

Grip Stabilization of Novel Objects using Slip Prediction

Filipe Veiga, Jan Peters and Tucker Hermans

Abstract—Controlling contact with arbitrary, unknown objects defines a fundamental problem for robotic grasping and in-hand manipulation. In real-world scenarios, where robots interact with a variety of objects, the sheer number of possible contact interactions prohibits acquisition of the necessary models for all objects of interest. As an alternative to traditional control approaches that require accurate models, predicting the onset of slip can enable controlling contact interactions without explicit model knowledge. In this article, we propose a grip stabilization approach for novel objects based on slip prediction. Using tactile information, such as applied pressure and fingertip deformation, our approach predicts the emergence of slip and modulates the contact forces accordingly. We formulate a supervised-learning problem to predict the future occurrence of slip from high-dimensional tactile information provided by a BioTac sensor. This slip mapping generalizes across objects, including objects absent during training. We evaluate how different input features, slip prediction time horizons, and available tactile information channels, impact prediction accuracy. By mounting the sensor on a PA-10 robotic arm, we show that employing prediction in a controller’s feedback loop yields an object grip stabilization controller that can successfully stabilize multiple, previously unknown objects by counteracting slip events.

Index Terms—Slip Prediction, Object Stabilization, Robot Manipulation, Tactile Sensing, Machine Learning.

1 INTRODUCTION

Grasping and in-hand manipulation remain challenging tasks in robotics due to a variety of issues. For grasping, it is necessary to infer finger positions on the object and manage the force distribution onto multiple fingers while ensuring grip stability in the presence of uncertainty [1]–[3]. Manipulation additionally deals with the contact dynamics between objects and fingers while executing desired motions [4]. For many of these issues, the key problem is how to adapt robot actions in order to deal with undesired contact changes. Controlling the contact state based on meaningful feedback signals may provide a solution to this problem that can potentially generalize across a variety of objects. Here, tactile feedback is an attractive option, as it provides high frequency information directly from the interaction points. Modern, deformable tactile sensors such as the BioTac [5], offer many different measured quantities (e.g., pressure, high frequency vibrations, and temperature fluctuations) while interacting safely due to their compliance. These rich measurements of the local interactions allow the robot to predict the effects of its actions and to adapt them in order to reach the desired contact state. For example, while object stabilization is classically achieved by applying grasps that maximize measures such as form or force closure, we could alternatively accomplish the same goal by minimizing the predicted slip during grasp execution. The classical approach relies on rather accurate contact models, while the slip control method can be based only upon sensory input and prior experience.

Slip, i.e., the partial loss of contact between finger and held object [4], is known to be a key element of human manipula-



Fig. 1. A human-robot grip stabilization experiment where a human and a robot collaborate in order to preserve a stable grip on a deformable plastic cup. A detailed description of the experiment can be found in Section 3.7. Results show that the robot is able to respond to actions taken by the human in order to keep the object from slipping. The experimental results are discussed in Section 4.6.

tion [6] and may provide robots with the necessary feedback for maintaining grip stability during manipulation actions [7]. For example, such feedback can be used by a robot to reposition objects in its hand through controlled sliding [2]. In robotics, slip can be detected not only from tactile information [7]–[20], but also from vision [21], force-torque sensors [22], and laser-based range sensors [23]. Despite the extensive work, approaches based on the sense of touch either rely on large sensors [11], [17], are based on physical models of contact [8], [20], do not use slip information for control [7], [24], or do not evaluate the generalization capabilities of their approach [13], [25]. Further, the mere occurrence of slip is often associated with grasp or manipulation failure. Hence, endowing a robot with the ability to *not just measure slip, but to predict slip ahead of time*, allows it to react prior to the onset of slip. This differs substantially from approaches that focus on

Filipe Veiga and Jan Peters are with the Technische Universität Darmstadt, Darmstadt, Germany, FG Intelligent Autonomous Systems. e-mail: {veiga, peters}@ias.tu-darmstadt.de

Jan Peters is also with the Max-Planck-Institut für Intelligente Systeme, Tübingen, Germany.

Tucker Hermans is with the School of Computing at the University of Utah, Utah, USA. email: thermans@cs.utah.edu

sensing actively occurring slip [8]–[13], [19], [20] or the closely related problem of detecting the onset of slip, i.e., incipient slip [7], [14]–[18], [26]. Correlating the occurrence of slip with object stabilization failure, has the advantage of allowing a grip stability problem to be formalized as a slip avoidance problem. Unlike traditional approaches to grip stability analysis where the entire hand configuration is taken into account [27]–[31], avoiding slip takes a more local view, where the stability of the contact surface is assessed.

As a step towards robust in-hand manipulation, we focus on controlling the contact state using a tactile based prediction of slip. From the tactile information provided by the BioTac [5], we learn slip detectors and slip predictors from sampled data (Section 2.1). In contrast to other work where the Biotac was used [12], [18], we explore the sensors multiple information channels and autonomously extract the relevant information from each channel. This autonomous extraction of data allows for slip detection/prediction that is generalizable to previously unobserved objects. A similar approach is used in [13] only for detecting slip and with no analysis of generalization performance. Compared to physical modeling and analysis of slip, our approach does not require explicit knowledge on the friction properties or shape of the object.

We incorporate our learned slip detectors and predictors into a feedback controller to perform grip stabilization (Section 2) in two scenarios where objects are pinned by a single finger, either jointly with a human as illustrated in Figure 1 or against a table as illustrated in Figure 2. We compare slip detection to slip prediction and show that stabilizing controllers based on prediction achieve better stabilization performance (Section 3). The training is performed using common household objects, where the robot collects labeled data by either pinning the objects or sliding on their surface. We evaluate the generalization capabilities of our controllers by purposely leaving the test object data out of the object data set that is used for training the slip predictors.

This article extends our preliminary research results published in [32]. Here, we collect higher time-resolution data, where we autonomously label slip based on fingertip position and applied force. We examine a wider set of feature functions, analyze how each BioTac channel influences slip prediction accuracy and extend the robotic experiments to give further insights into the impact of slip prediction in stabilization control tasks and how well it generalizes to unknown objects.

2 GRIP STABILIZATION CONTROL THROUGH TACTILE-BASED SLIP PREDICTION

This section gives an overview of our approach to stabilize the grip on objects by predicting slip events. Our approach begins by extracting slip prediction signals from high dimensional tactile data, provided by the BioTac [5]. We formalize the slip prediction problem as a supervised classification problem in Section 2.1. Section 2.2 then describe the tactile data provided by the sensor and Section 2.3 details the tactile features we use as input for prediction. We give an overview of the classification methods used in Section 2.4, and describe the controller used for collecting the training data in Section 2.5.

Following our learning approach to predict slip, we describe a controller that takes advantage of the prediction signals in order to stabilize grips on multiple objects in Section 2.6. This includes objects that are not used in training the slip predictors. The

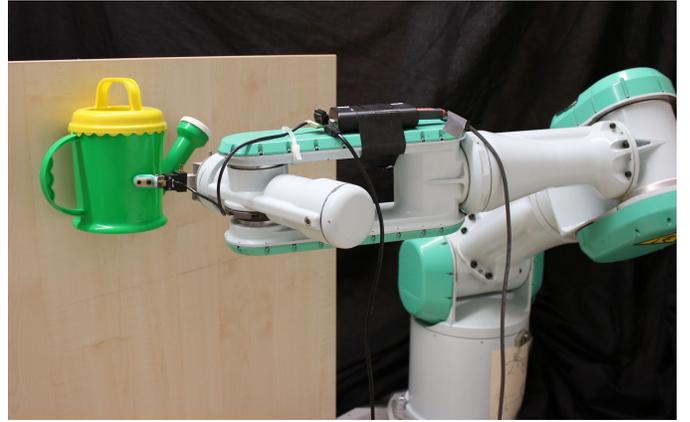


Fig. 2. The experimental setup used for our robot in the grip stabilization experiments. We use a Mitsubishi PA-10 robot arm with seven degrees of freedom and a BioTac tactile sensor mounted on the arm as a single finger end effector.

controller attempts to avoid slip events by increasing the contact normal force in response to predicted slip. This control concept has been previously applied in other controllers [11], [13], [17], [18], however, only slip detection was used without prediction of future slip events. Additionally these methods either used much larger tactile sensors or did not analyze the generalization properties of the controller across objects.

2.1 Learning to Predict Tactile Slip

In previous work [32], we showed that it is possible to detect and predict slip based on rich tactile signals provided by a highly deformable fingertip sensor. Here, we thoroughly analyze the generalization properties of the best performing approach in [32] while maintaining the problem formalization as a supervised learning problem. Our formalization involves learning a classifier, $f(\cdot)$, that labels the state at time $t + \tau_f$ as slip or not slip,

$$c_{t+\tau_f} = f(\phi(\mathbf{x}_{1:t})), \quad (1)$$

where $c_{t+\tau_f} \in \{\text{slip}, \neg\text{slip}\}$ is the state class at time $t + \tau_f$ with t being the current time. The prediction horizon $\tau_f \geq 0$ specifies the future time step in which the predictor is assessing slip. In the case where $\tau_f = 0$, $f(\cdot)$ becomes a slip detector. The feature function, $\phi(\cdot)$, applies a transformation on the raw sensor data, $\mathbf{x}_{1:t}$, providing the input for the classifier. We provide a detailed description of the feature functions explored in this work in Section 2.3.

Using this learning formulation, our approach is able to predict slip prior to its onset. As we integrate these signals into the feedback loop of the grip stabilization controller described in Section 2.6, we assess how the prediction window size, $\tau_f > 0$, impacts the outcome of the stabilization control, emphasizing the comparison between prediction and detection (i.e. $\tau_f = 0$).

2.2 Tactile Sensor Data

The raw tactile data is extracted from the BioTac [5], a multi-channel tactile sensor whose design was inspired by the human fingertip. The sensor is comprised of a rigid core, enveloped by a deformable skin. The space between the core and the skin houses fluid, contributing to the skin’s deformability whenever pressure is applied. Inside, 19 impedance-sensing electrodes, distributed

across the core's surface, measure the local skin deformation, while a pressure transducer measures fluid pressure and a set of heaters coupled with a thermistor manage and measure the fluid temperature. The sensor output is composed of the electrode signals $E \in \mathbb{R}^{19}$, low frequency, $P_{dc} \in \mathbb{R}$, and high frequency, $P_{ac} \in \mathbb{R}$, pressure measurements, temperature $T_{dc} \in \mathbb{R}$, and temperature flow $T_{ac} \in \mathbb{R}$. All channels are sampled at a rate of 100 Hz. The P_{ac} is acquired by the sensor at a rate of 2.2 kHz, but is still sampled at 100 Hz, producing batches of 22 values every 10 ms. The resulting sensor state $\mathbf{x}_t \in \mathbb{R}^{44}$, is given by

$$\mathbf{x}_t = [E^T, P_{ac}^{(1)}, \dots, P_{ac}^{(22)}, P_{dc}, T_{dc}, T_{ac}]^T. \quad (2)$$

An example of the raw signals produced by the sensor in one of our data collection trials can be seen in Figure 3.

2.3 Tactile-based Feature Functions

We explore several feature functions $\phi(\cdot)$, each of them representing distinct assumptions about the predictability of $c_{t+\tau_f}$ from the accumulation of tactile information over time. If we assume the class label to be directly observable from the current sensor reading \mathbf{x}_t ,

$$\phi(\mathbf{x}_{1:t}) = \mathbf{x}_t, \quad (3)$$

represents the single element feature function. Incorporating the change in the sensor values from the previous time step $\Delta\mathbf{x}_t = \mathbf{x}_t - \mathbf{x}_{t-1}$, yields the delta feature function

$$\phi(\mathbf{x}_{1:t}) = \begin{bmatrix} \mathbf{x}_t \\ \Delta\mathbf{x}_t \end{bmatrix}. \quad (4)$$

While the feature functions above only use values directly acquired from the sensor, we consider two complimentary functions that perform frequency analysis on the P_{ac} component, using the extracted properties as features. Based on one of the features proposed by Chu et al. [33], we start by calculating the energy spectral density (ESD), of the P_{ac} signal. The statistics computed on the *ESD* are the total energy Ω_s , spectral centroid C_s , variance σ_s , skewness S_s , and kurtosis K_s . We refer to the original work [33] for a more detailed explanation of the feature extraction method. These statistics provide an overview of the time series information given by the P_{ac} channel, no longer viewing the time series values as independent features. The statistics are concatenated with the remainder of the BioTac channels, $\hat{\mathbf{x}}_t \in \mathbb{R}^{27}$, to produce the complex single element feature

$$\phi(\mathbf{x}_{1:t}) = \hat{\mathbf{x}}_t = [E, P_{dc}, T_{dc}, T_{ac}, \Omega_s, C_s, \sigma_s, S_s, K_s]. \quad (5)$$

As for the single element feature, we also assess how previous time step information can be incorporated in the complex single element feature by setting

$$\phi(\mathbf{x}_{1:t}) = \begin{bmatrix} \hat{\mathbf{x}}_t \\ \Delta\hat{\mathbf{x}}_t \end{bmatrix}. \quad (6)$$

The resulting feature denotes a complex delta feature.

A second set of features is based on the assumption that $c_{t+\tau_f}$ depends on past sensor readings leading to the current time step t . Multiple feature functions are represented through

$$\phi_\eta(\mathbf{x}_{1:t}) = \mathbf{x}_{t-\eta:t} \quad (7)$$

where η controls the size of the window of past data to be considered in each feature function. All sensor readings over the last η time steps are accumulated into an input buffer that is then

evaluated by the slip predictor. These features are denoted as the time window η feature functions.

Finally, we also show results when using all features introduced by Chu et al. [33]. By comparing to features that were originally designed with the goal of object property learning, we showcase the importance of having the relevant information for predicting slip directly extracted from the raw data by the classifiers.

2.4 Classification Methods

To implement our slip predictors, we use random forest classifiers [34], support vector machines [35], and a spectral slip classifier adapted from [36] solely based on high frequency pressure information. In our previous work [32], we show that random forests generally outperform SVMs for predicting tactile slip. As such our analysis mostly focuses on prediction using random forest classifiers.

Random forest classifiers are ensembles of randomly trained binary decision tree classifiers [34]. Each decision tree classifies a given test example independently. The result of the entire forest is obtained by averaging over the distributions of the leaves reached in each of the trees. The class with the highest probability is then selected as the corresponding class for the current sample. Each decision tree is a binary tree where all non-terminal nodes have an associated splitting function, which decides if the currently evaluated example should traverse down the tree following the left or right branch. Leaf nodes contain a probability distribution over the class labels of training examples which reach this node. Tree training consists of selecting the feature and threshold to split at each node. These values are selected through the optimization of a specific performance criterion.

Support vector machines (SVM) are discriminate classifiers that separate the training samples by partitioning the feature space using a single decision boundary [35]. Each partition of the feature space defined by the decision boundary represents a single class. The decision boundary is chosen with respect to the closest samples of each class referred to as support vectors. During training the decision function which maximizes the classification margin, defined as the sum over the distances to each support vector, is found. The resulting linear classifier evaluating feature vector \mathbf{z} takes the form

$$f(\mathbf{z}) = \sum_{i=1}^k \alpha_i (\mathbf{z}^T \mathbf{z}_i) + b, \quad (8)$$

where α_i is the weight associated with the i^{th} support vector, \mathbf{z}_i , and b is a constant offset term. The support vectors and weights can be found efficiently by using quadratic programming. Both implementations of the previous classifiers come from the scikit-learn library for Python [38].

The spectral slip classifier, adapted from [36], computes the total energy in the P_{ac} channel at each time step, after bandpass filtering the output from 30 to 200 Hz. This frequency cutoff band is specifically tailored for the BioTac, as shown in [36]. If the signal energy in this specific frequency band exceeds a threshold Ω_{thresh} , the classifier signals a slip event. We choose Ω_{thresh} by optimizing the classifiers performance over the training data.

2.5 Surface Surveying Control for Slip Data Acquisition

In order to train our slip predictors, we require tactile data that has been labeled for slip classification. To collect such data, we

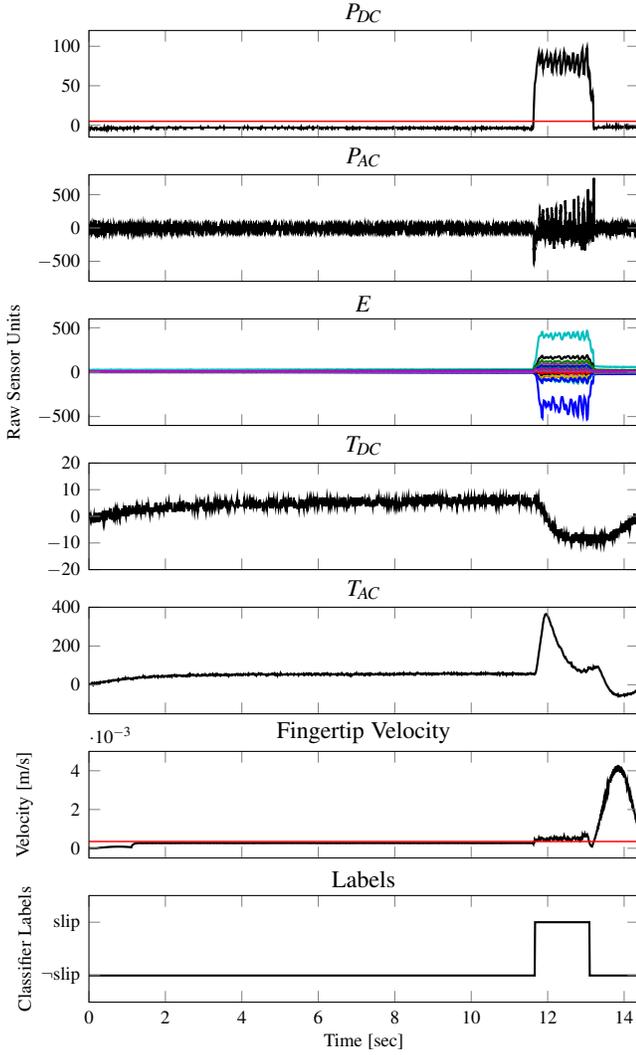


Fig. 3. Data traces for one of the data collection trials performed on the ball object. The data collection procedure is described in Section 3.2. From top to bottom, we show the low-frequency pressure variations, P_{dc} , high frequency pressure variations, P_{ac} , electrode responses in the fingertip, E , temperature, T_{dc} , temperature flow, T_{ac} , fingertip velocity, and ground truth labels. Sensor values are expressed in raw sensor units with no direct physical meaning. For a detailed description of each sensor channel output refer to [37]. The pressure and velocity thresholds used in the autonomous labeling procedure in Section 3.2 are shown in the respective plots.

perform exploratory actions along the surface of several objects. These actions are two-dimensional trajectories specified on the plane tangential to the contact point with the object and allow the robot to survey the objects surface in multiple directions, creating several different slip examples.

As the object surface is often not planar, a controller is used to ensure that the surveying trajectories are projected onto the object surface. This controller is a hybrid pressure-velocity controller that estimates the corrections that need to be applied to the predefined velocity trajectories, in order to keep the P_{dc} values of the sensor constant throughout the trajectory. Since the P_{dc} value is one dimensional, these corrections are projected onto the contact normal direction that is estimated using the electrode sensors of the BioTac, as proposed by Wettels et al. [39]. The estimate is obtained from the weighted average of the electrodes spatial normals, where the weights are the responses of each electrode.



Fig. 4. The objects comprising our data set. We selected objects covering a range of shapes and stiffness in order to adequately test classifier generalization. In the back we show a tape, followed by a row with, from left to right, a watering can, a box, a cup and a ball. In front we have a standard marker and behind it a measuring stick.

The magnitude of the resulting estimated vector is normalized, resulting in a unit vector in the direction of the applied contact force. To control the P_{dc} values, we calculate the pressure error

$$P_e = P_{dc}^D - P_{dc}, \quad (9)$$

where P_{dc}^D is the desired pressure and P_{dc} is the observed pressure. A PD controller is used to regulate the pressure error P_e by applying the necessary corrections to the predefined velocity trajectory

$$\mathbf{v} = \mathbf{v}_{des} + \mathbf{N}(K_p P_e + K_d \dot{P}_e), \quad (10)$$

where \mathbf{v}_{des} is the desired surveying velocity, \mathbf{N} is a unit vector representing the contact normal direction, K_p and K_d are the PD controller gains and \mathbf{v} is the applied task space velocity. The task space velocity \mathbf{v} is integrated in order to acquire the desired task space position and the respective desired joint positions using the robot's inverse kinematics. The surveying controller runs at 100 Hz which is the P_{dc} sampling frequency of the BioTac. The procedure in which the controller is used as well as the survey velocities and desired pressures are discussed in detail in Section 3.2.

2.6 Grip Stabilization Control using Slip Prediction

Taking advantage of the slip prediction, we design a highly reactive controller that avoids slip regardless of what object it is stabilizing. When slip is predicted to occur, the controller increases the desired task space velocity in the contact normal direction until the robot no longer predicts slip. By adjusting the desired velocity in the contact normal direction, the robot implicitly corrects the applied normal force. If we use the Coulomb friction model for contact, the implicit force adjustment ensures that the applied force remains within the friction cone of the contact location.

The stabilization control is triggered when the sensor touches the object since we can easily detect contact using thresholds on the sensor pressure values. The predictors run at 100 Hz classifying each sensor state as slip or -slip . If the state is labeled as slip, the controller, which also runs at 100 Hz, imposes a desired task space velocity $\mathbf{v}_N(t)$ in the contact normal direction. If the robot predicts no slip, the desired velocity is set to zero. The

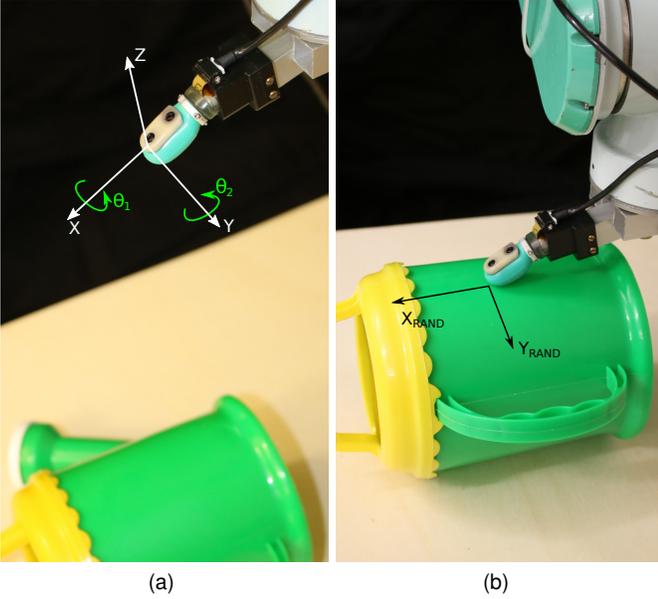


Fig. 5. The random components in the data collection procedure. In (a), two random rotations are applied on the initial position before the finger is lowered to establish contact with the object. In (b), the velocities along the two axes defining the tangential contact plane, are randomly chosen determining the surveying trajectory on the object surface.

imposed velocity $\mathbf{v}_N(t)$, corresponds to the contact normal $N(t)$ at time t weighted by a constant δ , i.e.,

$$\mathbf{v}(t+1) = \begin{cases} \delta N(t), & \text{if } c_t + \tau_f = \text{slip} \\ 0, & \text{otherwise} \end{cases}, \quad (11)$$

where δ is empirically defined based on a set of calibration trials where several objects of different weight were tested. The contact normal N is estimated using the method proposed by Wettels et al. [39], described in Section 2.5. The controller runs at 100 Hz, reacting to each prediction given by the classifiers.

It is important to notice that the controller does not keep track of the total applied normal force F_N , as it simply increases the velocity in the contact normal direction. As we show, even such a simple controller can perform well when using rich feedback, such as that provided by our slip classifiers.

3 EXPERIMENTAL PROCEDURE

In this section, we describe the experimental procedures necessary to realize and evaluate our slip prediction based grip stabilization controller on a real robot. After a short description of our hardware setup in Section 3.1, we describe how the surface surveying controller from Section 2.5 is used to collect the necessary data for training the predictors. The procedure, as well as the number of data collection trials, desired surveying velocities and desired pressures are all discussed in Section 3.2. Once the data has been acquired, it needs to be labeled and partitioned into training and test sets. This process is described in Section 3.3. With the data labeled, we can finally train the slip predictors. Training is done using three different strategies described in Section 3.4, for evaluating the effectiveness of the predictions when a single object, all objects, or all objects excluding the test object are used.

As our data set is unbalanced in the number slip examples, versus the number of \neg slip examples, we evaluate the predictors with the F_{score} metric. In Section 3.5, we describe and motivate the use of this metric in our analysis.

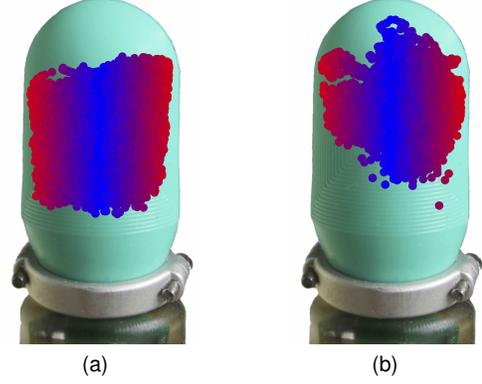


Fig. 6. Estimated contact locations on the fingertip sensor during (a) data collection and (b) the stabilization against a fixed plane. The points at the center of the BioTac are displayed in blue while peripheral contacts become progressively red as the distance to the center increases.

Finally, to evaluate the performance of the grip stabilization controller, we perform two sets of experiments. In the first set of experiments, described in Section 3.6, the robot stabilizes the gripped object between the finger and a fixed plane, evaluating the success rate of the stabilizations. The second set showcases the robustness of the stabilization by having the robot stabilize an object gripped jointly with a human experimenter. We describe this more dynamic experiment in Section 3.7.

3.1 Robotic Platform

All robotic experiments described in the following sections were performed using a Mitsubishi PA-10, a robotic arm with seven degrees of freedom. A BioTac tactile sensor, described in Section 2.2, was rigidly mounted to the arm as a single finger for manipulation. We directly control the arm's joint positions at a frequency of 100 Hz. The complete experimental setup can be seen in Figure 2.

3.2 Slip Data Collection Procedure

We acquired the data used to train our slip classifiers through two separate procedures specifically designed for collecting examples for slip states and static contact states.

In the slip procedure, the object was fixed to a table. The robot first moved to an initial position above the object. Two random rotations were applied to the initial position in each trial as shown in Figure 5a, ensuring different initial contact locations between the fingertip and the object across trials. Following the random rotations, the robot moved the fingertip down toward the object until achieving contact and a desired pressure P_{dc} was reached. Thereafter, a random velocity vector was defined in the plane tangential to the contact point as shown in Figure 5b. This velocity is denoted by \mathbf{v}_{des} in the surface surveying controller in Equation (10) and never exceeds 0.05 m/s. After performing the movement along the object's surface, the robot returned to its initial position. Each procedure was repeated ten times for each of the seven objects for three different target pressure values, making a total of 30 trials per object.

For the static contact data collection, the procedure was similar. The difference was that after the robot moved down and contact was established, there was no longer movement along the tangential contact plane. The robot effectively stays in static contact with the object for one second in each of the trials. This procedure was also repeated ten times for three target pressure

values resulting in another 30 trials for each of the seven objects. With both procedures combined, 60 data collection trials were performed per object for a total of 420 trials. The target values for P_{dc} used in the experiments were 20, 50, and 80. All sensor values were considered with respect to sensor baselines collected at a resting posture (no contact).

The random components introduced in the rotation of the initial position and the magnitude of the tangential surveying velocity, serve to cope with the multiple locations and velocities in which slip may occur on the fingertip. Figure 6 shows how finger contact locations collected during training compare to those produced during grip stabilization experiments performed on the robot. These stabilization experiments are described in detail in Sections 3.6 and 3.7.

3.3 Data Labeling and Partitioning for Training and Testing

Before the collected data can be used, it needs to be properly labeled. The labeling process was performed autonomously, relying on the robot's forward kinematics and the overall pressure on the fingertip. In a first stage, we removed data according to a contact threshold, P_{thresh} , on the P_{dc} values. If the finger pressure was below P_{thresh} , the finger was considered not in contact with the surrounding environment, and the corresponding time step was removed from the training data. The remaining data was then labeled as slip or static contact using a threshold, Δq_{thresh} , on the instantaneous end-effector velocity, Δq , estimated by calculating the difference in end-effector position between two consecutive time steps. If Δq was greater than Δq_{thresh} , the finger position was changing with respect to the object while contact was established and the data was labeled as slip. If Δq did not exceed the threshold the data was then labeled as static contact. Both thresholds P_{thresh} and Δq_{thresh} were tuned by an expert by observing the sensor data and the corresponding labels produced by each threshold pair. Increasing P_{thresh} resulted in more data points being removed due to being considered non-contact examples, while larger values of Δq_{thresh} implied that fewer contact examples were labeled as slip. We show the pressure values, instantaneous end-effector velocity, respective thresholds, and resulting labels in Figure 3 for one of the trials.

After the data was labeled, trials were partitioned into training and test sets. The training set was used for training the slip classifiers and was composed of seven of the trials for each object and P_{dc} value pairs. The remaining three trials per set were used to validate the classifiers after training. The same training and test sets were used for all experiments.

3.4 Training Strategies

We analyze slip detection and prediction according to the accuracy and generalization capabilities of our classifiers by introducing different training strategies. The first strategy ($S1$) involves training the classifiers on a single object and evaluating how they can classify slip on that same object. This strategy is denoted per object training and serves to evaluate how well slip can be classified on an object when only that object is known. Another training strategy, denoted all object training strategy ($S2$), involves using data from all objects during training and assessing the classification performance on each object individually. When making a comparison between these first two strategies, we can effectively assess how slip classification rates for a single object change when

multiple objects are known. If an increase in classification rate is observed for any single object when transitioning from $S1$ to $S2$, we can assume that relevant transfer of slip information is occurring between objects. Finally, in order to assess the generalization capabilities of the classifiers, a leave one out training strategy ($S3$) is used. In $S3$, classifiers are trained on data from all objects except one, that is used for testing the generalization to novel objects. If similar classification rates are achieved with $S2$ and $S3$, slip information transfer to between-object cases can be said to compensate for the absence of information regarding previously unobserved objects.

3.5 F-Score Metric

After performing the data collection and labeling procedures of Sections 3.2 and 3.3, we verified that the resulting data set was unbalanced in the number of class examples. Since the surveying motions only represent a small part of each data collection trial, the resulting data set is biased in the number of class samples, i.e., the slip examples are greatly outnumbered by the \neg slip examples.

Evaluating our classifiers based on classification accuracy, we observe very high accuracies (above 90%) from classifiers that simply label everything as \neg slip. Considering it is more important to detect when slip occurs than when it does not, a more helpful analysis of the classification results is performed by reporting the F_{score} instead of classification accuracy. The F_{score} is a harmonic mean of the precision and recall measures

$$F_{score} = 2 \frac{pr}{p+r}. \quad (12)$$

The precision, p , depicts the ratio between accurate positive classifications and total positive classifications

$$p = \frac{\text{true positives}}{\text{true positives} + \text{false positives}}. \quad (13)$$

In our case, precision evaluates the quality of the classifiers predictions by calculating the ratio of correct slip predictions with respect to the total number of slip predictions made by the classifier. The recall, r , is the ratio between accurate positive classifications and positive examples

$$r = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}}. \quad (14)$$

Here, recall represents how likely the classifiers are to miss slip instances by calculating the ratio between all instances where slip was predicted with respect to all instances where slip should have been predicted.

When the predictions are correct (high precision) and few slip instances are missed by the classifier (high recall), the F_{score} approaches its maximum value of one. If the classifier either predicts several instances of slip incorrectly (low precision) or misses several slip instances (low recall), the F_{score} approaches its minimum value of zero.

3.6 Stabilizing Objects against a Fixed Plane

Following the evaluation of our slip predictors, we describe the first experiment that showcases the relevance of the acquired slip prediction signals in the context of grip stabilization. The experiment involves stabilizing the grip on an object pinned between the fingertip sensor and a fixed table, as show in Figure 2. The grips are performed on unknown objects to test the generalization capabilities of our prediction based stabilizers.

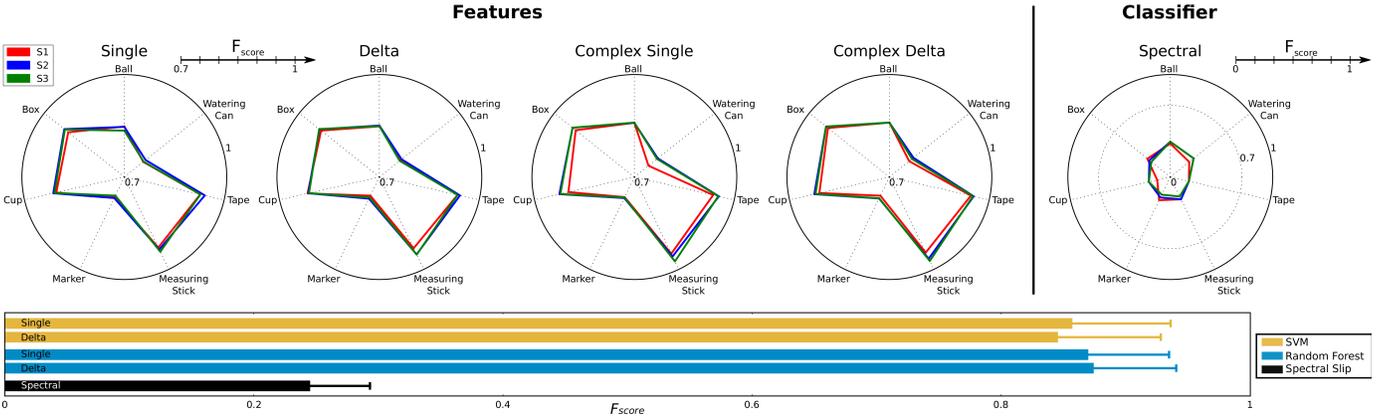


Fig. 7. Slip detection results showing (top) F_{score} achieved by random forest and spectral detectors for each feature and (bottom) mean and standard deviation across all objects for each of the classifiers. Plots for the individual features show results for each object, comparing across training strategies (single object $S1$, all objects $S2$ and leave one out $S3$) introduced in Section 3.4. The F_{score} for the proposed features shows similar patterns across all objects and all strategies. Note that the F_{score} minimum for the spectral slip radial plot is zero in contrast to 0.7 in the remaining plots. This is highlighted in the mean plots, where the random forests and spectral detectors respectively show the best and worst performances.

To perform grip stabilization, the previously trained random forest slip predictors are embedded in the feedback loop of the grip stabilization controller presented in Section 2.6. The random forests are trained with a leave one out strategy ($S3$), and a number of grip stabilization trials are performed on the object that was left out of the training set. Each trial consists of the robot initially pinning the object against a vertical plane. Once the object has been successfully stabilized against the plane, a random velocity is applied to move the robot away from the object. The three components of the exit velocity are sampled from different Gaussian distributions. The two lateral velocity components are sampled from Gaussian distributions with 0.0 m/s mean and 0.05 m/s standard deviation and the exit component is sampled from a Gaussian distribution with 0.05 m/s mean and 0.05 m/s standard deviation. As the robot moves away from the object, and as soon as slip is predicted, the stabilization controller becomes active, counteracting the exit motion and attempting to re-stabilize the object. The stabilization controller remains active until no slip is predicted during a period of 2 seconds or for a total of 10 seconds, after which the trial finishes. If the robot does not drop the object before the trial concludes, the trial is considered a success. We conduct ten trials per object for each feature function and prediction windows $\tau_f \in \{0, 5, 10, 15, 20\}$ and report the percentage of successful trials for each combination across all objects.

3.7 Human-Robot Joint Grip Stabilization

In order to assess the robustness of our slip prediction based grip stabilizers, we attempt to stabilize a grip on an object jointly with a human experimenter, by replacing the vertical plane with one of the experimenter's fingers.

Initially, the experimenter holds the object. As soon as the robot touches the object, the experimenter repositions his hand, and leaves only a single finger in contact with the object, as depicted in Figure 1. Simultaneously, the stabilization controller is activated. The experimenter makes a qualitative assessment of the robot's performance as it tries to compensate for his movements. This qualitative assessment is a relative preference between controllers, where each controller uses a different feature for predicting slip. The assessment is based on the responsiveness and stability (oscillations when attempting to keep the object still)

of each controller. Note that the grip stabilizers have to cope with the random noise introduced into the system by the experimenter. In addition, a thin plastic cup is used in this experiment. This object is not present in the previous object set, and is completely unknown to the slip predictors. The evaluation was repeated three times for each of the predictors and was performed by a single experimenter.

4 RESULTS

In this section we report the results of our evaluation for both the slip predictors and the grip stabilizers based on the predictors outputs. We begin with a brief comparison of results achieved with the different classification methods, Section 4.1, in order to confirm our previous findings [32], with the newly acquired data set. The classifier comparison is followed by an analysis of the individual performance of each BioTac channel for slip detection in Section 4.2, showcasing the importance of having the classifiers extract the relevant information from each of the BioTac channels. The performance of our slip detection and prediction classifiers as well as the relevance of the proposed features in both cases is reported in Sections 4.3 and 4.4 respectively. Finally, the success rates achieved for our grip stabilizers in the stabilization against a fixed plane experiment are reported in Section 4.5 and the results from the human-robot joint stabilization experiment are reported in Section 4.6.

4.1 Classifier Comparison

The average over $S1$ and $S2$ is used to compare the three classifiers described in Section 2.4. We examine the SVM and the random forests with two of the proposed features and use the P_{ac} signals as input to the spectral slip classifier. From the results shown in the bottom plot of Figure 7, we observe that SVM performance is on par with the performance achieved with random forests. The spectral slip classifier performance is much worse than the performance of the other two classifiers.

A more detailed comparison between the best and worst performing classifiers, respectively the random forest and the spectral slip classifiers, is shown on the radial plots of Figure 7, and in a comparison between the ground truth labels and the labels produced by both classifiers in Figure 8. The radial plots show similar F_{score} patterns for all training strategies and, in the case of

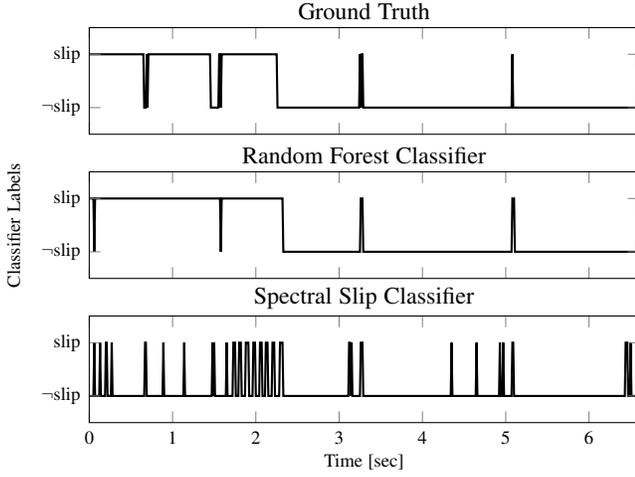


Fig. 8. Traces of the ground truth labels and the labels generated by the random forest and spectral slip detectors for one of the test trials. While the random forests very accurately reproduce the ground truth results, the spectral slip classifier is only able to detect the ground truth slip transitions, failing to detect continuous slip.

the random forests, across four of the proposed features. From the traces in Figure 8, it is clear that the spectral classifiers accurately detect the onset of slip, but fail to continuously label slip as the finger surveys the object’s surface. These findings suggest that spectral signals are prone to noise caused by motor induced vibrations during trajectory execution. Using all the BioTac channels, the impact of the noise is reduced, explaining the higher accuracy observed for the SVM and random forest classifiers throughout continuous slip phases. Following these observations, and considering that the random forests slightly outperform SVMs in terms of the mean value, the remaining results to be presented in this article are associated with classifiers trained using random forests.

4.2 BioTac Channel Relevance Analysis

In order to better understand the contribution of each BioTac channel for the detection of slip, we compare slip detectors that are separately trained with each channel as input using S_2 . The results are shown in Table 1. As a reference, results obtained by a slip detector trained with the single element feature are included in the right most column of Table 1.

The electrode information dominates the classification performance, achieving the best score on all objects except the marker. The pressure information from P_{ac} and P_{dc} are fairly

TABLE 1

F_{score} for different detectors where each of the BioTac channels is used as the only feature of the classifier. All results are for the random forest classifier trained on all objects (S_2). Bold values indicate the best F_{score} value obtained per object, excluding results obtained when using detectors trained using the single element feature function.

	E	P_{ac}	P_{dc}	T_{ac}	T_{dc}	\mathbf{x}_t
Ball	0.793	0.642	0.696	0.161	0.196	0.847
Box	0.876	0.683	0.702	0.099	0.184	0.925
Cup	0.809	0.614	0.596	0.171	0.072	0.913
Marker	0.718	0.559	0.802	0.249	0.003	0.768
Measuring Stick	0.905	0.701	0.665	0.099	0.087	0.936
Tape	0.886	0.617	0.706	0.088	0.053	0.942
Watering Can	0.667	0.637	0.634	0.172	0.402	0.780

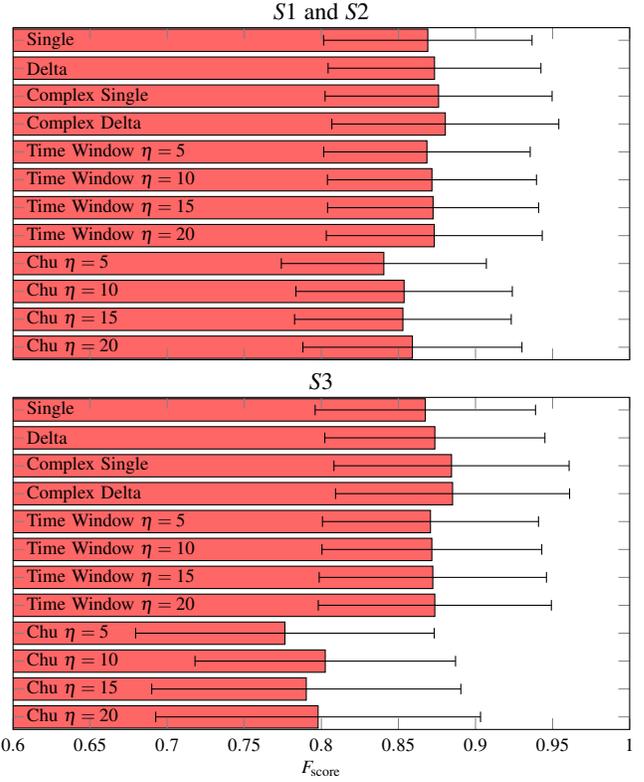


Fig. 9. The mean F_{score} and standard deviation obtained per feature for slip detection. These results show the average performance jointly over S_1 and S_2 (a) and when only considering S_3 (b). The memoryless and short term memory features outperform the long term memory features with respect to the mean F_{score} . When testing generalization with S_3 , a significant drop of performance is observed for the features of Chu et al. [33].

successful, with the P_{dc} alone achieving the best performance for the marker. Further, the temperature channels, T_{dc} and T_{ac} , provide no meaningful ability to detect slip. These results show that the deformation information present on the BioTac can be used to detect slip extremely well. This observation is especially relevant as deformation information is not traditionally associated with the detection of slip. Vibration information stored in the P_{ac} and P_{dc} offers good detection rates, while not having an average performance comparable to the electrodes. Finally, results achieved with the single element feature function show that the detectors are able to consolidate information from multiple channels, improving the slip detection rates.

4.3 Slip Detection and Feature Influence

From the results shown in the radial plots of Figure 7, we see that for four of the proposed features, the random forests can successfully classify slip, although different success rates are observed depending on the object. For objects such as the box, cup, tape and measuring stick, slip is classified quite accurately. On the remaining objects, despite lower classification rates than for previous objects, F_{score} values are still above 0.7. When comparing training strategies, we observe that classification performance remains mostly the same across all strategies, hinting that (1) not much knowledge is gained from additional objects once the target object is already known and (2) the classifiers can generalize quite well for previously unknown objects.

In Figure 9, we show results across all features when averaging over $S1$ and $S2$. While the time window η features perform quite well, there is no significant difference in performance with respect to the single element or the delta features. In fact the top performing feature is the complex delta feature. The complex versions of the single element and delta features seem to outperform their simpler counter parts. This suggests that condensing the time series information provided by the P_{ac} , as described in Section 2.3, benefits the detection rate. Finally, the features of Chu et al. [33], albeit being the worst performing features, still achieve good detection rates. This is surprising since they were designed with other tasks in mind. In fact these features cluster information over temporal windows which are large in comparison to the duration of slip events.

Figure 9 also shows average results over $S3$ for each feature. The same performance patterns are observed as in the previous experiment, with the complex delta feature still as the top performing feature. The features of Chu et al. [33] perform significantly worse than in earlier experiments. This suggests that the information retained by these features is specific to the objects observed and does not generalize well to novel objects. These findings support our choice of having the features built only under data aggregation assumptions, relying on the classification approach to balance the multi-channel information.

Considering the results shown so far and taking into account that we aim for real time performance on the real robot, we base the rest of our analysis on the single element and delta features and their complex counterparts. It is hard to justify the additional computation power required to process the time window η features in real time, as we do not observe any advantages in terms of classification performance when using these features.

4.4 Slip Prediction Offline Results

In this section, we analyze how the classifiers perform when $\tau_f > 0$, by training slip predictors with prediction horizons τ_f of 5, 10, 15, and 20 steps, equating to times of 0.005, 0.01, 0.015, and 0.02 seconds respectively. Following the same type of analysis as performed for slip detection, we observe that there is no significant difference between the results obtained for the average of $S1$ and $S2$ and the results obtained for $S3$, shown in Figure 10. We observe similar F_{score} values as with slip detection, clearly showing the feasibility of predicting slip. Nonetheless, a consistent drop in classification rates can be observed for $\tau_f = 15$. A more in-depth analysis reveals that, while prediction rates for most objects remain stable from $\tau_f = 10$ to $\tau_f = 15$, a decrease in rate can be observed for the ball and the watering can when using the single and complex single features. For the delta and complex delta features, a significant performance drop is observed only for the watering can. In terms of best performance, the delta feature overtakes the complex delta feature for $\tau_f \geq 10$. Further, complex features display a more accentuated drop in performance with increasing values of τ_f . This performance drop suggests that, despite having good discriminative and generalization properties, the complex features are not as suited for prediction due to their more compact representation of the P_{ac} signals that results in a loss of relevant information for prediction.

The results show that it is not only possible to predict slip but that it also possible to generalize the prediction of slip to previously unknown objects, producing the desired feedback signals for the controller described in Section 2.6.

4.5 Grip Stabilizing Control on the Real Robot

In order to assess the performance of our slip based grip stabilizers, Figure 11 reports the success rates for the stabilization of objects against a fixed plane experiment, described in Section 3.6. Results are show for each object separately, while comparing feature functions and τ_f values. By comparing different values of τ_f , we are assessing how earlier controller responses (larger prediction windows) affect the grip stabilization success rates. We observe an increase in the stabilization rates with τ_f for all objects, specifically for the cases where $\tau_f \geq 10$. Note that the increase in stabilization rate is particularly interesting considering the results shown in the previous section, where prediction accuracy dropped for $\tau_f \geq 10$. Although prediction accuracy is lower for this range of τ_f , the ability to predict farther into the future facilitates the stabilization task.

Spectral slip classifiers perform quite well independently of the value of τ_f and changes to τ_f do not seem to influence the performance of the spectral slip controller. Their performance is unexpectedly high, considering the results obtained in the offline evaluation. Careful observation of the stabilization trials confirms the results obtained in the offline experiments, where spectral slip predictors capture very accurately the transition from static contact to slip (transition from \neg slip to slip) but only as it is occurring. On the other hand, accurately detecting slip only as it occurs proves insufficient to stabilize all objects. For light objects, the brief response of the controller to the initial slip transition is enough to stabilize the object. This is the case for the ball, box, cup and marker. For objects such as the measuring stick or the watering can, the brief response generated by the controller during this first slip transition is insufficient to fully stabilize the object, as they are heavier and, in the case of the watering can, suffer from larger torsional slips. When comparing the predictors with the spectral slip classifiers, the latter outperform our approach for $\tau_f \leq 10$ but,

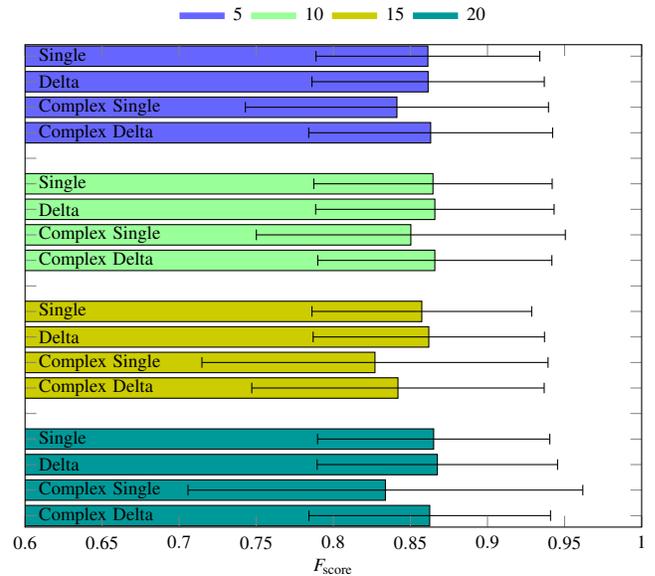


Fig. 10. The mean F_{score} and standard deviation obtained per feature for slip prediction with several prediction windows, τ_f . These results show the performance when only considering $S3$, hence testing how slip prediction generalizes to novel objects. While the complex features show better average results for low values of τ_f , the delta features show the top performance for $\tau_f > 10$, suggesting that the complex features are not as suited for prediction.

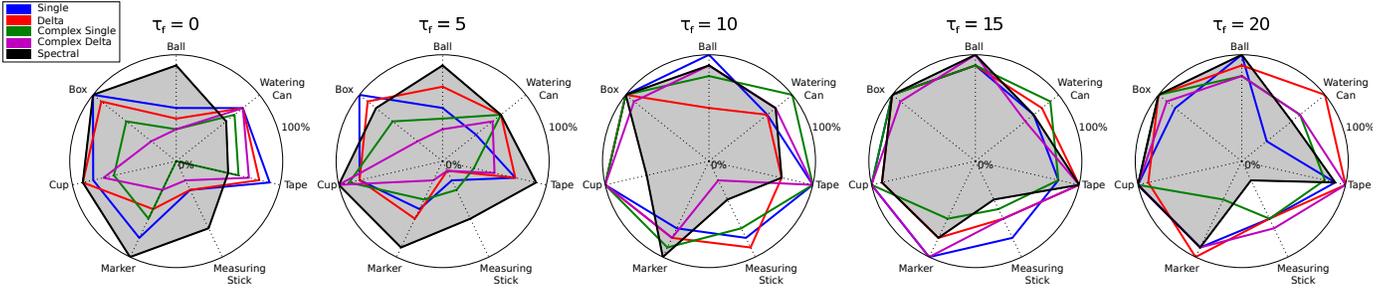


Fig. 11. Stabilization success rates of the slip prediction based controllers for each of the objects. The success rate represents the percentage of trials for which the robot successfully stabilized the object. All predictors are trained with a leave one out strategy S_3 . Different values of τ_f are considered results are shown for each feature. The filled region represents the performance of the controllers using the spectral slip classifier. It is clear that the performance of all proposed features increases with τ_f .

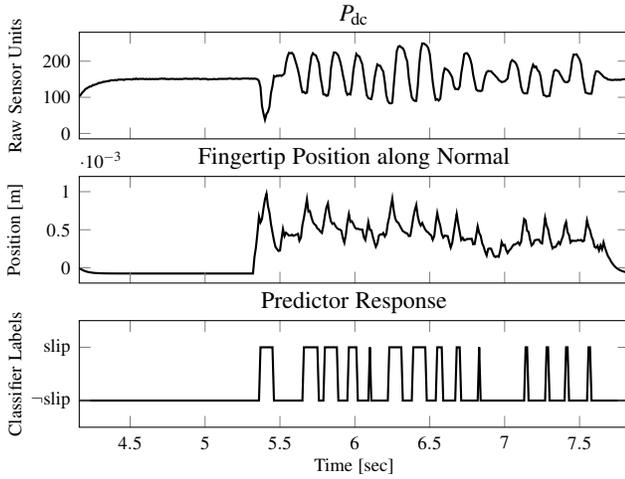


Fig. 12. Traces of the P_{dc} , the fingertip position along the contact normal and the predictor responses during a trial of the grip stabilization against a fixed plane experiment. The slip predictor used in this trial was trained with the delta features and a prediction window $\tau_f = 20$. After an initial perturbation, the grip stabilizers adjust the position of the fingertip, moving it towards the table (in the negative direction of the axis) whenever slip is predicted to occur.

as τ_f increases, controllers using the predictors eventually achieve the best stabilization rates for objects where previous performance was inferior to that of the spectral slip classifier.

Similar results can be observed in Figure 13, where the mean and standard deviation of the grip stabilization success rates across all objects are shown per feature and per value of τ_f . These results clearly show the consistent behavior of the controllers using spectral slip classification, unaffected by changes to the value of τ_f . For controllers relying on our proposed slip predictors, a clear pattern is observed, where stabilization performance increases with τ_f , eventually matching the performance of the spectral slip classifiers or even outperforming it for $\tau_f \geq 10$. In addition, it is also clear that despite the decrease in prediction rates observed for $\tau_f = 15$, there is an increase in stabilization success rate. The exception is the controller using the predictors trained with the complex single feature, which was the feature with the most significant drop in performance for $\tau_f = 15$. This confirms that, on average, the ability to predict slip further into the future has a bigger impact in controller performance than the resulting drop in prediction accuracy due to a larger value of τ_f .

We evaluate if there is a statistically significant difference between the success rates as a function of prediction time horizon. We conduct a separate test for each feature type. We perform

a Kruskal-Wallis H-test [40], a non-parametric version of the popular ANOVA test. The Kruskal-Wallis H-test is chosen since the variances of the distributions are not equal, a necessary assumption for the ANOVA test. The results show that there is no statistically significant differences between prediction horizons for the spectral and the single element features, with p-values of 0.9563 and 0.1386 respectively. On the other hand, for the delta, complex single, and complex delta features, there is a statistically significant differences between the distributions, with respective p-values of 0.04688, 0.0029 and 0.0091. For these three features, the results are not surprising, as they display the most accentuated increases in success rate.

Traces for the P_{dc} , fingertip position along the contact normal axis and predictor response are shown for a trial on the watering can using the delta feature with $\tau_f = 20$. After the initial perturbation occurring at 5.4s, whenever slip is predicted to occur, shifts in the negative contact normal direction and increases in the fingertip pressure can be observed. Following the 7.6s mark, the object remains stable for 2s and the trial concludes successfully.

4.6 Cooperative Grip Stabilization

By using controllers based on our slip predictors, the robot is able to successfully complete the human-robot joint grip stabilization task. While all controllers were able to jointly stabilize the objects, the experimenter found that as τ_f increased, the controller's response time was shorter, compensating for more sudden movements. The spectral slip method completely failed to stabilize the object with the human. The spectral slip's inability to classify continuous slip causes the controller to react only in short bursts as described in Sec. 4.5. For the movements produced by the human, where the duration and velocity of the movement is unknown, these extremely short responses are insufficient to keep the object stable.

To compensate for the accelerations introduced by the human, we replace a constant valued δ with a function $\delta(t)$ that increases linearly with time since slip was detected. The value of $\delta(t)$ returns to zero whenever $-slip$ is detected. For small movements, the robot does not apply the maximum response instantly, avoiding crushing the object or injuring the experimenter. As its response varies over time, the robot can cope with longer movements of variable velocity, by constantly readjusting its response.

5 CONCLUSION AND DISCUSSION

The proposed slip prediction based grip stabilization controllers were inspired by studies on human manipulation, specifically,

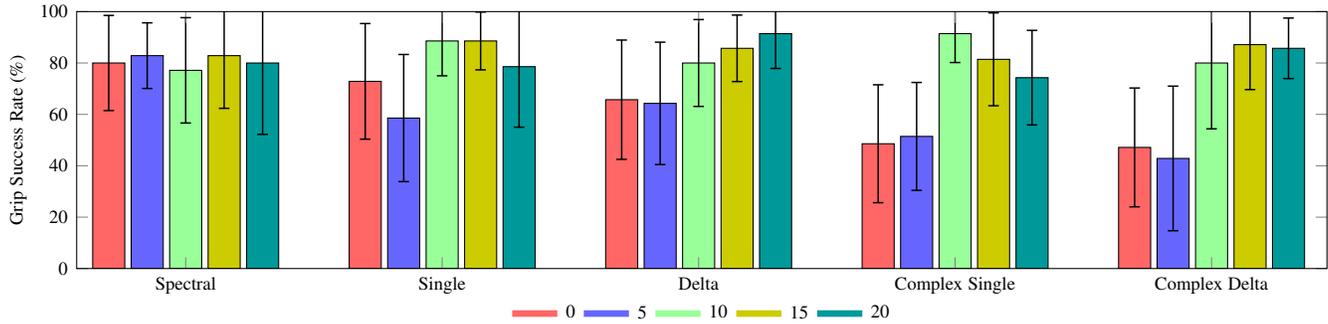


Fig. 13. Mean and standard deviation of the success rates for the grip stabilization experiments against a fixed plane experiments on the real robot. The success rate represents the percentage of trials where the robot successfully stabilizes the object, out of a total of 70 trials (10 per object). By varying the prediction window τ_f , we evaluate how the ability to predict slip farther into the future impacts the stabilization success rates of the controllers. While changes to τ_f have no effect on the controller using the spectral classifier, for all other controllers the success rates clearly increase for larger prediction windows.

neuroscientific studies suggesting that the human tactile system has a strong discrete feedback component, and relies on sensory prediction for control [6].

5.1 Summary of the Contribution

In this article, we have presented a learning based approach for predicting slip from high dimensional tactile information. Our slip predictors are integrated into the feedback loop of a grip stabilization controller allowing it to compensate for slip before its onset. Controllers based on predicted slip signals are shown to increase stabilization rates when compared to controllers solely relying on the detection of slip. In addition, we show the controllers to be highly generalizable to novel objects and sufficiently robust for the robot to stabilize objects jointly with a human. This robustness to severe perturbations, observed during the human-robot joint grip stabilization experiment, shows that such an approach can potentially be used for multi-fingered cases during in-hand manipulation.

5.2 Recognized Shortcomings

The proposed grip stabilization approach aims for the generalization of the slip sensation based control across a wide range of objects. To fulfill such a goal, a diverse training data set is required, in order to cover a broad set of interactions. Such data sets are not readily available, and have to be collected by experts in the field. Additionally, the data has to be labeled for the slip events, requiring highly accurate systems or several man hours to label the data manually. Our approach fails to compensate for rotational slip as the data collection procedures introduced in Section 3.2 were not designed to collect data for rotational events. Finally, training such methods with the required amounts of data is time consuming and sometimes renders them slow at execution time when compared to simple approaches such as the spectral slip classifier.

The controller proposed in Section 2.6 is highly affected by the heuristic used for the estimation of the normal contact direction. This estimation is very noisy and should be improved to solve more complex tasks.

5.3 Future Work

Our work has focused on merging multiple tactile sensing modalities for predicting slip and on the usefulness of these predictions

for grip stabilization with a single finger. A natural next step is the extension of this approach to multi-fingered manipulation tasks, where we wish to analyze how multiple fingers, each with its own tactile sensor, can be integrated in order to perform in-hand grip stabilization and possibly facilitate in-hand manipulation.

While we chose slip prediction as a first step, we see the possibility of using the proposed learning approach for the prediction of contact breaking (e.g., during a controlled release of an object) as in re-grasping or finger gaiting or for predicting the onset of a lift phase (i.e., when the held object leaves the supporting surface).

6 ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007/2013) under grant agreement 610967 (TACMAN).

REFERENCES

- [1] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis: A survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2014.
- [2] D. L. Brock, "Enhancing the dexterity of a robot hand using controlled slip," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1988, pp. 249–251.
- [3] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2000, pp. 348–353.
- [4] H. Yousef, M. Boukallel, and K. Althoefer, "Tactile sensing for dexterous in-hand manipulation in robotics: a review," *Sensors and Actuators A: Physical*, vol. 167, no. 2, pp. 171–187, 2011.
- [5] N. Wettels, J. A. Fishel, and G. E. Loeb, "Multimodal tactile sensor," in *The Human Hand as an Inspiration for Robot Hand Development*. Springer, 2014, pp. 405–429.
- [6] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," *Nature Reviews Neuroscience*, vol. 10, no. 5, pp. 345–59, May 2009.
- [7] M. Tremblay and M. Cutkosky, "Estimating friction using incipient slip sensing during a manipulation task," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1993, pp. 429–434 vol.1.
- [8] C. Melchiorri, "Slip detection and control using tactile and force sensors," *IEEE/ASME Transactions on Mechatronics*, vol. 5, no. 3, pp. 235–243, 2000.
- [9] K. Hosoda, Y. Tada, and M. Asada, "Anthropomorphic robotic soft fingertip with randomly distributed receptors," *Robotics and Autonomous Systems*, vol. 54, no. 2, pp. 104–109, 2006.
- [10] Y. Tada and K. Hosoda, "Acquisition of multi-modal expression of slip through pick-up experiences," *Advanced Robotics*, vol. 21, no. 5-6, pp. 601–617, 2007.

- [11] M. Shöpfer, C. Schürmann, M. Pardowitz, and H. Ritter, "Using a piezo-resistive tactile sensor for detection of incipient slippage," in *International Symposium on Robotics (ISR) and German Conference on Robotics (ROBOTIK)*, 2010.
- [12] B. Heyneman and M. Cutkosky, "Slip interface classification through tactile signal coherence," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, 2013, pp. 801–808.
- [13] J. Reinecke, A. Dietrich, and F. Schmidt, "Experimental Comparison of Slip Detection Strategies by Tactile Sensing with the BioTac on the DLR Hand Arm System," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2742–2748, 2014.
- [14] R. D. Howe and M. R. Cutkosky, "Sensing skin acceleration for slip and texture perception," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1989, pp. 145–150.
- [15] G. Canepa, R. Petrigliano, M. Campanella, and D. De Rossi, "Detection of incipient object slippage by skin-like sensing and neural network processing," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 28, no. 3, pp. 348–356, 1998.
- [16] I. Fujimoto, Y. Yamada, T. Morizono, Y. Umetani, and T. Maeno, "Development of artificial finger skin to detect incipient slip for realization of static friction sensation," in *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2003, pp. 15–20.
- [17] V. A. Ho, T. Nagatani, A. Noda, and S. Hirai, "What can be inferred from a tactile arrayed sensor in autonomous in-hand manipulation?" *2012 IEEE International Conference on Automation Science and Engineering (CASE)*, pp. 461–468, Aug. 2012.
- [18] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, "Force estimation and slip detection for grip control using a biomimetic tactile sensor," in *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, 2015, pp. 297–303.
- [19] S. Dong, W. Yuan, and E. Adelson, "Improved gelsight tactile sensor for measuring geometry and slip," *arXiv preprint arXiv:1708.00922*, 2017.
- [20] A. Ajoudani, E. Hocaoglu, A. Altobelli, M. Rossi, E. Battaglia, N. Tsagarakis, and A. Bicchi, "Reflex control of the pisa/iit softhand during object slippage," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1972–1979.
- [21] A. Ikeda, Y. Kurita, J. Ueda, Y. Matsumoto, and T. Ogasawara, "Grip force control for an elastic finger using vision-based incipient slip feedback," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, vol. 1, 2004, pp. 810–815.
- [22] F. Viña, Y. Bekiroglu, C. Smith, Y. Karayiannidis, and D. Kragic, "Predicting Slippage and Learning Manipulation Affordances through Gaussian Process Regression," in *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, 2013, pp. 462–468.
- [23] A. Maldonado, H. Alvarez, and M. Beetz, "Improving robot manipulation through fingertip perception," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, Oct. 2012, pp. 2947–2954.
- [24] N. Jamali and C. Sammut, "Slip prediction using hidden markov models: Multidimensional sensor data to symbolic temporal pattern learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 215–222.
- [25] X. Song, H. Liu, J. Bimbo, K. Althoefler, and L. D. Seneviratne, "A novel dynamic slip prediction and compensation approach based on haptic surface exploration," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*. IEEE, 2012, pp. 4511–4516.
- [26] X. A. Wu, D. L. Christensen, S. A. Suresh, H. Jiang, W. R. Roderick, and M. Cutkosky, "Incipient slip detection and recovery for controllable gecko-inspired adhesion," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 460–467, 2017.
- [27] Y. Bekiroglu, J. Laaksonen, J. A. Jørgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 616–629, 2011.
- [28] H. Dang and P. K. Allen, "Stable grasping under pose uncertainty using tactile feedback," *Autonomous Robots*, vol. 36, no. 4, pp. 309–330, 2014.
- [29] M. Madry, L. Bo, D. Kragic, and D. Fox, "ST-HMP: Unsupervised spatio-temporal feature learning for tactile data," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 2262–2269.
- [30] M. Li, Y. Bekiroglu, D. Kragic, and A. Billard, "Learning of grasp adaptation through experience and tactile sensing," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, 2014, pp. 3339–3346.
- [31] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. Kuchenbecker, "Human-Inspired Robotic Grasp Control With Tactile Sensing," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1067–1079, 2011.
- [32] F. Veiga, H. van Hoof, J. Peters, and T. Hermans, "Stabilizing novel objects by learning to predict tactile slip," in *IEEE/RSJ International*

Conference on Intelligent Robotics and Systems (IROS), 2015, pp. 5065–5072.

- [33] V. Chu, I. McMahon, L. Riano, C. G. McDonald, J. M. Perez-Tejada, M. Arrigo, N. Fitter, J. C. Nappo, T. Darrell, and K. J. Kuchenbecker, "Using robotic exploratory procedures to learn the meaning of haptic adjectives," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 3048–3055.
- [34] A. Criminisi, J. Shotton, and E. Konukoglu, "Decision forests for classification, regression, density estimation, manifold learning and semi-supervised learning," *Microsoft Research Cambridge, Tech. Rep. MSRTR-2011-114*, vol. 5, no. 6, 2011.
- [35] C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data mining and knowledge discovery*, vol. 2, no. 2, pp. 121–167, 1998.
- [36] J. A. Fishel, "Design and Use of a Biomimetic Tactile Microvibration Sensor with Human-Like Sensitivity and its Application in Texture Discrimination using Bayesian Exploration," Ph.D. dissertation, University of Southern California, 2012.
- [37] J. Fishel and G. Lin, "Biotac® product manual."
- [38] F. Pedregosa and G. Varoquaux, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [39] N. Wettels, A. R. Parnandi, J.-H. Moon, G. E. Loeb, and G. S. Sukhatme, "Grip control using biomimetic tactile sensing systems," *IEEE/ASME Transactions on Mechatronics*, vol. 14, no. 6, pp. 718–723, 2009.
- [40] W. H. Kruskal and W. A. Wallis, "Use of ranks in one-criterion variance analysis," *Journal of the American statistical Association*, vol. 47, no. 260, pp. 583–621, 1952.



Filipe Veiga is a PhD student at the Intelligent Autonomous System lab of Technische Universität Darmstadt under the supervision of Jan Peters. His PhD focuses on the development of biomimetic control approaches for robot grasping and manipulation using tactile information for improving the dexterous capability of robots. His work contributed for the success of the TACMAN Project. Before that, he received his M.Sc. in Electrical and Computer Engineering from Instituto Superior Técnico in Lisbon, Portugal. His master thesis focused on Robotic Grasp Optimization from Contact Force Analysis, and was developed as part of the HANDLE project.



Jan Peters is a full professor (W3) for Intelligent Autonomous Systems at the Computer Science Department of the Technische Universität Darmstadt and at the same time a senior research scientist and group leader at the Max-Planck Institute for Intelligent Systems, where he heads the interdepartmental Robot Learning Group. Jan Peters has received the Dick Volz Best 2007 US PhD Thesis Runner-Up Award, the 2012 Robotics: Science & Systems - Early Career Spotlight, the 2013 INNS Young Investigator Award, and the IEEE Robotics & Automation Societys 2013 Early Career Award. In 2015, he was awarded an ERC Starting Grant. Jan Peters has studied Computer Science, Electrical, Mechanical and Control Engineering at TU Munich and FernUni Hagen in Germany, at the National University of Singapore (NUS) and the University of Southern California (USC), and been an visiting researcher at the ATR Telecommunications Research Center in Japan. He has received four Masters degrees in these disciplines as well as a Computer Science PhD from USC.



Tucker Hermans is an assistant professor in the School of Computing at the University of Utah, where he is a founding member of the University of Utah Robotics Center. Previously he was a postdoctoral researcher in the Intelligent Autonomous Systems lab at TU Darmstadt in Darmstadt, Germany. He was at Georgia Tech from 2009 to 2014 in the School of Interactive Computing, where he earned his PhD in Robotics and an MSc in Computer Science. He received his AB in Computer Science and

German from Bowdoin College.