

Human-Guided Grasp Measures Improve Grasp Robustness on Physical Robot

Ravi Balasubramanian, Ling Xu, Peter D. Brook, Joshua R. Smith, Yoky Matsuoka

Abstract—Humans are adept at grasping different objects robustly for different tasks. Robotic grasping has made significant progress, but still has not reached the level of robustness or versatility shown by human grasping. It would be useful to understand what parameters (called grasp measures) humans optimize as they grasp objects, how these grasp measures are varied for different tasks, and whether they can be applied to physical robots to improve their robustness and versatility. This paper demonstrates a new way to gather human-guided grasp measures from a human interacting haptically with a robotic arm and hand. The results revealed that a human-guided strategy provided grasps with higher robustness on a physical robot even under a vigorous shaking test (91%) when compared with a state-of-the-art automated grasp synthesis algorithm (77%). Furthermore, orthogonality of wrist orientation was identified as a key human-guided grasp measure, and using it along with an automated grasp synthesis algorithm improved the automated algorithm's results dramatically (77% to 93%).

I. INTRODUCTION

Young healthy humans have a near-perfect success rate for grasping everyday objects, probably because they spend years as toddlers learning to grasp objects. But surprisingly little is known about the algorithms that humans employ for grasping objects. Currently, robots can achieve this level of success only in an industrial assembly line with grippers and environments specifically designed for a part's shape. If a versatile robotic hand can be programmed to grasp a variety of common objects at close to a 100% success rate, it would change the landscape for the industrial grippers and would also enable personal robotic assistance for the elderly or disabled in everyday environments.

To achieve a high grasping success rate for everyday objects using robots, a variety of features and metrics have been explored as “grasp measures” to judge and optimize grasp quality. Simulation based grasp measures from the prior work include: 1) computing grasp

strength and the largest disturbance-wrench magnitude “Epsilon” that the grasp can resist [15], [10], [16], [21], 2) finding independent contact regions and the distance of the grasp center from center of mass [23], and 3) finding quick grasps using just the object vertices [22] (see [2] for a comparison of these grasp measures using simulation and [29] for a survey of grasp synthesis techniques). Heuristics for hand preshaping and wrist orientation [32] and contact points [3] have also been used for grasp synthesis.

However, due to calibration errors in the physical robot and mismatches between simulated models and real objects/environments, it is critical that a grasp's success be measured when it is expressed on a physical robotic hand rather than just in simulation. The relationship between success in simulation and on a real robot is beginning to be explored. For example, an open-source software called GraspIt! [20] has been used to generate real-world grasps for different robotic hands given an object model. In this paper, we have carefully evaluated GraspIt!'s capability to produce real-world grasps and used it as a benchmark for automated grasp synthesis. While GraspIt! uses some of the grasp measures mentioned above to predict grasp success in the real world (also see [17] where objects with varied mass distributions were considered), some approaches use heuristic grasp measures such as hand grasp volume and hand symmetry [5] to quantify grasp quality on a physical robot. In addition, some approaches distinguish “contact robustness” from “grasp robustness” for grasp synthesis [24]. Finally, new approaches combining machine learning [25] and computer vision [27] have also been used for grasp synthesis. To our knowledge, the best reported success rate for grasping everyday objects using three or more robotic fingers is 60–80% for simple lifting¹ [5], [27].

While some of these strategies are now yielding higher success rates due to new and clever algorithms, we believe that people would abandon a robotic assis-

This work was supported by NIH grants and Intel Labs Seattle. R. Balasubramanian (bravi@cs.washington.edu), P. Brook, and Yoky Matsuoka are with the University of Washington and J. R. Smith with Intel Labs Seattle. L. Xu is at Carnegie Mellon University (work was done while she interned at the University of Washington).

¹The 80% result is from experiments with one novel object over five trials [27]. A higher success rate of 87.8% [27] has been reported based on more experiments with novel objects and a parallel-plate gripper, which has a highly limited task space compared to multifinger hands.

tant that drops an object one out of five times, just as unreliable or hard-to-use prosthetic limbs have little acceptance. In addition, we are interested in “robustness” of grasps; that is, we are not only interested in whether the object can be picked up, but whether the object can be held stably in the presence of disturbances and uncertainty in modeling and actuation. So our testing procedure includes shaking the robot vigorously after the object is grasped.

Another key aspect of grasping that humans can provide is task specificity. For example, one could grasp a stapler with the intention of lifting it up for transportation, handing over to someone else, or stapling papers. While automated grasp synthesis using task-specific forces and torques have been explored before [16], extracting position-based task-specific grasp strategy is currently not possible by other techniques.

In order to approach a near-perfect success rate even with perturbation and to provide us with task-specific strategies, we took an approach to let humans physically guide a robot to demonstrate the grasping strategy that they would choose for a given task. Our goal is to use the collected grasps to learn the rules used by humans for grasping and then generalize and improve automated grasping techniques. Please note that this paper is not claiming that this is the only technique to use over others or to say that this technique “as is” is scalable.

However, we *do* claim that the strategies humans use to grasp robustness have not been completely identified. Rather than guessing what grasp measures humans may use or what in general may be good for grasping, we observed and collected data from humans, extracted the grasp measures, and then evaluated their effect on an automated grasp synthesis algorithm. We believe that these human-based grasp measures can speed up and produce higher robustness results for existing grasp generators.

From human-subject data collected in one full day (described in section II), we demonstrate that a human guided grasping feature improves real-world robot grasping robustness significantly in section III. This is our first paper with this new human guided technique, and we discuss scalability, task-specific grasp measures, and how to achieve a near-perfect success rate in section IV.

II. METHODS

A. Human Haptic Interaction Environment

The Human Haptic Interaction Environment was a framework where a person could teach a robot different

grasps by being in the robot’s workspace and physically interacting with the robot. This interaction method allowed the person to guide the robot to specific wrist configurations and finger postures, and these grasps were called “human-guided grasps”. Such interactive robotic grasping with a human in the loop has been explored before by the GraspIt! group [6], but only wrist posture was controlled by the human and finger posture was controlled by GraspIt!. Their purpose was to demonstrate how GraspIt! goes through search iterations to generate a grasp for a given wrist position. The purpose of our experiment was to collect human-guided grasping strategies and identify features that may not be expressed properly in other methods. While more scalable approaches like simulation or video-based interactions could have been used to gather human strategies, as our first experiment with human subjects, we did not want to lose important grasp measures by placing people in an artificial environment that forced them to visualize the objects/robotic hands on a two-dimensional display. Haptic interaction was the most intimate way for the human subjects to know the object properties and to build their own internal model of robotic hand shape and capabilities. Others have considered motion-capture approaches to convert human movements to robot motions [26], but the difference in kinematics between the human body (particularly the human hand) and robot makes it extremely challenging to transfer movements.

The Human Haptic Interaction Environment used a robotic platform consisting of a seven degree-of-freedom Barrett Whole Arm Manipulator robotic arm and a three-fingered four degree-of-freedom Barrett-Hand [1]. The robotic hand was equipped with electric field sensors [31] which enabled the fingers to detect their proximity to objects. The electric field sensors enabled all three fingers to close in on the object simultaneously.

The Human Haptic Interaction Environment placed the robotic arm in a “gravity compensation” mode, where the arm had negligible weight and could be easily moved by a person. The object to be grasped was placed at a known location and orientation in the workspace. The robotic arm was reset to a neutral position in the workspace and the fingers of the hand were kept open. The grasp guidance process proceeded as follows: 1) The human subject guided the robot to an initial wrist pose at which the object could be grasped (see Figs. 1a and 1b). 2) Using electric-field sensing, the fingers closed into the object so that each fingerpad was approximately 5 mm from object surface.




Fig. 1. The experiment procedure of a human subject guiding the robot to grasp an object: (a), (b) approach the object, (c) adjust wrist orientation and finger spread, (d) fingers close in on the object, and (e) lift object.

At this point, the human subject guided the finger posture by haptically moving the spread angle of the fingers. Additionally, the subject could re-adjust wrist pose to better align the fingers with the object (see Fig. 1c). 3) When the subject was satisfied with this grasp pose, the robotic fingers were commanded to close around the object, completing the grasp teaching procedure. The final closure step was guided by the electric-field sensors so that all fingers contacted at the same time to not perturb the object (see Fig. 1d). 4) Subjects were then allowed to lift and shake the robotic arm to determine if they liked the grasp. If the subject did not like the grasp or if the object slipped out, the grasp was not considered a human guided grasp (see Fig. 1e). We eliminated such grasps because the key idea was to collect the best grasps that humans can provide. Since the subjects had little experience with the system, the procedure provided an opportunity for the subjects to review the grasps. It turned out in the experiment described in the next section that less than five percent of all the human guidance grasps were eliminated because the subject was not satisfied with the grasp. A valid grasp was represented as an eleven dimensional vector containing the seven degree-of-freedom robot arm joint angles and the four degree-of-freedom (one spread and three flexion) hand joint angles.

B. Human Experimental Paradigm

Seven subjects participated in the study approved by the University of Washington Human Subjects Division, and a total of 210 grasps were collected with the robot. Each subject was given five minutes practice time. Nine objects were used in the experiment: three small objects, three medium-sized objects, and three large objects (see Fig. 2). Each subject was asked to

perform three different tasks for an object, namely, lifting the object, handing the object over, and performing a function with the object. For the handing over task, the subject was asked to grasp the object such that there was room left for someone else to grasp it. The functional tasks depended on the object. For example, for the wine glass, the functional task was pouring and for the phone, the task was picking up to make a phone call. For each object, the subject was asked to perform two grasps for each of the three tasks for a total of six grasps, and the subjects were asked to vary the grasps if they could. Each subject was randomly assigned to five objects, while ensuring that we had an even distribution of grasps for each of the objects (each object was selected four times except for the soda can which was selected three times).

From the eight human-guided grasps for each object-task pair (six for the soda can), three were randomly chosen for testing on the robot (3 candidate grasps \times 3 tasks \times 9 objects), and each grasp was tested five times. Each time, the robot arm was commanded to the recorded arm joint angles with the fingers full opened. The robot hand was then commanded to the required spread angle. Finally, the fingers were commanded to close in quickly on the object, and the robot lifted the object and then executed a shaking procedure where the object was shaken four times in a continuous circular motion (absolute mean (peak) values for angular velocity: 2.74 (4.62) rad/s, linear velocity: 0.39 (0.62) m/s, angular acceleration: 2.22 (4.39) rad/s², linear acceleration: 0.33 (0.63) m/s²). If the object stayed in the hand after the shaking, it was considered a success (rated 1); otherwise, a failure (rated 0). Note that the grasp testing procedure was intentionally kept simple to maintain focus on grasp generation rather than grasp testing. Also the human subjects did not know that the grasp would be tested by shaking and were only providing grasps for the various tasks.

The success rate was computed for each grasp by averaging over the five trials. Hypothesis testing was performed with a p-value of 0.05, and standard errors were reported for all mean values.

C. Grasp Success Validation on Physical Robot

To compare human-guided grasping against a good simulation-based technique, we ran GraspIt! for thirty minutes with the intention of generating the top six grasps for each object (using the same procedure in [11]). While we expected to have a total of 54 grasps, we ended up with a total of 49 grasps because the grasp search yielded fewer than six grasps for some objects due to search complexity and time limit. In addition, an



Fig. 2. Objects used in the experiment.

inverse kinematics solution did not exist for some of the grasps because the robot was stationary relative to the table and object in the testing set-up. Specifically, three objects had fewer than six grasps (wine glass: 4, coil of wire: 5, one-liter bottle: 4), while the remaining objects had six grasps each, yielding a total of 49 grasps. Note that GraspIt! cannot provide task-specific grasps, and its grasps are intended for lifting tasks only. The GraspIt! grasps also were validated using the same process as the human guided grasps.

D. Analysis Using a Grasp Measure Set

Our goal was to understand human grasping strategies in order to improve the robustness of robotic grasping. To do so, we used grasp measures already used in literature [10], [20], [5], [28] and one new measure that became apparent during the human-subject experiments (see Table I). The new measure, orthogonality, measures the orientation of the wrist relative to the object's principal axis and its perpendiculars. Suppose the object principal axis (axis of longest dimension) is u , and the axis pointing out of the palm of the BarrettHand is v . The angle δ between u and v may be computed as $\delta = \arccos(u \cdot v)$. Then the orthogonality measure α is defined as:

$$\alpha = \begin{cases} \delta, & \text{if } \delta < \pi/4 \\ \pi/2 - \delta, & \text{if } \pi/4 < \delta < \pi/2 \\ \delta - \pi/2, & \text{if } \pi/2 < \delta < 3\pi/4 \\ \pi - \delta, & \text{if } \delta > 3\pi/4 \end{cases} \quad (1)$$

In Fig. 3, since the axis pointing out of the palm of the robotic hand in the bottle-lifting task is approximately parallel to the bottle's principal axis (vertical), the orthogonality measure α for that grasp is close to zero. In contrast, the GraspIt! grasp for the stapler would have an orthogonality measure α close to $\pi/4$ rad, the maximum possible value.

III. RESULTS

Fig. 3 shows a sample of the grasp strategy used by human subjects and GraspIt! for three objects. While human subjects varied grasping strategies for different tasks, the GraspIt! grasps are analogous to human

TABLE I

GRASP MEASURE SET

Grasp Measure	Description	Citation
Epsilon ¹	Minimum disturbance wrench that can be resisted	[10], [20]
Wrench space volume ¹	Volume of grasp wrench space	
Grasp energy ²	Hand-object proximity	
Point arrangement ¹	Proximity of fingertips being in a plane parallel to palm.	[5]
Grasp volume ¹	Volume enclosed by hand	
Hand flexion ²	Similarity of finger flexion	
Hand spread ²	Proximity of the finger spread to equilateral triangle.	
Finger limit ³	Extent of finger extensions	
Volume of object enclosed ¹	Object volume enclosed by hand normalized by object volume.	[28], [27]
Parallel symmetry ²	Distance between center of mass and contact point centroid along object principal axis.	
Perpendicular symmetry ²	Distance between center of mass and contact point centroid perpendicular to object principal axis.	
Orthogonality	See section II-D	

¹Larger⇒Better grasp; ²Smaller⇒Better grasp;

³Mid-range⇒Better grasp

grasps for the lifting task. That is, we believe that the grasp measures used by GraspIt! and humans are comparable. In contrast, for the handing-over task, the subjects likely prioritized leaving space on the object for a receiver as against optimizing for a robust grasp. For the functional task, subjects likely optimized the grasp for the subsequent functional movement. Thus, we compared GraspIt! grasps with human-guided lifting grasps in the analysis below.

A. Grasping Success Rate on Physical Robot

Table II presents the success rates for each object (after being shaken vigorously five times) for the human-guided lifting grasps and for the best GraspIt! grasps (a total of 370 trials). Across objects, the human-guided lifting strategy yielded a 91(3)% success rate while GraspIt! yielded 77(3)%. An outlier for the human lifting grasps was the one-liter bottle, without which the success rate for human-guided lifting grasps was 97(1)%.




Fig. 3. Example grasp postures generated by human subjects (for three different tasks) and GraspIt! for three objects. Note that the human subjects manually specified the grasps on the physical Barrett robotic hand, which were then visualized using the OpenRAVE program [7].

TABLE II
MEAN SUCCESS RATES FOR HUMAN-GUIDED GRASPING (LIFTING TASK) AND GRASPIT!

Object	Human guidance Lifting	GraspIt!
Wine glass	93 (7)	100 (0)
One-liter bottle	40 (13)	65 (26)
Soda can	93 (7)	90 (5)
Cereal box	93 (7)	90 (4)
Coil of wire	95 (4)	32 (11)
Phone	100 (0)	70 (6)
Pitcher	100 (0)	83 (5)
Soap dispenser	100 (0)	67 (11)
CD pouch	100 (0)	100 (0)
Overall	91 (3)*	77 (3)*
Number of grasps	25	49

* $p < 0.05$

B. Grasp Measures Used by Humans

Table III shows the range of values for the grasp measures for human-guided lifting and GraspIt! grasps. Four grasp measures, namely epsilon, grasp wrench-space volume, hand flexion, and orthogonality, were significantly different between the human-guided lifting and GraspIt! grasps. In addition, the energy measure showed borderline significant difference ($p = 0.05$) between human-guided lifting and GraspIt! but that was because of outliers. While larger epsilon and volume indicated better grasp quality theoretically, we noticed

TABLE III
GRASP MEASURE VALUES FOR HUMAN-GUIDED LIFTING AND GRASPIT!

Feature	Mean (Standard Error)	
	Human guidance Lifting	GraspIt!
Epsilon*	0.1 (0.02)	0.19 (0.01)
Wrench space volume*	0.19 (0.05)	0.42 (0.04)
Grasp energy	-1.33 (0.09)	3.95 (2.57)
Point arrangement	0.78 (0.02)	0.76 (0.02)
Grasp volume (cm ³)	281 (29)	259 (33)
Hand flexion*	0.05 (0.01)	0.19 (0.04)
Hand spread	0.39 (0.02)	0.37 (0.02)
Finger limit	0.70 (0.05)	0.76 (0.02)
Volume of object enclosed	0.06 (0.01)	0.05 (0.01)
Parallel symmetry	0.30 (0.05)	0.39 (0.03)
Perpendicular symmetry	0.33 (0.03)	0.28 (0.03)
Orientation*	5.2 (1.3)	23.2 (1.86)
Number of grasps	25	49

* $p < 0.05$

from the experiment that epsilon and volume were lower for the human grasps when compared with the GraspIt! grasps even though the human guided grasps have a higher success rate than the GraspIt! grasps. The hand-flexion measure indicated that humans use grasps which have similar finger flexion values when compared with the GraspIt! grasps.

The stand-out measure however was orthogonality. The orthogonality measure for the human grasps was significantly smaller than for the GraspIt! grasps, indicating that wrist orientation in the human grasps is much closer to the object's principal axis or its perpendiculars (see Fig. 3; the principal axis for the bottle and wine glass was vertical and phone horizontal). Fig. 4 shows the orientation box plots for three objects for all human and GraspIt! grasps.

C. GraspIt! Performance Improvement with Human Grasp Measures

Each grasp, whether from GraspIt! or human guidance, is stored as a eleven dimensional vector containing the seven robot arm angles and four hand joint angles. We divided all the grasps into two groups: Group 1 is the set of grasps obtained by merging the set of grasps from human-guided lifting and the set of grasps from GraspIt!. Group 2 consisted of GraspIt! grasps only. Fig. 5 shows the variation in success rates for the two groups of grasps, each split by an orientation angle threshold of 13 degrees. Grasps whose orthogonality measure was less than 13 degrees were considered orthogonal. This result showed that the success rate of GraspIt! grasps with low orientation value was significantly higher than GraspIt! grasps with a large orientation value (93(5)% compared with 77(3)%). In

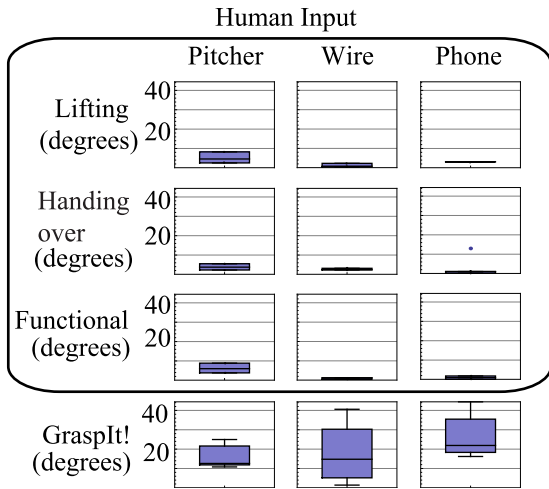


Fig. 4. Variation of the orthogonality measure for three objects. The first three rows represent the human grasps and the fourth row the GraspIt! grasps. Each box-plot shows the spread of the orthogonality measure, where the blue box spans from the 25% quantile to the 75% quantile. The whiskers represent the extremes of the data and the (blue) dot an outlier.

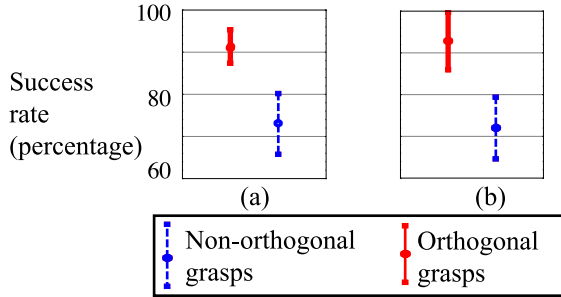


Fig. 5. Success rates for orthogonal (< 13 degrees) and non-orthogonal grasps from two groups: (a) Human lifting grasps combined with GraspIt! grasps (orthogonal and non-orthogonal grasps $n = 37$ each), and (b) GraspIt! grasps only (orthogonal grasps $n = 14$, non-orthogonal grasps $n = 35$).

contrast, when investigating the significance of the hand-flexion measure for grasping, we did not see a significant difference in grasp success for grasps with small hand-flexion measures when compared with grasps with large hand-flexion measures. This indicated that a low hand-flexion measure was likely not a reason for a better grasp.

D. Task-Dependent Variation in Human Performance

As seen in Fig. 3, humans varied the grasping strategy for different task requirements. Table IV shows the success rate for the handing over and functional tasks (a total of 265 trials). We note that the success rates for these tasks are lower than the success rates for the lifting task. Grasp measures used by the handing-over and functional tasks remained statistically indifferent from lifting task except for hand flexion feature ($p < 0.05$).

TABLE IV

MEAN SUCCESS RATES FOR HUMAN-GUIDED GRASPING FOR THE HANDING-OVER AND FUNCTIONAL TASKS

Object	Human guidance	
	Handing-over	Functional
Wine glass	33 (13)	100 (0)
One-liter bottle	67 (13)	93 (7)
Soda can	90 (9)	100 (0)
Cereal box	87 (10)	100 (0)
Coil of wire	100 (0)	60 (13)
Phone	75 (10)	50 (13)
Pitcher	100 (0)	100 (0)
Soap dispenser	87 (9)	100 (0)
CD pouch	100 (0)	67 (13)
Overall	82 (3)	86 (3)
Number of grasps	28	25

The lack of difference was surprising, and we need to possibly find more appropriate grasp measures (than those measures listed in Table I) and object-task pairs that are suitable for differentiating human task-specific strategies.

IV. DISCUSSION

A. Human-Guided Grasps and Their Robustness

We wanted to see if a human's near-perfect grasping performance with her own hand transferred to successful grasping using a real robot. Table II shows that the human guided grasps resulted in more robust grasps than GraspIt! when expressed on a real robot. While GraspIt! likely produces some of the best automated grasps, the mismatch between simulation models and the real world can cause automated grasp synthesis to fail. Furthermore, humans have an advantage from their strong sense of causal physicality for tool use [12].

The orthogonality of the wrist orientation may seem obvious when we think about how most of the objects in the world are designed with Cartesian coordinate frames. With these Cartesian objects, palm contact and finger placement may be improved when the wrist orientation is parallel to or perpendicular to the object's principal axis. Since the BarrettHand has a flat palm, orthogonal grasps are likely to generate more palm contact which creates a more robust grasp. Wrist orientation parallel to the ground has been used as a heuristic earlier for grasp synthesis [32], where it was claimed that grasp orthogonality likely comes from environmental constraints or the relative object location. It would be interesting to explore further if object shape influences grasp orthogonality as well. Finally, human motor control literature has shown that many motor neurons encode human movements in extrinsic Cartesian coordinate frames rather than intrinsic (muscle or joint) coordinate frames [13].

B. Implication for Automated Grasp Synthesis

One of goals of this work is to use human skill to identify key grasp measures that can speed up automated grasp synthesis and improve real-world grasp quality. Table III shows that the orientation feature has significantly different values for human grasps and GraspIt! grasps. Furthermore, Fig. 5 shows that orthogonal grasps have significantly higher success rate than non-orthogonal grasps. These results indicate that an automated search process can focus on grasps with small orientation values before exploring grasps with larger orientation values. This will likely result in better grasps faster for GraspIt! and other automated grasp synthesis methods.

This paper did not further analyze the grasp measures that produced similar results between human-guided grasps and GraspIt!. This is because our data only contained highly successful grasps and thus it could not be used to identify good and bad grasp measures, unless significant differences were found between human-guided and GraspIt! grasps. Also, the lack of correlation between epsilon and grasp wrench space volume with the high human-guided grasp success rates is worth investigating further to inform the grasp measures used by the grasping research community.

C. Achieving 100% Robustness

Human guidance has produced 91(3)% success rate for multi-fingered grasping with vigorous shaking and a 100% success rate for grasping without shaking. However, we believe that a robotic hand with 91(3)% success rate is still not good enough as a prosthetic, assembly line, or personal assistance device where a near-perfect success rate is desired. So how can we achieve even higher success rates?

We believe that we can make several changes to our experiment protocol to improve on this result. We collected data from subjects who had never seen or interacted with a robotic arm/hand before. It is possible that with more practice with the robot, a subject would provide better grasping strategies. Second, we asked human subjects to vary the grasping strategy every trial, if they could. In retrospect, we should not have forced people to devise different grasping strategies as we do not believe that there are always multiple optimal solutions. Third, the subjects were not informed of the vigorous shaking used in the robustness test. If the subjects had known, they may have chosen different grasps.

One outlier in the human guided grasps success rates is the success rate for the one-liter bot-

tle (only 40(13)%, see Table II). If this outlier is removed, the human grasping success rate is 97(1)% even with vigorous shaking. As seen in Fig. 1, subjects chose to grasp the bottle from the top, when most humans with their own hand would not grasp a filled bottle this way. This strategy was chosen when we instructed subjects to vary the grasps when they could. This technique did not work well on the bottle's slippery surface and large mass.

Finally, it is worth noting that this paper is based on experiments with the BarrettHand, which is widely used and is a great first tool for comparing results across the grasping research community. While highly reliable, the BarrettHand is not backdriveable and is not as anthropomorphic or versatile as some of the newer robotic hands. It would be interesting to quantify success rates with more compliant robotic hands such as the SDM hand [8] or more anthropomorphic hands such as the ACT hand [30], Robonaut hand [19], and DLR hand [4] and with additional sensing capabilities like computer vision and touch sensors.

D. Prediction of Grasp Success

An important goal for the grasping community is to predict real-world grasp quality for novel objects from simulation [27]. While the grasp measures in Table I have been used to compute grasp quality in simulation, these predictions do not always extend to the real world. Using parallel grippers and computer vision, the grasp success for novel objects was shown to be 87.8% [27]. We are interested in having equally high or higher success rate for multi-fingered grasps with our approach. Unfortunately, we do not have sufficient number of failed examples to build a grasp classifier, because our grasps were heavily biased towards success. With more grasp results and modified grasp measures based on human strategies, we can then build a grasp classifier and test our algorithm on novel objects.

E. Widening and Generalizing the Grasp Database

Data collected using the human haptic interaction approach has exceptional quality, but the process is labor intensive and not scalable for many objects, tasks, and robotic hands. We plan to compare the haptic interaction with other less labor-intensive modes of human-robot interaction for grasping purposes. These modes include remote control operation (with robot in sight) [9], remote control with a video camera [18], and remote control through a simulated environment [14]. This multi-modal approach will allow us to understand the fidelity required in gathering human data for robotic grasping and also collect additional data to complement

other grasp databases [11]. Ultimately, the more we understand human grasping strategy and how to bring a 91% success rate to 100%, the less accurate information we need to complete the database, capture a useful grasp measure set, and generalize grasps to new objects. Finally, while we have already collected unique and exciting data on task-specific grasps, future work involves analysis of the grasp measures that explain the variability of grasps with tasks.

V. ACKNOWLEDGMENT

We gratefully acknowledge the contribution of Louis LeGrand for interesting discussions on robotic grasping and Brian Mayton for help with the robot. We also thank the OpenRAVE community for help with using their software, Matei Ciocarlie for help with GraspIt!, and Eric Rombokas for reviewing this draft.

REFERENCES

- [1] <http://www.barrett.com/robot/index.htm>.
- [2] G. M. Bone and Y. Du. Multi-metric comparison of optimal 2d grasp planning algorithms. In *Proc. IEEE Internat. Conf. on Robotics and Automation*, 2001.
- [3] C. Borst, M. Fischer, and G. Hirzinger. A fast and robust grasp planner for arbitrary 3d objects. In *Proc. of the Int. Conf. on Robotics and Automation*, pages 1890–1896, 1999.
- [4] J. Butterfass, M. Grebenstein, H. Liu, and G. Hirzinger. DLR-Hand II: next generation of a dextrous robot hand. In *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, volume 1, pages 109–114 vol.1, 2001.
- [5] E. Chinellato, A. Morales, R. B. Fisher, and A. P. del Pobil. Visual quality measures for characterizing planar robot grasps. *IEEE Trans. Systems, Man, and Cybernet*, 35(1):30–41, 2005.
- [6] M. T. Ciocarlie and P. K. Allen. On-line interactive dextrous grasping. In *Proc. of Eurohaptics*, 2008.
- [7] R. Diankov and J. Kuffner. OpenRAVE: A planning architecture for autonomous robotics. Technical Report CMU-RI-TR-08-34, The Robotics Institute, Pittsburgh, PA, July 2008.
- [8] A. M. Dollar and R. D. Howe. Towards grasping in unstructured environments: Grasper compliance and configuration optimization. *Adv. Robotics*, 19(5):523–544, 2005.
- [9] C. Fernández, M. A. Vicente, C. Pérez, O. Reinoso, and R. Aracil. Learning to grasp from examples in telerobotics. In *Proc. Conf. on Artificial Intell. and Appl.*, 2003.
- [10] C. Ferrari and J. Canny. Planning optimal grasps. In *Proceeding of the IEEE International Conference on Robotics and Automation*, pages 2290–2295, 1992.
- [11] C. Goldfeder, M. Ciocarlie, H. Dang, and P. Allen. The columbia grasp database. In *Proc. Internat. Conf. Robotics and Automation*, pages 1710–1716, 2009. DOI 10.1109/ROBOT.2009.5152709.
- [12] S. H. Johnson-Frey. Whats so special about human tool use? *Neuron*, 2003.
- [13] S. Kakei, D. S. Hoffman, and P. L. Strick. Muscle and movement representations in the primary motor cortex. *Science*, 285(5436):2136–9, 1999.
- [14] U. Kartoun, H. Stern, and Y. Edan. *Advances in e-engineering and digital enterprise technology-1: Proc. Int. Conf. on e-Engineering and Digital Enterprise Technology*, chapter Virtual Reality Telerobotic System. John Wiley and Sons, 2004.
- [15] D. Kirkpatrick, B. Mishra, and C. K. Yap. Quantitative steinitz’s theorems with applications to multifingered grasping. In *ACM Symp. on Theory of Computing*, pages 341–351, 1990.
- [16] Z. Li and S. S. Sastry. Task-oriented optimal grasping by multifingered robot hands. *IEEE Journal of Robotics and Automation*, 4(1), 1988.
- [17] Y. Liu, G. Starr, J. Wood, and R. Lumia. Spatial grasp synthesis for complex objects using model-based simulation. *Industrial Robot: An International Journal*, 2005.
- [18] J. Lloyd, J. Beis, D. Pai, and D. Lowe. Model-based telerobotics with vision. In *Robotics and Automation, 1997. Proceedings., 1997 IEEE International Conference on*, volume 2, pages 1297–1304 vol.2, Apr 1997.
- [19] C. Lovchik and M. Diftler. The robonaut hand: a dextrous robot hand for space. In *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, volume 2, pages 907–912 vol.2, 1999.
- [20] A. Miller and P. K. Allen. GraspIt!: a versatile simulator for robotic grasping. *IEEE Robotics and Automation Magazine*, 2004.
- [21] A. T. Miller and P. K. Allen. Examples of 3D grasp quality computations. In *Proceedings IEEE International Conference on Robotics and Automation*, pages pp 1240–1246, 1999.
- [22] B. Mirtich and J. Canny. Easily computable optimum grasps in 2-D and 3-D. In *Proc. IEEE Internat. Conf. Robotics and Automation*, pages pp739–747, 1994.
- [23] J. Ponce and B. Faveqon. On computing three- finger force-closure grasps of polygonal objects. *IEEE Trans. on Robotics and Automation*, 1995.
- [24] D. Prattichizzo, J. K. S. Jr. and A. Bicchi. *Experimental Robotics IV: Lecture Notes in Control and Information Sci.*, chapter Contact and grasp robustness measures: Analysis and experiments, pages 83–90. Springer Berlin/Heidelberg, 1997. DOI 10.1007/BFb0035191.
- [25] N. Ratliff, J. A. D. Bagnell, and S. Srinivasa. Imitation learning for locomotion and manipulation. In *Proc. IEEE-RAS Internat Conf Humanoid Robotics*, 2007.
- [26] A. Safonova, J. K. Hodgins, and N. S. Pollard. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Trans. Graphics*, 2004.
- [27] A. Saxena, J. Driemeyer, and A. Y. Ng. Robotic grasping of novel objects using vision. *Int. J. Robotics Res.*, 27(2):157–173, 2008.
- [28] A. Saxena, L. L. S. Wong, and A. Ng. Learning grasp strategies with partial shape information. In *Proc. AAAI Conf Artificial Intell.*, 2008.
- [29] K. B. Shimoga. Robot grasp synthesis algorithms: A survey. *Internat. J Robotics Res*, 1996. DOI: 10.1177/027836499601500302.
- [30] M. Vande Weghe, M. Rogers, M. Weissert, and Y. Matsuoka. The ACT hand: design of the skeletal structure. In *Proc. IEEE Internat. Conf. Robotics and Automation*, volume 4, pages 3375–3379, 2004.
- [31] R. Wistort and J. R. Smith. Electric field servoing for robotic manipulation. In *Proc. IEEE/RSJ Internat Conf. Intell. Robots and Sys.*, 2008.
- [32] D. Wren and R. Fisher. Dextrous hand grasping strategies using preshapes and digit trajectories. In *Proc. IEEE Internat. Conf. Sys., Man, and Cybernet.*, volume 1, pages 910–915 vol.1, Oct 1995.