

DexH2R: Task-oriented Dexterous Manipulation from Human to Robots

Shuqi Zhao*, Xinghao Zhu*, Yuxin Chen, Chenran Li, Yichen Xie, Xiang Zhang, Mingyu Ding†, Masayoshi Tomizuka

Abstract—Dexterous manipulation remains a challenging problem due to the high-dimensional action space of robot hands. To solve this challenge, recent advances focus on transferring human hand motions to dexterous robot motions. However, the most straightforward method, retargeting, fails to ensure successful task execution due to the ignorance of hand-object interaction and environmental feedback, while other methods combining retargeting with human feedback significantly increase human labor and degrade motion quality. To address this limitation, we propose a novel framework that enhances retargeted motions with a residual policy based on reinforcement learning, allowing robots to autonomously correct action errors without human intervention. Specifically, our approach leverages human demonstrations to provide natural and task-relevant finger motions while incorporating object and environmental information to improve manipulation performance. Extensive experiments in multiple robot hand embodiments demonstrate that our method outperforms baseline approaches by approximately 40% in task success rates on both seen and unseen objects. Furthermore, our approach eliminates the need for real-time human correction, significantly reducing labor costs and facilitating scalable data collection for dexterous manipulation.

Index Terms—Robotics, Dexterous manipulation, Retargeting, Reinforcement learning, Learning from human demonstration

I. INTRODUCTION

HUMANS manipulate a wide range of objects with their hands in both industrial processes and daily life. This extraordinary ability is also desirable for dexterous robot hands to handle various objects in diverse manipulation tasks. However, further development in this field is hindered by the difficulty in the generation of successful trajectories for dexterous manipulation given the high dimensional action space of dexterous robot hands. To solve this challenge, human motions serve as an easily accessible guidance of dexterous manipulation controller in the wild. Therefore, recent advancement has focused on transferring human hand motions to dexterous robot hand motions [1], [2], which leads to ample development in retargeting. As a promising way, retargeting [3], [4] aims at maximizing the kinematics similarity between human hand and robot hand, which is defined by the relationship between their keypoints. However, retargeting alone often fails to yield effective manipulation outcomes in practice due to the ignorance of manipulation results like object information or self-collision of dexterous robot hand. For example, when grasping a mug, retargeting only maps the keypoints from human hand to robot hand, but doesn't consider

*Equal Contribution. †Corresponding author. All authors are with the Department of Mechanical Engineering, the University of California, Berkeley.

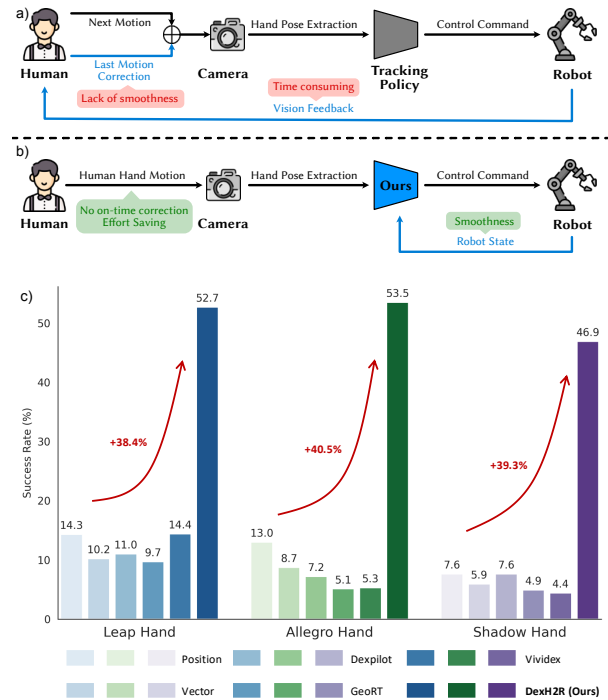


Fig. 1: Unlike traditional methods with human feedback [7]–[9] (Fig. a), DexH2R operates without real-time human intervention ((Fig. b)), significantly reducing human effort, ensuring smooth operation and outperforming baselines by approximately 40% on different robot hands (Fig. c)).

whether the mug is actually lifted and following the desired trajectory at all.

One common approach to address such problem is combining retargeting and human feedback together [5]. By doing so, human can iteratively correct the manipulation errors after performing retargeting at each step based on the feedback from environments, which mitigates the negative impact of retargeting errors in a close-loop manner [3], [6]. However, incorporating human feedback introduces two additional challenges: First, the reliance on human feedback makes manipulation process highly time-consuming and labor-intensive. Second, human tend to make sudden corrective movements in response to system errors, leading to unnatural and non-smooth robot motions. These two disadvantages significantly reduce the transferring efficiency and degrade the quality of robot motions compared to actual human behaviors.

To this end, we propose a novel method that automatically compensates for action errors by leveraging reinforcement learning, thus gets rid of the need for human feedback. Specifically, we develop a residual module to correct the robot hand motions generated by the retargeting module. Taking both environment and object information as inputs, this task-oriented residual model remains beneficial for the improvement of concrete manipulation performance rather than simply mapping the human hand motions. As illustrated in Figure 1, our method bridges the gap between human hand and dexterous robot hand, enables robots to successfully finish the trajectories without any human correction, thus eliminating the high cost of human labor and time consumption. Moreover, our method is compatible with different embodiments and can be directly applied on different robot hands, free from heavy hyper-parameter tuning or human hand manipulation dataset adjustment. Such benefits enable our method to obtain robot manipulation data with high quality in different kinds of dexterous robot hands, which can potentially serve as a scalable way to collect a large amount of dexterous robot data. Overall, contributions of our work are summarized as follows:

- We propose a novel human hand translator for dexterous manipulation, which enables dexterous hands to imitate human hand motions and finish the manipulation tasks at the same time. This algorithm can be generic and compatible with different robot hands.
- We develop a residual action policy based on reinforcement learning to boost the primitive actions generated through retargeting by leveraging both environment and object information.
- We conduct extensive experiments to evaluate our method on three different embodiments using both seen and unseen objects. It notably outperforms baselines that purely consider human hand motions by approximately 40% and demonstrates feasibility in real-world experiments.

II. RELATED WORK

A. Human-robot Motion Alignment

Due to the structural similarity between human and dexterous hands, human hand motions are considered as a good guidance of dexterous robot hand motions. Many previous works aim to obtain dexterous hand motions from human hand motions. The most straightforward method is retargeting [10]–[12], which aligns keypoints positions [3] or vectors [3], [4] between human hands and robot hands to kinematically maximize their similarity. To further improve the transferring performance, human feedback is often introduced to provide on-time correction in a closed-loop manner based on the feedback from environments. Specifically, visual-based teleoperation receives visual feedback from third-view [7], [13], in-hand [8], [14], [15], ego-centric camera [9], [16], [17], VR [6], [18]–[21], human eyes [22] and force-based teleoperation requires operators to wear special equipment to receive real-time force feedback from robot hands [23]. With sufficient environmental feedback, operators can adjust the hand gesture accordingly, thus ensuring the success task completion. However, the performance and

TABLE I: Overview of differences between our method and related work.

	Human hand guidance	Generalize in one policy	Imitate human motions	No human correction	Doesn't need robot data
Retargeting [3], [4], [58]	✓	✓	✓	×	✓
Dexcap [5]	×	✓	✓	×	✓
Dexmv [11]	✓	✓	×	✓	✓
GraspGF [45]	✓	✓	×	✓	×
PGDM [?]	×	×	×	✓	✓
ViVidex [59]	✓	×	×	✓	✓
Ours	✓	✓	✓	✓	✓

efficiency of methods with human feedback are both limited due to the cost and intensity of human labor.

Compared to aforementioned approaches, our method takes task completion into consideration and requires no human correction when transferring human motions to robot motions, thus improving both performance and efficiency.

B. Dexterous Hand Manipulation

Dexterous manipulation remains a big challenge due to its high-dimensional action space [24] and complicated mechanical structures among different types of dexterous hands [25]–[28] for diverse tasks such as pick-and-place, in-hand rotation, and bi-manual manipulation [29]–[31], [31]–[34]. As a result, directly applying a specially designed controller or calculating dexterous actions optimization can be limited and high-cost, thus leading to the exploration of learning methods which obtains better potential in generalization. Reinforcement Learning (RL) is widely introduced as a concise task-oriented approach [29], [35]–[40]. Focusing on grasping objects and following targeted object trajectories [41]–[43], PGDM [44] trains individual policy for each trajectory and emphasizes pre-grasp as a crucial feature in the object grasping phase. By leveraging real robot data, GraspGF learns a score-based grasping primitive action and trains a residual policy through RL as a task-oriented finetune [45]. Besides RL, Imitation Learning (IL) is also explored in dexterous hand [46]–[50]. By leveraging human videos [1], [51]–[53] or extracted human hand trajectories [2], [5], [54], [55], IL shows relatively convincing performances with less real robot data. Furthermore, [56] and [57] provide a reinforcement learning policy that enables objects to follow targeted trajectories after retrieving primitive actions through human demonstration.

Different from their methods, ours focuses on translating human motions to dexterous hand behaviors to realize successful manipulation by leveraging robot-object interaction information through simulation, thus replicating human hand motions as closely as possible without any requirements on real robot data. Table I demonstrates the differences between DexH2R and other work.

III. METHOD

Illustrated in Figure 2, our method that transfers human hand actions to robot hand actions consists of two key stages: acquiring primitive actions through retargeting and obtaining residual actions via reinforcement learning. In the first stage, we first enhance the diversity of human demonstrations via data augmentation. Then, the primitive actions are derived through retargeting optimization based on human hand motions, which

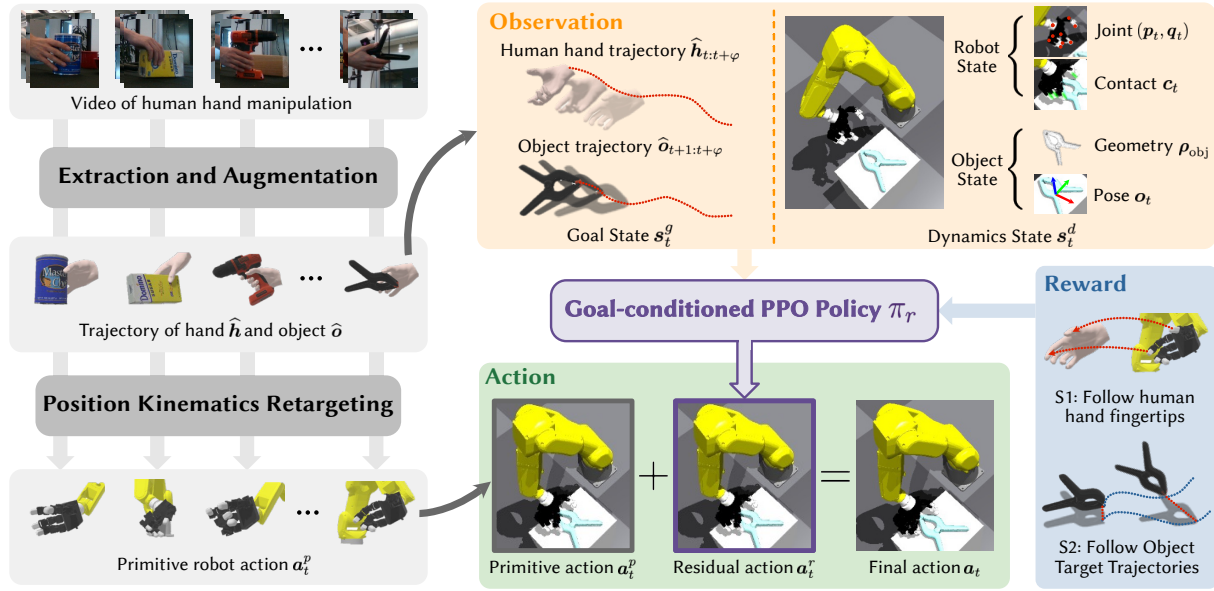


Fig. 2: Overview of DexH2R framework. After extracting and augmenting human hand and object information from demonstrations, we obtain a large amount of trajectories distributed over all workspace. We perform position kinematics retargeting to acquire primitive actions a_t^p . Afterwards, taking in both human hand and object trajectories, we learn a residual action a_t^r to equip our method with task completion information. The primitive and residual actions are combined together as the final actions.

provides a dexterous grasping solution at the kinematic level. In the second stage, we employ reinforcement learning to develop a residual policy to polish these primitive actions and produce the final actions.

Overall, our two-stage method delivers a series of task-accomplished actions transferred from human actions by taking in human demonstration trajectories as well as the current robot state and object state. Leveraging the residual policy to compensate lack of environment and object information when generating primitive actions through retargeting, our method can effectively transfer human hand motions to dexterous hand motions with successful task completion, narrowing the gap between human and robot hands.

A. Preparing Human Data: Extraction and Augmentation

To perform hand poses extraction from grasping images, we first identify the hand poses and keypoints using the MANO model [60], which derives hand keypoints from human hand images. To improve the generalization ability of policy with a limited set of human hand demonstrations, we perform data augmentation on the poses of both hands and objects [2]. Specifically, $\tau = \{(\mathbf{h}_1, \mathbf{o}_1), (\mathbf{h}_2, \mathbf{o}_2), \dots, (\mathbf{h}_N, \mathbf{o}_N)\}$ represents a trajectory composed of human hand poses \mathbf{h}_t and object poses \mathbf{o}_t . Here, human hand poses \mathbf{h}_t are represented as several human hand keypoints positions extracted by MANO. Then, the augmented trajectory τ' is derived as follows:

$$\mathbf{h}'_t = \mathcal{T}' \mathbf{h}_t, \quad \mathbf{o}'_t = \mathcal{T}' \mathbf{o}_t, \quad (1)$$

$$\tau' = \left\{ \left(\mathbf{h}'_1, \mathbf{o}'_1 \right), \dots, \left(\mathbf{h}'_t, \mathbf{o}'_t \right), \dots, \left(\mathbf{h}'_N, \mathbf{o}'_N \right) \right\}, \quad (2)$$

where \mathcal{T}' is the transformation matrix. It translates the trajectories in 3D space and rotates them along the gravity

direction. The augmentation parameters are chosen based on two criteria: (i) the object must lie within the robot's workspace, and (ii) the object should be distributed as broadly as possible across that workspace. The procedure is as follows. 1) We first sample object poses uniformly within the workspace and record the transformation matrix \mathcal{T} from each sampled pose to the original dataset pose. 2) We then apply \mathcal{T} to the robot end-effector pose in Cartesian space, which preserves the relative pose between the robot and the object and thus maintains physical feasibility. 3) If the transformed end-effector pose admits a valid inverse-kinematics solution, we compute the corresponding joint angles and check temporal smoothness against recent history.

Such criteria ensures the consistent and physically plausible hand-object relative motion, and further promotes a more uniform distribution of the augmented dataset over the workspace. In this case, we generate additional feasible demonstration trajectories with diverse object poses and reduce the requirement on the amount of human data, which is crucial for robust policy training.

B. Primitive Actions: Kinematic Retargeting

To facilitate policy training and minimize redundant exploration, we initialize the policy with retargeting solutions. The retargeting process maps human hand motions to robot's hand joints based on human demonstrations. However, due to different sizes and shapes of human and robot hands, direct keypoint mapping [3], [25] can result in infeasible robotic motions. To address this, we exert a scaling factor α to human manipulation trajectory to approximate the size of robot hand. This scaling factor is computed based on the ratio between fingertip span of

the robot and that of the human, ensuring approximate spatial alignment between corresponding keypoints. Subsequently, we formulate a non-linear optimization problem to minimize the topological discrepancies between the keypoints of human and robot hands as follows:

$$\min_{\mathbf{q}_t} \sum_{k=0}^N \left\| \alpha \left(\mathbf{h}_t^k - \mathbf{o}_{p,t} \right) - \left(f_k(\mathbf{q}_t) - \mathbf{o}_{p,t} \right) \right\|^2, \quad (3)$$

$$\text{s.t. } \|\mathbf{q}_t - \mathbf{q}_{t-1}\| \leq d. \quad (4)$$

where \mathbf{q}_t represents the robot joint angles at time step t . $\mathbf{o}_{p,t}$ is the position of the object. $\mathbf{h}_t^k, k = 1, 2, \dots, N$ denotes the position of the k -th human hand keypoint. The function $f_k(\mathbf{q}_t)$ computes the position of the k -th robot hand keypoint through forward kinematics. d is a threshold for joint angles difference between adjacent frames. N and α are fixed per embodiment. N denotes the number of hand keypoints used. With two keypoints per finger, we have $N = 2 \times n_{\text{fingers}}$; for example, Allegro and Leaphand each have four fingers, giving $N = 8$. α is the robot-human size scale, i.e., $\alpha = (\text{robot hand length scale})/(\text{human hand length scale})$. We follow standard values in practice: $\alpha = 1.6$ for Allegro/Leaphand and $\alpha = 1.2$ for Shadow. The spatial step d is determined by the joint maximum velocity v_{max} and the retargeting loop frequency f_z as $d = v_{\text{max}}/f_z$. In practice, we set d slightly larger than v_{max}/f_z to accommodate minor timing jitter and avoid oscillations.

After solving the optimization, we define the primitive action as $\mathbf{a}_t^p = \mathbf{q}_t - \mathbf{q}_{t-1}$ to represent the delta values of joint angles, which are used along with the residual actions.

C. Residual Actions: Goal-conditioned Reinforcement Learning

While the primitive action \mathbf{a}_t^p provides a solution for manipulation, its performance can be highly unstable in real-world deployment since retargeting focuses solely on mapping human hand movements to the robot hand without considering the specific task goals. This open-loop system lacks human correction that exists in previous works [3], [6], [61], so kinematic retargeting errors can accumulate along the trajectory. We observe that the retargeting performance is significantly affected by the gap between human and robot hand, keypoint detection inaccuracies and control errors during deployment, which often causes task failures.

To this end, we introduce a residual policy to refine the final actions by evaluating whether the objects are following the desired trajectories. By incorporating task-specific feedback, the policy compensates for errors in both keypoint extraction and retargeting, leading to improved stability and success rates during manipulation. Specifically, the residual policy is developed based on reinforcement learning whose state and reward are explained as below.

State The state \mathbf{s}_t at each time step consists of the dynamic state \mathbf{s}_t^d and goal state \mathbf{s}_t^g :

$$\mathbf{s}_t \triangleq (\mathbf{s}_t^d, \mathbf{s}_t^g), \quad (5)$$

$$\mathbf{s}_t^d \triangleq (\boldsymbol{\rho}_{\text{obj}}, \mathbf{o}_t^{\text{all}}, \mathbf{q}_t, \mathbf{p}_t - \mathbf{o}_{p,t}, \mathbf{c}_t), \quad (6)$$

$$\mathbf{s}_t^g \triangleq (\hat{\mathbf{h}}_{t:t+\varphi}, \hat{\mathbf{o}}_{t+1:t+\varphi}^{\text{all}} - \mathbf{o}_t^{\text{all}}), \quad (7)$$

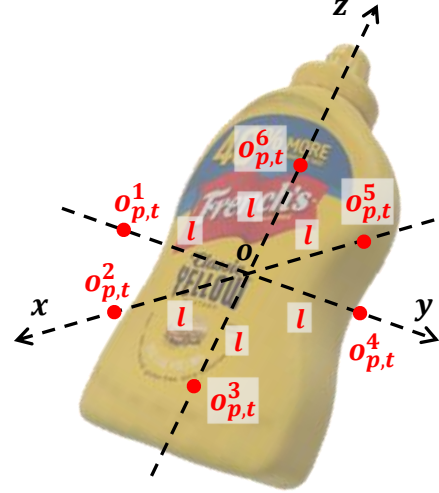


Fig. 3: Illustration of keypoints $\{\mathbf{o}_{p,t}^i\}_{i=1}^6$ that are used to define object poses.

The dynamic state \mathbf{s}_t^d represents the states decided only by the environment (including object and robot states). \mathbf{q}_t is current joint angles. $\mathbf{o}_t^{\text{all}} = \{\mathbf{o}_{p,t}^i\}_{i=1}^6$ stands for positions of six keypoints decided by the object coordinate frame, representing the objects poses. As shown in Figure 3, for each principal axis of the object coordinate frame (x , y , and z), we select two points that lies at a fixed distance l from the object center along both the positive and negative direction of that axis. $\mathbf{p}_t - \mathbf{o}_{p,t}$ emphasizes the relative positions between robot joints and objects. We use a binary variable \mathbf{c}_t to represent the contact state between each robot finger and object due to the inaccuracy and significant sim-to-real gap of contact forces measurement in simulation. $\boldsymbol{\rho}_{\text{obj}}$ denotes the object information including its scale and shape, enabling DexH2R to be generalizable among different objects. The goal state \mathbf{s}_t^g reflects the desired future states of both the human hands $\hat{\mathbf{h}}_{t:t+\varphi}$ and object $\hat{\mathbf{o}}_{t+1:t+\varphi}^{\text{all}}$ between frame t (or $t + 1$) and $t + \varphi$ extracted from human demonstrations, which guides current policy with future prospect to reduce ambiguity.

Reward. We propose a two-stage reward function to guide the manipulation process with different goals in each stage: (1) Following the human demonstration when the robot hand is away from the object, and (2) Tracking the object trajectory when the robot hand is close to the object. Concretely, the reward function is written as:

$$r_t = \begin{cases} r_t^{\text{hand}} & \text{if } t \leq t_0 \\ r_t^{\text{obj}} & \text{if } t > t_0 \end{cases} \quad (8)$$

When the robot hand is far from the object, r_t^{hand} encourages the robot hand to move closer to the object. This reward is designed to compensate for the inaccuracies of primitive actions caused by inevitable errors in optimization computation and robot control. Once the robot hand is close enough, r_t^{obj} makes the object follow the targeted trajectory, which is the primary objective of the manipulation task. We empirically set a pre-defined switching time t_0 to transit between these two stages,

which is 15 steps before the object is lifted in demonstrations. The detailed reward function for each stage is shown as follow:

$$\begin{aligned} r_t^{\text{hand}} &= \beta_{\text{hand}} * \exp\left(-\gamma_{\text{hand}} \sum_{m=0}^4 \|\widehat{\mathbf{h}}_t^m - \mathbf{p}_t^m\|^2\right), \quad (9) \\ r_t^{\text{obj}} &= \beta_{\text{obj}}^{\text{close}} * \exp\left(-\gamma_{\text{obj}}^{\text{close}} \sum_{m=0}^4 \|\mathbf{p}_t^m - \mathbf{o}_{p,t}\|^2\right) \\ &\quad + \beta_{\text{obj}}^{\text{follow}} * \exp\left(-\gamma_{\text{obj}}^{\text{follow}} \|\widehat{\mathbf{o}}_t^{\text{all}} - \mathbf{o}_t^{\text{all}}\|^2\right) \end{aligned} \quad (10)$$

where m represents the index for each fingertip and β, γ are hyperparameters. This staged reward function brings a smooth and accurate manipulation process from approaching the object to completing the task.

IV. EXPERIMENTS

A. Experiment Setup

Demonstration Dataset. For our human demonstration dataset, we utilize DexYCB [62], which focuses on grasping tasks. It contains 20 objects in total, with approximately 25 grasping trajectories for each object using diverse grasp poses. In our study, we leave out 4 objects for unseen test objects while the remaining objects are adopted in training. Due to the inherent differences between human hands and the hardware of dexterous robot hands, we filter the dataset based on two key criteria: 1) The trajectory must be achievable using only rigid body contact, without relying on deformable materials to provide additional friction or forces from multiple directions (e.g. tasks like lifting an upside-down bowl). 2) The object should be placed stably without the risk of falling or rolling away during manipulation. We further expand the dataset using the approach detailed in Section III-A to generate demonstrations that cover the entire workspace. Additionally, we interpolate the robot joint trajectories to match its limited velocity constraints for robots.

Baselines. We compare our method against three widely-used retargeting techniques that aim to transfer human hand motions to robot motions:

- **Position Retargeting** [3]: This method minimizes the distances between the keypoints of the human hand and the robot hand. It is exactly the same approach that generates the primitive actions in our method (Sec. III-B).
- **Vector Retargeting** [3]: This technique aligns the vectors from the wrist to the fingertips between human and robot embodiments. It also calculates the robot arm joint angles using IK based on the relative poses between the human wrist and the robot’s initial frame.
- **Dexpilot Retargeting** [4]: Similar to vector retargeting, this method aligns both the vectors from the wrist to the fingertips and the vectors between fingertips for each embodiment. The arm joint angles are also computed using IK.
- **GeoRT** [58]: GeoRT tackles retargeting by training a model with a designed loss that maps human finger keypoints to robot hand keypoints, achieving high speed and accuracy.
- **ViVidex-SP** [59]: ViVidex-SP denotes the state policy of ViVidex. Like our method, it uses RL with trajectory-guided rewards to learn state-based policies from human motions. However, it (1) trains a separate policy per trajectory, while our single model handles all trajectories; (2) does not take

human-hand motions as input, so it misses fine-grained hand cues; and (3) relies on precomputed retargeting only up to pre-grasp, whereas we condition directly on human-hand keypoints and use retargeted motions merely as primitives. Consequently, ViVidex-SP has to learn more from scratch and cannot correct retargeting errors at run time.

Evaluation Metrics. We quantitatively evaluate the performance of all methods with four key metrics: grasping success rate SR_{Grasp} , following success rate $\text{SR}_{\text{Follow}}$, position error E_p (m) and rotation error E_r (radian). SR_{Grasp} reflects successful grasping and lifting that occurs during the whole trajectory. $\text{SR}_{\text{Follow}}$ reflects successful trajectory following where objects must be constantly lifted without falling from the robot hand along the entire trajectory. Two success rates are counted among the entire test set. E_p and E_r are the average translation and rotation errors between current trajectories and targeted object trajectories. Those two metrics are only counted in successful trajectories after the objects are grasped. It is worth noticing that In our setup, the episode length mainly follows the human demonstration, so factors like initial human–object distance differ across demos and do not reflect algorithmic speed. To make evaluation fair, we use a fixed time window T a trajectory is counted successful only if it finishes within T and without re-grasps. Under this protocol, all methods have comparable inference time 10 s.

Implementation details We utilize Isaac Gym [63] for simulation. The robot’s control frequency is set to 10 Hz with a PD controller for low-level position control. PPO [64] is employed to train the residual policy and only one 3090 GPU is needed for policy training. Network structure of our residual policy is shown in Figure 6, where all features after encoders are concatenated and sent into actor-critic policy. For encoding features, all inputs are encoded with MLP except for the object shape, which is encoded with a pretrained PointNeXt [65] network to better process point cloud information. Unless otherwise specified, we apply the Leap Hand [25] with a FANUC robot arm in the experiments. Table II and III shows the detailed setting of reward parameters and training parameters.

The two-stage reward hyperparameters are chosen as follows. (1) In the first stage we jointly constrain fingertip and fingerdip positions to preserve finger directions and improve grasping, but only fingertip position is used for termination/reset since fingerdip alignment is not a hard constraint. (2) In the second stage, lifting and trajectory-following terms get the largest weights because they are the final goals, and the lifting term is linearized to better handle multi-finger coordination. (3) The decay coefficient is set so the exponential term is between e^{-2} and e^{-1} in the first episode before weighting. (4) The reset threshold is set so about half of the environments terminate in the first episode, which stabilizes training.

B. Results and Comparisons

We conduct extensive experiments to evaluate DexH2R in both simulation and real-robot environments comprehensively. Specifically, we unfold the results of the targeted experiments to address the following key questions:

- **Q1:** Can DexH2R improve the retargeting performance?

TABLE II: Reward parameter. Several things to notice: 1) Only exponent designed reward has decay coefficient, while linear reward are labeled as N/A in table. 2) N/A in threshold means there is no threshold. 3) lift up reward has two thresholds. Within this threshold, reward grows linearly based on the lift-up distance.

Reward category	weight	decay coefficient	threshold
fingertips	1.0	-0.4	0.1
fingerdips	0.5	-0.4	N/A
velocity smoothness	0.2	-5	0.5
lift up	10.0	N/A	[0.01,0.08]
hand reaching objects	1.0	-0.5	0.07
object following trajectory	10.0	-0.5	1.0




Fig. 4: Simulation results showing that DexH2R can generalize to both seen and unseen objects even for those with long-tailed and rare shapes in three different embodiments.

- **Q2:** Is DexH2R generic among different embodiments?
- **Q3:** Can we deploy DexH2R to a real-world robot?
- **Q4:** Which components are critical to DexH2R?
- **Q5:** What is the current limitation of DexH2R?

Q1: Can DexH2R improve the retargeting performance?

In Table IV, we report the quantitative results in the simulation environment. It is obvious that the retargeting baselines achieve poor performance when there is no real-time human correction. Without the capacity to adapt to task dynamics or environmental feedback, these baselines often fail to maintain stable manipulation, especially during long-horizon execution. Additionally, data-driven method also underperforms, which further confirms that our strategy of leveraging retargeted demonstrations as primitive actions, together with our reward design, is effective. DexH2R significantly improves the success rates of both grasping and trajectory following, showing a clear advantage in terms of robustness and adaptability. Especially, notable performance gains are witnessed in the SR_{follow} metric, which reflects the great task completion performance of DexH2R. Specifically, DexH2R can follow desired trajectories throughout the entire manipulation process without falling, which is crucial for practical applications with high robustness requirements. These improvements suggest that our method is capable of producing temporally consistent and stable behaviors by remaining contact between robots and




Fig. 5: Data distribution of original and augmented dataset.




Fig. 6: Network structure of DexH2R.

objects, which demonstrates that our method effectively bridges the gap between human demonstrations and dexterous robot execution, ensuring successful task completion across diverse scenarios.

We also provide some visualizations in Figure 4, where DexH2R can grasp both seen and unseen objects properly including those with complex geometrical shapes. Such results highlight the necessity of our proposed task-oriented residual policy, which takes the task-specific goal and environment information into consideration. This design enables the policy to focus on critical aspects of the manipulation task, such as contact stability and path precision, ultimately contributing to its generalizability and robustness.

It is worth mentioning that some baselines are comparable with DexH2R in tracking errors because they can only follow some trivial trajectories, and only successful trajectory-following cases are considered in the calculation of tracking errors. Specifically, Position and orientation tracking errors (E_p, E_r) are computed only on trajectories that finish successfully (i.e., the object does not slip or fall during the entire episode). Because our method achieves a much higher finishing success rate SR_{follow} , the set over which its errors are averaged is larger and contains more challenging cases (e.g., longer motions or larger pose changes), whereas several baselines contribute fewer and typically easier successful episodes. Nevertheless, improving tracking accuracy remains an important goal, and we will continue to reduce errors while maintaining high success.

Q2: Is DexH2R generic among different embodiments?

We claim DexH2R as a generic algorithm compatible with different embodiments with minimal hyperparameter adjustments

TABLE III: Encoder design and training parameter. Encoder parameter variations is caused by different number of fingers among embodiments.

	Encoder Network								Training setting		
	object shape	object scale	human hand trajectory	history joint	object poses	object follow distances	object-hand distances	contact	learning rate	batch size	horizon length
LeapHand	[1024,256,64]	[1,64]	[288,512,256]	[22, 128]	[30, 128]	[198, 128]	[48, 128]	[4, 128]	1e-5	32768	32
Allegro											
Shadow			[360,512,256]	[30,128]							

TABLE IV: Quantitative results for seen objects, unseen trajectories and unseen objects with our method.

Test set	Method	Leap Hand				Allegro Hand				Shadow Hand			
		SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓	SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓	SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓
Seen objects	Position	32.0%	13.7%	0.083	0.364	25.0%	12.3%	0.077	0.414	17.4%	5.9%	0.103	0.675
	Vector	32.7%	9.1%	0.108	0.649	25.3%	6.0%	0.096	0.632	22.5%	5.3%	0.131	0.884
	Dexpilot	32.8%	10.7%	0.111	0.523	23.3%	6.2%	0.092	0.578	22.5%	5.3%	0.131	0.884
	GeoRT	32.4%	12.3%	0.055	0.245	23.7%	7.3%	0.093	0.575	31.1%	8.1%	0.118	0.557
	ViViDex-SP	44.6%	20.0%	0.038	0.278	28.9%	7.9%	0.084	0.501	30.6%	8.7%	0.110	0.508
	Ours	69.8%	52.1%	0.048	0.452	69.8%	52.5%	0.054	0.483	65.4%	43.4%	0.084	0.616
Unseen objects	Position	29.0%	14.3%	0.091	0.389	23.9%	13.0%	0.074	0.375	13.4%	7.6%	0.124	0.568
	Vector	31.8%	10.2%	0.103	0.533	22.3%	8.7%	0.083	0.479	16.2%	5.9%	0.115	0.829
	Dexpilot	32.2%	11.0%	0.105	0.545	22.5%	7.2%	0.083	0.539	15.9%	7.6%	0.117	0.692
	GeoRT	29.9%	9.7%	0.063	0.221	19.9%	5.1%	0.095	0.449	24.6%	4.9%	0.105	0.488
	ViViDex-SP	33.1%	14.4%	0.043	0.309	21.5%	5.3%	0.106	0.652	17.0%	4.4%	0.103	0.500
	Ours	70.9%	52.7%	0.055	0.476	69.1%	53.5%	0.055	0.568	62.4%	46.9%	0.103	0.808

during policy training. In this part, we implement our policy in three distinct embodiments: the Leap Hand with a FANUC robot arm, the Allegro Hand with a Kinova robot arm, and the Shadow Hand with a UR10e robot arm. These embodiments vary significantly in terms of joint configurations and physical dynamics, making it a challenging testbed for generalization.

Despite these differences, quantitative results in Table IV indicate that our policy performs greatly in all of them and significantly outperforms baselines in both grasp success and trajectory-following metrics. Such consistency across embodiments suggests that DexH2R is robust to variations in kinematic structures and control latency, which are common challenges in real-world multi-robot deployment. This demonstrates the adaptability of our policy across robot systems as a generic solution that does not require architecture-specific re-design. Furthermore, our policy enables direct successful policy training in new hardware with only minimal hyperparameter adjustment, reducing the tuning effort traditionally required when adapting learned models to different robots. This highlights a significant advantage: by leveraging DexH2R, a single dataset of human demonstrations can be successfully transferred to various robot embodiments for task execution without the need for hardware-specific optimization, substantially enhancing the efficiency of robot data collection.

Q3: Can we deploy DexH2R to a real-world robot?

We deploy our policy to a real robot system composed of a FANUC LR Mate robot with a Leap Hand [25]. We evaluated our policy on four different objects, including two seen and two unseen from the YCB dataset [66]. Objects are scaled up to match their sizes in human demonstrations. We conduct 10 experiments for each object through open-loop execution and record the success rates of grasping SR_{Grasp} and trajectory following SR_{Follow} . The evaluation covers both seen and unseen objects with diverse poses and orientations, providing insight into the generalizability of DexH2R across different scenarios.

TABLE V: Performance of our method on different objects.

Objects		SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓
Seen objects	master chef can	84.7%	63.3%	0.035	0.259
	bleach cleanser	87.5%	75.9%	0.031	0.301
	scissors	31.2%	29.6%	0.066	0.598
	power drill	65.8%	34.8%	0.05	0.408
	Overall	69.8%	52.1%	0.048	0.452
Unseen objects	potted meat can	71.9%	57.3%	0.045	0.496
	mug	62.0%	47.0%	0.062	0.489
	Overall	70.9%	52.7%	0.055	0.476

As shown in Table VIII, our proposed method successfully bridges the sim-to-real gap by achieving relatively high performance in the success rates of both grasping and trajectory following. In most cases, the hand pose remains fairly consistent in the entire manipulation process, demonstrating its practical feasibility. We also notice a drop in the success rate of trajectory following for unseen objects due to the lack of closed-loop feedback and object-wise physical properties especially if the unseen objects have irregular shapes or the desired trajectories contain complex rotations.

We visualize some qualitative results in Figure 8, which shows that DexH2R can successfully grasp the objects and follow the desired trajectories with diverse objects in the real world. These visualizations further validate the adaptability of our approach and its potential for practical deployment in unstructured environments.

Q4: Which components are critical to DexH2R?

In this part, we investigate the role of each module in DexH2R. Specifically, we dig into the necessity of each stage in our pipeline in Table VI a). A significant performance drop is witnessed if we do not apply data augmentation in human data extraction since data augmentation can greatly boost the diversity of demonstrations in human data. This is because the augmented dataset covers a larger portion of the workspace. Figure 5 compares the original and augmented pose

TABLE VI: Ablation studies of two different stages. Experiments are conducted on unseen objects to make the differences more clear.

a) Ablation studies on staging pipeline.

Method	SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓
Ours	70.9%	52.7%	0.055	0.476
w/o data aug	54.7%	34.1%	0.079	0.557
w/o prim actions	66.3%	47.4%	0.055	0.553

b) Ablation studies on residual policy.

Method	SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓
Ours	70.9%	52.7%	0.055	0.476
w/o contacts	67.6%	47.5%	0.056	0.511
w/o dense reward	17.4%	2.7%	0.094	0.487

TABLE VII: DexH2R with different switching steps.

Switching step	SR _{Grasp} ↑	SR _{Follow} ↑	E_p ↓	E_r ↓
5	64.9%	46.8%	0.041	0.431
10	68.8%	52.0%	0.040	0.430
15	69.8%	52.1%	0.048	0.452
20	65.7%	48.7%	0.048	0.420
25	64.3%	44.9%	0.049	0.446

distributions of object 'master chef can', increasing number of data in workspace from **14** to **182**. Besides, the elimination of the primitive action slightly hurts the model performance. Although reinforcement learning can derive a reasonable policy even without primitive actions, it would take a long time to converge (see Table IX) and can hardly ensure the smoothness of the trajectories. This underscores the benefit of primitive actions in speeding up the learning process, as they provide a good initialization for the policy.

In Table VI b), we analyze the design of the residual in the second stage. The binary contact information in the observation can provide useful cues about the interaction between human hands and objects, which helps to lift the performance. We also emphasize the importance of the dense rewards. Without the guidance provided by human hand rewards, the exploration of the residual policy becomes almost unmanageable due to the high-dimensional action space, which aligns with the challenges described in [44]. This highlights the necessity of our two-stage reward to provide directional guidance during training, especially when explicit supervision is limited.

We also conducted experiments to assess how object properties affect performance. As demonstrated in Table V, the dexterous hand achieves higher success rates with objects that have regular shapes and suitable sizes. These objects present more feasible grasping solutions compared to objects with irregular shapes or extreme sizes because the consistent geometry and contact surface of these regular objects allow for more stable and repeatable contact interactions during manipulation. This finding highlights the importance of object properties in determining the effectiveness of manipulation tasks, suggesting a promising direction for future research. Additionally, the pre-defined switching step is initially selected based on prior experience and we also conducted an ablation over different switching steps shown in Table VII and Figure 7.

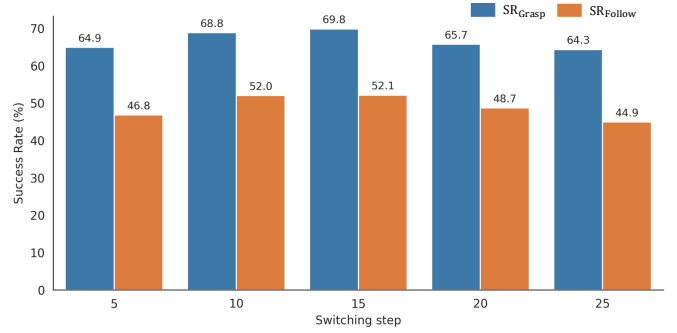


Fig. 7: DexH2R with different switching steps.

TABLE VIII: Success rates of real-world experiments on both seen and unseen objects.

Objects		SR _{Grasp} ↑	SR _{Follow} ↑
Seen objects	sugar box	10/10	9/10
	tomato soup can	9/10	7/10
Unseen objects	mug	8/10	5/10
	mustard bottle	8/10	4/10

The quantitative results indicate that, within a reasonable range, the switching step has limited impact on final performance.

Q5: What is the current limitation of DexH2R?

Although DexH2R improves success rates over retargeting baselines, we observe several method-specific failure modes that manifest consistently across simulation and real hardware.

Long-horizon drift without closed-loop correction. Despite error mitigation brought our residual design, small pose errors and contact noise can still accumulate over extended trajectories, leading to late-stage object drops or poses offsets during precise rotations.

Physics- and shape-dependent fragility. Irregular geometries, scales, and physical parameters of object increase sensitivity to small pose deviations. The effect is more pronounced on unseen objects and in sequences requiring sustained rotational alignment because of unseen shape and physical parameters. This is also amplified in the real robot experiments, where more out-of-distribution environment settings are involved such as friction, mass and contact.

Perception mismatch Because DexH2R conditions on human-hand keypoints and object encodings, their perception inaccuracies can shift the intended contact region at approach, elevating failure risk near first touch. Even though our reward design can mitigate such bad effect by focusing on object contact at second stage, these inaccuracies still reduce policy's sensitivity to this input channel, thereby degrading overall accuracy.

Overall, despite strong performance, DexH2R can still struggle on precise, long-horizon rotations without feedback, during contact-rich transitions on irregular objects, and under perception noise of human hand and object. These limitations arise from open-loop execution, lack of training data, and sensitivity to perception mismatch, suggesting promising directions for future work.




Fig. 8: Real-world experiments. DexH2R can generalize to various objects with each row as a manipulation action sequence.

TABLE IX: Training time and converge iterations for our policies and policy without primitive actions.

a) Time performance of our method across different embodiments.

Policy	Training Time	Iterations for converge
LeapHand	16.53h	2593
Allegro	21.98h	3504
Shadow	21.98h	3504

b) Time performance comparison between our method with and without primitive actions, where 'w prim' means with primitive actions and 'w/o prim' means without.

Policy	Training Time	Iterations for converge
w prim	16.53h	2593
w/o prim	21.98h	3504

V. CONCLUSIONS

In this work, we introduce DexH2R, a framework that transfers human hand trajectories to dexterous robot motions while achieving task completion without on-time human feedback. By decoupling our policy execution from human intervention, the system is more suitable for autonomous deployment when ensuring the successful task completion. Also, DexH2R is generic and compatible with different embodiments without hardware-specific adjustment for human demonstrations. Those two benefits enable DexH2R to possess large potential in becoming a scalable data collection method, lowering the demand on both hardware and human labor. With comprehensive experiments on three different embodiments, our method exhibits a notable performance gain compared with baselines on grasping tasks in both simulation and real-world environments.

In the future, we plan to extend DexH2R to more complex tasks, such as in-hand manipulation, to further test its

generality under richer contact interactions. Another direction is integrating tactile sensing to provide fine-grained contact information, enhancing the policy’s perception of subtle contact dynamics and enabling more precise manipulation in contact-rich scenarios. We also envision exploring improved training protocols and evaluation suites that stress long-horizon stability and robustness across diverse objects and embodiments.

REFERENCES

- [1] Y. Qin, Y.-H. Wu, S. Liu, H. Jiang, R. Yang, Y. Fu, and X. Wang, “Dexmv: Imitation learning for dexterous manipulation from human videos,” in *European Conference on Computer Vision*. Springer, 2022, pp. 570–587.
- [2] Z. Q. Chen, K. Van Wyk, Y.-W. Chao, W. Yang, A. Mousavian, A. Gupta, and D. Fox, “Dextranfer: Real world multi-fingered dexterous grasping with minimal human demonstrations,” *arXiv preprint arXiv:2209.14284*, 2022.
- [3] Y. Qin, W. Yang, B. Huang, K. Van Wyk, H. Su, X. Wang, Y.-W. Chao, and D. Fox, “Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system,” *arXiv preprint arXiv:2307.04577*, 2023.
- [4] A. Handa, K. Van Wyk, W. Yang, J. Liang, Y.-W. Chao, Q. Wan, S. Birchfield, N. Ratliff, and D. Fox, “Dexpilot: Vision-based teleoperation of dexterous robotic hand-arm system,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9164–9170.
- [5] C. Wang, H. Shi, W. Wang, R. Zhang, L. Fei-Fei, and C. K. Liu, “Dexcap: Scalable and portable mocap data collection system for dexterous manipulation,” *arXiv preprint arXiv:2403.07788*, 2024.
- [6] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang, “Open-television: teleoperation with immersive active visual feedback,” *arXiv preprint arXiv:2407.01512*, 2024.
- [7] Q. Gao, Z. Ju, Y. Chen, Q. Wang, and C. Chi, “An efficient rgb-d hand gesture detection framework for dexterous robot hand-arm teleoperation system,” *IEEE Transactions on Human-Machine Systems*, vol. 53, no. 1, pp. 13–23, 2022.
- [8] Z. Si, K. L. Zhang, Z. Temel, and O. Kroemer, “Tilde: Teleoperation for dexterous in-hand manipulation learning with a deltaxhand,” *arXiv preprint arXiv:2405.18804*, 2024.
- [9] Y. Qin, H. Su, and X. Wang, “From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 873–10 881, 2022.
- [10] D. Antotsiou, G. Garcia-Hernando, and T.-K. Kim, “Task-oriented hand motion retargeting for dexterous manipulation imitation,” in *Proceedings of the European conference on computer vision (ECCV) workshops*, 2018, pp. 0–0.

- [11] S. P. Arunachalam, S. Silwal, B. Evans, and L. Pinto, "Dexterous imitation made easy: A learning-based framework for efficient dexterous manipulation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5954–5961.
- [12] Z. Chen, S. Kiami, A. Gupta, and V. Kumar, "Genaug: Retargeting behaviors to unseen situations via generative augmentation," *arXiv preprint arXiv:2302.06671*, 2023.
- [13] H. Liu, S. Nasiriany, L. Zhang, Z. Bao, and Y. Zhu, "Robot learning on the job: Human-in-the-loop autonomy and learning during deployment," *The International Journal of Robotics Research*, p. 02783649241273901, 2022.
- [14] S. Li, J. Jiang, P. Ruppel, H. Liang, X. Ma, N. Hendrich, F. Sun, and J. Zhang, "A mobile robot hand-arm teleoperation system by vision and imu," in *2020 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 900–10 906.
- [15] S. Li, N. Hendrich, H. Liang, P. Ruppel, C. Zhang, and J. Zhang, "A dexterous hand-arm teleoperation system based on hand pose estimation and active vision," *IEEE Transactions on Cybernetics*, vol. 54, no. 3, pp. 1417–1428, 2022.
- [16] S. Li, X. Ma, H. Liang, M. Görner, P. Ruppel, B. Fang, F. Sun, and J. Zhang, "Vision-based teleoperation of shadow dexterous hand using end-to-end deep neural network," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 416–422.
- [17] R. M. Aronson and H. Admoni, "Gaze complements control input for goal prediction during assisted teleoperation," in *Robotics science and systems*, 2022.
- [18] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 5628–5635.
- [19] P. Ponomareva, D. Trinitatova, A. Fedoseev, I. Kalinov, and D. Tsetserukou, "Grasplook: a vr-based telemanipulation system with r-cnn-driven augmentation of virtual environment," in *2021 20th International Conference on Advanced Robotics (ICAR)*. IEEE, 2021, pp. 166–171.
- [20] S. P. Arunachalam, I. Güzey, S. Chintala, and L. Pinto, "Holo-dex: Teaching dexterity with immersive mixed reality," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5962–5969.
- [21] J. I. Lipton, A. J. Fay, and D. Rus, "Baxter's homunculus: Virtual reality spaces for teleoperation in manufacturing," *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 179–186, 2017.
- [22] Z.-H. Yin, C. Wang, L. Pineda, F. Hogan, K. Bodduluri, A. Sharma, P. Lancaster, I. Prasad, M. Kalakrishnan, J. Malik *et al.*, "Dexteritygen: Foundation controller for unprecedented dexterity," *arXiv preprint arXiv:2502.04307*, 2025.
- [23] H. Liu, X. Xie, M. Millar, M. Edmonds, F. Gao, Y. Zhu, V. J. Santos, B. Rothrock, and S.-C. Zhu, "A glove-based system for studying hand-object manipulation via joint pose and force sensing," in *2017 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6617–6624.
- [24] H. Mania, A. Guy, and B. Recht, "Simple random search provides a competitive approach to reinforcement learning," *arXiv preprint arXiv:1803.07055*, 2018.
- [25] K. Shaw, A. Agarwal, and D. Pathak, "Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning," *arXiv preprint arXiv:2309.06440*, 2023.
- [26] 'shadow hand'. [Online]. Available: <https://www.wonikrobotics.com/research-robot-hand>
- [27] 'allegro hand'. [Online]. Available: <https://ninjatek.com/shop/edge/>
- [28] Z. Si, K. Zhang, O. Kroemer, and F. Z. Temel, "Deltahands: A synergistic dexterous hand framework based on delta robots," *IEEE Robotics and Automation Letters*, 2024.
- [29] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.
- [30] T. Chen, J. Xu, and P. Agrawal, "A system for general in-hand object re-orientation," in *Conference on Robot Learning*. PMLR, 2022, pp. 297–307.
- [31] W. Huang, I. Mordatch, P. Abbeel, and D. Pathak, "Generalization in dexterous manipulation via geometry-aware multi-task learning," *arXiv preprint arXiv:2111.03062*, 2021.
- [32] A. Gupta, J. Yu, T. Z. Zhao, V. Kumar, A. Rovinsky, K. Xu, T. Devlin, and S. Levine, "Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6664–6671.
- [33] Y. Chen, T. Wu, S. Wang, X. Feng, J. Jiang, Z. Lu, S. McAleer, H. Dong, S.-C. Zhu, and Y. Yang, "Towards human-level bimanual dexterous manipulation with reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5150–5163, 2022.
- [34] S. Zhao, K. Yang, Y. Chen, C. Li, Y. Xie, X. Zhang, C. Wang, and M. Tomizuka, "Dexctrl: Towards sim-to-real dexterity with adaptive controller learning," *arXiv preprint arXiv:2505.00991*, 2025.
- [35] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar, "Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3651–3657.
- [36] K. Li, P. Li, T. Liu, Y. Li, and S. Huang, "Maniptrans: Efficient dexterous bimanual manipulation transfer via residual learning," *arXiv preprint arXiv:2503.21860*, 2025.
- [37] G. Khandate, S. Shang, E. T. Chang, T. L. Saidi, Y. Liu, S. M. Dennis, J. Adams, and M. Ciocarlie, "Sampling-based exploration for reinforcement learning of dexterous manipulation," *arXiv preprint arXiv:2303.03486*, 2023.
- [38] X. Liu, J. Adalibieke, Q. Han, Y. Qin, and L. Yi, "Dextrack: Towards generalizable neural tracking control for dexterous manipulation from human references," *arXiv preprint arXiv:2502.09614*, 2025.
- [39] Q. She, R. Hu, J. Xu, M. Liu, K. Xu, and H. Huang, "Learning high-dof reaching-and-grasping via dynamic representation of gripper-object interaction," *arXiv preprint arXiv:2204.13998*, 2022.
- [40] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1101–1112.
- [41] J. Bae, J. Won, D. Lim, C.-H. Min, and Y. M. Kim, "Pmp: Learning to physically interact with environments using part-wise motion priors," in *ACM SIGGRAPH 2023 Conference Proceedings*, 2023, pp. 1–10.
- [42] Y. Wang, J. Lin, A. Zeng, Z. Luo, J. Zhang, and L. Zhang, "Physhoi: Physics-based imitation of dynamic human-object interaction," *arXiv preprint arXiv:2312.04393*, 2023.
- [43] J. Braun, S. Christen, M. Kocabas, E. Aksan, and O. Hilliges, "Physically plausible full-body hand-object interaction synthesis," in *2024 International Conference on 3D Vision (3DV)*. IEEE, 2024, pp. 464–473.
- [44] S. Dasari, A. Gupta, and V. Kumar, "Learning dexterous manipulation from exemplar object trajectories and pre-grasps," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3889–3896.
- [45] T. Wu, M. Wu, J. Zhang, Y. Gan, and H. Dong, "Graspgf: learning score-based grasping primitive for human-assisting dexterous grasping," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023, pp. 22 132–22 150.
- [46] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [47] P. Englert and M. Toussaint, "Learning manipulation skills from a single demonstration," *The International Journal of Robotics Research*, vol. 37, no. 1, pp. 137–154, 2018.
- [48] 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," *arXiv preprint arXiv:2302.12422*, 2023.
- [49] W. Wan, H. Geng, Y. Liu, Z. Shan, Y. Yang, L. Yi, and H. Wang, "Unidexgrasp++: Improving dexterous grasping policy learning via geometry-aware curriculum and iterative generalist-specialist learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3891–3902.
- [50] H.-S. Fang, H. Yan, Z. Tang, H. Fang, C. Wang, and C. Lu, "Anydexgrasp: General dexterous grasping for different hands with human-level learning efficiency," *arXiv preprint arXiv:2502.16420*, 2025.
- [51] A. Sivakumar, K. Shaw, and D. Pathak, "Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube," *arXiv preprint arXiv:2202.10448*, 2022.
- [52] H. Bharadhwaj, A. Gupta, V. Kumar, and S. Tulsiani, "Towards generalizable zero-shot manipulation via translating human interaction plans," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6904–6911.
- [53] K. Shaw, S. Bahl, A. Sivakumar, A. Kannan, and D. Pathak, "Learning dexterity from human hand motion in internet videos," *The International Journal of Robotics Research*, vol. 43, no. 4, pp. 513–532, 2024.
- [54] J. Gu, S. Kirmani, P. Wohlhart, Y. Lu, M. G. Arenas, K. Rao, W. Yu, C. Fu, K. Gopalakrishnan, Z. Xu *et al.*, "Rt-trajectory: Robotic task generalization via hindsight trajectory sketches," *arXiv preprint arXiv:2311.01977*, 2023.


- [55] M. Xu, Z. Xu, C. Chi, M. Veloso, and S. Song, "Xskill: Cross embodiment skill discovery," in *Conference on Robot Learning*. PMLR, 2023, pp. 3536–3555.
- [56] Z. Chen, S. Chen, C. Schmid, and I. Laptev, "Vividex: Learning vision-based dexterous manipulation from human videos," *arXiv preprint arXiv:2404.15709*, 2024.
- [57] Z. Luo, J. Cao, S. Christen, A. Winkler, K. Kitani, and W. Xu, "Grasping diverse objects with simulated humanoids," *arXiv preprint arXiv:2407.11385*, 2024.
- [58] Z.-H. Yin, C. Wang, L. Pineda, K. Bodduluri, T. Wu, P. Abbeel, and M. Mukadam, "Geometric retargeting: A principled, ultrafast neural hand retargeting algorithm," *arXiv preprint arXiv:2503.07541*, 2025.
- [59] Z. Chen, S. Chen, E. Arlaud, I. Laptev, and C. Schmid, "Vividex: Learning vision-based dexterous manipulation from human videos," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 3336–3343.
- [60] J. Romero, D. Tzionas, and M. J. Black, "Embodied hands: Modeling and capturing hands and bodies together," *arXiv preprint arXiv:2201.02610*, 2022.
- [61] T. He, Z. Luo, W. Xiao, C. Zhang, K. Kitani, C. Liu, and G. Shi, "Learning human-to-humanoid real-time whole-body teleoperation," *arXiv preprint arXiv:2403.04436*, 2024.
- [62] Y.-W. Chao, W. Yang, Y. Xiang, P. Molchanov, A. Handa, J. Tremblay, Y. S. Narang, K. Van Wyk, U. Iqbal, S. Birchfield *et al.*, "Dexycb: A benchmark for capturing hand grasping of objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9044–9053.
- [63] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.
- [64] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [65] G. Qian, Y. Li, H. Peng, J. Mai, H. Hammoud, M. Elhoseiny, and B. Ghanem, "Pointnext: Revisiting pointnet++ with improved training and scaling strategies," *Advances in neural information processing systems*, vol. 35, pp. 23 192–23 204, 2022.
- [66] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.



Chenran Li is currently a researcher scientist in Nvidia. He received his Ph.D. degree from University of California, Berkeley in 2025 and his B.S. degree from Harbin Institute of Technology in 2020. His research interests include humanoid and reinforcement learning.



Yichen Xie received his B.E. degree in Computer Science from Shanghai Jiao Tong University in 2021. He is currently pursuing the Ph.D. degree in Mechanical Engineering at University of California, Berkeley, advised by Prof. Masayoshi Tomizuka. His research interest covers computer vision and robotics. He has published papers at conferences such as CVPR, ICCV, ICLR, and NeurIPS, and has served as a reviewer for CVPR, ECCV, NeurIPS, ICLR, and ICML.



Xiang Zhang received the B.Eng. degree in precision mechanical and instrument from the University of Science and Technology of China, Hefei, China, in 2019 and the Ph.D. degree in mechanical and control engineering from the University of California, Berkeley, CA, USA, in 2024. He is currently with FANUC Advanced Research Laboratory, FANUC America Corporation, USA. His research interests include robot modeling, control, and manipulation.



Mingyu Ding is an assistant professor of computer science at the University of North Carolina at Chapel Hill. Prior to joining the department, he was a postdoctoral fellow at BAIR@UC Berkeley with Masayoshi Tomizuka and a visiting scholar at CSAIL@MIT with Joshua Tenenbaum. He obtained a doctorate from the University of Hong Kong in 2022 with Ping Luo. Mingyu's research interests lie at the intersection of robotics, embodied AI, and vision. He received Baidu Fellowship, Microsoft Fellowship Nomination, ME Rising Star, WAIC Rising Star, and the Best Paper Award from ICRA 2024.


Shuqi Zhao is currently a Ph.D. student in Mechanical Engineering from UC Berkeley, advised by professor Masayoshi Tomizuka. She received the B.S. degree in Control Science and Engineering with an honor minor degree in Chu Kochen Honor College from Zhejiang University in 2023 and her research interests include dexterous manipulation, reinforcement learning, and sim-to-real transfer.



Xinghao Zhu holds a B.E. in Automation from Xi'an Jiaotong University (2018) and a Ph.D. in Mechanical Engineering from the University of California, Berkeley (2024). His research interests span control theory, robotics, machine learning, and optimization. He primarily investigates problems in contact-rich and dexterous manipulation, as well as loco-manipulation.



Yuxin Chen received his B.Sc. degree in Mechanical Engineering from Shanghai Jiao Tong University, Shanghai, China, in 2020, and both the B.S.E. degree in Aerospace Engineering and the M.S. degree in Robotics from the University of Michigan, Ann Arbor, MI, USA, in 2020 and 2022, respectively. He is currently pursuing a Ph.D. in Mechanical Engineering at the University of California, Berkeley. His research interests lie in control, optimization, and reinforcement learning, with a focus on mobile robotics, dexterous manipulation, and human-robot



interaction.



the Best Paper Award from ICRA 2024.

Masayoshi Tomizuka was born in Tokyo, Japan, in 1946. He received the B.S. and M.S. degrees in mechanical engineering from Keio University, Tokyo, in 1968 and 1970, respectively, and the Ph.D. degree in mechanical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1974. In 1974, he joined the Faculty of the Department of Mechanical Engineering, University of California at Berkeley, Berkeley, CA, USA, where he is currently the Cheryl and John Neerhout, Jr., Distinguished Professorship Chair and teaches courses in dynamic systems and controls. From 2002 to 2004, he was the Program Director of the Dynamic Systems and Control Program of the Civil and Mechanical Systems Division of NSF. His current research interests include optimal and adaptive control, digital control, signal processing, motion control, and control problems related to robotics, machining, manufacturing, information storage devices, and vehicles. Dr. Tomizuka was the Technical Editor of the ASME Journal of Dynamic Systems, Measurement and Control (J-DSMC) (1988–1993), Editor-in-Chief of IEEE/ASME TRANSACTIONS ON MECHATRONICS (1997–1999), and an Associate Editor for the Journal of the International Federation of Automatic Control, and Automatica. He was the General Chairman of the 1995 American Control Conference, and was the President of the American Automatic Control Council (1998–1999). He is a Life Fellow of the ASME and a Fellow of the International Federation of Automatic Control (IFAC) and the Society of Manufacturing Engineers. He was the recipient of the Best J-DSMC Best Paper Award (1995 and 2010), DSCD Outstanding Investigator Award (1996), Charles Russ Richards Memorial Award (ASME, 1997), Rufus Oldenburger Medal (ASME, 2002), John R. Ragazzini Award (AACC, 2006), Richard E. Bellman Control Heritage Award (AACC, 2018), Honda Medal (ASME, 2019), and Nathaniel B. Nichols Medal (IFAC, 2020). He is a member of the National Academy of Engineering.