Appeared in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation.*

Sensing Polygon Poses by Inscription*

Yan-Bin Jia   Michael Erdmann
The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Abstract

Industrial assembly involves sensing the pose (orientation and position) of a part. Efficient and reliable sensing strategies can be developed for an assembly task if the shape of the part is known in advance. In this paper we investigate the problem of determining the pose of a convex n-gon from a set of m supporting cones, i.e., cones with both sides supporting the polygon. An algorithm with running time $O(nm)$ which almost always reduces to $O(n + m \log n)$ is presented to solve for all possible poses of the polygon. As a consequence, the *polygon inscription* problem of finding all possible poses for a convex n-gon inscribed in another convex m-gon, can be solved within the same asymptotic time bound. We prove that the number of possible poses cannot exceed $6n$, given $m \geq 2$ supporting cones with distinct vertices. Experiments demonstrate that two supporting cones are sufficient to determine the real pose of the n-gon in most cases.

Our results imply that sensing in practice can be carried out by obtaining viewing angles of a planar part at multiple exterior sites in the plane. As a conclusion, we generalize this and other sensing methods into a scheme named *sensing by inscription.*

1 Introduction

In this paper we will study the problem of detecting the pose, i.e., the orientation and position, of a convex polygon in the plane by taking views of the polygon from multiple exterior sites. The shape of the polygon is assumed to be *known* in advance, but the pose of the polygon can be arbitrary. Each view results in a *cone* formed by the two outermost occluding rays starting from the viewing site; that cone conversely imposes a constraint on the possible poses of the polygon—it has to be contained in the cone and make contact with both its sides. A containment like this in which every edge of the containing object touches the contained object is called an *inscription*, so we shall say that the polygon is *inscribed* in the cone. Such constraints imposed by individual views together allow only a small number of possible poses of the polygon, which often reduces to one. For example, Figure 1 illustrates two views taken of

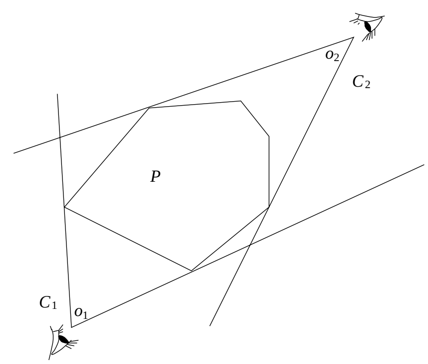

Figure 1: Sensing the pose of a polygon by taking two views.

a convex 6-gon P in some unknown pose from sites o_1 and o_2 respectively: The two cones C_1 and C_2 thus formed determine the real pose of P, and this pose can be solved using the algorithm presented later in Subsection 2.4.

The above sensing approach appears to be simple, but to make it efficient and to minimize the cost of sensing hardware we would like to have as few views taken as possible. This leads us to the main question of this paper: *How many views are sufficient in the general case in order to determine the pose of a convex n-gon?*[1]

*Support for this research was provided in part by Carnegie Mellon University, and in part by the National Science Foundation through the following grants: NSF Research Initiation Award IRI-9010686, NSF Presidential Young Investigator award IRI-9157643, and NSF Grant IRI-9213993.

[1] It should be noted that there exist cases in which the pose of a convex n-gon cannot be uniquely determined no matter how many views are taken. This happens only if the polygon preserves self-congruence over certain rotations. (It is not hard to see that in such a case the rotation angle must be a multiple of $\frac{2\pi}{k}$, where k is a positive integer such that $k|n$.) However, all congruent poses are usually considered as the same in real applications.

The answer to the above question is *two*, and to argue this answer we will go through several steps each of which occupies a separate section: Section 2 describes how to compute the set of possible poses for a convex polygon inscribed in multiple cones, and derives an upper bound on the number of possible poses for two-cone inscription in particular; Section 3 empirically demonstrates that, in spite of the upper bound, two cones have turned out to be sufficient in most cases to uniquely determine the pose of an inscribed polygon; Section 4 further discusses the extensions of this method and proposes a general sensing scheme—sensing by inscription.

1.1 Related Work

Canny and Goldberg [4] have introduced a Reduced Intricacy in Sensing and Control (RISC) paradigm that aims at improving the accuracy, speed and robustness of sensing by coupling simple and specialized hardware with fast, task-oriented geometric algorithms.

The *cross-beam sensing* method developed in [15] finds the orientation of a polygon (or polyhedron) by measuring its diameters along three different directions and comparing the measurements with the precomputed diameter function [9]; then it solves a vertex-line correspondence problem for the position of the polygon by least squares fitting. This method essentially determines the pose by inscribing the polygon in a hexagon constructed from the sensory data.

For the special case that the poses are finite, [14] presented an efficient method of placing a *registration mark* on the object so that the pose can be recognized by locating the mark position (with a simple vision system). For robustness to sensor imperfections, the marked point maximizes the distance between the nearest pair among its possible locations.

Model-based recognition and localization can often be regarded as a *constraint satisfaction* problem which searches for a consistent matching between sensory data (2-D) and model(s) (3-D) based on the geometric constraints between them. (See [10].) A variety of polygon shape descriptors [1], [12] have been analyzed theoretically and/or demonstrated experimentally to be efficient and robust to uncertainties.

Geometric algorithms for sensing unknown poses as well as unknown shapes have also been studied. Cole and Yap [6] considered "finger" probing a convex n-gon (n unknown) along directed lines, and gave a procedure guaranteed to determine the n-gon with $3n$ probes. This work was later extended in [7] which investigates the complexities of various models for probing convex polytopes in d-dimensional Euclidean space.

The polygon containment problem, that is, deciding whether an n-gon P can fit into an m-gon Q under translations and/or rotations, has been studied by various researchers in computational geometry. (See [2], [3], [5], [8].) In the case where Q is convex, the best known algorithm runs in time $O(m^2 n)$ when both translations and rotations are allowed [5]. Here we will deal with a special case of containment in which each edge of Q must touch P; this constraint causes a reduction of the running time to $O(mn)$, or $O(n + m \log n)$ in practical situations.

2 Multi-Cone Inscription

To simplify the presentation, let us *agree* throughout this section that all angles take values in the half-open interval $[0, 2\pi)$. In accordance with this agreement, intervals for angle values whose left endpoints are greater than right endpoints are allowed; for example, an interval $[\alpha, \beta]$, where $0 \le \beta < \alpha < 2\pi$, is understood as the interval union $[\alpha, 2\pi) \cup [0, \beta]$.

2.1 Sliding a Triangle in a Cone

We first deal with the case of a triangle in a cone, not only because it is the simplest but also because the case of a polygon, as we will see later, can be decomposed into subcases of triangles. Let $\triangle p_0 p_1 p_2$ be a triangle inscribed in an upright cone C with angle ϕ and vertex o, where $0 < \phi < \pi$, making contacts with both sides of the cone at vertices p_1 and p_2 respectively. What is the locus of vertex p_0 as edge $p_1 p_2$ slides against the two sides of the cone?

Two different situations can occur with this inscription: (a) p_0 is outside $\triangle p_1 o p_2$, and (b) p_0 is inside (only when $\angle p_1 p_0 p_2 \ge \phi$). (See Figure 2.) Assume that $\triangle p_0 p_1 p_2$ may

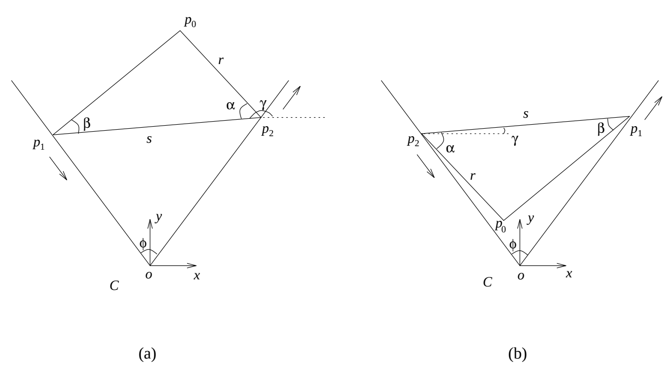

Figure 2: A triangle sliding in a cone. Vertices p_1 and p_2 move along the two sides of cone C. The locus of p_0 is an elliptic curve (possibly degenerating into a line segment) parameterized by the angle γ between directed edge $\overrightarrow{p_2 p_1}$ and the x axis. There are two different cases: (a) p_0 is above edge $p_1 p_2$; (b) p_0 is below edge $p_1 p_2$.

degenerate into any one of its edges but not a point; writing

$\alpha = \angle p_0p_2p_1$, $\beta = \angle p_0p_1p_2$, $r = |p_0p_2|$ and $s = |p_1p_2|$, this degeneracy is taken into account by the following constraints:

$$0 \le \alpha, \beta < \pi, \quad s \ge 0, \quad \text{and} \quad \begin{cases} r > 0, & \text{if } s = 0 \\ & \text{or } \alpha > 0; \\ 0 \le r \le s, & \text{otherwise.} \end{cases}$$

Let us set up a coordinate system with the origin at o and the y axis bisecting angle ϕ, as shown in Figure 2. Then the orientation of $\triangle p_0p_1p_2$ can be denoted by the angle γ between vector $\overrightarrow{p_2p_1}$ and the x axis. Note the range of valid γ values is a closed interval. For any valid γ, there exists a unique pose of the triangle in cone C; this allows us to parameterize the locus (x, y) of p_0 by γ.

As edge p_1p_2 slides in the cone, p_0 traces out an *elliptic* curve $\mathcal{C}$ with implicit equation

$$ax^2 \pm bxy + cy^2 = d,$$

where

$$\begin{aligned} a &= r^2 - rs\frac{\sin(\frac{\phi}{2} \mp \alpha)}{\sin\frac{\phi}{2}} + \frac{s^2}{2 - 2\cos\phi}; \\ b &= \frac{(2r\cos\alpha - s)s}{\sin\phi}; \\ c &= r^2 - rs\frac{\cos(\frac{\phi}{2} \mp \alpha)}{\cos\frac{\phi}{2}} + \frac{s^2}{2 + 2\cos\phi}; \\ d &= \left(r(r - s\frac{\sin(\phi \mp \alpha)}{\sin\phi})\right)^2. \end{aligned}$$

Here the notation '$\pm$' means '$+$' in case (a) and '$-$' in case (b) and the notation '$\mp$' means just the opposite. Furthermore, if the orientation γ changes monotonically within its range, p_0 moves *monotonically* along $\mathcal{C}$ except when $\mathcal{C}$ degenerates into a line segment.[2] If that degeneracy happens, p_0 may cross the same point twice.

2.2 One-Cone Inscription

Now consider the case that a convex n-gon P with vertices $p_0, p_1, \ldots, p_{n-1}$ in counterclockwise order is inscribed in a cone C. Let us choose the same coordinate system as used in the previous subsection. Then the pose of P is uniquely determined by the location of some vertex, say p_0, and the angle θ between the x axis and some directed edge, say $\overrightarrow{p_0p_1}$. Clearly any orientation θ gives rise to a *unique* pose of P; so we can compute the locus of p_0 as a function of θ over $[0, 2\pi)$. Let p_l and p_r be the two vertices currently incident on the left and right sides of cone C respectively. (See Figure 3.) As long as p_l and p_r remain incident on these two sides respectively, the problem reduces to the case of $\triangle p_0p_lp_r$ sliding in cone C except that the locus of p_0 (an elliptic curve) now needs to be parameterized by θ, instead of γ which we had used before.

[2] From the curve equation, the reader should have no difficulty verifying that $\mathcal{C}$ becomes a segment on a line through the cone vertex o when $r = 0$ or $\frac{r}{s} = \frac{\sin(\phi \mp \alpha)}{\sin\phi}$.

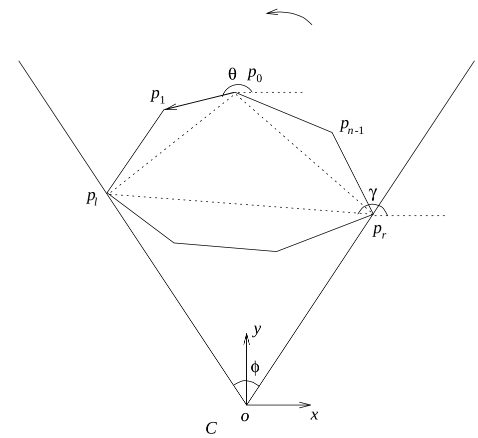

Figure 3: A convex polygon P rotating in a cone. The pose of P is denoted by the position of vertex p_0 and the angle θ between directed edge $\overrightarrow{p_0p_1}$ and the x axis. The space of orientations $[0, 2\pi)$ is partitioned into closed intervals, each defining an elliptic curve that describes the corresponding locus of p_0.

More observations show that the entire range $[0, 2\pi)$ of orientations can be partitioned into a sequence of closed intervals, within each of which the vertices p_l and p_r incident on cone C are invariant.

We present a linear-time algorithm that computes the above orientation intervals as well as the corresponding elliptic curves describing the locus of p_0. The algorithm rotates the polygon counterclockwise in the cone, generating a new interval whenever one (or both) of the incident vertices p_l and p_r changes; the new incident vertex (or vertices) is determined by a comparison between angle ϕ and the angle intersected by the two rays $p_{l-1}p_l$ and p_rp_{r-1}.

Let $[\theta_{\min}, \theta_{\max}]$ denote the current interval, and let φ_i denote the interior angle $\angle p_{i-1}p_ip_{i+1}$ for $0 \le i \le n-1$. (For convenience, arithmetic operations performed on the subscripts of vertices or internal angles are regarded as followed by a '$\bmod n$' operation; for example, p_{-1} is identified with p_{n-1} and p_n with p_0.) In the algorithm, Φ_{left} and Φ_{right} keep track of the angle between $\overrightarrow{p_lp_{l+1}}$ and $\overrightarrow{p_0p_1}$ and the angle between $\overrightarrow{p_rp_{r+1}}$ and $\overrightarrow{p_0p_1}$ respectively. The algorithm proceeds as follows:

Algorithm 1

Step 1 Start at the pose such that edge p_0p_1 aligns with the right side of C. Locate the vertex p_i in contact with the left side of C. Set $\Phi_{\text{left}} \leftarrow \sum_{j=l+1}^{n}(\pi - \varphi_j)$, $\Phi_{\text{right}} \leftarrow 0$, $\theta_{\min} \leftarrow \frac{\pi}{2} - \frac{\phi}{2}$, $l \leftarrow i$, and $r \leftarrow 0$.

Step 2 The current orientation interval has its left endpoint at $\theta_{\min}$. Output the elliptic curve (now parameterized by θ) resulting from sliding edge $p_l p_r$ in cone C.

Next determine the right endpoint $\theta_{\max}$ of the current interval. Let ψ be the angle intersected by rays $p_{l-1}p_l$ and $p_r p_{r-1}$; set $\psi \leftarrow nil$ if they do not intersect. There are three different cases:

Case 1 $\psi < \phi$ or $\psi = nil$. [Advance p_r clockwise to the next vertex.] Set $\Phi_{\text{right}} \leftarrow \Phi_{\text{right}} + \pi - \varphi_r$, $\theta_{\max} \leftarrow \Phi_{\text{right}} + \frac{\pi}{2} - \frac{\phi}{2}$, and $r \leftarrow r - 1$.

Case 2 $\psi > \phi$. [Advance p_l.] Set $\Phi_{\text{left}} \leftarrow \Phi_{\text{left}} + \pi - \varphi_l$, $\theta_{\max} \leftarrow \Phi_{\text{left}} + \frac{3\pi}{2} + \frac{\phi}{2}$, and $l \leftarrow l - 1$.

Case 3 $\psi = \phi$. [Advance both p_l and p_r.] Set $\Phi_{\text{left}} \leftarrow \Phi_{\text{left}} + \pi - \varphi_l$, $\Phi_{\text{right}} \leftarrow \Phi_{\text{right}} + \pi - \varphi_r$, $\theta_{\max} \leftarrow \Phi_{\text{left}} + \frac{3\pi}{2} + \frac{\phi}{2}$, $l \leftarrow l - 1$, and $r \leftarrow r - 1$.

Output the current interval $[\theta_{\min}, \theta_{\max}]$. Set $\theta_{\min} \leftarrow \theta_{\max}$ and repeat Step 2 until r changes from 1 to 0.

The number of intervals produced by the above algorithm cannot exceed $2n$, because each loop of Step 2 advances either p_r to p_{r-1}, or p_l to p_{l-1}, or both to p_{r-1} and p_{l-1} respectively, and because there are $2n$ vertices in total (n each for p_l and p_r) to advance before returning to the initial incident vertices p_0 and p_i.

We can easily apply the above algorithm for the general case in which the vertex of cone C is at an arbitrary point (x_0, y_0) and the axis of the cone forms an angle θ_0 with the y axis. Each generated interval $[\alpha, \beta]$ now needs to be right shifted to $[\alpha + \theta_0, \beta + \theta_0]$, and the corresponding locus of p_0 can be represented as $(a_x \cos\theta + b_x \sin\theta + x_0, a_y \cos\theta + b_y \sin\theta + y_0)$, where a_x, b_x, a_y and b_y are constants determined by θ_0, ϕ and $\triangle p_0 p_l p_r$. (The equations for these constants are given in [11].)

2.3 Upper Bounds

The preceding subsection tells us that the set of possible poses for a convex polygon inscribed in one cone can be described by a continuous and piecewise elliptic curve defined over orientation space $[0, 2\pi)$. We call this curve the *locus curve* for the inscription. This subsection will show that only finite possible poses exist for a convex n-gon P inscribed in two cones, so long as the vertices of the cones do not coincide. An upper bound on the number of possible poses can be obtained straightforwardly by intersecting two locus curves, each resulting from the inscription of P in one cone.

Claim 1 *There exist no more than $8n$ possible poses for a convex n-gon P inscribed in two cones C_1 and C_2 with distinct vertices.*

Proof. Let $p_0, \ldots, p_{n-1}$ be the vertices of P in counterclockwise order; then a pose of P can be represented by the location of p_0 as well as the angle θ between directed edge $\overrightarrow{p_0 p_1}$ and the x axis. Let $\mathcal{C}_1(\theta)$ and $\mathcal{C}_2(\theta)$ be the two locus curves, for the inscriptions of P in cones C_1 and C_2 respectively. We only need to show that $\mathcal{C}_1$ and $\mathcal{C}_2$ meet at most $8n$ times, that is, they pass through common points at no more than $8n$ values of θ.

It is known that each $\mathcal{C}_i$ consists of at most $2n$ elliptic curves defined over a sequence of intervals that partition $[0, 2\pi)$. Intersecting these two sequences of intervals gives a partition that consists of at most $4n$ intervals. Within each interval the possible orientations (hence the possible poses) of P can be found by computing where the corresponding pair of elliptic curves meet. According to the last subsection, this pair of curves may be written in the parameterized forms

$$(a_{ix} \cos\theta + b_{ix} \sin\theta + x_i,\ a_{iy} \cos\theta + b_{iy} \sin\theta + y_i),$$

for $i = 1, 2$. Here (x_i, y_i) is the vertex of cone C_i, and $a_{ix}, b_{ix}, a_{iy}, b_{iy}$ are constants determined by P and C_i. Using the condition $(x_1, y_1) \neq (x_2, y_2)$, we suppose $x_1 \neq x_2$ without loss of generality, and let $\Delta = \sqrt{(a_{1x} - a_{2x})^2 + (b_{1x} - b_{2x})^2}$. Then it is not hard to show that these two curves do not meet if $|x_1 - x_2| > \Delta$. Otherwise they may meet only at

$$\theta = \beta - \alpha \qquad \text{and} \qquad \theta = \pi - \beta - \alpha,$$

where $\alpha = \operatorname{atan}(\frac{a_{1x} - a_{2x}}{\Delta}, \frac{b_{1x} - b_{2x}}{\Delta})$ and $\beta = \sin^{-1}\frac{x_2 - x_1}{\Delta}$. □

The upper bound $8n$ is *not* tight: A lower one can be obtained even without using two cones to constrain the polygon. Notice in the proof above that the bound came from a partition of orientation space $[0, 2\pi)$ into up to $4n$ intervals which combined the individual pose constraints imposed by the two cones. Therefore an improvement on that bound must require a different partitioning of $[0, 2\pi)$. To see this, we regard each cone as the intersection of two half-planes and decompose its constraint on the polygon into two constraints introduced by the half-planes independently.

A polygon P is said to be *embedded* in a half-plane h if P is contained in h and supported by its bounding line. Thus P is inscribed in a cone if and only if it is embedded in the two half-planes defining the cone by intersection. Two cones with distinct vertices together provide three or four half-planes, of which no three have *concurrent* bounding lines, i.e., bounding lines that pass through a common point. Such three half-planes are indeed enough to bound the number of possible poses of P within $6n$.

Theorem 1 *There exist no more than 6n possible poses for a convex n-gon P embedded in three half-planes with non-*

concurrent bounding lines; furthermore, this upper bound is tight.

Proof. Let l_1, l_2 and l_3 be the bounding lines of the three half planes respectively. We can assume that these lines are not all parallel; otherwise it is easy to see that no feasible pose for P exists. So suppose l_1 and l_2 intersect; their corresponding half-planes form a cone in which P is inscribed. Let the orientation of P be represented by the angle θ between the x axis and some directed edge of P. Then orientation space $[0, 2\pi)$ is partitioned into at most $2n$ intervals, according to which pair of vertices are possibly on l_1 and l_2 respectively. In the mean time, it is also partitioned into exactly n intervals, according to which vertex is possibly on l_3. Intersecting the intervals in these two partitions yields a finer partition of $[0, 2\pi)$ that consists of at most $3n$ intervals, each containing orientations at which P is to be supported by l_1, l_2 and l_3 at the same three vertices.

Let us look at one such interval, and let p_i, p_j and p_k be the vertices of P on l_1, l_2 and l_3 respectively whenever a possible orientation exists in the interval. The possible orientations occur exactly where l_3 crosses an elliptic curve $\mathcal{C}(\theta)$ traced out by p_k when sliding p_i and p_j on l_1 and l_2 respectively. Now we prove that this interval contains at most two possible orientations. Note $\mathcal{C}$ does not degenerate into a point because the case $p_i = p_j = p_k$ will never happen, given l_1, l_2 and l_3 are not concurrent. Therefore $\mathcal{C}$ is either an elliptic segment monotonic in θ or a line segment that attains any point for at most two θ values. (See Subsection 2.1.) In both cases, it is clear that $\mathcal{C}$ crosses l_3 for no more than two θ values. Thus there are at most $6n$ possible poses in orientation space $[0, 2\pi)$.

[11] gives an example in which a polygon can actually have $6n$ poses when embedded in three given half-planes, thereby proving the tightness of this upper bound. □

Since any two cones with distinct vertices are formed by three or four half-planes with non-concurrent bounding lines, and since embedding a polygon in three half-planes with non-concurrent bounding lines is equivalent to inscribing it in any two cones determined by intersecting a pair of the half-planes, we immediately have

Corollary 1 *There exist at most $6n$ possible poses for a convex n-gon inscribed in two cones with distinct vertices, and this upper bound is tight.*

Would more half-planes (or cones) further reduce the number of possible poses for an embedded polygon to be asymptotically less than n? The answer is no. For example, an embedded regular n-gon will always have kn possible poses, where $1 \leq k \leq 6$, no matter how many half-planes are present. However, the experimental results in the next section will show that two cones (or four half-planes) are usually sufficient to determine a unique pose for a generic polygon.

2.4 An Algorithm for Inscription

With the results in the previous subsections, we here present an algorithm that computes all possible poses for a convex n-gon P to be inscribed in m cones $C_1, \ldots, C_m$, where $m \geq 2$. (The vertices of these cones are assumed to be distinct.) Let $p_0, \ldots, p_{n-1}$ be the vertices of P in counterclockwise order.

Algorithm 2

Step 1 [Compute an initial set of poses w.r.t. two cones.] Solve for all possible poses of P when inscribed in cones C_1 and C_2 (use Algorithm 1 and see the proof of Claim 1), and let set S consist of the resulting poses (already sorted by orientation). Set $i \leftarrow 3$.

Step 2 [Verify with the remaining cones.] If $i = m + 1$ or $S = \emptyset$ then terminate. Otherwise go to Step 3 if $|S| = 1$ or $|S| = 2$. Otherwise apply Algorithm 1 to generate the locus curve $\mathcal{C}_i(\theta)$ for the inscription of P in C_i. Sequentially verify whether each pose in S is on $\mathcal{C}_i(\theta)$, deleting from S those poses that are not. Set $i \leftarrow i + 1$ and repeat Step 2.

Step 3 [More efficiently verify one or two poses.] For each pose in S, let polygon P' be P in that pose and do the following: For $i \leq j \leq m$ construct the supporting cone C'_j of P' at the vertex of cone C_j; if there exists some C'_j that does not coincide with C_j, then delete the corresponding pose from S.

When the above algorithm terminates, set S will contain all possible poses for the inscription. Corollary 1 shows that there are at most $6n$ poses in S after Step 1. Since the supporting cone of P from a point can be constructed in time $O(\log n)$ using binary search [13], the running time of the algorithm is $O((k-1)n + (m-k+1)\log n)$, i.e., $O(mn)$ in the worst case, where k is the value of variable i when leaving Step 2. However, the experiments in Section 3 will demonstrate that $k = 3$ almost always holds, hence Step 2 will almost never get executed more than once, reducing the running time to $O(n + m \log n)$.

Since a convex m-gon Q is naturally the intersection of m cones, each with a vertex of Q as its vertex and with the internal angle at that vertex as its apex angle, we can use Algorithm 2 to compute all possible inscriptions of P in Q with the same time cost. This problem shall be called the *polygon inscription* problem, which can be regarded as another version of multi-cone inscription because the intersection of multiple cones is always a polygon (possibly unbounded or empty).

3 Experiments

The first set of experiments were conducted to find out how many possible poses usually exist for a polygon embedded in three half-planes with non-concurrent bounding lines, and the results are summarized in Table 1.

# tests	data source	# poly vertices		# possible poses	
		range	mean	range	mean
10000	10 sq.	3–9	5.9376	1–16	5.5436
10000	10 cir.	3–9	6.1208	2–16	5.7006
1000	100 sq.	6–19	11.917	2–16	7.242
1000	100 cir.	10–22	15.108	2–24	10.012
1000	1000 sq.	11–24	18.43	2–18	7.84
1000	1000 cir.	23–43	33.595	2–60	18.709
1000	cir. mar.	3–15	8.98	2–14	4.537

Table 1: Experiments on embedding a polygon in three half-planes.

Seven groups of convex polygons were tested as shown in the above table. The first six groups consisted of convex hulls generated over 10, 100 and 1000 random points successively, and for each number, in two kinds of uniform distributions—inside a square and inside a circle respectively. It can be seen in the table that the polygons in these groups had a wide range (3–43) of sizes, i.e., numbers of vertices, but their shapes were not arbitrary enough, approaching either a square or a circle when large number of random points were used. So we introduced the last group of data that consisted of polygons generated by a method called *circular march* which outputs the vertices of a convex polygon as random points inside a circle in counterclockwise order. The size of a polygon in this group was randomly chosen between 3 and 15. We refer the reader to [11] for details on how all convex polygon data were randomly generated.

Given a convex polygon, three supporting lines, each bounding a half-plane on the side of the polygon, were generated according to the uniform distribution; namely, with probability $\frac{\pi-\varphi_i}{2\pi}$ each line passed through vertex p_i with internal angle φ_i. An additional check was performed to ensure that these lines were not attached to the same vertex of the polygon. The number of possible poses for the polygon to be embedded in these generated half-planes was then computed, and the summarized results for all group are listed in the last two columns of the table.

Table 1 tells us that three half-planes are *insufficient* to limit all possible poses of an embedded polygon to a unique one, namely, the real pose; in fact the table suggests that linear (in the size of the polygon) number of possible poses will usually exist. We can see in the table that despite the appearances of cases with one or two possible poses, the ratio between the mean of numbers of possible poses and mean polygon size lies in the approximate range 0.43–0.93, decreasing very slowly as the mean polygon size in a group increases. These results tend to support a conjecture that in the average case there exist $O(n)$ possible poses for a convex n-gon embedded in three half-planes with non-concurrent bounding lines.

The above conjecture may be very difficult to prove. However, a plausible explanation for the experimental results can be sought. Recall, a polygon with three half-planes defines a partition of orientation space $[0, 2\pi)$ into at most $3n$ intervals, each containing orientations that would allow the same three incident vertices whenever a possible pose exists at that orientation. The feasible orientations in each interval occur when one supporting line crosses the locus curve of its associated incident vertex. The locus curve results from moving the other two incident vertices along their supporting lines. As these curves (for all intervals) may often cluster together, the likelihood that they get crossed $O(n)$ times in total by the first supporting line is quite large. This happened particularly often during the experiments when a vertex coincided with an intersection of two supporting lines. (See Figure 4.)

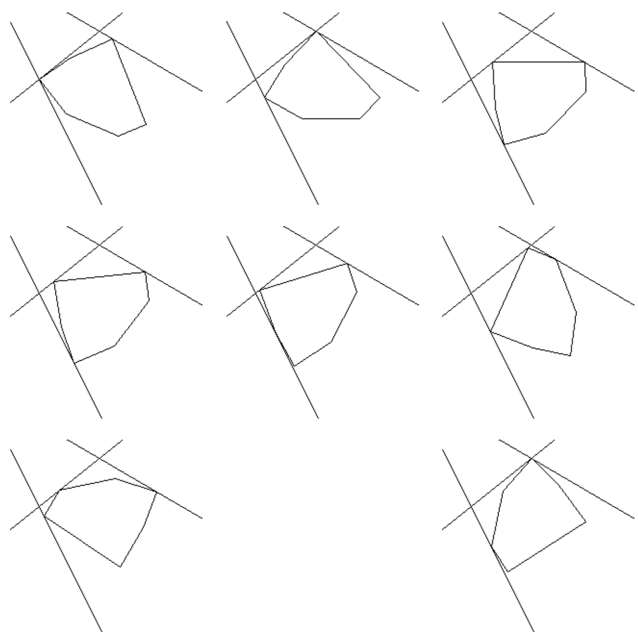

Figure 4: Eight possible poses for a convex 6-gon bounded by three supporting lines (as taken from a sample run). The first one represents the real pose whose supporting lines as shown were generated randomly; the remaining seven represent all other poses consistent with the supporting line constraints. Notice in this example that three of the eight poses occurred when a vertex of the polygon coincided with an intersection of two supporting lines.

The purpose of the second set of experiments was to study how many poses usually exist for a convex polygon inscribed into two or more cones with distinct vertices. We first tested with two cones using the same source of random data generated in the way we did for the first set of experiments, and the results are shown in Table 2. Since

# tests	data source	# poly vertices		# possible poses	
		range	mean	range	mean
10000	10 sq.	3–10	5.9493	1–2	1.036
10000	10 cir.	3–10	6.1113	1–2	1.0117
1000	100 sq.	6–18	12.002	1–2	1.01
1000	100 cir.	9–21	15.138	1–1	1
1000	1000 sq.	10–26	18.073	1–2	1.002
1000	1000 cir.	26–45	33.665	1–2	1.003
1000	cir. mar.	3–15	9.009	1–2	1.106

Table 2: Experiments on inscribing polygons with two cones.

a polygon was always generated inside a square (circle), the cone vertices were chosen as random points uniformly distributed between this boundary and a larger square (circle). The ratio between these two squares (circles) was set uniformly to be $\frac{1}{2}$ for all seven groups of data.

In contrast to Table 1, Table 2 tells us that two cones allow a *unique* pose of an inscribed polygon in most cases. In each group of tests, *only* cases with one pose or two poses occurred, and the mean of possible poses stayed very close to 1, independent of the mean polygon size. (It is not hard for us to see that the percentage of two-pose cases was very low in the range 0%–3.6% for the first six groups of data. The percentage 10.6% for the seventh group was a bit high but expected, because polygons generated by circular march were more likely to be in a certain shape that would often incur two possible poses, as we will discuss later.)

Tests were also conducted with 3–10 cones on reproduced data of four of the seven groups, while the other experiment parameters were kept the same. As shown in Table 3, the means of possible poses did not decrease dramatically as compared to those in Table 2. Finally, we repeated the first group of tests with two cones but chose their vertices from two wider ranges (with the previous ratio $\frac{1}{2}$ replaced by ratios $\frac{1}{5}$ and $\frac{1}{10}$ respectively), and the results are shown in the last two rows of the same table.

The experimental result that two non-incident cones usually allow a unique pose of an inscribed convex n-gon P has in fact a very intuitive explanation. As mentioned before, two such cones generally provide four half-planes, any three of which will limit the number of possible poses of P to at most $6n$. Let polygons $P_1, \ldots, P_m$, where $m \leq 6n$, represent P in all possible poses respectively when embedded in the first three half-planes; then those P_i corresponding to the final possible poses must be supported by the bounding

# tests	data source	# poly vertices		# possible poses	
		range	mean	range	mean
10000	10 sq.	3–10	5.95	1–2	1.0056
1000	100 sq.	6–19	11.901	1–1	1
1000	1000 sq.	10–27	18.26	1–1	1
1000	cir. mar.	3–15	8.96	1–2	1.022
10000	10 sq.	3–10	5.9587	1–2	1.1158
10000	10 sq.	3–10	5.9741	1–2	1.1738

Table 3: More experiments on inscribing polygons with cones. The first four groups of data were tested with a random number (between 3 and 10) of cones; the last two groups were tested with two cones whose vertices were chosen from wider ranges than the previous tests.

line l of the fourth half-plane. So the probability that a two-pose case occurs is no more than the probability that l passes through a vertex of P_i and a vertex of P_j, for $i \neq j$. Note the vertices of $P_1, \ldots, P_m$ together occupy $\Theta(mn)$ points in the plane in general, only one of which must lie on l. If no two of these vertices coincide, the probability that l passes through two vertices of different polygons is zero (assuming that l is independent of the other three bounding lines), which means that a two-pose case almost never occurs in this situation. Otherwise suppose two vertices of P_i and P_j respectively are at the same point p for some $i \neq j$, then the probability that l passes through p is $\Theta(\frac{1}{nm})$. This is $\Theta(\frac{1}{n^2})$ in the average case, given that $m = \Theta(n)$ as suggested by the first set of experiments. Since in the usual case there only exist a constant number of such coincident vertex pairs, the probability $\Theta(\frac{1}{n^2})$ is an approximate upper bound on how often two-pose cases occur. This bound turns out to be consistent with the percentages of two-pose cases in Table 2.

It was observed during the experiments that a large number of two-pose cases occurred when both cones happened to be supporting the polygon at the same pair of vertices. (See Figure 5.) The two possible orientations differed by

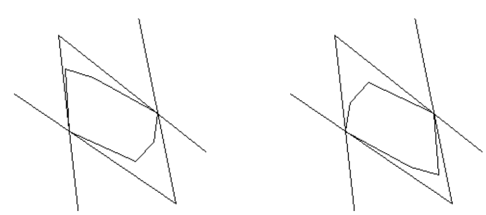

Figure 5: Two possible poses for a convex 6-gon inscribed in two cones (as taken from a sample run). The two cones are supporting the polygon at a pair of vertices.

π, and each supporting vertex in one pose coincided with the other in the other pose. This situation often happened when the distance between one pair of vertices of the polygon was much larger than the distance between any other pair of vertices, or when the sites were far away from the polygon (as evidenced by the high percentages of two-pose cases in the last two groups of tests in Table 3).

4 Conclusion

The analyses and experiments in this paper have laid out the basis for a general sensing scheme applicable to planar objects with known shapes. The scheme, termed *sensing by inscription*, determines the pose of an object by finding its inscription in a polygon of geometric constraints derived from the sensory data. A specific implementation of this scheme may employ certain combinations of simple and robust sensors to obtain the necessary constraints. In particular, two supporting cones are often enough to detect the real pose of a polygonal object. In real situations, if two (or more) possible poses arise from a two-cone inscription, they can be distinguished by probing at a point contained inside only one of the poses.[3]

Though only the inscription of a convex polygon is treated in this paper, the extensions to any arbitrary polygon and any polyhedron (with near constant cross-section along some direction) should be straightforward; but the extension to a closed and piecewise smooth curve needs further study. The technique can also be applied in object recognition: A finite set of polygons are generally distinguishable by inscription.

Future work will involve the design of specialized cone sensors or other sensors suited for inscription, as well as an investigation of a theoretical framework for incorporating sensing uncertainties into the inscription algorithms.

Acknowledgements We wish to thank Matt Mason for his insights and for helping us better interpret the experimental data, and Mel Siegel for suggesting a possible optical implementation using a linear CCD coupled with a diverging lens, and Steve Shafer for offering useful feedback on the presentation of the paper.

References

[1] E.M. Arkin, L.P. Chew, D.P. Huttenlocher, K. Kedem, and J.S.B. Mitchell. An efficiently computable metric for comparing polygonal shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(3):209–216, 1991.

[2] Francis Avnaim and Jean Daniel Boissonnat. Polygon placement under translation and rotation. Technical Report 889, INRIA Sophia-Antipolis, 1988.

[3] B.S. Baker, S.J. Fortune, and S.R. Mahaney. Polygon containment under translation. *Journal of Algorithms*, 7(4):532–548, 1986.

[4] John F. Canny and Kenneth Y. Goldberg. A "RISC" paradigm for industrial robotics. Technical Report ESRC 93-4, 1993.

[5] Bernard Chazelle. The polygon containment problem. In Franco P. Preparata, editor, *Advances in Computing Research*, pages 1–32. Jai Press Inc., 1983.

[6] Richard Cole and Chee K. Yap. Shape from probing. *Journal of Algorithms*, 8(1):19–38, 1987.

[7] D. Dobkin, H. Edelsbrunner, and C.K. Yap. Probing convex polytopes. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing*, pages 424–432. ACM Press, 1986.

[8] Steven Fortune. Fast algorithms for polygon containment. In *Automata, Languages and Programming*, pages 189–198. Springer-Verlag, 1986. LNCS 194.

[9] Kenneth Y. Goldberg and Matthew T. Mason. Bayesian grasping. In *Proceedings of 1990 IEEE International Conference on Robotics and Automation*, pages 1264–1269, 1990.

[10] W.E.L. Grimson. *Object Recognition by Computer: the Role of Geometric Constraints*. MIT Press, 1990.

[11] Yan-Bin Jia and Michael Erdmann. Sensing with known shapes. Technical Report CMU-RI-93-21, CMU Robotics Institute, 1993 (to appear).

[12] D. Mumford. The problem of robust shape descriptors. In *Proceedings of the First International Conference on Computer Vision*, pages 602–606, 1987.

[13] Franco P. Preparata and Michael Ian Shamos. *Computational Geometry: an Introduction*. Springer-Verlag, 1988.

[14] Anil S. Rao and Kenneth Y. Goldberg. Placing registration marks. In *Proceedings of 1993 IEEE International Conference on Robotics and Automation*, pages 1:161–167, 1993.

[15] Aaron S. Wallack, John F. Canny, and Dinesh Manocha. Object localization using crossbeam sensing. In *Proceedings of 1993 IEEE International Conference on Robotics and Automation*, pages 1:692–699, 1993.

[3] Notice that cases with more than two possible poses never occurred in our experiments on two-cone inscription.