

Pose Error Robust Grasping from Contact Wrench Space Metrics

Jonathan Weisz and Peter K. Allen

Abstract—Grasp quality metrics which analyze the contact wrench space are commonly used to synthesize and analyze preplanned grasps. Preplanned grasping approaches rely on the robustness of stored solutions. Analyzing the robustness of such solutions for large databases of preplanned grasps is a limiting factor for the applicability of data driven approaches to grasping. In this work, we will focus on the stability of the widely used grasp wrench space epsilon quality metric over a large range of poses in simulation. We examine a large number of grasps from the Columbia Grasp Database for the Barrett hand. We find that in most cases the grasp with the most robust force closure with respect to pose error for a particular object is not the grasp with the highest epsilon quality. We demonstrate that grasps can be reranked by an estimate of the stability of their epsilon quality. We find that the grasps ranked best by this method are successful more often in physical experiments than grasps ranked best by the epsilon quality.

I. INTRODUCTION AND RELATED WORK

Analyzing the robustness of a grasp in the presence of object localization errors and external force perturbations is an important and difficult problem. Grasp quality metrics which analyze the contact wrench space of a grasp are commonly used to synthesize and analyze grasps. For approaches which preplan a large database of grasps offline, quantifying the robustness is the critical component for ranking these grasps in order to select the best candidate in a particular environment.

Automating the construction of robust grasps has been studied by many researchers. A key criterion that has been used to define a stable grasp is that the contact points of the grasp are able to generate a wrench in an arbitrary direction, termed *force closure*. This property is considered fundamental in grasp analysis because it provides a simple, necessary condition for grasp stability[1]. This property is determined from the set of wrenches that can be feasibly produced from the contact points of the grasp, termed the *grasp wrench space (GWS)*.

Force closure provides a binary analysis of grasp robustness. In order to analyze the strength of the force closure, we commonly use the *epsilon quality* (ϵ_{GWS}). ϵ_{GWS} is defined as the radius of the largest ball around the origin that fits in the convex hull of the grasp's contact wrench space. This radius is the magnitude of the minimum norm wrench that will break the object free of the grasp[2].

In previous work, measures have been proposed to quantify the robustness of contact location based quality metrics,

This work has been funded by NSF Grant IIS-0904514 and DARPA ARM-S program

J. Weisz and P. Allen are with the Department of Computer Science, Columbia University, NY 10027, USA. E-mail: {jweisz, allen}@cs.columbia.edu

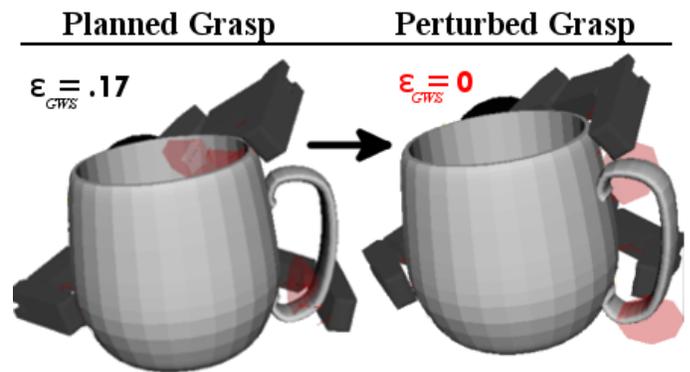


Fig. 1. An illustration of the effect of pose uncertainty on the force closure of a grasp. On the left is a planned grasp from a database of preplanned grasps with an ϵ_{GWS} of 0.17. On the right is the same grasp after a 20° clockwise rotation and 1cm translation. The contacts of the perturbed grasp have different normals and locations from the planned grasp. In this case, this perturbation results in a non-force closed grasp.

with particular focus on the ϵ_{GWS} . One well developed approach considers the effect of contact model or contact location uncertainty with respect to the object. The independent contact region approach (ICR) introduced in [3] considers the effect of contact location uncertainty on the ϵ_{GWS} of a grasp. This approach grows a convex set of acceptable contact locations for each planned finger contact on the object. Any set of contact locations chosen from this set produces a grasp whose quality metric is lower bounded by a preselected quality threshold. Faverjon et. al. developed analytical approaches to finding ICRs for complex 2D and simple 3D models[4][5][6]. In Roa and Suarez [7], this approach is extended to complex three dimensional shapes and minimal thresholds on the ϵ_{GWS} . Pollard et al.[8] followed a similar approach with a focus on larger numbers of original contacts and with extensions to task wrench space quality metrics, and also presented a method for generalizing sets of ICRs between objects. Zheng and Qian proposed a quality metric based on a similar region growing approach by measuring the radius of the smallest ball in the configuration space of the regions produced for the grasp parametrized by contact location, contact normal, or frictional coefficients[9].

All of these approaches consider the effects of contact location uncertainty independently, but do not directly address the effect of calibration error and object pose uncertainty on the ϵ_{GWS} . These approaches assume that uncertainty causes the contacts to land near the planned contact points on the object. However, simulating the entire grasping procedure often produces very different sets of contacts in the presence of uncertainty. In Fig. 1, we illustrate this problem and show

that it is not sufficient to consider the contact locations on the object in isolation of the full hand geometry and trajectory planning. In fact, none of the ICR approaches described are able to fully account for the effects of these resulting deviations from the planned contact locations.

Some work has addressed the effect of object pose uncertainty in grasping tasks for measures of grasp stability other than the ϵ_{GWS} . Analytical solutions to the robustness of spring-like finger models to pose error have been explored in 2D models[10] or simple 3D models with simple contact formulations[11]. In Balasubramanian et al. the authors derive a novel grasp measure for robust grasping from observations of human guided grasping [12]. Bekiroglu et al. proposed a method to learn haptic features of grasp stability to assess grasps during their execution[13].

The general problem of motion and manipulation planning robust to pose error was addressed by by Lozano-Perez et. al [14] using preimage-backchaining. In Berenson et. al. [15], this method was extended to hand target pose regions for grasping objects on a table top. In [16], Brost and Christiansen investigated the issue of pose error robustness for a 2D gripper using a sampling based approach on a fixed trajectory plan in physical experiments. More recently, Hsiao et al. investigated Bayesian approaches to grasp planning under object pose and identity uncertainty for a simple gripper using a sampling approach to select robust grasps as measured by a grasp quality metric in simulation [17].

In this paper we follow a sampling approach similar to Hsiao et al. extended to a more commonly used grasp quality measure on a more complex and dexterous hand using an uncertainty model similar to Berenson et. al. We apply these models to a simulated grasping pipeline to estimate the pose error robustness of the ϵ_{GWS} of a grasp.

There are two primary goals of this research. First, we show that for a reasonable error model, we can produce an extension of the ϵ_{GWS} by uniformly sampling the model and observing the behavior of contact wrench space metrics over this set of sampled poses. Second, we use these observations to re-prioritize grasps selected from the Columbia Grasp Database [18] and thus select more robust grasps. We validate this method by comparing its performance to the ϵ_{GWS} in physical experiments.

In the rest of section I, we discuss the ϵ_{GWS} metric and the grasp dataset used in this paper. In section II, we describe the implementation details necessary to evaluate grasps from a database in the presence of uncertainty. In section III, we demonstrate how uncertainty effects the ϵ_{GWS} under our noise model and grasping pipeline, and how we use these measurements as a quality metric. In section IV, we present the results of physical experiments using this quality metric.

A. THE ϵ_{GWS} -QUALITY METRIC

The ϵ_{GWS} metric of grasping is one of the most widely cited benchmark metrics in the grasp planning field[2]. Formally, for a set of contact wrenches $C \subset R^6$, one can define the neighborhood ball $B(\epsilon)$ and wrench space metric

ϵ_{GWS} as follows:

$$B(\epsilon) = \{x \in R^6 \mid \|x\|_2 < \epsilon\}$$

$$\epsilon_{GWS}(C) = \max_{\epsilon} [B(\epsilon) \subseteq \text{convexhull}(C)]$$

This quality metric is popular largely because it analytically addresses an intuitively necessary condition for grasp stability, the ability to resist force perturbations.

B. GRASP DATABASE

The grasps discussed in this paper were produced using the Eigengrasp planner presented by Ciocarlie et al.[19]. Briefly, this planner uses a stochastic optimization approach on the hand configuration space. The objective function is an approximation of the ϵ_{GWS} in which preselected contact points on the hand are projected along the finger normal. Those that make contact with the object are added to the contact wrench space weighted in proportion to their distance from the object and the angle of the projected contact relative to the object's surface normal at that point. In the limit, the contributions from points in or near contact with the object dominate and the approximation nearly agrees with the ϵ_{GWS} .

The grasps we have analyzed in this work are drawn from the Columbia Grasp Database. This is a large database of grasps compiled by Goldfeder et al.[18][20] in order to apply these off-line planning methodologies to the general problem of automated grasping. The CGDB was constructed by using the Eigengrasp planner on the Princeton Shape Benchmark [21], a set of 1814 models with class labels.

II. ANALYZING THE EFFECTS OF POSE ERROR IN SIMULATION

A. GRASPING PIPELINE

The CGDB describes grasps as pre-grasp and final-grasp pairs. The pre-grasp is a collision free pose and set of joint angles found by the simulated annealing planner. The final-grasp is determined by simulating a grasping heuristic for closing the hand around the object until contact is made or joint limits are reached. The pre-grasp is assumed to be a collision free pose, but in the presence of uncertainty this may not be the case for all grasps. In the case of uncertainty, it is necessary to generate a collision free initial pose.

In this work we consider a three-phase grasping heuristic.

- 1a. Move hand and fingers to an initial pose and initial posture far from the object along a pre-specified approach direction.
- 1b. Move to the planned pre-grasp hand pose from the initial pose along the approach direction, stopping if contact is made.
- 2a. Move fingers to their pre-grasp postures by linearly interpolating from initial posture, stopping each finger independently if it makes contact.
- 2b. (non-fingertip grasps only) For grasps with any contact points that are not on a fingertip of the hand, if no contact has occurred thus far, continue to approach until contact is made.

3. Close fingers until each joint is constrained by contact or joint limits.

We use the normal to the center of the palm as the approach directions for all of the grasps considered. To determine the finger positions for the initial posture, we open the fingers so that the hand in the initial pose can wrap around itself in the pre-grasp position. We illustrate this strategy in Fig. 2. We open the hand at a constant rate for all joints until the tips of each finger are outside of the projection of the second most distal link along the approach direction of the hand in the planned pre-grasp pose.

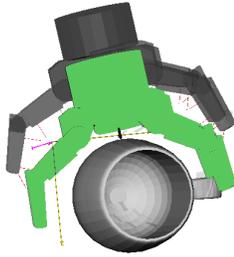


Fig. 2. Illustration of the initial pose and posture. The hand is withdrawn along its approach direction and all of the fingers are opened at the same constant velocity until each of the fingertips are outside of the knuckles of the hand in its pre-grasp configuration. The hand is then opened an additional ten percent of the remaining joint range. The pre-grasp configuration is shown in green.

B. MODEL OF POSE UNCERTAINTY

Because it is computationally intractable to densely sample the full six dimensional space of poses near the planned pregrasp pose, we consider a three dimensional error model representing an object on a support surface. We assume that each object is restricted to a set of stable poses on the surface.

This error model is parametrized by $[x, y, \theta]$ as illustrated by Fig. 3. We use the centroid of the planned contact locations as the origin of the object because parameterizing the rotation around the center of mass of the object will have a disproportionately large effect on grasps planned further from the center of mass. A reasonable range to explore in this parameter space should allow each parameter to move the relative position of the planned grasp's contact points by at least 10 mm. If we make the approximation that the projection of the contact points perpendicular to the axis of rotation lie on a circle with a 5 cm radius, a reasonable range for these parameters is $\theta \in [-20^\circ, 20^\circ]$ and $x, y \in [-10\text{mm}, 10\text{mm}]$. These bounds are motivated by our anecdotal experience in aligning a known object to a point cloud using common methodologies. Exploring this parameter space uniformly in increments of $1[\text{mm},^\circ]$ in simulation requires 18,081 simulations per grasp. The time required for each simulation is object and grasp dependent. Using the GraspIt! grasping simulator[22] on a Xeon 2.67GHZ processor we can perform this many kinematic simulations of the grasping pipeline in under two hours for the worst case. With cluster computing this approach can analyze large numbers of grasps from a database off-line.

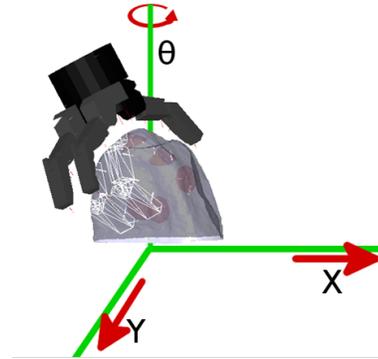


Fig. 3. Illustration of the tabletop error model used in this paper. The object is can be translated $[-10, 10]$ mm in both X and Y and rotated $[-20, 20]^\circ$ around θ . The red cones mark the planned contact points. The white wire frame indicates the planned end configuration of the fingers.

III. SIMULATION RESULTS

To produce a measure which analyzes the robustness of a grasp under uncertainty, we estimate the probability of attaining a force closure grasp (under our error model) for a large set of grasps and objects, denoted as $P(fc)$. We simulated the grasp pipeline and error model described in section II on a set of 480 grasps for the Barrett Hand using a friction coefficient of 2.0. This set of grasps is composed of the top five grasps by ϵ_{GWS} on 96 objects of the 'Tool' Princeton Shape Benchmark category in the CGDB. This set includes shovels, hammers, wrenches, screwdrivers, and other tools.

We define the probability of force closure $P(fc)$ as the probability of achieving a positive ϵ_{GWS} , which is calculated as $P(\epsilon_{GWS} > \delta)$. In order to be somewhat robust to small modeling errors, we must establish a lower threshold for the desired ϵ_{GWS} that is greater than zero. In this work, we choose a δ of 0.001.

The best grasp by $P(fc)$ has a $P(fc)$ score .19 higher than the $P(fc)$ of the best grasp by ϵ_{GWS} on average. This difference means under the error model described, choosing the highest $P(fc)$ grasp will result in a force closed grasp 19% percent more often than choosing the best grasp by ϵ_{GWS} . Using this criterion, we see that the ϵ_{GWS} is not sufficient to predict pose error robustness. In fact, for 72 of the 96 objects, ranking the grasps by ϵ_{GWS} does not produce the same grasp choice as ranking the grasps by $P(fc)$.

In Fig. 4, we have plotted the planned ϵ_{GWS} against the $P(fc)$ for each grasp. This figure clearly shows that over most of the range of ϵ_{GWS} , there is no correlation between ϵ_{GWS} and $P(fc)$. Although there appears to be some correlation in the region of the graph where $\epsilon_{GWS} > 0.2$, only five of the 96 objects have a grasp with this high an ϵ_{GWS} .

Although this method of estimating grasp robustness is not practical for analyzing grasps in real time, it does enable us to analyze a large database of grasps such as the CGDB off-line. The rationale for producing such a measure is further elucidated in the Discussion section.

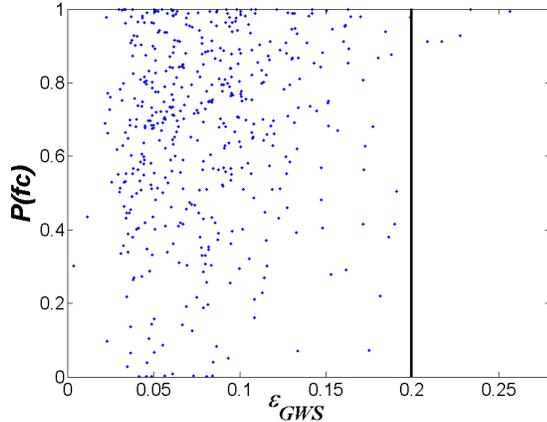


Fig. 4. ϵ_{GWS} plotted against $P(fc)$, which is equivalent to the probability of attaining a force closed grasp under our error model from a planned starting grasp configuration. This figure demonstrates that the ϵ_{GWS} is generally a poor predictor of pose error robustness as measured by the $P(fc)$. 480 grasps over 96 objects are presented in this figure. Note that while grasps with an $\epsilon_{GWS} > 0.2$ demonstrate a high probability of being robust, only five objects have a grasp in the database with an ϵ_{GWS} that high.

IV. PHYSICAL EXPERIMENTS

These simulation results show that using $P(fc)$ as a quality metric for ranking grasps from a database may select grasps which are more reliable in the face of object pose estimation error. To test whether these simulation results can be verified against physical experiments, we attempted to grasp ten objects for which we had high quality mesh models readily available courtesy of the the DARPA ARM-S project and the Willow Garage object database [23]. We added these models to our database, and calculated the $P(fc)$ for the top five grasps by ϵ_{GWS} value. To test these grasps on a real robotic grasping system we used a Barrett 280 model hand attached to a Staubli TX60L 6 degree of freedom arm to implement the grasping pipeline described in section II-A. To emulate the friction coefficient used in our simulations, we wrapped the contacting surfaces of the hand in rubberized shelf liner. One object attempted was a blue detergent bottle made of very smooth plastic. We added a rubberized tape to the contact locations on this object to increase its friction coefficient. For each object, we selected the best grasp ranked by ϵ_{GWS} and the best grasp ranked by $P(fc)$.

The objects were calibrated to the robot using a printed template of the bottom outline of the objects. We align the edge of each template to a known coordinate system in our robot’s workspace. We selected a perturbed position by sampling our error model and then applying the grasp pipeline as though the object’s origin was at the perturbed location. We tested up to ten randomly selected perturbed positions for each grasp. The same perturbations were applied to both of the tested grasps for each object.

We moved the arm and hand at very slow speeds to simulate quasi-static conditions. In the first two stages of the pipeline, we moved the arm at 1% of its maximum speed, and the fingers of the Barrett hand at 10% of their maximum speed. In the final closing step, initial contact had been

achieved or joint limits have been reached for each finger. We then raised the finger speed to 50% to drive the underactuated fingers to their final positions¹. After the fingers stopped moving, we lifted the object 3 cm perpendicular to the table’s surface. If any part of the object was still touching the table when the arm came to a stop, we graded the trial as a failure, otherwise we graded the trial as a success.

We find that the grasps ranked best by the $P(fc)$ are successful more often. The results of this experiment along with the grasps and objects analyzed are in Fig. 5 and Fig. 6. Overall, for these ten objects the $P(fc)$ ranking selected successful grasps on 85 out of 93 trials (91%), whereas the grasps selected by the ϵ_{GWS} ranking succeeded on 63 of 93 trials (67%).

For some objects, nearly all of the grasp attempts were successful irrespective of the $P(fc)$ score. We separated the objects in our experiment into two groups. One group is composed of of larger, more complex objects shown in Fig. 5 and the other group is composed of light, roughly cylindrical objects shown in shown in Fig. 6. We see that all of the objects in the roughly cylindrical category were grasped very robustly irrespective of $P(fc)$ score. For these objects, the Eigengrasp planner finds grasps which more or less encompass the object. In contrast, for the five objects which could not be enveloped by the hand, we see that the $P(fc)$ ranked grasp is successful in 35 out of 43 trials (81%), as compared to 16 out of 43 trials (37%) using the ϵ_{GWS} grasp. These results show that the $P(fc)$ ranking chooses more successful grasps; especially on the larger, more complex objects where the success rate is doubled.

V. DISCUSSION AND FUTURE WORK

In this work, we have shown that the planned ϵ_{GWS} of a planned grasp is not predictive of the probability of achieving a force closed grasp in the presence of uncertainty, neither in simulation nor physical experiments. To analyze this in simulation we generated the measure $P(fc)$ as an approximation of this probability. We then showed that the $P(fc)$ can itself be used to rank grasps from a preplanned grasp database and that this reranking predicts grasping success in physical experiments better than the ϵ_{GWS} .

Despite the fact that the analysis discussed in this paper is not applicable to on-line grasp analysis, it is still relevant for several reasons. For known objects this allows us to screen bad grasps out of our database in a way that appears to predict success in physical experiments. Additionally, one of the challenges in attempting to generalize grasps from a database to a novel object is finding an appropriate alignment to the novel object. If the grasps in the database are more robust to pose error, there is more latitude in this alignment.

There are many potential extensions of the $P(fc)$ to improve efficiency and accuracy. One potential extension

¹The Barrett Hand uses an underactuation mechanism in which the proximal and distal joints are coupled when the finger is moving freely. A large torque differential between the proximal and distal joint will decouple the two joints, allowing the distal joint to continue closing if the proximal joint is impeded.

	Best Grasp ranked by ϵ_{gws}			Best Grasp ranked by $P(fc)$		
	Simulated Grasp	Physical Grasp	Grasp Metrics	Simulated Grasp	Physical Grasp	Grasp Metrics
Flashlight			$\epsilon_{gws} = 0.18$ $P(fc) = 75\%$ Lift Test Success: 0/10			$\epsilon_{gws} = 0.15$ $P(fc) = 99\%$ Lift Test Success: 10/10
Drill			$\epsilon_{gws} = 0.13$ $P(fc) = 62\%$ Lift Test Success: 6/10			$\epsilon_{gws} = 0.12$ $P(fc) = 95\%$ Lift Test Success: 8/10
Detergent Bottle			$\epsilon_{gws} = 0.33$ $P(fc) = 79\%$ Lift Test Success: 2/10			$\epsilon_{gws} = 0.31$ $P(fc) = 97\%$ Lift Test Success: 5/10
Drink Carton			$\epsilon_{gws} = 0.38$ $P(fc) = 86\%$ Lift Test Success: 8/10			$\epsilon_{gws} = 0.37$ $P(fc) = 100\%$ Lift Test Success: 10/10
False Rock			$\epsilon_{gws} = 0.35$ $P(fc) = 75\%$ Lift Test Success: 0/3			$\epsilon_{gws} = 0.319$ $P(fc) = 98\%$ Lift Test Success: 2/3

Fig. 5. The results of physical experiments on five larger, more complex objects. In the left column are the best grasps for each object ranked by ϵ_{GWS} . An example of the simulated and physically realized grasp is illustrated for each grasp attempted. Additionally, the figure gives the $P(fc)$ and ϵ_{GWS} for each grasp attempted, as well as the fraction of attempts which succeeded in lifting the object by 3 cm in physical experiments. The higher $P(fc)$ grasp is successful on many more trials than the higher ϵ_{GWS} grasp for these objects.

would be to find the pose error space analog to the ϵ_{GWS} . This measure would find the magnitude of the smallest pose perturbation for which the grasp is not force closed. The algorithm for calculating this measure can terminate on the first non-force closed pose it finds, and is thus less computationally expensive.

VI. ACKNOWLEDGMENTS

The authors gratefully acknowledge the help of Corey Goldfeder and Hao Dang in pursuing this work.

REFERENCES

- [1] J. K. Salisbury, Jr., *Kinematic and force analysis of articulated hands*. New York, NY, USA: John Wiley & Sons, Inc., 1985, pp. 131–174.
- [2] C. Ferrari and J. Canny, “Planning optimal grasps,” in *IEEE Int. Conf. on Robotics and Automation*, August 1992, pp. 2290–2295.
- [3] V.-D. Nguyen, “Constructing force-closure grasps,” *IJRR*, vol. 7, pp. 3–16, June 1988.
- [4] J. Ponce and B. Faverjon, “On computing three-finger force-closure grasps of polygonal objects,” *Robotics and Automation, IEEE Transactions on*, vol. 11, no. 6, pp. 868–881, 1995.
- [5] J. Ponce, D. Stam, and B. Faverjon, “On computing two-finger force-closure grasps of curved 2d objects,” *IJRR*, vol. 12, no. 3, pp. 263–273, 1993.
- [6] J. Ponce, S. Sullivan, A. Sudsang, J.-D. Boissonnat, and J.-P. Merlet, “On Computing Four-Finger Equilibrium and Force-Closure Grasps of Polyhedral Objects,” *IJRR*, vol. 16, no. 1, pp. 11–35, Feb. 1997.
- [7] M. Roa and R. Suarez, “Computation of Independent Contact Regions for Grasping 3-D Objects,” *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 839–850, 2009.
- [8] N. S. Pollard, “Closure and quality equivalence for efficient synthesis of grasps from examples,” *IJRR*, vol. 23, no. 6, pp. 595–613, 2004.



Fig. 6. The results of physical experiments on five smaller, relatively simple objects. In the left column are the best grasps for each object ranked by ϵ_{GWS} . In the right column are the best grasps by $P(fc)$. An example of the simulated and physically realized grasp is illustrated for each grasp attempted. Additionally, the figure gives the $P(fc)$ and ϵ_{GWS} for each grasp attempted, as well as the fraction of attempts which succeeded in lifting the object by 3 cm in physical experiments.

[9] Y. Zheng and W.-H. Qian, "Coping with the Grasping Uncertainties in Force-closure Analysis," *IJRR*, vol. 24, no. 4, pp. 311–327, 2005.

[10] T. G. Sugar and V. Kumar, "Metrics for analysis and optimization of grasps and fixtures," *IJRA*, vol. 17, no. 1, pp. 28–37, 2002.

[11] T. Yamada, T. Koishikura, Y. Mizuno, N. Mimura, and Y. Funahashi, "Stability analysis of 3d grasps by a multifingered hand," in *IEEE Int. Conf. on Robotics and Automation*, vol. 3, 2001, pp. 2466 – 2473.

[12] R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka, "Human-guided grasp measures improve grasp robustness on physical robot," in *IEEE Int. Conf. on Robotics and Automation*, 2010, pp. 2294–2301.

[13] Y. Bekiroglu, J. Laaksonen, J. 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, June 2011.

[14] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *IJRR*, vol. 3, no. 1, pp. 3–24, 1984.

[15] D. Berenson, S. S. Srinivasa, and J. J. Kuffner, "Addressing pose uncertainty in manipulation planning using Task Space Regions," *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pp. 1419–1425, 2009.

[16] R. Brost and A. Christiansen, "Empirical verification of fine-motion planning theories," in *Experimental Robotics IV*, 1997, vol. 223, pp. 475–485.

[17] K. Hsiao, M. Ciocarlie, and P. Brook, "Bayesian grasp planning," in *ICRA 2011 Workshop on Mobile Manipulation*, May 2011.

[18] C. Goldfeder, M. Ciocarlie, H. Dang, and P. Allen, "The columbia grasp database," in *IEEE Int. Conf. on Robotics and Automation*, May, 2009, pp. 1710–1716.

[19] M. Ciocarlie and P. Allen, "Hand posture subspaces for dexterous robotic grasping," in *IJRR*, vol. 28, Jul. 2009, pp. 851–867.

[20] C. Goldfeder and P. Allen, "Data-driven grasping," *Autonomous Robots*, vol. 31, pp. 1–20, 2011.

[21] "The princeton shape benchmark," June 2004.

[22] A. Miller and P. Allen, "Graspit! a versatile simulator for robotic grasping," *Robotics Automation Magazine, IEEE*, vol. 11, no. 4, pp. 110 – 122, dec. 2004.

[23] M. Ciocarlie, G. Bradski, K. Hsiao, and P. Brook, "A dataset for grasping and manipulation using ros," in *IROS Workshop: RoboEarth - Towards a World Wide Web for Robots*, Taipei, Taiwan, Oct. 2010.