MODEL-BASED PROBING STRATEGIES FOR CONVEX POLYGONS

Eugene Joseph and Steven Sol Skiena

ABSTRACT

We prove that $n + 4$ finger probes are sufficient to determine the shape of a convex n-gon from a finite collection of models, improving the previous result of $2n + 1$. Further, we show that $n - 1$ are necessary, proving this is optimal to within an additive constant. For line probes, we show that $2n + 4$ probes are sufficient and $2n - 3$ necessary. The difference between these results is particularly interesting in light of the duality relationship between finger and line probes.

October 9, 2003

MODEL-BASED PROBING STRATEGIES FOR CONVEX POLYGONS

Eugene Joseph and Steven Sol Skiena

1. Introduction

Tactile sensing is an important paradigm in robotics, and for reasons of economy and robustness is often used instead of more sophisticated vision systems to explore unknown environments. Cole and Yap [1] introduced the notion of a finger probe to model a tactile sensor, where a finger probe measures the first point of contact between a directed line l and an object P. Since Cole and Yap's work, a significant literature in geometric probing has developed, which studies the power of different sensor models for reconstructing geometric objects. The most up-to-date collection of results in probing appear in [2].

We seek probing strategies which completely determine a convex polygon in as few probes as possible. Cole and Yap [1] proved that $3n$ finger probes are necessary and sufficient to determine an unknown convex n-gon, given only the position of some point in the interior of the polygon. Probing strategies for non-convex polygons have been developed by Alevizos, Boissonnat and Yvinec [3].

Since the environment of industrial robots is usually very restricted, we often have apriori knowledge of the class of objects the robot will manipulate. Thus in most tactile sensing applications, we are concerned with identifying an object and its orientation from a finite, pre-defined set of possible objects. Grimson and Lozano-Pérez [456] have studied model-based tactile sensing, and shown that heuristics can be effective to distinguish between models. More efficient probing strategies can result for model-based determination problems. Bernstein [7] proved that $2n+2$ finger probes are sufficient to determine a convex n-gon from a finite collection of models Γ, which is improved to $2n+1$ in [8]. Lyons and Rappaport [9] showed that $m-1$ probes are sufficient to identify a convex polygon from a set of m models, if each model has a particular edge aligned with a known reference plane. This is a severe restriction, which if relaxed leads to a $mn-1$ probe determination strategy. In this paper, we prove $n-1$ probes are necessary for model-based determination of convex polygons and that $n+4$ probes are sufficient. Therefore, our result is optimal to within an additive constant.

Other sensor models are also of interest. A line probe measures the first time of intersection between a line moving parallel to itself and an object. Thus the first line tangent to the object with a given slope is returned. Li [10] showed that $3n+1$ line probes are necessary and sufficient to determine an unknown convex n-gon. The problem of model-based determination with line probes was posed in [11], and previously no non-trivial bounds were known. In this paper, we also prove that $2n-3$ line probes are necessary and $2n+4$ probes sufficient, which are again tight to within an additive constant.

These results are particularly interesting in light of the duality relationship, discovered independently by Dobkin, Edelsbrunner, and Yap [12] and Greschak [13], that exists between line probes and finger probes which all pass through a single point, the origin. For all previous determination problems, the finger and line probing models have been identical in power to within one probe. However, our results show that line probes are significantly weaker than unrestricted finger probes for model-based determination.

Figure 1: Determining e_n, e_1 and ψ_1 using Bernstein's strategy.

Figure 2: Bernstein's strategy modified; three possible choices for e_2.

2. Model-based Results for Finger Probes

In the model-based probing problem, we are given a set of convex polygons Γ, and a point O which lies in the relative interior of an unknown convex polygon P from Γ. We seek to determine P and its orientation using as few probes as possible.

Our finger probing strategy is a refinement of Bernstein's strategy [7]. There are two aspects to this strategy. First, all the models are preprocessed to find an angle $\theta_{\min}$ small enough so that at most five probes through O, each inclined with respect to the previous probe by $\theta_{\min}$ as in Figure 1, will determine the first edge e_1 of P. Since P is convex, three collinear contact points determine an edge, and as we will show an appropriately small $\theta_{\min}$ can be computed from Γ which will guarantee three such points in five probes. If five probes are actually used, then two neighboring edges will be determined.

Bernstein also observed that if two probes (F_1, F_2) are aimed parallel to a previously determined edge e_i but at a height less than some $h_{\min}$, both probes will contact the next edge e_{i+1} of P, determining e_{i+1} and implicitly the vertex between e_i and e_{i+1}. Bernstein's strategy proceeds to walk around the polygon determining each edge in two probes, for a total of $2n+2$ probes.

We improve Bernstein's strategy by showing that vertex v_{i+1} between e_i and e_{i+1} can be determined from the initial angle ψ_1, e_i and Γ in exactly one probe. Determine e_n and e_1 in five probes using Bernstein's strategy; these labels are defined after probing. Overlay all models which possess the same initial angle ψ_1 between e_n and e_1, as in Figure 2. Relative to edge e_i (e_1 in Figure 2), we have a number of choices for the next edge e_{i+1}.

Figure 3: Starting configurations for Bernstein's strategy.

Aim a probe F_i (F_1 in Figure 2) that is parallel to and above e_i and is at a height below both the lowest model vertex (S_1 in Figure 2), and below any intersection of candidate edges (S_2 in Figure 2). Such a probe will intersect P at a point unique to only one candidate edge, which can be determined by substituting the coordinates of the collision point into the equations for the candidate edges, although the length of this new edge is still unknown. Thus vertex v_{i+1} and the orientation of edge e_{i+1} have been found at a cost of one probe, and we can walk around P determining each new vertex at the cost of a single probe.

It remains to be specified how to determine $\theta_{\min}$ [8]. For any point s in polygon $P \in \Gamma$, define β_s^P as the smallest angle spanned by any edge of P by a point s. Further, let $\beta_{\min}^P = Min\{\ \beta_s^P,\ s \in P\}$. In a convex polygon, the point which gives rise to the minimum angle must be on a vertex or edge of P. Finally define $\beta_{\min} = Min\ \{\beta_{\min}^P, | P \in \Gamma\}$ and $\theta_{\min} = \beta_{\min}/5$. The factor of 1/5 ensures that five probes, each inclined at $\theta_{\min}$ with respect to the previous one, will all remain within an angular sector of $\beta_{\min}$. Such an angular sector can cross only one vertex boundary. By testing each pair of edges for each model, $\theta_{\min}$ can be computed in $O(n^2 m)$ time.

Theorem 1: $n+4$ finger probes are sufficient to determine a convex polygon P from a set of models Γ.

Proof: The previous discussion demonstrated that it is possible to determine P in one probe per vertex, once the initial vertex v_1 has been determined. Figure 3 illustrates the possible results for our initial probes, each aimed at O at an angle of $\theta_{\min}$ often from the previous probe. Three collinear points determine e_1 after either three, four, or five probes. When five probes are required, the orientation of e_n and the angle between e_1 and e_n also results. When four probes suffice to determine e_1, a fifth probe can be sent at an angle $-\theta_{\min}$ relative to F_1, also determining e_n. When three probes suffice to determine e_1, Bernstein's strategy can be employed to determine e_n with two more probes. Thus we can determine two edges i.e. the first vertex v_1 in exactly five probes. As discussed above, the other $n-1$ vertices can be identified with one probe each, for a total of $n+4$ finger probes. □

To determine the time complexity of this strategy, consider that there are m models in the set, each of at most n sides. The initial stage finds the angle ψ_1 between edges e_n and e_1 with $O(n^2 m)$ invested in computing $\theta_{\min}$. All internal angles of the polygons that match ψ_1 are then superimposed to give the overlay diagram of Figure 2. Each polygon can have up to n angles matching ψ_1, so the overlay diagram consist of at most nm superimposed polygons. For the second phase, each probe is aimed below the lowest intersection point, which can be computed taking the minimum over all $\binom{nm}{2}$ possible intersection points. This process of finding the lowest intersection point must be repeated for each edge as we walk along the polygon, for a total time complexity of $O(n^3 m^2)$. We remark that a global value for the minimum intersection point can be precomputed in $O(m^2 n^2)$ time, so that each probe takes $O(1)$.

Figure 4: Models for finger probe lower bound.

Theorem 2: $n-1$ finger probes are necessary to determine a convex polygon P from a set of models Γ.

Proof: Γ will consist of two models, each regular $(n-1)$-gons with an additional vertex raised above a single edge e of each polygon. The raised vertex will be close to the center of e and infinitesimally above e, such that any line passing through two raised vertices intersects the interior of P. The raised vertices added to the two models are not identical, as in Figure 4.

In our lower bound proof, we assume that the position of the $n-1$ regular vertices are freely given to the prober, so that to complete determination only the position of the raised vertex must be found. Because the raised vertex lies only slightly above an $(n-1)$-gon edge, only one of the $n-1$ possible positions of the raised vertex can be tested with a single probe. Thus an adversary can adjust the orientation of the model so the $n-2$ non-raised edges will be probed before the location of the raised edge is known. Then another probe must be spent to distinguish between the two models. Thus $n-1$ probes are necessary to determine P. □

3. Model-based Results for Line Probes

There is a duality relationship between finger probes through O and line probes, which means the lower bound of Theorem 2 immediately dualizes to line probes. Although we might hope that the strategy of Theorem 1 can be adapted to line probes, this strategy aims probes close to edges, which in general will not pass through O. In this section, we prove that $\sim 2n$ line probes are necessary and sufficient for determination.

3.1. An Upper Bound for Line Probes

The structure of our strategy is based on developing constraints from superimposing all possible orientations of models, as in Theorem 1. We observe that the position of a vertex can be identified with a single line probe l, if it could be known that l was oriented in such a manner that no two candidate vertices define a line with the same slope as l. An additional probe may be necessary to confirm that two vertices define an edge of P.

We define a *diagonal* as a line segment joining two vertices of a polygon P. After an initialization procedure which determines diagonal AB of P, we may use AB as a reference to superimpose all models with diagonals of equal length, as in Figure 5. This divides the problem of determining P into two parts, determining the vertices of P above and below the diagonal. Each new vertex v_i of P defines two new diagonals, $v_i v_a$ and $v_i v_b$, where $v_a v_b$ is the current diagonal being probed. The algorithm recurs on each diagonal until determination is completed. The diagonals encountered during the execution of this procedure define P and its triangulation.

Figure 5: Three polygons P_1, P_2, and P_3 overlayed on diagonal AB.

Figure 6: Determining the first two vertices of P.

The initialization phase of our strategy determines the first vertex v_1 of P by sending line probes inclined at $\phi_{\min}$ with respect to the previous probe, until three probes pass through the same point, v_1. The angle $\phi_{\min} = \dfrac{\pi - \psi_{\max}}{5}$, where $\psi_{\max}$ is largest internal angle of all models, has the property that five successive line probes each inclined at $\phi_{\min}$ to the previous one can only cross a single edge boundary. Ideally, only three probes are necessary to identify an initial vertex. However similar to Theorem 1, in the worst case L_1, L_2 and L_3 do not pass through the same point; see Figure 6. By sending L_4 inclined at $+\phi_{\min}$ with respect to L_3 and L_5 inclined at $-\phi_{\min}$ with respect to L_1, we can identify two adjacent vertices, and the edge between them for the cost of five probes. No other edges can be crossed since it takes at least $5\phi_{\min}$ to cross an edge boundary. It will be shown that each vertex or edge can be confirmed for the cost of a single probe. Thus we can take the cost of an initial vertex to be three probes. The extra vertex and edge of Figure 6 will either be identified during initialization by L_1 and L_5, or be confirmed by two other probes later in the algorithm. By performing the initialization procedure twice, once from the top of the polygon and once from the bottom, we can determine two distinct initial vertices defining the initial diagonal AB.

Figure 7: Subproblem with known diagonal $v_a v_b$.

The strategy will recur on each diagonal, aiming probes at shallow-enough angles to determine new vertices if they exist. To find this angle, consider the situation of Figure 7, where (v_a, v_b) is the current diagonal. Let M be the set of all possible model vertices which lie above (v_a, v_b). For any points $v_i, v_j \in M \cup v_a$, let α_{ij} be the angle defined between lines (v_a, v_b) and (v_i, v_j). Let $\alpha_{\min} = Min(\alpha_{ij} > 0)$ for all $v_i, v_j \in M \cup v_a$. A probe aimed at an angle $\alpha_{\min}/2$ above (v_a, v_b) will contact exactly one vertex of $M \cup v_a$. If it contacts v_a, $v_a v_b$ must be an edge of P.

Theorem 3: $2n+4$ line probes are sufficient to determine a convex polygon P from a set of models Γ.

Proof: The first two vertices of P can be determined in three probes each as discussed above. Each second phase probe, aimed at $\alpha_{\min}/2$ with respect to a current diagonal $v_a v_b$, will either be incident upon one vertex of M or else pass through one of the vertices of $v_a v_b$. In the first case, we have determined a new vertex of P and defined two new diagonals, while in the second case we have confirmed that the two vertices of the diagonal are adjacent on P. Since there are $n-2$ vertices and n edges of P which must be confirmed, and each second phase probe confirms either an edge or vertex of P, $6+n+n-2=2n+4$ probes are sufficient to determine P. □

Each of the m convex n-gon models defines at most $O(n \log n)$ diameters of a given length [14]. Thus a given diagonal may define $O(mn^2 \log n)$ points in M. For any set of r points, the minimum slope $\alpha_{\min}$ can be determined in $O(r \log r)$ using the algorithm of Cole, Salowe, Steiger and Szemeredi [15]. Thus with $O(n)$ probes and $O(mn^2 log(n) log(mn^2 \log n))$ steps to determine $\alpha_{\min}$ for each probe, we have a time complexity of $O(mn^3 log(n) log(mn^2 \log n))$ for the algorithm. We remark that a global value for $\alpha_{\min}$ can be precomputed in $O(mn^3 log(mn^2 \log n))$ time, so that each probe takes $O(1)$ time.

3.2. A Lower Bound for Line Probes

A lower bound on the complexity of determination of an n-gon can be shown by specifying a set of models, and describing an adversary which forces any probing strategy to take a given number of probes to determine P from the given set of models. We shall prove a $\sim 2n$ lower bound on determination with line probes, which requires a more complicated set of models than the proof of Theorem 2.

Consider three regular $(n-1)$-gons, of diameters 1, $1+x$, and $1+2x$ where $0 < x \ll 1$, nested within each other as in Figure 8. v_i^k is the i th vertex of the k th largest $(n-1)$-gon. Observe that we can now construct 3^{n-1} distinct convex $(n-1)$-gons $(v_1^{k_1}, v_2^{k_2}, \cdots, v_{n-1}^{k_{n-1}})$ where $k_i \in \{1,2,3\}$. We now convert each of these $(n-1)$-gons to an n-gon, by adding a single raised vertex to some edge of the polygon. Now define a raised edge e as having associated with it a raised vertex v^r, a distance $0 < \varepsilon \ll x$ above the center of e. By raising each of the possible edges, each of these $(n-1)$-gons gives rise to $n-1$ distinct n-gons. Eliminating duplicates from the resulting set of $(n-1)3^{n-1}$ polygons defines the set of models for our lower bound proof.

Figure 8: Models for line probe lower bound.

Figure 9: The size of a raised edge, $0 < \theta < \phi$

What is the significance of the raised edge? A raised edge forces any probing strategy to probe both the edges and the incident vertices. The condition $0 < \varepsilon \ll x$ must hold because if ε were much larger as in Figure 9, a single probe L_i could determine if v_i^1 or v_i^2 existed, and whether the resulting edge e_i were raised.

For the given set of models, each of the $n-1$ major vertices of P has three possible positions, and any of the $n-1$ edges may be the raised edge. Thus any determination strategy can be considered as

Figure 10: Case of Lemma 4

solving a series of subproblems, each of which is of one of the following types:

- Two consecutive major vertices of P are known, but it is not determined whether this edge e is raised. Determine whether e is raised.
- One major vertex of P is known, but the adjacent major vertex and the adjoining edge of P are not known. Determine the unknown vertex and edge.
- Two consecutive major vertices and the connecting edge of P are not known. Determine the vertices and the edge.

The adversary will force any strategy to take one probe to determine each vertex and edge, after initialization. This involves showing that these subproblems require at least one, two, and three probes respectively. Since there are two possible models in the first case, and any probe which will determine whether e is raised is restricted to such a narrow range that this probe cannot help in determining any other vertices of P, at least one probe is necessary to test if e is raised. The second and third cases are resolved below.

Lemma 4: At least 2 line probes are required to determine an unknown edge and a single incident vertex.

Proof: We will actually consider the more restricted case where the unknown incident vertex has a choice between two, and not three, possible locations, as illustrated in Figure 10. Let v_{i-1} be the known vertex and v_i and e_i the unknown vertex and edge, with $v_i^{r_1}$ or $v_i^{r_2}$ the possibly raised vertex. The general case of three possible locations for the unknown vertex cannot be solved in fewer probes, and the restricted case here will be required in the proof of Lemma 5.

Table 1 summarizes the contact points returned by the adversary for probes of different orientations. Orientation angles are measured with respect to $v_{i-1}O$ as shown in Figure 10. Orientations > 90 degrees or < $90 - \phi$ degrees provide no information about the current edge e_i and thus need not be considered. Define θ_i^a as the angle $v_i^a v_{i-1} O$ where O is the origin of P.

Note that for the orientations of interest, if a second probe is not sent, we will not be able to distinguish whether the edge under consideration is normal or raised. Thus two probes are necessary to determine a single unknown vertex and its adjoining edge. □

Probe orientation and contact points			
1st Probe	*1st Contact*	*2nd Probe*	*2nd Contact*
$\geq \theta_i^1 - \phi_r$	v_{i-1}	$\geq \theta_i^2$	v_{i-1}
$\geq \theta_i^1 - \phi_r$	v_{i-1}	$< \theta_i^2$	v_i^2
$< \theta_i^1 - \phi_r$	v_i^1	$\geq \theta_i^1$	v_{i-1}
$< \theta_i^1 - \phi_r$	v_i^1	$< \theta_i^1$	v_i^1

Table 1: Adversary strategy for unknown edge and a single incident vertex.

Figure 11: Case of Lemma 5

Lemma 5: At least 3 line probes are required to determine an unknown edge and two incident vertices.

Probe orientation and contact point for 1st probe		
1st Probe	*1st Contact*	*Status*
$\geq \theta$	v_{i-1}^1	*discard* v_i^1, *reduces to case of Lemma 4 with unknown vertex locations* v_i^2, v_i^3
$< \theta$	v_i^1	*discard* v_{i-1}^1, *reduces to case of Lemma 4 with unknown vertex locations* v_{i-1}^2, v_{i-1}^3

Table 2: Adversary strategy for unknown edge and two incident vertices.

Proof: Consider Figure 11 and the adversary strategy of Table 2, which describes the response to the first probe when an edge and both incident vertices are unknown. Here the adversary returns one outermost vertex and discards the other. This reduces the problem to an unknown edge and a single incident vertex. From Lemma 4, we know that at least two more probes are needed to determine the remaining vertex and connecting edge, for a total of at least three probes. □

Note the requirement that the case of Lemma 5 reduce to that of Lemma 4 after the first probe forces us to use three instead of two nested $(n-1)$-gons to form our models. We are now in a position to prove the lower bound theorem.

Theorem 3: $2n-3$ line probes are necessary to determine a convex polygon P from a set of models Γ.

Proof: Each model in our adversary set contains $n-1$ major vertices and one raised vertex. The adversary can easily ensure that the raised edge of the particular model will be identified only after $n-2$ edges have been verified to be unraised. Thus $n-2$ edges need be verified, and the location of the raised edge would have then been located by elimination. The raised edge contributes a single vertex to the n-gon, the position of which can be inferred from the fact the edge is raised. However, each of the $n-1$ major vertices need to be verified. Since each verification requires at least a single probe, then at least $n-2+n-1=2n-3$ line probes are required for determination. □

4. Conclusions

We have proven bounds, tight within an additive constant, on the number of finger and line probes required for model-based determination for convex polygons. The disparity between these bounds is interesting in light of the duality relationship between them. Our lower bound proof for line probes required an exponential number of models. It would be interesting to know whether fewer models suffice.

The problem of model-based probing strategies remains open for more advanced models. The x-ray probe [16] returns the length of intersection between a line and an object. Since an x-ray probe can be simulated by two finger probes, Theorem 2 can be used to prove a lower bound of $\sim n/2$ for model-based determination. The question is whether this is achievable. Since an x-ray probe through a known point behaves as a finger probe, $\sim n$ x-ray probes suffice for model-based determination using the ideas in [16]. Since each x-ray probe passes through two different edges, perhaps a $\sim n/2$ probe strategy is possible. Further, the model-based problem is open for half-plane probes [17].

We thank the referees for thorough reports which improved the presentation of our results.

References

1. Richard Cole and Chee K. Yap, ''Shape from Probing,'' *Journal of Algorithms*, vol. 8, no. 1, pp. 19-38, 1987.

2. S. S. Skiena, ''Interactive Reconstruction from Projections,'' *Proceedings of IEEE, Special issue on Computational Geometry*, to appear.

3. P. D. Alevizos, J. D. Boissonnat, and M. Yvinec, ''Non-convex Contour Reconstruction,'' *J. Symbolic Computation*, to appear.

4. W. Eric L. Grimson and Tomas Lozano-Perez, ''Model-Based Recognition and Localization from Sparse Range or Tactile Data,'' *International Journal of Robotics Research*, vol. 3, pp. 3-35, Fall 1984.

5. Peter C. Gaston and Tomas Lozano-Perez, ''Tactile Recognition and Localization Using Object Models: The Case of Polyhedra on a Plane,'' *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, pp. 257-266, May 1984.

6. W. Eric L. Grimson and Tomas Lozano-Perez, ''Localizing Overlapping Parts by Searching the Interpretation Tree,'' *IEEE Transactions on Pattern Recognition and Machine Intelligence*, vol. PAMI-9, pp. 469-482, July 1987.

7. Herbert J. Bernstein, ''Determining the Shape of a Convex *n*-sided Polygon by Using *2n+k* Tactile Probes,'' *Information Processing Letters*, vol. 22, pp. 255-260, 1986.

8. Steven S. Skiena, *Geometric Probing*, Doctoral Dissertation, University of Illinois, Department of Computer Science, Urbana, IL, 1988.

9. Kelly Lyons and David Rappaport, *An Efficient Algorithm for Indentifying Objects Using Robot Probes*, Dept. of Computing and Information Science, Queen's University, Ontario Canada, January 1989.

10. S. Y. Robert Li, ''Reconstruction of Polygons from Projections,'' *Information Processing Letters*, vol. 28, pp. 235-240, 1988.

11. Steven S. Skiena, ''Problems in Geometric Probing,'' *Algorithmica*, vol. 4, pp. 599-605, 1989.

12. David P. Dobkin, Herbert Edelsbrunner, and Chee K. Yap, ''Probing Convex Polytopes,'' *Proc. 18th ACM Symposium on the Theory of Computing*, pp. 424-432, 1986.

13. John Paul Greschak, *Reconstructing Convex Sets*, Ph.D. Dissertation, MIT, Department of Electrical Engineering and Computer Science, Cambridge MA, 1985.

14. Z. Fü redi, ''The Maximum Number of Unit Distances in a Convex 0sfont]I]ize]u'n0rfont]-gon,'' *Manuscript*, 1987.

15. R. Cole, J. Salowe, R. Steiger, and E. Szemeredi, ''An Optimal-time Algorithm for Slope Selection ,'' *SIAM J. Computing*, 1989.

16. H. Edelsbrunner and S. S. Skiena, ''Probing Convex Polygons with X-rays,'' *SIAM J. Computing*, vol. 17, pp. 870-882, 1988.

17. Steven S. Skiena, ''Probing Convex Polygons with Half-Planes,'' *J. Algorithms*, vol. 12, pp. 359-374, 1991.