Going In Blind: Object Motion Classification using Distributed Tactile Sensing for Safe Reaching in Clutter

Rachel Thomasson[1], Etienne Roberge[2], Mark R. Cutkosky[1], and Jean-Philippe Roberge[2]

Abstract— **Robotic manipulators navigating cluttered shelves or cabinets may find it challenging to avoid contact with obstacles. Indeed, rearranging obstacles may be necessary to access a target. Rather than planning explicit motions that place obstacles into a desired pose, we suggest allowing incidental contacts to rearrange obstacles while monitoring contacts for safety. Bypassing object identification, we present a method for categorizing object motions from tactile data collected from incidental contacts with a capacitive tactile skin on an Allegro Hand. We formalize tactile cues associated with categories of object motion, demonstrating that they can determine with $> 90\%$ accuracy whether an object is movable and whether a contact is causing the object to slide stably (safe contact) or tip (unsafe).**

I. INTRODUCTION

When operating in human environments, robotic manipulators will encounter scenarios where obstacle avoidance is impractical. For instance, consider the task of retrieving an item from the back of a kitchen cabinet. It is likely that a collision-free path either does not exist or is intractable to visually identify due to occlusions. However, many of the obstacles in the scene will be movable and can therefore be pushed aside to create a path to the target. In situations such as these, humans are comfortable reaching towards the target and allowing, but not explicitly planning for, incidental contacts which rearrange the scene as necessary for task completion. They are able to sense and adapt to these contacts in order to maintain safety.

This work is aimed towards enabling a robotic manipulator to use incidental contacts during reaching for safe, passive object rearrangement. We do not focus on identifying the objects in contact and our goal is not to plan motions that push objects into a desired pose. Rather, we allow unplanned contacts to occur and aim to sense them and evaluate their safety. To minimize the need to reason about specific object properties, we propose that the safety of a contact can be inferred from the object motion it induces. The robot should

Fig. 1. When reaching through clutter, especially if the identity and mobility of objects are unknown, tactile sensors on surfaces such as the back of the hand can be used to prevent unsafe contact events, such as toppling an object.

be able to distinguish between contacts that stably reposition objects without knocking them out of their stable poses (e.g. sliding an object across the surface of a shelf) and those that may be destructive (e.g. toppling an object or continuing to push against an object that is fixed in the world). As cluttered scenes present a challenge for vision systems due to occlusions and light contact events may produce little visual change, we rely on distributed tactile sensing. Whereas most tactile sensors developed for robotics applications focus on the fingertips and nominal grasping surfaces of manipulators, incidental contacts may occur anywhere on the hand and indeed may occur most often on the back of the hand when reaching toward a target. Therefore, we present a tactile skin design that covers not only the inner grasping surfaces, but also the back and sides of the 16 degree of freedom Allegro Hand. We then explore techniques for locally classifying object motion at the contacts.

The contributions of this paper are:
- an approach for reaching in clutter that uses tactile data to identify broad classes of object behavior and tactile cues that do not require identifying objects or object classes to predict behavior;
- a high-coverage, flexible sensor network for the Allegro Hand and its use to collect data during incidental contacts;
- demonstration of tactile cues for distinguishing immovable and movable objects, and further differentiating sliding (safe) and tipping (unsafe) motion.

A. Related Work

Tactile Sensors and Sensory Skins: There has been extensive work on developing tactile sensors for robotic end-effectors, mainly aimed at providing information for grasping

Fig. 2. (**A**) The layers that comprise a networked sensor, shown for the back of the hand. (**B**) Sensorized Allegro Hand with electrodes exposed for clarity. (**C**) Signal change as force is applied to the back of hand sensor. Solid and dashed lines show the mean signal of 10 taxels during loading and unloading, respectively, and shaded regions show the standard deviation.

and manipulation. Yamaguchi et al. [1] provide a review of recent developments in tactile sensing technology for robotic manipulation. For the most part, these sensors cover nominal grasping surfaces, most commonly the fingertips [2]–[8], but occasionally also the proximal inner surfaces and palm (e.g. [9, 10]). A few sensors have been developed specifically for, or integrated with, the commercial Allegro Hand, a 16 degree of freedom (DOF) robotic gripper [10]–[12], but these do not cover the backside of the hand. Some tactile skins that can potentially cover large areas of robotic arms (e.g. [13]–[15]) have also been developed but are not specialized for a multi-fingered hand such as the Allegro Hand.

While optical tactile sensors offer remarkable resolution and sensitivity for manipulation [5]–[8], capacitive sensors are easier to distribute across the different surfaces of a robotic hand and offer adequate resolution. The sensor presented here is adapted from a design by Ruth et al. [16], and is placed on the front, back, and sides of an Allegro Hand (Fig. 2).

Reaching Through Obstacles: Previous studies have addressed challenges of manipulation in clutter by introducing a rearrangement subtask, where objects are repositioned using prehensile manipulations [17, 18], non-prehensile pushes [19]–[22], or a combination of both [23]. Zhong et al. [24] track contacts along multiple objects during rummaging with soft tactile sensors. Moll et al. [25] propose to avoid reasoning about specific interactions and note that, in many cases, controlling the final pose of every object in the scene is unnecessary. They present an open-loop method which allows pushing movable objects while avoiding collisions with fixed obstacles, assuming that the mobility of each object is known *a priori*. In this work, we use tactile sensing to detect object mobility, which is initially unknown and may change over time. Jain et al. [26] developed an MPC controller which enables a robotic manipulator with whole-arm tactile sensing to reach through obstacles while keeping contact forces below a threshold. We propose that using object motion as an indication of safety, rather than force, will allow sliding aside objects using forces which could topple objects in other scenarios.

Object Mobility Classification: The benefit of tactile sensing in detecting object motion within a grasp, most notably the onset of slip, has been well discussed in literature [1]. Recent work from Ma et al. [27] uses vision-based distributed

tactile sensing to estimate constraints to motion of a grasped object in contact with the environment. Using a novel sensor which provides both visual and touch information, Razaei et al. [28] predict the final resting pose of an object that contacts the sensor after being released from an unstable configuration. There has been comparatively little work using tactile data from incidental contact on a robotic manipulator to classify object motion. Bhattacharjee et al. [29, 30] use whole-arm tactile sensing to classify objects as Rigid-Fixed, Rigid-Movable, Soft-Fixed, and Soft-Movable through various data-driven approaches. Here, we expand on this work to differentiate between safe object motion (i.e. sliding) and potentially dangerous object motion (i.e. tipping) and propose a classification method which relies on designed tactile cues rather than a data-driven approach.

II. Distributed Tactile Sensor Design

We now consider the design requirements for sensory skins intended to provide information about incidental contacts, focusing on classifying object motions, and present a new distributed tactile sensor network for the Allegro Hand. The skin used in this work is a modified version of the fringe field mutual capacitance sensor presented in [16]. The elements comprising each sensor in the network are outlined in Fig. 2A. A flexible PCB contains 3x3 mm taxels with an interdigitated design arranged in a matrix. The electrodes are connected to the capacitance-to-digital converter of a Cypress PsoC 4200 microcontroller, which translates capacitance to signal counts. Pressure is transformed to a change in mutual capacitance through the addition of a textured dielectric which, when compressed, increases the effective dielectric constant as air is displaced. The textured dielectric is molded from urethane (SmoothOn Vytaflex-30) and features cylindrical pillars (0.8 mm diameter, 0.8 mm height). The dielectric is fixed to the electrode using a 1.6 mm thick adhesive-backed silicone encapsulation layer.

Contact localization is essential for determining how the robot should move in response to a contact. Especially when covering large areas, there is a trade-off between spatial resolution and sampling rate. In this work, we present experiments conducted with a back of hand sensor, which contains four taxels per cm^2 and has a sampling rate of 10 Hz. Figure 2C shows the sensor's response to loading and unloading up to 50 N as force is applied with a 12 mm

diameter probe. For a point contact (1.9 mm diameter probe), the sensor can detect forces as light as 0.06 N.

A. Design Considerations for Interaction in Cluttered Spaces

1) Coverage: Although contacts during grasping and manipulation may be constrained to specific areas of the hand, contacts during reaching often occur on the backside of the hand, especially if reaching with fingers slightly flexed. To avoid undetected contacts, the skin should ideally cover the entire hand. The coverage achieved by our sensor network is displayed in Fig. 2B.

2) Curvature: While flat tactile sensors are easier to manufacture, curvature is beneficial for capturing and tracking incidental contact events. As shown in Fig. 3, flat sensors provide large contact areas when aligned with a surface, but are prone to quickly losing contact when there is any misalignment, such as the onset of tip. Adding curvature to the sensor in the direction of gravity leads to contacts which travel down the surface of the sensor as an object tips.

We develop a 2D kinematic model for contact point motion as a curved sensor tips an object. This will later be used to determine appropriate tactile cues for tip events. For a sensor with radius of curvature R contacting a vertically uniform object of width w, one can determine the tip angle, θ, of the object by solving the vector loop equation set up in Fig. 3. We define two right-handed orthonormal bases: $\hat{n}_{xyz}$ and $\hat{b}_{xyz}$, fixed to the world and object, respectively. Unit vector $\hat{n}_x$ points along the ground and $\hat{n}_y$ points vertically upward and, initially, $\hat{b}_x = \hat{n}_x$ and $\hat{b}_y = \hat{n}_y$. The sensor contacts the object at height h and moves along $\hat{n}_x$ a distance of x, causing the object to tip by θ. We write a vector loop beginning and ending at P, the center of curvature of the sensor when it initially contacts the object, to obtain

$$x\hat{n}_x + R\hat{b}_x - y\hat{b}_y + w\hat{b}_x - (R+w)\hat{n}_x = \vec{0}. \quad (1)$$

Taking the dot product of eq. 1 with $\hat{n}_x$ and $\hat{n}_y$ produces two scalar equations which can be solved for θ. The distance we expect the contact to move in the gravity direction can then be calculated as $R\theta$. Smaller radii may require a larger tip angle before contact patch motion indicating tip is detectable but larger radii will lose contact with a tipping object faster. The sensor used in experiments features a gentle curvature with a radius of 372 mm.


Fig. 3. Curved surfaces are beneficial for sensing tipping. A schematic illustrates the vector loop in eq. 1.

III. TECHNICAL APPROACH

A. Environment Assumptions

We consider a robot performing a planar reaching task in the presence of obstacles. Robot trajectories lie in a horizontal plane and consist of piecewise linear motions. Operating speeds are low enough that a quasi-static assumption applies (i.e. there is no dynamic evolution of the scene that is not associated with robot motion). We assume contacts occur on a sensorized surface of the hand. Objects may vary in geometry, but we assume they are convex, rigid, and feature a stable support surface which is initially in contact with the surface of the world. We note that non-convex objects can often be treated as locally convex if considering only their sliding or tipping motion. The mobility of objects is initially unknown. They may be free to slide, constrained such that they are likely to tip, or immovable in certain directions. We consider any object which takes more than 20 N to move to be immovable.

B. Terminology

We use the following terms to formulate expectations about how tactile signals evolve in response to various object behaviors: **Tactile Feature:** a quantity extracted from tactile data at an instant. Examples include shape, size, location, and intensity of contact patches. **Tactile Cue:** a measure of how tactile features evolve over time – for instance, an increase in contact patch size. **Tactile Event:** the tactile response corresponding to an object behavior (tipping, sliding, or fixed contact events). A tactile event may correspond to one or more tactile cues.

C. Object Motion Categories

In general, object motion from incidental contact is complex and depends on the object's geometric and mass properties. For determining the safety of a contact interaction, it may suffice to broadly understand the object's motion category rather than estimating its trajectory. We consider the following categories:

- **Immovable** in the direction of applied force either because it is fixed in the environment, constrained, or considered too heavy to push.
- **Sliding** across the working surface while remaining in its stable pose.
- **Tipping** out of its stable pose. Excessive tip may lead to toppling, potentially causing damage or spills.

We hypothesize that these broad motion categories can be inferred from tactile cues that are not highly object-dependent. Therefore, using tactile signals to categorize object behavior could allow robots to reason about the safety of their actions without needing to identify the object itself. There also exist objects without stable poses (e.g. spheres), which we do not explicitly consider here. In addition to sliding, these objects may also roll, which would produce similar tactile signals to sliding and is also considered locally safe.

Prior to choosing potential cues which indicate motion category, we note that radial symmetry of an object impacts
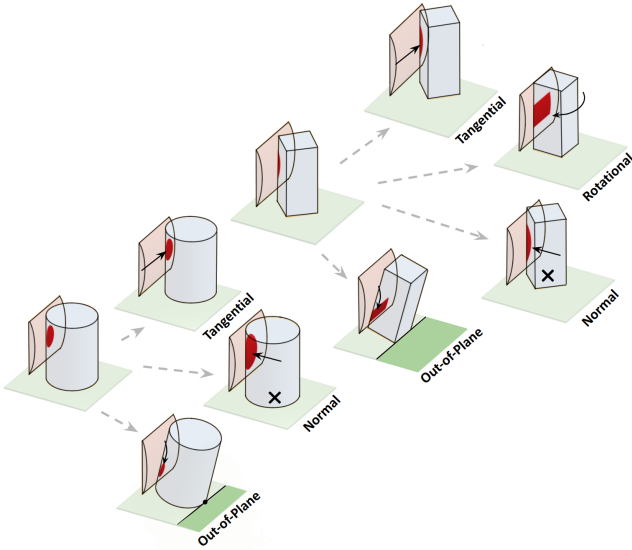
Fig. 4. We can distinguish two planar DoFs for radially symmetric objects (e.g. cylinders) and three for radially asymmetric objects (e.g. boxes). Out of plane motion (tip) provides another distinguishable DoF.

the tactile signals we expect. In their stable poses, objects that are radially symmetric about the vector normal to the floor have two detectable degrees of freedom: motion in the directions normal and tangential to the sensor surface. Radially asymmetric objects have three detectable degrees of freedom in their stable poses, as rotation in the plane will affect the sensor signal, as illustrated in Fig. 4.

D. Tactile Feature Extraction

Our candidate tactile features are: (1) [x,y] location of the contact patch center of pressure, (2) contact patch area, and (3) intensity of the activated taxels within the contact patch. They are identified using the following processing steps. First, we represent the normalized data at each timestep as a greyscale image. Signal intensity is extracted at this step. We then use bicubic interpolation to obtain sub-taxel resolution. Next, thresholding is used to isolate the contact patch. The same threshold value was used across experiments. Image moments are used to calculate the center of pressure of the contact patch, (C_x, C_y), using $C_x = \frac{M_{10}}{M_{00}}$ and $C_y = \frac{M_{01}}{M_{00}}$ where M_{ij} is the $(i,j)^{th}$ image moment. Finally, a contour which surrounds the patch is found. To accommodate cases where contact with a single object results in multiple disjoint patches (due to sensor noise or light, uneven contact), we take the convex hull of all contours to combine them into a single patch before calculating area as illustrated in Fig. 5A.

E. Tactile Cue Extraction

Candidate tactile cues such as an increase in signal intensity or a contact patch motion may be straightforward to identify conceptually, but reliably deriving them from raw tactile data is more involved. While theoretically, one could consider a tactile cue to be a discrete change in a tactile feature between two points in time, we observe that it is more reliable to identify trends over sequences of tactile

data. Indeed, due to real world factors such as sensor noise, uneven contact surfaces, and variations in friction between the objects and the table, discrete changes in the contact patch often instantaneously resemble tactile events that are not actually present. We perform a piecewise linear least squares fit on the evolution of extracted tactile features either over time or over the progression of a robot motion [31]. Two line segments are used to capture a change in object motion partway through the interaction (e.g. objects may initially slide then begin to tip). For each tactile feature, we consider the slopes and means over each line segment to be candidate tactile cues. Before performing the linear regression, the features are passed through the RANSAC algorithm to remove outliers caused by noise.

IV. EXPERIMENTS

We now describe the process for collecting incidental contact data and discuss the selection process for tactile cues. We then present classification results for contacts with immovable, tipping, and sliding objects.

A. Incidental Contact Data Collection

1) Experimental Setup: Incidental contact data were collected using the sensor presented in Sec. II on the back of the Allegro Hand, which was mounted on a UR16e robotic arm and stationed above a flat table. The hand was placed in a configuration suitable for reaching, with the fingers extended and the palm orthogonal to the table surface, as in Fig. 1. Future work will apply the approach demonstrated here on the back of the hand to the entire hand sensor network.

2) Object Sets: We collected data for contact with two sets of objects, shown in Fig. 6. The first set contains prototypical objects which are approximately rigid and geometrically simple (2 boxes and 3 cylinders). We expect motion categories for these objects to be separable using the tactile cues discussed in Sec. III-E. The second set consists of eight additional commonplace kitchen objects. These objects
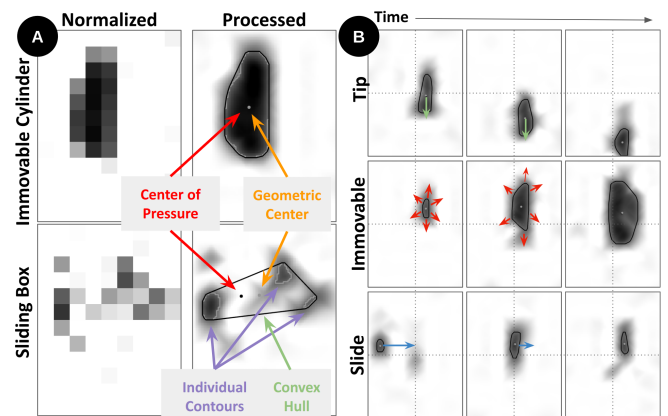

Fig. 5. (A) Snapshot tactile images taken from contact with a fixed cylinder and sliding box are shown at two stages of post-processing. Features extracted from the processed images are labeled. (B) Sequences of tactile images in which tactile events can be identified.

Fig. 6. Two sets of objects used for experiments. Object dimensions and weights are listed in Table I.

Object	Dimensions [mm]	Weight [g]
(1) Glass Vase	90 x 190	390
(2) Cardboard Cylinder	61 x 270	170
(3) Salt Shaker	68 x 208	910
(4) Wooden Box	100 x 165 x 230	420
(5) Cardboard Box	43 x 133 x 180	260
(6) Pancake Mix	60 x 155 x 208	960
(7) Mayonnaise	70 x 90 x 180	610
(8) Coconut Water	74 x 80 x 230	1060
(9) Lemonade	80 x 110 x 270	2000
(10) Sunflower Oil	76 x 76 x 240	950
(11) Balsamic Vinaigrette	45 x 103 x 204	300
(12) Syrup	70 x 106 x 255	900
(13) Coffee Grounds	90 x 140 x 180	620

TABLE I

OBJECT PROPERTIES

display a wider range of geometries than the prototypical object set which led to more varied object motions in response to contact. While our assumption was that objects are rigid, many of these are slightly deformable. One object featured concavities (lemonade) and three featured support surfaces that, while visually appearing flat, rested unevenly on the table leading to wobbly motion during sliding (pancake mix, coconut water, coffee beans). The purpose of the second object set is to investigate the robustness of our classification method to different object geometries and to identify and discuss failure cases.

3) Data Collection: Data were collected for objects individually. For each trial, an object was placed on the table in one of three conditions: (1) free-standing without constraints, (2) placed along the edge of a flat obstacle (a placemat) which promoted tipping, or (3) constrained against a wall. The arm was then commanded to move along pre-specified Cartesian paths which extended beyond the object, leading to a collision with the back of hand sensor. For each object, 20 trials were taken for each mobility condition – half for which the tangential component of motion was zero and half for which it was equal to the normal component. Ground truth class labels were assigned by hand. The robot was commanded to stop if the force in the direction normal to the sensor surface exceeded 30 N. This was monitored with the integrated F/T sensor at the UR16e wrist, but F/T data were

not used for object motion classification. Between trials, the initial pose of the objects was manually shifted such that the object and sensor came into contact at different locations. To limit inertial effects, we command the robot to move at a speed of 10 mm/s for all trials.

B. Tactile Cue Selection

Preliminary incidental contact data were used to identify which cues would be used for classification and to establish appropriate threshold values. To test the ability of eq. 1 to predict tipping based on contact patch motion, we collected empirical results for motion of an object with known width being tipped by contact at three heights. Center of pressure locations calculated using the feature extraction method described in Sec. III-D are plotted against analytical results in Fig. 7. To accommodate for a range of object widths, curvatures, and collision heights, we estimate that an object is tipping if the rate at which the contact patch moves in the direction of gravity as the robot moves into the object is between 1 and 5 (slopes shown in Fig. 7).

To identify contact with an immovable object, we hypothesized that signal intensity and contact patch growth could both serve as cues. However, it was observed from preliminary data that contact patch growth was highly variable between objects and trials, largely due to orientation changes of radially asymmetric objects. On the other hand, signal intensity, calculated from normalized tactile images as the mean value of taxels at the center of the contact patch, was found to be a good separator of contact with immovable and movable objects. As the signal for a mutual capacitance pressure sensor decreases with increasing pressure, we identify fixed contact if the average signal intensity (over one segment of the piecewise linear fit for signal intensity versus time data) is less than 65. Figure 8 illustrates the use of these cues for object motion classification, depicting the thresholds defined and examples of real data for each category.

C. Classification Results & Discussion

Figure 9 shows classification results for incidental contact with immovable, tipping, and sliding objects. For objects in the prototypical set, simple and object-independent cues
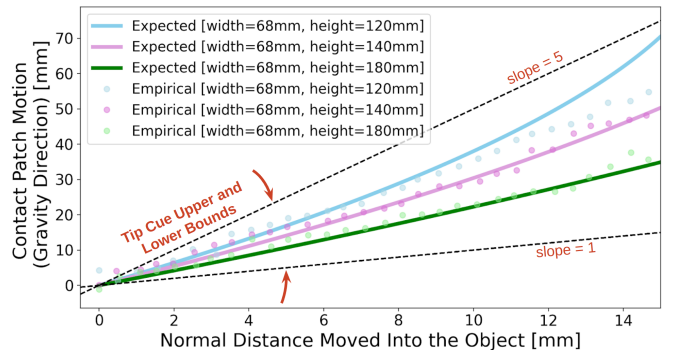

Fig. 7. Empirical tip data plotted over expected contact patch motion during tip at three heights, calculated as $R\theta$ where θ is found by solving eq. 1.
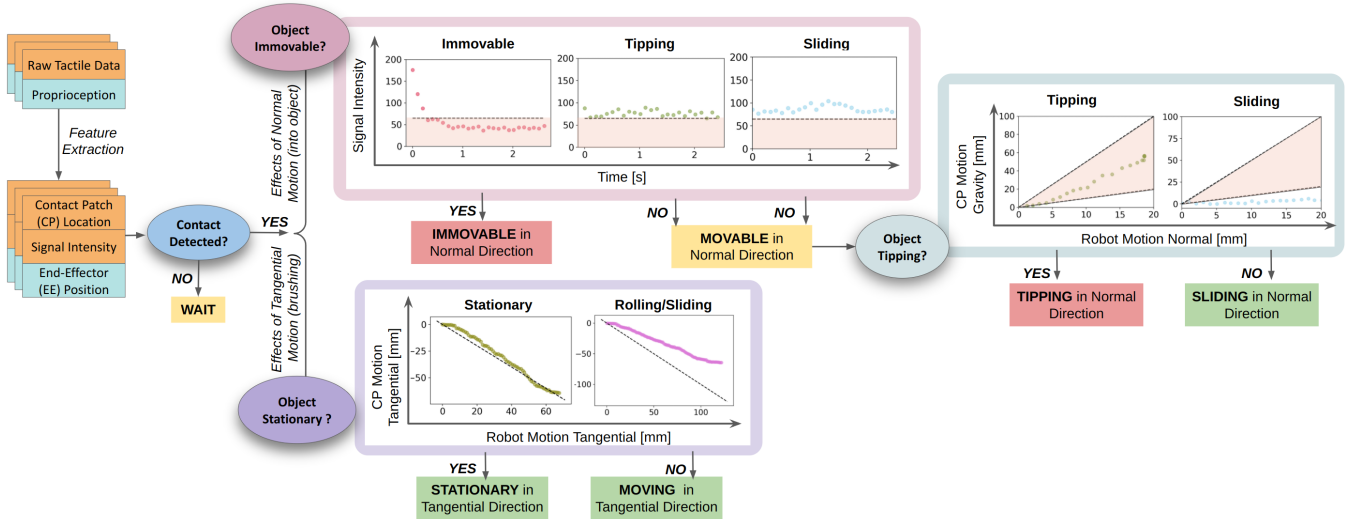
Fig. 8. A flowchart summarizing the steps for determining object motion category from tactile and proprioception data. Plots show real experimental data along with the conditions that signals must meet for tactile events to be predicted.

achieve 98% accuracy in separating motion classes. Though the misclassifications are few, we note an interesting failure case for box-like objects. In a few cases where a box was placed such that it was likely to tip about a corner rather than a flat edge, contact caused it to tip in a direction orthogonal to the normal force, which was undetectable by the sensor.

Results for the more geometrically diverse real-world objects showed a combined accuracy of 92%, where most confusion comes from mistaking tipping for sliding motion. 14 of the 18 misclassifications of tipping as sliding arose from two objects: coffee beans and lemonade. Deformable and granular coffee beans interacted with the sensor differently than rigid objects during tip events. During tip, the sensor would deform the object, causing a contact patch that often appeared nearly stationary rather than moving in the direction of gravity. When contact causing tip of the lemonade container occurred on areas featuring concavities, contacts did not travel smoothly down the length of the object, as we would expect with a convex object. Rather, the portion of the container that protruded the most would remain in contact with the sensor, causing the contact patch

to travel upwards along the sensor rather than downwards. For difficult cases such as these, where rigidity and convexity assumptions do not hold, it may be necessary to identify the object prior to estimating motion. A complementary video is provided to demonstrate the use of this classification method.

D. Analyzing Tangential Motion

While the focus of this work is classifying object motion induced by applying normal force, we note that useful information can be extracted from tangential components of the data as well. By comparing the horizontal contact patch motion to the robot's tangential motion, one can estimate whether an object is stationary or moving as a result of the robot brushing against it. If horizontal contact patch motion is equal and opposite to the robot's motion, this indicates a stationary object. Otherwise, the brushing contact is causing the object to roll or slide. Figure 8 shows this distinction for data collected as the robot moves tangentially while in contact with a cylindrical object placed against a wall (fixed in the normal direction) both when the object is stationary and when it exhibits a combination of rolling and sliding.

V. CONCLUSIONS AND FUTURE WORK

We present a method that uses incidental contact data from a tactile sensor on the back of a robotic hand to classify object motion. We show that our formulation of tactile cues is able to capture useful information about whether objects are movable or immovable and, if movable, whether contact results in safe sliding motion or potentially dangerous tipping. By using tactile data to understand the motions induced by incidental contact events, robotic manipulators reaching in clutter can rearrange objects without explicitly planning rearrangement motions or identifying the object in contact.

Future work involves applying this classification method to the full sensor network and developing reactive control methods which modify a robot's trajectory when a dangerous

Fig. 9. Confusion matrices containing the result of classifying incidental contact with immovable, tipping, and sliding objects. Numerical results indicate number of trials.

object motion is detected. A set of exploratory probing techniques could be designed to determine whether an object classified as immovable is completely fixed, or constrained only to move in certain directions. Additionally, methods for analyzing normal and tangential motion components – perhaps with shear as well as normal tactile data – could be combined to estimate objects' trajectories in more detail for higher level planning.

ACKNOWLEDGMENT

REFERENCES

[1] Akihiko Yamaguchi and Christopher G Atkeson. Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision? *Advanced Robotics*, 33(14):661–673, 2019.

[2] Nicholas Wettels, Jeremy A Fishel, and Gerald E Loeb. Multimodal tactile sensor. In *The Human Hand as an Inspiration for Robot Hand Development*, pages 405–429. Springer, 2014.

[3] Alexis Maslyczyk, Jean-Philippe Roberge, Vincent Duchaine, et al. A highly sensitive multimodal capacitive tactile sensor. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 407–412. IEEE, 2017.

[4] Heba Khamis, Raquel Izquierdo Albero, Matteo Salerno, Ahmad Shah Idil, Andrew Loizou, and Stephen J Redmond. Papillarray: An incipient slip sensor for dexterous robotic or prosthetic manipulation– design and prototype validation. *Sensors and Actuators A: Physical*, 270:195–204, 2018.

[5] Elliott Donlon, Siyuan Dong, Melody Liu, Jianhua Li, Edward Adelson, and Alberto Rodriguez. Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1927–1934. IEEE, 2018.

[6] Akhil Padmanabha, Frederik Ebert, Stephen Tian, Roberto Calandra, Chelsea Finn, and Sergey Levine. Omnitact: A multi-directional high-resolution touch sensor. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 618–624. IEEE, 2020.

[7] Daniel Fernandes Gomes, Zhonglin Lin, and Shan Luo. Geltip: A finger-shaped optical tactile sensor for robotic manipulation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9903–9909. IEEE, 2020.

[8] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, 2020.

[9] Barrett Heyneman and Mark R Cutkosky. Slip classification for dynamic tactile array sensors. *The International Journal of Robotics Research*, 35(4):404–421, 2016.

[10] Satoshi Funabashi, Tomoki Isobe, Fei Hongyi, Atsumu Hiramoto, Alexander Schmitz, Shigeki Sugano, and Tetsuya Ogata. Multi-fingered in-hand manipulation with various object properties using graph convolutional networks and distributed tactile sensors. *IEEE Robotics and Automation Letters*, 7(2):2102–2109, 2022.

[11] Branden Romero, Filipe Veiga, and Edward Adelson. Soft, round, high resolution tactile fingertip sensors for dexterous robotic manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4796–4802. IEEE, 2020.

[12] Nicolas Sommer and Aude Billard. Multi-contact haptic exploration and grasping with tactile sensors. *Robotics and autonomous systems*, 85:48–61, 2016.

[13] Alexander M Gruebele, Michael A Lin, Dane Brouwer, Shenli Yuan, Andrew C Zerbe, and Mark R Cutkosky. A stretchable tactile sleeve for reaching into cluttered spaces. *IEEE Robotics and Automation Letters*, 6(3):5308–5315, 2021.

[14] Andrea Cirillo, Fanny Ficuciello, Ciro Natale, Salvatore Pirozzi, and Luigi Villani. A conformable force/tactile skin for physical human–robot interaction. *IEEE Robotics and Automation Letters*, 1(1):41–48, 2015.

[15] Aimee Goncalves, Naveen Kuppuswamy, Andrew Beaulieu, Avinash Uttamchandani, Katherine M Tsui, and Alex Alspach. Punyo-1: Soft tactile-sensing upper-body robot for large object manipulation and physical human interaction. *arXiv preprint arXiv:2111.09354*, 2021.

[16] Sara Rachel Arussy Ruth, Vivian Rachel Feig, Min-gu Kim, Yasser Khan, Jason Khoi Phong, and Zhenan Bao. Flexible fringe effect capacitive sensors with simultaneous high-performance contact and non-contact sensing capabilities. *Small Structures*, 2(2):2000079, 2021.

[17] Siddharth Srivastava, Eugene Fang, Lorenzo Riano, Rohan Chitnis, Stuart Russell, and Pieter Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *2014 IEEE international conference on robotics and automation (ICRA)*, pages 639–646. IEEE, 2014.

[18] Adithyavairavan Murali, Arsalan Mousavian, Clemens Eppner, Chris Paxton, and Dieter Fox. 6-dof grasping for target-driven object manipulation in clutter. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6232–6238. IEEE, 2020.

[19] Mehmet Dogar and Siddhartha Srinivasa. A framework for push-grasping in clutter. *Robotics: Science and systems VII*, 1, 2011.

[20] Weihao Yuan, Johannes A Stork, Danica Kragic, Michael Y Wang, and Kaiyu Hang. Rearrangement with nonprehensile manipulation using deep reinforcement learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 270–277. IEEE, 2018.

[21] Mohamed Hasan, Matthew Warburton, Wisdom C Agboh, Mehmet R Dogar, Matteo Leonetti, He Wang, Faisal Mushtaq, Mark Mon-Williams, and Anthony G Cohn. Human-like planning for reaching in cluttered environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7784–7790. IEEE, 2020.

[22] Huang Huang, Marcus Dominguez-Kuhne, Vishal Satish, Michael Danielczuk, Kate Sanders, Jeffrey Ichnowski, Andrew Lee, Anelia Angelova, Vincent Vanhoucke, and Ken Goldberg. Mechanical search on shelves using lateral access x-ray. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2045–2052. IEEE, 2020.

[23] Jinhwi Lee, Changjoo Nam, Jonghyeon Park, and Changhwan Kim. Tree search-based task and motion planning with prehensile and non-prehensile manipulation for obstacle rearrangement in clutter. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8516–8522. IEEE, 2021.

[24] Sheng Zhong, Nima Fazeli, and Dmitry Berenson. Soft tracking using contacts for cluttered objects to perform blind object retrieval. *arXiv preprint arXiv:2201.10434*, 2022.

[25] Mark Moll, Lydia Kavraki, Jan Rosell, et al. Randomized physics-based motion planning for grasping in cluttered and uncertain environments. *IEEE Robotics and Automation Letters*, 3(2):712–719, 2017.

[26] Advait Jain, Marc D Killpack, Aaron Edsinger, and Charles C Kemp. Reaching in clutter with whole-arm tactile sensing. *The International Journal of Robotics Research*, 32(4):458–482, 2013.

[27] Daolin Ma, Siyuan Dong, and Alberto Rodriguez. Extrinsic contact sensing with relative-motion tracking from distributed tactile measurements. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11262–11268. IEEE, 2021.

[28] Sahand Rezaei-Shoshtari, Francois Robert Hogan, Michael Jenkin, David Meger, and Gregory Dudek. Learning intuitive physics with multimodal generative models. *arXiv preprint arXiv:2101.04454*, 2021.

[29] Tapomayukh Bhattacharjee, James M Rehg, and Charles C Kemp. Inferring object properties from incidental contact with a tactile-sensing forearm. *ArXiv*, 2014.

[30] Tapomayukh Bhattacharjee, James M Rehg, and Charles C Kemp. Haptic classification and recognition of objects using a tactile sensing forearm. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4090–4097. IEEE, 2012.

[31] Charles F. Jekel and Gerhard Venter. *pwlf: A Python Library for Fitting 1D Continuous Piecewise Linear Functions*, 2019.