

Zero-Shot Sim-to-Real Transfer of Tactile Control Policies for Aggressive Swing-Up Manipulation

Thomas Bi , Carmelo Sferrazza , and Raffaello D'Andrea 

Abstract—This letter aims to show that robots equipped with a vision-based tactile sensor can perform dynamic manipulation tasks without prior knowledge of all the physical attributes of the objects to be manipulated. For this purpose, a robotic system is presented that is able to swing up poles of different masses, radii and lengths, to an angle of 180° , while relying solely on the feedback provided by the tactile sensor. This is achieved by developing a novel simulator that accurately models the interaction of a pole with the soft sensor. A feedback policy that is conditioned on a sensory observation history, and which has no prior knowledge of the physical features of the pole, is then learned in the aforementioned simulation. When evaluated on the physical system, the policy is able to swing up a wide range of poles that differ significantly in their physical attributes without further adaptation. To the authors' knowledge, this is the first work where a feedback policy from high-dimensional tactile observations is used to control the swing-up manipulation of poles in closed-loop.

Index Terms—Force and tactile sensing, modeling, control, and learning for soft robots, dexterous manipulation.

I. INTRODUCTION

TACTILE sensors aim to provide robots with a sense of touch that captures information from their environment through physical contact. In this letter, the vision-based tactile sensor presented in [1] is deployed in order to demonstrate that it can provide robots with a dexterity akin to that of humans in dynamic manipulation tasks. For this purpose, a robotic system that performs swing-up maneuvers for different poles is presented (see Fig. 1). The robotic system consists of a parallel gripper, mounted to a linear motor, with two tactile sensors acting as fingers. Thereby, three key capabilities enabled by the artificial sense of touch provided by the tactile sensor are demonstrated: (i) The system is able to adapt its motion and successfully swings up poles that differ in their physical attributes (e.g. mass, length, and radius) without prior knowledge of these attributes. (ii) The system does not rely on external visual sensing; instead, the pose and attributes of the pole in contact are implicitly inferred from the tactile observations alone. (iii)

Manuscript received January 7, 2021; accepted May 7, 2021. Date of publication May 28, 2021; date of current version June 14, 2021. This letter was recommended for publication by Associate Editor Y. Pan and Editor Dan Popa upon evaluation of the reviewers' comments. (Corresponding author: Thomas Bi.)

The authors are with the Members of the Institute for Dynamic Systems and Control, ETH Zurich, Zürich 8092, Switzerland (e-mail: thomas@novobit.ch; csferrazza@ethz.ch; rdandrea@ethz.ch).

This letter has supplementary downloadable material available at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LRA.2021.3084880

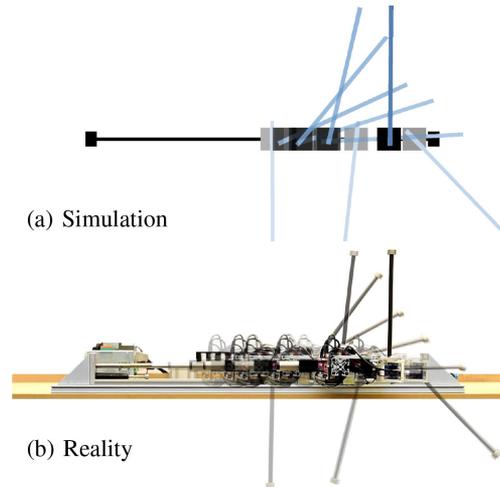


Fig. 1. In this work, tactile control policies for the swing-up manipulation of poles are learned in a simulation of the tactile sensor that runs faster than real-time. When transferred to the real-world system, the policy achieves the desired behaviour without further adaptation, and swings up poles to an upright position without prior knowledge of the physical attributes of the pole.

The tactile observations can be processed in real-time and act as feedback for closed-loop control at 60Hz. As a result, highly dynamic swing-up manipulations are achieved without the need for a previous in-hand exploration of the pole.

The three components that enable such adaptive dynamic swing-up manipulation are presented here. First, the high-dimensional force distribution acting on the sensor surface is directly inferred from the sensor camera images using an efficient convolutional network, which is trained on purely simulated contact interactions of the sensor with different poles. Second, a novel simulator is developed that accurately models the behaviour of the soft sensor surface when interacting with a rigid cylindrical pole. This simulation is based on combining the finite element method with state-of-the-art semi-implicit time-stepping schemes for contact resolution and runs at 360Hz on a single core of an Intel Core i7-7700 k processor. Third, a framework for learning adaptive feedback policies conditioned on a history of sensory observations is proposed. Deep reinforcement learning is utilized to train a single policy, entirely in simulation, that is able to swing up different poles with unknown physical attributes. Thereby, various strategies that facilitate the sim-to-real transfer of policies learned in simulation are employed, namely *dynamics randomization* [2], and *privileged learning* [3], [4].

A. Related Work

As reviewed in [5] and [6], many works have demonstrated how robots can leverage the sense of touch in dexterous manipulation tasks in closed-loop. For example, in [7] tactile data is used to control both the grasping force and slippage of a tactile gripper. Other examples include [8], where a pair of grippers are used to pick one end of a cable and follow it to the other end. Each gripper contains tactile sensors from which the current pose and friction forces acting on the cable can be estimated in real-time, enabling the approach to generalize to cables of different thicknesses and materials, based on a model learned from real data. A similar approach was proposed in [9]; a dual palm robotic system estimates the pose and the stick/slip behaviour of an object solely from tactile feedback, in order to manipulate an object on a planar surface to a desired position. In [10], a deep dynamics model is learned that can predict future tactile observations based on the previous observations and actions taken. Data for the training of the dynamics model is autonomously collected on the physical system. The learned model is then used in an MPC-framework to manipulate a ball, analog stick, and 20-sided die to a desired configuration. Other learning-based approaches rely on deep reinforcement learning to find optimal control policies directly on the physical hardware. Examples include a 5-DoF arm that learns to reorient objects using a latent representation of the tactile data [11], and a robotic system that learns to type on a Braille keyboard [12].

While these approaches demonstrate robustness against external disturbances and changes in object properties, the manipulation tasks they solve generally do not require a high degree of dynamicism. For more aggressive manipulation tasks such as the swing-up manipulation of poles, feedback control based on tactile data has proven to be challenging and thus differing methods have been proposed. In [13], the sensing and manipulation are separated into two steps. First, the physical features of different poles are learned by shaking and tilting the pole in-hand and observing the tactile feedback. The learned features are then used to optimize an open-loop trajectory of a robotic arm that dynamically swings the pole up to a desired angle. The learning of the physical features, as well as the trajectory optimization, are performed end-to-end using models trained on a physically collected dataset. In [14], tactile sensing and visual tracking are combined to pivot an object to a desired angle by adjusting the gripping force exerted by a two-finger gripper. This fusion of visual and tactile information was also employed on a robotic hand in [15] to perform highly dynamic tasks such as pen spinning, ball dribbling, and ball throwing.

In this work, a unified approach is presented where aggressive swing-up maneuvers can be achieved in closed-loop from high-dimensional tactile feedback, without relying on a visual tracking system or prior in-hand exploration of the pole. Moreover, instead of relying on data collection on the physical system, as is done in the learning-based methods mentioned above, the feedback control policy is learned entirely in simulation. This removes the cost of collecting data on the physical system, which can be highly time-consuming. Additionally, challenging motions that may lead to unsafe behaviours by the physical

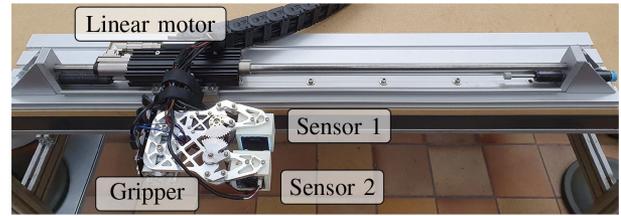


Fig. 2. The robotic system presented in this letter consists of a parallel gripper comprising two tactile sensor, and a linear motor to which the gripper is mounted.

hardware can be first explored without repercussions. This data can then be utilized to train the policy to satisfy the safety constraints that are present on the physical system. While simulators for the behaviour of tactile sensors have been developed (see e.g. [16]–[23]), the authors are not aware of any work where a simulation from first principles is utilized to learn tactile feedback control policies. Rather, the mentioned works focus on gathering supervised datasets of tactile images in simulation to train deep neural networks that can predict object position and rotation ([16], [21]), the force distribution acting on the sensor surface ([22], [23]), or the three-dimensional mesh of the object in contact ([17]).

B. Outline

The hardware employed for the experiments is presented in Section II. In Section III, the proposed methods are described. This includes the sensing approach of the tactile sensor in Section III-A, the design of the tactile simulator in Section III-B, and the synthesis of the swing-up control policy in Section III-C. Results from employing the learned policy on the real-world system are presented in Section IV. Finally, Section V draws conclusions and gives an outlook on future work. In the remainder of this letter, vectors are expressed as tuples for ease of notation, with dimension and stacking clear from the context.

II. HARDWARE

The robotic system considered in this letter consists of three main parts; a two-finger robotic gripper where each finger comprises a tactile sensor, a linear motor (stator/slider), to which the gripper is mounted, and finally, embedded computing systems that process the sensing data and send commands to the actuators. The linear motor and tactile gripper are pictured in Fig. 2.

A. Tactile Gripper

To enable the high-resolution gripping capability of the system, a custom 1-DoF parallel two-finger gripper was built in-house, see Fig. 3(a). A Dynamixel MX-28R servo motor is used to control the distance between the two fingers with a resolution of 0.06 mm. Two tactile sensors, placed opposite each other, act as fingers for the gripper. The sensing principle employed in this letter is based on [1]. Three soft silicone layers are poured on top of an RGB fisheye camera (ELP USBFHD06H), surrounded by LEDs. The base layer (ELASTOSIL RT 601 RTV-2, mixing

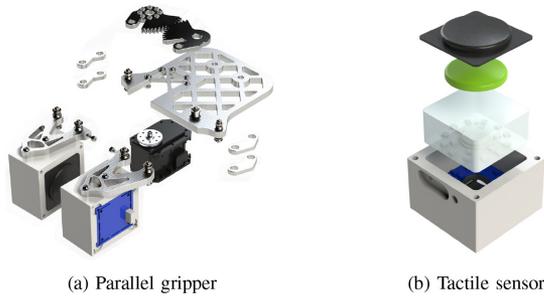


Fig. 3. This figure shows exploded views of the gripper (a) and the tactile sensor that acts as finger (b). A Dynamixel MX-28R controls the opening and closing of the two fingers, each of them equipped with a tactile sensor.

ratio 7:1, shore hardness 45 A) is stiff and transparent, and serves as a spacer. The middle layer (ELASTOSIL RT 601 RTV-2, ratio 25:1, shore hardness 10 A) is soft and transparent, and embeds a spread of randomly distributed fluorescent green particles. A black top layer (made of the same material as the middle layer) completes the sensor and shields it from external light disturbances. The soft sensor's surface is slightly curved to provide a more anatomical grasping surface. An exploded view of the sensor layers is shown in Fig. 3(b).

B. Linear Motor

In order to achieve the translational motion of the gripper, a linear motor comprising a *stator* and a *slider* is employed. The stator (LinMot P01-23x160H-HP-R) contains the motor windings, bearings for the slider, position capture sensors and a microprocessor, and is thus able to generate motion with respect to the slider. The slider (LinMot PL01-12x850/810-HP) is a stainless steel tube and is fixed to a table so that the stator is the only moving part. The gripper is then mounted to the stator through the use of a motor flange (LinMot PF02-23x120). A motor drive (LinMot C1100-GP-XC-0S-000) controls the motion of the stator.

C. Embedded Systems

Two embedded devices are used to control the system. First, a Raspberry Pi (RPI) runs two low-level controllers, one for the linear motor and one for the gripper. The linear motor controller tracks commanded acceleration setpoints, while the gripper controller tracks the distance between the two fingers. Second, an NVIDIA Jetson TX2, a compact embedded device with a built-in GPU, obtains the camera images from the tactile sensor at 60Hz, pre-processes the images, and infers the force distribution. Note that this pipeline is only executed for sensor 1 (see Fig. 2). This is motivated by the fact that due to the planar nature of the system, the forces acting on sensor 2 can be assumed to be symmetrical to those acting on sensor 1. Furthermore, this reduces the computational complexity of the pipeline.

The Jetson also receives the current actuator states from the RPI. Control actions are then inferred using the proposed control policy and are communicated to the low-level controllers on the RPI that execute the commands.

III. METHOD

The proposed method can be divided into three different parts. First, the vision-based tactile sensor estimates the force distribution acting on its surface from its camera images. Second, a simulator for the dynamics of a pole and the given robotic system is developed. Third, a tactile feedback control policy for the swing-up manipulation is learned in the simulation using reinforcement learning.

A. Vision-Based Tactile Sensing

The tactile sensor employed in this letter follows the same sensing principle as introduced in [1]. When the soft sensing surface is subject to force, the material deforms and displaces the particles tracked by the camera. This motion generates different patterns in the images. The material deformation at any point in time can thus be described by two camera images, one where no loads are applied and the material is at rest, and another at the current deformed state.

In [22], a method to generate such images in simulation is presented to train a supervised learning architecture that aims to accurately estimate the real-world 3D contact force distribution. The same approach to generate training data is employed here, using finite-element simulations of the sensor surface under various contact conditions, where hyperelastic material models for the sensor's soft materials are employed. The details of this procedure can be found in [22]. In addition to the two mentioned images per datapoint, the polar coordinates of each pixel are encoded here as two additional image channels. Explicitly incorporating such spatial location features has previously been shown to significantly improve accuracy where the location of image features is relevant for the task at hand [24]. Using a fully convolutional neural network based on ShuffleNet V2 [25], the resulting four-channel image is then mapped to accurate contact force distribution labels (see Fig. 4), with ground truth also obtained from finite element simulations [26].

On the real-world sensor, camera images are preprocessed to match those of the simulated training dataset as described in [22] and [23]. Specifically, images are converted to gray-scale and remapped using the real-world camera model (obtained via a state-of-the-art calibration technique [27]) to images of the same scene as if they were taken in the simulated world. A circular mask is then applied to remove any irrelevant image information. The results of this preprocessing procedure are illustrated in Fig. 5. On the given hardware, the force distribution for a given preprocessed camera image can be inferred in real-time in 2.5 ms.

B. Tactile Simulation

In order to achieve a fast simulator, essential for training reinforcement learning algorithms in a reasonable amount of time, a few model simplifications are introduced here. First, the material of the sensor is assumed to be linearly elastic. Second, the forces acting on the sensor surface are decoupled into two components: the forces arising due to the material deformation in the z direction, and the lateral friction forces resulting from

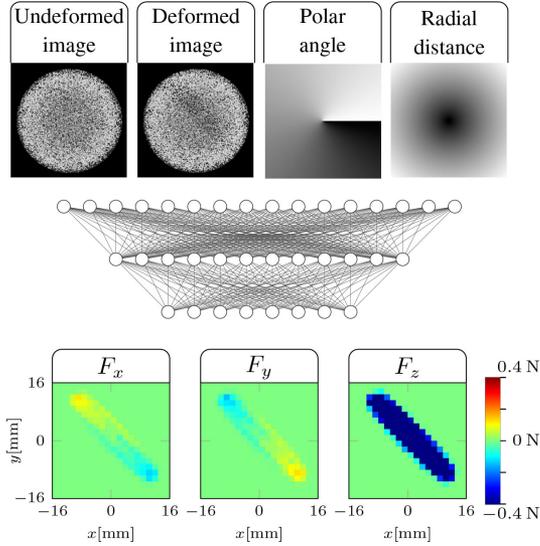


Fig. 4. The features and labels of our supervised learning task. The features are the images of the particles in the undeformed and in the deformed states. Two additional channels provide the polar coordinates of the pixels. For the labels, F_x , F_y and F_z denote the x , y and z components of the discretized force distribution. The force distributions shown in the figure were collected with a pole pivoting on the sensor (as in the experiments discussed in Section IV). They show how the shape and pose of the pole manifest themselves correctly in the force distribution readings. Additionally, the lateral forces (F_x and F_y) on either side of the center of rotation point in opposite directions, which can be deduced from their change in color.

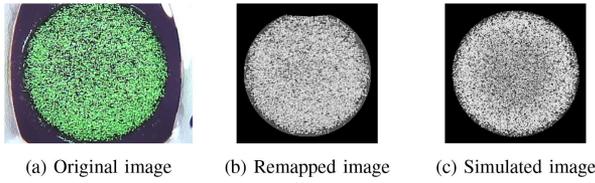


Fig. 5. Camera images from the real-world sensor are preprocessed converting the original image (a) to gray-scale, and remapping the image using a calibrated camera model (b). The last image (c) corresponds to a simulated image that is used for training and is provided for comparison.

the relative motion of the pole with respect to the sensor. A real-time finite element approach is then employed to compute the two mentioned force components.

1) *Problem Statement:* A sketch of the system considered in this work is shown in Fig. 6, where a single coordinate system is defined. The x -axis is aligned to the moving axis of the linear motor, denoted in the following as the cart. The y -axis points in the opposite direction to gravity, i.e. upwards. Finally, the z -axis is chosen such that all the points on sensor 1 exhibit a negative z -coordinate. The reference position $\mathbf{s} = (x_s, y_s, z_s)$ is chosen with the point on the curved surface of sensor 1 that is closest to the x - y plane (at rest). While there are two tactile sensors, their positions are symmetrical about the x - y plane. Hence, it suffices to only consider the position of a single sensor. Next, the orientation of the pole is defined by the angle ϕ . Lastly, $\mathbf{p} = (x_p, y_p, z_p)$ denotes the position of the center of mass of the pole. Note that the sensor is fixed in the y -direction ($y_s = 0$) and the pole is fixed in the z -direction ($z_p = 0$). These definitions

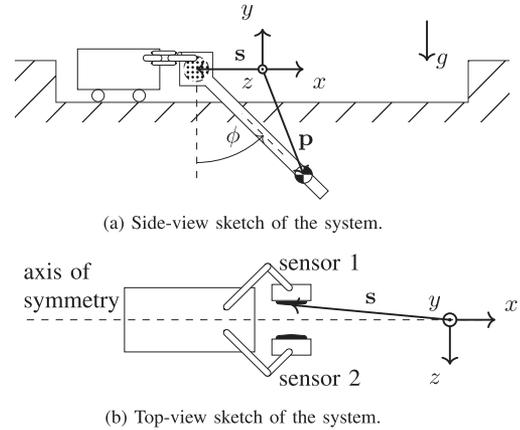


Fig. 6. The system can be described as a cart-pole augmented with a parallel gripper that features the two sensors.

are illustrated in Fig. 6. The state vector \mathbf{x} is then defined as

$$\mathbf{x} = (x_s, \dot{x}_s, z_s, x_p, \dot{x}_p, y_p, \dot{y}_p, \phi, \dot{\phi}) \quad (1)$$

Since only the static behavior of the sensor material is analyzed, \dot{z}_s is not considered.

The inputs to the system are the cart acceleration and the increment in z_s between two subsequent timesteps:

$$\mathbf{u} = (\ddot{x}_s, \Delta z_s) \quad (2)$$

Note that on the real system, the servo commands are mapped to Δz_s with a linear mapping identified from data.

Next, the pole is characterized as a rigid cylinder. Its radius is given by r_p , the mass by m_p and its moment of inertia about the center of mass and along the z -axis by I_p . The length of the pole above its center of mass is given by $l_{p,u}$ and the length below the center of mass by $l_{p,l}$.

Given these definitions, the goal is to model the state evolution over time, i.e., $\mathbf{x}(k+1) = f(\mathbf{x}(k), \mathbf{u}(k))$, where k is the discrete time index, and f describes a functional dependency. In the following, the time index (k) will be omitted, and variables at time ($k+1$) will be denoted by a $+$ superscript, e.g. x^+ .

2) *Equations of Motion:* The pole is modeled as a free-body constrained to move in the x - y plane, meaning that its motion is governed by the force $F_p = (F_{p,x}, F_{p,y}, 0)$ and torque $T_p = (0, 0, T_{p,z})$ acting on its center of mass. Using a semi-implicit integration scheme ([28], [29]) the equations of motion are then given by

$$\dot{x}_s^+ = \dot{x}_s + \Delta t \ddot{x}_s^+ = x_s + \Delta t \dot{x}_s^+ \quad (3)$$

$$\dot{x}_p^+ = \dot{x}_p + \Delta t \frac{F_{p,x}}{m_p} x_p^+ = x_p + \Delta t \dot{x}_p^+ \quad (4)$$

$$\dot{y}_p^+ = \dot{y}_p + \Delta t \frac{F_{p,y}}{m_p} y_p^+ = y_p + \Delta t \dot{y}_p^+ \quad (5)$$

$$\dot{\phi}^+ = \dot{\phi} + \Delta t \frac{T_{p,z}}{I_p} \phi^+ = \phi + \Delta t \dot{\phi}^+ \quad (6)$$

$$z_s^+ = z_s + \Delta z_s \quad (7)$$

In the following, the derivation of F_p and T_p is presented.

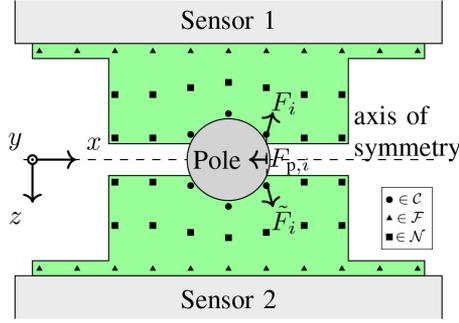


Fig. 7. The nodes of the finite element mesh are assigned to three sets: \mathcal{C} if in contact with the pole, \mathcal{F} if constrained to not move, \mathcal{N} otherwise.

Both sensors are discretized using an identical mesh of $N = 576$ finite elements (nodes). Hereafter, all quantities introduced will refer to sensor 1, where the corresponding counterparts of sensor 2 are clear from the symmetrical context and are denoted using a tilde, i.e. $\tilde{\cdot}$. Then, for node i of the mesh, let (x_i, y_i, z_i) be its coordinates, and F_i the force acting on the node. Each node i in contact with the pole leads to a planar reaction force

$$F_{p,i} = -(F_i + \tilde{F}_i) \Rightarrow F_{p,i,x:y} = -2F_{i,x:y} \quad (8)$$

where the implication follows from symmetry, with the $x:y$ subscript denoting the stacked x and y components of the three-dimensional vector. Next, the gravitational force acting on the pole is denoted by $F_g = (0, -m_p g, 0)$, where $g = 9.81 \text{ms}^{-2}$.

Defining $\mathbf{r}_i := (x_i - x_p, y_i - y_p, 0)$, the total force and torque acting on the pole are then

$$F_p = F_g + \sum_{i \in \mathcal{S}} F_{p,i} \quad T_p = \sum_{i \in \mathcal{S}} \mathbf{r}_i \times F_{p,i}. \quad (9)$$

where \mathcal{S} is the set of all nodes on the surface of the sensor.

As mentioned above, the contact forces are postulated to be the superposition of forces F_i^0 arising from the normal indentation of the pole into the sensor, and the lateral friction forces F_i^f , that is, $F_i = F_i^0 + F_i^f$, which implies

$$F_{p,i} = - \underbrace{(F_i^0 + \tilde{F}_i^0)}_{=: F_{p,i}^0} - \underbrace{(F_i^f + \tilde{F}_i^f)}_{=: F_{p,i}^f} \quad (10)$$

3) *Forces Arising From Normal Indentation:* The forces F_i^0 are derived using the finite element theory for linearly elastic materials. Let U_i^0 be the deformation of a node i . Then, a linear relationship between the external forces and deformations is found by the finite element method as:

$$F^0 = KU^0 \quad (11)$$

where $F^0 := (F_1^0, \dots, F_N^0)$, $U^0 := (U_1^0, \dots, U_N^0)$, and K is the global stiffness matrix, obtained in this work in Abaqus/Standard. The system of equations in (11) can be solved by introducing the following sets, displayed in Fig. 7:

- \mathcal{C} : The set of all nodes that are in contact with the pole. It is the intersection of the set of nodes at the surface of the sensor (i.e., the set \mathcal{S}) and the set of nodes whose positions at

rest collide with the pole, based on the geometric properties of the pole considered. The nodes in this set are assumed here to translate only in z -direction, and their deformation is obtained by finding the appropriate z -coordinate that intersects with the surface of the pole (see Fig. 7).

- \mathcal{F} : The set of all nodes that are in contact with the base layer of the sensor. Since the base layer's stiffness is much larger than the stiffness of the sensor surface, the nodes of this set are assumed to be rigid. Therefore, their deformation is set to zero, i.e., $U_i^0 = 0, \forall i \in \mathcal{F}$.
- \mathcal{N} : The set of nodes that are neither in contact with the base layer nor in contact with the pole. No external forces are acting on these nodes, i.e., $F_i^0 = 0, \forall i \in \mathcal{N}$.

Therefore, for a node i , once a corresponding set is identified, either the force F_i^0 or the deformation U_i^0 is known. The system (11) is then solved by using the UMFPAK library, and $F_{p,i}^0$ computed as in (10).

Note that here the current approach exploits the cylindrical geometry of the poles, rendering a mathematically simple intersection problem, which enables a highly efficient identification of the aforementioned sets. The extension to objects of various geometries may still be addressed efficiently by employing algorithms tailored to solve the intersection problem for generic polygons, e.g., based on the Weiler-Atherton clipping algorithm [30].

4) *Lateral Friction Forces:* In order to find the lateral friction forces for the nodes in contact, first the case where only the friction at a single node is unknown is considered. From the solution of this case, an iterative method is utilized to solve for all friction forces in the multi-contact case.

First, it is assumed that all the friction forces except for the one at node i are known, i.e. F_j^f is known for all $j \neq i$. Let

$$\mathbf{v}_{i,\text{rel}} := \begin{pmatrix} \dot{x}_i \\ \dot{y}_i \end{pmatrix} - \begin{pmatrix} \dot{x}_s \\ 0 \end{pmatrix} = \begin{pmatrix} \dot{x}_p \\ \dot{y}_p \end{pmatrix} + \begin{pmatrix} -\mathbf{r}_{i,y} \\ \mathbf{r}_{i,x} \end{pmatrix} \dot{\phi} - \begin{pmatrix} \dot{x}_s \\ 0 \end{pmatrix}$$

be the relative planar velocity of the point on the pole which is in contact with the node i at time k . Then, by plugging in the equations of motion (3)-(7), the relative velocity at the next timestep is found to be

$$\mathbf{v}_{i,\text{rel}}^+ = \mathbf{v}_{i,\text{rel}} - \Delta t \begin{pmatrix} \ddot{x}_s \\ 0 \end{pmatrix} + \mathbf{J}_{ii} \underbrace{F_{p,i,x:y}}_{F_{p,i,x:y}^0 + F_{p,i,x:y}^f} + \sum_{j \neq i} \mathbf{J}_{ij} F_{p,j,x:y} \quad (12)$$

where, for generic indices a and b ,

$$\mathbf{J}_{ab} = \Delta t \begin{bmatrix} \frac{1}{m_p} + \frac{\mathbf{r}_{a,y}\mathbf{r}_{b,y}}{I_p} & \frac{-\mathbf{r}_{a,y}\mathbf{r}_{b,x}}{I_p} \\ \frac{-\mathbf{r}_{a,x}\mathbf{r}_{b,y}}{I_p} & \frac{1}{m_p} + \frac{\mathbf{r}_{a,x}\mathbf{r}_{b,x}}{I_p} \end{bmatrix}. \quad (13)$$

In this work, Coulomb friction is assumed, and two cases are identified, where μ indicates both the static and kinetic friction coefficients. First, for the static friction case, consider the force $F_{p,i}^{\text{f,static}}$ that takes on exactly the value to prevent motion at node i . This force can be found by setting $\mathbf{v}_{i,\text{rel}}^+ = \mathbf{0}$ in (12) and solving

for $F_{p,i,x:y}^f$. If this force satisfies the friction cone constraint, i.e.

$$\left\| F_{p,i}^{f,\text{static}} \right\| \leq 2\mu |F_{i,z}^0|, \quad (14)$$

then $F_{p,i,x:y}^f = F_{p,i,x:y}^{f,\text{static}}$. The z -component is set to zero as it would eventually cancel out when considering both fingers.

If (14) is not satisfied, friction is not sufficient to prevent the motion at node i . In this case, kinetic friction is present, where the force is opposite to the direction of the velocity and is proportional to the normal component of the force. Since the velocity $\mathbf{v}_{i,\text{rel}}^+$ and the friction force $F_{p,i}^f$ are coupled, an approximation of the subsequent velocity is employed as

$$\hat{\mathbf{v}}_{i,\text{rel}}^+ := \mathbf{v}_{i,\text{rel}} - \Delta t \begin{pmatrix} \ddot{x}_s \\ 0 \end{pmatrix} + \mathbf{J}_i F_{p,i,x:y}^0 + \sum_{j \neq i} \mathbf{J}_j F_{p,j,x:y}$$

which is the relative velocity at the subsequent step when the effects of friction at node i are ignored. If the number of nodes is sufficiently large, this approximation is close to the true value, since the effect of the force at the single node i is small compared to the combined effect of the remaining forces at nodes $j \neq i$. Using this approximation, the kinetic friction is set to

$$F_{p,i,x:y}^f = -2\mu |F_{i,z}^0| \hat{\mathbf{v}}_{i,\text{rel}}^+ / \|\hat{\mathbf{v}}_{i,\text{rel}}^+\| \quad (15)$$

Given this solution to the single contact problem, the multi-node contact case is solved by repeatedly iterating over all nodes in contact until convergence, and updating the friction force at node i using the above solution, given the values at nodes $j \neq i$ of the current iteration. Then, $F_{p,i}$ is obtained as in (10) from $F_{p,i}^0$ and $F_{p,i}^f$, and finally F_p and T_p can be computed as in (9).

C. Learning Tactile Control Policies

Given a simulation of a robotic system, deep reinforcement learning (deep RL) algorithms have been successfully applied to learn sophisticated behaviours [32]. These algorithms typically depend on the Markov property of the system, i.e. they assume that the state and physical parameters that fully describe the system at a given time are available to the policy. However, for the experiments presented in Section IV, the physical parameters of the pole, e.g. the length, are unknown to the policy and the Markov property no longer holds. In deep RL, such problems are typically dealt with by using the history of observations and parametrizing the policy using recurrent neural networks which, however, can be challenging to train.

The approach employed here exploits the fact that in simulation the state and physical parameters are known. In a first stage, an *expert* policy π^e is learned that has access to the state as well as the simulation parameters (satisfying the Markov property). In a second stage, a *student* policy π^s that only has access to the observations that are available on the real system is learned by imitating the behaviour of the expert policy (see Fig. 8). This is also referred to as privileged learning [3], [4].

1) *State-Feedback Expert Policy*: In order to achieve the swing-up with a feedback policy that adapts to different poles, the expert policy is conditioned on the state $\mathbf{x}(k)$ (defined in Section III-B), as well as the pole's physical parameters which

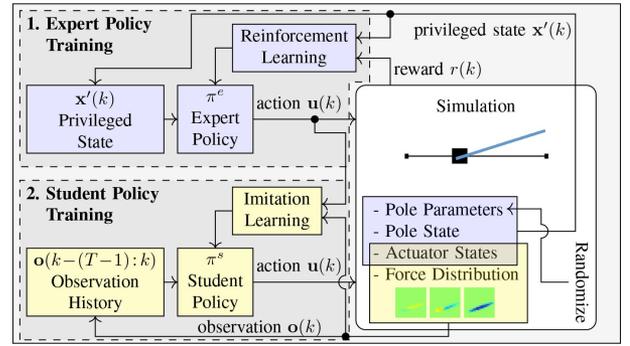


Fig. 8. The figure depicts the privileged learning approach. The expert and student policy training take place in two separate steps, both performed entirely in simulation. Further, both policies are parametrized using two-layer fully connected neural networks. For the reinforcement learning of the expert policy, the SAC [31] method is employed to find the optimal policy according to (17). The student policy is then deployed to the real-world system without further adaptation.

may vary. This yields the augmented state

$$\mathbf{x}'(k) := \left(\mathbf{x}(k), r_p, m_p, I_p, l_{p,u}, l_{p,l}, \mu \right) . \quad (16)$$

As a result, the policy may choose different control actions based on the features of the pole.

The goal is then to find a policy, $\pi : \mathbf{x}'(k) \rightarrow \mathbf{u}(k)$, that is optimal in the sense of maximizing the expected sum of future discounted rewards, i.e.

$$\pi^e = \max_{\pi} \mathbb{E} \left(\sum_k \gamma^k r(\mathbf{x}'(k), \pi(\mathbf{x}'(k))) \right) , \quad (17)$$

where γ is the discount factor. The reward function r is shaped to encourage low slippage and pole orientations that are close to 180° . The policy is learned using deep RL, namely the SAC [31] algorithm with the stable-baselines3 implementation [33]. The discount factor is set to $\gamma = 0.995$ while the remaining hyperparameters, as well as the policy network architecture, correspond to the default ones proposed in [31]. During training, the pole parameters r_p , m_p , I_p , $l_{p,u}$, $l_{p,l}$ and μ are randomly sampled at each new episode such that the policy learns the correct behaviour for different poles and friction ratios. This *dynamics randomization* [2] also greatly aids in the successful transfer from simulation to reality.

2) *Tactile Student Policy*: The expert policy is conditioned on privileged knowledge, only available in simulation, and can thus not be deployed on the real system, where the pole's pose and physical attributes can only indirectly be observed through the available force distribution measurements. As a result, the student policy must be able to reason over time and implicitly recover the missing state information. First, in order to condense the sensory information into a compressed representation, an estimate of the pole's orientation $\hat{\phi}(k)$ is obtained by computing the force magnitude at each bin, thresholding the magnitudes to obtain a binary image, and finally applying a Hough line transform [34]. In addition, the total sensed normal force $F_z^{\text{tot}}(k)$ is extracted by summing the z -distribution at all bins. This is motivated by the fact that the normal force yields direct

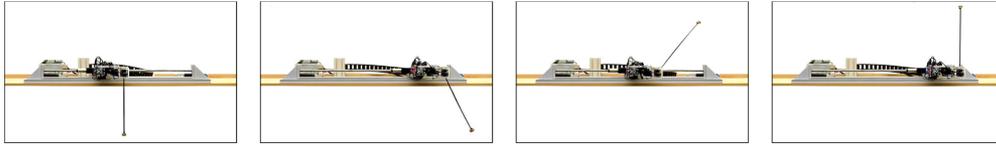


Fig. 9. This figure shows a trajectory which results from employing the learned feedback control policy on the robotic system. As can be seen, the pole is dynamically swung up to an upright position.

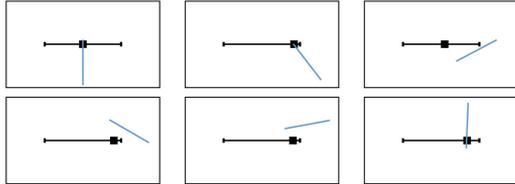


Fig. 10. Using deep reinforcement learning, robust policies can be learned to achieve various tasks. Here, the reward function is shaped to encourage the throwing and catching of the pole after a rotation of 360° .

information about the friction and slippage, while the angle of the pole is the main quantity to be controlled. Then, a student policy conditioned on a history of condensed representations of the observations is learned by imitating the behaviour of the expert policy.

A condensed observation at time k is given by

$$\mathbf{o}(k) = \left(x_s(k), z_s(k), \hat{\phi}(k), F_z^{\text{tot}}(k) \right). \quad (18)$$

Note that $x_s(k)$ and $z_s(k)$ are known for the real system, and the velocity of the cart $\dot{x}_s(k)$ is not included, since it can implicitly be derived from the history of $x_s(k)$ observations.

The student policy π^s is then parametrized by a neural network that maps the history of the last T condensed observations $\mathbf{o}(k - (T - 1) : k)$ to the control action $\mathbf{u}(k)$, where $T = 12$ is the fixed history length. The same stochastic network as proposed in [31] is used, which outputs a squashed Gaussian distribution over the control actions. This stochasticity accomplishes a desirable smoothing of the policy. The imitation of the expert policy is then posed as a supervised learning task that minimizes the negative log-probability

$$\mathcal{L} := -\log \Pr(\pi^s(\mathbf{o}(k - (T - 1) : k)) = \pi^e(\mathbf{x}'(k))).$$

In this work, the DAGGER [35] method is employed, where the dataset is continuously aggregated with the incoming data from the training rollouts of the student policy. Labels are obtained by querying the expert policy for the visited states. At each training iteration, the student policy is updated by performing an optimization step with batches sampled from the aggregated dataset.

IV. RESULTS

The validity of the methods presented is verified on the physical system, where the learned feedback policy is deployed to swing up different poles.

Feedback is crucial for this task for three reasons: i) the control actions to perform a successful swing up greatly depend on the

physical parameters of the specific pole, which are assumed to be unknown to the policy in this work, ii) these control actions depend on the initial position and orientation of the pole, which is likely to differ across trials on the real system, iii) even when the physical parameters and starting pose of the pole are well known, and a trajectory is generated in simulation for such a configuration, in the authors' experience this led to swing-ups with an offset in the final angle due to slight model mismatches. Feedback is thus needed to precisely control the final angle.

Throughout the following experiments, the initial grasping of the pole is achieved by a human holding the pole between the two tactile sensors. The gripper then slowly closes its fingers until the total force applied on the sensor by the pole reaches a user-defined threshold.

The student feedback control policy is evaluated on the real-world robotic system on four different poles with masses ranging from 20 g to 38 g, lengths from 20 cm to 35 cm, and radii from 2.5 mm to 5 mm. For each pole, the control policy is run ten times and the error from 180° in the final estimated angle $\hat{\phi}$ is recorded. Experiments show that all four poles are successfully swung up to an upright position, and a mean absolute error of 4.3° is achieved. A detailed analysis of the experimental results is provided in an experimental report [36]. These results demonstrate how a single policy is able to adapt the robot's motion to perform swing-up maneuvers for a wide range of different poles without any prior knowledge of the pole's physical features, based on the feedback provided by the tactile sensor. The resulting behaviour of the policy for one of the listed poles is depicted in Fig. 9. The supplementary video¹ contains the trajectories for the remaining poles. It is vital to note that the pole shown in Fig. 9 is not contained in the distribution of poles that is used while learning either the teacher or the student policy.

Moreover, the policy is transferred directly from the simulation to the real system with no adaptation needed. This further asserts the robustness of the policy as it is able to adapt to the real system that exhibits dynamics that are not modeled in the simulation (such as dynamic effects of the sensor material, unmodeled dynamics of the actuators, and delays of the actuator commands).

V. CONCLUSION

In this letter, a strategy has been presented to transfer tactile control policies for the swing-up manipulation of different poles from simulation to a physical robotic system. As the simulator has been shown to closely match the dynamics of the real system,

¹<https://youtu.be/In4jkaHzJLc>

the policy learned in simulation generalizes to the real-world robotic system with no adaptation needed.

Note that the system presented here does neither exploit a mechanically fixed pivot point nor directly control the rotational degree-of-freedom of the pole, but it can achieve the desired motion only through the presence of friction, whose modeling was crucial in enabling a realistic simulation.

This constitutes an important step towards a general framework to learn a wide variety of tactile manipulation tasks safely in simulation. Yet, current results have only been demonstrated for a single task on a single robotic system. Future work will focus on several aspects to further extend the generalizability of this work. In a first step, the proposed framework could be utilized to learn other pole manipulation skills on the given system, e.g. the throwing and catching of a pole. While such a policy was already successfully learned in simulation (see Fig. 10), the transfer to the physical system requires further work due to non-idealities of the hardware. For instance, when the pole is thrown in the air, it may leave the plane to which the motion is assumed to be constrained. In fact, instead of relying on the planar nature of the manipulation task, as was done in this work, the suggested simulator could be extended to handle non-planar tasks. As a result, manipulation skills for grippers that can be controlled in six degrees of freedom could also be learned. Moreover, in this letter, hand-engineered features are extracted from the tactile observations, i.e. the orientation and total normal force acting on the pole. These features may not be relevant for other tasks, where learning such features end-to-end with the policy, e.g. using autoencoders, may further generalize the proposed framework.

ACKNOWLEDGMENT

The authors would like to thank M. Egli and M. Mueller for their contribution to the development of the system.

REFERENCES

- [1] C. Sferrazza and R. D'Andrea, "Design, motivation and evaluation of a full-resolution optical tactile sensor," *Sensors*, vol. 19, no. 4, p. 928, 2019.
- [2] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 1–8.
- [3] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl, "Learning by cheating," in *Proc. Conf. Robot Learn.*, 2020, pp. 66–75.
- [4] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Sci. Robot.*, vol. 5, no. 47, 2020.
- [5] H. Yousef, M. Boukallel, and K. Althoefer, "Tactile sensing for dexterous in-hand manipulation in robotics-a review," *Sensors Actuators A: Phys.*, vol. 167, no. 2, pp. 171–187, 2011.
- [6] Z. Kappassov, J.-A. Corrales, and V. Perdureau, "Tactile sensing in dexterous robot hands," *Robot. Auton. Syst.*, vol. 74, pp. 195–220, 2015.
- [7] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *IEEE Trans. Robot.*, vol. 27, no. 6, pp. 1067–1079, 2011.
- [8] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable manipulation with a tactile-reactive gripper," in *Robotics: Sci. Syst.*, 2020, pp. 1–8.
- [9] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, "Tactile dexterity: Manipulation primitives with tactile feedback," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020.
- [10] S. Tian *et al.*, "Manipulation by feel: Touch-based control with deep predictive models," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 818–824.
- [11] H. Van Hoof, N. Chen, M. Karl, P. van der Smagt, and J. Peters, "Stable reinforcement learning with autoencoders for tactile and visual data," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 3928–3934.
- [12] A. Church, J. Lloyd, R. Hadsell, and N. F. Lepora, "Deep reinforcement learning for tactile robotics: Learning to type on a braille keyboard," *IEEE Robot. Automat. Lett.*, vol. 5, no. 4, pp. 6145–6152, 2020.
- [13] C. Wang, S. Wang, B. Romero, F. Veiga, and E. Adelson, "Swingbot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 5633–5640.
- [14] Y. Karayiannidis, C. Smith, D. Kragic *et al.* "Adaptive control for pivoting with visual and tactile feedback," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 399–406.
- [15] T. Senoo, Y. Yamakawa, S. Mizusawa, A. Namiki, M. Ishikawa, and M. Shimajo, "Skillful manipulation based on high-speed sensory-motor fusion," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009, pp. 1611–1612.
- [16] Z. Ding, N. F. Lepora, and E. Johns, "Sim-to-real transfer for optical tactile sensing," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020.
- [17] Y. Wang, W. Huang, B. Fang, and F. Sun, "Elastic Interaction of Particles for Robotic Tactile Simulation," *arXiv:2011.11528*, 2020.
- [18] Z. Kappassov, J.-A. Corrales-Ramon, and V. Perdureau, "Simulation of tactile sensing arrays for physical interaction tasks," in *Proc. IEEE/ASME Int. Conf. Adv. Intell. Mechatronics*, 2020, pp. 196–201.
- [19] J. A. Joergensen, L.-P. Ellekilde, and H. G. Petersen, "RobWorkSim-an open simulator for sensor based grasping," in *Proc. 41st Int. Symp. Robot. ROBOTIK. VDE*, 2010, pp. 1–8.
- [20] S. Moio, B. León, P. Korkealaakso, and A. Morales, "Model of tactile sensors using soft contacts and its application in robot grasping simulation," *Proc. Robot. Auton. Syst.*, 2013, pp. 1–2.
- [21] M. Bauza, E. Valls, B. Lim, T. Sechopoulos, and A. Rodriguez, "Tactile object pose estimation from the first touch with geometric contact rendering," in *Proc. Conf. Robot Learn.*, 2020.
- [22] C. Sferrazza and R. D'Andrea, "Sim-to-real for high-resolution optical tactile sensing: From images to 3D contact force distributions," *arXiv:2012.11295*, 2020.
- [23] C. Sferrazza, T. Bi, and R. D'Andrea, "Learning the sense of touch in simulation: A sim-to-real strategy for vision-based tactile sensing," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2020.
- [24] M. Ghafoorian *et al.*, "Location sensitive deep convolutional neural networks for segmentation of white matter hyperintensities," *Sci. Rep.*, vol. 7, no. 1, pp. 1–12, 2017.
- [25] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet V2: Practical guidelines for efficient CNN architecture design," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 116–131.
- [26] C. Sferrazza, A. Wahlsten, C. Trueeb, and R. D'Andrea, "Ground truth force distribution for learning-based tactile sensing: A finite element approach," *IEEE Access*, vol. 7, pp. 173 438–173 449, 2019.
- [27] D. Scaramuzza, A. Martinelli, and R. Siegwart, "A toolbox for easily calibrating omnidirectional cameras," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2006, pp. 5695–5701.
- [28] J. Hwangbo, J. Lee, and M. Hutter, "Per-contact iteration method for solving contact dynamics," *IEEE Robot. Automat. Lett.*, vol. 3, no. 2, pp. 895–902, 2018.
- [29] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2012, pp. 5026–5033.
- [30] K. Weiler and P. Atherton, "Hidden surface removal using polygon area sorting," *ACM SIGGRAPH*, vol. 11, no. 2, pp. 214–222, 1977.
- [31] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning With a Stochastic Actor," in *Proc. Int. Conf. on Mach. Learn.*, vol. 80. PMLR, 2018, pp. 1856–1865.
- [32] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *the int. J. Robot. Res.*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [33] A. Raffin, A. Hill, M. Ernestus, A. Gleave, A. Kanervisto, and N. Dormann, "Stable baselines3," <https://github.com/DLR-RM/stable-baselines3>, 2019.
- [34] R. O. Duda and P. E. Hart, "Use of the hough transformation to detect lines and curves in pictures," *Commun. ACM*, vol. 15, no. 1, pp. 11–15, 1972.
- [35] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2011, pp. 627–635.
- [36] T. Bi, C. Sferrazza, and R. D'Andrea, "Zero-Shot Sim-to-Real Transfer of Tactile Control Policies for Aggressive Swing-Up Manipulation - Experimental Report!" <https://arxiv.org/src/2101.02680/anc/ExperimentalReport.pdf>.