Efficient Sim-to-real Transfer of Contact-Rich Manipulation Skills with Online Admittance Residual Learning

Anonymous Author(s)
Affiliation
Address
email

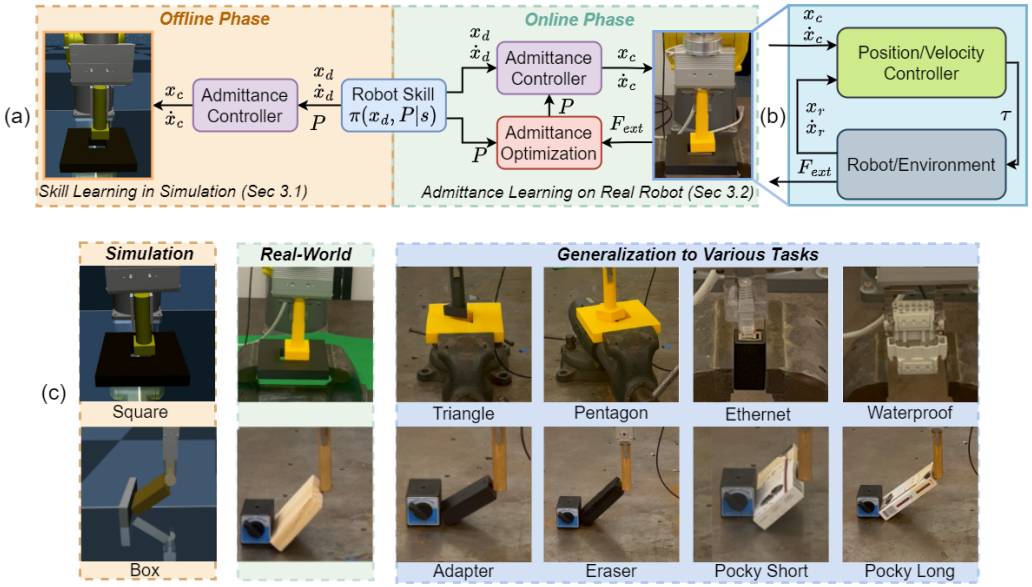

Figure 1: As shown in (a), we propose a robust contact-rich manipulation skill learning framework that offline learns the robot motion and compliance control parameters in the simulation and online adapts to the real world. The structure of the admittance controller is depicted in (b). Our framework demonstrates robustness in sim-to-real transfer and generalizability to diverse real-world tasks in (c)

Abstract: Learning contact-rich manipulation skills is essential to robotic applications. Such skills require the robots to interact with the environment with feasible manipulation trajectories and suitable compliance control parameters to enable safe and stable contact. However, learning these skills is challenging due to data inefficiency in the real world and the sim-to-real gap in simulation. In this paper, we introduce a hybrid offline-online framework to learn robust manipulation skills. We employ model-free reinforcement learning for the offline phase to obtain the robot motion and compliance control parameters in simulation. Subsequently, in the online phase, we learn the residual of the compliance control parameters to maximize robot performance-related criteria with force sensor measurements in real-time. To demonstrate the effectiveness and robustness of our approach, we provide comparative results against existing methods for assembly and pivoting tasks. Videos are available at https://sites.google.com/view/admitlearn.

Keywords: Contact-rich Manipulation, Admittance Control

Submitted to the 7th Conference on Robot Learning (CoRL 2023). Do not distribute.

1 Introduction

Contact-rich manipulation is common in a wide range of robotic applications, including assembly [1, 2, 3, 4, 5], manufacturing, object pivoting [6, 7, 8], and pushing [9, 10]. To accomplish these tasks, robots need to learn both the manipulation trajectory and the force control parameters. The manipulation trajectory guides the robot toward completing the task while physically engaging with the environment, whereas the force control parameters regulate the contact force. Incorrect control parameters can lead to oscillations and excessive contact forces that may damage the robot or the environment.

Past works have tackled the contact-rich skill-learning problem in different ways. First, the majority of previous works [6, 2, 7, 3, 11] focus on learning the manipulation trajectories and rely on human experts to manually tune force control parameters. While this simplification has demonstrated remarkable performance in many applications, letting human labor tune control parameters is still inconvenient. Furthermore, the tuned parameter for one task may not generalize well to other task settings with different kinematic or dynamic properties. For example, assembly tasks with different clearances will require different control parameters. Another line of work deals with this problem by jointly learning the robot's motion and force control parameters [12, 13, 14, 15, 16, 17, 4, 18]. Such learning processes can be conducted in both real-world and simulation. However, learning such skills on real robots is time-consuming and may damage the robot or environment. Learning in simulation is efficient and safe, but the learned control parameters are difficult to transfer to real robots due to the sim-to-real gap, and directly deploying the learned control parameters may cause damage to the robot.

In this paper, we focus on transferring the learned robot skill in simulation to the real world. We notice that the manipulation trajectory is more related to the kinematic properties, such as size and shape, which have a smaller sim-to-real gap and can be transferred directly, as demonstrated by previous works [6, 2, 7, 3]. However, simulating the contact dynamics proves to be challenging, primarily due to its sensitivity to various parameters, including surface stiffness and friction [19]. This sensitivity will result in a large sim-to-real gap and affects the learned compliance control parameters. Inspired by the above analysis, we propose a framework to learn robot manipulation skills that can transfer to the real world. As depicted in Fig. 1(a), the framework contains two phases: skill learning in simulation and admittance adaptation on the real robot. We use model-free RL to learn the robot's motion with domain randomization to enhance the robustness for direct transfer. The compliance control parameters are learned at the same time and serve as an initialization to online admittance learning. During online execution, we iteratively learn the residual of the admittance control parameters by optimizing the future robot trajectory smoothness and task completion criterion. We conduct real-world experiments on two typical contact-rich manipulation tasks: assembly and pivoting. Our proposed framework achieves efficient transfer from simulation to the real world. Furthermore, it shows excellent generalization ability in tasks with different kinematic or dynamic properties, as shown in Fig. 1(c). Comparison and ablation studies are provided to demonstrate the effectiveness of the framework.

2 Related Works

2.1 Sim-to-real Transfer in Robot Contact-Rich Manipulation

Contact-rich manipulation tasks involve the interaction between robots and the environment through physical contact. In recent years, there has been a growing trend in utilizing simulation environments such as MuJoCo [20], Bullet [21], and IsaacGym [11] to learn and train robots for these tasks. These simulation environments offer advantages in terms of safety, scalability, and cost-effectiveness. Nevertheless, the sim-to-real gap remains a significant challenge.To address the gap, various approaches have been explored, including system identification, transfer learning, and domain randomization. System identification approaches [22, 23] involves the calibration of simulation parameters to improve accuracy and align the simulation with real-world dynamics. Transfer

learning methods [24, 17] aim to fine-tune skills learned in simulation for application in real-world scenarios. Domain randomization techniques [7, 6, 17, 25] are employed to create diverse environments with varying properties, enabling the learning of robust skills for better generalization.

2.2 Learning Variable Impedance/Admittance Control

Compliance control [26], such as impedance and admittance control, enables robots to behave as a mass-spring-damping system. Tuning the compliance control parameters is crucial for stabilizing the robot and accomplishing manipulation tasks. However, manual tuning can be time-consuming. To address this issue, learning-based approaches have been applied to automatically learn the control parameters. Previous methods have focused on learning compliance control parameters either from expert demonstrations [12, 13, 14] or through reinforcement learning (RL) [15, 16, 17, 4, 18] to acquire gain-changing policies. [13, 4, 12, 15, 17] propose to directly collect data in the real world. However, it is time-consuming to collect the data. Authors in [14, 18] have demonstrated success on direct transferring the learned control parameters from the simulation to the real world. Nevertheless, their applications are limited to simple tasks, such as waypoint tracking and whiteboard wiping.

3 Proposed Approach

We focus on learning robust contact-rich manipulation skills that can achieve efficient sim-to-real transfer. We define the skill as $\pi(x_d, P|s)$, which generates both the robot desired trajectory x_d and the compliance control parameters P given the current state s.

We use *Cartesian space admittance control* as the compliance controller. As shown in Fig. 1(b), the admittance control takes in the desired trajectory $[x_d, \dot{x}_d, \ddot{x}_d] \in \mathbb{R}^{18}$, and the external force/torque $F_{ext} \in \mathbb{R}^6$ measured on the robot end-effector and outputs the compliance trajectory $[x_c, \dot{x}_c, \ddot{x}_c] \in \mathbb{R}^{18}$ to the position/velocity controller according to the mass-spring-damping dynamics [26]:

$$M(\ddot{x}_c - \ddot{x}_d) + D(\dot{x}_c - \dot{x}_d) + K(x_c - x_d) = F_{ext} \tag{1}$$

where M, K, D are the robot inertia, stiffness, and damping matrices, respectively. We assume M, K, D is diagonal for simplicity and $P = \{M, K, D\}$ as the collection of all control parameters.

To achieve this goal, We propose an offline-online framework for learning contact-rich manipulation skills as depicted in Fig. 1(a). In the offline phase, we employ the model-free RL with domain randomization to learn the robot motion and the initial guess of compliance control parameters from the simulation (Section 3.1). In the online phase, we execute the offline-learned motions on the real robot and learn the residual compliance control parameters by optimizing the future robot trajectory smoothness and task completion criteria(Section 3.2).

3.1 Learning offline contact-rich manipulation skills

We utilize model-free RL to learn contact-rich manipulation skills in MuJoCo simulation [20]. The problem is modeled as a $MDP := \{S, A, R, T, \gamma\}$, where S is the state space, A is the action space, R is a reward function, T denotes the state-transition function, and γ is the discount factor. Our goal is to learn a policy $\pi(a|s)$ that maximizes the expected future return $\mathbb{E}\left[\sum_t \gamma^t r_t\right]$. The details of RL training tasks are summarized below:

Assembly Task: The goal is to align the peg with the hole and then insert it.

State space: The state space $s \in \mathbb{R}^{18}$ contains peg pose $s_p \in \mathbb{R}^6$ (position and Euler angles) relative to the hole, peg velocity $v_p \in \mathbb{R}^6$, and the external force measured on the robot wrist $F_{ext} \in \mathbb{R}^6$.

Action space: The action $a \in \mathbb{R}^{12}$ consists of the end-effector velocity command $v_d \in \mathbb{R}^6$ and the diagonal elements of the stiffness matrix $k \in \mathbb{R}^6$. To simplify the training, the robot inertia M is fixed to $diag(1, 1, 1, 0.1, 0.1, 0.1)$, and the damping matrix $D = diag(d_1, \cdots, d_6)$ is computed according to the critical damping condition $d_i = 2\sqrt{m_i k_i}, i = \{1, 2, 3, 4, 5, 6\}$.

Reward function: The reward function is defined as $r(s) = 10^{(1-\|s_{pos}-s^d_{pos}\|_2)}$, where $s_{pos} \in \mathbb{R}^3$ is the peg position and $s^d_{pos} \in \mathbb{R}^3$ is the nominal hole location. The exponential function encourages successful insertion by providing a high reward.

Pivoting Task: The goal is to gradually push the object to a stand-up pose against the wall.

State space: The state $s \in \mathbb{R}^{12}$ consists of the robot pose $s_p \in \mathbb{R}^6$ and the external force $F_{ext} \in \mathbb{R}^6$.

Action space: For simplicity, we consider a 2d pivoting problem: the robot can only move in the X, Z direction. The robot action $a \in R^4$ contains the velocity command in X, Z direction and the corresponding stiffness parameter.

Reward function: We use the rotational distance between the goal orientation R^{goal} and current object orientation R as the cost and define the reward function as $r = \frac{\pi}{2} - d$, with $d = \arccos\big(0.5(\mathrm{Tr}(R^{goal}R^T) - 1)\big)$, which computes the distance of two rotation matrices between R and R^{goal}. The constant term $\frac{\pi}{2}$ simply shifts the initial reward to 0. This reward encourages the robot to push the object to the stand-up orientation.

We use domain randomization on the robot's initial pose and contact force to improve the robustness of the learned skills. The implementation details can be found in Appendix. B.1.

3.2 Online Optimization-Based Admittance Learning

We have learned a policy that can perform contact-rich manipulation tasks in simulation. However, the sim-to-real gap may prevent us from directly transferring the learned skills to the real world. Our goal is to adapt the offline learned skills, especially the admittance control parameters, with online data. Instead of retraining skills with real-world data, we propose locally updating the control parameters using the latest contact force measurements during online execution. We formulate online learning as an optimization problem that optimizes the residual control parameters to achieve smooth trajectory and task completion criteria while respecting the interaction dynamics between the robot and the environment. We will describe the optimization constraints, objective function, and overall online learning algorithm in Section 3.2.1, 3.2.2, and 3.2.3, respectively.

3.2.1 Optimization Constraints

Robot dynamics constraint: Admittance control enables robot to behave as a mass-spring-damping system as shown in Eq. 1. We consider the robot state $x = [e, \dot{e}]$, where $e = x_c - x_d$, and we can obtain the robot dynamics constraint in the state space form:

$$\dot{x} = \begin{bmatrix} \dot{e} \\ \ddot{e} \end{bmatrix} = f(x, F_{ext}, u) = \begin{bmatrix} \dot{e} \\ -M^{-1}D\dot{e} - M^{-1}Ke + M^{-1}F_{ext} \end{bmatrix} \tag{2}$$

where the optimization variable $u = [m_1^{-1}, \ldots, m_6^{-1}, k'_1, \ldots, k'_6, d'_1, \ldots, d'_6]^T$ is the diagonal elements of M^{-1}, $K' = M^{-1}K$, $D' = M^{-1}D$. $e, \dot{e}$ are the robot states that can be directly accessed, and F_{ext} is the external force that should be modeled from the environment dynamics.

Contact force estimation: Modeling the contact force explicitly is difficult because the contact point and mode can change dramatically during manipulation. Therefore, we propose to estimate the contact force online using the force/torque sensor measurements. In our experiments, we utilize a simple but effective *record & replay* strategy, where we record a sequence of force information $\{F^0_{ext}, \ldots, F^T_{ext}\}$ within a time window $[0, T]$ and replay them during the optimization.

There are other approaches for force estimation, such as using analytical contact models [27, 28] or numerically learning the contact force by model fitting. However, we found the *record & replay* strategy is better by experiments. We provide analyses in the Appendix. C.2.

Stability constraint: To ensure stability for admittance control, we need the admittance parameters to be positive-definite. Therefore, we constrain the optimization variable u to be positive.

3.2.2 Objective Function Design

We want to optimize the admittance parameters to establish stable contact and successfully achieve the task. Previous work [29] introduces the FITAVE objective $\int_0^T t|\dot{e}(t)|dt$ to effectively generate smooth and stable contact by regulating the robot's future velocity error. In addition, the ITAE objective $\int_0^{+\infty} t|e(t)|dt$ in [30] minimizes the position error to ensure the robot tracking the desired trajectory and finishes the task. We combine those two functions as our objective:

$$C(x) = \int_0^T t[w|e(t)| + (1-w)|\dot{e}(t)|]dt \tag{3}$$

where $w \in \mathbb{R}$ is a weight scalar to balance the trajectory smoothness and task completion criterion.

3.2.3 Online admittance learning

The optimization formulation is shown in Eq. 4. We optimize the residual admittance parameters, with u_{init} obtained from the offline learned skill.

$$\begin{aligned} \min_{\delta u} \quad & C(x) \\ \text{s.t.} \quad & \dot{x} = f(x, F_{ext}, u_{init} + \delta u) \\ & F_{ext} \leftarrow \textit{record \& replay} \\ & u_{init} + \delta u > 0 \end{aligned} \tag{4}$$

We illustrate the online admittance learning procedures in Alg. 1. In the online phase, we execute the skill learned offline on the real robot and recorded the contact force at each time step. Every T seconds, the online optimization uses the recorded force measurements, the current robot state and the admittance parameters learned offline to update the admittance parameter residual. The process runs in a closed-loop manner to complete the desired task robustly.

$$u = u_{init} + \delta u^*, M = diag\{m_1, \cdots, m_6\}, K = M \cdot diag\{k'_1, \cdots, k'_6\}, D = M \cdot diag\{d\prime_1, \cdots, d'_6\} \tag{5}$$

Algorithm 1: Online Admittance Residual Learning

Require: u_{init} from the offline policy $\pi(a|s)$, current robot state x

```
while task not terminated do
    if every T seconds then
        δu* ← admittance optimization in (4)
        M, K, D ← Recover admittance parameters from (5)
    end if
    {F_ext} ← record force sensor data
end while
```

4 Experiment Results

We conduct experiments on two contact-rich manipulation tasks, peg-in-hole assembly and pivoting, to evaluate: 1) the robustness of sim-to-real transfer and 2) the generalizability of different task settings. We provide comparison results with two baselines: 1) **Direct Transfer**: directly sim-to-real transfer both the learned robot trajectory and the control parameters [18], 2) **Manual Tune**: transfer learned trajectory with manually tuned control parameters [2]. We consider three metrics for evaluation: 1) **success rate** indicates the robustness of transfer, 2) **completion time** for successful trials denotes the efficiency of the skills, and 3) **max contact force** shows the safety. In addition, the effect of weight selection of the proposed online optimization objective is discussed.

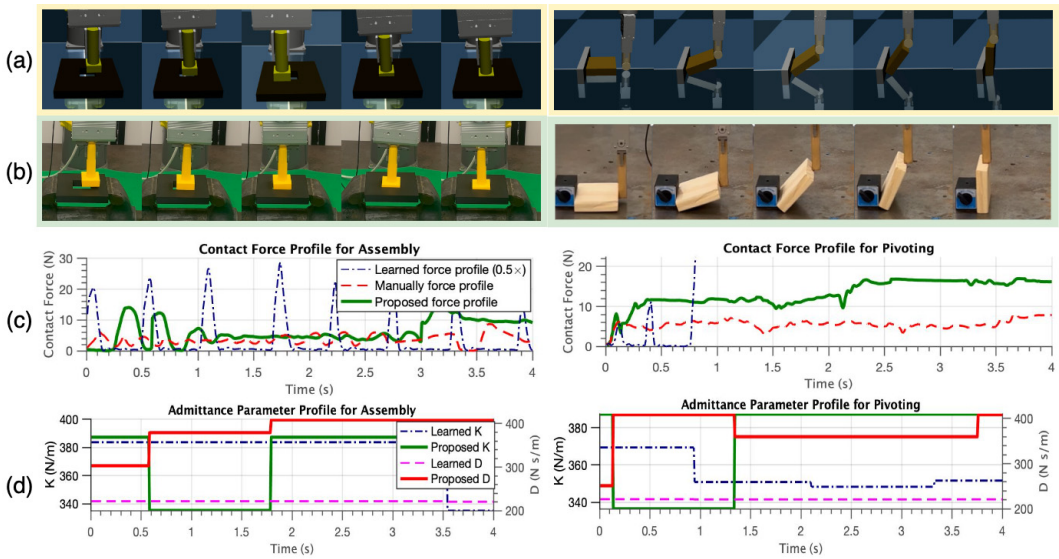

Figure 2: (a) shows the snapshots of the learned policy in simulation.(b) demonstrates the snapshots using the proposed approach for sim-to-real transfer. (c)(d) illustrate the forces and control parameters profiles for both the learned and proposed approach in the real world. The proposed approach can real-time adjust the parameters to get the best performance.

4.1 Skill Learning in Simulation

We use Soft Actor-Critic [31] to learn manipulation skills in simulation. During the evaluation, the learned assembly and pivoting skills both achieved a 100% success rate. Fig. 2(a) shows the snapshots of the learned assembly skills. The robot learns to search for the exact hole location on the hole surface with a learned variable admittance policy and smoothly inserts the peg into the hole. For the learned pivoting skill, the robot pushes the object against the wall and gradually pivots it to the target pose with suitable frictional force.

4.2 Sim-to-Real Transfer

We evaluate the sim-to-real transfer performance on the same task. In the real world, the task setup, such as the object and robot geometry, is identical to the simulation. We mainly focus on evaluating the effect of the sim-to-real gap on robot/environment dynamics.

We first apply the offline learned skill to the real world. From the experiments, we notice that the *Direct Transfer* baseline fails to produce safe and stable interactions. As depicted in Fig 2(c), for peg-in-hole assembly tasks, the peg bounces on the hole surface and generates large contact forces, making the assembly task almost impossible to complete. Similarly, in the pivoting task, the robot cannot make stable contact with the object and provide enough frictional force for pivoting.

Then we examine whether the learned robot motion is valid with manually tuned control parameters. As shown in Tab. 1, the *Manual Tune* baseline can achieve a 100% success rate for both tasks. This supports our hypothesis and previous works that the manipulation trajectory is directly transferable with suitable control parameters to address the sim-to-real gap.

However, manual tuning requires extensive human labor. We want to evaluate whether the proposed online admittance learning framework can perform similarly without any tuning. Table 1 presents an overview of the sim-to-real transfer results. For all experiments, the weight parameter w of the proposed approach was consistently set to 0.4. Notably, our proposed method achieves a 100% success rate in the assembly task, along with a 90% success rate in the pivoting task. Furthermore, it achieves these results while exhibiting shorter completion times than the other two baselines.

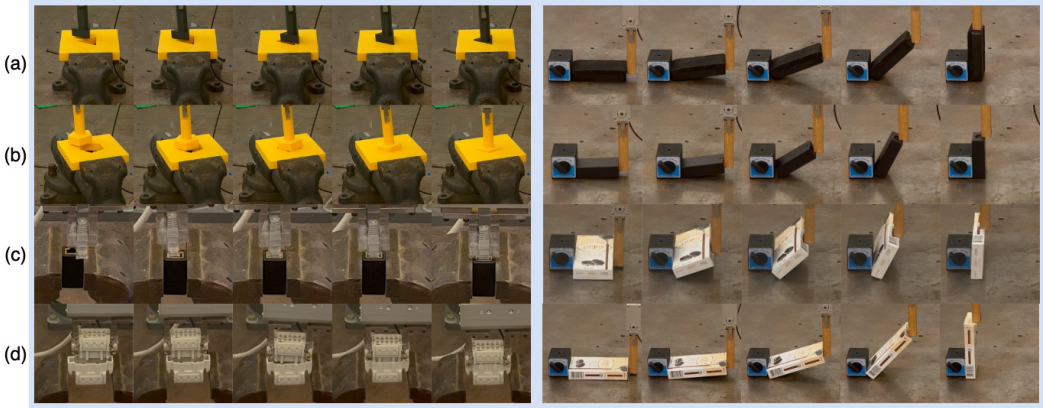

Figure 3: Snapshots of using the proposed approach to generalize to various task settings. The snapshots and videos of the baseline methods are available on our website.

We also investigate the contact force and the adjusted admittance parameters during the manipulation, shown in Fig. 2(c)(d). Initially, the robot establishes contact with the environment using the offline learned parameters, resulting in a large applied force. In the subsequent update cycle, the proposed method effectively adjusts the parameters by decreasing K and increasing D, enabling the robot to interact smoothly with the environment and reduce the contact force. Later, it increases K and decreases D to suitable values to finish the task more efficiently.

	Assembly Task			Pivoting Task		
	Succ. Rate	Time (s)	Max F (N)	Succ. Rate	Time (s)	Max F (N)
Proposed	$10/10$	19.0 ± 11.2	23.6 ± 6.3	$9/10$	25.6 ± 2.1	20.1 ± 4.1
Manual	$10/10$	28.1 ± 8.6	10.3 ± 2.2	$10/10$	25.3 ± 3.6	9.2 ± 0.6
Direct	$3/10$	39.0 ± 12.8	63.7 ± 6.8	$0/10$	N/A	30.7 ± 4.6

Table 1: Success rate evaluation in real-world experiments.

4.3 Generalization to Different Task Settings

The aforementioned experiments highlight the ability of the proposed framework to achieve sim-to-real transfer within the same task setting. In this section, we aim to explore the generalization capabilities of the proposed approach across different task settings, which may involve distinct kinematic and dynamic properties. For two baselines, we directly use the manually tuned or learned control parameters of the training object to new tasks.

	Triangle		Pentagon		Ethernet		Waterproof	
	Succ. Rate	Time (s)	Succ. Rate	Time (s)	Succ. Rate	Time (s)	Succ. Rate	Time (s)
Proposed	$10/10$	15.9 ± 6.2	$10/10$	20.1 ± 8.9	$9/10$	42.1 ± 13.7	$9/10$	37.8 ± 17.7
Manual	$8/10$	$43. \pm 17.0$	$9/10$	38.0 ± 18.0	$1/10$	78.0 ± 0.0	$0/10$	N/A
Direct	$0/10$	N/A	$1/10$	7.0 ± 0.0	$0/10$	N/A	$0/10$	N/A

	Adapter [L=8.8 cm, w=69g]		Eraser [L=12.2 cm, w=36g]		Pocky Short [L=7.9 cm, w=76g]		Pocky Long [L=14.8 cm, w=76g]	
	Succ. Rate	Time (s)	Succ. Rate	Time (s)	Succ. Rate	Time (s)	Succ. Rate	Time (s)
Proposed	$8/10$	25.0 ± 4.8	$9/10$	28.4 ± 2.7	$8/10$	12.9 ± 1.7	$7/10$	31.8 ± 11.0
Manual	$0/10$	N/A	$10/10$	30.0 ± 1.0	$1/10$	19.0 ± 0.0	$1/10$	40.0 ± 0.0
Direct	$0/10$	N/A	$0/10$	N/A	$0/10$	N/A	$0/10$	N/A

Table 2: Generalization performance to different assembly tasks (**Top**) and pivoting tasks (**Below**).

Peg-in-hole assembly: We test various assembly tasks, including polygon-shaped peg-holes such as triangles and pentagons, as well as real-world socket connectors like Ethernet and waterproof connectors. These tasks are visualized in Fig. 1(c). The outcomes of our experiments are outlined in Table 2. Our proposed method achieves 100% success rates on the polygon shapes and a commendable 90% success rate on Ethernet and waterproof connectors. Moreover, the completion time of the proposed method is much shorter than other baselines. The *Manual Tune* baseline also achieves

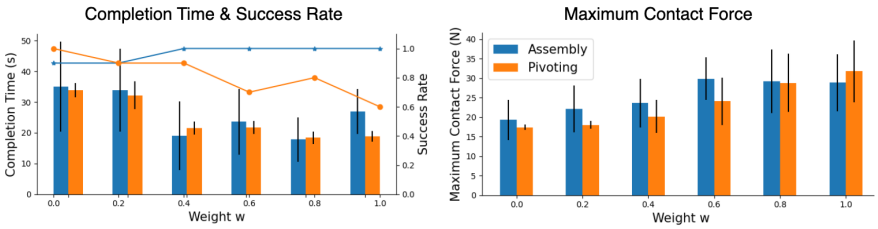

Figure 4: Ablation on the weight parameter. The left figure shows the completion time and success rate with respect to different w, and the right figure shows the contact force.

decent success rates on the polygon shapes as it is similar to the scenario in that we tune the parameters. However, for the socket connectors, due their tighter fit and irregular shapes, substantial force is required for insertion (approximately 15N for Ethernet and 40N for waterproof connectors) and *Manual Tune* baseline cannot accomplish these two tasks.

Pivoting: Similarly, we conduct a series of pivoting experiments on various objects, encompassing diverse geometries and weights as shown in Table 2. Remarkably, our proposed approach exhibits robust generalization capabilities across all tasks, achieving a success rate exceeding 70%. However, when relying solely on manually tuned parameters, the ability to pivot an object is limited to the eraser that has a similar length to the trained object and is the lightest object in the test set. As the object geometry and weight diverge significantly, the manually tuned parameters often fail to establish stable contact with the object and exert sufficient force to initiate successful pivoting.

4.4 Ablation: Objective Weight Selection

In this subsection, we would like to study the effect of weight selection. We evaluated different weight parameters on both the assembly and pivoting tasks. The results are depicted in Figure 4. For the assembly task, all proposed method variations achieve a 100% success rate except for $w <= 0.2$. Smaller weight parameters tend to prioritize trajectory smoothness, which may not provide sufficient contact force for successful insertion. On the other hand, in pivoting tasks, larger weight values led to a decrease in the success rate. It is because larger weight values prioritize task completion, potentially leading to a failure in establishing a stable initial contact for pivoting. These observations align with the objective design motivation. Based on our findings, selecting the parameter 0.4 strikes a good balance between both objectives and yields the best overall performance.

5 Conclusion and Limitations

This paper proposes a contact-rich skill learning framework for sim-to-real transfer. It consists of two main components: skill learning in simulation during the offline phase and admittance learning on the real robot during online execution. These components work together to enable the robot to acquire the necessary skills in simulation and optimize admittance control parameters for safe and stable interactions with the real-world environment. We evaluate the performance of our framework in two contact-rich manipulation tasks: assembly and pivoting. Our approach achieves promising success rates in both tasks and demonstrates great generalizability across various tasks.

However, there are some limitations of our proposed framework: 1) We assume the task settings are similar from training to testing. When there are significant differences in geometry between the simulation and the testing scenario, the current framework may struggle to transfer the learned policy to the real world solely through online admittance learning. 2) We employ the *record & replay* strategy to estimate the contact force during online optimization. This approach may pose challenges when the contact force pattern undergoes significant changes in the subsequent time window. 3) Our method refines the policy during execution, which means that initially, a sub-optimal policy is used to make contact with the environment. As a result, this scheme may not be suitable for contact with fragile objects.

References

[1] T. Inoue, G. De Magistris, A. Munawar, T. Yokoya, and R. Tachibana. Deep reinforcement learning for high precision assembly tasks. In *2017 IEEE/RSJ Int. Conf. on Intelligent Robots and Syst. (IROS)*, pages 819–825. IEEE, 2017.

[2] X. Zhang, S. Jin, C. Wang, X. Zhu, and M. Tomizuka. Learning insertion primitives with discrete-continuous hybrid action space for robotic assembly tasks. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 9881–9887. IEEE, 2022.

[3] N. Vuong, H. Pham, and Q.-C. Pham. Learning sequences of manipulation primitives for robotic assembly. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4086–4092. IEEE, 2021.

[4] J. Luo, E. Solowjow, C. Wen, J. A. Ojea, A. M. Agogino, A. Tamar, and P. Abbeel. Reinforcement learning on variable impedance controller for high-precision robotic assembly. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3080–3087. IEEE, 2019.

[5] S. Jin, X. Zhu, C. Wang, and M. Tomizuka. Contact pose identification for peg-in-hole assembly under uncertainties. In *2021 American Control Conference (ACC)*, pages 48–53. IEEE, 2021.

[6] W. Zhou and D. Held. Learning to grasp the ungraspable with emergent extrinsic dexterity. In *ICRA 2022 Workshop: Reinforcement Learning for Contact-Rich Manipulation*, 2022. URL https://openreview.net/forum?id=Zrp4wpa9lqh.

[7] X. Zhang, S. Jain, B. Huang, M. Tomizuka, and D. Romeres. Learning generalizable pivoting skills. *arXiv preprint arXiv:2305.02554*, 2023.

[8] Y. Shirai, D. K. Jha, A. Raghunathan, and D. Romeres. Chance-constrained optimization in contact-rich systems for robust manipulation. *arXiv preprint arXiv:2203.02616*, 2022.

[9] J. Stüber, C. Zito, and R. Stolkin. Let's push things forward: A survey on robot pushing. *Frontiers in Robotics and AI*, page 8, 2020.

[10] K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. Riedmiller. Learning an embedding space for transferable robot skills. In *International Conference on Learning Representations*, 2018.

[11] Y. Narang, K. Storey, I. Akinola, M. Macklin, P. Reist, L. Wawrzyniak, Y. Guo, A. Moravanszky, G. State, M. Lu, et al. Factory: Fast contact for robotic assembly. *arXiv preprint arXiv:2205.03532*, 2022.

[12] L. Peternel, T. Petrič, and J. Babič. Human-in-the-loop approach for teaching robot assembly tasks using impedance control interface. In *2015 IEEE int. conf. on robotics and automation (ICRA)*, pages 1497–1502. IEEE, 2015.

[13] F. J. Abu-Dakka, L. Rozo, and D. G. Caldwell. Force-based learning of variable impedance skills for robotic manipulation. In *2018 IEEE-RAS 18th Int. Conf. on Humanoid Robots (Humanoids)*, pages 1–9. IEEE, 2018.

[14] X. Zhang, L. Sun, Z. Kuang, and M. Tomizuka. Learning variable impedance control via inverse reinforcement learning for force-related tasks. *IEEE Robotics and Automation Letters*, 6(2):2225–2232, 2021.

[15] J. Buchli, F. Stulp, E. Theodorou, and S. Schaal. Learning variable impedance control. *The Int. J. of Robotics Research*, 30(7):820–833, 2011.

[16] J. Rey, K. Kronander, F. Farshidian, J. Buchli, and A. Billard. Learning motions from demonstrations and rewards with time-invariant dynamical systems based policies. *Autonomous Robots*, 42(1):45–64, 2018.

[17] C. C. Beltran-Hernandez, D. Petit, I. G. Ramirez-Alpizar, and K. Harada. Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach. *Applied Sciences*, 10(19):6923, 2020.

[18] R. Martín-Martín, M. A. Lee, R. Gardner, S. Savarese, J. Bohg, and A. Garg. Variable impedance control in end-effector space: An action space for reinforcement learning in contact-rich tasks. *arXiv preprint arXiv:1906.08880*, 2019.

[19] M. Parmar, M. Halm, and M. Posa. Fundamental challenges in deep learning for stiff contact dynamics. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5181–5188. IEEE, 2021.

[20] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ Int. Conf. on Intelligent Robots and Syst.*, pages 5026–5033. IEEE, 2012.

[21] E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. http://pybullet.org, 2016–2019.

[22] L. Ljung. *System identification*. Springer, 1998.

[23] V. Lim, H. Huang, L. Y. Chen, J. Wang, J. Ichnowski, D. Seita, M. Laskey, and K. Goldberg. Real2sim2real: Self-supervised learning of physical single-step dynamic actions for planar robot casting. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8282–8289, 2022. doi:10.1109/ICRA46639.2022.9811651.

[24] R. Chitnis, S. Tulsiani, S. Gupta, and A. Gupta. Efficient bimanual manipulation using learned task schemas. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1149–1155. IEEE, 2020.

[25] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8973–8979. IEEE, 2019.

[26] C. Ott, R. Mukherjee, and Y. Nakamura. Unified impedance and admittance control. In *2010 IEEE international conference on robotics and automation*, pages 554–561. IEEE, 2010.

[27] N. Doshi, O. Taylor, and A. Rodriguez. Manipulation of unknown objects via contact configuration regulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2693–2699. IEEE, 2022.

[28] J. Zhou, M. T. Mason, R. Paolini, and D. Bagnell. A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation. *The International Journal of Robotics Research*, 37(2-3):249–265, 2018.

[29] C. Wang, X. Zhang, Z. Kuang, and M. Tomizuka. Safe online gain optimization for cartesian space variable impedance control. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 751–757. IEEE, 2022.

[30] F. G. Martins. Tuning pid controllers using the itae criterion. *International Journal of Engineering Education*, 21(5):867, 2005.

[31] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

[32] J. J. Craig. *Introduction to robotics*. Pearson Educacion, 2006.

[33] Rail-Berkeley. Rail-berkeley/rlkit: Collection of reinforcement learning algorithms. URL https://github.com/rail-berkeley/rlkit.