

# Evetac: An Event-based Optical Tactile Sensor for Robotic Manipulation

Niklas Funk<sup>1</sup>, Erik Helmut<sup>1</sup>, Georgia Chalvatzaki<sup>1,2</sup>, Roberto Calandra<sup>3,4</sup>, Jan Peters<sup>1,2,5</sup>

**Abstract**—Optical tactile sensors have recently become popular. They provide high spatial resolution, but struggle to offer fine temporal resolutions. To overcome this shortcoming, we study the idea of replacing the RGB camera with an event-based camera and introduce a new event-based optical tactile sensor called Evetac. Along with hardware design, we develop touch processing algorithms to process its measurements online at 1000 Hz. We devise an efficient algorithm to track the elastomer’s deformation through the imprinted markers despite the sensor’s sparse output. Benchmarking experiments demonstrate Evetac’s capabilities of sensing vibrations up to 498 Hz, reconstructing shear forces, and significantly reducing data rates compared to RGB optical tactile sensors. Moreover, Evetac’s output and the marker tracking provide meaningful features for learning data-driven slip detection and prediction models. The learned models form the basis for a robust and adaptive closed-loop grasp controller capable of handling a wide range of objects. We believe that fast and efficient event-based tactile sensors like Evetac will be essential for bringing human-like manipulation capabilities to robotics. The sensor design is open-sourced at <https://sites.google.com/view/evetac>.

**Index Terms**—Touch Sensing, Optical Tactile Sensor, Event-based Camera

## I. INTRODUCTION

**M**ANIPULATION of mechanical objects is essential for real-world robotic applications ranging from industrial assembly [1], [2] to household robots [3]–[5]. Physical manipulation of objects includes making and breaking contact between the robotic manipulator and the object of interest. In particular, the manipulator should apply sufficient forces for stable grasping, while preventing any loss of contact and damage to the objects by applying excessive force. For the goal of having reactive, adaptive, reliable, and efficient dexterous manipulation skills also in unstructured environments with little or no prior knowledge available, direct sensing of contacts, i.e., tactile sensing, is of crucial importance [6].

Due to their huge potential, there exists a long history in developing tactile sensors for robotics [6]–[9]. While a wide range of sensing technologies have been proposed, recently, especially RGB optical tactile sensors have received




Fig. 1: Two Evetac sensors installed in the ROBOTIS RH-P12-RN(A) gripper holding a pen. In the bottom left and right, we depict a magnified version of the sensors’ measurements. Evetac is an open-source event-based optical tactile sensor for robotic manipulation. Its main components are an illuminated, dotted, soft silicone gel that interacts with the environment. Changes in gel configuration are captured by a high-resolution event-based camera inside the sensor as shown in the bottom left & right.

increased attention [10], [11]. They are also known as vision-based tactile sensors, as their functioning principle relies on an RGB camera capturing an elastomer’s deformation. The advantages of RGB optical tactile sensors are their small form factors, compatibility with standard interfaces, low cost, and very high, human-like, or even superhuman spatial resolution. Yet, compared to other tactile sensing technologies and the human sense of touch, they typically lack temporal resolution. Alternative sensing technologies such as capacitive [12], piezoresistive [13], [14], magnetic [15], or pressure-based [16] tactile sensors are contrarily typically fast, but lack resolution.

Inspired by human tactile sensing capabilities and their different mechanoreceptors, which offer a high temporal resolution of resolving vibrations of at least up to 400 Hz [7], [17], as well as a high spatial resolution through approximately 17000 tactile sensing units per hand [17]–[19], with higher densities in the fingertips, in this paper, we propose a novel event-based optical tactile sensor called Evetac. While event-based optical tactile sensors have been presented previously

<sup>1</sup>Department of Computer Science, Technical University of Darmstadt, Darmstadt, Germany  
{niklas, georgia, jan}@robot-learning.de

<sup>2</sup>Hessian Centre for Artificial Intelligence

<sup>3</sup>LASR Lab, Technische Universität Dresden, Germany  
roberto.calandra@tu-dresden.de

<sup>4</sup>The Centre for Tactile Internet with Human-in-the-Loop (CeTI), Germany

<sup>5</sup>German Research Center for AI (DFKI), Research Department: SAIROL  
This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

[20]–[22], it still remains an underexplored area.

The proposed Evetac sensor is largely inspired by popular RGB optical tactile sensors such as GelSight [23], TacTip [24], and DIGIT [25]. The sensor’s functioning principle thus also relies upon a camera capturing the deformation of a soft silicone gel. The gel has imprinted markers which provide natural features for the interaction between sensor and object. What makes our Evetac sensor substantially different from common RGB optical tactile sensors is that we replace the RGB camera with an event-based camera. Event-based cameras recently gained lots of attraction due to their properties of high temporal resolution, high pixel bandwidth, and low energy consumption [26]. This camera replacement fundamentally changes the sensor’s properties as it allows for achieving a significantly increased temporal resolution. In our case, we obtain touch measurements at 1000 Hz, while at the same time pertaining a spatial resolution of 640x480 pixels. The improved temporal resolution can be beneficial for reliably detecting fast contact-related phenomena such as vibrations. Moreover, the sparse signal returned from event-based cameras, which only return signal upon intensity changes at the pixel locations, also enables real-time signal processing despite the high readout rate. This opens the door to closing tactile control loops at high frequencies.

From a design point of view, our proposed sensor attempts to maximize the re-use of existing components to reduce the entry barrier into the field of event-based optical tactile sensing. Apart from the housing, which has to be 3D-printed, all other components are commercially available. In addition to the hardware design, we develop the necessary software interfaces for reading out the sensor online, in real-time at 1000 Hz. We also present a novel, gradient-based algorithm for real-time tracking of the dots imprinted in the gel. The algorithm allows to keep track of the gel’s global state despite the sensor’s sparse outputs. The dots’ displacement can be used to reconstruct the shear forces acting on the sensor. Finally, Evetac’s raw sensor output and the information from the dot tracking provide the main features for developing slip detection and prediction models. In particular, we integrate Evetac into a commercially available robotic parallel gripper as shown in Figure 1 and present a data-driven approach for slip detection. For this task, we train and compare different neural network architectures that benefit from the expressive, low-dimensional features. The models can be evaluated online at 1000 Hz, and form the basis for the design of a closed-loop grasp controller operating at 500 Hz, capable of stably grasping a wide range of objects with different masses and materials.

In summary, our contributions are the design of a novel, open-source event-based optical tactile sensor called Evetac. The sensor design aims to maximize re-use of existing, commercially available components to mitigate the manufacturing barrier and incentivize reproducibility. Besides sensor design, we demonstrate the sensor’s high temporal resolution by sensing vibrations up to 498 Hz and showcase improved sensing efficiency w.r.t. data rate. Despite Evetac’s high sensing frequency, it generates significantly fewer data compared to RGB optical tactile sensors. Moreover, we provide real-

time touch processing algorithms. In particular, we devise an algorithm for tracking the dots imprinted in the gel with 1000 Hz and show its effectiveness for reconstructing the shear forces acting upon Evetac’s gel. Lastly, we showcase Evetac’s effectiveness for robotic grasping by training efficient data-driven neural networks for slip detection and prediction. The models allow integration into high-frequency feedback control loops for achieving robust and reliable grasping across a wide range of household objects. Importantly, we are able to show generalization of the slip detectors across objects and the controller’s adaptiveness w.r.t. object mass and reactivity upon grasp perturbation.

## II. RELATED WORKS

Tactile sensing [8] has a huge potential for robotics. Contact information is a crucial source of information to e.g., recover object properties [6], capture haptic information [27], or react to undesirable contact configurations for preventing slip and achieving stable grasping [9], [28]. In the following section, we will mainly focus on the latter aspect as progress in this direction holds the promise to improve the performance of every robotic manipulation system regarding reliability, robustness, and generalization to a wider range of objects.

**Vibration-based Tactile Sensing.** For stable grasping, humans make use of fast-adapting receptors to detect small localized slips that allow adaption of grasping force prior to gross slippage [29], [30]. Inspired by these fast human mechanoreceptors, several tactile sensors have been developed, offering high temporal resolutions [7], [31]–[33]. The corresponding touch processing algorithms for slip detection investigate the energy of potential vibrations [16], [29], [34], frequency-domain features in combination with neural networks [33], [35], signal coherence analysis [36], and data-driven approaches using the raw sensor data [37]. Regarding slip timing, [34] showed that their approach can detect slip more than 30 ms before an IMU accelerometer picked it up. While all these works present promising approaches to slip detection, almost all of them rely on special hardware, which is difficult to access and requires substantial manufacturing knowledge. Additionally, these tactile sensing technologies lack spatial resolution.

**Optical Tactile Sensors** contrarily offer very high spatial resolutions and have recently become relatively cheap to acquire, thereby, significantly reducing the entry barrier into the field. As standard RGB cameras are typically significantly slower than the previously presented sensors, different slip detection criteria have been developed. They include analysis of the marker displacement field [38]–[40], model-based criteria analyzing the inhomogeneity of the force field [41], and data-driven slip detectors [42] eventually combined with closed-loop feedback control [43] and multi-fingered hands [44]. Overall, these works on slip detection using RGB optical sensors rather focus on sensing displacements than high-frequency phenomena such as vibrations. Particularly, these RGB optical tactile sensors cannot offer the temporal resolutions found in human fast-adapting type II mechanoreceptors, which are sensitive to mechanical vibrations of at least up to 400 Hz [7], [17].

**Event-based Tactile Sensors**, on the contrary, are able to operate at human-like temporal resolutions. Moreover, their asynchronous functioning principle, i.e., only returning information whenever the quantity of interest changes, significantly reduces the overall sensor data rate and could potentially facilitate operating large tactile skins. Examples of event-based tactile sensing include [13], [45]–[47]. Yet, these tactile sensors cannot keep up with the spatial resolution offered by optical tactile sensors.

**Event-based Optical Tactile Sensing.** Recent progress in event-based cameras has led to the development of event-based optical tactile sensors that can eventually provide both high spatial and temporal resolution. Ward et al. [22] introduced the NeuroTac, an event-based optical tactile sensor based on TacTip [24]. They showed the sensor’s effectiveness on texture classification tasks and reported spatiotemporal features to be beneficial. A miniaturized sensor version was presented in [48]. In similar efforts, [49]–[51] used an event-based camera behind a piece of silicone inside a parallel gripper, and investigated its effectiveness on the task of force reconstruction and material classification for robotic sorting [52]. Compared to this paper, these works investigated different tasks. In particular, they do not cover exploiting the sensors for high-frequency closed-loop feedback control for robotic manipulation.

More closely related are [20], [21]. They both investigate the task of slip detection from event-based optical tactile readings, however, without considering closed-loop robotic manipulation. The authors of [21] presented a marker-based tactile sensor that is read out using a very high temporal resolution of 0.5 ms, and presented a proof of concept for slip detection using a hand-defined threshold. Their proof of concept only included a single object, compared to 18 objects investigated herein. Rigi et al. [20] placed an event-based camera behind a transparent silicone. They integrated the events for 10 ms, thereby operating at a 10 times reduced temporal resolution compared to this paper. They also employed a hand-defined threshold on the change in contact area between sensor and object to detect the onset of slip. They evaluated their approach on five different objects, which are all dark in color, as this benefits their approach.

Closest to this paper is the work of Muthusamy et al. [53]. Their tactile sensor consists of an event-based camera placed behind transparent plexiglass, creating an event-based version of the Fingervision sensor [54]. For detecting slippage, they investigated two approaches: one based on the raw spike count, and the other one based on the count of edge and corner features extracted from integrated images. They also combined their model-based slip detectors with control. Contrary to our proposed Evetac, their sensor is transparent and comes without any imprinted markers. Therefore, their model-based slip detectors rely on the manipulated objects having sufficient texture, making them less general. Additionally, the gel’s transparency results in their sensor capturing not only contact-related phenomena but also events triggered, for instance, by moving background. This might be disadvantageous as the background events are essentially noise when considering the task of slip detection.

From a methodological point of view, all previous approaches for slip detection using event-based optical tactile sensors leveraged hand-designed, model-based criteria. Herein, we take a different approach and use a model-free approach for learning slip detection and prediction models solely relying on labeled experimental data. Our approach is thus not focused on one specific, pre-defined criterion for detecting slip. Instead, during training, the neural network models are refined to automatically extract the most important information from the input features. For automatic data labeling without compromising temporal resolution, we develop a new criterion based on optical flow. Regarding input features, this work proposes a novel method for tracking the markers imprinted in the gel. The marker tracking is capable of providing information about the gel’s global deformation, complementing the raw sensor’s sparse and local measurements. We showcase that this information benefits slip detection.

In terms of sensor properties, Evetac offers a significantly higher spatial resolution of 640x480 pixels, compared to a maximum resolution of 240x180 pixels offered by previous event-based optical tactile sensors. The improved spatial resolution closes the gap to standard RGB optical tactile sensing, as it is similar to, for instance, the DIGIT sensor [25]. Additionally, we open-source the entire sensor design, and all other components, such as connectors, used during the experiments. All the custom components can be recreated using an off-the-shelf 3D printer. Exemplarily, Evetac can be seamlessly integrated into a commercially available robotic gripper as shown in Figure 1. In contrast, for the other event-based optical tactile sensors, the design files are not openly available. Moreover, Evetac’s design is built around the idea of modularity and re-use of existing components. Evetac uses a commercially available event-based camera, and, contrary to prior works, also exploits a commercially available soft silicone elastomer from GelSight Inc. Lastly, to the best of our knowledge, this is the first work quantitatively comparing an event-based optical tactile sensor with an established RGB optical tactile sensor in terms of data rate and shear force reconstruction.

### III. BACKGROUND - EVENT-BASED CAMERAS

Event-based cameras work fundamentally differently than standard cameras. While in standard, frame-based cameras, every pixel is read out at a constant frequency, in event-based cameras, every pixel is independent and only reacts to brightness changes at its location, making the sensor asynchronous [26]. Given the brightness of a pixel at position  $x, y$ , i.e., its log photocurrent  $L(x, y, t) = \log(I(x, y, t))$  at time  $t$ , the pixel is sensitive to changes in intensity  $\Delta L(x, y, t) = L(x, y, t) - L(x, y, t_k)$  w.r.t. reference value  $L(x, y, t_k)$ . If this change in intensity reaches either the positive or negative threshold ( $\pm C$ ) at time  $t_{k+1}$ , this pixel triggers a new event  $e_{k+1} = (x, y, t_{k+1}, p_{k+1})$ . Every event is a tuple and contains all the information about the event’s location  $(x, y)$ , timing  $(t_{k+1})$ , and polarity  $p_{k+1} \in \{-1, 1\}$ , which signals whether the brightness increased or decreased. Subsequently, the pixel’s reference brightness value is adapted




Fig. 2: Exploded view of the proposed Evetac sensor. From left to right, it consists of the following components: A) DVXplorer Mini, event-based camera, B) 3D printed camera housing, C) LED stripe to illuminate the sensor from the inside, and D) GelSight Mini dotted gel. The 3D printed housing is designed such that the camera's distance from the gel can be adjusted to ensure that the camera is in focus. Moreover, it allows mounting Evetac to an external gripper as visualized in Fig. 1. For an overview of the components also see Table I. The total dimensions of the assembled sensor are 32x33x65mm (width x height x length).

TABLE I: Evetac Hardware Components.

Component	Specifications
Camera	DVXplorer Mini from Inivation, 640 x 480px resolution.
Housing	Custom Design, 3D printed.
Camera Screws	1/4", thread length 6.3mm e.g., available at [55].
LED stripe	LED COB band 4000K, height 5mm, width 2.2mm, 12V, e.g., available at [56].
Gel	GelSight Mini Marker Gel [57].

to  $L(x, y, t_{k+1}) = L(x, y, t_k) + p_{k+1}C$ , and from then on, the pixel is sensitive to changes w.r.t. the updated value. Overall, the camera's output is this stream of events.

This asynchronous functioning principle offers many appealing properties, such as reduced power consumption, high dynamic range, lower latencies, and high temporal resolution [26], which we, herein, aim to exploit in the context of touch sensing, processing, and robotic manipulation.

#### IV. THE EVETAC SENSOR

We now introduce our proposed Evetac sensor, a new event-based optical tactile sensor. The sensor mainly consists of off-the-shelf components and a 3D printed case aiming to reduce the entry barrier into the field. This section focuses on the sensor itself by introducing its mechanical design and hardware components, followed by a description of Evetac's raw sensory output.

##### A. Hardware

Figure 2 depicts the main components of our proposed Evetac sensor. It consists of an event-based camera capturing the deformation of a soft silicone gel. The sensor is held together by a 3D printed housing and the illumination is provided by a white LED stripe surrounding the gel. The housing is designed such that the distance between camera and gel can be adapted through two 1/4" camera screws. This allows to position the camera such that the gel is in focus. The housing further offers a mounting mechanism on the bottom, allowing for sensor integration into parallel grippers or a single-finger configuration as shown in Figure 1 & Figure 5b. We also designed a similar casing for the standard GelSight Mini such that we can easily create a setup with one standard GelSight

Mini and one Evetac inside a parallel gripper (cf. Figure 7). All files for 3D printing are open-sourced on our website. The two main design goals of Evetac were to ensure modularity and maximize reuse of existing components. Apart from the 3D printed housing, all the components are commercially available. The hardware components are also summarized in Table I. The event-based camera is a DVXplorer Mini from Inivation. The soft silicone gel is the same dotted gel that is also used in the standard GelSight Mini. The choice of this non-transparent, dotted gel (cf. Figure 9) ensures that the camera cannot see through the gel. It thus naturally focuses on capturing contact-related phenomena between sensor and object with minimum distraction. This also implies that the sensor is agnostic w.r.t. the manipulated object being textured or not. Moreover, the black dots that are imprinted in the gel can provide important information about the gel's current global deformation. The dots provide natural features that can be captured by the camera and used in tasks such as shear force estimation or slip detection [38], [39], [41], as we will later also show for Evetac in Sec. VI & Sec. VII.

Besides the modular design, the event-based camera is the key component that differentiates Evetac from standard, classical optical tactile sensors such as GelSight [58], DIGIT [25], or TacTip [24]. The asynchronous functioning principle has many desirable properties. In this work, we particularly want to build upon the camera's sparse output, high temporal resolution, and low latency for high-frequency, real-time touch sensing, processing, and feedback control.

##### B. Raw Sensor Output

As mentioned in Sec. III, on the lowest level, event-based cameras return single events characterized by their location  $x, y$ , timing  $t$ , as well as polarity  $p$ , i.e.,  $e = (x, y, t, p)$ . Yet, when reading out event-based cameras with standard computers, the events are typically accumulated on the camera before they are sent via USB to the computer for further processing. Herein, we configure the event-based camera such that the events are accumulated for 1 ms before they are sent to the computer. Therefore, every millisecond we receive the set of  $N_E$  events  $\mathcal{S}_E(t_i) = \{e_k, k \in N_E\}$ , that have been created within the previous millisecond, i.e.  $\forall e_k = (x_k, y_k, t_k, p_k) \in \mathcal{S}_E(t_i) : t_i - 1 \text{ ms} = t_{i-1} \leq t_k < t_i$ . Evetac's current measurements, i.e., the received set of events, can also be visualized in image form, e.g., as shown in Figure 3a. For the image visualization, all pixel locations where no events occurred are colored in gray. The white pixels correspond to locations where on-events have been triggered (i.e., the pixel intensity increased), and the pixels in black correspond to locations of off-events. As the sensor's output is very sparse, for visualization purposes, we actually show the combination of the last five measurements in all figures throughout the paper. Thus, the images depict the events triggered within the past 5 ms. The images have the same resolution as provided by the camera, which is 640x480 pixels. We want to point out that the information received from Evetac and its event-based camera is highly time-dependent. When there is no change in the gel's state, no events are triggered, and we do not receive any information. If the gel is






Fig. 3: Both pictures show the current contact configuration (left) & Evetac’s output in image form (right). As mentioned in Sec. IV-B, Evetac returns the raw events accumulated for 1 ms. Since Evetac’s raw output is extremely sparse, for visualization purposes, we actually illustrate the combination of the last 5 measurements, i.e., the events triggered within the last 5 ms. In the pictures that show Evetac’s raw output, all pixels in gray correspond to locations where no events have been triggered, while the white pixels illustrate locations of on-events, and the black pixels are the locations of off-events.

moved (e.g., as shown in Figure 3a & Figure 3b), many events are triggered and returned to the computer. While it would be possible to configure Evetac’s event-based camera to return event sets even more frequently, we believe that our choice of grouping the events for time intervals of 1 ms, thereby receiving measurements at 1000 Hz, provides a good tradeoff between temporal resolution and computational feasibility of additional signal processing algorithms. In terms of software implementation, we build on top of the dv-processing library provided by Inivation [59]. We extend their code with our touch processing methods presented in the next section and also implement a ROS interface, facilitating data recording and visualization.

## V. EVETAC TOUCH PROCESSING

In its standard configuration, Evetac returns event sets at 1000 Hz (cf. Sec. IV-B). While it would be possible to directly exploit Evetac’s raw sensory output for solving downstream tasks such as slip detection, in this section, we will introduce additional touch processing algorithms operating on top of Evetac’s output. These algorithms and their output aim to provide meaningful intermediate representations. In particular, they also further compress the dimensionality of Evetac’s raw output, which benefits meeting the computational requirements and ultimately achieving real-time feedback control. Moreover, the dot tracking algorithm, which will be introduced next, is designed for keeping track of Evetac’s global configuration, i.e., the configuration of the gel, despite the sensor’s sparse measurements triggered by local changes.

### A. Signal Processing - Dot Tracking

Evetac’s raw sensory output solely contains local, relative information. This is fundamentally different from the raw output of RGB optical tactile sensors, which, at every sensor measurement, return a complete image of the gel. While the former might be more efficient from the perspective of only retrieving information when changes occur, it also comes at the disadvantage of not being able to reconstruct the gel’s




Fig. 4: Illustrating how the movement of a black dot in front of a white, bright background triggers events. In the leftmost picture, at time  $t_i$ , the dot is moving to the right, resulting in a slightly shifted position at time  $t_{i+1}$  shown in the middle frame. This slight dot movement between  $t_i$  and  $t_{i+1}$  causes events  $S_E(t_{i+1})$  that are visualized in pictorial form in the rightmost frame. The pixels colored in white in the right frame correspond to the locations where on-events were triggered. The brightness of all these pixels changed from being occupied by the black dot at  $t_i$  to being occupied by the white background at  $t_{i+1}$ . Due to this intensity change, events have been triggered at these locations. For the pixels colored in black, the opposite happened, i.e., the brightness changed from the white background to the black dot. At all the remaining pixel locations, no events have been triggered. They are thus colored in neutral gray.

global configuration given a single measurement. When no changes are happening across the gel surface, Evetac’s measurements will not contain any events, and, therefore, provide no information. As mentioned earlier, Evetac uses a gel with imprinted markers, i.e., dots (cf. Figure 9), to visualize gel movements. Their positioning provides information about the current gel deformation. Since this global information might be important for some tasks, such as shear force reconstruction or slip detection, in the following, we provide an algorithm for tracking the dots’ positions.

Our proposed dot tracking algorithm builds upon the work of [60], who presented a model-based tracker for event-based cameras. The algorithm is based on the assumption that all triggered events are caused by the movement of the object that is to be tracked. Since event-based cameras register changes in lighting intensity, it is the edges of uniformly colored objects that trigger events. The two leftmost frames of Figure 4 illustrate the scenario of a black dot (i.e., the object) moving to the right between times  $t_i$  and  $t_{i+1}$ . As illustrated in the rightmost frame of Figure 4, in the direction of movement, pixels change from the white, bright background to the black color of the moving dot. This change, i.e., decrease, in brightness triggers negative off events. On the other side, i.e., opposite to the dot’s moving direction, pixels change from the dot’s black color to the bright background, yielding positive on events. Exploiting this insight that an object’s edges trigger events upon movement, the model-based tracking algorithm consists of two main steps. First, finding correspondences between the triggered events and model edge points. Second, updating the model’s pose estimate given the established correspondences.

Without loss of generality, in the following, we present the algorithm assuming a single event  $e_k(x_k, y_k, t_k, p_k) = (x_k, t_k, p_k)$  that has already been assigned to one of the gel’s dots. In the first step of correspondence matching, we find the closest model edge point  $x_j^m \in \mathbb{R}^2$  that could have caused this event, given the model’s current pose described by rotation  $R(\theta_i) \in SO(2)$  and translation vector  $c_i = [c_{x_i}, c_{y_i}]^T \in \mathbb{R}^2$ . We create an assignment between event and model point  $k \rightarrow j$  minimizing  $d(k, j) = \|x_k, x_j^m\|_2$ . Given the correspondence, we update the model’s pose

through rotation matrix  $R(\theta) \in SO(2)$  and translation vector  $\mathbf{c} = [c_x, c_y]^T$ , attempting to correct the model pose such that it explains the observation. This is done by minimizing objective function  $f = \|\mathbf{x}_k - (R(\theta)\mathbf{x}_{j \rightarrow k}^m + \mathbf{c})\|_2^2$ . While it would be possible to find the unique optimal solution to minimize  $f$ , the authors of [60] proposed to apply a gradient-based update, i.e.,  $\theta_{i+1} = \theta_i - \alpha \nabla_{\theta} f|_{\theta=0, \mathbf{c}=[0,0]^T}$ , and  $\mathbf{c}_{i+1} = \mathbf{c}_i - \alpha \nabla_{\mathbf{c}} f|_{\theta=0, \mathbf{c}=[0,0]^T}$ . One motivation for the gradient-based update is that the event data is inherently sparse. Also, in case of multiple events, the gradients can simply be accumulated. Therefore, the gradient magnitude changes with the number of events. Exemplary, in scenarios with only a few events, the gradient-based update only slightly adjusts the object pose, instead of greedily converging to the best pose. This is advantageous since not all events are triggered by object movements. Especially in low event scenarios it might happen that the majority of events come from sensor noise.

For our particular case of dot tracking, we can simplify the update rule and correspondence matching. First, due to the rotational symmetry of the model, we can omit optimizing its orientation. Second, while [60] define model edges through a discrete set of points, for tracking dots, we can analytically calculate the closest model point  $\mathbf{x}_{j \rightarrow k}^m$  for every event at location  $\mathbf{x}_k$ . Using geometry, the closest model point has to lie at dot radius  $r$  away from the dot's center point, in the direction of the event, i.e.,  $\mathbf{x}_{j \rightarrow k}^m = r\mathbf{x}_k / \|\mathbf{x}_k\|_2$ . Note that this assumes that the event's coordinate  $\mathbf{x}_k$  is already given with respect to the dot's current center location  $\mathbf{c}_i$ . Subsequently, the objective for adapting the dot's location equates to

$$f = \|\mathbf{x}_k - (\mathbf{x}_{j \rightarrow k}^m + \mathbf{c})\|_2^2 = \|\mathbf{x}_k - (r\mathbf{x}_k / \|\mathbf{x}_k\|_2 + \mathbf{c})\|_2^2. \quad (1)$$

The gradient can be obtained in closed form as

$$\nabla_{\mathbf{c}} f|_{\mathbf{c}=[0,0]^T} = -2(\mathbf{x}_k - r\mathbf{x}_k / \|\mathbf{x}_k\|_2). \quad (2)$$

In conclusion, we obtain a closed-form solution for updating the dot's center coordinate, given an event at location  $\mathbf{x}_k$ . If more than one event is associated to a certain dot, i.e., when having to deal with a set of events  $\mathcal{S}_E(t_{i+1}, \mathbf{c}_i)$ , the resulting update is the sum of the updates for every individual event

$$\nabla_{\mathbf{c}} f(\mathcal{S}_E(t_{i+1}, \mathbf{c}_i))|_{\mathbf{c}=[0,0]^T} = \sum_{l=1}^{|\mathcal{S}_E|} -2(\mathbf{x}_l - r\mathbf{x}_l / \|\mathbf{x}_l\|_2). \quad (3)$$

Initial testing of this gradient-based tracking scheme showed that it worked well in conditions where the dots' movement caused the majority of the events. Yet, the algorithm is prone to losing track when an external object moves really fast over the gel. In such situations, the stream of events caused by the edges of the moving object can outweigh the events triggered by the dots' movement, as shown in Figure 3b. This might result in our tracking algorithm following the stream of events caused by the external object, thereby losing track. This issue is especially problematic when considering that Evetac itself only returns sparse, relative measurements. It is, therefore, very difficult to recover the dots' poses once track has been lost since single Evetac measurements are not sufficient. To

counteract the issue and attempt to prevent losing track at first place, we propose the following addition to the algorithm. The addition is a regularizing term, emerging from the idea that despite the gel deforming, the dots should still roughly stay within their original grid-like structure. In other words, the dots' positions relative to each other should not change too drastically. Without loss of generality, we now consider a pair of neighboring dots at locations  $\mathbf{c}_i^1$  and  $\mathbf{c}_i^2$  which have been at an initial distance  $d_{1,2} = \|\mathbf{c}_0^1 - \mathbf{c}_0^2\|_2^2$  at the start of the tracking for which we assume that nothing is pressing against Evetac. Exemplarily, if we want to regularize the position update  $\mathbf{c}$  of dot 1 based on its initial distance to dot 2, the updated objective with weighting factor  $w_{\text{dist}}$  equates to

$$\begin{aligned} f_{\text{reg}} &= f + w_{\text{dist}} f_{\text{dist}} \\ &= f + w_{\text{dist}} (\|(\mathbf{c}_i^1 + \mathbf{c}) - \mathbf{c}_i^2\|_2^2 - d_{1,2})^2. \end{aligned} \quad (4)$$

The additional term basically penalizes moving too far or too close to the other dot, taking the initial distance as reference. Calculating the gradient for the regularizing term results in

$$\nabla_{\mathbf{c}} f_{\text{dist}}|_{\mathbf{c}=[0,0]^T} = 4(\mathbf{c}_i^1 - \mathbf{c}_i^2)(\|\mathbf{c}_i^1 - \mathbf{c}_i^2\|_2^2 - d_{1,2}). \quad (5)$$

While this example only considered a pair of dots, we add the same regularization term for all the direct neighbors of a dot. Herein, a dot can have at maximum 8 neighbors (i.e., along the horizontal, vertical, and the two diagonals), and at minimum three for the dots in the corner. Due to these differences in the number of neighbors,  $w_{\text{dist}}$  is scaled by 8 divided by the dot's actual number of neighbors such that the regularizing term is of similar magnitude for all dots. Lastly, we want to point out that a dot's location is only updated if events are triggered at its location. Without this additional constraint, the regularizing term might cause a dot to move without any events being present at its location.

## B. Evetac Touch Features

Before proceeding with the next sections, in which we will use Evetac to infer contact-related phenomena, we provide a short overview of the features that we consider in this work. While the sensor's raw measurement (cf. Sec. IV-B) provides a set of events  $\mathcal{S}_E(t_i)$  at current time  $t_i$  at a spatial resolution of 640 by 480 pixels, in the next sections, we will mainly exploit lower dimensional features. Building on top of these lower dimensional features facilitates meeting Evetac's real-time requirements, as any further processing also has to be able to handle the data at a frequency of 1000 Hz. In the remainder of the paper, we will consider the following features:

- Overall, raw number of events,  $N_E(t_i) = |\mathcal{S}_E(t_i)|$ .
- Raw number of events per dot, i.e., for dot at location  $\mathbf{c}_i$ ,  $N_E(t_i, \mathbf{c}_i) = |\mathcal{S}_E(t_i, \mathbf{c}_i)|$ . This quantity relies on the dot tracking algorithm estimating the dot's center location  $\mathbf{c}_i$ . An event is associated with a dot if the event's location is less than 20 pixels away from the dot's center.
- Number of events for all the  $N_D$  dots  $E_C(t_i) = \{N_E(t_i, \mathbf{c}_i^l), l \in N_D\}$
- Positions of all of the  $N_D$  dots  $P_C(t_i) = \{\mathbf{c}_i^l, l \in N_D\}$
- Displacement of the dot at current location  $\mathbf{c}_i$  w.r.t. its initial location  $\mathbf{c}_0$ , i.e.,  $d_{c_i} = \|\mathbf{c}_i - \mathbf{c}_0\|_2$




Fig. 5: Experimental setups for the Evetac benchmarking experiments of sensing vibrations (cf. Sec. VI-A, **(a)**) and shear force reconstruction (cf. Sec. VI-D, **(b)**). For both experiments, we have Evetac mounted in the end-effector of a Franka Panda robot. In Fig. 5a, Evetac presses against a speaker which is set to generate a tone with a desired frequency. Through making contact with the speaker, Evetac can perceive the vibrations of the speaker and reconstruct the vibration frequency as shown in Fig. 6. In Fig. 5b, Evetac presses against an object mounted on top of a F/T sensor. Through moving the robot, we shear the gel. By combining our proposed dot tracking algorithm with a model, we attempt to recover the shear forces acting upon the gel as shown in Fig. 8.

- Displacements of all of the dots  $D_C(t_i) = \{d_{c_l}^l, l \in N_D\}$

Our reasoning for this choice of features is that the raw number of events might be beneficial for resolving very fast phenomena, while the displacement features will provide more information about the gel’s current deformation. We nevertheless want to point out that this list only covers a very small portion of possible features that could be extracted from Evetac’s raw output. It might well be that there exist more powerful and informative features that would actually benefit and improve Evetac’s performance on downstream tasks.

## VI. BENCHMARKING EVETAC

In this section, we present four experiments to showcase Evetac’s properties. First, we use Evetac to sense vibrations up to 498 Hz and validate its high temporal resolution. Second, we compare the data rate of Evetac and a RGB optical tactile sensor in a grasping and slipping experiment, highlighting the advantages of Evetac’s sparse output. Third, we investigate the effectiveness of the regularizing term for dot tracking. For this purpose, we compare the unregularized and regularized version of the dot tracker w.r.t. losing track of the dots when interacting with the gel using our fingertips and tools. Lastly, we evaluate the quality of the proposed regularized dot tracking method. We do so by considering the task of shear force reconstruction, solely using the dot’s positions as input. We again also compare against an RGB optical tactile sensor. Additional videos are available on our website.

### A. Sensing Vibrations

This experiment aims to verify Evetac’s high temporal resolution of receiving sensor feedback at 1000 Hz by




Fig. 6: Frequency spectrum of the vibration sensing experiment. The plot shows the normalized frequency spectrum (i.e., amplitudes normalized to  $[0,1]$ ) measured by Evetac when pressing against a speaker which is set to different frequencies  $f_d$  (cf. Fig. 5a). The spectrum is obtained by performing a Fourier transform on the total number of events measured per millisecond for a duration of  $T_w=10$  s. Evetac recovers the main vibration frequency of the speaker ( $f_d$ ) reliably up to frequencies of 498 Hz as the peaks in the measured frequency spectrum align with  $f_d$ . This confirms the sensor’s high temporal resolution and its readout frequency of 1000 Hz.

TABLE II: Numerical results for the vibration sensing experiment (cf. Fig. 5a & Fig. 6). The table reports the success percentage and successful detections of the speaker’s frequency  $f_d$  considering different time windows  $T_w$ . The results underline Evetac’s high temporal resolution. Vibrations up to 498 Hz can be detected reliably.

Speaker Frequency	Detected Frequency		
	$T_w = 10$ s	$T_w = 1$ s	$T_w = 1$ s
50 Hz	100% (10/10)	100% (50/50)	99% (99/100)
100 Hz	100% (10/10)	98% (49/50)	96% (96/100)
150 Hz	100% (10/10)	100% (50/50)	100% (100/100)
200 Hz	100% (10/10)	92% (46/50)	69% (69/100)
250 Hz	100% (10/10)	100% (50/50)	100% (100/100)
300 Hz	100% (10/10)	100% (50/50)	99% (99/100)
350 Hz	100% (10/10)	100% (50/50)	100% (10/10)
400 Hz	100% (10/10)	100% (50/50)	98% (98/100)
450 Hz	100% (10/10)	100% (50/50)	100% (100/100)
498 Hz	100% (10/10)	100% (50/50)	100% (100/100)
Avg.	100% (100/100)	99% (495/500)	96% (961/1000)

measuring vibrations generated by a speaker. According to Nyquist–Shannon sampling theorem [61], given sampling frequency  $f_s$ , perfect signal reconstruction is possible for bandwidth  $B < f_s/2$ , i.e.,  $B < 500$  Hz for our case. The experimental setup is shown in Figure 5a. We attach Evetac to a 3D-printed mount that can be screwed to the mounting flange of the Franka Panda 7DoF robot. The robot presses Evetac against a commercially available Bluetooth speaker (Anker Soundcore mini). The tone frequency played by the speaker  $f_d$  is controlled through the mobile phone app “Tone Generator” which generates pure sine wave tones at the desired frequency. The speaker is placed upside down and set to maximum volume. Evetac presses against the bottom of the speaker. We ensure that Evetac directly touches the speaker’s metal housing and not the rubbery ring on the edge, as the rubber will dampen the vibrations. On the other side, the speaker’s housing is in direct contact with a piece of wood. This experimental setup ensures that there is no dampening material between the metal housing of the speaker and the piece of wood / Evetac. Therefore, the speaker will vibrate when generating the sound.

Since Evetac is touching the speaker, the vibrations will be transmitted to the gel and make the black dots within the gel move. Evetac's event-based camera will capture this movement of the dots, and events will be generated proportional to the velocity of the dots or changes in their size. We, therefore, expect the number of events to oscillate with the same frequencies that are triggered by the speaker's vibrations. The most excited frequency component should coincide with the tone frequency  $f_d$  that the speaker is tasked to play. For recovering  $f_d$  from the sensor readings of Evetac, we propose to apply a Fourier transform to the absolute number of events ( $N_E(t_i)$ ) over a time series of length  $T_w$ . Given the spectrum in the frequency domain, we remove all frequency components smaller than 25 Hz, as we are only interested in sensing higher frequency vibrations. Subsequently, we select the frequency component that exhibits maximum amplitude. If this frequency is within  $\pm 1$  Hz of the tone frequency that we set the speaker to, we label this as a correct detection of the vibration.

To evaluate the performance of this procedure, we set the speaker to play tones with 10 different frequencies ranging from 50 Hz to 498 Hz. For each frequency, we create a recording for 100 s. The result for extracting and analyzing one measurement of length  $T_w=10$  s for all frequencies is illustrated in Figure 6. As can be seen, the peaks in the measured frequency spectrum coincide with the frequency we set the speaker to ( $f_d$ ). For a more thorough analysis, we divide the entire recording into pieces of length  $T_w$  and apply our procedure to recover  $f_d$  to all of the sequences individually. The results are presented in Table II. Overall, the sensor and processing pipeline are capable of sensing the speaker's tone through the vibrations with high success rates. Considering longer time series, i.e.,  $T_w=10$  s, it is possible to perfectly recover the vibration frequency. For  $T_w=2$  s, i.e., dividing the trajectory into 50 segments of length 2 seconds each, only in 5 out of 500 segments, the vibration frequency of the speaker ( $f_d$ ) was only the second most excited frequency component, while half of the tone frequency ( $f_d/2$ ) was the most excited one. Considering shorter time windows such as  $T_w=1$  s, the performance of identifying  $f_d$  through the most excited frequency component again slightly decreases. In the shorter time windows, it happens more frequently that either half or double of the tone frequency  $f_d$  are the most excited components. Nevertheless, the speaker's vibration frequency is still always amongst the three most excited components. Overall, we conclude that Evetac is able to accurately sense high-frequency vibrations. We thereby also validate Evetac's high temporal resolution. The highest frequency that was able to be recovered was 498 Hz, which is very close to the Nyquist frequency of our setup, i.e., 500 Hz.

### B. Data Rate Experiments

Next, we compare the sensor data rate of our proposed Evetac sensor and a standard RGB optical tactile sensor (GelSight Mini [57]), considering the sensors' raw outputs. For tactile sensing applications, sensor data rate is crucial as we ideally want to deploy tactile sensors all around a robot's surface. Less sensor data rate thus also correlates with less

information that needs to be processed. As experimental setup, we consider a parallel gripper, where one of the fingers is equipped with Evetac, and the other one with a GelSight Mini (cf. Figure 7). As shown in the small pictures of Figure 7 and in the supplementary videos, for comparing the data rates, we execute the following maneuver. We start with the gripper fully open and close it to make contact with the object. Next, we establish a predefined grasping force, resulting in a stable grasp. After stable grasping, we mimic perturbations that might occur during manipulation by pressing onto the object. Since the sensors' gel is elastic, this will make the object oscillate within the fingers, however, without breaking contact. Finally, we open the gripper, which will cause the object to slip.

For the data rate comparison, we make the following assumption. Since we used a modified version of the gels, which will later be introduced in more detail (cf. Sec. VII), we only consider a sensing area of 540 by 480 pixels for both of the sensors. Yet, we want to point out that the reduced sensing area has no effect on the results. The RGB optical tactile sensor, i.e., GelSight Mini, returns 3 color values per pixel. Each of them is in the range of 0-255 and can be represented by 1 byte. Thus, one output from the RGB optical tactile sensor has  $540 \times 480 \times 3 = 777600$  bytes. For Evetac, we receive a set of events with varying size every millisecond (cf. Sec. IV-B). Every event can be represented by 5 bytes. In particular, we require 2 bytes to encode the event's x coordinate and 2 bytes for the y coordinate, together with 1 byte signaling its polarity.

Figure 7 illustrates the two sensors' outputs in bytes for one of the exemplary trajectories. As can be seen, while the synchronous RGB optical tactile sensor always returns a fixed-size output, Evetac is more selective. Evetac's output exhibits clear spikes whenever something happens at the contact location, as its output size correlates with the number of events. At all other times, almost no sensor output is generated. Numerical results are presented in Table III. The table reports the ratio between data rate (i.e., bytes/s) used by Evetac w.r.t. data rate generated by the RGB optical tactile sensor for different parts of the trajectory. Although Evetac is read out at a 40 times higher frequency (1000 Hz vs 25 Hz), it only produces 1.7% of the RGB optical tactile sensor's data rate considering the entire trajectories. Even when only considering a 0.5 s interval around the moment of slippage, which is the point in time when Evetac returns most information, it still only generates around 11.9% of the RGB optical tactile sensor's data.

This experiment underlines the efficiency of Evetac and the underlying event-based camera. It only returns information when something at the contact is changing and thus produces considerably less data compared to a standard, RGB optical tactile sensor that always returns fixed-size measurements irrespective of changes in the image and contact configuration.

### C. Effectiveness of Regularization in Dot Tracking

While the previous two experiments mainly considered the sensors' raw outputs, in this experiment, we investigate the effectiveness of the regularization in the dot tracking algorithm (cf. Sec. V-A). As mentioned previously, the dot



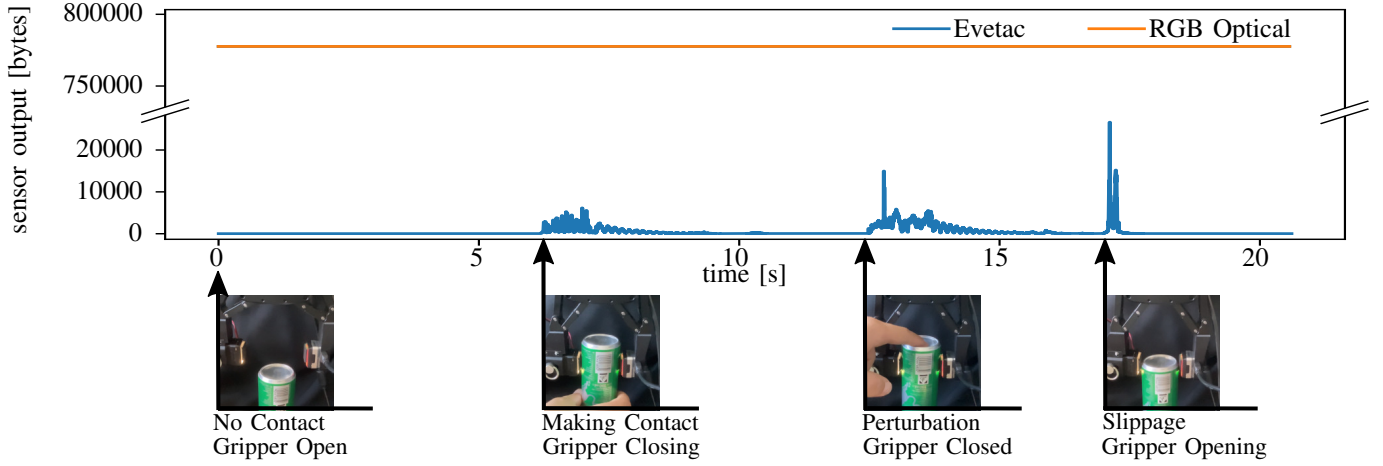


Fig. 7: Data rate comparison between Evetac and a RGB optical tactile sensor (GelSight Mini), considering a manipulation maneuver of grasping an object, perturbing the object, and finally forcing object slippage. As can be seen, Evetac’s output size correlates with changes in contact configuration, and, per measurement, is of significantly smaller size. For numerical results on sensor data rate, see Table III.

TABLE III: Data rate comparison between Evetac and a RGB optical tactile sensor averaged over 5 trajectories (reporting mean & standard deviation). See Fig. 7 for one of them. The relative data rate of Evetac w.r.t. the RGB optical tactile sensor is reported. Since Evetac’s output is correlated with changes in contact configuration, different time intervals are considered. The ratio is reported for 1) the entire trajectories, 2) the timespan between making contact and object slippage, and 3) a 0.5s interval around the moment of slippage. Despite Evetac running at a significantly higher frequency (1000 vs 25 Hz), regarding the overall trajectory, it only consumes 0.017 of the data rate of the RGB optical sensor.

data rate ratio	entire trajectory	making - breaking contact	slip only
Evetac / RGB Optical Tactile Sensor	<b>0.017</b> (0.004)	0.028 (0.008)	0.119 (0.006)

TABLE IV: Evaluating the effectiveness of the regularized dot tracking. We report on how many out of 10 trajectories the regularized and unregularized version of the dot tracking algorithm successfully completed the dot tracking, i.e., did not lose track of any of the dots. We also report the average number of dots for which track was lost in the unsuccessful trajectories ( $N_{lt}$ ).

Tracking Algorithm Version	Success Rate	$N_{lt}$
Unregularized	20% (2/10)	9.5
Regularized	<b>80%</b> (8/10)	<b>1.5</b>

tracking algorithm assumes that dot movements trigger all events. While this holds true for some tactile interactions as shown in Figure 3a, in other situations, the edges of an object sliding over the gel might trigger many events that might even outweigh the events triggered by the dots (cf. Figure 3b). These additional events, which are not triggered by the dot movement, contradict the assumptions of the dot tracking and might cause it to lose track of the dots. To counteract the dot predictions following the events triggered by the moving object, we introduced a regularization term, regularizing the dots’ positions relative to each other based on their initial distance.

In this experiment, we compare the two versions of the dot tracking algorithm, i.e., the unregularized and the regularized one. We consider 10 trajectories in which we either make contact with Evetac using a finger or the handle of a scissor. Snapshots from two trajectories are shown in Figure 3a and Figure 3b. After making contact, we performed several

movements to shear the gel. Most importantly, we also ensure that there is slippage, i.e., relative motion between the sensor and object, during which also the object’s edges and shape will trigger events. All the trajectories start and end with nothing touching the gel. Thus, the dots should be at their equilibrium position at the start and end. For comparing the two dot tracking versions, we report the percentage of trajectories for which the dots end up close to their initial position, i.e., for which the tracking was successful throughout the entire trajectory. For this, all dots have to end up within a radius of 20 pixels from their initial position. Additionally, for the unsuccessful trajectories, we report the mean number of dots for which track was lost ( $N_{lt}$ ).

Table IV shows the results and underlines the effectiveness of the regularized dot tracking. The regularizer helps in increasing the number of successful trajectories by a factor of 4, from two to eight. Moreover, even in the unsuccessful cases, the number of dots for which track was lost is substantially decreased. We also provide supplementary videos comparing the two trackers across the trajectories. They illustrate that the regularized tracker does not compromise tracking quality. We nevertheless want to point out that the trajectories in this experiment rather cover extreme cases. As we will show later, when an object slips, normally, most of the events are still triggered by the dots. We still think that a robust dot tracker is important as Evetac and the touch processing algorithms should be capable of dealing with a wide range of scenarios. Thus, throughout all of the following experiments, we use the regularized version of the tracker.

#### D. Dot-based Shear Force Reconstruction

Due to their sensing principle, event-based cameras naturally only return relative, sparse information. While this property is beneficial with respect to sensor data rate, it might be a limitation in scenarios where knowledge about the gel’s global deformation is important. In this section, we, therefore, evaluate the previously proposed regularized dot tracking algorithm (cf. Sec. V-A) through the task of shear force reconstruction solely using the dots’ displacements




Fig. 8: Shear force reconstruction from tracking the dots of Evetac (top) and a RGB optical tactile sensor (bottom). For the experimental setup, see Fig. 5b. GT corresponds to ground truth measurements of a F/T sensor, LR to the linear regression model, and NN to the neural network model. Qualitatively, the shear force reconstruction leads to similar quality despite the different sensor types. Since the reconstruction is solely based on the dots’ displacement, we conclude that our presented regularized dot tracking tracks the dots with good accuracies, despite the sparse and relative output of Evetac. Numerical results are presented in Table V.

TABLE V: Numerical Results for Shear Force Reconstruction experiment (cf. Fig. 8).

Method	Sensor	MAE [N]	
		$F_x$	$F_y$
Linear Regression	Evetac	0.53 (0.41)	0.41 (0.29)
	RGB Optical	0.53 (0.49)	0.44 (0.29)
Neural Network	Evetac	0.33 (0.25)	0.22 (0.18)
	RGB Optical	0.35 (0.30)	0.37 (0.26)

$D_C(t_i)$  as input. We will compare against the quality of shear force estimation based on tracking the dots using a standard RGB optical tactile sensor (GelSight Mini).

For conducting the experiment, we mount the two tactile sensors on the end effector of a 7DoF Franka Panda robot. Figure 5b shows the setup for Evetac. We designed a similar mount allowing to attach GelSight Mini to the robot in the same pose. As shown, the robot presses the respective sensor against a flat, 3D-printed object, which in turn is mounted on top of a *SCHUNK FTCL-50-40* force torque sensor (F/T sensor). Upon the robot establishing contact between tactile sensor and the 3D printed object, it executes a trajectory of predefined waypoints that result in shearing the gel. The trajectories cover shear forces in the range of  $\pm 10$  N. The tactile sensors are aligned such that the direction in which they have 9 dots, i.e., the longer side, aligns with the x-axis of the F/T sensor. During manipulation, we track the position of the dots and record the readings of the F/T sensor. For Evetac, we use the previously proposed regularized dot tracker Sec. V-A. For the RGB optical tactile sensor, we use the optical-flow-based tracker from the GelSight repo [62]. We end up with one training dataset per sensor, which consists of 4 trajectories, and one separate, previously unseen trajectory for evaluation. The trajectories differ in that the sequence of

the waypoints is randomized. Additionally, we randomized the relative translation in x- and y-direction between sensor and 3D printed object within  $\pm 1$  cm. Each trajectory contains 200 datapoints. Each datapoint contains the displacement of all of the 63 dots of the gel ( $D_C(t_i)$ ), as well as the shear force readings of the F/T sensor.

To evaluate the dot-tracking quality, we attempt to reconstruct the shear forces based on the displacement of the tracked dots. For this experiment, we investigate two models.

**Linear Regression.** Inspired by [23], we attempt to fit a linear model, mapping from the overall displacement of the dots in the x and y direction to the shear forces in the respective direction. For fitting the model, we use the training trajectories and do a least squares minimization.

**Neural Network.** The second model is a fully connected neural network. As input, this model takes the displacement of every individual dot in x- and y-direction. Given we have 63 dots, this forms a 126 dimensional input. The network consists of 2 fully connected layers with 128 neurons, each, using the ReLU activation function. We also add 25% dropout. On the last layer, we map to two outputs, i.e.,  $F_x$  and  $F_y$ . For training, we use the same trajectories we use to fit the linear model. However, we additionally divide these 4 trajectories in 3, which will be used for training, and 1 for testing. We train the model for 50 epochs and select the model with the lowest loss on the test dataset.

The results when evaluating the trained models on the previously unseen evaluation trajectory are presented in Table V and Figure 8. As shown in Figure 8, qualitatively, both of the models are able to reconstruct the shear force throughout the trajectory. Quantitatively (cf. Table V), considering the Linear Regression models, the force reconstruction through the dot tracking from Evetac, or the RGB optical tactile sensor yield similar results. The fact that the error for the x-direction is slightly higher might be related to the fact that there are more measurements at higher forces for the x-direction. Using the more powerful neural network model helps to improve the results, reducing the error in estimating  $F_x$  by around 37% for both sensors, and for the y-direction by around 46% and 16% for Evetac and the RGB optical tactile sensor, respectively. Overall, we conclude that reconstructing the shear forces from tracking the dots’ locations is possible with good accuracies. Moreover, the shear force reconstruction is of similar quality for both sensors. Small variations in the results might be related to the fact that the trajectories for recording the data with the different sensors might be slightly different. Nevertheless, these results are in line with our expectations, and they confirm that our proposed regularized dot tracking from Evetac’s raw output is of similar quality than performing dot tracking from images of the RGB optical tactile sensors.

## VII. SLIP DETECTION USING EVETAC

After successful validation of Evetac’s basic properties, we next focus on a more practically relevant task - slip detection. Reliable detection of slip is a crucial task in robotics and has received lots of attention. The task is especially important as any slippage is related with unstable contacts between

TABLE VI: Overview of the training and testing objects considered for the slip detection and grasp control experiments. In the materials column, G represents glass, M metal, P plastic, and Pa paper. For cylindric objects, only their diameter is provided (as width).



















	Training Objects.									Testing Objects.								
																		
Obj ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Mass g	341	325	70	102	236	273	100	114	119	72	108	38	116	166	56	74	248	550
Height mm $\updownarrow$	145	145	145	145	140	135	115	105	125	185	140	150	118	85	102	80	190	190
Width mm $\leftrightarrow$	74	70	35	27	36	52	26	41	33	50	48	10	40	71	25	42	60	60
Depth mm	61				36		63									35		
Material	G	G	M	P	G	M	P	Pa	G	P	M	M	P	P	P	Pa	G	G
Comment																	empty	filled



Fig. 9: Illustrating the original gel (left) in comparison with the modified cut version (right), providing the view from the outside and the inside. The modified gel enables seeing the object's movement from within the sensor (cf. Fig. 10). This allows running the slip classifier on the same image, which also contains the tactile measurements and circumvents any additional delays (cf. Sec. VII-B). The part of the gel was removed using a simple box cutter, cutting vertically until reaching the plexiglass and subsequently scraping off the small part.

finger, i.e., sensor, and object. For achieving stable grasping, any slippage requires quick corrective actions in order to prevent dropping the grasped object. In the following, we will introduce, our data-driven, model-free approach for slip detection. We particularly want to learn the slip detector from data as we want to avoid pre-specifying any slip criteria. We want the neural network models to automatically learn and focus on the most important information based on labeled training data. The following section is structured as follows. First, we describe the experimental setup and our procedure for labeled data collection. Second, we present the classifier for labeling the data. Third, we provide the training procedure and the architectures of the neural network models for slip detection. Last, we provide an evaluation of the trained models on previously unseen data, first considering the training objects, followed by an evaluation on novel objects. One section later, we also investigate the effectiveness of the slip detection models for stable object grasping, by integrating them into a real-time, reactive feedback control loop.

#### A. Experimental Setup and Data Collection Procedure

For collecting the labeled data, we use the setup shown in the leftmost picture of Figure 10. The setup consists of 2 Evetac sensors that are installed in ROBOTIS RH-P12-RN(A) 2 finger gripper. The two Evetac sensors are synchronized before data collection, and both read out at 1000 Hz. The gripper is controlled at 500 Hz. The training objects and their properties are presented in Table VI. As shown in the supplementary videos, for collecting the data, the objects will be grasped by the gripper. Subsequently, we force slip by opening the gripper using current control. Starting from the currently applied current, when the object is held stably, we adapt the target current every timestep by randomly sampling a value from within  $[-0.025, -0.0025]$  mA, i.e. considering one second, the target current is adapted between  $[-12.5, -1.25]$  mA. This




Fig. 10: Series of pictures illustrating Evetac's measurements during labeled data collection. Left, we see Object 8 held in the gripper, shortly (30 ms) before the object is going to slip. Upon further opening the gripper, the object starts to slip (middle). Note how, especially in the lower, window region of the gel, we can now clearly see events related to the texture moving. The number of events triggered by the moving texture increases further upon the object accelerating (right).

way, we collect 40 trajectories per object, and 360 in total, that we will use for training. We split them into 315 trajectories used to train the model and 45 trajectories that form the test dataset during training.

#### B. Data Labelling - Slip Classifier

One of the most important components for learning a slip detector from labeled data is the classifier for actually labeling the data. While most previous works rely on different sensors for data labeling, such as external cameras tracking Aruco markers [44], OptiTrack markers [13], or using IMUs attached to the object [34], in this work, we take a different approach. We aim to recover the moment of slip from the readings of Evetac. For this purpose, we cut the original gel and removed a small piece using a box cutter as shown in Figure 9. Due to the transparent plexiglass remaining in the cut region, we can essentially see through the sensor at this location. Now, looking through the transparent part of the sensor, i.e., the region without any gel remaining, we can determine when the object starts moving using an optical flow based criteria. While this choice comes at the disadvantage of losing part of the contact area of the tactile sensor (roughly 15%), it comes with the big advantage that the signal that is used to determine the moment of onset of object slippage is inherently aligned with the raw tactile measurements. Any errors due to synchronization can thus be eliminated, and the classifier naturally provides the same temporal resolution as the tactile readings. We nevertheless want to point out that we had to run the slip classification procedure offline, i.e., after having collected the data, thereby not having to meet any real-time requirements. The slip classifier relies on representing Evetac's raw measurements as images and the

objects exhibiting texture in the window region, as shown in Figure 10. We label Evetac’s raw measurements as slip or non-slip samples by computing the optical flow between the current measurement and the measurements 4 ms ahead. We motivate this forward-oriented flow calculation (i.e., calculating the flow of the current measurement w.r.t. a measurement in the future) with the fact that for events to be triggered, the object has to be in motion. Thus, if we have optical flow between consecutive Evetac measurements, we know that the object changed its position, but also that the motion started already in the initial frame. Moreover, this choice aims to mitigate any delays in slip classification and is enabled by the fact that we label the measurements offline, after having collected the data. For computing the optical flow, we use the OpenCV implementation of the Gunnar Farneback [63] method. We compute the flows for both regions of the image separately, i.e., the flow for the region with the markers and the window region. If the relative flow between the window and the marker region exceeds a certain threshold, then we mark the current measurement as belonging to the slip class. We consider the relative flow between the two regions since in situations where the flow in the tactile region and the transparent window region are equal, but both nonzero, we still want to obtain a non-slip label, as the elastic gel and object are moving in accordance and not relative to each other. Thus, there is no slip. In general, we found this analysis of optical flow more robust compared to a simpler approach that would only analyze the number of events in the respective regions, since optical flow is more invariant to the specific texture of the object. While the previously described procedure for determining the onset of slip requires that the object has texture, we want to point out that textured objects are only needed during data collection as they enable automatic data labeling. The trained models will only operate on the tactile data without having access to the cut, window region. Thus, during deployment, object texture is not necessary.

### C. Model Architectures

For training the models for slip detection, similar as in [40], we make use of Neural Networks. Apart from fitting the data as good as possible, we want to keep the network inference times low, such that we can later evaluate the models online in real time with 1000 Hz. As features, we will, therefore, mainly focus on two quantities. First, we make use of the current displacement  $d_{c_i}$  of the dots, i.e., their distance to their initial location. As shown in Sec. VI-D, this information can be used to reconstruct shear forces acting on the sensor and thus provide information about the global gel configuration. Second, we consider the number of events per dot  $N_E(t_i, c_i)$ , i.e., the current number of events triggered in the vicinity of each of the dots. From Sec. VI-A, we know that the number of events is effective for resolving high frequency phenomena and sensing vibrations. We want to point out that information from the transparent cut region is not available to the classifiers. They only have the information from the remaining  $7 \times 8 = 56$  dots. While we collected data for both Evetac sensors in the parallel gripper, herein, we will focus on slip detection

and grasp control using a single sensor only. We leave slip detection and grasp control using multiple Evetacs for future work.

The general architecture used for the slip detection models is depicted in Figure 11. In the later sections, we investigate different configurations of this architecture that differ in the input features that are available (cf. Table VII). As input, the models either receive the dots’ displacements, the number of events per dot, or both. In case of both, we first concatenate the dot displacement  $\mathcal{F}_D$  and dot event features  $\mathcal{F}_E$ . For almost all architectures, we consider a time series of measurements, i.e., we provide a history of previous measurements in addition to the current measurement. Given the resulting input vector per dot with dimension  $1 \times l_i$ , we first pass it through a two-layered fully connected neural network for encoding. This procedure is repeated for all of the dots, and we use the same weights for all dots. Next, we spatially combine these initial per-dot embeddings, in the same way as the dots are placed relative to each other in the gel. This allows us to subsequently use two convolutional layers, taking the spatial information and topology of the gel into account. We use convolutions, as slip is a local phenomenon, i.e., it is likely that an entire region of the gel is slipping. Thus, if an object is slipping, the slip signal should typically be sensed at multiple locations of the gel. Lastly, we flatten the features and pass them through two fully connected layers to receive the output, which is a scalar between  $[0, 1]$  and can be interpreted as a slip probability.

### D. Model Training & Data Selection

Apart from the model architecture, careful selection of the training data is also important. If we consider a single trajectory that is recorded as described in Sec. VII-A, it first contains many measurements where the object is held stably, then follows object slippage, and slightly after, the object has slipped completely and lost contact with the sensor. As also discussed in [44], it is important to determine, how many measurements to consider after the classifier detected slip for the first time. This choice is crucial since slightly after the onset of object slippage, it is likely that there is no contact between sensor and object. Thus, it does not make sense to classify such a measurement as belonging to the slip class. In line with these considerations, we also want to stress the importance of the models correctly detecting the moment of onset of slip, or even predicting it, as this greatly influences the timing of the corrective actions to prevent and minimize slippage. Empirically, we found best performance, when cutting the trajectories 15 ms after the onset of object slippage (i.e., after the classifier detected slip for the first time) for all models with a history of 10 or less, cutting after 20 ms for models with histories of 20, and cutting after 50 ms for the remaining architectures.

Cutting the trajectories after few samples of slip, however, also introduces data imbalance. We have way more datapoints of non-slip data. To counteract model bias, we use a modified version of the dataloader during training. In particular, we sample from three dataloaders at the same time. One of them solely contains slip data, and the other two non-slip data.



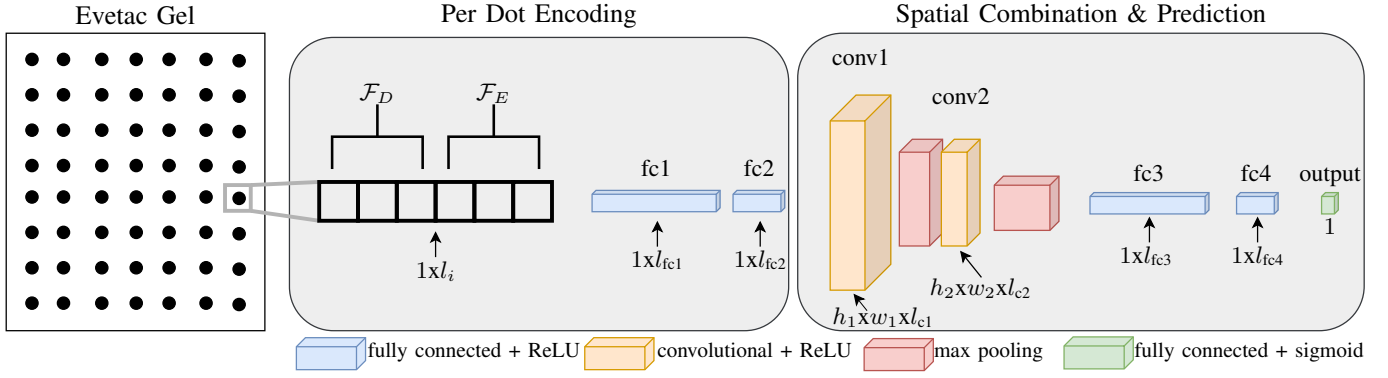


Fig. 11: General neural network architecture used for slip detection. First, per dot, a feature vector containing the current dot displacement  $\mathcal{F}_D$  (eventually combined with a sequence of past measurements) is extracted. The same is repeated for the number of events per dot, yielding  $\mathcal{F}_E$ . Second, the feature vectors are concatenated and embedded using two fully connected layers, sharing weights between the dots. Third, the embedded features are spatially combined, in the same way as the dots are placed relative to each other, and processed through two convolutional layers. Finally, the result is flattened, processed using two last fully connected layers, and yields the slip predictions.

TABLE VII: Parameters of the different model configurations used for slip detection. Fig. 11 shows the general architecture.

Name	$\mathcal{F}_D$	$\mathcal{F}_E$	$l_i$	$l_{fc1}$	$l_{fc2}$	$h_1, w_1, l_{c1}$	$h_2, w_2, l_{c2}$	$l_{fc3}$	$l_{fc4}$
no hist	$[d_c(t_i)]$	$[N_E(t_i, c)]$	2	10	4	7,6,16	3,3,32	32	10
hist 10	$[d_c(t_i), d_c(t_{i-1}), \dots, d_c(t_{i-9})]$	$[N_E(t_i, c), N_E(t_{i-1}, c), \dots, N_E(t_{i-9}, c)]$	20	12	4	7,6,16	3,3,32	32	10
events only									
hist 10		$[N_E(t_i, c), N_E(t_{i-1}, c), \dots, N_E(t_{i-9}, c)]$	10	8	4	7,6,16	3,3,32	32	10
disp only									
hist 10	$[d_c(t_i), d_c(t_{i-1}), \dots, d_c(t_{i-9})]$		10	8	4	7,6,16	3,3,32	32	10
hist 20	$[d_c(t_i), d_c(t_{i-1}), \dots, d_c(t_{i-19})]$	$[N_E(t_i, c), N_E(t_{i-1}, c), \dots, N_E(t_{i-19}, c)]$	40	20	8	7,6,16	3,3,32	32	10
hist 50	$[\sum_{l=0}^4 0.2d_c(t_i - l), \sum_{l=0}^4 0.2d_c(t_{i-5} - l), \dots, \sum_{l=0}^4 0.2d_c(t_{i-45} - l)]$	$[\sum_{l=0}^4 0.2N_E(t_i - l, c), \sum_{l=0}^4 0.2N_E(t_{i-5} - l, c), \dots, \sum_{l=0}^4 0.2N_E(t_{i-45} - l, c)]$	20	12	4	7,6,16	3,3,32	32	10
down 5									
fast slow	$[d_c(t_i), d_c(t_{i-1}), \dots, d_c(t_{i-9}), \sum_{l=0}^9 0.1d_c(t_i - l), \sum_{l=0}^9 0.1d_c(t_{i-10} - l), \dots, \sum_{l=0}^9 0.1d_c(t_{i-40} - l)]$	$[N_E(t_i, c), N_E(t_{i-1}, c), \dots, N_E(t_{i-9}, c), \sum_{l=0}^9 0.1N_E(t_i - l, c), \sum_{l=0}^9 0.1N_E(t_{i-10} - l, c), \dots, \sum_{l=0}^9 0.1N_E(t_{i-40} - l, c)]$	30	15	8	7,6,16	3,3,32	32	10
hist 50									

The dataloaders for sampling the non-slip data are different in that one of them contains measurements where the number of events exceeds a certain threshold, i.e., 25, and the other one contains the remaining non-slip measurements. The distinction across the non-slip frames attempts to separate measurements in which really nothing is happening, from measurements that do contain some events and information. The latter measurements are especially important to label correctly, as the event patterns have to be differentiated from the actual slip data. We train using a batch size of 72, where 32 samples are drawn from the slipping frames, 32 samples from the non-slipping frames above the threshold, and 8 frames from the non-slipping frames that are below the threshold. One episode consists of fully looping through the dataset with the non-slipping frames above the threshold. To additionally improve the models' robustness, we add data augmentation. Since all the training data was captured with the objects being grasped from top, with probability 50%, we rotate the input features within  $[90, 180, 270^\circ]$  to mimic different grasp poses. For the cases of 90 and 270°, we crop the measurement and pad it such that it is compatible with the networks' usual input size.

Using the training data, we train each of the models for 70 epochs, using stochastic gradient descent with a learning rate of 0.001 and the binary cross entropy loss. We log the models every 10 epochs.

### E. Evaluation Procedure & Metrics

Table VII shows the different model configurations we considered herein. For model evaluation, we will mainly consider

2 metrics. First of all, we investigate the point in time when the model first detects slip  $t_{s0}^m$ . This point in time is very important, as in a control task, action needs to be taken whenever slip is detected. It is thus highly undesirable if this point in time is not aligned well with the object really starting to slip, as determined by the classifier ( $t_{s0}^c$ ). In the following, we will refer to this as the slip timing criterion, which we define as follows. If the first instance of slip detected by the models  $t_{s0}^m$  is within the interval of  $t_{s0}^c - 50 \text{ ms} \leq t_{s0}^m \leq t_{s0}^c + 20 \text{ ms}$ , then we label this trajectory as one in which the slip timing was identified correctly, i.e., "slip corr". In other words, the models are allowed to identify slippage at most 50 ms prior to the first instance of slip detected by the classifier, and at latest 20 ms afterwards. We choose this rather long period before slip occurs, as it might be that the classifier learns to detect features related to incipient slip, which have been reported to be found up to 30 ms beforehand [34]. We, therefore, do not want to penalize potential detections of incipient slip and choose this time interval regarding the slip timing criterion. We refer to the remaining two cases as "slip detected too early", i.e.,  $t_{s0}^m < t_{s0}^c - 50 \text{ ms}$ , or as "too late", i.e.,  $t_{s0}^m > t_{s0}^c + 20 \text{ ms}$ . Note that the "too late" slip detection also includes trajectories where no slip is detected at all. As second metric, we consider the F1 score, which is the geometric mean of recall and precision. Recall is the ratio of true positives (TP) w.r.t. the sum of true positives and false negatives (FN), i.e.,  $\text{recall} = \frac{TP}{TP+FN}$ . Precision is the ratio of true positives w.r.t. the sum of true positives and false positives (FP), i.e.,  $\text{precision} = \frac{TP}{TP+FP}$ . Recall gets lower whenever positive samples are erroneously






Fig. 12: Evaluating different slip detection models on previously unseen trajectories (90 in total) using the training objects. The results show the mean and standard deviation averaged across all objects and five seeds per model configuration. The figure shows performance w.r.t. slip timing criterion (left) and F1 score (right).

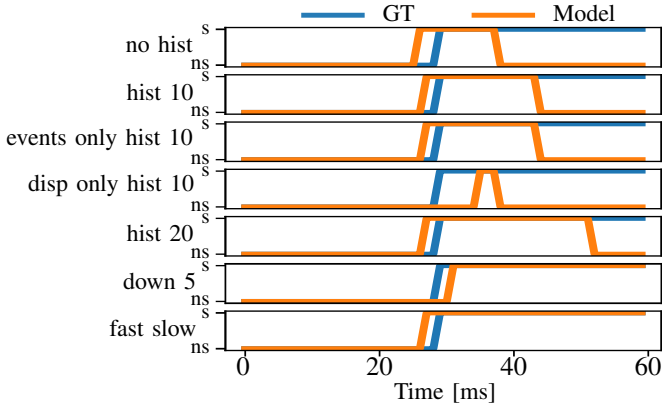


Fig. 13: Evaluating the different models on the task of slip detection considering a single previously unseen trajectory with Object 3. s corresponds to slip, and ns to no slip.

predicted as negative by the models, i.e., slip is not detected. Precision drops when non-slipping samples are erroneously labeled as slip. While this metric is related to the previous one, it has a clearer focus on temporal consistency, i.e., only detecting the moment in time of onset of slip correctly, will still result in a bad F1 score, as the recall would be low. Thus, the F1 score considers the entire trajectory. In line with our previous definition for the slip timing criterion, we cut the trajectories 20 ms after the classifier detected slip for the first time when calculating the F1 scores.

Given that the models' output can, at first glance, be interpreted as slip probability and is continuous, but our metrics require binary labels, we have to convert the models' output into a binary signal. This is done through thresholding. To determine a good threshold, we use the previously introduced dataset of 45 test trajectories and perform a grid search, sweeping the threshold value from 0 and 1 with increments of 0.025. We use the three last checkpoints of every model and perform the threshold optimization on all of them. Subsequently, we select the best combination of model checkpoint and threshold value, w.r.t. maximizing the combination of slip timing criterion and F1 score.

#### F. Evaluation on Training Objects

Figure 12 shows the results when evaluating the trained models on new, previously unseen trajectories using the training objects. For each object, we recorded 10 trajectories.

Overall, this evaluation considers 90 trajectories. For each model configuration (cf. Table VII), we trained 5 models using different seeds. The black bars in Figure 12 depict the standard deviation.

As illustrated, using the models without any history (i.e., no hist) already yield quite good results when considering the slip timing criterion. The onset of slip is detected correctly (i.e., "slip corr") in 87% of all trajectories. However, they only achieve an F1 score of 0.59. Increasing the history length to 10 (hist 10) slightly improves the rate of correct initial slip detections to 89.5% and the F1 score significantly to 0.73. For this model configuration with histories of 10, we also show an ablation, comparing models that either only have the number of events (events only) or the dots displacement (disp only) as features available. The models having access to the number of events per dot outperform the models only having access to the displacements. Overall, they nevertheless still perform slightly worse than the models that have access to both, especially when considering the F1 score. We therefore conclude that both features are crucial. Increasing the history length, i.e., to 20 ms (hist 20), further improves the rate of correct initial slip detection to 92.5%, the F1 score to 0.815. Using an even longer history, at the cost of downsampling the signal by a factor of 5 (i.e., the down 5 model), which mimics a sensor running at a 5 times reduced frequency, i.e., 200 Hz performs slightly worse in the correct initial slip detection with on average 88.2%. Lastly, we also investigate an architecture, which on the one hand has direct access to the last 10 measurements, i.e., the same input as the hist 10 model, but additionally also access to the last 50 measurements, which are downsampled by a factor of 10. We call this model fast slow, as it has direct, unfiltered access to the recent history, as well as knowledge about the signal evolution over a longer horizon. Due to the downsampling of the longer history, it overall still has a smaller input size than the model with a history of 20 (cf. Table VII). These fast slow models perform on par with the hist 20 model, having a correct initial slip detection rate of 92.6% and an average F1 score of 0.816.

These findings are also underlined in Figure 13, where we see the models' predictions over time for Object 3. The model without any history gets the moment of initial slip correct, however, has difficulties classifying slip states correctly afterwards, which explains the lower F1 scores. The model




Fig. 14: Evaluating two model configurations on the task of slip detection and prediction for  $\Delta T_{pred}=10$  ms or 20 ms in advance using the previously unseen testing objects. On the left, we show the slip timing criterion, and on the right, the F1 scores. Mean and standard deviation are reported, averaging across 80 trajectories, and five seeds per model configuration.

with a history of 10 classifies slip correctly over a longer time horizon and thus achieves better temporal consistency and higher F1 scores. For this trajectory, the events only model performs similarly, while the model that only has the displacements available exhibits a delayed slip detection. This might be due to the fact that there is a delay between initial events signaling slip and significant dot movement. Moreover, it also struggles achieving temporal consistency. Lastly, the models with longer histories are capable of capturing the timely evolution substantially better. We can also see that the downsampled model (down 5) is slightly delayed, considering initial slip detection. This can be explained by the fact that it can only make predictions every 5 ms, thus, this phenomenon is due to temporal discretization.

In the following experiments, we will present a closer investigation of the two best performing models, i.e., the model with 20 ms history (hist 20), as well as the fast slow model.

#### G. Evaluating on unseen Testing Objects & Slip Prediction

Using the previously best performing models, i.e., the models with a history of 20, as well as the fast slow models, we next present an evaluation on 8 previously unseen testing objects (Objects 9-17, cf. Table VI). Again, for every object, we record 10 trajectories, yielding 80 trajectories. Additionally, we now also investigate the effectiveness of training the models on the task of slip prediction. In particular, we train both model configurations on data where we shift the classifier signal by  $\Delta T_{pred}=10$  ms and  $\Delta T_{pred}=20$  ms forward in time. The respective models are thus tasked to detect slip 10 ms or 20 ms before the classifier detected the onset of slip. Predicting slip is beneficial, as it increases the time window to react and counteract the slippage. For consistency, the F1 scores for the slip prediction models are computed using the shifted labels. The slip prediction models are trained on the same data as the other models and we again consider five different seeds.

The results are shown in Figure 14. Considering the prediction qualities w.r.t. slip timing criterion, we observe that for the model with a history of 20, the percentage of correct initial slip detections (slip corr) continuously decreases (i.e., 86%, 84%, 81%) when training the models for the task of slip prediction. In line with this decrease, the percentage of trajectories where slip is detected too early increases. Thus, when training the hist 20 models for slip prediction, it seems that the models adjust

their features, which, however, also results in detecting slip too early more frequently. For the fast slow models, the results are slightly different. Here, the models that have been trained on the task of detecting slip 10 ms in advance yield the highest percentage of correct slip detections (88%), while the models trained on the original signal and on the 20 ms shifted one perform equally (around 84%), with the ones tending towards detecting slip too late, while the others again rather detect slip too early, respectively. It thus seems that the fast slow models which can access the evolution of Evetac’s measurements over a longer horizon, however, downsampled, as well as the most recent measurements, are capable of extracting discriminative slip prediction features while avoiding an increase in detecting slip too early for the task of predicting slip 10 ms ahead. Considering F1 scores, the models for just detecting slip  $\Delta T_{pred}=0$  ms perform comparable. For the task of predicting slip 10 ms ahead in time, the fast slow architecture achieves similar scores compared to the baseline slip detection models, while the hist 20 models perform worse on average. For the task of predicting slip 20 ms in advance, the hist 20 models perform better than the fast slow architecture, however, as analyzed previously, these models come at the cost of significantly lower performance concerning the timing of the first slip detection.

In this experiment, we found a tradeoff between attempting to predict slip, and the models detecting slip way too early. We nevertheless find that the fast slow architecture offers slightly beneficial performance in that the tendency to detect slip too early is not as prominent while attempting to detect slip 10 ms ahead. In fact, this model still performs on par with the non-predictive models, considering F1 score. We will thus use the fast slow model architecture with predicting slip 10 ms ahead in all the following experiments.

#### H. Evaluating Slip Timing

Finally, we analyze the timing of the first slip detection  $t_{s0}^m$  of the previously selected fast slow models predicting slip 10 ms ahead, on the previously unseen testing objects. Table VIII shows the results, again averaging across the five seeds. In line with the configuration the model has been trained in, most of the initial slip detections, i.e., 57%, are within the interval of 10 ms ahead to the actual moment of onset of slip (i.e., 10-0). The second most frequent interval of initial slip

TABLE VIII: Evaluating the timing of the fast slow models' initial slip detection  $t_{s0}^m$  trained for the task of predicting slip 10 ms in advance using the testing objects. Reported are mean and standard deviation considering 5 models trained with different seeds.

Obj	Avg	10	11	12	13	14	15	16	17
Slip too early	0.07 (0.02)	0.1 (0.0)	0.0 (0.0)	0.1 (0.06)	0.0 (0.0)	0.0 (0.0)	0.18 (0.07)	0.0 (0.0)	0.18 (0.07)
$t_{s0}^m < t_{s0}^c - 50$ ms									
50-30	0.04 (0.02)	0.0 (0.0)	0.0 (0.0)	0.14 (0.08)	0.07 (0.04)	0.02 (0.04)	0.0 (0.0)	0.0 (0.0)	0.06 (0.05)
$t_{s0}^c - 50$ ms $\leq t_{s0}^m$ $< t_{s0}^c - 30$ ms									
30-10	0.22 (0.01)	0.1 (0.0)	0.24 (0.08)	<b>0.48</b> (0.12)	0.38 (0.04)	0.04 (0.05)	0.22 (0.04)	0.08 (0.1)	0.26 (0.14)
$t_{s0}^c - 30$ ms $\leq t_{s0}^m$ $< t_{s0}^c - 10$ ms									
10-0	<b>0.57</b> (0.03)	<b>0.76</b> (0.08)	<b>0.72</b> (0.1)	0.08 (0.04)	<b>0.53</b> (0.04)	<b>0.74</b> (0.1)	<b>0.32</b> (0.04)	<b>0.9</b> (0.09)	<b>0.48</b> (0.12)
$t_{s0}^c - 10$ ms $\leq t_{s0}^m$ $< t_{s0}^c - 0$ ms									
0-20	0.05 (0.01)	0.0 (0.0)	0.0 (0.0)	0.04 (0.05)	0.02 (0.05)	0.16 (0.05)	0.14 (0.05)	0.0 (0.0)	0.02 (0.04)
$t_{s0}^c - 0$ ms $\leq t_{s0}^m$ $< t_{s0}^c + 20$ ms									
Slip too late	0.05 (0.01)	0.04 (0.08)	0.04 (0.08)	0.16 (0.05)	0.0 (0.0)	0.04 (0.05)	0.14 (0.05)	0.02 (0.04)	0.0 (0.0)
$t_{s0}^m > t_{s0}^c + 20$ ms									




Fig. 15: Closed-loop grasp control experiments. We mount the gripper equipped with Evetacs (cf. Fig. 1) on a Franka Panda 7 DoF robot, for stably grasping and lifting previously unseen objects. OptiTrack markers are attached to the object to determine how much the object moved relative to the gripper. (a) depicts the initial situation in which the gripper establishes light contact with Object 17 (cf. Table VI), which is insufficient for lifting. Upon the robot starting to lift, object slip is detected by the slip detector, resulting in the closed-loop grasp controller adjusting the grasping force. (b) shows the end of the successful lifting phase. After object lift, we aim to minimize the applied grasping force, by identifying the gripper opening width that is just sufficient for grasping, and at the boundary to the object slipping. This so-called balancing phase typically includes small object slippages, and (c) depicts the object pose at the end of this phase. The successful completion of this experiment, also demonstrates that our slip detector does not require any object texture, as the object is transparent at the grasping locations.

detection is between 30 to 10 ms ahead in time. Averaged across all objects, the success percentages of detecting slip correctly, i.e., the interval between 50 ms prior to 20 ms after the classifier detects slip for the first time, is high, with 88%. However, for Objects 12 and 15, slip is often predicted too early or too late (26% and 32%, respectively). Potential explanations for these findings could be that Objects 12 and 15 are lighter and smaller compared to the objects that the model has been trained on (cf. Table VI). Their different weight and size will deform the gel differently, which might be one reason for the slightly reduced performance.

## VIII. CLOSED-LOOP GRASP CONTROL USING EVETAC

In this last experimental section, we investigate the effectiveness of using Evetac in combination with the previously introduced slip detection and prediction models for reactive

robotic grasping. In particular, we equip a ROBOTIS RH-P12-RN(A) gripper with two Evetac sensors (cf. Figure 1) and mount it as the end effector onto a Franka Panda robot arm, as shown in Figure 15. As also mentioned previously, herein, we only consider the signal of one Evetac and use it in combination with the best model architecture from the previous section, i.e., the fast slow model that has been trained on predicting slip 10 ms ahead. The slip prediction model is integrated into a real-time grasp control loop. In the following, we first provide the experimental setup and description of the control strategy, followed by closed-loop pickup and grasping experiments using the previously unseen testing objects. We end the section by investigating the robustness of the grasp controller w.r.t. changing from top-down to sideways grasps, and adding additional disturbances by dropping weights onto the grasped objects. Again, we provide videos on our website.

### A. Evaluation Procedure & Control Strategy

The experiments deal with the situation in which a robot has to pick up an a priori unknown object. To solve this task, first, the robot has to position the gripper such that the fingers will make contact with the object upon closing. Subsequently, sufficient grasping force has to be applied to lift the object stably. Finally, if having to hold the object for longer, it will be beneficial to apply minimal grasping forces for improved efficiency. Throughout the whole procedure, it is desirable that the grasping force is adaptive w.r.t. the object that is to be lifted. The controller should apply less force when dealing with lighter objects, and vice versa. Moreover, the grasp controller should be reactive w.r.t. any disturbances. For control, we use the real-time position control interface of the gripper. Given the gripper's current opening width  $x_g$ , we specify the reference opening width  $x_{ref}$ . Together with the control gain  $K_p = 50$ , the control law yields  $u_c = K_p(x_e + u_{ff})$  with  $x_e = x_{ref} - x_g$ . The reason for using the position control interface is that it offers a higher resolution and, therefore, finer control compared to the current control interface.

To investigate the effectiveness of our trained models for online grasp control in the previously described scenario, we propose the following procedure. We first move the robot to a suitable pre-grasp pose, assuming knowledge about the object's pose. Next, we close the gripper to make light contact with the object to avoid any damage. This is achieved by setting the feedforward signal to a small constant value  $u_{ff} = -2$ , which is just sufficient to make the gripper move, while keeping  $x_e = 0$ . The chosen, small control signal will make the fingers stop upon making very light contact. When attempting to lift the object starting from this initial light grasp configuration, the object would just slip and remain on the table surface, as the applied forces are not sufficient.

**Lifting Phase.** To counteract this slippage during the lift, we use our slip detector in an online fashion, essentially controlling the width of the gripper. Upon detecting slippage, we attempt to further close the gripper until the slip stops. In particular, we now have a time-dependent feedforward term  $u_{ff}(t) = u_{ff}(t-1) + u_i(t)$ , with the increment  $u_i(t) = -1$ , if slip was detected ( $s(t)=1$ ) between the last and the current call




Fig. 16: Closed-loop grasp control experiments using the previously unseen testing objects (cf. Table VI). For every object, we perform 10 trials. While the top row shows the success rates, the middle row shows the change in gripper opening applied by the controller for lifting the object and optimizing the grasp force, i.e., opening width, during the subsequent balancing phase. The bottom row shows how much the objects moved relative to the gripper during the whole procedure of pickup, lift, and balance. For an exemplary trajectory, see Fig. 15.

to the controller, and  $u_i(t)=0.01$ , if no slip was detected ( $s(t)=0$ ). The reason for the different increments is that whenever slip has been detected, we want to react fast, while we want to reduce the grasping force gently in case of no slip, as the reduction of grasping force might lead to new slippage. Moreover, to limit the amount of force the gripper can exert onto the object, we clip  $u_{ff}(t)$  to stay within  $[-5, 0]$ . Additionally, in the switching moments from slip to no slip (i.e.  $s(t-1)=1$  and  $s(t)=0$ ), we set the reference gripper position to the current gripper opening width  $x_{ref}=x_g(t)$ . This is the first gripper closing width realizing a stable grasp without slippage, and should thus be the desirable setpoint. Before the first occurrence of this switching moment, we leave  $x_e$  zero. Note that this control law is reactive, i.e., if there are multiple occurrences where a switch between slip and no slip happens, the reference is adapted accordingly. Yet, the grasping force can only increase. Since the slip detector is running at 1000 Hz, while the control loop is operating at half the frequency, we choose a shared memory to pass information between the slip detector and control loop asynchronously. We create one integer variable that is incremented whenever slip occurs. In the controller, we access this variable and check whether the slip count increased (slip), or is unchanged (no slip) w.r.t. the previous function call.

**Balancing Phase.** After successful object lift, we aim to minimize grasping force while holding the object. We thus attempt to open the gripper, i.e. decrease the grasping force until slip is detected, and then counteract it, similar as in the previous maneuver. However, since now, gravity is also acting on the object, we use a bigger increment to counteract slippage. In particular, we use the same time-varying feedforward term. Now, we open the gripper through  $u_i(t)=0.01$  as long as no slip is detected, and close it through  $u_i(t)=-2$  upon detecting slip. Moreover, the first moment of detecting slip

(i.e.,  $s(t-1)=0$  and  $s(t)=1$ ), activates the reference position of minimal force for holding the object ( $x_{ref}=x_g(t-1)$ ). Before,  $x_e$  has been set to zero. Additionally, as initially, we want to open the gripper, we clip  $u_{ff}(t)$  within  $[-5, 2]$ . However, upon first slip detection, we adapt it to  $[-5, 0]$ , as we do not want to open the gripper beyond the reference position, which is the last stable position.

For evaluating the performance of our closed-loop grasp controller, we employ the following metrics. First, we consider the success rates of the individual phases. Success lift represents successful lifting of the object by around 10 cm, i.e., lifting the object without it slipping completely through the fingers. Successful balance means successful completion of the second phase, i.e., opening the gripper until slippage is detected, catching the object, and subsequently holding the object applying as little force as possible. This phase is successful, if it still ends with the object in between the fingers. The two phases take 10 and 20 s, respectively. Please note that the balance success rate only counts the trials for which the lifting was successful. Overall success, is the amount of trajectories for which both phases were successful. As shown in Figure 15, we also attached Optitrack markers to the objects to measure by how much the object moved relative to the gripper throughout the entire maneuver. Lastly, we will also report the change in gripper opening width, which provides insight on the grasping strength applied to the objects. As we do not have any force sensing capabilities available, similar as in [44], we assume the grasp strength to be proportional to the gripper closing width.

#### B. Evaluating the Closed-Loop Pickup and Grasp Controller

Figure 16 shows the experimental results for the testing objects, i.e., the objects that have not been seen during model training (cf. Table VI). For every object, we conduct 10 trials,



and we use the same fast slow slip prediction model, predicting slip 10 ms ahead in time, and control strategy. To ensure that the initial closing of the gripper for establishing object contact does not erroneously result in already establishing sufficient forces for successful lifting, for all objects, we also did 5 repetitions in closing the gripper, however, then attempting to lift the object without any further control. This resulted in 0% lifting successes. Therefore, slip detection and appropriate closed-loop control are crucial for successful completion. As shown in the first row of Figure 16, across all objects, in total, we have a success rate of 92%. Only for Objects 12 & 15 we have one failure during lifting, and for Object 14, in the nominal configuration, we have 4 failures during balance control. Regarding the two lifting failures, the initial contact forces might have been too light, and thus slip has not been detected. Moreover, they occur with the two lightest objects that are lighter than all of the training objects (cf. Table VI). These failed lifts are also in line with the results from the offline slip detection & prediction experiments, where for both objects (12 & 15), we had more than 10% probability of not detecting slip at all. For the cone-shaped Object 14, the 4 balancing failures might be due to the fact that we had to grasp the object upside down, which makes the task of stably grasping and balancing more difficult once the object has slipped and accelerated. As can also be seen in the last column of the figure, we repeated the experiment for Object 14 and doubled the grasp control gain ( $K_p$ ) during the balancing phase. This resulted in a reduced number of balancing failures. This hints at the fact that the failures in balancing might not only be due to the slip detection, but the interplay with control is also crucial. However, we believe that, in general, one should try to keep the control signals rather small in order to avoid exerting excessive forces, which is especially problematic when dealing with more delicate objects. For all the other six objects, we did not observe any failures. Looking at the second row, which shows the change in gripper opening width, we can actually see that the grasping force during lifting and the overall maneuver is really adaptive w.r.t. the object that is grasped. Since object properties such as their surface might also play a role in the required grasping forces, the first two columns of the figure provide a good comparison as they compare the grasping efforts for the same object (i.e., a bottle), that is once filled (Obj 18) and once empty (Obj 17). We can clearly observe that for the heavier, filled object, the closed-loop grasping control pipeline applies more grasping strength, during lifting and also balance. Comparing the filled and unfilled object, the overall change in gripper opening width for grasping and stabilizing is  $-13.4$  and  $-7.6$ , on average, respectively. This underlines the adaptiveness of the proposed grasp controller. The filled bottle is also roughly twice as heavy as the empty one. In the figure, we can also see, that the balancing phase is effective in that the grasping strength can be reduced across all the objects, on average by 43%. Lastly, when investigating the distance that the object traveled, we see that during lifting, in almost all attempts, the objects only move by a couple of millimeters. Considering the overall maneuver, for most experiments, the objects move less than 1 cm. As it is also shown in the supplementary

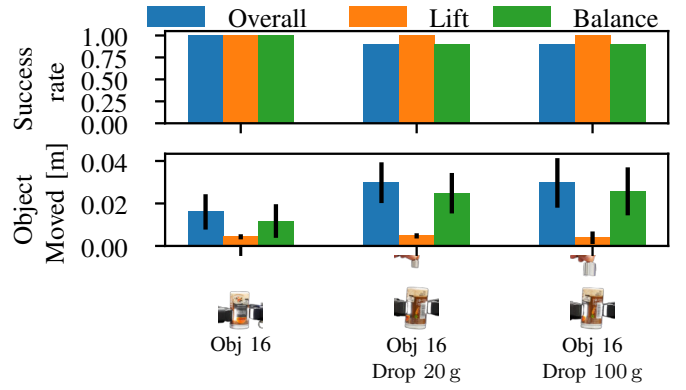


Fig. 17: Grasp controller robustness evaluation. Contrary to data collection and the previous experiments, Object 16 is now grasped sideways. The left column presents the results for the standard maneuver as done in Fig. 15, while the other two columns consider cases where the grasp is further perturbed during the balancing phase by dropping weights with 20 and 100 g onto the object after having established an initial stable grasp. As metrics, success rate and the distance that the object moved relative to the gripper are reported.

videos, some objects move more than once in the balancing phase. This on the one hand illustrates the reactivity of our controller and also underlines that the gripper opening width has to be carefully chosen. It can occur that the first desired setpoint is still too open, which might potentially be related with the finger's velocity during the opening movement. The slip detector and grasp controller are run jointly, and in real-time on a single desktop PC with 128 GB RAM, *NVIDIA GeForce RTX 3090* GPU, and *AMD Ryzen 9 5950X 16-Core* CPU. All components from reading the sensor, tracking the dots, and evaluating the neural network slip detector (mean inference time of 95  $\mu$ s) were run at 1000 Hz.

### C. Evaluating Controller Robustness w.r.t. Grasp Orientation & External Disturbances

In the last experiments, we show the effectiveness of the proposed pipeline for scenarios in which objects might be grasped differently, i.e., from the side. Such data has not been included in the raw training data. It is only covered through our data augmentation. The sideways grasps result in a different orientation of Evetac w.r.t. the object. We also investigate the reactivity of the proposed control laws by dropping weights of either 20 g or 100 g onto the object, after successful initial stabilization in the balancing phase. As dropping 100 g onto the object is a substantial perturbation, in this configuration, we again double the control gain ( $K_p$ ) to maintain high success rates. We perform 10 trials per experiment configuration.

As can be seen in Figure 17, grasping the object sideways still results in ten out of ten successful trials for Object 16. Also, the distance that the object moves is only slightly increased and still comparable to the previous experiment. Considering the scenarios where additional weights are dropped onto the object upon grasp stabilization, we see that the distances that the object moves increase. The results also show that the closed-loop control pipeline is reactive as the balancing success rates remain high at 90%. Only in one trial there is one balancing failure for both configurations.



Additionally, the experiments with dropping the 20 g weight object indicate that the grip controller does not apply excessive forces. The addition of 20 g is already sufficient to destabilize the grasp and make the object move inside the gripper, as can also be seen in the supplementary videos.

## IX. DISCUSSION

In the previous experimental sections, we investigated Evetac's properties, demonstrated different models for offline slip detection and prediction, and finally showed Evetac's effectiveness in grasping different known and previously unseen objects with different surfaces, materials, and weights. While Evetac – in combination with the corresponding touch processing and control algorithms – yields good performance across all the tasks and satisfies the desiderata of high-frequency sensing, processing, and control, we also discovered some limitations. In the grasp control experiments, our approach had the most failures with the cone-shaped Object 14. We hypothesize that one reason for these difficulties is that Evetac's gel surface is planar and cannot adapt well to the geometry of the grasped object. One way to overcome this limitation might be to employ a different gel with more curvature, as used in the BioTacs [32] or TacTip [24]. Another limitation of Evetac is the sensor's dimension, which makes integration with dexterous robotic hands challenging. For this purpose, we hope that in the near future smaller event-based cameras will be released, as the camera's size is currently the major limiting factor. Regarding the slip detection and prediction experiments, it would be interesting to consider additional network types, such as spiking neural networks, in the future. They hold the potential to further decrease the latency of the overall pipeline, as they can directly operate on the sensor's asynchronous output. Lastly, the slip timing experiments indicated that it is object-dependent. Future work should try to identify the causes for these findings. They could provide essential information regarding feature selection or sensor material choice.

## X. CONCLUSION

This work introduced a new event-based optical tactile sensor called Evetac. The sensor design aims to maximize re-use of existing components and solely requires 3D printing of a housing that connects together the event-based camera, the soft silicone gel, and the lighting. In addition to the sensor design, this work also presented the necessary software to read out the sensor in real-time at 1000 Hz, as well as suitable touch processing algorithms running at the same frequency. In particular, we devised a novel algorithm for tracking the dots imprinted in the gel and proposed a set of tactile features that were exploited for learning efficient neural network-based slip detectors from collected data. The thorough experimental section first demonstrated the sensor's natural properties of being able to detect tactile vibrations of up to 498 Hz, providing significantly reduced data rates compared to RGB optical tactile sensors despite Evetac's high temporal resolution, and exploiting the dot displacements for shear force reconstruction. To showcase Evetac's practical relevance, we also evaluated and compared different neural network architectures for the

task of slip detection and prediction from Evetac's data on a wide range of objects. The models formed the basis for designing an effective real-time grasp control loop, achieving high success rates of 92%, robustness to perturbations, and adaptiveness w.r.t. object mass. We hope that our proposed open-source Evetac sensor, together with the touch processing algorithms and closed-loop grasping controller, will encourage further research in the field of event-based tactile sensing, and contribute to the efforts for achieving the still unparalleled human manipulation capabilities.

## ACKNOWLEDGMENTS

We thank Luca Dziarski for his help with the force reconstruction experiments. This work is supported by the AICO grant by the Nexplore/Hochtief Collaboration with TU Darmstadt. Calculations for this research were conducted on the Lichtenberg high performance computer of the TU Darmstadt. This work was partly supported by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden, and by Bundesministerium für Bildung und Forschung (BMBF) and German Academic Exchange Service (DAAD) in project 57616814 (SECAI, School of Embedded and Composite AI). This work was also supported by the German Research Foundation (DFG) Emmy Noether Programme (CH 2676/1-1).

## REFERENCES

- [1] Korbinian Nottensteiner, Arne Sachtler, and Alin Albu-Schäffer. Towards autonomous robotic assembly: Using combined visual and tactile sensing for adaptive task execution. *Journal of Intelligent & Robotic Systems*, 2021.
- [2] Francesco Grella, Giulia Baldini, Roberto Canale, Keerthi Sagar, Si Ao Wang, Alessandro Albini, Michal Jilich, Giorgio Cannata, and Matteo Zoppi. A tactile sensor-based architecture for collaborative assembly tasks with heavy-duty robots. In *ICAR*, 2021.
- [3] Matei Ciocarlie, Kaijen Hsiao, Edward Gil Jones, Sachin Chitta, Radu Bogdan Rusu, and Ioan A Şucan. Towards reliable grasping and manipulation in household environments. In *Experimental Robotics: The 12th International Symposium on Experimental Robotics*, 2014.
- [4] Tapomayukh Bhattacharjee, Henry M Clever, Joshua Wade, and Charles C Kemp. Multimodal tactile perception of objects in a real home. *IEEE Robotics and Automation Letters*, 2018.
- [5] Luca Lach, Niklas Funk, Robert Haschke, Severin Lemaignan, Helge Joachim Ritter, Jan Peters, and Georgia Chalvatzaki. Placing by touching: An empirical study on the importance of tactile sensing for precise object placing. In *IROS*, 2023.
- [6] Qiang Li, Oliver Kroemer, Zhe Su, Filipe Fernandes Veiga, Mohsen Kaboli, and Helge Joachim Ritter. A review of tactile information: Perception and action through touch. *IEEE T-RO*, 2020.
- [7] Ravinder S Dahiya, Giorgio Metta, Maurizio Valle, and Giulio Sandini. Tactile sensing—from humans to humanoids. *IEEE T-RO*, 2009.
- [8] Zhanat Kappasov, Juan-Antonio Corrales, and Véronique Perdureau. Tactile sensing in dexterous robot hands. *Robotics and Autonomous Systems*, 2015.
- [9] Wei Chen, Heba Khamis, Ingvars Birznies, Nathan F Lepora, and Stephen J Redmond. Tactile sensors for friction estimation and incipient slip detection—toward dexterous robotic manipulation: A review. *IEEE Sensors Journal*, 2018.
- [10] Akihiko Yamaguchi and Christopher G Atkeson. Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision? *Advanced Robotics*, 2019.
- [11] Alexander C Abad and Anuradha Ranasinghe. Visuo-tactile sensors with emphasis on gelsight sensor: A review. *IEEE Sensors Journal*, 2020.
- [12] Alexander Schmitz, Marco Maggiali, Lorenzo Natale, Bruno Bonino, and Giorgio Metta. A tactile sensor for the fingertips of the humanoid robot icub. In *IROS*, 2010.
- [13] Tasbolat Taunyazov, Weicong Sng, Brian Lim, Hian Hian See, Jethro Kuan, Abdul Fatir Ansari, Benjamin Tee, and Harold Soh. Event-Driven Visual-Tactile Sensing and Learning for Robots. In *RSS*, 2020.

- [14] Gereon H Büscher, Risto Koiva, Carsten Schürmann, Robert Haschke, and Helge J Ritter. Flexible and stretchable fabric-based tactile sensor. *Robotics and Autonomous Systems*, 2015.
- [15] Raunaq Bhirangi, Tess Hellebrekers, Carmel Majidi, and Abhinav Gupta. Reskin: versatile, replaceable, lasting tactile skins. In *Conference on Robot Learning*, 2021.
- [16] Jeremy A Fishel. *Design and use of a biomimetic tactile microvibration sensor with human-like sensitivity and its application in texture discrimination using Bayesian exploration*. PhD thesis, USC, 2012.
- [17] Roland S Johansson and J Randall Flanagan. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 2009.
- [18] Ake B Vallbo, Roland S Johansson, et al. Properties of cutaneous mechanoreceptors in the human hand related to touch sensation. *Hum neurobiol*, 1984.
- [19] Giulia Corniani and Hannes P Saal. Tactile innervation densities across the whole body. *Journal of Neurophysiology*, 2020.
- [20] Amin Rigi, Fariborz Baghaei Naeini, Dimitrios Makris, and Yahya Zweiri. A novel event-based incipient slip detection using dynamic active-pixel vision sensor (davis). *Sensors*, 2018.
- [21] Kenta Kumagai and Kazuhiro Shimonomura. Event-based tactile image sensor for detecting spatio-temporal fast phenomena in contacts. In *IEEE WHC*, 2019.
- [22] Benjamin Ward-Cherrier, Nicholas Pestell, and Nathan F Lepora. Neuro-tac: A neuromorphic optical tactile sensor applied to texture recognition. In *ICRA*, 2020.
- [23] Wenzhen Yuan, Siyuan Dong, and Edward H Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 2017.
- [24] Benjamin Ward-Cherrier, Nicholas Pestell, Luke Cramphorn, Benjamin Winstone, Maria Elena Giannaccini, Jonathan Rossiter, and Nathan F Lepora. The tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies. *Soft robotics*, 2018.
- [25] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, 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*, 2020.
- [26] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [27] Yaonan Zhu, Shukrullo Nazirjonov, Bingheng Jiang, Jacinto Colan, Tadayoshi Aoyama, Boris Belousov, Kay Hansel, and Jan Peters. Visual tactile sensor based force estimation for position-force teleoperation. In *IEEE CBS*, 2022.
- [28] Rocco A Romeo and Loredana Zollo. Methods and sensors for slip detection in robotics: A survey. *IEEE Access*, 2020.
- [29] Marc R Tremblay, WJ Packard, and Mark R Cutkosky. Utilizing sensed incipient slip signals for grasp force control. In *Japan-USA Symposium on Flexible Automation*, 1992.
- [30] Roland S Johansson and Göran Westling. Tactile afferent signals in the control of precision grip. In *Attention and performance XIII*. 2018.
- [31] Robert D Howe and Mark R Cutkosky. Sensing skin acceleration for slip and texture perception. In *ICRA*, 1989.
- [32] Jeremy A Fishel, Veronica J Santos, and Gerald E Loeb. A robust microvibration sensor for biomimetic fingertips. In *2nd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics*, 2008.
- [33] Carsten Schürmann, Matthias Schöpfer, Robert Haschke, and Helge Ritter. A high-speed tactile sensor for slip detection. *Towards Service Robots for Everyday Environments: Recent Advances in Designing Service Robots for Complex Tasks in Everyday Environments*, 2012.
- [34] Zhe Su, Karol Hausman, Yevgen Chebotar, Artem Molchanov, Gerald E Loeb, Gaurav S Sukhatme, and Stefan Schaal. Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor. In *Humanoids*, 2015.
- [35] Matthias Schöpfer, Carsten Schürmann, Michael Pardowitz, and Helge Ritter. Using a piezo-resistive tactile sensor for detection of incipient slippage. In *ISR*, 2010.
- [36] Barrett Heyneman and Mark R Cutkosky. Slip interface classification through tactile signal coherence. In *IROS*, 2013.
- [37] Filipe Veiga, Herke Van Hoof, Jan Peters, and Tucker Hermans. Stabilizing novel objects by learning to predict tactile slip. In *IROS*, 2015.
- [38] Wenzhen Yuan, Rui Li, Mandayam A Srinivasan, and Edward H Adelson. Measurement of shear and slip with a gelsight tactile sensor. In *ICRA*, 2015.
- [39] Siyuan Dong, Wenzhen Yuan, and Edward H Adelson. Improved gelsight tactile sensor for measuring geometry and slip. In *IROS*, 2017.
- [40] Jianhua Li, Siyuan Dong, and Edward Adelson. Slip detection with combined tactile and visual information. In *ICRA*, 2018.
- [41] Siyuan Dong, Daolin Ma, Elliott Donlon, and Alberto Rodriguez. Maintaining grasps within slipping bounds by monitoring incipient slip. In *ICRA*, 2019.
- [42] Jasper W James, Stephen J Redmond, and Nathan F Lepora. A biomimetic tactile fingerprint induces incipient slip. In *IROS*, 2020.
- [43] Jasper Wollaston James, Nicholas Pestell, and Nathan F Lepora. Slip detection with a biomimetic tactile sensor. *IEEE Robotics and Automation Letters*, 2018.
- [44] Jasper Wollaston James and Nathan F Lepora. Slip detection for grasp stabilization with a multifingered tactile robot hand. *IEEE T-RO*, 2020.
- [45] Wang Wei Lee, Sunil L Kukreja, and Nitish V Thakor. A kilohertz kilotaxel tactile sensor array for investigating spatiotemporal features in neuromorphic touch. In *BioCAS*, 2015.
- [46] Anupam K Gupta, Rohan Ghosh, Aravindh N Swaminathan, Balakrishna Deverakonda, Godwin Ponraj, Alcimar B Soares, and Nitish V Thakor. A neuromorphic approach to tactile texture recognition. In *ROBIO*, 2018.
- [47] Andrei Nakagawa-Silva, Nitish V Thakor, John-John Cabibihan, and Alcimar B Soares. A bio-inspired slip detection and reflex-like suppression method for robotic manipulators. *IEEE Sensors Journal*, 2019.
- [48] Benjamin Ward-Cherrier, Jörg Conradt, Manuel G Catalano, Matteo Bianchi, and Nathan F Lepora. A miniaturised neuromorphic tactile sensor integrated with an anthropomorphic robot hand. In *IROS*, 2020.
- [49] Fariborz Baghaei Naeini, Aamna M AlAli, Raghad Al-Husari, Amin Rigi, Mohammad K Al-Sharman, Dimitrios Makris, and Yahya Zweiri. A novel dynamic-vision-based approach for tactile sensing applications. *IEEE Transactions on Instrumentation and Measurement*, 2019.
- [50] Fariborz Baghaei Naeini, Dimitrios Makris, Dongming Gan, and Yahya Zweiri. Dynamic-vision-based force measurements using convolutional recurrent neural networks. *Sensors*, 2020.
- [51] Fariborz Baghaei Naeini, Sanket Kachole, Rajkumar Muthusamy, Dimitrios Makris, and Yahya Zweiri. Event augmentation for contact force measurements. *IEEE Access*, 2022.
- [52] Xiaodian Huang, Rajkumar Muthusamy, Eman Hassan, Zhenwei Niu, Lakmal Seneviratne, Dongming Gan, and Yahya Zweiri. Neuromorphic vision based contact-level classification in robotic grasping applications. *Sensors*, 2020.
- [53] Rajkumar Muthusamy, Xiaodian Huang, Yahya Zweiri, Lakmal Seneviratne, and Dongming Gan. Neuromorphic event-based slip detection and suppression in robotic grasping and manipulation. *IEEE Access*, 2020.
- [54] Akihiko Yamaguchi and Christopher G Atkeson. Implementing tactile behaviors using fingervision. In *Humanoids*, 2017.
- [55] SMALLRIG 1/4" Screws. [https://www.amazon.de/gp/product/B01AZ558WE/ref=ppx\\_yo\\_dt\\_b\\_asin\\_title\\_o01\\_s00?ie=UTF8&psc=1](https://www.amazon.de/gp/product/B01AZ558WE/ref=ppx_yo_dt_b_asin_title_o01_s00?ie=UTF8&psc=1). [Accessed 28-08-2023].
- [56] LED Band 12V. [https://www.amazon.de/gp/product/B09WHY4Q33/ref=ppx\\_yo\\_dt\\_b\\_asin\\_title\\_o00\\_s00?ie=UTF8&th=1](https://www.amazon.de/gp/product/B09WHY4Q33/ref=ppx_yo_dt_b_asin_title_o00_s00?ie=UTF8&th=1). [Accessed 28-08-2023].
- [57] GelSight Mini Tracking Marker Replacement Gel. <https://www.gelsight.com/product/tracking-marker-replacement-gel/>. [Accessed 28-08-2023].
- [58] Micah K Johnson, Forrester Cole, Alvin Raj, and Edward H Adelson. Microgeometry capture using an elastomeric sensor. *ACM TOG*, 2011.
- [59] Inivation - dv-processing. <https://gitlab.com/inivation/dv/dv-processing>. [Accessed 28-08-2023].
- [60] Zhenjiang Ni, Sio-Hoi Ieng, Christoph Posch, Stéphane Régnier, and Ryad Benosman. Visual tracking using neuromorphic asynchronous event-based cameras. *Neural computation*, 2015.
- [61] Claude E Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 1949.
- [62] GelSight Robotics Software. <https://github.com/gelsightinc/gsrobotics>. [Accessed 28-08-2023].
- [63] Gunnar Farneback. Two-frame motion estimation based on polynomial expansion. In *Image Analysis: 13th Scandinavian Conference*, 2003.