Pursuing Projections: Keeping a Robot on Path*

Karen T. Sutherland
Computer Science Department
University of Minnesota
Minneapolis, MN 55455

William B. Thompson
Department of Computer Science
University of Utah
Salt Lake City, UT 84112

Abstract

For an autonomous robot navigating in an unstructured outdoor environment, staying close to a path is crucial to successfully reaching its goal. Although the degree of accuracy with which it estimates its own location affects its ability to stay on the path, accuracy in estimate of lateral distance from the path is far more important for successful navigation than accuracy in estimate of position along the path. Utilizing methods based only on relative angular measurements between landmarks in the environment, we draw from techniques used in statistical pattern recognition to show how landmarks can be chosen for localization which will not only give good estimate of location in spite of the measurement error, but will also keep the robot on the path. We demonstrate how identical landmark configurations can produce very different results in localizing to a path and show how simple heuristics can be used to choose the best configuration for path localization.

1 Introduction

A robot navigating with the aid of a map must be able to estimate its own position as accurately as possible. We have previously shown [10, 11] how heuristics can be used to choose landmarks for localization which will reduce the error in location estimate. However, it is just as important, if not more so, to localize to a path. An autonomous robot following a path could err in estimating where it is on the path and still successfully reach its goal, but a poor estimate of lateral distance from that path might have severe consequences. This is particularly true in an outdoor, unstructured environment where straying from a path could easily lead to vehicle destruction.

Dependable localization methods in this type of environment are few. Actual distances to landmarks are often impossible to estimate. Absolute bearings, accurately registered to a map, are frequently unavailable. For these reasons, we are addressing this "path" localization problem using methods based only on relative angular measurements between landmarks.

It is known [5, 6, 8, 9] that, when landmarks have been identified and matched to a map and landmark order in the view is known, accurate measurement of visual angle (i.e., the angle formed by the rays from the viewpoint to a pair of point landmarks) using three point landmarks in the environment is sufficient to determine the viewpoint unless all landmarks plus the viewpoint lie on a single circle. When an error in measurement is introduced, the estimated viewpoint is constrained to an area on the terrain surrounding the true viewpoint. We call this area the *area of uncertainty*. In this case, [10, 11] the error in estimate of location is significantly affected by the landmarks chosen for localization.

We now draw from techniques used in statistical pattern recognition to show how heuristics can be used to choose landmark configurations which will produce good path localization.

2 Statistical Projection Pursuit

A popular method of analyzing multivariate data is to find a low-dimensional projection which will reveal characteristics which are not apparent in the higher dimension. This process is aptly termed "projection pursuit" because the direction in which the projection is to be made must be determined [1, 2, 3, 4]. For many high dimensional data clouds, the majority of the lower dimensional projections are approximately Gaussian. In this case, the interesting information in a pattern recognition problem is obtained by projecting the data onto an axis so that a single Gaussian distribution is *not* produced. However, the opposite can also be true. Diaconis and Freedman [1] showed that there are classes of data sets where the interesting projections are the ones which are close to Gaussian. It is this type of projection which we will pursue.

*This work was supported by National Science Foundation grant IRI-9196146, with partial funding from the Advanced Research Projects Agency.

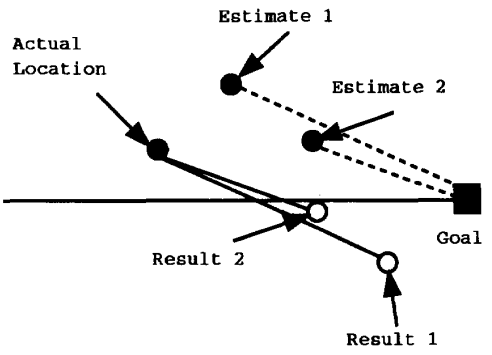

Figure 1: The dashed lines show computed paths. The solid lines show actual paths. The resulting locations are shown by empty circles. If the goal is to stay as close to the path as possible, Estimate 2 is preferable to Estimate 1.

The error in location estimate can be divided into two components: error in lateral distance from the path and error in distance along the path. As shown in Figure 1, one estimate of location can be closer to the true location than another but have greater error in lateral distance from the path. Distance and direction of movement are based on that estimate. The navigator will remain closer to the path when the error in estimate of lateral distance from the path is minimal. To find good configurations for path localization, the two-dimensional area of uncertainty can be projected into one-dimensional space. Knowledge of one-dimensional distribution is then used to chose a triple of landmarks which will provide good path localization. Whereas in classical projection pursuit, the two-dimensional data cloud is given and an appropriate one-dimensional projection must be found, this "inverse" problem can be described as finding a configuration so that the resulting two-dimensional data cloud produces the desired projection onto a given axis.

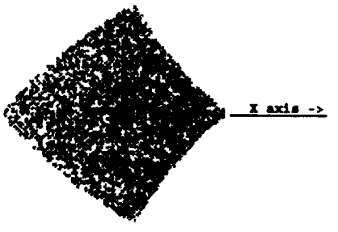

Figure 2: A two-dimensional area of uncertainty. The actual viewpoint is at the large black point.

3 Pursuing Projections

Consider the area of uncertainty shown in Figure 2. Assuming a uniform distribution of error in the visual angle measure, each small point represents 1 of 10,000 iterations. The true viewpoint is at the large black point. Although there is no central tendency in two-dimensional distribution, the distribution of points along the x-axis as shown in the graph on the left of Figure 3 shows a definite central tendency.

However, if the robot is heading at a 45° angle counter-clockwise from the positive x-axis, the distribution of its location, as shown in the graph on the right of Figure 3, is close to uniform across the path. Thus, whether or not there is any central tendency in a one-dimensional sense depends on the direction of the one-dimensional slice.

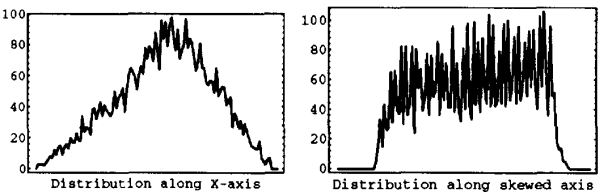

Figure 3: Distribution of points along x-axis and along an axis rotated 45° counterclockwise from the x-axis.

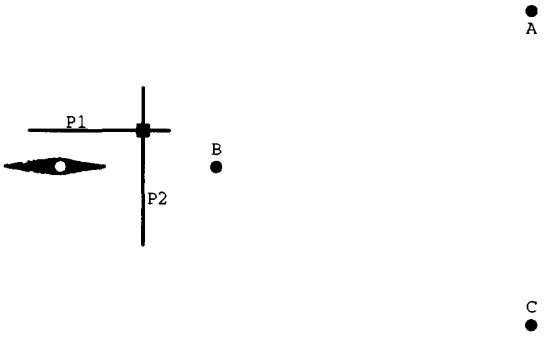

Figure 4: A navigator located at the white dot and moving to the black square will have better path localization to path P1 than to path P2.

A second example in Figure 4 shows an elongated area of uncertainty. The landmarks used for localization are at A, B, and C. Error bound was ±10% of visual angle measure. The actual viewpoint is at the white dot surrounded by the area of uncertainty. When the area is projected onto an axis perpendicular to path P1, a much smaller variance results than when

it is projected onto an axis perpendicular to path P2. Although the actual error distribution is the same, a navigator would have better localization to path P1 than to path P2.

Our goal: to find an area of uncertainty which has a Gaussian distribution with small variance when projected onto an axis perpendicular to path direction.

There are two geometric properties which must be considered when pursuing projections:

- The smaller angle between the path and the axis of the orientation region, "axis" being the line passing through landmark B and the midpoint of line segment AC. This angle can be measured exactly with knowledge of landmark configuration and path location.
- The ratio W/L where L = the diameter[1] of the area of uncertainty in direction of landmark B and W = the diameter of the area of uncertainty perpendicular to direction of landmark B. This ratio gives a measure of the "fatness" of the area of uncertainty. An area such as that in Figure 4 is "thin". In most cases, the ratio W/L can be estimated with knowledge of landmark configuration, path location and orientation region in which the viewpoint lies.

The properties are not independent. The goodness of a particular configuration in terms of localization to path depends on both.

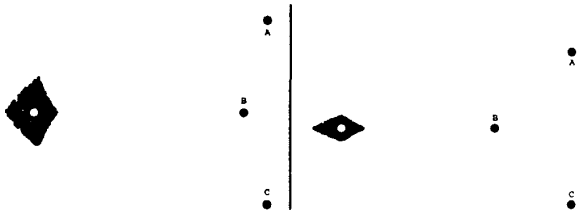

Figure 5: Configurations used to show how orientation to path affects path localization. Landmarks are at A, B and C. Viewpoint is at the white dot. Orientations for the left configuration are shown in Figure 6 and for the right configuration in Figure 7.

3.1 Effect of path-axis angle

To show how the angle between the path and axis of the orientation region affects path localization, two sets of simulations were run, the first using the configuration on the left and the second using the configuration on the right of Figure 5. In each set of simulations, trials of 1000 runs each were done. Visual angles to all configurations were the same. Location was estimated and a move made toward the goal.

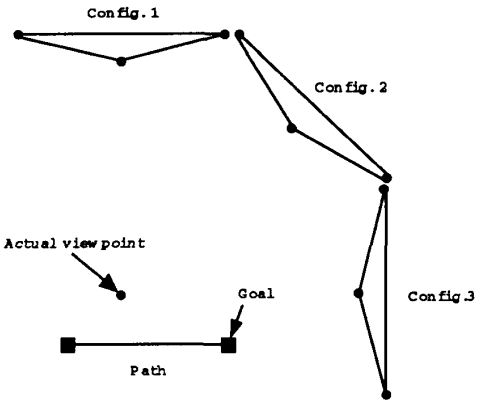

Figure 6: Center landmark was 225 linear units and line joining outer landmarks was 250 linear units from the actual viewpoint for all three configurations. Angle from the viewpoint to all outer landmark pairs was 43.6°. As shown in Table 1, path localization is best for Configuration 1.

Config.	Path/Axis Angle	Error in Angle	Mean Dist. to Goal	Mean Dist. to Path
1	0°	10%	18.27	9.43
		20%	36.58	19.27
		30%	55.00	29.80
2	45°	10%	17.99	11.74
		20%	36.00	23.54
		30%	54.11	35.40
3	90°	10%	18.17	12.90
		20%	36.35	25.72
		30%	54.58	38.43

Table 1: Results of simulated trials using the three configurations shown in Figure 6. Distances are given in linear units.

The only parameter which varied was the angular relationship of the configuration to path direction. A uniform error of 10%, 20% and 30% of visual angle measure was introduced using an implementation of the Wichmann-Hill Algorithm [12]. Throughout this analysis, we have assumed a uniform error distribution in visual angle measure. By assuming a large error, uniformly distributed, whatever error does occur is likely to be a subset of that which was assumed. The decision to use a multiplicative rather than an additive error bound was based on the fact that, although errors due to factors such as quantization are additive, those due to lens distortion and camera movement are

[1] "Diameter" is defined as length of the longest line segment joining two points on the boundary of the area of uncertainty in a specified direction.

multiplicative. The latter will in most cases dominate, particularly for large visual angles. The results for the first set of simulations are shown in Table 1. The positions of the configurations are shown in Figure 6. Although, as expected, mean distance to goal does not differ as the configuration is rotated, Configuration 1 provides better path localization than Configuration 3. Results for the second set of simulations are shown in Table 2. Configuration positions are shown in Figure 7. In this case, Configuration 3 provides better path localization than does Configuration 1.

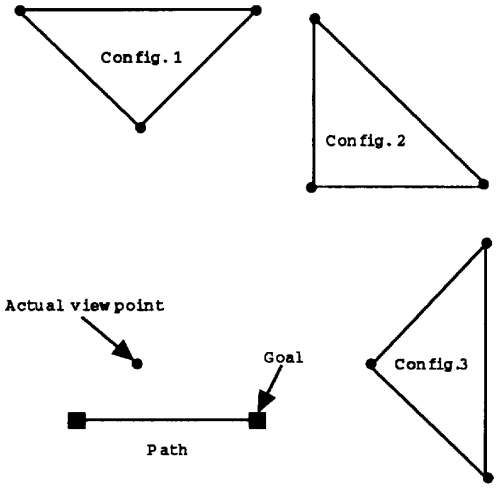

Figure 7: Center landmark was 200 linear units and line joining outer landmarks was 300 linear units from the actual viewpoint for all three configurations. Angle from the viewpoint to all outer landmark pairs was 36.6°. As shown in Table 2, path localization is best for Configuration 3.

Config.	Path/Axis Angle	Error in Angle	Mean Dist. to Goal	Mean Dist. to Path
1	0°	10%	13.62	10.89
		20%	27.45	21.95
		30%	41.72	33.32
2	45°	10%	13.47	8.87
		20%	27.16	17.88
		30%	41.29	27.19
3	90°	10%	13.57	5.46
		20%	27.31	11.02
		30%	41.43	16.83

Table 2: Results of simulated trials using the three configurations shown in Figure 7. Distances are given in linear units.

3.2 Effect of W/L ratio

Changing the angle between path and axis of the orientation region changes path localization but how it is changed depends on the shape of the area of uncertainty, which can be estimated by the ratio W/L. In most cases, an estimate of this ratio can be obtained in the following way: for any landmark triple A, B and C, the circle through A, B and the viewpoint V intersects the circle through B, C and V at points B and V. It is well known [7] that when two circles intersect, there is only one angle of intersection, the same at both intersection points. Thus, the angle at V equals the angle at B. Levitt and Lawton [6] called the lines joining pairs of landmarks *landmark pair boundaries* or LPB's. We will use the LPB's joining AB and BC (and intersecting at B) to put a bound on the angle of intersection of the circles passing through A, B and V and B, C and V.

If landmark B lies closer to the viewpoint than does the line segment joining landmarks A and C, as shown in Figure 8a, the angle of intersection of the circles, equal to the angle of intersection of the tangents (dashed lines) at that point, cannot be larger than the angle of intersection of the LPB's. This is due to the fact that the limits of the slopes of the chords AB and BC as A and C approach B equal the slopes of the tangents at B. In this case, which can easily be determined when orientation region is known, an upper bound is placed on that angle. Thus, the angle of intersection of the circles at V, which is unknown, is bounded above by the angle of intersection of the LPB's, which is known.

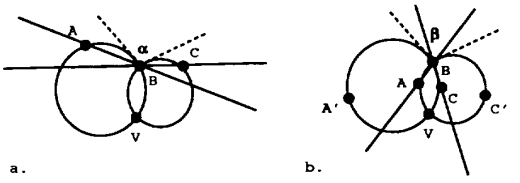

Figure 8: In a., the angle α of intersection of the circles cannot be greater than the angle of intersection of the LPB's. In b., the angle β of intersection of the circles cannot be less than the angle of intersection of the LPB's.

If landmark B lies further from the viewpoint than does the line segment joining landmarks A and C, A and C can lie on the inner circular arcs as shown in Figure 8b, but they could also lie on the outer circular arcs, as do A′ and C′ in Figure 8b. When they lie on the inner circular arcs, the angle of intersection of the circles cannot be less than the angle of intersection

of the LPB's. In this case, a lower bound is placed on that angle. If they lie on the outer circular arcs, the angle of intersection of the LPB's cannot be used as a bound on the angle of intersection of the circles. The tangents are not limits of the chords because the chords are on one circle and the tangents are on the other.

To determine if A and C lie on the inner circular arcs, resulting in a lower bound for the angle of intersection of the circles, consider a circle passing through the three landmark points, as shown in Figure 9. The measure of the inscribed angle γ can be easily computed from landmark positions. For any viewpoint inside the circle, such as V_1 in the figure, the visual angle to landmarks A and C is greater than γ. For any viewpoint outside the circle, such as V_2 in the figure, the visual angle to A and C is less than γ. If the viewpoint was on the circle, the configuration would be single circle (i.e., all landmarks and viewpoint on one circle). We are assuming, as will be shown in Section 4, that anything close to a single circle configuration has been eliminated as too error prone to use. It follows that the estimate of the visual angle to A and C will either be significantly less than or significantly greater than γ. Once the viewpoint location in relation to the circle is determined, the inner-outer circular arc question is also answered. As shown in Figure 9, a circle through A, B and V_1 is smaller than the circle through A, B and C, with the result that C is then on the outer circular arc. A circle through A, B and V_2 is larger than the circle through A, B and C, with the result that C is then on the inner circular arc. The same argument holds for the circle through B, C and either V_i.

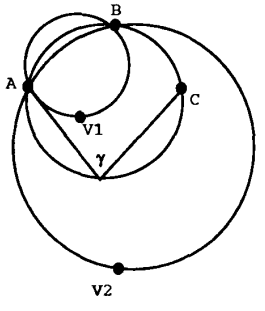

Figure 9: Visual angle to A and C is larger than γ at V_1 and smaller than γ at V_2.

The angle of intersection of the circles at V affects the ratio W/L because the area of uncertainty is formed by intersecting thickened rings surrounding the circles, the thickness of the rings determined by the amount of the error [9]. As shown in Figure 10, the ratio W/L is proportional to the angle of intersection of the circles. Landmarks are at A, B and C. The area of uncertainty surrounds the actual viewpoint V. If the angle of intersection of the circles is small, the area of uncertainty will be "thin". If the angle is large, the area will be "fat". Thus, the limits imposed by the LPB's provide a heuristic for estimating the shape of the area of uncertainty. A small upper bound produces a small angle of intersection of the circles which, in turn, produces a "thin" area of uncertainty.

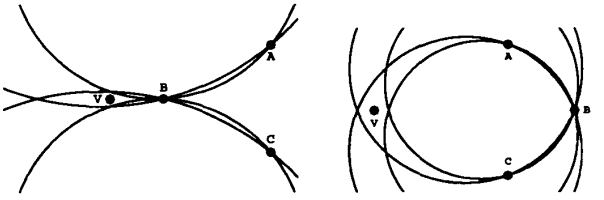

Figure 10: The angle of intersection of the circles affects the shape of the area of uncertainty.

To summarize, we can determine, using landmark configuration alone, whether or not the LPB's provide a bound for the angle of intersection of the circles. When they provide an upper bound, the more acute the angle of intersection of the LPB's, the smaller the ratio W/L will be, providing good path localization for a path heading toward landmark B. When they provide a lower bound, the more obtuse the angle of intersection of the LPB's, the larger the ratio W/L will be, providing good path localization for a path heading perpendicular to landmark B. This information will be used to choose landmark configurations which will give good path localization.

4 Choosing good configurations

The function given below (and described in detail in [11]) was first used to rank configurations for general goodness, based on the size of the area of uncertainty. This function approaches zero as the configuration approaches single circle, the one case when the viewpoint cannot be uniquely determined. The smaller the area of uncertainty, the larger the function value. In Figure 11, let $A = (Ax, Ay, Az)$, $B = (Bx, By, Bz)$, $C = (Cx, Cy, Cz)$, $V = (Vx, Vy, Vz)$ be the projections of the landmark points and V_0 on a horizontal plane. Let I be point of intersection of the line through V and B with the circle through A, C, and V; L be point of intersection of the line through A and C with the line through V and B; and $d(p, q)$ be distance between any two points p and q.

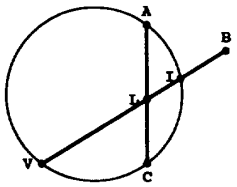

Figure 11: Simple geometric relations can be used to rank landmark configurations.

Then:

$$G(A, B, C, V_0) = h + f$$

where

$$h = \left(\left|\frac{\frac{(Az+Bz+Cz)}{3} - Vz}{Vz}\right| + 1\right)^{-1}$$

$$f = \begin{cases} \frac{2(d(V,B)-d(V,I))}{d(V,I)*\pi} * \frac{d(A,C)}{d(V,L)} & if\ d(V,B) \geq d(V,I) \\ \frac{d(V,I)-d(V,B)}{d(L,I)} * \frac{d(A,C)}{d(V,L)} & if\ d(V,L) \leq d(V,B) < d(V,I) \\ \frac{d(A,C)}{d(V,B)} & if\ d(V,B) < d(V,L) \end{cases}$$

This function consists of two parts. The h function weighs the elevation of the landmarks compared to the elevation at point V_0. It is non-negative and attains its maximum of 1 when the average elevation of the landmarks is equal to the elevation at V_0. The f function, also non-negative and defined piecewise, has the major effect on the goodness measure. It is based on the size of the area of uncertainty for the projected points. Note in Figure 11 that as B approaches the circle, the measure approaches zero. If B lies on line AC, the measure is the ratio $\frac{d(A,C)}{d(V,L)}$. The function increases in value as B is pulled away from the circle and the estimated viewpoint. The factor of $\frac{2}{\pi}$ in the first piece of f causes this increase to occur at a rate such that when B is pulled back to the point that the area of uncertainty is the same size as for a straight line configuration, the function value is the same. The function increases in value as B moves nearer the viewpoint.

We then augmented the original function with an additive factor p to weight path goodness. This factor was added only if the goodness measure for point localization was above a given threshold. In this way, the path goodness was not able to add weight to a configuration which was generally poor for localization. This factor is determined in the following way.

If an upper bound exists:

$$p = k * \left(\frac{1}{\alpha * \beta + 1} - \frac{1}{\frac{\pi^2}{2} + 1}\right) * \frac{1}{1 - \frac{1}{\frac{\pi^2}{2}+1}}$$

If a lower bound exists:

$$p = k * \frac{2.0 * \alpha * \beta}{\pi^2}$$

where α is the angle between landmarks A and C with vertex B and β is the computed angle between the axis of the orientation region and the path. This factor ranges from 0 to 1, with 0 being poor for path localization and 1 being good. The constant k should be set to a weighting factor based on how important path localization is. If path localization is not desired, $k = 0$.

5 Experimental results

A sequence of simulations were run, using U. S. Geological Survey 30 meter Digital Elevation Map (DEM) data with a goal of keeping the robot navigator as close to the path as possible. The contour map in Figure 12 shows the location of one set of runs. The area is approximately 21 kilometers in the east-west direction by 14 kilometers in the north-south direction just northeast of Salt Lake City, Utah. UTM coordinates of the southwest corner are 426180E, 4511040N. The path, shown on the map, runs along the banks

Figure 12: Contour map of area in which simulations were run. The path is shown by a heavy black line. Start position is at left. Goal is at star.

of the creek through City Creek Canyon. Start position and each point at which a new image was taken are shown by squares. The goal is marked by a star. Landmarks are marked by filled black circles. The assumption was made that the point landmarks (in this case, mountain peaks) had been identified and

matched to the map. The landmarks used for localization were chosen only from those visible at any given location. The rugged terrain coupled with the proximity of path to creek makes this a classic example of the type of situation where path localization is important. The total length traveled is about 8000 meters or 5 miles. In each run, the navigator takes a new reading at approximately 1 mile intervals. Uniform error bounds of $\pm 5\%$, $\pm 10\%$ and $\pm 20\%$ in visual angle measure were introduced. Due to the spread between most landmarks, these limits produced errors in angle measure which were quite large in angular value. For this particular set of experiments, 60% of the possible configurations were bounded either above or below. Only those configurations which had a goodness measure greater than 1 were considered for the additional path weighting. Although this threshold worked well in our simulations, it would also be possible to choose, for example, the configurations which were in the top 30% of the rankings for point localization, eliminating the need for a numeric threshold. The constant k was set to 5. Again, although this worked well in all our simulations, its value could also be chosen based on the range of goodness measures for point localization. The results are shown in Table 3. Adding the heuristic for path localization did not change the total distance traveled. Final distance from goal averaged 10% better when the heuristic was used.

Error in Angle	Heuristic for Choosing Landmark Triple		
	None used	Point only	Point and Path
5%	160 m.	90 m.	70 m.
10%	270 m.	205 m.	162 m.
20%	426 m.	408 m.	393 m.

Table 3: Results of runs through City Creek Canyon. Mean distance to path (in meters) is given for three different bounds in angular measure error. Fifty trips were recorded with each error bound.

6 Conclusions

We have inverted the classic statistical projection pursuit problem to show that simple heuristics can be used to keep a robot which is navigating using only relative angle measure closer to its path. Such heuristics are easy to implement and need only be applied if the configuration of landmarks has already been judged good for point localization, keeping the added computation to a minimum. Our experiments using USGS 30 meter DEM data have shown that, even with only a few landmarks from which to choose configurations, path localization can be improved by pursuing good projections.

References

[1] Persi Diaconis and David Freedman. Asymtotics of graphical projection pursuit. *The Annals of Statistics*, 12(3):793–815, September 1984.

[2] Jerome H. Friedman. Exploratory projection pursuit. *Journal of the American Statistical Association*, 82(397):249–266, March 1987.

[3] Nathan Intrator. Feature extraction using an unsupervised neural network. *Neural Computation*, 4:98–107, 1992.

[4] Nathan Intrator. On the use of projection pursuit constraints for training neural networks. In C. L. Giles, S. J. Hanson, and J. D. Cowan, editors, *Advances in Neural Information Processing Systems*, volume 5, pages 3–10. Morgan Kaufmann, 1993.

[5] Eric Krotkov. Mobile robot localization using a single image. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 978–983. IEEE, 1989.

[6] Tod S. Levitt and Daryl T. Lawton. Qualitative navigation for mobile robots. *Artificial Intelligence*, 44(3):305–360, August 1990.

[7] Daniel Pedoe. *A Course of Geometry*. Cambridge University Press, 1st edition, 1970.

[8] Kokichi Sugihara. Some localization problems for robot navigation using a single camera. *Computer Vision, Graphics, and Image Processing*, 42:112–129, 1988.

[9] Karen T. Sutherland. Sensitivity of feature configuration in viewpoint determination. In *Proc. DARPA Image Understanding Workshop*, pages 315–319, January 1992.

[10] Karen T. Sutherland. Landmark selection for accurate navigation. In *Proc. DARPA Image Understanding Workshop*, pages 485–490, April 1993.

[11] Karen T. Sutherland and William B. Thompson. Inexact navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1–1 – 1–7. IEEE, May 1993.

[12] B. A. Wichmann and I. D. Hill. An efficient and portable pseudo-random number generator. *Applied Statistics*, 31:188–190, 1982.