The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

Grasp Planning by Alignment of Pairwise Shape Descriptorss

Amer Agovic and Nikos Papanikolopoulos
University of Minnesota, Minneapolis, MN 55455

Abstract— **The majority of work related to grasp planning has centered on the understanding of what constitutes a good grasp. However, to reach a good grasp we must first find the relative position of the gripper, its approach vector, and finger configuration. This search problem is the focus of our paper. We propose an on-line method that uses pairwise shape descriptors to quickly find good alignments between the gripper contact surface and the target. Having found a good fit, we then evaluate how alignment quality relates to grasp quality and what can be done to speed up the exploration of the DOF space.**

I. INTRODUCTION

Our interest in grasp planning is motivated by the application of robotics in surgical environments. In our case the robot is required to assist the surgeon during microsurgery. Microsurgery procedures are common in the field of Otolaryngology. For the sake of precision and timing the surgeon needs to keep constant eye contact with the working area. Because of that the instruments are handled by a specially trained technician who delivers them when requested. To automate this instrument delivery process, a set of delicate grasping tasks has to be performed. Beyond the usual complications such as obstacles in the environment, this particular case also demands that grasp planning be recast as planning a grasp on an instrument already being grasped by the surgeon. In effect the grasp planning that is required is one where the target surface is only partially reachable and graspable.

We propose a method for finding correspondence between contact points on the target and the gripper. Unlike most existing approaches, we utilize shape alignment to find the correspondence between the gripper and the target in an efficient manner. We use pairwise shape descriptors invariant to rotation and translation to determine partial or full alignments between the gripper and the target. We hypothesize that good alignments (as defined later) should lead to good grasps. The end-result is an algorithm that finds grasps based on on-line evidence as opposed to using off-line precomputed results and that executes the planning process in time, when needed.

In the rest of the paper we initially review the existing work in grasping, object alignment, and shape descriptors. We then formulate the problem, describe our method in detail, and finally present an empirical evaluation of our approach.

A. Related Work

For an overview of grasp planning we refer the reader to [1] and [2]. While grasp planning overlaps with other fields such as haptics, one can define it as the search for the best way to grasp an object. This problem is complicated by the geometry of the object, modeling of the grasp contact, number of contacts, nature of external forces, and the parameterization of the gripper. The earliest attempts were based on the idea of "form closure" which tries to determine when an object becomes immovable. Form closure is mostly determined by the geometry but it requires more contacts than are absolutely necessary. If one considers friction as part of the contact model, we can find grasps with fewer contacts. Such grasps are based on force closure.

Work done in [3] and [4] is arguably the basis for many efforts in this area. It starts by considering point contacts with friction. The friction cone is approximated by a pyramid. The overal impact of external wrenches or forces is combined into a convex hull which contains the origin if we have achieved force closure. The distance of the origin to the border of such a convex hull is usually used as the quality metric. Another example of analytic grasp synthesis can be found in [5]. The authors formulate a quality metric based on the Q-distance which is differentiable and therefore allows simple gradient descent to be used in the search process. An attempt to overcome non-uniformity of the geometry is presented in [6]. Here a grasp quality measure is derived which approximates the grasp wrench space via spherical shapes that account for the worst-case disturbances. Another approach to grasp quality functions is shown in [7]. However the assumption here is that the geometry is smooth and numerically well behaved. For a more realistic contact model soft contacts have to be considered. The initial attempt at soft contacts addressed the issue of sliding [8]. Soft contact modeling demands the application of the more elaborate Hertzian analytic model [1][9]. According to this model moments and forces at a contact are coupled. The friction cone of the Coulomb model is replaced with a friction limit surface. This appears to be the primary reason why soft contact modeling has not found wider usage. More recently in [10] an approach was presented which treats limit surfaces in a similar fashion as friction cones in [3]. The limit surface is approximated by a convex polyhedron.

Even if the contacts are understood and a quality metric is formulated, there still remains the issue of gripper parameterization. The problem is one of correspondence and the question is how we position, orient, and articulate the gripper in order to maximize the grasp quality. One interesting approach has been proposed in [11][12]. The idea is to find the gripper parameters if one knows the contact points on the target which yield a stable grasp. Both sources settle on an optimization scheme to find the best parameterization. The dilemma here is whether contacts generate a gripper

978-1-4244-3804-4/09/$25.00 ©2009 IEEE

configuration or if it is the other way around. As a purely analytic method, the need for optimization means that real-time application is limited. In this paper we assume that the gripper configuration determines contacts.

The impact of geometry is considered by the authors in [13]. The shape is assumed concave and convex decomposition is used to divide and conquer the problem. Along similar lines use of shape primitives has been proposed [14]. In fact the authors pursue a more heuristic methodology. Besides using shape primitives for geometry representation, the search for grasps occurs over a predetermined set of gripper preshapes. The grasp quality evaluation happens on-line by means of a grasping simulator. Building on the heuristic method of using simple geometric preshapes the authors in [15] propose a more generic means of object representation. They use hierarchical decomposition trees in terms of quadrics. One disadvantage of the approach is the tree decomposition which is a clustering problem coupled with the actual fitting of superquadrics. Quadrics were chosen over other decompositions because they encode normals naturally. Similarly in [16] authors present work on a prosthetics system to enable human teleoperation of robotic grasping. The system achieves real-time performance by reducing the dimensionality of the gripper DOF space. While this simplifies the search, it yields suboptimal results which requires an on-line validation step to fine-tune the grasp. Instead of looking for the optimal grasp analytically, the aim of the authors was to reduce the search space while still preserving most of the good grasps.

Having considered geometry, contact modeling, grasp quality, and gripper parameterization the remaining problem is reachability. The search for a good grasp will inevitably yield multiple possible grasps [1]. In [17] the authors examine a scoring method as a way to rank the identified grasps. In particular this approach looks at how nearby clutter might affect a grasp planning algorithm. The method is based on force closure.

Alignment of $3D$ objects is a research field in its own right. The field is mainly driven by media content retrieval problems and registration of medical data [18]. The methods can be roughly split into optimization-based approaches where object geometry is directly used [19], and methods which use a variety of local and global features usually invariant to certain transformations. An interesting idea is the use of spherical harmonics [20]. For a solid review of specific pose recovery methods the reader is referred to [21]. The problem with the majority of these methods is that they try to solve the alignment problem between whole objects. Partial matching is evidently more difficult.

Shape descriptors have found wide usage in computer vision where they have lead to decent results in classification [22]. An early example of shape driven grasp planning is presented in [23]. Here we observe the use of antipodal grasps. The underlying basis is found in force closure and as such this can be considered part of the more general approaches complicated by visual sensing.

Two works that are very similar to this paper are [24] and [25]. In [24] we find a method for cluttered environments. The goal is to find high probability grasps which is a very attractive way to simplify and improve grasping. While it strives to find utility for unstructured environments, it starts out with a set of precomputed preshapes which are then adapted to the situation. This work highlights the cost of computing pose and the importance of an efficient solution to this problem. Unlike our method it uses a nonlinear optimizer combined with SVD decomposition to solve it. The work in [25] comes from the field of animation. The authors recognize and use invariant features very much like this work but the approach here is to start with a database of preshapes which are then queried by using geometric similarity. The problem of pose recovery is addressed by using three correspondences and clustering in high dimensional space as compared to our approach which uses one correspondence and 3D-space for clustering. In addition it does not examine if shape matching leads to good grasps and in which situation.

II. Problem Statement

In this section we present our problem formulation. In addition we outline conventions and assumptions which we use throughout the paper.

Given a gripper $G(\vec{\pi})$ and a target object T we seek to find the set of optimal parameterizations $\vec{\pi}$ that yield good grasps. We divide the problem into a number of subproblems. First and foremost we need to define the quality metric of a given valid grasp. By itself however the metric by which optimality of a DOF parameterization is evaluated is not enough. Equally crucial is the problem of correspondence and the closely related pose recovery given the correspondence. In this light it could be said that we find the optimal DOF parameterization by finding the pose which is valid and yields the best grasp quality.

The gripper is modeled as a triangulated mesh and parameterized by $\vec{\pi}$ whose dimension equals the number of degrees of freedom in addition to the three dimensions for wrist position and three for wrist orientation ($6D + DOF$). Furthermore the gripper is subdivided into "N" fingers with each finger having "M" links. The mesh surface of each link is partitioned into active facets and inactive facets. The active facets are commonly on the inside of the fingers and only they form contacts with the target. Please note that commonly we would expect the finger tips to be included in the contact surface. We do not do that here because we use uniform area sampling of the contact surface. We expect that good alignment will produce good grasps because of improved surface contacts not point contacts. Along these lines we observe that the active facets usually subtend convex shapes [26] also known as grasping convex (Figure 1).

The target is also modeled as a triangulated mesh. The mesh is partitioned into occluded and free facets. It is the task of the sensor to paint the facets according to this partitioning.

What constitutes a good grasp depends on the application. Frequently a stable grasp is a good grasp but a good grasp might be the one that is sufficiently stable while providing

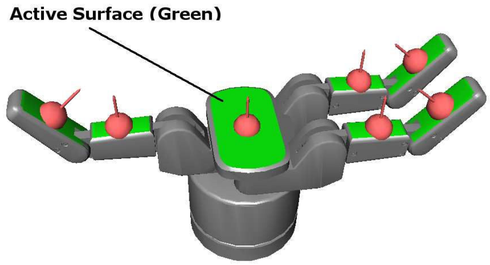

Fig. 1. Gripper Sampling

the surgeon the opportunity to also grasp the object in a stable manner. We define an objective function $Q(G(\pi), T)$ that is at maximum for an optimal value of π. Given the grasp quality function our primary goal can be reformulated as

$$\hat{\vec{\pi}} = \arg\max_{\pi} Q(G(\vec{\pi}), T). \tag{1}$$

During a grasp active facets of the gripper are in contact with facets on the target. To limit the complexity we sample a set of "P" contacts $\{p_1, \ldots, p_P\}$ evenly over the facets. The set of these P samples will be called a constellation. For each of these contacts the set of valid forces within the friction cone can be expressed as an approximation over the "R" edge vectors $\{d(p_i)_1, \ldots, d(p_i)_R\}$ of the cone:

$$f(p_i) = \sum_{j}^{R} \alpha_{ij} d_j(p_i) \tag{2}$$

$$1 \geq \sum_{j=1}^{R} \alpha_{ij}, \tag{3}$$

where $f(p_i)$ denotes a contact force at point p_i, and α_{ij} are non-negative coefficients. If the sum of all forces acting on a body is zero, that body will not move but it might rotate. We therefore consider both forces and torques. The contact wrench is defined as

$$w(p_i) = \left[f(p_i), p_i \times f(p_i)\right]^T. \tag{4}$$

The overall wrench acting on the target is then

$$w_T = \sum_{i=1}^{P} \sum_{j=1}^{R} \alpha_{ij} w_{ij} \tag{5}$$

where $w_{ij} = [d_j(p_i), p_i \times d_j(p_i)]^T$.

With the above derivation which follows [5] the target is in force closure if

$$0 \in interior(ConvexHull\,[w_{11}, \ldots, w_{PR}]). \tag{6}$$

One common quality measure is the minimal distance of the origin to the surface of the convex hull.

$$Q = \min_{\vec{w} \in \mathbb{W}} ||\vec{w}||. \tag{7}$$

We will make use of it and in addition consider only painted facets on the target as reachable. Finally in the end a ranking metric can be used to order and consolidate alignment and grasp metrics.

III. Proposed Method

A robotic assistant in microsurgery must be able to grasp objects held by a surgeon. All other issues aside, we need a grasping method that will run in real-time and is capable of planning grasps over partially covered objects.

Our proposed method is inspired by the use of pairwise features in computer vision. We observe that the search for a good grasp necessitates an efficient identification of correspondence (which finger goes where). The same problem is also present in computer vision and object alignment. Another inspiration has been the work done with the GraspIT simulator. By studying it, one realizes that their use of preshapes and heuristics to obtain gripper approach vectors is an attempt to simplify the problem of how to align the gripper to the target to obtain a stable grasp. This is where we contribute novel work. For a given DOF parameterization of the gripper we find the optimal geometric alignment between the gripper and the target. By maximizing the contact surface between the two, we should obtain a subset of poses leading to a good grasp. In order to match $3D$ object with different parameterization we utilize pairwise feature vectors. These features are invariant to translation and rotation and therefore simplify the pose estimation problem. Please notice that unlike in vision we do not need more invariance and so designing the features is easier.

Given an "N" finger gripper our method finds partial and full grasp alignments and works even if certain parts of the target object are occluded (i.e., surgeons hands). Direct geometry alignment is an expensive proposition. Partial matching makes it even more difficult. Our method lifts the representation of the target and the gripper into an invariant space. The only requirement is that the facets we match are normalized to have approximately equal size. Such unit facet will ensure that we test the target face multiple times if there is room for the finger to move.

The end result of alignment is to obtain the orientation $\vec{R}$ and the position $\vec{P}$ of the gripper wrist. Since multiple results are possible, we rank the results by means of an alignment quality metric. The quality metric we choose considers average distance within the estimated pose and the number of votes or correspondences reporting it:

$$QA = \frac{VoteCount}{1 + DistanceVariance}. \tag{8}$$

We now describe the construction of our feature vectors. As illustrated by Figure 2, we consider two unit facets sampled over the active surface of the gripper. Each facet will have a position p_i and a normal v_i. From this information we compute three quantities. The distance d_{ij} between two facet positions:

$$d_{ij} = ||\vec{p}_i - \vec{p}_j||. \tag{9}$$

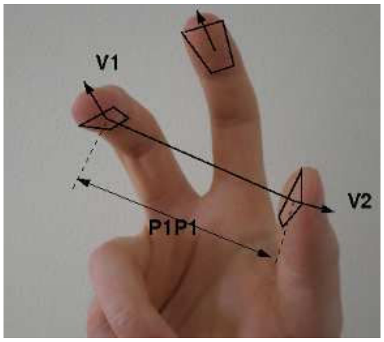

Fig. 2. Design of Features

The cosine similarity ϕ_{ij} between the normals $\vec{v}_i$ and $\vec{v}_j$:

$$\phi_{ij} = \vec{v}_i \cdot \vec{v}_j. \tag{10}$$

As the third quantity we compute the cosine similarity between the normals and the line between the points p_i and p_j:

$$\beta_{ij} = \frac{\vec{v}_i \cdot (\vec{p}_i - \vec{p}_j)}{||\vec{v}_i|| ||\vec{p}_i - \vec{p}_j||}. \tag{11}$$

We now construct our feature vectors as:

$$e_{ij}(\pi) = (d_{ij}, \phi_{ij}, \alpha_{ij} + \beta_{ji})^T. \tag{12}$$

Note we could have kept β_{ij} and β_{ji} separate which would give us a directional feature. However when they become equal we lose the directionality and have to deal with ambiguity, so we assume the ambiguity is there by design.

To find all possible combinations of “N” contacts of the gripper over the geometry of the target having “V" vertices including the partial matches, we consider both gripper and target as point sets (centered on unit facets). Then the size of the search space is given by

$$NC = \sum_{i=1}^{N} \binom{V}{i} \tag{13}$$

if we sample the gripper active surface once per finger (i.e., tip of the finger).

Our method obtains $\frac{V(V-1)}{2}$ pairs for a target object with “V" unit facets. For a gripper sampled at “P" unit facets across all fingers we obtain $\frac{P(P-1)}{2}$ pairs.

One way to do pairwise matching is all against all. A more efficient way is to build a volume hierarchy of the target using Axis Aligned Boxes (AAB). For a well balanced binary tree the size of the search space becomes:

$$C = \left(\ln \frac{V(V-1)}{2}\right) \times \frac{P(P-1)}{2}. \tag{14}$$

A. Pose Recovery

Once pairwise correspondences are found, we estimate the orientation first and then the position. This is in line with mainstream methods described in [21] where the same approach is taken. In our case finding these quantities consecutivelly is benefitial because of memory constraints even if it has stability implicaions. The main reason for developing a novel approach lies in the formulation of these methods. They all solve the problem by least square error of an objective function. The problem with such an approach is that all points are considered at once which makes them succeptible to outliers. In addition the matrix from which rotation is obtained by means of SVD or eigenvector decomposion are as big as the point set. In contrast our method works like Hough Transform. The benefits are geared towards resource limited embedded systems. Our method has much smaller memory footprint because it accumulates results, it is amenable to filtering as a way to combat outlier isses and because it does not rely on matrix operations it is significantly faster.

When we say we have correspondence it also means we have two unit vectors $\vec{A}$ and $\vec{B}$ that are related by a rotation.

$$\vec{B} = Rq(\theta)\vec{A}. \tag{15}$$

A naive approach for orientation recovery would be to pick the three vectors a correspondence gives us (two normals and a difference) and invert the rotation matrix. We can obtain two possible rotations this way because we do not have point correspondence. However this requires matrix inversion and provides no means of dealing with singular cases (i.e., where two or more vectors are co-linear). Our approach is more efficient in terms of computation and stability. We find six rotations from a single correspondence of which three should agree.

It helps to visualize what we mean with Figure 3. Any rotation can be expressed using unit quaternions or versors. A versor encodes an axis of rotation and an angle. In Figure 3 the axis of rotation goes through the south and north pole $\vec{N}$. The angle of rotation is depicted as a thick arc on the equator. One would think that to recover the rotation we could simply take the cross product of $\vec{A}$ and $\vec{B}$. The problem is that an entire set of of quaternions can rotate A into B. Any quaternion lying on the great circle that travels half way between $\vec{A}$ and $\vec{B}$ will do the job. The versor having axis $\vec{N}$ rotates our points using the smallest angle. Versor $\vec{AB}$ is the other extreme and rotates by 180 degrees. A pairwise correspondence gives us three vectors. For each vector we can register a great circle on a spherical map and then pick intersections with the most votes. The problem is that spherical to cartesian mapping is not uniform. In addition we would need two coordinates for the axis and one for the angle requiring a $3D$ accumulator and it is well known that Hough transform suffers from artifacts when the bins are too large. As an alternative we go a step further and actually find intersections between two great circles. Accounting for directional ambiguity pairwise correspondence gives us six vector pairings. Using vector pairings we can find exact versor agreements. In case of singular cases which are unconstrained we revert to using the axis with the least angle (cross product).

First we observe that the set of all versors rotating A into B can be expressed as:

$$\vec{n} = cos(t)\vec{N} + sin(t)\vec{AB}. \quad (16)$$

The intersection of two such great circles is given by:

$$\vec{m_{ij}} = (\vec{N_i} \times \vec{A_iB_i}) \times (\vec{N_j} \times \vec{A_jB_j}). \quad (17)$$

To test if two circles really intersect at that point, we must also check that the angle amounts are in agreement. For a given intersection point we calculate two angles one for each A,B pair:

$$\begin{aligned} \theta_i &= \angle(\vec{A_i} - (\vec{A_i} \cdot \vec{m}_{ij})\vec{m_{ij}}), (\vec{B_i} - (\vec{B_i} \cdot \vec{m}_{ij})\vec{m_{ij}}) \\ \cos(\theta_i) &= \frac{[\vec{m}_{ij} \times (\vec{m}_{ij} \times \vec{A_i})] \cdot [\vec{m}_{ij} \times (\vec{m}_{ij} \times \vec{B_i})]}{|(\vec{m}_{ij} \times \vec{A_i})||(\vec{m}_{ij} \times \vec{B_i})|}. \end{aligned} \quad (18)$$

Please notice that two great circle intersect at two antipodal locations and both must be considered as valid. This also impacts how we determine if two angles are in agreement since the same rotation results if angles and axis are negated. We can eliminate the axis ambiguity by choosing only the upper hemisphere and flipping the angle appropriately. After that the sign of the angle only depends on the angle between the found axis $\vec{m}_{ij}$ and the cross axis $\vec{N_i}$.

The position of wrist is obtained in a second pass from the midpoint of the correspondence (i.e. constellation points). Given an identified rotation $\vec{q} = (\theta_i, \vec{m_{ij}})^T$we rotate the midpoint forward in the gripper frame and then subtract it from its position in the target frame:

$$\vec{P} = \vec{P_T} - Rq(\theta)\vec{P_G}. \quad (19)$$

In the end we could theoretically end up with more than C possible rotations. If that happens, we most likely do not have a good alignment. If we find correspondences for x contacts, then $3[x(x-1)/2]$ rotations should agree on the same pose.

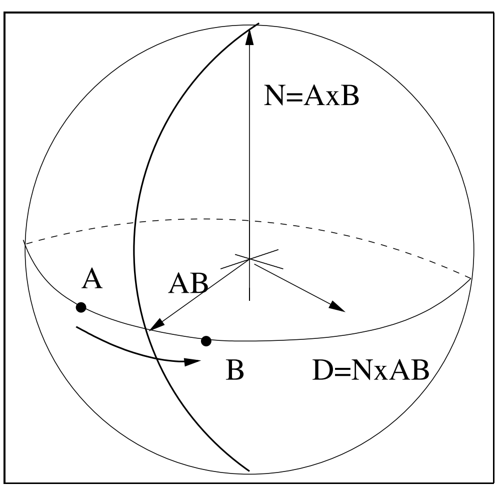

Fig. 3. Pose Recovery

```
Tesselate Target into unit faces
PT=ComputePairs(T)
BVT=ComputeBoundingHierarchy(PT)
Tesselate Gripper active surface into unit faces
foreach pi in (DOFSpace)
    PG=ComputePairs(G(pi))
    foreach gPair in (PG)
        bPairs=gPair intersect BVT
        foreach bPair in (bPairs)
            q=EstimatePose(bPair, gPair)
            p=RecoverPosition(q, bPair, gPair)
            VoteFor(q, p) in Q

RankPose(Q)
PrunePenetration(Q)
ComputeGraspQuality(Q)
```

Fig. 4. Pairwise Feature Alignment Planner

B. Practical Consideration

In Figure 4 we present the pseudocode for our algorithm that summarizes our method.

The structure of the selected pairwise feature is where a lot of the application specific engineering occurs. For example directionality can be desirable. Another desirable property is that the facets we try to align are of approximately the same shape and size (compatibility). Measures such as facet area or facet aspect ratio come to mind. For this paper we preprocess both surfaces by tessellating them into approximately equal triangles which we call unit facets.

The next practical issue is the distance metric used to detect intersection of features, intersection of orientations, and intersection of positions. For the recovery of position given an orientation we use the Euclidean distance. The intersection of features is a bit more tricky. The feature elements could have different domains such as distances and angles. We chose to use a weighted distance metric and picked the weighting parameters by manual introspection. For orientation distance measure we opted for the angle measurements. For example the rotation $(\theta_i, \vec{m_{ij}})^T$ is close to $(-\theta_i, -\vec{m_{ij}})^T$ and so is $(\pi, \vec{m_{ij}})^T$and $(-\pi, \vec{m_{ij}})^T$.

We present the algorithm as a nested loop to emphasize that it is local and sequential in nature and amenable to parallelism. In practice however we converted the loops into three passes. Pass one would accumulate correspondences, pass two would accumulate rotation agreements, and pass three would recover position given rotation.

Post-processing involves actual validation of the alignment against the grasping kinematics and dynamics. Besides the actual grasping simulation we also perform interpenetration testing. We picked a very simple feature to capture shape characteristics and it would not guard against all eventualities. Figure 5 shows an alignment that is impossible for a rigid gripper and a rigid body.

IV. Results

The experimental validation of our method was performed in octave using OpenRAVE as the back-end server. To obtain the results we run experiments by sampling evenly the four DOF dimensions and picking the highest alignment at each step. The gripper was sampled once per finger link giving

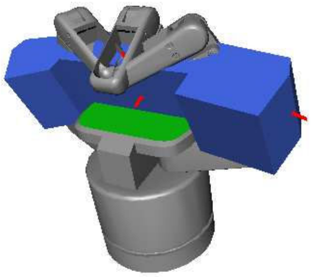

Fig. 5. Case for Penetration Test

us a seven point constellation. The DOF space was sampled evenly every 0.25 radians over all four DOF dimensions. The target was sampled differently depending on the polygon count. After collecting best alignments for each DOF step we then evaluated the grasp quality. The grasp quality evaluation is based on (7). Configurations where gripper surface penetrates the object were disqualified. The penetration test uses ray intersection and thresholding. For the rest we would perform a fine tuning step during which we slightly perturb the wrist position and DOF values (within the sampling step). During the finetuning the best grasp quality is retained. The majority of our computation occurred on the client side but collision checks and link transformation states were queried from the simulator. Because of that timings would give a skewed picture. Figure 6 shows the type of targets we used for experimentation. They are arbitrary objects but with an increasing polygonal count. We selected them because they are of approximately the same size as the Barrett Hand. The Barrett Hand was deemed a good testbed because its parametrization is simple enough to conceptualize yet it can verify the characteristics of the algorithm. We manually performed the target facet painting where it was feasible. For example the flask has a region which is unreachable at the top of the neck.

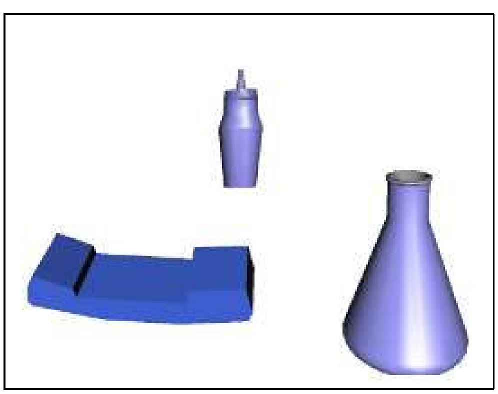

Fig. 6. Selected Targets of Varying Polygon Count

We have already discussed algorithmic complexity and here we present two results that we feel support the premise of this paper.

First in Figure 7 we show a projection of the gripper and target constellation in feature space. We have plotted side by side constellations of a number of grasp configurations (case 2-case 6) and contrast that with the constellation of the target (case 1). We used Multi Dimensional Scaling (MDS) to visualize the data as it is commonly done in machine learning. The intent of this experiment is to demonstrate that the gripper constellation has a degree of separability from the target constellation and that this degree depends to some extent on the parameterization. Simply put we have proposed here a method that quickly finds gripper to object alignment but we still have the problem of sampling the DOF space and searching it. As Figure 7 demonstrates there is separability, the search process can thus be optimized. If we can't reduce complexity, then at least we can ensure that the least amount of time is spent at each point.

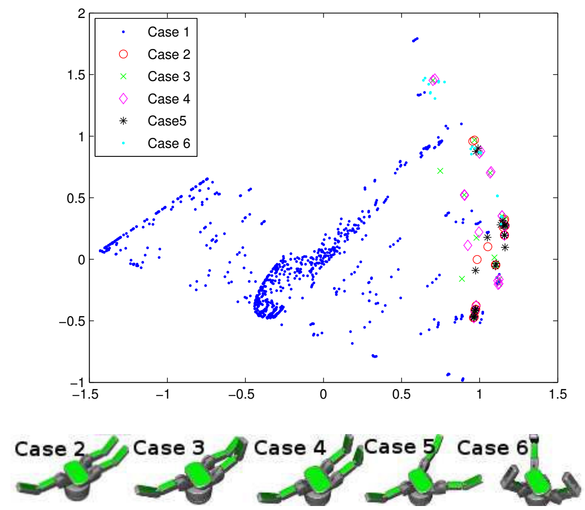

Fig. 7. MDS Projection of the Gripper and the Target in Feature Space

The second result follows in Figure 8.

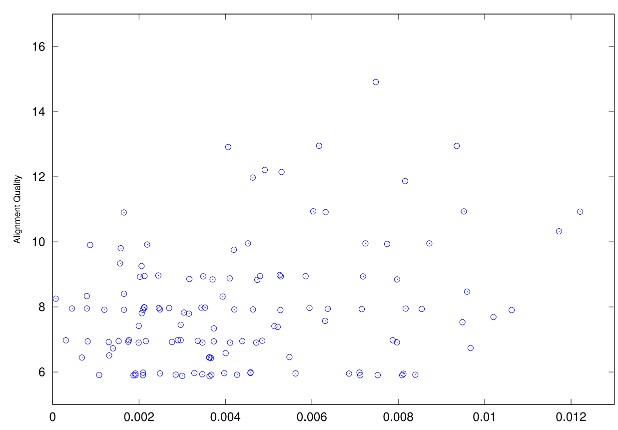

Fig. 8. Quality of Alignment Vs. Quality of Grasp

It establishes by empirical means the relationship between gripper-target alignment and the resulting grasp quality. The reasoning here is that if there exists proportionality between the two, we might be able to cast part of one problem in terms of the other as we have attempted to do here. While this transitivity might not have any benefits, it does provide one extra tool in pursuing the grasping problem. As can be seen from this figure, that relationship is not as clear cut as we initially assumed. Instead of clear proportional relation we get a more measured behavior. It appears that a good grasp need not be optimally aligned but at the same time if we have a good alignment our chances of obtaining a good grasp are better. The results seem to suggest that a good alignment could indeed be used as a heuristic in the search for a good grasp.

Lastly, Figures 9-11 present a qualitative analysis of our work. They show the top grasps found by the method described here on the three sample targets. As we can see, each grasp indeed agrees with the intuitive expectations of a good grasp. We can observe that in some cases there is some minimal penetration. This is barely visible but explains why a finetuning is necessary to adapt the proposed grasp to the situation. For the same reason Figure 10 appears suboptimal.

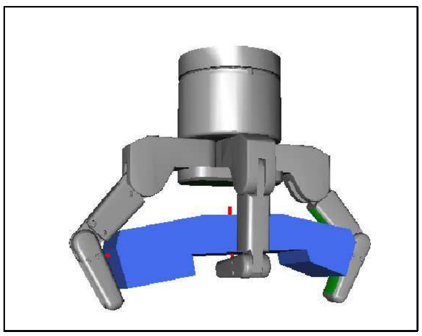

Fig. 9. An Example of an Alignment Grasp

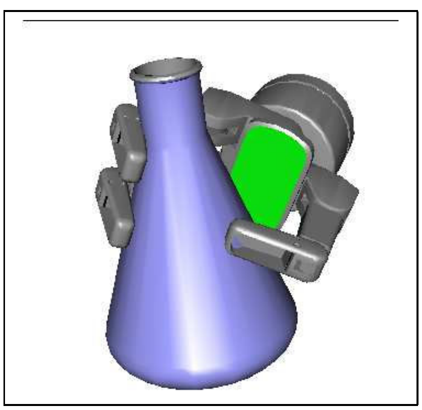

Fig. 10. An Example of an Alignment Grasp

V. Conclusions and Future Work

We have presented here a method that addresses a sub-problem of the grasp planning problem. The sub-problem arises from the high dimensional search space in which a grasp planner must operate to find the approach direction and position of the gripper wrist relative to a target. Our initial assumption was that good grasps require good alignment. While we have not tried to reduce the search space spanned

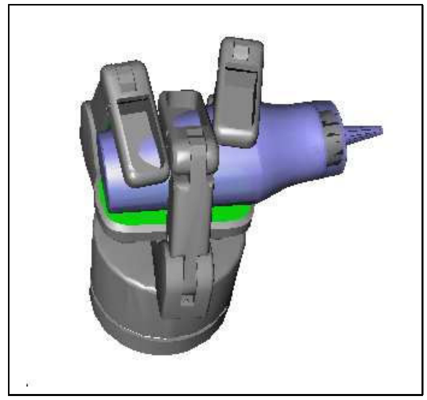

Fig. 11. An Example of an Alignment Grasp

by the Degrees of Freedom (DOF) we have explored the separability of a gripper and a target in the feature space. The idea is that absent any means of reducing the DOF space, we could construct a classifier that might quickly eliminate certain regions.

This research was performed within the specific context of a Human Robot Interface (HRI) needed in medical robotics. We believe that the majority of grippers used today including the Barrett Hand utilized here are rigid tools and as such more suitable for industrial tasks. Even if soft rubber fingers are attached, the active surface is used passively and is blind. Robots capable of human interaction should feature soft, deformable active surfaces. For one the softness prevents damage to the human subject (hand). Additionally the soft surface provides for embedding of the sensor matrix needed for the increased dexterity requirements. It also generates shape conforming surface grasps. In this light even “over-grasps” such as Figure 5 can be considered as mildly valid.

Future work will involve direct field studies in the operating room during which we intend to record how the surgeon interacts with an assistant at the haptic level (since his eyes are on the microscope). Besides collecting evidence for a touch based communication protocol, this field study will also examine how the instrument is obscured (partial coverage) during the interaction. The results here and from such field studies will allow us to complete a novel shape conforming gripper for HRI being developed in our lab.

VI. Acknowledgments

This research is supported by a grant from the Medical Devices Center at the Institute for Engineering and Medicine, University of Minnesota. Advice and input provided by Dr. Ahmed Tewfik and Dr. Samuel Levine were invaluable during conception and implementation of this study. Finally this paper would not have happened without open access to software such as GraspIT and OpenRAVE.

References

[1] A. Bicchi and V. Kumar, “Robotic grasping and contact: A review,” 2000, pp. 348–353.

[2] M. A. R. Garzon and R. S. Feijoo, “Grasp synthesis for 3D objects,” *Technical University of Catalonia*, July 2006.

[3] C. Ferrari and J. Canny, “Planning optimal grasps,” *International Confererence on Robotics and Automation*, May 1992.
[4] B. Mirtich and J. Canny, “Easily computable optimum grasps in 2-D and 3-D,” in *In IEEE International Conference on Robotics and Automation*, 1994, pp. 739–747.
[5] X. Zhu and J. Wang, “Automatic grasp planning using shape primitives,” *IEEE Transactions on Robotics and Automation*, vol. 19, August 2003.
[6] C. Borst, M. Fischer, and G. Hirzinger, “Grasp planning: How to choose a suitable task wrench space,” *Proceedings of the International Confererence on Robotics and Automation*, pp. 319–325, April 2004.
[7] G. Liu, J. Xu, X. Wang, and Z. Li, “On quality functions for grasp synthesis, fixture planning, and coordinated manipulation,” *IEEE Transactions on Automation Science and Engineering*, October 2004.
[8] R. D. Howe, I. Kao, and M. R. Cutkosky, “The sliding of robot fingers under combined torsion and shearloading,” *Proceedings of the International Confererence on Robotics and Automation*, vol. 1, pp. 103–105, April 1988.
[9] Y. Li and I. Kao, “A review of modeling of soft-contact fingers and stiffness control for dextrous manipulation in robotics,” *Proceedings of the International Confererence on Robotics and Automation*, vol. 3, pp. 3055–3060, May 2001.
[10] M. Ciocarlie, C. Lackner, and P. Allen, “Soft finger model with adaptive contact geometry for grasping and manipulation tasks,” in *WHC ’07: Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 219–224.
[11] C. Borst, M. Fischer, and G. Hirzinger, “Calculating hand configurations for precision and pinch grasps,” in *Proc. of the 2002 IEEE/RSI/GI International Conference on Intelligent Robots and Systems*. IEEE, 2002, pp. 1553–1559.
[12] C. Rosales, J. M. Porta, R. Suarez, and L. Ros, “Finding all valid hand configurations for a given precision grasp,” *Technical University of Catalonia*, May 2008.
[13] E. L. Damian, D. Sidobre, and R. Alami, “Grasp planning for non-convex objects,” *36th International Symposium on Robotics (ISR’2005)*, December 2005.
[14] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, “Automatic grasp planning using shape primitives,” *Proceedings of the International Confererence on Robotics and Automation*, vol. 2, pp. 1824–1829, September 2003.
[15] C. Goldfeder, P. K. Allen, C. Lackner, and R. Pelossof, “Grasp planning via decomposition trees,” in *2007 IEEE International Conference on Robotics and Automation*, April 2007.
[16] M. T. Ciocarlie and P. K. Allen, “On-line interactive dexterous grasping,” in *EuroHaptics ’08: Proceedings of the 6th International Conference on Haptics*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 104–113.
[17] D. Berenson, R. Diankov, K. Nishiwaki, S. Kagami, and J. Kuffner, “Grasp planning in complex scenes,” in *IEEE-RAS International Conference on Humanoid Robots (Humanoids07)*, December 2007.
[18] B. Bustos, D. Keim, D. Saupe, T. Schreck, and D. Vranic, “An experimental comparison of feature-based 3D retrieval methods,” in *3DPVT ’04: Proceedings of the 3D Data Processing, Visualization, and Transmission, 2nd International Symposium*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 215–222.
[19] M. Novotni and R. Klein, “A geometric approach to 3D object comparison,” in *SMI ’01: Proceedings of the International Conference on Shape Modeling & Applications*. Washington, DC, USA: IEEE Computer Society, 2001, p. 167.
[20] H. Huang, L. Shen, R. Zhang, F. Makedon, B. Hettleman, and J. Pearlman, “Surface alignment of 3D spherical harmonic models: Application to cardiac mri analysis,” *Medical Image Computing and Computer-Assisted Intervention*, September 2005.
[21] D. W. Eggert, A. Lorusso, and R. B. Fisher, “Estimating 3-D rigid body transformations: a comparison of four major algorithms,” *Mach. Vision Appl.*, vol. 9, no. 5-6, pp. 272–290, 1997.
[22] M. Leordeanu, M. Hebert, and R. Sukthankar, “Beyond local appearance: Category recognition from pairwise interactions of simple features,” *IEEE Conference on Computer Vision and Pattern Recognition*, June 2007.
[23] M. Taylor, A. Blake, and A. Cox, “Visually guided grasping in 3D,” *International Confererence on Robotics and Automation*, vol. 1, pp. 761–766, May 1994.
[24] D. Berenson and S. Srinivasa, “Grasp synthesis in cluttered environments for dexterous hands,” in *IEEE-RAS International Conference on Humanoid Robots (Humanoids08)*, December 2008.
[25] Y. Li and J. L. Fu, “Data-driven grasp synthesis using shape matching and task-based pruning,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 4, pp. 732–747, 2007.
[26] K. Harada, K. Kaneko, and F. Kanehiro, “Fast grasp planning for hand/arm systems based on convex model,” *IEEE International Conference on Robotics and Automation*, May 2008.