

DexMV: Imitation Learning for Dexterous Manipulation from Human Videos

Yuzhe Qin*, Yueh-Hua Wu*, Shaowei Liu, Hanwen Jiang,
Ruihan Yang, Yang Fu, and Xiaolong Wang

University of California San Diego, La Jolla 92093, USA




Fig. 1: We record human videos on manipulation tasks (1st row) and perform 3D hand-object pose estimations from the videos (2nd row) to construct the demonstrations. We have a paired simulation system providing the same dexterous manipulation tasks for the multi-finger robot (3rd row), including *relocate*, *pour*, and *place inside*, which we can solve using imitation learning with the inferred demonstrations.

Abstract. While significant progress has been made on understanding hand-object interactions in computer vision, it is still very challenging for robots to perform complex dexterous manipulation. In this paper, we propose a new platform and pipeline DexMV (**D**exterous **M**anipulation from **V**ideos) for imitation learning. We design a platform with: (i) a simulation system for complex dexterous manipulation tasks with a multi-finger robot hand and (ii) a computer vision system to record large-scale demonstrations of a human hand conducting the same tasks. In our novel pipeline, we extract 3D hand and object poses from videos, and propose a novel demonstration translation method to convert human motion to robot demonstrations. We then apply and benchmark multiple imitation learning algorithms with the demonstrations. We show that the demonstrations can indeed improve robot learning by a large margin and solve the complex tasks which reinforcement learning alone cannot solve. Code and videos are available at <https://yzqin.github.io/dexmv>

Keywords: Dexterous Manipulation; Learning from Human Demonstration; Reinforcement Learning

* Equal Contribution

1 Introduction

Dexterous manipulation of objects is the primary means for humans to interact with the physical world. Humans perform dexterous manipulation in everyday tasks with diverse objects. To understand these tasks, in computer vision, there is significant progress on 3D hand-object pose estimation [28,79] and affordance reasoning [13,72]. While computer vision techniques have greatly advanced, it is still very challenging to equip robots with human-like dexterity. Recently, there has been a lot of effort on using reinforcement learning (RL) for dexterous manipulation with an anthropomorphic robot hand [52]. However, given the high Degree-of-Freedom joints and nonlinear tendon-based actuation of the multi-finger robot hand, it requires a *large amount* of training data with RL. Robot hands trained using only RL will also adopt *unnatural* behavior. Given these challenges, can we leverage humans’ experience in the interaction with the physical world to guide robots, with the help of computer vision techniques?

One promising avenue is imitation learning from human demonstrations [60,65]. Particularly, for dexterous manipulation, Rajeswaran *et al.* [60] introduces a simulation environment with four different manipulation tasks and paired sets of human demonstrations collected with a Virtual Reality (VR) headset and a motion capture glove. However, data collection with VR is relatively high-cost and not scalable, and there are only 25 demonstrations collected for each task. It also limits the complexity of the task: It is shown in [58] that RL can achieve similar performance with or without the demonstrations in most tasks proposed in [60]. Instead of focusing on a small scale of data, we look into increasing the difficulty and complexity of the manipulation tasks with diverse daily objects. This requires large-scale human demonstrations which are hard to obtain with VR but are much more available from human videos.

In this paper, we propose **a new platform and a novel imitation learning pipeline** for benchmarking complex and generalizable dexterous manipulation, namely DexMV (**D**exterous **M**anipulation from **V**ideos). We introduce new tasks with the multi-finger robot hand (Adroit Robotic Hand [42]) on diverse objects in simulation. We collect real human hand videos performing the same tasks as demonstrations. By using human videos instead of VR, it *largely reduces the cost* for data collection and allows humans to perform more *complex and diverse* tasks. While the video demonstrations might not be optimal for perfect imitation (e.g., behavior cloning) to learn successful policies, the diverse dataset is beneficial for augmenting the training data for RL, which can learn from both successful and unsuccessful trials.

Our DexMV platform contains a paired systems with: (i) A computer vision system which records the videos of human performing manipulation tasks (1st row in Figure 1); (ii) A physical simulation system which provides the interactive environments for dexterous manipulation with a multi-finger robot (3rd row in Figure 1). The two systems are aligned with the same tasks. With this platform, our goal is to bridge 3D vision and robotic dexterous manipulation via a novel imitation learning pipeline.

Our DexMV pipeline contains three stages. First, we extract the 3D hand-object poses from the recorded videos (2nd row in Figure 1). Unlike previous imitation learning studies with 2-DoF grippers [80,70], we need the human video to guide the 30-DoF robot hand to move each finger in 3D space. Parsing the 3D structure provides critical and necessary information. Second, a **key contribution** in our pipeline is a novel demonstration translation method that connects the computer vision system and the simulation system. We propose an optimization-based approach to convert 3D human hand trajectories to robot hand demonstrations. Specifically, the innovations lie in 2 steps: (i) Hand motion retargeting approach to obtain robot hand states; and (ii) Robot action estimation to obtain the actions for learning. Third, given the robot demonstrations, we perform imitation learning in the simulation tasks. We investigate and benchmark algorithms which augment RL objectives with state-only [58] and state-action [31,60] demonstrations.

We experiment with three types of challenging tasks with the YCB objects [15]. The first task is to *relocate* an object to a goal position. Instead of relocating a single ball as [60], we increase the task difficulty by using diverse objects (first 5 columns in Figure 1). The second task is *pour*, which requires the robot to pour the particles from a mug into a container (Figure 1 from column 6). The third task is *place inside*, where the robot hand needs to place an object into a container (last 2 columns in Figure 1). In our experiments, we benchmark different imitation learning algorithms and show human demonstrations improve task performance by a large margin. More interestingly, we find the learned policy with our demonstrations can even generalize to unseen instances within the same category or outside the category. We highlight our contributions as follows:

- DexMV platform for learning dexterous manipulation using human videos. It contains paired computer vision and simulation systems with multiple dexterous complex manipulation tasks.
- DexMV pipeline to perform imitation learning, with a key innovation on demonstration translation to convert human videos to robot demonstrations.
- With DexMV platform and pipeline, we largely improve the dexterous manipulation performance over multiple complex tasks, and its generalization ability to unseen object instances.

2 Related Work

Dexterous Manipulation. Manipulation with multi-finger hands is one of the most challenging robotics tasks, which has been actively studied with optimization and planning [63,10,50,20,3,7]. Recently, researchers have started exploring reinforcement learning for dexterous manipulation [52,51]. However, training with only RL requires huge data samples. Different from a regular 2-DoF robot gripper, the Adroit Robotic Hand used in our experiment has 30 DoFs, which greatly increases the optimization space.

Imitation Learning from Human Demonstrations. Imitation learning is a promising paradigm for robot learning. It is not limited to behavior

cloning [55,8,62,11,76] and inverse reinforcement learning [64,49,1,31,23,5,77,45], but also augmenting the RL training with demonstrations [54,21,78,60,58]. For example, Rajeswaran *et al.* [60] propose to incorporate demonstrations with on-policy RL. However, all these approaches rely on expert demonstrations collected using expert policies or VR, which is not scalable and generalizable to complex tasks. In fact, it is shown in [58] that RL can achieve similar performance and sample efficiency with or without the demonstrations in most dexterous manipulation tasks in [60] besides *relocate*. Going beyond a VR setup, researchers have recently explored imitation learning and RL with videos [66,17,65,69,70,80]. However, all tasks are relatively simple (e.g., pushing a block) with a parallel gripper. In this paper, we propose new challenging dexterous manipulation tasks. We leverage pose estimation for providing demonstrations that largely improve imitation learning performance, while RL alone fails to solve these tasks.

Hand-Object Interaction. Hand-object interaction is a widely studied topic. One line of work focus on object pose estimation [39,57,79,73,53,33,29] and hand pose estimation [81,34,71,25,6,12,28,41,46]. For example, Berk *et al.* [15] propose YCB dataset with real objects and their 3D scans. Beyond object poses, datasets and methods for joint estimation of the hand-object poses are also proposed [24,26,46,18]. Another line of work focus on hand-object contact reasoning [13,72], which can be used for functional grasps [35,14,47]. Recently, researchers explored to use hand pose estimation to control a robot hand [27,4,44]. They require motion retargeting to map the hand pose into robot motion. These work focus on teleoperation where no objects are involved. In DexMV, we consider the demonstration translation setting, which converts a sequence of hand-object poses into robot demonstrations. We show that the translated demonstration is beneficial for imitation learning for dexterous manipulation.

3 Overview

We propose DexMV for **D**exterous **M**anipulation with imitation learning from human **V**ideos including:

(i) **DexMV platform**, which offers a paired computer vision system recording human manipulation videos and physical simulators conducting the same tasks. We design three challenging tasks in the simulator and use computer vision system to collect human demonstrations for these tasks in real world. The videos can be efficiently collected with around *100 demonstrations per hour*.

(ii) **DexMV pipeline**, a new pipeline for imitation learning from human videos. Given the recorded videos from the computer vision system, we perform pose estimation and motion retargeting to translate human video to robot hand demonstrations. These demonstrations are then used for policy learning where multiple imitation learning and RL algorithms are investigated. By learning from demonstrations, our policy can be deployed in environments with *different goals and object configurations*, instead of just following one trajectory.

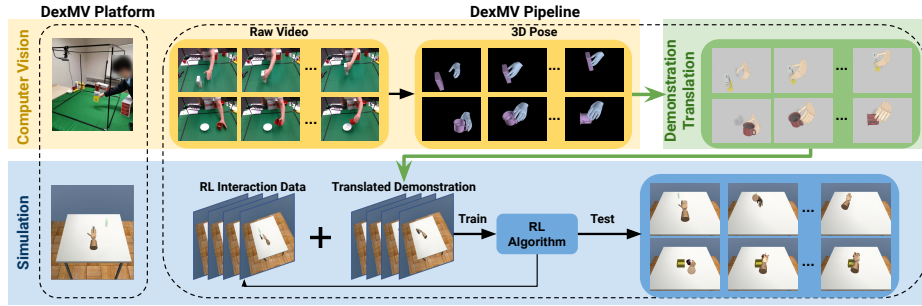


Fig. 2: **DexMV platform and pipeline overview.** Our platform is composed of a computer vision system (yellow), a simulation system (blue), and a demonstration translation module (colored with green). In computer vision system, we collect human manipulation videos. In simulation, we design the same tasks for the robot hand. We apply 3D hand-object pose estimation from videos, followed by demonstration translation to generate robot demonstrations, which are then used for imitation learning.

4 DexMV Platform

Our DexMV platform is shown in Figure 2. It is composed of a computer vision system and a simulation system.

Computer Vision System. The computer vision system is used to collect human demonstration videos on manipulating diverse real objects. In this system, we build a cubic frame (35 inch³) and attach two RealSense D435 cameras (RGBD cameras) on the top front and top left, as shown on the top row of Figure 2. During data collection, a human will perform manipulation tasks inside the frame, e.g., relocate sugar box. The manipulation videos will be recorded using the two cameras (from a front view and a side view).

Simulation System. Our simulation system is built on MuJoCo [74] with the Adroit Hand [42]. We design multiple dexterous manipulation tasks aligned with human demonstrations. As shown in the bottom row of Figure 2, we perform imitation learning by augmenting the Reinforcement Learning with the collected demonstrations. Once the policy is trained, it can be tested on the same tasks with different goals and object configurations.

Task Description. We propose 3 types of manipulation tasks with different objects. We select the YCB objects [15] with a reasonable size for human to manipulate. For all tasks, the state is composed of robot joint readings, object pose. For *relocate*, we also include target position. The action is 30-d control command of the position actuator for each joint.

Relocate. It requires the robot hand to pick up an object on the table to a target location. It is inspired by the hardest task in [60] where the robot hand needs to grasp a sphere and move it to the target position. We further increase the task difficulty by using 5 complex objects as shown in the table above and visualized in the first 5 columns of Figure 1. The transparent green shape repre-

sents the goal, which can change during training and testing (goal-conditioned). The task is successfully solved if the object reaches the goal, without the need for a specific orientation. We train one policy for relocating each object.

Pour. It requires the robot hand to reach the mug and pour the particles inside into a container (Figure 1 column 6). The robot needs to grasp the mug and then manipulate it precisely. The evaluation criteria is based on the percentage of particles poured inside the container.

Place Inside. It requires the robot to pick up an object (e.g., a banana) and place it into a container. The robot needs to rotate the object to a suitable orientation and approach the container carefully to avoid collisions. The evaluation is based on how much percentage of the object mesh volume is inside the container.

Besides testing on the trained objects, we also evaluate the generalization ability of the trained policies on unseen object instances, within the training categories like can, bottle, mug, and also outside the training categories such as camera from ShapeNet dataset [16].

5 Pose Estimation

5.1 Object Pose Estimation

We use the 6-DoF pose to represent the location and orientation of objects, which contains translation $T \in \mathbf{R}^3$ and rotation $R \in \mathbf{SO}(3)$. For each frame t in the video, we use the PVN3D [29] model trained on the YCB dataset [15] to detect objects and estimate 6-DoF poses. By taking both the RGB image and the point clouds as inputs, the model first estimates the instance segmentation mask. With dense voting on the segmented point clouds, the model then predicts the 3D location of object key points. The 6-DoF object pose is optimized by minimizing the PnP matching error.

5.2 Hand Pose Estimation

We utilize the MANO model [61] to represent the hand that consists of hand pose parameter θ_t for 3D rotations of 15 joints and root global pose r_t , and shape parameters β_t for each frame t . The 3D hand joints can be computed using hand kinematics function $j_t^{3d} = \mathbf{J}(\theta_t, \beta_t, r_t)$.

Given a video, we use the off-the-shelf skin segmentation [40] and hand detection [68] models to obtain a hand mask M_t . We use the trained hand pose estimation models in [46] to predict the 2D hand joints j_t^{2d} and the MANO parameters θ_t and β_t for every frame t using RGB image. We estimate the root joint r_t using the center of the depth image masked by M_t . Given the initial estimation θ_t, β_t, r_t of each frame t and the camera pose Π , we formulate the 3D hand joint estimation as an optimization problem,

$$\theta_t^*, \beta_t^*, r_t^* = \arg \min_{\theta_t, \beta_t, r_t} \|\Pi J(\theta_t, \beta_t, r_t) - j_t^{2d}\|^2 + \lambda \|M_t \cdot (\mathbf{R}(\theta_t, \beta_t, r_t) - D_t)\|^2 \quad (1)$$

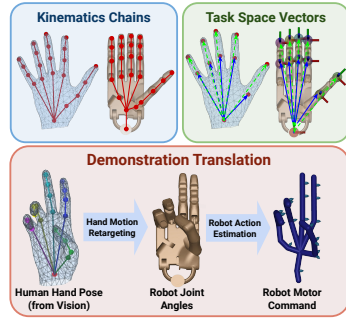


Fig. 3: **Top row:** Kinematic Chains and Task Space Vectors (TSV). The TSV (dash arrows) are ten vectors to be matched between both hands. **Bottom row:** (i) Hand motion retargeting; (ii) Action Estimation.




Fig. 4: Visualization of hand motion retargeting results. **Top row:** 3D hand-object poses. **Mid row:** Retargeting results using Finger Tip Mapping as optimization objective. **Bottom row:** Retargeting results using the proposed optimization objective.

where $\lambda = 0.001$ and \mathbf{R} is a depth rendering function [38] and D_t is the corresponding depth map in frame t . We minimize the re-projection error in 2D and optimize the hand 3D locations from the depth map. Equation 1 can be further extended to a multi-camera setting by minimizing the objective from different cameras with calibrated extrinsics. We use two cameras by default to handle occlusion. Additionally, we deploy a post-processing procedure with minimizing the difference of j_t^{3d} and β_t between frames.

6 Demonstration Translation

Common imitation learning algorithm consumes state-action pairs from expert demonstrations. It requires the state of the robot and the action of the motor as training data but not directly using the human hand pose. As shown in the kinematics chains in the blue box of Figure 3, although both robot and human hands share a similar five-finger morphology, their kinematics chains are different. The finger length of each knuckle is also different.

In this paper, we propose a novel method to translate demonstrations from *human-centric pose estimation* result into *robot-centric imitation data*. Specifically, there are two steps in demonstration translation (as shown in the red box of Figure 3): (i) **Hand motion retargeting** to align the human hand motion to robot hand motion, which are with different DoF and geometry; (ii) **Predicting robot action**, i.e. torque of robot motor: Without any wired sensors, we need to recover the action from only pose estimation results.

6.1 Hand Motion Retargeting

In computer graphics and animation, motion retargeting is used to retarget a performer’s motion to a virtual character for applications such as movies

and cartons [2,30]. While the animation community emphasizes realistic visual appearance, we consider more on the physical effect of the retargeted motion in the context of robotics manipulation. That is, the robot motion should be executable, respecting the physical limit of motors, e.g. acceleration and torque.

Formally, given a sequence of human hand pose estimated $\{(\theta_t, \beta_t, r_t)\}_{t=0}^T$ from a video, the hand pose retargeting can be defined as computing the sequence of robot joint angles $\{q_t\}_{t=0}^T$, where t is the step number in a sequence of length $T + 1$. We formulate the hand pose retargeting as an optimization problem.

Optimization with Task Space Vectors. In most manipulation tasks, human and multi-finger robot hand contacts the object using finger tips. Thus in fingertip-based mapping [4,27], retargeting is solved by preserving the Task Space Vectors (TSV) defined from finger tips position to the hand palm position (green arrow) as shown in the green box of Figure 3, so that the human and robot hands will have same finger tip position relative to the palm. However, only considering finger tip may results in unexpected optimization result as shown in Figure 4. Although the tip position is preserved, the bending information of hand fingers are lost, leading to penetrating the object after retargeting. It becomes more severe when the joint angles are closed to the robot singularity [48]. To solve this issue, we include the TSV from palm to middle phalanx (blue arrow) as shown in the green box of Figure 3. The optimization objective is,

$$\min_{q_t} \sum_{i=0}^N \|\mathbf{v}_i^{\mathbf{H}}(\theta_t, \beta_t, r_t) - \mathbf{v}_i^{\mathbf{R}}(q_t)\|^2 + \alpha \|q_t - q_{t-1}\|^2, \quad (2)$$

where the forward kinematics function for human $\mathbf{v}_i^{\mathbf{H}}(\theta, \beta, r)$ computes the i -th TSV for human hand and the robot forward kinematics function $\mathbf{v}_i^{\mathbf{R}}(q)$ computes i -th TSV for the robot. The first term in Equation 2 is to find the best $\{q_t\}_{t=0}^T$ that matches the finger tip-palm TSV and finger tip-phalanx TSV from both hands. Note the human hand TSV $\{v_i^{\mathbf{H}}(\theta_0, \beta_0, r_0)\}_{t=0}^T$ is first processed by a low pass filter before optimization for better smoothness. We also add an L2 normalization (second term in Equation 2) to improve the temporal consistency. For optimization, we use Sequential Least-Squares Quadratic Programming in NLOpt [36], where $\alpha = 8e - 3$ in implementation.

When retargeting hand pose, the objective in Equation 2 is optimized for each t from $t = 0$ to $t = T$. For $t \geq 1$, we initialize q_t using the optimization results q_{t-1} from last step to further improve the temporal smoothness.

6.2 Robot Action Estimation

Hand motion retargeting provides temporal-consistent translation from human hand poses to robot joint angles $\{q_t\}_{t=0}^T$ in different time steps. But the action, i.e. joint torque or control command, is unknown. In robotics, the joint torque τ can be computed via robot inverse dynamics function $\tau = f_{inv}(q, q', q'')$. To use this function, we first fit the sequence of joint angles into a continuous joint trajectory function $\mathbf{q}(t)$. Then, the first and second order derivative $\mathbf{q}'(t)$ and $\mathbf{q}''(t)$ can be used to compute the torque $\tau(t) = f_{inv}(q(t), q'(t), q''(t))$.

One important property of $\mathbf{q}(t)$ is: The absolute value of third-order derivative $\mathbf{q}'''(t)$, which also refers as jerk, should be as small as possible. We call it minimum-jerk requirement. The reason is two-fold: (i) First, physiological study [22] shows that the motion of human hand is a minimum jerk trajectory, where the lowest effort is required during the movement. This requirement will resemble the behavior of human, which in turn leads to more natural robot motion. (ii) Second, minimizing jerk can ensure low joint position errors for motors [43] and limit excessive wear on the physical robot [19]. Thus we also expect the fitted trajectory function $\{q_t\}_{t=0}^T$ to be minimum-jerk. In implementation, we use the model proposed in [75] to achieve this.

7 Imitation Learning

We perform imitation learning using the translated demonstrations. Instead of using behavior cloning, we adopt imitation learning algorithms that incorporate the demonstrations into RL.

Background. We consider the Markov Decision Process (MDP) represented by $\langle S, A, P, R, \gamma \rangle$, where S and A are state and action space, $P(s_{t+1}|s_t, a_t)$ is the transition density of state s_{t+1} at step $t+1$ given action a_t . $R(s, a)$ is the reward function, and γ is the discount factor. The goal of RL is to maximize the expected reward under policy $\pi(a|s)$. We incorporate RL with imitation learning. Given trajectories $\{(s_i, a_i)\}_{i=1}^n$ from demonstrations π_D , we optimize the agent policy π_θ with $\{(s_i, a_i)\}_{i=1}^n$ and reward R . We evaluate imitation learning under two settings: state-action imitation and state-only imitation.

7.1 State-action imitation learning

We will introduce two algorithms. The first one is the Generative Adversarial Imitation Learning (GAIL) [32], which is the SOTA IL method that performs occupancy measure matching to learn policy. Occupancy measure ρ_π [56] is a density measure of the state-action tuple on policy π . GAIL aims to minimize the objective $d(\rho_{\pi_E}, \rho_{\pi_\theta})$, where d is a distance function, π_θ is the policy parameterized by θ . GAIL use generative adversarial training to estimate the distance and minimize it with the objective as,

$$\min_{\theta} \max_{w, s, a \sim \rho_{\pi_\theta}} \mathbb{D} [\log D_w(s, a)] + \mathbb{E}_{s, a \sim \rho_{\pi_D}} [\log(1 - D_w(s, a))], \quad (3)$$

where D_w is a discriminator. To equip GAIL with reward function, we adopt the approach proposed by [37], denoted as GAIL+ in this paper.

The second algorithm is Demo Augmented Policy Gradient (DAPG) [59]. The objective function for policy optimization at each iteration k is as follow.

$$g_{aug} = \sum_{(s, a) \in \rho_{\pi_\theta}} \nabla_{\theta} \ln \pi_{\theta}(a|s) A^{\pi_{\theta}}(s, a) + \sum_{(s, a) \in \rho_{\pi_D}} \nabla_{\theta} \ln \pi_{\theta}(a|s) \lambda_0 \lambda_1^k, \quad (4)$$

where A^{π_θ} is the advantage function [9] of policy π_θ , and λ_0 and λ_1 are hyper-parameters to determine the advantage of state-action pair in demonstrations.

7.2 State-only imitation learning

Besides state-action imitation, we also evaluate the State-Only Imitation Learning (SOIL) [58] algorithm which does using action information from demonstrations. SOIL extends DAPG to the state-only imitation setting by learning an inverse model h_ϕ with the collected trajectories when running the policy. The inverse model predicts the missing actions in demonstrations and the policy is trained as DAPG. The policy and inverse model are optimized simultaneously.

8 Experiment

Experiment settings. We adopt TRPO [67] as our RL baseline. We benchmark imitation learning algorithms SOIL, GAIL+, and DAPG. All these algorithms incorporate demonstrations with TRPO of same hyper-parameters. Thus all the algorithms are comparable. All policies are evaluated with the same three random seeds. By default, the hand pose is estimated using two cameras.

8.1 Experiments with Relocate

Main comparisons. We benchmark four methods: SOIL, GAIL+, DAPG, and RL on the *relocate* tasks. The *success rate* is shown in Table 1 and training curves are in Figure 5. A trial is counted as success only when the final position of the object (after 200 steps) is within 0.1 unit length to the target. The initial object position and target position are randomized. In figure 5, the x-axis is training iterations and the y-axis is normalized average-return over three seeds.

Both Table 1 and Figure 5 show that the imitation learning methods outperform RL baseline for all five tasks. For mustard bottle and sugar box, a pure RL agent is not able to learn anything. SOIL performs the best for sugar box while DAPG achieves comparable or better performance on mug, mustard bottle, and tomato soup can. *Relocate* with sugar box is most challenging since it is tall and thin with a flat surface. Once it drops on the table, it is hard to grasp again. We conjecture training the inverse dynamics model online in SOIL helps obtain more accurate action for *Relocate*.

Ablation on motion retargeting method. Figure 6 (a) illustrates the performance of SOIL with respect to demonstrations generated by different retargeting method, on relocating a tomato soup can. It shows that our retargeting design (blue curve) can improve the imitation learning performance over fingertip mapping (red curve) based retargeting baseline. Besides, without L2 norms in Equation 2 (green curve), the performance shows a huge drop. It indicates the importance of smooth robot motion when used as demonstrations.

Ablation on the number of demonstrations. Figure 6 (b) shows the performance of SOIL with respect to different number of demonstrations, on




Fig. 5: Learning curves of the four methods on the relocate task with respect to five different objects. The x-axis is training iterations. The shaded area indicates standard error and the performance is evaluated with three individual random seeds.

Model	Task - Relocate				
	Mustard	Sugar Box	Tomato Can	Clamp	Mug
SOIL	0.33 ± 0.42	0.67 ± 0.47	0.98 ± 0.02	0.89 ± 0.15	0.71 ± 0.35
GAIL+	0.06 ± 0.01	0.00 ± 0.00	0.66 ± 0.47	0.52 ± 0.39	0.53 ± 0.37
DAPG	0.93 ± 0.05	0.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00
RL	0.06 ± 0.01	0.00 ± 0.00	0.67 ± 0.47	0.51 ± 0.37	0.49 ± 0.36

Table 1: Success rate of the evaluated methods on *Relocate* with five different objects. Success is defined based on the distance between object and target, evaluated via 100 trials for three seeds.

relocating a tomato soup can. SOIL can achieve better sample efficiency, performance and the variance is reduced when more demonstrations are given. Table 2 illustrates the success rate of the different policies trained with a different number of demonstrations. *This shows the importance of our efficient platform to scale up the number and diversity of demonstrations.*

Ablation on environmental conditions. Figure 6 (c) and (d) illustrates that our demonstrations can transfer to different physical environments. Given the same demonstrations on *Relocate* the tomato soup can, we train policies with objects scaling into different size from $\times 0.75$ to $\times 1.125$ (Figure 6 (c)) and applying different frictions (Figure 6 (d)) from $\times 0.8$ to $\times 1.2$. With same demonstrations, we can still perform imitation learning in different environments and achieve consistent results. All policies achieve much better performance than pure RL. *This experiment shows the robustness of our pipeline against the gap between simulation (environment) and real (demonstration).*

Ablation on hand pose estimation. We choose 4 different settings for hand pose estimation from video captured by: (i) single camera; (ii) single camera plus post-processing (iii) dual camera; (iv) dual cameras plus the post-processing (mentioned in Section 5.2) Since we have no ground-truth pose annotation in our dataset, we evaluate the performance of hand pose estimation approaches on the DexYCB [18] dataset, which provides the pose ground-truths. We follow DexYCB to use Mean Per Joint Position Error (MPJPE, smaller the better). We experiment with SOIL on relocating a tomato soup can. The object pose remains




Fig. 6: **Ablation Study:** Learning curves of SOIL on *Relocate* with tomato soup can. The x-axis is training iterations. We ablate: (a) hand pose estimation methods; (b) number of demonstrations used to train SOIL; (c) scaling of object size; (d) friction of the relocated object. The demonstrations are kept the same for all conditions.

#Demo	400 Iter.	600 Iter.	800 Iter.
SOIL (100)	0.36 ± 0.13	0.70 ± 0.25	1.00 ± 0.00
SOIL (50)	0.19 ± 0.23	0.40 ± 0.43	0.59 ± 0.39
SOIL (10)	0.04 ± 0.05	0.17 ± 0.11	0.36 ± 0.21
RL (0)	0.00 ± 0.00	0.00 ± 0.00	0.13 ± 0.19

Table 2: Success rate with different number of demonstrations on *Relocate* task with a tomato soup can, evaluated via 100 trials for three random seeds at 400, 600, and 800 training iterations.

Setting	MPJPE	Success (%)
1 Cam	41.7	66.7 ± 57.7
1 Cam Post	36.2	69.7 ± 33.3
2 Cam	36.6	84.7 ± 25.7
2 Cam Post	32.5	93.3 ± 11.5

Table 3: Hand pose error and the success rate of learned policy for 4 hand pose estimation settings. Smaller MPJPE means better hand pose estimation performance.





Fig. 7: *Pour* and *Place Inside*. (a) learning curves of *Pour*; (b) success rate of *Pour*; (c) learning rate of *Place Inside*; (d) Inside score of *Place Inside*.

the same for all four settings. Table 3 shows better pose estimation in general corresponds to better imitation, except the comparison on 1-Cam Post and 2-Cam. It seems 2-Cam are better than 1-Cam in imitation, even pose estimation result is close. We conjecture 2-Cam can provide more smooth trajectories.


8.2 Experiments with Pour

The *Pour* task involves a sequence of dexterous manipulations: reaching the mug, holding and stably moving the mug, and pouring water into the container. We benchmark four imitation learning algorithms in Figure 7 (a) and (b). We observe that DAPG converges to a good policy with much fewer iterations: 27.2% of the particles are poured into the container on average. Without the demonstrations, pure RL will only have a very small chance to pour a few particles



Setting	SOIL	GAIL+	DAPG	RL
Novel Object Instances, Same Category				
tomato.→can	62.0±27.3	30.7±23.5	83.6±4.7	15.9±16.7
mustard.→bottle	0.0±0.0	0.0±0.0	68.5±7.0	0.0±0.0
mug.→mug	51.4±33.2	33.8±36.1	79.4±11.9	27.7±38.8
Novel Category				
tomato.→cam.	27.1±13.5	5.9±6.3	47.2±2.5	3.8±2.1
mustard.→cam.	2.6±2.8	0.0±0.0	49.9±8.5	0.1±0.1
mug.→cam.	25.8±18.8	17.3±23.1	33.4±3.9	7.5±9.5

Fig. 8: *Relocate* generalization. Left: visualization of *Relocate* with unseen ShapeNet objects from 4 categories. Right: success rate of *Relocate*. A→B denotes using the policy trained on YCB object A to evaluate the performance on ShapeNet category B.



(a) Visualization of Two Tasks

Model	Success (%)	Category	Inside Score
SOIL	3.8 ± 3.2	SOIL	21.6 ± 16.1
GAIL+	2.1 ± 0.8	GAIL+	10.9 ± 4.6
DAPG	23.3 ± 10.4	DAPG	26.9 ± 17.3
RL	0.5 ± 0.3	RL	1.8 ± 1.5
(b) Pour		(c) Place Inside	

Fig. 9: Generalization with *Pour* and *Place Inside*. (a) visualization of *Pour* (top row) with ShapeNet mug, and *Place Inside* with ShapeNet cellphone; (b) success rate of *Pour* with ShapeNet Mug; (c) inside score of *Place Inside* with ShapeNet cellphone.

into the container. State-action method DAPG performs better than state-only method SOIL. It is challenging to learn inverse model with water particles in this task for SOIL, while the analytical inverse dynamics function still provides reasonable actions for DAPG. See our videos and website for visualization.

8.3 Experiments with Place Inside

The *Place Inside* task requires the robot hand to first pick up a banana, rotate it to the appropriate orientation, and place it inside the mug. We define a metric *Inside Score*, which is computed based on the volume percentage of the banana inside of the mug. We benchmark the imitation learning algorithms in Figure 7 (c) and (d). We find that DAPG outperforms other approaches whereas RL hardly learns to manipulate the object. Between state-only method (SOIL) and state-action method (DAPG), we again observe that computing action offline analytically achieves better results, due to the complexity of the task.

8.4 Generalization on Novel Objects and Category

Generalization with Relocate. We benchmark the generalization ability of the trained policy with different imitation learning algorithms. We directly

deployed the trained policies for the *Relocate* task using our demonstrations on: (i) novel unseen objects within the same category; (ii) object instances from a new category. For evaluation on unseen object instances within the same category, we choose the can, bottle, and mug categories from ShapeNet [16] dataset, which corresponds to the tomato soup can, mustard bottle, and mug in our demonstrations. We omit evaluation on the clamp and sugar box policies since there are no corresponding categories in ShapeNet. For the evaluation on new category, we choose the camera category from ShapeNet. We select 100 object instances for each category and pre-process the objects to fit the size of robot hand. More details about data pre-processing can be found in supplementary. We report the results in Figure 8. The success rate is computed as a mean accuracy over all object instances within a category. We find DAPG generalizes well on unseen objects in both same category and novel category experiments, significantly outperforming other approaches. We can also observe generalization to novel categories brings a larger challenge.

Generalization with Pour. We benchmark the generalization performance of the trained policies from all four methods on new mug instances for *Pour* task. We choose 100 object instances from the ShapeNet mug category for deployment. The top row of Figure 9 (a) visualizes the task with different mugs and Figure 9 (b) shows the success rate. In this benchmark, DAPG still performs the best across all methods. Interestingly, the performance is not much worse than *Pour* with the trained object (by comparing Fig. 7 (b)). This shows the robustness of imitation learning using our pipeline.

Generalization with Place Inside. For *Place Inside*, we replace the YCB banana using the objects from cellphone category in ShapeNet as visualized in the bottom row of Figure 9 (a). Figure 9 (c) shows the generalization performance when directly deploying the policy trained with banana. We find the policy performance is also close to Place Inside with banana (by comparing Fig. 7 (d)). This indicates the robustness of policy against small shape variations and human demonstration on single object can also benefit the tasks on new objects.

9 Conclusion

To the best of our knowledge, DexMV is the first work to provide a platform on computer vision/simulation systems and a pipeline on learning dexterous manipulation from human videos. We propose a novel demonstration translation module to bridge the gap between these two systems. We hope DexMV opens new research opportunities for benchmarking imitation learning algorithms for dexterous manipulation.

Acknowledgement. This work was supported, in part, by grants from DARPA LwLL, NSF CCF-2112665 (TILOS), NSF 1730158 CI-New: Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI), NSF ACI-1541349 CC*DNI Pacific Research Platform, and gifts from Meta, Google, Qualcomm and Picsart.

References

1. Abbeel, P., Ng, A.Y.: Apprenticeship learning via inverse reinforcement learning (2004)
2. Aberman, K., Wu, R., Lischinski, D., Chen, B., Cohen-Or, D.: Learning character-agnostic motion for motion retargeting in 2d. *arXiv preprint arXiv:1905.01680* (2019)
3. Andrews, S., Kry, P.G.: Goal directed multi-finger manipulation: Control policies and analysis. *Computers & Graphics* (2013)
4. Antotsiou, D., Garcia-Hernando, G., Kim, T.K.: Task-oriented hand motion retargeting for dexterous manipulation imitation. In: *ECCV Workshops* (2018)
5. Aytar, Y., Pfaff, T., Budden, D., Paine, T., Wang, Z., de Freitas, N.: Playing hard exploration games by watching youtube. In: *NeurIPS* (2018)
6. Baek, S., Kim, K.I., Kim, T.K.: Pushing the envelope for rgb-based dense 3d hand pose estimation via neural rendering. In: *CVPR* (2019)
7. Bai, Y., Liu, C.K.: Dexterous manipulation using both palm and fingers (2014)
8. Bain, M., Sammut, C.: A framework for behavioural cloning. In: *Machine Intelligence* (1995)
9. Baird III, L.C.: Advantage updating. Tech. rep. (1993)
10. Bicchi, A.: Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on robotics and automation* (2000)
11. Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., Zhang, J., et al.: End to end learning for self-driving cars. *arXiv* (2016)
12. Boukhayma, A., Bem, R.d., Torr, P.H.: 3d hand shape and pose from images in the wild. In: *CVPR* (2019)
13. Brahmbhatt, S., Ham, C., Kemp, C.C., Hays, J.: Contactdb: Analyzing and predicting grasp contact via thermal imaging. In: *CVPR* (2019)
14. Brahmbhatt, S., Handa, A., Hays, J., Fox, D.: Contactgrasp: Functional multi-finger grasp synthesis from contact. *arXiv* (2019)
15. Calli, B., Walsman, A., Singh, A., Srinivasa, S., Abbeel, P., Dollar, A.M.: Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols. *arXiv* (2015)
16. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015)
17. Chang, M., Gupta, A., Gupta, S.: Semantic visual navigation by watching youtube videos. In: *NIPS* (2020)
18. Chao, Y.W., Yang, W., Xiang, Y., Molchanov, P., Handa, A., Tremblay, J., Narang, Y.S., Van Wyk, K., Iqbal, U., Birchfield, S., et al.: Dexycb: A benchmark for capturing hand grasping of objects. In: *CVPR* (2021)
19. Craig, J.J.: Introduction to robotics: mechanics and control, 3/E. Pearson Education India (2009)
20. Dogar, M.R., Srinivasa, S.S.: Push-grasping with dexterous hands: Mechanics and a method (2010)
21. Duan, Y., Chen, X., Houthoofd, R., Schulman, J., Abbeel, P.: Benchmarking deep reinforcement learning for continuous control (2016)
22. Flash, T., Hogan, N.: The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of neuroscience* **5**(7), 1688–1703 (1985)

23. Fu, J., Luo, K., Levine, S.: Learning robust rewards with adversarial inverse reinforcement learning. *arXiv* (2017)
24. Garcia-Hernando, G., Yuan, S., Baek, S., Kim, T.K.: First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. *CVPR* (2018)
25. Ge, L., Ren, Z., Li, Y., Xue, Z., Wang, Y., Cai, J., Yuan, J.: 3d hand shape and pose estimation from a single rgb image. In: *CVPR* (2019)
26. Hampali, S., Rad, M., Oberweger, M., Lepetit, V.: Honnotate: A method for 3d annotation of hand and object poses. In: *CVPR* (2020)
27. Handa, A., Van Wyk, K., Yang, W., Liang, J., Chao, Y.W., Wan, Q., Birchfield, S., Ratliff, N., Fox, D.: Dexpilot: Vision-based teleoperation of dexterous robotic hand-arm system. In: *ICRA* (2020)
28. Hasson, Y., Varol, G., Tzionas, D., Kalevatykh, I., Black, M.J., Laptev, I., Schmid, C.: Learning joint reconstruction of hands and manipulated objects. In: *CVPR* (2019)
29. He, Y., Sun, W., Huang, H., Liu, J., Fan, H., Sun, J.: Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. *CVPR* (2020)
30. Hecker, C., Raabe, B., Enslow, R.W., DeWeese, J., Maynard, J., van Prooijen, K.: Real-time motion retargeting to highly varied user-created morphologies. *ACM Transactions on Graphics (TOG)* **27**(3), 1–11 (2008)
31. Ho, J., Ermon, S.: Generative adversarial imitation learning. In: *NeurIPS* (2016)
32. Ho, J., Ermon, S.: Generative adversarial imitation learning. In: *NeurIPS* (2016)
33. Hu, Y., Hugonot, J., Fua, P., Salzmann, M.: Segmentation-driven 6d object pose estimation. *CVPR* (2019)
34. Iqbal, U., Molchanov, P., Breuel Juergen Gall, T., Kautz, J.: Hand pose estimation via latent 2.5 d heatmap regression. In: *ECCV* (2018)
35. Jiang, H., Liu, S., Wang, J., Wang, X.: Hand-object contact consistency reasoning for human grasps generation. *arXiv* (2021)
36. Johnson, S.G.: The nlopt nonlinear-optimization package (2014)
37. Kang, B., Jie, Z., Feng, J.: Policy optimization with demonstrations. In: *ICML* (2018)
38. Kato, H., Ushiku, Y., Harada, T.: Neural 3d mesh renderer. In: *CVPR* (2018)
39. Kehl, W., Manhardt, F., Tombari, F., Ilic, S., Navab, N.: Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. *ICCV* (2017)
40. Khaled, S.M., Islam, M.S., Rabbani, M.G., Tabassum, M.R., Gias, A.U., Kamal, M.M., Muctadir, H.M., Shakir, A.K., Imran, A., Islam, S.: Combinatorial color space models for skin detection in sub-continental human images. In: *IVIC* (2009)
41. Kulon, D., Guler, R.A., Kokkinos, I., Bronstein, M.M., Zafeiriou, S.: Weakly-supervised mesh-convolutional hand reconstruction in the wild. In: *CVPR* (2020)
42. Kumar, V., Xu, Z., Todorov, E.: Fast, strong and compliant pneumatic actuation for dexterous tendon-driven hands. In: *ICRA* (2013)
43. Kyriakopoulos, K.J., Saridis, G.N.: Minimum jerk path generation. In: *Proceedings. 1988 IEEE international conference on robotics and automation*. pp. 364–369. *IEEE* (1988)
44. Li, S., Ma, X., Liang, H., Görner, M., Ruppel, P., Fang, B., Sun, F., Zhang, J.: Vision-based teleoperation of shadow dexterous hand using end-to-end deep neural network. In: *ICRA* (2019)
45. Liu, F., Ling, Z., Mu, T., Su, H.: State alignment-based imitation learning. In: *ICLR* (2020)
46. Liu, S., Jiang, H., Xu, J., Liu, S., Wang, X.: Semi-supervised 3d hand-object poses estimation with interactions in time. In: *CVPR* (2021)

47. Mandikal, P., Grauman, K.: Dexterous robotic grasping with object-centric visual affordances. *arXiv* (2020)
48. Nakamura, Y., Hanafusa, H.: Inverse kinematic solutions with singularity robustness for robot manipulator control (1986)
49. Ng, A.Y., Russell, S.J., et al.: Algorithms for inverse reinforcement learning. (2000)
50. Okamura, A.M., Smaby, N., Cutkosky, M.R.: An overview of dexterous manipulation. In: *ICRA* (2000)
51. OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., Zhang, L.: Solving rubik’s cube with a robot hand. *arXiv* (2019)
52. OpenAI, Andrychowicz, M., Baker, B., Chociej, M., Józefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., Schneider, J., Sidor, S., Tobin, J., Welinder, P., Weng, L., Zaremba, W.: Learning dexterous in-hand manipulation. *arXiv* (2018)
53. Peng, S., Liu, Y., Huang, Q.-X., Bao, H., Zhou, X.: Pynet: Pixel-wise voting network for 6dof pose estimation. *CVPR* (2019)
54. Peters, J., Schaal, S.: Reinforcement learning of motor skills with policy gradients. *Neural networks* (2008)
55. Pomerleau, D.A.: Alvin: An autonomous land vehicle in a neural network. In: *NeurIPS* (1989)
56. Puterman, M.L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc. (1994)
57. Rad, M., Lepetit, V.: Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. *ICCV* (2017)
58. Radosavovic, I., Wang, X., Pinto, L., Malik, J.: State-only imitation learning for dexterous manipulation. *IROS* (2021)
59. Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., Levine, S.: Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv* (2017)
60. Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., Levine, S.: Learning complex dexterous manipulation with deep reinforcement learning and demonstrations (2018)
61. Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. *ToG* (2017)
62. Ross, S., Bagnell, D.: Efficient reductions for imitation learning. In: *AISTATS* (2010)
63. Rus, D.: In-hand dexterous manipulation of piecewise-smooth 3-d objects. *The International Journal of Robotics Research* (1999)
64. Russell, S.: *Learning agents for uncertain environments* (1998)
65. Schmeckpeper, K., Rybkin, O., Daniilidis, K., Levine, S., Finn, C.: Reinforcement learning with videos: Combining offline observations with interaction. *arXiv* (2020)
66. Schmeckpeper, K., Xie, A., Rybkin, O., Tian, S., Daniilidis, K., Levine, S., Finn, C.: Learning predictive models from observation and interaction. *arXiv* (2019)
67. Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P.: Trust region policy optimization. In: *ICML* (2015)
68. Shan, D., Geng, J., Shu, M., Fouhey, D.: Understanding human hands in contact at internet scale. In: *CVPR* (2020)
69. Shao, L., Migimatsu, T., Zhang, Q., Yang, K., Bohg, J.: Concept2robot: Learning manipulation concepts from instructions and human demonstrations. In: *RSS* (2020)

70. Song, S., Zeng, A., Lee, J., Funkhouser, T.: Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations. *Robotics and Automation Letters* (2020)
71. Spurr, A., Song, J., Park, S., Hilliges, O.: Cross-modal deep variational hand pose estimation. In: *CVPR* (2018)
72. Taheri, O., Ghorbani, N., Black, M.J., Tzionas, D.: Grab: A dataset of whole-body human grasping of objects. In: *ECCV* (2020)
73. Tekin, B., Sinha, S.N., Fua, P.: Real-time seamless single shot 6d object pose prediction. *CVPR* (2018)
74. Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: *IROS* (2012)
75. Todorov, E., Jordan, M.I.: Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements. *Journal of Neurophysiology* (1998)
76. Torabi, F., Warnell, G., Stone, P.: Behavioral cloning from observation. *arXiv* (2018)
77. Torabi, F., Warnell, G., Stone, P.: Generative adversarial imitation from observation. *arXiv* (2018)
78. Večerík, M., Hester, T., Scholz, J., Wang, F., Pietquin, O., Piot, B., Heess, N., Rothörl, T., Lampe, T., Riedmiller, M.: Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv* (2017)
79. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv* (2018)
80. Young, S., Gandhi, D., Tulsiani, S., Gupta, A., Abbeel, P., Pinto, L.: Visual imitation made easy. *arXiv* (2020)
81. Zimmermann, C., Brox, T.: Learning to estimate 3d hand pose from single rgb images. In: *CVPR* (2017)