John Canny
Eric Paulos

Department of Electrical Engineering
and Computer Sciences
University of California, Berkeley
Berkeley, CA 94720
jfc@cs.berkeley.edu
paulos@cs.berkeley.edu

Optimal Probing Strategies

Abstract

Probing is a common operation employed to reduce the position uncertainty of objects. This paper demonstrates a technique for constructing provably near-optimal probing strategies for precisely localizing polygonal parts. This problem is shown to be dual to the well-studied grasping problem of computing optimal finger placements as defined by B. Mishra and others. A useful quality metric of any given probing strategy can easily be computed from simple geometric constructions in the displacement space of the polygon. The approach will always find a minimal set of probes that is guaranteed to be near optimal for constraining the position of the polygon. The size of the resulting set of probes is within $O(1)$ of the optimal number of probes and can be computed in $O(n \log^2 n)$ time, whereas the exact optimal solution is in NP-hard. The result of this work is a probing strategy useful in practice for refining part poses.

KEY WORDS—Pose refinement, probing, fine positioning, RISC robotics, optical sensing, flexible manufacturing

1. Introduction

In industrial manufacturing and automated assembly, accuracy is extremely important. Attaining and maintaining high precision can increase the cost of fixturing and feeding several-fold (Nevins and Whitney 1978). The meaning of high versus low precision depends on the application, but for typical mechanical assembly, low precision tooling might provide accuracy in the tens of mils whereas high precision would be around one mil or less (one mil $= 10^{-3}$ in. $= 25.4$ μm). This paper studies the pose refinement problem. In pose refinement, sensing is used as an inexpensive route to high-precision part pose, assuming the pose is already known at low precision. Most research to date in computer vision and reduced intricacy in sensing and control (RISC) (Canny and Goldberg 1994) sensing addresses the pose acquisition problem, where pose is determined with no knowledge of initial pose. The result of pose refinement is a high-precision estimate, but it differs from the problem of high-precision pose acquisition. Because initial pose is approximately known in pose refinement, it can be used to make judicious choices about sensor placement. The same accuracy can be achieved with fewer or less expensive sensors for pose refinement as compared to pose acquisition, which must deal with all possible poses.

Initial motivation for tackling this problem arose after visits to several state-of-the-art manufacturing companies, especially Productivity Technologies Inc., Adept, and Hewlett Packard Labs. In typical industrial work cells, it was pose refinement rather than pose estimation that was the dominant sensing task. There are two reasons for this:

1. **Feeder economics**: Vibratory feeders are an inexpensive way to provide many part types in known (albeit low-precision) pose. Small parts can also be fed on tape, which is more expensive (a couple of cents per part) but still costs far less than a high-precision pallet. So the initial and ongoing costs of achieving low-precision pose without sensing are small.

2. **Multistep manufacturing**: In typical manufacturing, there are not one but several sequential stages, including assembly stages, testing, and packaging. A single step might mate two parts whose poses are known at high precision. But the assembly step itself introduces a small amount of uncertainty, and it is expensive to transport the partial assembly at high precision to the next assembly stage. A more economical solution is to use pose refinement at the next stage. So although there might be one pose acquisition step per part to get an approximate initial pose, there will be several pose refinement steps for that part that start with the low precision output from the previous step and feeder, and increase the precision as needed for the next step.

The International Journal of Robotics Research
Vol. 20, No. 8, August 2001, pp. 694-704,

The arguments for pose refinement are that it (1) replaces the most expensive (high-precision) fixturing and feeding steps and (2) replaces the most frequent fixturing and feeding steps in multistep assembly.

This paper focuses primarily on the use of simple light-beam sensors that act as line probes in three dimensions, or point probes in the object's projection onto a horizontal surface. A point light source and receiver define a line in space that is broken and unbroken by an object as it moves relative to the beam (see Fig. 1). The positions of the object when the beam breaks give position readings, and three or more of these determine pose. Those readings are subject to error, and the pose estimate accuracy is limited by those errors and the sensor placement. This research provides algorithms for choosing the probe placements to achieve near-optimal accuracy with a fixed number of probes, or to find a near-minimal number of probes to achieve a specified accuracy. An important assumption of this work is that individual probing operations do not disturb or alter the part's pose in any way. Common industrial optical sensors easily satisfy this constraint by being contactless. The problem is best summarized with the following problem statement.

PROBLEM STATEMENT

Given. A polygonal part geometry and an initial pose estimate of $O_i = (O_{ix}, O_{iy}, O_{i\theta})$ within $\sigma = \|O - O_i\|_m$ of the exact actual pose of the object, O. Here, $\|\cdot\|_m$ is defined to be the m-norm.

Assumption. A probing operation leaves a part's pose unchanged.

Solve. Find the optimal set of point probes defined as the minimal number of probes and their placement necessary to reduce the uncertainty in the position of an object to better than some acceptable level. The probes are defined by a set of fixed points and vectors denoting the direction of travel for each probe. A probe returns a real value. That value is the time or position that a simple binary sensor changes state. The error associated with this value is at most ϵ, where $\epsilon \ll \sigma$, based on the presence or absence of an object's edge at a particular point along the path of the probe.

With enough time, one could simply perform a large number of probes as shown in Figure 2. However, in real industrial robotic assembly work cell design, throughput is a heavily weighted criterion. Therefore, the goal is to produce the best possible probing strategies that conform to the imposed constraints. The probing strategies that result can be used by any point probe of the object's two-dimensional projection. There are natural generalizations to higher dimensions, although they are not as efficient. A typical probe, shown in Figure 1, consists of a simple reflective light beam sensor that can easily detect the presence or absence of an object. The algorithm also allows for construction of specialized optimal probing strategies such as those for a scanning array of probes as shown in Figure 3.

These near-optimal probings are obtained with a small number of actual probes by maximizing the utility of each sensor probe placed on the object. This in turn makes the problem tractable for a real robot in a high-throughput automation system. These strategies are within a constant factor of the optimal probing strategy and can be solved in $O(n \log^2 n)$ time, whereas the exact optimal solution is in NP-hard (Das and Joseph 1990).

2. Previous and Related Work

We first discuss related work in part probing. Our main result is to show that probing is dual to the grasping problem, and to give an optimal probing algorithm based on set coverings. We then discuss related work on grasping and on set-covering algorithms.

2.1. Work in Probing

The importance of probing in terms of localizing and identifying objects with probes has been explored by several individuals. Cole and Yap (1987) and Bernstein (1986) developed algorithms for choosing probes to obtain the geometry of an unknown two-dimensional convex object. A generalization of this strategy for higher dimensions was presented by Dobkin, Edelsbrunner, and Yap (1986), whereas a nonconvex version was developed by Boissonnat and Yvinec (1992). Also, Lindenbaum and Bruckstein (1990) described similar probing strategies for a geometric probe composed of two line probes rotating about a common axis point.

The development of efficient algorithms for scanning objects with probes for the purpose of identification and localization was studied by Wallack, Canny, and Manocha (1993) and Wallack and Canny (1994) . Likewise, point-probing strategies were developed for insertion operations by Paulos and Canny (1994).

Jia and Erdmann (1995) demonstrated an elegant technique for choosing placements of simple binary sensors to discriminate objects in the plane. In fact, they also employed recent work on hitting sets and set coverings in solving their problem. The work in this paper differs mainly in the type of problem that is solved. Jia and Erdmann chose fixed probes to discriminate individual object poses from a large set of possible poses. The problem tackled in this paper is how to best choose moving probes to refine the pose of a known object.

2.2. Work in Grasping

The need for good grasp-planning algorithms for arbitrary shapes has always been important for robotics and industrial automation. The problem of optimal finger placement was

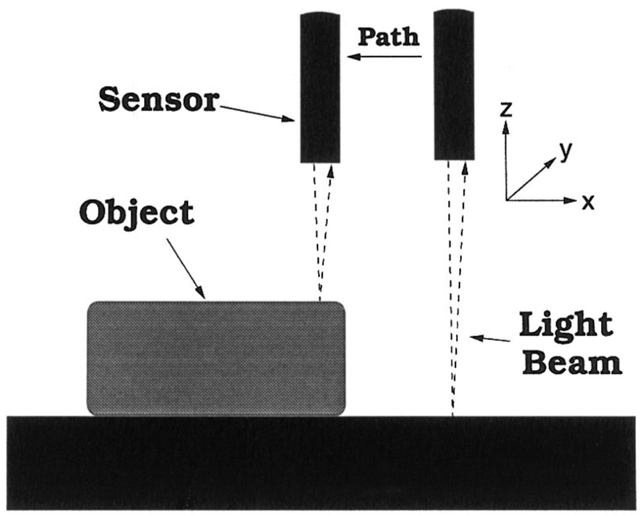

Fig. 1. A typical simple reflective sensor used for probing.

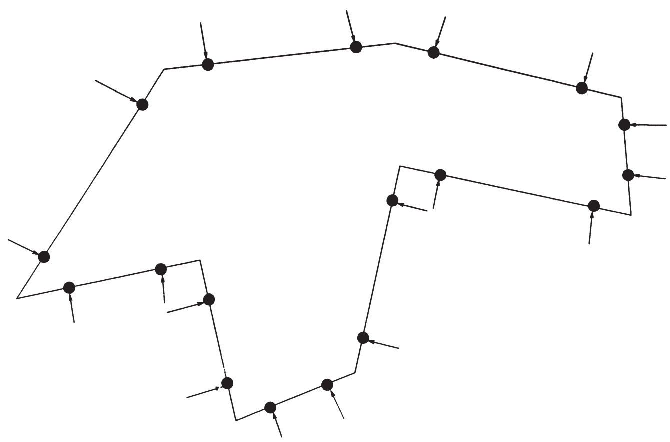

Fig. 2. Typical initial probe placement along the edges of an object.

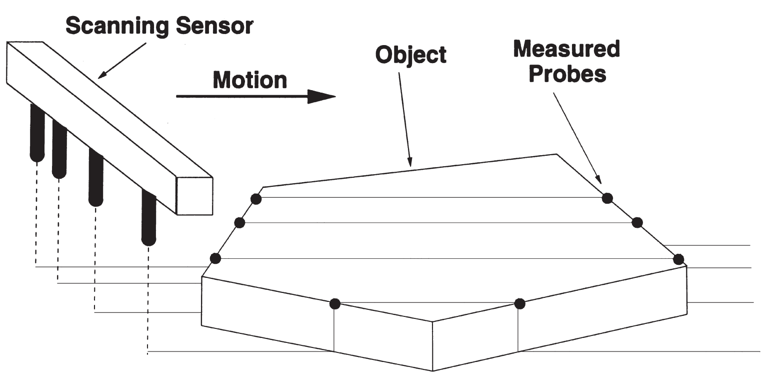

Fig. 3. Example of a fixed-array scanning beam sensor.

addressed by Mishra, Schwartz, and Sharir (1987), who defined easily computable quality metrics for grasps. Markenscoff and Papadimitriou (1989) chose to optimize the grasp with respect to minimizing the forces needed to balance the object's weight through friction. Ponce and Faverjon (1991) fixed the number of fingers and solved a system of linear constraints in the positions of the fingers to optimally position them along the polygonal edges. A similar technique for three-dimensional polyhedral objects was developed by Ponce et al. (1991). Goldberg (1993) also detailed a method for choosing grasps with a parallel jaw gripper when the initial pose of the object is unknown. Other optimal grasping techniques based on simple geometric constructions were developed by Brost (1988) and, later, Mirtich, Canny, and Ferrari (1992, 1994).

2.3. Work in Set Coverings

This paper will prove that finding the minimal set of probes is equivalent to solving the convex-set-covering problem. This problem was discussed by Clarkson (1993), who described a $O(cn \log^{O(1)} n)$ time-randomized algorithm for finding covering sets of cardinality within $O(\log c)$ of the optimal set covering c.

More recent results by Brönnimann and Goodrich (1994) on the dual problem of finding minimal hitting sets improve on these bounds. They demonstrate an $O(n \log^2 n)$ algorithm that finds a hitting set of size $O(1)$ from the optimal set size. They employed work by Matoušek (1990) using ϵ-nets.

3. RISC Robotics

RISC robotics (Canny and Goldberg 1994) is an attempt to fuse automation and robotics technologies. The RISC acronym, borrowed from computer architecture, suggests the parallels between the two technologies. RISC robotics performs complex manufacturing operations by composing simple elements. A synonymous phrase to describe this theme is simply "minimalist robotics."

RISC robotics can be applied to many areas of manufacturing. For example, RISC grasping uses simple two- and three-fingered grippers with traditional fixturing devices such as clamps and vices (Wallack and Canny 1994). RISC sensing employs simple but precise sensor elements that can be combined to form complete systems for localizing and recognizing arbitrary objects from a library (Wallack, Canny, and Manocha 1993; Wallack and Canny 1993).

RISC robotics systems inherently consist of few degrees of freedom and low-dimensional sensor spaces. This results in algorithms for manipulation and sensing that are simple, highly accurate, and very fast.

4. Defining Optimality

When probing an object, the objective is to choose point probes that allow the minimum variation of the object pose. Point probes inherently contain some known error, so it is not enough to take k independent measurements to constrain k degrees of freedom. The placement of the probes affects the worst-case object displacement. Therefore, it is the relationship between the object displacements and the corresponding probe displacements that is of interest. The goal is to find a set of probe placements that minimizes the potential worst-case object displacement.

4.1. Object Pose Definition

We define O as the actual pose of the object in two dimensions where

$$O = (O_x, O_y, O_\theta).$$

Our aim in this work is not to locate an object but, rather, to refine the position of a known object whose pose is known to some reasonable degree of accuracy. Our approach relies on this initial coarse accuracy pose information. We define the assumed initial pose as O_i and quantify a bound on the worse-case displacement of the assumed pose from the actual pose as

$$\|O_i - O\|_m \leq \sigma,$$

where $\|\cdot\|_m$ is defined to be the m-norm. At this point, we run into the usual problem of defining a metric on a space with distance and angular coordinates. There may be application-specific ways to weight the angular component, but a good default is to weight the angular component by the object's radius (i.e., the largest distance from any point in the object to its coordinate origin). With this choice, the metric bounds the maximum distance between any two corresponding points on the object at O and O_i. A typical value for σ would be tens of mils. Finally, although we are considering an m-norm for generality, the 2-norm would seem to be the most natural choice.

We will be using point probes to refine the position of the object. Therefore, for a given set of probe measurements, there will also be a set of valid poses for the object consistent with those sensor readings. We denote this object pose as $\bar{O}$ and define it to be an object pose chosen by an adversary consistent with some sensor readings given the object is at O. We define the difference between the actual object position and the adversary's choice as o

$$o = O - \bar{O}.$$

Recall that we are attempting to refine the position of the object so that σ will always be at least an order of magnitude larger than $\|o\|$

$$\sigma \gg \|o\|.$$

In terms of linear displacement, the initial pose uncertainty, σ, is typically on the order of tens of mils whereas the uncertainty under probing is on the order of one mil or less.

To summarize, we have O as the actual pose of the object, O_i as our initial pose estimate, and $\bar{O}$ as a pose that an adversary can choose that does not violate our sensor readings. That is, we cannot determine from the sensors whether the object is at $\bar{O}$. We will clarify later exactly how $\bar{O}$ is defined.

4.2. Probe Placement

We construct probes along the perimeter of the object and denote them as

$$p = (p_x, p_y) \qquad\qquad l = (l_x, l_y),$$

where p is the point where the probe touches the object when the object is at O_i and l is the direction of motion of the probe. We must guarantee that no matter where the object actually is, this probe always contacts the same edge. Assuming that the object radius was used to weight the angular component of the pose metric, this can be accomplished in a simple manner: construct a strip about the probe path l whose boundaries are parallel to l, and at distance σ from it; this strip represents the possible relative positions of the probe for various actual object poses O. If the edge we are probing crosses the entire strip, it will always be probed correctly. If the edge crosses only part of the strip, then there is a possible O such that this probe misses the edge completely. From now on, we assume that all probes are chosen so that their σ-strips touch only the edge of interest.

The initial probe placement consists of placing a pair of probes on each edge. Each probe is placed as near as possible to an endpoint of the edge, but subject to the strip constraint above. As we shall see, there is no loss of generality with this step because the probes at the edge endpoints always provide the most constraining measurements. Our algorithm will choose a subset of these initial probes as the near- optimal probe set. Our initial choice is based on the fact that we receive the most accurate pose information by probing near the vertices of an object. Observe that probes near the vertices give rise to large sensor displacements as a result of small rotational perturbations, whereas position information is the same anywhere on the edge. Figure 4 demonstrates how moving a set of probes with a given error out toward the vertices of an edge shrinks the size of allowable displacements for that edge.

This set of probes is guaranteed to contain the optimal probe placement. Any edge-interior probes would provide only redundant information in the worst case, and our probe choice is based on a worst-case analysis. A typical initial probe placement example is shown in Figure 2. The remaining problem is to determine a subset of these probes that still provide a substantial gain in object pose accuracy.

4.3. Probing Function

We place a coordinate system at the center of mass (COM) of the object. In addition, we define the rotational displacement of the object to be about this COM axis. In Figure 5, we depict the construction of the corresponding probe displacement for a given object displacement. This will define the probing function. In Figure 5 , n is a unit normal to the edge being probed, p is the initial probe location, and p' is the probe location after the displacement O from the origin.

Recall that σ is very small, allowing us to take small angle approximations and write

$$p'_k \approx p_k + (O_x, O_y) + p_k^{\perp} O_\theta,$$

where $p^{\perp} = (x, y)^{\perp} = (-y, x)$ and k denotes the kth probe. It follows that the change in probe position is

$$\begin{aligned} \Delta p_k &= p'_k - p_k \\ &= (O_x, O_y) + p_k^{\perp} O_\theta. \end{aligned}$$

The probe only gives us useful position information normal to the edge being probed. We could freely displace the object along the edge without changing that probe reading. Therefore, the change in probe position along the edge normal n_k can be written as $n_k \cdot \Delta p_k$. Observe that even if we approach an edge at an angle, when we detect the edge, we can only claim that some point of the edge must intersect the detected point. This is equivalent to the information we receive if we approach normal to the edge. Therefore, the two probe approach techniques are equivalent, and thus the choice of edge approach is independent and left as a final implementation detail. It *does* affect the σ-strip described earlier, and the amount of clearance from the edge endpoint needed to ensure the correct edge is detected.

We are now ready to define the probing function $P : \Re^3 \rightarrow \Re^k$ to be a real valued function that maps object positions into ideal probe outputs of the form $(P_1, P_2, \ldots, P_k)$. We define each element to be

$$P_k(O) = \Delta p_k \cdot n_k. \tag{1}$$

Our probes will have a sensor error ϵ, typically one mil or so. We define the measured probes as $\bar{P} \in \Re^k$. Given a sensor error of ϵ, we observe that the measured probe values $\bar{P}$ must be consistent with the ideal probes given object pose O:

$$\|\bar{P} - P(O)\|_\infty \leq \epsilon. \tag{2}$$

Similarly, any possible object position $\bar{O}$ that the adversary chooses must have all measured probes within ϵ of the given measurements:

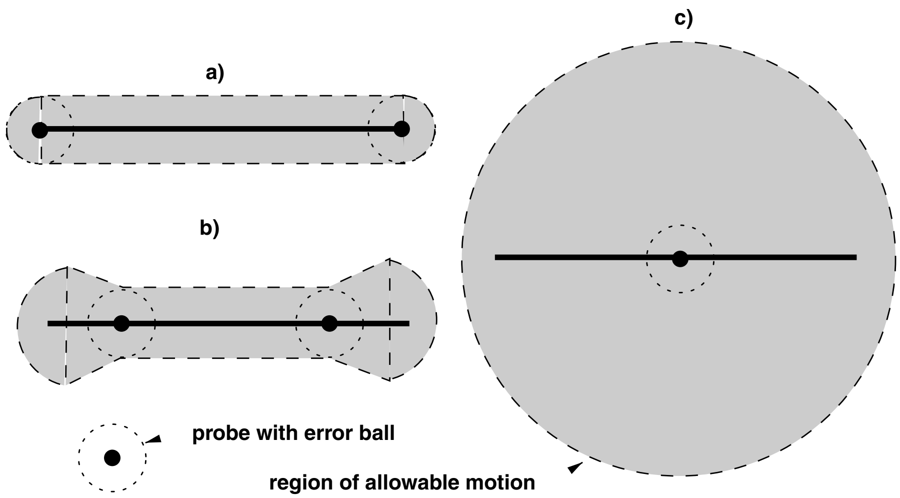

Fig. 4. Relationship of placement of probes (solid points) with error balls (dashed circles) along an edge (solid line) and the resulting displacement region (shaded): (a) probes at endpoints of the edge, (b) probes moved inward from the edge resulting in a larger region of allowable motion, (c) probes coincident at the center of the edge, allowing a maximum displacement region.

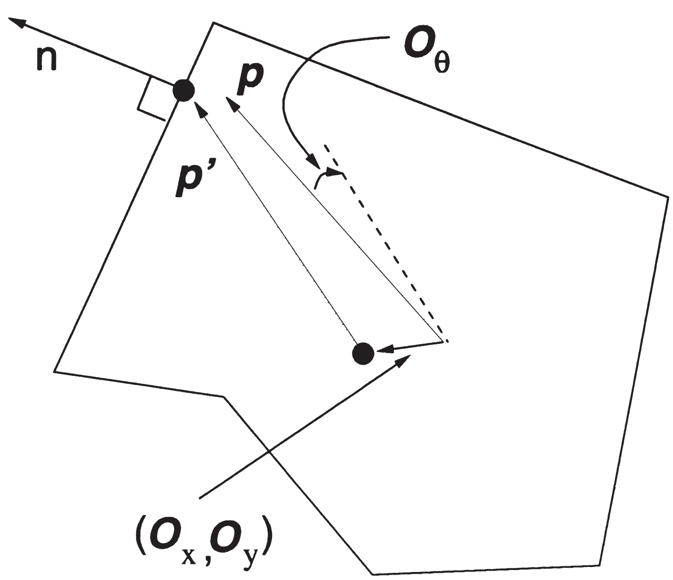

Fig. 5. Original probe p and resulting probe p' after an object displacement (O_x, O_y, O_θ).

$$\|P(\bar{O}) - \bar{P}\|_\infty \leq \epsilon. \tag{3}$$

Using the triangle inequality on these last two expressions, we find that the O and $\bar{O}$ satisfy

$$\|P(O) - P(\bar{O})\|_\infty \leq 2\epsilon \tag{4}$$

and observe that for any $\bar{O}$ satisfying this inequality, there is a $\bar{P}$ satisfying eqs. (2) and (3). Thus, the combined bound is tight. We will call the set of object displacements $\bar{O}$ that satisfy this inequality $\mathcal{K}$.

Recall that the actual position of the object is defined as O and the possible interpreted object position for some sensor reading is $\bar{O}$, where $\bar{O}$ is any $\bar{O}$ satisfying eq. (4). We want to constrain the distance between the interpreted object position and the actual object position to be as small as possible. This in turn minimizes the worst-case distance between the actual and measured poses, which is the ultimate goal of pose refinement. We represent the former quantity as

$$\|O - \bar{O}\|_m. \tag{5}$$

We employ an adversarial argument and note that if an adversary is allowed to move the object to some valid $\bar{O}$ consistent with the sensor readings, it will always choose the $\bar{O}$ such that the quantity in eq. (5) is maximized. We express this as

$$\sup_{\bar{O} \in \mathcal{K}} \|O - \bar{O}\|_m. \tag{6}$$

However, we are allowed to choose the set of probes P. Furthermore, we desire a set of probes that will output drastically different values for different nearby object poses, thus allowing us to identify different poses easily. Essentially, we would like to eliminate the possibilities of obtaining identical or near-identical sensor readings for an object in two different poses. We can write this as

$$\max_P \|P(O) - P(\bar{O})\|_\infty, \tag{7}$$

or, because $P(O)$ is linear in O, we can make the substitution to

$$\max_P \|P(O - \bar{O})\|_\infty \tag{8}$$

or, rewriting,

$$\min_P \frac{1}{\|P(O - \bar{O})\|_\infty}. \tag{9}$$

Equation (6) scales linearly with $O - \bar{O}$, whereas eq. (9) scales as the reciprocal of $O - \bar{O}$. It is natural to combine them as a product that is then independent of the magnitude of $O - \bar{O}$:

$$\min_P \left(\max_{\bar{O} \in \mathcal{K}} \frac{\|O - \bar{O}\|_m}{\|P(O - \bar{O})\|_\infty} \right). \tag{10}$$

From this, we can arrive at our final optimality criterion and probe quality measurement Q:

$$Q(P) = \min_P \left(\max_{\bar{O} \in \mathcal{K}} \frac{\|O - \bar{O}\|_m}{\|P(O - \bar{O})\|_\infty} \right). \tag{11}$$

5. Displacement Space

Working in displacement space, we observe that there is a simple geometric construction of the optimality criterion as given in eq. (11). Displacement space, denoted $\mathbb{D} \in \Re^3$, is the space of all displacements in (x, y, θ) of the object O to be probed. Each probe sensor that we introduce imposes constraints on the allowable set of displacements of the object without violating the probe value.

Equation (4) from the previous section defines a pair of half-spaces in displacement space $\mathbb{D}$ for each probe p_k:

$$\begin{aligned} \|P_k(O) - P_k(\bar{O})\|_\infty &\leq 2\epsilon \\ \|P_k(O - \bar{O})\|_\infty &\leq 2\epsilon \\ \|P_k(o)\|_\infty &\leq 2\epsilon \\ \|n_x o_x + n_y o_y + (p^\perp \cdot n) o_\theta\|_\infty &\leq 2\epsilon, \end{aligned}$$

where $o = O - \bar{O}$. These two half-spaces can be written as

$$n_x o_x + n_y o_y + (p^\perp \cdot n) o_\theta - 2\epsilon \leq 0 \tag{12}$$

$$n_x o_x + n_y o_y + (p^\perp \cdot n) o_\theta + 2\epsilon \geq 0. \tag{13}$$

The intersection of all $2k$ half-spaces constructed from k probes by definition represents a convex polytope in $\mathbb{D}$. We name this polytope S, with the definition

$$S = \cap_{h \in \mathcal{H}(P)} h,$$

where $\mathcal{H}(P)$ is the family of $2k$ half-spaces defined by the set of k probes P.

In displacement space, this polytope S will have the farthest outlying point, which will occur in the nondegenerate case at a vertex of S. This farthest outlying point represents the largest object displacement from the assumed pose that still satisfies the given probe measurements. More formally, we define this point as

$$\Gamma(S) = \sup_{q \in S} \|q\|_m.$$

The distance to $\Gamma(S)$ is exactly the optimality criteria as defined in eq. (11). To see this, assume that the denominator of eq. (11) has been fixed to some constant λ. Because the denominator scales with $O - \bar{O}$, we can always do this. The constraint $\|P(O) - P(\bar{O})\|_\infty = \lambda$ defines a polytope in displacement space (choice of O), and if we set $\lambda = 2\epsilon$, it defines exactly the polytope S. With its denominator constrained, maximizing (11) means maximizing its numerator $\|O - \bar{O}\|$, which is exactly what is specified by $\Gamma(S)$.

We assume that P is fixed, so there is only optimization by the adversary over q. Recall that an adversary can choose the actual sensor readings $\bar{P}$ such that the object displacement $\Gamma(S)$ is a valid interpretation of $\bar{P}$. Hence, this is the largest displacement of the object undetectable by the given probing strategy.

For illustrative purposes, we work through a simple example without rotation. In Figure 6, we show three probes on a triangle, an admittedly simple case, but enough to demonstrate our method. Note that each probe in real space gives rise to a pair of parallel half-spaces in displacement space, $\mathbb{D}$. If we remove probe p_2, the area of polygon S in displacement space increases, which represents the additional translational freedom that the object can undergo and still remain consistent with the remaining two sensor readings. Therefore, the added sensor p_2 is a useful addition because it decreases the area of S and reduces the distance to $\Gamma(S)$.

We are interested in probing strategies P' that have approximately the same quality metric P. Remember that every probe we remove from P removes a pair of half-spaces in $\mathbb{D}$. This in turn changes the shape of S but not always the point $\Gamma(S)$, which defines the optimality criteria. Therefore, we would like to find other optimal probes with fewer probes. In particular, we would like to find

$$\min_{P' \subset P} |P'| : Q(P') \approx Q(P),$$

where $|\cdot|$ is simply the cardinality of the set P'.

We define S' to be the polytope defined by the intersection of the half-spaces defined by the probes P'. Observe that

$$S' \supseteq S,$$

which implies that when we remove a probe, hence two half-spaces, we expect the farthest outlier to remain where it is or increase in distance from the origin, giving

$$|\Gamma(S')| \geq |\Gamma(S)|.$$

Rather than remove half-planes in an ad hoc manner such that $Q(S')$ remains essentially unchanged, we will dualize and solve for a minimal convex set covering for the corresponding points in the dual. This minimal set of points will be exactly dual to the minimal set of half-spaces in displacement space by definition of the minimal convex-set-covering problem. These resulting half-spaces in $\mathbb{D}$ correspond to a minimal set of probes, as desired. We discuss this dualization in the next section.

Observe that the production of any such probing strategy is independent of the error, 2ϵ. This is true because we are interested in optimizing the ratio shown in eq. (11). One can also note that topology of the polytope S of the solution space is independent of ϵ, which serves only as a scaling factor. That is, when we double ϵ, we get the same polytope at twice the linear size (eight times the volume). Therefore, without loss of generality, we set ϵ to one for the duration of the paper.

6. Displacement Space Dual

A strong relationship to grasping is shown in this section. We show that finding the optimal k probe placements is equivalent to finding the optimal push-pull grasp for a set of k fingers. A push-pull grasp is defined as a grasp that employs fingers capable of exerting a pushing or pulling force at the contact.

We define $\mathbb{D}^D$ to be the dual of $\mathbb{D}$. We define the dual exactly in Table 1. In this mapping, we show how points in $\mathbb{D}$ map to planes in $\mathbb{D}^D$ and how planes in $\mathbb{D}$ map to points in $\mathbb{D}^D$. We note that by definition, the dual of $\mathbb{D}^D$ is $\mathbb{D}$; hence, the duality operation is symmetric.

Observe that a polytope S defined as the intersection of a set of half-spaces h_k becomes the polytope S^D. We define Bound(h_k) to be the plane on the boundary of the half-space h_k. The polytope S^D can also be expressed as the convex hull of the union of dual points Bound$(h_k)^D$

$$S^D = \text{Conv}(\{\text{Bound}(h_i)^D \mid i = 1 \ldots, k\}).$$

Let $r \in S$. The distance of r from the origin in $\mathbb{D}$ is simply

$$|r| = \sqrt{r_x^2 + r_y^2 + r_\theta^2}.$$

The dual plane r^D in $\mathbb{D}^D$ by definition is represented as

$$r_x x + r_y y + r_\theta \theta = 1.$$

This distance of the closest point on this plane to the origin in $\mathbb{D}^D$ is given by

$$|r^D| = \frac{1}{\sqrt{r_x^2 + r_y^2 + r_\theta^2}}.$$

Setting $\alpha = \sqrt{r_x^2 + r_y^2 + r_\theta^2}$, we discover that the distance of this point r from the origin in $\mathbb{D}$ is α and the minimal distance of the dual plane r^D from the origin in $\mathbb{D}^D$ is $\frac{1}{\alpha}$. Therefore,

$$|r| = \frac{1}{|r^D|}.$$

Let f_c be the plane closest to the origin of $\mathbb{D}^D$ not intersecting Int(S^D). The distance to f_c is the same as the distance to the closest point u_c in the boundary of S^D (which is contained in f_c). And it is easy to see that $\Gamma(S)^D = f_c$, where $\Gamma(S)$ was defined earlier as the farthest outlying point in S.

The closest point to the origin in the boundary of a polytope lies on the largest inscribed sphere centered at the origin. Observe that $\Gamma(S)$ lies on the smallest circumscribing sphere of S in $\mathbb{D}$. Therefore, finding the smallest circumscribing sphere Σ for a polytope S is equivalent to finding the largest inscribed sphere Σ^D of the dual polytope S^D. This follows from the relationship

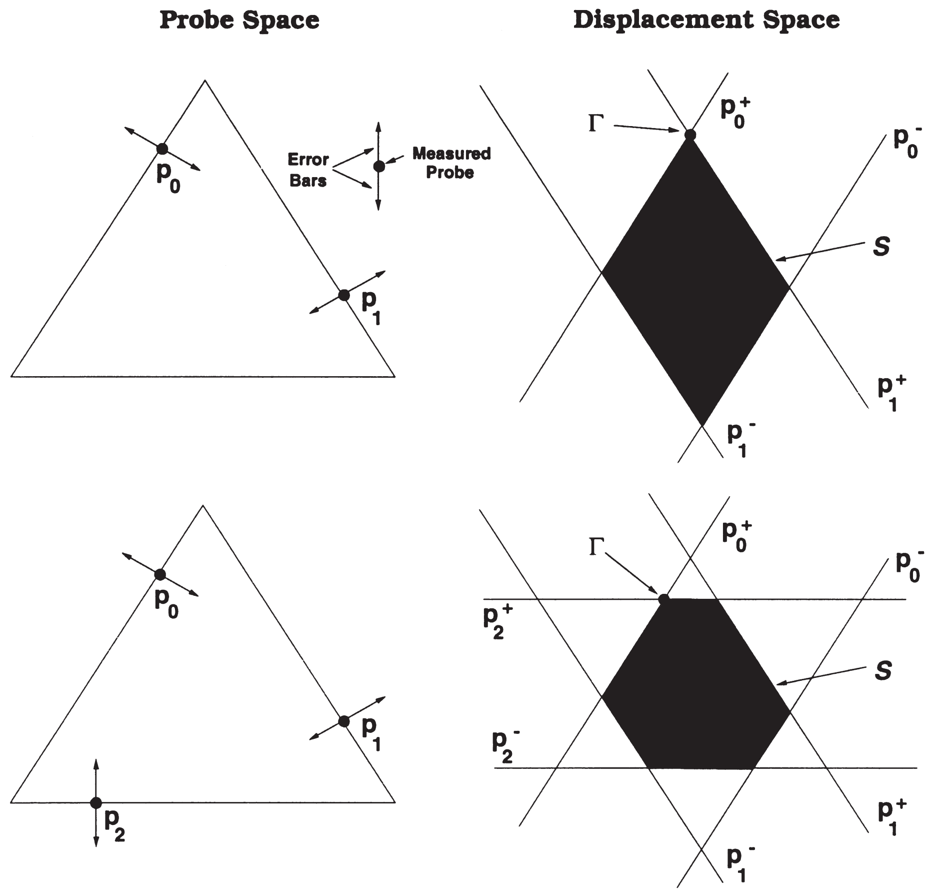

Fig. 6. Simple example of probe space and displacement space without considering rotation. Note the removal of one probe and the resulting change of the polygon S in the displacement space.

Table 1. Duality Mappings

$\mathbb{D}$		$\mathbb{D}^D$
$p : (p_x, p_y, p_\theta)$	$\leftrightarrow$	$p^D : p_x x + p_y y + p_\theta \theta = 1$
$f : ax + by + c\theta = 1$	$\leftrightarrow$	$f^D : (a, b, c)$
$S =$ polytope	$\leftrightarrow$	$S^D = \{f^D : f \cap \text{Int}(S) = \emptyset\}$

$$\mathrm{rad}(\Sigma) = \frac{1}{\mathrm{rad}(\Sigma^D)},$$

where $\mathrm{rad}(\Sigma)$ is the radius of the sphere Σ.

The planes $\mathrm{Bound}(h)$ through the half-spaces in $\mathbb{D}$ dualize to the points

$$\mathrm{Bound}(h_k)^D = (n_x, n_y, p_k^\perp).$$

These points are equivalent to the wrenches due to unit pull finger forces acting at p. In the probing problem, we obtain a pair of half-spaces for each probe. Hence, the optimal probing problem is equivalent to the optimal push-pull grasping problem. We use the optimal grasping criteria as defined by Mishra, Schwartz, and Sharir (1987) and others (Ferrari and Canny 1992; Mirtich and Canny 1994), which are the set of finger placements such that the 1-norm of the finger forces can resist the largest externally applied wrench on the object. We also define the optimal probe placement as the minimal number of probes and their placement necessary to reduce the uncertainty in the position of an object to better than some acceptable level. Using these metric definitions, we obtain the following result.

THEOREM 1. Finding the optimal placement of k probes is equivalent to finding the optimal push-pull frictionless grasp for a set of k fingers.

7. Hitting Sets and Set Covers

Recall from the previous section that the quality of the probing strategy is given directly by the radius of the maximally inscribed sphere in S^D. We would like to remove some vertices of S^D such that the radius of the maximally inscribed sphere does not decrease by much.

This problem can be posed as a convex-set-covering problem. The set-covering problem is stated for arbitrary sets L and U in $\Re^d$. The problem is to find $C \subseteq U$ with $L \subseteq \mathrm{Conv}(U)$. Here, L is the sphere of desired radius. This problem has been studied by several individuals. Recently, Clarkson (1993) described a randomized algorithm for computing the three-dimensional convex point set covering from an initial set of n points to within $O(\log c)$ of the optimal covering of c points. Clarkson's algorithm had a running time of $O(cn \log^{O(1)} n)$.

Recent work by Brönnimann and Goodrich (1994) improved on both the running time and approximation to the optimal convex set covering. Brönnimann and Goodrich's deterministic algorithm solved the equivalent problem of finding a minimal hitting set, where a hitting set is a subset $H \subseteq X$ such that H has a nonempty intersection with every set R in a collection of subsets of X. The algorithm employs work by Matoušek (1990) on ϵ-nets to obtain a hitting set in $O(n \log^2 n)$ time that is within $O(1)$ of the optimal size hitting set. This set corresponds exactly to the optimal probe placement, which we define as a set of c probes that reduce the uncertainty in the position of an object to at least some necessary level for the operation to be performed.

In our optimal construction, we obtain pairs of half-spaces, hence, pairs of points in the dual. However, in the Brönnimann and Goodrich (1994) algorithm, they are treated as two completely unrelated elements. This will result in near-optimal set sizes that are in the worst case twice as large as we could achieve by grouping the pairs. Alternatively, we can group the pairs to obtain the near-optimal hitting set at a slight running time cost. This performance slowdown is a result of an increase in the VC dimension (Vapnik and Červonenkis 1971) as a result of our pairing.

The VC dimension (Vapnik and Červonenkis 1971) is defined for a range space (X, R), with $P \subseteq X$ as the cardinality of the largest set P that is shattered by R. A set P is shattered by R if $\Pi_R(P)$ is the power set of P, where $\Pi_R(P)$ denotes the set of all intersections of P with sets in R.

To obtain an optimal probing strategy for an array of scanning sensors as shown in Figure 3, we identify the colinear points in the displacement space and assign them labels such that the hitting set algorithm will include all or none of a set of colinear points in the probe optimization selection. This also results in an increase in the VC dimension that affects the running time but still finds a hitting set within $O(1)$ of the optimal one.

Our algorithm successfully handles other variations similar to the colinear constraint for the scanning sensor without major modification. This makes it well adapted to situations in which optimal probing strategies under special constraints are needed and not intuitive to observe.

The theorem below summarizes much of the results of this paper.

THEOREM 2. A near-optimal set of c point probes can be found for any polygonal object in $O(n \log^2 n)$ time. Furthermore, the size of the set c will be within $O(1)$ of the size of the optimal set of c point probes.

8. Conclusion

This paper has demonstrated a method by which optimal probe placements can quickly be obtained for any known polygonal object. More important, the solutions it generates are guaranteed to be within a constant of the actual optimal number of probes necessary. These probing strategies refine the position of an object whose pose is approximately known. Furthermore, it is this pose refinement problem that is a real and frequently encountered challenge in industrial manufacturing. The constraint of requiring probes that leave a part's pose unchanged after each probing operation is easily satisfied by employing optical contactless sensors commonly found in industry. This paper also shows that the problem of optimal probe placement is dual to the well-studied push-pull grasping problem of positioning frictionless fingers on an object.

References

Bernstein, H. 1986. Determining the shape of a convex n-sided polygon by using $2nk+k$ tactile probes. *Information Processing Letters*, pp. 225–260.

Boissonnat, J., and Yvinec, M. 1992. Probing a scene of nonconvex polyhedra. *Algorithmica*, pp. 321–342.

Brönnimann, H., and Goodrich, M. 1994. Almost optimal set covers in finite VC-dimension. *Proceedings of the 10th ACM Symposium on Computational Geometry (SCG)*, pp. 293–302.

Brost, R. C. 1988. Automatic grasp planning in the presence of uncertainty. *International Journal of Robotics Research* 7(1):3–17.

Canny, J., and Goldberg, K. 1994. RISC for industrial robotics: Recent results and open problems. *IEEE International Conference on Robotics and Automation*, Vol. 3, pp. 1951–1958.

Clarkson, K. L. 1993. Algorithms for polytope covering and approximation. *Proceedings of the 3rd Workshop on Algorithms and Data Structure*, Vol. 709, pp. 246–252.

Cole, R., and Yap, C. 1987. Shape from probing. *Journal of Algorithms* 8:19–38.

Das, G., and Joseph, D. 1990. The complexity of minimum nested polyhedra. *Canadian Conference on Computational Geometry.*

Dobkin, D., Edelsbrunner, H., and Yap, C. 1986. Probing convex polytopes. *18th Annual ACM Symposium on Theory of Computing*, pp. 424–432.

Ferrari, C., and Canny, J. 1992. Planning optimal grasps. *IEEE International Conference on Robotics and Automation*, pp. 2290–2295.

Goldberg, K. 1993. Orienting polygonal parts without sensors. *Algorithmica* 10:201–225.

Jia, Y.-B., and Erdmann, M. 1995. The complexity of sensing by point sampling. *Algorithmic Foundations of Robotics*, pp. 283–298.

Lindenbaum, M., and Bruckstein, A. 1990. Reconstructing a convex polygon from binary perspective projections. *Pattern Recognition*, pp. 1343–1350.

Markenscoff, X., and Papadimitriou, C. H. 1989. Optimum grip of a polygon. *International Journal of Robotics Research* 8(2):17–29.

Matoušek, J. 1990. Construction of ϵ-nets. *Discrete Computational Geometry* 5:427–448.

Mirtich, B., and Canny, J. 1994. Easily computable optimum grasps in 2D and 3D. *IEEE International Conference on Robotics and Automation.*

Mishra, B., Schwartz, J. T., and Sharir, M. 1987. On the existence and synthesis of multifingered positive grips. *Algorithmica* 2:541–558.

Nevins, J., and Whitney, D. 1978. Computer controlled assembly. *Scientific America* 238(2):62–74.

Paulos, E., and Canny, J. 1994. Accurate insertion strategies using simple optical sensors. *IEEE International Conference on Robotics and Automation*, pp. 1656–1662.

Ponce, J., and Faverjon, B. 1991. On computing three-finger force closure grasps of polyhedral objects. *International Conference on Advanced Robotics*, pp. 1018–1023.

Vapnik, V. N., and Červonenkis, A. Y. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probability Applications* 16:264–280.

Wallack, A., and Canny, J. 1993. A geometric matching algorithm for beam scanning. *SPIE International Society for Optical Engineering*, Vol. 2060, Vision Geometry II, pp. 143–160.

Wallack, A., and Canny, J. 1994. Planning for modular and hybrid fixtures. *IEEE International Conference on Robotics and Automation.*

Wallack, A., Canny, J., and Manocha, D. 1993. Object localization using crossbeam sensing. *IEEE International Conference on Robotics and Automation*, pp. 692–699.