On the Variability of Tactile Signals During Grasping

Qian Wan* and Robert D. Howe*†

*Harvard School of Engineering and Applied Sciences, Cambridge, USA

†Centre for Intelligent Systems Research, Deakin University, Geelong, Australia

{qwan,howe}@seas.harvard.edu

Abstract

Robotic manipulation in unstructured environments must handle a wide range of objects despite errors in visual perception. Tactile sensing is presumed to provide essential information in this context, but there has been little work examining the tactile sensor signals produced during realistic manipulation tasks. This paper presents tactile sensor data from grasping a generic object in hundreds of trials. Position error between the hand and object was varied to model the uncertainty in real-world grasping. Results show that tactile signals are highly variable despite good repeatability in grasping conditions. The observed variability appears to be intrinsic to the grasping process, due to the mechanical coupling between fingers as they contact the object in parallel, and due to numerous factors such as frictional effects and inaccuracies in the robot hand. These results have implications for improved tactile sensor system design and signal processing methods.

1 Introduction

Grasping is essential for many real-world applications of robotics. Tactile sensing is presumed to be a necessary component of autonomous grasping systems, because it provides information about the finger-object contact state that determines grasp success, and this information cannot be obtained through other sensing modalities like vision. Tactile sensing has the potential to enable robots to autonomously grasp and manipulate a wide range of objects in unstructured environments like homes and workplaces.

Although there has been considerable progress in the development of tactile sensors, the principles and techniques for integrating tactile sensors into real-time control of grasping remains a major challenge. Work on tactile signal processing has produced theoretical analysis of finger-object mechanical interactions to estimate contact location and object shape from tactile arrays [Fearing and Hollerbach, 1985], [Fearing, 1990], and control algorithms for robot fingers that use tactile signals to produce desired object motions [Maekawa *et al.*, 1995]. There has, however, been little experimental work characterizing and analyzing real tactile signals produced during grasping tasks. Recent work on using tactile sensors with machine learning has involved system-level experimental testing with diverse everyday objects, but this work has focused on learning methods without consideration of the details of the tactile signals [Bekiroglu *et al.*, 2011a], [Bekiroglu *et al.*, 2011b], [Dang *et al.*, 2011], [Dang and Allen, 2014].

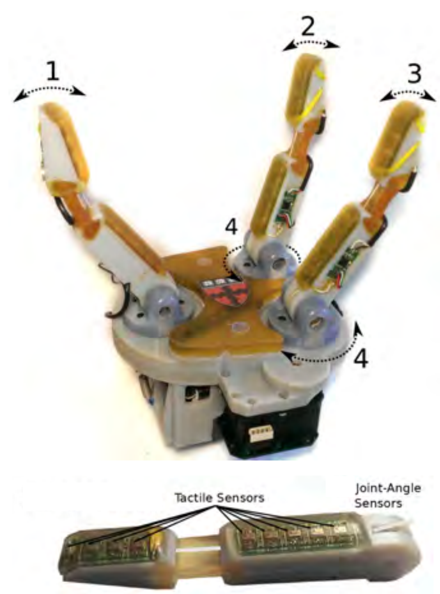

Figure 1: Top: The underactuated hand used here has four actuated degrees of freedom – three for finger flexion, one for coupled rotation of two fingers to transition between wrap grasps and pinch grasps. Bottom: Tactile sensors molded into finger contact surface.

This study provides the first detailed look at tactile sensor signals during realistic grasping tasks. In this setting, the relationship between the hand and target object is not perfectly regulated, due to the lack of *a priori* knowledge of object properties, errors in visual guidance, and inaccuracies in robot hand control. In this study, the approach is to limit the experiments to a single generic object, with a large number of trials under repeatable environmental conditions. We systematically vary hand-object positioning error, to model the uncertainty in real-world grasping tasks. In the following, we first describe the compliant underactuated robot hand and tactile sensing suite that we developed for unstructured environments. We then present sensor signals from hundreds of grasping trials, and analyze the results in terms of variation with task properties. We conclude with a discussion of the implications for the design tactile sensing systems and signal processing techniques, including machine learning.

2 Experimental methods

2.1 Underactuated hand and tactile sensors

The robot hands used in this experiment is a slightly simplified version of the iHY Hand [Odhner *et al.*, 2014] (Reflex Hand, RightHand Robotics, Inc., Cambridge, USA). This is a compliant, underactuated hand with three fingers (Fig. 1). Each identical finger has two joints, with a simple revolute pin joint with a return spring between the palm and proximal link, and an elastomer flexture joint between the links. Each finger is actuated by a tendon that passes over both joints to a pulley in the palm that is connected to a geared DC servo motor (Dynamixel RX-28, Robotis, South Korea). The motor is driven by a local torque-limited proportional-derivative position control loop. Encoders on the motor-driven spool that pulled the tendons as well as on each finger base joint provide proprioceptive motion sensing. The combination of spring-loaded joints and a single tendon allows the fingers to passively adapt to object shape as the fingers close, without the need for elaborate sensing and control.

A fourth motor provides coupled rotation of two of the fingers about their base. At one limit of rotation, the fingers articulate in parallel and oppose the third finger to perform power or wrap grasps; this is the configuration shown in Fig. 1. At the other limit, the two fingers rotate to oppose each other for precision fingertip grasps. Intermediate configurations are possible as well; in the experiments reported here, the fingers are rotated so that all three are equally spaced at 120 degrees from each other. To provide a convenient naming convention, the non-rotating finger is referred to as the thumb, and the other two fingers as the index and middle fingers, in analogy with the human right hand. Previous work has shown that iHY hand is capable of grasping a large range of objects despite significant positioning errors. Please see [Odhner *et al.*, 2014] for further details of the hand design and performance.

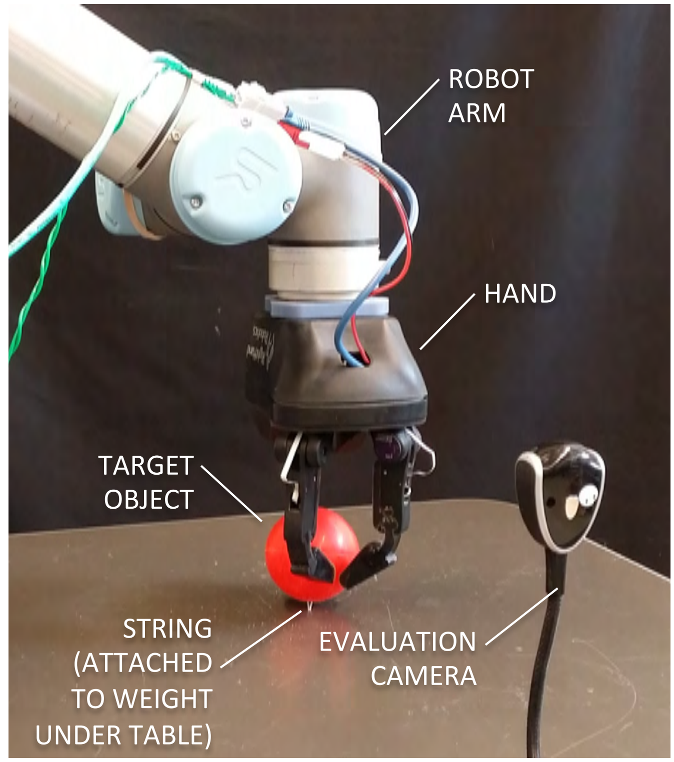

Figure 2: Experimental set up.

A row of tactile sensors is embedded in each link of each finger, with five sensors in the proximal link and four in the distal link, including one in the tip. These sensors are based on MEMS barometer sensors and include a silicon micromachined pressure sensor, precision instrumentation amplifier, high-resolution analog-to-digital converter, a microcontroller, and a standard bus interface. By leveraging the engineering investment in these sensors (which were developed for high-volume mobile phone applications), the resulting tactile sensor system has excellent performance, with 0.02 N sensitivity, approximately 100:1 signal-to-noise ratios, minimal hysteresis, good linearity, and fast sample rates. Details of the tactile sensor system fabrication and performance are provided in [Tenzer *et al.*, 2014].

The hand is mounted on a 6 dof robot arm (UR5, Universal Robots, Odense, Denmark), which positions the hand using the grasping controller described below. Hand and arm motion and all sensor processing and logging are performed under ROS by a personal computer running Linux.

2.2 Procedure

The goal here was to create an experimental protocol that enabled execution of a large number of trials, with good repeatability of the environmental conditions, particularly the relationship between the target object and the fingers. This can enable, for the first time, quantification of the variability of the tactile sensor signals as the grasping task parameters are varied. This goal was accomplished here by using a single, generic target object, a solid rubber ball approximately 65 mm in diameter (Fig. 2). To automatically return it to the same location for repeated grasping trials, it was attached to a thin string that passed through a 2 mm hole in the table top, and connected to a 200 g weight suspended only by the string. Following each grasp, the hand released the ball, and the weight pulled it into the same location on the table top for the next trial. The mean variation of the ball's location between trials was approximately 1.1 mm.

At the start of each trial, the hand moved from its starting location directly downwards. The hand descended until it reaches a preprogrammed fixed height with the finger tips just above the table top, with no rotation of the wrist. The controller then began to close the fingers, which stopped upon detecting contact with the object. The threshold for contact detection from the tactile sensors on the distal link was set to approximately 0.12 N normal force, which was found in preliminary tests to reliably detect contact. (Tactile sensors on the proximal link were not used in this study as this link made only occasional contact with the target object during grasping.) Each finger could stop independently; if contact was not detected, that finger continued closing until it reached a predetermined joint limit, with the fingers flexed to approximately 100 degrees. After all fingers stopped, the controller tightened all the tendons by approximately 3 mm; this step increased the grasp force beyond the low-force contact detection level to provide enough grasp force to lift the ball.

The arm then attempted to lift the ball. Once the arm reached a fixed height of approximately 20 cm, it stopped and an evaluation camera took a photograph to record the presence or absence of the ball, which was the criterion for success or failure of the grasping trial. The ball was colored orange to permit simple and reliable segmentation from the dark fingers and background in the photograph (Fig. 2). The controller then opened the fingers to release the ball; the weight pulled the ball back to the starting position, and the arm moved the hand to the location of the next trial. It took on average approximately 3.5 sec from start of finger closing to start of lifting (depending on how far the fingers move before stopping), and about another 10 sec from end of finger closing until the ball is lifted, photographed, and released. This setup enabled completely automatic and repeatable execution of grasping trials without human intervention, which is essential for acquiring the large number of trials for sensor characterization.

A common source of grasping variability in unstructured environments is errors in object position estimation due to visual perception limitations and robot hand and arm inaccuracies, which produce errors in positioning the hand with respect to the target object. To study the ability of tactile sensing to detect and correct this type of error, we performed two experiments. The first preliminary experiment was designed to characterize the overall grasping ability of the hand system and control algorithm, and determine the range of hand-object offsets that results in successful grasps. For this experiment, the horizontal position of the hand was varied over a range of 8 cm in both x and y directions. The hand was initially roughly centered over the ball manually to establish the center of the test offset grid, then the controller moved the hand in 1 cm increments, executing one trial at each of the 8x8 locations. The controller repeated the entire grid 16 times, for a total of 1024 trials.

The second experiment was intended to explore in greater detail the repeatability of tactile sensing signals during grasping. The hand was initially placed in a location where it was nearly centered over the ball, then moved laterally in seven 1 mm increments. At each position increment the system performed 20 grasping trials as described above. The locations were selected on the basis of the first experiment so that in the initial location, the hand successfully grasped the ball in every trial, and at the final location, the hand was unsuccessful in grasping the ball in every trial. Intermediate locations showed decreasing success rates from initial to final locations.

2.3 Results

First experiment

In the 8x8 test grid, the system achieves 100% success when the hand is positioned within a central band approximately 3 cm wide in x and 6 cm in y, where the positive y direction goes from the thumb to the opposing index and middle fingers. Most locations immediately adjacent to this central zone have intermediate success rates, while at greater distances from centered the grasping success rates generally go to zero. This provides a variety of success rates for subsequent analysis of tactile sensor signals.

Fig. 3 shows the variation of the sensor signals from the hand as a function of time during typical trials for four cases, ranging from uniform failure to uniform success. For each case, the upper three plots are the tactile pressure signals for the four sensors on the distal links, with one plot for each finger. The fourth plot in each

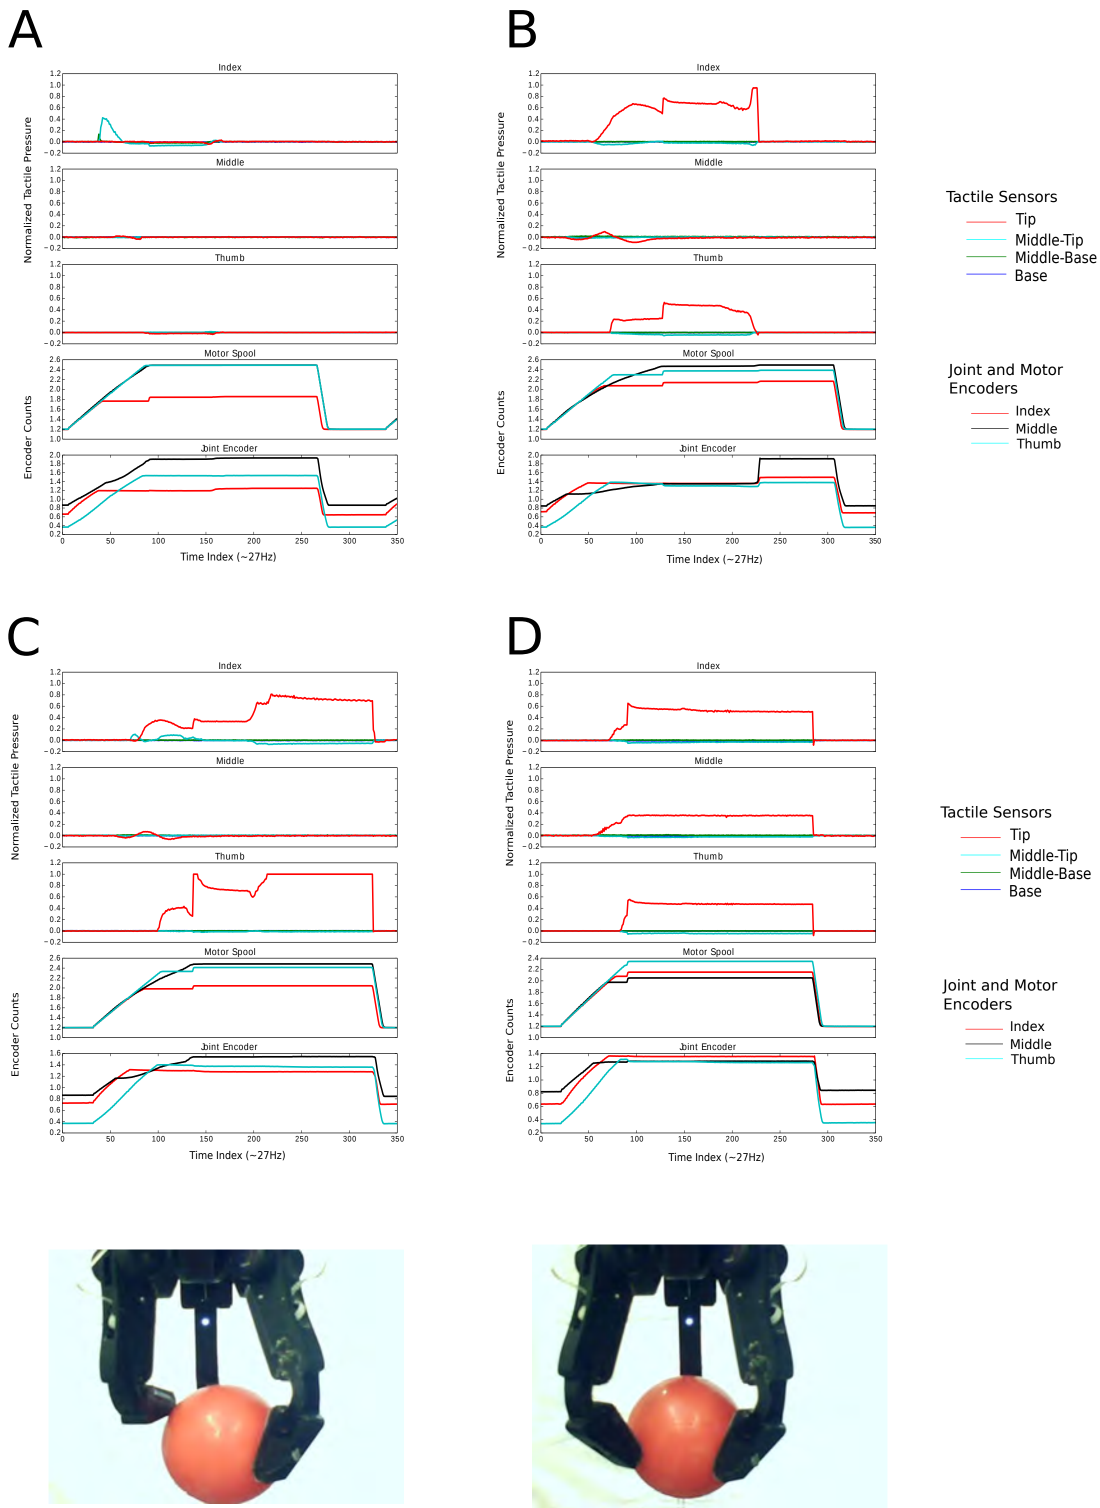

Figure 3: Sensor signals for typical examples of four cases. A: Clear failure case; B: Marginal failure case; C: Marginal success case; D: Clear success case. End-of-trial evaluation photographs are shown for successful cases C and D to illustrate hand-object configuration.

case shows the overall tendon lengths of each finger, as measured by the encoders on the motor spools, and the bottom plot shows the joint angle of the base joint of each finger, as measured by the joint encoder. The joint angle of the distal flexture joint can be estimated by the difference between the tendon length and the base joint angle, but for simplicity we elected to work directly with the raw encoder signals in this study. In each case, the plot begins with the fingers starting to close, which the controller executes by rotating the spools to shorten the tendons, producing the observed ramps in the motor spool signals. The base joint angles follow this ramp trajectory as well, unless contact with the target object deflects the spring-loaded finger.

In the clear failure case (Fig. 3B), the upper plot shows that at about 40 samples after the start of the trial, the middle-tip tactile sensor on the index finger registers a contact, and the lower plots show that the index finger stops closing, while the other fingers continue to close. After about 60 samples, however, the middle finger joint encoder and tactile sensors show small perturbations, presumably due to glancing contact with the offset target object. At this time the tactile signal from the index finger also decreases, as the object is apparently pushed aside by the other finger. At about 80 samples the middle finger and thumb have reached the limit of travel without detection of a contact, so the controller stops their motion. The controller then applies the slight tightening of the tendons (visible in the motor spool signal for the index finger) to apply sufficient grasp force for lifting, and raises the arm. Because contact was not established on the fingers, the ball is not lifted.

Similar sequences of events are visible in the other cases. In the clear success case (Fig. 3D), the tip tactile sensor on each finger records steady pressures from contact with the object, which leads to successful execution of the grasp-and-lift. The marginal cases (Fig. 3B,C) present more complex signals. Strong contact pressure signals are recorded on some fingers, but their magnitudes vary greatly as the fingers presumably push the target object around between the three fingers. In the marginal failure case (Fig. 3B), contact is apparently achieved on all three fingers because the the ball is initially lifted, but the ball slips from the fingers at about 230 samples, causing the tactile pressures to drop to zero and the joint angles to jump forward as the ball leaves the hand. Note that a consistent pressure signal is not obtained on the middle finger after an initial transient, although the finger must have provided enough force to enable lifting the ball. Similar observations apply to the marginal success case (Fig. 3C).

Some of the tactile sensor signals show negative responses during the trials, e.g. middle-tip sensor on the index finger around 100 samples in Fig. 3A and middle-tip sensor on the index finger around 70 samples in Fig. 3C. These signals denote negative stresses in the rubber fingertip at the location of the embedded tactile sensor, due to shear forces on the finger tip surface. These negative signals had been noted years ago in the context of object shape estimate using tactile sensors but their relevance for grasping and manipulation had not been apparent [Fearing and Hollerbach, 1985], [Fearing, 1990].

Second experiment

Fig. 4 shows the tactile sensor signals for the 1 mm incremental displacement experiment. A 3x3 matrix of plots is shown for each of the seven 1 mm position increments. Columns show the tactile signals for the sensors on each finger (left=index, center=middle, right=thumb), and rows show the three distal sensors (top=middle-base, center=middle-tip, bottom=tip); the base sensor is omitted as it is almost uniformly zero. The success rate progresses from all-successful trials (shown in blue) at the initial location to all-failure trials (shown in red) at the final location.

All of the plots show considerable variability at each location, despite the reasonably good repeatability of test conditions. Even at the initial location, where the grasping system achieves a successful in grasp every trial, the tactile signals are not consistent. The tip sensors on all three fingers, which have the greatest response, show initial transients of different heights and at different times, as the ball is pushed between the fingers.

In the intermediate locations, which have a mix of successes and failures, there is no clear difference between the signals from successful grasps and those from failures. For example, the plots for the location 3 mm from the start (Fig. 4 top row, right), have success and failure traces intermingled, with nearly identical traces in each category.

As a preliminary means of investigating whether there are aspects of the tactile sensor signals that change with the variation in location and success rate, Fig. 5 shows basic statistical measures at each of the 1 mm incremental locations. For each finger, the mean of all samples for all 20 trials for the signals from the three distal tactile sensor is plotted, along with the standard deviations. While strong conclusions are not reasonable to draw from this limited analysis, the plots show an increase in the mean signal for the index finger and decrease for the middle finger as the hand moves laterally. The standard deviations also show increases and decreases in parallel with the means for these fingers.

3 Discussion

This study aimed to characterize the variability of tactile sensor signals during grasping tasks. By using a single

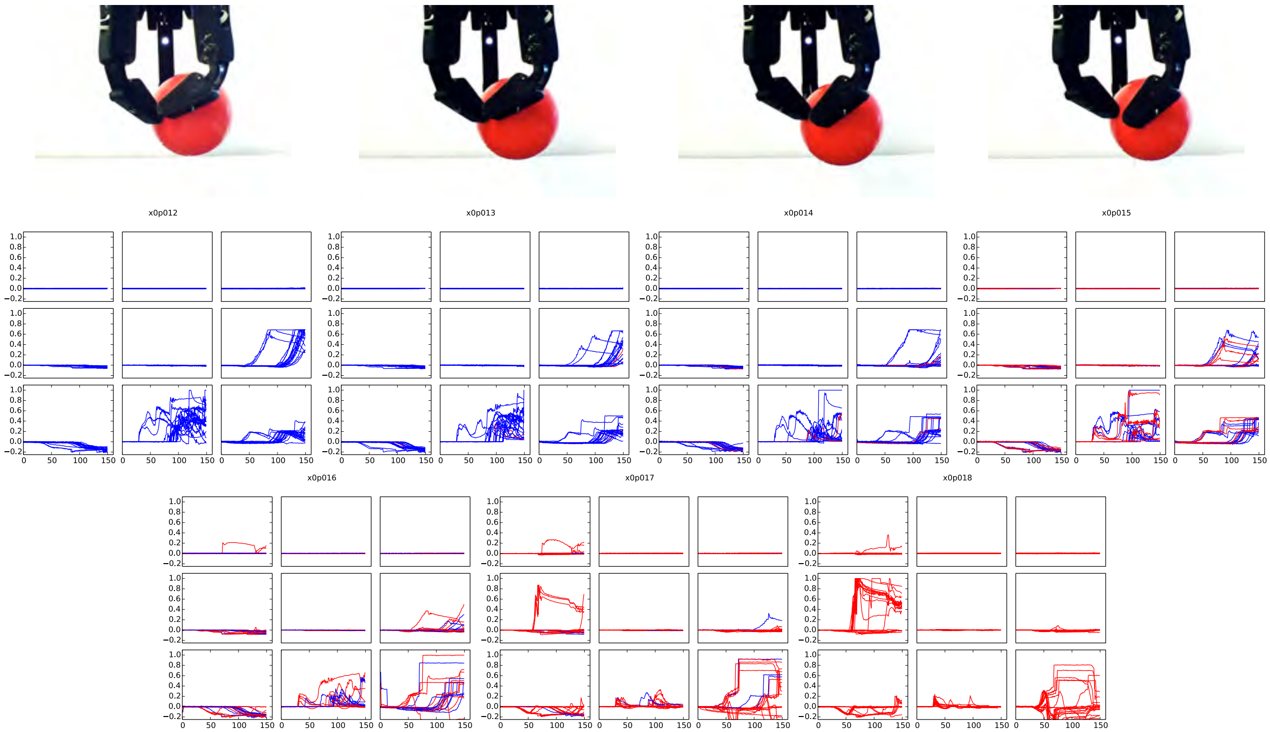

Figure 4: Tactile sensor signals vs. time as the hand moves in 1 mm increments, with 20 trials at each location. The 3x3 matrix of subplots at each location show the fingers as columns (index, middle, thumb) and the three distal sensors on each link as rows (middle-base, middle-tip, tip). Each subplot includes the sensor signals vs. time traces for all 20 trials at that location. Blue traces are successful grasps, red traces are failed grasps. Photographs show hand-object relationship at conclusion of grasp for the first four locations.

generic object and controlling the relative position of the hand and object, many repetitions of the grasping process under similar conditions were executed, so trial-to-trial variation could be examined. The results show that even under these constraints, tactile signals showed great variability. Furthermore, there was little apparent difference between the characteristics of tactile signals for successful grasping trials and failures.

There are likely a number of factors leading to this high variability. In terms of the experimental procedure, the relative position of the hand and object was not perfectly controlled. One significant source of variation was due to the target ball return string, which accounted for about 1 mm of position variation between trials. Another likely source was the robot hand's fingers, which are compliantly mounted and include an elastomer flexture joint. This compliance permits the hand to passively adapt to object geometry without active sensing and control, but it also means the fingertip positions are not uniquely determined by the motor positions. The use of 20 sequential trials repeated at identical time intervals at each location should, however, minimize this source of variability. In addition, the data showed significant difference in tactile signals and success rate at each of the 1 mm position increments, which implies that the variability in the experimental setup was not so large that the effects of mm-scale displacements were swamped.

This implies that the observed tactile signals variability is largely due to factors in the grasping process itself. From detailed observation, it appears that a major factor is the mechanical coupling between the fingers through the grasped object. Because fingers act in parallel mechanically, interactions at one finger necessarily perturb the contacts at the other fingers. In addition, the non-linear friction of polymer fingertips leads to transients as fingers stick and slip on the target object. A further complication is that the shear forces generated by friction can be confounded with normal forces to greatly affect the stress levels at the locations of the tactile sensing elements within the polymer coverings of the finger surfaces. While the ball-return string adds an "unnatural" constraint, an unconstrained object could also displace as the fingers make contact and apply forces. Similarly, stiff finger transmissions would reduce finger positioning inaccuracy, but experiments with underactuated hands have shown lower interaction forces in unstructured en-

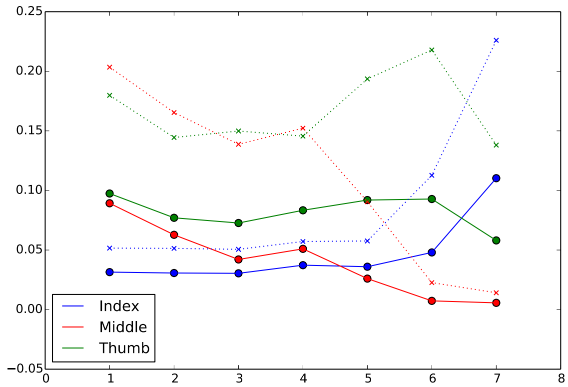

Figure 5: Mean (solid curves) and standard deviations (dashed curves) of all tactile sensors on each finger, at each 1 mm incremental location.

vironments.

3.1 Historical context

The main role of this study is to define and bring attention to a fundamental challenge in the quest to develop autonomous grasping and manipulation systems. It is perhaps surprising that no one has looked at tactile signals during realistic manipulation tasks in the past. Until recently, each of the areas associated with grasping worked largely independently: hand designers were principally focused on the challenge of designing anthropomorphic hands; tactile sensor researchers worked to find good transducers and integrate them into finger-level sensors; while systems developers used parallel-jaw grippers (particularly the PR-2) and employed tactile signals only for the most basic validation of grasp.

We have worked over many years to develop control algorithms that use tactile sensors signals to enable robust grasping in unstructured environments. In this effort we constantly fought tactile sensor signal "noise." Then a few years ago we developed MEMS barometer-based tactile sensors [Tenzer *et al.*, 2014]. These sensors have high-quality analog instrumentation and A-to-D conversion within the sensor chip, so the resulting digital signals are very clean by conventional measures, i.e. excellent signal-to-noise ratios, high linearity, low hysteresis, etc. Nonetheless, when we succeeded in integrating these high-quality sensors with our compliant robot hands, we still observed large variability in tactile signals. This study aims to quantify these effects and to draw the attention of other researchers to this issue.

3.2 Potential solutions

These results underscore the need for tactile sensor-based control methods that are immune to high variability. One approach is the contact-based grasp control approach used here, where discrete events are extracted from the continuous sensor signals, which provides some immunity from signal variations. For more sophisticated control needs, such as predicting grasp stability or selecting corrective actions to prevent dropping objects, more information must be extracted from the signals. Machine learning seems appropriate here because tactile signals are high-dimensional and noisy, and accurate models that can enable control despite noisy signals have proved difficult to define. The coherent change in simple statistical measures as the hand changes location (Fig. 5) suggests that helpful information is present in the tactile signals.

The human model may provide guidance for hand-system design. Human finger tips have only moderate coefficients of friction, but do not generally have the stick-slip behavior of soft polymers, which can lead to transients in surface loading. In addition, human finger pads are very soft, which minimizes changes in contact force levels as the fingers interact through the object. Unfortunately, this makes tactile sensing particularly challenging: sensors must be compliant for mounting in the "skin" surface, as sensors beneath a highly compliant layer will suffer from low sensitivity and poor spatial resolution [Fearing and Hollerbach, 1985], [Fearing, 1990].

Improving tactile sensor systems will require a shift in research focus, away from the development of transducers that has been the main emphasis until now, to a focus on system-level integration with hands to provide the best signals during grasping and manipulation. This necessarily includes consideration of the finger surface materials and sensor placement within the fingertip. Testing of sensor systems in realistic manipulation tasks will be an essential part of this process.

3.3 Conclusions

These studies are the first to use highly-repeatable grasping tasks with large numbers of trials to enable the study of variability in tactile sensor signals. The results demonstrate that tactile signals are messy - not due to limitations in the sensors themselves, but due to the high variability of hand-object interactions in the real world. Making tactile sensing effective in autonomous grasping and manipulation will require better designs that account for integration of sensors and hands, and new signal processing and control methods such as machine learning that can deal with high dimensionality and high variability in the signals.

Acknowledgments

Leif Jentoft was instrumental in the development of the grasping system used in these experiments. Ryan Adams provided helpful insight into tactile signal characterization. The second author, Robert Howe, has a financial interest in RightHand Robotics, Inc., the company that manufactured the hand used in these experiments. Funding support for this study was provided by the US National Science Foundation and by Deakin University.

References

[Bekiroglu *et al.*, 2011a] Yasemin Bekiroglu, Renaud Detry, and Danica Kragic. Learning tactile characterizations of object-and pose-specific grasps. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 1554–1560. IEEE, 2011.

[Bekiroglu *et al.*, 2011b] Yasemin Bekiroglu, Janne Laaksonen, Jimmy Alison Jørgensen, Ville Kyrki, and Danica Kragic. Assessing grasp stability based on learning and haptic data. *Robotics, IEEE Transactions on*, 27(3):616–629, 2011.

[Dang and Allen, 2014] Hao Dang and Peter K Allen. Stable grasping under pose uncertainty using tactile feedback. *Autonomous Robots*, 36(4):309–330, 2014.

[Dang *et al.*, 2011] Hao Dang, Jonathan Weisz, and Peter K Allen. Blind grasping: Stable robotic grasping using tactile feedback and hand kinematics. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 5917–5922. IEEE, 2011.

[Fearing and Hollerbach, 1985] Ronald S Fearing and John M Hollerbach. Basic solid mechanics for tactile sensing. *The International journal of robotics research*, 4(3):40–54, 1985.

[Fearing, 1990] Ronald S Fearing. Tactile sensing mechanisms. *The International Journal of Robotics Research*, 9(3):3–23, 1990.

[Howe and Cutkosky, 1993] Robert D Howe and Mark R Cutkosky. Dynamic tactile sensing: Perception of fine surface features with stress rate sensing. *Robotics and Automation, IEEE Transactions on*, 9(2):140–151, 1993.

[Hsiao *et al.*, 2010] Kaijen Hsiao, Sachin Chitta, Matei Ciocarlie, and E Gil Jones. Contact-reactive grasping of objects with partial shape information. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 1228–1235. IEEE, 2010.

[Maekawa *et al.*, 1995] Hitoshi Maekawa, Kazuo Tanie, and Kiyoshi Komoriya. Tactile sensor based manipulation of an unknown object by a multifingered hand with rolling contact. In *Robotics and Automation, 1995. Proceedings., 1995 IEEE International Conference on*, volume 1, pages 743–750. IEEE, 1995.

[Odhner *et al.*, 2014] Lael U Odhner, Leif P Jentoft, Mark R Claffee, Nicholas Corson, Yaroslav Tenzer, Raymond R Ma, Martin Buehler, Robert Kohout, Robert D Howe, and Aaron M Dollar. A compliant, underactuated hand for robust manipulation. *The International Journal of Robotics Research*, 33(5):736–752, 2014.

[Schiebener *et al.*, 2012] David Schiebener, Julian Schill, and Tamim Asfour. Discovery, segmentation and reactive grasping of unknown objects. In *Humanoids*, pages 71–77, 2012.

[Tenzer *et al.*, 2014] Yaroslav Tenzer, Leif P Jentoft, and Robert D Howe. The feel of mems barometers: inexpensive and easily customized tactile array sensors. *Robotics & Automation Magazine, IEEE*, 21(3):89–95, 2014.