Vision-based three-finger grasp synthesis constrained by hand geometry

Antonio Morales[1], Pedro J. Sanz[1], Angel P. del Pobil[1] and Andrew H. Fagg[2]

Abstract

This paper presents a practical approach for synthesizing two and three-finger grasps on planar unknown objects by means of vision. We make use visual perception to reduce the uncertainty and to obtain relevant information about the objects. Moreover, we focus on non-modeled planar extruded objects, that can represent most real-world objects. Finally, we also take into account the particular mechanical constraints of the hand, with particular application to the three-fingered Barrett Hand.

Our approach contains three main components. First, a vision processing module that extracts from object images the relevant information for the grasp synthesis: a contour description and a set of feasible grasp regions. A second component is a set of algorithms for synthesizing two and three-finger grasps having into account force-closure and contact stability conditions, with a low computational effort.

Finally, the third and most important component consists of a procedure for constraining these results to the kinematics of the particular hand, something that has not been considered by most authors. In addition, a set oh heuristic metrics for assessing the quality of the computed grasps is described.

As a conclusion, we developed a complete grasping system that implements the components described, that is able to execute any of the grasps computed. Experimental results using the Barrett Hand are shown and discussed.

Keywords

Vision-based grasping, multi-finger grasp synthesis, robotic perception, Barrett hand, dexterous manipulation

[1]Department of Engineering and Computer Science, University Jaume I of Castellón, Spain. E-mail: {morales,sanzp,pobil}@icc.uji.es

[2]Department of Computer Science, University of Massachusetts. E-mail: fagg@cs.umass.edu

1 Introduction

Extensive research in robotic grasping and dexterous manipulation over the last twenty years has established a theoretical framework for grasp analysis, simulation and synthesis. However, there remains a gap between theoretical promise and practical delivery (Okamura et al., 2000). In a survey Shimoga (1996) concluded stating that: “None of the synthesis algorithms have been implemented in real time thus far... Scanty experimental reports have altogether avoided the use of such algorithms.” This is the case for a variety of reasons including prohibitive computational complexity, lack of precise sensing capabilities and, importantly, the difficulty in modeling real-life objects.

These model-based theoretical approaches have dominated the field, and due to the vast existing literature, an attempt to even summarize the state of the art is beyond the purpose of this paper. The reader is, therefore, referred to recent and classical surveys: Pertin-Troccaz (1989); Shimoga (1996); Bicchi (2000); Bicchi and Kumar (2000); Okamura et al. (2000). The first fact that accounts for the lack of practical working systems is that the implementation of model-based algorithms typically requires solving linear or non-linear programming problems in real time. This results in slow multi-finger systems, particularly when contact locations must be determined, which is the problem focused on this paper.

In addition, most algorithms rely on some critical assumptions. The first one is the availability of a complete geometric model of the object to be manipulated, either two-dimensional or three-dimensional. This underlying assumption supports most of current work in the field. Usually a polygonal model -sometimes polyhedral- is the input to algorithms for selecting optimal force-closure contact locations, and these locations are the starting point for grasp analysis and dexterous manipulation methods. Then, they optimize some desired properties in terms of dexterity, equilibrium, stability, dynamic behavior. The only exception is the recent area called “exploration through manipulation” ideally using tactile sensors (Grupen et al., 1995) or proximity devices (Teichmann and Mishra, 2000). Similarly, most of the literature that adopts a friction model for the contacts assumes that friction coefficients (including torsional friction for soft fingers) are known *a priori*.

Whereas the hand configuration is perfectly known in most applications, only in controlled industrial environments the above-mentioned data-object model, friction coefficients, exact object locations etc- can be known *a priori*. In unstructured service scenarios these assumptions are no longer reasonable and the use of vision and tactile sensing for dealing with essentially unknown objects is a must. Since there will be unavoidable errors in positioning and orienting the end-effector, it is important to choose a grasp so that the system performance is robust under these positioning errors. To date, the robustness of a grasp to errors in positioning the effector has not been widely addressed in the literature.

A different trend, motivated by studies of human grasping, has been the so-called knowledge-based approach. This approach tries to simplify the choice of a grasp by reasoning at a more abstract level, similarly as humans. Instead of geometric CAD models, these approaches use more symbolic description of the objects, either based on volumetric primitives (Liu et al., 1989; Stansfield, 1991), or precomputed models of the objects, matched through vision (Kragic, 2001). Alternatively, they use a previous knowledge of grasp prototypes (Mackenzie and Iberall, 1994), that are adapted to particular object by using predefined rules or sensor feedback (Rao et al., 1989; Bekey et al., 1993; Miller et al., 2003).

Grasp synthesis and vision-based contact selection

Summarizing, most of the work that addresses the problem of grasp synthesis take as input a model of the object contour (2D) or surface (3D). According to the model of contact (frictionless, hard, or soft) a static and kinematic description of the problem is built, and a quality metric which describes the appropriateness of a grasp regarding to some desired property is often defined. An analytical method is then employed to optimize the quality criteria. The first work to propose a purely geometrical interpretation for different contact models is Nguyen (1988), who obtained results on polygonal planar objects. Later, Faverjon and Ponce (1991) proposed an algorithm on computing force-closure two-finger grasp on planar curved objects. Simultaneously, (Park and Starr, 1992) and Ponce and Faverjon (1995) proposed a synthesis algorithm for force-closure three-finger grasps for polygonal objects. Other relevant work that has developed this approach are Markenscoff et al. (1990); Francois et al. (1991); Ferrari and Canny (1992); Chen and Burdick (1992); Guo et al. (1992); Mishra and Teichmann (1994); Mirtich and Canny (1994). More recently there has been a considerable effort in the design of faster algorithms: Tung and Kak (1996); Liu (1998); Borst et al. (1999) and Smith et al. (1999).

Some work on planar grasping have relied on vision for extracting the object description. Jarvis (1988) and Hauck et al. (1999) based their methods on the analysis of the object skeleton, while Stanley et al. (1999) performed a quadtree-resolution expansion of the image in order to extract the contour. Bendiksen and Hager (1994) systematically tested in the image plane all the possible orientations and positions of a parallel-jaw gripper while minimizing a quality metric. Later, Sanz et al. (1998) sped up the search by associating thresholds to a set of grasp quality conditions, and relied on a heuristic that made use of symmetry information to find the best grasp. Kamon et al. (1996) used visually-computed features to select the grasping points on the contour of the object, and eventually learned the values of these features that lead to better (more stable) grasps. In a different approach Taylor et al. (1994) tracked the apparent contours while moving the camera to find the best view point of the object. Cipolla and Hollinghurst (1997) used stereo vision for searching planar facets in simple shapes, but they did not deal with the problem of finding stable grasps. Finally, Coelho Jr. et al. (2000) uses visual features in combination with feedback from force-driven contact controllers to learn the best orientation of the hand for grasping an object for the given visual input.

Most of this work involves two-finger grasps, in fact, there has been little, if any, effort to use real images of objects as input for determining three-finger grasps. In this line, two key contributions are Park and Starr (1992) and Ponce and Faverjon (1995). Both proposed different strategies for finding three-finger grasps for 2D polygonal objects using the geometric interpretation of force-closure. Ponce and Faverjon (1995) centered their work on the characterization of three-fingers on force-closure grasps and in the development of an algorithm for reducing the complexity of the linear constraints that describe the problem. On the other hand, Park and Starr (1992) took into account not only those grasps with contact points lying on planar edges but also those with contact points placed on vertexes and concavities of the contour. This second work also proposed a series of algorithms for any of the possible combinations of the different types of contact points. Neither approach, however,considered visual input but instead rely on predefined geometric models of the objects.

Gripper geometric constraints

Other important problems that we face in this paper are the constraints imposed by the geometry of the gripper. Since dexterous manipulation is object-centered - i.e., it is formulated in terms of the arrangement of contacts, forces and interactions with respect to the geometry of the object to be manipulated - little attention has been given to the kinematics of the hand (Bicchi and Kumar, 2000). In fact, force-closure is strongly affected by the mechanical structure and physical restraints of the gripper, and this is often insufficiently considered during the grasp synthesis and analysis stages.

Several authors have developed grasp synthesis methods that consider the particular kinematic constraints of the Barrett Hand. Strandberg (2002); Miller et al. (2003) embedded the hand kinematics in the grasp synthesis algorithm, pruning out the search space. In a different trend Bowers and Lumia (2003) used a set of predefined object shapes and hand configurations that map to a unmodeled object through a fuzzy-logic system. Similarly to our work, Bowers and Lumia presented a complete grasping system.

Miller et al. (2003) focused their work on three-dimensional object models, and validated their approach in a virtual environment. They did not use visual information, and did not make any consideration about uncertainties and robustness

Strandberg (2002) followed a two-dimensional approach similar to ours. The main difference between our approach and that presented by Strandberg is that it was designed *ad hoc* for the Barrett Hand. Despite it argued that some of the ideas could be applied to other grippers, the methodology developed in the work was strongly suited to the characteristics of the Barrett hand. On the contrary, our two-phases approach, first, computing general planar two and three-finger grasps, and then, constraining to the particular kinematics of the gripper, is much more generalizable.

Finally, an interesting approach was followed by Borst et al. (1999), though they did not use the same gripper. They proposed an algorithm that randomly generates the initial contact point for a three-finger grasp, and then computed the rest of the contact points according to the structure of the hand.

Paper goals and outline

The purpose of our paper is to present a complete grasping system that is able the grasp an object having an image of it as only input. The main contributions are several algorithms that, through a geometric analysis, compute the two and three-finger grasps of an object. This analysis takes into account global and local stability conditions. The input is the contours of the objects to be grasped obtained from the images of the objects. In order to reduce the computational cost of the algorithms the analysis of the contour will be focused on certain segments that accomplish the local stability condition of low curvature (Montana, 1991).

Another main contribution is the development of a procedure to constrain the results to the particular kinematics of the hand that will also be introduced.

With this approach we face two of the practical issues that real grasping systems must deal with: uncertainty in the shape of the objects, the computational complexity of the search process, and the mechanical limitations of the components involved.

The paper is organized as follows. The next section presents the assumptions, and the characterization of local and global stability used through the rest of the paper. In section 3 we define the concept of grasp regions and introduce the procedure for computing them. Section 4 contains the description of the algorithms for computing pairs and triplets of grasp points using grasp regions. In section 5 the procedure for adapting these techniques to a particular gripper, namely the Barrett Hand[1], is described. Finally, in section 6 a experimental setup that implements the proposed techniques and some results obtained are described.

2 General assumptions and restrictions

In such a complex problem as grasp synthesis some restrictions must be established to limit its complexity. These restrictions deal with the assumptions we make about the physical characteristics of the objects, the physical and mechanical properties of the gripper, and the physical model of contact between the gripper and the objects:

I Visually-guided grasping. As little knowledge as possible about the objects to be manipulated is assumed before hand. The necessary information about the object shape and pose is obtained from visual input.

II 2D grasp synthesis. This work is focused on two-dimensional grasp synthesis. This means that our methodology is assuming planar representations of the objects and grasps.

III Planar objects. Two important assumptions are made about objects. First, the objects are planar, that is, their height is constant, and they do not contain hidden

[1] See `http://www.barretttechnology.com/`

holes or occluded features. Second, the mass is homogeneously distributed. This allows us to reason about the mass properties of the object, in particular, the location of the center of masses, using exclusively visual perception.

IV **Two and three-finger grippers.** The number of fingers in the grippers in the bibliography varies from two to four. We only focus on two and three-finger gripper because they were the only kind of available grippers for the practical development of this paper. Moreover two and three-finger grippers are the most commonly used in robotic manipulation applications

V **Contacts with friction.** We assume friction contacts as well as the Coulomb friction model. Despite, the friction coefficient between the fingers and the surface of the object is not known beforehand the algorithms proposed assume the existence of some friction.

2.1 System architecture

In addition to the assumptions introduced above, some other assumptions about the organization of the grasping system are necessary. This includes basically a description of the different modules of the grasping system and the connections between them.

The grasping system is composed of four modules:

i) **Vision** A camera placed over the workspace takes a top picture of the object. An image processing system analyzes the image and extracts the contour of the object, along with its location.

ii) **Grasp synthesis** This module uses the contour provided by the previous module to compute a list of feasible two and three-finger grasps. The results are filtered to meet minimum stability conditions.

iii) **Grasp selection** In this stage, an intelligent algorithm would select one of the candidate grasps resulting from the previous phase. This selection may be conditioned by their appropriateness to the task to be executed.

iv) **Execution** Finally, a control module will execute the selected grasp. This execution consists in approaching the manipulator to the object location and gripping it. This takes into account all the necessary substeps like preshaping the gripper, and controlling the applied forces.

2.2 Contact and object stability

Our approach to stability characterization is decomposed into object and contact stability. Contact stability includes those aspects that involve the contact of a single finger against the surface of the object. Object stability focuses in the set of forces and torques exerted by the fingers on the object in order to counteract external forces or to exert forces and torques on the object. This differentiation is equivalent to the distinction between *spatial grasp stability* and *contact grasp stability* presented by Montana (1991).

A critical issue is the type of contact. A typical assumption in analytical approaches is that of hard fingers making point contacts (see e.g. Ponce and Faverjon (1995)). However, it is very convenient that the actual robot fingers are to some extent soft and with a polygon contact. In this case, the contact allows rotational friction, i.e., a torque can be exerted through the contact in order to balance an external torque. It must be noted that Nguyen's force-closure condition for a planar body assumes that all the external forces are coplanar with the object plane. A torque resulting from a force perpendicular to the object plane, such as gravity, could not be balanced by two hard fingers. A soft finger is characterized by the fact that its surface is deformed as it enters in contact with the object. Human fingers are highly deformable, which partly accounts for the dexterity and stability of two-finger human grasps. Robot soft fingers are usually designed by enclosing them in a rubber covering. This is the case in our experimental approach, the rubber has been chosen so that it is compliant and exhibits good friction properties. In this way, we take advantage of non-point contacts, considering the fact that the stability of a grasp position is directly related to the surface of contact between fingers and object.

An important study about the forces that appears in the contact between two objects is developed by Mason and Salisbury (1985). In our case the most relevant lesson from this work is the definition of *non-sliding condition*. This condition is based on the *Coulomb's Law*: "the tangential force of friction during sliding is directed opposite to the direction of motion, with magnitude proportional to the normal force". The constant of proportionality is the *coefficient of static friction* μ, and depends on the nature of the contacting materials.

As indicated by Montana (1992), another important aspect when talking about contact stability is the curvature of the surfaces in contact. Higher curvature produces higher instability in the contact. It is important to note that the curvature not only refers to curvature of the object surface, but also the curvature of the finger.

Finally, in the framework of grasping, *object stability* is defined as the property of a grasp that is able to maintain the object unaffected by external disturbances. A key concept in object stability when friction is involved is *force-closure*. As stated by Nguyen (1988) force-closure is achieved by a set of contact points, when any external force/torque pair can be counteracted by the forces and torques exerted through the contact points.

3 Contour analysis: Determination of grasp regions

This section is focused on the application of the contact stability assumptions. This will lead to the introduction of the concept of grasp region, a key abstraction that reduces the necessary computations, and a practical tool in the further stages of the grasp synthesis. First we need define a practical criterion that encompasses the contact stability constraints.

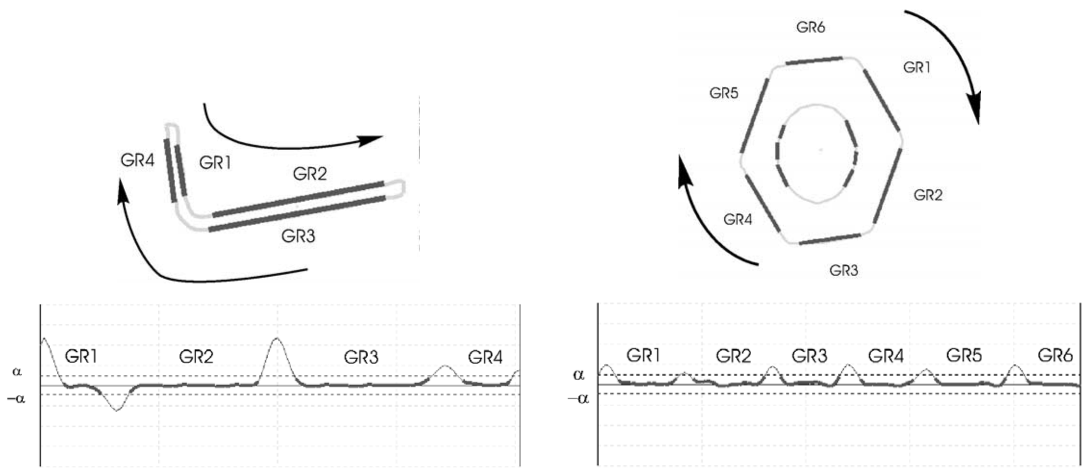

Figure 1: Grasp regions on an Allen key's contour (left) and a nut (right). Based on the analysis of the curvature functions (below), some grasp regions are identified. They are labeled and highlighted with solid lines on the contours (above).

3.1 Finger adaptation criterion

As mentioned in the previous section, the stability of the contact of a finger at a point in the contour is directly related to the curvature of the surface at that point (Montana, 1992). Considering this fact, we establish the *finger adaptation criterion.* It evaluates the fitting of the gripper's finger against the object surface in terms of the estimated contact area.

Our implementation of the criterion establishes that the object's surface curvature must not exceed a threshold $\hat{\alpha}$ (*curvature threshold*) for being considered stable.

3.2 Determination of grasp regions

The contiguous regions on the contour that comply with the finger adaptation criterion are called *grasp regions*, i.e.: those having curvature under $\hat{\alpha}$ for all their points.

Grasp regions are computed by means of the *curvature function.* Given a discrete contour $C(u)$ the curvature function $V(u)$ computes the curvature at each point u of the contour. To measure the curvature of digital contours we have used the *k-angular bending* notion defined by Rosenfled and Johnston (1973). A value of zeros means a planar point in the contour; positive values indicate convexities, and negative values indicate concavities (see the examples in Fig. 1 and 2).

The grasp regions are determined by grouping consecutive points with absolute curvature below the curvature threshold $\hat{\alpha}$. The size of a grasp region is lower-bounded, since, being too small, they would not allow for imprecision in the positioning of the robot fingers. Figure 1 illustrates the use of the curvature function to find grasp regions on the object's contour.

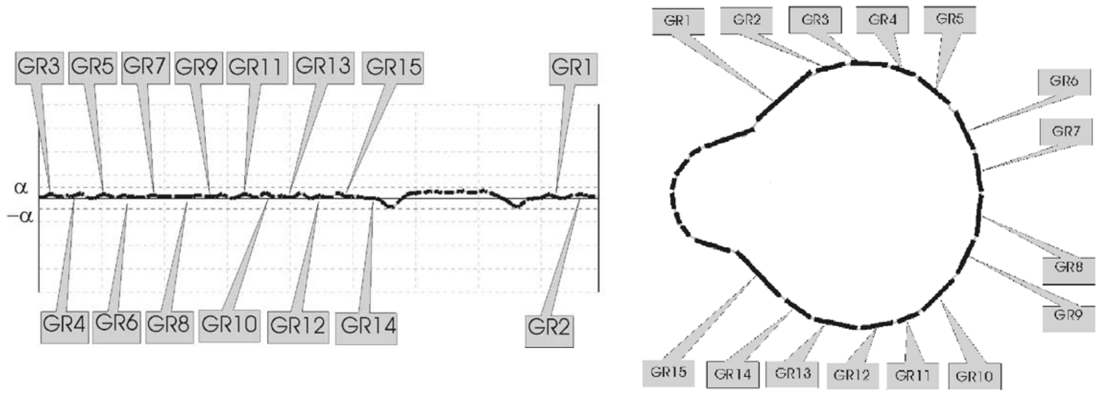

Figure 2: Grasping regions on a circular contour. The use of the accumulated curvature produces a set of small regions, instead of a large one covering most of the contour.

Finally, each grasp region is abstracted as a straight segment called *grasp region segment*. These segments are computed from the line defined by the minimal square fitting of the points of the region, bounded by the first and last points of the region. The inward normal direction (pointing to the interior of the object) of the segment is used as representation of the inward normals of all the points of the contour included in the grasp region. This representation is compact and appropriate for the computation in the further stages.

Ideally, it should be possible to approximate these regions as straight lines (see samples in Fig. 1). However, in some objects, such as the one shown in figure 2, large contiguous portions of the contour lie under the curvature threshold and produce large grasp regions that can not be approximated as straight segments. A way to avoid this effect is to break long regions in smaller pieces that can be approximated as straight lines. With this purpose, the accumulated curvature along a grasp region is used to limit its size.

The use of such an abstraction of grasp regions has important practical implications . First the use of grasp regions discards those points with a high curvature, that is, with an unstable contact. And second, the use of grasp regions reduces the complexity of the problem of converting a contour composed, potentially, of hundreds of points, to a few grasp regions. This allows for the combinatorial techniques of the further stages to perform their computations in a reasonable amount of time.

4 Grasp synthesis

In this section we address the characterization of object stability for two and three grasp contacts. We also provide two practical procedures that, taking the set of grasp regions extracted from the contour of the object, produce a list of pairs and triplets of grasp contacts.

The scheme that we follow in the explanation of both procedures is similar. First we define which conditions made a two or three-finger grasp stable. With this purpose, we define a *force-closure criterion* for each case.

Next, we describe a procedure for computing pairs and triples of grasp points that meet this criterion. Both procedures take as inputs the grasp regions provided by the contour analysis. The first part of the procedures, the *selection of compatible regions*, are aimed at finding pairs or triples of grasp regions that contain pairs or triples of grasp points that meet the force-closure criterion. We define, for each case, several conditions that must meet a set of regions to be compatible.

The second step in both procedures is the *grasp refinement*. Given a pair or triplet of grasp regions, it computes a pair or triplet of contact points, one per region, that meets the force-closure condition.

The results of both procedures is, respectively, a list of pairs and triples of contact points. In section 5 we describe how this pairs and triplets can be constrained to the kinematics of the particular hand.

4.1 Selection of grasp pairs

We define the *force-closure criterion* with the goal of characterizing force-closure grasps using two fingers. The purpose of this criterion is to ensure that the gripper does not cause the object to slide due to a torque when it closes its fingers to grasp it. The evaluation of this criterion is based on the concept of *friction cone* (see subsection 2.2).

Theoretically, as stated by Nguyen (1988), force-closure with two friction contact points is achieved when the grasping line (the line that joins the contact points) lies inside both friction cones. These angles should not exceed a threshold $\hat{\beta}$ (*angular threshold*). The value of this threshold is directly related to the friction coefficient μ according to $\hat{\beta} = \arctan \mu$. The value is selected experimentally in a conservative way. The criterion selects that grasps whose angles β_1 and β_2 (see figure 3) are bellow the threshold $\hat{\beta}$ and discards the others.

Step 1: Selection of compatible regions

Once the grasp regions have been found, the next step is to build a list of pairs of grasp regions where the robot's fingers can be placed to grasp the object. These regions are characterized by containing two points –one per region– such that, when the robot's fingers are placed at them, the grasp complies with the force-closure criterion. The regions in each pair will be termed hereinafter *compatible regions*.

The compatibility between two grasp regions is verified through the following conditions:

1. The angle between the normal vectors to each region is in the range $180.0 \pm 2\hat{\beta}$ degrees.
2. The projection of each region, in the direction of its normal, intersects with the other region, that is, the regions are faced each other.

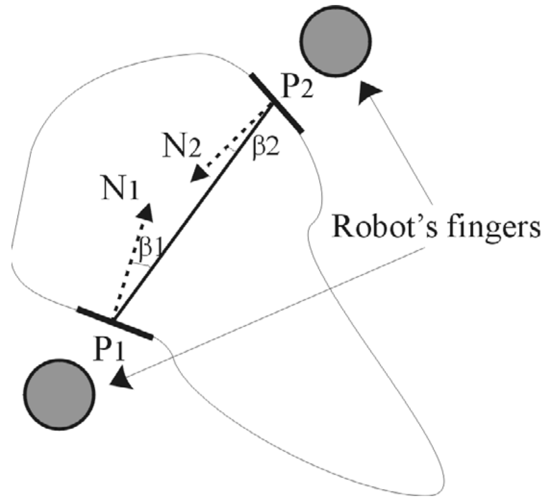

Figure 3: Geometric interpretation of the quality criterion used for the selection of grasps. The contact points are P_1 and P_2, $\overline{P_1P_2}$ represents the grasp line. N_1 and N_2 are the normal directions to the surface in the contact points. Finally β_1 and β_2 are the angles formed between the normal directions (the axis of the friction cones) and the grasping line.

With regard to the first condition, the angular threshold ($\hat{\beta}$) is considered for tolerance in the angle between both vectors; therefore, the allowed range between these vectors is $180.0 \pm 2\hat{\beta}$ degrees. Figure 4 shows two grasp regions and how each region is projected perpendicularly on the line defines by the opposite region. The intersection between the projection of a region and the other region, if exists, is called *feasible region*, If both feasible regions exists and the first condition above is met, these two regions are said to be compatible.

This test is not performed only between regions within the same contour, but also between regions of different contours, so all combinations of regions are checked. This produces pairs of compatible regions such that the normals to each region points to the opposite region, and pairs in which the normals point toward outside this space. The first kind of pairs correspond to squeezing grasps, which could be executed by closing the gripper fingers on the object and are the only ones considered by many approaches to grasp determination. The others correspond to expansion grasps.

Step 2: Grasp refinement

This second step tries to find the contact points within each of the regions in a compatible pair. The procedure is simple and relies in some of the results obtained in the previous step.

The projection of a region defines in the opposite region a subregion called *feasible region*. These subregions are especially interesting since any point within any of these regions has a correspondence in the opposite region thanks to the projection procedure.

The refinement procedure chooses the center of each feasible region as the contact point

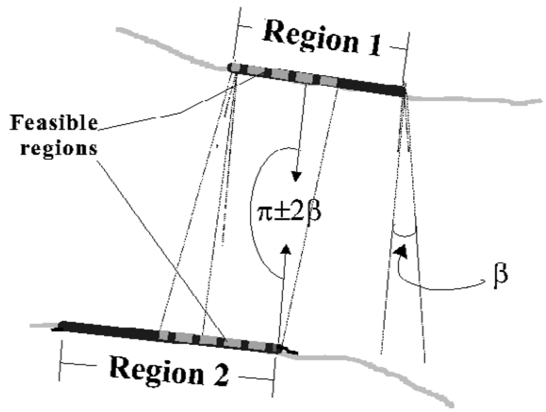

Figure 4: Compatibility test between grasp regions.

in this region for the two-finger grasp between both grasp regions.

Finally, the result of the whole procedure is the list of all the pairs of contact points, hereinafter two-finger grasps, found from the pairs of compatible regions. These grasps meet the *force-closure criterion*, and both contact points meet the *finger adaptation criterion*.

4.2 Selection of grasp triplets

A similar approach is performed for the three-finger case. In this case a force-closure criterion is also defined, however its definition is slight different.

Ponce and Faverjon (1995, Proposition 3) state that a sufficient condition for reaching force-closure with three fingers is that the intersection of the friction cones of the three contact points is not empty, and that the unit normal vectors to the surface defined by the contact points *positively spans the plane*. Three vectors positively span the plane $\Re^2$ if any of them can be written as a positive combination of the other two (see fig. 5). In a more formal way n vectors $u_i, i = 1 \ldots n$ spans the plane if and only if

$$\forall i, i = 1 \ldots n, \exists\ c_j > 0, u_i = \sum_{j=1, i \neq j}^{n} c_j u_j \tag{1}$$

Similar conditions are stated by Park and Starr (1992, Theorem 1). This theorem introduces the term *force focus point* as a point "such that, for each contact point, a line connecting the focus point and the contact point is located within the convex cone of the contact wrench system"

Summarizing, force-closure by triplets of contact points is achieved when two conditions are met:

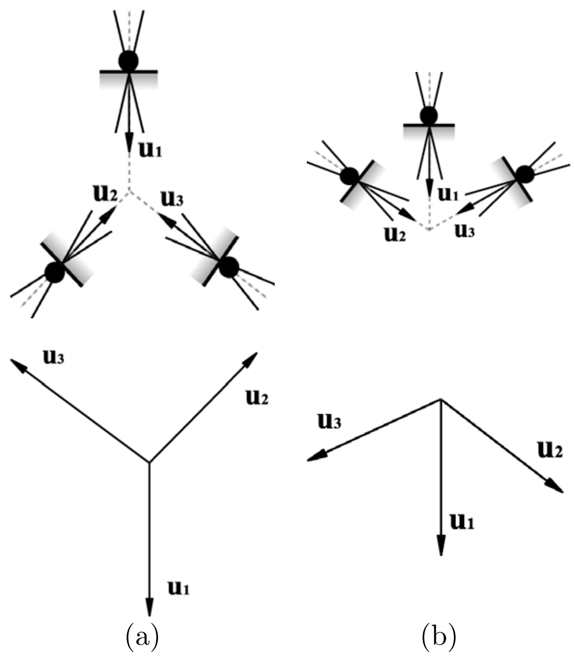

Figure 5: The grasp (a) achieves force closure, whereas the grasp (b) does not

1. The intersection of the regions defined by the friction cones of the points in the three grasp regions is not empty. (see fig. 6). In our approach the width of the friction cone is defined by the threshold $\hat{\beta}$ (*angular threshold*), which is the same defined for the two-finger case.
2. The unit normal vector to the surfaces defined by the grasp regions *positively span the plane*. That is, the three vectors are not contained in the same half space (see fig. 5).

Step 1: Region suitability test

Once all the grasp regions that appear on the contour have been determined, we have to find which triplets of regions could allow a three-finger grasp.

Each point in a grasp region could be a contact point and defines a friction cone. The interior cone, the half facing the interior of the object, defines a region in the plane. We will apply the term *interior projection of a grasp region* to the union of the interior cones of each of the points of that grasp region (see Fig. 6a). If the intersection of the interior projections of three grasp regions exists there will be at least three points, one per grasp region, for which the intersection of the friction cones will not be empty (see fig. 6b).

Taking into account the force-closure-criterion it is possible to state that a set of three regions contains a three-finger grasp if two conditions are meet:

1. The normal vectors of the three regions positively span the plane.

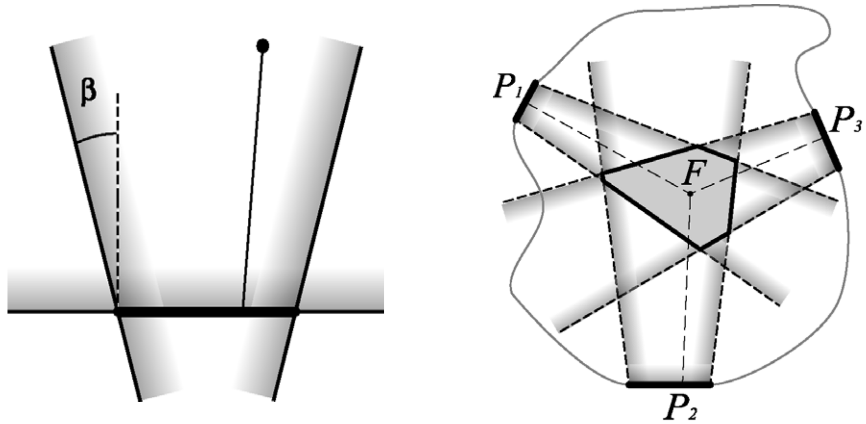

(a) (b)

Figure 6: (a) The projection subspace is build from the intersection of the three half spaces defined by each grasp region. Notice too, how the projection of a point on the regions is modulated according to the values of the $\hat{\beta}$ threshold. (b) the intersection of the projections of the three regions is computed. If not empty, the grasp focus point F is defined as the centroid of the intersection area and its projection on the regions results in the grasp points P_1, P_2 and P_3.

2. The intersection of the interior projections of each of the grasp regions is not empty.

The first step of the grasp determination process searches for combinations of three grasp regions meeting the two above conditions. In practice, the first one is easy to check, but the second one requires a deeper explanation.

Each interior projection can be defined by the set of solutions of three linear inequalities each of them defining a half-space on the plane (see fig. 6). The intersection, if it exists, of the interior projections of three grasp regions is defined by nine linear inequalities. This set of inequalities or constraints defines an standard LP problem called *the vertex enumeration problem* (Chávtal, 1983). The goal of these kind of problems is to compute the vertexes, if any exist, of the convex polytope defined by a set of linear inequalities[2].

Step 2: Determination of grasp points

The last stage is to determine which points in each grasp region will be the contact points. This can be done as an indirect result of the operations performed in the previous phase. When determining the existence of the intersection of the interior projection of the grasp regions, its boundaries are also determined. This intersection is a convex region whose centroid can be easily computed. This center is defined as the grasp focus, as it was introduced in Park and Starr (1992).

The last step is to project this focus onto each of the grasp regions. This projection

[2]For more information about this subject the reader is referred to Chávtal (1983) and Fukuda (2000).

must take into account the deviation due to the tolerance introduced by the $\hat{\beta}$ threshold (see figures 6a and b). In theory any point that belongs to this intersection could be reprojected over the regions, however the centroid has the property of being the furthest point from the boundaries of that region. So it is the point less sensitive to possible inaccuracies.

4.3 Possible improvements

The algorithms are complete at this point, however several aspects could be improved. Among them, one specially relevant is related wit the fact that both two and three-finger algorithms assume that a stable contact point is such that the surface curvature is planar. However, it is a known fact that concavities can increase the stability of contact, under some assumptions about the shape of the finger. Park and Starr (1992) indicate how this problem can be addressed. Indeed, Boivin et al. (2004) describe a working grasping system that implement the use of corners as contact regions. This former work is an improvement of the approach we describe in this paper

5 Constraining the gripper kinematics

The sets of candidate grasp, pairs and triplets, resulting from the preceding strategies are in agreement with model-based criteria for force-closure and local stability, but do not include any consideration of hand geometry. In this section we show how to adapt the results of the procedures shown in the previous section to the characteristics of a particular gripper, the Barret Hand.

A Barrett hand is composed of three-fingers with a total of four controllable degrees of freedom. Each finger consists of two flexion degrees of freedom which are driven by a single motor. The fourth motor controls abduction/adduction of two of the fingers which are symmetrically placed on either side of the remaining finger, the *thumb*. When adducted, the two fingers flex in parallel to the thumb; when fully abducted, the two fingers flex in opposition to the thumb. This design is very compact, allowing for the motors and control CPU to be embedded into the wrist, but limits the possible placement of finger contacts. In this paper we will work with a planar geometric top-view description of the hand shown in Figure 7.

For a given instance of a previously-computed candidate grasp pair or triplet, our problem is to find a set of feasible hand configurations that correspond to that particular candidate. It would be also desirable to provide a quality measure for each configuration. A *hand configuration* is, therefore, given by the set of four controllable parameters of the Barrett Hand (α, d_1,d_2, d_3 in figure 7) and it *corresponds* to a particular grasp pair or triplet as long as the three fingers make contact with the object contour at the same contact points as defined by the initial candidate.

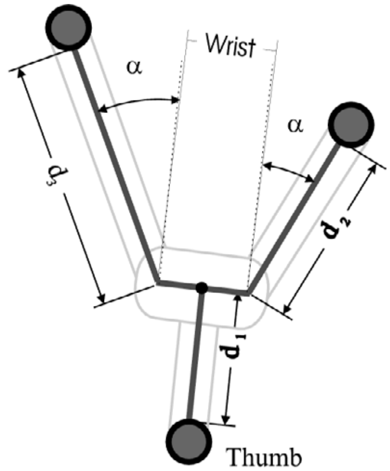

Figure 7: On the right, a picture of the Barrett Hand, and on the left a drawing showing the hand configuration schematics. The hand is composed of a thumb, and two opposite fingers that can be spread symmetrically along the axis defined by the thumb. "Wrist" is the separation between the fingers at the palm. The four controllable degrees of freedom are the spread α, and the finger extension d_1 for the thumb and d_2 and d_3 for the other two fingers.

5.1 Constraining pairs of candidate grasp points.

Although the Barrett Hand is three-fingered, its design allows for two-finger grasps, one of the fingers being a *virtual finger* (Mackenzie and Iberall, 1994) made up of the two symmetric opposite fingers. In this case a two-finger grasp can be used to obtain up to two different configurations.

This consists of a procedure that recovers information produced for the two-finger grasp synthesis, in particular the grasp regions where both contact points of the grasp are located[3]. One of the contact points is selected to be the thumb. The other contact point represents the virtual finger. The real contacts should be placed symmetrically on both sides of the virtual contact along the object contour. In order to keep the contact stability properties the locations of the contact points are constrained to be within the same grasp region of the virtual contact.

In practice, the only aspect that must be checked is whether there is enough space in the grasp region around a virtual contact point for placing the two opposite fingers (Fig. 8). Due to the geometry of the hand, there exists a minimal separation between the two opposite fingers, given by the wrist width. Given a two-finger grasp, the space around both contact points is checked, if both regions are too small the grasp pair is discarded.

Once the placement of the three contact points is selected, the values of the parameters of the hand configuration d_1, d_2 and d_3 are selected so the wrist position W_x, W_y is located halfway between the two contacts of the original two-finger grasp. That is, $d_1 = d_2 = d_3$, and, obviously $\alpha_h = 0°$. The orientation W_θ is easily deduced from this construction.

[3]See section 4.1 and figure 3 as a reminder.

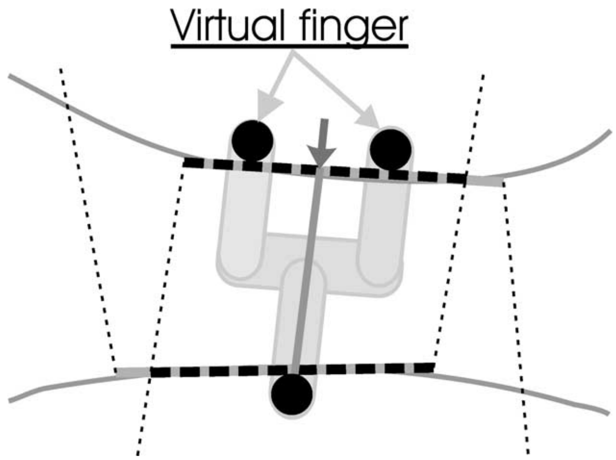

Figure 8: An adaptation of a two-finger grasp. In the top grasp region there is room for placing the two opposite fingers.

5.2 Constraining triplets of candidate grasp points.

The goal when constraining a three-finger grasp is to find a configuration of the hand such that the positions of the three fingers is the same as in the original grasp, and the orientation of forces exerted by the fingers are similar. This problem can be very complex since the description of a configuration is composed by seven independent parameters (d_1, d_2, d_3, α_h, W_y, W_y, and W_θ) that have to be determined. The goal is then to develop a procedure able to obtain a hand configuration from a initial three-finger grasp with the least possible computational complexity.

The key of our procedure consists of fixing the contact points of the hand configuration in the same locations as the three-finger grasp (Fig. 9). This, along with the constrains imposed with the hand kinematics, reduces the problem to a more tractable one.

Given a thumb position and wrist orientation W_θ, there is one solution most of the time for α_h, d_1,d_2 and d_3. This can be deduced from the geometrical construction shown in figure 10. Due to the mechanical constraints of the hand (the spread angle of the opposite fingers is symmetric), the triangles defined by segments a and b, $\widehat{ab}$,and a' and b', $\widehat{a'b'}$, have to be proportional. Then a pair of trivial equations can be stated:

$$\left.\begin{array}{c} tan(\alpha_h) = \dfrac{a}{b} = \dfrac{a'}{b'} \\ b = b' + h \end{array}\right\} \tag{2}$$

Since the position of the fingers, the orientation of the thumb and the width of the wrist are known, the parameters a, a' and h are also known. Moreover a and a' are positive due to the particular kinematics of the hand, and h is also chosen to be positive (by choosing $\widehat{ab}$ to be the larger triangle). The system of equations (2) can be solved:

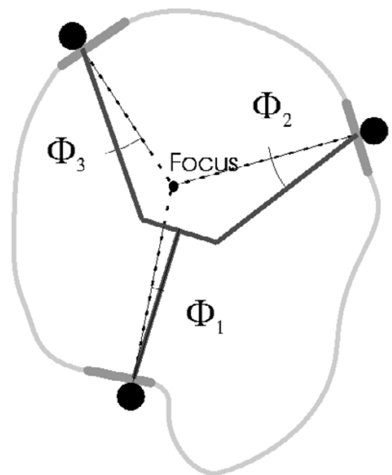

Figure 9: Hand configuration (solid lines) for a three-finger grasp (dashed lines). Φ_1, Φ_2 and Φ_3 are the differences in the angles between the directions of the desired forces, ideally normal to the surface, and the actual directions of the forces exerted by the finger flexion.

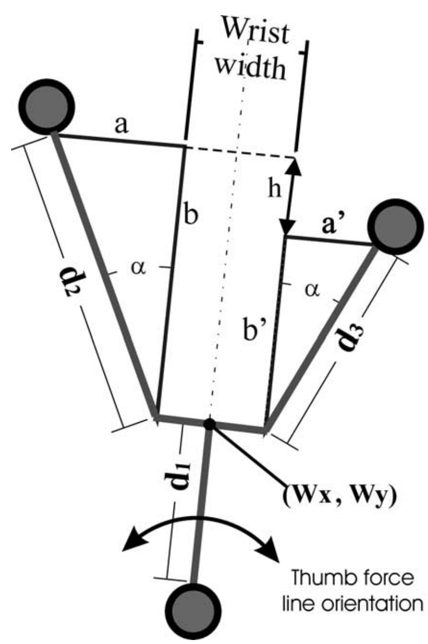

Figure 10: This figure shows all the variables involved in the geometrical construction of a configuration.

$$b' = \frac{a'h}{a - a'} \quad (3)$$

$$b = \frac{ah}{a - a'} \quad (4)$$

The expressions (3) and (4) are the solution of the 2 an their values are unique unless $a = a'$. For interpreting this last case we must focus on the value of h. If $h = 0$, the triangle formed by the three grasp points is either isosceles or equilateral. In that case the only valid orientation of the force line of the thumb is the symmetry axis between the fingers. In the case that $h \neq 0$, it can be interpreted just by noticing that

$$tan(\alpha_h) = \frac{a - a'}{h} = 0$$

This is only possible if $\alpha_h = \{0, \pi\}$. This last result means that the solution, the position of the wrist, is located at an infinite distance along the thumb axis. This is kinematically impossible, so this case, this orientation of thumb line, is not relevant to our analysis.

Once the the values h and h‘ are known, the four parameters of the hand can be easily computed. And as a main consequence of this analysis, there will be a unique solution (α_h, d_1,d_2 and d_3) for a given thumb position W_θ and orientation that satisfies the geometric constraints imposed by the kinematic structure of the Barrett hand as well as the physical bounds of those four parameters. Moreover the position of the wrist (W_x, W_y) can be also easily computed from the above parameters.

With this geometrical construction the only undetermined parameter yet is W_θ. The values of W_θ are restricted by the mechanics of the gripper to a range of values. The configuration axis has to lie between the the opposite fingers[4] and W_θ cannot has a value out of the range $[0 \ldots \pi]$. The lower value is mechanically fixed, while the upper is a consequence of the positive spanning condition that was exposed extensively in section 4.2. All the solutions such that W_θ has a value out of the mentioned range are discarded.

5.3 Quality criteria for a hand configuration

Several criteria can be used to measure the quality of a hand configuration corresponding to a grasp triplet. Once a triplet of contact points has been following the procedure described in previous section the following quality criteria can be defined for it.

Feasibility criterion. If for the three placements of the thumb and all orientations no feasible solution for (α, d_1,d_2, d_3) exists within their physical bounds, then the triplet is discarded as infeasible.

[4] Actually, this is not completely true. It is possible that a flexing finger cross the axis defined by thumb. But these cases are so extreme that they are simply discarded

Force Line criterion (FL). Since no knowledge about the actual value of the friction coefficient can be assumed, the grasp stability critically depends on the actual direction of the forces exerted by the fingers. We propose as a quality measure the sum of the square differences between the directions of the finger forces and those of the desired grasp. The desired force lines are given by the lines joining the contact points and the grasp focus:

$$FL = {\Phi_1}^2 + {\Phi_2}^2 + {\Phi_3}^2$$

where the meaning of Φ_1, Φ_2 and Φ_3 is shown in Figure 9. Lower values of FL correspond to better configurations. It is worthwhile to note that the finger force lines intersect in a different point or configuration focus which might even be out of the intersection region of figure 6 (in fact another valid measure would be the distance between the two foci). In addition, the closer the new focus is to the boundary of the region, the less robust the grasp will be with respect to errors in finger positioning.

Finger Extension criterion (FE). Though much literature has been devoted to planar grasping, planar objects simply do not exist, or rather, a true planar object –such as a paper cut-out– just cannot be grasped. Then, the term *planar object* should be properly understood as the 3-dimensional object resulting from the extrusion of a certain planar shape. Such an object will obviously have a finite height, and consequently a finger contact point will also have a certain height. Since the three fingers of the Barret hand are three copies of the same finger design, it is apparent from Fig. 7 that two fingertips will not have the same height unless their extension parameters d are identical, regardless of the value of α. Consequently, a quality criterion can be naturally introduced as:

$$FE = (d_1 - d_2)^2 + (d_1 - d_3)^2 + (d_2 - d_3)^2$$

If we consider the plane defined by the three fingertips, the larger FE is, the larger the torque that is applied out of the horizontal plane.

Focus Centering Criterion (CD) This criterion is designed to obtain stable grips with respect to wrenches generated by gravitational and inertial forces. These wrenches are minimum when the center of the grip (configuration focus)[5] is closer to the the mass center of the object. The center of the grip is point where the force lines of the three fingers meet. In general, this point may not exist, but within the particular kinematics of the Barrett hand it always exists. Thus, the quality value is measured as the distance D_G between such a point and the shape centroid.

$$CD = D_G$$

Reachability criterion. Finally, any hand posture must be reachable, which amounts to finding a plan to move the hand safely to its final configuration. In its general formulation this is a complex problem in reach planning for grasping. We will only consider here some simple criteria discussed below.

[5]Do not confuse this focus with the force focus of the candidate triplet defined in section 4.2.

Reachability and feasibility criteria are used as discarding criteria. If a computed configuration does not pass the tests defined by these criteria it is discarded.

The other criteria can be used either as metrics in the 1-dimensional search for finding the best thumb orientation for the constraining of a grasp triplet or as quality heuristics for comparing and ranking different configurations. To use them together we defines and aggregation rule of the normalized values of the criteria:

$$Q = \overline{FC} + \overline{FE} + \overline{CD}$$

The criteria are normalized in order to have the acceptable values within the range [0..1] being 0 the optimal best value, and any value larger than 1 a bad quality value. All the criteria are normalized by dividing their values by a particular constant. FC is divided by $N_{FC} = 3 * (\frac{\beta}{2})^2$ where $\hat{\beta}$ is the angular threshold defined in subsection 4.1. In a similar way $N_{FE} = \eta^2$, where η is the optimal finger extension in which the contact area with the vertical surface of the sides of the object is maximized (in the case of the Barrett Hand this value is about 101 mm). Finally, CD is divided by $N_{CD} = \frac{\phi_{min} + \phi_{max}}{2}$ where ϕ_{min} and ϕ_{max} are respectively the minimal and maximal inertia axis of the contour of the object. These normalization values are somewhat arbitrary, and only with experimental validation can their values can be better adjusted. For additional discussion of these criteria see Chinellato et al. (2003) and Chinellato et al. (2004).

6 Implementation and experimental results

The procedures described in the previous sections have been implemented and tested within a complete grasping system, the UMass torso. In this section we describe the experimental setup and the results obtained from the implementation of the strategy.

6.1 Experimental Setup: the UMass Torso

Our experiments have been implemented using the UMass Torso (Platt et al., 2003). This humanoid robot (Fig. 11) consists of two Whole Arm Manipulators (WAMs) from Barrett Technologies, two Barrett Hands, a BiSight stereo head, and a quadraphonic audio system. The WAMs are seven degree-of-freedom manipulators, whose kinematics are roughly anthropomorphic. The BiSight head consists of two cameras and ten controllable degrees of freedom (four mechanical and six optical).

A high-level *control basis* provides a set of look/reach/grasp primitives upon which our reach and grasp sequences are implemented (Coelho and Grupen, 1996; Platt et al., 2002). The stereo vision system estimates the two-dimensional location of the target object on the table, and provides a monocular image for surface curvature analysis. Once a grip is selected (consisting of contact locations and a hand posture), the hand is preshaped and positioned above the object.

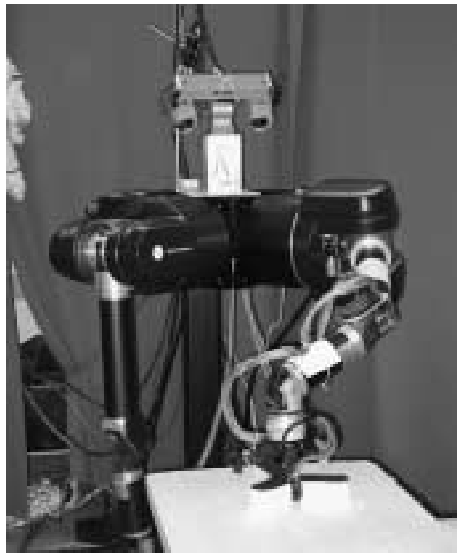

Figure 11: The UMASS Torso. A humanoid robotic system developed in the Laboratory for Perceptual Robotics in the University of Massachusetts.

A tactile-driven grasp controller in then engaged in order to establish a stable grasp upon the object. The grasp controller uses sensor-derived estimates of the net force and moment applied to the object in order to further refine the contact locations (Coelho and Grupen, 1996; Platt et al., 2002). This latter stage is critical in addressing any misalignments that might occur due to errors in visuo-motor calibration or to movements of the target object during the grasping process. Once a stable grasp is established, the object is lifted and transported to a designated location.

6.2 Algorithmic results

The procedures shown in the previous sections have been implemented tested in different stages. In this subsection we focus on the results obtained by the procedures for generating pairs and triplets of grasp candidates. Then, in the next subsection we will describe some results observed during the experimental tests of the grasping system.

The implementation of the results has been written in C++. Single gray scale images of real objects have been used for testing the synthesis procedures.

Figures 12 and 13 show some examples of the results obtained by the grasp synthesis procedures. Tables 1 and 2 contain the times employed by an ordinary PC for computing the results shown in the figures. These performance times show that our strategies are able to compute the grasp pairs and triplets in a reduced time, allowing them to be used on-line. Indeed, most of the processing time is dedicated to image processing. The time to compute the grasps is significant only when the number of grasp regions found is large.

[6] Synthetic figure adapted from the one that appears in Faverjon and Ponce (1991)

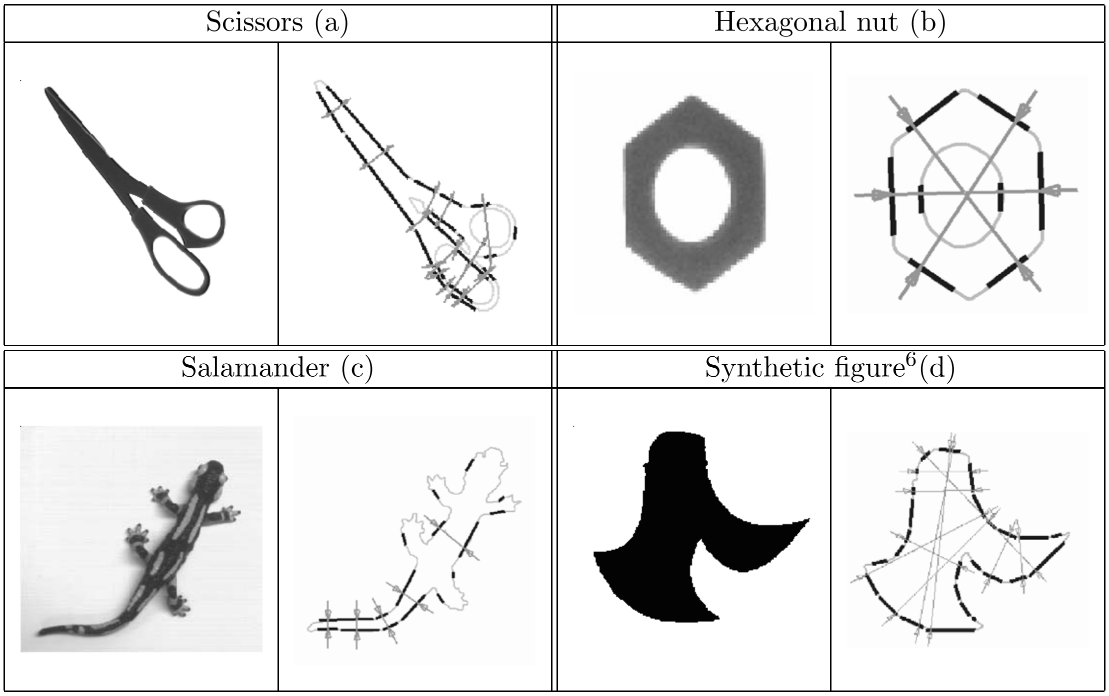

Figure 12: Some examples of the pairs of contact points computed for several objects. For each case, the original image of the object is showed and all the pairs of contact points are showed(each pair joined by a thin line), as well as grasp regions (solid black lines on each contour).

	Image processing time[7]	Region extraction		Pair computation	
		Region	Time	Pair	Time
Scissors	0.030 sec.	19	0.0044 sec.	20 out of 171	0.0007 sec.
Hexagonal Nut	0.009 sec.	8	0.0024 sec.	6 out of 28	0.0001 sec.
Toy Salamander	0.038 sec.	16	0.0050 sec.	5 out of 120	0.0005 sec.
Synthetic figure	0.020 sec.	26	0.0056 sec.	26 out of 325	0.0006 sec.

Table 1: Time measurements obtained while computing the candidate pairs of contact points. The second, fourth and sixth columns contain the times employed in the image processing (binarization, contour extraction and grouping), the determination of the grasp regions and the search of for candidate pairs. The third column shows the number of grasp regions found in all of the contours of the object. Finally the fifth column indicates the number of valid pairs found out of the number of possible pairs of grasp regions.

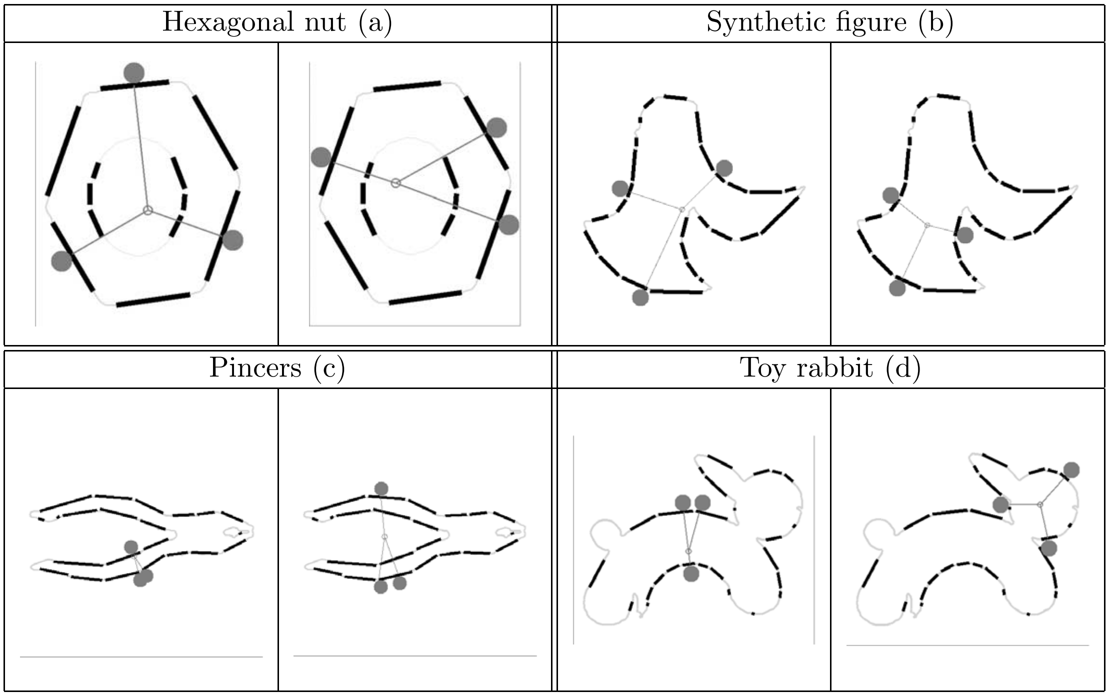

Figure 13: Some examples of the triplets of contact points computed for several objects. Only a couple of triplets is shown for each object. Moreover all grasp regions are shown. The contour of each object is extracted from input images as those shown in Fig. 12.

6.2.1 Experimental results

A set of real objects has been built for this experiment. Their main features are that they are planar objects with a constant height made of an homogeneous material. Moreover, the colors of the object have been selected to simplify the image processing. In this subsection not only the pairs and triplets of contact points are computed but also the hand configurations.

Figures 14 and 15 show the configurations computed from the images of three different objects. Figure 14 shown the first best three configurations with the quality values computed

[7] All the times that are shown in tables 1 and 2, have been obtained on Intel Celeron 800 MHz processor based computer, with 256 Mb RAM, running a 2.4 Linux kernel.

	Image processing time	**Regions extraction**		**Triplets computation**	
		Regions	**Time**	**Triplets**	**Time**
Hexagonal nut	0.020 sec.	12	0.0035 sec.	32 out of 70	0.0236 sec.
Synthetic figure	0.020 sec.	26	0.0056 sec.	26 out of 690	0.1825 sec.
Pincers	0.019 sec.	22	0.0030 sec.	24 out of 430	0.1295 sec.
Toy Rabbit	0.022 sec.	26	0.0056 sec.	28 out of 696	0.2092 sec.

Table 2: Time measurements obtained while computing the candidate triplets of contact points. The meaning of the data in the columns is equivalent to those in table 1.

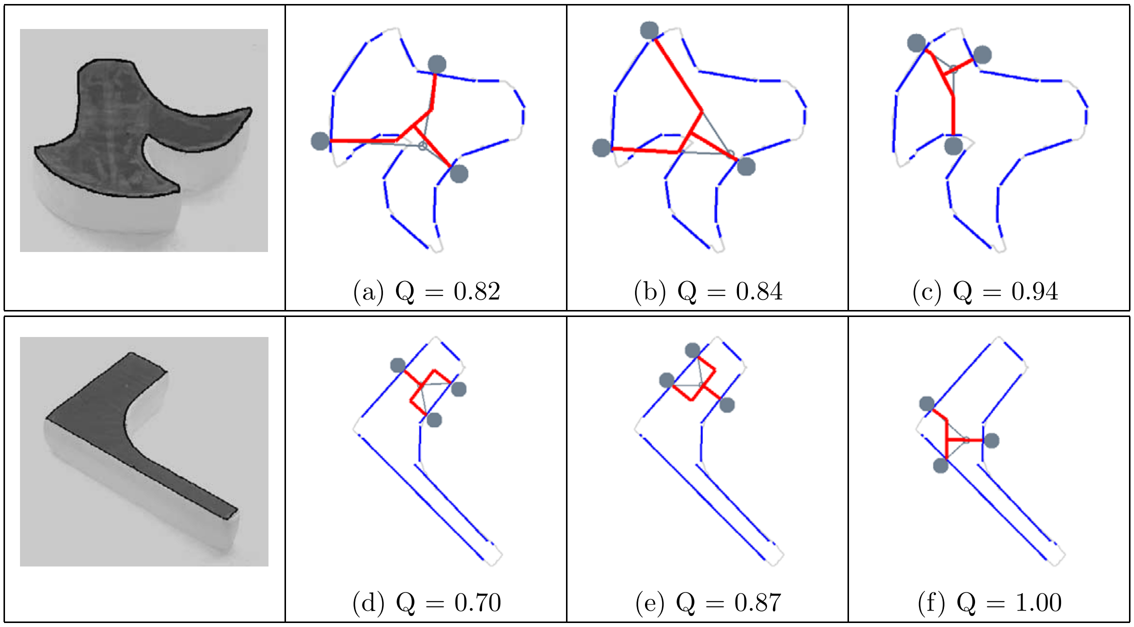

Figure 14: The three best configurations found for the objects shown in the images. The contours of the objects appear highlighted on these images.

for them. It is interesting to notice that in the configuration (c) one of the fingers lies within a concavity. Depending on the strategy implemented to reach the configurations, this could not be reached due to collision of the finger with other parts of the object. In this former case, despite the configuration scores a high quality value it is discarded by the reachability criterion, and as a consequence, it is not considered for execution.

Figure 15 shows some of the configurations found for an object similar to the shape presented in the classical work of Ponce and Faverjon (1995). In this figure we present only the quality estimated for the configurations that correspond to the best seven according to Ponce et al.; our system finds all the triplets reported in that paper. However the ranking changes as it was a,d,g, i,h. Now Only e, a, and d cases remain among the best configurations, while the quality for the rest is worse. Ponce et al. did not consider the particular kinematics of a gripper for ranking their results. With this example we show that if kinematics are taken into account the classification of the the best grasps changes dramatically.

All the configurations shown in figures 14 and 15 have been executed on the real system. An exhaustive experimental validation has been carried out. The experiments consisted of trying to lift the object from the table using the computed configuration. To do this, the location of the object was estimated using the images obtained from the stereo system, and in open-loop the humanoid moved the hand over the object and closed the hand. When a contact is detected by the force sensors situated in each finger, the object is lifted. These observations are not strict, nevertheless the results are interesting as guide for future developments.

From the configuration shown in 14, case c is unreachable. With the other two config-

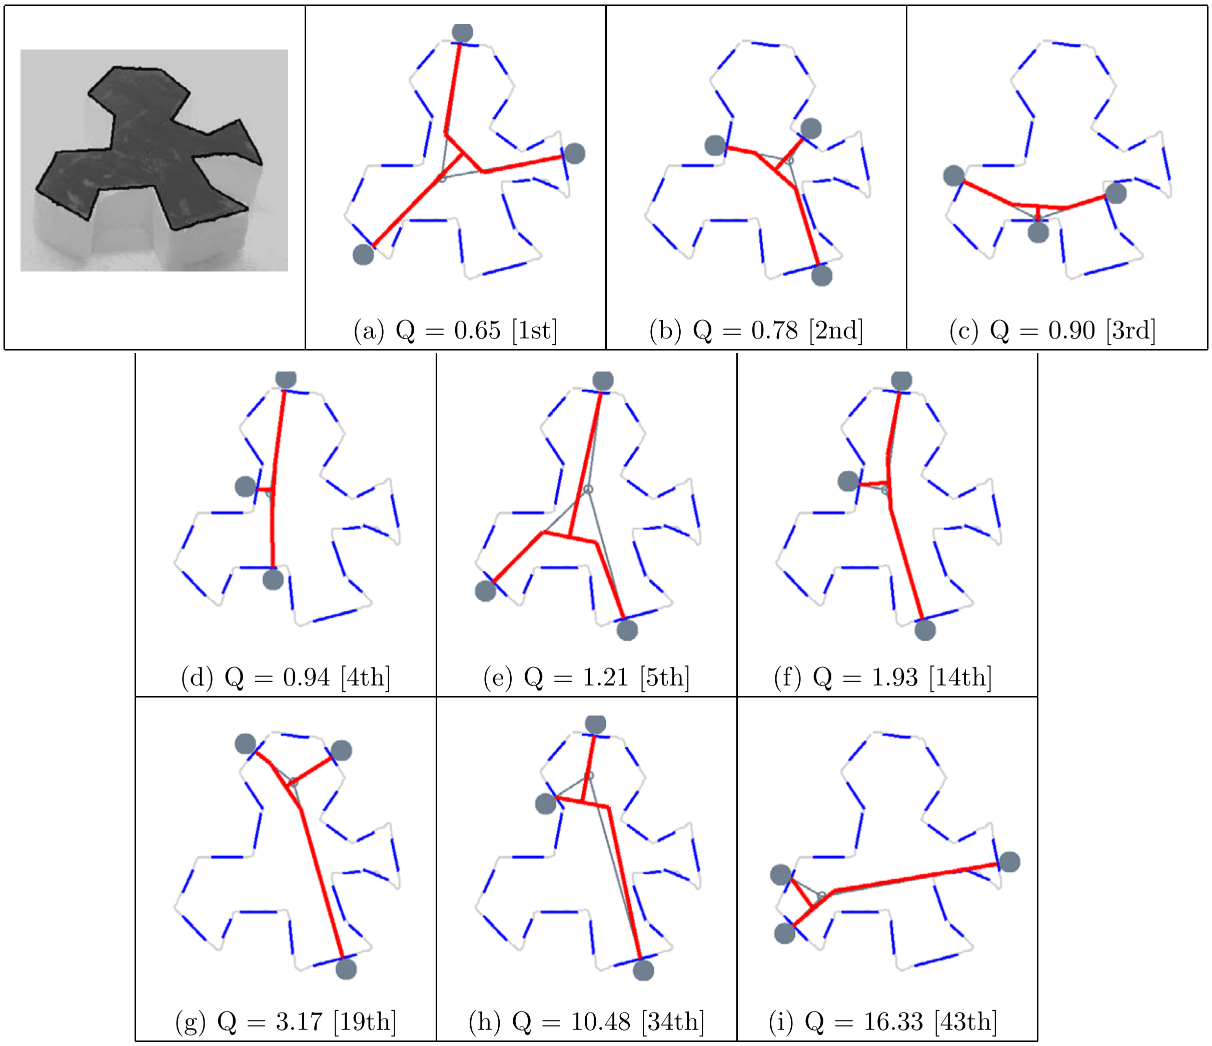

Figure 15: Nine different configurations for the same object. The object is inspired by the one that appears in Ponce and Faverjon (1995).

urations (a and b) the object is lifted without any problem. Cases d and e, were rejected because the contact points are too far from the center of gravity, these could indicate that the weighting of criterion CD is not adequate. Case f is executed perfectly.

Among the configurations presented in figure 15 cases c and h are unreachable. The rest of the configurations behave more or less as predicted by the quality measure. After several trials case b has obtained better performance on average than case a, which has the problem of extending the fingers to almost its limits. Cases d and f had problems due to the inaccuracy in positioning the hand and the fingers exactly in the computed contact locations. In these cases the fingers tend to land in unstable locations in the surroundings of the expected ones with an unpredictable, often bad, result.

Finally, in `http://www.robot.uji.es/people/morales/experiments` a sample of over sixty experiments can be found including video recordings for all of them.

7 Discussion and conclusions

The contributions of this paper are four. First, we have presented a pair of grasp synthesis algorithms for two and three-finger planar grasps. Both algorithms rely on the concept of grasp regions which has proved to be very useful since it reduces the complexity of the analysis and the search for the feasible grasps for a given object. As we have shown in subsection 6.2 both algorithms discover a large amount of feasible grasps from complex shapes with a small computation time.

Second, we have provided a simple method for constraining the candidate grasps computed by the initial algorithms to the particular kinematics of the Barrett hand. An interesting feature of this procedure is that it is able to cast two-finger grasps with a three-finger hand. This makes explicit use of the concept of *virtual finger*. One more interesting feature is that it would be possible to adapt the whole system to the use a different mechanical hand by only modifying this single procedure.

Thirst, we have suggested a set of criteria for assessing the quality of a particular configuration. These criteria can be used for selecting the best grasp or a subset of configurations that are over a given quality threshold. It is important to note that the computation of these criteria is exclusively based on information obtained from the input images. Only a few assumptions about the physical characteristics of the objects are taken into account.

Finally, the design of all the procedures has been approached from a practical point of view. This can be observed in several aspects of the design of the graping system. First, the information about the objects is obtained from images, with the problems of noise, inaccuracy and complexity in the processing that it implies. Second, the performance of the processing is critical, so it is kept as low as possible by designing fast algorithms. Finally, ans additional aspect is the constraining to the characteristics of a particular gripper. This not only affects the procedure for adapting general two or three-finger grasps, but also the criteria defined to evaluate the final configurations. The inclusion of a feasibility and a reachability criteria to the set of criteria is one more issue of this practical approach.

Several improvements can be done to the algorithms presented here, indeed Boivin et al. (2004) implement several modifications to the algorithms presented in our preliminary works (Morales et al., 2002a,b). These improvements consist of the use of corners to increase stability; an modification of the three-finger algorithms to obtain optimal solutions; and a combined use of quality criteria for optimizing hand configurations.

The final feature of the work presented here is the execution of several of the grasps found for the objects. These experiments, though not have been carried out in an exhaustive way, have illustrated several important issues. First of all, the quality rule predicts the performance of the grasps within certain range. Some problems of the grasp, caused mainly by the inaccuracy in placing the fingers in the precomputed location, are not taken into account by the criteria presented here. New criteria that focus in these issues have to be defined and introduced in the measurement. A preliminary study of this issue has already been done by Chinellato et al. (2003).

An immediate extension of this work is the inclusion of the new criteria in the evaluation of the quality. It still remains the problem of how to combine these criteria. Here, we have suggested a simple addition, but this is somewhat arbitrary. A more systematic way of building these quality function will be through more extensive experimentation. This will require the design design of an experimental protocol that accurately evaluates accurately the performance of an executed configuration. An off-line analysis of the collected data could indicate the combination of criteria that best fits the observed performances. Indeed a more ambitious approach will be the design of a procedure that incrementally *learns* this function. This procedure will select to execute, within an exploration protocol, that configuration that best increases the knowledge of the hidden function (Morales et al., 2004).

8 Acknowledgments

This work describes research partially carried out at the Robotic Intelligence Laboratory of Universitat Jaume I. Support for this research is provided in part by the Ministerio de Ciencia y Tecnología under project DPI2001-3801, by the Generalitat Valenciana under projects inf01-27, GV01-244, CTIDIA/2002/195, by the Fundació Caixa-Castelló under project P1-1B2001-28, P1-1A2003-10. The authors wish to thank to Eris Chinellato, Gabriel Recatalá and Raül Marín for their advice and help in the development of this work.

An important part of the work was carried out during a stay at the Perceptual Robotics Laboratory in the Computer Science Department of the University of Massachusetts. The fisrt author was funded during this stay by a studentship of the Spanish government under the program "Ayudas para la realización de estancias breves en el extranjero". The research in this laboratory has been funded by the grants NSF/CISE/CDA-9703217 and IRI-9704530, DARPA MARS DABT63-99-1-0004, and NASA/RICIS. The authors wish to thank Roderic Grupen, David Wheeler, Robert Platt and Danny Radhakrishnan who have provided much of the foundation and support with which the robot experiments were possible.

References

G.A Bekey, H. Liu, R. Tomovic, and W.J. Karplus. Knowledge-based control of grasping in robot hands using heuristics from human motor skills. *IEEE Transactions on Robotics and Automation*, 9(6):709–722, 1993.

A. Bendiksen and G. Hager. A viosn-based grasping system for unfamiliar planar objects. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2844–2849, San Diego, California, May 1994.

A. Bicchi. Hand for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on Robotics and Automation*, 16(6):652–662, Dec 2000.

A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *IEEE Intl. Conf. on Robotics and Automation*, April 2000.

E. Boivin, I. Sharf, and M. Doyon. Optimun grasp of planar and revolute objects with gripper geometry constraints. In *IEEE Intl. Conf. on Robotics and Automation*, pages 326–332, New Orleans, Louisina, April 2004.

Ch. Borst, M. Fischer, and G. Hirzinger. A fast and robust grasp planner for arbitrary 3d objects. In *IEEE Intl. Conf. on Robotics and Automation*, pages 1890–1896, Detroit, Michigan, May 1999.

D.L. Bowers and R. Lumia. Manipulation of unmodeled objects using intelligent grasping schemes. *IEEE Transactions on Fuzzy Systems*, 11(3):320–329, June 2003.

V. Chávtal. *Linear Programming.* W.H. Freeman, New York, 1983.

I.M. Chen and J.W. Burdick. Finding antipodal point grasps on irregularly shaped objects. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2278–2283, May 1992.

E. Chinellato, R.B Fisher, A. Morales, and A.P. del Pobil. Ranking planar grasp configurations for a three-finger hand. In *IEEE Intl. Conf. on Robotics and Automation*, Taipei, Taiwan, September 2003.

E. Chinellato, A. Morales, R.B. Fisher, and A.P. del Pobil. Visual features for characterizing robot grasp quality. *IEEE Transactions on Systems, Man and Cybernetics (Part C)*, 2004. In Press.

R. Cipolla and N. Hollinghurst. Visually guided grasping in unstructured environments. *Robotics and Autonomous Systems,*, 19:337–346, 1997.

J.A. Coelho and R.A. Grupen. Online grasp synthesis. In *IEEE Intl. Conf. on Robotics and Automation*, volume 3, pages 2137–2142, Minneapolis, Minessota, April 1996.

J.A. Coelho Jr., J. Piater, and R. Grupen. Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot. In *International Conference on Humanoid Robots*, Cambridge, Massachusetts, 2000.

B. Faverjon and J. Ponce. On computing two-finger force-closure grasps of curved 2D objects. In *IEEE Intl. Conf. on Robotics and Automation*, pages 424–429, 1991.

C. Ferrari and J. Canny. Planning optimal grasps. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2290–2295, Nice, France, May 1992.

C. Francois, K. Ikeuchi, and M. Hebert. A three-figer gripper for manipulation in unstructured environments. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2261–2264, Sacramento, California, April 1991.

K. Fukuda. FAQ in polyhedral computation. `http://www.ifor.math.ethz.ch/~fukuda/polyfaq/polyfaq.html`, Oct. 2000.

R. Grupen, M. Huber, J. A. Coelho Jr., and K. Souccar. Distributed control of manipulation tasks. *IEEE Expert, Special track on Intelligent Robotic Systems*, 10(2):9–14, April 1995.

G. Guo, W.A. Gurver, and Q. Zhang. Optimal grasps for planar multifinger robotic hands. *IEEE Transactions on Systems, Man and Cybernetics*, 22(1):193–198, 1992.

A. Hauck, J. Rüttinger, M. Sorg, and G. Färber. Visual determination of 3D grasping points on unknown objects with a binocular camera system. In *IEEE Intl. Conf. on Robotics and Automation*, pages 272–278, Detroit, Michigan, May 1999.

R. A. Jarvis. Automatic grip site detection for robotics manipulators. *Australian Computer Science Communications*, 10(1):346–356, 1988.

I. Kamon, T. Flash, and S. Edelman. Learning to grasp using visual information. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2470–2476, Minnesota, April 1996.

D. Kragic. *Visual Servoing for manipulation: robustness and integration issues*. PhD thesis, Univ. of Stockholm, 2001.

H. Liu, T. Iberall, and G.A. Bekey. The multi-dimensional quality of task requirements for dextrous robot hand control. In *IEEE Intl. Conf. on Robotics and Automation*, pages 452–457, May 1989.

Y.H Liu. Computing n-finger force-closure grasps on polygonal objects. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2734–2739, Leuven, Belgium, May 1998.

C. Mackenzie and T. Iberall. *The Grasping Hand*. North Holland, 1994.

X. Markenscoff, L. Li, and C.H. Papadimitriou. The geometry of grasping. *International Journal of Robotics Research*, 9(1):61–74, 1990.

M.T. Mason and J.K. Salisbury. Robot hands and the mechanics of manipulation. In P. H. Winston and M. Brady, editors, *The MIT Press Series in Artificial Intelligence*. MIT Press, 1985.

A.T. Miller, S. Knopp, H.I. Chirsitensen, and P.K Allen. Automatic grasp planning using shape primitives. In *IEEE Intl. Conf. on Robotics and Automation*, Taipe, Taiwan, Septmeber 2003.

B. Mirtich and J. Canny. Easily computable optimun grasps in 2D and 3D. In *IEEE Intl. Conf. on Robotics and Automation*, pages 739–747, May 1994.

B. Mishra and M. Teichmann. Three finger optimal planar grasp. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, Grenoble, France, 1994.

D.J. Montana. The condition for contact grasp stability. In *IEEE Intl. Conf. on Robotics and Automation*, pages 412–417, Sacramento, California, 1991.

D.J. Montana. Contact stability for two-finger grasps. *IEEE Transactions on Robotics and Automation*, 8(4):421–430, 1992.

A. Morales, A.P. del Pobil, and A.H. Fagg. An active learning framework for assessing the relibility of grasps by means of visual features. In preparation, 2004.

A. Morales, P.J. Sanz, and A.P. del Pobil. Vision-based computation of three-finger grasps on unknown planar objects. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 1693–1698, Lausanne, Switzerland, October 2002a.

A. Morales, P.J. Sanz, A.P del Pobil, and A.H. Fagg. An experiment in constraining vision-based finger contact selection with gripper geometry. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 1711–1716, Lausanne, Switzerland, October 2002b.

V.D. Nguyen. Constructing force-closure grasps. *International Journal of Robotics Research*, 7(3), 1988.

A.M. Okamura, N.S. Smaby, and M.R Cutkosky. An overview of dextereous manipulation. In *IEEE Intl. Conf. on Robotics and Automation*, pages 255–260, San Francisco, California, April 2000.

Y.C. Park and G.P. Starr. Grasp synthesis of polygonal objects using a three-fingered robot hand. *International Journal of Robotics Research*, 11(3):163–184, 1992.

J. Pertin-Troccaz. Grasping: A state of the art. In O. Khatib, J. Craig, and T. Lozano-Perez, editors, *The Robotics Review*, pages 71–98. MIT Press, Cambridge, Massachusetts, 1989.

R. Platt, O. Brock, A.H Fagg, D. Karupiah, M. Rosenstein, J. Coelho, M. Huber, J. Piater, D. Wheeler, and R. Grupen. A framework for humanoid control and intelligence. In *IEEE International Conference on Humanoid Robots (Humanoids 2003)*, Karlsruhe, Germany, October 2003.

R. Platt, A.H. Fagg, and R. Gruppen. Nullspace composition of control laws for grasping. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 1717–1723, Lausanne, Switzerland, 2002.

J. Ponce and B. Faverjon. On computing three-finger force-closure grasps of polygonal objects. *IEEE Transactions on Robotics and Automation*, 11(6):868–881, 1995.

K. Rao, G. Medioni, H. Liu, and G.A. Bekey. Shape description and grasping for robot hand-eye coordination. *IEEE Control Systems Magazine*, 9:22–29, 1989.

A. Rosenfled and E. Johnston. Angel detection on digital curves. *IEEE Transations on Computers*, C–22:875–878, September 1973.

P.J. Sanz, A.P. del Pobil, J.M. Iñesta, and G. Recatalá. Vision-guided grasping of unknown objects for service robots. In *IEEE Intl. Conf. on Robotics and Automation*, pages 3018–3025, Leuven, Belgium, 1998.

K.B. Shimoga. Robot grasp synthesis: A survey. *International Journal of Robotics Research*, 3(15):230–266, June 1996.

G. Smith, E. Lee, K. Goldberg, K. Böhringer, and J. Craig. Computing parallell-jaw grips. In *IEEE Intl. Conf. on Robotics and Automation*, pages 1897–1903, Detroit, Michigan, May 1999.

K. Stanley, J. Wu, A. Jerbi, and W.A. Gruver. A fast two dimensional image based grasp planner. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 266–271, Kyongju, Korea, 1999.

S.A. Stansfield. Robotic grasping of unknown objects: A knowledge-based approach. *International Journal of Robotics Research*, 10(4):314–326, Aug 1991.

M. Strandberg. A fast grasp planner for a three-fingered hand based on 2d contours. In *33rd International Symposium on Robotics*, Stockholm, Sweden, October 2002.

M. Taylor, A. Blake, and A. Cox. Visually guided grasping in 3D. In *IEEE Intl. Conf. on Robotics and Automation*, pages 761–766, San Diego, California, May 1994.

M. Teichmann and B. Mishra. Reactive robotics I: Reactive grasping with a modified gripper and multi-fingererd hands. *International Journal of Robotics Research*, 19(7):697–708, 2000.

C.P. Tung and A. C. Kak. Fast construction of force-closure grasps. *IEEE Transactions on Robotics and Automation*, 12(4):615–626, August 1996.