

Fine Manipulation Using a Tactile Skin: Learning in Simulation and Sim-to-Real Transfer

Ulf Kasolowsky and Berthold Bäuml

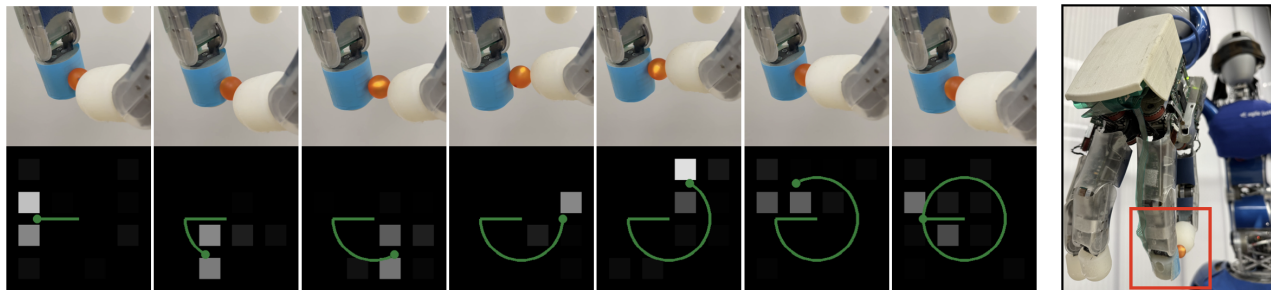


Fig. 1. Time sequence of a fine manipulation task ($\Delta t = 1$ s between images). The marble has to be rolled between two fingers along a specified trajectory (here a circle) using the feedback from the tactile skin (blue fingertip, see also Fig. 2, Right). The lower row shows the actual tactile image with 4×4 taxels. The task is executed on the torque-controlled DLR-Hand II [1] which is mounted on the humanoid robot DLR Agile Justin [2] (on the right). Only the hand's position and torque sensors and the tactile skin are used, but no, e.g., visual input.

Abstract—We want to enable fine manipulation with a multi-fingered robotic hand by using modern deep reinforcement learning methods. Key for fine manipulation is a spatially resolved tactile sensor. Here, we present a novel model of a tactile skin that can be used together with rigid-body (hence fast) physics simulators. The model considers the softness of the real fingertips such that a contact can spread across multiple taxels of the sensor depending on the contact geometry. We calibrate the model parameters to allow for an accurate simulation of the real-world sensor. For this, we present a self-contained calibration method without external tools or sensors. To demonstrate the validity of our approach, we learn two challenging fine manipulation tasks: Rolling a marble and a bolt between two fingers. We show in simulation experiments that tactile feedback is crucial for precise manipulation and reaching sub-taxel resolution of < 1 mm (despite a taxel spacing of 4 mm). Moreover, we demonstrate that all policies successfully transfer from the simulation to the real robotic hand. Website: aidx-lab.org/skin/iros24

I. INTRODUCTION

The dextrous manipulation of small objects is a key skill required for all kinds of industrial tasks like sorting or assembling. Humans excel in fine manipulation using their dextrous hands in combination with their spatially resolved sense of touch. In robotics, only recently, the advent of modern learning methods based on excessive training in simulation enabled breakthroughs in dextrous manipulation with multi-fingered hands, e.g., in-hand manipulation. However, for learning fine manipulation tasks, in addition a precise but computationally efficient simulation of tactile sensors is needed.

The authors are with the Learning AI for Dextrous Robots Lab (aidx-lab.org), Technical University of Munich, Germany, and the DLR Institute of Robotics and Mechatronics, German Aerospace Center. Email: {ulf.kasolowsky|berthold.baeuml}@tum.de

In this work we use a tactile skin glued to the soft fingertip of a multi-fingered hand (see Fig. 1 and Fig. 2). Although camera-based sensors like GelSight [3] and DIGIT [4] provide a higher spatial resolution and are more widely used, they are bulky and can not be easily added to an existing robot structure (e.g., the complete fingertip has to be replaced).

Important for all sensors is that they are soft so that in contact, the object can penetrate to some extent, resulting in a contact area and not only an uninformative point contact (Fig. 2). Regardless of the used sensor type, the challenge in simulating the spatially resolved tactile response is the need to model this soft contact. Vanilla rigid body simulators, which are the standard in reinforcement learning because of their high computational efficiency, do not suffice: in these simulators contacts between two colliding objects result only in a single (or few) contact point(s), i.e., no spatially resolved tactile image is provided.

A. Related Work

In this section, we discuss the existing work on tactile simulation and robotic fine manipulation.

Tactile Simulation: The most accurate way for simulating the tactile sensor response would be to incorporate softness and deformations directly in the physics simulator. Narang et al. [5, 6] used Finite element methods (FEM) for the (soft) BioTac sensor. However, FEM-based simulations are infeasibly slow to be used in deep reinforcement learning. Moio et al. [7] simulate softness by evaluating a penalty-based contact model at a fixed set of points on a taxel-based sensor. The computed reaction forces at each point are used to compute the response of the tactile sensor but also to simulate the physics. Xu et al. [8] use a similar model to simulate the contact dynamics and sensor response of a vision-based sensor. Another custom soft-contact model for an optical tactile sensor is presented in Ding et al. [9].




Fig. 2. **Left:** Tactile response due to soft contact. The top row shows two contact situations of a bolt with a tactile sensor on a cylindrical fingertip. Only if the contact is soft, the bolt will penetrate the fingertip resulting in a contact area causing multi-taxel response instead of a single contact point. The second row depicts the pressure distribution on the fingertip and the bottom row the resulting response for the individual taxel.

Right: Tactile sensor. The top row shows the sensor without and with the rubber glove covering (for better grip). The sensor has 4×4 taxels, which each measure $2.5 \text{ mm} \times 2.5 \text{ mm}$ and are arranged on a 4 mm grid, leading to a 1.5 mm in between. At the bottom, details of the sensor and contact modeling are depicted as needed in Section II-A. The discretization of the sensor is shown for an exemplary resolution of 0.5 mm per tactile point. Red tactile points correspond to the taxel regions \mathcal{T}_j , whereas blue points lie within the gaps and purple points form a margin around the sensor. These additional points are needed to correctly model the contact at each point of the sensor.

Many recently developed tactile simulators (mainly for vision-based sensors) follow a more efficient approach by decoupling the physics simulation from the actual simulation of the tactile response. Typically, a fast rigid body simulation engine is used which feeds its output (object poses, penetration depth, forces, ...) to a more or less approximated soft contact model for computing the tactile response. Gomes et al. [10] assume a known penetration of a manipulated object into the sensor, i.e., there is no dependence on the contact force. Other works assume compliant contacts [11, 12, 13], but rely directly on the penetration depth from the rigid body simulation to compute the tactile response. While compliant contact forces linearly relate the penetration to the applied force, the penetration also depends on the number of contact points detected by the rigid-body simulator. This can lead to sudden jumps and makes the model unreliable. The TACTO simulator [14] avoids this problem, as the mapping between the force and penetration is separated from the physics simulation. This allows the use of non-linear mappings. However, this makes the penetration only dependent on the force but not on the contact geometry. The penetration is observed by a camera within the simulated sensor and often an illumination model and/or a learned domain-adaptation is used to generate realistic images.

For taxel-based sensors, commonly multiple rigid bodies are used to simulate the tactile forces. One approach is to introduce softness by using compliant contacts for each

taxel [15, 16]. However, the simulation is not accurate enough such that the response of each taxel is reduced to a binary signal losing sub-taxel precision. Alternatively, the softness can be introduced by using a mass-spring-damper systems [17]. Using multiple bodies has the drawback that the contact geometry is altered leading to unnatural dynamics.

Fine Manipulation: There exists only little work on learning fine manipulation of small objects in simulation with a successful sim-to-real transfer. In Lambeta et al. [4], a multi-fingered hand rolls a marble between two vision-based DIGIT sensors to a target position on the fingertip. This task is closely related to our’s (see Fig. 1) but the surfaces of the DIGITS are flat in contrast to the round fingertips in our setting, making it significantly easier to keep a stable contact. Moreover, the marble is rolled to single target positions while we aim to follow a continuous trajectory. Furthermore, they do not learn the rolling in simulation but train a dynamic model with real-world data and apply model predictive control. This approach was first used in Tian et al. [18] but for a simpler setting: rolling a marble between a GelSight-style tactile sensor and a fixed flat surface making the task significantly easier. Church et al. [11] learn this easier task with simulation-based reinforcement learning and demonstrate a successful sim-to-real transfer. Do et al. [19] focus especially on the inter-finger manipulation of small objects like screws, which are picked up from a bowl between two vision-based tactile sensors and are reoriented for a better classification. However, the reorientation is not learned but follows a handcrafted strategy.

B. Contributions

Our main contributions are as follows:

- A novel model for a tactile skin with soft contacts coupled to a regular rigid-body simulator that does not use the incorrect penetration depths from the simulator but computes for each tactile image the penetration depth from a force equilibrium.
- A self-contained calibration procedure to tune the parameters of the tactile skin model. No external sensors or a test bench setup is needed, but only a cap with an indenter has to be mounted on a fingertip.
- Two fine manipulation tasks, which highlight the need for tactile feedback, are presented: rolling a marble between two fingers on a specified trajectory and rolling an arbitrary oriented bolt along its normal direction. The tasks are solved via reinforcement learning using straightforward rewards.
- A detailed analysis with simulation experiments demonstrating the advantages of tactile feedback for fine manipulation and also that our skin model captures important effects allowing, e.g., sub-taxel precision.
- Experiments on the real DLR-Hand II demonstrate the successful zero-shot sim2real transfer. Although both policies directly use the raw tactile signal, the sim2real transfer works, proving the validity of our tactile skin simulation model.

- To our best knowledge, the two shown fine manipulation tasks are the most challenging ever reported on a real robotic system.

II. SIMULATING A TACTILE SKIN

First, we present our general model idea. Afterward, we detail the implementation specific for the used TekScan Sensor and the fingertips of the DLR-Hand II.

A. Skin Model

General Model Idea: First of all, the surface of the sensor is discretized into a set of N_t tactile points

$$\mathcal{T} = \{t_i = (\vec{p}_{t,i}, \vec{n}_{t,i}, a_{t,i}) : i = 1, \dots, N_t\}, \quad (1)$$

where $\vec{p}_{t,i} \in \mathbb{R}^3$ describes the Cartesian position, $\vec{n}_{t,i} \in \mathbb{R}^3$ the surface normal vector, and $a_{t,i} \in \mathbb{R}$ the surface area depending on the discretization grid (see Fig. 3, top left). We assume that the indenter penetrates the fingertip by a maximum distance ϵ_{\max} in the direction of the contact normal \vec{n}_c . As a result, we can define a local penetration ϵ_i for every tactile point t_i (see Fig. 3, top right). We now assume a purely linear relationship between the local penetration and sensor response $\sigma_i = E\epsilon_i$, where the factor E is a tunable parameter of the model. The notation was chosen to resemble Hooke's law: σ_i describes the stress at the fingertip in the direction of the contact normal, and E is comparable to a modulus of elasticity.

Each taxel of the sensor patch consists of a subset of all tactile points:

$$\mathcal{T}_j \subset \mathcal{T}, \quad j = 1, \dots, N_T, \quad (2)$$

where N_T is the number of taxels of the sensor. Therefore, the force applied to the region of each taxel can be calculated by summing the normal stresses σ_i of all the respective tactile points:

$$F_j = \sum_{t_i \in \mathcal{T}_j} \sigma_i a_{t,i} \max(\vec{n}_c \cdot \vec{n}_{t,i}, 0), \quad j = 1, \dots, N_T. \quad (3)$$

The additional factor $\max(\vec{n}_c \cdot \vec{n}_{t,i}, 0)$ projects the area of every tactile point onto the normal direction. The lower bound of zero excludes tactile points from the sum that face in the opposite direction of the contact normal. Moreover, we assume that the sensor only senses forces in its local normal direction but no shear forces. If the indenter is rigid, the elastic fingertip will bend around it. The stresses σ_i act in the direction of the contact normal but the normal of the fingertip surface is orientated as the surface normal of the indenter (see Fig. 3, bottom left). Therefore, an additional projection factor is introduced. For every tactile point, the stress is projected on the respective surface normal: $(-\vec{n}_{s,i} \cdot \vec{n}_c)$. Finally, the taxel values T_j of the real sensor lie between 0 and 255. Hence, an additional, taxel-specific scaling parameter S_j is introduced. Eventually, the sensor model to compute the value of every taxel is given as

$$T_j = S_j \sum_{t_i \in \mathcal{T}_j} E\epsilon_i a_{t,i} (\vec{n}_c \cdot \vec{n}_{t,i}) (-\vec{n}_{s,i} \cdot \vec{n}_c), \quad j = 1, \dots, N_T. \quad (4)$$

Interplay with Rigid-Body Simulator: As mentioned above, the model is designed to work on top of a regular rigid-body simulator. Hence, the maximum penetration depth is not modeled correctly but must be computed separately. The simulator provides the relative position of the objects, the contact normal \vec{n}_c , and the normal force F_n . The local penetrations ϵ_i depend solely on the maximum penetration ϵ_{\max} , which can be varied by moving the relative position along the direction of the provided contact normal. The maximum penetration is chosen such that the sum of all local stresses matches the reported normal force:

$$F_n \stackrel{!}{=} \sum_{t_i \in \mathcal{T}} E\epsilon_i(\epsilon_{\max}) a_{t,i} (\vec{n}_{t,i} \cdot \vec{n}_c). \quad (5)$$

Obtaining the Local Penetrations: The local penetration depths ϵ_i are computed using an operation called ray casting, which determines the distance in a certain direction between a point in space and a mesh object (see Fig. 3, bottom right). First of all, the contact between the fingertip and the indenter is resolved by moving the tactile points along the contact normal by a large enough constant D so that they all lie outside of the indenter:

$$\vec{s}_i = \vec{p}_{t,i} - D\vec{n}_c, \quad i = 1, \dots, N_t. \quad (6)$$

From these shifted points, a ray cast in the direction of the contact normal determines the distances d_i to the indenter. The offset δ between the starting points and the deepest point of the mesh can be computed as $\delta = \min d_i$. The local penetrations are then given by subtracting the measured distances d_i , corrected by the offset δ :

$$\epsilon_i(\epsilon_{\max}) = \max(\epsilon_{\max} - (d_i - \delta), 0), \quad i = 1, \dots, N_t. \quad (7)$$

The maximum operation is required to account for tactile points that are not penetrated by the indenter, which would otherwise lead to negative penetration depths. The ray casting itself is independent of the maximum penetration as it only occurs in the formula for the local penetrations. Hence, the expensive ray cast operation must only be executed once before searching the maximum penetration satisfying the equilibrium defined in (5).

B. Model Implementation

The tactile skin is bent around the cylindrical part of the fingertip. The position can be described by three parameters:

- y : The offset along the cylinder's longitudinal axis, measured as a distance.
- β : The offset along the perimeter of the cylinder, measured as an angle around the longitudinal axis.
- α : The orientation of the sensor on the cylinder, measured as an angle around the surface normal of the cylinder.

The discretization of the sensor is shown in the bottom right of Fig. 2. For our implementation we use a resolution of 0.25 mm per taxel and a margin of 8 mm leading to 14 884 tactile points. We incorporate our model into the rigid-body simulator PyBullet [20], where we simulate two fingers of




Fig. 3. Visualization of the tactile skin model as detailed in Section II-A.

the DLR-Hand II. For computing the ray casts, we rely on the Python library Trimesh [21]. Finding the correct ϵ_{\max} is done via a line search with a step size of 0.1 mm. While the physics simulation runs at a rate of 1000 Hz, the soft contact model for the tactile images is evaluated every 100th step. The simulation of 0.1 s takes in total 56 ms, where the computation of the tactile images takes 26 ms; i.e., the simulation runs with a real-time factor of 1.8 on a small single-core AMD EPYC 7B13 @ 3 GHz. The bottleneck is the expensive ray cast operation as the line search takes only 0.3 ms.

III. MODEL CALIBRATION

To transfer control policies learned in simulation to the real-world system, the simulation must resemble reality as accurately as possible. Therefore, this chapter presents a calibration procedure for the model parameters. A perfect calibration is not required for a successful sim-to-real transfer, but identifying reasonable parameter intervals for domain randomization is sufficient.

A. Self-Contained Data Generation

A cap with a spherical indenter is mounted to the thumb and the indenter is pressed against the tactile skin attached to the forefinger at a set of desired contact points (see Fig. 4). At each point, the desired contact force is sampled uniformly between 1 N and 4 N and is applied by commanding the required torques to the fingers. The measured contact points $\vec{p}_{c,k}$ can be calculated with the forward kinematics of the fingers and the measured joint angles. Moreover, the real contact force $F_{n,k}$ can be reconstructed by the measured joint torques. We want to calibrate the skin at sub-taxel precision. Hence, we keep only samples where at least two taxels were active in the measured tactile image \mathbf{T}_k during contact. Eventually, the data set includes $N_P = 24$ valid samples including the measured contact positions, contact forces, and tactile images:

$$\mathcal{D} = \{(\vec{p}_{c,k}, F_{n,k}, \mathbf{T}_k) : k = 1, \dots, N_P\}. \quad (8)$$

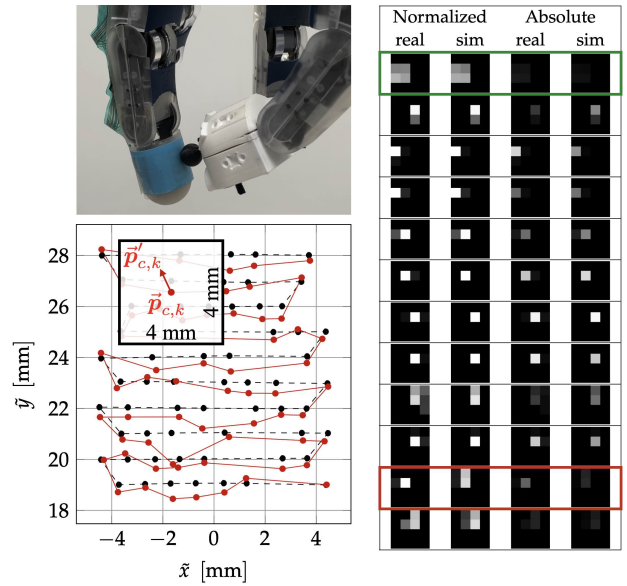


Fig. 4. **Left:** Self-Contained data generation. The top image shows the cap with an spherical indenter of radius 6 mm. The bottom image shows the set of 60 desired contact points, randomly selected from a uniform grid with a resolution of 1 mm. While the desired contact points are shown in black, the measured contact points are shown in red. Moreover, the slack introduced for each measured contact point is visualized.

Right: Results of the micro variable optimization for an exemplary set of macro variables. The green box marks a positive result where the images match well between simulation and reality. The red box marks a negative example that can be caused if the slack of ± 2 mm is not sufficient.

B. Model Fitting as an Optimization Problem

The goal of the calibration is to optimize the model parameters so that the simulated tactile images best match the measured ones:

$$\arg \min_{E, S, y, \beta, \alpha} \frac{1}{16N_V} \sum_k \underbrace{\|\mathbf{T}(\vec{p}_{c,k}, F_n, E, S, y, \beta, \alpha) - \mathbf{T}_k\|_2^2}_{\mathcal{L}_{\text{mse}}}, \quad (9)$$

where each simulated image is generated for the measured contact position and force.

However, small errors in the measured contact position due to an imperfect kinematic model of the fingers can cause significantly different tactile images. Although we calibrated the kinematic model [22], the remaining uncertainty of a few millimeters is still too high as the taxels measure only 2.5 mm. Hence, we additionally introduce the estimated contact positions $\vec{p}'_{c,k}$ to the optimization problem that allow a slack of ± 2 mm around each measured contact point $\vec{p}_{c,k}$ (see Fig. 4, bottom left).

To solve the optimization problem, it is divided into two stages as introduced in Fig. 5. For a given set of macro variables, the estimated contact positions are optimized using the normalized tactile images, because the information on the contact position lies purely in the ratio of the taxel values but not in their absolute values. The optimal scaling can then be obtained analytically by solving the least squares problem for the absolute images. The results of an exemplary micro variable optimization are shown on the right side of Fig. 4.

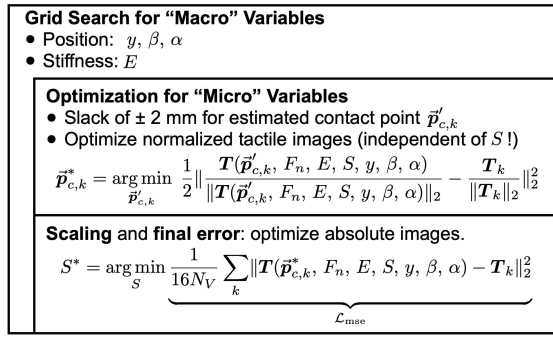


Fig. 5. Calibration scheme.

Due to the slack introduced by the micro variables, the grid search results do not have a sharp minimum. Instead they show a valley with a loss of $\mathcal{L}_{mse} \leq 25$, i.e., a root mean square error of less than 5 per taxel. We take the boundaries of this valley as the intervals for the calibration result:

- $y \in [21.75 \text{ mm}; 25.25 \text{ mm}]$,
- $\beta \in [0^\circ; 5.03^\circ]$,
- $\alpha \in [-15^\circ; 10^\circ]$,
- $E \in [236 \text{ MPa m}^{-1}; 848 \text{ MPa m}^{-1}]$.

The scaling factors corresponding to all macro variable combinations with $\mathcal{L}_{mse} \leq 25$ form a coherent cluster. Again, we take the boundary values to obtain the interval of reasonable values:

- $S \in [61 \text{ N}^{-1}; 76 \text{ N}^{-1}]$.

IV. APPLICATION TO LEARNING FINE MANIPULATION

We want to learn two challenging fine-manipulation tasks in simulation. The first is rolling a marble along arbitrary trajectories between the fingers. The second is rolling a bolt following a sinus motion in different orientations. The policy has to detect the orientation and roll the bolt in the according direction. To learn these tasks, we first introduce the simulation setup before specifying the learning details.

A. Simulation and Environment Setup

The relevant fingers of the DLR-Hand II and the manipulated object are loaded into the PyBullet simulator. The model and simulation of the torque-controlled fingers follows the work of Sievers et al. [23]. The fingers start already in a pinch grasp configuration and the object is placed between the finger tips. The focus of the learning is to utilize the tactile signal and not finding the correct grasp. For the bolt rolling task, the initial orientation α_b is randomly selected from the set $\{-45^\circ, 0^\circ, 45^\circ, 90^\circ\}$ at the beginning of each episode. The manipulated object is initially kept in place with a position constraint. This constraint is removed when the fingers grasped the object successfully, i.e., both fingers are in contact. When the object is dropped, i.e., one of fingers lost contact, the current episode is terminated in failure, the simulation is reset to its initial state, and a new episode begins. Furthermore, we limit the duration of each episode to 6s. After this timeout, the environment is reset without sending a termination signal to the reinforcement learning algorithm. The control rate of the real robotic system is

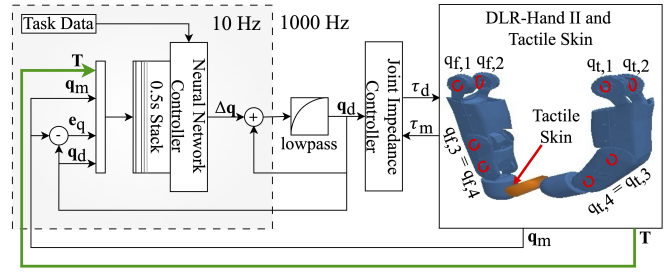


Fig. 6. Embedding of the policy into the overall control architecture. The input of the neural network is a subset of the stack of the last five observations ($\mathbf{T} \in \mathbb{R}^{16}$ and $\mathbf{q}_m, \mathbf{q}_d, \mathbf{e}_q \in \mathbb{R}^6$), and task specific data. The output is a six-dimensional change of the desired joint angles.

$f_{sys} = 1000 \text{ Hz}$. Hence, we chose the same rate for the physics simulation. In contrast, the control policy is evaluated every 100th environment step leading to a controller sample rate of $f_{cont} = 10 \text{ Hz}$. The overall control architecture is summarized in Fig. 6. As the soft contact model for the tactile images, does not influence the physics simulation but only provides the tactile images, it is evaluated at the same rate as the policy.

Action Space: The final output of the controller must be a vector of desired joint angles $\mathbf{q}_d \in \mathbb{R}^6$ that are commanded to the underlying joint impedance controller of the thumb and forefinger. The output of the policy $\mathbf{a} \in [-1; 1]$ is added to the latest desired joint angles. This corresponds to commanding a joint velocity and, hence, the output \mathbf{a} is scaled to not exceed the maximum joint velocity \dot{q}_{max} . Moreover, the output is clipped to stay within the joint limits:

$$\mathbf{q}_{out} = \text{clip}(\mathbf{q}_d + \Delta\mathbf{q}, \mathbf{q}_{min}, \mathbf{q}_{max}), \quad \Delta\mathbf{q} = \mathbf{a} \frac{\dot{q}_{max}}{f_{cont}}. \quad (10)$$

Finally, to avoid jumps in the commanded signal, the commanded joint angles are smoothed by a first-order low-pass filter running at f_{sys} .

Observation Space: The joints of the DLR-Hand II are torque-controlled. Hence, a difference between the desired \mathbf{q}_d and measured joint angles \mathbf{q}_m exists that includes the information on the current torque: $\boldsymbol{\tau} = \mathbf{K}(\mathbf{q}_d - \mathbf{q}_m)$, where \mathbf{K} is the stiffness of the impedance controller. Therefore, the control error $\mathbf{e} = (\mathbf{q}_d - \mathbf{q}_m)$ includes already some tactile information. Most robotic hands, however, do not have torque sensors. Hence, their controller can only set the desired joint angles. Any elasticity results purely from the mechanical structure but cannot be measured. This can be imitated by passing only the desired joint angles to the policy. We train two policies for each task with (*Tactile*) and without (*Desired*) the skin feedback \mathbf{T} in addition to the desired joint angles. To demonstrate the effectiveness of the tactile skin, we compare it to a third policy (*Joint Angles*) that can leverage the torque information, which is specifically available for the DLR-Hand II. The different observation spaces are summarized in Table I. To further enrich the information, the aforementioned observations are stacked over the last five iterations so that the policy sees a history of 0.5s. Moreover, all policy receive the desired

TABLE I
OBSERVATION SPACE OF DIFFERENT POLICIES.

Observation		Desired	Joint Angles	Tactile
Tactile image	\mathbf{T}			✓
Desired joint angles	\mathbf{q}_d	✓	✓	✓
Measured joint angles	\mathbf{q}_m		✓	
Control error	\mathbf{e}_q		✓	
Target position	$\vec{\mathbf{p}}_{d,k}$	✓	✓	✓
Target movement	$\Delta\vec{\mathbf{p}}_{d,k}$	✓	✓	✓

position of the manipulated object on the sensor $\vec{\mathbf{p}}_{d,k}$ and the difference to the previous iteration $\Delta\vec{\mathbf{p}}_{d,k} = \vec{\mathbf{p}}_{d,k} - \vec{\mathbf{p}}_{d,k-1}$. For the marble, this is a two-dimensional vector as it can be an arbitrary position on the fingertip. In contrast, for the bolt it is one-dimensional as it only specifies the position of the up-down movement but the orientation must be figured out by the policy itself.

B. Task Definition

To formalize the two tasks, we have to define a reward function for the reinforcement learning algorithm. The individual components are presented below.

Safety Critical Rewards: To prevent self-destructive behavior, we introduce penalty terms for getting close to the physical limitations of the hand. For the penalty terms we use the following function:

$$\text{pen}(x; \Delta x) = \begin{cases} 0, & x < 0 \\ \frac{1}{2} \left(1 - \cos\left(\frac{x}{\Delta x}\pi\right)\right), & 0 \leq x \leq \Delta x \\ 1, & x > \Delta x. \end{cases} \quad (11)$$

We penalize running into joint, joint velocity, and torque limits, where $x_{q/\dot{q}/\tau}$ describe how much the values exceed the soft margins, which are $\Delta x_{q/\dot{q}/\tau}$ away from the hard constraints:

$$p_{q/\dot{q}/\tau} = \lambda_{q/\dot{q}/\tau} \text{pen}(x_{q/\dot{q}/\tau}; \Delta x_{q/\dot{q}/\tau}). \quad (12)$$

Holding Rewards: The focus of the learning is to utilize the skin signal. Hence we define two rewards pointing the policy towards the correct grasp. The first just gives a binary reward for touching the object with the tip of the thumb and the cylindrical part of the forefinger:

$$r_g = \begin{cases} 0, & \text{no or incorrect contact} \\ \lambda_g, & \text{correct contact.} \end{cases} \quad (13)$$

The second reward is granted for holding the object with a normal force F_n of roughly 1.5 N: $r_f = \lambda_f \text{rew}(|F_n - 1.5 \text{ N}|; 1.5 \text{ N})$, where the general shape of the reward function is given as

$$\text{rew}(x; \Delta x) = \begin{cases} 0, & x < 0 \\ \frac{1}{2} \left(1 + \cos\left(\frac{x}{\Delta x}\pi\right)\right), & 0 \leq x \leq \Delta x \\ 1, & x > \Delta x. \end{cases} \quad (14)$$

Marble Rolling: To learn rolling the marble along arbitrary trajectories, we reward the policy for keeping the measured contact point $\vec{\mathbf{p}}_k$ of the marble on the fingertip close to the desired position $\vec{\mathbf{p}}_{d,k}$:

$$r_p = \lambda_p \text{rew}(\|\vec{\mathbf{p}}_{d,k} - \vec{\mathbf{p}}_k\|_2, 5 \text{ mm}). \quad (15)$$

TABLE II
DOMAIN RANDOMIZATION.

Parameter	Unit	Distribution	Mode	
Joint offsets ¹	q_{off}	rad	$\mathcal{U}(-0.04, 0.04)$	per episode
Joint noise	q_{noise}	rad	$\mathcal{N}(0, 0.02)$	per step
	y	mm	$\mathcal{U}(21.5, 25.5)$	per episode
Sensor position	β	°	$\mathcal{U}(-12, 12)$	per episode
	α	°	$\mathcal{U}(-15, 15)$	per episode
Elasticity	E	MPa m ⁻¹	$\mathcal{U}(236, 848)$	per episode
Scale	S_j	N ⁻¹	$\mathcal{U}(61, 76)$	per episode
Taxel offsets	T_{off}	1	$\mathcal{U}(-5, 5)$	per episode
Noise range	σ_{taxel}	1	$\mathcal{U}(0, 5)$	per episode
Taxel noise	T_{noise}	1	$\mathcal{N}(T_{\text{off}}, \sigma_{\text{taxel}})$	per step

The trajectories of the desired positions are generated randomly during training and move within a 6 mm × 6 mm box. The positions are defined relative to the sensor position. Eventually, the final reward for rolling the marble is given as

$$r = r_p + r_f + r_g - p_q - p_{\dot{q}} - p_{\tau}. \quad (16)$$

Bolt Rolling: For the bolt rolling task, we reuse the location reward from the marble rolling. However, the desired position provided to the policy describes only the position along one-dimensional sinus movement. Hence, the two-dimensional position vectors required by the reward are computed using the initial orientation of the bolt α_b . The motions have an amplitude of 6 mm and the frequency is sampled between 0.25 Hz and 0.5 Hz. Again the positions are specified relative to the sensor position. Moreover, an additional reward for keeping the bolt in its initial orientation relative to the sensor is introduced:

$$r_{\alpha} = \lambda_{\alpha} \text{rew}(|\alpha_b - \alpha_{b,k}|, 15^{\circ}). \quad (17)$$

Hence, the final reward for the bolt rolling task is given as

$$r = r_p + r_{\alpha} + r_f + r_g - p_q - p_{\dot{q}} - p_{\tau}. \quad (18)$$

C. Domain Randomization

For a successful Sim2Real transfer we apply domain randomization during simulation. See Table II for the most important parameters and the website for a complete list.

Hand Parameters: To introduce the uncertainties of the kinematic model we sample joint offsets q_{off} at the beginning of each episode that are added to the measured joint angles to imitate a constant error. Moreover, Gaussian noise is added to the measured joint angles in each step to simulate, for instance, slack in the joint gears. Furthermore, domain randomization is applied to the controller parameters.

Object parameters: While the bolt keeps a constant radius of 6 mm, we vary the radius of the marble from 4 mm to 8 mm. Domain randomization is also applied to the dynamic parameters such as mass or friction. Moreover, noise is added to the initial position and orientation of the object at the beginning of each episode.

¹Just before publication, we noticed that we sampled the joint offsets not independently – but this amplifies the domain randomization even more.

Skin parameters: The intervals for the domain randomization of the scaling and elasticity are taken from the calibration results. The sensor position is, however, based on the precision that can be achieved by manually gluing the sensor to the fingertip, as it is repeatedly removed between experiments. To allow a fair comparison, the sensor position is not randomized for policies that do not use the tactile feedback but fixed at the middle of the fingertip. The target trajectories are defined relative to the sensor position and without tactile feedback the policies cannot know the skin position. Hence, it is kept constant. Moreover, noise is added to the tactile images to better resemble the real sensor. At the beginning of each episode, constant offsets and noise ranges are sampled for each taxel. During the episodes, Gaussian noise with a mean of the offset and a standard deviation of the noise range is added to the tactile images.

D. Reinforcement Learning Training

To implement the simulation environment, we stick to the interface defined by the Gymnasium API standard [24]. This makes the environment compliant with most reinforcement learning frameworks. We use the Soft Actor Critic [25] algorithm implemented in Stable Baselines3 [26] to train the policies. We rely on simple multilayer perceptrons with two hidden layers for the actor and critic networks. Each hidden layer has 256 neurons, and ReLU activation functions are used. All network inputs are normalized, with the means and variances of all input dimensions obtained as a running average during training. The networks are trained with the Adam optimizer [27]. We use asymmetric observations [28, 29, 23, 30] to provide additional information to the critic networks. We train the marble policies for 800 000 and the bolt policies for 600 000 environment steps, while using 8 parallel environments.

V. RESULTS

We first highlight the importance of tactile feedback by analyzing the marble rolling policy in simulation. Afterward, we present the successful sim-to-real transfer of the policies.

A. Importance of Tactile Feedback

Marble: To evaluate the performance of the different marble rolling policies, a circular trajectory with a radius of 5 mm and a duration of 6 s is commanded after moving from the center to the side for 2 s. During the evaluation the domain randomization, such as the sensor position, is turned off to allow a fair comparison between the policies with and without tactile feedback. Just the randomization of the joint offset is left active to keep the effect of the kinematic uncertainties. The results are analyzed qualitatively in Fig. 7 and quantitatively in Fig. 8. Recall that the taxels measure 2.5 mm, i.e., the marble is manipulated at sub-taxel precision with tactile feedback as the tracking error is less than a millimeter. The feedback of the measured joint angles and the control error that include the tactile information of the torque sensors only marginally improves the performance compared to the policy receiving only the desired joint angles. Hence,

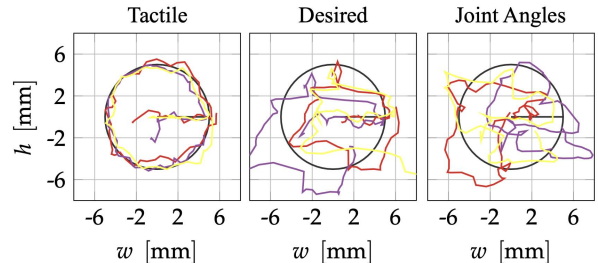


Fig. 7. Qualitative analysis of marble rolling. Three exemplary runs are shown for each policy. While the circle is tracked smoothly with tactile feedback, the trajectories without it follow the target position poorly and the circular pattern can only be recognized roughly.

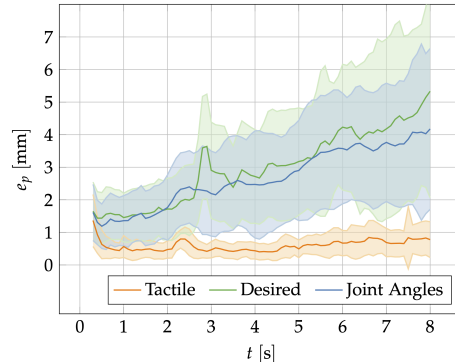


Fig. 8. Quantitative analysis of marble rolling. The plot shows the tracking error e_p over time for each policy. With tactile feedback it reduces from initially 1.5 mm to approximately half a millimeter. After 4.5 s, the error marginally increases as the manipulation makes the pinch grasp more and more unstable over time. After 8 s the error is approximately 0.8 mm. For the policies w/o tactile feedback, the tracking error monotonically increases from 1.5 mm to 5 mm (*Desired*) and 4 mm (*Joint Angles*) after 8 s.

tactile feedback is essential for the precise fine manipulation of small objects.

Bolt: To roll the bolt in the correct direction, the policy must estimate its orientation. During training we observe that this is only possible with tactile feedback but not with the desired joint angles alone. With access to the measured joint angles the policy can infer the orientation by ”wiggling motions”. However, this is limited to special robotic hands with precise torque control such as the DLR-Hand II.

B. Simulation to Reality Transfer²

For detailed results of the sim-to-real transfer, please refer to the accompanying video.

Marble: To test the marble rolling policy we command the same trajectory as for the simulation analysis except we add two more circles so that one run takes 20 s. Each policy was executed three times for a marble with a radius of 5 mm and 7 mm, respectively. Although all three policies achieved a success rate of 100%, the tactile policy tracked the circle significantly smoother than the ones without tactile feedback. An exemplary run is shown in Fig. 1. To push the tactile policy to its limits, we choose a trajectory that depicts the initials of the German Aerospace Centre (DLR). The letters are ”written” on the sensor by the marble sequentially. This

²To compensate for sensitivity loss due to wear and tear, before the shown experiments an overall scaling factor for the taxel response was determined by rolling the bolt on the real robot and in simulation.




Fig. 9. Writing letters with the marble in reality. The desired trajectory of each letter is shown in black and the estimated position of the marble is shown in orange. The estimated position is obtained by training a neural network to predict the marble position based on the tactile images. As no real-world test data is available, the result is only qualitative. The run is shown in the accompanying video.

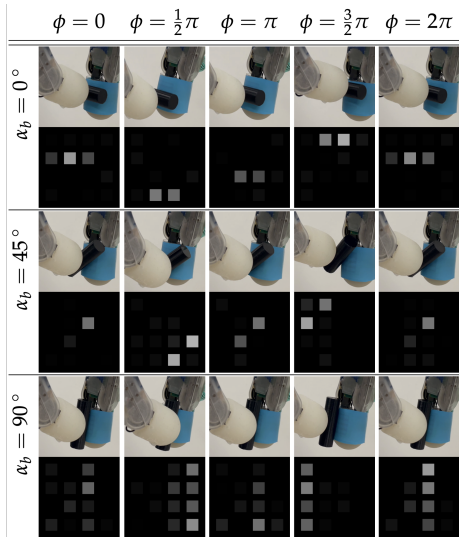


Fig. 10. Bolt rolling in reality for three different bolt orientations. ϕ denotes the phase of the sinusoidal rolling motion, which has a period of 3 s.

requires manipulation longer than 32 s, almost covering the entire sensor surface and including sharp corners. The results are presented in Fig. 9.

Bolt: Moreover, we show the successful sim-to-real transfer of the tactile bolt rolling policy. For every of the four possible orientations the policy was tested three times. In all runs, the bolt was rolled in the correct direction. In 11 out of the 12 runs the policy succeeded to run until the timeout after 24 s. The failed run dropped the bolt after 18 s, as it slipped out the fingers. Exemplary image sequences of the rolling motion are shown in Fig. 10.

VI. CONCLUSION

We have shown that tactile skin-based fine manipulation with a robotic hand can be learned in simulation and successfully transferred to the real world. For two challenging fine manipulation tasks, we performed detailed simulation experiments and proved the necessity of the tactile feedback. For the marble rolling, the tracking precision was significantly increased from about 5 mm without to less than 1 mm with the tactile skin, reaching clearly sub-taxel precision as the taxel spacing is 4 mm. The task of rolling a bolt along its normal direction even failed completely without the tactile feedback. Both outcomes prove that the policies depend on the details of the raw tactile signal. In experiments on the real robotic DLR-Hand II, the very same in simulation trained policies performed seamlessly and resembled the

same qualitatively behavior in both tasks.

To achieve these results, a novel model for the skin was introduced that can be used together with a regular rigid-body physics simulator. To allow a successful sim-to-real transfer, the model was calibrated with real-world data. The calibration process is self-contained and does not require any external sensors. To our knowledge, this is the first time that complex fine manipulation of such small objects with robotic fingers was successfully performed on a real robot system – at least for the case that it was trained purely in simulation.

REFERENCES

- [1] J. Butterfaß et al., “DLR-Hand II: Next generation of a dextrous robot hand,” in *Proc. IEEE International Conference on Robotics and Automation*, 2001.
- [2] B. Bäuml et al., “Agile justin: An upgraded member of dlr’s family of lightweight and torque controlled humanoids,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2014.
- [3] W. Yuan et al., “Gelsight: High-resolution robot tactile sensors for estimating geometry and force,” *Sensors*, vol. 17, no. 12, 2017.
- [4] M. Lambeta et al., “Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, 2020.
- [5] Y. S. Narang et al., “Interpreting and predicting tactile signals for the syntouch biotac,” *The International Journal of Robotics Research*, vol. 40, no. 12-14, 2021.
- [6] Y. Narang et al., “Sim-to-real for robotic tactile sensing via physics-based simulation and learned latent projections,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2021.
- [7] S. Moio et al., “Model of tactile sensors using soft contacts and its application in robot grasping simulation,” *Robotics and Autonomous Systems*, vol. 61, no. 1, 2013.
- [8] J. Xu et al., “Efficient tactile simulation with differentiability for robotic manipulation,” in *Proc. Conf. on Robot Learning*, 2023.
- [9] Z. Ding et al., “Sim-to-real transfer for optical tactile sensing,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2020.
- [10] D. F. Gomes et al., “Generation of gelsight tactile images for sim2real learning,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, 2021.
- [11] A. Church et al., “Tactile sim-to-real policy transfer via real-to-sim image translation,” in *Proc. Conf. on Robot Learning*, 2022.
- [12] Y. Lin et al., “Tactile gym 2.0: Sim-to-real deep reinforcement learning for comparing low-cost high-resolution robot touch,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, 2022.
- [13] —, “Bi-touch: Bimanual tactile manipulation with sim-to-real deep reinforcement learning,” *IEEE Robotics and Automation Letters*, 2023.
- [14] S. Wang et al., “Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, 2022.
- [15] Z. Ding et al., “Sim-to-real transfer for robotic manipulation with tactile sensory,” in *Proc. IEEE/RSS Int. Conf. on Intelligent Robots and Systems*, 2021.
- [16] L. Yang et al., “Tactgn: Learning tactile-based in-hand manipulation with a blind robot using hierarchical graph neural network,” *IEEE Robotics and Automation Letters*, vol. 8, no. 6, 2023.
- [17] A. Habib et al., “Skinsim: A simulation environment for multimodal robot skin,” in *Proc. IEEE Int. Conf. on Automation Science and Engineering*. IEEE, 2014.
- [18] S. Tian et al., “Manipulation by feel: Touch-based control with deep predictive models,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2019.
- [19] W. K. Do et al., “Inter-finger small object manipulation with densetact optical tactile sensor,” *IEEE Robotics and Automation Letters*, vol. 9, no. 1, 2023.
- [20] E. Coumans and Y. Bai, “Pybullet, a python module for physics simulation for games, robotics and machine learning,” 2016.
- [21] Dawson-Haggerty et al., “trimesh,” 2019. [Online]. Available: <https://trimesh.org/>
- [22] J. Tenhumberg et al., “Self-contained and automatic calibration of a multi-fingered hand using only pairwise contact measurements,” in *Proc. IEEE-RAS Int. Conf. on Humanoid Robots*, 2023.
- [23] L. Sievers et al., “Learning purely tactile in-hand manipulation with a torque-controlled hand,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2022.
- [24] G. Brockman et al., “Openai gym,” *arXiv preprint arXiv*, 2016.
- [25] T. Haarnoja et al., “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *Int. Conf. on Machine Learning*, 2018.
- [26] A. Raffin et al., “Stable-baselines3: Reliable reinforcement learning implementations,” *The Journal of Machine Learning Research*, vol. 22, no. 1, 2021.
- [27] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Int. Conf. on Learning Representations*, 2015.
- [28] M. Andrychowicz et al., “Learning dexterous in-hand manipulation,” *The International Journal of Robotics Research*, vol. 39, no. 1, 2020.
- [29] I. Akkaya et al., “Solving rubik’s cube with a robot hand,” *arXiv preprint*, 2019.
- [30] J. Pitz et al., “Dextrous tactile in-hand manipulation using a modular reinforcement learning architecture,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2023.