

A System for General In-Hand Object Re-Orientation

by

Tao Chen

S.M., Carnegie Mellon University (2019)

S.B., Shanghai Jiao Tong University (2016)

Submitted to the Department of Electrical Engineering
and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
January 3, 2022

Certified by.....
Pulkit Agrawal
Assistant Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

A System for General In-Hand Object Re-Orientation

by

Tao Chen

Submitted to the Department of Electrical Engineering
and Computer Science
on January 3, 2022, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

In-hand object reorientation has been a challenging problem in robotics due to high dimensional actuation space and the frequent change in contact state between the fingers and the objects. We present a simple model-free framework that can learn to reorient objects with both the hand facing upwards and downwards. We demonstrate the capability of reorienting over 2000 geometrically different objects in both cases. The learned policies show strong zero-shot transfer performance on new objects. We provide evidence that these policies are amenable to real-world operation by distilling them to use observations easily available in the real world.

Thesis Supervisor: Pulkit Agrawal

Title: Assistant Professor of Electrical Engineering and Computer Science

Acknowledgments

First, I would like to thank my advisor, Professor Pulkit Agrawal, for all his generous support throughout my master's program and encouragement in my research. His unique insights and perspectives really helped me think about problems in a different way and had a deep impact on my research. I would also like to thank all my collaborators, who made my research experience a pleasant journey. Furthermore, I would like to thank all the members of the Improbable AI lab for their help with the research and life. Finally, I would like to thank my family for their unconditional support and love.

Contents

1	Introduction	13
2	Method	17
2.1	Learning the teacher policy	18
2.2	Learning the student policy	19
2.2.1	Training student policy from low-dimensional state	19
2.2.2	Training student policy from vision	20
3	Experimental Setup	23
4	Results	25
4.1	Reorient objects with the hand facing upward	25
4.2	Reorient objects with the hand facing downward	27
4.2.1	Reorient objects on a table	28
4.2.2	Reorient objects in air with hand facing downward	28
4.3	Zero-shot policy transfer across datasets	30
4.4	Object Reorientation with RGBD sensors	30
5	Related Work	33
6	Discussion and Conclusion	35
A		43
A.1	Evidence indicating transfer to real-world	43
A.2	Environment Setup	45

A.2.1	State definition	45
A.2.2	Dataset	45
A.2.3	Camera setup	46
A.3	Experiment Setup	47
A.3.1	Network architecture	47
A.3.2	Training details	48
A.3.3	Dynamics randomization	49
A.3.4	Gravity curriculum	51
A.4	Supplementary Results	51
A.4.1	Hand faces upward	51
A.4.2	Hand faces downward (in the air)	53
A.4.3	Success rate on each type of YCB objects	53
A.4.4	Torque analysis	54
A.4.5	Vision experiments with noise	55

List of Figures

2-1	Visual policy architecture. MK stands for Minkowski Engine. q_t is the joint positions and a_t is the action at time step t	20
3-1	Examples of initial poses of the hand and object. (a) : hand faces upward. (b) , (c) , (d) : hand faces downward. (b) : both the hand and the object are initialized with random poses . (c) : there is a table below the hand. (d) : the hand and the object are initialized from the lifted poses.	24
A-1	We learn policies that can reorient many objects in three scenarios respectively: (1) hand faces upward, (2) hand faces downward with a table below the hand, (3) hand faces downward without any table. The extra object in each figure shows the desired orientation.	45
A-2	First row: examples of EGAD objects. Second row: examples of YCB objects.	46
A-3	Examples of EGAD objects. The first and third row shows the visual mesh of the objects. The second and fourth row show the corresponding collision mesh (after V-HACD decomposition).	47
A-4	Camera positions	47
A-5	Learning curves of the MLP and RNN policies on the EGAD (Left) and YCB datasets (Middle). The Right plot shows that using the full state information speeds up the policy learning compared to only using the reduced state information.	52

A-6	Examples of failure cases. (a) and (b): objects are too small. (c): the object is reoriented close to the target orientation, but not close enough. (d): the object is too big and initialized around the palm border. . . .	53
A-7	An example of reorienting a cup with the hand facing downward. From left to right, top to bottom, we show the some moments in an episode.	54
A-8	An example of reorienting a sponge with the hand facing downward. From left to right, top to bottom, we show the some moments in an episode.	55
A-9	Reorientation success rates for each object in the YCB dataset when the hand faces upward.	57
A-10	Reorientation success rates for each object in the YCB dataset when the hand faces downward without a table.	58
A-11	Distribution of the <i>maximum</i> absolute joint torque values on all joints for all the time steps. We exclude a few outliers in the plot, i.e., we only plot the data distributions up to 99% quantile.	59

List of Tables

4.1	Success rates (%) of policies tested on different dynamics distribution. $\bar{\theta} = 0.1\text{rad}$. DR: domain randomization and observation/action noise. X→Y: distill policy X into policy Y. The full table is in Table A.5.	27
4.2	Performance of the student policy when the hand faces upward and downward	30
4.3	Zero-shot RNN policy transfer success rates (%) across datasets. U . (D.) means hand faces upward (downward). FS (RS) means using full-state (reduced-state) information.	30
4.4	Vision policy success rate ($\bar{\theta} = 0.2\text{rad}$, $\bar{d}_C = 0.01$)	31
A.1	Full state $s_t^{\mathcal{E}} \in \mathbb{R}^{134}$ information. Orientations are in the form of quaternions.	45
A.2	Hyperparameter Setup	50
A.3	Mesh Parameters	50
A.4	Dynamics Randomization and Noise	50
A.5	Success rates (%) of policies tested on different dynamics distribution. $\bar{\theta} = 0.1\text{rad}$. DR stands for domain randomization and observation/action noise. X→Y: distill policy X into policy Y.	53
A.6	Success rates (%) of policies trained with hand facing downward and to reorient objects in the air. Due to the large number of environment steps required in this setup, we fine-tune the model trained without DR with randomized dynamics instead of training models with DR from scratch.	54

A.7 Vision policy success rates when the policy is trained and tested
with/without data augmentation ($\bar{\theta} = 0.2 \text{ rad}, \bar{d}_C = 0.01$) 56

Chapter 1

Introduction

A common maneuver in many tasks of daily living is to pick an object, reorient it in hand and then either place it or use it as a tool. Consider three simplified variants of this maneuver shown in Figure ???. The task in the top row requires an upward-facing multi-finger hand to manipulate an *arbitrary* object in a random orientation to a goal configuration shown in the rightmost column. The next two rows show tasks where the hand is facing downward and is required to reorient the object either using the table as a support or without the aid of any support surface respectively. The last task is the hardest because the object is in an intrinsically unstable configuration owing to the downward gravitational force and lack of support from the palm. Additional challenges in performing such manipulation with a multi-finger robotic hand stem from the control space being high-dimensional and reasoning about the multiple transitions in the contact state between the finger and the object. Due to its practical utility and several unsolved issues, in-hand object reorientation remains an active area of research.

Past work has tackled the in-hand reorientation problem via several approaches: (i) The use of analytical models with powerful trajectory optimization methods [1, 2, 3]. While these methods demonstrated remarkable performance, the results were largely in simulation with simple object geometries and required detailed knowledge of the object model and physical parameters. As such, it remains unclear how to scale these methods to real-world and generalize to new objects. Another line of work has employed (ii)

model-based reinforcement learning [4, 5]; or (iii) model-free reinforcement learning with [6, 7, 8, 9] and without expert demonstrations [10, 11, 12, 13]. While some of these works demonstrated learned skills on real robots, it required use of additional sensory apparatus not readily available in the real-world (e.g., motion capture system) to infer the object state, and the learned policies did not generalize to diverse objects. Furthermore, most prior methods operate in the simplified setting of the hand facing upwards. The only exception is pick-and-place, but it does not involve any in-hand re-orientation. A detailed discussion of prior research is provided in Section 5.

In this paper, our goal is to study the object reorientation problem with a multi-finger hand in its general form. We desire (a) manipulation with hand facing upward or downward; (b) the ability of using external surfaces to aid manipulation; (c) the ability to reorient objects of novel shapes to arbitrary orientations; (d) operation from sensory data that can be easily obtained in the real world such as RGBD images and joint positions of the hand. While some of these aspects have been individually demonstrated in prior work, we are unaware of any published method that realizes all four. Our main contribution is building a system that achieves the desiderata. The core of our framework is a model-free reinforcement learning with three key components: teacher-student learning, gravity curriculum, and stable initialization of objects. Our system requires no knowledge of object or manipulator models, contact dynamics or any special pre-processing of sensory observations. We experimentally test our framework using a simulated Shadow hand. Due to the scope of the problem and the ongoing pandemic, we limit our experiments to be in simulation. However, we provide evidence indicating that the learned policies can be transferred to the real world in the future.

A Surprising Finding: While seemingly counterintuitive, we found that policies that have no access to shape information can manipulate a large number of previously unseen objects in all the three settings mentioned above. At the start of the project, we hypothesized that developing visual processing architecture for inferring shape while the robot manipulates the object would be the primary research challenge. On the contrary, our results show that it is possible to learn control strategies for

general in-hand object re-orientation that are shape-agnostic. Our results, therefore, suggest that visual perception may be less important for in-hand manipulation than previously thought. Of course, we still believe that the performance of our system can be improved by incorporating shape information. However, our findings suggest a different framework of thinking: a lot can be achieved without vision, and that vision might be the icing on the cake instead of the cake itself.

Chapter 2

Method

We found that *simple* extensions to existing techniques in robot learning can be used to construct a system for general object reorientation. First, to avoid explicit modeling of non-linear and frequent changes in the contact state between the object and the hand, we use model-free reinforcement learning (RL). An added advantage is that model-free RL is amenable to direct operation from raw point cloud observations, which is preferred for real-world deployment. We found that better policies can be trained faster using *privileged* state information such as the velocities of the object/fingertips that is easily available in the simulator but not in the real world. To demonstrate the possibility of transferring learned policies to the real world in the future, we overcome the need for privileged information using the idea of *teacher-student training* [14, 15]. In this framework, first, an expert or *teacher* policy ($\pi^{\mathcal{E}}$) is trained using privileged information. Next, the *teacher* policy guides the learning of a *student* policy ($\pi^{\mathcal{S}}$) that only uses sensory inputs available in the real world. Let the state space corresponding to $\pi^{\mathcal{E}}$ and $\pi^{\mathcal{S}}$ be $\mathbb{S}^{\mathcal{E}}$ and $\mathbb{S}^{\mathcal{S}}$ respectively. In general, $\mathbb{S}^{\mathcal{E}} \neq \mathbb{S}^{\mathcal{S}}$.

We first trained the teacher policy to reorient more than two thousand objects of diverse shapes (see Section 2.1). Next, we detail the method for distilling $\pi^{\mathcal{E}}$ to a student policy using a reduced state space consisting of only the joint positions of the hand, the object position, and the difference in orientation from the goal configuration (see Section 2.2.1). However, in the real world, even the object position and relative orientation must be inferred from sensory observation. Not only does this process

require substantial engineering effort (e.g., a motion capture or a pose estimation system), but also inferring the pose of a symmetric object is prone to errors. This is because a symmetric object at multiple orientations looks exactly the same in sensory space such as RGBD images.

To mitigate these issues, we further distill $\pi^{\mathcal{E}}$ to operate directly from the point cloud and position of all the hand joints (see Section 2.2.2). We propose a simple modification that generalizes an existing 2D CNN architecture [16] to make this possible.

The procedure described above works well for manipulation with the hand facing upwards and downwards when a table is available as support. However, when the hand faces downward without an underlying support surface, we found it important to initialize the object in a stable configuration. Finally, because gravity presents the primary challenge in learning policies with a downward-facing hand, we found that training in a curriculum where gravity is slowly introduced (i.e., *gravity curriculum*) substantially improves performance. These are discussed in Section 4.2.

2.1 Learning the teacher policy

We use model-free RL to learn the teacher policy ($\pi^{\mathcal{E}}$) for reorienting an object ($\{O_i | i = 1, \dots, N\}$) from an initial orientation α_0^o to a target orientation α^g ($\alpha_0^o \neq \alpha^g$). At every time step t , the agent observes the state s_t , executes the action a_t sampled from the policy $\pi^{\mathcal{E}}$, and receives a reward r_t . $\pi^{\mathcal{E}}$ is optimized to maximize the expected discounted return: $\pi = \arg \max_{\pi^{\mathcal{E}}} \mathbb{E}[\sum_{t=0}^{T-1} \gamma^t r_t]$, where γ is the discount factor. The task is successful if the angle difference $\Delta\theta$ between the object’s current (α_t^o) and the goal orientation (α^g) is smaller than the threshold value $\bar{\theta}$, i.e., $\Delta\theta \leq \bar{\theta}$.

To encourage the policy to be smooth, the previous action is appended to the inputs to the policy (i.e., $a_t = \pi^{\mathcal{E}}(s_t, a_{t-1})$) and large actions are penalized in the reward function. We experiment with two architectures for $\pi^{\mathcal{E}}$: (1) an MLP policy π_M , (2) an RNN policy π_R . We use PPO [17] to optimize $\pi^{\mathcal{E}}$. More details about the training are in Section A.3.1 and Section A.3.2 in the appendix.

Observation and action space: We define $\mathbb{S}^\mathcal{E}$ to include joint, fingertip, object, and goal information as detailed in Table A.1 in the appendix. Note that $\mathbb{S}^\mathcal{E}$ does not include object shape or information about friction, damping, contact states between the fingers and the object, etc. We control the joint movements by commanding the relative change in the target joint angle (q_t^{target}) on each actuated joint (action $a_t \in \mathbb{R}^{20}$): $q_{t+1}^{target} = q_t^{target} + a_t \times \Delta t$, where Δt is the control time step. We clamp the action command if necessary to make sure $|\Delta q_t^{target}| \leq 0.33$ rad. The control frequency is 60 Hz.

Dynamics randomization: Even though we do not test our policies on a real robot, we train and evaluate policies with domain randomization [18] to provide evidence that our work has the potential to be transferred to a real robotic system in the future. We randomize the object mass, friction coefficient, joint damping and add noise to the state observation s_t and the action a_t . More details about domain randomization are provided in Table A.4 in the appendix.

2.2 Learning the student policy

We distill the teacher $\pi^\mathcal{E}$ into the student policy $\pi^\mathcal{S}$ using Dagger [19], a learning-from-demonstration method that overcomes the covariate shift problem. We optimize $\pi^\mathcal{S}$ by minimizing the KL-divergence between $\pi^\mathcal{S}$ and $\pi^\mathcal{E}$:

$$\pi^\mathcal{S} = \arg \min_{\pi^\mathcal{S}} D_{KL} (\pi^\mathcal{E}(s_t^\mathcal{E}, a_{t-1}) || \pi^\mathcal{S}(s_t^\mathcal{S}, a_{t-1}))$$

Based on observation data available in real-world settings, we investigate two different choices of $\mathbb{S}^\mathcal{S}$.

2.2.1 Training student policy from low-dimensional state

In the first case, we consider a non-vision student policy $\pi^\mathcal{S}(s_t^\mathcal{S}, a_{t-1})$. $s_t^\mathcal{S} \in \mathbb{R}^{31}$ includes the joint positions $q_t \in \mathbb{R}^{24}$, object position $p_t^o \in \mathbb{R}^3$, quaternion difference between the object’s current and target orientation $\beta_t \in \mathbb{R}^4$. In this case, $\mathbb{S}^\mathcal{S} \subset \mathbb{S}^\mathcal{E}$,

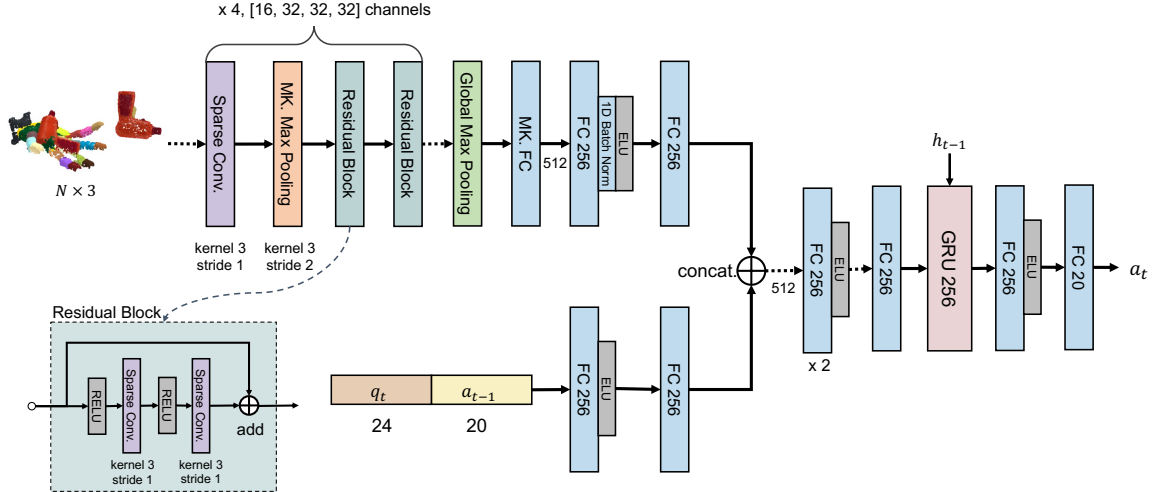


Figure 2-1: Visual policy architecture. MK stands for Minkowski Engine. q_t is the joint positions and a_t is the action at time step t .

and we assume the availability of object pose information, but do not require velocity information. We use the same MLP and RNN network architectures used for $\pi^{\mathcal{E}}$ on $\pi^{\mathcal{S}}$ except the input dimension changes as the state dimension is different.

2.2.2 Training student policy from vision

In the second case, $\mathbb{S}^{\mathcal{S}}$ only uses direct observations from RGBD cameras and the joint position (q_t) of the robotic hand. We convert the RGB and Depth data into a colored point cloud using a pinhole camera model [20]. Our vision policy takes as input the voxelized point cloud of the scene W_t , q_t , and previous action command a_{t-1} , and outputs the action a_t , *i.e.*, $a_t = \pi^{\mathcal{S}}(W_t, q_t, a_{t-1})$.

Goal specification: To avoid manually defining per-object coordinate frame for specifying the goal quaternion, we provide the goal to the policy as an object point cloud rotated to the desired orientation W^g , *i.e.*, we only show the policy how the object should look like in the end (see the top left of Figure 2-1). The input to $\pi^{\mathcal{S}}$ is the point cloud $W_t = W_t^s \cup W^g$ where W_t^s is the actual point cloud of the current scene obtained from the cameras. Details of obtaining W^g are in Section A.3.2.

Sparse3D-IMPALA-Net: To convert a voxelized point cloud into a lower-dimensional feature representation, we use a sparse convolutional neural network. We

extend the IMPALA policy architecture [16] for processing RGB images to process colored point cloud data using 3D convolution. Since many voxels are unoccupied, the use of regular 3D convolution substantially increases computation time. Hence, we use Minkowski Engine [21], a PyTorch library for sparse tensors, to design a 3D version of IMPALA-Net with sparse convolutions (*Sparse3D-IMPALA-Net*)¹. The Sparse3D-IMPALA network takes as input the point cloud W_t , and outputs an embedding vector which is concatenated with the embedding vector of (q_t, a_{t-1}) . Afterward, a recurrent network is used and outputs the action a_t . The detailed architecture is illustrated in Figure 2-1.

Mitigating the object symmetry issue: $\pi^{\mathcal{E}}$ is trained with the the ground-truth state information $s_t^{\mathcal{E}}$ including the object orientation α_t^o and goal orientation α^g . The vision policy does not take any orientation information as input. If an object is symmetric, the two different orientations of the object may correspond to the same point cloud observation. This makes it problematic to use the difference in orientation angles ($\Delta\theta \leq \bar{\theta}$) as the stopping and success criterion. To mitigate this issue, we use Chamfer distance [22] to compute the distance between the object point cloud in α_t^o and the goal point cloud (i.e., the object rotated by α^g) as the evaluation criterion. The Chamfer distance is computed as $d_C = \sum_{a \in W_t^o} \min_{b \in W^g} \|a - b\|_2^2 + \sum_{b \in W^g} \min_{a \in W_t^o} \|a - b\|_2^2$, where W_t^o is the object point cloud in its current orientation. Both W_t^o and W^g are scaled to fit in a unit sphere for computing d_C . We check Chamfer distance in each rollout step. If $d_C \leq \bar{d}_C$ (\bar{d}_C is a threshold value for d_C), we consider the episode to be successful. Hence, the success criterion is $(\Delta\theta \leq \bar{\theta}) \vee (d_C \leq \bar{d}_C)$. In training, if the success criterion is satisfied, the episode is terminated and used for updating $\pi^{\mathcal{S}}$.

¹We also experimented with a 3D sparse convolutional network based on ResNet18, and found that 3D IMPALA-Net works better.

Chapter 3

Experimental Setup

We use the simulated Shadow Hand [23] in NVIDIA Isaac Gym [24]. Shadow Hand is an anthropomorphic robotic hand with 24 degrees of freedom (DoF). We assume the base of the hand to be fixed. Twenty joints are actuated by agonist–antagonist tendons and the remaining four are under-actuated.

Object datasets: We use the EGAD dataset [25] and YCB dataset [26] that contain objects with diverse shapes (see Figure A-2) for in-hand manipulation experiments. EGAD contains 2282 geometrically diverse textureless object meshes, while the YCB dataset includes textured meshes for objects of daily life with different shapes and textures. We use the 78 YCB object models collected with the Google scanner. Since most YCB objects are too big for in-hand manipulation, we proportionally scale down the YCB meshes. To further increase the diversity of the datasets, we create 5 variants for each object mesh by randomly scaling the mesh. More details of the object datasets are in Section A.2.2.

Setup for visual observations: For the vision experiments, we trained policies for the scenario of hand facing upwards. We place two RGBD cameras above the hand (Figure A-4). The data from these cameras is combined to create a point cloud observation of the scene [20]. For downstream computation, the point cloud is voxelized with a resolution of 0.003 m.

Setup with the upward facing hand: We first consider the case where the Shadow Hand faces upward and is required to reorient objects placed in the hand (see




Figure 3-1: Examples of initial poses of the hand and object. **(a)**: hand faces upward. **(b)**, **(c)**, **(d)**: hand faces downward. **(b)**: both the hand and the object are initialized with random poses. **(c)**: there is a table below the hand. **(d)**: the hand and the object are initialized from the lifted poses.

Figure 3-1a). We use the coordinate system where the z -axis points vertically upward and the xy -plane denotes the horizontal plane. The object pose is initialized with the following procedure: xy position of the object’s center of mass (COM) $p_{0,xy}^o$ is randomly sampled from a square region of size $0.09\text{ m} \times 0.09\text{ m}$. The center of this square is approximately located on the intersection of the middle finger and the palm so that the sampling region covers both the fingers and the palm. The z position of the object is fixed to 0.13 m above the base of the hand to ensure that the object does not collide with the hand at initialization. The initial and goal orientations are randomly sampled from the full $SO(3)$ space.

Setup with the downward facing hand: Next, we consider the cases where the hand faces downward. We experiment with two scenarios: with and without a table below the hand. In the first case, we place a table with the tabletop being 0.12 m below the hand base. We place objects in a random pose between the hand and the table so that the objects will fall onto the table. We will describe the setup for the second case (without a table) in Section 4.2.2.

Evaluation criterion: For non-vision experiments, a policy rollout is considered a success if $\theta \leq \bar{\theta}$. $\bar{\theta} = 0.1\text{ rad}$. For vision experiments, we also check $d_C \leq \bar{d}_C$ as another criterion and $\bar{\theta} = 0.2\text{ rad}$, $\bar{d}_C = 0.01$. The initial and goal orientation are randomly sampled from $SO(3)$ space in all the experiments. We report performance as the percentage of successful episodes when the agent is tasked to reorient each training object 100 times from arbitrary start to goal orientation. We report the mean and standard deviation of success rate from 3 seeds.

Chapter 4

Results

We evaluate the performance of reorientation policies with the hand facing upward and downward. Further we analyze the generalization of the learned policies to unseen object shapes.

4.1 Reorient objects with the hand facing upward

Train a teacher policy with full-state information We train our teacher MLP and RNN policies using the full state information using all objects in the EGAD and YCB datasets separately. The progression of success rate during training is shown in Figure A-5 in Appendix A.4.1 . Figure A-5 also shows that using privileged information substantially speeds up policy learning. Results reported in Table 4.1 indicate that the RNN policies achieve a success rate greater than 90% on the EGAD dataset (entry B1) and greater than 80% on the YCB dataset (entry G1) without any explicit knowledge of the object shape¹. This result is surprising because apriori one might believe that shape information is important for in-hand reorientation of diverse objects.

The visualization of policy rollout reveals that the agent employs a clever strategy that is invariant to object geometry for re-orienting objects. The agent throws the object in the air with a spin and catches it at the precise time when the object's

¹More quantitative results on the MLP policies are available in Table A.5 in the appendix.

orientation matches the goal orientation. Throwing the object with a spin is a dexterous skill that automatically emerges! One possibility for the emergence of this skill is that we used very light objects. This is not true because we trained with objects in the range of 50-150g which spans many hand-held objects used by humans (e.g., an egg weighs about 50g, a glass cup weighs around 100g, iPhone 12 weighs 162g, etc.). To further probe this concern, we evaluated zero-shot performance on objects weighing up to 500g² and found that the learned policy can successfully reorient them. We provide further analysis in the appendix showing that forces applied by the hand for such manipulation are realistic. While there is still room for the possibility that the learned policy is exploiting the simulator to reorient objects by throwing them in the air, our analysis so far indicates otherwise.

Next, to understand the failure modes, we collected one hundred unsuccessful trajectories on YCB dataset and manually analyzed them. The primary failure is in manipulating long, small, or thin objects, which accounts for 60% of all errors. In such cases, either the object slips through the fingers and falls, or is hard to be manipulated when the objects land on the palm. Another cause of failures (19%) is that objects are reoriented close to the goal orientation but not close enough to satisfy $\Delta\theta < \bar{\theta}$. Finally, the performance on YCB is lower than EGAD because objects in the YCB dataset are more diverse in their aspect ratios. Scaling these objects by constraining $l_{\max} \in [0.05, 0.12]\text{m}$ (see Section 3) makes some of these objects either too small, too big, or too thin and consequently results in failure (see Figure A-6). A detailed object-wise quantitative analysis of performance is reported in appendix Figure A-9. Results confirm that sphere-like objects such as tennis balls and orange are easiest to reorient, while long/thin objects such as knives and forks are the hardest to manipulate.

Train a student policy with a reduced state space The student policy state is $s_t^S \in \mathbb{R}^{31}$. In Table 4.1 (entries E1 and J1), we can see that π_R^S can get similarly high success rates as π_R^E . The last two columns in Table 4.1 also show that the policy is

²We change the mass of the YCB objects to be in the range of [0.3, 0.5]kg, and test π_R^E from the YCB dataset on these new objects. The success rate is around 75%.

Table 4.1: Success rates (%) of policies tested on different dynamics distribution. $\bar{\theta} = 0.1\text{rad}$. DR: domain randomization and observation/action noise. $X \rightarrow Y$: distill policy X into policy Y. The full table is in Table A.5.

Exp. ID	Dataset	State	Policy	1	2	3
				Train without DR		Train with DR
				Test without DR	Test with DR	Test with DR
B	EGAD	Full state	RNN	95.95 ± 0.8	84.27 ± 1.0	88.04 ± 0.6
E		Reduced state	RNN→RNN	91.96 ± 1.5	78.30 ± 1.2	80.29 ± 0.9
G	YCB	Full state	RNN	80.40 ± 1.6	65.16 ± 1.0	72.34 ± 0.9
J		Reduced state	RNN→RNN	81.04 ± 0.5	64.93 ± 0.2	65.86 ± 0.7

more robust to dynamics variations and observation/action noise after being trained with domain randomization.

4.2 Reorient objects with the hand facing downward

The results above demonstrate that when the hand faces upwards, RL can be used to train policies for reorienting a diverse set of geometrically different objects. A natural question to ask is, does this still hold true when the hand is flipped upside down? Intuitively, this task is much more challenging because the objects will immediately fall down without specific finger movements that stabilize the object. Because with the hand facing upwards, the object primarily rests on the palm, such specific finger movements are not required. Therefore, the hand facing downwards scenario presents a much harder exploration challenge. To verify this hypothesis, we trained a policy with the downward-facing hand, objects placed underneath the hand (see Figure 3-1b), and using the same reward function (Equation (A.1)) as before. Unsurprisingly, the success rate was 0%. The agent’s failure can be attributed to policy needing to learn to both stabilize the object under the effect of gravity and simultaneously reorient it. Deploying this policy simply results in an object falling down, confirming the hard-exploration nature of this problem.

4.2.1 Reorient objects on a table

To tackle the hard problem of reorienting objects with the hand facing downward, we started with a simplified task setup that included a table under the hand (see Figure 3-1c). Table eases exploration by preventing the objects from falling. We train π_M^ϵ using the same reward function Equation (A.1) on objects sampled from the EGAD and YCB datasets. The success rate using an MLP policy using full state information for EGAD and YCB is $95.31\% \pm 0.9\%$ and $81.59\% \pm 0.7\%$ respectively. Making use of external support for in-hand manipulation has been a challenging problem in robotics. Prior work approach this problem by building analytical models and constructing motion cones [27], which is challenging for objects with complex geometry. Our experiments show that model-free RL provides an effective alternative for learning manipulation strategies capable of using external support surfaces.

4.2.2 Reorient objects in air with hand facing downward

To enable the agent to operate in more general scenarios, we tackled the re-orientation problem with the hand facing downwards and without any external support. In this setup, one might hypothesize that object shape information (e.g., from vision) is critical because finding the strategy in Section 4.1 is not easy when the hand needs to overcome gravity and stabilize the object while reorienting it. We experimentally verify that even in this case, the policies achieve a reasonably high success rate without any knowledge of object shape.

A good pose initialization is what you need: The difficulty of directly training the RL policies when the hand faces downward is mainly because of the hard-exploration issue in learning to catch the objects that are moving downward. However, catching is not necessary for the reorientation. Even for human, we only reorient the object after we grasp it. More specifically, we first train an object-lifting policy to lift objects from the table, collect the ending state (joint positions q_T , object position p_T^o and orientation α_T^o) in each successful lifting rollout episode, and reset the hand and objects to these states (velocities are all 0) for the pose initialization in

training the reorientation policy. The objects have randomly initialized poses and are dropped onto the table. We trained a separate RNN policy for each dataset using the reward function in Section A.3.2. The success rate on the EGAD dataset is 97.80%, while the success rate on the YCB dataset is 90.11%. Note that objects need to be grasped first to be lifted. Our high success rates on object lifting also indicate that using an anthropomorphic hand makes object grasping an easy task, while many prior works [28, 29] require much more involved training techniques to learn grasping skills with parallel-jaw grippers. After we train the lifting policy, we collect about 250 ending states for each object respectively from the successful lifting episodes. In every reset during the reorientation policy training, ending states are randomly sampled and used as the initial pose of the fingers and objects. With a good pose initialization, policies are able to learn to reorient objects with high success rates. $\pi_R^\mathcal{E}$ trained on EGAD dataset gets a success rate more than 80% while $\pi_R^\mathcal{E}$ trained on YCB dataset gets a success rate greater than 50%. More results on the different policies with and without domain randomization are available in Table A.6 in the appendix. This setup is challenging because if at any time step in an episode the fingers take a bad action, the object will fall.

Improving performance using gravity curriculum: Since the difficulty of training the re-orientation policy with the hand facing downward is due to the gravity, we propose to build a gravity curriculum to learn the policy $\pi^\mathcal{E}$. Since $\pi^\mathcal{E}$ already performs very well on EGAD objects, we apply gravity curriculum to train $\pi^\mathcal{E}$ on YCB objects. Our gravity curriculum is constructed as follows: we start the training with $g = 1 \text{ m/s}^2$, then we gradually decrease g in a step-wise fashion if the evaluation success rate (w) is above a threshold value (\bar{w}) until $g = -9.81 \text{ m/s}^2$. More details about gravity curriculum are in Section A.3.4. In Table A.6 (Exp Q and T) in the appendix, we can see that adding gravity curriculum (g -curr) significantly boost the success rates on the YCB dataset.

Table 4.2: Performance of the student policy when the hand faces upward and downward

Dataset	Upward	Downward (air)	Downward (air, g -curr)
EGAD	91.96 ± 1.5	74.10 ± 2.3	/
YCB	81.04 ± 0.5	45.22 ± 2.1	67.33 ± 1.9

Table 4.3: Zero-shot RNN policy transfer success rates (%) across datasets. **U.** (**D.**) means hand faces upward (downward). **FS** (**RS**) means using full-state (reduced-state) information.

	EGAD \rightarrow YCB	YCB \rightarrow EGAD
U.FS	68.82 ± 1.7	96.41 ± 1.2
U.RS	59.64 ± 1.8	96.38 ± 1.3
D.FS	62.73 ± 2.2	85.45 ± 2.9
D.RS	55.30 ± 1.3	77.91 ± 2.1







4.3 Zero-shot policy transfer across datasets

We have shown the testing performance on the same training dataset so far. How would the policies work on a different dataset? To answer this, we test our policies across datasets: policies trained with EGAD objects are now tested with YCB objects and vice versa. We used the RNN policies trained with full-state information and reduced-state information respectively (without domain randomization) and tested them on the other dataset with the hand facing upward and downward. In the case of the hand facing downward, we tested the RNN policy trained with gravity curriculum. Table 4.3 shows that policies still perform well on the untrained dataset.

4.4 Object Reorientation with RGBD sensors

In this section, we investigate whether we can train a vision policy to reorient objects with the hand facing upward. As vision-based experiments require much more compute resources, we train one vision policy for each object individually on six objects shown in Table 4.4. We leave training a single vision policy for all objects to future work. We use the expert MLP policy trained in Section 4.1 to supervise the vision policy. We also

Table 4.4: Vision policy success rate ($\bar{\theta} = 0.2 \text{ rad}$, $\bar{d}_C = 0.01$)

	Object	Success rate (%)
	025_mug	89.67 ± 1.2
	065-d_cups	68.32 ± 1.9
	072-b_toy_airplane	84.52 ± 1.4
	073-a_lego_duplo	58.16 ± 3.1
	073-c_lego_duplo	50.21 ± 3.7
	073-e_lego_duplo	66.57 ± 3.1

performed data augmentation on the point cloud input to the policy network at each time step in both training and testing. The data augmentation includes the random translation of the point cloud, random noise on the point positions, random dropout on the points, and random variations on the point color. More details about the data augmentation are in Section A.4.5. We can see from Table 4.4 that reorienting the non-symmetric objects including the toy and the mug has high success rates (greater than 80%). While training the policy for symmetric objects is much harder, Table 4.4 shows that using d_C as an episode termination criterion enables the policies to achieve a success rate greater than 50%.

Chapter 5

Related Work

Dexterous manipulation has been studied for decades, dating back to [30, 31]. In contrast to peg insertion [32], parallel-jaw grasping, pushing, pivoting [33], or pick-and-place, dexterous manipulation typically involves continuously controlling force to the object through the fingertips of a robotic hand [34]. Some prior works used analytical kinematics and dynamics models of the hand and object, and used trajectory optimization to output control policies [1, 2, 35] or employed kinodynamic planning to find a feasible motion plan [36]. However, due to the large number of active contacts on the hand and the objects, model simplifications such as simple finger and object geometries are usually necessary to make the optimization or planning tractable. Sundaralingam and Hermans [35] moved objects in hand but assumes that there is no contact breaking or making between the fingers and the object. Furukawa et al. [37] achieved a high-speed dynamic regrasping motion on a cylinder using a high-speed robotic hand and a high-speed vision system. Prior works have also explored the use of a vision system for manipulating an object to track a planned path [38], detecting manipulation modes [39], precision manipulation [40] with a limited number of objects with simple shapes using a two-fingered gripper. Recent works have explored the application of reinforcement learning to dexterous manipulation. Model-based RL works learned a linear [4, 7] or deep neural network [5] dynamics model from the rollout data, and used online optimal control to rotate a pen or Baoding balls on a Shadow hand. However, when the system is unstable, collecting informative

trajectories for training a good dynamics model that generalizes to different objects remains challenging. Another line of works uses model-free RL algorithms to learn a dexterous manipulation policy. For example, OpenAI et al. [11] and OpenAI et al. [12] learned a controller to reorient a block or a Rubik’s cube. Van Hoof et al. [41] learned the tactile informed policy via RL for a three-finger manipulator to move an object on the table. To reduce the sample complexity of model-free learning, [9, 13, 8, 42, 6] combined imitation learning with RL to learn to reorient a pen, open a door, assemble LEGO blocks, etc. However, collecting expert demonstration data from humans is expensive, time-consuming, and even incredibly difficult for contact-rich tasks [8]. Our method belongs to the category of model-free learning. We use the teacher-student learning paradigm to speed up the deployment policy learning. Our learned policies also generalize to new shapes and show strong zero-shot transfer performance. To the best of our knowledge, our system is the first work that demonstrates the capabilities of reorienting a diverse set of objects that have complex geometries with both the hand facing upward and downward.

Chapter 6

Discussion and Conclusion

Our results show that model-free RL with simple deep learning architectures can be used to train policies to reorient a large set of geometrically diverse objects. Further, for learning with the hand facing downwards, we found that a good pose initialization obtained from a lifting policy was necessary, and the gravity curriculum substantially improved performance. The agent also learns to use an external surface (i.e., the table). The most surprising observation is that information about shape is not required despite the fact that we train a single policy to manipulate multiple objects. Perhaps in hindsight, it is not as surprising – after all, humans can close their eyes and easily manipulate novel objects into a specific orientation. Our work can serve a strong baseline for future in-hand object reorientation works that incorporate object shape in the observation space.

While we only present results in simulation, we also provide evidence that our policies can be transferred to the real world. The experiments with domain randomization indicate that learned policies can work with noisy inputs. Analysis of peak torques during manipulation (see Figure A-11 in the appendix) also reveals that generated torque commands are feasible to actuate on an actual robotic hand.

Finally, Figure A-9 and Figure A-10 in the appendix show that the success rate varies substantially with object shape. This suggests that in the future, a training curriculum based on object shapes can improve performance. Another future work is to directly train one vision policy for a diverse set of objects. A major bottleneck

in vision-based experiments is the demand for much larger GPU memory. Learning visual representations of point cloud data that can ease the computational bottleneck is, therefore, an important avenue for future research.

Bibliography

- [1] Igor Mordatch, Zoran Popović, and Emanuel Todorov. Contact-invariant optimization for hand manipulation. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 137–144, 2012. 13, 33
- [2] Yunfei Bai and C Karen Liu. Dexterous manipulation using both palm and fingers. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1560–1565. IEEE, 2014. 13, 33
- [3] Vikash Kumar, Yuval Tassa, Tom Erez, and Emanuel Todorov. Real-time behaviour synthesis for dynamic hand-manipulation. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6808–6815. IEEE, 2014. 13
- [4] Vikash Kumar, Emanuel Todorov, and Sergey Levine. Optimal control with learned local models: Application to dexterous manipulation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 378–383. IEEE, 2016. 14, 33
- [5] Anusha Nagabandi, Kurt Konolige, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. In *Conference on Robot Learning*, pages 1101–1112. PMLR, 2020. 14, 33
- [6] Abhishek Gupta, Clemens Eppner, Sergey Levine, and Pieter Abbeel. Learning dexterous manipulation for a soft robotic hand from human demonstrations. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3786–3793. IEEE, 2016. 14, 34

- [7] Vikash Kumar, Abhishek Gupta, Emanuel Todorov, and Sergey Levine. Learning dexterous manipulation policies from experience and imitation. *arXiv preprint arXiv:1611.05095*, 2016. 14, 33
- [8] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017. 14, 34
- [9] Ilija Radosavovic, Xiaolong Wang, Lerrel Pinto, and Jitendra Malik. State-only imitation learning for dexterous manipulation. *arXiv preprint arXiv:2004.04650*, 2020. 14, 34
- [10] Mayur Mudigonda, Pulkit Agrawal, Michael Deweese, and Jitendra Malik. Investigating deep reinforcement learning for grasping objects with an anthropomorphic hand. 2018. 14
- [11] OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020. 14, 34, 44
- [12] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019. 14, 34
- [13] Henry Zhu, Abhishek Gupta, Aravind Rajeswaran, Sergey Levine, and Vikash Kumar. Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3651–3657. IEEE, 2019. 14, 34
- [14] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco

- Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47), 2020. 17
- [15] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Conference on Robot Learning*, pages 66–75. PMLR, 2020. 17
- [16] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning*, pages 1407–1416. PMLR, 2018. 18, 21
- [17] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 18, 48
- [18] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017. 19
- [19] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011. 19, 48
- [20] Alex M Andrew. Multiple view geometry in computer vision. *Kybernetes*, 2001. 20, 23
- [21] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019. 21

- [22] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 21
- [23] ShadowRobot. ShadowRobot Dexterous Hand, 2005. URL <https://www.shadowrobot.com/dexterous-hand-series/>. 23
- [24] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021. 23, 49
- [25] Douglas Morrison, Peter Corke, and Jürgen Leitner. Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation. *IEEE Robotics and Automation Letters*, 5(3):4368–4375, 2020. 23
- [26] Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *2015 international conference on advanced robotics (ICAR)*, pages 510–517. IEEE, 2015. 23
- [27] Nikhil Chavan-Daffe, Rachel Holladay, and Alberto Rodriguez. In-hand manipulation via motion cones. *arXiv preprint arXiv:1810.00219*, 2018. 28
- [28] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2901–2910, 2019. 29
- [29] Jeffrey Mahler, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio Ojea, and Ken Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. 2017. 29
- [30] J Kenneth Salisbury and John J Craig. Articulated hands: Force control and kinematic issues. *The International journal of Robotics research*, 1(1):4–17, 1982.

- [31] Matthew T Mason, J Kenneth Salisbury, and Joey K Parker. Robot hands and the mechanics of manipulation. 1989. 33
- [32] Tao Chen, Adithyavairavan Murali, and Abhinav Gupta. Hardware conditioned policies for multi-robot transfer learning. In *Advances in Neural Information Processing Systems*, pages 9355–9366, 2018. 33
- [33] Yiannis Karayiannidis, Christian Smith, Danica Kragic, et al. Adaptive control for pivoting with visual and tactile feedback. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 399–406. IEEE, 2016. 33
- [34] Nikhil Chavan Daffe, Alberto Rodriguez, Robert Paolini, Bowei Tang, Sidhartha S Srinivasa, Michael Erdmann, Matthew T Mason, Ivan Lundberg, Harald Staab, and Thomas Fuhlbrigge. Extrinsic dexterity: In-hand manipulation with external forces. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1578–1585. IEEE, 2014. 33
- [35] Balakumar Sundaralingam and Tucker Hermans. Relaxed-rigidity constraints: kinematic trajectory optimization and collision avoidance for in-grasp manipulation. *Autonomous Robots*, 43(2):469–483, 2019. 33
- [36] Daniela Rus. In-hand dexterous manipulation of piecewise-smooth 3-d objects. *The International Journal of Robotics Research*, 18(4):355–381, 1999. 33
- [37] Noriatsu Furukawa, Akio Namiki, Senoo Taku, and Masatoshi Ishikawa. Dynamic regrasping using a high-speed multifingered hand and a high-speed vision system. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 181–187. IEEE, 2006. 33
- [38] Berk Calli, Andrew Kimmel, Kaiyu Hang, Kostas Bekris, and Aaron Dollar. Path planning for within-hand manipulation over learned representations of safe states. In *International Symposium on Experimental Robotics*, pages 437–447. Springer, 2018. 33

- [39] Berk Calli, Krishnan Srinivasan, Andrew Morgan, and Aaron M Dollar. Learning modes of within-hand manipulation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3145–3151. IEEE, 2018. 33
- [40] Berk Calli and Aaron M Dollar. Vision-based model predictive control for within-hand precision manipulation with underactuated grippers. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2839–2845. IEEE, 2017. 33
- [41] Herke Van Hoof, Tucker Hermans, Gerhard Neumann, and Jan Peters. Learning robot in-hand manipulation with tactile features. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 121–127. IEEE, 2015. 34
- [42] Rae Jeong, Jost Tobias Springenberg, Jackie Kay, Daniel Zheng, Yuxiang Zhou, Alexandre Galashov, Nicolas Heess, and Francesco Nori. Learning dexterous manipulation from suboptimal experts. *arXiv preprint arXiv:2010.08587*, 2020. 34
- [43] Khaled Mamou, E Lengyel, and AK Peters. Volumetric hierarchical approximate convex decomposition. In *Game Engine Gems 3*, pages 141–158. AK Peters, 2016. 46
- [44] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015. 48
- [45] Alexander Sergeev and Mike Del Balso. Horovod: fast and easy distributed deep learning in TensorFlow. *arXiv preprint arXiv:1802.05799*, 2018. 48
- [46] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 48

Appendix A

A.1 Evidence indicating transfer to real-world

Due to the scope of the problem and the ongoing pandemic, we limit our experiments to be in simulation. However, we provided evidence indicating that the learned policies can be transferred to the real world in the future in the paper. We summarize this evidence as follows.

Convex decomposition The objects after the convex decomposition still have geometrically different and complex geometries as shown in Figure A-3. The objects in the EGAD dataset are 3D printable. The YCB objects are available in the real world.

Action space We control the finger joints via relative position control as explained in Section 2.1. This suffers less sim-to-real gap compared to using torque control on the joints directly.

Student policies We designed two student policies and both of them use the observation data that can be readily acquired from the real world. The first student policy only requires the joint positions and the object pose. Object pose can be obtained using a pose estimation system or a motion capture system in the real world. Our second student policy only require the point cloud of the scene and the joint positions. We can get the point cloud in the real world by using RGBD cameras such

as Realsense D415, Azure Kinect, etc.

Domain randomization We also trained and tested our policies with domain randomization. We randomized object mass, friction, joint damping, tendon damping, tendon stiffness, etc. Table A.4 lists all the parameters we randomized in our experiments. We also add noise to the state observation and action commands as shown in Table A.4. For the vision experiments, we also added noise (various ways of data augmentation including point position jittering, color jittering, dropout, etc.) to the point cloud observation in training and testing as explained in Section A.4.5.

The results in Table A.5, Table A.6, and Table 4.4 show that even after adding randomization/noise, we can still get good success rates with the trained policies. Even though we cannot replicate the true real-world setups in the simulation, our results with domain randomization indicates a high possibility that our policies can be transferred to the real Shadow hand. Prior works [11] have also shown the domain randomization can effectively reduce the sim-to-real gap.

Torque analysis We also conducted torque analysis as shown in Section A.4.4. We can see that the peak torque values remain in an reasonable and affordable range for the Shadow hand. This indicates that our learned policies are less likely to cause motor overload on the real Shadow hand.

A.2 Environment Setup




Figure A-1: We learn policies that can reorient many objects in three scenarios respectively: (1) hand faces upward, (2) hand faces downward with a table below the hand, (3) hand faces downward without any table. The extra object in each figure shows the desired orientation.

A.2.1 State definition

The full state space $\mathcal{S}^{\mathcal{E}}$ includes joint, fingertip, object and goal information detailed in Table A.1. To compute the angle difference between two quaternion orientations α_1 and α_2 , we first compute the difference rotation quaternion: $\beta = \alpha_1\alpha_2^{-1}$. Then the angle difference (distance between two rotations) $\Delta\theta$ is computed as the angle of β from the axis-angle representation of β .

Table A.1: Full state $s_t^{\mathcal{E}} \in \mathbb{R}^{134}$ information. Orientations are in the form of quaternions.

Parameter	Description	Parameter	Description	Parameter	Description
$q_t \in \mathbb{R}^{24}$	joint positions	$v_t^f \in \mathbb{R}^{15}$	fingertip linear velocities	$\alpha^g \in \mathbb{R}^4$	object goal orientation
$\dot{q}_t \in \mathbb{R}^{24}$	joint velocities	$w_t^f \in \mathbb{R}^{15}$	fingertip angular velocities	$v_t^o \in \mathbb{R}^3$	object linear velocity
$p_t^f \in \mathbb{R}^{15}$	fingertip positions	$p_t^o \in \mathbb{R}^3$	object position	$w_t^o \in \mathbb{R}^3$	object angular velocity
$\alpha_t^f \in \mathbb{R}^{20}$	fingertip orientation	$\alpha_t^o \in \mathbb{R}^4$	object orientation	$\beta_t \in \mathbb{R}^4$	$\alpha_t^o(\alpha^g)^{-1}$

A.2.2 Dataset

We use two object datasets (EGAD and YCB) in our paper. To further increase the diversity of the datasets, we create 5 variants for each object mesh by randomly

scaling the mesh. The scaling ratios are randomly sampled such that the longest side of the objects’ bounding boxes l_{\max} lies in $[0.05, 0.08]$ m for EGAD objects, and $l_{\max} \in [0.05, 0.12]$ m for YCB objects. The mass of each object is randomly sampled from $[0.05, 0.15]$ kg. When we randomly scale YCB objects, some objects become very small and/or thin, making the reorientation task even more challenging. In total, we use 11410 EGAD object meshes and 390 YCB object meshes for training.

Figure A-2 shows examples from the EGAD and YCB dataset. We can see that these objects are geometrically different and have complex shapes. We also use V-HACD [43] to perform an approximate convex decomposition on the object meshes for fast collision detection in the simulator. Figure A-3 shows the object shapes before and after the decomposition. After the decomposition, the objects are still geometrically different.




Figure A-2: First row: examples of EGAD objects. Second row: examples of YCB objects.

A.2.3 Camera setup

We placed two RGBD cameras above the hand, as shown in Figure A-4. In ISAAC gym, we set the camera pose by setting its position and focus position. The two cameras’ positions are shifted from the Shadow hand’s base origin by $[-0.6, -0.39, 0.8]$ and $[0.45, -0.39, 0.8]$ respectively. And their focus points are the points shifted from the Shadow hand’s base origin by $[-0.08, -0.39, 0.15]$ and $[0.045, -0.39, 0.15]$ respectively.




Figure A-3: Examples of EGAD objects. The first and third row shows the visual mesh of the objects. The second and fourth row show the corresponding collision mesh (after V-HACD decomposition).




Figure A-4: Camera positions

A.3 Experiment Setup

A.3.1 Network architecture

For the non-vision policies, we experimented with two architectures: The MLP policy π_M consists of 3 hidden layers with 512, 256, 256 neurons respectively. The RNN policy

π_R has 3 hidden layers (512 – 256 – 256), followed by a 256-dim GRU layer and one more 256-dim hidden layer. We use the exponential linear unit (ELU) [44] as the activation function.

For our vision policies, we design a sparse convolutional network architecture (*Sparse3D-IMPALA-Net*). As shown in Figure 2-1, the point cloud W_t is processed by a series of sparse CNN residual modules and projected into an embedding vector. q_t and a_{t-1} are concatenated together and projected into an embedding vector via an MLP. Both embedding vectors from W_t and (q_t, a_{t-1}) are concatenated and passed through a recurrent network to output the action a_t .

A.3.2 Training details

All the experiments in the paper were run on at most 2 GPUs with a 32GB memory. We use PPO [17] to learn $\pi^{\mathcal{E}}$. Table A.2 lists the hyperparameters for the experiments. We use 40K parallel environments for data collection. We update the policy with the rollout data for 8 epochs after every 8 rollout steps for the MLP policies and 50 rollout steps for the RNN policies. A rollout episode is terminated (reset) if the object is reoriented to the goal orientation successfully, or the object falls, or the maximum episode length is reached. To learn the student policies $\pi^{\mathcal{S}}$, we use Dagger[19]. While Dagger typically keep all the state-action pairs for training the policy, we do Dagger in an online fashion where $\pi^{\mathcal{S}}$ only learns from the latest rollout data.

For the vision experiments, the number of parallel environments is 360 and we update policy after every 50 rollout steps from all the parallel environments. The batch size is 30. We sample 15000 points from the reconstructed point cloud of the scene from 2 cameras for the scene point cloud W_t^s and sample 5000 points from the object CAD mesh model for the goal point cloud W^g .

We use Horovod [45] for distributed training and Adam [46] optimizer for neural network optimization.

Reward function for reorientation: For training $\pi^{\mathcal{E}}$ for the reorientation task,

we modified the reward function proposed in [24] to be:

$$r(s_t, a_t) = c_{\theta_1} \frac{1}{|\Delta\theta_t| + \epsilon_\theta} + c_{\theta_2} \mathbb{1}(|\Delta\theta_t| < \bar{\theta}) + c_3 \|a_t\|_2^2 \quad (\text{A.1})$$

where $c_{\theta_1} > 0$, $c_{\theta_2} > 0$ and $c_3 < 0$ are the coefficients, $\Delta\theta_t$ is the difference between the current object orientation and the target orientation, ϵ_θ is a constant, $\mathbb{1}$ is an indicator function that identifies whether the object is in the target orientation. The first two reward terms encourage the policy to reorient the object to the desired orientation while the last term suppresses large action commands.

Reward function for object lifting: To train the lifting policy, we use the following reward function:

$$r(s_t, a_t) = c_{h_1} \frac{1}{|\Delta h_t| + \epsilon_h} + c_{h_2} \mathbb{1}(|\Delta h_t| < \bar{h}) + c_3 \|a_t\|_2^2 \quad (\text{A.2})$$

where $\Delta h_t = \max(p_t^{b,z} - p_t^{o,z}, 0)$ and $p_t^{b,z}$ is the height (z coordinate) of the Shadow Hand base frame, $p_t^{o,z}$ is the height of the object, \bar{h} is the threshold of the height difference. The objects have randomly initialized poses and are dropped onto the table.

Goal specification for vision policies: We obtain W^g by sampling 5000 points from the object’s CAD mesh using the Barycentric coordinate, rotating the points by the desired orientation, and translating them so that these points are next to the hand. Note that one can also put the object in the desired orientation right next to the hand in the simulator and render the entire scene altogether to remove the need for CAD models. We use CAD models for W^g just to save the computational cost of rendering another object while we still use RGBD cameras to get W_t^s .

A.3.3 Dynamics randomization

Table A.4 list all the randomized parameters as well the state observation noise and action command noise.

Comparing the Column 1 and Column 2 in Table A.5, we can see that if we directly

Table A.2: Hyperparameter Setup

Hyperparameter	Value	Hyperparameter	Value	Hyperparameter	Value
Num. batches	5	Entropy coeff.	0.	Num. pts sampled from W_t^s	15000
Actor learning rate	0.0003	GAE parameter	0.95	Num. pts sampled from W^g	5000
Critic learning rate	0.001	Discount factor	0.99	Num. envs	40000
Num. epochs	8	Episode length	300	Num. rollout steps per policy update (MLP/RNN)	8/50
Value loss coeff.	0.0005	PPO clip range	0.1	c_{θ_1}	1
c_{θ_2}	800	c_3	0.1	ϵ_θ	0.1
$\bar{\theta}$	0.1rad	c_{h_1}	0.05	ϵ_h	0.02
\bar{h}	0.04	c_{h_2}	800		

Table A.3: Mesh Parameters

Parameter	Range
longest side of the bounding box of EGAD objects	[0.05, 0.08]m
longest side of the bounding box of YCB objects	[0.05, 0.12]m
mass of each object	[0.05, 0.15]kg
No. of EGAD meshes	2282
No. of YCB meshes	78
No. of variants per mesh	5
Voxelization resolution	0.003 m

Table A.4: Dynamics Randomization and Noise

Parameter	Range	Parameter	Range	Parameter	Range
state observation	$+\mathcal{U}(-0.001, 0.001)$	action	$+\mathcal{N}(0, 0.01)$	joint stiffness	$\times\mathcal{E}(0.75, 1.5)$
object mass	$\times\mathcal{U}(0.5, 1.5)$	joint lower range	$+\mathcal{N}(0, 0.01)$	tendon damping	$\times\mathcal{E}(0.3, 3.0)$
object static friction	$\times\mathcal{U}(0.7, 1.3)$	joint upper range	$+\mathcal{N}(0, 0.01)$	tendon stiffness	$\times\mathcal{E}(0.75, 1.5)$
finger static friction	$\times\mathcal{U}(0.7, 1.3)$	joint damping	$\times\mathcal{E}(0.3, 3.0)$		

$\mathcal{N}(\mu, \sigma)$: Gaussian distribution with mean μ and standard deviation σ .

$\mathcal{U}(a, b)$: uniform distribution between a and b . $\mathcal{E}(a, b) = \exp^{\mathcal{U}(\log(a), \log(b))}$.

+: the sampled value is added to the original value of the variable. \times : the original value is scaled by the sampled value.

deploy the policy trained without domain randomization into an environment with different dynamics, the performance drops significantly. If we train policies with domain randomization (Column 3), the policies are more robust and the performance only drops slightly compared to Column 1 in most cases. The exceptions are on C3 and H3. In these two cases, the π_M^S policies collapsed in training during the policy distillation along with the randomized dynamics.

A.3.4 Gravity curriculum

We found building a curriculum on gravity helps improve the policy learning for YCB objects when the hand faces downward. Algorithm 1 illustrates the process of building the gravity curriculum. In our experiments, we only test on the training objects once (one random initial and goal orientation) to get the surrogate average success rate w on all the objects during training. $\bar{w} = 0.8, g_0 = 1 \text{ m/s}^2, \Delta g = -0.5 \text{ m/s}^2, K = 3, L = 20, \Delta T_{\min} = 40$.

Algorithm 1 Gravity Curriculum

```
1: Initialize an empty FIFO queue  $Q$  of size  $K, \Delta T = 0, g = g_0$ 
2: for  $i \leftarrow 1$  to  $M$  do
3:    $\tau = \text{rollout\_policy}(\pi_\theta)$  ▷ get rollout trajectory
4:    $\pi_\theta = \text{optimize\_policy}(\pi_\theta, \tau)$  ▷ update policy
5:    $\Delta T = \Delta T + 1$ 
6:   if  $i \bmod L = 0$  then
7:      $w = \text{evaluate\_policy}(\pi_\theta)$  ▷ evaluate the policy, get the success rate  $w$ 
8:     append  $w$  to the queue  $Q$ 
9:     if  $\text{avg}(Q) > \bar{w}$  and  $\Delta T > \Delta T_{\min}$  then
10:       $g = \max(g - \Delta g, -9.81)$ 
11:       $\Delta T = 0$ 
12:    end if
13:  end if
14: end for
```

A.4 Supplementary Results

A.4.1 Hand faces upward

Learning curves Figure A-5 shows the learning curve of the RNN and MLP policies on the EGAD and YCB datasets. Both policies learn well on the EGAD and YCB datasets. The YCB dataset requires much more environment interactions for the policies to learn. We can also see that using the full-state information can speed up the learning and give a higher final performance.




Figure A-5: Learning curves of the MLP and RNN policies on the EGAD (**Left**) and YCB datasets (**Middle**). The **Right** plot shows that using the full state information speeds up the policy learning compared to only using the reduced state information.

Testing performance - Teacher The testing results in Table A.5 show that both the MLP and RNN policies are able to achieve a success rate greater than 90% on the EGAD dataset (entries A1, B1) and greater than 70% on the YCB dataset (entries F1, G1) without any explicit knowledge of the object shape. This result is surprising because intuitively, one would assume that information about the object shape is important for in-hand reorientation.

Testing performance - Student We experimented with the following three combinations: (1) distill $\pi_M^\mathcal{E}$ into $\pi_M^\mathcal{S}$, (2) distill $\pi_M^\mathcal{E}$ into $\pi_R^\mathcal{S}$, (3) distill $\pi_R^\mathcal{E}$ into $\pi_R^\mathcal{S}$. The student policy state is $s_t^\mathcal{S} \in \mathbb{R}^{31}$. In Table A.5 (entries C1-E1, H1-J1), we can see that $\pi_R^\mathcal{E} \rightarrow \pi_R^\mathcal{S}$ gives the highest success rate on $\pi^\mathcal{S}$, while $\pi_M^\mathcal{E} \rightarrow \pi_M^\mathcal{S}$ leads to much worse performance (36% drop of success rate in EGAD, and 47% drop in YCB). This shows that $\pi^\mathcal{S}$ requires temporal information due to reduced state space. The last two columns in Table A.5 also show that the policy is more robust to dynamics variations and observation/action noise after being trained with domain randomization.

Failure cases Figure A-6 shows some example failure cases. If the objects are too small, thin, or big, the objects are likely to fall. If objects are initialized close to the hand border, it is much more difficult for the hand to catch the objects. Another failure mode is that the objects are reoriented close to the goal orientation but not close enough to satisfy $\Delta\theta \leq \bar{\theta}$.

Table A.5: Success rates (%) of policies tested on different dynamics distribution. $\bar{\theta} = 0.1\text{rad}$. DR stands for domain randomization and observation/action noise. $X \rightarrow Y$: distill policy X into policy Y.

Exp. ID	Dataset	State	Policy	1	2	3
				Train without DR		Train with DR
				Test without DR	Test with DR	Test with DR
A	EGAD	Full state	MLP	92.55 ± 1.3	78.24 ± 2.4	91.92 ± 0.4
B			RNN	95.95 ± 0.8	84.27 ± 1.0	88.04 ± 0.6
C		Reduced state	MLP→MLP	55.55 ± 0.2	25.09 ± 3.0	23.77 ± 1.8
D			MLP→RNN	85.32 ± 1.2	68.31 ± 2.6	81.05 ± 1.2
E			RNN→RNN	91.96 ± 1.5	78.30 ± 1.2	80.29 ± 0.9
F	YCB	Full state	MLP	73.40 ± 0.2	54.57 ± 0.6	66.00 ± 2.7
G			RNN	80.40 ± 1.6	65.16 ± 1.0	72.34 ± 0.9
H		Reduced state	MLP→MLP	34.08 ± 0.9	12.08 ± 0.4	5.41 ± 0.3
I			MLP→RNN	69.30 ± 0.8	47.38 ± 0.6	53.07 ± 0.9
J			RNN→RNN	81.04 ± 0.5	64.93 ± 0.2	65.86 ± 0.7




Figure A-6: Examples of failure cases. (a) and (b): objects are too small. (c): the object is reoriented close to the target orientation, but not close enough. (d): the object is too big and initialized around the palm border.

A.4.2 Hand faces downward (in the air)

Testing performance For the case of reorienting objects in the air with the hand facing downward Table A.6 lists the success rates of different policies trained with/without domain randomization, and tested with/without domain randomization.

Example visualization We show an example of reorienting a cup in Figure A-7 and an example of reorienting a sponge in Figure A-8. More examples are available at <https://taochenshh.github.io/projects/in-hand-reorientation>.

A.4.3 Success rate on each type of YCB objects

We also analyzed the success rates on each object type in the YCB dataset. Using the same evaluation procedure described in Section 3, we get the success rates of each object using $\pi_R^{\mathcal{E}}$. Figure A-9 shows the distribution of the success rates on YCB objects with the hand facing upward while Figure A-10 corresponds to the case of

Table A.6: Success rates (%) of policies trained with hand facing downward and to reorient objects in the air. Due to the large number of environment steps required in this setup, we fine-tune the model trained without DR with randomized dynamics instead of training models with DR from scratch.

Exp. ID	Dataset	State	Policy	1	2	3
				Train without DR		Finetune with DR
				Test without DR	Test with DR	Test with DR
K	EGAD	Full state	MLP	84.29 ± 0.9	38.42 ± 1.5	71.44 ± 1.3
L			RNN	82.27 ± 3.3	36.55 ± 1.4	67.44 ± 2.1
M		Reduced state	MLP→RNN	77.05 ± 1.6	29.22 ± 2.6	59.23 ± 2.3
N			RNN→RNN	74.10 ± 2.3	37.01 ± 1.5	62.64 ± 2.9
O	YCB	Full state	MLP	58.95 ± 2.0	26.04 ± 1.9	44.84 ± 1.3
P			RNN	52.81 ± 1.7	26.22 ± 1.0	40.44 ± 1.5
Q		RNN + <i>g</i> -curr	74.74 ± 1.2	25.56 ± 2.9	54.24 ± 1.4	
R		MLP→RNN	46.76 ± 2.5	25.49 ± 1.4	34.14 ± 1.3	
S		Reduced state	RNN→RNN	45.22 ± 2.1	24.45 ± 1.2	31.63 ± 1.6
T		RNN + <i>g</i> -curr→RNN	67.33 ± 1.9	19.77 ± 2.8	48.58 ± 2.3	




Figure A-7: An example of reorienting a cup with the hand facing downward. From left to right, top to bottom, we show the some moments in an episode.

reorienting the objects in the air with the hand facing downward. We can see that sphere-like objects such as tennis balls and orange are easiest to reorient. Long or thin objects such as knives and forks are the hardest ones to manipulate.

A.4.4 Torque analysis

We randomly sampled 100 objects from the YCB dataset, and use our RNN policy trained without domain randomization with the hand facing downward to reorient each of these objects 200 times. We record the joint torque values for each finger joint




Figure A-8: An example of reorienting a sponge with the hand facing downward. From left to right, top to bottom, we show the some moments in an episode.

at each time step. Let the joint torque value of i^{th} joint at time step j in k^{th} episode be J_{ij}^k . We plot the distribution of $\{\max_{i=\llbracket 1,24 \rrbracket} |J_{ij}^k| \mid j \in \llbracket 1, T \rrbracket, k = \llbracket 1, 20000 \rrbracket\}$, where $\llbracket a, b \rrbracket$ represents $\{x \mid x \in [a, b], x \in \mathbb{Z}\}$. Figure A-11 shows that the majority of the maximum torque magnitude is around 0.2 N m.

A.4.5 Vision experiments with noise

We also trained our vision policies with noise added to the point cloud input. We added the following transformations to the point cloud input.

We applied four types of transformations on the point cloud:







- *RandomTrans*: Translate the point cloud by $[\Delta x, \Delta y, \Delta z]$ where $\Delta x, \Delta y, \Delta z$ are all uniformly sampled from $[-0.04, 0.04]$.
- *JitterPoints*: Randomly sample 10% of the points. For each sampled point i , jitter its coordinate by $[\Delta x_i, \Delta y_i, \Delta z_i]$ where $\Delta x_i, \Delta y_i, \Delta z_i$ are all sampled from a Normal

distribution $\mathcal{N}(0, 0.01)$ (truncated at -0.015m and 0.015m).

- *RandomDropout*: Randomly dropout points with a dropout ratio uniformly sampled from $[0, 0.4]$.
- *JitterColor*: Jitter the color of points with the following 3 transformations: (1) jitter the brightness and rgb values, (2) convert the color of 30% of the points into gray, (3) jitter the color contrast. Each of this transformation can be applied independently with a probability of 30% if *JitterColor* is applied.

Each of these four transformations is applied independently with a probability of 40% for each point cloud at every time step. Table A.7 shows the success rates of the vision policies trained with the aforementioned data augmentations until policy convergence and tested with the same data augmentations. We found that adding the data augmentation in training actually helps improve the data efficiency of the vision policy learning even though the final performance might be a bit lower. As a reference, we show the policy performance trained and tested without data augmentation in Table A.7. For the mug object, adding data augmentation in training improves the final testing performance significantly. Without data augmentation, the learned policy reorients the mug to a pose where the body of the mug matches how the mug should look in the goal orientation, but the cup handle does not match. Adding the data augmentation helps the policy to get out of this local optimum.

Table A.7: Vision policy success rates when the policy is trained and tested with/without data augmentation ($\bar{\theta} = 0.2\text{rad}$, $\bar{d}_C = 0.01$)

	Object	Without data augmentation (noise)	With data augmentation (noise)
	025_mug	36.51 ± 2.8	89.67 ± 1.2
	065-d_cups	79.17 ± 2.3	68.32 ± 1.9
	072-b_toy_airplane	90.25 ± 3.7	84.52 ± 1.4
	073-a_lego_duplo	62.16 ± 3.7	58.16 ± 3.1
	073-c_lego_duplo	58.21 ± 3.9	50.21 ± 3.7
	073-e_lego_duplo	76.57 ± 3.6	66.57 ± 3.1

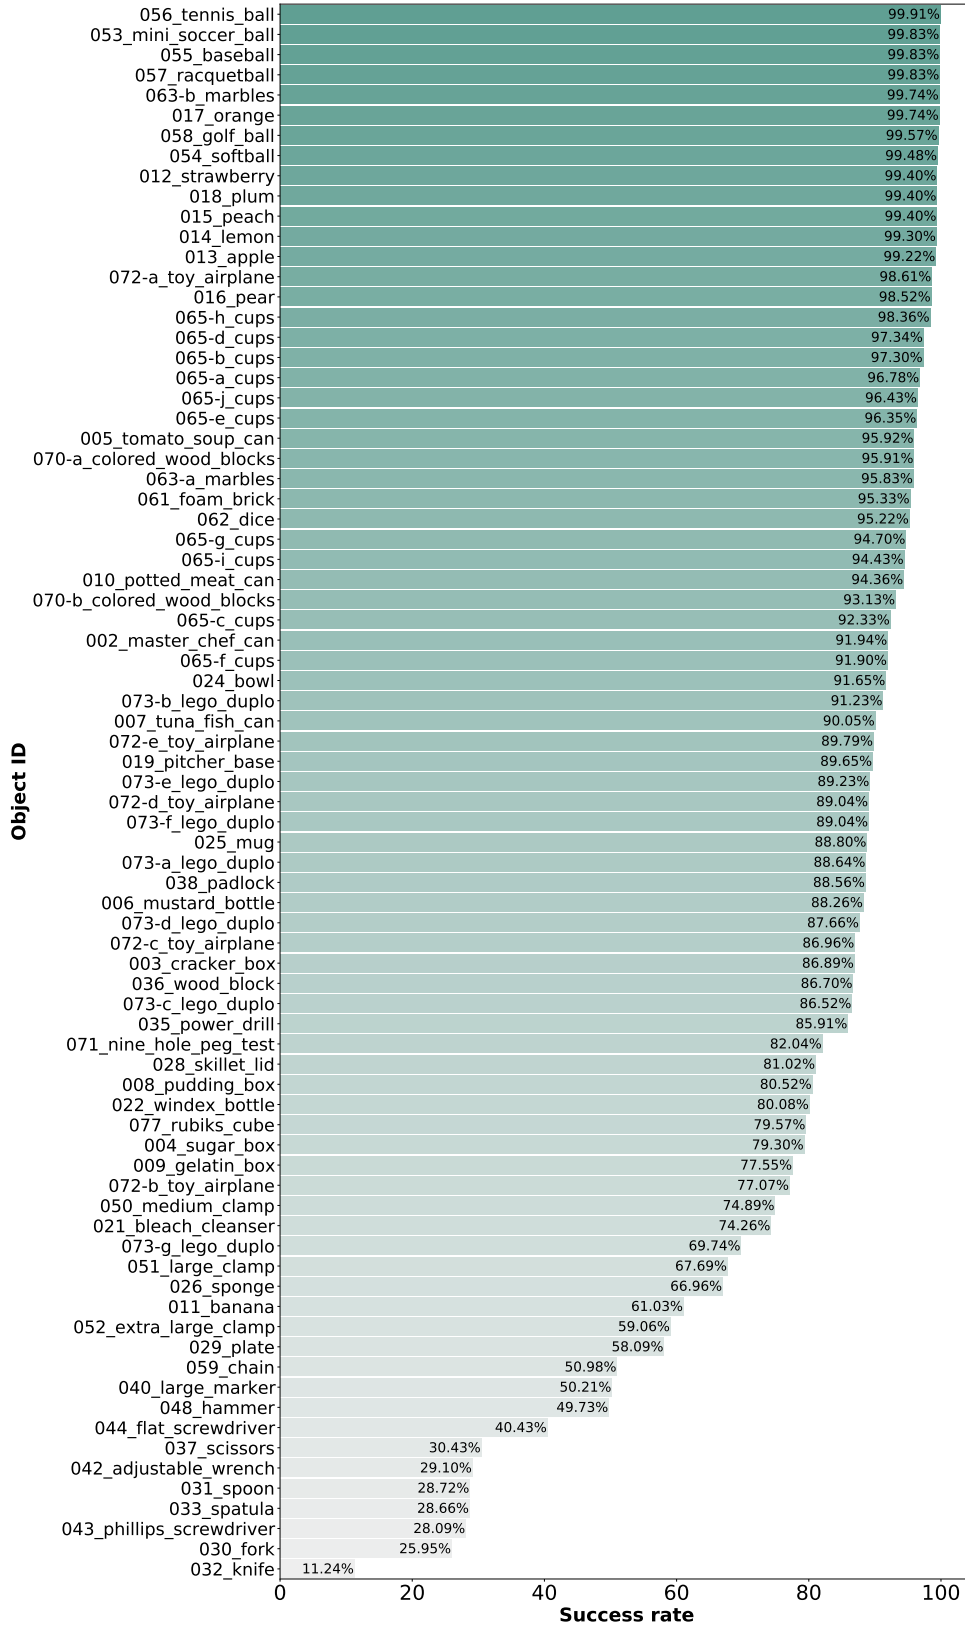


Figure A-9: Reorientation success rates for each object in the YCB dataset when the hand faces upward.




Figure A-10: Reorientation success rates for each object in the YCB dataset when the hand faces downward without a table.




Figure A-11: Distribution of the *maximum* absolute joint torque values on all joints for all the time steps. We exclude a few outliers in the plot, i.e., we only plot the data distributions up to 99% quantile.