Generalizable Point Cloud Policy Learning for Sim-to-Real Dexterous Manipulation

Anonymous Author(s)
Affiliation
Address
email

Abstract: We propose a sim-to-real framework for dexterous manipulation which can generalize to new objects of the same category in the real world. The key of our framework is to train the manipulation policy with point cloud inputs and dexterous hands. We propose two new techniques to enable joint learning on multiple objects and sim-to-real generalization: (i) using imagined hand point clouds as augmented inputs; and (ii) designing novel contact-based rewards. We empirically evaluate our method using an Allegro Hand to grasp novel objects in both simulation and real world. To the best of our knowledge, this is the first policy learning-based framework that achieves such generalization results with dexterous hands. Our anonymous project page is available at http://dexpc.github.io/.

Keywords: Dexterous Manipulation, Policy Learning, Point Clouds, Sim-to-Real

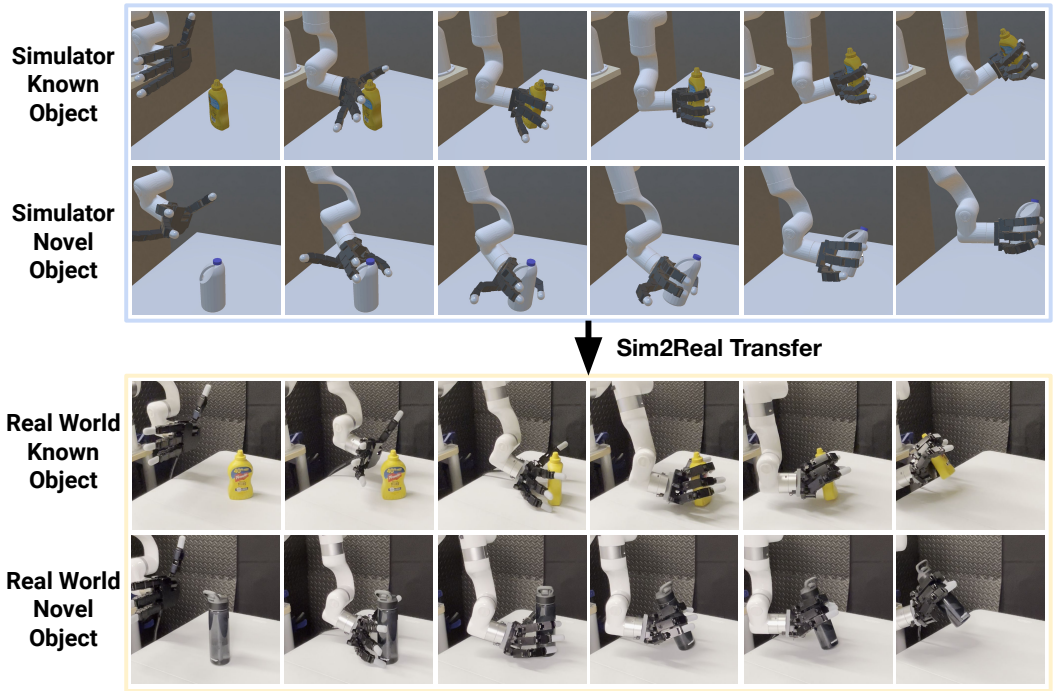

Figure 1: We introduce a policy learning method which takes the point cloud as input for dexterous manipulation tasks. By introducing several techniques in the policy learning process, our point cloud based policy can successfully generalize to novel objects and transfer from simulation to real-world.

Submitted to the 6th Conference on Robot Learning (CoRL 2022). Do not distribute.

1 Introduction

Dexterous manipulation has remained to be one of the most challenging problems in robotics [1]. While multi-finger hands create ample opportunities for our robots to flexibly manipulate objects in our daily life, the nature of the high degree of freedom and high-dimensional action space creates significant optimization challenges for both search-based planning algorithms and policy learning algorithms. Recent efforts using model-free Reinforcement Learning have achieved encouraging results on complex manipulation tasks [1, 2]; however, it still faces many challenges in generalizing to diverse objects and being deployed on real robots.

For example, the dexterous manipulation framework proposed by OpenAI et al. [1] can solve in-hand manipulation of Rubik's Cube with RL and transfer to the real robot hand. However, the policy is only trained with one particular object and it is not able to *generalize to diverse objects*. To achieve cross-object generalization, recent efforts proposed to learn robust 3D point cloud representations [3, 4, 5, 6] with diverse objects using RL in simulation. While point cloud input has also been shown easier for Sim2Real transfer [7] given its focus on geometry instead of texture, the assumption on the access of complete object point clouds and ground truth states limit the transferability of above methods to the *real robot deployment*. Among these works, Chen et al. [3] showed that cross-object generalization is achievable in simulation without knowing the shape of the object to grasp, but its requirement of real-time access to the object states is itself a very challenging robot perception problem, especially under large occlusions caused by real robot hands.

In this paper, we provide a sim-to-real policy learning framework for dexterous manipulation of unseen objects of the same category, using the grasping task with the Allegro Hand[1] as the platform. We perform our studies with two types of learning algorithms: (i) train a point cloud based reinforcement learning policy jointly with multiple objects; (ii) first train a state-based RL policy jointly with multiple objects, and then perform visual distillation (on point clouds) with teacher-student behavior cloning. With both approaches, we list the three discoveries of our framework for learning generalizable point cloud policy below:

(i) We justified that it is possible to achieve *direct sim-to-real transfer* for a dexterous grasping policy with category-level generalizability when we use point cloud as the data representation;

(ii) Raw point clouds captured by sensors often come with heavy occlusions and noise. For example, a very small portion of the points from the observation are representing the robot fingers. We propose to *imagine* the complete robot hand point clouds according to the internal robot states and use them to augment the occluded real point cloud observations. We find that explicitly augmenting the input by *imagined* points can help achieve better robustness and sample efficiency for policy learning.

(iii) Different from existing works that add contact information to the input of RL, we design a novel reward using contact pair information without adding contact to the observation. This practice remarkably improves sample efficiency as well as learning stability and avoids the dependency on contact sensor that is often unavailable for real robot models.

2 Related Work

Dexterous Manipulation. Dexterous manipulation aims to enable robotic hands to achieve human-level dexterity for grasping and manipulating objects. Analytical methods have been proposed to adopt planning to solve this problem [8, 9, 10, 11, 12, 13, 14]. However, they rely on detailed object models, which are not accessible when testing on unseen objects. To mitigate this issue, researchers propose a learning-based method for dexterous manipulation [15]. Some methods [16, 17, 1, 2, 3, 5] take a reinforcement learning (RL) approach, while another line of work [18, 19, 20, 21] also proposes to learn from demonstration using imitation learning (IL) to acquire the control policy. Recently, the combination of RL and IL [22, 23, 24, 25, 4] has also shown encouraging dexterous manipulation results in simulation. However, most methods are either learning policy for one single

[1] http://www.simlab.co.kr/Allegro-Hand.htm

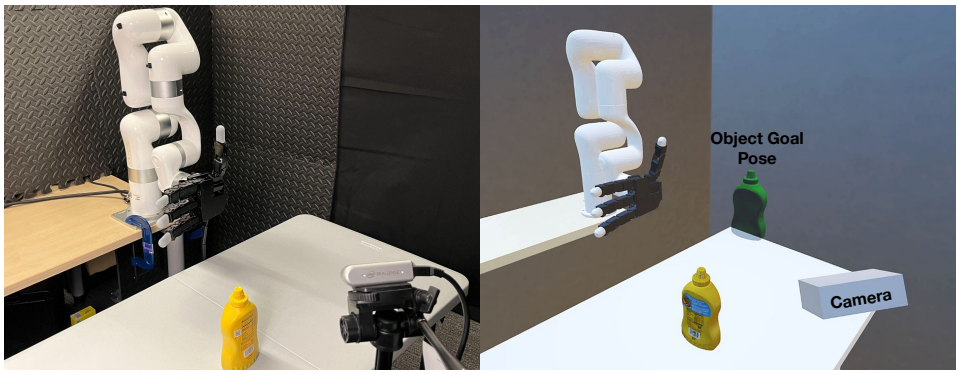

Figure 2: **Our Dexterous Manipulation System.** Left: Real-world setup with an Allegro Hand attached on an XArm6 and a RealSense D435. Right: Simulator setup with the transparent green object representing goal pose.

object or assume access to object states produced by perfect detector which increases challenges to Sim2Real transfer. In this paper, we surpass these limitations by introducing training on multiple objects and applying point cloud inputs for control.

Point Cloud in Robotic Manipulation. Point cloud representation has been widely applied in the robotics community. Researchers have studied matching the observed point clouds to an object in a grasping dataset [26] and executing the corresponding grasping action [27, 13]. For learning-based approaches, one line of research has focused on first estimating grasp proposals or affordance given the point cloud input and then planning accordingly for manipulation [28, 29, 30, 31, 32, 33, 34, 35]. While these methods are designed for parallel-jaw grippers, recent advancements have also been made for grasping with dexterous hands with similar approaches [36, 37, 38, 39]. However, this line of methods requires feedback from motion planning to estimate whether a grasp is plausible. Oftentimes, a stable grasp pose is proposed but it is not achievable by planning. To achieve efficient and flexible manipulation, a Reinforcement Learning policy with point cloud inputs is proposed [3, 5, 4], which allows the robot hand to flexibly adjust its pose while interacting with the object. However, these approaches still face challenges in transferring to the real robot, given the noisy, occluded point clouds in the real world, and training RL policy with noisy point clouds increases the optimization difficulty. In this paper, we propose to use imaginary point clouds, contact information, and multi-object training to combat these problems and achieve Sim2Real generalization.

Using Contact Information for Manipulation. Humans can manipulate objects purely from tactile sensing without seeing them. This biological fact inspires researchers to integrate contact and tactile information into the learning pipeline [40, 41, 42, 43, 44, 45, 46]. For example, both tactile and visual information inputs are combined together in [45] for decision making. The tactile information is also utilized as inputs for RL-based manipulation policies [47, 48, 49, 23, 50]. For example, visual-tactile sensor is used with an Allegro hand in simulator for playing piano [50]. Instead of using contacts as inputs, some methods also use contact and tactile information as reward to encourage exploration and boost policy learning [48, 23]. Motivated by these works, we provide a novel design of reward based on each contact link of the robot hand, which encourages more reasonable grasping behavior. Our design does not require a real tactile sensor when training in simulation or deploying in the real robot.

3 Approach

Our objective is to train a generalizable point cloud policy on a dexterous robot hand-arm system that is able to grasp a wide range of objects. As is discussed before, training such a policy comes with numerous technical challenges, including reward design and imperfect point cloud information. In this work, we propose to novel reward design based on contact and imagined point cloud to deal with these challenges. In the following sections, we first introduce the setup in section 3.1. We

introduce reward design and imagined point cloud in section 3.2 and 3.3. Finally, we briefly discuss the implementation in section 3.4.

Preliminaries: We model the dexterous manipulation problem as a Partially Observable Markov Decision Process (POMDP) $\mathcal{M} = (\mathcal{O}, \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \mathcal{U})$. Here, $\mathcal{O}$ is the observation space, $\mathcal{S}$ is the underlying state space, $\mathcal{A}$ is the action space, $\mathcal{R}$ is the reward function, $\mathcal{T}$ is the transition dynamics, and $\mathcal{U}$ generates agent's observation. At timestep t, the enviroment is at the state $s_t \in \mathcal{S}$. The agent observes $o_t \sim \mathcal{U}(\cdot|s_t) \in \mathcal{O}$. The agent takes action a_t and receives reward $r_t = \mathcal{R}(s_t, a_t)$. The environment state at timestep $t+1$ then transit to $s_{t+1} \sim \mathcal{T}(s_t, a_t)$. The objective of the agent is to maximize the return $\sum_{t=0}^{T} \gamma^t r_t$, where γ is a discount factor.

3.1 Environment Setup

System Setup: Many previous works on learning-based dexterous manipulation attach the hand to a fixed platform to simplify the experiment environment in the real world. In this work, we create a more flexible and powerful dexterous manipulation system which includes both the robotic hand and the arm (Figure 2). Concretely, we attach the Allegro Hand to an XArm6 robot. Allegro Hand is a 16-DoF anthropomorphic hand with four fingers and XArm6 is a 6-DoF robot arm. We place a RealSense D435 camera at the right front of the robot to capture the point cloud. This setup brings additional challenges to RL exploration and Sim2Real deployment. We use SAPIEN [51] platform to build the simulation environment of the whole system. The simulation time step is $0.005s$ to ensure stable contact simulation. Each control step lasts for $0.05s$.

Tasks and Objects: In this paper, we expect our robot to perform the grasping task over a diverse set of objects. In this grasping task, we first select a random object from an object dataset and place it on the table. The robot is then required to move it to a target pose. Moreover, the robot should be able to generalize to different initial states, so we randomize both the initial pose and the goal pose for each trial. In simulation experiments, we use bottles and cans from both ShapeNet [52] dataset and YCB [53] dataset. In real world experiments, we only use bottles from YCB. More details about the object set can be found in appendix.

Observation Space: The observation contains both visual and proprioceptive information with four modalities: (1) Observed point cloud provided by the camera; (2) Proprioception signals of the robot including joint positions and end-effector position; (3) Imagined hand point cloud proposed in Sec. 3.3; (4) Object goal position provided in each trial. All the information is accessible on the real robot. The dimension of each observation modality is shown in Figure 3.

Action Space: The action is responsible for controlling both the 6-DOF robot arm and the 16-DOF hand. It has $6 + 16 = 22$ dimension in total. The robot arm is parametrized by the 6D translation and rotation of the end-effector relative to a reference pose. We use the damped least square inverse kinematics solver with a damping constant $\lambda = 0.05$ to compute the joint motion. Each finger joint of the Allegro hand is controlled by a position controller. Both robot arm and hand are controlled by PD controllers.

Network Architecture: The network architecture is visualized in Figure 3,

3.2 Reward Design with Oracle Contact

Since we aim to solve the dexterous manipulation problem with pure RL, the reward design is central to the method. We need a good reward function to ensure proper interaction between the robotic hand and the object. The whole interaction process consists of two phases. The first phase is to simply reach the object. The second phase is to grasp the object and move it to the target, which is more challenging. For the first phase, we encourage reaching with the following reaching reward:

$$r_{\text{reach}} = \sum_{\text{finger}} \frac{1}{\epsilon_r + d(\mathbf{x}_{\text{finger}}, \mathbf{x}_{\text{obj}})}. \tag{1}$$

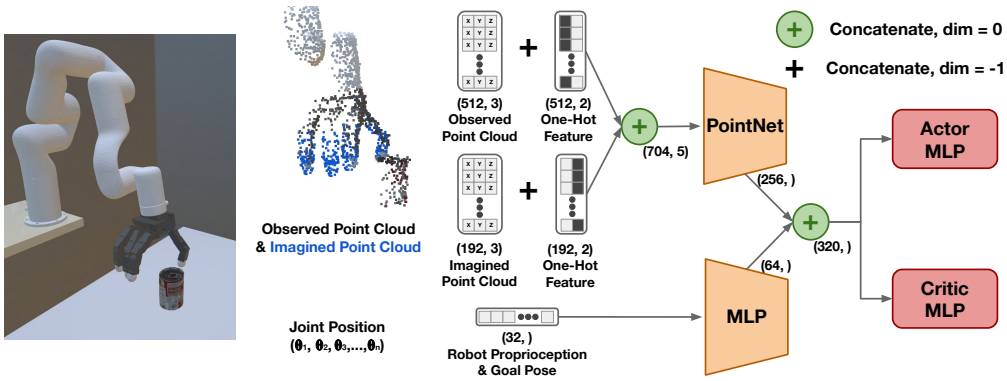

Figure 3: **Architecture:** our feature extractor takes the observed point cloud, imagined point cloud, robot proprioception, and goal pose as input to output a feature embedding. Both actor and critic take the same feature to predict action and value. The blue point represented the imaged point cloud. Note that our network does not require RGB information.

Here, $\mathbf{x}_{\text{finger}}$ and $\mathbf{x}_{\text{obj}}$ are the Cartesian position of each fingertip and the target object. Note that $\mathbf{x}_{\text{obj}}$ is available when we perform training in simulation. However, using this reward alone cannot ensure proper contact for grasping. For example, the robot can touch the object with the back of the hand rather than the palm and then get stuck in this local minimum. Therefore, we introduce a novel contact reward to guarantee meaningful contact behavior:

$$r_{\text{contact}} = \textbf{IsContact}(\text{thumb}, \text{object}) \textbf{ AND } \left(\sum_{\text{finger}} \textbf{IsContact}(\text{finger}, \text{object}) \geq 2 \right). \quad (2)$$

This contact reward function outputs a boolean value in $\{0, 1\}$. It outputs 1 only if the thumb is in contact with the object and there are more than one finger in contact with the object. Intuitively, it encourages the robot to cage the object within fingers. In this case, the robot can quickly find out stable grasping and lift the object to the target location. The lifting behavior is encouraged by

$$r_{\text{lift}} = r_{\text{contact}} \textbf{Lift}(\mathbf{x}_{\text{obj}}, \mathbf{x}_{\text{target}}). \quad (3)$$

The **Lift** function is basically in the form of Equation 1 and the main difference is that it will return a large reward value upon task completion. The overall reward function is a weighted combination of the terms above plus a control penalty:

$$\mathcal{R} = w_{\text{reach}} r_{\text{reach}} + w_{\text{contact}} r_{\text{contact}} + w_{\text{lift}} r_{\text{lift}} + w_{\text{penalty}} r_{\text{penalty}}. \quad (4)$$

We refer readers to the appendix for the detailed calculation of each term and the hyper-parameters.

3.3 Imagined Hand Point Cloud

The usage of point cloud comes with two challenges. The first challenge is occlusion, which may occur to both the object under manipulation and the hand itself. When the robot hand is interacting with an object, the fingers may be occluded by the object. Since we do not assume tactile sensors in this work, this occlusion problem can be serious. The second challenge is the low point cloud resolution during RL training, where we can only use a limited number of points due to the memory limit. In this case, the number of points from hand finger may not be adequate to precisely capture the spatial relationship between the robot and the object. We propose a simple yet effective method to handle both issues in a unified manner. Our idea is to use an imagined hand point cloud in the observation to help the robot to "see" the interaction.

We provide one example in Figure 3. Black points indicate the point cloud captured by the camera, in which some important details of fingers are missing. These missing details provide crucial information of the interaction. Though such interaction information can also be inferred by combining

(a) (b) (c) (d)

Figure 4: **Training Curves.** The left two plots show the single-object and multi-object training curve of (a) bottle category and (b) can category. The right two plots show the ablation results on the (c) bottle category and (d) can category. The x-axis is the training episode and y-axis is the success rate of the corresponding task. The shaded area indicates standard error and the performance is evaluated on five random seeds.

Settings		Bottle		Can	
		Known Obj.	Novel Obj.	Known Obj.	Novel Obj.
RL	Single Obj. Training	0.81 ± 0.09	0.60 ± 0.06	0.96 ± 0.04	0.63 ± 0.18
	Multi Obj. Training	$\mathbf{0.83 \pm 0.16}$	$\mathbf{0.81 \pm 0.15}$	$\mathbf{0.93 \pm 0.07}$	$\mathbf{0.68 \pm 0.09}$
BC	Single Obj. Training	$\mathbf{0.85 \pm 0.06}$	0.44 ± 0.16	$\mathbf{0.89 \pm 0.10}$	0.37 ± 0.13
	Multi Obj. Training	0.78 ± 0.12	$\mathbf{0.74 \pm 0.15}$	0.81 ± 0.14	$\mathbf{0.71 \pm 0.09}$

Table 1: **Comparison for training with single and multiple objects.** We evaluate the policy trained with single and multiple objects on bottle (upper two rows) and can (bottom two rows) categories. We train policy via two different approaches: (1) training the point cloud RL directly, and (2) distilling an oracle policy trained on oracle-states to point cloud RL policy. We test them on both known or novel objects. The results are reported on 5 seeds.

the information from both the proprioception and visual input, we find that the best way is to synthesize these missing details. Concretely, we can compute the pose for each finger link via forward kinematics given the joint position from the robot joint encoder and the robot kinematics model. Then, we synthesize the imagined point cloud (blue points in Figure 3) by sampling the points from the mesh of each finger link. This process is possible in both simulation and the real world.

3.4 Implementation

We adopt Proximal Policy Optimization(PPO) [54] to train the agent in simulation, and the policy is then deployed to real. The network architecture is illustrated in Figure 3. Both value and policy networks share the same visual feature extraction backbone. We concatenate the observed point cloud with the imagined hand point cloud together as the input to the feature extractor. We also attach a one-hot encoding to each point which indicates whether it is observed or an imagined point. The detailed implementation can be found in appendix.

4 Experiments

4.1 Experimental Setup

Point Cloud Preprocessing: To enable smooth transfer from simulation to real-world, we apply the same data preprocessing procedure to the point cloud captured by the camera. It involves four steps: (i) Crop the point cloud to the work region with a manually-defined bounding-box; (ii) Downsample the point cloud uniformly to 512 points; (iii) Add a distance-dependent Gaussian noise to the simulated point cloud to improve the sim2real robustness; (iv) Transform point cloud from the camera frame to the robot base frame using camera pose. In simulation, we use the ground-truth camera pose with multiplicative noise for frame transformation. In the real-world, we perform hand-eye calibration to get the relative transformation from camera to robot base.

Evaluation Criterion: We evaluate the performance of a policy by its success rate. In simulation, a task is considered a success if $d_{ot} < 0.05m$, where d is the distance between object position and

Settings	Bottle		Can	
	Known Obj.	Novel Obj.	Known Obj.	Novel Obj.
w/o Imagined PC.	0.60 ± 0.46	0.56 ± 0.51	0.91 ± 0.17	0.63 ± 0.07
w/o Contact Rew.	0.00 ± 0.00	0.00 ± 0.00	0.06 ± 0.08	0.03 ± 0.06
w/o Both	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Ours	$\mathbf{0.83 \pm 0.16}$	$\mathbf{0.81 \pm 0.15}$	$\mathbf{0.93 \pm 0.07}$	$\mathbf{0.68 \pm 0.09}$

Table 2: **Comparison for different settings on bottle and can categories:** we study the influence of contact-based reward design and imaged point cloud. We evaluate the trained policy on both known and novel objects. The left figure is the learning curve and the right table is the success rate. *PC.* means point cloud and *Rew.* means reward.

goal position. In real world, measuring this distance is very inconvenient. We find that moving to the target position is not hard as long as the object is lifted up. Therefore, we consider the task as successful if the object is grasped and lifted up by $0.2m$ above the table. For each experiment, we train policies with **5 different random seeds** and report the mean and standard deviation.

Training: We train RL in two settings: (i) training on a single object (ii) training on multiple objects jointly. For the single object setting, we perform experiments on both "tomato soup can" and "mustard bottle" from YCB. For the multi-object training, we choose 10 objects from the can or bottle categories of ShapetNet. We randomly choose one object to train for each episode in the multi-object training. Here, we use can and bottle as experimental subjects since they represent two different basic grasping patterns [55] for anthropomorphic hand: precision grasp and power grasp. More training details can be found in appendix.

4.2 Comparison of Single-object and Multi-object Training

We plot the training curve of RL in Figure 4. In general, our method can learn to grasp and move the object to the goal pose within 600 episodes, where each episode contains 20K environment steps. Then we evaluate the policy trained on both known and novel objects, and the results are shown in Table 1. We run 100 trials to compute the average success rate. The policies trained on multiple objects can outperform the policies trained on a single object by a large margin. In particular, policies trained with multiple objects can generalize much better to novel objects compared with those trained on a single object.

Besides the end-to-end RL training, we also experiment an alternative approach: training RL on oracle states, and distill the orcale-state-based policy into the point cloud policy. The oracle state is composed of robot proprioception, target pose, object pose, the distance vector between the object and each fingertips, and a finger contact vector indicating whether which finger is in contact with the object. Note that the contact vector can act like a haptic sensor, which enables the robot to manipulate the object even without vision.

We perform single-object and multi-object training for this alternative approach. The results are shown in Table 1. On grasping known objects, we find that agent trained on single-object outperforms the agent trained on multi-object by a small margin. However, agent trained on multi-objects does much better at grasping novel objects. Our results suggest that using multiple object during training is important, and is of great importance for novel object generalization.

4.3 Ablation Results in Simulation

We ablate two key innovations of the work: the reward design with oracle contact and the imaged hand point cloud. We perform experiments on four different variants: (i) without imagined point cloud; (ii) without contact-based reward design; (iii) without both imagined point cloud and contact based reward design; (iv) our standard approach with both techniques. Note that the variant (iii) is an approximation of [3] in our environments and tasks. We compare both the learning curve and the

evaluation success rate of these four variants. Figure 4 (c) and (d) show the results on bottle and can categories and Table 2 shows the success rate. Our findings can be summarized as follows.

First, we find that contact reward information is of vital importance for training the point cloud RL policy on the multi-finger robot hand. Without using contact reward, the agent can hardly learn anything (red and green curve in the figure) and get nearly zero success rate during evaluation for both bottle and can categories. By encouraging the contact between fingers and object, the RL agent can avoid getting stuck in local minimums and learn meaningful manipulation behavior.

Second, the imagined point cloud can also improve the training and test performance for both categories, though it is not so important as the contact reward. As is shown in the learning curves of Figure 4 (c) and (d), the policy utilizing the imagined point cloud as input can learn faster in the early stage of the training and show much smaller variances. An interesting fact is that imagined point cloud is more beneficial for the bottle category than the can category. One possible reason is that grasping a bottle requires multi-finger coordination to perform a power grasp. Such coordination is highly dependent on the detailed finger information provided by the imagined point cloud. The imagined point cloud can help the agent to better see the fingers even if they are occluded by the object, e.g., fingers behind the object.

4.4 Real-World Evaluation

Test Object	Single Obj. Training	Multi Obj. Training
Mustard Bottle	$\mathbf{0.90 \pm 0.14}$	0.88 ± 0.14
Bleach Cleanser	0.90 ± 0.08	$\mathbf{0.93 \pm 0.04}$
Wide Bottle	$\mathbf{0.93 \pm 0.04}$	0.70 ± 0.14
Thin Bottle	0.77 ± 0.12	$\mathbf{0.83 \pm 0.09}$

Table 3: **Real-World Experiments**: Here we choose four real-world objects to evaluate with our trained policy. The left table shows the success rate on different objects and the right figure shows the objects we use for evaluation. We choose the best three seeds for each experiment based on the performance in the simulator and compute the mean and variance.

We perform sim2real experiments to evaluate the performance of our method in the real world. As is shown in Figure 3.1, we attached an Allegro hand onto an XArm-6 robot arm to grasp the object on the front table. We apply the same data-preprocessing steps for both simulated environment and real-world as mentioned in Sec.4.1. The task execution sequence is visualized on the bottom row of Figure 1. We run 10 independent trials seeds for each object-policy pair.

The real-world evaluation results are shown in Table 3. The policy trained in simulator can directly transfer to the real-world without fine-tuning. Our policy can even grasp objects that has never been seen during the training. We find that for grasping a mustard bottle, training on the same object can ensure better performance while for novel objects like bleach cleanser and water bottle II, training with more diverse objects can improve the performance.

5 Conclusion and Limitation

Limitations In our experiments, we only train and test our method on grasping task, which limit the scope of the proposed method. More dexterous manipulation tasks will be our future research direction. Another potential improvement of our method is to use Recurrent Neural Network and temporal information for policy network. It will enable use to do long-horizon tasks.

Conclusion To the best of our knowledge, our approach is the first work to train a dexterous manipulation policy with point cloud inputs that can transfer to the real world. We justified that direct sim-to-real-real transfer is possible for dexterous grasping with point cloud representation. We are committed to **releasing our code and environment**.

References

[1] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.

[2] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar. Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3651–3657. IEEE, 2019.

[3] T. Chen, J. Xu, and P. Agrawal. A system for general in-hand object re-orientation. In *Conference on Robot Learning*, pages 297–307. PMLR, 2022.

[4] Y.-H. Wu, J. Wang, and X. Wang. Learning generalizable dexterous manipulation from human grasp affordance. *arXiv preprint arXiv:2204.02320*, 2022.

[5] W. Huang, I. Mordatch, P. Abbeel, and D. Pathak. Generalization in dexterous manipulation via geometry-aware multi-task learning. *arXiv preprint arXiv:2111.03062*, 2021.

[6] T. Mu, Z. Ling, F. Xiang, D. Yang, X. Li, S. Tao, Z. Huang, Z. Jia, and H. Su. Maniskill: Generalizable manipulation skill benchmark with large-scale demonstrations. *arXiv preprint arXiv:2107.14483*, 2021.

[7] X. Zhang, R. Chen, F. Xiang, Y. Qin, J. Gu, Z. Ling, M. Liu, P. Zeng, S. Han, Z. Huang, et al. Close the visual domain gap by physics-grounded active stereovision depth sensor simulation. *arXiv preprint arXiv:2201.11924*, 2022.

[8] J. K. Salisbury and J. J. Craig. Articulated hands: Force control and kinematic issues. *The International journal of Robotics research*, 1(1):4–17, 1982.

[9] L. Han and J. C. Trinkle. Dexterous manipulation by rolling and finger gaiting. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 730–735. IEEE, 1998.

[10] D. Rus. In-hand dexterous manipulation of piecewise-smooth 3-d objects. *The International Journal of Robotics Research*, 1999.

[11] D. Prattichizzo and J. C. Trinkle. Grasping. In *Springer handbook of robotics*, pages 955–988. Springer, 2016.

[12] A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 348–353. IEEE, 2000.

[13] J. Bohg, A. Morales, T. Asfour, and D. Kragic. Data-driven grasp synthesis—a survey. *IEEE Transactions on robotics*, 30(2):289–309, 2013.

[14] M. R. Dogar and S. S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. 2010.

[15] O. Kroemer, S. Niekum, and G. D. Konidaris. A review of robot learning for manipulation: Challenges, representations, and algorithms. *Journal of machine learning research*, 22(30), 2021.

[16] V. Kumar, E. Todorov, and S. 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.

[17] S. Levine, N. Wagener, and P. Abbeel. Learning contact-rich manipulation skills with guided policy search. In *IEEE International Conference on Robotics and Automation, ICRA*, pages 156–163. IEEE, 2015.

[18] A. Gupta, C. Eppner, S. Levine, and P. Abbeel. Learning dexterous manipulation for a soft robotic hand from human demonstrations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3786–3793. IEEE, 2016. doi:10.1109/IROS.2016.7759557. URL https://doi.org/10.1109/IROS.2016.7759557.

[19] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5628–5635, 2018. doi:10.1109/ICRA.2018.8461249.

[20] I. Radosavovic, X. Wang, L. Pinto, and J. Malik. State-only imitation learning for dexterous manipulation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7865–7871. IEEE, 2020.

[21] E. Johns. Coarse-to-fine imitation learning: Robot manipulation from a single demonstration. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4613–4619. IEEE, 2021.

[22] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *Robotics: Science and Systems*, 2018.

[23] S. Christen, S. Stevšić, and O. Hilliges. Guided deep reinforcement learning of control policies for dexterous human-robot interaction. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2161–2167. IEEE, 2019.

[24] R. Jeong, J. T. Springenberg, J. Kay, D. Zheng, A. Galashov, N. Heess, and F. Nori. Learning dexterous manipulation from suboptimal experts. In J. Kober, F. Ramos, and C. J. Tomlin, editors, *Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 915–934. PMLR, 2020.

[25] Y. Qin, Y.-H. Wu, S. Liu, H. Jiang, R. Yang, Y. Fu, and X. Wang. Dexmv: Imitation learning for dexterous manipulation from human videos. *arXiv preprint arXiv:2108.05877*, 2021.

[26] F. Pomerleau, F. Colas, and R. Siegwart. A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends in Robotics*, 4(1):1–104, 2015.

[27] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. *Robotics: Science and System*, 2017.

[28] A. ten Pas, M. Gualtieri, K. Saenko, and R. Platt. Grasp pose detection in point clouds. *The International Journal of Robotics Research*, 36(13-14):1455–1473, 2017.

[29] H. Liang, X. Ma, S. Li, M. Görner, S. Tang, B. Fang, F. Sun, and J. Zhang. Pointnetgpd: Detecting grasp configurations from point sets. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3629–3635, 2019. doi:10.1109/ICRA.2019.8794435.

[30] A. Mousavian, C. Eppner, and D. 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.

[31] Y. Qin, R. Chen, H. Zhu, M. Song, J. Xu, and H. Su. S4g: Amodal single-view single-shot se(3) grasp detection in cluttered scenes. In *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 53–65. PMLR, 2020.

[32] P. Ni, W. Zhang, X. Zhu, and Q. Cao. Pointnet++ grasping: Learning an end-to-end spatial grasp generation algorithm from sparse point clouds. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3619–3625, 2020. doi:10.1109/ICRA40945.2020.9196740.

[33] C. Wu, J. Chen, Q. Cao, J. Zhang, Y. Tai, L. Sun, and K. Jia. Grasp proposal networks: An end-to-end solution for visual learning of robotic grasps. *Advances in Neural Information Processing Systems*, 33:13174–13184, 2020.

[34] W. Wei, Y. Luo, F. Li, G. Xu, J. Zhong, W. Li, and P. Wang. Gpr: Grasp pose refinement network for cluttered scenes. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4295–4302. IEEE, 2021.

[35] W. Wei, D. Li, P. Wang, Y. Li, W. Li, Y. Luo, and J. Zhong. Dvgg: Deep variational grasp generation for dextrous manipulation. *IEEE Robotics and Automation Letters*, 2022.

[36] S. Andrews and P. G. Kry. Goal directed multi-finger manipulation: Control policies and analysis. *Computers & Graphics*, 2013.

[37] J. Varley, J. Weisz, J. Weiss, and P. Allen. Generating multi-fingered robotic grasps via deep learning. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4415–4420. IEEE, 2015.

[38] S. Brahmbhatt, A. Handa, J. Hays, and D. Fox. Contactgrasp: Functional multi-finger grasp synthesis from contact. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2386–2393. IEEE, 2019.

[39] Q. Lu, K. Chenna, B. Sundaralingam, and T. Hermans. Planning multi-fingered grasps as probabilistic inference in a learned deep network. In *Robotics Research*, pages 455–472. Springer, 2020.

[40] Y. Chebotar, O. Kroemer, and J. Peters. Learning robot tactile sensing for object manipulation. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3368–3375, 2014. doi:10.1109/IROS.2014.6943031.

[41] W. Yuan, S. Dong, and E. H. Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12):2762, 2017.

[42] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine. More than a feeling: Learning to grasp and regrasp using vision and touch. *IEEE Robotics and Automation Letters*, 3(4):3300–3307, 2018.

[43] A. Murali, Y. Li, D. Gandhi, and A. Gupta. Learning to grasp without seeing. In *International Symposium on Experimental Robotics*, pages 375–386. Springer, 2018.

[44] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, 2020.

[45] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg. Making sense of vision and touch: Learning multimodal representations for contact-rich tasks. *IEEE Transactions on Robotics*, 36(3):582–596, 2020.

[46] R. Bhirangi, T. Hellebrekers, C. Majidi, and A. Gupta. Reskin: versatile, replaceable, lasting tactile skins. *arXiv preprint arXiv:2111.00071*, 2021.

[47] H. Van Hoof, T. Hermans, G. Neumann, and J. 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.

[48] N. Vulin, S. Christen, S. Stevšić, and O. Hilliges. Improved learning of robot manipulation tasks via tactile intrinsic motivation. *IEEE Robotics and Automation Letters*, 6(2):2194–2201, 2021.

[49] S. H. Huang, M. Zambelli, J. Kay, M. F. Martins, Y. Tassa, P. M. Pilarski, and R. Hadsell. Learning gentle object manipulation with curiosity-driven deep reinforcement learning. *arXiv preprint arXiv:1903.08542*, 2019.

[50] H. Xu, Y. Luo, S. Wang, T. Darrell, and R. Calandra. Towards learning to play piano with dexterous hands and touch. *arXiv preprint arXiv:2106.02040*, 2021.

[51] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su. SAPIEN: A simulated part-based interactive environment. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[52] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.

[53] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *2015 international conference on advanced robotics (ICAR)*, pages 510–517. IEEE, 2015.

[54] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[55] J. R. Napier. The prehensile movements of the human hand. *The Journal of bone and joint surgery. British volume*, 38(4):902–913, 1956.