pick-and-place of deformable objects, precise insertions, and complex object sorting—all selected to challenge the limits of contemporary manipulation systems. Crucially, our methodology includes a robust data collection protocol to identify and mitigate failure modes through targeted self-correction data, which we show is vital for improving robustness and policy generalization.

Intuitive teleoperation
for dexterous robotic hands



Data collection and model training



Task variety and self-correction
capabilities

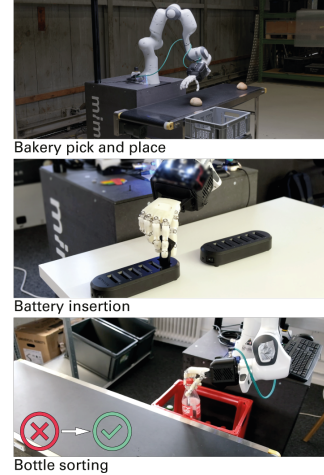


Figure 1: **mimic-one** is a scalable model recipe and data collection protocol for general purpose dexterous manipulation. Our system features a newly designed 16-DoF tendon-driven hand, a 7-DoF Franka Emika Panda robot arm, and a VR teleoperation system for data collection. We showcase the framework across difficult, dynamic real-world manipulation scenarios, which highlight the system’s high frequency fine-motor skills, visual generalization, and ability to recover from failures.

40 We demonstrate that mimic-one policies are not only capable of completing these tasks with high
 41 success rates, but also natively exhibit recovery from failed grasps and adaptive corrections. Our
 42 ablation studies reveal key data curation and representation decisions that dramatically impact per-
 43 formance. Importantly, we uncover scaling trends that highlight how policy generalization and
 44 reliability improve downstream of data diversity and curation.

45 2 Methodology

46 2.1 The mimic dexterous robotic hand

47 With the ultimate goal of demonstrating highly dexterous, human-like manipulation, we developed a
 48 novel “mimic” dexterous robotic hand (Fig. 4.c). The hand features 20 joints with 16 DoF, actuated
 49 by an antagonistic tendon-driven mechanism with slack compensation. The rigid articulated bodies
 50 are padded with a soft silicone skin at relevant contact surfaces. The system successfully demon-
 51 strated lifting capacity for payloads exceeding 7 kg with 5-fingered power grasps. The hand also
 52 features two wide-angle RGB wrist cameras (Fig. 4.b), with the aim of providing rich, low-latency
 53 feedback from diverse viewpoints. An emphasis is placed on abduction and adduction for all fingers
 54 and the anatomically accurate opposability of the thumb. This design choice allows an intuitive user
 55 experience for teleoperation, narrowing the embodiment gap in between the robot and the human
 56 operator. Other, less relevant DoF of human biomechanics are simplified with joint coupling and
 57 underactuated mechanisms to reduce system complexity and weight.

58 2.2 Robot setup and teleoperation

59 Our robot station setup consists of a Franka Emika Panda robot arm, a mounted mimic dexterous
 60 hand, and one additional external “workspace” RGB camera (Fig. 4.b). Demonstration data is col-
 61 lected via teleoperation using one of two methods: 1. Manus gloves for hand motion capture and
 62 SpaceMouse control for wrist poses; or 2. Apple Vision Pro hand and wrist tracking (Fig. 4.a). The
 63 Franka Emika Panda robot arm runs a low-level Cartesian impedance controller from [1], following
 64 a high-level target pose commanded by the teleoperator. We apply a retargeting algorithm to the
 65 hand finger joints, as detailed in App. B.1.

66 2.3 Model recipe and data collection protocol

67 Our recipe relies on several key design choices for state and action representation that significantly
68 boost performance and generalization for dexterous policies. These design choices mainly involve:
69 **1. Using cartesian end effector poses** as states and actions; **2. Using a “relative” end effector**
70 **state and action representation**; **3. Using absolute hand joint angles** state and actions; **4. Using**
71 **6D rotations**. We detail the model recipe choices in further detail in App. C.1.

72 At the same time, we observe that, irrespective of scale, obtaining high success rates with imita-
73 tion learning policies is facilitated when following a specific **data collection protocol**, involving:
74 **1. Data collection** with randomized object positions, distractors and backgrounds; **2. Data labeling**
75 for success/failure; **3. Data curation**, filtering for high-quality trajectories; **4. Policy training**; **5. It-**
76 **erative self-correction data collection** for the common policy failure modes. We further detail the
77 protocol in App. C.2.

78 2.4 Policy architecture

79 Our goal is to train a policy $\pi(a_{t:t+H_a} | o_{t-H_o:t})$ that maps from a sequence of observations $o_{t-H_o:t}$ to
80 future actions $a_{t:t+H_a}$ for the dexterous hand plus arm. Our observations $o_i = (l_i, I_i)$ are composed
81 of low-dimensional sensory readings l_i , such as proprioceptive robot joint angles and end effector
82 poses, together with RGB images I_i from the cameras mounted on the workstation and the wrist. We
83 predict an entire “action chunk” of length H_a , as this makes it easier to learn and generate smooth,
84 high-frequency motions that are temporally consistent and exhibit lower compounding error.

85 Our chosen approach for simple but effective end-to-end high-frequency generative control is to use
86 a UNet Diffusion Policy architecture [2], receiving as conditioning input an observation horizon with
87 three RGB images from the wrist-mounted and overhead cameras and proprioceptive state inputs
88 (end-effector pose, joint angles). The low-dimensional proprioceptive state inputs are projected to
89 the embedding dimension with a linear layer, while the RGB images are fed into CLIP [3] pre-
90 trained ViT-B/16 [4] encoders with duplicated weights for each camera stream. The embeddings are
91 concatenated and passed to the UNet Diffusion Policy backbone, which samples the action chunks
92 by denoising a vector of sampled Gaussian noise (Fig. 5.c).

93 3 Results

94 We evaluate our model recipe on three challenging dexterous manipulation benchmarks designed
95 to test generalization, precision, and self-correction capabilities (Fig. 6). These tasks intentionally
96 move beyond simple pick-and-place scenarios to probe the limits of dexterous control.

97 **General pick-and-place (Bread Picking):** This task involves picking deformable bread loaves from
98 varied surfaces (static table, moving conveyor belt) and placing them into a box. Key challenges
99 include handling the deformable object gently and adapting to randomized loaf positions and diverse
100 surfaces (varied table colors, conveyor). **Evaluation:** Assessed on unseen table/background settings
101 with arbitrary loaf starting positions. Success requires the loaf to be gently placed into the target
102 container, with the policy allowed to self-correct errors.

103 **Complex placement (Bottle Sorting):** This task mimics a recycling scenario, requiring the robot to
104 grasp plastic bottles (potentially slippery, empty or partially filled) sideways on a conveyor belt and
105 insert them precisely into a bottle rack with variable slot occupancy. Challenges include achieving
106 a stable sideways grasp and accurate placement into a constrained slot. **Evaluation:** Performed on
107 the training conveyor belt but with unseen background settings, arbitrary bottle positions, and ran-
108 domized rack occupancy. Success requires placing a bottle into an empty slot, permitting recovery
109 from initial misplacements (e.g., re-orienting a poorly slotted bottle).

110 **Precise insertion (Battery Insertion):** This task requires high precision, involving picking a battery
111 from one rack, transporting it, and inserting it fully into a specific slot in a second rack, including
112 a final “punch” motion to ensure it’s connected. The main challenges lie in the precise alignment

needed for insertion (which depends on the battery’s pose in hand, not just the end-effector pose) and executing the dynamic push. **Evaluation:** Conducted on an unseen table/background with randomized rack positions. Success requires the battery to be picked, correctly inserted into the target slot, and fully ”punched” in.

Quantitative evaluation. We assessed policy performance by varying the amount of training data (Fig. 2). Success rates improved significantly with more data across all tasks: Bread Pick (22.5%@20% data to 93.3%@100%), Bottle Sort (25.0%@50% to 75.0%@100%), and Battery Insertion (12.5%@50% to 37.5%@100%). These results, achieved with relatively modest dataset sizes, highlight the effectiveness of our data collection protocol.

Crucially, incorporating self-correction trajectories markedly improved robustness (Fig. 2, dashed vs. solid bars). Success rates increased substantially when allowing for recovery behaviors: Bread Pick (+26.6%), Bottle Sort (+33.3%), and Battery Insertion (+25.0%), demonstrating the value of targeted failure recovery data.

Ablation study. We validated key components of the **mimic-one** recipe via ablations on the Bread Pick-and-Place task (Fig. 3). Using sub-optimal action representations (i, ii), limiting data diversity (iii), omitting data filtering (iv), or excluding self-correction data (v) all led to significantly lower success rates (ranging from 5.0% to 56.7%) compared to the full recipe (vi, 93.3% success rate). This confirms the importance of each element in our proposed methodology, particularly the relative action representation with the correct base frame, data diversity/curation, and inclusion of self-correction trajectories.

4 Conclusion

We presented **mimic-one**, a scalable recipe for achieving general-purpose robotic dexterity using a novel 16-DoF humanoid hand and a diffusion-based generative control policy. Our approach leverages efficient teleoperation for data collection, incorporating a crucial protocol for targeted self-correction data, and relies on specific design choices like relative Cartesian action representations for effective end-to-end imitation learning. Experiments demonstrated the system’s capability on challenging real-world tasks, showcasing high success rates, smooth control, generalization, and emerging self-corrective behaviors. Ablations and scaling results validated our methodology’s key components. The **mimic-one** framework offers a practical and reproducible path towards bridging the gap between generative models and the demands of robust, real-world dexterous manipulation.

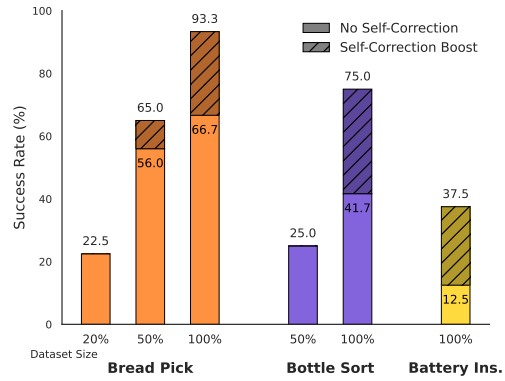


Figure 2: Task success rates vs. dataset scale. Dashed bars include self-corrected successes; solid bars count any error as failure.

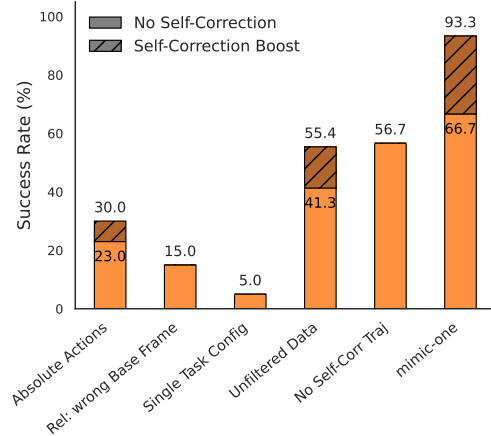


Figure 3: Ablation study on Bread Pick-and-Place, comparing the full **mimic-one** recipe (vi) against variants removing key components: (i) Absolute actions, (ii) Incorrect relative base frame, (iii) Single task configuration, (iv) Unfiltered data, and (v) No self-correction data.

Appendix

A Related Work

Dexterous manipulation hardware. Robotic hand design has progressed toward replicating human-level dexterity via either high-fidelity anthropomorphic mechanisms or simplified, task-driven alternatives. Foundational systems such as the humanoid hand in [5], DLR/HIT Hand II [6], and the DLR Hand-Arm System [7] emphasized high-DoF kinematics and multisensory integration. Later designs like [8], [9], and [10] incorporated biomimetic actuation and compliant skin, targeting contact-rich interaction. Single-print fabrication of monolithic hands enhanced rapid prototyping [11]. However, these designs often lack integrated perception and scalability/reliability for learning. In contrast, our design balances dexterity and simplicity through a tendon-driven, 16-DoF architecture with wide-FOV wrist cameras, engineered specifically for scalable imitation learning.

Imitation learning and generative control. End-to-end imitation learning approaches have recently emerged as a promising paradigm for training robot manipulation policies, enabling robots to learn complex skills directly from demonstration data. This surge has been driven by the adoption of generative modeling techniques in robotics, typically applied to large-scale datasets, predominantly involving two-finger gripper platforms. End-to-end autoregressive Vision-Language-Action models (VLAs) [12, 13, 14, 15] pioneered the application of autoregressive transformers for action generation, demonstrating the ability to learn diverse skills with a single model, at the cost of slow inference and suboptimal action tokenization. To benefit from the strengths of autoregressive VLA models for the high-frequency, dexterous manipulation case, tokenization schemes alternative to naive action binning have been proposed (e.g., FAST [16]).

Alternatively to autoregression, learning high-frequency control policies from demonstrations with generative model techniques can be achieved with the use of action chunking. The first of such approaches was the VAE-based ACT [17]. Further development in methods for fast inference at high control frequency leveraged Diffusion Models [18, 19] or Flow Matching [20] over an action chunk, such as Diffusion Policy [2] or PointFlowMatch [21]. π_0 [22] illustrates that flow matching-based models can be successfully scaled up and combined with Vision-Language Model (VLM) backbones to enable diverse, language instruction-driven robot behaviors.

Generative control for scalable (humanoid) dexterity. Several systems have been introduced to enable scalable data collection and policy training for robots of varying dexterity. Universal Manipulation Interface (UMI) [23] introduced a low-cost, portable data collection and policy interface for two-finger gripper robots, agnostic to robotic arms and compatible with static and mobile robot systems [24]. Follow-up works [25] have investigated the data quantity and diversity scaling laws of the UMI Diffusion Policy. Other works [26] have applied the Diffusion Policy model to lower-DoF humanoid dexterous hands. Several works have demonstrated the benefits of pre-training on human video data [27, 28, 29].

B Robot Setup

B.1 Robot hand pose retargeting

Agnostic to the hand capture method used for data collection, our method requires re-targeting of human hand poses to robot hand joint angles. To do so, we employ the key-vector-based re-targeting method from [30, 27] (Fig. 4.d), in which the “keyvectors” v_i^h and v_i^r between palms and fingertips for the human and robot hand are used to compute an energy loss function $E((\beta_h, \theta_h), q) = \sum_{i=1}^{15} \|v_i^h - (c_i \cdot v_i^r)\|$, which minimizes the distance between human hand poses (β, θ) and robot hand poses q , with c_i being a scale parameter.

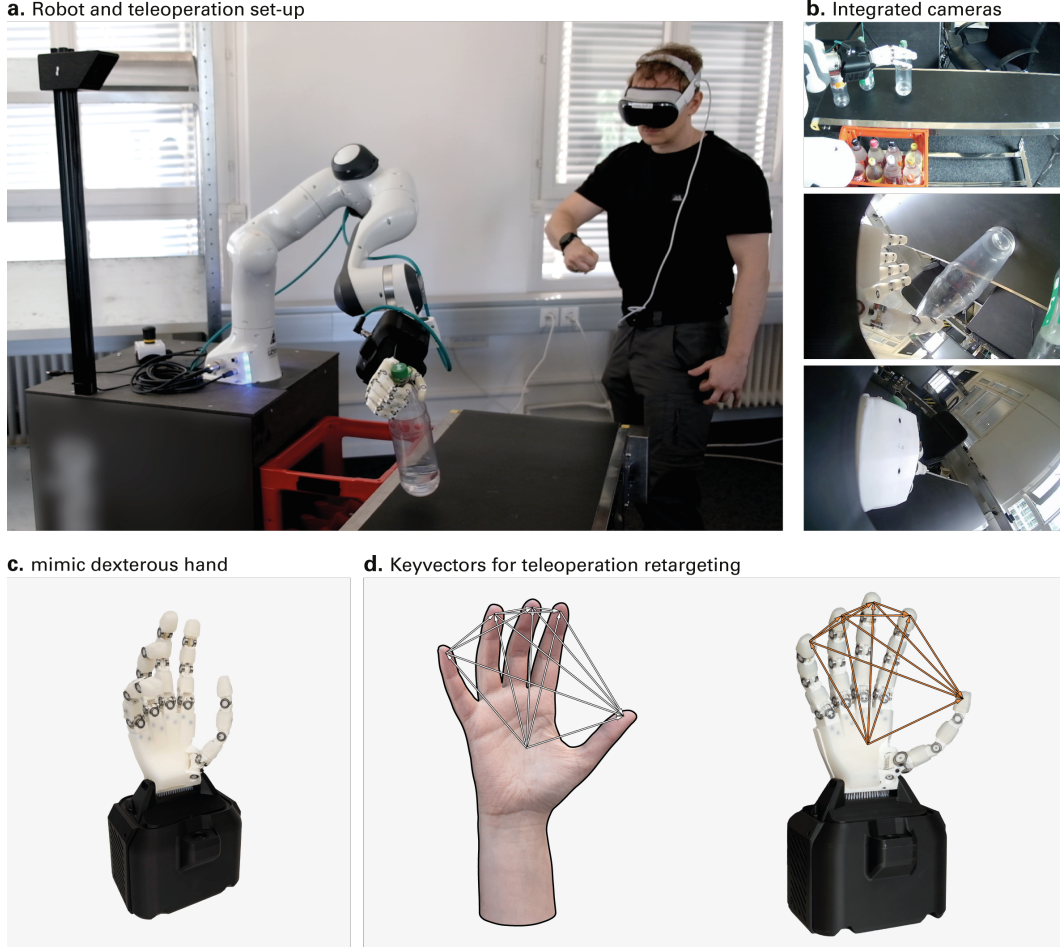


Figure 4: System overview and teleoperation setup. (a) On the left, the robot station, with a mimic robotic hand mounted on a 7-DoF Franka Emika Panda arm. On the right, a data collector using an Apple Vision Pro for teleoperation. (b) Synchronized visual inputs from wrist-mounted fisheye cameras (below and above the wrist) and an external overhead camera. (c) mimic hand, featuring 16 DoFs with soft skin contact surfaces and tendon actuation. (d) Visualization of keyvector representations used for retargeting human hand poses to robot joint configurations during teleoperation.

C Model recipe and Data Collection Protocol

C.1 Key ingredients for generative control for highly dexterous manipulation

Our recipe relies on several key design choices for state and action representation that significantly boost performance and generalization for dexterous policies:

Cartesian target end effector pose as action. Following [23], we use the Cartesian target pose for the end effector as the action, commanded to the low-level impedance controller. This improves robustness to object position changes and promotes arm-agnostic policies. Crucially, the action is the *target* pose from the teleoperator, not the currently observed pose. The policy essentially learns to act as a drop-in replacement for the teleoperator.

Cartesian proprioceptive end effector pose as state. The end effector state uses the Cartesian pose derived from robot proprioception.

“Relative” end effector action representation. Choosing the right reference frame for Cartesian end effector actions is crucial. Common options include **Absolute** (poses in a global frame), **Delta** (each pose relative to the immediately preceding one), and **Relative** (all poses in the action chunk

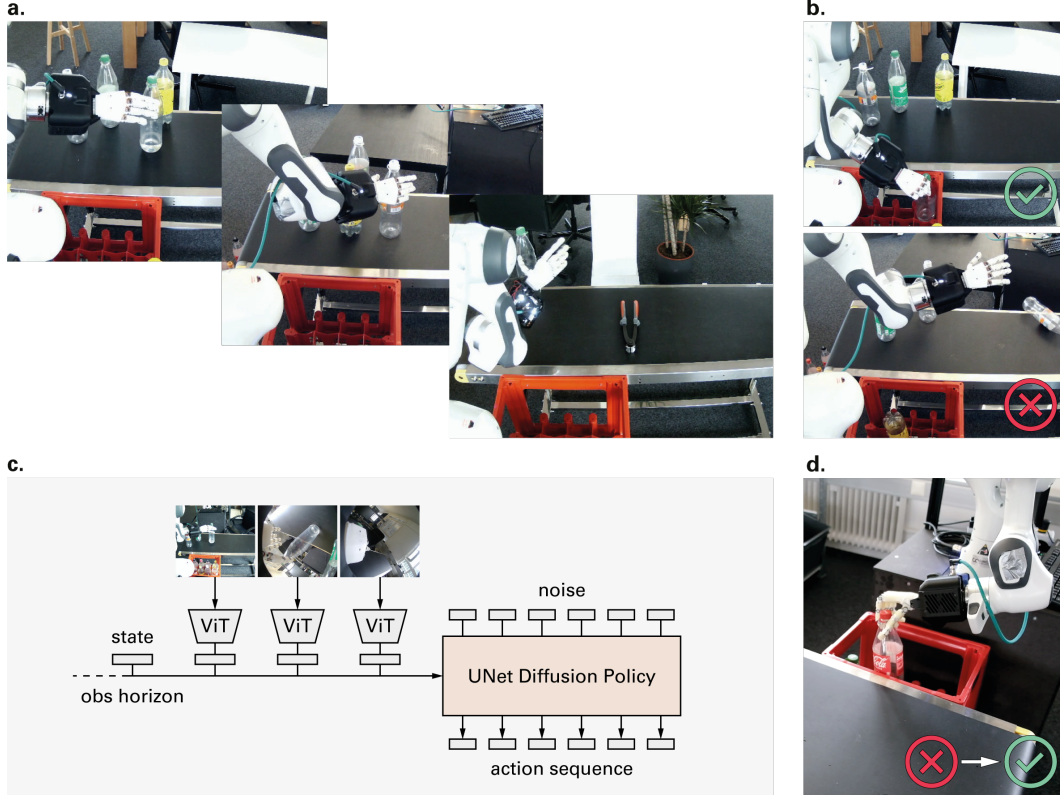


Figure 5: The **mimic-one** data collection protocol and policy recipe. (a) Teleoperation data collection. Variation in the setting involve randomize object positions, robot starting pose, “task config”, and distractors. (b) Data labeling and filtering (removing failures, non-stable grasps and suboptimal completions). (c) The diffusion policy model architecture, receiving as conditioning input an observation horizon with encoded RGB images and proprioceptive state inputs, predicting the future action chunk. (d) Self-correction trajectory collection based on common failure modes. Robot and workspace are positioned in a failure state, and correction episodes are collected.

relative to a single ‘base pose’) [17, 23]. For dexterous manipulation requiring adaptation to varying object positions, we find the **Relative** representation superior for generalization. It makes the planned trajectory segment largely independent of the robot’s absolute starting position. A critical implementation detail for correctness and stability is selecting the correct base pose: it must be the *last observed proprioceptive end effector pose* (i.e., the final state in the input observation horizon $o_{t-H_o:t}$). Using the first commanded target pose from the current action chunk as the base frame would mismatch the policy’s conditioning at inference time, which only includes observed states.

“Relative” end effector state representation. Similarly, the end effector state horizon is represented relative to the last observed proprioceptive end effector pose.

Absolute hand joint angles as action and state. For the dexterous hand, we use absolute joint angles for both state and actions. This allows the policy to learn meaningful grasp configurations directly, while this information would be lost if using relative angles.

6D rotations. All end effector rotations are represented using the 6D upper-triangular rotation representation [31], which is well-suited for neural network learning due to preserving continuity.

C.2 A data collection protocol to boost success rates

Existing works often focus on either limited, in-distribution demonstrations, or attempt to achieve generalization and high task success rates by simply scaling up data collection to a very high degree. We observe that irrespective of scale, obtaining high success rates for general purpose tasks with

approaches purely based on Imitation Learning is strongly facilitated when following a specific **data collection protocol**, structured as follows:

(1) Data collection: Initially, data collection is performed with teleoperation, attempting to perform the entire task successfully in each episode (Fig. 5.a). While doing this, one should adhere to the following: **(1.1)** For each different episode, randomize the starting pose of the robot within the workspace and the position of any movable task object. **(1.2)** Add distractor objects (objects not involved in the task) in the workspace in approximately 50% of the episodes. **(1.3)** We define a “Task config” as the overall setting in which an episode is collected. Specifically, we consider as a different task config a situation in which one of **(a)** table color or type; **(b)** Background; **(c)** Lighting conditions; has changed. Following the observations from [25], we apply the scaling-law optimal ratio of different task configs to the overall number of collected episodes, which means that we change task config approximately every 100 episodes collected.

(2) Data labeling: Once the first data collection step is completed, the original data collector must manually label each collected episode as either a success or a failure, depending on whether the task is completed (Fig. 5.b).

(3) Data curation: A different person conducts further filtering, eliminating trajectories involving non-stable grasps, erratic motions, and ambiguous task completions.

(4) Policy training: A first policy is trained and evaluated in-distribution. Common failure modes are identified and documented.

(5) Repeat — collect self-correction trajectories: based on the common failure modes. The robot and workspace are positioned in a failure state before starting the episode collection (Fig. 5.d). After a new batch of self-correction trajectories has been collected, one can begin training again.

Overall, we collect episodes over 3 main task families. We select tasks based on the desire to test for: 1. Overall capability for pick and place in diverse, real-world settings. 2. Completion of dexterous tasks that are not easily achievable with two-finger gripper solutions. 3. Precision in tasks such as assembly and insertion. 4. Illustration of self-correction behaviors.

D Training Details

Observation horizon. Each training sample consists of an observation horizon $H_o = 2$ and an action chunk of $H_a = 48$ time steps, corresponding to 3.2 s of future prediction at a control rate of 15 Hz.

Model architecture. The image encoders are CLIP [3] ViT-B/16 [4] encoders, kept unfrozen at training time. The action denoising UNet uses the architecture from Diffusion Policy [2], with an embedding dimension of 128, kernel size = 5, number of groups = 8, and a training-time number of inference step of 16. The UNet is conditioned with a concatenated embedding from the observation horizon. The visual embeddings from RGB camera observations are taken from the CLS token of the ViT encoder. As for the low-dimensional observation, they are not projected, but are directly concatenated with the visual embeddings, before being used for conditioning the UNet.

Observation and Action Representations The representation of the robot arm’s end effector Cartesian poses, for both observed states and predicted actions, incorporates specific strategies to enhance learning and generalization. Let $P_k = (T_k, R_k)$ denote a Cartesian pose at time step k , where $T_k \in \mathbb{R}^3$ is the translation and $R_k \in SO(3)$ is the rotation matrix. The current time step is t .

- **“Relative” End Effector Pose Representation.** A key aspect of our methodology is the use of a “Relative” representation for end effector poses. This means that sequences of poses are expressed relative to a common reference frame, specifically the most recent

Table 1: Dataset statistics.

Task Name	# Success Eps	# Filtered Eps	# Task Configs
Bread Pick	1830	550	5
Bottle Sort	2420	882	8
Battery Insertion	1192	528	5

283 observed proprioceptive pose of the end effector. Let $P_t^{obs} = (T_t^{obs}, R_t^{obs})$ be the latest
 284 observed proprioceptive end effector pose at the current time t . This P_t^{obs} serves as the
 285 “base pose”.

- **State Representation:** The input observation horizon for the end effector poses, denoted as $\mathcal{O}_{ee} = \{P_{t-H_o}^{obs}, \dots, P_{t-1}^{obs}, P_t^{obs}\}$, consists of $H_o + 1$ observed poses. Each pose P_k^{obs} in this sequence (for $k \in [t - H_o, t]$) is transformed into a relative pose, $P_k^{state_rel}$, with respect to the base pose P_t^{obs} :

$$P_k^{state_rel} = (T_k^{state_rel}, R_k^{state_rel})$$

where

$$\begin{aligned} T_k^{state_rel} &= (R_t^{obs})^T (T_k^{obs} - T_t^{obs}) \\ R_k^{state_rel} &= (R_t^{obs})^T R_k^{obs} \end{aligned}$$

286 The sequence $\{P_{t-H_o}^{state_rel}, \dots, P_t^{state_rel}\}$ (where $P_t^{state_rel}$ is the identity pose $(0, I)$)
 287 forms the state input to the policy concerning the end effector.

- **Action Representation:** The policy predicts an “action chunk,” which is a sequence of H_a future target end effector poses. These are also represented relative to the same base pose P_t^{obs} . The policy outputs:

$$\mathcal{A}_{pred} = \{P_t^{act_rel}, P_{t+1}^{act_rel}, \dots, P_{t+H_a-1}^{act_rel}\}$$

Each $P_j^{act_rel} = (T_j^{act_rel}, R_j^{act_rel})$ for $j \in [t, t + H_a - 1]$ is a pose relative to P_t^{obs} . To obtain the absolute target pose $P_j^{target} = (T_j^{target}, R_j^{target})$ that is commanded to the low-level robot controller, this relative pose is combined with the base pose P_t^{obs} :

$$\begin{aligned} T_j^{target} &= T_t^{obs} + R_t^{obs} T_j^{act_rel} \\ R_j^{target} &= R_t^{obs} R_j^{act_rel} \end{aligned}$$

- **6D Rotation Representation** All 3D rotations $R \in SO(3)$ for the end effector poses, whether in the observed states or the predicted actions (i.e., $R_k^{state_rel}$ and $R_j^{act_rel}$), are represented using a continuous 6D format as proposed by [31]. If a rotation matrix R is given by its column vectors:

$$R = [\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3], \quad \text{where } \mathbf{r}_i \in \mathbb{R}^3$$

The 6D representation, $\mathbf{R}_{6D} \in \mathbb{R}^6$, is formed by concatenating the first two column vectors:

$$\mathbf{R}_{6D} = \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix}$$

288 This representation avoids discontinuities and gimbal lock issues associated with other rota-
 289 tion parametrizations like Euler angles or quaternions when used directly in neural network
 290 outputs.

291 **Data augmentation.** We apply random cropping with a 0.95 ratio to the RGB images, together with
 292 a color jitter random filter. No augmentations are applied to proprioceptive inputs or actions.

293 **Normalization.** All proprioceptive inputs and actions are normalized to fit within limits of +1 and
 294 -1, using statistics computed over the training set. RGB images are normalized according to the
 295 ImageNet normalization used for the ViT encoder.

296 **Training.** Policies are trained with the standard denoising diffusion loss (mean squared error be-
 297 tween denoised and ground truth action chunk), using the AdamW optimizer with a learning rate of
 298 3×10^{-4} , weight decay of 1×10^{-6} , a cosine decay schedule and a warmup of 2000 steps. Training
 299 runs for 120 epochs over the task dataset, which has been observed to be enough for all individual
 300 test tasks.

301 **Inference.** During inference, the policy runs in real-time at 15Hz, predicting new action chunks with
 302 a 16-step DDIM sampler. We use a dynamic action scheduler, playing back actions when they are
 303 ready, according to their scheduled timestamp. We re-compute the action chunk before its complete
 304 execution, after the execution of the 10th action of 48.

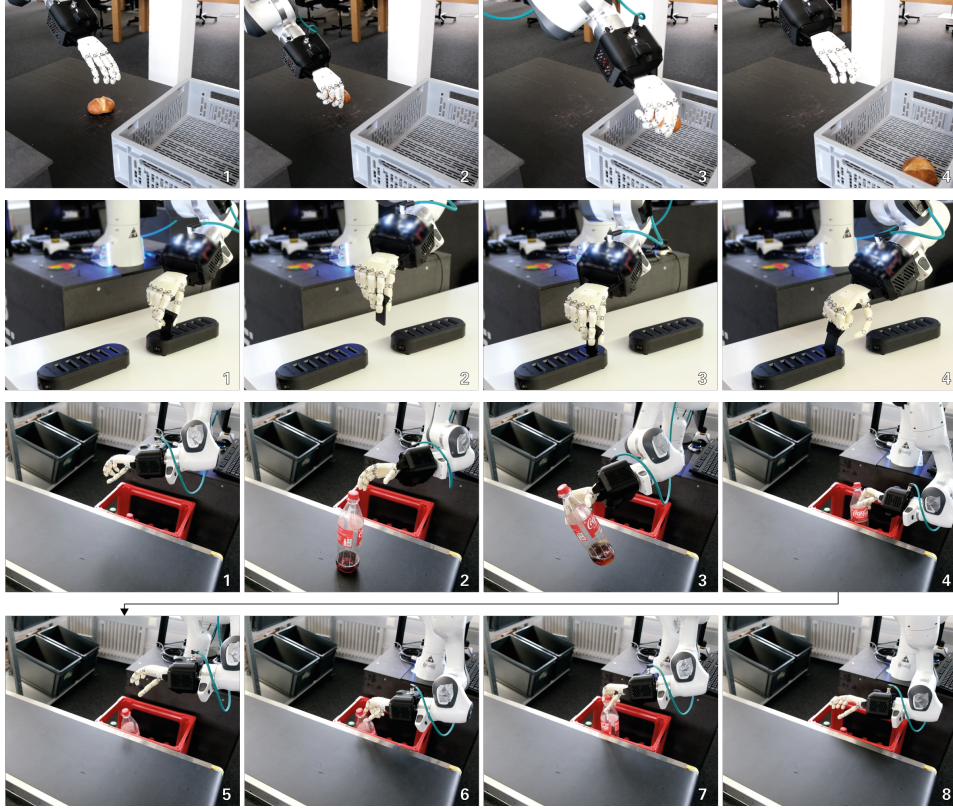


Figure 6: Representative rollout sequences across three benchmark tasks. The mimic-one policy demonstrates smooth, self-correcting behaviors over **bread pick-and-place**, **battery insertion** and **bottle sorting**. We thoroughly illustrate self-correcting behavior for the bottle sorting task: the hand re-oriens the bottle post-grasp to achieve successful insertion.

E Task Success Criteria

We define success criteria for each task as follows:

- **Bread Pick-and-Place:** The bread loaf must be fully inside the target container and not crushed or dropped during placement.
- **Bottle Sorting:** The bottle must be inserted into a free slot in the rack and remain upright. Corrections after misplacement are permitted.
- **Battery Insertion:** The battery must be inserted into the designated slot and visually confirmed to be flush with the rack (fully punched in).

F Failure Mode Taxonomy

To facilitate targeted self-correction collection, we categorize failure modes as follows:

1. **Unstable grasp:** Object is grasped but not securely held; results in drop or slippage during placement.
2. **No grasp:** Missed grasp attempts due to poor object localization or approach trajectory.
3. **Misalignment:** Object is correctly picked but misaligned for target insertion or placement.

These failure modes can happen either during data collection itself, or during policy inference. During data collection, the protocol is to stop the recording of the current episode, and start a new recording involving performing a correction of the current error. If happening during inference, a

descriptive annotation of the error is made and used to create custom “reset scenes” during the next round of data collection, recording custom recovery demonstrations.

G Limitations

While mimic-one demonstrates significant progress in real-world dexterous manipulation, several limitations should be acknowledged:

- **Dependence on imitation learning:** The core approach relies on imitation learning, which inherently limits the policy’s capabilities to the quality, diversity, and optimality of the human demonstrations. Performance is fundamentally upper-bounded by the teleoperator’s skill and the fidelity of the teleoperation system.
- **Data collection scalability and cost:** Generating the required demonstration data, including the crucial self-correction trajectories, necessitates access to the specific robot hardware (mimic hand, Franka arm, cameras) and human operator time for teleoperation and data curation.
- **Single-task focus in evaluation:** Although the proposed recipe is suitable for multi-task learning, the current empirical validation focuses on training and evaluating policies for individual task families. Training separate policies may also be less data-efficient than multi-task learning frameworks that promote representation sharing.
- **Hardware specificity:** The presented policies are trained specifically for the 16-DoF mimic hand and its sensor configuration (wrist cameras, proprioception). Transferring these policies directly to different robotic hands, would likely require retraining or domain adaptation. The system currently lacks tactile sensing, which could further enhance robustness and enable more contact-rich manipulation tasks.

References

- [1] J. Luo, Z. Hu, C. Xu, Y. L. Tan, J. Berg, A. Sharma, S. Schaal, C. Finn, A. Gupta, and S. Levine. SERL: A Software Suite for Sample-Efficient Robotic Reinforcement Learning, Mar. 2025. URL <http://arxiv.org/abs/2401.16013>. arXiv:2401.16013 [cs].
- [2] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion, Mar. 2024. URL <http://arxiv.org/abs/2303.04137>. arXiv:2303.04137 [cs].
- [3] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020 [cs]*, Feb. 2021. URL <http://arxiv.org/abs/2103.00020>. arXiv: 2103.00020.
- [4] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, June 2021. URL <http://arxiv.org/abs/2010.11929>. arXiv:2010.11929 [cs].
- [5] H. Liu, K. Wu, P. Meusel, G. Hirzinger, M. Jin, Y. Liu, S. Fan, T. Lan, and Z. Chen. A dexterous humanoid five-fingered robotic hand. In *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, pages 371–376, Aug. 2008. doi:10.1109/ROMAN.2008.4600694. URL <https://ieeexplore.ieee.org/document/4600694>. ISSN: 1944-9437.
- [6] H. Liu, K. Wu, P. Meusel, N. Seitz, G. Hirzinger, M. Jin, Y. Liu, S. Fan, T. Lan, and Z. Chen. Multisensory five-finger dexterous hand: The DLR/HIT Hand II. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3692–3697, Sept. 2008. doi:10.

- 1109/IROS.2008.4650624. URL <https://ieeexplore.ieee.org/document/4650624>.
ISSN: 2153-0866.
- [7] M. Grebenstein, A. Albu-Schäffer, T. Bahls, M. Chalon, O. Eiberger, W. Friedl, R. Gruber, S. Haddadin, U. Hagn, R. Haslinger, H. Höppner, S. Jörg, M. Nickl, A. Nothhelfer, F. Petit, J. Reill, N. Seitz, T. Wimböck, S. Wolf, T. Wüsthoff, and G. Hirzinger. The DLR hand arm system. In *2011 IEEE International Conference on Robotics and Automation*, pages 3175–3182, May 2011. doi:10.1109/ICRA.2011.5980371. URL <https://ieeexplore.ieee.org/abstract/document/5980371>. ISSN: 1050-4729.
- [8] Z. Xu and E. Todorov. Design of a highly biomimetic anthropomorphic robotic hand towards artificial limb regeneration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3485–3492, May 2016. doi:10.1109/ICRA.2016.7487528. URL <https://ieeexplore.ieee.org/abstract/document/7487528>.
- [9] U. Kim, D. Jung, H. Jeong, J. Park, H.-M. Jung, J. Cheong, H. R. Choi, H. Do, and C. Park. Integrated linkage-driven dexterous anthropomorphic robotic hand. *Nature Communications*, 12(1):7177, Dec. 2021. ISSN 2041-1723. doi:10.1038/s41467-021-27261-0. URL <https://www.nature.com/articles/s41467-021-27261-0>. Publisher: Nature Publishing Group.
- [10] Y. Toshimitsu, B. Forrai, B. G. Cangan, U. Steger, M. Knecht, S. Weirich, and R. K. Katzschmann. Getting the Ball Rolling: Learning a Dexterous Policy for a Biomimetic Tendon-Driven Hand with Rolling Contact Joints. In *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, pages 1–7, Dec. 2023. doi:10.1109/Humanoids57100.2023.10375231. URL <http://arxiv.org/abs/2308.02453>. arXiv:2308.02453 [cs].
- [11] T. J. K. Buchner, S. Rogler, S. Weirich, Y. Armati, B. G. Cangan, J. Ramos, S. T. Twiddy, D. M. Marini, A. Weber, D. Chen, G. Ellson, J. Jacob, W. Zengerle, D. Katalichenko, C. Keny, W. Matusik, and R. K. Katzschmann. Vision-controlled jetting for composite systems and robots. *Nature*, 623(7987):522–530, Nov. 2023. ISSN 1476-4687. doi:10.1038/s41586-023-06684-3. URL <https://www.nature.com/articles/s41586-023-06684-3>. Publisher: Nature Publishing Group.
- [12] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich. RT-1: Robotics Transformer for Real-World Control at Scale, Dec. 2022. URL <http://arxiv.org/abs/2212.06817>. arXiv:2212.06817 [cs].
- [13] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid, Q. Vuong, V. Vanhoucke, H. Tran, R. Soricut, A. Singh, J. Singh, P. Sermanet, P. R. Sanketi, G. Salazar, M. S. Ryoo, K. Reymann, K. Rao, K. Pertsch, I. Mordatch, H. Michalewski, Y. Lu, S. Levine, L. Lee, T.-W. E. Lee, I. Leal, Y. Kuang, D. Kalashnikov, R. Julian, N. J. Joshi, A. Irpan, B. Ichter, J. Hsu, A. Herzog, K. Hausman, K. Gopalakrishnan, C. Fu, P. Florence, C. Finn, K. A. Dubey, D. Driess, T. Ding, K. M. Choromanski, X. Chen, Y. Chebotar, J. Carbajal, N. Brown, A. Brohan, M. G. Arenas, and K. Han. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. In *Proceedings of The 7th Conference on Robot Learning*, pages 2165–2183. PMLR, Dec. 2023. URL <https://proceedings.mlr.press/v229/zitkovich23a.html>. ISSN: 2640-3498.
- [14] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine. Octo: An Open-Source Generalist Robot Policy, May 2024. URL <http://arxiv.org/abs/2405.12213>. arXiv:2405.12213 [cs].

- [15] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. OpenVLA: An Open-Source Vision-Language-Action Model, Sept. 2024. URL <http://arxiv.org/abs/2406.09246>. arXiv:2406.09246 [cs].
- [16] K. Pertsch, K. Stachowicz, B. Ichter, D. Driess, S. Nair, Q. Vuong, O. Mees, C. Finn, and S. Levine. FAST: Efficient Action Tokenization for Vision-Language-Action Models, Jan. 2025. URL <http://arxiv.org/abs/2501.09747>. arXiv:2501.09747 [cs].
- [17] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware, Apr. 2023. URL <http://arxiv.org/abs/2304.13705>. arXiv:2304.13705 [cs].
- [18] Y. Song and S. Ermon. Generative Modeling by Estimating Gradients of the Data Distribution. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/3001ef257407d5a371a96dcd947c7d93-Abstract.html>.
- [19] J. Ho, A. Jain, and P. Abbeel. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html>.
- [20] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow Matching for Generative Modeling, Feb. 2023. URL <http://arxiv.org/abs/2210.02747>. arXiv:2210.02747 [cs].
- [21] E. Chisari, N. Heppert, M. Argus, T. Welschhold, T. Brox, and A. Valada. Learning Robotic Manipulation Policies from Point Clouds with Conditional Flow Matching, Sept. 2024. URL <http://arxiv.org/abs/2409.07343>. arXiv:2409.07343 [cs].
- [22] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky. pi0: A Vision-Language-Action Flow Model for General Robot Control, Nov. 2024. URL <http://arxiv.org/abs/2410.24164>. arXiv:2410.24164 [cs].
- [23] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song. Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots, Mar. 2024. URL <http://arxiv.org/abs/2402.10329>. arXiv:2402.10329 [cs].
- [24] H. Ha, Y. Gao, Z. Fu, J. Tan, and S. Song. UMI on Legs: Making Manipulation Policies Mobile with Manipulation-Centric Whole-body Controllers, July 2024. URL <http://arxiv.org/abs/2407.10353>. arXiv:2407.10353 [cs].
- [25] F. Lin, Y. Hu, P. Sheng, C. Wen, J. You, and Y. Gao. Data Scaling Laws in Imitation Learning for Robotic Manipulation, Oct. 2024. URL <http://arxiv.org/abs/2410.18647>. arXiv:2410.18647 [cs].
- [26] T. Lin, Y. Zhang, Q. Li, H. Qi, B. Yi, S. Levine, and J. Malik. Learning Visuotactile Skills with Two Multifingered Hands, May 2024. URL <http://arxiv.org/abs/2404.16823>. arXiv:2404.16823 [cs].
- [27] K. Shaw, S. Bahl, and D. Pathak. VideoDex: Learning Dexterity from Internet Videos, Dec. 2022. URL <http://arxiv.org/abs/2212.04498>. arXiv:2212.04498 [cs, eess].
- [28] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar. MimicPlay: Long-Horizon Imitation Learning by Watching Human Play, Feb. 2023. URL <http://arxiv.org/abs/2302.12422>. arXiv:2302.12422 [cs].

- 462 [29] S. Kareer, D. Patel, R. Punamiya, P. Mathur, S. Cheng, C. Wang, J. Hoffman, and D. Xu.
463 EgoMimic: Scaling Imitation Learning via Egocentric Video, Oct. 2024. URL [http:
464 //arxiv.org/abs/2410.24221](http://arxiv.org/abs/2410.24221). arXiv:2410.24221 [cs].
- 465 [30] A. Handa, K. V. Wyk, W. Yang, J. Liang, Y.-W. Chao, Q. Wan, S. Birchfield, N. Ratliff, and
466 D. Fox. DexPilot: Vision Based Teleoperation of Dexterous Robotic Hand-Arm System, Oct.
467 2019. URL <http://arxiv.org/abs/1910.03135>. arXiv:1910.03135 [cs].
- 468 [31] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li. On the Continuity of Rotation Repre-
469 sentations in Neural Networks, June 2020. URL <http://arxiv.org/abs/1812.07035>.
470 arXiv:1812.07035 [cs].