Variability and Predictability in Tactile Sensing During Grasping

Qian Wan, Ryan P. Adams, Robert D. Howe

***Abstract*— Robotic manipulation in unstructured environments requires grasping a wide range of objects. 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 thousands of trials. Position error between the hand and object was varied to model the uncertainty in real-world grasping, and a grasp outcome prediction was done using only tactile sensors. 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, as well as numerous factors such as frictional effects and inaccuracies in the robot hand. Using a simple machine learning algorithm, grasp outcome prediction based purely on tactile sensors is not reliable enough for real-world responsibilities. These results have implications for improved tactile sensor system and controller design, as well as signal processing and machine learning methods.**

I. Introduction

Tactile sensing is essential for robots to perform autonomous grasping and manipulation. Haptic sensing provides information about the finger-object contact state that cannot be measured through other sensing modalities such as vision. A central example of the critical role of tactile sensing is in estimating grasp stability. If the robot system can predict that the fingers have not achieved a stable grasp, then the system can regrasp the object before it is dropped.

The signals from tactile sensors, however, are high-dimensional and complex, reflecting the many variables that affect the contact state [1]. To deal with this complexity, a number of recent studies have used machine learning (ML) methods to extract pertinent information from tactile sensing to predict grasp stability [2], [3], [4], [5]. Results to date have achieved limited correct prediction of grasp stability on a small range of objects, which is too low for deployment in unstructured environments such as homes and workplaces.

There are many potential explanations for the lack of success in learning grasp stability. A major issue for ML in experimental robotics in general is the burden of obtaining the large data sets required for adequate training, due to the high dimensionality of the sensor signals and the large number of parameters, which for grasping includes object size, shape, compliance, and friction, and the positions of the fingers with respect to the object. A related issue is finding the best ML method and feature set for this application, considering limited training set size and evolution of tactile signals throughout the grasping process. There are also fundamental questions about the nature of the signals from tactile sensors. Current experimental grasping systems have tactile sensors on only a portion of the finger surface, and often the sensor arrays are flat and stiff. Whether these sensors provide sufficient information for determining grasp stability in unstructured environments is, in general, unknown.

In this paper we explore the issue of tactile information content during the grasping process by examining tactile signals in detail. The goal is to characterize sensor response, particularly signal variability and predictability, and relate this to grasp stability prediction. Experimentally, we limit the parameter space that must be characterized by looking at a single spherical object and a single grasp configuration. This allows the experiments to focus on the role of positioning errors between the hand and the target object. It also enables execution of thousands of trials to examine the effects of training set size.

In the next section of this paper, we begin by describing the experimental grasping system, followed by characterization of the tactile sensor signals. We then use a support vector machine (SVM) to predict grasp stability, and achieve the approximately 90% prediction success rate seen in previous studies. Analysis of sensor signals, however, shows that poor signal quality that limits ML methods is correlated with large errors in hand-object position. This implies that a conservative strategy that predicts grasp success only if the signals are of high quality will avoid incorrect predictions of stability that can results in dropped objects.

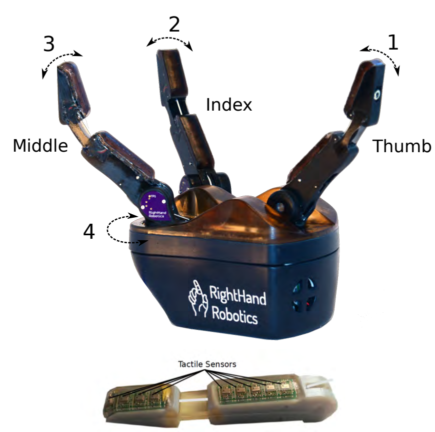

Fig. 1. Top: The underactuated hand used here has four actuated degrees of freedom – three for finger flexion, one for coupled rotation of two fingers to transition between wrap grasps and pinch grasps. Bottom: Barometer-based tactile sensors molded into finger contact surface.

All authors are with the Harvard John A. Paulson School of Engineering and Applied Sciences, Cambridge, USA
{qwan,rpa,howe}@seas.harvard.edu

II. Methods and Materials

A. Underactuated Hand and Tactile Sensors

The system used to characterize tactile signals during grasping consisted of a three-fingered hand mounted on a position-controlled robot arm, and a single target object with a return system that automatically repositioned it with high repeatability. The robot hand (Fig. 1) is a version of the compliant, underactuated iHY Hand [6] (Reflex Hand, RightHand Robotics, Inc., Cambridge, USA). Each identical finger has two joints, with a 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. The motor is driven by a local torque-limited proportional-derivative position control loop. 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, shown in Fig. 1. 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 the iHY hand is capable of grasping a large range of objects despite significant positioning errors [6].

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. These sensors are based on MEMS barometer sensors and each surface-mount package contains a pressure sensor, amplifier, analog-to-digital converter, microcontroller, and standard bus interface. The resulting tactile sensor system has excellent performance, with 0.02 N sensitivity, approximately 100:1 signal-to-noise ratio, minimal hysteresis, excellent linearity, and fast sample rate [7].

The hand is mounted on a 6 degrees-of-freedom robot arm (UR5, Universal Robots, Odense, Denmark), which has a positional repeatability specification of 0.1 mm. Hand and arm motion and all sensor processing and logging are performed under ROS by a computer running Linux. The sampling rate of the tactile sensors is at approximately 27 Hz.

B. Experimental Protocol

The experimental protocol created here focused on enabling execution of a large number of trials with good repeatability of the environmental conditions, particularly the spatial relationship between the target object and the fingers. The grasping target is a generic object — a hollow rubber ball 65 mm in diameter and 51.5 g in weight. The ball is sufficiently stiff that it does not deform appreciably under the grasping forces used here. A thin string is attached to the ball and passed through a small hole (1 mm) in the table top. A 200 g weight is suspended below the table on the

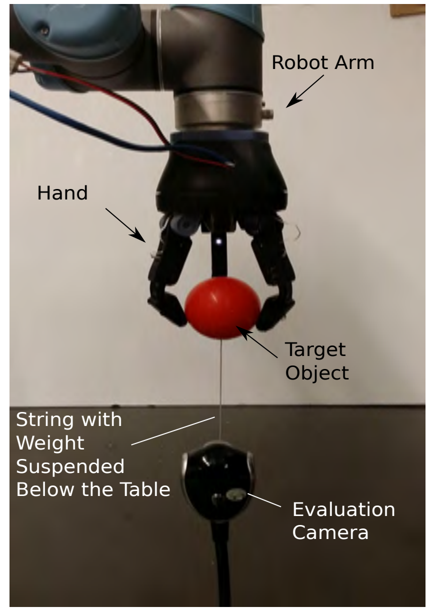

Fig. 2. Experimental setup. A 200 g weight hangs below the table on the string so that the ball automatically returns to the original position once released. Thousands of grasp trials with well-controlled position offsets can be collected using this setup without human supervision.

string. This setup gives the ball the freedom to move laterally or to be lifted off the table, but the suspended weight will automatically restore the ball to its original position after each release. Careful optical measurements showed that the ball position is highly repeatable, with mean distance of the ball from the center of the hole of 0.2 mm. This setup enabled completely automatic and repeatable execution of grasping trials without human intervention, which is essential for acquiring the large number of trials.

The grasp sequence starts with the hand descending vertically from its starting location until it reaches a preprogrammed height with the finger tips just above the table top, with no wrist rotation. The fingers then close slowly and stop upon detecting contact with the object. The tactile sensor contact detection threshold is approximately 0.12 N normal force, which preliminary tests showed to reliably detect contact. Each finger can stop independently; if contact is not detected, that finger continues closing until flexes to approximately 100 degrees. Once all the fingers stop moving, the controller tightens all the finger tendons by an additional 3 mm to increase the grasp force beyond the low-force contact detection level. The arm then attempts to lift the ball. Once the arm reaches a fixed height of approximately 20 cm, it stops, holds the ball for one second while an evaluation camera take a photograph. The presence or absence of the ball is the criterion for success or failure of the trial. The ball is then released as the weight pulled the ball back to the starting position, and the arm moves the hand to the location of the next trial. Trials take an average of about 13.5 seconds.

Grasp success vs. hand position. A common source of grasping variability in unstructured environments is error in hand-object positioning due to visual perception limitations and robot inaccuracies. To simulate such variation, the starting positions of the hand were shifted in random small offsets around the ball, within a space of 9×9 cm surrounding the ball. A total of 1400 trials were collected, out of which 35% were successful grasps. In addition, a large number of grasps were attempted at three locations separated by 1 mm at the boundary between successful and failure grasps to examine tactile signal variation in marginal locations in detail.

Rigid object. Our preliminary results suggested that tactile signals are highly variability despite repeatable hand-object positioning. This variability could be due object motion as the fingers first make contact. To investigate this hypothesis, trials were conducted with the ball rigidly fixed to the table by screws at its center. The grasping algorithm was the same as above, but the lifting segment was omitted. The hand starting location was $(x, y) = (0,\ 1.6\ \text{cm})$, where the positive x direction goes from the middle finger and extends toward the index finger, and the positive y direction goes from index and middle fingers and extends towards the opposing thumb. This location was determined to be near the limit of successful grasps and therefore exhibited high signal variation. The same procedure was also repeated at the same location, with the ball freely moving on the string, in order to ascertain the effects of object motion on grasping signals.

SVM prediction. To gain insight into the information content of the tactile signals, we used a simple linear support vector machine (SVM) to classify grasp success prior to liftoff. In this study, we used the following notation:

- $\mathscr{D} = [o_i], i = 1, \ldots, N$ denotes a data set with N trials.
- $o_i = [x_t^i], t = 1, \ldots, T_i$ is a trial with T_i samples.
- $x_t^i = [\mathbf{p}_t^i]$, where $\mathbf{p} \in \mathbb{R}^9$ is pressure from tactile sensors

Pressure signals from the tactile sensors were normalized to the maximum magnitude found in all of the sensors across datasets. A binary label was used to indicate the success and failure of each grasp. Because trial lengths vary depending on the timing of the contacts, we aligned the time index of the trials by setting the start of the lifting process as t_{100}.

A classifier was trained at every 10 time steps using both a linear SVM and a kernelized SVM with a radial basis function kernel (RBF). To classify a grasp outcome at a given time point t, all the x_t^i prior to t are concatenated to form the feature vector $\mathbf{f}_t^I = [\mathbf{p}_0^I, \ldots, \mathbf{p}_t^I]$. The soft margin parameter C and RBF parameter γ was optimized through a 10-fold cross-validation. Different portions of trials were taken out of the training sets in order to study the effect of training set size and learning results. The classifiers were tested on 100 trials that were randomly selected from the testing set of 400.

III. Results

Examples of unfiltered raw signals from two trials are shown in Fig. 3. For each case, the upper three plots are the tactile pressure signals for the three sensors on the distal links, with one plot for each finger. We executed only fingertip-grasps, so sensors on the proximal link were not used because the object rarely made contact there. The fourth plot in each case shows the overall tendon length of each finger, as measured by the encoders on the motor spools, and the bottom plot shows the joint angle of the base joint

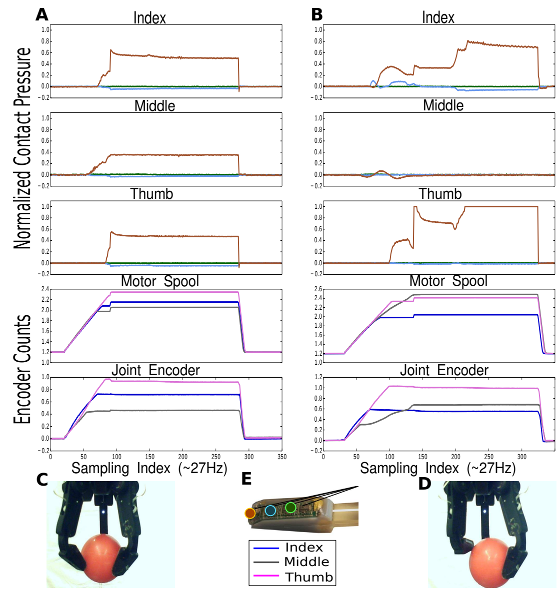

Fig. 3. Sensor signals. A: Clear success case. B: Marginal success case. C&D: corresponding end-of-trial evaluation photographs. E: top: color legend for the tactile sensors in A&B; bottom: encoder and spool legends.

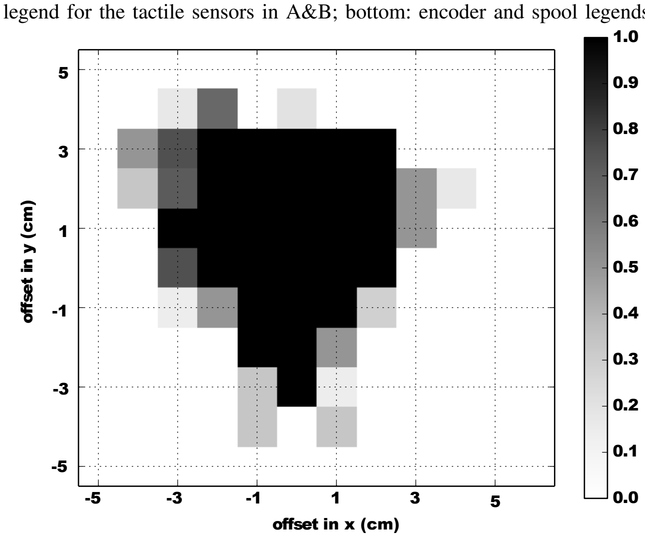

Fig. 4. Success rate at each hand offset position. The hand is approximately centered over the ball at (0,0). The color scale shows % of successful grasps out of a total trials executed in the 1cm×1cm block. The hand is oriented so that the thumb points towards positive y, and the fingers point toward the lower corners.

of each finger, as measured by the joint encoder. In each case, the plot begins with the fingers starting to close, which the controller executes by rotating the spools to shorten the tendons, increasing the angle of flexion, and 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.

Both of the examples shown are taken from successful grasp trials, though the signals are significantly different. The case where the ball is enclosed by fingers symmetrically (Fig. 3A,C), the tactile sensors on each finger record steady pressures from contact with the object, which leads to successful execution of the grasp-and-lift. For the case where the

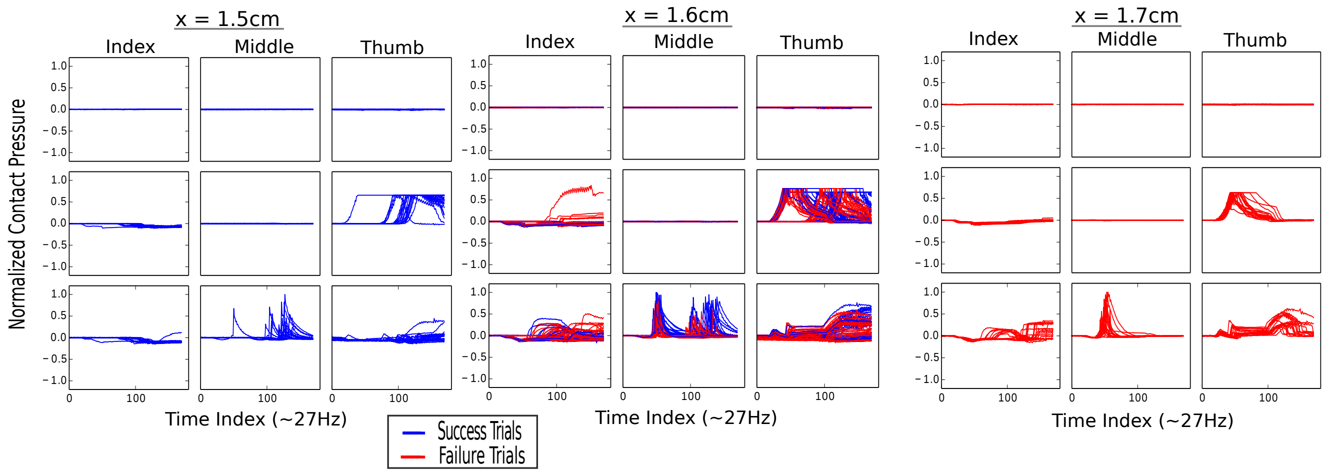

Fig. 5. Tactile sensor signals at the margin. These 3×3 plot matrices are raw sensor signals taken at y=0, x={1.5 cm,1.6 cm,1.7 cm}, with 100 trials at each location. The subplots at each location show the fingers in columns (index, middle, thumb) and the three distal sensors on each finger as rows (middle-base, middle-tip, tip). Each subplot includes the sensor signals vs. time traces for all 100 trials at that location. $t = 0$ marks the start of fingers closing. Blue traces are successful grasps, red traces are failed grasps.

fingers did not result in symmetrical enclosure due to hand-object position offsets (Fig. 3B,D), the sensors presented more complex signals. Strong contact pressure signals are recorded on some fingers, but their magnitudes fluctuated as the fingers push the target object around between the fingers. The final grasp configuration was stable but the ball was resting against the side of one of the compliant fingers.

Some tactile sensor signals show negative responses, e.g. middle sensors in all fingers in Fig. 3A. This denotes negative stresses in the rubber fingertip at the location of the embedded tactile sensor, due to shear forces on the fingertip surface. These negative signals have been noted in previous work on tactile sensing, beginning with Fearing's work on object shape estimation using tactile sensors [8].

Grasp success vs. hand position. Fig. 4 shows the success rate at each location in the 9×9 grid of hand locations. The fingers fail to contact the ball in regions beyond this grid, so grasp failure is certain. Within the grid, 100% of grasps are successful when the hand is positioned within a central band approximately 5 cm wide in x and 6 cm in y. It is approximately symmetric in the x direction and heavily biased in the y towards the thumb, because the ball can get wedged between the paired middle-index fingers.

The marginal regions between the success and failure regions in Fig. 4 show that there are locations where millimeter shifts in hand-object position offset produce different outcomes. Fig. 5 examines this situation in detail. These plots show tactile signals from one hundred trials at each of three locations that span 3 mm, where the grasp outcome drops from 100% success to 100% failure, with the intermediate location showing a mix of successes and failures. More importantly, although there are clear differences in the tactile signal patterns between the 100% success and 100% failure locations, at the intermediate location there is no clear difference between the signals from successful grasps and those from failures: success and failure trial signals are overlapping, with similar trajectories.

Rigid object. The trials examining contact sensor signals at the initial stages of grasping showed that when the ball is rigidly fixed to the table and cannot move due to finger contact, the contact signals are repeatable (Fig. 6B). All traces have very similar time courses, with clear step transitions when the fingers make and break contact with the ball. There is some variability in signal magnitude, which may in part be attributed to re-zeroing the sensors before each trial. In contrast, in trials where the ball is free to move (Fig. 6C), the fluctuation of signals is obvious, with much greater variability between the trials. Although contact is made on all three fingers when the ball is free to move, we see the clean step signal in only a handful of trials.

SVM prediction. The previous results demonstrated that both grasp success rates and tactile sensor signals can be highly variable from trial-to-trial despite good repeatability of the hand-object relative position. Furthermore, relationships between grasp success and tactile signals is not readily apparent in marginal locations (Fig. 5). One of the advantages of machine learning methods, however, is that they can uncover higher-order relationships among signals.

Figure 7 shows the SVM grasp stability prediction as a function of time. The success of the prediction at each time point was quantified in terms of the Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC). The prediction success plateaued at 0.80 at the time of lift. Classifiers that were trained using few data points (e.g. 50 trials) reached only 0.75. Training with significantly more data (e.g. 800 trials) did not perform better than training with 200 trials. The nonlinear RBF kernel did improve the prediction performance to 0.90, but also plateaued regardless of increased training samples.

A. *Using variability and location to predict stability*

The above results establish that tactile signals are variable in general, and highly variable at marginal locations. These are the locations where the hand is several cm away from

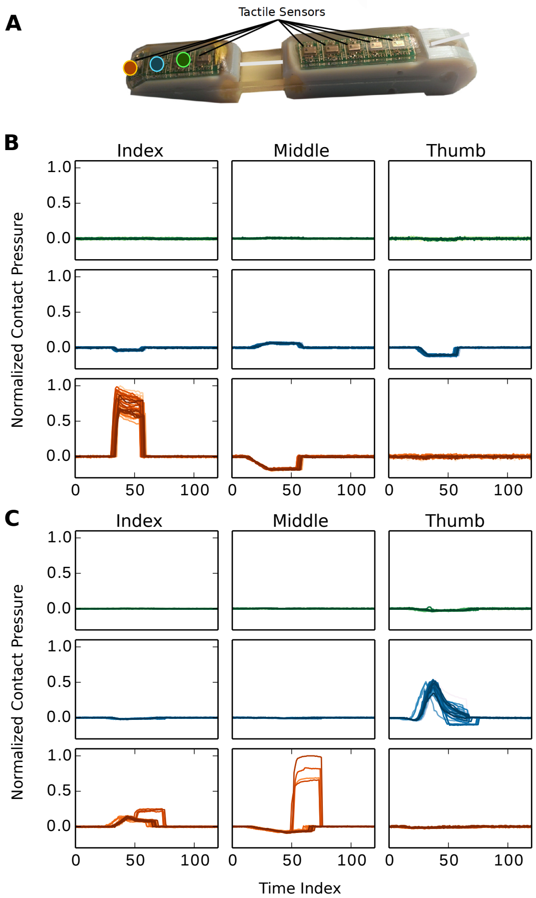

Fig. 6. Sensor variability at the onset of grasping. A: Color legend for sensor plots B&C. B: Signals from onset of grasping when the ball is fixed. C: Signals when the ball is tethered by the string and weight. The columns in B&C represent the fingers, and the rows represent the three distal tactile sensors, with the tip sensors on the bottom.

centered over the object and the grasp outcome transitions from 100% success to 100% failure (Fig. 4). In the following section we discuss the reasons for this behavior, but the data clearly supports the conclusion that both grasp stability and tactile signals are unreliable away from the center.

This suggests that an alternative strategy for predicting grasp stability from tactile signals. Rather than attempting to predict stability for grasps anywhere in the workspace, we can detect when the tactile signals are strong and consistent. Fig. 3A shows that for strong grasps, the distal sensors on all three fingers produce large, clear signals. For successful grasps that do not make consistent contact with the fingers (Fig. 3B), finger configuration in those grasps are usually not ideal (Fig. 3D). Therefore, if we only trust tactile sensors when the signals are loud and clear, we automatically eliminate the contorted grasps.

This is a conservative strategy in the sense that it cannot detect every stable grasp. Fig. 4 shows that successful grasps occur up to 5 cm from the centered location. However, This conservative strategy minimizes the likelihood of a false positive (incorrectly predicting a grasp is stable), which can lead to dropping the object. The concomitant increase in false negatives (incorrectly predicting a grasp is unstable) leads to

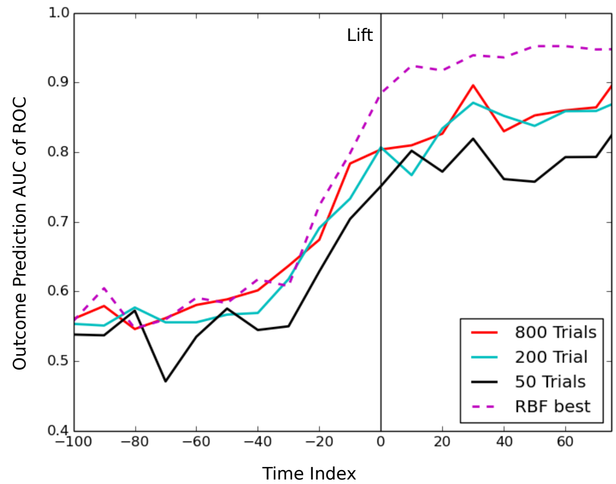

Fig. 7. Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC) for grasp outcome prediction. A prediction is made at every 10 sample intervals. Solid lines are results from a linear SVM trained using different numbers of samples. Dotted magenta line is the plateau reached by SVM with radial basis function kernel. T_0 is the time of lift for all trials. The baseline is at 0.65 due to bias in sample groups.

unnecessary regrasping, but this is presumably preferable to false positives in most applications.

We examined a variety of parameters derived from the sensor signals, and found strong correlation between locations with 100% grasp success and the total pressure measured at the distal link of each finger. Fig. 8 plots the spatial distribution of the smallest sum of tactile pressures on one finger, $\min_j\{\sum_{i=1}^{3} p_{i,j}\}$, where $p_{i,j}$ is the pressure at the ith sensor on the jth finger. This parameter is strongly peaked near the location where the grasp is symmetric and the tactile sensors received consistent and strong signals.

We can create a very simple deterministic classifier that labels a grasp as stable if the minimum-pressure-sum is above a threshold, which is set so that the included regions in Fig. 8 fall within the region of 100% successful grasp trials in Fig. 4. For demonstration purposes we selected a threshold of 1.0, corresponding to the top contour in Fig. 8. The resulting prediction success is plotted in Fig. 9, along with the SVM results from above for comparison. This deterministic prediction provides a low 72% correct prediction rate, compared to approximately 90% for the SVM classifier. The more important false positive rate, however, is only 0.3%, while the rate for SVM is about 5%. This suggests that even a simple strategy that takes into account the quality of the tactile signals can drastically improve grasp stability prediction.

IV. Discussion

A. Machine Learning of Grasp Stability

Previous attempts to use ML-based techniques to assess grasp stability achieved only limited success, typically around 90% or less [2], [3], [4], [5]. These studies were based on the assumption that tactile signals contain sufficient information to discriminate grasp success from failure. The

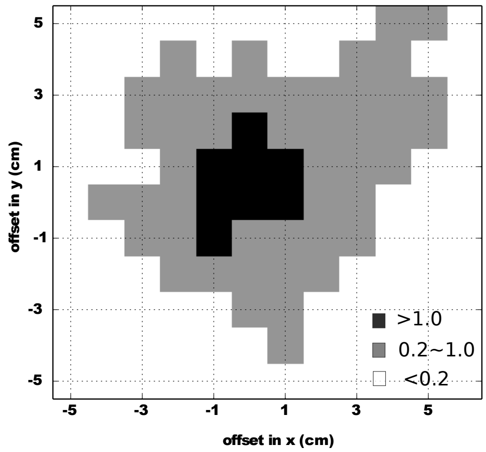

Fig. 8. Deterministic prediction parameter. The sum of the sensors on the distal link is calculated on all three fingers, with the lowest of the three fingers plotted here. A threshold of 1.0 ensures that sensors values above this level correspond to 100% success trials locations in Fig. 4.

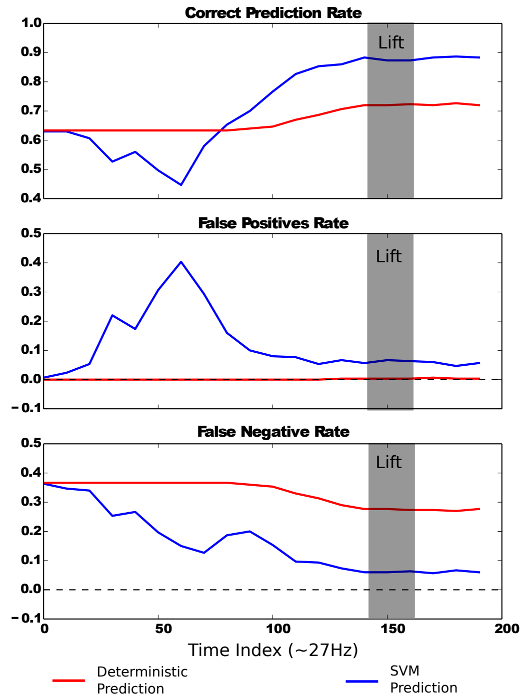

Fig. 9. Grasp outcome prediction vs. time. A prediction is made at every 10 sample intervals. Blue lines are results from a linear SVM classifier. Red line is prediction made by setting a threshold based on the minimum-pressure-sum. T_0 is the start of fingers closing. The gray band is the approximate times of when lift occurs.

results presented here suggest that tactile signals can be so variable that they cannot serve this function, at least not for all grasping situations in unstructured environments. In particular, the results in Fig. 5 show that the tactile signals produced in grasps with marginal stability are not well-correlated with grasp success. ML methods are thus unable to learn to discriminate stable from unstable grasps for these cases, and their inclusion in training and testing data reduces overall prediction rates below the very high level required for applications in unstructured environments.

One solution to this problem is to limit prediction of stability to cases where the tactile signals are unambiguous. In the limited context considered here, high-quality tactile signals occur in the regions where grasp stability is excellent. This means that a simple strategy of detecting good signal levels on all three fingers can be used to predict that the grasp will succeed, with excellent accuracy. While this deterministic type of prediction method reduced the overall prediction accuracy when compared to SVM classification, the false positive rate dropped to 0.3% — much closer to application requirements.

This approach, however, applies more generally than the specific experimental situation examined here. In unstructured environment, grasps should be conservative, with a large safety margin to allow for unmodeled and unsensed perturbations. Our experiment shows that the traditional "grabbed or dropped" definition of grasp stability may not be a good evaluation metric. For example, the grasp shown in Fig. 3B,D is stable, but the object is restrained by the side of one finger, rather than all three finger tips as in Fig. 3A,C. This grasp has asymmetric contact locations and relies on high friction to achieve force closure, which could lead to instability in subsequent manipulation. Additionally, most robot hands are designed with tactile sensors located on the main grasping surfaces, so that signals are only trustworthy when contacts are on the main surface. Therefore a more sophisticated success metric should discount grasps where contacts are made at where the sensors are not credible.

A key question for these results is their generality: to what extent do they depend on the specific hardware and experimental protocol used here? The problem of obtaining good tactile signals while grasping in unstructured environments is certainly not limited to this grasping system. Many robot hands have flat links on the intended grasping surface (e.g. the Barrett and Schunk Hands), with flat, stiff tactile array sensors. Aligning these finger surfaces with object surfaces to achieve good grasp and strong tactile signals can be challenging. To some extent this can be mitigated by other elements of the grasping system, e.g. better object localization by the vision system, better shape modeling by the perceptual system, etc. Recent studies have begun to consider the problem of finger re-positioning based on tactile signals [3], [4]. All of these approaches can help to generate the high-quality tactile signals that are needed to validate grasp quality.

B. Variation in Grasping and Tactile Signals

One of the main results of this study is the characterization of the variability of tactile sensor signals during grasping tasks. By using a single generic object, tightly controlling the relative position of the hand and object, and executing many repetitions of the grasping process under similar conditions, we were able to isolate the fundamental variability of the grasping process. The results show that even if the hand and object pose were repeatable to within 1 mm, there are marginal locations where grasp success or failure were equally likely. Furthermore, tactile signals showed great variability, and in the marginal locations there was no straightforward means to distinguish the signals of success trials from failure trials.

This suggests that the observed tactile signal variability is largely due to physical factors in the grasping process itself. It appears that a major factor is the mechanical coupling between the fingers through the grasped object. The fixed-ball experiment (Fig. 5) shows that because fingers act in parallel mechanically, interactions at one finger perturb the contacts at the other fingers, and mm-scale differences in object position cause significant changes in the resulting interaction forces and tactile signals. Seen from the viewpoint of forces applied to the object and the resulting behavior, these results are not specific to the hand hardware used here; other hand designs that execute a similar grasp strategy would likely produce the same object behavior. While the compliance of the fingers in this experiment allowed fingers to deflect during grasping which may have increased variability, stiff fingers would tend to perturb the object position and potentially increase variability as well.

More sophisticated machine learning methods than SVM can no doubt be effective in dealing with the higher-dimensional parameter space of diverse object in diverse settings; however, to train a system to handle all the variations that exist in unstructured environment would require experiments with a much wider variety of objects and grasping tasks. Because real-world robotic data takes considerable time and effort to obtain, exhaustive data will be challenging to obtain. These results also raise questions about the role of simulation in grasping; can and should simulations capture the intrinsic variability and the effects of mm-scale displacements which were shown to greatly affect grasp success and tactile signals?

C. Implications for hands and tactile sensing systems

To the best of our knowledge, this is the first study to systematically examine tactile signals during realistic grasping tasks. This study provides strong motivation for interweaving sensor design, controller and manipulator design, and system integration. The barometer-based tactile sensors used here have high-quality analog-to-digital conversion within the sensor chip, so the resulting digital signals are very clean by the conventional measures that have been used in the literature to evaluate tactile sensors (i.e. excellent signal-to-noise ratios, high linearity, low hysteresis, etc.). The fixed-ball experiment demonstrated good signal reproducibility when the complexity of the mechanical interaction is removed, confirming the functionality of the sensors. Nonetheless, the results here showed large variability in tactile signals once the sensors were integrated into a hand and used in realistic tasks, due to the high variability of hand-object interactions in the real world. Hence benchtop tests are not sufficient for predicting tactile sensor performance in integrated grasping systems. Experimental testing of sensor systems in realistic manipulation tasks should be an essential part of the characterization process going forward.

V. Conclusions

This study demonstrated that under realistic scenarios, grasping is intrinsically variable. Fingers are coupled to the target object and thus to each other, and friction is nonlinear and difficult to predict. This means that small changes in contact conditions can lead to different grasping outcomes. Concomitantly, the signals from tactile sensors are also highly variable. A simple SVM implementation extracted important grasping information from the tactile signals, but even in this carefully-constrained task, the rate of successful grasp stability prediction only reached about 90% accuracy at the point of lift. A more conservative strategy of predicting stability only when tactile signals are strong on all fingers achieved far lower false positive predictions, but requires regrasping of a significant fraction of nominally stable grasps. More sophisticated learning algorithms may resolve the complexities in environments with few sources of variation, but given the wide variety of objects and tasks that a grasping system must face in the unstructured real-world, a fundamental understanding of variability sources in the tasks and grasping system can inform the balance between designing better sensor systems, investigating more complex learning algorithms, and developing variation-resistant controllers.

Acknowledgments

Leif Jentoft was instrumental in the development of the grasping system used here. The senior author, Robert Howe, has a financial interest in RightHand Robotics, Inc., the company that manufactured the hand used in these experiments. Funding support was provided by the US National Science Foundation under grants IIP-1500178 and IIP-1448975.

References

[1] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *ICRA*. Citeseer, 2000, pp. 348–353.

[2] Y. Bekiroglu, J. Laaksonen, J. A. Jørgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *Robotics, IEEE Transactions on*, vol. 27, no. 3, pp. 616–629, 2011.

[3] H. Dang and P. K. Allen, "Stable grasping under pose uncertainty using tactile feedback," *Autonomous Robots*, vol. 36, no. 4, pp. 309–330, 2014.

[4] M. Li, Y. Bekiroglu, D. Kragic, and A. Billard, "Learning of grasp adaptation through experience and tactile sensing," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. Ieee, 2014, pp. 3339–3346.

[5] J. Schill, J. Laaksonen, M. Przybylski, V. Kyrki, T. Asfour, and R. Dillmann, "Learning continuous grasp stability for a humanoid robot hand based on tactile sensing," in *Biomedical Robotics and Biomechatronics (BioRob), 2012 4th IEEE RAS & EMBS International Conference on*. IEEE, 2012, pp. 1901–1906.

[6] L. U. Odhner, L. P. Jentoft, M. R. Claffee, N. Corson, Y. Tenzer, R. R. Ma, M. Buehler, R. Kohout, R. D. Howe, and A. M. Dollar, "A compliant, underactuated hand for robust manipulation," *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 736–752, 2014.

[7] Y. Tenzer, L. P. Jentoft, and R. D. Howe, "The feel of mems barometers: inexpensive and easily customized tactile array sensors," *Robotics & Automation Magazine, IEEE*, vol. 21, no. 3, pp. 89–95, 2014.

[8] R. S. Fearing and J. M. Hollerbach, "Basic solid mechanics for tactile sensing," *The International journal of robotics research*, vol. 4, no. 3, pp. 40–54, 1985.