

Holo-Dex: Teaching Dexterity with Immersive Mixed Reality

Sridhar Pandian Arunachalam
New York University

Irmak Güzey
New York University

Soumith Chintala
Meta AI Research

Lerrel Pinto
New York University

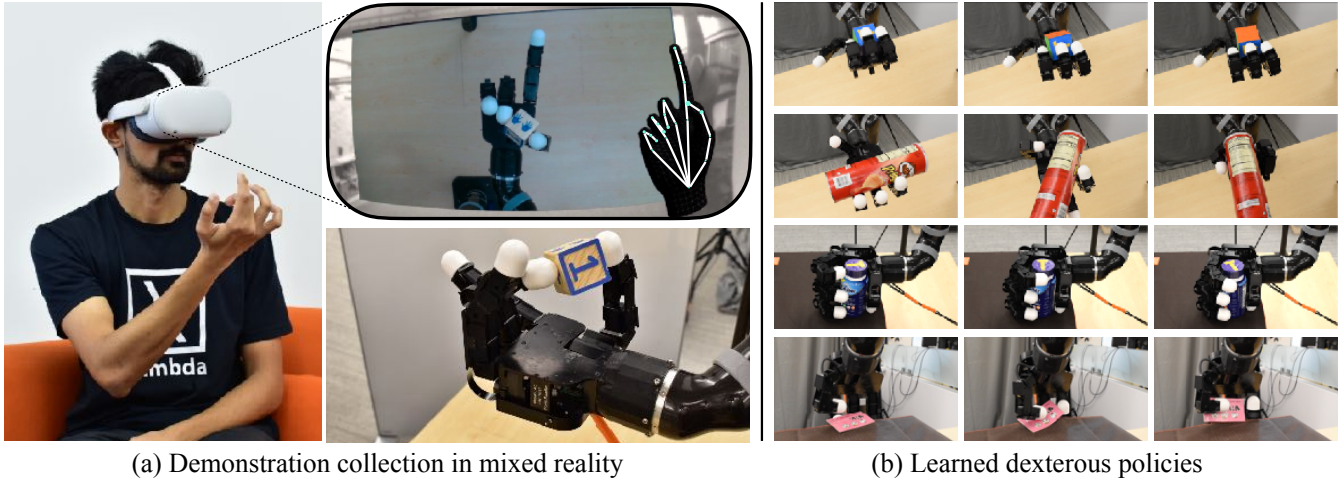


Fig. 1: We present HOLO-DEX, a framework that (a) collects high-quality demonstration data by placing human teachers in an immersive mixed reality world, and then (b) learns visual policies from a handful of these demonstrations to solve dexterous manipulation tasks.

Abstract—A fundamental challenge in teaching robots is to provide an effective interface for human teachers to demonstrate useful skills to a robot. This challenge is exacerbated in dexterous manipulation, where teaching high-dimensional, contact-rich behaviors often require esoteric teleoperation tools. In this work, we present HOLO-DEX, a framework for dexterous manipulation that places a teacher in an immersive mixed reality through commodity VR headsets. The high-fidelity hand pose estimator onboard the headset is used to teleoperate the robot and collect demonstrations for a variety of general-purpose dexterous tasks. Given these demonstrations, we use powerful feature learning combined with non-parametric imitation to train dexterous skills. Our experiments on six common dexterous tasks, including in-hand rotation, spinning, and bottle opening, indicate that HOLO-DEX can both collect high-quality demonstration data and train skills in a matter of hours. Finally, we find that our trained skills can exhibit generalization on objects not seen in training. Videos of HOLO-DEX are available on <https://holo-dex.github.io/>.

I. INTRODUCTION

Learning-based methods have had a transformational effect in robotics on a wide range of domains from manipulation [1, 2], locomotion [3, 4, 5], and aerial robotics [6, 7, 8]. Such methods often produce policies that input raw sensory observations and output robot actions. This circumvents challenges in developing state-estimation modules, modeling object properties and tuning controller gains, which requires significant domain expertise. Even with the steep progress in robot learning, we are still long way off from dexterous robots that can solve arbitrary robot tasks akin to methods

in game play [9, 10], text generation [11, 12] or few-shot vision [13, 14].

To understand what might be missing in robot learning, we need to ask a central question: How do we collect training data for our robots? One option is to collect data on the robot through self-supervised data collection strategies. While this results in robust behaviors [15, 16, 17, 18], they often require extensive real-world interactions in the order of thousands of hours even for relatively simple manipulation tasks [19]. An alternate option is to train on simulated data and then transfer to the real robot (*Sim2Real*). This allows for learning complex robotic behaviors multiple orders of magnitude faster than on-robot learning [20, 21, 22]. However, setting up simulated robot environments and specifying simulator parameters often requires extensive domain expertise [23, 24].

A third, more practical option to collect data is by asking human teachers to provide demonstrations [25, 26]. Robots can then be trained to quickly imitate the demonstrated data. Such imitation methods have recently shown promise in a variety of challenging dexterous manipulation problems [27, 28, 29]. However, there lies a fundamental limitation in most of these works – collecting high-quality demonstration data for dexterous robots is hard! They either require expensive gloves [30], extensive calibration [28], or suffer from monocular occlusions [29].

In this work, we present HOLO-DEX, a new framework to collect demonstration data and train dexterous robots. It uses VR headsets (e.g. Quest 2) to put human teachers in an immersive virtual world. In this virtual world, the teacher can view a robotic scene from the eyes of a robot, and control it using their hands through inbuilt pose detectors. HOLO-

Correspondence to sridhar@nyu.edu.

DEX allows humans to seamlessly provide robots with high-quality demonstration data through a low-latency observational feedback system. HOLO-DEX offers three benefits: (a) Compared to self-supervised data collection methods, it allows for rapid training without reward specification as it is built on powerful imitation learning techniques; (b) Compared to Sim2Real approaches, our learned policies are directly executable on real robots since they are trained on real data; (c) Compared to other imitation approaches, it significantly reduces the need for domain expertise since even untrained humans can operate VR devices.

We experimentally evaluate HOLO-DEX on six dexterous manipulation tasks that require performing complex, contact-rich behavior. These tasks range from in-hand object manipulation to single-handed bottle opening. Across our tasks, we find that a teacher can provide demonstrations at an average of 60s per demonstration using HOLO-DEX, which is $1.8\times$ faster than prior work in single-image teleoperation [29]. On 4/6 tasks, HOLO-DEX can learn policies that achieve $> 90\%$ success rates. Surprisingly, we find that the dexterous policies learned through HOLO-DEX can generalize on new, previously unseen objects.

In summary, this work presents HOLO-DEX, a new framework for dexterous imitation learning with the following contributions. First, we demonstrate that high-quality teleoperation can be achieved by immersing human teachers in mixed reality through inexpensive VR headsets. Second, we experimentally show that the demonstrations collected by HOLO-DEX can be used to train effective, and general-purpose dexterous manipulation behaviors. Third, we analyze and ablate HOLO-DEX over various decisions such as the choice of hand tracker and imitation learning methods. Finally, we will release the mixed reality API, demonstrations collected, and training code associated with HOLO-DEX on <https://holo-dex.github.io/>.

II. RELATED WORK

Our framework builds upon several important works in robot learning, imitation learning, teleoperation and dexterous manipulation. In this section, we briefly describe prior research that is most relevant to ours.

A. Methodologies for Teaching Robots

There are several approaches one can take to teach robots. Reinforcement Learning (RL) [31, 32, 33] can train policies to maximize rewards while collecting data in an automated manner. This process often requires a roboticist to specify the reward function along with ensuring safety during self-supervised data collection [16, 15]. Furthermore, such approaches are often sample-inefficient and might require extensive simulation training for optimizing complex skills.

Simulation to Real (*Sim2Real*) approaches focus on training RL policies in simulation, followed by transferring to the real robot [23, 34, 35]. Such a methodology of robot training has received significant success owing to the improvements in modern robot simulators. Sim2Real still requires significant human involvement as every task needs to be carefully

modeled in the simulator. Moreover, even during training special techniques are required to ensure that the resulting policies can transfer to the real robot [36, 37, 38, 39].

Imitation learning approaches focus on training policies from demonstrations provided by an expert. Behavior Cloning (BC) is an offline technique that trains a policy to imitate the expert behavior in a supervised manner [25, 40, 41, 42]. Recently, non-parametric imitation approaches have shown promise in learning from fewer demonstrations [43, 44, 29]. Another set of imitation learning is Inverse Reinforcement Learning (IRL) [45, 26, 46]. Here, a reward function is inferred from demonstrations, followed by using RL to optimize the inferred reward. While HOLO-DEX is geared towards offline imitation, the demonstrations we collect are compatible with IRL approaches as well.

B. Dexterous Teleoperation Frameworks

To effectively use imitation learning for dexterous manipulation we need to obtain accurate hand poses from a human teacher. There are several approaches to gather demonstrations for dexterous tasks. Using a custom glove to measure a user's hand movements such as CyberGlove [30, 47] or Shadow Dexterous Glove [48] has been a popular solution. However, although such gloves have high accuracy, they can be expensive and require significant calibration effort. Vision-based hand pose detectors have shown promise for dexterous tasks. Some examples include using multiple RGBD [28], single depth [49], RGB [29], and RGBD [39] images. However, such methods either require custom calibration procedures [28] or suffer from occlusion-related issues when using single cameras [29]. Recently, a new generation of VR headsets has enabled advanced multi-camera hand pose detection [50] that gave promising results in [51, 52]. This enhancement provides a robust solution that is significantly cheaper compared to CyberGlove and requires little calibration. While VR tools have been used to collect demonstrations [53, 54] for low-dimensional end-effector control, HOLO-DEX shows that the VR headsets can be used for high-dimensional control in augmented reality.

C. Dexterous Manipulation

Due to its high-dimensional action space, learning complex skills with dexterous multi-fingered robot hand has been a longstanding challenge [55, 56, 57, 58]. Model-based RL and control approaches have demonstrated significant success on tasks such as spinning objects and in-hand manipulation [59, 60]. Similarly, model-free RL approaches have shown that Sim2Real can enable impressive skills such as in-hand cube rotation and Rubik's cube face turning [20, 36]. However, both learning approaches requires hand-designing reward functions along with system identification [59] or task-specific training procedures [36]. Coupled with long training times, often requiring weeks [20, 36], they make dexterous manipulation difficult to scale for general tasks.

To address the poor sample efficiency of prior learning-based methods, several works have looked at imitation learning [61, 62]. Here, given a handful of demonstrations, sim-

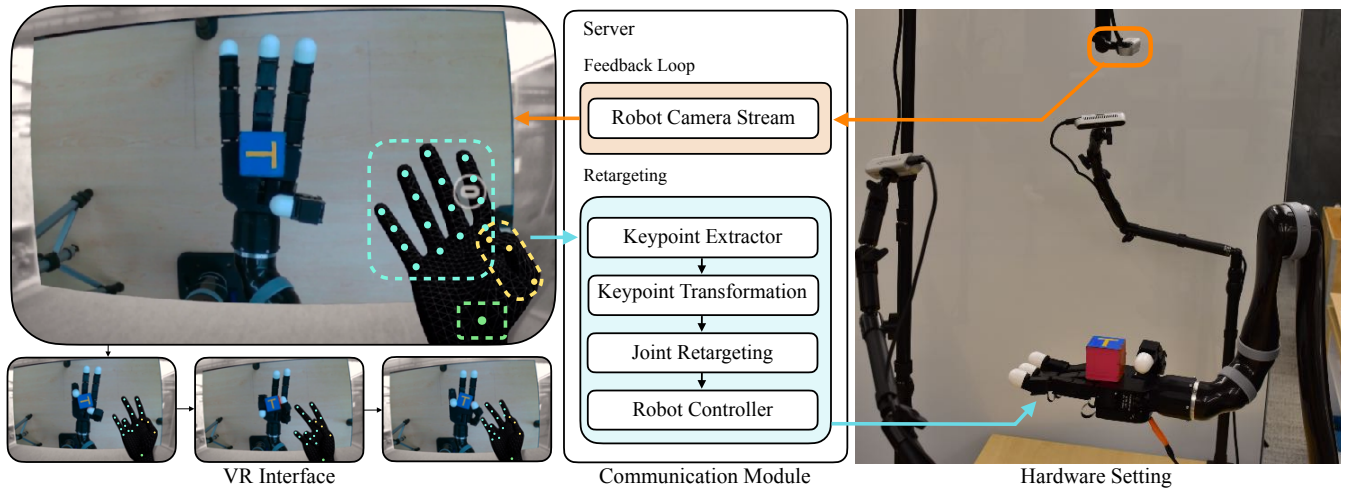


Fig. 2: Overview of HOLO-DEX’s teleoperation module. Given a hand pose in the VR interface, the controller streams the keypoint data to the robot’s server which transforms and retargets the human hand key points to the Allegro Hand. Visual feedback of the teleoperated Hand is then provided back to the VR Headset for real-time feedback.

ulated policies can be trained in a few hours. More recently, such imitation-based approaches have shown success on real robot hands [29]. HOLO-DEX takes this idea further by improving the teaching process and demonstrating its utility on a variety of in-hand manipulation tasks.

III. BACKGROUND ON VISUAL IMITATION LEARNING

To understand the imitation learning framework used in HOLO-DEX, we first formalize and describe important background work in self-supervised learning and non-parametric imitation. Together, these enable efficient imitation learning from high-dimensional visual observations.

A. Visual Self-Supervised Learning

Self-Supervised Learning (SSL) focuses on obtaining low-dimensional embeddings z from high-dimensional observations o [63]. Operationally, the observations (e.g. RGB images) are fed into an encoder f_θ , where θ denotes the weights of a parametric deep network. While there are several methods to train f_θ , the central principle in many works is to predict one ‘view’ of the observation given a different ‘view’ of the same observation. One example of such a learning scheme is data-augmented SSL. Here, the observation o is augmented by applying visual augmentations such as color jitter or random grayscale. Given two augmented views of this observation o^1 and o^2 , the corresponding embeddings would be $z^1 \equiv f_\theta(o^1)$ and $z^2 \equiv f_\theta(o^2)$. The training objective for f_θ amounts to maximizing the mutual information between the two embeddings $I(z^1, z^2)$.

To optimize this objective, we use the BYOL [64] training scheme, which amounts to predicting $z^2 \leftarrow g_\phi(z^1)$ through a small deep model g_ϕ called the ‘projector’. This scheme for learning embeddings has had significant success in a variety of domains ranging from computer vision, audio processing, and robotics. Given its simplicity, we use BYOL to obtain concise embeddings from our demonstrated data.

B. Non-Parametric Imitation Learning

In our framework for imitation learning we have access to expert demonstrations in the form of $\mathcal{D}^E \equiv \{(o_{t'}^E, s_{t'}^E, a_{t'}^E)\}$, where $o_{t'}^E$ represents the sensory observation at time t' , $s_{t'}^E$ represents the robot state, and $a_{t'}^E$ denotes the robot action taken. Note that $s_{t'}^E$ does not contain information about the object that is being manipulated. Hence, object information needs to be inferred from observations $o_{t'}^E$. Given these demonstrations, we would like to learn a policy $\pi(a_t|o_t)$ that follows the expert behavior \mathcal{D}^E . While there are several strategies for optimizing π , we resort to non-parametric approaches given their superior performance in low-data regimes [44, 29].

Our non-parametric control framework follows VINN [44], where given the observations o^E from the expert demonstration dataset \mathcal{D}^E , a BYOL encoder f_θ is trained. Next, the observations in the dataset are all converted to embeddings, i.e. $\{o_{t'}^E\} \xrightarrow{f_\theta} \{z_{t'}^E\}$. During run time, when the robot receives an observation o_t it is embedded to z_t . Then the Nearest-Neighbor (NN) example in $\{(z_{t'}^E, s_{t'}^E, a_{t'}^E)\}$ is selected to be imitated. We denote this NN example as $\{(z_{t^*}^E, s_{t^*}^E, a_{t^*}^E)\}$. Given a small dataset \mathcal{D}^E , which is often the case in robotic applications, this NN-based imitation learning provides effective learning compared to parametric approaches such as BC.

IV. HOLO-DEX

As seen in Fig. 1, HOLO-DEX operates in two phases. In the first phase, a human teacher uses a Virtual Reality (VR) headset to provide demonstrations to a robot. This phase consists of creating a virtual world for teaching, estimating hand poses from the teacher, retargeting the teacher’s hand pose to the robot’s hand and finally controlling the robot hand. After a handful of demonstrations are collected in phase one, the second phase of HOLO-DEX learns visual policies to solve the demonstrated tasks. In this section, we will describe each sub-component in detail.




Fig. 3: Demonstration collection process for the Object Flipping task. The first row shows the user’s perspective inside the VR Headset and the second row shows the corresponding robot pose.

A. Placing an Operator in a Virtual World

We use the Meta Quest 2 VR headset to place human teachers in a virtual world. The headset surrounds the human in a virtual environment at a resolution of 1832×1920 and a refresh rate of 72 Hz. The base version of this headset is affordable at \$399 and is relatively light at 503g. These features allow for comfortable operation by the teacher. Importantly, the API interface of the Quest 2 allows for creating custom mixed reality worlds that visualizes the robotic system along with diagnostic panels in VR. Examples of virtual scenes are depicted in Fig. 2 and Fig. 3.

B. Hand Pose Estimation with VR Headsets

In contrast to prior work on dexterous teleoperation, using VR headsets provides three benefits with regard to hand pose estimation of the human teacher. First, since the Quest 2 uses 4 monochrome cameras, its hand-pose estimator [50] is significantly more robust compared to single camera estimators [65]. Second, since the cameras are internally calibrated, they do not require specialized calibration routines that are needed in prior multi-camera teleoperation frameworks [28]. Third, since the hand pose estimator is integrated into the device, it can stream real-time poses at 72Hz. As noted in prior work [28, 29], a significant challenge in dexterous teleoperation is obtaining hand poses at both high accuracy and a high frequency. HOLO-DEX significantly simplifies this problem by using commercial-grade VR headsets.

C. Human to Robot Hand Pose Retargeting

Once we have extracted the teacher’s hand pose from VR, we will need to retarget it to the robotic hand. This is done by first computing the individual hand joint angles in the teacher’s hand. Given these joint angles, a straightforward method of retargeting is to directly command the robot’s joints to the corresponding angle. In practice, this works well for all fingers except the thumb. The thumb presents a unique challenge to our Allegro robot hand since its morphology does not match a human’s hand. To address this, we map the spatial coordinates of the teacher’s thumb fingertip to the robot’s thumb fingertip. The joint angles of the thumb are then computed through an inverse kinematics solver.

Since the Allegro hand does not have a pinky finger, we ignore the teacher’s pinky joints. The overall pose retargeting procedure does not require any calibration or user-specific tuning to collect demonstrations. However, we find that

thumb retargeting can be improved by finding user-specific maps from their thumb to the robot’s thumb. This entire procedure is computationally inexpensive and can stream desired robot hand poses at 60 Hz.

D. Robot Hand Control

Our Allegro Hand is controlled asynchronously over a ROS [66] communication framework. Given desired robot joint positions that were computed from the retargeting procedure, we use a PD controller to output desired torques at 300Hz. To reduce steady-state error, we use a gravity compensation module to compute offset torques. On latency tests, we find that when the VR headset is on the same local network as the robot hand, we achieve latency under 100 milliseconds. Having a low error and latency is crucial for HOLO-DEX since it allows for intuitive teleoperation of the robot hand by the human teacher.

As the human teacher controls the robot hand, they can see the robot change in real time (60Hz). This allows the teacher to correct execution errors in the robot. During the teaching process, we record observational data from three RGBD cameras and the action information of the robot at 5Hz. We had to reduce the recording frequency due to the large data footprint and associated bandwidth required from recording multiple cameras.

E. Imitation Learning with HOLO-DEX Data

Once data is collected in the first phase of HOLO-DEX, we now proceed to the second phase, where visual policies are trained on top of this data. We employ the Imitation with Nearest Neighbors (INN) algorithm for learning. Background details of INN is present in Section III-B. In prior work, INN was shown to produce state-based dexterous policies on the Allegro hand [29]. HOLO-DEX takes this a few steps further and demonstrates that these visual policies can generalize to novel objects in a variety of dexterous manipulation tasks.

To select the learning algorithm for obtaining low-dimensional embeddings (see Section III-A), we experiment with several state-of-the-art self-supervised learning algorithms [64, 67, 68, 69] and find that BYOL [64] provides the best nearest neighbour results. Hence we select BYOL as our base self-supervised learning method. Once BYOL is trained on the collected demonstrations, we can run dexterous manipulation policies on the robot by performing nearest neighbor action retrieval (see Section III-B) to get the closest example in the trainset $(z_{t^*}^E, a_{t^*}^E)$. To account for the slower




Fig. 4: Successful rollouts of policies trained through HOLO-DEX on our six dexterous tasks.

rate of demonstration collection, we set the action to the difference between the succeeding state to the closest neighbor and the current state, i.e. $a_t = (s_{t^*+k}^E) - (s_t^E)$. Here k represents the number of states skipped during our recording of demonstrations. Note that directly commanding $a_{t^*}^E$ would fail due to our asynchronous data storage framework.

V. EXPERIMENTAL EVALUATION

Our experiments and tasks are designed to answer the following questions:

- How long does it take HOLO-DEX to collect demonstrations?
- How successful are policies trained by HOLO-DEX?
- How general are the skills learned by HOLO-DEX?
- How many demonstrations from HOLO-DEX are required to successfully solve dexterous tasks?

A. Dexterous Manipulation Tasks

We study six dexterous manipulation tasks that require contact-rich, multi-fingered control for successful completion. Details of these tasks are described below.

- 1) *Planar Rotation*: Given an object placed at a random position on the palm of the robot hand, the goal is to rotate the object in the counter-clockwise direction along the palm normal vector. The task is considered a success if the object is rotated by 90° under a minute.
- 2) *Object Flipping*: Given an object placed at a random position on the palm of the robot, the goal is to flip the object in the hand’s direction. The task is considered a success if the object is flipped within a minute.
- 3) *Can Spinning*: Given a large can placed horizontally on the palm of the robot hand, the task is to spin the can in the counter-clockwise direction along the palm normal vector. The task is considered a success if the can spins by 90° under 30 seconds.
- 4) *Bottle Opening*: Given a bottle on a desk, the task is to single-handedly grab the bottle and open it. The task is considered a success when the cap rotates 360° .

TABLE I: Average time taken in seconds to collect a single demonstration on our Allegro hand using HOLO-DEX and DIME [29]

Task	DIME		HOLO-DEX	
	Expert	New User	Expert	New User
Planar Rotation	60	150	30	125
Object Flipping	6	21	5	6
Can Spinning	15	76	10	68
Bottle Opening	N/A	N/A	30	48
Card Sliding	N/A	N/A	150	N/A
PostIt Note Sliding	N/A	N/A	120	N/A

5) *Card Sliding*: Given a card on a desk in front of the hand, the task is to grab the card by sliding it off the table and picking it up. The task is considered a success when the card lifts off from the table and is stably grasped by the hand.

6) *PostIt Note Sliding*: This task is similar to card sliding, but instead of a card we use a thicker post-it note pad as the object to pick from the desk.

For the *Planar Rotation* task we collect 120 expert demonstrations, while for all the other tasks we collect 30 demonstrations. Additional demonstrations for *Planar Rotation* are collected to account for the relative difficulty of this task and experimentation with different dataset sizes.

B. How long does it take to collect demonstrations?

The closest dexterous teleoperation work to ours that uses commodity sensors is single-image teleoperation (e.g. DIME [29]), where hand poses are detected via RGB images to get robot joint angles. Despite its simplicity, such single-image pose estimation [65] suffers from hand occlusions, which results in poor teleoperation performance on challenging manipulation tasks [28]. In Table I we show that HOLO-DEX can collect successful demonstrations $1.8\times$ faster compared to DIME. For 3/6 tasks that require precise 3D movements, we find that single-image teleoperation is insufficient to collect even a single demonstration.

To demonstrate the versatility of HOLO-DEX we ask five



Fig. 5: On the left, we depict the start and end states of objects present in demonstration data. On the right, we depict the states produced by running our policies on objects not seen in demonstrations. Green boxes denote a successful rollout, while red boxes denote a failure.

TABLE II: Success rates on our Allegro hand using HOLO-DEX

Task	BC	BC - Rep	VINN	VINN (New Objects)
Planar Rotation	0/10	0/10	10/10	35/50
Object Flipping	0/10	0/10	9/10	50/50
Can Spinning	0/10	0/10	10/10	45/50
Bottle Opening	0/10	0/10	10/10	N/A
Card Sliding	0/10	0/10	3/10	N/A
PostIt Note Sliding	0/10	0/10	3/10	N/A

untrained users to collect demonstrations for each of our tasks. Unlike prior work [29, 28], no user-specific calibration was done for this evaluation. In Table I, we see that these users are successfully able to solve 4/6 tasks on their first try, failing only on more difficult sliding tasks. We also find that training on this system is quite important as it yields a nearly $2.6\times$ speedup in demonstration collection.

C. How successful are policies trained by HOLO-DEX?

We examine the performance of various imitation learning policies on all dexterous tasks. The imitation learning algorithms include Behavior Cloning (BC) [25], Behavior Cloning from pretrained representations (BC-Rep) [70] and VINN [44]. Table II shows the success rates of each task with different policies. We find that VINN outperforms both Behavior Cloning algorithms on all tasks. This is in line with prior work [44, 29] and showcases the effectiveness of non-parametric imitation with few demonstrations. However, we find that for the two tasks that involve sliding and picking, the performance of VINN is quite low at 30%. We believe this is due to our robot’s inability to sense touch, which limits our vision-only model from performing precise actions.

D. How general are the policies learned by HOLO-DEX?

Since our policies are vision-based and do not require explicitly estimating the states of objects, they are compatible with objects not seen in training. To understand the generalization capabilities of our models, we evaluate three of our policies (Planar Rotation, Object Flipping, and Can Spinning) on 10 visually and geometrically distinct objects, with 5 rollouts for every object in different initial positions. Results for this experiment are in Table II and are visualized in Fig. 5. Interestingly, for all three tasks, we find high success rates without any additional demonstration collection or training. This means that although we collect

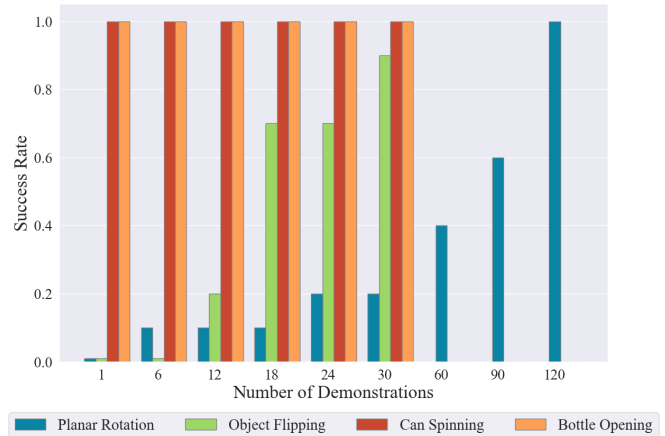


Fig. 6: Effect of varying demonstration data with HOLO-DEX.

demonstrations from HOLO-DEX on a single object, the learned policies can generalize in a zero-shot manner.

E. How many demonstrations are needed to solve our tasks?

In Fig. 6 we visualize the performance on four of our tasks across different dataset sizes. To decouple the effects of representation learning with action prediction, we use the same encoder (trained on all task data) for different dataset splits. We find that for Can Spinning and Bottle Opening, a single demonstration is sufficient to achieve high performance, while for Planar Rotation and Object Flipping we see steady gains in performance as we increase the amount of demonstration data.

VI. LIMITATIONS AND DISCUSSION

We have presented HOLO-DEX, a framework that takes some of the first steps towards immersive teaching of dexterous robots through VR. There are currently two limitations of this work. First, we find that for the harder manipulation tasks, such as *sliding a card* our learned policies achieve poor performance. Integrating tactile sensing to HOLO-DEX could remedy this issue. Second, our retargeting procedure only applies to robots that can map to human joints. This limits its applicability to robots with different morphologies (e.g. aerial robots, quadrupeds, etc.). Future research on UX design and retargeting mechanisms can enable mapping VR control to more complex end-effectors.

REFERENCES

- [1] C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Iterative residual policy: for goal-conditioned dynamic manipulation of deformable objects," *arXiv preprint arXiv:2203.00663*, 2022.
- [2] A. E. Tekden, A. Erdem, E. Erdem, M. Imre, M. Y. Seker, and E. Ugur, "Belief regulated dual propagation nets for learning action effects on groups of articulated objects," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10 556–10 562.
- [3] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control," *IEEE Transactions on Robotics*, 2022.
- [4] Y. Ma, F. Farshidian, T. Miki, J. Lee, and M. Hutter, "Combining learning-based locomotion policy with model-based manipulation for legged mobile manipulators," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2377–2384, 2022.
- [5] L. Smith, I. Kostrikov, and S. Levine, "A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning," *arXiv preprint arXiv:2208.07860*, 2022.
- [6] T. Zhang, G. Kahn, S. Levine, and P. Abbeel, "Learning deep control policies for autonomous aerial vehicles with mpc-guided policy search," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 528–535.
- [7] D. Gandhi, L. Pinto, and A. Gupta, "Learning to fly by crashing," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3948–3955.
- [8] J. Hwangbo, I. Sa, R. Siegwart, and M. Hutter, "Control of a quadrotor with reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2096–2103, 2017.
- [9] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, *et al.*, "Dota 2 with large scale deep reinforcement learning," *arXiv preprint arXiv:1912.06680*, 2019.
- [10] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [11] J. Guo, S. Lu, H. Cai, W. Zhang, Y. Yu, and J. Wang, "Long text generation via adversarial training with leaked information," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [12] M. Huang, F. Li, W. Zou, and W. Zhang, "Sarg: A novel semi autoregressive generator for multi-turn incomplete utterance restoration," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 14, 2021, pp. 13 055–13 063.
- [13] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [14] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4367–4375.
- [15] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," *ICRA*, 2016.
- [16] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *ISER*, 2016.
- [17] A. Singh, L. Yang, K. Hartikainen, C. Finn, and S. Levine, "End-to-end robotic reinforcement learning without reward engineering," *arXiv preprint arXiv:1904.07854*, 2019.
- [18] M. Vecerik, O. Sushkov, D. Barker, T. Rothörl, T. Hester, and J. Scholz, "A practical approach to insertion with variable socket position using deep reinforcement learning," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 754–760.
- [19] G. Dulac-Arnold, N. Levine, D. J. Mankowitz, J. Li, C. Paduraru, S. Goyal, and T. Hester, "An empirical investigation of the challenges of real-world reinforcement learning," *arXiv preprint arXiv:2003.11881*, 2020.
- [20] OpenAI, M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba, "Learning dexterous in-hand manipulation," *arXiv*, 2018.
- [21] T. Chen, J. Xu, and P. Agrawal, "A system for general in-hand object re-orientation," *Conference on Robot Learning*, 2021.
- [22] 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.
- [23] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *IROS*, 2017.
- [24] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *Advances in Artificial Life*, F. Morán, A. Moreno, J. J. Merelo, and P. Chacón, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 704–720.
- [25] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," in *NeurIPS*, 1989, pp. 305–313.
- [26] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *ICML*, 2004, p. 1.
- [27] 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," *RSS*, 2018.
- [28] 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)*, 2020, pp. 9164–9170.
- [29] S. P. Arunachalam, S. Silwal, B. Evans, and L. Pinto, "Dexterous imitation made easy: A learning-based framework for efficient dexterous manipulation," *arXiv preprint arXiv:2203.13251*, 2022.
- [30] M. Caeiro-Rodríguez, I. Otero-González, F. A. Mikic-Fonte, and M. Llamas-Nistal, "A systematic review of commercial smart gloves: Current status and applications," *Sensors*, 2021.
- [31] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: A survey," *JAIR*, 1996.
- [32] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint*, 2015.
- [33] D. Yarats, R. Fergus, A. Lazaric, and L. Pinto, "Mastering visual continuous control: Improved data-augmented reinforcement learning," *arXiv preprint arXiv:2107.09645*, 2021.
- [34] F. Sadeghi and S. Levine, "Cad2rl: Real single-image flight without a single real image," *arXiv preprint arXiv:1611.04201*, 2016.
- [35] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, "Asymmetric actor critic for image-based robot learning," *RSS*, 2018.
- [36] OpenAI, I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Welinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang, "Solving rubik's cube with a robot hand," *arXiv*, 2019.
- [37] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to manipulate deformable objects without demonstrations,"

arXiv preprint, 2019.

- [38] K. Rao, C. Harris, A. Irpan, S. Levine, J. Ibarz, and M. Khansari, “RI-cycleGAN: Reinforcement learning aware simulation-to-real,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 157–11 166.
- [39] Y. Qin, H. Su, and X. Wang, “From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation,” *arXiv preprint arXiv:2204.12490*, 2022.
- [40] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” 2021.
- [41] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, “End to end learning for self-driving cars,” *arXiv:1604.07316*, 2016.
- [42] S. Young, D. Gandhi, S. Tulsiani, A. Gupta, P. Abbeel, and L. Pinto, “Visual imitation made easy,” 2020.
- [43] S. Schaal and C. Atkeson, “Robot juggling: implementation of memory-based learning,” *IEEE CSM*, 1994.
- [44] J. Pari, N. M. Shafiqullah, S. P. Arunachalam, and L. Pinto, “The surprising effectiveness of representation learning for visual imitation,” 2021.
- [45] U. Syed and R. E. Schapire, “A game-theoretic approach to apprenticeship learning,” in *NerulPS*, J. Platt, D. Koller, Y. Singer, and S. Roweis, Eds., 2007.
- [46] S. Haldar, V. Mathur, D. Yarats, and L. Pinto, “Watch and match: Supercharging imitation with regularized optimal transport,” *arXiv preprint arXiv:2206.15469*, 2022.
- [47] V. Kumar and E. Todorov, “Mujoco haptix: A virtual reality system for hand manipulation,” in *Humanoids*, 2015.
- [48] 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/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 900–10 906.
- [49] 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.
- [50] S. Han, B. Liu, R. Cabezas, C. D. Twigg, P. Zhang, J. Petkau, T.-H. Yu, C.-J. Tai, M. Akbay, Z. Wang, A. Nitzan, G. Dong, Y. Ye, L. Tao, C. Wan, and R. Wang, “Megatrack: Monochrome egocentric articulated hand-tracking for virtual reality,” 2020.
- [51] T. Hentschel and J. A. Neuhöfer, “Steady hands - an evaluation on the use of hand tracking in virtual reality training in nursing,” in *Proceedings of the 21st Congress of the International Ergonomics Association (IEA 2021)*, N. L. Black, W. P. Neumann, and I. Noy, Eds. Cham: Springer International Publishing, 2022, pp. 643–649.
- [52] M. Salvato, N. Heravi, A. M. Okamura, and J. Bohg, “Predicting hand-object interaction for improved haptic feedback in mixed reality,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3851–3857, 2022.
- [53] Z. Gharaybeh, H. Chizeck, and A. Stewart, “Telerobotic control in virtual reality,” in *OCEANS 2019 MTS/IEEE SEATTLE*, 2019, pp. 1–8.
- [54] 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 *ICRA*, 2018.
- [55] V. Kumar, Y. Tassa, T. Erez, and E. Todorov, “Real-time behaviour synthesis for dynamic hand-manipulation,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 6808–6815.
- [56] I. Mordatch, Z. Popović, and E. Todorov, “Contact-invariant optimization for hand manipulation,” in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA ’12. Goslar, DEU: Eurographics Association, 2012, p. 137–144.
- [57] R. Deimel and O. Brock, “A novel type of compliant and underactuated robotic hand for dexterous grasping,” *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 161–185, 2016.
- [58] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, “Learning ambidextrous robot grasping policies,” *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [59] V. Kumar, A. Gupta, E. Todorov, and S. Levine, “Learning dexterous manipulation policies from experience and imitation,” *arXiv*, 2016.
- [60] A. Nagabandi, K. Konoglie, S. Levine, and V. Kumar, “Deep dynamics models for learning dexterous manipulation,” *arXiv*, 2019.
- [61] 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,” in *RSS*, 2018.
- [62] I. Radosavovic, X. Wang, L. Pinto, and J. Malik, “State-only imitation learning for dexterous manipulation,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 7865–7871.
- [63] L. Ericsson, H. Gouk, C. C. Loy, and T. M. Hospedales, “Self-supervised representation learning: Introduction, advances, and challenges,” *IEEE Signal Processing Magazine*, vol. 39, no. 3, pp. 42–62, 2022.
- [64] J.-B. Grill, F. Strub, F. Althé, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar, *et al.*, “Bootstrap your own latent—a new approach to self-supervised learning,” *NeurIPS*, 2020.
- [65] F. Zhang, V. Bazarevsky, A. Vakunov, A. Tkachenka, G. Sung, C.-L. Chang, and M. Grundmann, “Mediapipe hands: On-device real-time hand tracking,” 2020.
- [66] Stanford Artificial Intelligence Laboratory *et al.*, “Robotic operating system.” [Online]. Available: <https://www.ros.org>
- [67] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” *arXiv preprint*, 2020.
- [68] X. Chen, S. Xie, and K. He, “An empirical study of training self-supervised vision transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9640–9649.
- [69] A. Bardes, J. Ponce, and Y. LeCun, “Vicreg: Variance-invariance-covariance regularization for self-supervised learning,” *arXiv preprint arXiv:2105.04906*, 2021.
- [70] S. Young, J. Pari, P. Abbeel, and L. Pinto, “Playful interactions for representation learning,” *arXiv preprint arXiv:2107.09046*, 2021.