

GenDOM: Generalizable One-shot Deformable Object Manipulation with Parameter-Aware Policy

So Kuroki¹, Jiaxian Guo¹, Tatsuya Matsushima¹, Takuya Okubo¹, Masato Kobayashi², Yuya Ikeda¹,
Ryosuke Takanami¹, Paul Yoo¹, Yutaka Matsuo¹, Yusuke Iwasawa¹

Abstract—Due to the inherent uncertainty in their deformability during motion, previous methods in deformable object manipulation, such as rope and cloth, often required hundreds of real-world demonstrations to train a manipulation policy for each object, which hinders their applications in our ever-changing world. To address this issue, we introduce *GenDOM*, a framework that allows the manipulation policy to handle different deformable objects with only a single real-world demonstration. To achieve this, we augment the policy by conditioning it on deformable object parameters and training it with a diverse range of simulated deformable objects so that the policy can adjust actions based on different object parameters. At the time of inference, given a new object, *GenDOM* can estimate the deformable object parameters with only a single real-world demonstration by minimizing the disparity between the grid density of point clouds of real-world demonstrations and simulations in a differentiable physics simulator. Empirical validations on both simulated and real-world object manipulation setups clearly show that our method can manipulate different objects with a single demonstration and significantly outperforms the baseline in both environments (a 62% improvement for in-domain ropes and a 15% improvement for out-of-distribution ropes in simulation, as well as a 26% improvement for ropes and a 50% improvement for cloths in the real world), demonstrating the effectiveness of our approach in one-shot deformable object manipulation.

<https://sites.google.com/view/gendom/home>.

I. INTRODUCTION

Deformable object manipulation constitutes a salient aspect of numerous routine tasks, encompassing domestic chores such as folding laundry, culinary endeavors, and the weaving of rope [1], [2]. Despite the promising potential, learning a policy for it is still challenging. It often necessitates numerous real-world demonstrations due to the complex dynamics of deformable objects, which encompass unpredictable deformability, elasticity, and plasticity during motion [1], [2]. Nonetheless, the extensive volume of data required to train these policies, typically gathered through real robots, can often be costly and impractical [3], [4].

Recent research [5]–[9] tried to alleviate this problem via Simulation-to-Real (Sim2Real) transfer, wherein policies are trained within simulation environments with unlimited data and subsequently applied to real-world tasks [10]. Despite these efforts, mitigating Sim2Real gap continues to represent a protracted and complex problem within the robotics field [5]–[9], [11]–[14]. This challenge is exacerbated, especially when dealing with deformable objects situated in un-

predictable dynamic environments. A common approach to these complexities is to train a distinct policy for each object, each possessing unique dynamics [15]–[17]. This approach usually requires hundreds of real-world demonstrations for each object, which are both time-consuming and costly to collect. This requirement for real-world data thus becomes a significant bottleneck for the scalability of robotics tasks, restricting the potential for broader application [16], [17].

In order to bridge the Sim2Real gap in deformable object manipulation, prior approaches such as Real2Sim2Real [18] are proposed, which initially carry out simulation identification to align the simulation with real-world dynamics, a process known as Real-to-Simulation (Real2Sim). Subsequently, they leverage these aligned simulations to assist in policy learning, thereby reducing the need for extensive real-world data [15], [18], [19]. Despite these advances, they still require a considerable volume of real-world demonstrations for effective simulation identification. Moreover, they are typically specialized for certain objects or tasks, which limits their generalizability to a broader spectrum of applications.

In this paper, we propose a framework *GenDOM*, which allows the trained policy to manipulate different objects—even those with out-of-distribution dynamics—using just a single real-world demonstration, *i.e.*, one-shot generalization [20]. Our key idea is to condition the policy with specific deformable object parameters that can affect dynamics. Specifically, we select Young’s modulus and Poisson’s ratio [21]–[23] which are commonly utilized in modelling the deformation [21]–[23]. While these parameters hold significant potential in determining an object’s deformability, their usage in policy training remains largely uncharted territory due to the inherent difficulties in estimating these values. Our research represents an initial effort to integrate these parameters into manipulation policy training, thereby engendering a policy that is not only aware of object-specific parameters but is also attuned to dynamics. As a result, the policy can adjust its action output for various deformable objects. To train such an augmented manipulation policy, we randomly sample different object parameters in the simulation and add them as the input of the policy network. This inclusion allows the policy to encompass a broader range of information on various dynamics parameters, utilizing a single network for deformable object manipulation (Figure 1).

Furthermore, we introduce a novel, efficient gradient-based Real2Sim technique specifically crafted for diverse deformable objects within a unified setup. Leveraging the capabilities of differentiable physics, we propose a gradient-

¹The University of Tokyo, Japan

²Osaka University, Japan

Email: so.kuroki1931@gmail.com




Fig. 1: Task visualization. **Rope reaching:** A robot arm guides a rope end towards a release coordinate, after which the rope moves gravitationally towards goal coordinate g . **Rope casting:** A robot arm performs a momentary action in mid-air, guiding the rope towards coordinate (g_x, g_y) . **Cloth spreading:** A robot arm handles a cloth’s center and acts with acceleration to spread it over a table.

based technique [24]–[30] to estimate the object parameters with only a single real-world demonstration. Our approach seeks to minimize the discrepancy of point clouds between the simulation and real-world environments. Contrary to prior Real2Sim2Real methodologies that utilized estimated deformable object parameters as simulation environmental parameters to learn the policy [15], our method uses these estimated parameters directly as the condition of the policy, enabling it to generalize to objects with differing dynamics.

In summary, our contributions are as follows: (1) We propose a policy augmented with parameters that account for deformable object characteristics (Young’s Modulus and Poisson’s ratio). We show that these parameters are effective in controlling deformable objects such as rope and cloth, and these can be estimated via designed gradient-based optimization given only a single real-world demonstration. (2) Empirical validation on both in-domain and out-of-distribution parameters of the dynamics demonstrate that GenDOM excels in terms of its generalization capability across dynamics, even with only a single real-world demonstration. Specifically, GenDOM provides a 62% improvement for in-domain ropes and a 15% improvement for out-of-distribution ropes in simulation, as well as a 26% improvement for ropes and a 50% improvement for cloths in the real world, highlighting its superior adaptability and potential for broad applicability.

II. RELATED WORK

A. Deformable Object Manipulation

Deformable object manipulation remains a challenging research area due to its inherent complexity. Each object possesses unique dynamics and specific degrees of freedom that demand careful consideration, differing notably from their rigid counterparts[1], [21], [31]. Previous research has succeeded in manipulating various types of deformable objects, such as rope[32], cloth[33], elasto-plastic[34], and liquid[35]. Despite this progress, these methods are plagued with significant limitations. The primary concerns are their substantial reliance on a large number of real-world demonstrations for effective learning and generalization capabilities for different scenarios [1], [2], [16], [17], [21]. Various learning methods have been proposed to operate deformable objects with limited data [36], [37]. In this study, our method offers extensive generalization capabilities and benefits from one-shot learning, significantly reducing the dependency on numerous real-world demonstrations. It also shows comparable performance to policies trained with over 100 real-world

datasets.

B. Sim2Real

Numerous studies focus on filling the gap between simulation and real-world, domain randomization [5]–[8], hierarchical learning [9], [11], and data augmentation [12]. In pursuit of dynamics generalization, [38] integrated both explicit and implicit dynamic parameters into their policy in a rigid-object push manipulation task. In the realm of deformable object manipulation, [39]–[41] leverage differential physics to collect demonstrations on deformable object manipulation and subsequently train a policy through imitation learning to Sim2Real. Nonetheless, these models for deformable object manipulation have limitations of generalization ability for different objects due to their difficulty in learning complex dynamics. To achieve this, we use parameters that account for deformable object characteristics to train our policy.

C. Real2Sim

The demand for collecting realistic deformable object data in simulation scenarios is escalating. Much research is limited due to using specific devices or sensors [23], [42]. While learning-based methods to estimate parameters of deformable objects in simulation have been proposed [43], these methods rely on huge datasets. Thanks to the significant improvements in differentiable simulation [27]–[30], several gradient-based methods have been introduced. For example, [24] proposes differentiable point cloud sampling from mesh states. [25] employed differentiable rendering methods. In this study, we aim to estimate deformable object parameters. To ensure that the parameter captures the physical characteristics of objects, we select Young’s modulus and Poisson’s ratio to estimate, which is commonly used for modeling deformation as a linear and isotropic and homogeneous deformation model [21]–[23].

D. Real2Sim2Real

Some research has been undertaken as Real2Sim2Real. [18] reconstructs high-quality meshes in simulation from real-world point clouds to achieve a robust policy in the real world. [19] bridges the Sim2Real gap by optimizing parameters for tensegrity robots on a simulation platform using real data via gradient descent. [15] estimates parameters with Differential Evolution and executes a Sim2Real transfer with a policy trained by the dataset collected with the estimated parameters. While the preceding research advocates for a Real2Sim2Real pipeline, they focus on dealing with specific

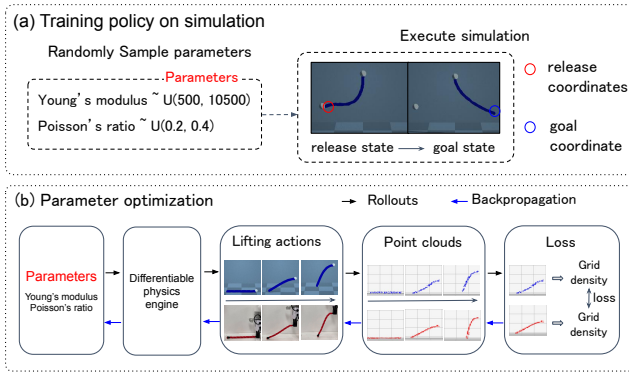


Fig. 2: Overview of our Pipeline: (a) We randomly set parameters in a simulation to collect release and goal coordinates. Using this dataset, we trained a simple MLP policy. (b) For inference, parameters are optimized from real-world objects by comparing point cloud grid densities between simulation and reality through differentiable physics.

dynamics, and few studies pursue the generalization ability for several dynamics.

III. PROPOSED METHOD

To achieve generalizable deformable manipulation, we propose GenDOM, which consists of two primal parts: (1) *training augmented policy with diverse parameters on simulator* and (2) *estimating deformable object parameters during the inference*. The entire GenDOM pipeline is illustrated in Figure 2. We will first provide the task description, followed by a detailed presentation of our framework.

A. Task Description

In this paper, we propose a generalizable framework for manipulating a diverse range of deformable objects. To evaluate the generalizability of our framework, we utilize different deformable objects, specifically ropes and cloths. We consider three tasks showcasing dynamic deformation, including Rope reaching, Rope casting, and Cloth spreading, as depicted in Figure 1. Due to the space limitation, we use Rope reaching as an example to explain the problem setup. In this task, we are given a goal coordinate g and aim to train a manipulation policy that can control a robot arm to handle the end-point of this rope and move it to a releasing coordinate r in mid-air so that the rope can move to the given goal coordinate g under the influence of gravity during its initial swing once released from r . The other types of deformable object manipulation are given in Section IV-A.

B. Training Policy with Diverse Parameters in Simulator

The main challenge for learning a universally applicable deformable object manipulation policy arises due to the variation in physical mechanics across objects. These attributional differences significantly influence the dynamics and the optimal robot control. To address this issue and learn such a generalizable deformable object manipulation policy, our key idea is to augment the policy condition with deformable object parameters. This augmentation enables the policy to be aware of the dynamic variations and adjust its actions accordingly for objects with different dynamics, thus

enhancing its generalization ability for deformable object manipulation. Specifically, we propose to further condition policy on Young's modulus and Poisson's ratio [21]. These two factors serve as critical determinants of a deformable object's behavior, representing stiffness and deformation characteristics, respectively. Both factors are frequently employed in the modeling of deformation [22], [23], and our paper is the first work to condition manipulation policy π on these two factors to achieve generalizable deformable object manipulation. As iteration cost in simulated environments is low, we opt to train our parameter-aware policy on simulation first. For each demonstration, we randomly sample Young's modulus parameter p^y and Poisson's ratio p^l . In this study, we utilize imitation learning for policy formulation as below.

$$\mathcal{L}_\theta = \frac{1}{N} \sum_{i=1}^N (X - \pi_\theta(Y, p_i^y, p_i^l))^2 \quad (1)$$

where θ is the trainable parameter of the manipulation policy π , and N is batch size. Depending on the desired task, by designing the policy's output X and input Y , e.g., the release coordinate is X , and the goal coordinate is Y in Rope reaching task, the proposed framework can adapt to various deformable objects and various tasks.

C. Gradient-based Real2Sim Parameter Optimization

After obtaining our augmented manipulation policy trained across an assortment of deformable object parameters, it remains challenging to apply this policy to real-world tasks due to the unavailability of specific object parameters during real-world deployment. To mitigate this obstacle, we introduce the Real2Sim method, which estimates these parameters using a single real-world demonstration. We begin by transforming a single real-world demonstration of a deformable object into point clouds, taking advantage of the inherent versatility of point clouds in representing diverse shapes and accommodating significant deformations, which makes them an ideal choice for simulating deformable objects. Utilizing differentiable physics [27], [44], we adjust and estimate the deformable object parameters in the simulation to minimize the divergence between the grid density of point clouds from real-world demonstrations and simulations. As depicted in Figure 2 (b), our system is constructed on the framework of PlasticineLab [27], where a differentiable simulator enables gradient-based trajectory optimization through particle dynamics. We perform predefined actions from the same initial states in both the simulated and real-world environments; the robot vertically lifts the object at a constant speed from either the end of the rope or the corner of the cloth. To capture the real-world demonstration, we employ two RGB-D cameras to record the trajectory and preserve its point clouds. At the same time, we also preserve the point clouds of the trajectory in simulation. However, due to the difference in the sequence and number of points between the simulator and real-world environments, directly calculating the loss between point clouds of simulation and real-world demonstration is infeasible.




Fig. 3: (a) Three different type ropes: cotton rope, polyester rope, and polyethylene rope. (b) Two different types of cloth: cotton cloth and rubber cloth. (c) Our experiments setup for Real2Sim: two Intel Realsense Depth Camera D435i and xArm 7.

To bypass this hurdle, we represent the point clouds from both simulation and real-world demonstration into a 3-dimensional grid density space (discretized into $64 \times 64 \times 64$ in our study) to synchronize them, regardless of their differing quantities and orders. Explicitly, each cloud’s position is computed within the grid, assigning a weight to each grid cell based on its fractional position f_x for superior density matching. Particularly, these weight assignments are attenuated based on cloud’s distance to the cell’s center, where the closer to the center, the higher the weight, thereby promoting a smooth distribution of the particle’s mass across neighboring grid cells. This smooth distribution of the particle’s mass enables successful alignment between the real-world and simulated scenarios, thereby utilizing the prior knowledge in the differentiable physics simulation to ensure an accurate loss that converges even with a single demonstration. By minimizing the L1 loss between the grid densities of the simulator and the real-world environments, with Young’s modulus and Poisson’s ratio serving as trainable variables, we can indirectly estimate these parameters through the differential simulator. We then condition these parameters directly on the pre-trained manipulation policy π in Section III-B to execute the manipulation task.

IV. EXPERIMENTS

In this section, we experimentally evaluate our method in both simulation and real world. Our experiments seek to address the following questions:

- Can our gradient-based optimization using a designed loss function estimate the physical parameters of ropes (Young’s Modulus and Poisson’s ratio) given a limited demonstration? Is it better than a learning-based method? (IV-C and IV-D)
- Does the estimated physical parameters help our parameter-aware policy to generalize diverse setups (in-domain, out-of-domain, and real-world deployment)? (IV-E, IV-F, and IV-G)
- Can our gradient-based optimization method and parameter-aware policy be deployed for cloth? (IV-H)

A. Tasks

As Figure 3 (a) and (b) show, we use three different types of ropes: cotton, polyester and polyethylene ropes, and two different cloths: cotton and rubber cloths to evaluate the effectiveness of our method on different deformable objects. Furthermore, to demonstrate our framework’s capability to handle deformable object manipulation with varying dynamics, we evaluated it on the following tasks:

Rope reaching: We randomly sample release coordinate r with the parameters. Given the sampled rope parameters, we record the maximum reach distance of the bottom tip of the rope as the goal coordinate g . Utilizing g as hindsight for Y , the policy is trained to predict r for X .

Rope casting: Given a goal coordinate (g_x, g_y) , a manipulation policy is trained to operate a robot arm to handle the top end of this rope and take one momentary action a in mid-air. The rope is expected to move towards (g_x, g_y) during its initial swing. We randomly sampled an action a that moves instantaneously and horizontally relative to the ground with the parameters. Given the sampled rope parameters, we record the highest reach coordinate of the bottom tip of the rope as the goal coordinate (g_x, g_y) . Utilizing (g_x, g_y) as hindsight for Y , we train our policy π to predict a for X .

Cloth spreading: The objective is to spread a cloth over a table. A manipulation policy is trained to control a robot arm to handle the center of the cloth and take actions with acceleration in mid-air. The robot moves 16cm in 16 steps. For each step i , the distance traveled for each step is defined as: $1 - (4.5 - 0.5i) \times A$, for $i = 1, \dots, 8$ and $1 + (0.5i - 4.0) \times A$, for $i = 9, \dots, 16$. If A is large, the robot’s movement range shifts from small to large, resulting in a greater acceleration. We randomly sampled A and cloth parameters, and recorded the minimum value of A that resulted in a spread of at least 75% of the cloth area being smooth and wrinkle-free. The policy is trained to predict A as X with the parameters. Y is unset for this task.

B. Environments Setup

Deformable Parameter Optimization: To bridge the gap between the behavior of deformable objects in the real world and in simulation, we use gradient-based optimization to estimate both Young’s modulus and Poisson’s ratio. Our hardware configuration, depicted in Figure 3 (c), incorporates two Intel Realsense Depth Camera D435i to obtain point clouds and xArm 7 to manipulate rope and cloth. Two camera poses are calibrated to the xArm-world coordinate. We take a demonstration for each object by executing the predefined lifting actions. Both real-world and simulation point clouds are transferred into a 3-dimensional grid density of size (64, 64, 64). We use PointNet++ [51] as the baseline. We adopt PointNet++ for its proficiency in processing point clouds, known for its versatility in representing diverse shapes and handling deformations. We randomly assign parameters to the simulation and execute the predefined actions to obtain the last state of either the rope or cloth within the simulation. Using the 998 sets of parameters and the last state, PointNet++ is trained to output the parameters from the last state. The results analysis are given in Section IV-C and IV-D.

Policy Training: We train our policy on a synthesized dataset that covers diverse physical parameters for deformable objects. Specifically, we randomize Young’s modulus and Poisson’s ratio within the ranges [1500, 8200] and [0.34, 0.36] for the simulation test, and we randomize them in [500, 10500] and [0.2, 0.4] for the real world test. We use the recent Real2Sim2Real method (R2S2R) [15]

TABLE I: Summary of our methods and baselines in deformable object manipulation

| Category & Method | | Policy | | Adaptation | | | |
|-------------------|-----------------------|-----------------|----------------|------------|-------|---------------|---------|
| Category | Method | Parameter-aware | Data | Type | Speed | #demo in real | S2R gap |
| Real | Many [4], [45], [46] | ✗ | Real (single) | None | N/A | Many | None |
| Real | Many [47], [48] | ✗ | Real (diverse) | None | N/A | Many | None |
| Sim2Real | Many [10], [49], [50] | ✗ | Sim (diverse) | None | N/A | 0 ~ | High |
| Real2SimReal | R2S2R [15] | ✗ | Sim (Single) | Train | Low | 20 ~ | Low |
| Real2SimReal | GenDOM (Ours) | ✓ | Sim (diverse) | Condition | High | 1 ~ | Low |

TABLE II: Comparison between our gradient-based method and the baseline (PointNet++) regarding the difference between estimated and ground-truth parameters and the Chamfer distance within a Sim2Sim setup.

| Method | Young’s modulus Mean of diff | Poisson’s ratio Mean of diff | Chamfer Distance |
|------------------|---------------------------------|---------------------------------|------------------------|
| Ours (all steps) | 11.69 ± 8.28 | 0.0032 ± 0.0074 | 0.0004 ± 0.0005 |
| Ours (last step) | 13.74 ± 6.75 | 0.0067 ± 0.0098 | 0.0004 ± 0.0004 |
| Pointnet++ | 173.51 ± 200.14 | 0.029 ± 0.024 | 0.0025 ± 0.0016 |

as the baseline where it estimates parameters and executes a sim2real transfer with a policy trained by the dataset collected with estimated parameters. We follow R2S2R’s setting and prepare two datasets to train it in simulation and the real world, respectively. Two datasets are (1) A dataset covering all three rope parameters, and (2) a dataset only covering the red rope parameter. **We collect 2,271 demonstrations for simulation evaluation and 40,000 for real-world evaluation for all datasets to ensure a fair comparison between the proposed method and baselines.** Furthermore, in the comparison in the real world, we prepare a policy trained on only the real-world dataset (π_{RD}): (1) 300 demonstrations for three ropes together and (2) 100 demonstrations with only the red rope’s parameters. Initializing policy weights with those trained in simulation, we also developed an adapted policy by fine-tuning it using 300 real-world demonstrations. The results analysis are given in Section IV-F. We summarize the baseline methods and the relevant information in Table I.

C. Sim2Sim Performance

To verify the performance of the proposed method in estimating the ground truth Young’s modulus and Poisson’s ratio parameters, we conducted evaluations on the simulation. Table II compares our gradient-based and learning-based methods (Pointnet++) on the accuracy of sim2sim parameter optimization for rope. We sampled five sets of ground-truth parameters and reported the mean and standard deviation. In addition, we calculate the Chamfer distance between point clouds of the last state with ground-truth parameters and estimated parameters. For the gradient-based method, we sample parameters as an initial state for each set of ground-truth parameters and proceed with optimization. PointNet++ estimates parameters from each last state after executing the predefined actions. We also test gradient-based optimization only with the last state loss for a fair comparison. From the results, the gradient-based method yields a **significantly smaller** mean and standard deviation than PointNet++. When comparing the gradient-based method, the gradient-based with the last step provides nearly identical performance with

TABLE III: Comparison of Real2Sim rope parameter optimizations on the optimized parameters and Chamfer distance.

| Rope | Method | Young’s modulus | Poisson’s ratio | Chamfer Distance |
|--------------|------------|------------------------|---------------------|---------------------|
| Cotton | Ours | 1779.38 ± 7.24 | 0.35 ± 0.04 | 0.034 ± 0.00 |
| | Pointnet++ | 509.83 | 0.30 | 0.092 |
| Polyester | Ours | 3276.12 ± 3.71 | 0.346 ± 0.03 | 0.060 ± 0.00 |
| | Pointnet++ | 4660.12 | 0.30 | 0.067 |
| Polyethylene | Ours | 8000.31 ± 48.04 | 0.36 ± 0.00 | 0.053 ± 0.00 |
| | Pointnet++ | 7305.96 | 0.30 | 0.057 |

the gradient-based with all steps. We utilize the gradient-based with all step loss in subsequent experiments since it shows slightly more accuracy.

D. Real2Sim Performance

To assess the efficacy of our proposed parameter estimation technique in the real world, we conducted an evaluation by estimating parameters from a single rope demonstration. Table III shows the estimated parameters and Chamfer distance between the last state of the real world and the simulation with estimated parameters. We utilize Chamfer distance since real-world ground truths aren’t accessible. The gradient-based method optimizes parameters from five randomly sampled parameters as the initial state for each object. Pointnet++ uses point clouds of the last state to estimate parameters. For all ropes, Chamfer distance of the gradient-based method **shows a smaller value** than PointNet++. The estimated Young’s modulus and Poisson’s ratio are [1779.38, 0.35], [3276.12, 0.346], and [8000.31, 0.36] for cotton, polyester, and polyethylene ropes, respectively. Our method can directly estimate parameters **without time-consuming data collection**. A supplemental video shows trajectory comparisons between real and simulated objects with estimated parameters.

E. Ablation Study

To discern the contribution of each parameter to the tasks, we evaluate the accuracy by modulating the inputs to the network for rope reaching and rope casting in simulation. This assessment is conducted with four distinct input configurations: both Young’s modulus and Poisson’s ratio, only Young’s modulus, only Poisson’s ratio, and neither of these parameters. The policy for each task is trained with 40,000 and 20,000 demonstrations, respectively. The parameters are in the range of [500, 10,500] for Young’s modulus and [0.2, 0.4] for Poisson’s ratio. We randomly sample 30 new goal coordinates and parameters. Table IV shows the best performance for the policy trained on both parameters, suggesting that the **two parameters in conjunction are informative** for the tasks.

TABLE IV: Ablation study results. For training, we utilize four variations of inputs. We report the mean and standard deviation on 30 new goal coordinates.

| | Both parameters | Young’s modulus | Poisson’s ratio | No parameters |
|---------------|-----------------------------|----------------------|----------------------|----------------------|
| Rope reaching | 0.024 ± 0.021 | 0.034 ± 0.040 | 0.062 ± 0.058 | 0.062 ± 0.059 |
| Rope casting | 0.065 ± 0.031 | 0.067 ± 0.033 | 0.077 ± 0.041 | 0.076 ± 0.042 |

TABLE V: Comparison of generalization for ID and OOD rope parameters (Calculated on new 40 goal coordinates).

| Method | ID | OOD |
|---------------------|-------------------------------------|-------------------------------------|
| Ours | 0.031 \pm 0.050 | 0.039 \pm 0.042 |
| R2S2R (Three ropes) | 0.082 \pm 0.081 | 0.096 \pm 0.104 |
| R2S2R (Cotton rope) | 0.047 \pm 0.033 | 0.046 \pm 0.044 |

F. Generalization Test

To test generalizability to various ropes, we evaluate the policy on in-domain (ID) and out-of-distribution (OOD) parameters in simulation. For the ID test, we randomly sample 40 new goal coordinates with parameters of three ropes. For the OOD test, we randomly sample 40 new goal coordinates with parameters range [500, 1500] and [0.3, 0.33] for Young’s modulus and Poisson’s ratio, respectively. Table V summarizes our findings. (1) Our method performs 62% better than R2S2R (trained on three ropes) for ID. (2) Our method performs well even in the OOD setup. Specifically, we compare our method with R2S2R (trained on the red rope) for OOD. Even though the baseline is trained on the red rope’s parameter, which is close to the OOD parameters’ range, our method performs 15% better than R2S2R (trained on the red rope). This suggests our parameter-conditioned policy has a **stronger generalization ability** against variable deformable object dynamics.

G. Real Robot Experiments

Finally, we deploy our parameter-aware policy in the real world using the estimated parameters for each rope. We set a goal conditioned from [27.5, 30.0, 32.5, 35, 37.5, 40.0, 42.5] (cm), as represented by the range depicted in Fig 1, and tested the policy three times for each goal coordinate. We sum up the difference for each goal point and calculate its mean and standard deviation on three times trials, with results in Table VI. For red rope, our method without adaptation performs 10% better than π_{RD} (trained on the red rope). This shows that our method, only using a single real-world demonstration, achieves better performance than the real-world policy trained with 100 demonstrations. With adaptation, our method (sim+real) performs 32% better than π_{RD} and 44% better than R2S2R (trained on the red rope).

H. Extension to Cloth

To demonstrate the versatility of our proposed method for other deformable objects, we conducted parameter optimization and real-world deployment tests using cloth.

Parameter optimization We use the same predefined action and setup for the rope parameter optimization. The robot grasps one of the four corners of the cloth on the desk

TABLE VI: Results of real-world rope reaching, summarizing the differences for goal coordinates with the mean and standard deviation from three trials.

| Rope reaching | Ours | Ours | π_{RD} | π_{RD} | R2S2R | R2S2R |
|---------------|-----------------------------|-----------------------------|----------------------|----------------------|----------------------|----------------------|
| | sim+real | adaptation | Three ropes | Cotton rope | Three ropes | Cotton rope |
| Cotton | 0.112 0.007 | 0.148 ± 0.012 | 0.426 ± 0.017 | 0.164 ± 0.012 | 0.318 ± 0.004 | 0.200 ± 0.021 |
| Polyester | 0.076 ± 0.025 | 0.106 ± 0.005 | 0.273 ± 0.056 | 0.512 ± 0.039 | 0.343 ± 0.017 | 0.464 ± 0.038 |
| Polyethylene | 0.193 ± 0.009 | 0.588 ± 0.070 | 0.297 ± 0.113 | 0.401 ± 0.060 | 0.364 ± 0.061 | 0.473 ± 0.024 |

TABLE VII: Real-world Cloth spread results with mean and standard deviation on three times trials. Additionally, we provided the proportion in relation to A without parameters as acceleration ratio.

| Cloth spreading | Both parameters | | No parameters | |
|-----------------|-----------------------------------|--------------------|-----------------|--------------------|
| | Extended ratio | Acceleration ratio | Extended ratio | Acceleration ratio |
| Cotton | 0.78 \pm 0.08 | 1.70 | 0.52 \pm 0.03 | 1.0 |
| Rubber | 0.98 \pm 0.01 | 0.26 | 0.97 \pm 0.01 | 1.0 |

and lifts it vertically at a constant speed. The gradient-based method optimizes parameters from five randomly sampled parameters as the initial state for each cloth. The estimated Young’s modulus and Poisson’s ratio are [584.21, 0.25] and [6828.33, 0.27] for cotton and rubber cloth, respectively.

Real-world deployment To evaluate the contribution of the parameters, we compare the accuracy by modulating the inputs to the network for cloth spreading task in the real world. We train our policy using an 800 dataset with Young’s modulus and Poisson’s ratio within the ranges [500, 10500] and [0.2, 0.4]. In addition, we fine-tuned the policy using two demonstrations for two cloths. This assessment is conducted with the policy trained with both Young’s modulus and Poisson’s ratio and one by neither of these parameters. We conducted five trials for each type of cloth and evaluated how much it could spread in its expanded state. In addition, based on the A by the policy without parameters, we normalized other A as the acceleration ratio. The results are presented in Table VII. Without parameters, the same acceleration is output for both cotton and rubber cloths. When using the policy with parameters, it enhances the spreading of cotton by 50%, applying higher acceleration. In contrast, for rubber, the policy achieves spreading with an acceleration reduced by 64%. To properly execute the task, it can adjust the speed, and it also suppresses risks such as malfunctions.

V. CONCLUSION

In this paper, we present a framework designed for one-shot manipulation of deformable objects, specifically ropes and cloths. By combining a pre-trained, parameter-conditioned policy with a gradient-based parameter optimization technique, we show that our method outperforms the baselines in both in-domain and out-of-distribution settings across simulation and real-world rope manipulation tasks, as well as in real-world deployment for both ropes and cloths.

ACKNOWLEDGEMENT

This research was supported in part by KIOXIA Corporation.

REFERENCES

- [1] H. Yin, A. Varava, and D. Kragic, "Modeling, learning, perception, and control methods for deformable object manipulation," *Science Robotics*, vol. 6, no. 54, eabd8803, 2021. DOI: [10.1126/scirobotics.abd8803](https://doi.org/10.1126/scirobotics.abd8803).
- [2] V. E. Arriola-Rios, P. Guler, F. Ficuciello, D. Kragic, B. Siciliano, and J. L. Wyatt, "Modeling of deformable objects for robotic manipulation: A tutorial and review," *Frontiers in Robotics and AI*, vol. 7, 2020, ISSN: 2296-9144. DOI: [10.3389/frobt.2020.00082](https://doi.org/10.3389/frobt.2020.00082). [Online]. Available: <https://www.frontiersin.org/articles/10.3389/frobt.2020.00082>.
- [3] J. Wang, Y. Song, T. Leung, *et al.*, "Learning fine-grained image similarity with deep ranking," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1386–1393.
- [4] A. Nair, D. Chen, P. Agrawal, *et al.*, "Combining self-supervised learning and imitation for vision-based rope manipulation," in *2017 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2017, pp. 2146–2153.
- [5] B. Mehta, M. Diaz, F. Golemo, C. J. Pal, and L. Paull, "Active domain randomization," in *Conference on Robot Learning*, PMLR, 2020, pp. 1162–1176.
- [6] C. Rabinovitz, N. Grupen, and A. Tamar, "Unsupervised feature learning for manipulation with contrastive domain randomization," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 10 153–10 159.
- [7] G. Leibovich, G. Jacob, S. Endrawis, G. Novik, and A. Tamar, "Validate on sim, detect on real-model selection for domain randomization," in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 7528–7535.
- [8] S. James, P. Wohlhart, M. Kalakrishnan, *et al.*, "Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 627–12 637.
- [9] O. Nachum, M. Ahn, H. Ponte, S. Gu, and V. Kumar, "Multi-agent manipulation via locomotion using hierarchical sim2real," *arXiv preprint arXiv:1908.05224*, 2019.
- [10] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*, PMLR, 2018, pp. 734–743.
- [11] S. W. Abeyruwan, L. Graesser, D. B. D' Ambrosio, *et al.*, "I-sim2real: Reinforcement learning of robotic policies in tight human-robot interaction loops," in *Conference on Robot Learning*, PMLR, 2023, pp. 212–224.
- [12] A. Pashevich, R. Strudel, I. Kalevatykh, I. Laptev, and C. Schmid, "Learning to augment synthetic images for sim2real policy transfer," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2019, pp. 2651–2657.
- [13] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *Advances in Artificial Life: Third European Conference on Artificial Life Granada, Spain, June 4–6, 1995 Proceedings 3*, Springer, 1995, pp. 704–720.
- [14] J. Mahler, J. Liang, S. Niyaz, *et al.*, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *arXiv preprint arXiv:1703.09312*, 2017.
- [15] V. Lim, H. Huang, L. Y. Chen, *et al.*, "Real2sim2real: Self-supervised learning of physical single-step dynamic actions for planar robot casting," in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 8282–8289.
- [16] R. Lee, M. Hamaya, T. Murooka, Y. Ijiri, and P. Corke, "Sample-efficient learning of deformable linear object manipulation in the real world through self-supervision," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 573–580, 2022. DOI: [10.1109/LRA.2021.3130377](https://doi.org/10.1109/LRA.2021.3130377).
- [17] A. Wang, T. Kurutach, K. Liu, P. Abbeel, and A. Tamar, "Learning robotic manipulation through visual planning and acting," *CoRR*, vol. abs/1905.04411, 2019. arXiv: [1905.04411](https://arxiv.org/abs/1905.04411). [Online]. Available: <http://arxiv.org/abs/1905.04411>.
- [18] L. Wang, R. Guo, Q. Vuong, Y. Qin, H. Su, and H. Christensen, "A real2sim2real method for robust object grasping with neural surface reconstruction," *arXiv preprint arXiv:2210.02685*, 2022.
- [19] K. Wang, W. R. Johnson III, S. Lu, *et al.*, "Real2sim2real transfer for control of cable-driven robots via a differentiable physics engine," *arXiv preprint arXiv:2209.06261*, 2022.
- [20] D. Rezende, I. Danihelka, K. Gregor, D. Wierstra, *et al.*, "One-shot generalization in deep generative models," in *International conference on machine learning*, PMLR, 2016, pp. 1521–1529.
- [21] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, and Y. Mezouar, "Robotic manipulation and sensing of deformable objects in domestic and industrial applications: A survey," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 688–716, 2018.
- [22] A. Sengupta, R. Lagneau, A. Krupa, E. Marchand, and M. Marchal, "Simultaneous tracking and elasticity parameter estimation of deformable objects," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 10 038–10 044.
- [23] B. Frank, R. Schmedding, C. Stachniss, M. Teschner, and W. Burgard, "Learning the elasticity parameters of deformable objects with a manipulation robot," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1877–1883. DOI: [10.1109/IROS.2010.5653949](https://doi.org/10.1109/IROS.2010.5653949).
- [24] P. Sundaresan, R. Antonova, and J. Bohgl, "Diffcloud: Real-to-sim from point clouds with differentiable simulation and rendering of deformable objects," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2022, pp. 10 828–10 835.
- [25] P. Ma, T. Du, J. B. Tenenbaum, W. Matusik, and C. Gan, "Risp: Rendering-invariant state predictor with differentiable simulation and rendering for cross-domain parameter estimation," *arXiv preprint arXiv:2205.05678*, 2022.
- [26] S. Le Cleac'h, H.-X. Yu, M. Guo, *et al.*, "Differentiable physics simulation of dynamics-augmented neural objects," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2780–2787, 2023.
- [27] Z. Huang, Y. Hu, T. Du, *et al.*, "Plasticinelab: A soft-body manipulation benchmark with differentiable physics," *The International Conference on Learning Representations*, 2021.
- [28] K. M. Jatavallabhula, M. Macklin, F. Golemo, *et al.*, "Gradsim: Differentiable simulation for system identification and visuomotor control," *arXiv preprint arXiv:2104.02646*, 2021.
- [29] S. Chen, Y. Xu, C. Yu, *et al.*, "Daxbench: Benchmarking deformable object manipulation with differentiable physics," in *The Eleventh International Conference on Learning Representations*.
- [30] E. Heiden, M. Macklin, Y. Narang, D. Fox, A. Garg, and F. Ramos, "Disect: A differentiable simulation engine for autonomous robotic cutting," *arXiv preprint arXiv:2105.12244*, 2021.
- [31] D. Seita, P. Florence, J. Tompson, *et al.*, "Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4568–4575. DOI: [10.1109/ICRA48506.2021.9561391](https://doi.org/10.1109/ICRA48506.2021.9561391).
- [32] Y. Deng, C. Xia, X. Wang, and L. Chen, "Deep reinforcement learning based on local gnn for goal-conditioned deformable object rearranging," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 1131–1138. DOI: [10.1109/IROS47612.2022.9981669](https://doi.org/10.1109/IROS47612.2022.9981669).
- [33] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to manipulate deformable objects without demonstrations," *ArXiv*, vol. abs/1910.13439, 2019.
- [34] X. Lin, C. Qi, Y. Zhang, *et al.*, "Planning with spatial-temporal abstraction from point clouds for deformable object manipulation," in *Conference on Robot Learning*, 2022.
- [35] G. Narasimhan, K. Zhang, B. Eisner, X. Lin, and D. Held, "Self-supervised transparent liquid segmentation for robotic pouring," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 4555–4561. DOI: [10.1109/ICRA46639.2022.9812000](https://doi.org/10.1109/ICRA46639.2022.9812000).
- [36] P. Sundaresan, J. Grannen, B. Thananjeyan, *et al.*, "Learning rope manipulation policies using dense object descriptors trained on synthetic depth data," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 9411–9418.
- [37] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," in *Conference on robot learning*, PMLR, 2017, pp. 357–368.
- [38] T. Murooka, M. Hamaya, F. von Drigalski, K. Tanaka, and Y. Ijiri, "Exi-net: Explicitly/implicitly conditioned network for multiple environment sim-to-real transfer," in *Conference on Robot Learning*, PMLR, 2021, pp. 1221–1230.
- [39] S. Kuroki, T. Matsushima, J. Arima, *et al.*, "Collective intelligence for 2d push manipulations with mobile robots," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2820–2827, 2023.

- [40] C. Qi, X. Lin, and D. Held, “Learning closed-loop dough manipulation using a differentiable reset module,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9857–9864, 2022.
- [41] X. Lin, C. Qi, Y. Zhang, *et al.*, “Planning with spatial-temporal abstraction from point clouds for deformable object manipulation,” in *6th Annual Conference on Robot Learning*, 2022.
- [42] F. Liu, E. Su, J. Lu, M. Li, and M. C. Yip, “Robotic manipulation of deformable rope-like objects using differentiable compliant position-based dynamics,” *IEEE Robotics and Automation Letters*, vol. 8, no. 7, pp. 3964–3971, 2023. DOI: [10.1109/LRA.2023.3264766](https://doi.org/10.1109/LRA.2023.3264766).
- [43] S. Yang, J. Liang, and M. C. Lin, “Learning-based cloth material recovery from video,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 4383–4393.
- [44] Y. Hu, L. Anderson, T.-M. Li, *et al.*, “DiffTaichi: Differentiable programming for physical simulation,” *The International Conference on Learning Representations*, 2020.
- [45] R. Lee, M. Hamaya, T. Murooka, Y. Ijiri, and P. Corke, “Sample-efficient learning of deformable linear object manipulation in the real world through self-supervision,” *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 573–580, 2021.
- [46] K. Shivakumar, V. Viswanath, A. Gu, *et al.*, “Sgtm 2.0: Autonomously untangling long cables using interactive perception,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2023, pp. 5837–5843.
- [47] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, “Learning predictive representations for deformable objects using contrastive estimation,” in *Conference on Robot Learning*, PMLR, 2021, pp. 564–574.
- [48] M. Yan, Y. Zhu, N. Jin, and J. Bohg, “Self-supervised learning of state estimation for manipulating deformable linear objects,” *IEEE robotics and automation letters*, vol. 5, no. 2, pp. 2372–2379, 2020.
- [49] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, “Learning to manipulate deformable objects without demonstrations,” *arXiv preprint arXiv:1910.13439*, 2019.
- [50] X. Lin, Y. Wang, J. Olkin, and D. Held, “Softgym: Benchmarking deep reinforcement learning for deformable object manipulation,” in *Conference on Robot Learning*, PMLR, 2021, pp. 432–448.
- [51] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” *Advances in neural information processing systems*, vol. 30, 2017.