Haptic Classification and Faulty Sensor Compensation for a Robotic Hand

Hannah Stuart, Paul Karplus, Habiya Beg
Department of Mechanical Engineering, Stanford University

Abstract
Currently, robots operating in unstructured environments rely heavily on visual feedback from cameras to perform tasks. However, in many situations visual feedback is not available, for example if the robot must work in the dark or grab an object from a bag. In these environments, haptic feedback in the robot's hands can augment visual feedback to help the robot identify object size, weight, and pull force.

We implemented machine learning algorithms on data taken from a robotic hand to predict the size of the object grasped, pull force on the object, and direction of the pull force. In addition, the robotic hand we tested had several faulty sensor samples (such as noisy electrical connections). We used Linear Regression, Multivariate Naïve Bayes, Multiclass Logistic Regression, Multiclass SVM Regression, Gradient Ascent Algorithm, and K-Means Clustering to determine which sensors were faulty and compensate if possible or chose an algorithm robust to noise and error.

1 Introduction

Currently, robotic control in unstructured environments relies heavily on visual perception because it is unobtrusive and reliable in many circumstances [1]. However, haptic exploration and manipulation allows robots to interact successfully with dynamic and complex environments when visual feedback fails (for example extracting a specific object from inside a duffle bag or manipulating a tool while vision is obstructed). However, tactile and position sensors can become noisy or erroneous due to physical damage in these potentially treacherous contact situations. We seek to identify and compensate for faulty haptic and position sensory data. Also, the process of estimating what we expect the hand to "feel" could become a part of the automatic feedback control of the hand – changing actuation to achieve a specific type of grasp. Given unreliable sensor feedback, robust error compensation algorithms may be particularly important, for example, in human-robot interaction safety.

Research has been conducted on replicating human haptic learning in humanoid robots to improve the cohesion and performance of robots in real world situations [2]. Machine learning algorithms have been employed to classify objects, solely from haptic sensor data without explicitly modeling the object shape [3].

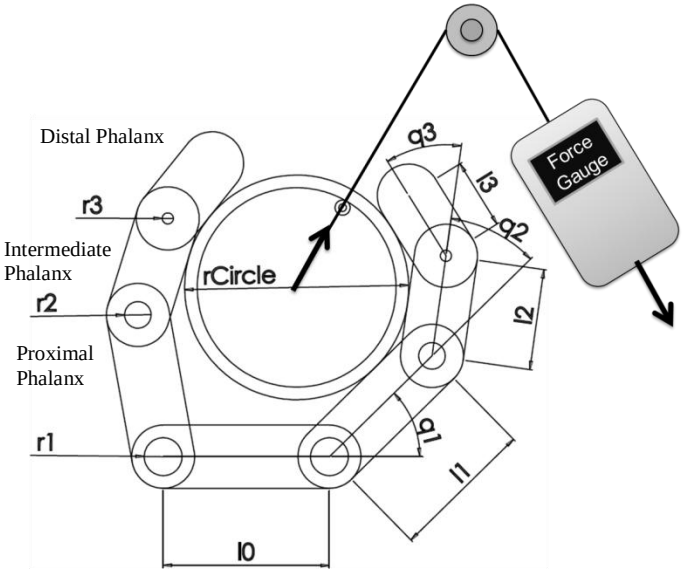

Figure 1: Experimental Setup.

2 Experimental Setup

We tested the ARM-H SRI hand, which has four cable-driven fingers, one motor per finger, and one tendon per finger – this hand is under-actuated because it has more degrees of freedom than actuators. During testing, two opposing fingers gripped a 3.5" or 4.5" diameter PVC, 1.5" long tube. Then, the tube was slowly pulled out by its center of mass (by using a low friction slider on the inside edge of the tube) at angles varying from 0 to 45 degrees (0 to pi/4 radians) in 15 degree increments as show in Figure1 and 2. Grasp force (same as tendon force) was also varied from 10 to 25 Newtons in steps of

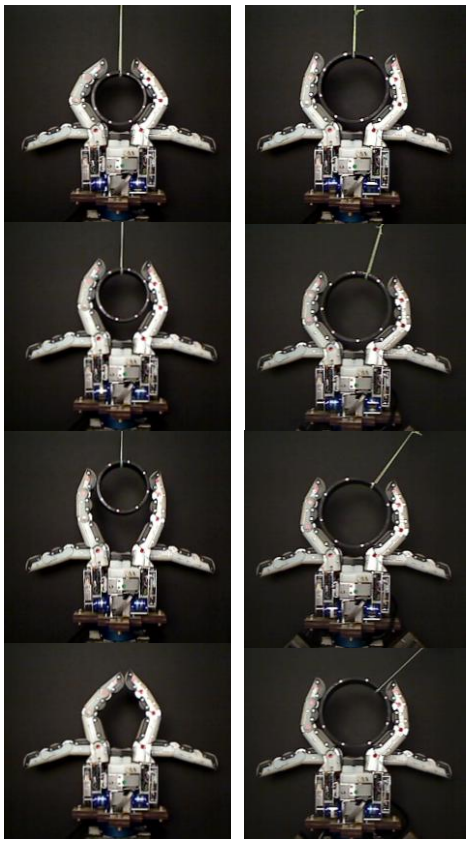

Figure 2: Examples of (Left) one pull trial with a 3.5" tube size and (Right) different pull angles with 4.5" tube size.

5N, but was kept constant during each trial.

Knuckle angle, finger pad pressure, video, gauge pull force, and tendon tension were recorded as a function of time using the Robot Operating System (ROS). There is one capacitive encoder per knuckle and 18 pressure sensors distributed through the finger pads of each finger (36 total). The hand was mounted on a 6-axis JR3 100-N load cell for gauge force redundancy. A webcam recorded images of the scene to collect displacement information and was calibrated using the image toolbox included in ROS.

Frictional effects between the object and finger pads are pose dependent, so it was desirable to place the object in roughly the same initial position every test. Therefore, the object was shaken in the grasp to allow the hand to settle to a more optimal starting condition (which also happens to indicate local stability in the initial grasp). The object's pull cable was routed through a pulley placed at approximately the desired pulling angle, and connected at the other end to a force gauge (see Figure 1). This provided easy force measurement along the direction of the pull cable, so will be easier to analyze. Video images and JR3 readings were not utilized for this experiment, but future work could incorporate them.

3 Machine Learning Experiments

Machine learning methods have been employed to classify pull angle, object force, and object size in the hand's grasp. Finger pad sensors were calibrated using Linear Regression and used to train Naïve Bayes, Support Vector, and Logistic Regression models for predicting pull angle and object force. Encoders were calibrated using K-means clustering and used to train an SVM model for object size classification.

The pull experiments were conducted 137 times, with a sampling rate that provided 8704 samples total. Because of the large data sets, all accuracies are calculated using a hold-out cross validation method where the model is trained on 70% of the data (randomly selected) and tested on the remaining 30%.

3.1 Finger Pad Force Sensor Characterization

The raw data was calibrated such that, given a reading from the 36 sensors you could estimate the normal force location and amplitude on each phalanx (3 per finger, 6 total). Calibration data for the finger sensors was collected by placing a printed strip of dots over the surface of each finger with a repeatable locating jig and equidistant test locations (5mm apart). Slowly pressing the force gauge, by hand up to 30N, directly on a known location of the pad surface, we recorded raw force values from the tactile sensors on each phalanx. During this preliminary test, 14088 pressure samples and 6296 gauge force samples were collected over the same time period, but with different sample rates. Therefore, finger pad readings were linearly interpolated to match the time stamps of the gauge.

Linear regression was used to formulate a hypothesis for applied force (F), length along the finger (Y), and location along the width of the finger (X) using the array of sensors on each phalange of one finger. The normal equation design matrix was dimensioned 6296xn, where n was the number of sensors on that specific phalange. Testing error was calculated for each phalange model:

Phalanx	Proximal	Intermediate	Distal
# of sensors, n	6	4	8
F error	12.68	12.60	12.40
Y error	0.89	0.88	0.88
X error	2.25	1.12	2.69

These errors are largely due to the mechanical shortcomings of the fingers. There are locations on the finger pads that do not have sensing material underneath. For example, large errors occur primarily in sample locations which are not located directly above sensors – locations where the sensors are only detecting strain propagations through the polyurethane skin. There is also mechanical play where finger pad connections do not have locating features. "No load" output from each individual sensor may also change due to flexing of the plastic base. Considering these factors, this model error is reasonable, however Linear Regression is susceptible to noise.

Figure 3 illustrates the hysteresis of the one working sensor on the intermediate phalanx. In this case, slow force application gives the more linear region of skin readings while faster force release decreases non-linearly, most likely due to the dynamic characteristics of the materials encasing the sensors. However, this

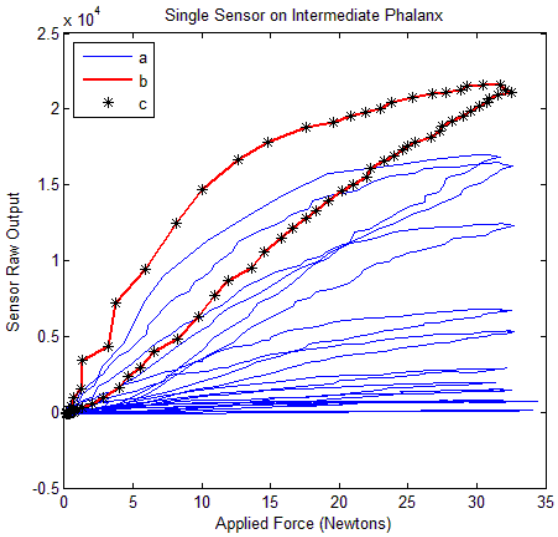

Figure 3: (a)Typical outputs from one working sensor – 12 locations are tested on this phalax, resulting in 12 hysteresis loops. One highlighted test location (b) appears to be centered on the sensor. (c) Individual samples (note change of force rate).

nonlinearity could be due to human error also. On the other finger, the sensors were more erratic and unreliable due to a faulty electrical connection. Therefore, symmetry is used to assumed that this calibration process can represent both fingers when they work properly (this is an important source of error during later model development). This highlights a main challenge for roboticists – unreliable sensors.

3.2 Pull Angle Classification Using Pad Sensors

Pull angle was classified using multivariate Naïve Bayes (NB), Support Vector Machine (SVM), and Logistic Regression (LogReg). These algorithms were run using all raw data (36 input variables) and the Y and F outputs from the calculated linear regression hypothesis (12 input variables). The variance of the data is high at low object pull forces, so each algorithm was retrained a second time excluding pull forces below 5N. NB was run using kernelized multivariate Naïve Bayes capabilities. Multivariate LogReg and SVM were conducted using the liblinear-1.92 library provided by the Machine Learning Group at National Taiwan University.

Figure 4 summarizes and compares the resulting test error for each training method. LogReg and SVM show similar trends, and performed best when trained on all raw data, resulting in 86.63% and 86.94% accuracy respectively.

Of the misclassified cases, when using NB on all raw data, most were within one class value of the correct angle (only 15 degrees off). Only 5.6% of predictions were significantly wrong (off by 30 degrees or more). Figure 5 shows the NB distribution of predictions, given each actual pull angle class. This trend is demonstrative of all training method trends.

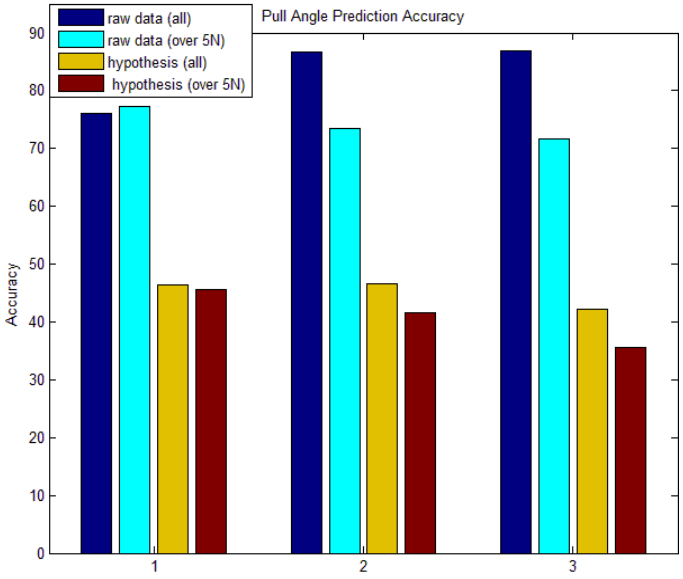

Figure 4: Percent accuracy for pull angle.
(1) Naïve Bayes, (2) Logistic Regression, (3) SVM.

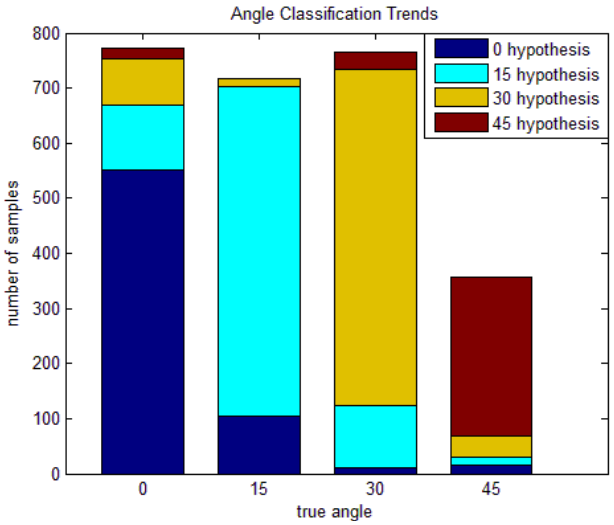

Figure 5: Angle classification distribution for NB (all raw data).

In this classification circumstance it is advantageous to keep low force data in training. Not only does LogReg and SVM work better when no data is excluded, but it gives the robot designer more information about small force object interactions. Also, it appears that the error of the Linear Regression hypothesis only hurts the accuracy of all three models.

3.3 Pull Force Classification Using Pad Sensors

Pull force was also classified using multivariate NB, SVM, and LogReg. Initial tests were run on all raw data with classes created from the gauge force reading rounded to the nearest integer for a total of 60 classes.

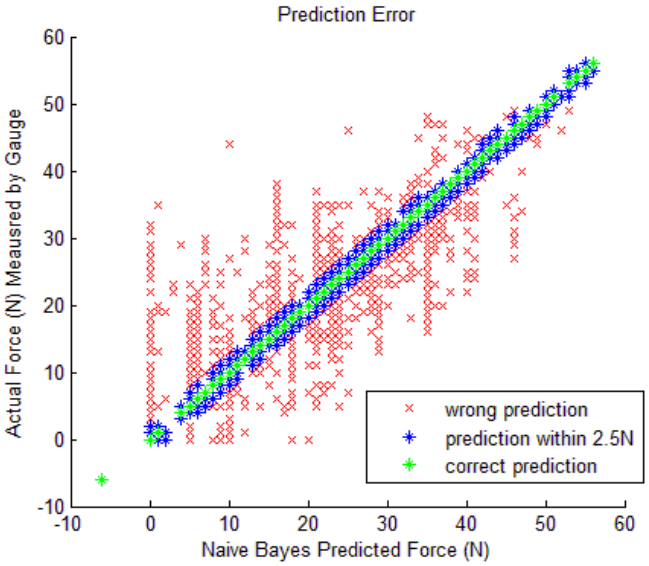

Figure 7: Naïve Bayes errors in prediction for 60 classes. The blue is within a 5N range around the correct hypothesis. Note: error seems to diminish at very high applied object forces. High forces usually correspond to a moving object. Therefore, it seems static frictional forces, or unchanging hand pose, may affect the reliability of pressure sensor readings.

Then seeking to improve tolerance, the force classes were binned into 5N ranges for a total of 13 classes from -5N to 55N. Both schemes were also trained using the Linear Regression. See Figure 6 for a summary of test results. Naïve Bayes gave the best result, 48.64% testing error, using all raw data with the 5N range binning scheme. Figure 7 shows the benefits of increasing class range. Figure 8 shows the NB distribution of predictions, given each actual pull force class.

3.3 Faulty Knuckle Angle Sensor Compensation

The knuckle angle sensors in the robotic hand were custom research prototypes. They were an entirely new design that worked by measuring a variable capacitance. They were designed to be absolute encoders but due to symmetry in the design they could be off by multiples of pi/4 radians.

Before we could use the knuckle angle sensor data to predict object size, we had to remove this sporadic offset. We could not simply subtract the initial value from all our trial runs because then we would lose information about the relationship between the different knuckle sensors.

In the end, we used a K-Means algorithm to match a set of means that were offset by pi/4 radians to the raw data. Then we normalized by these means to remove the offset.

The green dots in figure 9 show the knuckle data for knuckles one, two, and three from one finger grasping the 4.5in tube. The green line shows the density of points. It can be seen that the data is clustered around multiples of pi/4 radians.

The blue dots in figure 9 indicate the means that were fit to the raw data. A gradient ascent algorithm was used to find the best fit of the means to the data.

The original data was then normalized by the matched means. The red dots and line in figure 9 show the normalized data and the density of points respectively.

Figure 10 shows the knuckle angle data versus time for all trials before and after compensation of all three knuckles of the two fingers. Green lines show the angle of the first knuckle, the red lines show the second knuckle, and the blue lines show the third knuckle. As you can see, the measurements are much more consistent after compensation.

3.4 Object Size Prediction

After we compensated for the knuckle angle sensor offsets, we used an SVM algorithm to try to differentiate between when the hand is grasping the 3.5” versus the 4.5"diameter tube. We split the knuckle angle data into two bins: 70% for training, 30% for testing. Since the poses of the two fingers are roughly symmetric and the distal knuckle does not bend very much, we chose to

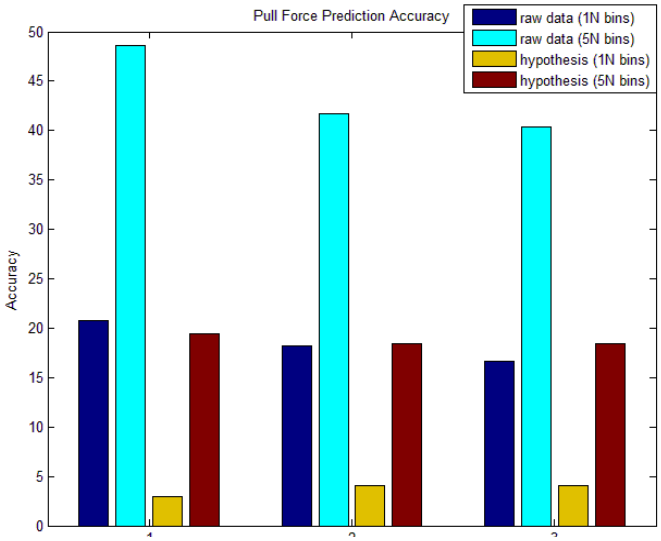

Figure 6: Percent accuracy for pull force.
(1) Naïve Bayes, (2) Logistic Regression, (3) SVM.

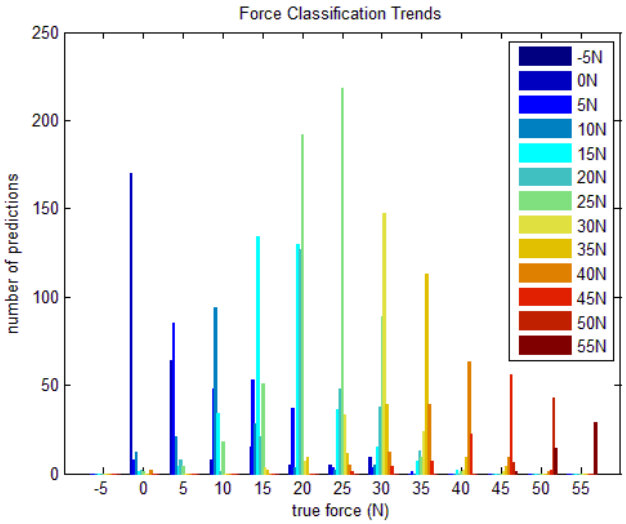

Figure 8: Force classification distribution for NB (5N bins).

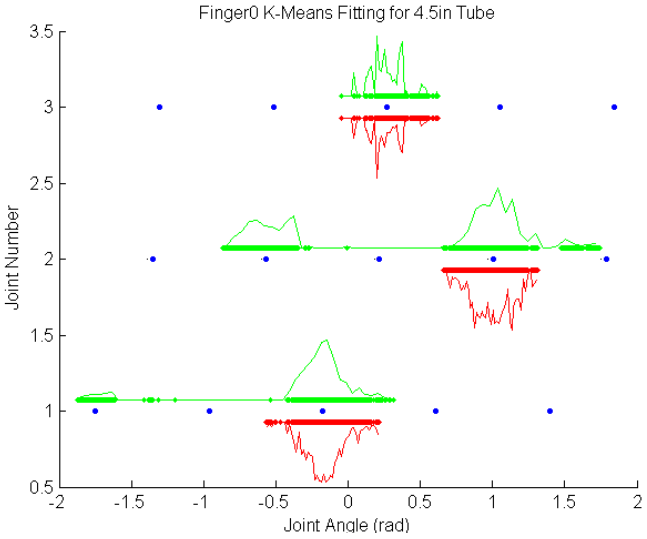

Figure 9: K-Means and Gradient Descent to Compensate for Fixed Knuckle Angle Sensor Offsets

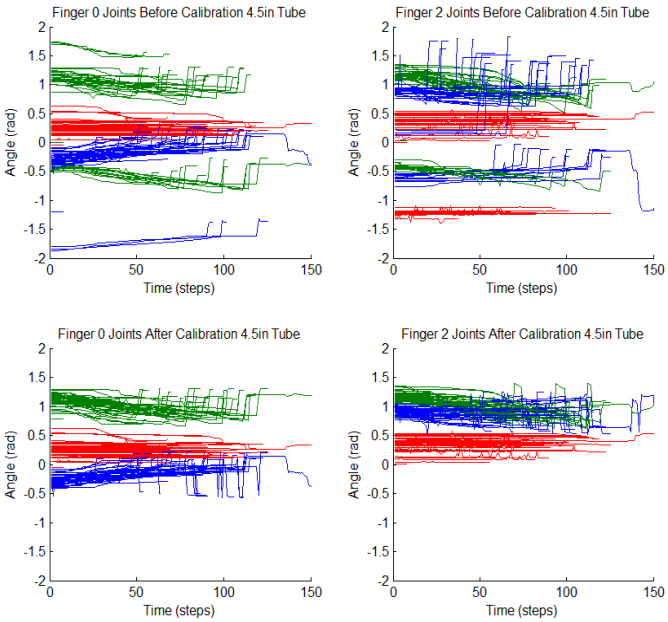

Figure 10: Original knuckle Angle Data (top) and Compensated knuckle Angle Data (bottom)

only use the knuckle angle data from the proximal and middle knuckle of the first finger to train our SVM.

The SVM results are shown in figure 11. The prediction accuracy was 94% for the test data. This is an exciting result because it suggests that a robotic hand could be trained to identify objects of different sizes with knuckle angle feedback alone.

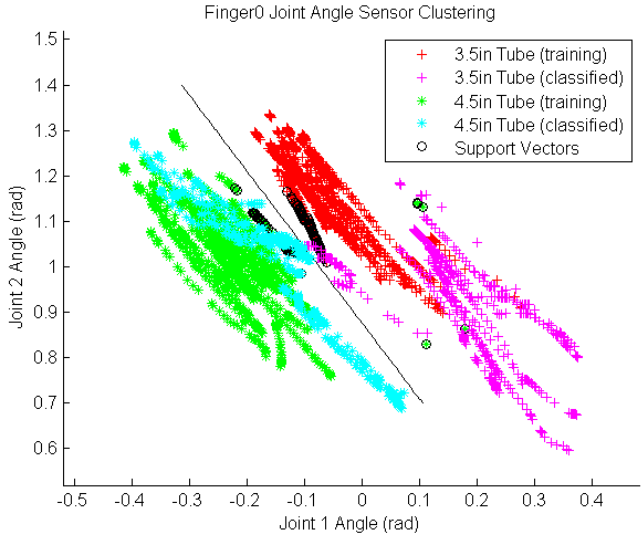

Figure 11: SVM of Joint Angle for 3.5" and 4.5" Diameter Tubes

6 Future Work

6.1 Data Collection

During data collect, human factors affected the pad sensors and should be more controlled. Improving data collection could lead classifying when pressure sensor data becomes faulty to adjust the hypothesis locally.

It would also be interesting to test and classify more object sizes and shapes. It is important to note that these machine learning methods have been trained for a specific test setup. For robots to function in unstructured environments, ongoing experimentation is vital.

6.2 Learning Algorithms

More leaning algorithms could be employed to improve fit. For example, the EM Algorithm could take into account pad sensor reliability.

Also, finding an adequate method to condition incoming finger pad data to make an accurate hypothesis of applied force and location on each phalanx may prove to cancel out specified noise to create a superior learning algorithm. For example using locally weighted regression over time may reduce the effects of erratic electrical connections.

7 Summary and Conclusion

The models developed in this project can be used to allow a robot to haptically explore its environment without visual feedback. For example, identifying an object in a bag by touch or estimating the weight of a object in the dark. Object size, pull force, and pull angle from knuckle angle sensor and finger pad pressure sensor data can be successfully predicted using machine learning algorithms. Even in situation where a faulty sensor might disrupt accuracy, the robustness of the algorithms still allows for reasonably reliable results.

Exciting future work could build on these techniques to help increase the capabilities of robots to perform in unstructured environments.

8 Acknowledgements

We would like to thank BDML members Mark Cutkosky, Dan Aukes, John Ulmen, and Barrett Heyneman for helping with the experimental setup and making this project possible. Thank you to SRI, DARPA, and NSF for past and present support. Finally, thank you to all the CS229 staff for all their hard work this quarter.

9 Works Cited

[1] Dipert, B.; Shoham, A.; , "Eye, Robot: Embedded Vision, the Next Big Thing in Digital Signal Processing," *Solid-State Circuits Magazine, IEEE* , vol.4, no.2, pp.26-29, June 2012

[2] Jefferson Coelho, Justus Piater, Roderic Grupen, "Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot," Robotics and Autonomous Systems, Volume 37, Issues 2–3, 30 November 2001, Pages 195-218.

[3] Gorges, N.; Navarro, S.E.; Göger, D.; Wörn, H.; , "Haptic object recognition using passive joints and haptic key features," *Robotics and Automation (ICRA), 2010 IEEE International Conference on* , vol., no., pp.2349-2355, 3-7 May 2010