

Learning Switching Criteria for Sim2Real Transfer of Robotic Fabric Manipulation Policies

Satvik Sharma^{*1}, Ellen Novoseller^{*1}, Vainavi Viswanath¹, Zaynah Javed¹, Rishi Parikh¹, Ryan Hoque¹, Ashwin Balakrishna¹, Daniel S. Brown¹, Ken Goldberg¹

Abstract—Simulation-to-reality transfer has emerged as a popular and highly successful method to train robotic control policies for a wide variety of tasks. However, it is often challenging to determine when policies trained in simulation are ready to be transferred to the physical world. Deploying policies that have been trained with very little simulation data can result in unreliable and dangerous behaviors on physical hardware. On the other hand, excessive training in simulation can cause policies to overfit to the visual appearance and dynamics of the simulator. In this work, we study strategies to automatically determine when policies trained in simulation can be reliably transferred to a physical robot. We specifically study these ideas in the context of robotic fabric manipulation, in which successful sim2real transfer is especially challenging due to the difficulties of precisely modeling the dynamics and visual appearance of fabric. Results in a fabric smoothing task suggest that our switching criteria correlate well with performance in real. In particular, our confidence-based switching criteria achieve average final fabric coverage of 87.2-93.7% within 55-60% of the total training budget. See <https://tinyurl.com/lsc-case> for code and supplemental materials.

I. INTRODUCTION

Training robot control policies in simulation for deployment in physical environments and automation tasks has seen increasing success in a broad variety of robotics applications such as robot grasping [22, 28], deformable manipulation [9, 33], and legged locomotion [15, 18]. Simulated training enables the collection of diverse experiences for policy learning that may be infeasible to collect on a physical system, creating potential for strong policy generalization and robustness. Additionally, learning policies purely in simulation can limit the risk of encountering suboptimal behaviors on physical robots and can reduce hardware wear-and-tear. However, transferring policies learned in simulation to the real world is challenging, as simulators are imperfect. As a result, learning algorithms must overcome a domain shift [37] between the dynamics and visual appearance of simulations and the physical world. Prior work has addressed a range of methods to enable robust domain adaptation of policies trained in simulation to physical experiments [7, 11, 18, 32, 35, 40]. Less attention has been given to the critical challenge of determining when a policy trained in simulation is ready for physical deployment. While deploying policies trained with too little simulation data can lead to low performance on the physical system, policies that are trained excessively in simulation risk overfitting to simulator artifacts. This motivates principled methods to determine




Fig. 1: **Physical experiment and simulator setup.** We study sim-to-real switching in a fabric smoothing task, in an environment consisting of an ABB YuMi robot with a single tweezer gripper. An overhead Photoneo Phoxi Camera captures grayscale images. The manipulation workspace border is marked with blue tape, and the fabric is located within the workspace border. The physical workspace is designed to visually emulate the Gym-Cloth simulator [33], as shown on the left; the top two lefthand images show example starting and ending configurations from an oracle smoothing policy, while the bottom two images show the same observations, processed to resemble the grayscale images taken by the Photoneo Phoxi Camera.

when policies trained in simulation are likely most ready for physical deployment.

Determining when a policy is ready for physical deployment requires reasoning about its expected performance in physical trials when the policy has only been trained in simulation. One common way to achieve this is to periodically evaluate learned policies in the physical world during simulated training to determine when sufficient performance has been achieved [3, 16, 24]. However, this poses a number of practical challenges, as performing physical rollouts during training costs significant time and engineering effort, and furthermore may pose safety risks. This motivates using policy rollouts only from simulation to determine when to deploy policies on a physical robot. While this setting is more practical, it is also significantly more difficult, since it is challenging to gain information about expected performance in physical trials at deployment time without access to the

¹Authors are with the University of California, Berkeley, USA.

^{*} These authors contributed equally.

physical system prior to deployment.

Our key insight is that the problem of determining when to transfer robotic control policies from simulation to reality bears similarities to the early stopping problem, a well-known technique for reducing overfitting in machine learning [23, 29, 39]. For instance, Prechelt [29] investigates 14 different early stopping criteria with cross validation, which assesses generalization by measuring performance on a validation dataset unseen by the learner. Meanwhile, other approaches do not rely on the existence of a validation set: Mahsereci *et al.* [23] derive an early stopping rule based on local gradient statistics and Yao *et al.* [39] propose a rule based on the bias-variance tradeoff. While early stopping is well-studied, however, few have explored applying the technique to sim2real transfer. One such work is Muratore *et al.* [27], who propose early stopping in simulation based on an upper confidence bound on the optimality gap during domain randomization.

This work, meanwhile, studies four different mechanisms for switching from simulation to physical experiments and evaluates the degree to which they are predictive of performance during physical deployment. Specifically, we focus on robot fabric manipulation, a domain that has seen increasing interest [9, 20, 33] and applies across many domains such as home robotics, senior care, laundry folding, and manufacturing. In particular, we study how our proposed sim2real switching criteria can be used to determine when to deploy fabric smoothing policies on a physical robotic system.

This paper contributes (1) a detailed empirical study of which metrics are predictive of robust policy transfer from simulation to the physical world and (2) an evaluation of the proposed method for transferring fabric smoothing policies from simulation to deployment on a physical robot. Fabric smoothing experiments suggest that our switching criteria correlate with performance in real, achieving an average final fabric coverage of 88.79% over all switching conditions. Furthermore, given a training budget of 200 learning iterations, our confidence-based switching criteria utilize only 55-60% of total training resources.

II. RELATED WORK

A. Sim2Real Transfer for Robot Learning

There is significant prior work on learning policies in simulation and facilitating transfer to physical experiments. The most common approach is domain randomization [22, 36], which varies dynamical and/or visual properties such as friction, lighting, camera angle, colors, and textures in simulation to enable zero-shot transfer to the physical world. Li *et al.* [18] apply dynamics randomization according to a curriculum to achieve physical bipedal robot walking. Similarly, Valassakis *et al.* [37] investigate dynamics domain randomization via injecting random forces into simulations. Domain randomization has achieved transfer in a variety of applications such as robot legged locomotion [18], fabric manipulation [9], and robotic grasping [22, 36]. Recent work has also explored adaptive methods and curricula for domain randomization, such as automatic domain randomization

for dexterous manipulation of a Rubik’s cube [1] and active domain randomization [26]. Other proposed approaches for crossing the reality gap include learning a canonical intermediate representation [11], domain adaptation via a small amount of real data [35], and “real2sim2real” tuning of simulation parameters based on real data [2, 19, 30]. Finally, Kadian *et al.* [13] propose a sim2real correlation coefficient for predicting sim2real performance but require real data to compute the coefficient. Despite the significant body of work on simulated training procedures to enable physical deployment, there has been very little research on automatically determining precisely when to stop training in simulation and deploy learned policies in the physical world without any real evaluation. The closest prior work to ours may be Muratore *et al.* [27], who propose early stopping in simulation based on an upper confidence bound on the optimality gap during domain randomization. However, while this work is specific to the reinforcement learning paradigm, we apply our approach to the imitation learning setting, in which algorithms learn from demonstrations. Furthermore, our proposed framework can readily extend to other settings such as reinforcement learning; this could be an interesting direction for future study.

B. Fabric Manipulation

Fabric manipulation is a challenging open problem for robots, as fabric has an infinite-dimensional configuration space, exhibits complex dynamics, and often self-occludes. There has been significant prior work in leveraging simulators and task demonstrations to learn fabric smoothing and folding policies for zero-shot transfer to the physical world [5, 6, 12, 25, 33]. Other recent works relax the assumption that task demonstrations are available but still leverage simulation to enable zero-shot transfer to real world experiments [9, 21, 38]. Another approach is to learn fabric manipulation policies entirely from real world interaction [4, 17]. Our work learns fabric manipulation policies purely in simulation for deployment in the physical world; however, unlike prior work that performs zero-shot transfer to physical experiments after some fixed amount of simulated training, we study how to automatically determine when to stop training policies in simulation to facilitate robust deployment on the physical robot.

III. PROBLEM STATEMENT

This work considers learning a policy π for some task given access to only a computational simulation. In particular, we seek to identify the optimal time at which to terminate learning in simulation and deploy the policy in the physical world. We model the task as a Partially Observable Markov Decision Process (POMDP) [14] $\mathcal{M}_{\text{sim}}^\phi$, parametrized by simulation parameters ϕ capturing variables such as visual appearance and simulator dynamics. The objective is to achieve high task performance when this policy is deployed in the physical world, which we model as POMDP $\mathcal{M}_{\text{real}}$.

More formally, we assume that $\mathcal{M}_{\text{sim}}^\phi$ and $\mathcal{M}_{\text{real}}$ have a shared state space \mathcal{S} , observation space \mathcal{O} , action space \mathcal{A} ,

reward function $R : \mathcal{S} \rightarrow \mathbb{R}$, initial state distribution μ , and episode time horizon H , but may have different mappings from states $s \in \mathcal{S}$ to observations $o \in \mathcal{O}$. In this work, we consider grayscale image observations ($\mathcal{O} = \mathbb{R}^{H \times W \times C}$, $C = 1$), but consider settings (as is typical in practice) where observations corresponding to a specific state may have different visual appearances between simulation ($\mathcal{M}_{\text{sim}}^\phi$) and reality ($\mathcal{M}_{\text{real}}$). We additionally consider settings in which the state transition dynamics associated with $\mathcal{M}_{\text{sim}}^\phi$, denoted by $P_{\text{sim}} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, may be different from those associated with $\mathcal{M}_{\text{real}}$, denoted by $P_{\text{real}} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$. This reflects the inability of computational simulations to precisely model the dynamics of the physical world.

We consider a setup in which robot policy $\pi : \mathcal{O} \rightarrow \mathcal{A}$ is first learned with some policy search algorithm in simulation, using T_{sim} total transitions in simulated environment $\mathcal{M}_{\text{sim}}^\phi$. During this phase, we aim to optimize the following objective based on the attained rewards in simulation:

$$J_{\text{sim}}^\phi(\pi) = \mathbb{E}_{\pi, \mathcal{M}_{\text{sim}}^\phi} \left[\sum_{t=1}^H R(s_t) \right], \quad (1)$$

where the expectation is with respect to observation-action trajectories sampled from policy π in POMDP $\mathcal{M}_{\text{sim}}^\phi$. The objective in this work is to identify the optimal stopping time T_{sim} such that when policy π is evaluated in the physical world, it maximizes the following objective based on the attained rewards in physical trials:

$$J_{\text{real}}(\pi) = \mathbb{E}_{\pi, \mathcal{M}_{\text{real}}} \left[\sum_{t=1}^H R(s_t) \right], \quad (2)$$

where analogously, the expectation is with respect to observation-action trajectories sampled from π in $\mathcal{M}_{\text{real}}$.

IV. SWITCHING CRITERIA FOR IMITATION LEARNING

In this work, we consider switching from simulation to real in the case of imitation learning, in which an algorithm learns a policy from sequentially-collected demonstrations. We propose four criteria to determine when to switch from simulated training to physical deployment. In Section IV-A, we discuss a general framework for evaluating switching criteria, and in Section IV-B, we will describe the two evaluation metrics for switching (reward in simulation and epistemic uncertainty) studied in this work. Then, Section IV-C will propose two stopping conditions for identifying switching points (absolute threshold and gradient-based). Pairing the two evaluation metrics and two stopping conditions results in four possible switching criteria, each a combination of a particular switching metric and stopping condition.

A. Switching to Real in Imitation Learning

To determine when to switch from learning in simulation to deploying policies in physical experiments, we learn a switching criterion $\psi : \pi \rightarrow \{0, 1\}$. If $\psi(\pi) = 1$, we terminate simulated training and deploy π in physical experiments, while otherwise, we continue training π . The switching function ψ identifies a switching time T_{sim} as a

function of the evolving robot policy π . While in this work, we specifically study switching criteria in the context of learning from demonstrations, future work includes exploring similar ideas for other possible policy learning algorithms, including model-based planning with system identification and model-free reinforcement learning.

We consider the imitation learning setting, in which the objective is to learn a robot policy π that emulates task demonstrations collected in simulation from some demonstration policy π_D . In principle, π_D could be a human controlling the robot in simulation via a teleoperation interface or an algorithmic controller defined using privileged state information present only in simulation; we consider the latter. We train π in simulation with *behavior cloning* [8, 31, 33], which aims to minimize the following loss function to encourage the robot's policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ to match that of the demonstrator (π_D) on a set of demonstration trajectories \mathcal{D} :

$$J_{BC}(\pi) = \mathbb{E}_{(s_t, \pi_D(s_t)) \sim \mathcal{D}} [\mathcal{L}(\pi(s_t), \pi_D(s_t))], \quad (3)$$

where $\mathcal{L}(\pi(s), \pi_D(s))$ is an action discrepancy measure between $\pi(s)$ and $\pi_D(s)$ (e.g., the squared loss or 0-1 loss; our experiments use the squared loss).

The key objective in this work is to determine how many training iterations in simulation are required to train π before it should be deployed on a physical system. To this end, we train π to minimize the loss in Equation (3) in simulation on trajectories in \mathcal{D} and compute a stopping criterion $\psi : \pi \rightarrow \{0, 1\}$. If $\psi(\pi) = 1$, we terminate simulated training and deploy π in physical experiments. Otherwise, we query π_D for an additional task demonstration in simulation, aggregate this demonstration into \mathcal{D} , and continue.

The learning process is described in Algorithm 1 and illustrated in Figure 2. The procedure alternates between 1) collecting a batch of demonstrations from the demonstration policy π_D , 2) appending the demonstrations to the dataset \mathcal{D} , 3) performing a model update epoch to minimize the loss function $J_{BC}(\pi)$ in Equation (3), and 4) checking whether the switching criterion $\psi(\pi)$ is satisfied. If $\psi(\pi) = 1$, then the process is halted and the final learned policy π is deployed on the physical robot system. In practice, our experiments collect a batch of $K = 10$ demonstrations per iteration, where we use $K = 10$ to obtain sufficient granularity of information for determining the switching point. Each model update epoch performs 400 gradient steps on the loss function $J_{BC}(\pi)$ with a minibatch size of 64. We utilize an Adam optimizer with a learning rate of 2.5×10^{-4} and L2 regularization of 10^{-5} .

B. Switching Metrics

We propose two switching metrics, which are computed after every model update step in Algorithm 1.

Simulation Reward: Here we evaluate π in simulation and compute the average total reward over L rollouts to approximate $J_{\text{sim}}(\pi)$ from Equation (1). These trials also serve as the cross-validation dataset $\mathcal{D}_{\text{cross}}$ for tuning our reward-based switching conditions via cross-validation in simulation, as detailed below. In practice, our experiments

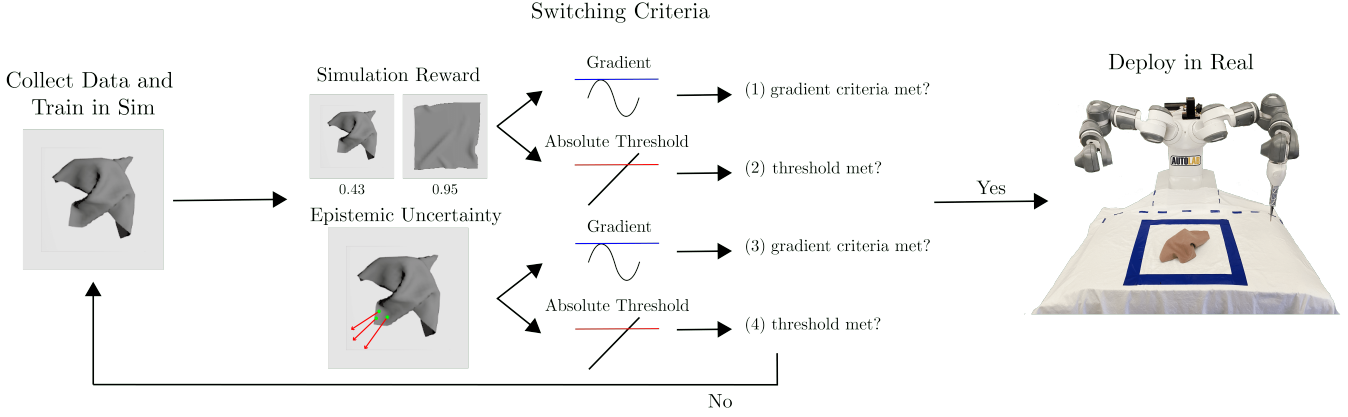


Fig. 2: **System Overview.** At each step, our algorithm pipeline collects a new batch of simulation data, performs a model update epoch, and then checks whether a switching condition is satisfied. If the switching criterion is met, then the model is ready to be deployed in real. Otherwise, we continue collecting simulation data to further update the model. We test four switching criteria, which utilize metrics based on a) reward when evaluated in simulation and b) epistemic uncertainty as estimated via an ensemble of policy networks; these metrics are paired with each of two stopping conditions based on 1) absolute thresholding or values, and 2) gradients.

Algorithm 1 Learned Switching Criteria for Sim2Real

Require: Maximum number of learning iterations N ; number of demonstrations K collected between each model update and calculation of the stopping criterion $\psi(\pi)$; demonstration policy π_D ; initialized buffer for demonstration data $\mathcal{D} \leftarrow \emptyset$; randomly initialized robot policy π

```

1: for  $i \in \{1, \dots, N\}$  do
2:   Collect  $\tau_j^{\text{demo}} = ((s_t^{(j)}, a_t^{(j)})_{t=1}^T)$ ,  $j \in \{1, \dots, K\}$  from  $\pi_D$ 
3:    $\mathcal{D} = \mathcal{D} \cup \tau_1^{\text{demo}} \cup \dots \cup \tau_K^{\text{demo}}$ 
4:    $\pi \leftarrow \arg \min_{\pi} \mathbb{E}_{(s_t, \pi_D(s_t)) \sim \mathcal{D}} [\mathcal{L}(\pi(s_t), \pi_D(s_t))]$ 
5:   if  $\psi(\pi)$  then
6:     Terminate learning and deploy on physical system
7:   end if
8: end for

```

use $L = 5$, which we found to provide sufficiently small standard error. Intuitively, performing well in simulation suggests that the policy may transfer well to real.

Epistemic Uncertainty: Here we approximate the epistemic uncertainty of π to characterize the policy’s confidence in the actions that it predicts. We estimate epistemic uncertainty by training an ensemble of E policies $\{\pi^{(i)}\}_{i=1}^E$ on bootstrapped minibatches of the training data \mathcal{D} . We then estimate the epistemic uncertainty over a holdout demonstration set $\mathcal{D}_{\text{cross}}$, which is not present during model training. (Note that $\mathcal{D}_{\text{cross}}$ has different definitions for the simulation reward and epistemic uncertainty metrics.) Intuitively, the more consistent (low-variance) the policy’s predictions on $\mathcal{D}_{\text{cross}}$ across ensemble members, the higher the probability of policy convergence. Let $\{a_i(s)\}_{i=1}^E \in \mathbb{R}^M$ denote the corresponding actions when each ensemble member is queried at observation s . Further, let $a_{ij}(s)$ denote the j^{th} component of the action predicted by ensemble member i . The epistemic uncertainty is estimated via:

$$\mathbb{E}_{s \in \mathcal{D}_{\text{cross}}} \left[\frac{1}{M} \sum_j \text{Var}_i(a_{ij}(s)) \right], \quad (4)$$

where $\text{Var}_i(a_{ij}(s))$ denotes the variance over ensemble member $i \in \{1, 2, \dots, E\}$ in action component j , and the expectation is taken over observation-action pairs in $\mathcal{D}_{\text{cross}}$.

C. Stopping Conditions

We now leverage the two metrics in Section IV-B to propose two stopping conditions for sim2real switching.

Value-Based: The first stopping condition is an absolute threshold, A , which we tune using the cross-validation set, $\mathcal{D}_{\text{cross}}$. For the simulation reward metric, A is set to approximately the maximum reward attained in $\mathcal{D}_{\text{cross}}$ (on average over the L rollouts in $\mathcal{D}_{\text{cross}}$); notably, A is tuned using $\mathcal{D}_{\text{cross}}$, an offline dataset in which rewards plateau to stable performance in simulation, so that future online learning can halt early using A . For epistemic uncertainty, A is the highest epistemic uncertainty over $\mathcal{D}_{\text{cross}}$ that falls within a preset margin of the minimum epistemic uncertainty value. Specific values of A are provided in the supplemental website [34]. At each evaluation episode i , a spline $f(x)$ is fit to the previous $i - 1$ data points. We use a spline—rather than raw data—to mitigate noise in the reward and uncertainty-based switching metrics, as seen in Fig. 4. Then, the spline is evaluated at the current episode and compared with the predefined absolute threshold A to determine whether the algorithm should stop training and deploy the policy in real.

Gradient-Based: The second stopping condition is gradient-based, and determines when the change over time in an evaluation metric falls below a threshold. Similarly to the absolute threshold, we identify the threshold by fitting a spline to the previous $i - 1$ data points. Then, finite differences are taken for two points evaluated on the spline to approximate the current gradient of the evaluation metric. If the gradient is within the range $[-\epsilon, \epsilon]$ (tuned for each switching metric as described below), the gradient is considered to be sufficiently close to zero and two counters are

potentially incremented: w_{consec} keeps track of the number of consecutive episodes with a sufficiently-small gradient, while w_{total} keeps track of the total number of episodes with a below-threshold gradient. The stopping condition is triggered if either $w_{\text{consec}} > U$ or $w_{\text{total}} > V$, where U is the maximum number of consecutive low-gradient episodes before stopping and V is the maximum total number of (potentially-nonconsecutive) low-gradient episodes. The parameter V allows training to stop even if there are some aberrations that would reset w_{consec} to zero. However, the V condition alone is insufficient, since if the gradient frequently oscillates near zero due to noise in the switching metric, w_{total} would be incremented even though we may not have converged to acceptable performance, in which case the policy can still learn and should not yet be deployed in real.

Fig 2 illustrates this process. In practice, the parameters (ϵ, U, V) of the stopping conditions are tuned via cross-validation in simulation, using the dataset $\mathcal{D}_{\text{cross}}$. In particular, ϵ is empirically set to the largest value such that not many early iterations (which have comparatively larger gradients) are classified as having a below-threshold gradient. U and V are tuned in conjunction so that the stopping condition is triggered as early as possible but not too early (e.g., for an overly-low simulation reward). Specific values for hyperparameters are provided in the supplemental website [34].

To summarize, our four proposed sim2real switching methods are: 1) simulation reward with the value-based stopping condition, 2) simulation reward with the gradient-based stopping condition, 3) epistemic uncertainty with the value-based stopping condition, and 4) epistemic uncertainty with the gradient-based stopping condition.

V. EXPERIMENTS

The four switching criteria—the two evaluation metrics coupled with each of the two stopping conditions—are each evaluated on a fabric smoothing task, in which we use behavior cloning to learn a policy from demonstrations.

A. Experimental Setup: Fabric Smoothing

1) *Sequential Fabric Smoothing Problem Statement:* We consider the specific manipulation task of sequential fabric smoothing, a challenging open problem in robotics that has received significant recent interest [5, 9, 33]. As described in prior work [33], the objective in sequential fabric smoothing is to find a sequence of robot actions to maximally smooth a fabric from an initially crumpled configuration. Concretely, we consider a square crop of fabric with initial configuration (state) s_0 and configuration s_t at time t . In both simulation and reality, the algorithm can access only top-down grayscale image observations of the workspace, $\mathbf{o}_t \in \mathcal{O} = \mathbb{R}^{H \times W}$, where H and W are the image height and width respectively. Following prior work [10], we assume that each side of the fabric is monochromatic and that the two sides are colored differently, where the colors are distinguishable in grayscale (see supplement [34]). At each timestep t , the robot executes a 4D pick-and-place action a_t parameterized by $(x_t, y_t, \Delta x_t, \Delta y_t)$, where (x_t, y_t) is the pick point in pixel

space and $(x_t + \Delta x_t, y_t + \Delta y_t)$ is the place point. Through a known pixel-to-world transform, the robot picks and lifts the top layer of fabric at (x_t, y_t) , translates by $(\Delta x_t, \Delta y_t)$, and releases. The robot seeks to learn a policy $\pi : \mathcal{O} \rightarrow \mathcal{A}$ that maximizes $R(s_T)$, where T denotes the final step of a fabric smoothing episode and $R(\cdot)$ gives the *coverage* or 2D area covered by the fabric (i.e., the area of fabric visible in the workspace). Following prior work [33], we assume that each smoothing episode terminates when the fabric reaches a coverage threshold of at least 92% or a limit of 10 actions (whichever happens first).

2) *Fabric Smoothing Simulation Environment:* To collect demonstrations, we use the oracle corner-pulling policy in the Gym-Cloth simulator [33], which is ideal for fabric manipulation tasks. Gym-Cloth is an OpenAI-Gym-style fabric manipulation environment, which simulates the cloth using structural, shear, and flexion springs. We use this simulator to collect data for training and testing our policies in simulation before switching to deployment in real. The cloth is represented as a 25x25 grid of point masses and is pictured with one light blue side and one dark blue side. The images generated for learning are 224x224 px, of which a completely smooth cloth would occupy 164x164 px. This ratio is reflected in our physical setup. Lastly, when using the Gym-Cloth-simulated data to train fabric smoothing policies, we perform color shifts and augmentations. Specifically, the colors of the background and both sides of the cloth are darkened, converted to gray scale, and blurred by a constant amount to match the images in real. In Fig 3, rollouts in simulation and real show the simulation images adapted to closely resemble real images. We collect a demonstration dataset \mathcal{D} of 2,000 episodes of the oracle corner-pulling policies in Gym-Cloth for training behavior cloning policies.

3) *Fabric Smoothing Physical Environment:* The workspace, depicted in Fig. 1, contains a bimanual ABB YuMi robot with a single tweezer gripper on its left arm; we do not use the right arm. Tweezer grippers are ideal for grasping the cloth as they are able to apply fine point pressure, unlike many standard grippers [33]. The manipulation surface is white and foam-padded to avoid end effector damage during any workspace collisions. The manipulation workspace is marked with blue tape (see Fig. 1), so that its size with respect to the cloth size resembles the setup in the Gym-Cloth simulator. We use a double sided 25 cm x 25 cm cloth with one light brown and one dark brown side. The workspace has an overhead PhotoNeo Phoxi Camera that captures grayscale images of resolution 732 x 1142 px, which are later cropped to specifically display the manipulation workspace, enclosed by the blue tape. To ease cloth perception, we also normalize the physical camera images based on a constant affine transformation to create greater contrast between the cloth and background. We further blur the images to remove noise, making them more closely resemble the smooth images seen in the Gym-Cloth simulator.

Our physical experiments evaluate four repetitions of each policy; for repeatability, these utilize a four fixed initial fabric




Fig. 3: **Example rollouts in simulation and physical experiments.** The top row (left to right) depicts a sample trajectory in simulation, while the bottom row similarly depicts a sample trajectory in real.

configurations, pictured on the supplemental website [34]. During physical experiments, we project fabric pick points onto a color-segmented mask of the fabric, and project place points within the workspace. Similarly, when evaluating learned policies in the simulator, we project pick points onto a mask of the fabric. Additionally, in simulation, episodes are terminated early if the fabric leaves the workspace.

B. Evaluation Metrics

Our experiments evaluate the following switching criteria introduced in Sec. IV to determine when the behavior cloning policy should stop training and be deployed in real:

- 1) (Reward value) Simulation reward with a value-based stopping condition,
- 2) (Reward gradient) Simulation reward with a gradient-based stopping condition,
- 3) (Confidence value) Epistemic uncertainty with a value-based stopping condition, and
- 4) (Confidence gradient) Epistemic uncertainty with a gradient-based stopping condition.

We evaluate the performance of these switching criteria via the following metrics, calculated during deployment in physical experiments after switching (averaged over 4 episodes in each case):

- 1) Final coverage
- 2) Improvement ratio (final coverage / initial coverage)
- 3) Number of actions per episode.

C. Experimental Results

Fig. 4 illustrates the switching calculation, displaying the evaluation metrics and stopping point for each method. Out of a maximum of 200 learning iterations, the stopping points are 111 and 117 for the confidence-based metrics, while they are 171 and 153 for the reward-based metrics. As seen in the figure, the confidence-based switching points occur after fewer learning iterations because the epistemic uncertainty plateaus earlier than the average simulation reward.

Table I summarizes the experimental results. We see that the average final coverage achieved across all four stopping criteria is 88.79%, which exceeds the average initial coverage by a factor of 1.78. Thus, all four switching conditions

Method (iters)	Imp. ratio	Final	Actions
Rew val (171)	1.732 ± 0.104	0.875 ± 0.051	8.75 ± 1.08
Rew grad (153)	1.568 ± 0.102	0.868 ± 0.037	8.25 ± 1.24
Conf val (111)	1.964 ± 0.070	0.937 ± 0.004	6.00 ± 1.27
Conf grad (117)	1.819 ± 0.103	0.872 ± 0.032	8.75 ± 1.08
Final (200)	1.858 ± 0.040	0.880 ± 0.048	7.50 ± 1.30

TABLE I: **Physical experiment results.** We report the final fabric coverage and improvement ratio (final / initial coverage) achieved in physical fabric smoothing experiments by the various switching policies: reward value (rew val), reward gradient (rew grad), confidence value (conf val), and confidence gradient (conf grad). We also compare with the performance after 200 learning iterations, the maximum training budget considered. Results are mean \pm standard error over 4 episodes in each case, and the initial coverage averaged 0.500 ± 0.026 standard error over all cases reported in the table. We also report the average number of actions (where each trajectory terminates upon reaching 10 actions or 92% coverage, whichever occurs first).

are well-correlated with performance in real. Plotting fabric coverage over time in Figure 5, we also see that 200 learning iterations—the maximum number considered—are not necessary to achieve competitive performance. Rather, our switching criteria identify earlier stopping times that yield similar final coverage to the 200-iteration comparison, saving training time.

Comparing results across the switching criteria, we see that the confidence-based metric yields significantly earlier stopping points than the simulation reward metric, which plateaus more slowly. Yet, the reward-based switching points do not result in improved physical fabric coverage relative to the confidence-based switching points. Thus, it may be that the simulation reward metric continues to fit to the simulator beyond the point where such fitting benefits physical task performance. We also see that the epistemic uncertainty metric yields higher final coverage than the simulation reward for both the value-based and gradient-based stopping conditions. Similarly, the value-based stopping condition shows higher performance than the gradient stopping condition for both the epistemic uncertainty and simulation reward metrics.




Fig. 4: **Determining stopping points for various switching criteria.** On all graphs, the dark blue curves are splines fit to the data to mitigate noise when approximating the gradient and determine if the value-based threshold has been met. **Left two graphs:** The simulation reward comes from evaluating the policy in the Gym-Cloth simulation environment and determining the final fabric coverage in each episode; curves are averaged over 5 episode rollouts in Gym-Cloth. For comparison with real, the orange points correspond to mean performance in real of the behavior cloning policy selected at that iteration. The error bars correspond to the standard error across four runs. The stopping points (red point) are determined to be at iteration 171 for reward value and 153 for the reward gradient. **Right two graphs:** The epistemic uncertainty is calculated at each iteration over a holdout set of 200 demonstration episodes and with five ensemble members. The confidence value determines the stopping point to be 111, while the confidence gradient determines it to be 117.

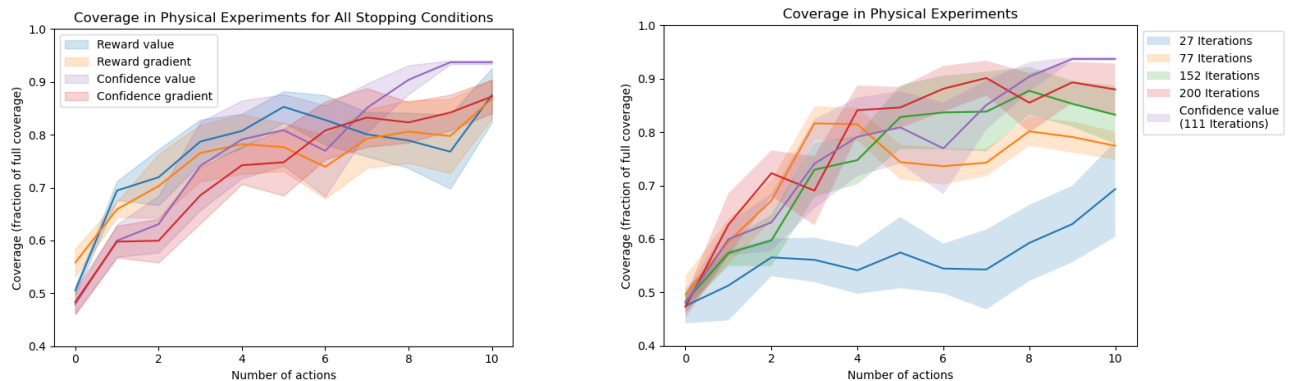


Fig. 5: **Performance of learned policies in physical fabric smoothing experiments at various stopping points.** Left: Final physical fabric coverage achieved for each of the four stopping conditions. Right: Comparing the final physical fabric coverage for the confidence value stopping condition to various checkpoints. We see that the stopping conditions are largely competitive with 200 iterations (the maximum iteration number considered), but require significantly less training. Plots show mean \pm standard error over 4 episodes. Note that to plot episodes that reach the target coverage of 92% in fewer than 10 actions, we repeat the final achieved coverage for the remainder of the 10-action budget.

However, it is possible that these discrepancies partially reflect noise in the behavior cloning training process, rather than an inherent difference between the switching criteria.

Further comparing the methods, Figure 4 suggests that the epistemic uncertainty exhibits less iteration-to-iteration noise than the simulation reward, while the simulation reward is easier to tune and interpret and only requires training a single behavior cloning policy (rather than a policy ensemble). Meanwhile, the value-based stopping condition is easier to tune than the gradient-based method. However, because the gradient-based stopping condition uses a time-based trend rather than individual values, it may yield more stable values and improved generalization compared to the value-based stopping condition. Testing generalizability of the stopping conditions is an interesting direction for future work.

VI. DISCUSSION AND FUTURE WORK

This work proposes and evaluates strategies for determining when to switch from training a learning algorithm in simulation to deploying it on a physical robot. We consider metrics based on simulation performance and epistemic uncertainty and suggest stopping conditions based on absolute thresholding and gradient tracking. We apply the method to sim2real transfer for behavior cloning in a fabric smoothing task. Results suggest that our switching criteria correspond well with performance in real and that the proposed switching criteria can save training time by early stopping, which can save compute resources and help to deploy robots in the real world faster. There are many directions for future work. In particular, we will study how the switching criteria generalize across multiple simulators and tasks. We will also consider switching to real based on policy performance

when transferring across simulators. Finally, we are also excited about switching to real in other algorithmic contexts, such as model-based planning and model-free reinforcement learning, as well as two-way switching between learning in simulation and in physical experiments.

ACKNOWLEDGMENT

The authors would like to thank Mike Danielczuk and Justin Kerr for helpful advice. This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, the CITRIS “People and Robots” (CPAR) Initiative, and the Real-Time Intelligent Secure Execution (RISE) Lab. The authors were supported a CRA Computing Innovation Fellowship and by equipment grants from PhotoNeo, NVidia, and Intuitive Surgical.

REFERENCES

- [1] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Wlinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang, “Solving Rubik’s cube with a robot hand,” *arXiv preprint arXiv:1910.07113*, 2019.
- [2] Y. Du, O. Watkins, T. Darrell, P. Abbeel, and D. Pathak, “Auto-tuned sim-to-real transfer,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1290–1296.
- [3] B. Eysenbach, S. Chaudhari, S. Asawa, S. Levine, and R. Salakhutdinov, “Off-dynamics reinforcement learning: Training for transfer with domain classifiers,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [4] C. Finn and S. Levine, “Deep Visual Foresight for Planning Robot Motion,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [5] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, J. Grannen, M. Hwang, R. Hoque, J. E. Gonzalez, N. Jamali, K. Yamane, S. Iba, and K. Goldberg, “Learning to smooth and fold real fabric using dense object descriptors trained on synthetic color images,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [6] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, R. Hoque, J. E. Gonzalez, and K. Goldberg, “MMGSD: Multi-modal Gaussian shape descriptors for correspondence matching in 1D and 2D deformable objects,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [7] S. Höfer, K. Bekris, A. Handa, J. C. Gamboa, M. Mozifian, F. Golemo, C. Atkeson, D. Fox, K. Goldberg, J. Leonard, et al., “Sim2Real in robotics and automation: Applications and challenges,” *IEEE transactions on automation science and engineering*, vol. 18, no. 2, pp. 398–400, 2021.
- [8] R. Hoque, A. Balakrishna, E. Novoseller, A. Wilcox, D. S. Brown, and K. Goldberg, “ThriftyDagger: Budget-aware novelty and risk gating for interactive imitation learning,” in *Conference on Robot Learning (CoRL)*, 2021.
- [9] R. Hoque, D. Seita, A. Balakrishna, A. Ganapathi, A. K. Tanwani, N. Jamali, K. Yamane, S. Iba, and K. Goldberg, “Visuospatial foresight for multi-step, multi-task fabric manipulation,” in *Robotics: Science and Systems (RSS)*, 2020.
- [10] R. Hoque, A. Balakrishna, C. Putterman, M. Luo, D. S. Brown, D. Seita, B. Thananjeyan, E. Novoseller, and K. Goldberg, “LazyDagger: Reducing context switching in interactive imitation learning,” in *IEEE Conference on Automation Science and Engineering (CASE)*, 2021, pp. 502–509.
- [11] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, “Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 12619–12629.
- [12] R. Jangir, G. Alenyà, and C. Torras, “Dynamic cloth manipulation with deep reinforcement learning,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4630–4636.
- [13] A. Kadian, J. Truong, A. Gokaslan, A. Clegg, E. Wijmans, S. Lee, M. Savva, S. Chernova, and D. Batra, “Sim2real predictivity: Does evaluation in simulation predict real-world performance?” *IEEE Robotics and Automation Letters (RA-L)*, 2020.
- [14] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.
- [15] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “RMA: Rapid motor adaptation for legged robots,” in *Robotics: Science and Systems (RSS)*, 2021.
- [16] A. Kumar, A. Singh, S. Tian, C. Finn, and S. Levine, “A workflow for offline model-free robotic reinforcement learning,” in *Conference on Robot Learning (CoRL)*, 2021.
- [17] R. Lee, D. Ward, A. Cosgun, V. D’Souza, P. Corke, and J. Leitner, “Learning Arbitrary-Goal Fabric Folding with One Hour of Real Robot Experience,” in *Conference on Robot Learning (CoRL)*, 2020.
- [18] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, “Reinforcement learning for robust parameterized locomotion control of bipedal robots,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [19] V. Lim, H. Huang, L. Y. Chen, J. Wang, J. Ichnowski, D. Seita, M. Laskey, and K. Goldberg, “Planar robot casting with real2sim2real self-supervised learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [20] X. Lin, Y. Wang, Z. Huang, and D. Held, “Learning visible connectivity dynamics for cloth smoothing,” in *Conference on Robot Learning (CoRL)*, 2022, pp. 256–266.
- [21] X. Lin, Y. Wang, J. Olkin, and D. Held, “SoftGym: Benchmarking deep reinforcement learning for deformable object manipulation,” in *Conference on Robot Learning (CoRL)*, 2020.
- [22] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, “Dex-Net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics,” in *Robotics: Science and Systems (RSS)*, 2017.
- [23] M. Mahserreci, L. Balles, C. Lassner, and P. Hennig, “Early stopping without a validation set,” *ArXiv*, vol. abs/1703.09580, 2017.
- [24] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, “What matters in learning from offline human demonstrations for robot manipulation,” in *Conference on Robot Learning (CoRL)*, 2021.
- [25] J. Matas, S. James, and A. J. Davison, “Sim-to-Real Reinforcement Learning for Deformable Object Manipulation,” in *Conference on Robot Learning (CoRL)*, 2018.
- [26] B. Mehta, M. Diaz, F. Golemo, C. J. Pal, and L. Paull, “Active domain randomization,” in *Conference on Robot Learning (CoRL)*, 2019.
- [27] F. Muratore, M. Gienger, and J. Peters, “Assessing transferability from simulation to reality for reinforcement learning,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, pp. 1172–1183, 2021.
- [28] P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection,” in *International Journal of Robotics Research (IJRR)*, 2017.
- [29] L. Prechelt, “Automatic early stopping using cross validation: Quantifying the criteria,” *Neural Networks*, vol. 11, no. 4, pp. 761–767, 1998.
- [30] F. T. Ramos, R. Possas, and D. Fox, “Bayessim: Adaptive domain randomization via probabilistic inference for robotics simulators,” in *Robotics: Science and Systems (RSS)*, 2019.
- [31] S. Ross, G. Gordon, and D. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011, pp. 627–635.
- [32] A. A. Rusu, M. Večerík, T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell, “Sim-to-real robot learning from pixels with progressive nets,” in *Conference on Robot Learning (CoRL)*, 2017, pp. 262–270.
- [33] D. Seita, A. Ganapathi, R. Hoque, M. Hwang, E. Cen, A. K. Tanwani, A. Balakrishna, B. Thananjeyan, J. Ichnowski, N. Jamali, et al., “Deep imitation learning of sequential fabric smoothing from an algorithmic supervisor,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9651–9658.
- [34] *Supplemental website*. <https://sites.google.com/view/learning-switching-criteria>.

- [35] A. K. Tanwani, “DIRL: Domain-invariant representation learning for sim-to-real transfer,” in *Conference on Robot Learning (CoRL)*, 2020.
- [36] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 23–30.
- [37] E. Valassakis, Z. Ding, and E. Johns, “Crossing the gap: A deep dive into zero-shot sim-to-real transfer for dynamics,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5372–5379.
- [38] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, “Learning to manipulate deformable objects without demonstrations,” in *Robotics: Science and Systems (RSS)*, 2020.
- [39] Y. Yao, L. Rosasco, and A. Caponnetto, “On early stopping in gradient descent learning,” *Constructive Approximation*, vol. 26, no. 2, pp. 289–315, Aug. 2007.
- [40] W. Zhao, J. P. Queralta, and T. Westerlund, “Sim-to-real transfer in deep reinforcement learning for robotics: A survey,” in *IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.