

Efficient Tactile Simulation with Differentiability for Robotic Manipulation

Anonymous Author(s)

Affiliation

Address

email

1 **Abstract:** Efficient simulation of tactile sensors can unlock new opportunities for
2 learning tactile-based manipulation policies in simulation and then transferring the
3 learned policy to real systems, but fast and reliable simulators for dense tactile normal
4 and shear force fields are still under-explored. We present a novel approach
5 for efficiently simulating both the normal and shear tactile force field covering
6 the entire contact surface with an arbitrary tactile sensor spatial layout. Our simulator
7 also provides analytical gradients of the tactile forces to accelerate policy
8 learning. We conduct extensive simulation experiments to showcase our approach
9 and demonstrate successful zero-shot sim-to-real transfer for a high-precision peg-
10 insertion task with high-resolution vision-based GelSlim tactile sensors.

11 **Keywords:** Tactile Simulation, Tactile Manipulation, Sim-to-Real

12 1 Introduction

13 Just as humans heavily rely on the rich and precise tactile cues for dexterous grasping and in-hand
14 manipulation tasks, robots can also utilize tactile cues as an important source of sensing for inter-
15 acting with the surrounding environments, especially when the visual information is unavailable
16 or occluded. With the recent development of various tactile sensors capable of generating dense
17 normal or shear load information [1, 2, 3, 4], researchers have been exploring how to leverage this
18 important mode of information for robotic manipulation tasks. With the dense tactile *normal* load
19 field, the static spatial relation between the object and the robot manipulators can easily be inferred,
20 which is useful for tasks such as edge following [5], pose estimation [6], object reconstruction and
21 recognition [7, 8]. On the other hand, the dense tactile *shear* force feedback more readily gives
22 rich information about the dynamic tangential motions between the object and the manipulators,
23 and thus can be utilized in tasks such as stable grasp [9], precise insertion [10, 11], and slip de-
24 tection [12, 13, 14, 15]. However, most of the tactile manipulation work still requires significant
25 amount of human effort on real hardware system for collecting data, cleverly building automatic
26 resetting mechanism, and carefully designing the learning strategy [10]. Such manual work can be
27 time-consuming, cost expensive, and more importantly unsafe during policy exploration.

28 Due to its capability to replicate the real world with high fidelity and low cost, physics-based simu-
29 lation has become a powerful recipe for learning robotic control policies [16, 17, 18, 19]. Previous
30 work has demonstrated that the policy can be efficiently learned in simulation and successfully
31 transferred to real robots via proper sim-to-real techniques [20, 21, 22]. Despite the prevalence of
32 simulation and the importance of tactile sensory in robotics, physics-based simulation is still under-
33 explored to efficiently simulate dense tactile normal and shear force fields for robotic applications.
34 Most popular simulators [18, 19] only support force-torque sensors which are attached to each robot
35 link, only producing the contact force values at a few points on each body. Although one can ac-
36 quire a dense tactile force field via attaching many small cuboids to the robot body and querying
37 the force sensor on each small cuboid from simulation [23, 24], the obtained tactile force values are
38 usually sparse and are unable to match the uniform force distribution on a real elastic tactile sensor
39 such as GelSlim [1]. While researchers have also tried simulating realistic tactile feedback via pure

40 geometric methods [5, 25], such methods typically only compute the normal tactile force and can-
41 not simulate the tactile effects in shear directions. On the other hand, the tactile shear forces have
42 been successfully simulated via finite element method (FEM) [26, 27, 28] or data-driven approach
43 [29], but these simulators suffer from expensive computation costs, and cannot be easily used for
44 data-hungry policy-learning approaches such as reinforcement learning (RL).

45 We present a novel tactile simulator that can efficiently and reliably simulate both normal and shear
46 tactile force fields covering the entire contact surface. We build upon rigid body dynamics formula-
47 tion and develop a fast penalty-based tactile model which can run at 1000 frames/s on a single core
48 of Intel i7-9700 CPU. Our tactile model can reasonably approximate the soft contact nature of soft
49 tactile sensor material such as the elastomer used in GelSlim [30], generate dense tactile force fields
50 (*e.g.*, the dense marker array on GelSlim), and is compatible with arbitrary tactile sensor spatial lay-
51 out (*i.e.*, flat plane, hemisphere, etc.). Furthermore, our compact tactile formulation is differentiable,
52 which allows the simulator to provide fast analytical gradients for the entire dynamics chain. We
53 conduct extensive experiments in simulation to demonstrate the capabilities of our tactile simulator,
54 including policy learning with reinforcement learning algorithms and gradient-based algorithms. We
55 also conduct a zero-shot sim-to-real experiment for a high-precision tactile-based peg-insertion task,
56 demonstrating that our simulator provides realistic tactile simulation.

57 2 Related Work

58 While there have been many physics-based simulators developed to simulate various types of robots,
59 efficiently and reliably simulating dense tactile sensing fields is less explored. As mentioned above,
60 most robotics simulators such as MuJoCo [18] and PyBullet [19] only support force-torque sen-
61 sors that are attached to each robot link. While it is possible to augment these simulators with
62 high-resolution tactile forces, they become computationally cumbersome. In order to acquire more
63 realistic and dense tactile forces, Narang et al. [26, 27] and Ding et al. [28] use soft materials to
64 model the tactile sensors and apply finite element method (FEM) to simulate the deformation and
65 force fields of the tactile sensors. Despite its high fidelity of the simulated tactile feedback, these
66 simulators suffer from expensive computation cost and are primarily used to collect supervised tac-
67 tile dataset instead of learning policies which are typically data-hungry and requires fast simulations.
68 Vision-based tactile sensors produce high-resolution tactile feedback. To simulate vision-based sen-
69 sors, Wang et al. [25] and Church et al. [5] use PyBullet [19] and render the intersecting part between
70 the object and the tactile manipulator as depth images, from which tactile information is generated.
71 However such purely geometry-based approaches cannot simulate the tactile effects in the shear
72 directions such as the marker displacements of GelSlim. Si and Yuan [29] compute the marker
73 displacement field by presenting a superposition method to approximate the FEM dynamics. While
74 they are able to simulate the tactile shear effects, the speed of the simulation is still slow, and no con-
75 trol tasks are demonstrated. Bi et al. [31] build an efficient simulation specialized for a tactile-based
76 pole swing-up task with a customized vision-based tactile sensor, but the proposed technique is not
77 readily extensible to simulate other types of tasks and tactile sensor types. Similarly to our work,
78 Habib et al. [32] use a spring-mass-damper model to simulate tactile normal forces, and Moio
79 et al. [33] use a soft bristle deflection model for simulating the tactile forces. However, there is no
80 control tasks demonstrated to be learned with the presented simulators and no gradients information
81 is available. In contrast to these previous works, we present a generic simulator with analytical gra-
82 dients for tactile forces by leveraging the penalty-based rigid body dynamics, and we demonstrate
83 that our simulation is efficient enough for policy learning, and simulated tactile force field can be
84 successfully used for a sim-to-real task on the high-resolution vision-based GelSlim sensor.

85 3 Method

86 We now present our approach to simulate tactile forces for real-world tactile sensors. In §3.1, we
87 introduce our flexible representation for tactile sensors. In §3.2-3.3, we present our penalty-based
88 tactile model for simulation and derive the analytical gradients of the dynamics. In §3.4, we describe
89 our intermediate tactile signal representation for the sim-to-real transfer of the policies.

90 3.1 Tactile Sensor Representation

91 Each tactile point i on a sensor pad is represented by a tuple $\langle \mathbf{B}_i, \mathbf{E}_i, \xi_i \rangle$
 92 as shown in Fig. 1. \mathbf{B}_i is the rigid body the tactile point is attached to,
 93 and $\mathbf{E}_i \in \text{SE}(3)$ is the position/orientation of the point in the local coord-
 94 inate frame of the body, with the x_i and y_i axes in the shear-direction
 95 plane and the z_i axis along the normal tactile direction. (These axes are
 96 the same for all points for a *planar* sensor pad.) Finally, ξ_i are the sim-
 97 ulation parameters of the penalty-based tactile model, which will be introduced later in §3.2. Our
 98 representation for tactile points is flexible, allowing us to specify any number of points in arbitrary
 99 geometry layouts on a robot, and each tactile sensor can have its individual configuration parameters.




Figure 1: Tactile Sensor Representation.

100 3.2 Penalty-based Tactile Model

101 We use a penalty-based tactile model to characterize the force on each tactile point. For each point
 102 $\langle \mathbf{B}_i, \mathbf{E}_i, \xi_i \rangle$, we use the following contact model [34] to obtain the contact force at the tactile point’s
 103 location represented in the local coordinate frame \mathbf{B}_i . (For brevity, we drop the subscript i .)

$$\mathbf{f}_n = (-k_n + k_d \dot{d}) \mathbf{n}, \quad \mathbf{f}_t = -\frac{\mathbf{v}_t}{\|\mathbf{v}_t\|} \min(k_t \|\mathbf{v}_t\|, \mu \|\mathbf{f}_n\|), \quad (1)$$

104 where \mathbf{f}_n is the contact force at the tactile point along the contact normal direction \mathbf{n} , and \mathbf{f}_t is the
 105 contact friction force in the plane tangential to the contact normal direction. The scalar d (nonposi-
 106 tive) is the penetration depth between the point and the collision object, and \dot{d} is its time derivative.
 107 The vector \mathbf{v}_t is the relative velocity at the contact point along the contact tangential direction.
 108 Scalars k_n, k_d, k_t, μ are contact stiffness, contact damping coefficient, friction stiffness, and coeffi-
 109 cient of friction respectively, and they together form the simulation parameters ξ of the tactile point:
 110 *i.e.*, for the i^{th} tactile point, $\xi_i = \{k_n^i, k_d^i, k_t^i, \mu^i\}$. After the frictional contact force is computed for
 111 each point as $\mathbf{f} = \mathbf{f}_n + \mathbf{f}_t$, we transform this force into the local coordinate frame of the tactile
 112 point to acquire the desired *shear* and *normal* tactile force magnitudes:

$$T_{sx} = \mathbf{f}^\top \mathbf{x}, \quad T_{sy} = \mathbf{f}^\top \mathbf{y}, \quad T_n = \mathbf{f}^\top \mathbf{z}, \quad (2)$$

113 where $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are the axes of frame \mathbf{E} .

114 Our penalty-based tactile model can be integrated into any simulator as long as the required values,
 115 such as the world-frame location of the tactile points, the contact normal, the collision penetration
 116 depth and its time derivatives, can be acquired from the simulator. We implement our tactile model in
 117 C++ and integrate it into differentiable RedMax (DiffRedMax) [34, 35] since DiffRedMax is open-
 118 source and readily provides all the required information for our computation, and more importantly,
 119 its differentiability allows us to make our tactile simulation differentiable with a moderate amount
 120 of modifications to its backward gradients computation.

121 3.3 Differentiable Tactile Simulation

122 Since we use an implicit time integration scheme for forward dynamics, the core step of gradients
 123 computation is to differentiate through the nonlinear equations of motion. We start by formulating a
 124 finite-horizon tactile-based policy optimization problem:

$$\text{minimize}_{\theta} \mathcal{L} = \sum_{t=1}^H \mathcal{L}_t(\mathbf{u}_t, \mathbf{q}_t, \mathbf{v}_t(\mathbf{q}_t)) \quad (3a)$$

$$\text{s.t. } g(\mathbf{q}_{t-1}, \dot{\mathbf{q}}_{t-1}, \mathbf{u}_t, \mathbf{q}_t) = 0 \quad (\text{Equations of Motion}) \quad (3b)$$

$$\mathbf{u}_t = \pi_{\theta}(\tilde{\mathbf{q}}_{t-1}, \tilde{\mathbf{v}}_{t-1}(\mathbf{q}_{t-1}), T_{t-1}(\mathbf{q}_{t-1}, \dot{\mathbf{q}}_{t-1})). \quad (\text{Policy Execution}) \quad (3c)$$

125 Here, H is the task horizon, \mathcal{L}_t is a step-wise task-dependent reward function, \mathbf{u} is the action
 126 (*e.g.*, joint torque), \mathbf{q} is the simulation state (*i.e.*, joint angles), and \mathbf{v} is the derived auxiliary sim-
 127 ulation variables (*e.g.*, fingertip positions) which themselves are a function of \mathbf{q} . Eq. 3b describes
 128 the nonlinear equations of motion (§A.1). Eq. 3c represents the inference of the control policy π_{θ} to
 129 obtain the desired action given the partial observation of the simulation state $\tilde{\mathbf{q}}$, partial observation
 130 of the simulation computed variables $\tilde{\mathbf{v}}$, and the tactile force values T from Eq. 2.

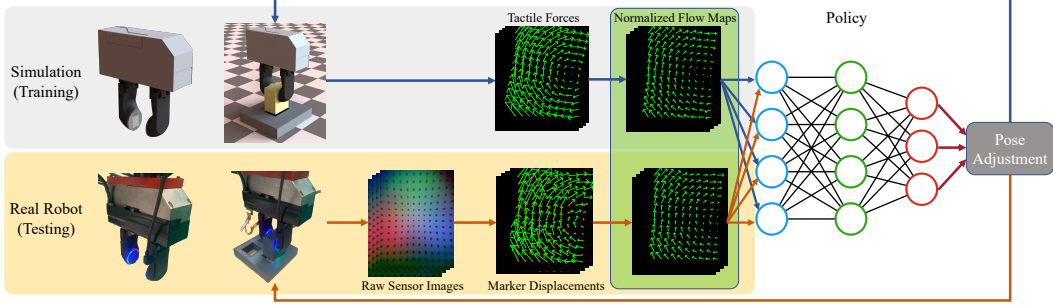


Figure 2: Sim-to-Real Pipeline for Insertion Task (§4.5). *Gray Box:* During training, we convert the tactile force output from the simulator into the normalized flow map representation (shaded in green). *Yellow Box:* When executing the policy on a real system, we convert the sensor output into the same normalized flow map. This intermediate representation is then treated as the observation input to a neural network policy to output the pose adjustment for the next attempt. Here we only visualize the tactile output from one tactile sensor pad.

131 We compute the gradients $d\mathcal{L}/d\theta = (d\mathcal{L}/d\mathbf{u}_t)(d\mathbf{u}_t/d\theta)$ for policy optimization. We embed our
 132 simulator as a differentiable layer into the PyTorch computation graph and use reverse mode differ-
 133 entiation to backward differentiate through dynamics time integration. The first gradient, $d\mathcal{L}/d\mathbf{u}_t$,
 134 which includes the tactile derivatives, is derived analytically, as shown below. The second gradient,
 135 $d\mathbf{u}_t/d\theta$, is computed by PyTorch’s auto-differentiation.

136 At each time step t , we derive $\partial\mathcal{L}/\partial\mathbf{u}_t$ given the analytically computed gradients with respect to the
 137 system states, auxiliary variables, and tactile forces ($\partial\mathcal{L}/\partial\mathbf{q}_t, \partial\mathcal{L}/\partial\mathbf{v}_t, \partial\mathcal{L}/\partial T_t$; see §A.2-A.3):

$$\frac{d\mathcal{L}}{d\mathbf{u}_t} = \underbrace{\frac{\partial\mathcal{L}_t}{\partial\mathbf{u}_t}}_{\mathbf{a}} + \underbrace{\left(\frac{\partial\mathcal{L}}{\partial\mathbf{q}_t} + \frac{\partial\mathcal{L}}{\partial\mathbf{v}_t} \frac{\partial\mathbf{v}_t}{\partial\mathbf{q}_t} + \frac{\partial\mathcal{L}}{\partial T_t} \left(\frac{\partial T_t}{\partial\mathbf{q}_t} + \frac{\partial T_t}{\partial\dot{\mathbf{q}}_t} \frac{\partial\dot{\mathbf{q}}_t}{\partial\mathbf{q}_t} \right) \right)}_{\mathbf{b}} \underbrace{\frac{\partial\mathbf{q}_t}{\partial\mathbf{u}_t}}_{-A^{-1}D}. \quad (4)$$

138 The right-most derivative can be computed by applying the implicit function theorem on Eq. 3b,
 139 which gives us $\partial\mathbf{q}_t/\partial\mathbf{u}_t = -(\partial g/\partial\mathbf{q}_t)^{-1}(\partial g/\partial\mathbf{u}_t)$. Writing this as $\partial\mathbf{q}_t/\partial\mathbf{u}_t = -A^{-1}D$ and
 140 combining with Eq. 4, we first solve the linear system $A^\top \mathbf{c} = -\mathbf{b}^\top$ for \mathbf{c} , and then we compute the
 141 final gradient as $d\mathcal{L}/d\mathbf{u}_t = \mathbf{a} + \mathbf{c} \cdot D$ using the adjoint approach.

142 3.4 Normalized Tactile Flow Map for Sim-to-Real

143 We use the GelSlim 3.0 sensor [4], which utilizes small markers to track motions in the shear direc-
 144 tion, to demonstrate the sim-to-real capability (§4.5). There is an unavoidable sim-to-real gap be-
 145 tween our simulator emulating tactile forces and the physical GelSlim sensor that relies on imaging.
 146 In this section, we demonstrate how we overcome this gap by constructing a common intermediate
 147 tactile representation for policy input observation. We assume that the stiffness of the sensor along
 148 different shear directions is isotropic and that there exists a linear relationship between the displace-
 149 ment and the contact *shear* forces (T_{sx} and T_{sy} from Eq. 2) at each tactile point. We connect these
 150 two different sensor output formats via a unitless normalized tactile flow map representation.

151 Specifically, we use the raw tactile sensor images from the past k steps from the n tactile sensor pads
 152 on a real robot as our policy observation $T_{\text{image}}^{\{1:k,1:n\}}$. As shown in Fig. 2, we first detect and identify
 153 the marker positions in each image, and obtain the marker displacement field $T_{\text{displacement}}^{\{1:k,1:n\}} \in \mathbb{R}^{r \times c \times 2}$
 154 by subtracting the marker positions in the rest configuration from their positions in the deformed
 155 configuration, where r and c are the rows and columns of the tactile marker array in each sensor
 156 pad, with each marker giving us the x and y displacement information. Then we normalize the
 157 displacement field so that the maximal length of the marker displacement across all tactile points
 158 and images is of unit length; *i.e.*,

$$T_{\text{normalized}}^{\{1:k,1:n\}} = \frac{T_{\text{displacement}}^{\{1:k,1:n\}}}{\max(\max_{k,n,r,c}(\|T_{\text{displacement}}^{\{k,n\}}(r,c)\|), \epsilon)}, \quad (5)$$

159 where ϵ ensures that the output is zero when there is no any displacement on the markers (*i.e.*, no
 160 contact). We concatenate these flow maps into a single tensor $T_{\text{normalized}} \in \mathbb{R}^{k \times n \times r \times c \times 2}$, which is

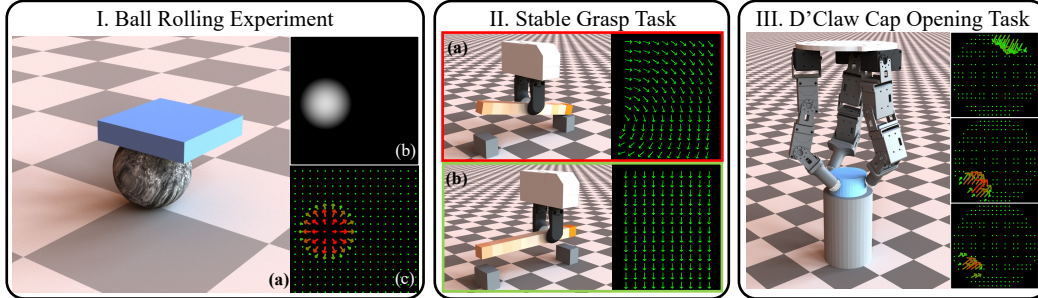


Figure 3: **(I) Ball rolling experiment:** The tactile sensors are installed on the lower surface of the pad. The depth map of the tactile normal forces is shown in (b). The tactile force field is shown in (c) with the arrow denoting the shear forces and the color denoting the magnitude of the normal force. **(II) Stable Grasp Task:** The bar composed of 11 blocks with random densities (the deeper the color, the heavier the block). (a) An unsuccessful grasp results in rotational patterns in the tactile force field and (b) a successful grasp requires the gripper to adjust the grasp location to the center of mass of the bar. **(III) D’Claw Cap Opening Task:** The tactile sensors (white dots) are installed at the three hemisphere fingertips of the hand. We map each tactile point at one fingertip onto a 2D image plane and visualize the tactile forces field of three fingertips on the right.

161 our normalized tactile flow map representation. For the tactile shear forces $\{T_{sx}^{\{1:k,1:n\}}, T_{sy}^{\{1:k,1:n\}}\}$
 162 acquired from the simulation, we conduct the same normalization process as Eq. 5.

163 Intuitively speaking, the normalized tactile flow map provides the directional information about the
 164 relative motion of the markers induced by the contact forces, and it also keeps the relative tactile
 165 load magnitude relationships among different sensors and different time steps so as to preserve the
 166 meaningful spatial and temporal information about the contact. For our sim-to-real experiments,
 167 we only use the shear directional information from the sensor, but the same technique can also be
 168 applied to the normal directional information via normalizing the depth map of the contact surface
 169 reconstructed from the GelSlim image [4] across different frames and different sensor pads.

170 4 Experiments

171 We conduct extensive tactile-based experiments to demonstrate the capability of our approach.¹
 172 We investigate the following questions: (§4.1, §4.2) Can our simulator reliably simulate the high-
 173 resolution tactile force field at a high speed for RL algorithms? (§4.3) Does the differentiability
 174 of our simulator provide advantages in policy learning? (§4.4) Is our tactile sensor representation
 175 flexible enough for sensors with arbitrary geometrical layouts? (§4.5) How does our simulated
 176 tactile force field compare to the tactile feedback from real sensors, and does our normalized tactile
 177 flow map representation help to transfer the policies learned in the simulator to a real robot?

178 4.1 Speed and Reliability: High-Resolution Tactile Ball Rolling Experiment

179 We design a ball rolling experiment to show the efficacy of the tactile force field generated by our
 180 simulator and to test the simulation speed. The simulation setup is shown in Fig. 3(I). A high-
 181 resolution tactile pad (200×200 markers) touches the marble ball and moves it around. The simu-
 182 lation step size $h = 5$ ms, and we compute the tactile force field every 5 steps (*i.e.*, 40 Hz). Fig. 3(I)
 183 also shows the normal tactile force (represented by a depth map) and the tactile shear forces ac-
 184 quired from our simulator. For this example, our simulation runs at 1050 frames per second (FPS)
 185 on a single core of Intel Core i7-9700K CPU. The simulation speed can be further accelerated by
 186 simply parallelizing it across multiple CPU cores, as we do in the RL experiments.

187 4.2 RL Training: Tactile-Based Stable Grasp Task

188 Our tactile simulator provides the shear force information on the contact surfaces, which is criti-
 189 cal for many manipulation tasks. Inspired by the setup in [9], we show the usage of shear force
 190 information for control and the effectiveness of our tactile simulator in a parallel-jaw grasping
 191 task. As shown in Figure 3(II), the task requires a WSG-50 parallel-jaw gripper to stably grasp a bar

¹See also the supplementary video.

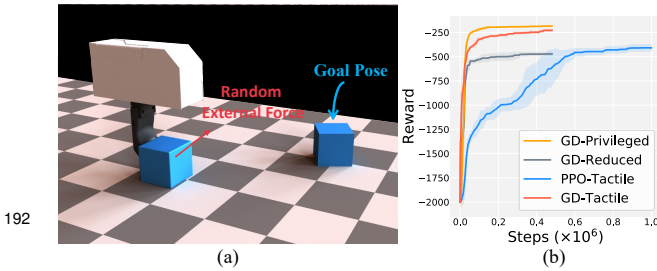


Figure 4: Tactile-based box pushing task. (a) The goal of the gripper policy is to use its tactile feedback to push a box to a randomized target position/orientation. A time-varying external force is randomly applied on the box during the task. (b) the training curve for each policy variation is averaged from the five independent runs with different random seeds.

	POS. ERROR	ORI. ERROR
<i>GD-Privileged</i>	$0.037 \pm 0.002m$	$0.043 \pm 0.003^\circ$
<i>GD-Reduced</i>	$0.126 \pm 0.009m$	$0.255 \pm 0.021^\circ$
<i>PPO-Tactile</i>	$0.123 \pm 0.034m$	$0.241 \pm 0.123^\circ$
<i>GD-Tactile</i>	$0.058 \pm 0.003m$	$0.074 \pm 0.020^\circ$

Table 1: Metrics comparison on box pushing task. We compute the final position/orientation errors of the best policy in each run and average the metrics from five runs for each policy variation. *GD-Privileged* gives a reference of the best possible metrics, and without the privileged state information of the box, our *GD-Tactile* achieves much better position error and rotation error than other two variations.

with *unknown mass distribution* in fewer than 10 attempts. The gripper has two tactile sensors with a tactile marker resolution of 13×10 . The bar is composed of 11 blocks where the density of each block is randomized. The total mass of the bar ranges in $[45, 110]$ g. We consider a grasp to be a failure if the bar tilts more than 0.02 rad after the gripper grasps a bar.

We use RL to train a control policy that determines the grasp location. The initial grasp location is the geometric center of the bar. Based on the tactile sensor readings (the only observation input to the policy), the policy outputs a delta change in the grasping location. The policy is a shallow CNN (*Conv-ReLU-MaxPool-Conv-ReLU-FC-FC*) that takes as input the two tactile sensor readings. We train the policies with PPO [36] using 32 parallel environments with 20K environment steps in total. We train the policies with 3 different random seeds and test them 320 times. The success rate is $93.2 \pm 1.6\%$. The average number of attempts taken to stably grasp the bars is 2.1, meaning the policy can compute the correct grasping location after a single failed attempt in most cases.

4.3 Differentiability: Tactile-Based Box Pushing Task

In this experiment, we design a box pushing task similar to [5] to demonstrate how we can leverage the provided analytical gradients to help learn tactile-based control policies better and faster.

Task Specification As shown in Fig. 4(a), the task here is to use the same WSG-50 parallel jaw gripper as §4.2 (with only one finger kept) to push the box to a randomly sampled goal location and orientation. The initial position of the box is randomly disturbed. A random external force is applied continually on the box, which changes every 0.25 s. More details of the task are in §B.3.

Comparing Policy Learning Algorithms We train the control policies through four different combinations of learning algorithms and observation spaces. *GD-Privileged*: This variation uses the gradient-based optimizer Adam by utilizing the analytical policy gradients computed from our differentiable simulation. The policy observation contains all the privileged state information of the gripper, the box, and the goal. This policy provides an upper-bound performance reference. *GD-Reduced*: Similar to *GD-Privileged*, except that the observation space only contains the state information that can be acquired on a real system such as the gripper state and the goal. *GD-Tactile*: Other than the state information used in *GD-Reduced*, we also include the tactile sensor readings in the policy input. The policy is trained using the analytical policy gradients. *PPO-Tactile*: Similar to *GD-Tactile*, but the policy is trained by PPO.

All the policies are trained to maximize the same reward function (§B.3). We run each variation five times with different random seeds, and plot their training curves averaged from five runs in Fig. 4(b). We also randomly sample 300 goal poses and measure the final position and orientation errors between the box and the goal pose of the best policy from each run, and report the average metrics across five seeds in Table 1. The results show that when neither state information of the box or tactile information is available (*i.e.*, *GD-Reduced*), the policy cannot reliably push the box to the target location since the gripper has no any clue when the box goes outside of the control of

229 the gripper due to the random initial box position perturbation and the random external forces. With
230 tactile information feedback (*i.e.*, *PPO-Tactile* and *GD-Tactile*), the gripper has this tactile infor-
231 mation to keep the gripper touching the box and allowing it to push the box to the goal effectively.
232 However, the high dimensional tactile observation space results in higher computational cost with
233 PPO which relies on stochastic samples to estimate the policy gradients. In contrast, with the help
234 of our differentiable tactile simulation, *GD-Tactile* makes use of the analytical policy gradients and
235 leads to faster policy learning and better policy performance.

236 4.4 Flexibility: Tactile Sensor Simulation on Curved Surfaces

237 To demonstrate that our method supports tactile sensors on curved surfaces, we train a D’Claw [37]
238 tri-finger hand to open a cap on a bottle. We put the tactile sensors on the three rounded fingertips
239 as shown in Fig. 3(III). The sensor layout on each fingertip is a hemisphere, and we use evenly-
240 spaced 302 tactile markers. We build a coordinate mapping to project the marker positions on the
241 3D surface into a 20×20 2D array (with some empty values around the boundary). Fig. 3(III) shows
242 that our simulation can produce reliable and realistic tactile sensor readings on a rounded fingertip.

243 The task is to open a cap using the D’Claw hand. The position and the radius of the cap are random-
244 ized and unknown. There is also unknown random damping between the cap and the bottle. The task
245 is considered a success if the cap is rotated by 45° . The only observation data that the policy gets
246 are the angles of each joint, fingertip positions, and the tactile sensor readings. This task is similar
247 to how we open caps by just using proprioception sensory data and tactile feedback on the fingers
248 without knowing the exact size and location of the cap. The policy outputs the delta change on the
249 joint angles. We again use PPO to train the policy (a shallow CNN) using 32 parallel environments.
250 To show that the tactile sensors are useful in this task, we also train a baseline policy (a simple MLP
251 policy) where the policy only takes as input the joint angles and fingertip positions. With tactile
252 sensor readings, policies learn significantly faster and achieve an 87.3% success rate, while policies
253 only achieve a 59.7% success rate when tactile sensor information is unavailable. More details about
254 task setup and results are provided in §B.4.

255 4.5 Zero-Shot Sim-to-Real: Tactile RL Insertion Task

256 In this experiment, we show the quality of the simulated tactile feedback when compared to the
257 tactile sensing obtained from the real system, and demonstrate how to do zero-shot transfer for the
258 policies learned in simulation to the real robot via our normalized tactile flow map representation.

259 **Task Specification** We experiment on the tactile-RL insertion task similar to [10]. In this task, a
260 gripper (same as in §4.2) is controlled to insert a cuboid object into a rectangle-shaped hole with
261 a random initial pose misalignment. The insertion process is modeled as an episodic policy that
262 iterates between open-loop insertion attempts followed by insertion pose adjustments (shown in
263 Fig. 2). The robot has up to 15 pose correction attempts, and the robot only has access to tactile
264 feedback from the sensors installed on both gripper fingers. For the real robot system, we use a
265 6-DoF ABB IRB 120 robot arm with a WSG-50 parallel jaw gripper. On each side of the gripper
266 finger, we mount a GelSlim 3.0 tactile sensor that captures the tactile interaction between the fingers
267 and the grasped object as a high resolution tactile image. More details are in §B.5.

268 This task is more challenging than the stable grasp task (§4.2) because it not only needs to recognize
269 the rotational pattern of the tactile field when the object contacts the front/back edges of the hole, but
270 also needs to leverage the different magnitude relationship of two sensors’ outputs to tell whether
271 the object hits the left or the right hole edge (Fig. 5). It becomes even more challenging when the
272 object hits the hole at four corners, because the robot must recognize nuances in the tactile pattern
273 to decide the insertion pose adjustment. Therefore, this task requires a high-quality simulated tactile
274 force field in order to transfer the learned policy to the real system successfully.

275 **Policy Learning via RL** We train the control policies with PPO [36] for three types of misalign-
276 ments. *Rotation*: The object is initialized at the hole’s center and has random rotation misalignment
277 around the vertical axis. The action of the policy is the angle adjustment in the next insertion at-
278 tempt. *Translation*: The object has randomly initialized translation misalignment to the hole and no




Figure 5: Comparison of the normalized tactile flow maps. The flow maps in the top blue boxes are from simulation (with noise added), while the flow maps in the bottom red boxes are produced from the real GelSlim sensor. In each box, the two flow maps (left and right) are for two tactile pads on the two gripper fingers.

279 rotation misalignment. The action space in this case is two dimensional for the translational correc-
 280 tion on the x - y plane. *Rotation & Translation*: The object has both rotation and translation initial
 281 misalignments. The action space of the policy is three dimensional.

282 During training, we convert the simulated tactile force field into our normalized tactile flow map rep-
 283 resentation (§3.4), and treat the resulting tactile flow map as a 13×10 flow “image” with 4 channels
 284 (2 sensors and 2 shear components of tactile forces). We model the policy by a convolutional RNN
 285 to leverage more information from previous attempts. For better sim-to-real performance, we also
 286 apply the domain randomization technique [38] on contact parameters, tactile sensor parameters,
 287 grasp forces, grasp height and tactile readings, to increase the robustness of the learned policies.
 288 More details of policy learning are provided in §B.5.

289 **Experiment Results** We first qualitatively compare the normalized
 290 tactile flow maps generated by simulation and by real GelSlim sen-
 291 sors. We plot the normalized tactile flow maps at four representative
 292 contact configurations (*i.e.*, object contacts at different edges of the
 293 hole) [10] in Fig. 5, which shows that our simulation is able to pro-
 294 duce highly realistic patterns in those contact configurations. We then
 295 deploy the policies learned in simulation on real hardware and quan-
 296 titatively test its zero-shot performance by conducting 100 insertion
 297 experiments under different initial pose misalignments. As reported in Table 2, our zero-shot pol-
 298 icy transfer achieves 100% success rates on *Rotation* and *Translation* tasks. We also calculate the
 299 average number of pose corrections for the successful experiments. The average number of pose
 300 corrections is 1.53 for the *Rotation* task and 2.33 for the *Translation* task, which means that the pol-
 301 icy is able to successfully infer the pose misalignment after just one or two failed attempts in most
 302 experiments. Given the policies being purely trained with simulated tactile data, the high success
 303 rates indicate that our simulation is able to produce normalized tactile flow maps with highly real-
 304 istic tactile pattern and magnitude to help the gripper to infer the exact adjustment. For challenging
 305 *Rotation & Translation* task, our zero-shot transferred policy also achieves 83% success rate and
 306 4.81 pose corrections in average. For comparison, Dong et al. [10] achieve 89.6% success rate and
 307 5.42 times pose adjustments for the cuboid object, with a policy trained directly on the real hardware
 308 from a pre-trained policy and with a carefully designed task curriculum. Our policy is trained from
 309 scratch only in simulation without observing any real-world data.

TASK	SUCCESS	ATTEMPTS
<i>R</i>	100%	1.53
<i>T</i>	100%	2.33
<i>R&T</i>	83%	4.81

Table 2: Zero-shot sim-to-real performance of the tactile RL insertion policies.

310 5 Limitations and Future Work

311 We presented an efficient differentiable simulator that can handle dense tactile force fields with
 312 both normal and shear components. When the tactile pad is *very soft* (*e.g.*, TacTip), its dynamics
 313 cannot be well approximated by our penalty-based approach. An interesting direction to explore is
 314 how to efficiently simulate such soft tactile sensors. We demonstrated with the box pushing task
 315 (§4.3) the potential advantage of differentiable simulators. However, how to effectively leverage
 316 analytical gradients for more complex tactile-based tasks is still an open question and it may require
 317 more advanced policy learning algorithms [39]. Furthermore, in the sim-to-real experiment (§4.5),
 318 the zero-shot success rate of *Rotation & Translation* is not perfect. This is probably due to some
 319 intricacies of the real hardware that are difficult to model in our simulation. We believe that further
 320 fine-tuning the learned policies with a few shots on the real hardware will likely lead to improved
 321 performance.

References

- 322
- 323 [1] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez. Gelslim: A high-resolution,
324 compact, robust, and calibrated tactile-sensing finger. In *2018 IEEE/RSJ International Con-*
325 *ference on Intelligent Robots and Systems (IROS)*, pages 1927–1934. IEEE, 2018.
- 326 [2] S. Sundaram, P. Kellnhofer, Y. Li, J.-Y. Zhu, A. Torralba, and W. Matusik. Learning the
327 signatures of the human grasp using a scalable tactile glove. *Nature*, 569(7758):698–702,
328 2019.
- 329 [3] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos,
330 A. Byagowi, G. Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution
331 tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*,
332 5(3):3838–3845, 2020.
- 333 [4] I. Taylor, S. Dong, and A. Rodriguez. Gelslim3. 0: High-resolution measurement of shape,
334 force and slip in a compact tactile-sensing finger. *arXiv preprint arXiv:2103.12269*, 2021.
- 335 [5] A. Church, J. Lloyd, R. Hadsell, and N. F. Lepora. Optical tactile sim-to-real policy transfer
336 via real-to-sim tactile image translation. *arXiv preprint arXiv:2106.08796*, 2021.
- 337 [6] M. Bauza, A. Bronars, and A. Rodriguez. Tac2pose: Tactile object pose estimation from the
338 first touch. *arXiv preprint arXiv:2204.11701*, 2022.
- 339 [7] E. Smith, R. Calandra, A. Romero, G. Gkioxari, D. Meger, J. Malik, and M. Drozdal. 3d shape
340 reconstruction from vision and touch. *Advances in Neural Information Processing Systems*, 33:
341 14193–14206, 2020.
- 342 [8] S. Suresh, Z. Si, J. G. Mangelson, W. Yuan, and M. Kaess. Shapemap 3-d: Efficient shape
343 mapping through dense touch and vision.
- 344 [9] R. Kolamuri, Z. Si, Y. Zhang, A. Agarwal, and W. Yuan. Improving grasp stability with rotation
345 measurement from tactile sensing. In *2021 IEEE/RSJ International Conference on Intelligent*
346 *Robots and Systems (IROS)*, pages 6809–6816. IEEE, 2021.
- 347 [10] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez. Tactile-rl for inser-
348 tion: Generalization to objects of unknown geometry. In *2021 IEEE International Conference*
349 *on Robotics and Automation (ICRA)*, pages 6437–6443. IEEE, 2021.
- 350 [11] S. Kim and A. Rodriguez. Active extrinsic contact sensing: Application to general peg-in-hole
351 insertion. *arXiv preprint arXiv:2110.03555*, 2021.
- 352 [12] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson. Measurement of shear and slip with a
353 gelsight tactile sensor. In *2015 IEEE International Conference on Robotics and Automation*
354 *(ICRA)*, pages 304–311. IEEE, 2015.
- 355 [13] S. Dong, W. Yuan, and E. H. Adelson. Improved gelsight tactile sensor for measuring geom-
356 etry and slip. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*
357 *(IROS)*, pages 137–144. IEEE, 2017.
- 358 [14] F. Veiga, H. Van Hoof, J. Peters, and T. Hermans. Stabilizing novel objects by learning to
359 predict tactile slip. In *2015 IEEE/RSJ International Conference on Intelligent Robots and*
360 *Systems (IROS)*, pages 5065–5072. IEEE, 2015.
- 361 [15] M. Li, Y. Bekiroglu, D. Kragic, and A. Billard. Learning of grasp adaptation through expe-
362 rience and tactile sensing. In *2014 IEEE/RSJ International Conference on Intelligent Robots*
363 *and Systems*, pages 3339–3346. Ieee, 2014.
- 364 [16] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba.
365 Openai gym, 2016.

- 366 [17] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin,
367 A. Allshire, A. Handa, et al. Isaac gym: High performance gpu based physics simulation for
368 robot learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets
369 and Benchmarks Track (Round 2)*, 2021.
- 370 [18] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In
371 *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–
372 5033. IEEE, 2012.
- 373 [19] E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics
374 and machine learning. 2016.
- 375 [20] O. . M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki,
376 A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder,
377 L. Weng, and W. Zaremba. Learning dexterous in-hand manipulation. *The International Jour-
378 nal of Robotics Research*, 39(1):3–20, 2020.
- 379 [21] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter. Learn-
380 ing agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), 2019.
- 381 [22] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke. Sim-
382 to-real: Learning agile locomotion for quadruped robots. In *Robotics: Science and Systems*,
383 2018.
- 384 [23] Z. Ding, Y.-Y. Tsai, W. W. Lee, and B. Huang. Sim-to-real transfer for robotic manipulation
385 with tactile sensory. In *2021 IEEE/RSJ International Conference on Intelligent Robots and
386 Systems (IROS)*, pages 6778–6785. IEEE, 2021.
- 387 [24] A. Melnik, L. Lach, M. Plappert, T. Korthals, R. Haschke, and H. Ritter. Using tactile sensing
388 to improve the sample efficiency and performance of deep deterministic policy gradients for
389 simulated in-hand manipulation tasks. *Frontiers in Robotics and AI*, page 57, 2021.
- 390 [25] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra. Tacto: A fast, flexible, and open-source
391 simulator for high-resolution vision-based tactile sensors. *IEEE Robotics and Automation
392 Letters*, 7(2):3930–3937, 2022.
- 393 [26] Y. S. Narang, K. Van Wyk, A. Mousavian, and D. Fox. Interpreting and predicting tactile
394 signals via a physics-based and data-driven framework. *arXiv preprint arXiv:2006.03777*,
395 2020.
- 396 [27] Y. Narang, B. Sundaralingam, M. Macklin, A. Mousavian, and D. Fox. Sim-to-real for robotic
397 tactile sensing via physics-based simulation and learned latent projections. In *2021 IEEE
398 International Conference on Robotics and Automation (ICRA)*, pages 6444–6451. IEEE, 2021.
- 399 [28] Z. Ding, N. F. Lepora, and E. Johns. Sim-to-real transfer for optical tactile sensing. In *2020
400 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1639–1645. IEEE,
401 2020.
- 402 [29] Z. Si and W. Yuan. Taxim: An example-based simulation model for gelsight tactile sensors.
403 *IEEE Robotics and Automation Letters*, 2022.
- 404 [30] D. Ma, E. Donlon, S. Dong, and A. Rodriguez. Dense tactile force estimation using gelslim
405 and inverse fem. In *2019 International Conference on Robotics and Automation (ICRA)*, pages
406 5418–5424. IEEE, 2019.
- 407 [31] T. Bi, C. Sferrazza, and R. D’Andrea. Zero-shot sim-to-real transfer of tactile control policies
408 for aggressive swing-up manipulation. *IEEE Robotics and Automation Letters*, 6(3):5761–
409 5768, 2021.

- 410 [32] A. Habib, I. Ranatunga, K. Shook, and D. O. Popa. Skinsim: A simulation environment for
411 multimodal robot skin. In *2014 IEEE International Conference on Automation Science and*
412 *Engineering (CASE)*, pages 1226–1231. IEEE, 2014.
- 413 [33] S. Moio, B. León, P. Korkealaakso, and A. Morales. Model of tactile sensors using soft
414 contacts and its application in robot grasping simulation. *Robotics and Autonomous Systems*,
415 61(1):1–12, 2013.
- 416 [34] J. Xu, T. Chen, L. Zlokapa, M. Foshey, W. Matusik, S. Sueda, and P. Agrawal. An End-to-
417 End Differentiable Framework for Contact-Aware Robot Design. In *Proceedings of Robotics:*
418 *Science and Systems*, Virtual, July 2021. doi:10.15607/RSS.2021.XVII.008.
- 419 [35] Y. Wang, N. J. Weidner, M. A. Baxter, Y. Hwang, D. M. Kaufman, and S. Sueda. RED-
420 MAX: Efficient & flexible approach for articulated dynamics. *ACM Trans. Graph.*, 38(4), July
421 2019. ISSN 0730-0301. doi:10.1145/3306346.3322952. URL [https://doi.org/10.1145/](https://doi.org/10.1145/3306346.3322952)
422 [3306346.3322952](https://doi.org/10.1145/3306346.3322952).
- 423 [36] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization
424 algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 425 [37] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar. Robel:
426 Robotics benchmarks for learning with low-cost robots. arxiv e-prints, page. *arXiv preprint*
427 *arXiv:1909.11639*, 2019.
- 428 [38] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization
429 for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ*
430 *international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017.
- 431 [39] J. Xu, V. Makoviychuk, Y. Narang, F. Ramos, W. Matusik, A. Garg, and M. Macklin. Accelerated policy learning with parallel differentiable simulation. In *International Conference on Learning Representations*, 2021.