Modeling the Effects of Contact Sensor Resolution on Grasp Success

Qian Wan[1], and Robert D. Howe[1,2]

***Abstract*—This paper presents a quantitative method for analyzing the effect of sensor resolution on grasp stability prediction. Resolution limits for contact sensors are expressed as a range of contact locations and contact surface normals that cannot be disambiguated by the sensor. Grasp quality is assessed at the limits of this range to determine whether the uncertainties caused by sensor resolution lead to uncertainty in grasp outcome prediction. The analysis also enables calculation of the specific contact locations on an object where the tactile sensors are trustworthy and where the object is reliably graspable. Our approach lays the foundation for quantitative evaluation in design tradeoffs in sensor choices and sensor layout, as well as finger shapes and materials.**

***Index Terms*—Grasping, Force and Tactile Sensing, Perception for Grasping and Manipulation**

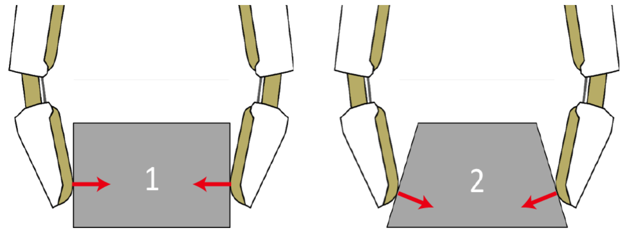

Fig. 1: Robot hand using two-fingered pinch grasp. With a moderate coefficient of friction objects 1 and 2 can both be stably grasped. If, however, the contact sensors in the robot fingertip lack sufficient resolution, then the surface normal direction for object 2 may be so uncertain that grasp stability cannot be reliably predicted.

I. Introduction

TACTILE sensors are crucial for providing feedback at the contact during grasping and manipulation tasks. The quality of these sensors has direct impact on the effectiveness of the signals and ultimately the task performance. Inevitably all sensors have limitations, such as limited range of coverage, sensitivity, and resolution. Understanding how these limitations effect task performance is crucial for achieving desired performance.

Take the task of predicting grasp stability as an example: Highly reliable grasping is essential in many real-world robotics applications. A household robot that achieves only 99.9% success in grasping will still drop many objects each week, making it unacceptable for most potential users. Assessing the stability of a grasp before the object is lifted is key to avoiding the costly outcome of dropping the object. Making that assessment reliable enough to meet the demanding requirement for commercial products will likely push the current sensor technology limits. However, little has been done to investigate the effect of tactile sensor's limitations on the robot's ability to evaluate grasp stability.

In this paper, we examine the effect of contact sensor's spatial resolution on grasp stability prediction, and develop a method that produces a quantitative relationship between the spatial resolution of the sensors and the reliability of the grasp stability prediction.

In theory, there are many well-developed analytical methods for predicting grasp stability [1], [2]. Most involve determining whether the forces exerted on the object by the fingers, the environment (particularly gravity), and the task are in equilibrium. For simple lifting tasks using fingertip precision grasps, stability can be calculated from contact locations, object surface normals or forces, coefficient of friction, the object's mass, and center of mass.

The implementation of these theories usually calls for perfect sensing, which is never the case in the real world. Real sensors have limitations such as finite spatial resolution, inherent sensor noise, and bounded sensitivity. These limitations can lead to uncertainties in the parameters used for calculating grasp stability, which eventually can compromise the stability prediction (Fig.1). Understanding how sensor uncertainties propagate through the prediction algorithm is key in knowing the reliability of the predictions. It will also help identifying the sensor improvements needed to reach desired performance.

Using grasp theory as a guide, we will relate the tactile sensors' spatial resolution limitation to uncertainties in contact surface normal and location — parameters that are essential in calculating the force and moment equilibrium of a grasp — and draw quantitative conclusions regarding the accuracy of stability prediction as a function of tactile sensor resolution.

Our analysis provides, for the first time, a quantitative relationship between a tactile sensor quality parameter (spatial resolution) and performance in grasping and manipulation tasks. The insights gained here can be used to improve both physics-based and data-based stability prediction, hand and sensor design, control, and planning. While we focused on

Manuscript received: September 10, 2017; Revised January 2, 2018; Accepted February, 2018.

This paper was recommended for publication by Editor Han Ding upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by National Science Foundation under grant IIP-1500178.

[1]Qian Wan and Robert D. Howe are with the Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA, 02138 qwan,howe@seas.harvard.edu

[2]Robert D. Howe is with RightHandRobotics, Inc. Somerville, MA, USA 02143

Digital Object Identifier (DOI): see top of this page.

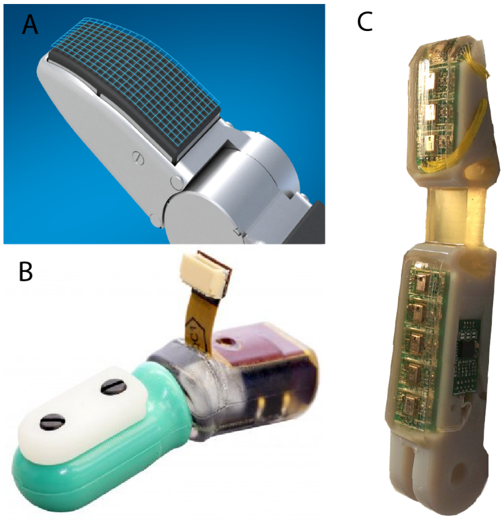

Fig. 2: Examples of integrated tactile sensors on robotic fingers. (A) Tactile array for the Schunk Hand [3], (B) Conductive fluid-filled multimodal fingertip sensor [4] (C) Baraometer-based tactile array on a Reflex Hand finger [5].

spatial resolution, the same procedures can also be applied to other sensor specifications, such as sensitivity and noise. It could also be extended to other modes of sensors, such as kinematic sensors.

We begin by reviewing the models for tactile sensing and grasp stability. We lay out the relationship between sensor resolution and stability prediction accuracy in section III, and discuss the implication of our analysis in Section IV.

II. Models

We used classic grasp theory to guide our investigation into the relationship between tactile sensor spatial resolution and grasp stability. In this section, we describe the contact model and the physical theory that underlies our analysis, as well as the model used to parameterize the sensors and their spatial resolution.

A. Grasp Contact and Success Model

Many models of grasping mechanics and metrics for grasp quality have been proposed [1], [2], and most could be employed for the purposes of the present analysis. We will follow the widely-cited analysis by Ferrari and Canny [6], which assumes that all contacts are point-contact-with-friction. While this contact model may not be the most accurate for many fingers, it is the most conservative choice for the many tactile sensors that cannot reliably differentiate a point-contact and a soft-contact.

Let the force exerted at the i^{th} contact be $\boldsymbol{F}_i$, then all the forces that do not cause slip form a friction cone. The friction cone is defined by a normal force $\boldsymbol{f_n}$ in the local surface normal direction $\boldsymbol{n}_i$ and an orthogonal tangential force $\boldsymbol{f_t}$, where $\boldsymbol{f_t} \leqslant \mu \boldsymbol{f_n}$, with μ being the coefficient of friction. The same friction cone can also be written as a convex linear combination of a set of edge vectors $\boldsymbol{f_j}$'s so that the friction

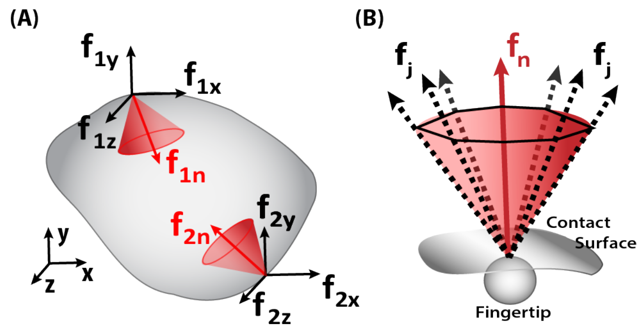

Fig. 3: (A) The point-contact-with-friction model for grasp quality estimation. The force that can be exerted without slipping by each contact is contained within a friction cone. (B) Friction cone is defined by the normal force vector f_n and coefficient of friction μ , but can be approximated by a polygon defined by convex combinations of vectors f_j.

cone is approximated by a polygon (Fig.3). Using a total of m vectors to span the possible force at contact i allows the force to be expressed as

$$\boldsymbol{F}_i = \sum_{j=1}^{m} c_{i,j} \boldsymbol{f}_{i,j}$$

where weights $c_{i,j} \geqslant 0$ and the total force is normalized to express the actuation limit, so that $\sum_{j=1}^{m} c_{i,j} \leqslant 1$.

The corresponding wrench is

$$\boldsymbol{\omega}_i = (\boldsymbol{F}_i, \boldsymbol{\tau}_i)^T$$

where with $\boldsymbol{r}_i$ is the moment arm for contact i,

$$\boldsymbol{\tau}_i = \sum_{j=1}^{m} c_{i,j} (\boldsymbol{r}_i \times \boldsymbol{f}_{i,j})$$

The reference point around which moments are calculated can be arbitrarily selected. One convenient choice is the center of mass (COM) of the object, so that gravity force do not produce additional torque in the system.

Solving for the wrench equilibrium of any grasp absent task forces is equivalent to finding a non-trivial solution to

$$\boldsymbol{\omega} = \sum_{i=1}^{n} c_i \boldsymbol{\omega}_i = 0 \qquad c_i \geqslant 0$$

This is equivalent to determining if the origin is inside or on the surface of the convex hull, $\mathcal{G}$

$$\mathcal{G} = ConvexHull(\bigoplus_{i=1}^{n} \{\boldsymbol{\omega_{i,1}}, ...\boldsymbol{\omega_{i,m}}\})$$

Grasp quality ϵ is defined as the distance from origin to the closest hyperplane of the convex hull. It can be thought of as the minimum amount of perturbation required to push the object out of the fingers. The wrench equilibrium exists and the grasp is stable if $\mathcal{G}$ contains the origin. If the solution is on the surface of the hull, then it will have $\epsilon = 0$, which will be defined as unstable under this particular definition. See [6] for further details.

This model aims to determine whether a stable grasp is possible at the configurations and contacts [6]. It assumes necessary actuation forces can be generated as needed and does not account for finger kinematics limitations. Different contact and stability models can be substituted, and task forces and kinematic constrains can all be added to suit specific systems and tasks.

B. Sensor Model

There are numerous tactile sensor designs using a variety of transducers and measuring many features of contacts (Fig.2) [7], [8]. Of all the parameters needed to calculate grasp stability according to the model described in the previous section, we will focus on sensors that provide contact location and local surface normals estimates. Sensors and algorithms for estimating parameters such as friction, mass, and COM are an important open question but are usually distinct from contact sensors that report spatial and geometric information. We will assume these parameters will be measured independently. The effect of uncertainty of those parameters are discussed in section IV.

Regardless of design and algorithm details, most sensors have a spatial resolution limit, where contacts within that limit cannot be distinguished from each other. Determining the exact resolution of the sensors depends on the sensing mechanism and signal processing algorithm. For example, some sensor arrays overlap the receptive field of individual sensors, so that the overall spatial resolution is subtaxel [9]. In some fluid-filled sensors, the relationship between the signal and contact location is a highly nonlinear function of how much fluid has been displaced between the electrodes [10]. However, despite having completely different sensing mechanisms and signal processing algorithms, as long as there exists a region on the receptive surface where small variation in contacts cannot be disambiguated, then the following analysis is applicable.

We will define the sensor spatial resolution as $\boldsymbol{\alpha}$, which represents the angle difference between the two extreme possible surface normal vectors within an unit resolution segment (Fig. 4). Contact anywhere on the arc covered by $\boldsymbol{\alpha}$ will result in the same estimates. Consequently, if the sensors estimate a particular segment to be active, then the average surface normal and location of that segment will be used for subsequent calculations, even though the true surface normal vector can be anywhere on that segment and may be rotated up to $\pm\alpha/2$ from the reported normal. The limited spatial resolution effectively quantizes the signals and attaches an uncertainty range for any signal received.

As long as the contact sensor's spatial resolution is known, the same definition above can be applied to fingers regardless of material, shape, and sensor layout. For example, many robotic fingers have rubber-like fingerpads that conform to the contact surface as a function of the magnitude and direction of force exerted and the properties of the material. Soft contacts can resist higher friction force and moment, and this behavior can be taken into account in signal processing, often using mechanical models, to estimate contact location and surface normals [11]. Additionally, $\boldsymbol{\alpha}$ depends heavily on the curvature of the contact surface. Sensors of the same receptive field size but facing a flatter contact surface will have a smaller $\boldsymbol{\alpha}$, and the corresponding finger will have a finer resolution. Regardless, most sensors will likely have a quantifiable spatial resolution limitation that translates to uncertainties in perceived surface normal vector and contact location. In the following section we will show how uncertainties in these parameters can produce uncertainties in orientation and apex angles of friction cones, which propagate into uncertainty in grasp stability calculation.

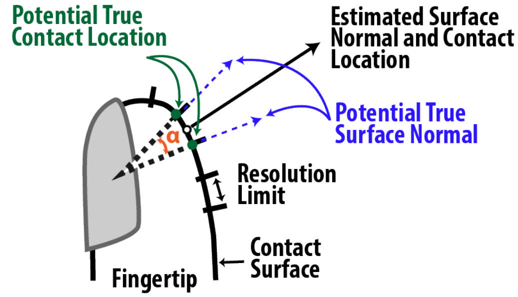

Fig. 4: The sensor resolution is described by a segment of an active area that extends over angle α, where contacts on the arc covered by α cannot be differentiated. The actual contact surface normal can differ from the estimated normal vector by as much as $\alpha/2$, and the contact location can vary up to half of the length of the receptive field.

III. Analysis

A. Stability Prediction Uncertainty Analysis

In the sensor model described in section II-B, the boundary between slip and stable forces is represented by a friction cone. The coefficient of friction, μ, dictates the apex angle of the friction cone, 2β, where $\beta = tan^{-1}\mu$; the surface normal vector determines the orientation of the friction cone. A sensor that has resolution angle of α can result in a perceived normal rotated up to $\alpha/2$ from the true normal direction. The uncertainty in contact location due to spatial resolution is considered negligible here. Contact location effects only the calculation of moment balance, and for grasps where the objects are significantly larger than the receptive area of a tactile sensor, the minuscule shift in contact location would have negligible effect on the moment. As a result, the surface normal direction is the main source of the stability prediction uncertainty.

A grasp stability prediction is reliable if for a given perceived surface normal vector, none of its corresponding potential true normal vectors can produce a prediction that is different from the perceived one. Hence, to check the reliability of a prediction is to check the stability measure calculated by all the potential true normal vectors within the $\alpha/2$ of the perceived normal.

Recall the underlying formulation of convex hull can be thought as looking among all the combinations of force vectors — one from each contact and inside the friction cones — to see if there exist a combination that can achieve wrench equilibrium. Therefore, the process of using friction cones

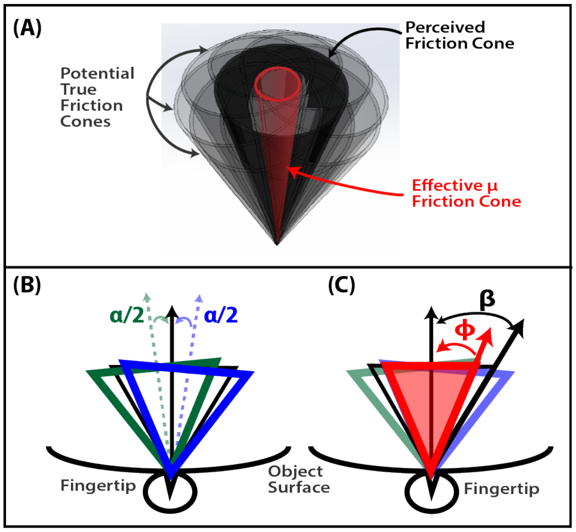

Fig. 5: (A) Perceived, actual, and effective friction cones in 3D. (B) The cross-section of (A), where uncertainty in contact normal results in a range of potential friction cones, bounded by the green and blue cones. (C) the overlap between the green and blue cones, shown in red, is the cone containing the force vectors that can be reliably exerted at the given point; the angle that spans this "effective friction cone" can be calculated by $\phi = \beta - \alpha/2$.

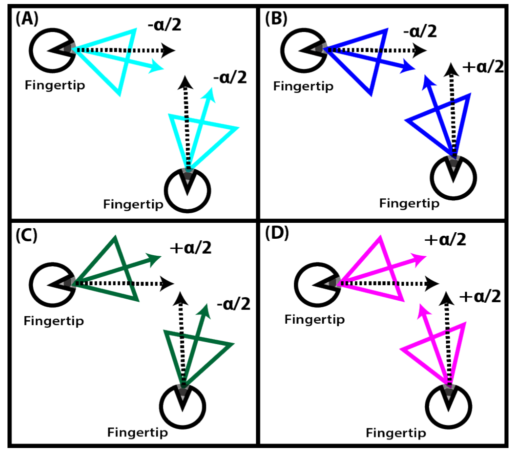

Fig. 6: There are two possible extreme contact normal vectors for a single planar contact: $+\alpha/2$ and $-\alpha/2$ surrounding the true contact normal. This results in four combinations of the extreme friction cones for a 2-finger planar grasp.

to calculate ϵ means all the combinations of force vectors within the friction cones are automatically tested for equilibrium. Consequently, to ensure all possible combinations of force vectors at each contact are tested, we simply have to calculate the grasp stability measure using a few friction cones that collectively span all potential force vectors at the given contact. In most cases, the friction cones defined by a set of extreme surface normal vectors that can approximate the full range(Fig. 5).

Alternatively, when using a perceived surface normal that is skewed from the actual normal, the stability calculation may be relying on force vectors that will cause slip in reality. Therefore, to be safe, instead of using all the force vectors inside the friction cone defined by the perceived normal, only the force vectors that are simultaneously inside all potential friction cones will be used. For a contact with apex angle of $\alpha \leqslant 2\beta$, all the non-slipping force vectors form a new cone with apex angle of 2ϕ, where $\phi = \beta - \alpha/2$ (Fig. 5). This is equivalent to assuming a more conservative coefficient of friction, which inevitably produces more conservative stability estimates and smaller stable graspable regions on objects. We will refer to the coefficient of friction that corresponds to the new conservative friction cone defined by ϕ as the effective coefficient of friction, or μ_{eff}(Fig. 5. If $\alpha > 2\beta$, then it would be impossible to guarantee that only non-slip forces are used in force equilibrium calculations using tactile sensors alone.

B. Example Objects

To demonstrate how uncertainty in spatial resolution propagates into grasp stability calculations, we apply our method to a few objects of known geometry, so that we can compare the stability measure calculated from the perceived normals to the true one, and see how tactile sensor uncertainty is manifest on actual objects.

We will also make a few simplifying assumptions for ease of illustration, though all of the assumptions can be modified as needed to match specific systems. The magnitude of total force here is normalized to 1 to represent actuation limits, the coefficient of friction is also 1 unless otherwise noted, and the COM coincides with the geometric centroid of the object. We also assume no external task forces, so that force equilibrium in this case can be thought as whether the object will slide against the fingers when squeezed. Task forces and gravity can be added in a straightforward fashion using the original model described by [6]. Also recall that the contact model used here is the rigid point-with friction contact model, so that surface normal on the finger at contact is equal and opposite to the surface normal of the object.

Many objects have spherical or rectangular profiles, so two-fingered planar grasps of these objects can be simplified to be represented in 2D as grasping of a circle or a square for ease of illustration. At each configuration, we calculate the stability predictions using the full range of possible normals — which can be shifted by $\pm\alpha/2$ from the perceived normal— and then compare them against the true stability. The overall reliability of grasp prediction for a given object can then be evaluated after the reliability of prediction at each configuration has been examined.

Two contact locations and the center of the sphere determines a great circle. To examine all possible contact configurations on a circle, we fixed one finger at one point on the circle and scanned the second finger along the perimeter (Figure 7A). The angle between the two fingers is parameterized by $\theta = [0, 2\pi)$. The potential ϵ at each point of circle are plotted in Figure 7B for two sensor resolutions: $\alpha = \pi/6(\approx 30^{\circ})$ and $\alpha = \pi/12(\approx 15^{\circ})$. In both plots, the black lines represent the true ϵ values at each point on the perimeter. In two finger grasps where both fingers have the same sensor resolutions,

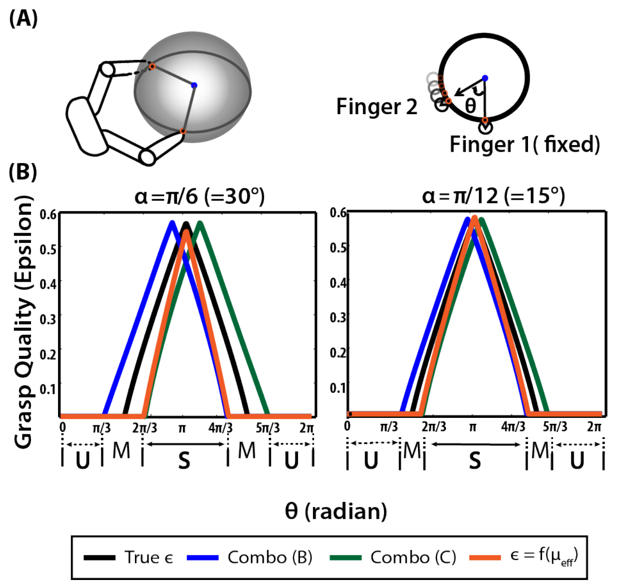

Fig. 7: (A) The ϵ values for all contact locations calculated by fixing one finger at the bottom of the circle, and scanning the second finger along the perimeter. (B) the variability of ϵ increases with lower sensor spatial resolution. The regions of definite stable(S), definite unstable(U), and ambiguous(M) contact locations are labeled along the bottom. Values where $\epsilon < 0$ are unstable and are plotted as $\epsilon = 0$.

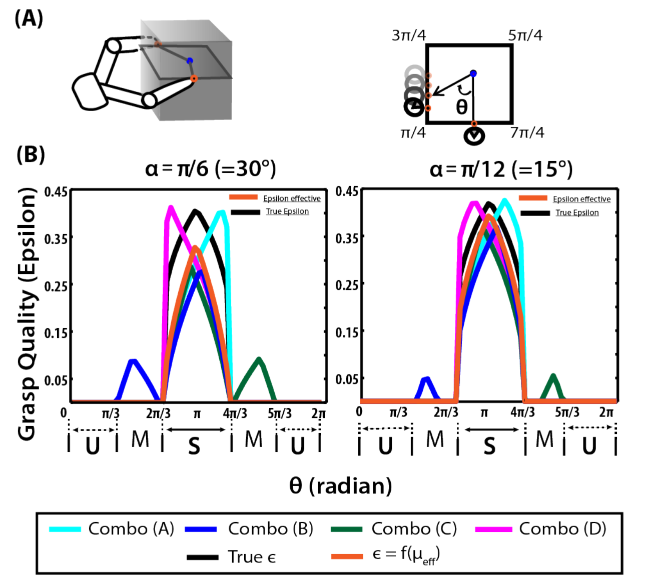

Fig. 9: (A) The ϵ values for all contact locations calculated by fixing one finger at the bottom of the square, and scanning the second finger along the perimeter. (B) ϵ calculated by using the four combinations of extreme friction cones shown in fig. 6. The regions of definite stable(S), definite unstable(U), and ambiguous(M) contact locations are labeled along the bottom. Values where $\epsilon < 0$ are unstable and are plotted as $\epsilon = 0$.

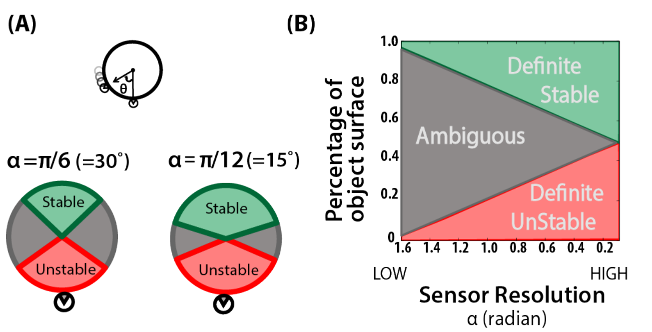

Fig. 8: (A) The placement of the second finger that produces stable, unstable, and ambiguous regions on the circle, (B) As the resolution increases, the size of regraspable region (definitely unstable+ambiguous) decreases.

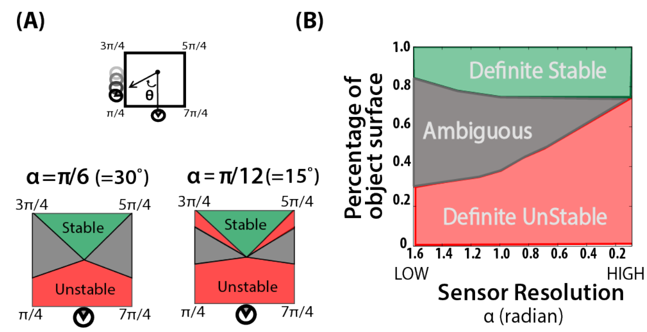

Fig. 10: (A) The grasp is only definitely stable when the fingers are on opposite sides of the square, (B) As the resolution increases, the size of definite stable region remains stable, and the size of ambiguous regions reduces.

there are four combinations of the extreme friction cones (See Figure 6), and all four should be calculated. If all the four ϵ values have the same sign, then the grasp prediction at this specific configuration is reliable. These four combinations of extreme cones are all plotted. However, for a circle, two of the combinations result in the same ϵ values as the other two, and they are shown in the plot as blue and green lines. ϵ calculated using μ_{eff} is also plotted in red.

For a given configuration where the finger 2 contact is at θ, the corresponding ϵ may vary depending on the perceived contact normal vector used for calculation. If the signs of all the possible ϵ values are all positive, then the grasps corresponding to these θs are definitely stable despite resolution-induced uncertainties(S). When they are all zeros, then the corresponding grasps are unmistakably unstable(U). But if they disagree and have different signs(M), then the predictions corresponding to these θs can go either way, and the outcomes is indistinguishable by tactile sensors. Grasp configurations that land in region S are definite successes, and those that in either region U or M require regrasping.

For objects with a circular cross-section, the plots indicate that when the fingers are close together, the grasps will always fail, and when the fingers are on the opposite sides of the object, the grasps will succeed. The ambiguous regions are sandwiched between the definitely stable and definitely unstable region, and the sizes of those areas decrease as the tactile resolution increases (Fig. 8). Hence, higher tactile resolution allows more area on the object to be reliably graspable.

Figure 7 also shows that using the μ_{eff} outputs the same sign as the minimum grasp quality values throughout the grasping surface. While the exact values are not the same, the limits of θ that separates the reliably stable region from the rest is consistent. Therefore μ_{eff} can be used as a shortcut to calculate stability, as long as the extreme friction cones have overlapping regions.

The analysis for an object with a square cross-section is similar, with the added complexity of corners, where there is a range of true surface normals of $90°$, which implies the perceived surface normal can have a range of $90° + \alpha$. As a result, when calculating the grasp qualities at a corner, the two extreme friction cones may not overlap. Therefore in order to cover all the possible force vectors on the corner, a friction cone in addition to the two extreme ones are used to evaluate potential ϵ. The additional cone has a normal that bisects the corner angle.

The results are plotted in Figure 9. Similar to a circle, grasps where the fingers are on the opposite sides of the square are always successful, and grasps with fingers on the same side will always fail. If the friction coefficient is high enough, it is possible to achieve success with fingers pinching the square from two adjacent sides. The ambiguous regions are similar to the circle, and μ_{eff} also correctly marks the θ limits. The variation of ϵ is the highest at the two opposing corners, for they could produce ϵ values ranging anywhere from 0 to almost $\max(\epsilon)$.

This analysis reveals that limitation in tactile spatial resolution can drastically change the ratio of definite-stable, definite-unstable, and ambiguous regions. For objects with a circular cross-section, higher resolution linearly increases the sizes of graspable region and definite-failure region, as well as linearly reducing the size of ambiguous region (Fig. 8). For planar grasps on objects with square cross-sections, the guaranteed success region remains relatively constant regardless of resolution, but the ambiguous region shrinks as the resolution increases (Fig.10). The decrease in ambiguous region is mostly the result of distinguishing when the pinch grasp is no longer obtainable due to lack of moment balance for points closer to the upper corners.

IV. Discussion

In this paper, we determined the effect of tactile sensors spatial resolution on grasp stability prediction accuracy. Using classic grasp analysis, the uncertainties introduced by limited spatial resolution on estimates of contact surface normal and location are propagated into the grasp stability calculation, so that the reliability of the stability prediction can be quantitatively correlated with tactile sensor resolution.

Historically, task performance has been difficult to relate to sensor parameters. This is the first time to our knowledge that task performance is directly linked to sensor quality quantitatively. This framework allows sensor design parameters to be set by desired performance, as well as more accurate task performance evaluation by including sensor quality in the calculations. Understanding the relationship between sensor metrics and task performance is key in optimizing the system for expected tasks.

A. Implication for Sensor and System Design

Our approach offers quantitative solutions to sensor design decisions for grasping tasks. For example, assuming a sensor layout where the resolution is the same as each tactile sensor's coverage, then when using a moderately conservative coefficient of friction $\mu = 0.5$, each sensor must have a coverage angle, α, of at most $2tan^{-1}\mu$, or approximately $53°$. In a hemispherical fingertip with a diameter of 2cm, $53°$ covers an arc of 9mm, or approximately a 9mm-diameter circular patch in 3D. Then to cover the whole hemisphere would require at least 10 such sensors. Figures 8 and 10 show that a coverage of $53°$ or 0.9rad still leaves significant ambiguous regions on an object with circular or square cross-sections. An even more conservative estimate for coefficient of friction of $\mu = 0.2$ would decrease the α to $22.6°$, or the patch to 4mm in diameter. It would increases the minimum number of sensors needed to 50, but reduce the size of the ambiguous region by approximately half.

Our analysis suggests that while the ability to predict grasp stability through tactile sensors is based on local contact surfaces and independent of global object geometry, the benefit of using tactile sensors to ensure stability is more prominent for some geometry than others. In rectangular objects, where there are large parallel surfaces and discrete contact normals, the system can achieve reliable stable grasps as long as it guarantees the fingers are on opposite surfaces and close to the center of mass, in which case simple vision accompanying parallel grippers could be enough. For objects with curved surfaces and continuous contact normal directions, the prediction accuracy is then directly correlated with the sensor resolution as analyzed. Therefore, the requirement for tactile sensors may also depend on the variety of object shapes exist in the task, where objects with non-parallel surfaces are more likely to benefit from tactile sensors with high resolution.

B. Beyond Ferrari and Canny and Spheres

The method proposed here is applicable to a wide range of grasping systems because it is agnostic to the sensor and finger mechanics and the choice of grasp theory employed. The key concept is that physics-based grasp prediction focuses sensor evaluation into a well-formulated physics problem, and the parameters essential for solving the equations are the information we need from the sensors. The range of possible sensor values due to resolution limits are propagated through the physical model to determine their effects on grasp prediction.

Ferrari and Canny defined a stable grasp by two conditions: (1) the ability to achieve force and moment balance, or force-closure, at the given grasp configuration, and (2) all resultant forces are inside their corresponding friction cones, so that there is no slipping at the contacts. To establish the friction cones at each contact, we need the coefficient of friction and contact surface normal; to calculate force and moment balance, we need contact locations plus the friction cones sizes and orientations. These parameters are what one should look for in the sensors.

The Ferrari and Cannys stability conditions translate to a minimal set of "essential parameters" — parameters that are necessary to solve for the stability condition. Other grasp stability definitions may have a different set. Some stability definitions use the curvature of the local surface [12], in which case the finger and object curvature becomes essential. For models assuming soft contact, the material deformation property is key because the area, pressure distribution, and friction of a contact are all functions of the material properties [13]. Some of these additional essential parameters can be obtained *a priori*, such as finger geometry and material properties. But parameters such as local object curvature must be delegated to sensors to collect in real time, and are therefore subjected to sensor limitations.

Furthermore, information obtained by other sensor modalities, such as kinematic and dynamics sensors, can also be mapped onto the same essential parameters. For example, joint angles are needed to calculate the fingertip position in space, which when combined with finger geometry and tactile sensors, can determine the surface contact normals and locations in robot coordinates. If the encoder count is low, then the uncertainty in fingertip position can result in enough uncertainty in the contact normal to affect the task performance. The bounding values of the encoder uncertainty can be propagated through the stability calculation — as above with the tactile sensor data — to determine the resulting limitation on grasp stability prediction. Hence, the needed accuracy of joint angle encoders can also be determined based on desired performance in the same way.

The above analysis on the effect of tactile sensor spatial resolution is decoupled from the issue of obtaining other essential parameters such as friction, mass, and COM. In some cases, heuristically estimating them from geometric information and past experience has been shown to work well enough. Nevertheless, sensors for measuring friction, mass, and COM, as well as dynamic events such as slip will be important for reliable task performance and real-time control in unstructured environments. For further discussions of friction, slip, and inertia detection, see [14], [15].

While we illustrated our concepts using 2-finger planar grasps for the ease of explanation and visualization, the models on which we based our analysis are already in 3D, and applicable for multi-contact grasps. Therefore the extension to 3D is straightforward. More work may be needed to make the algorithm more computationally tractable as more dimension and contacts are added.

C. Implication for Machine Learning

Real world grasping has complex interaction dynamics and significant noise and uncertainties. Recently, a number of researchers have employed a data-driven, machine learning approach [16], [17], [18]. However, despite experimenting with different algorithms, machine learned results so far are not convincingly generalizable, with most obtaining stability prediction accuracies of approximately 80% to 90% when tested on unknown objects [19].

The reasons behind the plateau may be twofold: 1) high dimensionality of grasping and high cost of obtaining real data, leading to a training dataset that is too sparse; and 2) sensors are not capturing enough information. Our analysis gives some evidence for (2). More specifically, grasp configuration in the ambiguous region may have the same sensor signal representing both failed and successful grasps. Data collected from ambiguous grasps would have seemingly randomly labeled data. The algorithm trained on ambiguous data will produce the same random labels for future grasps in the same region. Hence the best the resulting algorithm can do is predicting all the definitely-stable and definitely-unstable regions perfectly, and responding to grasps in the ambiguous region by chance.

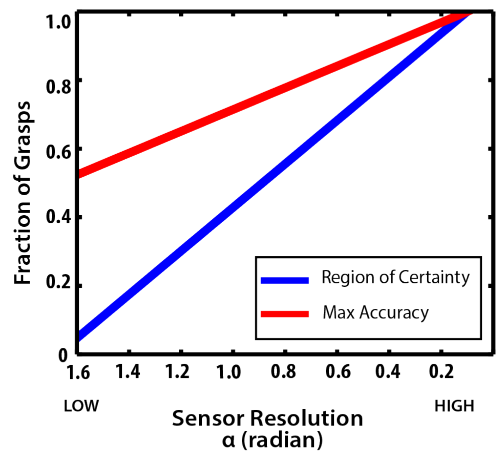

Fig. 11: Implication for Machine Learning. Spatial resolution limits lead to a fraction of the training data being ambiguous, which sets the ceiling for maximum accuracy if those signals are used to train a grasp prediction algorithm. Example here is for spherical objects.

For example, as roughly calculated in section IV-A, a sensor with resolution $\alpha = 22.6°$ covers a patch of approximately 4mm on a 2cm-diameter fingertip — a resolution that is on par with some sensorized robotic hands [5] — will produce an algorithm that guesses the outcome of almost 20% of grasps on spherical objects (Fig. 8), assuming that training data spans the average of possible contact locations. Unless more accurate surface normal and contact location information are embedded in other sensors, the stability prediction based on tactile sensors alone will not perform above 90% accuracy.

While better differentiation between unstable and ambiguous region does not necessarily change the control strategies, for they both require regrasping, it will affect sample labeling in data-driven strategies. Larger ambiguous regions means more randomly labeled samples in the dataset, which will fix a ceiling of stability prediction accuracy that is correlated with the size of the ambiguous regions (Fig. 11).

Our analysis suggests that tactile sensor spatial resolution limits, and potentially other inherent sensor limitations, can be one of the culprits in the plateauing of success reached by data-driven stability prediction algorithms. If the needed information is embedded in the sensors but in complex patterns, then machine learning can be potentially helpful in extracting them, but even the most sophisticated algorithms cannot compensate for gaps in relevant information signals. Therefore it is imperative to understand the information content of the sensors for each setup on case-by-case basis in order to have reasonable expectations for the respective learning algorithms.

V. Conclusion

Understanding the relationship between sensor quality and task performance is essential for designing better grasping systems that achieve desired performance. This paper presented a method to quantitatively relate the spatial resolution of tactile sensors to the system's ability to guarantee grasp stability. Anchored by grasp theory, our analysis laid out how tactile sensor spatial resolution introduces uncertainties in surface normal and contact location, and how it ultimately influences grasp stability calculation. The approach here can be extrapolated to a wide range of tasks that have relatively reliable physical models by (1) identifying the essential parameters for task performance; (2) mapping the sensor signals onto the essential parameters; and then (3) quantifying sensor uncertainty and propagating errors to the task, and thereby determining the potential performance.

Acknowledgments

We thank Matei Ciocarlie, Matt Mason, and Doug Perrin for insightful conversations during the development of this work. Funding was provided by the National Science Foundation under grant IIP-1500178.

References

[1] M. A. Roa and R. Suárez, "Grasp quality measures: review and performance," *Autonomous Robots*, vol. 38, no. 1, pp. 65–88, 2015.

[2] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *International Conference on Robotics and Automation*, vol. 1. IEEE, 2000, pp. 348–353.

[3] "Schunk Hand," http://www.schunk-modular-robotics.com, accessed:2017-Aug-02.

[4] "the BioTac©, SynTouch, LLC," https://www.syntouchinc.com/sensor-technology/, accessed:2017-Aug-02.

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

[6] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE International Conference on Robotics and Automation*. IEEE, 1992, pp. 2290–2295.

[7] H. R. Nicholls and M. H. Lee, "A survey of robot tactile sensing technology," *The International Journal of Robotics Research*, vol. 8, no. 3, pp. 3–30, 1989.

[8] P. S. Girão, P. M. P. Ramos, O. Postolache, and J. M. D. Pereira, "Tactile sensors for robotic applications," *Measurement*, vol. 46, no. 3, pp. 1257–1271, 2013.

[9] N. F. Lepora, U. Martinez-Hernandez, M. Evans, L. Natale, G. Metta, and T. J. Prescott, "Tactile superresolution and biomimetic hyperacuity," *IEEE Transactions on Robotics*, vol. 31, no. 3, pp. 605–618, 2015.

[10] N. Wettels, V. J. Santos, R. S. Johansson, and G. E. Loeb, "Biomimetic tactile sensor array," *Advanced Robotics*, vol. 22, no. 8, pp. 829–849, 2008.

[11] 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. [Online]. Available: https://doi.org/10.1177/027836498500400304

[12] W. S. Howard and V. Kumar, "On the stability of grasped objects," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 6, pp. 904–917, 1996.

[13] M. Ciocarlie, C. Lackner, and P. Allen, "Soft finger model with adaptive contact geometry for grasping and manipulation tasks," in *EuroHaptics Conference, and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics. Second Joint*. IEEE, 2007, pp. 219–224.

[14] C. Melchiorri, "Slip detection and control using tactile and force sensors," *IEEE/ASME transactions on mechatronics*, vol. 5, no. 3, pp. 235–243, 2000.

[15] V. A. Ho, T. Nagatani, A. Noda, and S. Hirai, "What can be inferred from a tactile arrayed sensor in autonomous in-hand manipulation?" in *IEEE International Conference on Automation Science and Engineering*, Aug 2012, pp. 461–468.

[16] H. Dang and P. K. Allen, "Learning grasp stability," in *IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 2392–2397.

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

[18] M. Madry, L. Bo, D. Kragic, and D. Fox, "St-hmp: Unsupervised spatio-temporal feature learning for tactile data," in *IEEE International Conference on Robotics and Automation*. IEEE, 2014, pp. 2262–2269.

[19] J. Laaksonen, V. Kyrki, and D. Kragic, "Evaluation of feature representation and machine learning methods in grasp stability learning," in *IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2010, pp. 112–117.