

Under Pressure: Learning to Detect Slip with Barometric Tactile Sensors

Abhinav Grover*, Christopher Grebe, Philippe Nadeau*, and Jonathan Kelly†

Abstract—Despite the utility of tactile information, tactile sensors have yet to be widely deployed in industrial robotics settings. Part of the challenge lies in identifying slip and other key events from the tactile data stream. In this paper, we present a learning-based method to detect slip using barometric tactile sensors. Although these sensors have a low resolution, they have many other desirable properties including high reliability and durability, a very slim profile, and a low cost. We are able to achieve slip detection accuracies of greater than 91% while being robust to the speed and direction of the slip motion. Further, we test our detector on two robot manipulation tasks involving common household objects and demonstrate successful generalization to real-world scenarios not seen during training. We show that barometric tactile sensing technology, combined with data-driven learning, is potentially suitable for complex manipulation tasks such as slip compensation.

I. INTRODUCTION

During grasping and manipulation, tactile signals provide vital information about slip—relative motion at the contact interface between the hand and an object—faster than any exteroceptive perception method. Slip can be disastrous (e.g., when transporting a fragile object) or advantageous (e.g., when moving an object without lifting it) depending on the context and the task [1]. In robotics, the well-studied “handover” task, which involves passing an object from a robot hand to a human hand, requires control of the gripping force with accuracy and speed to avoid significant slip [2]. The requisite feedback can only be provided through tactile sensing [3] and, consequently, the detection and control of slip events is fundamental to the completion of handovers and other relevant tasks [4].

Many new tactile sensors have become available over the last two decades (see [5] for a review). Each tactile sensor has inherent characteristics such as fragility, bulk, resolution, nonlinearity, hysteresis, and production cost. Consideration of these characteristics and how they relate to varying task constraints leads to the best choice of sensor for a specific job. For example, the BioTac fingertip [6] is equipped with impedance-based tactile receptors, hydrophones, and thermistors to provide rich multimodal tactile information, but is very expensive. Optical sensors, such as the GelSight [7], GelSlim [8], and DIGIT [9], infer tactile information from the visually-sensed deformation of the sensor contact surface.

All authors are with the Space & Terrestrial Autonomous Robotics Systems (STARS) Laboratory at the University of Toronto Institute for Aerospace Studies, Toronto, Ontario, Canada. <firstname>.<lastname>@robotics.utias.utoronto.ca

*Abhinav Grover and Philippe Nadeau were supported in part by the Vector Institute Scholarship in Artificial Intelligence.

†Jonathan Kelly is a Vector Institute Faculty Affiliate. This research was supported in part by the Canada Research Chairs program.




Fig. 1. We perform in-hand slip detection using barometric tactile sensors. We train a temporal convolutional neural network (TCN) to recognize slip events directly from tactile data. The tactile sensors are mounted on the fingertips of a Robotiq 3-finger gripper. Inset: TakkTile sensor (support frame, barometer circuit board, and rubber matrix, left to right).

Slip can be detected by monitoring changes to shear force at the surface, although vision-based sensors struggle in some situations due to a limited camera frame rate [10]. Optical sensors also tend to be relatively bulky, preventing them from being attached at positions other than the fingertips.

In contrast to other tactile-sensing devices, barometric sensors are compact and can be easily distributed over the fingertips and on the palm of a hand [11]. We investigate the potential of low-cost barometric tactile sensors, in combination with a state-of-the-art neural network, to perform slip detection. The sensors (the TakkTile model from RightHand Labs) are built from an array of commercial MEMS barometers fixed to a PCB backplane, with a thin rubber matrix forming the contact surface. These sensors have a very slim profile, are mechanically robust, exhibit a consistent linear response and no hysteresis, and are easy to integrate with existing end-effectors [12].¹ Each TakkTile sensor provides pressure and temperature data at a sampling rate of 100 Hz, which is higher than that of most vision-based tactile sensors. Our aim is to show that, despite having relatively low resolution, barometric tactile sensors are capable of recognizing important tactile events and should be considered as a valuable part of the tactile robotics ‘toolbox.’

The complex spatiotemporal signature of pressure changes during slip is difficult to model analytically. Instead, we take a data-driven approach by training a temporal convolutional neural network (TCN) to classify the time-series data produced by the tactile sensors as either static or slipping.

¹The MEMS TakkTile barometers (NXP MPL115A2) cost ~US\$2 each.




Fig. 2. TakkTile sensor [12] mounted on the fingertip of our Robotiq gripper and actuated automatically to slide over a 3D-printed spherical surface during training data collection.

Data-driven approaches are increasingly being used for grasp stability assessment and slip detection [8], [13]–[15]. Slip is often detected by analyzing the vibration pattern induced when an object moves within the gripper, where the specific frequencies involved depend on the composition of the surfaces in contact [16]. In our work, we propose a method that uses static pressure readings only to determine whether an object is slipping, relaxing the need (to some extent) for high-frequency sampling. We choose to use a TCN, rather than a recurrent neural network (RNN) (e.g., [10], [17]) or similar architectures, because recent work [18] suggests that TCNs can outperform RNNs for a diverse set of sequence modeling tasks. We make the following contributions:

- an algorithm for slip detection using barometric sensors that achieves an accuracy of over 90% on average;
- a comparison of our TCN approach with two other slip-detection methods that rely on vibration data;
- a preliminary analysis of the sensitivity and robustness of the TCN detector to factors related to surface properties and slipping motion;
- experimental results for in-hand slip detection involving objects with various surface and material properties.

II. SLIP DATA ACQUISITION

We desire a general model for slip detection, which necessitates training with diverse data. To this end, we ensure diversity by collecting slip measurements using a UR10 robotic arm equipped with the Robotiq three-finger gripper (shown in Figure 1); the TakkTile sensors replace the original fingertips of the gripper. Using this actuation setup, we are able to move a fingertip across fixed test surfaces to generate a variety of customized slip scenarios (see Figure 2). We determine the ground-truth ‘slip’ label by computing the planar velocity of the fingertip using the arm’s proprioceptive feedback. The barometer data and the velocity data are then time-synchronized and recorded for training.

For a given material, we identified three factors that influence tactile sensor feedback during slip events: surface curvature, slip speed, and slip direction. We chose to collect slip data over surfaces with three different types of curvature (spherical, cylindrical, and planer), at three speeds (0.05 m/s, 0.075 m/s, and 0.1 m/s), and in eight different directions. This range of motion covers a majority of practical slip scenarios. We 3D-printed a section of a cylinder and a sphere in ABS plastic. For a planer surface, we used the lid of a smooth plastic box, which is also used for the

robot experiments discussed in Section V. We programmed specific arm maneuvers to cover the range of parameters and automate the collection process. The complete dataset collected includes over 45 minutes of TakkTile samples at a rate of 100 Hz, with 143,584 data points belonging to the static class and 122,918 data points belonging to the slip class. The data in the slip class is evenly distributed across slip speed, slip direction, and surface curvature, and includes both translational and rotational motions.

III. LEARNING TO DETECT SLIP

We train a TCN to detect slip events. As shown in Figure 3, the input to the TCN is a set of sensor measurements that constitute the last T_k seconds of tactile data. For our 100 Hz tactile sensors, we capture the last one second of data, which makes the input an array of size 6×100 . We use a Keras implementation of the TCN to extract temporal features from the array signals. The TCN layers are similar to the ones described in [18], except that we use layer normalization instead of weight normalization and SELU activations instead of the ReLU activations. These architectural changes are motivated by our experimental tests, which indicated better performance on the tactile dataset.

The network is trained using the Adam optimizer [20] with a learning rate of 0.002 and a cross-entropy loss function. We use 20% dropout (for each layer including the fully connected layers), a mini-batch of size 256, and *He normal* kernel initialization for network regularization. Moreover, at the beginning of each epoch we equalize the class distribution of the entire training dataset through random undersampling to prevent classification bias. The dataset was split into training (90%) and validation (10%) subsets, and the network was trained for 800 epochs.

The MEMS barometers, which are the primary transduction component of the barometric tactile sensors, form a 2×3 array on each fingertip module. We exploit the axial symmetry of the array to augment our data in order to reduce overfitting. Before each epoch, we randomly apply one of three transformations: an x -axis flip, a y -axis flip, or a 180° rotation, to every data point with a probability of 25% (for each transformation). At the same time, we add a small amount of random gaussian noise to the network inputs with the intention of ensuring robustness to sensor noise.

IV. SLIP DETECTION PERFORMANCE

In this section, we compare the performance of our slip detector with two popular approaches from the literature. Although slip detection using barometric tactile sensors has not previously been attempted, there is prior work that employs other pressure transduction technologies. Holweg et al. [16] determine the power spectral density (PSD) of the vibrations induced in a piezoresistive rubber material by slip; the sum of the high-frequency PSD components is compared against a threshold to infer that slip has occurred. Meier et al. [15] train a CNN to classify slip events using data from an array of piezoresistive pressure sensors, where the CNN processes ‘frequency images’ from the array. For




Fig. 3. The TCN [18] architecture used for slip detection, where temporal features are extracted using dilated convolution layers (DCC) [19]. For each DCC, d represents the dilation and k represents the kernel size. The next two layers are fully connected, followed by a softmax output for classification.

comparison purposes, we implemented our own versions of the algorithms in [16] and [15] and tested them on our dataset. While both of the learning approaches (the frequency CNN and ours) demonstrate promising results, slip detection using the PSD threshold exhibits poor performance. The TCN outperformed the frequency CNN, with improvement of more than 5% on each metric shown in Table I.

Additionally, we evaluated the performance of the detector for possible combinations of slip type, slip speed, slip direction, and surface curvature, with the intention of characterizing the sensitivity of the detector. Table III summarizes the performance of the detector, where we use the F1-score statistic. Performance of the detector appears to be correlated with the number of barometers that are strongly stimulated by the surface at a given time. We note that performance improves as slip speed increases for almost all curvatures and slip directions, since higher speeds induce larger surface deformations and lead to more prominent temporal features within the one-second input time window. Performance is also affected by slip direction: slip along the oblique axes of the sensor is better detected than along the primary axes. The detector produces good results ($> 80\%$ F1-score) for rotational data as well, with the exception of performance on spherical surfaces.

V. ROBOT EXPERIMENTS

In addition to establishing the performance and sensitivity of the slip detector on test data, we determined its performance on two real-world manipulation tasks: slip detection of an in-hand object under an externally-applied impulsive force (the mallet tap test), and slip detection while lifting an object with insufficient grasping force (the object lift test). The test objects for the experiments (see Figure 4) were selected with the intention of varying the properties of the contact surfaces, such as curvature, roughness and deformability.

TABLE I

PERFORMANCE COMPARISON BETWEEN TCN AND FREQUENCY-BASED METHODS ON TEST DATA. THE WEIGHTED AVERAGE OF THE CLASSES IS USED TO COMPUTE EACH METRIC.

Method	Accuracy	Precision	Recall	F1-Score
PSD Thresh. [16]	57.4%	57.9%	57.4%	57.5%
Freq. CNN [15]	86.0%	86.0%	86.0%	86.0%
TCN (ours)	91.3%	91.4%	91.4%	91.4%

TABLE II
SLIP DETECTION RESULTS FOR REAL-WORLD EXPERIMENTS.

Object	Mallet Tap	Object Lift
Plastic Ball	90%	35%
Plastic Box	85%	100%
Cardboard Can	90%	85%
Football Sleeve	90%	85%
Foam Sleeve	80%	35%
Metal Can	90%	95%
Avg. Success	87.5%	72.5%

The goal of the mallet tap test was to evaluate the performance of the slip detector for objects held with a constant grasping force. For the test, the gripper was equipped with a single tactile sensor; a mallet was used to disturb the object and induce slip. The slip state was registered only if the detection network produced a slip label for two consecutive time steps. In a given trial, each object was tapped on the top and on its side; the trial was deemed successful if slip was registered when the object was tapped and the detector output returned to the nominal (static) value when the object stopped moving. For the object lift test, the gripper was programmed to hold the object with limited force while still maintaining contact with the surface. This was followed by an attempt to lift the object while increasing the gripping force if slip was registered. We added weights to the lighter objects (cardboard can, football sleeve, foam sleeve) to increase the likelihood of slip. A trial was considered successful if slip



Fig. 4. Objects used for our robot experiments, from left to right (objects, surface properties): *plastic box* (planer, smooth, rigid), *plastic ball* (spherical, smooth, rigid), *football sleeve* (spherical, rough, deformable), *foam sleeve* (cylindrical, smooth, deformable), *metal can* (cylindrical, rough, rigid), and *cardboard can* (cylindrical, smooth, rigid).

TABLE III

TCN PERFORMANCE, MEASURED BY F1-SCORE, WITH VARIATION IN SLIP TYPE, SLIP SPEED, SLIP DIRECTION, AND SURFACE CURVATURE.

	Max. Speed	Planer	Spherical	Cylindrical (y-axis aligned)	Cylindrical (x-axis aligned)	Curvature Independent
Translation (Primary Axes)	5 cm/s	92.5%	88.8%	78.7%	94.0%	89.7%
	7.5 cm/s	94.9%	88.0%	90.0%	90.3%	91.6%
	10 cm/s	95.7%	84.5%	87.8%	92.1%	93.1%
Translation (Oblique Axes)	5 cm/s	94.0%	88.9%	87.1%	93.2%	91.5%
	7.5 cm/s	96.6%	89.3%	93.7%	92.2%	93.1%
	10 cm/s	96.4%	88.6%	93.5%	93.4%	93.8%
Rotation	1 rad/s	84.6%	74.5%	84.0%	86.1%	81.5%
Motion-Independent		94.0%	88.9%	88.0%	91.9%	

was detected and compensated for, that is, if the object was stably grasped and lifted.

We conducted 20 trials for each object on each test; the experimental results are shown in Table II. For the mallet tap test, failure modes included a nearly-equal number of false negatives and false positives, which indicates an absence of detection bias. On the other hand, the object lift test was particularly challenging for the detector because the normal contact force remained low throughout the lifting phase, despite the added weight. As a general trend, the success rate was greatest for longer objects, which we observed to be the result of reaction time: there is more time to ‘catch’ longer objects. When lifting, two prominent failure modes were observed for specific objects: failure due to slow reaction time (plastic ball) and failures in the form of false negatives (foam sleeves).

We note that these experiments demonstrate an ability to generalize from single-material training to real-world, multi-material slip detection involving different objects, a key contribution of our work.

VI. CONCLUSION

We have presented a learning-based method for slip detection using a temporal convolution network in conjunction with inexpensive barometric tactile sensors. The network achieved a detection accuracy of greater than 90% on a diverse dataset. We compared our approach with an existing classical and a learning-based method, and found significant performance improvement. Our detector is sensitive to slip type (translation direction, rotation) and surface curvature while being relatively robust to slip speed and direction. Finally, we demonstrated the utility of our slip detection method on two real-world robotic manipulation tasks, showcasing generalization to unseen material types. We believe that the slip detection performance of barometric tactile sensors can be further improved by increasing the spatial density of the MEMS barometers at each grasp point.

REFERENCES

- [1] C. Wang, S. Wang, B. Romero, F. Veiga, and E. Adelson, “SwingBot: Learning Physical Features from In-hand Tactile Exploration for Dynamic Swing-up Manipulation,” in *IEEE/RSJ Intl. Conf. Intelligent Robots and Systems (IROS)*, 2020, pp. 5633–5640.
- [2] V. Ortenzi, A. Cosgun, T. Pardi, W. Chan, E. Croft, and D. Kulić, “Object Handovers: a Review for Robotics,” *arXiv preprint arXiv:2007.12952*, 2020.
- [3] W. P. Chan, C. A. Parker, H. M. Van der Loos, and E. A. Croft, “Grip forces and load forces in handovers: implications for designing human-robot handover controllers,” in *Proc. ACM/IEEE Intl. Conf. Human-Robot Interaction*, 2012, pp. 9–16.
- [4] R. A. Romeo and L. Zollo, “Methods and sensors for slip detection in robotics: a survey,” *IEEE Access*, April 2020.
- [5] M. R. Cutkosky and W. Provancher, “Force and tactile sensing,” in *Springer Handbook of Robotics*. Springer, 2016, pp. 717–736.
- [6] R. S. Johansson, G. E. Loeb, N. Wettels, D. Popovic, and V. J. Santos, “Biomimetic tactile sensor for control of grip,” Feb. 1 2011, US Patent 7,878,075.
- [7] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson, “Measurement of shear and slip with a GelSight tactile sensor,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2015, pp. 304–311.
- [8] S. Dong, D. Ma, E. Donlon, and A. Rodriguez, “Maintaining Grasps within Slipping Bounds by Monitoring Incipient Slip,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2019, pp. 3818–3824.
- [9] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, et al., “DIGIT: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [10] Y. Zhang, W. Yuan, Z. Kan, and M. Y. Wang, “Towards Learning to Detect and Predict Contact Events on Vision-based Tactile Sensors,” in *Conf. Robot Learning (CORL)*, vol. 100. PMLR, 2020, pp. 1395–1404.
- [11] R. Koirava, T. Schwank, G. Walck, M. Meier, R. Haschke, and H. Ritter, “Barometer-based Tactile Skin for Anthropomorphic Robot Hand,” in *IEEE/RSJ Intl. Conf. Intelligent Robots and Systems (IROS)*, 2020, pp. 9821–9826.
- [12] Y. Tenzer, L. P. Jentoft, and R. D. Howe, “The Feel of MEMS Barometers: Inexpensive and Easily Customized Tactile Array Sensors,” *IEEE Robot. Autom. Mag.*, vol. 21, no. 3, pp. 89–95, 2014.
- [13] E. Hyttinen, D. Kragic, and R. Detry, “Learning the tactile signatures of prototypical object parts for robust part-based grasping of novel objects,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2015, pp. 4927–4932.
- [14] D. Goeger, N. Ecker, and H. Woern, “Tactile sensor and algorithm to detect slip in robot grasping processes,” in *IEEE Intl. Conf. Robotics and Biomimetics*, 2009, pp. 1480–1485.
- [15] M. Meier, F. Patzelt, R. Haschke, and H. J. Ritter, “Tactile Convolutional Networks for Online Slip and Rotation Detection,” in *Intl. Conf. Artificial Neural Networks (ICANN)*, 2016, pp. 12–19.
- [16] E. Holweg, H. Hoeve, W. Jongkind, L. Marconi, C. Melchiorri, and C. Bonivento, “Slip detection by tactile sensors: algorithms and experimental results,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, vol. 4, 1996, pp. 3234–3239.
- [17] F. Veiga, J. Peters, and T. Hermans, “Grip Stabilization of Novel Objects Using Slip Prediction,” *IEEE Trans. Haptics*, vol. 11, no. 4, pp. 531–542, Oct 2018.
- [18] S. Bai, J. Z. Kolter, and V. Koltun, “An empirical evaluation of generic convolutional and recurrent networks for sequence modeling,” *arXiv preprint arXiv:1803.01271*, 2018.
- [19] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “WaveNet: A Generative Model for Raw Audio,” in *ISCA Speech Synthesis Workshop*, 2016, pp. 125–125.
- [20] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.