

Learning Tool Morphology for Contact-Rich Manipulation Tasks with Differentiable Simulation

Mengxi Li, Rika Antonova, Dorsa Sadigh and Jeannette Bohg

Abstract—When humans perform contact-rich manipulation tasks, customized tools are often necessary to simplify the task. For instance, we use various utensils for handling food, such as knives, forks and spoons. Similarly, robots may benefit from specialized tools that enable them to more easily complete a variety of tasks. We present an end-to-end framework to automatically learn tool morphology for contact-rich manipulation tasks by leveraging differentiable physics simulators. Previous work relied on manually constructed priors requiring detailed specification of a 3D object model, grasp pose and task description to facilitate the search or optimization process. Our approach only requires defining the objective with respect to task performance and enables learning a robust morphology through randomizing variations of the task. We make this optimization tractable by casting it as a continual learning problem. We demonstrate the effectiveness of our method for designing new tools in several scenarios, such as winding ropes, flipping a box and pushing peas onto a scoop in simulation. Additionally, experiments with real robots show that the tool shapes discovered by our method help them succeed in these scenarios.

I. INTRODUCTION

Humans are distinct from other species in that tool use is a defining and universal characteristic [1]. This suggests that in the pursuit of equipping robots with human-like dexterity, tools may play an important role. Robots are already using various tools in a range of contact-rich manipulation tasks. For example, to make knots, robots can use a tri-needle to maintain the loop [2]. For cooking, robots use spatulas to flip pancakes [3] and skewers to pick food for assistive feeding [4], [5]. While tools greatly influence how robots interact with the environment in these contact-rich tasks, most works focus on learning how to use existing tools. Little attention is paid to optimal tool design. Rather than forcing robots to use pre-defined tools, we aim to intelligently adapt the tools to the tasks, thus helping robots become more effective.

In this work, we aim to develop a general framework for learning robust tool morphology for contact-rich tasks. Relevant to our setting, works on aerodynamic design [6] and vehicle component design also tackle the problem of finding an optimal shape. However, these prior works do not focus on tool design for contact-rich tasks and thus do not require complex contact modeling that is necessary in our scenarios. Another relevant line of research investigates robot gripper design. Some of these works aim to discover

The authors are with the Department of Computer Science at Stanford University (contact: mengxili@stanford.edu, rika.antonova@stanford.edu, dorsa@cs.stanford.edu, bohg@stanford.edu). This project was supported in part by FANUC. Rika Antonova is supported by the National Science Foundation grant No.2030859 to the Computing Research Association for the CIFellows Project.

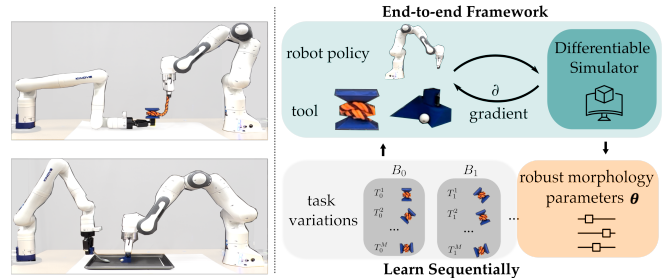


Fig. 1: We build an end-to-end framework for learning tool morphologies suitable for contact-rich tasks. Our goal is to learn a tool morphology for a given scenario that is robust to task variations. We achieve this with a method based on continual learning that trains on a sequence of task variations.

gripper designs for grasping a wide range of objects [7] or executing re-orientation primitives [8]. However, they do not provide a way to optimize the gripper shape for more complex tasks. Recent work [9] proposes a new approach to optimize morphology for a given task objective. However, as we will show in our experiments, it does not necessarily generalize to task variations, such as different initial object poses. We aim to automatically design tools for a given task objective, such that these tools are robust to task variations.

To obtain an optimal shape that minimizes a task-relevant objective, previous work on gripper design employed heuristics for guiding their search [10], or needed Monte Carlo estimation of the gradients to attempt gradient-based optimization [11]. When a robot uses a tool in a contact-rich manipulation task and knows the corresponding dynamics model, it is possible to directly get the exact gradients by differentiation; this enables better numerical stability and faster convergence [12]. To this end, we leverage recent advances in differentiable physical simulators [13], [14], [15], [9] and build an end-to-end framework for learning tool morphology suitable for contact-rich tasks. To design an effective tool for a specific scenario while maintaining the ability to handle task variations, we optimize tool morphology over a distribution of task variations. This presents two challenges. First, the training process for learning to handle the entire distribution can be computationally expensive. Second, due to the complex dynamics of contact-rich tasks, the underlying optimization landscape is highly non-linear and rugged, which makes it more difficult for optimization to converge to a good solution [16], [17]. To tackle these challenges, we propose an approach based on continual learning that samples task variations and conducts optimization in a sequential manner. Our insight is that we can re-interpret continual learning as a robust optimization framework for problems with challenging loss landscapes. Compared to

prior work, we broaden morphology optimization to tool design, thereby opening up a path for tackling a wider array of tasks and objects – including deformables. Furthermore, we demonstrate that the tools obtained with our approach are effective at completing the given tasks with real robots.

II. RELATED WORK

Morphology Optimization in Manipulation. In this work, we explore the promise of fully differentiable end-to-end optimization of tool morphology with the help of differentiable physics. *Gripper design* is one related problem. Some classic works provide guidelines for manually designing grippers guided by practical insights [18]. Recent works aim to learn an optimal design within a given design space. Evolutionary strategies are employed in [19] to optimize both the robot morphology and the controller. In [20], the authors optimize gripper quality using a set of manually designed metrics. Another popular paradigm for gripper customization is imprint-based methods [21], [22]. However, these methods only produce customized finger geometry, where an object 3D model, grasp pose, and task description are specified by the user. A survey in [23] reviews other examples of automated finger design. Further recent examples include a gradient-free method [24], and a gradient-based shape generation method with non-differentiable simulation for training [25]. Most closely related to our work is DiffHand [9], where a differentiable simulator is developed to enable co-design of robot morphology and control by optimizing task-specific objectives. However, this work only considers specific object initial states, and thus does not learn a morphology that would generalize over various initial object states.

Tool design is related to gripper design but unlike grippers and fingers, tools are typically not rigidly attached to robots. This opens future possibilities to study a) re-grasping of tools to improve versatility and dexterity of manipulation robots, and b) the use of multiple tools at the same time (e.g. for dual-arm manipulation). While we do not focus on these aspects in our current work, our formulation facilitates exploring them in the future. Works that consider selecting tools and optimizing policies for tool use are common in robotics literature, e.g. [26], [27], where [27] employs a differentiable simulator. However, the vast majority of these works do not optimize tool shape. Related work for tool *morphology optimization* includes MacGyvering [28]. However, it assumes access to a ‘reference tool’, and aims to construct a tool from a set of available parts. Instead, we consider the problem of evolving tool morphology from an initial shape without assuming prior knowledge about the optimal tool shape. Furthermore, we use end-to-end differentiable simulation, and leverage differentiability at all levels of the optimization pipeline.

Differentiable Simulation. Instead of using heuristics or search algorithms for optimizing morphology, we leverage differentiable simulation for directly obtaining analytical gradients with respect to the final task objective. Several differentiable simulation frameworks have been developed recently, e.g. [13], [14], [15], [9]. In this work, we adopt the

differentiable simulator and morphology representation from DiffHand [9]. However, our framework is agnostic to the choice of differentiable simulator. Differentiable simulation already enabled some impressive results: e.g. system identification (real-to-sim) and control optimization for cutting [29]; solving a dynamic ball-in-cup task in 4 minutes on a real robot [30]. These works consider advanced phenomena, such as modeling deformation. However, they do not address the aspect that could be especially challenging for computing gradients – making and breaking contacts in contact-rich tasks. Contact-rich scenarios bring a new level of complexity. They yield sharp changes in the loss landscapes and could make gradient-based optimization difficult [16], [17]. Our focus on tool morphology requires us to address this challenge because tools interact with objects in the scene. Hence, we propose a method that not only employs differentiability, but also leverages a continual learning formulation of the problem to tackle the optimization challenges that arise in contact-rich scenarios.

Continual Learning. Continual Learning considers the problem of learning to solve a sequence of tasks (or task variations), with the objective to perform well on the current task without forgetting what has been learned from previous tasks [31], [25], [32]. The tasks are usually provided as a stream, i.e. data from previous tasks is usually not retained due to memory limitations. There are different categories of continual learning methods. For example, replay methods [33], [34] usually store a fixed number of samples in the replay buffer and use these to construct a distillation loss between previous and current model predictions. This distillation loss encourages the model not to forget what it previously learned. Parameter isolation methods [35], [36] divide model parameters into different subsets and fix the subset of parameters learned with previous tasks when learning a new task. In this way, they prevent the model from forgetting previous tasks and also improve training stability. We build upon these works and leverage the insight that continual learning could be re-interpreted as a framework for robust optimization suitable for problems with non-smooth dynamics that yield challenging loss landscapes, such as optimizing the morphology of manipulation tools.

III. END-TO-END FRAMEWORK WITH DIFFERENTIABLE SIMULATION

In this section, we describe how we learn tool morphology in an end-to-end manner. In Sec. III-A, we first formalize the task where the tool will be used. With the deformation based morphology parameterization described in Sec. III-B, we obtain a low-dimensional design space. Finally, we show the end-to-end morphology parameter learning pipeline for a single task variation in Sec. III-C.

A. Problem Statement

We formulate the overall manipulation task as a discrete-time Markov Decision Process (MDP) with state space \mathcal{S} , action space \mathcal{A} , reward r , discount factor γ , and ρ_0 as the distribution of the initial state \mathbf{s}_0 : $(\mathcal{S}, \mathcal{A}, \mathcal{T}_\theta, r, \gamma, \rho_0)$. We




Fig. 2: An example of morphology parameter vector θ and shape deformation. Here, the morphology is parameterized by $\theta = [\theta_1, \theta_2, \theta_3, \theta_4]^T$, where $\theta_1, \dots, \theta_4$ represent the lengths of the four segments highlighted in blue. Upon updating morphology parameters θ to θ' , we first get the deformed cage vertices, then get the full mesh, as described by Eqn. (1).

parameterize the transition function $\mathcal{T}_\theta(s, a)$ by θ , which denotes a vector of tool morphology parameters as will be detailed in Sec. III-B. Different tool morphologies influence how the robot has to interact with an object to maximize reward. Imagine we are pushing peas onto a scoop with a tool that has a simple rectangular shape. It is likely that the peas will roll off to the side requiring the robot to re-orient the tool while pushing. However, if the tool is shaped to have a concave recess in the center, the peas would be less likely to roll off and escape while being pushed. This illustrates that different θ s yield different MDPs, since they change the transition function $\mathcal{T}_\theta(s, a)$. Our goal is to learn the optimal morphology parameters θ^* that maximize the task reward when the robot executes a control policy π . In principle, this framework is general enough to allow joint learning of morphologies and control policies. In practice, in this work we focus on learning tool morphology¹. To be consistent with notation in the literature on differentiable simulation, instead of reward maximization we describe the optimization problem as loss minimization.

B. Morphology Parameterization using Cage-Based Deformation

Cage-based deformation techniques are a common tool for deforming meshes in graphics applications [37] and has previously been used to learn optimal hand morphology in a robotics context [9]. We adopt the same morphology parametrization. Typically, a cage is a closed, low resolution mesh that envelopes the high-resolution mesh of the object we want to deform. Given an initial mesh \mathcal{M} of the object and the cage \mathcal{C} around it, we use $\mathbf{m}_i \in \mathcal{M}$ to denote the position of the i^{th} mesh vertex, and $\mathbf{c}_j \in \mathcal{C}$ to denote the position of the j^{th} vertex in the cage. Cage-based deformation establishes a linear mapping from the cage vertices $\mathbf{c}_j \in \mathcal{C}$ to each mesh point $\mathbf{m}_i \in \mathcal{M}$ by computing deformation weights w_{ij} :

$$\mathbf{m}_i = \sum_{j, \mathbf{c}_j \in \mathcal{C}} w_{ij} \mathbf{c}_j, \quad \sum_{j, \mathbf{c}_j \in \mathcal{C}} w_{ij} = 1, \quad \forall i, \mathbf{m}_i \in \mathcal{M} \quad (1)$$

¹In this work, we focus on morphology learning and regard policy learning as a separate line of work. However, our framework is compatible with policy learning as well. We show in the supplement video that existing methods, e.g. DiffHand [9] have significant difficulties with joint optimization of policy and morphology when presented with task variations. Hence, this is still an open problem for future work. Please see the Appendix in our arXiv version (arXiv 2211.02201) and the video for more details on the policy training.

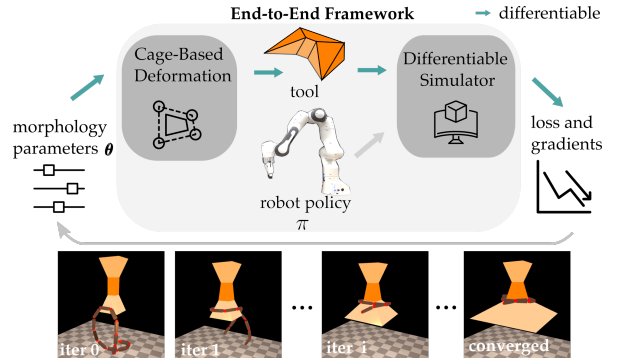


Fig. 3: Visualization of our end-to-end pipeline for learning tool morphology. We first obtain the tool parametrized by morphology parameters θ with cage-based deformation. Then we execute the policy π with this tool in the simulator and output the loss. Since all of these operations are differentiable, we can get the analytical gradients to update θ . We also show the evolving shape for a winding tool. The bottom part gradually becomes larger, which prevents the rope from slipping off.

In this work, we adopt the Mean Value Coordinate method [38], [9] for computing the deformation weights w_{ij} . Using these weights, we can manipulate the low-resolution cage vertices to deform the high-resolution mesh. Suppose we deform the cage \mathcal{C} and obtain new cage vertices $\mathbf{c}'_j \in \mathcal{C}'$. Then, we can compute the deformed mesh vertices as $\mathbf{m}'_i \in \mathcal{M}'$, with $\mathbf{m}'_i = \sum w_{ij} \mathbf{c}'_j, \forall j, \mathbf{c}'_j \in \mathcal{C}'$.

To get an even more compact representation of the tool morphology, we further extract the high-level morphology attributes (e.g. tool segment length, height, width) and denote these morphology parameters as $\theta \in \mathbb{R}^d$. Fig. 2 shows a basic example. When we change the morphology parameters corresponding to segment width from θ to θ' , we first map θ' to a deformed cage \mathcal{C}' , obtaining vertices $\mathbf{c}'_j \in \mathcal{C}'$. Then, we compute the deformed tool mesh $\mathbf{m}'_i \in \mathcal{M}'$ using Eqn. (1). The mappings from θ' to cage \mathcal{C}' and that from cage \mathcal{C}' to object mesh \mathcal{M}' are both linear, so the overall mapping from θ' to mesh \mathcal{M}' is linear as well.

C. End-to-end Pipeline Based on Differentiable Simulation

With the morphology parameterized by Cage-Based Deformation, we can now build our end-to-end tool morphology learning framework with differentiable simulation. In this work, we adopt DiffHand [9] as our simulator and build our pipeline upon it. However, our framework is agnostic to the choice of differentiable simulators and thus compatible with other differentiable simulators as well.

We visualize our end-to-end pipeline in Fig. 3. With the cage-based deformation, we effectively parametrize the tool with morphology parameter θ . We then instantiate the task with the deformed tool, initialize it with state \mathbf{s}_0 , and execute the robot policy π in the differentiable simulator. We then compute the gradients of the task-dependent loss $\mathcal{L}(\pi; \mathbf{s}_0, \theta)$ with respect to θ :

$$\frac{\partial \mathcal{L}(\pi; \mathbf{s}_0, \theta)}{\partial \theta} = \sum_{m_i \in \mathcal{M}, c_j \in \mathcal{C}} \frac{\partial \mathcal{L}(\pi; \mathbf{s}_0, \theta)}{\partial m_i} \cdot \frac{\partial m_i}{\partial c_j} \cdot \frac{\partial c_j}{\partial \theta}. \quad (2)$$

$\frac{\partial \mathcal{L}(\pi; \mathbf{s}_0, \theta)}{\partial m_i}$ is provided by the differentiable simulator. $\frac{\partial m_i}{\partial c_j}$ and $\frac{\partial c_j}{\partial \theta}$ are derived from the cage-based deformation in

Sec. III-B. We then take a gradient step on the tool morphology parameter θ and repeat the process until convergence.

IV. CONTINUAL LEARNING FOR ROBUST TOOL MORPHOLOGY

In this section, we describe how we learn a robust tool morphology using the end-to-end pipeline from Sec. III. We first elaborate on the challenges of learning robust tool morphology for contact-rich scenarios from several perspectives in Sec. IV-A, and then formulate the learning problem as continual learning and introduce our proposed method in Sec. IV-B.

A. Challenges for Learning Robust Tool Morphology

We aim to learn a tool morphology that is customized for a specific contact-rich manipulation task, while also being robust and generalizable to task variations, e.g. different start states for the task. Instead of being able to handle just one initial state, as in [9], we propose to learn a tool morphology parameter vector θ that generalizes across a distribution of initial states $\mathbf{s}_0 \sim \rho_0(\mathbf{s})$. Since evaluating the expectation of the loss would be intractable, we instead sample a set of N initial states, $S = \{\mathbf{s}_0^i | \mathbf{s}_0^i \sim \rho_0(\mathbf{s}), i = 1 \dots N\}$ and optimize the empirical expectation of the loss:

$$\begin{aligned} \theta^* &= \operatorname{argmin}_{\theta} \frac{1}{N} \sum_{\mathbf{s}_0^i \in S} [\mathcal{L}(\pi; \mathbf{s}_0^i, \theta)], \\ S &= \{\mathbf{s}_0^i | \mathbf{s}_0^i \sim \rho_0(\mathbf{s}), i = 1 \dots N\} \end{aligned} \quad (3)$$

Directly optimizing this loss with our pipeline in Sec. III by unrolling episodes in the simulator with different initial states $\mathbf{s}_0^i \in S$ presents several challenges. First of all, compared to tasks that do not have complex dynamics, contact-rich manipulation tasks usually have a highly non-convex optimization landscape, which causes gradient-based optimization algorithms to easily get stuck in local minima. Fig. 4 illustrates this by comparing the contact-rich task of *Winding* a rope on a spool to a free-space *Reaching* task. Therefore, for contact-rich tasks, directly optimizing Eqn. (3) might converge to local optima leading to suboptimal θ values that determine the tool morphology. Second, there is a trade-off when choosing N , the number of samples in our empirical estimate of the loss in Eqn. (3): Small values of N correspond to an insufficient number of samples, and would fail to capture the distribution $\rho_0(\mathbf{s})$. Large values of N are computationally expensive, since evaluating even one gradient step requires unrolling N episodes in the differentiable simulator.

B. Continual Learning Based Algorithm

Given the challenges of learning a robust tool morphology with our end-to-end pipeline in Sec. III, we draw inspiration from continual learning [32], [33], [35].

Formally, we define a *scenario* as an MDP with an initial state distribution (see Sec. III). Within a given scenario, we define a *task variation* T^i as a restricted MDP with one initial state \mathbf{s}_0^i instead of a distribution over initial states. Therefore in this work, task variations are MDPs with




Fig. 4: We visualize the 2D slices of the loss landscape for the contact-rich *Winding* scenario and the no-contact *Reaching* scenario. For *Winding*, the goal is to prevent the rope from falling. For *Reaching*, the goal is to optimize the arm so that the end effector can reach the green dots. The variables we optimize over are illustrated in the figure as x, y : for *Winding* these are the lengths of two sides of the tool base; for *Reaching* these are the lengths of two of the robot links. The details for the loss functions are given in the Appendix. It is clear that the optimization landscape of the contact-rich task is significantly more complex than that of the no-contact task.

different initial states. In general, task variations could be defined in other ways, e.g. variations of physical parameters or goal states. With sampling task variation T^i using $\mathbf{s}_0^i \sim \rho_0(\mathbf{s})$, optimizing Eqn. (3) is equivalent to minimizing the average loss on the sampled task set $T = \{T^1 \dots T^N\}$:

$$\begin{aligned} \theta^* &= \operatorname{argmin}_{\theta} \frac{1}{N} \sum_{T^i \in T} [\mathcal{L}(\pi; \mathbf{s}_0^i, \theta)], \\ T &= \{T^1 \dots T^N\} \end{aligned} \quad (4)$$

Morphology Optimization over a Sequence of Tasks. We can now formulate the problem of learning the morphology for a set of initial states as a continual, multi-task learning problem, where the morphology learned for one initial state can inform the optimization of morphology in other task variations with different initial states. To tackle the issue of this optimization being intractable for large values of N , we draw inspiration from continual learning and solve the problem for the sampled task set T in a sequential manner. Task variations are solved sequentially in batches B_0, B_1, \dots , where each batch B_t contains M task variations, $B_t = \{T_t^{(1)} \dots T_t^{(M)} | T_t^{(i)} \in T, i = 1 \dots M\}$. We select $M \ll N$ so that optimizing for each batch becomes tractable. By processing the batches sequentially, we address the trade-off outlined in Sec. IV-A between computational efficiency and sufficient coverage of task variations.

Constructing the Loss with Knowledge Distillation Regularization. We denote the morphology parameter vector we obtain after optimizing over a sequence of batches $B_1 \dots B_{t-1}$ by θ_{t-1} . At timestep t , our method aims to learn θ_t using an incremental learning strategy by processing the current batch B_t . Intuitively, we aim to minimize the task loss L_t^{task} on the current batch B_t : $L_t^{\text{task}}(\theta) = \frac{1}{M} \sum_{T_t^{(i)} \in B_t} \mathcal{L}(\pi; T_t^{(i)}, \theta)$. To avoid forgetting what has already been learned, we construct a regularization term L_t^{distill} for distilling previous knowledge. To this end, we maintain a distillation task set \mathcal{D} , which is obtained by randomly sampling M task variations from the previously seen ones, i.e. from $\bigcup\{B_1, \dots, B_{t-1}\}$. When optimizing for θ_t on the current batch, we still want the updated parameters θ_t to perform similarly to θ_{t-1} on the previous task variations in distillation set \mathcal{D} . Thus, we define the regularization term L^{distill} as the squared error between

the simulated trajectories generated when using θ_t vs θ_{t-1} :

$$L_t^{\text{distill}}(\theta_t) = \frac{1}{M} \sum_{T^{(i)} \in \mathcal{D}} \left(\text{sim}(\pi; T_t^{(i)}, \theta_{t-1}) - \text{sim}(\pi; T_t^{(i)}, \theta_t) \right)^2.$$

We could compute the distillation loss on a part of the trajectory that matters for the task, e.g. height of the rope for the *Winding* scenario (Appendix gives further details).

Finally, we get the overall loss by combining L_t^{task} and L_t^{distill} with a regularization coefficient α : $L^t(\theta_t) = L_t^{\text{task}}(\theta_t) + \alpha L_t^{\text{distill}}(\theta_t)$.

Simplifying Optimization with Dimensionality Reduction.

As visualized in Fig. 4, the contact-rich tasks considered in this paper result in a complex optimization landscape that is difficult to optimize over. To alleviate this issue, we draw inspiration from Coordinate Descent [39] and propose to simplify the optimization problem by reducing the dimensionality of the decision variable. For a new batch B_t , we evaluate the gradient with respect to the currently optimal θ_{t-1} and select $d' < d$ dimensions of parameters θ_{t-1} with the largest gradient magnitude. Then we only update these $d' < d$ dimensions of parameters θ_{t-1} to get θ_t . Once a dimension has been optimized in previous batches, we do not select it anymore. When all dimensions have been optimized, we restart (i.e. mark all dimensions as available for optimization) and repeat the process, decaying the learning rate with $\epsilon = e^{-1}$.

V. EXPERIMENTS

To test the performance of our proposed algorithm, we learn the tool morphology for three scenarios in simulation and evaluate the learned tool both in simulation and on a real robot manipulation setting. We first discuss the aspects that are the same for the three scenarios we consider, and then further elaborate on each of the scenarios separately.

Setup. We learn the tool morphology for all three scenarios with the differentiable simulator DiffHand [9] and use L-BFGS-B as our optimizer implemented using Scipy [40]. The implementation details including hyperparameters are given in the Appendix.

Scenarios. We consider three scenarios shown in Fig. 5.

- 1) *Winding*: In this scenario, a rope is wound around a tool. If the rope does not slip off the tool for a range of tool orientations, we consider this a success. For *Winding*, task variations correspond to the initial orientation of the rope and tool sampled from the uniform distribution over the space of 3D rotations i.e. $SO(3)$. As a robot policy, we use a simple circular motion. We compute the loss after the rope is wound and the end of it is left hanging free. Unless supported by the tool, the rope will fall off. Since the simulator we use in this work can only model rigid objects, we use a chain of cuboids to approximate the rope in this task.
- 2) *Flipping*: In this scenario, we aim to learn a robot arm morphology that can flip a box by 90° . For control, here

we first learn a basic flipping policy π (described in Appendix). This scenario is similar to that in previous work [9], except we consider a distribution of initial box poses. In prior work, the box is placed at one fixed position and fixed orientation facing the robot arm. Here, the cube position is randomly sampled from a square region. The cube orientation is sampled uniformly between $[-90^\circ, 90^\circ]$ around the z-axis.

- 3) *Pushing*: This scenario is based on bimanual scooping used for food acquisition as in [5]. We aim to learn the morphology of a pusher that can push a pea on a table into a scoop. Peas are modelled as spheres. Task variations T^i are instantiated with different initial pea positions sampled uniformly from a square region. For the trajectory of the pusher we use a forward motion with a zig-zag shape, which can make it challenging for the pusher to prevent the peas from rolling away.

We selected these scenarios because they include (i) various types of objects (rope, box, spherical peas), (ii) objects with different physical properties (boxes with sharp edges in scenario *Flipping*, smooth spheres in scenario *Pushing*), and (iii) varying contact conditions including a rope sliding on a winding tool, a robot holding on to a sharp box edge for pivoting, and spheres rolling over a tool’s surface.

Baselines. We implemented two baselines for comparison:

- 1) *Baseline-DiffHand*: we optimize θ for only one batch of tasks: $\theta = \text{argmin} L_0^{\text{task}}(\theta)$. This baseline is a direct extension of DiffHand [9] obtained by replacing the single initial state with one batch of initial states and optimizing until convergence, without bringing in the continual learning aspect.
- 2) *Simple-Continual*: we optimize θ by minimizing the task loss sequentially for batches B_1, \dots, B_t, \dots as we would in a continual learning setting, making this a stronger baseline. At batch B_t , we start from θ_{t-1} and obtain θ_t using $\theta_t = \text{argmin} L_t^{\text{task}}(\theta)$. After optimizing for all the batches, we get the final morphology parameter $\theta_{t=N/M}$. With the continual learning setting introduced, the difference between *Ours* and this baseline is that *Ours* uses knowledge distillation regularization in addition to the task loss, and simplifies the optimization with dimensionality reduction.

A. Results & Analysis in Simulation

As shown in Fig. 5, the simulation results demonstrate that the tool morphology learned with our algorithm consistently outperforms both *Baseline-DiffHand* and *Simple-Continual* in terms of a lower test loss and higher test success rate. Across all scenarios, *Baseline-DiffHand* demonstrates the worst performance with the largest variance. This is because, compared to optimizing over sequential batches of task variations, approximating the task distribution with only one batch will inherently be noisier. In scenario *Flipping* and *Pushing*, the baseline *Simple-Continual* outperforms the *Baseline-DiffHand*. This also indicates that learning from a sequence of task variations instead of one batch is beneficial. In

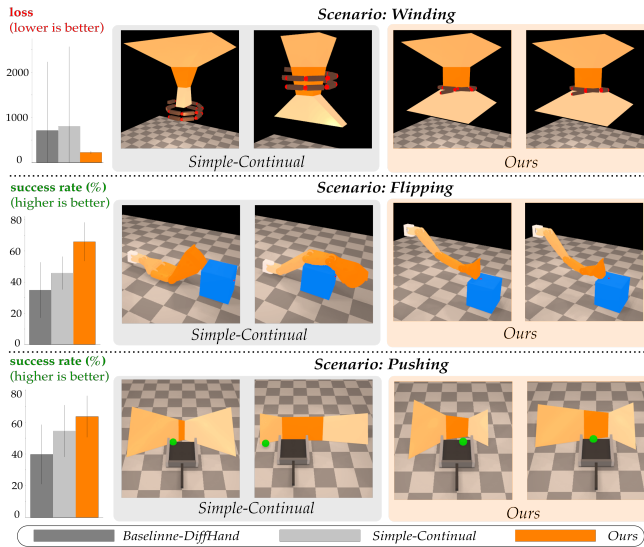


Fig. 5: Evaluation results on scenarios *Winding* (top row), *Flipping* (middle row) and *Pushing* (bottom row). For each scenario, we implement *Baseline-DiffHand*, *Simple-Continual* and our algorithm (*Ours*) in simulation. For quantitative evaluation, we report the mean and standard deviation over ten runs for the task loss for *Winding* and test accuracy for *Flipping* and *Pushing*. For qualitative evaluation we visualize examples of two optimized morphologies from baseline *Simple-Continual* and our algorithm (*Ours*).

scenario *Winding*, baseline *Simple-Continual* and *Baseline-DiffHand* show poor performance with a high average loss and large standard deviation, while *Ours* achieves a low loss and small variance (visualized as standard deviation in the bar plots). For some of the runs, *Simple-Continual* and *Baseline-DiffHand* have an exceptionally high loss due to completely failing to hold the rope, so the rope quickly falls down. For both baselines, the performance deteriorates significantly for some batch sequences by converging to a suboptimal morphology parameter. In contrast, our method is not sensitive to how batches are sampled.

We visualize two of the learned tool morphologies for *Ours* and baseline *Simple-Continual* for each scenario in Fig. 5. In scenario *Winding*, one of the learned shapes from *Simple-Continual* has one pointed end and one flat end, making the tool not effective for a certain range of orientations. This shows that the shape of the tool learned with baseline *Simple-Continual* sometimes does not perform well across all task variations. In scenario *Flipping*, in contrast to the wide arm tip learned by the *Simple-Continual* baseline, our method learns a finer triangular tip, which can exert enough pressure close to the edge of the box to cause it to flip, thus achieving a higher success rate. In scenario *Pushing*, our method learns a tool morphology with widening ends and a middle segment has a similar width as the scoop. The baseline *Simple-Continual* either learns a narrow middle segment or a wide one. An appropriate width for the middle segment can better facilitate pushing the peas all the way into the scoop. Specifically, a too narrow middle segment might result in a gap between the pusher and the scoop and thus have a hard time pushing the peas all the way into the scoop, while a too wide middle segment can result in the

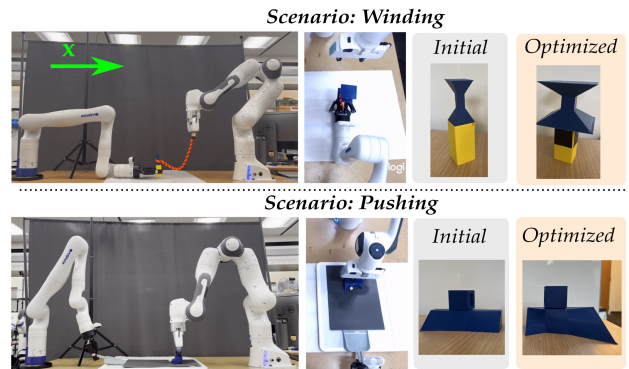


Fig. 6: Real world experiment set up for *Winding* and *Pushing*. For each scenario, we 3D print the tool with the initial shape and the optimized shape obtained by our approach. For scenario *Winding*, we first wind the rope around the tool and then rotate the tool around the highlighted x axis for 360 degrees to test whether the rope drops. For videos of experiments please see <https://sites.google.com/stanford.edu/learning-tool-morphology>

pea rolling away. In summary, across the three scenarios, we observe that the learned morphology from our algorithm demonstrates better performance.

B. Evaluating Learned Tools in the Real World

We manufactured the tools learned with our method for scenario *Winding* and *Pushing* as shown in Fig. 6. The tools were 3D printed on an Ender 3 Pro 3D printer using PLA filament. We tested the functionality of the tool in these two scenarios. In scenario *Winding*, we execute a simple control sequence that winds the rope around the tool. After winding, we rotate the tool in the air with a full circle of 360 degrees around the x axis as shown in Fig. 6. A trial is counted as success if the rope successfully stays wrapped around the tool after winding and rotating. We run 5 episodes with the initial shape and the optimized shape. Results show that the optimized shape achieves a 100% success rate while the initial shape keeps the rope on the tool only 20% of the time. For scenario *Pushing*, as shown in Fig. 6, we use the Franka Panda robot arm on the right to hold the pusher and try to push the white sphere to the flat dustpan held by a Kinova Gen3 arm on the left. The arm follows the waypoints of a pre-defined zig-zag trajectory. We also sample 4 initial locations for the pea and run 5 episodes for each initial location, which leads to 20 episodes for one tool. The experiments show that the optimized shape achieves a 70% success rate across the 20 trials while the initial shape achieves a 15% success rate.

VI. CONCLUSION

To summarize, in this work, we approach the problem of custom tool design for robot manipulation in an end-to-end manner by leveraging the advantages of differentiable simulation. To learn versatile tool morphologies, we propose a continual learning approach that enables optimization over task variations, e.g. by varying initial object poses. We show that tools optimized with our method help to achieve an improved task performance compared to baselines in simulation. We also demonstrate that these tools enable successful task completion in reality.

REFERENCES

- [1] K. R. Gibson, K. R. Gibson, and T. Ingold, *Tools, language and cognition in human evolution*. Cambridge University Press, 1994.
- [2] M. Saha and P. Isto, "Manipulation planning for deformable linear objects," *IEEE Transactions on Robotics*, vol. 23, no. 6, pp. 1141–1150, 2007.
- [3] T. Tsuji, J. Ohkuma, and S. Sakaino, "Dynamic object manipulation considering contact condition of robot with tool," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 3, pp. 1972–1980, 2015.
- [4] P. Sundaresan, S. Belkhal, and D. Sadigh, "Learning visuo-haptic skewering strategies for robot-assisted feeding," in *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022.
- [5] J. Grannen, Y. Wu, S. Belkhal, and D. Sadigh, "Learning bimanual scooping policies for food acquisition," in *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022.
- [6] Z. Lyu, Z. Xu, and J. Martins, "Benchmarking optimization algorithms for wing aerodynamic design optimization," in *Proceedings of the 8th International Conference on Computational Fluid Dynamics, Chengdu, Sichuan, China*, vol. 11, 2014.
- [7] H. Zhang, P. Long, D. Zhou, Z. Qian, Z. Wang, W. Wan, D. Manocha, C. Park, T. Hu, C. Cao, *et al.*, "Dorapicker: An autonomous picking system for general objects," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2016, pp. 721–726.
- [8] N. Rojas, R. R. Ma, and A. M. Dollar, "The gr2 gripper: An underactuated hand for open-loop in-hand planar manipulation," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 763–770, 2016.
- [9] J. Xu, T. Chen, L. Zlokapa, M. Foshey, W. Matusik, S. Sueda, and P. Agrawal, "An End-to-End Differentiable Framework for Contact-Aware Robot Design," in *Robotics: Science and Systems*, 2021.
- [10] P. Gorce and J. G. Fontaine, "Design methodology approach for flexible grippers," *Journal of Intelligent and Robotic Systems*, vol. 15, no. 3, pp. 307–328, 1996.
- [11] R. Saravanan, S. Ramabalan, N. G. R. Ebenezer, and C. Dharmaraja, "Evolutionary multi criteria design optimization of robot grippers," *Applied Soft Computing*, vol. 9, no. 1, pp. 159–172, 2009.
- [12] F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter, "End-to-end differentiable physics for learning and control," *Advances in neural information processing systems*, vol. 31, 2018.
- [13] J. Degraeve, M. Hermans, J. Dambre, *et al.*, "A differentiable physics engine for deep learning in robotics," *Frontiers in neurorobotics*, p. 6, 2019.
- [14] Y. Hu, L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand, "DiffTaichi: Differentiable Programming for Physical Simulation," *ICLR*, 2020.
- [15] K. Werling, D. Omens, J. Lee, I. Exarchos, and C. K. Liu, "Fast and feature-complete differentiable physics for articulated rigid bodies with contact," in *Robotics: Science and Systems*, 2021.
- [16] H. J. Suh, M. Simchowitz, K. Zhang, and R. Tedrake, "Do differentiable simulators give better policy gradients?" in *International Conference on Machine Learning*. PMLR, 2022, pp. 20668–20696.
- [17] R. Antonova, J. Yang, K. M. Jatavallabhula, and J. Bohg, "Rethinking Optimization with Differentiable Simulation from a Global Perspective," in *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022.
- [18] G. C. Causey and R. D. Quinn, "Gripper design guidelines for modular manufacturing," in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No. 98CH36146)*, vol. 2. IEEE, 1998, pp. 1453–1458.
- [19] A. Meixner, C. Hazard, and N. Pollard, "Automated design of simple and robust manipulators for dexterous in-hand manipulation tasks using evolutionary strategies," in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2019, pp. 281–288.
- [20] A. Wolniakowski, J. A. Jorgensen, K. Miatliuk, H. G. Petersen, and N. Krüger, "Task and context sensitive optimization of gripper design using dynamic grasp simulation," in *2015 20th International Conference on Methods and Models in Automation and Robotics (MMAR)*. IEEE, 2015, pp. 29–34.
- [21] L. C. M. W. Schwartz, A. Wolniakowski, A. Werner, L.-P. Ellekilde, and N. Krüger, "Designing fingers in simulation based on imprints," in *SIMULTECH*, 2017.
- [22] V. Velasco and W. S. Newman, "Computer-assisted gripper and fixture customization using rapid-prototyping technology," in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No. 98CH36146)*, vol. 4. IEEE, 1998, pp. 3658–3664.
- [23] M. Honarpardaz, M. Tarkian, J. Ölvander, and X. Feng, "Finger design automation for industrial robot grippers: A review," *Robotics and Autonomous Systems*, vol. 87, pp. 104–119, 2017.
- [24] X. Pan, A. Garg, A. Anandkumar, and Y. Zhu, "Emergent hand morphology and control from optimizing robust grasps of diverse objects," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7540–7547.
- [25] H. Ha, S. Agrawal, and S. Song, "Fit2Form: 3D generative model for robot gripper form design," in *Conference on Robotic Learning (CoRL)*, 2020.
- [26] L. Shao, T. Migimatsu, and J. Bohg, "Learning to scaffold the development of robotic manipulation skills," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 5671–5677.
- [27] M. A. Toussaint, K. R. Allen, K. A. Smith, and J. B. Tenenbaum, "Differentiable physics and stable modes for tool-use and manipulation planning," 2018.
- [28] L. Nair, J. Balloch, and S. Chernova, "Tool macgyvering: Tool construction using geometric reasoning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5837–5843.
- [29] E. Heiden, M. Macklin, Y. S. Narang, D. Fox, A. Garg, and F. Ramos, "Disect: A differentiable simulation engine for autonomous robotic cutting," in *Robotics: Science and Systems*, 2021.
- [30] M. Lutter, J. Silberbauer, J. Watson, and J. Peters, "Differentiable physics models for real-world offline model-based reinforcement learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4163–4170.
- [31] M. Delange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars, "A continual learning survey: Defying forgetting in classification tasks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [32] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, and N. Díaz-Rodríguez, "Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges," *Information fusion*, vol. 58, pp. 52–68, 2020.
- [33] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert, "icarl: Incremental classifier and representation learning," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017, pp. 2001–2010.
- [34] D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [35] A. Mallya and S. Lazechnik, "Packnet: Adding multiple tasks to a single network by iterative pruning," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 7765–7773.
- [36] J. Serra, D. Suris, M. Miron, and A. Karatzoglou, "Overcoming catastrophic forgetting with hard attention to the task," in *International Conference on Machine Learning*. PMLR, 2018, pp. 4548–4557.
- [37] J. R. Nieto and A. Susín, "Cage based deformations: a survey," in *Deformation models*. Springer, 2013, pp. 75–99.
- [38] T. Ju, S. Schaefer, and J. Warren, "Mean value coordinates for closed triangular meshes," in *ACM Siggraph 2005 Papers*. Proceedings of SIGGRAPH, 2005, pp. 561–566.
- [39] S. J. Wright, "Coordinate descent algorithms," *Mathematical Programming*, vol. 151, no. 1, pp. 3–34, 2015.
- [40] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, *et al.*, "Scipy 1.0: fundamental algorithms for scientific computing in python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.