

Synthesizing Grasps from Generalized Prototypes

Nancy S. Pollard
Artificial Intelligence Laboratory*
Massachusetts Institute of Technology

Abstract

This paper introduces a grasp synthesis algorithm that can use any grasp prototype as the starting point in a search for a good grasp. The algorithm makes the given prototype more effective by generalizing it for a specific task. This generalization step expands the range of application of the prototype to a wide variety of target object geometries, while ensuring that the resulting grasps are appropriate for the intended task. An example whole-hand grasp synthesized for the Salisbury hand is shown at the end of the paper.

1 Introduction

It has been very difficult to successfully automate the processes of acquiring and manipulating objects using today's robotic hands. Because many successful grasps are whole-hand, or enveloping grasps, one tempting solution is to develop prototypes, such as a generic cylinder grasp, to be applied to common objects, an approach motivated by studies of human grasping (e.g. see Napier [17], and Jeannerod [10]). Prototype grasps benefit from requiring very little computation, but they do not seem sufficiently robust to work in complex situations. In addition, it is difficult to determine, without trying a grasp, whether it will work at all.

At the opposite extreme, a thorough analysis of each grasping or manipulation task would allow more robust grasps to be designed. This approach has the benefit of accuracy, but a thorough analysis of a task requiring a many-contact grasp of a complex three-dimensional object is a problem too complex to be solved in a realistic amount of computation time.

This paper presents a novel, hybrid solution to the problem of grasp synthesis. It shows how a grasp prototype can be analyzed offline to determine how it can

be adapted to a wide variety of target object geometries, while maintaining the effectiveness of the grasp for the intended task. This is achieved by focusing on the wrench spaces of the grasp and the task rather than on the appearance of the grasp. This technique has been successfully applied to synthesize whole-hand grasps of complex three-dimensional objects.

2 Previous Work

Work in the area of prototype grasps includes Bard and Troccaz [2] who explore a grasping strategy for ellipsoidal objects, Cutkosky and Howe [6], who develop an expert system for selecting grasps for manufacturing applications, and Brock [3], who develop tools for sensor-based grasping strategies.

In the area of grasp analysis, work on the analysis of stable grasps can be found in Mishra et al. [16], Cutkosky [5], Kerr and Roth [11], Salisbury [21], and Salisbury and Craig [22].

Representative of work on grasp optimization functions, or grasp quality measures, are Li, Hsu, and Sastry [13], where the optimality criterion is the ability to manipulate a grasped object; Li and Sastry [14], Kirkpatrick and Yap [12], and Ferrari and Canny [8], where the ability to apply task forces and torques to a grasped object is measured; and Ponce et al. [20], Faverjon and Ponce [7], and Nguyen [18], where the optimality criterion is robustness of a grasp to errors in contact placement.

Computationally efficient grasp optimization techniques for two-dimensional objects can be found in Markenscoff and Papadimitriou [15], Chen and Burdick [4], Baker et al. [1], Hanafusa and Asada [9]. The efficiency of these techniques is a result of constraints on the dimensionality of the object, the task to be performed, and the number of contacts of the grasp.

The algorithm outlined in this paper pulls ideas from all of these areas of work, demonstrating a grasp synthesis technique that involves offline analysis of a grasp prototype. This offline analysis both generalizes

*Current mailing address: Monitor Company, 25 First Street, Cambridge, MA 02141

that prototype and makes it possible to rapidly identify high quality grasps of new target objects based on that prototype. Identification of grasps of new target objects can be done at runtime.

3 Problem Statement

The problem addressed in this paper is to find the grasp of a given target object that is most suitable for achieving a given task. A grasp is defined as a set of contact points (e.g. see Figure 1). A grasp prototype is used as a starting point, to reduce the computational requirements of the grasp synthesis problem.

Assumptions: The assumptions required for the algorithm described in this paper are as follows.

- A geometric model of a target object, a grasp prototype, and a task are provided as inputs.
- Only non-singular, frictionless point contacts are used. For two-dimensional objects, contact is restricted to point contacts on edges of the target object and line contacts at vertices of the target object. This assumption allows the quality of a grasp to be uniquely determined from the location of the contact points of that grasp, and thus reduces the grasp synthesis problem to one of finding a good set of contact points on a target object.

Notation:

$$\begin{aligned}
 \underline{w}_i &= \text{wrench at contact } i = \begin{bmatrix} \underline{f}_i \\ \underline{\tau}_i \end{bmatrix} \\
 \underline{f}_i &= \text{applied force at contact } i \\
 \hat{n}_i &= \text{unit outward normal at contact } i \\
 \underline{\tau}_i &= \text{torque at contact } i, = \lambda (\underline{f}_i \times \underline{d}_i) \\
 \underline{d}_i &= \text{vector to object center of mass} \\
 \lambda &= \text{torque to force conversion factor} \\
 c &= \text{number of contacts}
 \end{aligned}$$

4 A Grasp Prototype

A grasp prototype is defined as an example object and a high-quality grasp of that object. Figure 1 illustrates the grasp prototype that is used as an example in this paper. The five-contact grasp shown in the figure is represented only by its contact points. Because frictionless point contacts are assumed, unit contact wrenches can be obtained directly from the




Figure 1: An example grasp prototype, represented by forces applied at the numbered contacts.

target object geometry and the coordinates of the contact points with respect to the grasp origin, which is the object center of mass.

5 A Grasp Quality Measure

A grasp quality measure is an estimate of the suitability of a grasp for the task to be performed. All grasps that are synthesized will be evaluated according to this measure. In order to precisely define grasp quality, good descriptions of grasps and tasks, and a mechanism for comparing the two are needed. This section draws heavily from Kirkpatrick and Yap [12], Ferrari and Canny [8], and Li and Sastry [14].

The grasp wrench space region (GS): A grasp can be characterized based on the set of wrenches that can be applied to the target object from the contacts of a grasp, given certain limitations on applied forces. In this paper, the scaled grasp wrench space region κGS is defined by limiting the sum of magnitudes of contact forces to κ :

$$\kappa GS = \{ \underline{w} \mid \underline{w} = \sum_{i=1}^c \alpha_i \underline{w}_i, \alpha_i \geq 0, \sum_{i=1}^c \alpha_i \leq \kappa, |\underline{f}_i| = 1 \}. \quad (1)$$

In other words, GS is bounded by the convex hull of the contact wrenches formed from unit applied forces at the contacts of the grasp.

The task wrench space region (TS): A task can be characterized similarly as the space of wrenches that must be applied to the object by the robot in order to complete the task objective. One useful task wrench space region, TS_{ball} , is used when the task does not favor any particular subset of wrench space:

$$TS_{ball} = \{ \underline{w} \mid |\underline{w}| \leq 1 \} \quad (2)$$

The grasp quality measure (Q): Intuitively, the grasp quality measure is the amount the robot has to squeeze the object in order to be capable of resisting all task wrenches while maintaining the grasp. The grasp quality measure used in this paper can be defined as follows:

$$\kappa_m = \min(\kappa) \mid TS \in \kappa GS \quad (3)$$

$$Q = \left(\frac{1}{\kappa_m} \right) \quad (4)$$

In other words, grasp quality is the reciprocal of the amount by which GS must be scaled so that it just contains TS .

The grasp quality measure for the grasp shown in Figure 1 is 0.35 for task TS_{ball} . This means that if a task wrench must be countered in a direction in which the grasp is weak, the wrenches that must be applied to the object through the contacts of the grasp will be approximately three times the magnitude of that task wrench.

6 Generalizing a Grasp Prototype (Offline)

If a grasp prototype is to be applied to a range of target objects that are geometrically different from the target object, it must first be made more general. This section describes how a grasp prototype can be generalized to represent a *set* of grasps that are appropriate for a specified task. This generalization step can be performed offline and the resulting generalized grasp prototype stored in a grasp library for access at runtime.

The generalized grasp prototype $GGP(TS, Q)$ is defined based on the grasp wrench space region of the prototype GS_{proto} as follows:

- \underline{w}_i = contact wrench i of the grasp prototype,
- g_j = facet j of convex hull bounding GS_{proto} ,
- \hat{n}_j = unit outward pointing normal of facet g_j .
- γ_j = the size of TS in the \hat{n}_j direction

Let G_i index the set of facets containing \underline{w}_i :

$$G_i = \{j : g_j \text{ contains } \underline{w}_i\}.$$

A grasp falls within $GGP(TS, Q)$ iff each wrench of the new grasp \underline{w}'_i is placed so that it meets the following conditions:

$$\forall j \in G_i, \left[(\underline{w}'_i \cdot \hat{n}_j) \geq \gamma_j Q \right]. \quad (5)$$




Figure 2: A 2D wrench space example: (L) Contact wrenches w_i of a grasp prototype with wrench space origin O and convex hull bounding GS for this grasp prototype. (R) Wrench space regions w'_i that comprise $GGP(TS_{ball}, Q)$. A grasp having a grasp quality measure greater than or equal to Q can be formed by selecting one contact wrench from each of the four regions. As required quality Q shrinks, these regions grow.

Regions \underline{w}'_i for a hypothetical two-dimensional wrench space and task TS_{ball} are shown in Figure 2. This figure can also be used to illustrate the following:

1. All grasps within $GGP(TS, Q)$ have grasp quality measure greater than or equal to Q for task TS .
2. If q is less than or equal to the grasp quality measure of the grasp prototype for task TS , then $GGP(TS, q)$ contains the grasp prototype.
3. $GGP(TS, Q)$ is described in terms of independent constraint sets for each contact c of a grasp. Contacts of a grasp can be chosen in parallel, while guaranteeing a quality measure of at least Q for the grasp as a whole. This allows grasp synthesis algorithms to be developed that have complexity linear in the number of contacts.

7 Contact Constraint Sets

The contact constraint set for a given grasp prototype, task TS , grasp quality threshold Q , and contact i is the set of contact wrenches \underline{w}'_i that satisfies Equation 5. The contact constraint sets for the hypothetical wrench space of Figure 2 are the shaded regions in the figure.

The contact constraint sets for a grasp of a two-dimensional object lie within a three-dimensional wrench space, but because all force vectors are normalized (see Equation 1), they fall on a unit radius




Figure 3: Contact constraint sets for the five contact grasp prototype, task TS_{ball} , and $Q \geq 0.75Q_{proto}$, superimposed on the object boundary curve of the example polygon. The object boundary curve indicates what contact wrenches are possible for this object. Actual prototype contact wrenches are the numbered points.

cylinder in this wrench space. These contact constraint sets can be plotted on a two-dimensional graph by unrolling this cylinder.

Figure 3 shows the contact constraint sets for the grasp prototype of Figure 1, task TS_{ball} , and $Q \geq 0.75Q_{proto}$. The x-axis represents orientation of the force vector in radians, and the y-axis represents the scaled torque vector. The shaded regions illustrate wrenches having contact quality measures of at least 75% of the grasp quality measure of the grasp prototype. The numbered points within these regions represent the contact wrenches of the grasp prototype. Each contact wrench falls within its corresponding contact constraint set.


8 Contact Quality Measures

Each contact wrench \underline{w}'_i within a contact constraint set i of Figure 3 will satisfy Equation 5 for a range of grasp quality thresholds Q . The *contact quality measure* for a contact wrench \underline{w}'_i is defined as the maximum grasp quality threshold for which Equation 5 is satisfied.

This is a very powerful measure. When contacts are selected for a new grasp, the quality of each contact wrench can be independently optimized. A lower bound on grasp quality can then be obtained by minimizing the contact quality measures over the contacts i of a grasp:

$$Q_{min} = \min_i \left(\min_{j \in G_i} \left[\frac{(\underline{w}'_i \cdot \hat{n}_j)}{\gamma_j} \right] \right). \quad (6)$$

An approach of independently optimizing each contact drastically reduces problem combinatorics and lends



Quality = 0.35

Figure 4: Independent contact regions that guarantee a grasp quality measure of 75% of the grasp quality measure of the original grasp. Only point-edge contacts are shown. The region corresponding to contact three overflows onto the edge containing contact one.

itself to the implementation of parallel algorithms for grasp synthesis. The ability to rigorously define a lower bound grasp quality measure ensures that grasps formed in this manner will be suitable for the task to be performed.

9 The Object Boundary Curve

A contact wrench selected arbitrarily from one of the contact constraint sets of Figure 3 may not be achievable through frictionless point contact with a given target object. Superimposing the *object boundary curve* of the target object upon the contact constraint sets of Figure 3 allows us to map contact wrenches to their corresponding contact points on the target object.

The object boundary curve is the set of normalized wrenches that represent frictionless point contact with a given object. Figure 3 shows the object boundary curve for the prototype object of Figure 1. The vertical lines in the plot represent contact with edges of the target object. Curved lines represent contact with vertices of the target object.

The combined plot of object boundary curve and contact constraint sets tells us a great deal about the robustness of a grasp to errors in contact placement. Consider Figure 4, for example. From the plot in Figure 3 we see that as long as point contacts are placed within the highlighted regions of Figure 4, the resulting grasp will have a grasp quality measure of at least 75% of the example grasp.

Note that contact regions may involve not only a

point of the robot in contact with an edge of the object, but also an edge of the robot in contact with a vertex of the object. For example, Figure 3 shows that the contact constraint set of contact 4 can be satisfied by one of two discontinuous contact regions: the region on the edge of the target object shown in Figure 4, or a separate region of edge contact on the nearby target object vertex.

The effects arising from a variety of changes in the grasp can be read from the plot in Figure 3. For example, uncertainty in contact placement displaces contact points on the object boundary curve. In addition, uncertainty in object orientation causes the curve to slide on the orientation axis, and changes in object geometry cause the object boundary curve to warp. We can get an idea of how robust a grasp might be by noting how easily such changes can cause the contact points to slip out of the good contact regions. This figure provides a way of visualizing the problems of designing an appropriate grasp for a given object or designing an object that can be easily grasped.

10 Synthesizing Optimal Grasps (Runtime)

Because a high quality grasp can be created by selecting contact points within *regions* defined by the contact constraint sets, a given $GGP(TS, Q)$ is sufficiently flexible to be applied to objects of varying target geometry. A brute force process for finding an optimal grasp of any target object from a given grasp prototype and task can be summarized as follows:

1. Given a grasp prototype (e.g. Figure 1), extract the contact wrenches \underline{w}_i of that grasp prototype. This grasp prototype may have been designed with careful thought given to representing an ideal grasp for the intended task. (*Offline*)
2. Given the task TS , compute the generalized grasp prototype $GGP(TS, Q)$ for that task as a function of grasp quality threshold Q (Equation 5). (*Offline*)
3. Given a new target object, find an optimal set of contacts on that target object as follows: (*Runtime*)
 - (a) Sample all orientations of the target object. (Specifically, the orientation of the target object with respect to the coordinate system of the grasp prototype must be defined.)

- (b) For each orientation, compute the objective function at that orientation as follows:
 - i. Compute the maximum contact region size for each contact (e.g. the sizes of the regions in Figure 4), or the maximum contact quality measure that can be obtained for each contact (maximize the inner expression of Equation 6).
 - ii. Minimize the results over the contacts i to obtain the value of the objective function at this orientation.
- (c) Maximize the objective function over all orientations.
- (d) Place each contact at the center of the largest region found for that orientation or at the point of maximum contact quality for that orientation. This is an optimal set of contacts for the given object based on the generalized grasp prototype $GGP(TS, Q)$.

Figure 5 shows an optimal grasp found for a test object by maximizing the size of independent contact regions. The grasp quality threshold used in measuring the sizes of the valid contact regions was set at 25% of the quality of the grasp prototype, but note that the quality of the grasp (0.27) is much better than the guaranteed lower bound of 0.0875. It is very near the quality of the grasp prototype, which is 0.35. A variety of two-dimensional objects were tested with similar results (Pollard [19]). High quality grasps were found for all objects tested, although as the objects became less like the prototype object the contacts tended to move onto object vertices.

11 Synthesizing Grasps of 3D Objects

When we move from two-dimensional to three-dimensional objects, the techniques used to construct contact constraint sets and contact quality measures and to synthesize optimal grasps from these constructs do not change. The contact types and the required intermediate constructs do become more complex, however. Important changes are listed below:

- Legal contacts between polyhedral objects are frictionless point-face contacts and frictionless edge-edge contacts.
- The wrench space is six-dimensional.
- There is a two-dimensional object boundary surface rather than a one-dimensional object boundary curve.




Figure 5: The best match of this new target object to the example grasp has a grasp quality measure of 0.27, at an orientation of π with respect to the orientation illustrated in the top figure. (Flip it upside down and compare the selected contact wrenches with those of the prototype.) The contact wrenches of the grasp are indicated in the plot, and independent contact region sizes for the quality threshold of 0.0875 can be estimated by examining the plot.

- The target object orientation can vary in a three-dimensional orientation space, rather than a one-dimensional orientation space.

The increase in complexity in moving to three-dimensional objects means that the brute force algorithm described in Section 10 will tend to be slow (especially step 3a, which would now sample a three-dimensional orientation space). In practice, geometric fitting of the target object to the prototype can be used to identify interesting portions of orientation space. This is especially true for any task that is asymmetrical (e.g. the task of drinking from a glass). Brute force techniques on a parallel SIMD machine and numerical routines on a PC have been used to rapidly perform step 3b (several seconds for either approach on objects like the toy plane of Figure 6).

The grasp synthesis techniques described in this paper were tested for a seven contact cylindrical grasp prototype and a model of the Salisbury three-fingered hand [21]. Figure 6 shows one of the grasps that were synthesized based on the cylindrical prototype. All grasps tested in this way were found to be very robust.

Note that in order to form a whole hand grasp of a three-dimensional object, it is not sufficient to find




Figure 6: Two views of the Salisbury hand grasping the toy plane.

an optimal set of contact points using this type of algorithm, because the constraints of the robot hand kinematics and geometry must also be satisfied. In Pollard [19], this was done using a parallel algorithm that first found good configurations for links of the robot hand independently and then chained these links together in kinematically feasible ways to form high-quality collision-free grasps of the given target object. The ability to optimize contacts independently and the existence of regions of good contacts were crucial to the success of this algorithm.

12 Summary

This paper presented a novel technique for synthesizing grasps by generalizing a single example grasp, or grasp prototype, and applying this generalized prototype to a new target object. Some features of the grasp synthesis technique presented in this paper are reviewed below.

Global Region Information: Lower bound quality values provided over a space of possible grasps can be used to determine the robustness of a grasp to errors in contact placement, and to provide the flexibility in contact placement that is needed to adapt to

the constraints of the robot hand kinematics.

Offline Computation: A grasp can be optimized for a particular task offline, generalized, and stored in a grasp library for later access.

Wide Applicability: There is no limit to the number of contacts in a grasp. In fact, independent contact region sizes increase dramatically as the number of contacts increases. This approach can be applied to three-dimensional, curved objects and any given task wrench space.

Speed: The run-time process of testing the quality of any potential grasp is very fast. It is reduced from the problem of constructing a grasp wrench space convex hull and finding the scale of the largest task wrench space that fits within that convex hull to the problem of evaluating a few dot products for each contact of the grasp. The quality measure for each contact can be independently calculated, allowing the placement of each contact to be independently optimized, and resulting in an algorithm that is easily parallelizable.

Acknowledgments

This paper describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-91-J-4038 and in part by the University Research Initiative under Office of Naval Research contract N00014-92-J-1814.

References

- [1] B. S. Baker, S. J. Fortune, and E. H. Grosse. Stable prehension with a multi-fingered hand. In *Proc. IEEE ICRA*, St. Louis, Missouri, March 1985.
- [2] C. Bard and J. Troccaz. Automatic preshaping for a dextrous hand from a simple description of objects. In *Proc. IEEE Intl. Workshop on Intelligent Robots and Systems*, Tsuchiura, Ibaraki, Japan, July 1990.
- [3] D. L. Brock. *A Sensor Based Strategy for Automatic Robotic Grasping*. PhD thesis, MIT AI Lab., 1993.
- [4] I. Chen and J. W. Burdick. Finding antipodal point grasps on irregularly shaped objects. In *Proc. IEEE ICRA*, Nice, France, May 1992.
- [5] M. R. Cutkosky. Mechanical properties for the grasp of a robotic hand. Technical Report CMU-RI-TR-84-24, Carnegie Mellon, 1984.
- [6] M. R. Cutkosky and R. D. Howe. Human grasp choice and robot grasp analysis. In S.T. Venkataraman and T. Iberall, editors, *Dextrous Robot Hands*. Springer-Verlag, New York, 1990.
- [7] B. Faverjon and J. Ponce. On computing two-finger force-closure grasps of curved 2d objects. In *Proc. IEEE ICRA*, Sacramento, California, 1991.
- [8] C. Ferrari and J. Canny. Planning optimal grasps. In *Proc. IEEE ICRA*, Nice, France, May 1992.
- [9] H. Hanafusa and H. Asada. Stable prehension by a robot hand with elastic fingers. In *Proc. 7th Int. Symp. Industrial Robots*, Tokyo, Japan, October 1977.
- [10] M. Jeannerod. The timing of natural prehension movements. *Journal of Motor Behavior*, 16(3), 1984.
- [11] J. Kerr and B. Roth. Analysis of multifingered hands. *IJRR*, 4(4), Winter 1986.
- [12] D. G. Kirkpatrick, B. Mishra, and C. K. Yap. Quantitative steinitz's theorems with applications to multifingered grasping. In *Proc. 20th ACM Symposium on Theory of Computing*, Baltimore, Maryland, May 1990.
- [13] Z. Li, P. Hsu, and S. Sastry. Grasping and coordinated manipulation by a multifingered robot hand. *IJRR*, 8(4), August 1989.
- [14] Z. Li and S. Sastry. Task-oriented optimal grasping by multifingered robot hands. *IEEE Journal of Robotics and Automation*, 4(1), February 1988.
- [15] X. Markenscoff and C. H. Papadimitriou. Optimum grip of a polygon. *IJRR*, 8(2), April 1989.
- [16] B. Mishra, J. T. Schwartz, and M. Sharir. On the existence and synthesis of multifinger positive grips. *Algorithmica*, 2(4), 1987.
- [17] J. Napier. The prehensile movements of the human hand. *Journal of Bone and Joint Surgery*, 38B(4), 1956.
- [18] V. Nguyen. The synthesis of stable force-closure grasps. Technical Report AI-TR-905, MIT AI Lab., July 1986.
- [19] N. S. Pollard. Parallel methods for synthesizing whole-hand grasps from generalized prototypes. Technical Report AI-TR-1464, MIT AI Lab., January 1994.
- [20] J. Ponce, S. Sullivan, J. Boissonnat, and J. Merlet. On characterizing and computing three- and four-finger force-closure grasps of polyhedral objects. In *Proc. IEEE ICRA*, Atlanta, Georgia, 1993.
- [21] J. K. Salisbury. *Kinematic and Force Analysis of Articulated Hands*. PhD thesis, Department of Mechanical Engineering, Stanford University, 1982.
- [22] J. K. Salisbury and J. J. Craig. Articulated hands: Force control and kinematic issues. *IJRR*, 1(1), Spring 1982.