Evaluating the Efficacy of Grasp Metrics for Utilization in a Gaussian Process-Based Grasp Predictor

Alex K. Goins[1], Ryan Carpenter[1], Weng-Keen Wong[2], and Ravi Balasubramanian[1]

Abstract— **With the goal of advancing the state of automatic robotic grasping, we present a novel approach that combines machine learning techniques and rigorous validation on a physical robotic platform in order to develop an algorithm that predicts the quality of a robotic grasp before execution. After collecting a large grasp sample set (522 grasps), we first conduct a thorough statistical analysis of the ability of grasp metrics that are commonly used in the robotics literature to discriminate between good and bad grasps. We then apply Principal Component Analysis and Gaussian Process algorithms on the discriminative grasp metrics to build a classifier that predicts grasp quality. The key findings are as follows: (i) several of the grasp metrics in the literature are weak predictors of grasp quality when implemented on a physical robotic platform; (ii) the Gaussian Process-based classifier significantly improves grasp prediction techniques by providing an absolute grasp quality prediction score from combining multiple grasp metrics. Specifically, the GP classifier showed a 66% percent improvement in the True Positive classification rate at a low False Positive rate of 5% when compared with classification based on thresholding of individual grasp metrics.**

I. INTRODUCTION

Developing automatic algorithms that enable robots to grasp objects robustly is fundamentally important to the field of robotics, since it would pave the way for the use of robots in domestic and outdoor environments and not just in structured industrial settings. Recognizing this need, a variety of approaches based on physics force modeling [1], [2], machine-learning based techniques [3], and human-inspired grasping [4] have been developed for the automatic generation and prediction of robotic grasp success prior to execution. While significant progress has been made, recent results show that even the best of these autonomous grasp generation methods has a failure rate of 23% when implemented on a physical robot [5].

Such a high failure rate shows the complexity of the robotic grasping problem. This may be attributed to the difficulty in modeling non-linear effects such as contact friction, slip, compliance, and object movement due to disturbances during grasping.

In order to overcome the challenges of modeling these effects, researchers have developed metrics with the intention of capturing the properties that make a grasp secure and robust even in the presence of such uncertainty. For example, the physics-based grasp metrics "epsilon" and "volume" were developed using grasp wrench-space computations based on the magnitude and direction of generalized forces applied by the gripper to evaluate the grasp stability [2]. Another example is "grasp energy", which measures the average distance between potential gripper contact points and the object to determine the extent to which the object is enveloped by the hand [6].

Surveys of grasping literature [7], [8], [9], [10], [11], [12], [13] list as many as 24 grasp metrics which have been developed, mostly based on kinematic models (see Table I for a partial list of some of them). While some metrics, like finger spread, apply only to three finger grippers, the majority of metrics are applicable to other multifinger grippers [14], [15] and even the human hand [16]. However, each grasp metric individually captures only a small aspect of what makes a good grasp. As was found in [17], [18], slight variations in hand placement relative to the object can significantly change the metric value and grasp performance. Furthermore, the metrics are also highly correlated since they are often calculated from dependent variables (such as finger contact location) which are based on the independent variables (such as hand pose, orientation, finger spread, and object type). Adjusting one independent variable could affect multiple dependent variables causing correlation among the various metrics.

In order to capture broader aspects of grasping and

[1]Affiliated with the School of Mechanical, Industrial, and Manufacturing Engineering, Oregon State University, Corvallis, OR 97331, USA ravi.balasubramanian@oregonstate.edu

[2]Affiliated with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331, USA

potentially improve grasp prediction performance, researchers have also developed aggregate grasps metrics that merge the evaluation signals from several individual metrics up to as many as nine metrics [19]. For example, weighted sums of epsilon, volume, and energy have been used simultaneously as a quality measure in the open source grasp planning and evaluation software GraspIt! [20] (also see [1], [21], [22], [23], [17] and Table I for other examples).

However, there are three key problems with the state of the art. First, most of the grasp metrics have been evaluated through simulation only [5], with limited validation of these metrics on physical robots [24], [19], [25]. Second, current methods have largely failed to account for the interactions or correlations between the grasp metrics [19], [16] which can lead to erroneous grasp quality prediction if unaccounted for. Third, most metrics only provide a measure of relative grasp quality, thus making it difficult to assess the grasp performance in absolute terms prior to execution.

Given the state of grasp generation and grasp quality prediction algorithms, this paper uses machine learning techniques and rigorous validation on a physical robotic platform to develop an absolute grasp quality prediction algorithm. This paper's key contributions are: (i) An evaluation of individual grasp metrics commonly used in the robotics literature. (ii) The development of a data-driven approach to use a state-of-the-art classification algorithm to predict grasp quality and quantitatively compare its performance with prediction using current grasp metrics individually.

TABLE I: Grasp Metrics

Metric	Description
Contact Point Equilateralness [7]	Equilateralness of the triangle made by the contact points of the finger tips
Grasp Volume	Volume of the triangular prism consisting of the finger tips and the palm
Finger Extension	Average finger flexion
Finger Spread	Amount of spread of the fingers
Finger Limit	Total flexion of all the fingers
Parallel Symmetry [27]	Distance between center of mass of object and contact point parallel to the object principal axis
Perpendicular Symmetry	Distance between center of mass of object and contact point perpendicular to the object principal axis
Object Volume Enclosed	Normalized volume of the object enclosed by the hand
Skewness [5]	Alignment of the hand principal axis parallel to the object principal axis
Grasp Wrench (Epsilon) [2], [20]	Minimum disturbance wrench that can be resisted
Grasp Wrench Volume	Volume of grasp wrench space
Grasp Energy	Distance of hand sample points to object

II. BACKGROUND

In this research, we use a Gaussian Process [26] as our machine learning algorithm because it can model the non-linear relationship among the grasp metrics as well as create a non-linear decision surface between good and bad grasps. In addition, Gaussian Processes also provides the variance of its predictions, thereby providing a measure of the confidence or uncertainty regarding the prediction. This learning algorithm allows us to estimate the absolute grasp quality, rather than a relative quality measure which current techniques provide. Other machine learning methods that can deal with the non-linear nature of the grasp space could also be used, but exploring all of them is not within the scope of this paper.

A. Gaussian Process

A Gaussian Process (GP) is a non-parametric model that can be used for supervised learning [26]. Specifically, given a set of n training samples $D = \{(x_1, y_1), \ldots, (x_n, y_n)\}$, where x_i is a feature vector and y_i is the output value, the algorithm learns a non-linear function $f(x)$ that generalizes from the training data in order to predict the output value y for some new data instance x.

GPs may be thought of as a generalization of a multivariate Gaussian distribution to infinite dimensions, such that any finite subset of the components of this infinite-dimensional vector is jointly Gaussian. Rather than just modeling a single function $f(x)$, a GP is a stochastic process that models a distribution over functions $f(x)$.

In our work, each data instance x_i is a grasp, which has k features that correspond to k grasp metrics used to represent it. Table I shows the $k = 12$ grasp metrics used in this paper. We use the GP to predict a continuous output value between 0 to 1 that represents the probability of the grasp being successful. We use an open-source GP package known as GPML[1] which was implemented in Matlab[2].

III. EXPERIMENTAL METHODS

Our approach includes a combination of grasp generation and evaluation on a physical robotic platform and

[1] http://www.gaussianprocess.org/gpml/code/matlab/doc/
[2] http://www.mathworks.com/

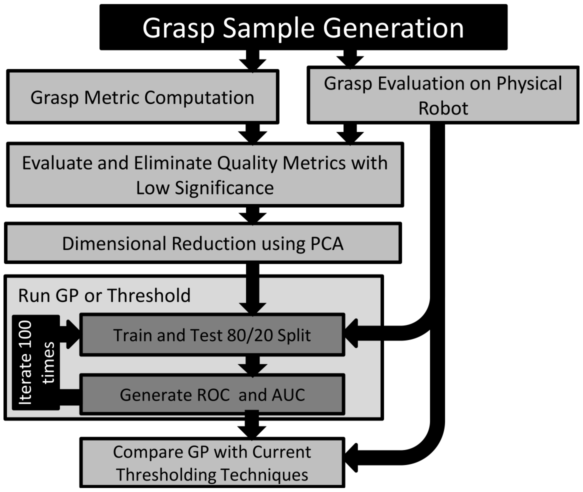

Fig. 1: Flow chart of experimental procedure

(a) (b)

Fig. 2: (a) Nine everyday objects used for grasp generation (b) Shake test setup using WAM and marked reference location for object placement

machine learning techniques to develop an algorithm for grasp quality prediction. An overview of the process used to develop the algorithm, including intermediate steps to perform dimensionality reduction on the data, is shown in Fig. 1.

A. *Grasp Metric Selection and Evaluation*

We selected twelve of the most common kinematic based metrics for evaluation and testing (see Table I). Other metrics which depend on having force or contact sensors were not included in this study since our Barrett manipulator system does not have the capabilities to support them (see Fig. 2b). While we did not analyze the other metrics, they can easily be included using the same procedure outlined below to increase the performance with grasping systems that have more capabilities.

B. *Collection of the Grasp Sample Set*

Twenty two human subjects were recruited to provide a total of 522 robotic grasp examples across nine everyday objects (see Fig. 2a) using a simulation environment developed in OpenRAVE [28]. Each human subject commanded the position, orientation, finger spread, and grasp closure of the virtual BarrettHand [29] robotic hand, and had the option of viewing the grasp from several angles. Subjects used one of three common human-robot interfaces to grasp and pick up an object, a gamepad controller, a three-dimensional mouse, and an "interactive marker" display [30]. Different human-robot interfaces were utilized to ensure that the grasp sample set was diverse and that one particular interface did not skew the grasp examples. Also, the robot hand's starting location was randomized between the objects so that a subject would not repeat the same grasp across multiple objects. When the user finished grasping the object and was satisfied with the final grasp, both the robot hand's posture relative to the object's coordinate frame and the computed metric scores were recorded. The human-subject experiment procedure was approved by Oregon State University's Human Subjects Division.

C. *Evaluation of the Grasp Sample Set*

To determine the quality of the grasps provided by the human subjects, we tested the example grasps on a Barrett WAM and BarrettHand equipped with standard rubber fingertips. This process was done in order to validate the predictive capability of the metrics, as well as provide ground truth data for the machine learning algorithm. Foam spacers were added during testing to the box and the soda can to prevent crushing but allow for minor flexing. The test procedure involved placing an object on a table at specific reference locations that were accurately measured and marked on the table (see Fig. 2b). These locations had a series of evenly spaced radial and axial lines such that the object centroid could be placed accurately on the reference point in the correct position and orientation.

Extra care was taken to ensure that all fingers would make contact simultaneously and the final grasp would closely resemble simulation. This was performed by computing the pre-grasp finger posture for each grasp which would place the fingers at a uniform distance away from the object's surface but at the desired finger spread. This is important because if the fingers did not make contact simultaneously, they would push away the object resulting in grasps and metric values that do not match those planned in simulation. We did this

TABLE II: End-Effector Shake Test Magnitudes

Type	Peak	Mean
Angular Velocity (rad/s)	41.93	4.67
Linear Velocity (m/s)	87.05	0.49
Angular Acceleration (rad/s^2)	40.76	3.49
Linear Acceleration (m/s^2)	86.94	0.44

to minimize such effects and ensure that the physical testing results were closely associated to the generated metric values so that the efficacy of the metrics could be validated. Thus, most of the errors in the grasping process were due to the precision in positioning of the object and robotic hand, rather than due to perception error. In this work, we are not investigating the problem of object perception, but rather are focusing only on the grasp quality prediction problem.

When grasping the object, the grasp controller used was the default controller provided by Barrett which closes all of the fingers simultaneously and stops each finger when a force or torque threshold is exceeded. After the robot hand closed on the object, the object was lifted and subjected to a series of rigorous disturbances. The disturbance was created by rotating the each of the three wrist joints sequentially from the current joint position to the furthest joint limit and then back to its starting position. This was done so that the object would be subjected to forces in all of the gripper's primary axes and would experience translational as well as rotational forces. The acceleration and velocity magnitudes created by the disturbances are provided in Table II and are comparable or greater in magnitude to disturbances used in evaluation procedures in prior work [5].

Each grasp was tested ten times for a total of 5220 trials, and a binary score (success or fail) was recorded for each test. A specific grasp execution was considered a failure if the object fell or slipped and hit the table during the shake process. The success and failure binary scores from the ten trials were averaged to compute a mean performance score for each grasp. A grasp sample was labeled "good" if it had a performance score greater than or equal to 80%, and labeled "bad" otherwise. This 80% threshold was based on a realistic consideration of the state of the art in automatic robotic grasp generation, where one in four automatically generated grasps failed even in ideal laboratory conditions [5]. However, our algorithms can easily be extended to higher thresholds of performance.

D. *Quantitative Evaluation of Grasp Metrics*

The grasp metric data was normalized to a mean of 0 as

$$x_{(m,sph)} = \frac{x_{(m,n)} - \overline{x}_m}{\sigma_m}, \quad (1)$$

where for a given metric m and n data points, $x_{(m,sph)}$ is the deviation value for the observation $x_{(m,n)}$ with sample mean $\overline{x}_m$ and sample standard deviation σ_m. Normalizing data is important when using dimensionality reduction techniques such as PCA so that raw metric values with large ranges do not skew the analysis. Most importantly, normalization does not alter the ability of each metric to predict grasp quality.

A two-tailed t-test (p-value $\leq$ 0.05) was used to determine if the grasp metric's values were significantly different between good and bad grasps. A metric that showed a statistically significant difference between good and bad grasps was considered to be a good metric which will benefit a grasp planner. In addition, a simple classifier was built based on thresholding over the grasp metric value to determine if a grasp was good or bad. Specifically, if the metric was greater or less than a desired threshold value, the grasp was considered a good grasp. These two methods help provide a baseline of how discriminative a grasp metric is. This simple classifier was compared with the GP based classifier (see section III-F).

E. *Dimensionality Reduction Using Principal Component Analysis and Statistical Testing*

Even though multiple grasp metrics are utilized to describe the grasp, it is possible that the grasp sample data may have smaller intrinsic dimensionality due to (i) strong correlations between the grasp metrics and (ii) poor predictive ability of some grasp metrics. In order to deal with the correlated metrics, we use Principal Component Analysis (PCA) to perform a dimensionality reduction of our data by reducing the data to only a few dimensions in the full dimensional space [31].

First, those metrics that did not show statistical significance in the t-tests between good and bad grasps were eliminated (see section III-D). Then PCA was applied to all the remaining dimensions and the data variance captured by the different principal components was analyzed to determine if some principal components contributed more to the data variance than others.

F. Building a Gaussian Process-based Classifier for Grasp Quality Prediction

The high complexity of the grasp space makes it prohibitively difficult to manually develop a custom, composite metric, and is ideally suited for a machine learning algorithm such as GP to merge the information provided by each metric. In this work, we utilize a GP with a squared exponential covariance function with an Automatic Relevance Determination distance measure[3]. Once the desired grasp metrics and principal components from PCA were selected (see section III-E), a cross-validation technique using an randomized 80/20 split, where 80% of the grasp sample set was randomly chosen to train the GP classifier and the remaining 20% of the grasp sample set was used to test the classifier [31]. This process was repeated one hundred times and the average performance of the GP-based classifier using a threshold was recorded.

The GP-based classifier's prediction was used to create a receiver operating characteristic (ROC) curve to analyze performance trade-offs. ROC is a common tool used in the machine learning community for evaluating a classifier's performance [32]. The ROC curve's shape indicates how effective the classifier is at keeping False Positive Rates (FPR) low and True Positive Rates (TPR) high. The TPR represents the success rate of correctly labeling successful grasps and FPR the rate of incorrectly labeling unsuccessful grasps as successful. After one hundred iterations of training and testing, the area under the curve (AUC) for all the iterations was averaged and the TPR at values of 5%, 10%, and 15% FPR were found. The AUC value represents the classifier's robustness by showing its probability to correctly classify a grasp. An AUC value of 1 indicates perfect performance, and an AUC value of 0.5 indicates random classification. To benchmark the GP classifier, we completed a similar ROC analysis for the simple classifiers based on thresholding on the grasp metrics (see section III-D).

IV. RESULTS

Of the 522 grasps in the dataset, 376 (72%) grasps were good (average success greater than 80%) and the remaining 146 were bad (28%).

A. Discriminative Ability of Individual Grasp Metrics

Table III shows the results of each grasp metric in terms of two aspects: (i) The statistical significance of each metric to discriminate between good and bad grasps based on t-tests, (ii) The performance of a simple classifier built by thresholding on each grasp metric. The table's rows are sorted based on increasing t-test p-values, which indicate that only nine of the twelve grasp metrics can individually differentiate between good and bad grasps for this set of grasps at the $p = 0.05$ statistical significance level. In Fig. 3, the ROC curves (mean$\pm$standard error over one hundred trials) for the best classifiers built by thresholding individual grasp metrics and the best GP-based classifier are shown. It is evident that the GP-based classifier performs better than classification using individual grasp metrics in the regions of low FPR values. Furthermore, classification based on all individual grasp metrics, except energy, is only marginally better than random guessing as shown by the low AUC values and low TPR values in Table III.

TABLE III: Individual Grasp Metric Evaluation

Grasp metric	**t-test p-value**	**AUC value**	**TPR at 10% FPR**
*Finger Extension	4.62e-13	0.65	0.24
*Skewness	2.78e-11	0.65	0.21
*Grasp Energy	1.67e-10	0.79	0.43
*Object Volume Enclosed	1.12e-8	0.65	0.24
*Parallel Symmetry	1.63e-6	0.62	0.14
*Perpendicular Symmetry	1.80e-6	0.56	0.15
*Point Arrangement	1.14e-5	0.57	0.13
*Finger Spread	2.56e-4	0.56	0.13
*Finger Limit	4.56e-3	0.61	0.12
Triangle Size	0.28	0.51	0.05
Epsilon	0.79	0.53	0.12
Grasp Wrench Volume	0.97	0.52	0.02

$*$p-value $<$ 0.05, which indicates strong discriminative power

B. Principal Component Analysis of the Grasp Sample Set

The results from performing principal component analysis on all twelve dimensions of the grasping data showed that there is significant information in all of the components. Specifically, the cumulative variance explained by each additional principal component increases almost linearly (correlation to a 45° slope line is 0.97). However, comparing the AUC values for a GP classifier using varying numbers of principal components (PC), the AUC increased from 0.76 with one PC to 0.82 with four PCs, after which there was no further improvements for adding additional PCs. While the variance explained data implies that there is significant information in each PC, the AUC values from the GP shows that more than half of the PCs

[3]http://www.gaussianprocess.org/gpml/code/matlab/doc/

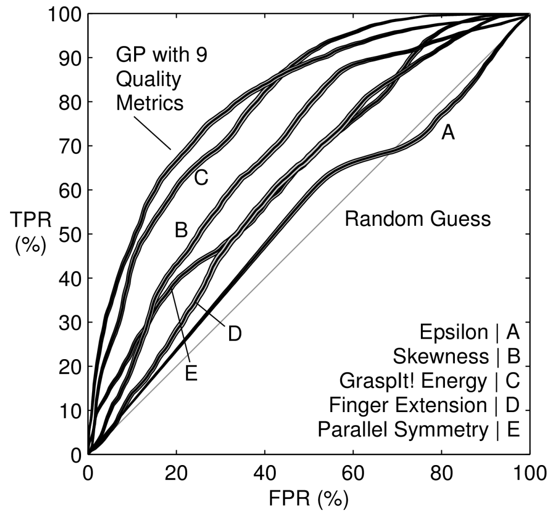

Fig. 3: ROCs of several representative grasp metrics and GP classifier (mean±standard error over one hundred train/test cycles).

can be excluded without affecting the performance of the GP classifier. However, testing would need to be done on a case by case basis for each new data set to confirm that some of the PCs could be excluded since the variance explained is insufficient alone to account for this discrepancy.

C. Performance of the GP-based Classifiers

Table IV shows the results from building and testing GP-based classifiers using all the grasp metrics and using all the principal components derived from subsets of the statistically significant grasp metrics. The results show that decreasing the number of grasp metrics used in the PCA process (but using all the principal components) based on the t-test performance significantly improves the TPR values of GP-based classifiers at a FPR of 5%. However, at the 10% and 15% FPR values, the data shows that using nine grasp metrics provides the best TPR values. Additionally, comparing Table III to the 10% FPR column of Table IV shows the significantly improved performance of the GP classifier over simple thresholding of the individual metrics.

Figure 4 presents a visualization of a two-dimensional projection of the classification surface the GP creates for evaluating grasp quality. This particular GP is built using all principal components of the top nine grasp metrics from Table III. Despite the non-linearities, it is clear that the GP has been successful in finding a boundary that divides the good and bad grasp region.

TABLE IV: GP performance using PCA on different number of grasp metrics: TPR and AUC values

Number of Grasp Metrics Used	TPR			AUC
	FPR =5%	FPR =10%	FPR =15%	
1	0.11	0.22	0.31	0.65
2	0.08	0.22	0.33	0.71
3	0.20	0.37	0.46	0.78
9	0.38	0.50	0.58	0.81
12	0.32	0.47	0.56	0.80

∗All scores are statistically different ($p < 0.05$)

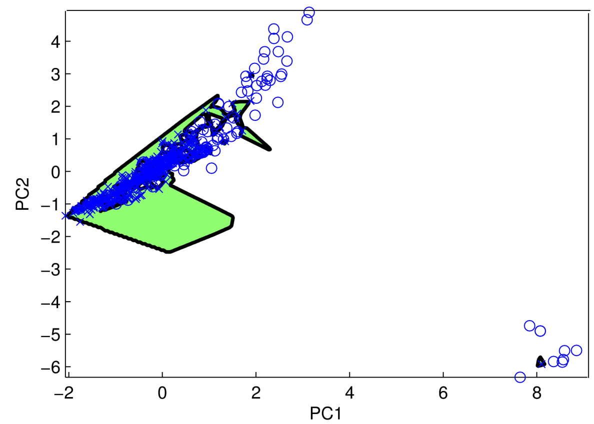

Fig. 4: Visualization of a two dimensional projection of the a nine-dimensional surface that the GP creates to predict grasp quality. The "x" indicates good grasps and "o" indicates bad grasps from the grasp sample set. The filled area represents the "good" grasp region with success rate greater than 83% and a 10% FPR classification level.

V. DISCUSSION

Accurately predicting grasp quality is a challenging problem, given the significant amount of uncertainty in the grasping process and lack of clarity in which grasp metrics correctly predict grasp performance. Table III shows that many of the grasping metrics commonly used in the robotics literature are weak predictors of grasp quality. However, the t-test procedure proved to be a good method for determining which grasp metrics were important and may be used to build a classifier that combines the metrics to improve classification performance.

Using the grasp metrics in Table I, the GP-based classifier (TPR=0.38) significantly improved over a classifier based on simple thresholding of individual grasp metrics (energy TPR=0.23) resulting in a 66% improvement in the TPR rate at an FPR of 5%.

This was because the GP-based classifier non-linearly merged the signals from multiple metrics. The key finding was that the grasp metrics which have low discriminative ability only serve to introduce noise into the classifier and make it more difficult for GP to learn the grasp quality function. These metrics had low discriminative power due to the users' preference of power grasps over precision grasps. Specifically, some of the objects and grasps were such that the fingers wrapped around the object, but there was no palm contact when the users created the grasp simulation. This resulted in grasps which did not have force closure, and thus had small grasp wrench values and zero epsilon. However, when executed, these grasps performed very well as they were able to fully enclose some or all of the object. Overall, the successful grasps had widely varying grasp wrench and epsilon scores, resulting in their low classification performance and exclusion from the GP. If different objects and grasps were selected, then these metrics could prove significant and be reintroduced into the GP.

The results from performing PCA implies that a linear dimension reduction technique may not be sufficient, however it was utilized to reduce the number of dimensions to decrease the GP training time. Looking at the PCA plot (Fig. 4), it is clear that the current data set's spread can be improved, given the clustering of grasp samples in the $(-2 < \mathrm{PC}_1 < 2, -1 < \mathrm{PC}_2 < 2)$ range. Also, the grasp data set was skewed towards higher performing grasps, which could influence the GP classification performance in the region of bad grasps. Future work will include non-linear dimensional reduction methods and expanding the grasp data set to include larger regions of the grasp space.

Similar experiments have been performed in prior work but usually on smaller data sets. Specifically, ninety grasps were generated across four planar objects and tested a total of 920 times and were able to achieve an average prediction success rate of about 76% [19]. Another group tested thirteen novel 3-D objects across 150 trials and achieved a prediction success rate of 81% across all objects [25]. In our work, our experiment used 522 grasps on nine objects a total of 5220 trials and was able to achieve a higher TPR at low FPR levels and an overall success rate of 88% (at 5% FPR). A key advantage of our work is the ability to select a desired FPR level for the prediction performance. However, given the complexity of the grasping problem, more grasp examples and validation over more platforms are needed to improve the grasp predictor's performance and to find regions of strong or weak performance.

One key advantage the GP-based classifier offers over the individual metrics is the significantly higher TPR at low FPR values. This will significantly reduce the online computation time required by reducing the number of rejected good grasps. For example, the online grasp planner GraspIt! searches about seventy five grasps a second in order to provide about thirty valid grasps [5]. With GP's higher TPR, this number can be increased to forty valid grasps which can improve the grasp performance especially in constrained environments where typical grasps are not possible. Alternatively, the computation time could be reduced for the same number of candidate grasps, resulting in better performance for real-time robotics.

One limitation of this research is the lack of including the dynamics of the grasping process in the metric computation. While great care was taken to ensure that the object moved negligibly during the grasping process, additional development to include grasping dynamics in grasp quality prediction would further improve the results as well as open up a new field for making grasp predictions of flexible and compliant objects. Second, we used a robotic platform commonly used for both research [33] and development[4] to make the results broadly applicable. However, more testing is needed to transfer the results to other robotic platforms with differing capabilities. One possibility is to add learning from tactile information [34], which would be required for grasping a new object for which no model information exists, or if the object cannot be identified properly due to sensor error or occlusion.

REFERENCES

[1] B. Kehoe, D. Berenson, and K. Goldberg, "Toward cloud-based grasping with uncertainty in shape: Estimating lower bounds on achieving force closure with zero-slip push grasps," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 576 – 583, 2012.

[2] C. Ferrari and J. Canny, "Planning optimal grasps," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 3, pp. 2290 – 2295, 1992.

[3] R. Pelossof, A. Miller, P. Allen, and T. Jebara, "An svm learning approach to robotic grasping," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2004, no. 4, pp. 3512 – 3518, 2004.

[4] D. R. Faria, R. Martins, J. Lobo, and J. Dias, "Extracting data from human manipulation of objects towards improving autonomous robotic grasping," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 396 – 410, 2012.

[4] www.thearmrobot.com

[5] R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka, “Physical human interactive guidance: Identifying grasping principles from human-planned grasps,” *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 899 – 910, 2012.

[6] M. Ciocarlie, C. Goldfeder, and P. Allen, “Dimensionality reduction for hand-independent dexterous robotic grasping,” *IEEE International Conference on Intelligent Robots and Systems*, pp. 3270 – 3275, 2007.

[7] E. Chinellato, A. Morales, R. B. Fisher, and A. P. del Pobil, “Visual quality measures for characterizing planar robot grasps,” *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews*, vol. 35, no. 1, pp. 30 – 41, 2005.

[8] Q. Lin, J. W. Burdick, and E. Rimon, “Stiffness-based quality measure for compliant grasps and fixtures,” *IEEE Trans. Rob.*, vol. 16, no. 6, pp. 675 – 688, 2000.

[9] R. Suárez, J. Cornellà, and M. R. Garzón, *Grasp quality measures*. Institut dÓrganització i Control de Sistemes Industrials, 2006.

[10] K. B. Shimoga, “Robot grasp synthesis algorithms: A survey,” *The International Journal of Robotics Research*, vol. 15, no. 3, pp. 230–266, 1996.

[11] S. El-Khoury, A. Sahbani, and P. Bidaud, “3d objects grasps synthesis: A survey,” *13th World Congress in Mechanism and Machine Science*, pp. 573 – 583, 2011.

[12] E. Lopez-Damian, D. Sidobre, and R. Alami, “A grasp planner based on inertial properties,” *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2005, pp. 754 – 759, 2005.

[13] J. Bohg, A. Morales, T. Asfour, and D. Kragic, “Data-driven grasp synthesis—a survey,” *IEEE Transactions on Robotics*, vol. PP, no. 99, pp. 1–21, 2013.

[14] B. Bounab, D. Sidobre, and A. Zaatri, “Central axis approach for computing n-finger force-closure grasps,” *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 1169 – 1174, 2008.

[15] S. El-Khoury and A. Sahbani, “On computing robust n-finger force-closure grasps of 3d objects,” *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 2480 – 2486, 2009.

[16] B. Leon, J. L. Sancho-Bru, N. J. Jarque-Bou, A. Morales, and M. A. Roa, “Evaluation of human prehension using grasp quality measures,” *International Journal of Advanced Robotic Systems*, vol. 9, 2012.

[17] E. Chinellato, R. B. Fisher, A. Morales, and A. P. Del Pobil, “Ranking planar grasp configurations for a three-finger hand,” *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 1, pp. 1133 – 1138, 2003.

[18] J. Weisz and P. K. Allen, “Pose error robust grasping from contact wrench space metrics,” *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 557 – 562, 2012.

[19] A. Morales, E. Chinelalto, P. Sanz, A. Del Pobil, and A. H. Fagg, “Learning to predict grasp reliability for a multifinger robot hand by using visual features,” *Proceedings of the Eighth IASTED International Conference on Artificial Intelligence and Soft Computing*, pp. 249 – 254, 2004.

[20] A. T. Miller and P. K. Allen, “Graspit: A versatile simulator for robotic grasping,” *IEEE Robotics and Automation Magazine*, vol. 11, no. 4, pp. 110 – 122, 2004.

[21] Z. Li and S. Sastry, “Task-oriented optimal grasping by multifingered robot hands.” *IEEE journal of robotics and automation*, vol. 6, no. 2, pp. 32 – 44, 1988.

[22] D. Kirkpatrick, B. Mishra, and C.-K. Yap, “Quantitative steinitz’s theorems with applications to multifingered grasping,” *Discrete & Computational Geometry*, vol. 7, no. 1, pp. 295–318, 1992.

[23] A. T. Miller and P. K. Allen, “Examples of 3d grasp quality computations,” *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2, pp. 1240 – 1246, 1999.

[24] A. Morales, E. Chinellato, A. H. Fagg, and A. P. Del Pobil, “Experimental prediction of the performance of grasp tasks from visual features,” *IEEE International Conference on Intelligent Robots and Systems*, vol. 4, pp. 3423 – 3428, 2003.

[25] A. Saxena, L. L. Wong, and Y. Andrew, “Learning grasp strategies with partial shape information,” *Proceedings of the National Conference on Artificial Intelligence*, vol. 3, pp. 1491 – 1494, 2008.

[26] C. E. Rasmussen, *Gaussian processes for machine learning*. MIT Press, 2006.

[27] A. Saxena, L. Wong, M. Quigley, and A. Y. Ng, “A vision-based system for grasping novel objects in cluttered environments,” *Springer Tracts in Advanced Robotics*, vol. 66, no. STAR, pp. 337 – 348, 2010.

[28] R. Diankov, “Automated construction of robotic manipulation programs,” Ph.D. dissertation, Carnegie Mellon University, Robotics Institute, August 2010.

[29] W. Townsend, “The barretthand grasper-programmably flexible part handling and assembly,” *Industrial Robot: An International Journal*, vol. 27, no. 3, pp. 181–188, 2000.

[30] D. Gossow, A. Leeper, D. Hershberger, and M. Ciocarlie, “ROS topics: Interactive markers: 3-d user interfaces for ROS applications,” *IEEE Robotics and Automation Magazine*, vol. 18, no. 4, pp. 14 – 15, 2011.

[31] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2009.

[32] J. A. Hanley and B. J. McNeil, “The meaning and use of the area under a receiver operating characteristic (ROC) curve.” *Radiology*, vol. 143, no. 1, pp. 29–36, 1982.

[33] S. S. Srinivasa, D. Ferguson, C. J. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. V. Weghe, “HERB: a home exploring robotic butler,” *Autonomous Robots*, vol. 28, no. 1, pp. 5–20, 2010.

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