

Robot Localization without Depth Perception

Erik D. Demaine¹, Alejandro López-Ortiz², and J. Ian Munro²

¹ MIT Laboratory for Computer Science, 200 Technology Square,
Cambridge, MA 02139, USA, edemaine@mit.edu

² Department of Computer Science, University of Waterloo, Waterloo, Ontario
N2L 3G1, Canada, {alopez-o, imunro}@uwaterloo.ca

Abstract. Consider the problem of placing reflectors in a 2-D environment in such a way that a robot equipped with a basic laser can always determine its current location. The robot is allowed to swivel at its current location, using the laser to detect at what angles some reflectors are visible, but no distance information is obtained. A polygonal map of the environment and reflectors is available to the robot. We show that there is always a placement of reflectors that allows the robot to localize itself from any point in the environment, and that such a reflector placement can be computed in polynomial time on a real RAM. This result improves over previous techniques which have up to a quadratic number of *ambiguous points* at which the robot cannot determine its location [1, 9]. Further, we show that the problem of optimal reflector placement is equivalent to an art-gallery problem within a constant factor.

1 Introduction

Problem: Robot localization. For a mobile robot to plan its motion, it requires both knowledge of its surrounding environment and accurate information of its current location in this environment. However, the robot's motion is imprecise from such effects as friction, unevenness of the terrain, and inertia, so the robot's location becomes uncertain. Consequently, robots often perform corrective measurements that allow them to rehome their current position (e.g. [5, 3, 8, 4]). Thus the problem of *robot localization* arises: determine the current location of the robot in its surrounding environment. The basic approach to localization is for the robot to sense its immediate surroundings, and then match this local image against an internal model or *map* of the entire environment. Common sensing devices include vision, radar, sonar, and ladar (laser radar).

Highly detailed information about the environment can be obtained only at the expense of a complex vision system, as well as collection and processing time for the data gathered. An efficient low-cost method of localization would thus allow more accurate motion control for the robot. This paper investigates robot localization with particularly cheap and limited vision systems.

Model: Reflecting Landmarks. Typically robots use landmarks to identify their position [8, 3]. These landmarks can either be naturally present (such as a wall or door) or be artificially introduced (magnetic markers, reflectors, beacons). In this paper we follow the model of Sugihara [9], using mutually indistinguishable reflective markers (*reflectors*) that provide angular measurements. This

can be realized in a simple and inexpensive form by placing reflective strips or mirrored cylinders in selected positions in the environment. The robot shines a laser in a 360-degree scan and records the angular magnitude of those directions at which a reflection was detected. The result is a star of rays that the robot must match against its given data. What makes this problem difficult is that the robot does not know the distance at which the reflection occurred, nor which reflector caused the reflection, and nor does the robot know of a preferred direction or “true north.”

Connection to Art Galleries.

Sugihara [9] observed that the reflector-placement problem is a generalization of an art-gallery problem. In the classic art-gallery problem, the goal is to choose fixed locations for *guards* (points) such that every point in the environment is visible from at least one guard; equivalently, at least one guard is visible from every point. If we think of reflective strips as guards, certainly the robot needs to see at least one reflective strip at all times. Thus any solution to the reflector-placement problem is also a solution to the art-gallery problem. We establish a connection in the reverse direction.

Previous Work: Ambiguities. Sugihara [9] showed that it is possible to mark the environment in such a way that the robot can localize itself from all but a finite number of *ambiguous points* [9]. Pairs of ambiguous points have the property that the angle readings are the same from either point in the pair, and hence if the robot is placed at either point, it cannot determine at which of the two points it is located. Avis and Imai [1] proved that the total number of degenerate positions for n reflectors, k of which are visible from the robot, is in the worst case $\Theta(n^2/k)$ [1]. Hence, by placing $k = O(n^2)$ reflectors the number of ambiguous points can be reduced to at most a constant number.

More recently, González-Banos and Latombe [5] considered the related problem of finding a minimum set of identifiably distinct reflectors in a given polygon subject to incidence and range constraints. They propose a randomized algorithm which returns, with high probability, a set of guards which is a small non-constant factor away from the minimum number of guards required. The incidence and range constraints model real-life limitations of reflector resolution and sensing devices. However, in either application, ambiguities are never fully resolved, so the robot cannot be guaranteed to be able to localize itself.




Fig. 1. The robot R knows only the cyclic sequence of angles between visible reflectors (drawn as circles), here placed along the boundary of the polygon. The visibility region is shaded.

Our Results. In this paper, we show that any polygon can be unambiguously marked using at most ten reflective strips per guard using a particular instance of the well-studied family of art-gallery problems. Next, we show that at least four reflectors per guard are needed in the worst case. Lastly, we study changes in the complexity of the localization task when we consider more powerful localization primitives, such as a compass (true north) or a 3-D environment with or without a preferred “up” position.

2 Marking A Single Wall

A simple but important subproblem is when there are no obstacles, and the robot R can be placed anywhere in the plane except on top of one of the reflectors. In this context, we show that two reflectors limit the robot’s position to a one-dimensional curve, and three reflectors limit the robot’s position to a finite number of points. For two reflectors A, B , the robot’s laser scan measures two angles, $\angle ARB$ and $2\pi - \angle ARB$. From elementary geometry it follows that the locus of points forming a fixed angle with two points is an arc of a circle passing through those points together with the reflection of that arc through the line joining the two points. See Figure 2.




Fig. 2. Two reflectors A, B on a line, and the crescent of points R with fixed angle $\angle ARB$.

Lemma 1. [6] *Given two distinct points A and B on a circle, the interior angle $\angle ARB$ is the same for all R on either of the open arcs connecting A and B [Euclid’s Proposition III.21]. Furthermore, if C is the center of the circle, then $\angle ARB = \frac{1}{2}\angle ACB$ for R on the longer arc and $\angle ARB = \pi - \frac{1}{2}\angle ACB$ for R on the shorter arc [Euclid’s Proposition III.33].*

There are two circles with center C and C' such that $\angle AC'B = \angle ACB = \theta$ for any $0 < \theta \leq \pi$, and these circles are reflections of each other through AB . Thus, the longer arcs of these circles correspond to angles θ satisfying $0 < \theta \leq \pi/2$, and the shorter arcs correspond to angles θ satisfying $\pi/2 \leq \theta < \pi$. Hence, there are precisely two arcs corresponding to each angle θ . Together these arcs are called the θ -crescent of A and B .

Lemma 2. *Given an angle $0 < \theta < \pi$, the θ -crescent of points A and B is precisely the locus of points R satisfying $\angle ARB = \theta$.*

For this lemma to hold for $\theta = 0$ and π as well, there are two additional special cases, corresponding to the points along the line AB which we have so far ignored. The points R strictly between A and B satisfy $\angle ARB = \pi$, and the other points R (except A and B) satisfy $\angle ARB = 0$. Thus, we define the π -crescent of A and B to be the open line segment between A and B , and the 0 -crescent to be the line AB minus the closed line segment between A and B .

In particular, two reflectors certainly do not suffice to uniquely determine the position of the robot: they leave every point in the plane ambiguous by an uncountably infinite amount.

We now turn to the case of three reflectors A, B, C in an arbitrary position in the plane. At first it might seem that three angles suffice to uniquely determine a position. Indeed, this would be the case if the robot knew the correspondence between reflectors and reflection angles. Because this information is not known, however, we can cyclicly shift this correspondence, compute the corresponding crescents, and take their intersection, as shown in Figure 3.

More precisely, in this figure, we proceed as follows from the triangle ABC and the point R chosen arbitrarily. We draw the circular arc with aperture $\angle ARB$ subtending AC . We repeat this procedure with $\angle BRC$ subtending AB . These two circular arcs intersect at R' . It follows then by construction that $\angle ARC = \angle BR'C$, $\angle ARB = \angle AR'C$, and $\angle BRC = \angle AR'B$, and hence R and R' cannot be distinguished. Therefore, three reflectors in general position do not uniquely determine the position of the robot. However, three reflectors do restrict the position to one of a finite number of locations, at most for each cyclic shift of the angles. The difficulty is to make this number of possible positions equal to exactly one.




Fig. 3. Positions R and R' cannot be distinguished because the labels of A, B, C are unknown and hence can be shifted. Thin lines show the half-crescents, medium lines show the lines of sight, and thick lines show the triangle.

3 Localization in Polygons

We now focus on the case in which robot motion is limited to a polygon, possibly with polygonal holes. Reflectors are placed on the walls, this means that the robot cannot collide with the reflector and that it must be on the same side of the wall as the reflector. Three reflectors along a common wall will always be seen in the same order so the results of the previous section apply and three reflectors suffice for a convex polygon.

For nonconvex polygons, we convert a particular type of art-gallery guard placement into a collection of reflectors by expanding each guard into a tight cluster of reflectors. For this to work, we consider a slight variation on the art-gallery problem. A *wall guard* is a positive-length (short) interval along an edge of the polygon. A *wall-guard placement* must satisfy that every point in the polygon can see an entire wall guard (strong visibility), and the goal of the art-gallery problem is to minimize the number of wall guards. It is easy to show that the worst-case bounds for wall guards are similar to standard vertex guards: $\Theta(n)$ guards suffice and are sometimes necessary to guard an n -vertex polygon. The main difficulty with obstacles is that (most likely) not all of the reflectors are visible from a given point. Thus a major challenge is to identify *which* reflectors are those seen. Such an encoding scheme will be the focus of this section. Once we have such a scheme, and we can identify three visible reflectors on a common

wall, then we can appeal to the previous section to determine the robot’s position relative to the reflectors. As a result, the entire reflector-placement problem will reduce to the art-gallery problem outlined above.

So we turn to the problem of encoding information in the reflectors so that the robot can tell which reflectors it can see. We make use of the *cross ratio* as suggested by Sugihara [9]. This fundamental concept in projective geometry allows us to store a number—represented by four collinear points (reflectors)—that is readable from any point not collinear with the four points.

3.1 The Cross Ratio of a Pencil

Let A, B, C , and D be four collinear points; a, b, c and d denote four concurrent lines and let O be the point of concurrency. We denote by ac the angle between lines a and c .

Definition 1. *The cross ratio $\{ABCD\}$ of four collinear points A, B, C, D is defined as the quotient $\frac{AC}{CB} / \frac{AD}{DB}$ where the magnitude of a segment is directed, i.e. $AC = -CA$. The cross ratio of a pencil of lines is defined as $\frac{\sin ac}{\sin cb} / \frac{\sin ad}{\sin db}$. The cross ratio of four lines is denoted by $\{abcd\}$. Given four points A, B, C, D and a point O , we denote by $O\{ABCD\}$ the cross ratio of the pencil defined by the lines $\overline{OA}, \overline{OB}, \overline{OC}$ and \overline{OD} .*

Theorem 1 ([7]). *Let $abcd$ be a pencil of lines passing through vertex O , and let L be a line transversal of the pencil not passing through O . Let A be the point of intersection of L and a , and analogously define points B, C , and D . Then $\{abcd\} = \{ABCD\}$. Conversely, let A, B, C , and D be four collinear points, and O be a point off the line $ABCD$. Then $\{ABCD\} = O\{ABCD\}$.*

In other words, the cross ratio of four collinear points is the same when viewed from any vantage point off the line.

We can use this principle to label walls in such a way that the label can be read from any robot position O . More precisely, let G_1, \dots, G_k denote the set of wall guards. Then we place four (collinear) reflectors along wall guard G_i so that those four points have integral cross ratio i . Thus, from any robot position O , the wall guarding guarantees that we see at least one integral cross ratio, and this cross ratio identifies the wall, in principle permitting localization.

Unfortunately, this approach does not suffice, because we do not know which of the visible reflectors form collinear quadruples from a common wall guard. For example, consider a situation in which the robot sees five reflectors, A, B, C, D, X ; the first four reflectors A, B, C, D correspond to a single wall guard; and the last reflector X corresponds to another guard whose quadruple of reflectors is partially occluded. Moreover, the robot happens to be positioned in such a way that both $O\{ABCD\}$ and $O\{BCDX\}$ are integers in $\{1, \dots, k\}$. In this case, the robot cannot in general distinguish which of the two sequences corresponds to a guard and which is spurious.

Indeed, this scenario is but one of several possible ambiguous configurations. To solve these problems, we use additional reflectors and more careful placements of guard reflectors to ensure that these ambiguities are fully resolved.

3.2 The Cross Ratio of Noncollinear Points

The following theorems from projective geometry [10] will help in the task of disambiguating a given set of angular measurements. First we need to characterize those points from which a given noncollinear quadruple forms an integer cross ratio in $\{1, \dots, k\}$:

Theorem 2 (Steiner’s Theorem [2, 10]). *Given four points A, B, C, D , not all collinear, and given a cross ratio r , the locus C_r of points O such that the cross ratio $O\{ABCD\}$ equals r — $C_r(ABCD) = \{O \mid O\{ABCD\} = r\}$ —is a conic curve. Conversely five points in the plane, not necessarily in general position, define a unique conic passing through them.*

The conic may be an ellipse, circle, parabola, hyperbola, or the degenerate cases of a point, a line, or two lines. Consider now four reflectors, three of which are collinear. The following two lemmas describe the robot locations O from which the four reflectors would appear to belong to a common wall guard.

Lemma 3. *Given three collinear points and cross ratio r , there is a fourth point on the line (including the projective point at infinity) realizing cross ratio r .*

Proof. We consider the four points to lie along the (projective) real line with C at the origin. Without loss of generality, let $\text{dist}(AC) = 1$. Define $b = \text{dist}(BC)$ and similarly $d = \text{dist}(BD)$ which implies $\text{dist}(AD) = 1 + b + d$. Hence $\frac{AC}{CB} / \frac{AD}{DB} = \frac{1}{-b} / \frac{1+b+d}{-d} = \frac{d}{b(1+b+d)} = c$. Solving for d we have $d = bc(1+b)/(1-bc)$. \square

Given an edge e of the polygon and three points $X, Y, Z \in e$, let I_{XYZ} be the set of points on e realizing an integer ratio, i.e. $I_{XYZ} = \{W \in e \mid \{XYZW\} \in \mathbb{Z}\}$. If, from a point Q , four reflectors not all collinear form a pencil with an integer cross ratio, then this quadruple is called *a spurious quadruple from Q* . If the location of Q is clear from the context, then we simply refer to the quadruple as *spurious*. In general, given four points A, B, C, D , not all collinear, we denote by $C_{\mathbb{Z}}(ABCD)$ the locus of points that make the quadruple A, B, C, D spurious. More formally, $C_{\mathbb{Z}}(ABCD) = \bigcup_{k \in \mathbb{Z}} C_k(ABCD)$.

Lemma 4. *Given three collinear points A, B, C , a point E not on the line, and a cross ratio r , the locus of points O such that $O\{ABCE\} = r$ is a line.*

Proof. Given the three collinear points A, B, C , let D be the point along that line such that $\{ABCD\} = r$ as per Lemma 3. Then given a point O along the line \overline{DE} we have that $O\{ABCE\} = O\{ABCD\} = r$. \square

From Steiner’s Theorem it follows that $C_{\mathbb{Z}}(ABCD)$ is composed of conics through the points A, B, C, D .

Claim 1. *Given an integer k , a point Q , and three reflectors, the set of fourth reflector points that forms a spurious quadruple from Q with cross ratio k is precisely a line.*




Fig. 4. Point D' such that $Q\{ABCD'\} = k$.




Fig. 5. The view from point Q is ambiguous.

Proof. Using Figure 4 as reference, we are given three reflectors A, B, C and a point Q . First construct the lines \overline{QA} , \overline{QB} , and \overline{QC} . Then use an auxiliary line l not passing through Q which creates the sequence $A'B'C'$ on l . There is a unique point D' on l creating a sequence of cross ratio k as per Lemma 3. Precisely the reflectors D on $\overline{QD'}$ form spurious quadruples with $Q\{ABCD\} = k$. \square

Definition 2. Given a reflector X , we denote by $C_Z(X)$ the union of all conics $C_Z(XYWZ)$ over all reflectors Y, W, Z forming a spurious quadruple. These conics are called *conics of ambiguity*.

Definition 3. A point Q is *ambiguous of degree k* if it belongs to the intersection of k conics of ambiguity for some reflectors X_1, \dots, X_k , i.e. $Q \in \bigcap_{j=1}^k C_Z(X_j)$.

3.3 A Reflector Placement Algorithm

Because four points alone do not in general uniquely determine the position of the robot, we need to increase the number of reflectors per guard. As we have shown, the existence of ambiguous points of degree 1 is unavoidable because there is a set of conics from which a quadruple becomes spurious. This fact extends to points of degree 2, as in general two conics of ambiguity corresponding to two different quadruples might intersect.

To distinguish from these ambiguous points, we place three quadruples along each guard G_i , realizing the cross ratios $3i$, $3i + 1$ and $3i + 2$, respectively. Denote by $q_j = \{A_j, B_j, C_j, D_j\}$ the quadruple of reflectors realizing cross ratio j . From any point in the region guarded by G_i , the robot can always see the three quadruples $q_{3i}, q_{3i+1}, q_{3i+2}$ with consecutive integer cross ratios. We arrange so that the points of ambiguity have degree at most 2. Thus, from any point in the polygon, the robot sees at most two spurious quadruples, and at least one consecutive sequence of three quadruples with cross ratios $3i$, $3i + 1$ and $3i + 2$

which is the actual set corresponding to guard G_i . The robot can localize itself by searching for a consecutive sequence of three quadruples that conclude a common location for the robot, which must then be the correct location because at least one of those three quadruples must not be spurious.

We provide an incremental construction, inserting the quadruples for each guard G_i for i from 0 to n . Denote by g_i the set of reflectors corresponding to guard G_i , i.e. $g_i = q_{3i} \cup q_{3i+1} \cup q_{3i+2}$. At iteration i of the algorithm, the robot can localize itself without ambiguity in any region guarded by a guard G_j for $j < i$. Moreover, as we introduce the reflectors of the set g_i , we ensure that no new ambiguities are introduced in previously guarded regions.

The algorithm maintains sets of relevant geometric objects:

1. A set \mathcal{M} of the reflectors placed thus far.
2. A set \mathcal{D} of all points of degree 2, i.e. $\mathcal{D} = \bigcap_{X \in \mathcal{M}} C_{\mathbb{Z}}(X)$.
3. As we insert A_{i+1} , B_{i+1} , C_{i+1} , and D_{i+1} , the sets of points $I_{B_i C_i D_i}$, $I_{C_i D_i A_{i+1}}$, $I_{C_i D_i A_{i+1}}$, and $I_{D_i A_{i+1} B_{i+1}}$ along edge $e_{\lceil i/3 \rceil}$ respectively.
4. Given a point $Q \in \mathcal{D}$, an integer k , and any three reflectors, as per Claim 1, there is a line along which placement of a reflector would increase the degree of Q . Let \mathcal{H} denote the set of these lines, over all $Q \in \mathcal{D}$, $k \in \mathbb{Z}$, and triples in \mathcal{M} . The algorithm avoids ambiguities by keeping track of the intersections of the lines with the boundary of the polygon ∂P .
5. After inserting the first two reflectors A_i, B_i of a quadruple q_i on edge $e_{\lceil i/3 \rceil}$, it computes the set of points $\mathcal{I} = \bigcup_{W, Y \in \mathcal{M}; X \in \mathcal{H} \cap e_{\lceil i/3 \rceil}} I_{WYX}$. This corresponds to undesirable potential locations for reflector C_i that would force D_i to coincide with a point $X \in \mathcal{H} \cap e_{\lceil i/3 \rceil}$, hence increasing the degree of some $Q \in \mathcal{D}$.

The following theorem gives an upper bound on the size of these sets.

Theorem 3 (Bezout's Theorem for Conics [2]). *Any two distinct conics intersect in at most four points.*

It follows then that there is only a countable set of points of degree 2. In fact, if we restrict ourselves to conics of cross ratio at most $3n$ then there are $O(n^4)$ points in H . With these sets in hand we can now describe the algorithm:

For each $i = 1, \dots, 3n$, insert one-by-one the reflectors A_i , B_i , and C_i of each quadruple while avoiding the sets I_{XYZ} in (3) above and the sets $\mathcal{H} \cap \partial P$ and \mathcal{I} . This can always be done as these sets are discrete and the positions of A_i , B_i and C_i are continuously varying. After this process, the location of D_i is fixed. Because of step (5) above, we know that for any triple XYZ not containing C_i , the conics of ambiguity through $XYZD_i$ do not increase the degree of points in \mathcal{D} . Nevertheless, it is indeed possible for a quadruple of the form $C_i D_i XY$, with $X, Y \in \mathcal{M}$, to increase the degree of a point in \mathcal{D} .

Let Q be the point in \mathcal{D} whose degree is three after inserting D_i (refer to Figure 5). Using the notation of this figure, we see that Q lies in the intersection of two conics (ellipses), and hence is of degree two. To remove this ambiguity, we must move C_i and D_i in such a way that the cross ratio $\{A_i B_i C_i D_i\} = i$ remains

constant yet the cross ratio $Q\{C_i D_i X Y\}$ is no longer an integer. First note that $Q\{C_i D_i X Y\} = \{C_i D_i X' Y'\}$. Without loss of generality, we introduce a real axis coordinate system on edge $e_{\lceil i/3 \rceil}$ such that A_i is the origin and $B_i A_i = 1$. Let c, d, x, y denote the position on the real axis of C_i, D_i, X', Y' respectively. Hence $\frac{1}{c-1} \Big/ \frac{d}{d-c-1} = 3 \implies d = \frac{c+1}{1-3c+3}$. Similarly $\frac{d-c}{x} \Big/ \frac{d-c+x+y}{y} = kx \implies d = \frac{yc-kxc+kx^2+kxy}{y-kx}$. Equating the two expressions for d , we obtain a quadratic expression on c with at most two solutions for each integer k . That is, the set of positions C_i that force a D_i to make Q ambiguous is a discrete set and hence it can be avoided by perturbing C_i by an ϵ amount.

Theorem 4. *If g is the number of wall guards required to guard a polygon P , then a robot can localize itself using at most $10g$ reflectors.*

Proof. In the previous algorithm, for a guard g_j we can identify $D_{3j} = A_{3j+1}$ and $D_{3j+1} = A_{3j+2}$, thus reducing reflectors per guard from 12 to 10. \square

4 Lower Bound

In this section we prove that, in the worst case, at least $4g - 2$ reflectors are required for unambiguous robot localization for a polygon guardable by g (wall) guards. This lower bound holds even if the reflectors are not on the walls of the polygon. The example is the standard *comb polygon* shown in Figure 6.




Fig. 6. Two types of ambiguous positions.

Theorem 5. *If g is the number of wall guards required to guard a polygon P , then we need at least $4g - 2$ reflectors to uniquely localize the position of a robot in P .*

Proof. It follows from the discussion in Section 2 that we need at least three reflectors per tine. Suppose that three or more tines had just three reflectors. Then either two of these tines have the three reflectors collinear, or two of the tines have the three reflectors noncollinear. If the three reflectors are collinear in at least two tines, then for a position arbitrarily close to the midpoint between two reflectors in such a tine, the angles observed are an almost- 180° angle and an almost-zero angle. This angular configuration can be realized in both of the tines with three collinear reflectors, and therefore it is ambiguous.

Alternatively, if at least two of the three-reflector sets are not collinear, then drop a perpendicular from the vertex opposite the longest edge of the triangle formed by three reflectors, and obtain a point for the robot that reads angles 90° , 90° , and 180° . This applies to both tines and therefore the position is ambiguous. \square

5 Other Localization Primitives

The localization problems may become a simpler task if the robot can benefit from alternative, independent orientation mechanisms. In real life, the robot moves on the floor, which is a 2-D surface embedded in a 3-D space. We can take advantage of this fact by using the third dimension to place reflectors. In this model, the robot can perform 360-degree scans along any given chosen plane through its current position. The robot might also perform a 2-D-like scan by performing a laser sweep along the horizontal plane of height zero. The robot is equipped with a device that indicates the “up” direction (defining the orientation of the floor plane) at all times. We omit proofs in this abstract.

Theorem 6. *If n is the number of vertices in a polygon P , then a robot aided by a compass indicating a North position at a point at infinity requires at least $n/4 - 8$ reflectors in the worst case to uniquely indentify its position in P .*

Theorem 7. *Consider a robot on a plane in a 3-D environment with walls, given an “up” direction and a map of the environment. Let g be the number of guards required to guard such environment. Then at least $4g$ reflectors and at most $6g$ reflectors are needed for the robot to localize itself in the environment.*

6 Conclusions

We have developed a method to remove ambiguities from Sugihara’s reflector model for robot localization. This model can be implemented economically both in terms of hardware (laser, reflectors) as well as computational requirements. We have given nearly matching upper and lower bounds on the number of reflectors needed per guard. We also considered alternative scenarios for localization with more general primitives and showed upper and lower bounds in these contexts.

Acknowledgments. We thank Martin Demaine for many helpful discussions.

References

1. D. Avis and H. Imai. Locating a robot with angle measurements. *Journal of Symbolic Computation*, 10(3–4):311–326, 1990.
2. M. Berger. *Geometry*, vol. II. Springer-Verlag, 1987.
3. M. Betke and L. Gurvits. Mobile robot localization using landmarks. In *Proc. IEEE International Conference on Robotics and Automation*, vol. 2, 135–142, 1994.
4. G. Dudek, K. Romanik, and S. Whitesides. Localizing a robot with minimum travel. In *Proc. 6th ACM-SIAM Symposium on Discrete Algorithms*, 1995.
5. H. Gonzalez-Banos and J. Latombe. A randomized art-gallery algorithm for sensor placement. In *Proc. 17th ACM Symposium on Computational Geometry*, 2001.
6. S. T. L. Heath. *The thirteen books of Euclid’s Elements translated from the text of Heiberg with introduction and commentary*. Dover Publications, New York, 1956.
7. L. S. Shively. *Introduction to Modern Geometry*. John Wiley & Sons, 1939.
8. R. Sim and G. Dudek. Mobile robot localization from learned landmarks. In *Proc. IEEE/RSJ Conference on Intelligent Robots and Systems*, 1998.
9. K. Sugihara. Some location problems for robot navigation using a single camera. *Computer Vision, Graphics, and Image Processing*, 42(1):112–129, 1988.
10. O. Veblen and J. W. Young. *Projective Geometry*, vol. 1. Ginn and Company, 1938.