

# Efficient Robot Skill Learning with Imitation from a Single Video for Contact-Rich Fabric Manipulation

Shengzeng Huo, Anqing Duan, Lijun Han, Luyin Hu, Hesheng Wang and David Navarro-Alarcon

**Abstract**—Classical policy search algorithms for robotics typically require performing extensive explorations, which are time-consuming and expensive to implement with real physical platforms. To facilitate the efficient learning of robot manipulation skills, in this work, we propose a new approach comprised of three modules: (1) learning of general prior knowledge with random explorations in simulation, including state representations, dynamic models, and the constrained action space of the task; (2) extraction of a state alignment-based reward function from a single demonstration video; (3) real-time optimization of the imitation policy under systematic safety constraints with sampling-based model predictive control. This solution results in an efficient one-shot imitation-from-video strategy that simplifies the learning and execution of robot skills in real applications. Specifically, we learn priors in a scene of a task family and then deploy the policy in a novel scene immediately following a single demonstration, preventing time-consuming and risky explorations in the environment. As we do not make a strong assumption of dynamic consistency between the scenes, learning priors can be conducted in simulation to avoid collecting data in real-world circumstances. We evaluate the effectiveness of our approach in the context of contact-rich fabric manipulation, which is a common scenario in industrial and domestic tasks. Detailed numerical simulations and real-world hardware experiments reveal that our method can achieve rapid skill acquisition for challenging manipulation tasks.

**Index Terms**—Skill learning, one-shot imitation learning, imitation from observations, contact-rich tasks, safe policy search.

## I. INTRODUCTION

**R**EINFORCEMENT learning (RL) algorithms are becoming more and more popular since they allow robots to learn by trial and error, avoiding troublesome modeling of specific situations [1]. However, it is difficult to directly transfer these RL methods to robots since these algorithms generally require millions of interactions [2]. Previous works have leveraged imitation learning to speed up the learning process, such as pouring [3], nonprehensile manipulation [4], and contact-rich manipulation [5]. Most of them require a large database comprising state-action tuples and additional explorations for policy learning. However, these two requirements are difficult to satisfy for a complex manipulation task in an unstructured environment since: (1) high-dimensional visuomotor policy is data-hungry [6]; (2) random explorations




Figure 1. Scenarios of contact-rich fabric manipulation. (a) Packing. (b) Covering. (c) Hanging.

in unstructured environments are dangerous for real robots [7]. An ideal imitation scenario is one which is similar to the way humans learn from others, e.g., a student observes and imitates a skill from a teacher [8] or even a video on the web [9]. Achieving this level of sensorimotor dexterity is precisely our goal in this paper. We develop a one-shot observational imitation approach that enables robots to deploy a task immediately.

One of the factors contributing to human beings' strong imitation ability is that we extract the temporal abstraction from demonstrations rather than precise state-action pairs [10]. Motivated by this understanding, our core idea of efficient robot skill learning is to follow the state sequences of a new scene in the demonstration as closely as possible. Two strategies are designed to achieve this goal: (1) learning general prior knowledge about the task family in a simulated scene; (2) taking long-term estimated resulting states under constraints into account with MPC. We emphasize that there are differences between sim-to-real RL methods [11] and our efficient skill-learning approach. The first one aims to create a simulation configuration that closely resembles the actual task before learning a policy for that specific activity. Compared with them, our goal is to enable robots to deploy a new scene in the task family, which does not appear in the simulation. The benefits of our formulation are as follows: (1) our policy can generalize to different scenes within the task family efficiently (with only a demonstration video); (2) we do not demand a very realistic simulation environment and additional sim-to-real skills.

Our approach is validated in the context of contact-rich fabric manipulation, which means the deformable fabric needs to interact with rigid objects and establish symbolic object-object relationships with each other. Fig. 1 presents various industrial and household scenarios of this task, revealing that it is an essential skill for robots. We choose this task for two main reasons. For one thing, it is difficult to solve with the conventional end-to-end imitation learning methods since

This work is supported by the Research Grants Council (RGC) under grant 15212721. *Corresponding author: D. Navarro-Alarcon.*

S. Huo, A. Duan and D. Navarro-Alarcon are with The Hong Kong Polytechnic University, Department of Mechanical Engineering, Hong Kong.

L. Hu is with The University of Hong Kong, Centre for Transformative Garment Production, NT, Hong Kong.

L. Han and H. Wang are with the Shanghai Jiaotong University, Department of Automation, Shanghai, China.

(1) a great scale of demonstrations is required due to the high-dimensional state (infinite degrees of freedom of fabrics) and action (bimanual manipulation) space; and (2) random explorations should be prohibited due to the potential for collisions in contact-rich environments. For another, the superiority of our approach can be illustrated with this specification since (1) transferring between various interaction scenarios provides strong evidence of the model’s inference ability; (2) the symbolic object-object association goal requires high-level planning instead of mathematical approximation. The original contributions of this work are:

- We formulate the efficient robot skill learning problem as the process of adapting from learned priors of the task family to a new scene.
- We design a state alignment-based reward function to extract temporal abstraction from a demonstration video.
- We leverage a sampling-based MPC method to generate promising and safe motions in imitation.
- We report a detailed experimental study to evaluate the performance of the proposed approach.

In contrast with most algorithms in the literature, our approach endows robots to perform a novel scene within the task family effectively. Specifically, the guidance of the new scene is only a single observational demonstration and robots can conduct it immediately without any explorations. The proposed approach could advance the development of efficient robot skill learning.

The rest of this paper is organized as follows. Sec. II introduces the related work. Sec. III states the problem formulation. Following that, we explain the details of our methodology in Sec. IV. Sec. V reports the validation experiments and the corresponding analysis. Finally, Sec. VI concludes this article.

## II. RELATED WORK

### A. Imitation Learning from Observation

Standard imitation learning methods, either behavior cloning (BC) or inverse reinforcement learning (IRL), assume both observations and actions are available in the demonstrations. In general, two methods are used to collect the proprioceptive information of the robots. For one thing, operators directly contact and guide the robot through kinesthetic teaching, which is not suitable for robots with multiple degrees of freedom. For the other, robots are teleoperated by remote control devices, such as VR [6] and mobile phones [12]. However, additional efforts of setting and training are required for the operator, affecting the practicality in real applications.

To simplify the demonstration procedure, imitation from observation (IfO) has garnered a great deal of attention, in which only the observation sequence is required. However, it brings additional challenges in imitation due to the lack of action information. A natural solution is recognizing the movement in the video to learn the policy in a conventional manner (e.g. [13], [14]). Other researchers [15], [16] establish auxiliary rewards for imitation through learning the representation and the correspondence of the observations. The main drawback of the above methods is the requirement of many demonstrations to guarantee generalizability. To learn a robot

skill efficiently, more and more researchers have switched their focus to one-shot imitation learning.

### B. One-shot Imitation Learning

One-shot imitation learning means robots are able to learn from a single demonstration of a given task and then generalize to new situations of the same task [17]. However, it is very challenging due to data scarcity. One solution is to incorporate imitation learning with RL to improve the robustness of the policy with fine-tuning. [18] integrates a task-specific inverse dynamic model into RL. [19] extracts the tool trajectory from the video and aligns the simulated environment with the video to initialize a policy for RL to learn. There are two main drawbacks of this kind of method: (1) training robots in the real world is challenging due to sample efficiency and safety concerns. (2) learning in simulation necessitates a high level of verisimilitude with the actual world and additional techniques to deal with the sim-to-real gap. Other researchers [20], [21] formulate this problem as visual servoing. [20] iterates the pose aligning between the robot and the human hand in the video; [21] obtains the approaching policy in a self-supervised manner through backward learning from the goal pose. However, the above methods hold strong task-related assumptions (e.g. known specific parameters [20], and proximity to goal pose [21]).

In addition, some researchers formulate one-shot imitation in a meta-learning manner [22], which means learning a general policy with few diverse demonstrations for different tasks, then adapting the policy for a new task given a single demonstration [3]. However, collecting such a big database is time-consuming and troublesome, especially in reality. Our work is built upon this formulation, in which we only learn prior knowledge in a single scene of the task family through random explorations and adapt to a novel scene according to a demonstration video.

### C. Contact-Rich Fabric Manipulation

Many researchers have addressed the robotic manipulation of soft materials, see [23], [24] for recent reviews. However, most existing methods (either model-based [25], [26] or model-free [27], [28]) only consider shaping tasks without complex contact with objects in structured environments. Contact-based manipulation of a deformable linear object towards a desired shape is considered in [29], while customized action primitives are required for the controller. Although fabric manipulation has gained interest from several researchers, most of them focused on laundry folding [30] and flattening [31]. [32] estimates the model utility of fabrics in unstructured environments, while the manipulation strategy is only validated in simulation. The method in [33] solves a bed-making task with strong assumptions about the initial configuration and customized placing policy. [11] exploits model-free reinforcement learning to learn a policy of hanging a towel over a hanger. However, it requires a high degree of verisimilitude in the simulation environment and a great demand for exploration resources.

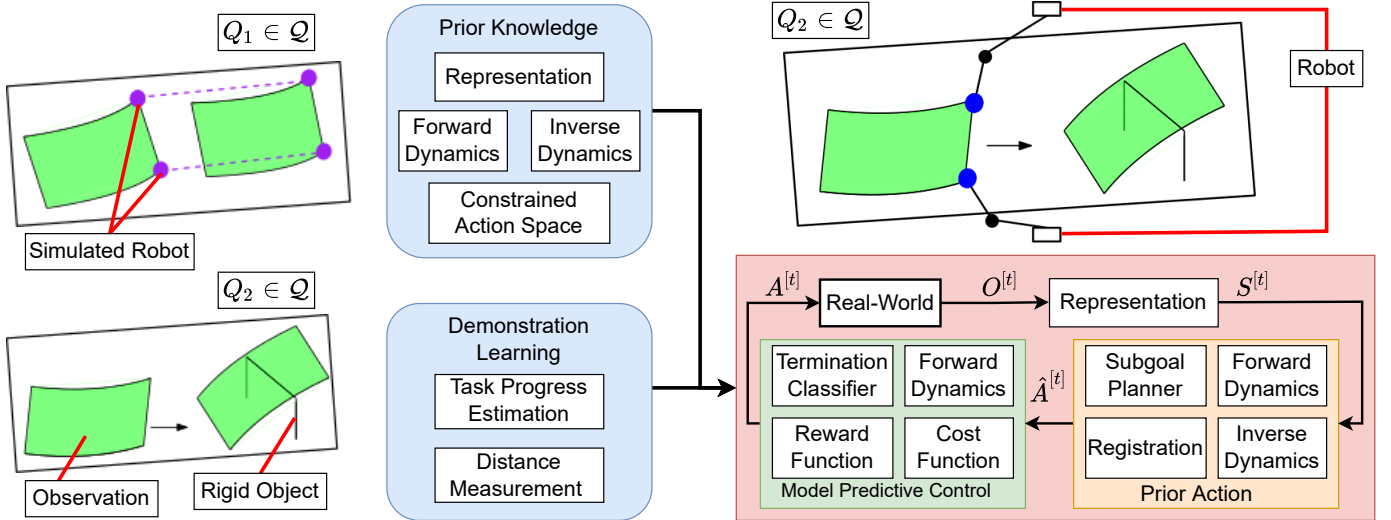


Figure 2. The overview of our efficient robot skill learning approach. With the prior knowledge learning in the scene  $Q_1 \in \mathcal{Q}$ , the model predictive control method generates motions for the novel scene  $Q_2 \in \mathcal{Q}$  guided by the reward function from the demonstration learning.

Due to the planning complexity of DOM, some researchers (e.g. [34], [35]) have leveraged imitation learning to improve the learning efficiency, while most of them require multiple demonstrations comprised of state-action transitions. Moreover, most of them only deal with simple tasks in a structured environment [34] or only implement the policy in simulation [35]. Hence, there are still open questions about how to acquire efficient robot skills for contact-rich fabric manipulation.

### III. PROBLEM FORMULATION

For a robotic task in real-world, we formulate the problem as a Partially Observable Markov Decision Process (POMDP) represented by a tuple  $\mathcal{M} = (\mathcal{S}, \mathcal{O}, \mathcal{A}, \mathcal{P}, R, \rho_0, \gamma)$ , where the state  $S^{[t]} \in \mathcal{S}$  is unknown and its corresponding observation is  $O^{[t]} \in \mathcal{O}$ . State transition function  $\mathcal{P}(S^{[t+1]}|S^{[t]}, A^{[t]})$  characterizes the probability of a switch from state  $S^{[t]}$  to  $S^{[t+1]}$  after executing action  $A^{[t]} \in \mathcal{A}$  and  $R(S^{[t]}, A^{[t]})$  is the reward function.  $\rho_0$  is the initial state distribution and  $\gamma$  is the discount factor.

For the reward function, we consider the sparse terminal configuration, which is a common situation in robot tasks [36]. This means that robots need to perform correct actions continuously and finally complete the task to obtain a positive reward. The specific definition is:

$$R(S^{[t]}, A^{[t]}) = \begin{cases} 1, & \text{if the task is completed} \\ 0, & \text{else} \end{cases} \quad (1)$$

During the learning process, robots probably enter dangerous states that cause catastrophic results, especially for safety-critical applications [37]. Hence, we induce a cost function to evaluate if an action leads to an unsafe state. Specifically, an action  $A^{[t]}$  is valid if its resulting state  $S^{[t+1]}$  is safe, denoted as:

$$C(S^{[t]}, A^{[t]}) = \begin{cases} 0, & \text{if } S^{[t+1]} \in \mathcal{S}_V \\ 1, & \text{else} \end{cases} \quad (2)$$

where  $\mathcal{S}_V \subseteq \mathcal{S}$  is the safe state subset. According to this definition, we constrain the feasible subset  $\mathcal{A}_V^{[t]}$  within the entire action space  $\mathcal{A}$  with respect to the current state  $S^{[t]}$ , denoted as  $\mathcal{A}_V^{[t]} \subseteq \mathcal{A}$ .

Constrained reinforcement learning is a general method for policy search under certain conditions [37], which aims to maximize the cumulative reward while limiting the cost incurred from constraint violations:

$$\pi^* = \arg \max_{\pi} [J_R(\pi) - J_C(\pi)] \quad (3)$$

where  $J_R(\pi) = \mathbb{E}_{\mathcal{T} \sim \pi} \left[ \sum_{t=0}^T \gamma^t R(S^{[t]}, A^{[t]}) \right]$  and  $J_C(\pi) = \mathbb{E}_{\mathcal{T} \sim \pi} \left[ \sum_{t=0}^T C(S^{[t]}, A^{[t]}) \right]$  denote the expected return of policy  $\pi$  with respect to the reward function  $R$  and the cost function respectively.

However, there are two main problems when applying RL in a real application: safe learning and sample efficiency. Human supervision and manual reset are necessary to avoid unacceptable catastrophes for robots, which are time-consuming and labor-costing for operators due to sample inefficiency. To deal with these issues, we formulate this problem as learning priors in a scene  $Q_1$  of the task family  $Q_1 \in \mathcal{Q}$  and then deploy a new scene of the same family  $Q_2 \in \mathcal{Q}$  with a demonstration provided, as shown in Fig. 2. We leverage the contact-rich fabric manipulation context to explain the idea of the formulation. We consider the interactions between a piece of fabric and different rigid objects as different scenes of the task family  $Q \in \mathcal{Q}$ . Specifically,  $Q_1$  in Fig. 2 is a contact-free scene in simulation without rigid objects, while  $Q_2$  requests the fabric to interact with a hanger to achieve the goal. The details about the policy transfer from a scene  $Q_1$  to another scene  $Q_2$  are delivered in Sec. IV.

Next, we introduce the specifications of the contact-rich fabric manipulation context, whose objective is to establish symbolic relationships between the deformable fabric and the rigid object. For example, the fabric covers the furniture

entirely and stably as shown in Fig. 1(b). The following assumptions are made for the context:

- The fabric is initially placed on a table with a flattened configuration and the rigid object is fixed on the table.
- The mask of the fabric in the scene is known.
- Dual arms rigidly grasp two ends of the fabric without any loose contact during manipulation.
- Dual arms are kinematically controlled and the coordination between them and the camera is calibrated.

We make the assumption about the fabric since there are a few existing works about unfolding [31] and the cloth region segmentation [38] is not our focus. Other assumptions are commonly used in deformable object manipulation [39].

Starting from a pre-grasped configuration, the manipulation action of our dual-arm robot  $A = [\mathbf{a}_L^T \ \mathbf{a}_R^T]$  consists of the relative movements of individual end-effectors, denoted as  $\mathbf{a}^T = [\Delta x \ \Delta y \ \Delta z]$ . When the termination conditions are reached, two end-effectors finish the movement and open the grippers to release the fabric.

In summary, our goal is to maximize the accumulated dense reward  $R_E(S^{[t]}, \{S_E^{[j]}\}_{j=1}^T)$  w.r.t. a  $H$ -length action sequence  $\mathcal{X} = (A^{[1]}, \dots, A^{[H]})$ , while avoiding breaking out the constraints  $C(S^{[t]}, A^{[t]})$ , denoted as:

$$\begin{aligned} \mathcal{X}^* = \arg \max_{A^{[1]}, \dots, A^{[H]}} \mathbb{E} \left[ \sum_{t=1}^H \gamma^t R_E \left( S^{[t]}, \{S_E^{[j]}\}_{j=1}^T \right) \right] \\ \text{s.t.} \quad S^{[t+1]} = \mathcal{P} \left( S^{[t]}, A^{[t]} \right) \quad \forall t \in \{1, \dots, H\} \\ C \left( S^{[t]}, A^{[t]} \right) = 0 \quad \forall t \in \{1, \dots, H\} \end{aligned} \quad (4)$$

## IV. METHODS

### A. Overview

Our objective is efficient robot skill learning for a new scenario, which eliminates the need for collecting a large amount of data in reality and additional explorations in the training phase following the demonstration. Fig 2 illustrates the overview of our proposed approach, which can be divided into three phases: **Prior**, **Demonstration**, and **Control**. In the following, we first introduce the procedures of data collection for **Prior** knowledge in simulation without human participation. Then, we explain how to extract the high-level temporal abstraction of a novel scenario from a **Demonstration** video. Finally, we introduce the details of our sampling-based model predictive **Control** that is able to achieve an efficient and safe policy in real-time.

### B. Prior Knowledge

The objective of learning priors is to equip agents with key knowledge about a task family  $\mathcal{Q}$ , which is generalizable and suitable for any scene within the task family  $Q \in \mathcal{Q}$ . In this paper, three kinds of prior knowledge are considered: (1) state representation; (2) dynamic models, including forward and inverse dynamics; (3) definition of the constrained action space.

To avoid specific modeling, we learn the descriptors and dynamics through neural networks in a data-driven manner.

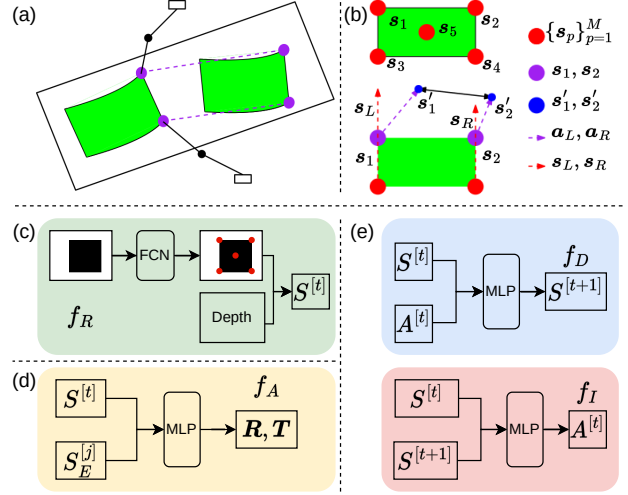


Figure 3. Graphical illustration of the priors learning procedure. (a) Bimanual manipulation in a contact-free scenario. (b) Details of the customized keypoints and the constrained action space. The architecture of the neural networks. (c) State representation  $f_R$ . (d) State registration  $f_A$ . (e) Forward dynamics  $f_D$  and inverse dynamics  $f_I$ .

We collect the required data in simulation instead of reality since (1) collecting real-world data is time-consuming and costly; (2) simulation provides direct access to the state  $S^{[t]}$  of the system, avoiding manual labeling. For our contact-rich fabric manipulation context, we learn priors in a contact-free simulation scenario as shown in Fig. 3(a), in which two virtual anchors act as grippers in a planar configuration. The reasons for choosing this kind of configuration are (1) simulating the complex fabric-rigid object interactions with a high degree of verisimilitude is challenging to achieve; (2) interactions between fabric and various rigid objects share little similarity resulting in a poor generalization. The details about the simulation specification are introduced in Sec. V-A.

A mask of fabric in an RGB image plays as the raw observation  $O^{[t]}$  in our context. Fig. 3(b) shows our customized state description with  $M$  keypoints  $S^{[t]} = \{s_p^{[t]}\}_{p=1}^M$ , which is given directly in simulation. According to this semantic descriptor, we define the constrained action space  $A^{[t]} \in \mathcal{A}_V^{[t]}$  with respect to an individual state  $S^{[t]}$  for the entire contact-rich task family  $\mathcal{Q}$ . The basic principles of the constrained action space are: (1) ensuring the safety of the robots; (2) avoiding the severe self-occlusion of the fabrics during the manipulation. In the following, we present the details of the constrained action space.

As shown in Fig. 3(b), the positions of the end-effectors after executing the movement are denoted as  $s'_1 = s_1 + \mathbf{a}_L$  and  $s'_2 = s_2 + \mathbf{a}_R$ , respectively. Also, we define the edge vectors  $s_L = s_1 - s_3$  and  $s_R = s_2 - s_4$  corresponding to the actions  $A = [\mathbf{a}_L^T \ \mathbf{a}_R^T]$ , respectively. The rules of avoiding approaching dangerous states after executing an action  $A^{[t]}$  are:

- Two grippers keep an appropriate distance from each other, avoiding over-stretching or severe self-occlusion:

$$\tau_C \times L < \|s'_1 - s'_2\|_2 < \tau_F \times L \quad (5)$$




Figure 4. (a) Graphical explanation of the state alignment-based reward function. The reward of a state  $S^{[t]}$  is measured through registration with the demonstration. The state  $S'^{[t]}$  that passes the final instance in the demonstration  $S_E^{[T]}$  is penalized in the reward function. (b) The paradigm of the acquisition of the prior action.

where  $L$  is the original length of the edge of the fabric,  $\tau_C$  and  $\tau_F$  are the ratio thresholds to measure the distance between two end-effectors with respect to  $L$ .

- The direction difference between the individual action  $(\mathbf{a}_L, \mathbf{a}_R)$  and its corresponding edge of the fabric  $(\mathbf{s}_L, \mathbf{s}_R)$  is constrained under the threshold  $\tau_D$ :

$$\max(\langle \mathbf{a}_L, \mathbf{s}_L \rangle, \langle \mathbf{a}_R, \mathbf{s}_R \rangle) < \tau_D \quad (6)$$

where  $\langle \mathbf{a}, \mathbf{b} \rangle$  computes the angle between two vectors and  $\tau_D$  is an angle limitation threshold.

- The fabric locates within the workspace.
- Two grippers do not collide with the fixed rigid objects in the environment (if any).

To avoid troublesome human intervention, we sample random movements within the constrained action space  $A^{[t]} \in \mathcal{A}_V^{[t]}$  at each time-step  $t$ . Through the random explorations, we acquire several transitions storing in a dataset  $\mathcal{D} = (O^{[t]}, S^{[t]}, A^{[t]}, O^{[t+1]}, S^{[t+1]})$ , where  $O^{[t]}$  is the masked image of the fabric and  $S^{[t]}$  is the corresponding customized keypoints. Based on the collected dataset  $\mathcal{D}$ , we train our data-driven models. Firstly, we learn the representation  $f_R(S^{[t]}|O^{[t]})$  with a Fully Convolution Network [29] (as shown in Fig. 3(c)), which outputs the pixel coordinates of the keypoints  $S^{[t]}$  in the mask image. Then, we obtain the 3D representation of the keypoints by querying the corresponding depth value. Finally, two fully-connected networks are established to approximate the forward  $f_D(S^{[t+1]}|S^{[t]}, A^{[t]})$  and the inverse dynamics  $f_I(A^{[t]}|S^{[t]}, S^{[t+1]})$  of the system on the basis of the predicted keypoints  $S^{[t]}$ , as shown in Fig. 3(e). All the models are trained with the mean square error (MSE).

### C. Demonstration Learning

Instead of learning an end-to-end imitation policy directly, we emphasize the resulting state of each action since (1) learning the state-action mapping is data hungry; (2) the

simulation-to-reality gap affects the performance of the end-to-end policy in transfer; (3) the suitable action with respect to a state to approach the ultimate goal is generally not unique, especially for a high-dimensional task. As a result, we convert the conventional end-to-end policy as searching for an optimal action whose forecast resulting state is closest to the state sequence in the demonstration. [40] holds a similar idea to us, while they are interested in an imitation scenario whose dynamics mismatch with the expert. Compared with them [40], we consider a more complex scenario (inference to a novel scene of the task family), and the imitation is conducted immediately following a single demonstration video without any additional explorations.

Given the  $T$ -length single observation-only demonstration  $\tau_E = \{O_E^{[j]}\}_{j=1}^T$ , we first obtain the corresponding states  $\{S_E^{[j]}\}_{j=1}^T$  with our representation model  $f_R(S^{[t]}|O^{[t]})$ . Since the goal of the contact-rich task is to establish the symbolic relationship between the fabric and the rigid object, we set up a coordinate frame associated with the pose of the rigid object and represent all the states in this frame. The advantage of this setting is that the imitation is platform-agnostic, regardless of the configuration of sensors or actuators. Since the 6-DOF pose recognition of a rigid object is not a focus of our work, it is achieved through detecting a marker in real-world experiments.

Fig. 4(a) explains the state alignment process. Specifically, there are two critical issues to determine in the reward function: (1) estimate the temporal progress in the imitation; (2) measure the spatial distance to the final goal.

For the first issue, we leverage the Intersection over Union (IoU) index to measure the similarity  $D_E(S^{[t]}, S_E^j)$  between the current state  $S^{[t]}$  and the states in the demonstration  $\{S_E^{[j]}\}_{j=1}^T$ .

$$D_E(S^{[t]}, S_E^j) = \frac{S^{[t]} \cap S_E^j}{S^{[t]} \cup S_E^j} \quad (7)$$

However, there is no overlapped area between the state  $S^{[t]}$  and the state sequence in the demonstration in some cases, which always happen at the early stage of the imitation. As a result, the estimator fails to distinguish between various states. Under this circumstance, we switch the similarity metric  $D_E(S^{[t]}, S_E^j)$  to the Euclidean distance of the centers. According to the above metrics, the temporal progress estimation is implemented with:

$$J = \arg \min_j D_E(S^{[t]}, S_E^j) \quad (8)$$

We assume the provided demonstration is optimal; thus the imitator should follow the demonstration step-by-step. For this sequential imitation, we take two aspects into consideration for the second issue (as shown in Fig. 4(a), whose details are: (1) **Positive**: the distance to the subgoal  $S_E^{[j+1]}$ , which is the next state of our temporal progress estimation result in the demonstration. (2) **Negative**: the distance to the final state  $S_E^{[T]}$  in the demonstration when the current state passes the last state in the demonstration  $S_E^{[T]}$ . The first term encourages the agents to approach the next stage in the demonstration, while the second term penalizes the actions when the resulting state

crosses the ultimate state in the demonstration. This crossing phenomenon ( $S^{[t]}$  in Fig. 4(b)) should be avoided since the manipulation in the opposite direction is required to return the state back to the common situation, which is not allowed in the constrained action space.

We implement registration between the current state  $S^{[t]}$  and the subgoal  $S_E^{[J+1]}$  to quantify their **Positive** similarity. A neural network  $f_A(\mathbf{R}, \mathbf{T} | S^{[t]}, S_E^{[J+1]})$  is utilized to output their relative transformation, including the rotation  $\mathbf{R}$  and the translation  $\mathbf{T}$ . The detailed architecture is shown in Fig. 3(d). This registration model is trained by taking the states in the prior dataset  $\mathcal{D}$  as the source and the demonstration  $\{S_E^{[j]}\}_{j=1}^T$  as the target, respectively, whose goal is to minimize the MSE between  $\{S_E^{[j]}\}_{j=1}^T$  and the transformed source  $\mathbf{R} \cdot S^{[t]} + \mathbf{T}$ :

$$\mathcal{L}_A = \|\mathbf{R} \cdot S^{[t]} + \mathbf{T} - S_E^{[j]}\| \quad (9)$$

Based on the trained model, the distance scalar  $D_P(S^{[t]}, S_E^{[J+1]})$  is obtained through integrating the rotation  $\mathbf{R}$  and the translation  $\mathbf{T}$ :

$$D_P(S^{[t]}, S_E^{[J+1]}) = |\mathbf{T}| + w_R \cdot |\mathbf{R}| \quad (10)$$

where  $\mathbf{R}, \mathbf{T} = f_A(\mathbf{R}, \mathbf{T} | S^{[t]}, S_E^{[J+1]})$  and  $w_R$  is the weight ratio to balance the rotation and the translation. Here, the rotation is represented with the Euler angle to correspond to the translation with different axes.

Compared with **positive** distance  $D_P(S^{[t]}, S_E^{[J+1]})$ , the **negative** distance  $D_N(S^{[t]}, S_E^{[T]})$  is triggered only when the current state passes the last state in the demonstration, whose definition is:

$$D_N = \begin{cases} \|S^{[t]} - S_E^{[T]}\| & J > T - 2 \wedge \\ & \|S^{[t]} - S_E^{[T]}\| < \|S^{[t]} - S_E^{[T-1]}\| \\ 0 & \text{else} \end{cases} \quad (11)$$

Taking the above elements into consideration, the state alignment-based reward is denoted as:

$$R_E \left( S^{[t]}, \{S_E^{[j]}\}_{j=1}^T \right) = J - w_P \cdot D_P(S^{[t]}, S_E^{[J+1]}) - w_N \cdot D_N(S^{[t]}, S_E^{[T]}) \quad (12)$$

where  $w_P$  and  $w_N$  are the weight ratio of the positive distance and the negative distance, respectively.

#### D. Model Predictive Control

Incorporating the prior knowledge from  $Q_1 \in \mathcal{Q}$ , we hope that our robots deploy a novel task  $Q_2 \in \mathcal{Q}$  immediately after a demonstration video is given without any explorations. To this end, we take the safety issues of real robots into consideration and formulate the policy search as a constrained optimization problem. MPC is an effective method to deal with this issue. For example, [41] considers the robotic visual servoing system's input and output constraints, while it depends on an analytical expression of the task. [37] studies the safe RL problem with sparse indicator signals for constraint violations. However, it requires a few violation budgets to explore the environment and only deals with a task with original dense rewards in simulation. In this paper, we extend this method

---

#### Algorithm 1: State Alignment-driven MPC

---

**Input:** Prior models ( $f_R, f_D, f_I$ ), Noise distribution  $\epsilon \sim \mathcal{N}(\mu, \Sigma)$

**Output:** Solution  $A^{[t]}$  with the highest reward  
Collect the observation-only demonstration  $\{O_E^{[j]}\}_{j=1}^T$

Obtain the state alignment-based reward function  
 $R(S^{[t]}, S_E) \leftarrow$  Eq. 12

**while** *is not terminated* **do**

    Get the action prior  $\hat{\mathcal{X}}^{[t]} \leftarrow$  Eq. 13

**while** *the convergence is not met* **do**

        Sample  $N$  solutions  $\Omega \in \{\mathcal{X}_i\}_{i=1}^N$

        Evaluate the accumulated reward  $\mathcal{R}(\mathcal{X}; S^{[t]})$   
        and the cost  $\mathcal{C}(\mathcal{X}; S^{[t]})$  for each solution.

        Select and sort the feasible set  $\Omega_V \subseteq \Omega$ .

        Update the parameters  $(\mu, \Sigma) \leftarrow$  Eq. 14

for efficient robot skill learning in reality in three aspects: (1) customize a proper reward function and a cost function for safe and robust imitation from observation; (2) acquire a prior action to improve the efficiency of the sampling-based controllers; (3) design a termination classifier that predicts if the imitation has been completed.

Taking the future expectation into consideration, MPC strives to obtain an  $h$ -length action sequences  $\mathcal{X}^{[t]} = \{A^{[t+k]}\}_{k=0}^{h-1}$  through solving an open-loop optimal control problem. To enable real-time control, it is important to have a high computation efficiency for each optimization step. The sampling-based MPC approach in [37] searches within the entire action space  $\mathcal{A}$ , resulting in the low efficiency of the iteration scheme for a high-dimensional control policy. To resolve this problem, we obtain a prior action sequence  $\hat{\mathcal{X}}$  acting as a baseline to narrow down the search space. The procedure of this paradigm is shown in Fig. 4(b). For the current state  $S^{[t]}$ , the subgoal planner  $f_G(S^* | S^{[t]}, \{S_E^{[j]}\}_{j=1}^T)$  outputs a target  $S^*$ , which is actually  $S_E^{[J+1]}$  computed with our temporal progress estimation model in Eq. (8). Then we acquire an estimated action  $\hat{A}^{[t]}$  at the current timestep  $t$  with the inverse dynamic model  $f_I$ .

$$\hat{A}_t = f_I(S^{[t]}, S^*), S^* = S_E^{[J+1]} \quad (13)$$

Since this imagined action will not be executed actually, we predict the achieving state  $\hat{S}^{[t+1]}$  with the forward dynamic model  $f_D$ . Finally, this estimated state  $\hat{S}^{[t+1]}$  will act as the new  $S^{[t]}$  to start another new loop. The above procedure iterates  $h$  steps to acquire a complete action sequence baseline  $\hat{\mathcal{X}} = \{\hat{A}^{[t+k]}\}_{k=0}^{h-1}$ .

Based on the prior action sequence  $\hat{\mathcal{X}}$ , each candidate of the sampling set  $\mathcal{X} \in \Omega$  for the constrained optimization is generated through integrating the noise with it  $\mathcal{X} = \hat{\mathcal{X}} + \epsilon$ . In particular,  $\epsilon$  is sampled from a  $n$ -dimensional factorized multivariate Gaussian distribution  $\epsilon \sim \mathcal{N}(\mu, \Sigma)$ , where  $\mu$  is the mean vector and  $\Sigma$  is a diagonal covariance matrix.

Next, the accumulated performance of each solution  $\mathcal{X} \in \Omega$  is evaluated with our customized reward function

$R_E(S^{[t]}, \{S_E^{[j]}\}_{j=1}^T)$  and the cost function  $C(S^{[t]}, A^{[t]})$ :

$$\begin{aligned} \mathcal{R}(\mathcal{X}; S^{[t]}) &= \sum_{k=1}^h \gamma^{k-1} R_E(S^{[t+k]}, \{S_E^{[j]}\}_{j=1}^T) \\ \mathcal{C}(\mathcal{X}; S^{[t]}) &= \sum_{k=1}^h C(S^{[t+k-1]}, A^{[t+k-1]}) \end{aligned} \quad (14)$$

where  $S^{[t+k]} = f_D(S^{[t+k-1]}, A^{[t+k-1]})$ ,  $\forall k \in \{1, \dots, h\}$  is predicted by the dynamic model  $f_D(S^{[t+1]}|S^{[t]}, A^{[t]})$ . Then, we select the feasible solutions of all the candidates  $\Omega_V \subseteq \Omega$  whose accumulated cost is zero. We sort these feasible solutions  $\Omega_V$  and select the top samples to update the noise distribution  $\epsilon \sim \mathcal{N}(\mu, \Sigma)$  for the next iteration  $i+1$  as:

$$\begin{aligned} \mu^{[i+1]} &\leftarrow (1 - \beta)\mu^{[i]} + \beta\mu_V \\ \Sigma^{[i+1]} &\leftarrow (1 - \beta)\Sigma^{[i]} + \beta\Sigma_V \end{aligned} \quad (15)$$

where  $\beta$  is a hyper-parameter to determine the update portion for each iteration,  $\mu_V$  and  $\Sigma_V$  are the mean and the covariance of the noise corresponding to the feasible set  $\Omega_V$ . The iteration scheme is stopped when the optimal solution is feasible  $C(\mathcal{X}; S^{[t]}) = 0$  and the variance of the noise  $\epsilon$  reaches the convergence:

$$\max \Sigma < \tau_C \quad (16)$$

As a receding horizon control method, agents only apply the first input in the solution while others are discarded.

The outline of the control law is presented in Alg. 1. After observing a single demonstration video, dual arms execute the generated motion from the controller at each time-step  $t$  to approach the ultimate goal continuously. An episode is terminated and the grippers are released if one of the situations occurs: (1) no feasible action under cost constraints is found; (2) exceeding the maximum exploration steps; (3) the goal is considered as finished from the terminal classifier. Since the goal specifications are diverse in the task family and challenge to model accurately in some complex instances, we instead exploit the similarity between the achieved state  $S^{[t+1]}$  and the final state in the demonstration  $S_E^{[T]}$  to classify the goal completion. Specifically, the task is considered finished when two conditions are satisfied simultaneously: (1) the IoU between  $S^{[t+1]}$  and  $S_E^{[T]}$  is larger than the pre-defined threshold  $\tau_I$ ; (2) the temporal progress estimation (Eq. (8)) of  $S^{[t+1]}$  matches with  $S_E^{[T]}$ .

## V. RESULTS

To assess our strategy comprehensively, we undertake statistical comparisons of our approach versus baselines and ablations in simulation. We begin by outlining the details of the simulation and a comparison study with several baselines. Then, an ablation study is implemented to evaluate the necessity of each component in our approach. Finally, we demonstrate the practicality of our method for efficient robot skill learning in real applications.

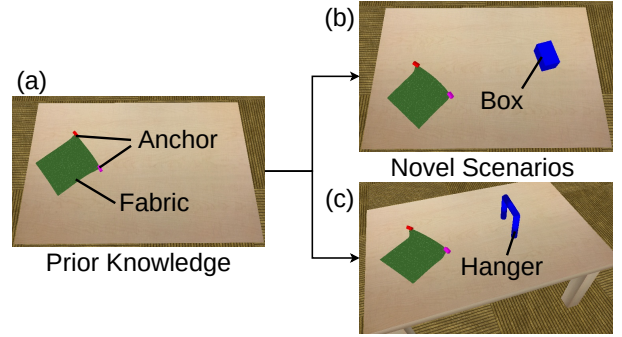


Figure 5. Screenshots of the simulation environment. (a) Contact-free configuration. (b) **Box**: Cover a box. (c) **Hang**: Hang up a Fabric.

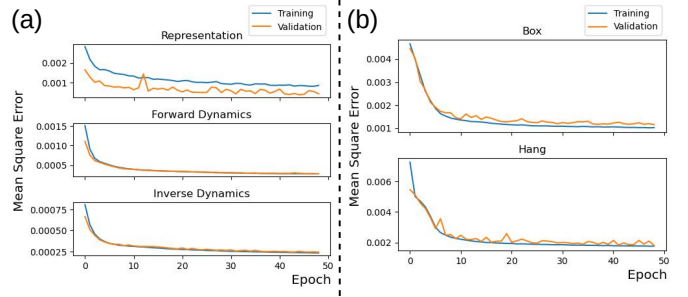


Figure 6. The training and validation errors for the prior models. (a) Representation, forward and inverse dynamics. (b) Registration.

### A. Simulation Setting

Fig. 5 illustrates our simulation environment based on DeDo [42], in which fabric is represented as a deformable grid mesh with 256 vertices. Two anchors attached to the corners of the fabric rigidly are used to render the grasping of two end-effectors.

To train the prior models in Sec. IV-B, we generate 10000 transitions  $(O^{[t]}, S^{[t]}, A^{[t]}, O^{[t+1]}, S^{[t+1]})$  through sampling random actions within the valid space  $\mathcal{A}_v$  at each time-step  $t$ , as shown in Fig. 5(a). Note that there are no rigid objects to interact with the fabric in this environment.

Two common contact-rich fabric manipulation scenarios, namely covering a box (**Box**) and hanging a towel (**Hang**), are established in simulation to evaluate the performance of the methods, as shown in Fig. 5(b) and Fig. 5(c), respectively. Note that these two environments are used to evaluate the algorithm instead of learning the specific models to transfer to the same task in reality. In the simulation, we provide the demonstration by controlling the anchors while the action information is not recorded. Rigid objects in both environments are fixed on the table. Each task is considered as successful if the symbolic relationship between the fabric and the rigid object is established after releasing the grippers. Starting from a pre-grasp configuration, the virtual grippers move and release at the end of an episode.

### B. Comparisons with Baselines

In this section, we first present the achieved results of our approach. Then, we compare our approaches with other baselines.




Figure 7. The scenarios of the contact-rich fabric manipulation tasks in simulation. (a) **Box**. (b) **Hang**.




Figure 8. The Chamfer distance comparison of several baselines. (a) **Box**. (b) **Hang**. The accumulated return comparison in the ablation study. (c) **Box**. (d) **Hang**.

The training and validation loss for the prior models (state representation and dynamic models) and the registration model for demonstration learning are visualized in Fig. 6(a) and (b) respectively. All of the training only requires 20 epochs. To analyze our approach in detail, we provide two typical successful examples of individual contact-rich tasks in the trials, visualized in Fig. 7(a) and (b) respectively. With a random initial configuration, the agents strive to align the states of the fabric to the state sequences in the demonstration. When the fabric is close to the rigid object, two virtual end-effectors follow the state sequences to lift up the object to avoid obstacle collision. The attachment interactions disappear when the termination classifier returns a positive value and the fabric falls down slowly due to its gravity.

The goal of our approach is efficient robot skill learning without access to the actual scenario in the training phase. Specifically, robots learn the control policy from an offline simulated dataset instead of explorations in the exact environment. Moreover, robots should distill the learned knowledge in a scene of the task family and adapt to a new scene guided by a single observational demonstration. To this end, our approach incorporates prior knowledge and a state-alignment reward function into the sampling-based MPC method to obtain a safe and robust control policy. To show the substantial improvements in our methods corresponding to the above arguments, we compare our method against various baselines,

including **Zero** [43], **CFM** [44] and **TCN** [45]. The details of individual baselines are:

- **Zero** learns a goal-conditioned skill policy with a forward consistency loss and then mimics the expert from a sequence of demonstration images.
- **CFM** optimizes the visual representation and the dynamics with contrastive learning and then implements a model-based one-step optimal predictive control.
- **TCN** disambiguates temporal changes in the demonstration videos to provide a reward function for agents’ policy search.

For a fair comparison, all the data-driven baselines share the same collected dataset  $\mathcal{D}$  in simulation with respect to the task  $Q_1 \in \mathcal{Q}$  and a demonstration video for a novel task  $Q_2 \in \mathcal{Q}$ . In addition, all the models are multi-layer perceptions (MLP) with two hidden layers of size 256 followed by ReLU activation functions.

Multiple trials are conducted to thoroughly assess the performance of the baselines. To evaluate the contact-rich fabric manipulation task performance quantitatively, we select Chamfer distance error [46] as a metric to compare the achieved observation  $O^{[t]}$  and the ultimate observation in the expert demonstration  $O_E^{[T]}$ , whose definition is:

$$d_{CD} = \frac{1}{O_1} \sum_{x \in O_1} \min_{y \in O_2} \|x - y\|_2^2 + \frac{1}{O_2} \sum_{y \in O_2} \min_{x \in O_1} \|y - x\|_2^2 \quad (17)$$

where  $O_1$  and  $O_2$  are two downsampled point clouds. Specifically, we acquire the 3D point cloud of the fabric by querying the corresponding depth value of the mask in the visual observation and down-sampling them to  $N = 200$  points equally.

Fig. 8(a) displays the Chamfer Distance curve for different baselines and Table I displays their success rates. Our method outperforms other baselines with respect to two individual tasks in success rate, while is comparable to **TCN** and **Zero** in terms of Chamfer distance in two tasks respectively. In the following, we present an analysis of this result. With an embedding model to distinguish different states in the space, the spatial and temporal connection between them is not considered in **CFM**, resulting in poor performance in this complex sequential manipulation task. We consider the failure of **TCN** is caused by the scarcity of data in a single demonstration, which hinders it to encode the temporal distance between the states and goal accurately. Without a long-term prediction, it is easy for **Zero** to trap in a local distance minimum to the ultimate goal. In addition, the comparable performance in terms of Chamfer Distance is mainly due to



Table I  
COMPARISONS BETWEEN DIFFERENT METHODS

Method	box			hanger		
	Rate %	$\mu_{IoU}$	$\sigma_{IoU}$	Rate %	$\mu_{IoU}$	$\sigma_{IoU}$
<b>Zero</b>	1.8	0.02	0.01	0.0	0.00	0.00
<b>CfM</b>	0.0	0.00	0.01	1.4	0.02	0.12
<b>TCN</b>	6.6	0.08	0.22	2.4	0.03	0.14
<b>No MPC</b>	0.0	0.00	0.01	3.0	0.03	0.17
<b>No Prior</b>	14.4	0.15	0.33	7.4	0.07	0.24
<b>No Cost</b>	3.0	0.03	0.16	2.2	0.02	0.14
<b>Our</b>	<b>81.9</b>	0.74	0.38	<b>72.7</b>	0.69	0.42

$\mu_{iou}$  and  $\sigma_{iou}$  are the mean and the variance of IoU, respectively.

the constraint awareness in our approach. Based on the state alignment reward function, our approach seeks to reduce the error between the achieved results and the state sequences in the demonstration video step-by-step in a safe way under the cost function. However, the above baselines generate motions to go straight to the final goal. Although they can achieve a relatively small distance from the target, the desired symbolic object-object relationship is not established, resulting in low success rates. In other words, Chamfer Distance is only an auxiliary value to indicate the imitation process. This is also the reason why we need a state alignment-based reward to guide our control policy for this symbolic task.

### C. Ablation Study

In order to highlight the necessity of each component in our proposed approach, we implement a detailed ablation study. We contrast three alternative methods, including (1) **No MPC**: Adopt the output from the inverse dynamic model. (2) **No Prior**: The MPC law iterates from scratch; (3) **No Cost**: The cost constraint function is removed in the MPC.

According to the POMDP formulation, we evaluate the performance of several ablations with the accumulated reward in this comparison. Note that agents will obtain a positive reward for episodic success and a negative warning reward due to the violation of the cost constraints. We compare the accumulated rewards of alternative ablations in Fig. 8(c)-(d) and compare their success rate of them in Table. I. These results show that our suggested MPC approaches get the best success rate among them, while others are affected severely.

In the following, we provide a detailed analysis of their performance. Without the predictive control, inverse dynamics is used in **No MPC** to output the control command. As a result, this controller usually traps in a local minimum to pursue the final goal while ignoring the long-term impacts. Without the action prior, **No Prior** generates random action sequences within the huge space  $\mathcal{A}$ . This setup prevents the algorithm from finding a solution that satisfies the requirements in a finite iteration period. Removing the cost function in **No Cost**, the solution from the optimization usually violates the constraints, which results in stopping an imitation episode in advance.

Although our approach is capable of handling the majority of the challenging tasks, there are some situations when it fails. Fig. 9 presents two typical failure examples. The failure case in Fig. 9(a) is mainly caused by the state alignment settings. The controller occasionally struggles in complex interactions




Figure 9. Typical failure cases in simulation. (a) The fabric is manipulated to pass the rigid object. (b) The fabric reaches a state where it is difficult for the controller to acquire a feasible action.

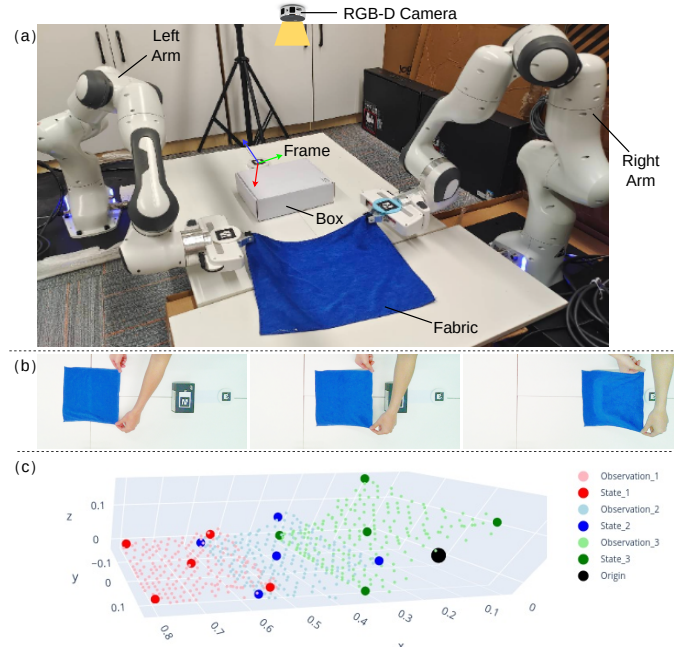


Figure 10. Physical Scenarios to validate our approach. (a) Experimental Setup. (b) The process of the demonstration video provision with human hands. (c) Graphical illustration of the sequential demonstration.

with the rigid object since they are not considered in the dynamics models of priors. As a result, agents can not reach a state that satisfies the termination classifier and have to iteratively generate motions until passing the final state in the demonstration regardless of the penalization term in the reward function. The failure in Fig. 9(b) is mainly caused by the computation efficiency of the sampling-based MPC method. To achieve real-time control, we define a finite searching period for each time-step, and our approach occasionally fails to find out an appropriate solution within this period. This phenomenon is more obvious when the fabric is close to the rigid object since the action space is more restricted.

### D. Physical Experiments

This section gives the results of our proposed approach on real-world experiments. Fig. 10(a) shows our robotic setup, which contains two Franka Emika Panda arms equipped with 2-fingered grippers and a visual feedback device. The RGB-D camera is fixed to provide a top-down perspective of the




Figure 11. The scenarios of various contact-rich fabric manipulation tasks in real experiments. (a) **Box-Small**. (b) **Hang**. (c) **Box-Big**.

Table II  
RESULTS IN PHYSICAL EXPERIMENTS

Task	Fabric (cm)	Rigidity (cm)	Success Rate %	
			Translation	Transformation
<b>Box-Small</b>	30×30	14×9×6	12/15	7/10
<b>Hang</b>	30×30	17×1×20	12/15	8/10
<b>Box-Big</b>	30×40	30×15×5	14/15	9/10




Figure 12. Typical failure cases in real-world experiments. (a) The fabric is manipulated to pass the rigid object. (b) The fabric reaches a state where it is difficult for the state representation model  $f_R(S^{[t]}|O^{[t]})$  to detect the customized keypoints.

manipulation space. For a novel task, an operator provides a demonstration video about the bimanual manipulation with the hands, as shown in Fig. 10(b). Note that we only collect the raw observations from the RGB-D camera and recognize the pose of the rigid object with its attached marker. Then, we process the collected observations with our representation model  $f_R(S^{[t]}|O^{[t]})$  to acquire the corresponding states, as shown in Fig. 10(c).

To validate the effectiveness and robustness of our approach, we perform three kinds of contact-rich fabric manipulation tasks under different configurations. The details of various tasks are introduced in Table. II and the individual successful scenarios are shown in Fig. 11. Specifically, we consider three kinds of tasks, including **Box-Small**, **Hang**, and **Box-Big**. The major difference between **Box-Small** and **Box-Big** is the size of the fabric, which is used to evaluate the robustness of

our algorithm in terms of the deformable object. Without any additional explorations, our dual-arm robot is able to execute the imitation for a novel task immediately with our efficient skill-learning approach. For each episode, dual arms start from a pre-grasp configuration and move the end-effectors when the control command is received. An episode is considered successful if the fabric establishes a corresponding symbolic relationship with the rigid object stably after releasing the grippers.

For each task, two levels of difficulty are analyzed: (1) **Translation**: Only change the position of the rigid object w.r.t. the situation in the demonstration. (2) **Transformation**: Change the pose of the rigid object arbitrarily within the workspace. The success rate for each individual contact-rich manipulation task is summarized in Table. II. Each episode lasts approximately 20s to 40s in real-time, depending on the spatial distance between the initial configuration and the ultimate goal. These quantitative results illustrate the superior performance of our proposed approach in terms of efficiency and robustness.

Fig. 12(a) shows a failure case in physical experiments, which is similar to the scenario of Fig. 9(a) in simulation. Since the exploration data is only collected in simulation, the transfer performance in reality is unavoidably affected due to their gap. Compared with the simulation results in Sec. IV-B, we acquire two distinct observations when analyzing the performance of the physical implementations. The first one is that the success rate is decreased when the rotation transform is included. This is mainly because the kinematic feasibility of the robotic arms limits the manipulation space and brings additional difficulties for robots to align with the expert demonstrations. Considering the manipulability of robot arms in simulation instead of simplifying them as virtual anchors is beneficial to alleviate this problem. The second one is that the occlusion phenomenon is unavoidable in robotic manipulations under grasping configuration, resulting in a lower keypoint detection accuracy. In the future, we consider leveraging temporal information to track the manipulated object during the entire manipulation.

## VI. CONCLUSION

In this paper, we propose an efficient robot skill learning approach through imitation from a single demonstration video. We learn general prior knowledge in one scene and a state alignment-based reward function for a new scene based on the provided demonstration video. Robots are able to deploy the new scenario immediately by incorporating the above issues into our sampling-based MPC method. Since we do not assume the exact dynamic consistency between training and evaluation, the data collection procedure can be implemented in simulation to avoid time-consuming and high-cost procedures in real scenarios. Furthermore, the need for random physical explorations of robots is obviated, thereby mitigating potential hazards in unstructured environments. The performance of the proposed approach is evaluated in the context of contact-rich fabric manipulation, in which two robotic arms need to cooperate in a constraint-aware manner. The results in both the simulation and real-world experiments show that robots are able to imitate the demonstration videos efficiently and safely.

In future work, we seek to leverage the temporal experience in the deployment of a task to improve the policy incrementally. In addition, integrating force sensing is beneficial to achieve a flexible and compliant control strategy for challenging robotic tasks.


## REFERENCES

- [1] X. Wang, S. Wang, X. Liang, D. Zhao, J. Huang, X. Xu, B. Dai, and Q. Miao, "Deep reinforcement learning: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2022.
- [2] K. Chatzilygeroudis, V. Vassiliades, F. Stulp, S. Calinon, and J.-B. Mouret, "A survey on policy search algorithms for learning robot controllers in a handful of trials," *IEEE Transactions on Robotics*, vol. 36, pp. 328–347, 2018.
- [3] D. Zhang, W. Fan, J. Lloyd, C. Yang, and N. F. Lepora, "One-shot domain-adaptive imitation learning via progressive learning applied to robotic pouring," *IEEE Transactions on Automation Science and Engineering*, pp. 1–14, 2022.
- [4] X. Sun, J. Li, A. V. Kovalenko, W. Feng, and Y. Ou, "Integrating reinforcement learning and learning from demonstrations to learn non-prehensile manipulation," *IEEE Transactions on Automation Science and Engineering*, pp. 1–10, 2022.
- [5] X. Liu, P. Huang, and Z. Liu, "A novel contact state estimation method for robot manipulation skill learning via environment dynamics and constraints modeling," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 3903–3913, 2022.
- [6] T. Zhang, Z. McCarthy, O. Jow, D. Lee, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–8, 2017.
- [7] P. Mitrano, D. McConachie, and D. Berenson, "Learning where to trust unreliable models in an unstructured world for deformable object manipulation," *Science Robotics*, vol. 6, 2021.
- [8] E. Valassakis, G. Papagiannis, N. Di Palo, and E. Johns, "Demonstrate once, imitate immediately (dome): Learning visual servoing for one-shot imitation learning," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8614–8621, 2022.
- [9] H. Zhang, P. Lai, S. Paul, S. Kothawade, S. Nikolaidis, H. Zhang, P. Lai, S. Paul, S. Kothawade, and S. Nikolaidis, "Learning collaborative action plans from youtube videos," in *International Symposium of Robotics Research*, 2019.
- [10] F. Liu, Z. Ling, T. Mu, and H. Su, "State alignment-based imitation learning," in *International Conference on Learning Representations*, 2020.
- [11] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*, pp. 734–743, PMLR, 2018.
- [12] A. Mandlkar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, et al., "Roboturk: A crowdsourcing platform for robotic skill learning through imitation," in *Conference on Robot Learning*, pp. 879–893, PMLR, 2018.
- [13] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," *ArXiv*, vol. abs/1805.01954, 2018.
- [14] S. Bahl, A. Gupta, and D. Pathak, "Human-to-robot imitation in the wild," *ArXiv*, vol. abs/2207.09450, 2022.
- [15] K. Zakka, A. Zeng, P. R. Florence, J. Tompson, J. Bohg, and D. Dwibedi, "Xirl: Cross-embodiment inverse reinforcement learning," in *Conference on Robot Learning*, 2021.
- [16] O. Mees, M. Merklinger, G. Kalweit, and W. Burgard, "Adversarial skill networks: Unsupervised robot skill learning from video," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4188–4194, 2019.
- [17] Y. Duan, M. Andrychowicz, B. Stadie, O. Jonathan Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-shot imitation learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [18] B. S. Pavse, F. Torabi, J. Hanna, G. Warnell, and P. Stone, "Ridm: Reinforced inverse dynamics modeling for learning from a single observed demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6262–6269, 2020.
- [19] K. Zorina, J. Carpentier, J. Sivic, and V. Petrik, "Learning to manipulate tools by aligning simulation to video demonstration," *IEEE Robotics Autom. Lett.*, vol. 7, pp. 438–445, 2021.
- [20] X. Deng, J. Liu, H. Gong, H. Gong, and J. Huang, "A human-robot collaboration method using a pose estimation network for robot learning of assembly manipulation trajectories from demonstration videos," *IEEE Transactions on Industrial Informatics*, pp. 1–9, 2022.
- [21] E. Johns, "Coarse-to-fine imitation learning: Robot manipulation from a single demonstration," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4613–4619, 2021.
- [22] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International conference on machine learning*, pp. 1126–1135, PMLR, 2017.
- [23] H. Yin, A. Varava, and D. Kragic, "Modeling, learning, perception, and control methods for deformable object manipulation," *Science Robotics*, vol. 6, no. 54, 2021.
- [24] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. LI, J. Pan, W. Yuan, and M. Gienger, "Challenges and outlook in robotic manipulation of deformable objects," *IEEE Robotics & Automation Magazine*, pp. 2–12, 2022.
- [25] F. Zhang and Y. Demiris, "Learning garment manipulation policies toward robot-assisted dressing," *Science Robotics*, vol. 7, 2022.
- [26] C. Wang, Y. Zhang, X. Zhang, Z. Wu, X. Zhu, S. Jin, T. Tang, and M. Tomizuka, "Offline-online learning of deformation model for cable manipulation with graph neural networks," *IEEE Robotics and Automation Letters*, vol. 7, pp. 5544–5551, 2022.
- [27] D. Navarro-Alarcon and Y.-H. Liu, "Fourier-based shape servoing: A new feedback method to actively deform soft objects into desired 2-d image contours," *IEEE Trans. on Robotics*, vol. 34, no. 1, pp. 272–279, 2018.
- [28] H. Wang, B. Yang, J. Wang, X. Liang, W. Chen, and Y. hui Liu, "Adaptive visual servoing of contour features," *IEEE/ASME Transactions on Mechatronics*, vol. 23, pp. 811–822, 2018.
- [29] S. Huo, A. Duan, C. Li, P. Zhou, W. Ma, H. Wang, and D. Navarro-Alarcon, "Keypoint-based planar bimanual shaping of deformable linear objects under environmental constraints with hierarchical action framework," *IEEE Robotics and Automation Letters*, vol. 7, pp. 5222–5229, 2022.
- [30] M. Lippi, P. Poklukar, M. C. Welle, A. Varava, H. Yin, A. Marino, and D. Kragic, "Enabling visual action planning for object manipulation through latent space roadmap," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 57–75, 2023.
- [31] H. Ha and S. Song, "Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding," in *Conference on Robotic Learning (CoRL)*, 2021.
- [32] D. McConachie, A. Dobson, M. Ruan, and D. Berenson, "Manipulating deformable objects by interleaving prediction, planning, and control," *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 957–982, 2020.
- [33] D. Seita, N. Jamali, M. Laskey, A. K. Tanwani, R. Berenstein, P. Baskaran, S. Iba, J. F. Canny, and K. Goldberg, "Deep transfer learning of pick points on fabric for robot bed-making," in *International Symposium of Robotics Research*, 2018.

- [34] G. Salhotra, I. Liu, M. Dominguez-Kuhne, and G. S. Sukhatme, "Learning deformable object manipulation from expert demonstrations," *IEEE Robotics and Automation Letters*, vol. 7, pp. 8775–8782, 2022.
- [35] D. Seita, P. R. Florence, J. Tompson, E. Coumans, V. Sindhwani, K. Goldberg, and A. Zeng, "Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4568–4575, 2021.
- [36] G. Wang, M. Xin, W. Wu, Z. Liu, and H. Wang, "Learning of long-horizon sparse-reward robotic manipulator tasks with base controllers," *IEEE transactions on neural networks and learning systems*, vol. PP, 2020.
- [37] Z. Liu, H. Zhou, B. Chen, S. Zhong, M. Hebert, and D. Zhao, "Safe model-based reinforcement learning with robust cross-entropy method," *ArXiv*, vol. abs/2010.07968, 2020.
- [38] J. Qian, T. Weng, L. Zhang, B. Okorn, and D. Held, "Cloth region segmentation for robust grasp selection," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9553–9560, 2020.
- [39] D. Navarro-Alarcon, H. M. Yip, Z. Wang, Y.-H. Liu, F. Zhong, T. Zhang, and P. Li, "Automatic 3-d manipulation of soft objects by robotic arms with an adaptive deformation model," *IEEE Transactions on Robotics*, vol. 32, no. 2, pp. 429–441, 2016.
- [40] F. Liu, Z. Ling, T. Mu, and H. Su, "State alignment-based imitation learning," *ArXiv*, vol. abs/1911.10947, 2019.
- [41] A. Hajiloo, M. Keshmiri, W.-F. Xie, and T.-T. Wang, "Robust online model predictive control for a constrained image-based visual servoing," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 4, pp. 2242–2250, 2016.
- [42] R. Antonova, P. Shi, H. Yin, Z. Weng, and D. Kragic, "Dynamic environments with deformable objects," in *NeurIPS Datasets and Benchmarks*, 2021.
- [43] D. Pathak, P. Mahmoudieh, G. Luo, P. Agrawal, D. Chen, Y. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell, "Zero-shot visual imitation," *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 2131–21313, 2018.
- [44] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, "Learning predictive representations for deformable objects using contrastive estimation," in *Conference on Robot Learning*, pp. 564–574, PMLR, 2021.
- [45] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, and S. Levine, "Time-contrastive networks: Self-supervised learning from video," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1134–1141, 2017.
- [46] H. Fan, H. Su, and L. J. Guibas, "A point set generation network for 3d object reconstruction from a single image," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2463–2471, 2016.



**Shengzeng Huo** received the B.S. degree in vehicle engineering from the South China University of Technology, Guangzhou, China, in 2019. He is currently pursuing the Ph.D. degree in the Department of Mechanical Engineering from The Hong Kong Polytechnic University, Hong Kong. His research interests include bimanual manipulation, deformable object manipulation, and robot learning.



**Anqing Duan** received his Ph.D. degree in robotics from the Italian Institute of Technology and the University of Genoa in 2021. He is currently Research Associate with the Robotics and Machine Intelligence Laboratory at The Hong Kong Polytechnic University. His research interests include robotics and learning-based control.



**Lijun Han** received the B.S. degree in automation from Tianjin University, Tianjin, China, in 2019. She is currently pursuing the Ph.D. degree in control technology and control engineering with the Department of Automation, Shanghai Jiao Tong University, Shanghai, China.

Her research interests include visual servoing, robot control, and surgical robots.



**Luyin Hu** received his BEng. and MPhil. degree in Mechanical Engineering from The Hong Kong Polytechnic University, KLN, Hong Kong in 2019 and 2022. He is currently a research assistant at the Centre for Transformative Garment Production, Hong Kong. His research interests include servomechanisms, control system design, and physics simulation.



**Hesheng Wang** (Senior Member, IEEE) received the M.Phil. and Ph.D. degrees in automation and computer-aided engineering from The Chinese University of Hong Kong, Hong Kong, in 2004 and 2007, respectively. He is currently a Professor with the Department of Automation, Shanghai Jiao Tong University, Shanghai, China. His current research interests include visual servoing, service robot, computer vision, and autonomous driving. Dr. Wang is an Associate Editor of IEEE Transactions on Automation Science and Engineering, IEEE Robotics and Automation Letters, Robotic Intelligence and Automation and the International Journal of Humanoid Robotics, a Technical Editor of the IEEE/ASME Transactions on Mechatronics, an Editor of Conference Editorial Board (CEB) of IEEE Robotics and Automation Society. He served as an Associate Editor of the IEEE Transactions on Robotics from 2015 to 2019. He was the General Chair of IEEE ROBIO 2022 and IEEE RCAR 2016, and the Program Chair of the IEEE ROBIO 2014 and IEEE/ASME AIM 2019. He will be the General Chair of IEEE/RSJ IROS 2025.



**David Navarro-Alarcon** (Senior Member, IEEE) received the Ph.D. degree in mechanical and automation engineering from The Chinese University of Hong Kong, NT, Hong Kong, in 2014. Since 2017, he has been with The Hong Kong Polytechnic University, where he is currently an Assistant Professor with the Department of Mechanical Engineering. His research interests include robotics, control, and nonlinear dynamics. He currently serves as an Associate Editor of the IEEE TRANSACTIONS ON ROBOTICS.