

Learning Grasp Stability

Hao Dang and Peter K. Allen

Abstract—We deal with the problem of *blind grasping* where we use tactile feedback to predict the stability of a robotic grasp given no visual or geometric information about the object being grasped. We first simulated tactile feedback using a soft finger contact model in *GraspIt!* [1] and computed tactile contacts of thousands of grasps with a robotic hand using the Columbia Grasp Database [2]. We used the K-means clustering method to learn a contact dictionary from the tactile contacts, which is a codebook that models the contact space. The feature vector for a grasp is a histogram computed based on the distribution of its contacts over the contact space defined by the dictionary. An SVM is then trained to predict the stability of a robotic grasp given this feature vector. Experiments indicate that this model which requires low-dimension feature input is useful in predicting the stability of a grasp.

I. INTRODUCTION

Many of the current grasp planning algorithms utilize visual information [3], [4], [5] or geometric data [6], [7] of the object to be grasped. To obtain such information, devices like cameras or range finders are necessary. This, to a large extent, limits the application of these methods across a wide range of different working situations, considering different lighting conditions and object occlusion.

Pushing this problem to an extreme, one would ask can a robotic hand still apply a stable grasp to an object when it is blind? An example that describes the situation could be a robotic hand reaching into a gym bag and getting a bottle of water out of it. As human beings, it is intuitive and straightforward for us to grasp objects even when we cannot see them. Lederman and Klatzky have shown [8] that humans have the ability to “blindly” recognize objects with a high degree of accuracy. Humans can also apply stable grasps on unknown objects in the total absence of any visual feedback. But this ability is sorely lacking with current robots that are performing grasping tasks. Using tactile exploration for grasping, Bierbaum *et al.* proposed a method to generate grasp affordances based on reconstructed faces of an object through tactile exploration [9]. Platt proposed a method to learn grasping strategies based on contact relative motions [10]. Hsiao used tactile sensing to reduce the localization uncertainty during grasping [11]. We are focusing on predicting the stability of a grasp assuming an exploration algorithm has already applied a grasp to an object which is unknown to the robot.

In our previous work [12], we used raw tactile readings and hand kinematic data together to predict the stability of

a grasp. This method achieved a low false positive rate on unstable grasps which is a very necessary feature in working environments where unstable grasping is risky and costly. Although the raw tactile and kinematic data carry a good amount of useful information to indicate the grasp stability, some properties of a grasp may not be easily extracted from the raw sensor data. One important property is the 3-dimensional spatial distribution of the contacts between the hand and the object, which to a very large extent influences the stability of a grasp. The feature we used in our previous work is based on the raw tactile and kinematic data. Thus, the dimension of a feature vector could be very large given a hand with a large number of tactile sensors or joints, which increases the learning complexity. This poses a potential limit to the learning method in practice. In this paper, we utilize a bag-of-words model to describe the 3-dimensional spatial distribution of the contacts of a given grasp and predict the stability of the grasp with this model, which reduces the dimension of the feature space and achieves a better accuracy. In addition, instead of using the epsilon quality of a grasp as the ground truth as in our previous work, we take into consideration the volume quality of a grasp and combine these two criteria together as the ground truth. Together, they can give a more comprehensive stability evaluation than the epsilon quality alone.

Another important work related to ours is by Bekiroglu *et al.* [13] and Laaksonen *et al.* [14]. They have a close look into assessing the grasp stability using static one-shot classification and dynamic time series assessment during the grasping procedure. Compared to their work, we propose a different feature based on tactile contacts that inherently encodes the information of the contact configuration extracted from the tactile data and finger joint values together. We trained our classifier on a larger data set of objects with more shape varieties compared to the work in [13] and achieved a better accuracy in a model-free classification scenario. In addition, we looked into how the grasping performance could be influenced by varying the surface friction coefficient and the object mass.

In Section II, we discuss the extraction of contacts from the tactile feedback and hand kinematics. In Section III, we propose a bag-of-words model for constructing feature vectors of grasps. We show the experiments in Section IV, followed by conclusions in Section V.

II. CONTACTS FROM TACTILE FEEDBACK

When visual information is not available, tactile feedback from the hand is crucial in object grasping and manipulation tasks. It usually gives information about the object’s local

This work has been funded by NSF Grant IIS-0904514 and the DARPA ARM-S program. All authors are with Department of Computer Science, Columbia University, 450 Computer Science Building, 1214 Amsterdam Avenue, New York, NY, 10027, USA, {dang, allen}@cs.columbia.edu

geometry which is difficult to obtain through vision alone. Tactile sensors play an important role in representing the contacts between the surface of the hand and the object that are touching each other. The output of the tactile sensors around each contact is characterized by the forces applied at each sensor cell. So, a reasonable contact model that approximates the contact region and the pressure distribution is necessary for simulating a reasonable tactile feedback. Pezzementi et al. [15] used a point spread function model to simulate the response of a tactile sensor system. In our work, we build our tactile simulation system based on a soft finger contact model proposed by Ciocarlie et al. [16]. Interested readers please refer to the original papers for more details [16], [12].

Tactile feedback indicates which sensor cells have contacts with the object and which do not. With the angle values for the joints of the hand, we can also use forward kinematics to determine both the location and the orientation of each sensor cell. So, we can utilize the tactile feedback to approximate the contact locations and orientations.

To represent the location and the orientation of a sensor cell, we want to use a coordinate system that is local to the hand and is consistent across different grasps. We chose the coordinate system attached to the palm as the reference coordinate system. Given a set of n joint angles of a grasp, $\mathcal{J} = [j_1, j_2, \dots, j_n]$, we write out the location and the orientation of the k^{th} sensor cell on the i^{th} link in a homogeneous transformation matrix as follows:

$$T_{palm}^{sensor_{ik}}(\mathcal{J}) = T_{palm}^{link_i}(\mathcal{J}) \cdot T_{link_i}^{sensor_k} \quad (1)$$

where $T_{palm}^{link_i}$ denotes the transformation between the link i and the palm; it is determined by the joint values \mathcal{J} and the hand kinematics for each grasp; $T_{link_i}^{sensor_k}$ is the transformation between the i^{th} link and the k^{th} sensor cell on this link; it is determined by the sensor cell configuration and is a constant for every grasp.

Using the location and the orientation of each sensor cell that is activated due to a contact, we can determine the configuration of the contacts involved in a grasp. It is worth noting that there is error in representing the actual contact locations and orientations using this method because each sensor cell has finite dimensions and any contacts residing within the same sensor cell will be indistinguishable.

III. A BAG-OF-WORDS MODEL

Bag-of-words models are widely used in natural language processing (NLP). They are also known as bag-of-features models in computer vision. In NLP, Bag-of-words models use a dictionary to represent a document without considering the order of the appearance of the words in the document. In computer vision, an image is treated similarly as a document, where the visual features of an image take the role of the words in a document.

By the same analogy, we can transfer this idea to the context of robotic grasping. A grasp contains a set of contacts just as a document consists of a number of words. If we treat a grasp as a document and a contact as a word, it is




Fig. 1. Cluster centers of contacts overlaid on a Barrett hand. Spheres are located at the centers of each cluster. The clusters contain 199,835 contacts collected from a training set of 24,640 grasps.

reasonable to use a bag-of-words model to describe a grasp in a similar way as a bag-of-words model does a document.

A. Building a Contact Dictionary

In order to use the Bag-of-words approach, we need to build a dictionary which represents the space of the potential contacts. It is impractical, if not impossible, to collect all the possible contacts that can appear in a grasp. Thus, we need a reasonable discretization of the space within the hand's coordinate system. Considering the hand kinematics of a robotic hand, we see there are some regions within the hand's local coordinate system that have larger potential than other regions for a contact to appear. Thus, using a set of representative contacts from these regions as a dictionary, we can enjoy both the statistically sound capability of representation of the contact space and a low dimensionality of the dictionary which determines the dimensionality of the feature of a grasp.

In order to learn a set of representative contacts, we first collected a training set of grasps and extracted the contacts out of them. Considering the fact that the reaching space for each finger of a Barrett hand intersects with each other very rarely, it is safe to think each contact location is associated with a contact orientation and vice versa. Based on this consideration, we only used the location part of a contact and applied a K-means clustering algorithm to all the contact locations to learn this set of representative contacts. Thus, the space we are clustering is a regular 3-dimensional Cartesian space. The distance function we chose in the K-means clustering algorithm measures the euclidean distance between two contact locations.

Figure 1 shows the cluster centers overlaid on a Barrett hand. The spread angle of the Barrett hand is set manually solely for giving a better idea of the hand's work space. The centers of the clusters outline the reaching space of each finger. In Figure 1(a), the contact spaces of finger 1 and 2 of the Barrett hand display very nice symmetry. This agrees with the symmetric mechanical design of the two fingers.

B. Grasp Feature Vectors

The set of cluster centers models the space of the contacts on the fingers and the palm in a highly discretized dimension. With this set of cluster centers, we use the distribution of the contacts among the contact cluster centers as feature vectors for grasps.

Formally, given a contact dictionary with of p cluster centers $\hat{C} = [\hat{c}_1, \hat{c}_2, \dots, \hat{c}_p]$ and a grasp \mathcal{G} which consists of

q contacts $\mathcal{C}_G = [c_1, c_2, \dots, c_q]$, we calculate the distribution vector of the contacts of grasp \mathcal{G} with respect to $\hat{\mathcal{C}}$ as follows:

$$\mathcal{D}(\mathcal{C}_G, \hat{\mathcal{C}}) = \sum_{i=1}^q \mathcal{H}(c_i, \hat{\mathcal{C}}) \cdot \frac{f_{c_i}}{S_{c_i}} \quad (2)$$

where f_{c_i} is the force value sensed at the tactile sensor cell corresponding to contact c_i , S_{c_i} is the total amount of forces sensed from all the sensors cells on the sensor pad which contact c_i is on, $\mathcal{H}(c_i, \hat{\mathcal{C}}) = [h_1, h_2, \dots, h_p]$ is a p -dimensional vector that stores the similarity values between contact c_i and each cluster center in $\hat{\mathcal{C}}$. It is computed as:

$$h_i = \exp\left(-\frac{\|c_i - \hat{c}_i\|^2}{\sigma^2}\right) \quad (3)$$

where c_i and \hat{c}_i are both 3-dimensional vectors storing the contact locations and σ is a parameter set manually. h_i measures the similarity between two contact locations. For a contact that is far from a cluster center, the corresponding euclidean distance is large resulting in a small similarity value. The parameter σ controls the speed the similarity values decrease as the distances go up. We want to use a reasonable σ that would maximize the range of h_i values in $[0, 1]$ and maximally distinguish the contacts by this similarity value. In our experiment, we analyzed the range of the contacts from all the grasps of a Barrett hand and set σ as 36.45.

Distribution vectors $\mathcal{D}(\mathcal{G}, \mathcal{C}_G)$ are built from the summation of the distributions of different numbers of contacts. Thus, we scale the distribution vector down by the number of sensor pads which have tactile contacts. Mathematically, for a grasp \mathcal{G} , the scaled distribution vector is calculated as:

$$\hat{\mathcal{D}}_i = \frac{\mathcal{D}_i}{|\mathcal{P}_G|} \quad (4)$$

where \mathcal{D}_i denotes the i^{th} element of the vector $\mathcal{D}(\mathcal{G}, \mathcal{C})$ and $|\mathcal{P}_G| \in \{1, 2, 3, 4\}$ is the number of sensor pads with tactile contacts. This scaled distribution vector is the final feature vector we use to describe a grasp in our work.

IV. EXPERIMENT

Following the explanation of the model we use to generate features of grasps, we now discuss the way we generate a set of grasp samples and the experiments.

A. Grasp Dataset

Our grasp data is from the Columbia Grasp Database (CGDB) [2]. This database contains hundreds of thousands of grasps constructed from several robotic hands and thousands of object models. To build this database, we used *GraspIt!* [1], a robotic grasping simulator to plan grasps and evaluate their qualities. Although the CGDB provides grasps for all these object models in the PSB, instead of using the full set of grasps in the CGDB, we choose to select grasps computed on a smaller set of objects that are more frequently grasped and manipulated by us in our everyday life. In total,

we collected about 36,960 robotic grasps from 704 objects across 19 different classes. We used $\frac{2}{3}$ of them for training and the remaining $\frac{1}{3}$ for testing.

B. Grasp Quality Measurements

Different measurements could be used to evaluate the quality of a grasp, such as its stability, feasibility, and dexterity [17]. In our work, we focus on two quality measurements related to the stability of a grasp: the epsilon quality and the volume quality [18], which are based on the grasp wrench space (GWS) generated by the grasp. In our experiment, we modeled the material of the surface of the hand as rubber and the material of the object as wood and set the friction coefficient between the finger and the object as 1.0.

The epsilon quality, ϵ , refers to the minimum relative magnitude of the outside disturbances that could destroy the grasp. So, when we take into account the limit of the maximum forces a robotic hand can apply, a grasp would be less stable if it has a smaller epsilon quality. This is because the smaller epsilon quality indicates that a relatively smaller outside disturbance can break this grasp even when the robotic hand has already applied the maximum forces it supports. Another consideration is from the perspective of the environment uncertainty. Due to the uncertainty of the environment, objects may move away slightly from their original position during a grasp execution. A fragile grasp may fail to fully grasp the object in this situation while a stable one may display its robustness and still succeed in grasping the object in the perturbation. We have experimentally found a strong correlation between this robustness and the epsilon quality. We have found that grasps with epsilon quality $\epsilon > 0.07$ tend to be more robust in uncertain object perturbations.

The volume quality, v , measures the volume of the potential wrench space generated by the grasp given unit contact normal forces. A grasp with a larger potential wrench space would require less forces at each contact than grasps with smaller potential wrench spaces. This indicates that the larger the volume quality is, the stronger the grasp could be.

C. Labeling Grasps in the Dataset

Given all the grasps in the grasp dataset, we observed that their epsilon qualities and volume qualities are not well correlated. In addition, Li and Sastry pointed out the epsilon quality measure is not invariant to the choice of torque origin [19], so we use the volume quality as an invariant average case quality measure for the grasp. Each of these measures has its own benefits. It makes sense that we combine them together and form up a new evaluation criterion. Based on our experimental results, we use thresholds $t_\epsilon = 0.07$ and $t_v = 0.1$ as the boundaries for epsilon and volume qualities to label a grasp $grasp_i$ as a good (1) if and only if their corresponding epsilon and volume qualities are above the threshold.

D. Simulation Experiment

In our experiment, the dictionary was learned based on the training set mentioned in Section IV-A. Experimentally,




Fig. 2. Accuracy analysis: X-axis is the group names for each object class. Y-axis is the percentage (%) of the overall false predictions (green bars) or the false positive predictions (yellow bars) per object class. As shown in the graph above, the percentages for false positive predictions per object class (yellow bars) are relatively low, which is necessary for blind grasping.

TABLE I
LEARNING PERFORMANCE IN SIMULATION

Data set	grasps	stable grasps	unstable grasps	accuracy
Training	24,640	11,849 (48.09%)	12,791	-
Test	12,320	5,914 (48.00%)	6,406	81.0%

we chose $k = 64$ to create 64 different clusters. So, the dictionary contains 64 contacts and a feature vector of a grasp is 64-dimensional. Figure 1 shows these 64 cluster centers within a Barrett hand’s coordinate system. We then used *libsvm* [20] to train an SVM based on the training data which contains about 24,640 grasps and tested the SVM on the remaining 12,320 grasps. Table I summarizes the number of stable grasps and unstable grasps in both training and test datasets. Figure 2 shows the classification result in more detail. The overall accuracy across all the classes of objects is 81.0%. Considering the context of blind grasping, the percentage of the false positive predictions is a crucial evaluation criterion. This is because a false positive prediction will guide the robot to apply an unstable grasp and use it as if it is a stable one. In most working conditions, this action is very risky and even unacceptable. In Figure 2, we show the percentages of error and false positive predictions for each object class. The percentages of false positive predictions illustrate the probabilities of the situation when an unstable grasp is incorrectly classified as a stable one during a grasping task. As can be seen, for most of the object classes, the probabilities of false positive predictions are relatively low. The overall false positive prediction is 8.6%.

E. Physical Experiment

To evaluate the performance of the classifier trained with simulated data in physical grasping scenarios, we did some experiments with a Barrett hand on six everyday objects: a pencil cup, a mug, a candle box, a paper wipe box, a canteen, and a decorative rock as shown in Figure 3(a). Only the mug belongs to a object class that is included in the simulated training data. The pencil cup, the candle box, and the paper wipe box are objects to some extent similar to the bottle class in the training set. The canteen and the decorative rock are two objects that are very different from other objects in the training set.

The Barrett hand is equipped with four tactile pads, one for

the distal link of each finger and one for the palm, resulting in a 96-sensor system. We place an object at a pre-defined location on a table that is in front of the robot. During a grasping experiment, the robot approaches the pre-defined location with different spread angles. When the robot hits the object, it stops approaching and closes the fingers. Tactile data and joint angles are then collected for grasp stability estimation using the classifier trained from Section IV-D. The arm lifts up the object once a stable grasp is perceived. A trial is considered to be a failure when the robot is not able to grasp the object stably, i.e., an object falls out of the hand when the robot tries to lift the object up.

110 trials were performed on six different objects, including a canteen of different weight and surface material, and a mug filled with different weights. In Table II, we summarize the experiment results. Overall, the success rate is 84.6% across all the objects in our experiments. The canteen without its fabric cover and the decorative rock are two objects that are very slippery and difficult to grasp. Compared to other objects, there are much fewer stable grasps on the canteen without its fabric cover and the decorative rock. In addition, the decorative rock is convex and a very large proportion of the surface is facing upwards to some extent, thus frictional force becomes the main source of forces to overcome the gravity during grasping, making the decorative rock an even more difficult target. In Section IV-F, we will discuss in detail the potential limitation of our method and the lessons we learned from failure cases.

The classifier was trained based on some general assumptions. In the physical experiment, the general success rate is reasonably good. Yet we also looked into two object properties that, we believe, have potential impact on the grasping performance: surface friction coefficient and mass.

1) *Different surface friction coefficient*: Since the quality of the grasps in our training data is calculated considering the closure properties of the GWS, surface friction coefficient is an important factor in the learning process. In generating the training data, we used a unit friction coefficient. Considering different surface materials of objects in everyday life, this is not ideal in practice. Although we see the grasping performed reasonably well using this assumption, we are still interested in seeing the influence brought by the difference in surface materials. Thus, we performed a quantitative experiment on grasping objects with different surface materials, mimicking different frictional coefficients. In our experiment, we chose a canteen with and without its fabric cover to mimic different surfaces. The canteen is made of plastic surface and it is slippery without a fabric cover. But with its fabric cover, the surface friction coefficient can be considered much larger than before. From Table II, we can see the grasping success rate on the canteen with its fabric cover is higher than on the canteen without its fabric cover. This also indicates that designing tactile sensors with materials of higher friction coefficient would benefit the grasping procedure.

2) *Different weight*: In the general experiment, the mass of the objects vary from 0.1 to 0.5 kg. Since the torque each motor in the Barrett hand can exert is limited, the mass of

TABLE II
EXPERIMENT RESULTS ON SIX OBJECTS

Object	Mass (kg)	# of exp	Success	Success rate
Mug	0.43 - 0.93	30	28	93%
Paper wipe box	0.17	10	9	90%
Pencil cup	0.09	10	9	90%
Candle box	0.11	10	9	90%
Decorative Rock	0.28	10	6	60%
Canteen	0.5 - 0.75	40	32	80%
Total	0.09 - 0.93	110	93	84.6%

TABLE III
GRASP A CANTEEN WITH DIFFERENT SURFACE AND WEIGHT

Object	half full (0.50 kg)	full (0.75 kg)
canteen w/o cover	90%	60%
canteen w/ cover	100%	70%

an object will influence the grasping performance. To see how much the influence could be, we did some experiments with two selected objects: a canteen filled with different amount of water and a mug filled with different weights. The performance on grasping the canteen is summarized in Table III and the performance on grasping the mug is summarized in Table IV. From Table III, we can see the grasping performance on the canteen decreases as the weight increases. This indicates that when the weight of the object increases the grasping performance decreases. But we did not see this tendency on grasping the mug. We think one major reason is that the surface friction between the mug and the tactile sensor is very big and the frictional force is sufficient to overcome the influence brought by the increase in the weight. This result also indicates that the influence brought by the weight and the surface friction can be coupled.

F. Discussion

Based on our experiments, we discuss our method and its potential limitation further as follows:

1) *Surface Normal*: We made an assumption that the contact normals can be approximated by the surface normals of the fingertip and the palm. For many cases, this is reasonable considering the local contact surface is almost flat. However, when a contact is right on an edge or a curved area, the contact normals cannot be determined easily from the surface normals. For a contact on an edge, the normal could range between the normals of the two sides of the edge. for a contact on a curved surface, the normals can vary among all the surface normals. In these situations,

TABLE IV
GRASP A MUG WITH DIFFERENT WEIGHT

Object	0.43 kg	0.68 kg	0.93 kg
mug	100%	80%	100%

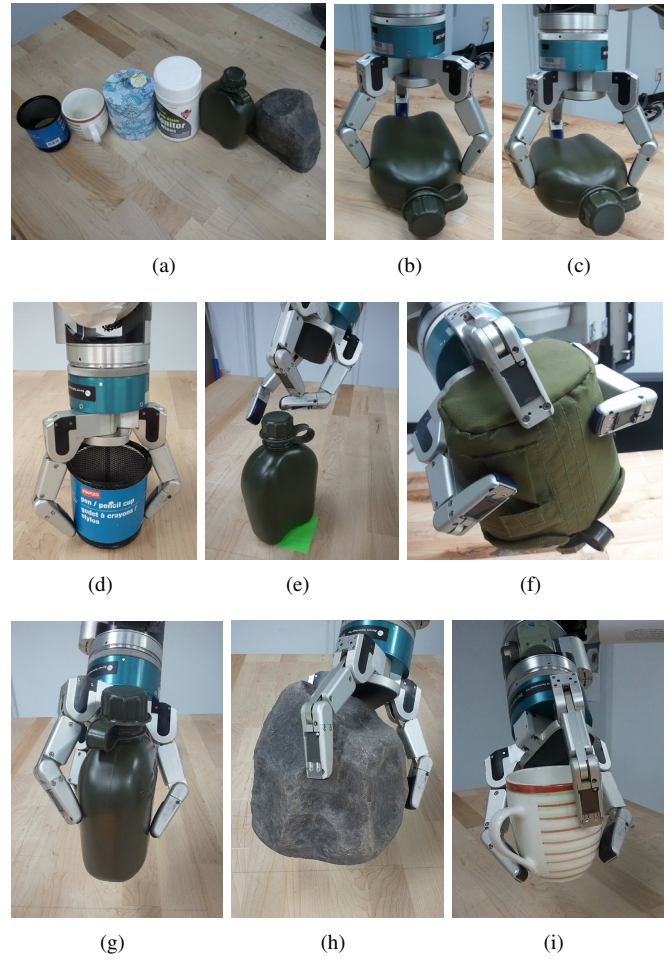


Fig. 3. Physical experiments. Figure 3(a) shows the six objects in the experiments: (left to right: pencil cup, mug, candle box, paper wipe box, canteen, decorative rock). Figure 3(b) and 3(c) show two grasps that have very similar set of contacts in hand's coordinate system while they have different true stability measurements. Figure 3(d) shows a grasp that is falsely classified as an unstable grasp due to incomplete coverage of the tactile sensors. Figure 3(e) is a false positive case resulted from misperceived tactile contacts while two fingers touch each other. Figure 3(f) shows a grasp that is falsely predicted as an unstable grasp due to the lack of knowledge of the object surface friction coefficient. Figure 3(g) to 3(i) show three stable grasps on three different objects in the experiment.

our assumption does not hold and this could contribute to false positive predictions. One example is in Figure 3(b). Compared with the grasp in Figure 3(c), the grasp in Figure 3(b) contains almost the same set of contacts in the hand's coordinate system. So, for the SVM, it would classify them as the same in stability. However, the contacts of the grasp in Figure 3(b) exist between the curved edge of the finger tip and the curved surface of the canteen while the grasp in Figure 3(c) has its contacts more on the paralleled sides of the canteen, making the former grasp unstable and the latter stable. We think it is beneficial to have tactile sensors with smaller sizes cover the edges so that we can distinguish contacts on the edge or to approximate the force normal at a contact directly from the sensor values around the contact.

2) *Sensor Coverage*: In our experiment, we see that the tactile pads do not cover all the places that could have contacts. Due to an incomplete coverage, some contacts of

a grasp that actually contribute to the grasp stability may not be captured. When this situation happens, our method could only take into consideration the contacts captured for the stability analysis. Thus, incomplete input information would increase the possibility of a false prediction. In our physical experiment, we found this contributed heavily to false negative predictions. Figure 3(d) shows a grasp on a pencil cup which is relatively small compared to the Barrett hand. This grasp is a stable grasp, but the classifier did not predict it as a stable one. The reason is that only the contacts on the finger tips are perceived and those contacts on the proximal links of the fingers are not captured. In fact, incomplete tactile coverage on the hand surface is also one of the primary differences between a human hand and a robot hand. We think a design of a tactile system with a more comprehensive coverage would alleviate this issue.

3) *Finger Collision*: When fingers touch each other, tactile contacts may mislead the classifier in stability estimation. Figure 3(e) shows one such grasp. In this grasp, only finger II and finger III were touching the canteen. Finger I touches finger III and the tactile pads perceived this contact. Thus, the classifier was confused and mistakenly considered this grasp as a stable one. To alleviate this issue, we believe a self-collision checking is helpful and necessary.

4) *Surface Friction*: We assumed a surface friction coefficient as 1.0 in Section IV-B. This turned out to be a limit of our method across different kinds of surfaces. In our physical experiment, we saw some false negative predictions resulted from this factor. In Figure 3(f), we show one example grasp that is classified as unstable. In reality, it is a stable grasp. One reason for the false prediction is the actual surface friction coefficient is more than what is assumed in our simulation. To alleviate this issue, we think one approach could be to use classifiers trained with different datasets generated by different friction coefficients and probabilistically unify them.

5) *Sensor Sensitivity*: In our experiment, we found that the sensitivity of the tactile sensors is also a critical factor in the final grasping performance. We observed that due to the lower sensitivity of the tactile sensors, signal values at some valid contacts cannot exceed the sensor noise range. Thus these contacts cannot be detected and captured, resulting in a false negative prediction. We believe that the grasping performance will benefit a lot from more sensitive sensors.

V. CONCLUSION AND FUTURE WORK

In this paper, we proposed a model for predicting the stability of a robotic grasp using only tactile and kinematic information. This model describes the contact spatial distribution within the hand's local coordinate system. As our ground truth criterion, we considered both the epsilon quality and the volume quality at the same time. This combined rule gives a more comprehensive evaluation of the stability of a grasp. Physical experiments were done and analyzed to evaluate the generalization capability of the classifier trained based on simulation data. Currently, we are working on designing a control algorithm to adjust the hand

pose during the autonomous exploration procedure for stable grasps instead of randomly re-orienting the palm to approach the object and shaping the spread joint.

VI. ACKNOWLEDGMENT

The authors gratefully acknowledge the help of Jonathan Weisz from Columbia University, Motilal Agrawal and Aravind Sundaesan from SRI International.

REFERENCES

- [1] A. T. Miller and P. K. Allen, "Graspi! a versatile simulator for robotic grasping," vol. 11, no. 4, pp. 110–122, 2004.
- [2] C. Goldfeder, M. Ciocarlie, H. Dang, and P. Allen, "The columbia grasp database," in *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pp. 1710–1716, may 2009.
- [3] A. Saxena, J. Driemeyer, J. Kearns, and A. Y. Ng, "Robotic grasping of novel objects," in *Advances in Neural Information Processing Systems 19* (B. Schölkopf, J. Platt, and T. Hoffman, eds.), pp. 1209–1216, Cambridge, MA: MIT Press, 2007.
- [4] M. Popovic, D. Kraft, L. Bodenhausen, E. Baseski, N. Pugeault, D. Kragic, T. Asfour, and N. Krüger, "A strategy for grasping unknown objects based on co-planarity and colour information," *Robotics and Autonomous Systems*, vol. 58, no. 5, pp. 551–565, 2010.
- [5] D. Bowers and R. Lumia, "Manipulation of unmodeled objects using intelligent grasping schemes," *Fuzzy Systems, IEEE Transactions on*, vol. 11, no. 3, pp. 320–330, 2003.
- [6] M. T. Ciocarlie and P. K. Allen, "Hand Posture Subspaces for Dexterous Robotic Grasping," *The International Journal of Robotics Research*, vol. 28, no. 7, pp. 851–867, 2009.
- [7] A. Morales, T. Asfour, P. Azad, S. Knoop, and R. Dillmann, "Integrated grasp planning and visual object localization for a humanoid robot with five-fingered hands," in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pp. 5663–5668, 2006.
- [8] S. J. Lederman and R. L. Klatzky, "Hand movements: A window into haptic object recognition," *Cognitive Psychology*, vol. 19, no. 3, pp. 342–368, 1987.
- [9] E. Bierbaum, M. Rambow, T. Asfour, and R. Dillmann, "Grasp affordances from multi-fingered tactile exploration using dynamic potential fields," in *Humanoids, IEEE Intl. Conference on*, pp. 168–174, 2009.
- [10] R. Platt, "Learning grasp strategies composed of contact relative motions," in *Humanoid Robots, 2007 7th IEEE-RAS International Conference on*, 29 2007.
- [11] K. Hsiao, L. Kaelbling, and T. Lozano-Perez, "Robust grasping under object pose uncertainty," *Autonomous Robots*, vol. 31, pp. 253–268, 2011. 10.1007/s10514-011-9243-2.
- [12] H. Dang, J. Weisz, and P. K. Allen, "Blind grasping: Stable robotic grasping using tactile feedback and hand kinematics," in *Robotics and Automation, 2011. ICRA 2011. Proceedings 2011 IEEE International Conference on*.
- [13] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *Robotics, IEEE Transactions on*, vol. 27, pp. 616–629, june 2011.
- [14] J. Laaksonen, V. Kyrki, and D. Kragic, "Evaluation of feature representation and machine learning methods in grasp stability learning," in *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on*, pp. 112–117, dec. 2010.
- [15] Zachary Pezzementi, Erica Jantho, Lucas Estrade, Gregory D. Hager, "Characterization and Simulation of Tactile Sensors," in *Haptics Symposium*, March 2010.
- [16] M. Ciocarlie, C. Lackner, and P. Allen, "Soft finger model with adaptive contact geometry for grasping and manipulation tasks," *World Haptics Conference*, vol. 0, pp. 219–224, 2007.
- [17] R. Suárez, M. Roa, and J. Cornella, "Grasp quality measures," tech. rep., Technical University of Catalonia, 2006.
- [18] A. Miller and P. Allen, "Examples of 3d grasp quality computations," pp. 1240–1246, 1999.
- [19] Z. Li and S. Sastry, "Task-oriented optimal grasping by multifingered robot hands," *Robotics and Automation, IEEE Journal of*, vol. 4, pp. 32–44, Feb. 1988.
- [20] C.-C. Chang and C.-J. Lin, *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.