Localizing in Unstructured Environments: Dealing with the Errors

Karen T. Sutherland and William B. Thompson, *Senior Member, IEEE*

Abstract— A robot navigating in an unstructured outdoor environment must determine its own location in spite of problems due to environmental conditions, sensor limitations and map inaccuracies. Exact measurements are seldom known, and the combination of approximate measures can lead to large errors in self-localization. The conventional approach to this problem has been to deal with the errors either during processing or after they occur. We maintain that it is possible to limit the errors before they occur. We analyze how measurement errors affect errors in localization and propose that a simple algorithm can be used to exploit the geometric properties of landmarks in the environment in order to decrease errors in localization. Our goal is to choose landmarks that will provide the best localization regardless of measurement error, determine the best areas in which to identify new landmarks to be used for further localization and choose paths that will provide the least chance of "straying." We show the result of implementing this concept in experiments run in simulation with USGS 30 m DEM data for a robot statically locating, following a path and identifying new landmarks.

I. INTRODUCTION

DETERMINING locations in the environment is a major task for an autonomous mobile robot navigating with a map. It must match the locations of specific objects on the map to sensed objects in the view, add the locations of new unmapped objects to that map and continually update its own location as it moves. An outdoor, unstructured environment presents unique challenges to a robot navigator, not the least of which involve these location determination tasks:

1) Errors in distance traveled can be significant and unpredictable, compounding as the robot moves. Whether the robot is wheeled [23] or legged [6], terrain surface conditions exacerbate these errors for land vehicles. Kosaka and Kak, in their very thorough treatment of indoor navigation techniques [7], state that occasional wheel slippage inconsistent with the mathematical model was the believed cause of failure in their hallway experiments. In an unstructured environment, this type of error occurs frequently.

2) The sensors commonly used in indoor navigation do not have a large enough range to be very useful outdoors. The accuracy of a compass or barometric altimeter

K. T. Sutherland was a graduate student at the University of Minnesota. She is now with the Department of Computer Science, University of Wisconsin-La Crosse, La Crosse, WI 54601 USA.
W. B. Thompson is with the Department of Computer Science, University of Utah, Salt Lake City, UT 84112 USA.
IEEE Log Number 9405742.

[23] is often affected by conditions such as magnetic fields or atmospheric pressures in the environment to the extent that these devices are not reliable. Thus, absolute bearings, registered to a map, may not be available. In that case, triangulation, a common method for determining locations that requires absolute bearings to two or more landmarks, cannot be used. The Navstar Global Positioning System (GPS) has often been promoted as an error free solution to the problem of self-localization. However, problems also exist with the use of GPS. The signals of the GPS satellites are corrupted by data noise, multipath errors, clock errors, atmospheric delays and instrument delays [29]. Averaging techniques such as the Extended Kalman filter are used to add differential corrections to the GPS measurements and improve accuracy [15]. When using the Kalman filter in this type of application, the dynamic and observation models that are assumed are often incorrect due to causes such as cycle slips, leading to significant errors or possible nonconvergence in the filter results [12]. Atmospheric conditions also affect GPS. The signal strength is a function of the thickness of the air mass that the signal passes through [1]. Experiments have been conducted [26] during periods of snowfall and when ice clouds were present in the atmosphere. The phase delays caused by the ice clouds and snow were calculated. These errors vary according to such difficult to measure parameters as ice crystal size, density and orientation or the "wetness" of the snow. Ice clouds and snowfalls can lie between the signalling satellite and the receiver when weather conditions around the receiver are good and landmarks are in clear view. The visible horizon when GPS is used for navigation on land is approximately 25° rather than the 5° assumed when a vehicle is at sea. This can be even worse in an urban or mountainous setting [13]. Duerr [3] showed that a GPS receiver located in a valley can have a significant loss of accuracy for long periods of time.

3) The objects that must be matched in an unstructured environment typically possess features that may be deceptively masked in the view. A common example of this is a mountain peak hidden by a subpeak. The errors caused by irregular terrain features such as this cannot be predicted or modeled.

4) Errors due to the inherent limits of the sensors themselves, such as discretization in images when a camera is the sensing device, are affected by the distance of the object from the camera and the focal length of the

lens. The wide-angle lens, often used in passive vision systems to acquire as large a view of the environment as possible, not only introduces distortion but, due to the large distance that each pixel represents, adds to discretization error. In addition, conditions such as fog surrounding a land navigator can cause image blurring and back-scattered noise much like particles suspended in the water around an undersea system [18]. This will adversely affect sensor accuracy.

5) Maps contain errors. As an example, the location of the highest point of a significant peak is 200 meters off on one of the United States Geological Survey (USGS) Digital Elevation Maps (DEM's) in our test data [24]. Navigation researchers working in small structured environments have a great deal of control over the maps they use. That is not the case for those working in large unstructured environments.

Traditionally, errors in localization have been dealt with either during processing or after the fact. Errors are modeled and/or estimated. Some type of filter or maximum likelihood function is implemented to find a good location estimate from the combined noisy measurements. Most of these methods assume a specific error distribution and/or initial solution estimate. As an example, the often used Extended Kalman filter (EKF) assumes white Gaussian noise. Although the Kalman filter is guaranteed to converge, the EKF is not. It can easily fall into a local minimum if a good estimate of the solution is not available in advance [17]. Due to the problems mentioned above, the initial solution estimates in an unstructured environment can be very poor. Even when convergence is possible, the number of measurements a filter requires to converge is often large due to the amount of error in those measurements. This factor is important when images must be taken and processed in real time.

A combination of all of these problems leads to error being a significant discriminant between navigating in structured and unstructured environments. If robot motion could be better modeled, if a variety of accurate sensors were available, if the environment surrounding the navigator was conducive to signal transmission and if the objects sensed were bar codes rather than flora covered terrain, outdoor navigation would be much easier. Unfortunately, that is not the case.

Our approach to navigating in an unstructured world has been to critically analyze the types of errors that will occur and, utilizing only those tools and techniques that are available in such environments, to exploit the geometric properties involved in the localization process so that steps can be taken to decrease the possibility of error before it occurs. The goal is to choose the landmarks that will provide the best localization regardless of measurement error, determine the best areas in which to identify new landmarks to be used for further localization and choose the paths that will provide the least chance of "straying." The result is less error and greater success for whichever filter is used to handle the error that does exist.

In the next section, we outline the problem that this work addresses, and review previous work dealing with both navigation in outdoor environments and the geometric properties that

we use. In Section III, we analyze the localization error with which our robot is faced. Section IV describes the methods that we have used to choose landmarks for localization and robot movement and shows the results of implementing these methods in navigation simulations using USGS 30 m DEM data. Section V summarizes our results.

II. OVERVIEW OF LOCALIZATION METHODOLOGY

In this paper, we are considering the problem of a robot determining its own location on a map using the method of relative bearing, which requires angular measures to three or more landmarks in the environment. We consider the terrain map features and avoid paths that would, for example, take the robot over a mountain peak. However, we have not addressed the local problem of actual physical traversability of a given path since this often requires more knowledge than a terrain map contains. Our focus has been on decreasing error in localization rather than on path or "route" planning.

In the previous section, we discussed which tools and measurements could not be dependably used. What do we have to work with?

1) We have a map and a camera. We presume that a terrain map of the environment is available and that the robot is somewhere on this map. For simplicity, the only sensing device we are assuming available is a CCD camera. This is not to say that other sensing methods could not (or should not) be used. Using a combination of several sensors may be the best approach, as long as the environment is conducive to their use. Whatever the sensing method, the result is a noisy measurement with which we must deal.

2) We have the landmarks. We are assuming that point landmarks have been identified in the view, matched to the map, and that their left to right order is known. One example of a point landmark is a mountain peak. Another example is a light beacon. Early work by Kuipers [9] showed that models of spatial memory based on sets of visible landmarks were useful in navigation. Although his TOUR model navigated in an outdoor structured setting, much of the later work in unstructured environments was based on his seminal ideas.

Levitt and Lawton [11], demonstrated that navigation can be performed in spite of poor or missing quantitative data. They define *landmarks* as distinctive visual events, meaning that a landmark must define a unique direction in three-dimensional space and be visually re-acquirable. (Landmarks stand out and don't move.) Their concept of using landmark pair boundaries (LPB's), formed by joining all pairs of known landmarks by straight lines to divide the terrain into orientation regions, allowed a robot that had identified and ordered visible landmarks to navigate qualitatively (i.e., without a single global coordinate system). As shown in Fig. 1, if the landmarks have been identified and the order of the landmarks in the view is known, the robot knows which orientation region it is in. This is due to the fact that a complete ordering of landmarks uniquely determines the orientation

region. Likewise, if it moves so that the landmark order changes, it knows that it has crossed one or more LPB's. It can thus move from one orientation region to the next as it proceeds toward a goal. This works particularly well when many landmarks have been identified, producing numerous relatively small orientation regions.

It is also a concept used by both human and insect navigators. Pick *et al.* [16] in their experiments with self-localization in human subjects found that, although people were very poor at judging actual measurements, not one subject self-located on the wrong side of an LPB. Likewise, Collett *et al.* [2] have shown that the desert ant, *Cataglyphis* spp., which does not lay pheromone trails, but navigates using landmarks, learns only which landmarks to pass and whether to pass on the left or on the right of each landmark. It is simply moving from one orientation region to another.

3) We have the visual angles. We define the *visual angle* from a navigator to two point features as the angle formed by the rays from the navigator to each feature. A perfect estimate of the visual angle measure between two points in three-dimensional space constrains the viewpoint to a surface of revolution somewhat resembling a torus [11]. When the points can be ordered with respect to the viewpoint position and *up* is known, the viewpoint is restricted to half the surface. When a two-dimensional approximation of the environment is assumed, the viewpoint is restricted to the perimeter of a slice of the torus. Likewise, knowing the order of the landmarks with respect to the viewpoint constrains the viewpoint to one arc of that perimeter. Fig. 2 gives an example of this constraint.

4) We have the geometry. The geometric relationships between the landmarks and the visual angle measures from the viewpoint to those landmarks can be combined to aid in robot localization and navigation. As shown in Fig. 3, knowing the landmark ordering and the exact visual angle measure from the viewpoint to three landmarks in two-dimensional space constrains the viewpoint to the intersection of two circular arcs. This intersection is a single point unless all three landmarks and the viewpoint lie on the same circle. Thorough analyses of error-free localization in a two-dimensional environment can be found in [8], [11], [19], [20]. Fig. 4 shows the comparable intersection of the two half torus-like figures for a three-dimensional environment. In this case, adding a third landmark does not uniquely determine the viewpoint, but constrains it to the curve of intersection shown on the right of the figure. Since it is highly unlikely that the terrain would intersect this curve in more than one point (a requirement for the location of a navigator traveling on terrain to be ambiguous), it can be assumed that knowing the visual angle to three landmarks will, in most cases, provide exact localization for a terrain-bound robot [21]. However, a vehicle not restricted to movement on the terrain may be unable to localize exactly with the use of only three landmarks.
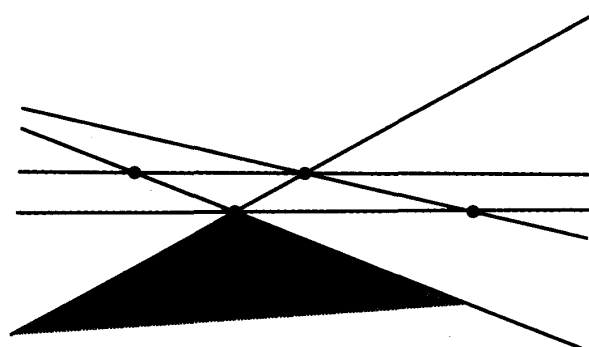

Fig. 1. Lines joining the landmarks divide space into orientation regions such as the shaded area in the foreground. Knowledge of the landmark order determines in which orientation region the robot lies.
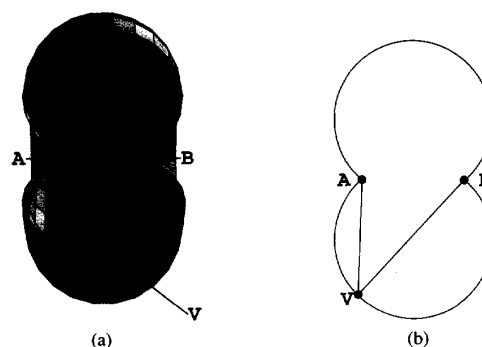

(a)　　　　　　　　　　　　　(b)

Fig. 2. Knowing the visual angle between two landmarks A and B in a three-dimensional environment restricts the viewpoint V to the surface of a torus-like figure as on the left. If a two-dimensional approximation of the environment is assumed, the viewpoint lies on the perimeter of a slice of the torus, as shown on the right. When the landmarks can be ordered with respect to the viewpoint position, the viewpoint is constrained to half the torus (if *up* is known) or to one of the circular arcs.

III. THE AREA OF UNCERTAINTY

If the visual angle measure is not exact but within a given range, the angle measures to three landmarks constrain the location to an area on the terrain or volume in three-dimensional space. We define the *area of uncertainty* to be the area in which the navigator may self-locate for a given error bound in the visual angle measure.

When a two-dimensional approximation of the environment is assumed, any given error bound in visual angle estimate will constrain the viewpoint to a thickened ring, as shown in Fig. 5(a). The thickness of the ring is determined by the amount of error [8], [11]. When three landmarks are used, any given error bound in estimate constrains the viewpoint to the intersection of two such rings [10], [11], [20].[1]

In Fig. 5(b), the visual angles from the viewpoint V to AB and BC are both 45°. The dark lines surround the area of uncertainty that represents an error bound of ±13.5° or ± 30% in both visual angles. Although the landmarks will not always be in a straight line, the visual angles will not always

[1] A third ring passing through the two landmarks lying at greatest distance from each other can be computed, but it does not affect area size.
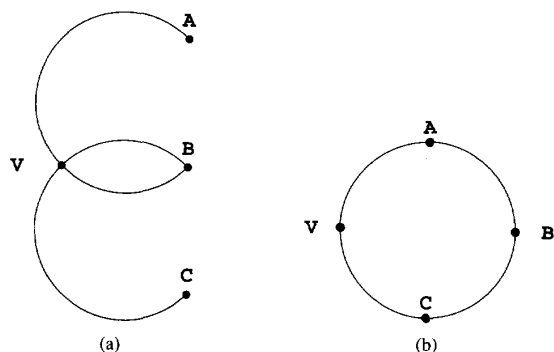
Fig. 3. Knowledge of the visual angles from the viewpoint V to three landmarks A, B and C in two-dimensional space uniquely determines the viewpoint, as shown on the left, unless all landmarks and the viewpoint lie on a single circle, as shown on the right.
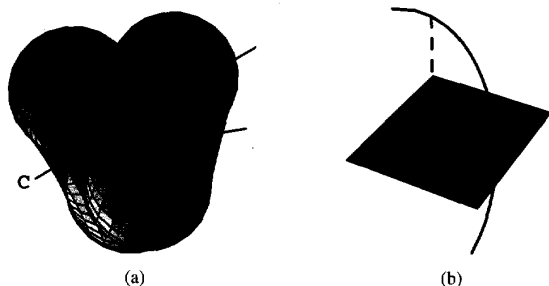

Fig. 4. Knowledge of the visual angles from the viewpoint to three landmarks in three-dimensional space constrains the viewpoint to the curve of intersection of the toroids on the left. This curve is projected to a unique point on a two-dimensional plane as on the right.
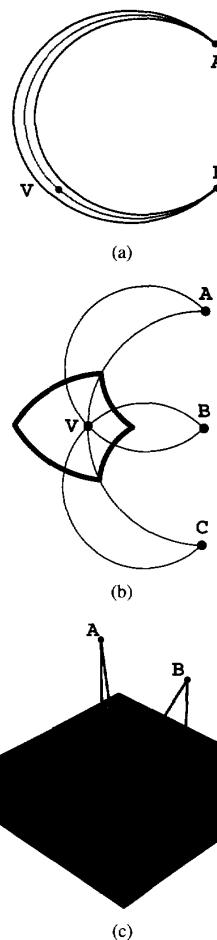

Fig. 5. In a two-dimensional approximation of the environment, the error in the visual angle estimate to two points constrains the viewpoint V to (a) a thickened ring. When three points are used, the viewpoint is constrained to (b) the intersection of two such rings. When traveling on terrain, (c), landmark elevation affects the size of this intersection.

be identical and the navigator will not always make the same error in estimate of each angle, the resulting area of uncertainty will always equal the intersection of these two rings.

For a robot traveling on terrain, the landmarks may differ in elevation from each other as well as from the navigator. This difference will affect the size and shape of the area of uncertainty. Fig. 5(c) shows an example of the difference that elevated landmarks can make in the area of uncertainty. The visual angles from the viewpoint V to AB and BC are both 45°. The smaller area on the plane is the area of uncertainty for planar angles of 45° and an error bound of ±10° or ±22% if the landmark points were at the same elevation as the viewpoint. The larger area is the actual area of uncertainty for this configuration given the same error bound in visual angle measure.

In addition to the amount of error and the landmark height, the landmark configuration alone can cause a significant change in the size and shape of the area of uncertainty. Fig. 6 shows four different configurations. The visual angles from the viewpoint to each landmark triple are identical in all four situations, as are the bounds on the measurement error. Yet, the areas differ greatly in both size and shape. If the robot had a choice of configurations to use for localization, the configuration in Fig. 6(b) should clearly be chosen. This difference is caused by the angle and location of the intersection of the two thickened rings. When all three

landmarks and the viewpoint lie on the same circle, the rings are coincident, leading to the worst possible localization.

The question arises as to whether or not better localization could be attained with more than the minimal three landmarks. Levitt and Lawton [11] showed experimentally that, in general, localization is sensitive to the number of landmarks used. However, they assumed that a large number of landmarks were available. Their results compare localization using up to 20 landmarks. Twenty landmarks provide 1140 different ordered triples. Indiscriminately adding more landmarks may eventually produce good localization, but the process of adding a landmark does not guarantee improvement. Fig. 7 shows two examples of situations where the benefit of adding more landmarks is questionable.

In the figure on the left, the innermost area corresponds to the most widely spaced landmarks, A′, B and C′, while the larger area results from localizing using the inner three landmarks, A, B and C. The error bounds were ± 30%of visual angle measure for both configurations. A multiplicative error
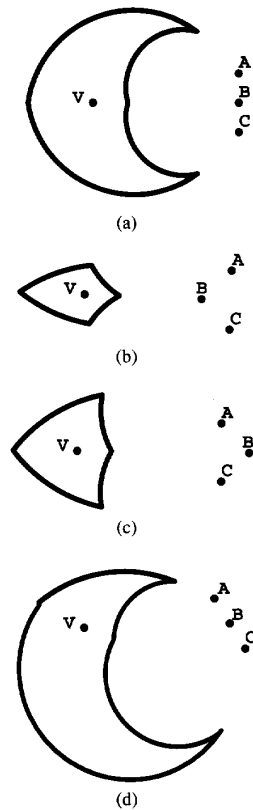
Fig. 6. The area of uncertainty will vary with landmark configuration: a) landmarks in a straight line b) center landmark pushed forward c) center landmark pushed backward d) observer off the line of symmetry of the configuration. All visual angles and error bounds are the same.

was used in order to show that the difference in area size is significant even for the most favorable error model. Because the most widely spaced landmarks, with corresponding larger visual angles, provided the best localization, an additive error model would make the difference even greater. The visual angles to A and B and to A' and B are $30°$ and $56.3°$, respectively. Thus, 30% of each produces additive error bounds of $9°$ and $16.9°$. If we used a $\pm 9°$ error bound, the smaller area, produced by A', B and C', would decrease. If we used a $\pm 16.9°$ error bound, the larger area, produced by A, B and C, would increase.

In the figure on the right, the area surrounded by 6 arcs corresponds to using landmarks A, B and C with the viewpoint inside the configuration. The area surrounded by 4 arcs corresponds to using landmarks A, B', and C. The error bounds were $\pm 12°$ of visual angle measure for both configurations. For this example, we used an additive error model. A multiplicative error would put the $AB'C$ area completely inside the ABC area. Even with using a conservative additive bound, combining estimates for the two configurations could easily lead to worse localization than using only the $AB'C$ area, particularly for a navigator attempting to stay on a narrow path while heading toward landmark B'. Whether the error is additive, multiplicative or (most likely) a combination of

the two, it is clear that some configurations will not improve localization.

To summarize, both the size and the shape of the area of uncertainty are affected by error amounts and landmark positions, lateral as well as vertical, relative to each other and to the viewpoint.

In addition to the differences in the area of uncertainty size and shape, the distribution of viewpoint estimates is not totally determined by the distribution of the error in angle measure. This is most easily seen by considering the distribution of viewpoint estimates within the area of uncertainty when the error in angle measure is uniformly distributed. As shown on the left of Fig. 8, the area is divided into sections, each representing a combination of errors in the estimations of the visual angles α, from the viewpoint to landmarks A and B, β, from the viewpoint to landmarks B and C, and γ, from the viewpoint to landmarks A and C such that $\gamma = \alpha + \beta$.

Table I shows which conditions hold in each section. For any landmark configuration, these sections are not of equal size, resulting in a non-uniform distribution of viewpoint estimates across the area of uncertainty. The relative size of each section is dependent on the landmark configuration as well as the bounds on the visual angle errors. However, the probability of a given estimate lying within a section is dependent only on the error bounds. If β ranges from $\beta - \epsilon_\beta$ to $\beta + \epsilon_\beta$ and α ranges from $\alpha - \epsilon_\alpha$ to $\alpha + \epsilon_\alpha$ where all $\epsilon_i > 0$, a joint probability distribution diagram, as shown on the right of Fig. 8, can be used to show how the probability of an estimated viewpoint lying within each of the numbered sections does not depend on the size of the sections in the figure. The numbers in the diagram correspond to the numbers of the sections in the figure. If er_i is the error in angle i, the top half of the square represents $er_\alpha > 0$, the right half of the square represents $er_\beta > 0$, and the diagonal line represents $er_\gamma = er_\alpha + er_\beta = 0$. Note that the diagram has been drawn to match the equal error bounds on α and β for this particular figure. When equality does not hold, the square becomes a rectangle. Sections 1 and 4 in the figure are significantly different in size. However, the corresponding areas in the diagram are of equal size. The probabilities can be easily computed as in the following examples:

$$\text{Prob}(\alpha, \beta, \gamma\,\text{overestimate}) = \frac{\text{Area1}}{\text{TotalArea}} = \frac{\epsilon_\beta * \epsilon_\alpha}{4 * \epsilon_\beta * \epsilon_\alpha} = \frac{1}{4}$$

or:

$$\text{Prob}(\beta, \gamma\,\text{under}, \alpha\,\text{over}) = \frac{\text{Area5}}{\text{TotalArea}} = \frac{\frac{\epsilon_\beta * \epsilon_\alpha}{2}}{4 * \epsilon_\beta * \epsilon_\alpha} = \frac{1}{8}$$

This analysis was corroborated by a number of trials, assuming a uniform error distribution in visual angle measure for several different configurations.[2] The result shown in Fig. 9 is for the same landmark configuration and error bounds used in Fig. 8. The true viewpoint is at the large black point. In the 10 000 iterations for each trial, the highest multiplicity of a single estimated viewpoint location was four. As predicted, sections 1 and 4, although quite different in

[2] The error amounts for all experiments were generated using an implementation of the Wichmann-Hill algorithm [28].
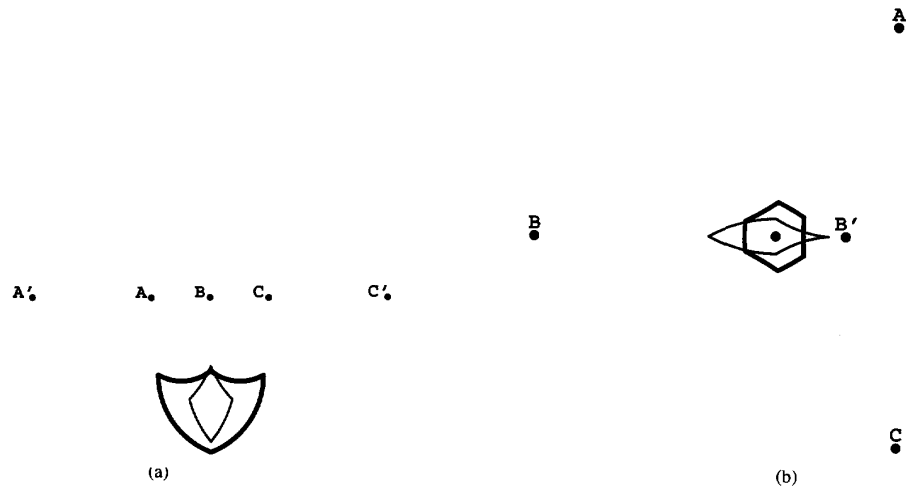
Fig. 7.   It is not necessarily the case that the more landmarks used for localization, the better. In the example on the left, the most widely spaced landmarks, A′, B and C′, produce the smallest area. In the example on the right, the area surrounded by 4 arcs resulted from using landmarks A, B′, and C.
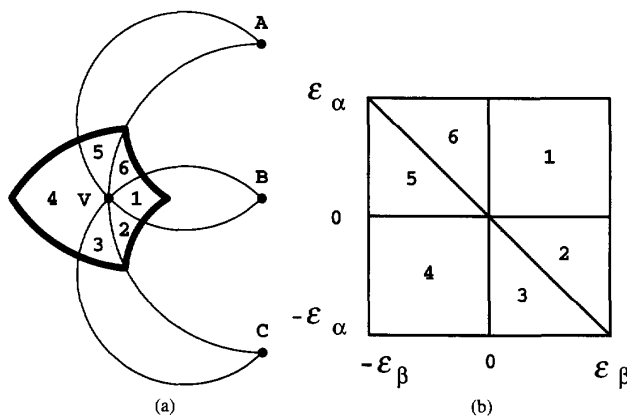
Fig. 8.   Uncertainty due to error in the visual angle measure: On the left, arcs of the circles divide the area of uncertainty into six sections. The landmarks are at A, B, and C. The viewpoint is at V. α is the visual angle from the viewpoint to A and B. β is the visual angle from the viewpoint to B and C. On the right is a diagram of the joint probability distribution. The numbers in the diagram correspond to the numbers in the figure on the left.

TABLE I
WHICH VISUAL ANGLES ARE UNDER- OR
OVERESTIMATED IN EACH OF THE SIX AREAS OF FIG. 8

Area	Overestimate	Underestimate
1	$\alpha, \beta, \gamma*$	—
2	β, γ	α
3	β	α, γ
4	—	α, β, γ
5	α	β, γ
6	α, γ	β

*Angle α is the angle from A to B with vertex V. Angle β is the angle from B to C with vertex V. Angle γ is the angle from A to C with vertex V.

size, each contained approximately one quarter of the points, leading to the less dense distribution in the left section of Fig. 9.
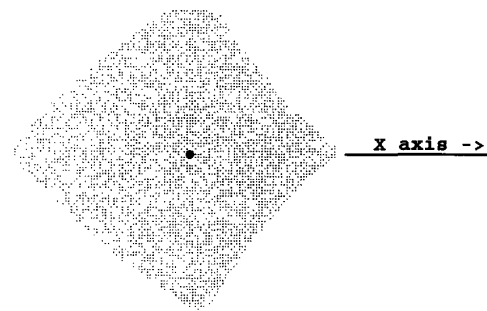
Fig. 9.   Assuming uniform distribution of error in visual angle measure, distribution of points within the areas of uncertainty is affected by the shape of that area. Each small point represents 1 of 10 000 iterations. The actual viewpoint is at the large black point. The example shown is for the area of uncertainty diagrammed in Fig. 8.

Although there is no central tendency in the two-dimensional distribution around the true viewpoint, one could assume that the viewpoint is at the origin with the x-axis horizontal through the middle of the area, as shown in Fig. 9. Consider the distribution of points along either the x or y axes, as shown in Fig. 10(a) and 10(b). There is a definite central tendency in point distribution relative to either axis. However, if the robot is heading at a 45° angle counterclockwise from the above assumed positive x-axis, the distribution of its location, as shown in Fig. 10(c), 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.

If a normal distribution of error in the visual angle measure is assumed, the distribution around the true viewpoint will exhibit a central tendency combined with the same skewing shown for the uniform distribution. In any case, if the amounts of under and overestimate in the angle measure are equal, the area in which an overestimate lies is smaller than is that for
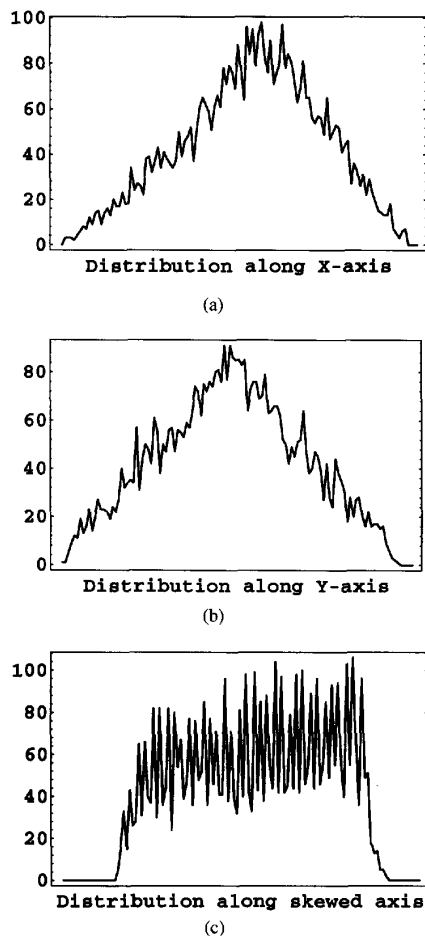
Distribution along X-axis

(a)

Distribution along Y-axis

(b)

Distribution along skewed axis

(c)

Fig. 10. The distribution of points along the x-axis, the y-axis and along an axis rotated 45° counterclockwise from the x-axis.

a corresponding underestimate, leading to the conclusion that it may be better to overestimate.

Throughout this analysis, we have assumed a uniform error distribution in the visual angle measure. This does not necessarily mean that we believe error to be uniformly distributed. A case could be made for running experiments in these environments and developing an error model such as those described by Mintz *et al.* [5], [14] for indoor use. However, many of the errors we must handle are due to environmental conditions. Weather, humidity levels, undersea thermal vents, fog, etc. will all affect the error. An error model developed under one set of conditions may not be appropriate when those conditions change. By assuming a large bound on a uniformly distributed error, we believe that whatever error occurs is likely to be a subset of that which we assumed.

IV. STRATEGIES FOR LOCALIZATION

Navigation, as defined by Gallistel [4], is the "process of determining and maintaining a course or trajectory from one place to another." This process can be viewed in terms of a cost function that includes such factors as the length of the path

or trajectory, the traversability of the path and localizability. It is the localizability (estimate of position and orientation) with which we are concerned. Using the method of relative bearing and assuming that we have a given number of visible landmarks that can be used to localize, our goal is to obtain the best localization possible at the least cost.

A part of that cost will be due to the method with which sightings are taken and processed. Although the development of active vision systems has progressed significantly in the last few years, (e.g., [27]), they remain expensive. Both active and multiple-camera passive systems are slowed down by the time required to process multiple images. A single camera system on a robot that changes orientation in order to take more views has the same problem, compounded by robot motion adding to positioning error. Our example in Section III showed that using landmarks surrounding the viewpoint to localize does not necessarily give better localization. For all of these reasons, we have chosen to deal with landmarks that are visible within a 90° angle to each side of the robot's heading.

A. Why Use Only One Landmark Triple?

As shown in Section III, the intersection of two areas of uncertainty is not always smaller than each of them. When processing sensor feedback in real time and having to possibly search for and identify new landmarks, the assumption of the more landmarks the better is not necessarily true. The navigator does not know what the intersection of the areas is. If it uses more than one landmark triple, it can only average its results. In preliminary work [21], we estimated the viewpoint in two different heading directions and averaged the results. We later found that localization with only one measure taken from a "good" landmark configuration was, in most cases, better than the average of the two.

We then ran a sequence of experiments comparing the results of using multiple ordered triples from a set of landmarks and taking only one measurement using a good triple. All experiments were run in simulation using topographic map data. The outcome of a typical run is in Table II. Triples of landmarks were picked from a group of five. A random error with uniform distribution within the given bounds was added to the angle measure for 100 iterations. The table shows the number of iterations (as a percent) in which the given combination of triples produced an estimated viewpoint closest to the actual one. This does not, of course, account for the cost of processing the additional sensor readings. In this particular run, the good triple of landmarks was also included in the 2 Triple and 5 Triple sequences. When the triples in these sequences were not very good for localization, the results were even more biased toward using the good triple alone. Fig. 11 shows the view and the map with the landmarks marked that were used for the run summarized in Table II.[3] The navigator is located at the star in the foreground. The landmarks creating the "good" triple are marked by triangles.

It is important to keep in mind the fact that, although in a probabilistic formulation adding more data should produce

---

[3] The figure shows a contour map of an area southeast of Salt Lake City, UT, generated from USGS 30 m DEM data.

TABLE II
PERCENTAGE OF 100 ITERATIONS FOR WHICH GIVEN
COMBINATION OF LANDMARK TRIPLES PRODUCED AN
ESTIMATED VIEWPOINT CLOSEST TO THE ACTUAL VIEWPOINT

Angle error	"Good" triple	2 triples*	5 triples	10 triples
10%	71%	20%	1%	8%
20%	72%	22%	2%	4%
30%	67%	21%	1%	11%

*In all cases, the 2 Triple and 5 Triple columns included the "Good" Triple.
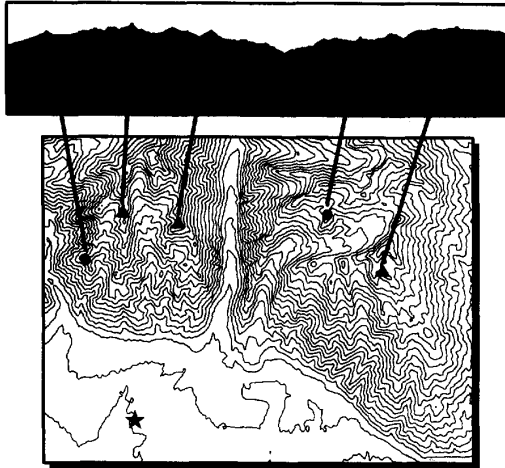

Fig. 11. The five marked mountain peaks in the view correspond to points on the map. The navigator is at the star in the foreground. The good triple is marked by triangles.

better results, unless there is a large amount of data to be added, which is most often not the case, one poor estimate can easily skew results. With the single assumption being that the error in the visual angle measure is bounded by a given amount, the only available knowledge is that the estimated viewpoint lies within some area of uncertainty, the size and shape of which is unknown. If, as on the left of Fig. 7, one area was nested inside another and viewpoints were estimated using both triples, an estimated viewpoint using the configuration that produced the smaller area would have to be somehow combined with the viewpoint estimated using the configuration that produced the larger area. Assuming a uniform distribution in the error in the visual angle measure, the estimate lying somewhere within the outer area is most likely *outside* the inner area and will negatively affect the accuracy of any estimated viewpoint within that inner area.

These results convinced us that exploiting the geometric properties of a few identified landmarks rather than assuming that numerous landmarks had been identified would help minimize our cost, particularly since the number of significant landmarks in outdoor environments is usually quite small and identifying landmarks in large-scale space is a difficult and time consuming task [25].

B. What Makes a Configuration "Good?"

If the navigator does take only one measurement using a "good" configuration of landmarks, how does it pick that

configuration without knowing its own location? We claim that knowledge of the landmark map location and the landmark viewing order, together with knowing that the viewpoint is located somewhere on the map, is sufficient for choosing good configurations. As was stated in Section II, when the robot passes from one orientation region to another, the landmark order changes. Since the computation of the area of uncertainty is dependent on the landmark order, the area will always be bounded by the orientation region formed by those three landmarks used for localization. Our algorithm begins by picking the triple of landmarks that produces the smallest orientation region on the map. (See Fig. 1.) An estimate, V, of the viewpoint is made using this triple.

Incorporating the following constraints, based on the error analysis of Section III, we have developed a "goodness" function to weight configurations.

1) If all triples produce the same orientation region (e.g., all landmarks lie on a straight line), the most widely spaced landmarks should be chosen. The effect that this spread has on the size of the area of uncertainty is illustrated in Fig. 7.
2) It follows from the previous point that the spread of the outer two landmarks relative to their distance to the viewpoint affects area size in all cases, not only when the landmark configuration is a straight line.
3) The closer a configuration is to a single circle (i.e., all landmarks plus the viewpoint on one circle), the greater the error in localization. An ongoing rule of thumb is to avoid anything near a single circle configuration.
4) The relative distance of the center landmark to the viewpoint compared to the distance of a line passing through the outer two landmarks to the viewpoint will affect the area size. The closer the center landmark is to the viewpoint, the better the localization.
5) Localization improves as the center landmark moves further back (away from the viewpoint) from the circle passing through the viewpoint and the outer two landmarks.
6) If the robot is moving rather than locating statically, the direction in which it chooses landmarks for localizing can affect the distance that it "strays" laterally from its path. The effect of the direction in which a one dimensional slice of the two-dimensional distribution is taken was shown in Figs. 9 and 10.

This heuristic function uses the locations of landmarks A, B, C and the estimated viewpoint V. If the robot is moving and wishes to minimize lateral distance from the path, the next point on the path, P, is also used. The larger the function value, the better the configuration. Although V is not necessarily the true viewpoint, our experiments have shown that this function discriminates in such a way that a good configuration to be used for localization can be determined using this estimate. Let $A = (Ax, Ay, Az)$, $B = (Bx, By, Bz)$, $C = (Cx, Cy, Cz)$, $V = (Vx, Vy, Vz)$, $P = (Px, Py, Pz)$ be the projections of the landmark points, the estimated viewpoint V and the next path point P on a horizontal plane. Let I be the point of intersection of the line through V and B with the circle
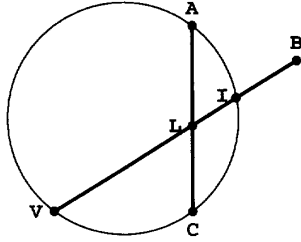
Fig. 12. Simple geometric relations can be used to rank landmark configurations.

through A, C, and V; L be the point of intersection of the line through A and C with the line through V and B; and $d(p, q)$ be the distance between any two points p and q. (See Fig. 12.)

Then:

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

where

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

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

$$w = a \left| \frac{(d(V, P)^2 + d(V, B)^2 - d(P, B)^2)^2}{2 * d(V, P)^2 * d(V, B)^2} - 1 \right|$$

The function consists of three parts:

1) The h function weighs the elevation of the landmarks compared to the elevation at point V. It is nonnegative and attains its maximum of 1 when the average elevation of the landmarks is equal to the elevation at V. An example was given in Section III showing how landmarks at elevations higher than the viewpoint produce an area of uncertainty larger than that produced if they were all at close to the same elevation. As the difference between the average elevation of the landmarks and the elevation of the viewpoint increases, the value of the h function decreases. So that it does not go to zero too quickly, the value of the constant k must be determined by the units in which the elevations are measured and the maximum difference in elevations in the area. The latter information is available from the map. Due to the fact that our elevation data was in meters, k was set to .005 in the experiments we describe here.

2) The f function, nonnegative and piecewise continuous, has the major effect on the goodness measure. It is based on the size of the area of uncertainty for the projected points. The continuity of the function was maintained so that small changes in positions would not produce jumps in function value. Note that, for any landmark triple and viewpoint position, a diagram such as that shown in Fig. 12 can be drawn. A circle passes through

the outer two landmarks, A and C, and the estimated viewpoint, V. A line joins the outer two landmarks. A second line passes through the estimated viewpoint and the center landmark. The center landmark, B, can be located anywhere along that line. We labeled the point where that line intersects the AC line as L and the point where it intersects the circle as I.

We split the line joining V and B into three sections, each corresponding to a piece of the piecewise defined function f. The first definition holds if B is on or outside the circle. The second definition holds if B is inside the circle and either on the AC line or further from V than that line. The third definition holds if B lies on the same side of the AC line as V.

Holding to the third constraint to avoid anything close to a single circle configuration, the factor $d(V, B) - d(V, I)$ produces 0 for the value of f when B is on the circle. This factor also appears (reversed) in the second piece of the f function. Thus, as B approaches I from either direction, the value of f goes to 0. This factor was not necessary in the third piece of the function since B cannot be on the circle if it is on that section of the line.

Holding to the second constraint that the wider the spread of the outer two landmarks relative to their distance from the estimated viewpoint, the smaller the area of uncertainty, the factor $\frac{d(A,C)}{d(V,L)}$ was included in the first two pieces of f. This factor was modified for the third piece of the function with $d(V, B)$ replacing $d(V, L)$ in the denominator. Thus, f would not only increase in value as the outer two landmarks moved further away, but also increase as B moved closer to V. This satisfies the fourth constraint.

If B lies on the AC line, $f = \frac{d(A,C)}{d(V,L)}$. The function then increases in value as the distance between the outer two landmarks increases relative to the distance from the landmark line to V. This satisfies the first constraint.

The fifth constraint is satisfied by the same factor, $d(V, B) - d(V, I)$, which goes to zero as B approaches the circle. As B moves away from the circle and the area of uncertainty decreases in size, this factor will increase. However, as shown in Fig. 6, B must move further back from the circle to have an area of uncertainty similar in size to that obtained when it moves a given distance from the circle toward V. Although the relationship is not linear, we have found that for the range of configurations that occurred in our trials, the factor $\frac{2}{\pi}$ does a good job of equally rating two configurations that produce similar-sized areas of uncertainty where one has B positioned toward V and the other has B outside of the circle.

3) The w function, which addresses the sixth constraint, is used by a moving navigator wishing to minimize the lateral distance from the path. It is nonnegative, continuous, attains its maximum of a when the angle from V to the heading and the center landmark is either 0° or 90° and equals zero when the angle from V to the heading and the center landmark is 45°. To obtain this function, we used the cosine of twice the angle from V to the heading and the center landmark. Since, for any

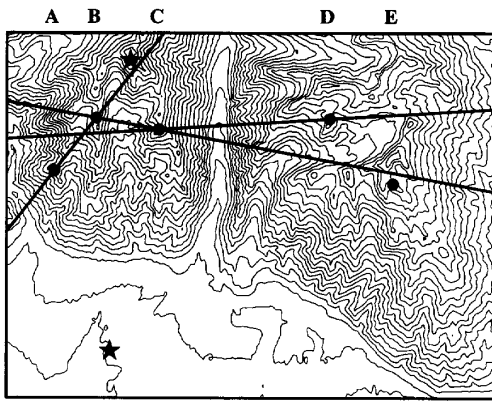Fig. 13. The algorithm begins by picking landmarks that produce the small-est orientation region on the map.

TABLE III
BASED ON THESE RESULTS, THE NAVIGATOR IN THE FOREGROUND OF THE MAP IN FIG. 13 WILL ESTIMATE ITS LOCATION USING CONFIGURATION BCE

Configuration	Rating	Configuration	Rating
ABC	.61	ADE	.46
ABD	.4	BCD	1.13
ABE	.58	BCE	1.41
ACD	1.13	BDE	.42
ACE	1.29	CDE	.14

Fig. 14. Scatterplots of 100 viewpoint estimates for the viewpoint shown in the foreground of Fig. 13. The plot on the left is the result of localizing using a good landmark configuration. The plot on the right is the result of choosing configurations at random. The map boundaries were 427890E, 4482180N by 452160E, 4498320N UTM coordinates. The maps are shown with east toward the top.

angle θ, $\cos 2\theta$ ranges from 1 to -1 as θ ranges from $0°$ to $90°$ and takes the value 0 when $\theta = 45°$, $|\cos 2\theta|$ attains a maximum of 1 at $0°$ and $90°$ and a minimum of 0 at $45°$

By a trigonometric identity:

$$|\cos 2\theta| = |2 * \cos^2 \theta - 1|$$

By the Law of Cosines, this equals:

$$\left| \frac{(d(V,P)^2 + d(V,B)^2 - d(P,B)^2)^2}{2 * d(V,P)^2 * d(V,B)^2} - 1 \right|$$

We set $a = 1$ for our moving navigator experiments. This function is sensitive to the one-dimensional distribution across the area of uncertainty described in Section III.

Fig. 13 shows the same map and landmarks as shown in Fig. 11 with the identified set of landmarks marked by large dots and three of the ten LPB's drawn on the map. The landmark labels run across the top of the map directly above the corresponding landmarks. A navigator located at the star in the foreground of the map views the landmarks in the order A, B, C, D, E. A navigator located at the star near the top views the landmarks in order D, E, C, B, A.

The algorithm is implemented in the following way:

Step I: The navigator in the foreground will choose landmarks A, B, and C (in that order) to determine the estimated viewpoint V. The navigator, knowing landmark order, knows its own orientation region and therefore knows that these three landmarks, when used alone, produce the smallest orientation region. (Note that the smallest region for the navigator at the top of the map is the one also formed by those landmarks, but in the order C, B, A.)

Step II: The navigator then estimates the viewpoint V using these landmarks. This estimated viewpoint is used in the goodness function. Table III shows the results. The navigator will estimate its location using configuration BCE. This is the "good" configuration that was used for the run summarized in Table II.

Our heuristic function is not guaranteed to always provide the best localization. However, it is easy to implement, intu-

itive in its method of rating, computationally inexpensive and, in general, does a good job. We compared the performance of our algorithm with indiscriminate choice of landmarks to be used for localization both for a robot statically locating and self-locating while moving on a path. We assumed at all times that the robot was terrain-bound. All landmarks used were actual mountain peaks with locations and elevations taken from USGS 30 m DEM data.

C. Static Localization

The set of identified landmarks for the static localization experiments was five, providing ten different combinations of ordered landmarks. The identified landmarks for each experiment were chosen based on location relative to the viewpoint. All landmarks could theoretically be seen from the viewpoint location. Multiple runs of 100 iterations each were made at different locations and with different groups of five landmarks. A random error with a uniform distribution bounded by $\pm 20\%$ of visual angle measure was added to all angle measures. Fig. 14 shows scatterplots for the run that was done using the configuration shown in Fig. 13 with the actual viewpoint at the star in the foreground. The plot on the left is the result of using a good landmark configuration to localize. The plot on the right is the result of choosing configurations randomly. Table IV shows mean distance and standard deviation in meters from the actual viewpoint for a sampling of runs done on one map, including that plotted in

TABLE IV
RESULTS OF A SAMPLING OF LOCALIZATION
RUNS USING THE MAP SHOWN IN FIG. 14*

	Good choice		Random choice	
Location	Mean distance	Standard deviation	Mean distance	Standard deviation
430560E,4494240N	1264.89	513.12	3237.74	2726.13
438840E,4486290N	402.73	202.35	688.74	478.83
430620E,4488210N	1142.44	670.81	3136.97	2520.50
445830E,4487340N	1368.04	1136.07	2876.42	3288.25
443460E,4496760N	987.45	358.96	1185.76	819.96

*Each run was of 100 iterations. The map boundaries were 427890E, 4482180N by 452160E, 4498320N UTM coordinates.

Fig. 14. The wise choice of landmarks consistently performed better than a random choice.
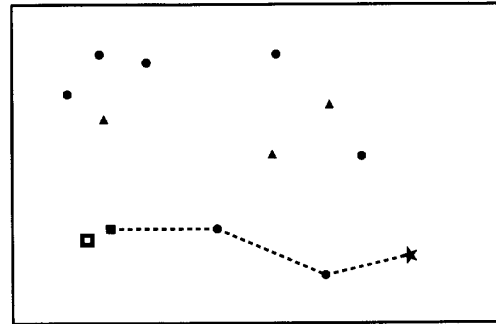

Fig. 15. The eight points at the top of the figure represent the eight landmarks used for localization. Both navigators start at the solid square on the lower left. The viewpoint is estimated (empty square) using the three landmarks (triangles) that produce the smallest orientation region. The desired path is shown as a dotted line. The goal is marked by a star.

D. Localization for a Moving Robot

In our experiments with simulated runs of a moving robot, we assumed that the navigator knew the map locations of points that defined the path as well as the landmarks and the landmark order with respect to the initial navigator location. Results for one example are shown in Fig. 15 and in the sequence of frames in Fig. 16. Each frame in this example represents an area approximately 18 by 12 kilometers with the lower left corner corresponding to UTM coordinates 427020E, 4497780N, southeast of Salt Lake City, UT. North is to the left of each frame, and east is toward the top. As in the static localization experiments, all landmarks were visible from the given path. The eight landmarks used for these trials provided 56 different combinations of ordered landmark triples.

The two navigators begin by using their knowledge of landmark order to determine the smallest orientation region in which they are located. They use the landmarks that form that region to estimate their initial location. Those three landmarks are shown as triangles in Fig. 15. The estimated location (the same for both navigators) is shown by the empty square. The desired path is shown by a dotted line. The goal is marked by a star. The sequence of frames in Fig. 16 shows each step as the navigators progress toward the goal. A configuration of three landmarks to use for localization (triangles) is chosen. The viewpoint (empty square) is estimated and a move is made toward the next path point (line ending in a solid square). The sequence on the left shows a wise choice of landmarks. The landmarks are chosen randomly in the sequence on the right.

The landmarks used by the navigator on the right in the first frame are not as widely spaced as those used on the left. In addition, the center landmark lies behind (with respect to the navigator) the line joining the outer two landmarks whereas the center landmark on the left lies in front of that line. These conditions result in a larger area of uncertainty for the configuration on the right and somewhat poor localization. This error is made up for in the second frame, but a large error in estimation occurs in the last frame. The reason for this is that the actual navigator location (from which the estimate was made) and the three landmarks chosen are very close to being on a single circle. The visual angles themselves in the

corresponding third frames are quite similar: 28° and 45° on the left and 42° and 28° on the right.[4]

The three pairs of frames in Fig. 17 show navigator positions for 50 trials, assuming uniform distribution of error within ±20% in visual angle measure and no error in movement, error within ±30% in visual angle measure and no error in movement, and error within ±20% in both visual angle and direction and distance of move.[5] The clustering around the path points is quite marked on the left, the result of using our algorithm to choose landmark configurations.

Table V gives the results for all three cases after 100 trials each. All distances have been rounded to the nearest meter. "Mean Extra Distance Traveled" is the average number of meters ± total path length that each navigator traveled. Due to the fact that paths in unstructured environments are seldom straight, the total distance traveled does not necessarily reflect how well the navigator stayed on the desired path. For that reason, we also recorded the distance of each path segment of the desired path to that of the actual path taken. The perpendicular distance of the midpoint of the desired path segment to the actual path segment taken was computed for each segment. The average of all these distances is given in the table as "Mean Distance to Path." This gives an indication of the lateral distance of each navigator to the desired path. "Mean Distance to Goal" is the average distance to the goal. The navigator that used our algorithm traveled less, remained closer to the path and ended closer to the goal than the second navigator. It is important in this type of environment that, when better localization at the goal is needed, the navigator is close enough to that goal to exploit local constraints. The navigator who chose landmarks wisely should be able to use local constraints in all three sets of trials. It is questionable if the second navigator, averaging a minimum of two miles away from the goal, will be able to take advantage of such constraints.

[4]The landmark elevation affects the visual angle measure. That is why the sums of the angles are not equal even though the outer landmarks are the same.

[5]A point was picked from a uniform distribution within a circle of radius 20%$ of the path segment length around the desired path point.
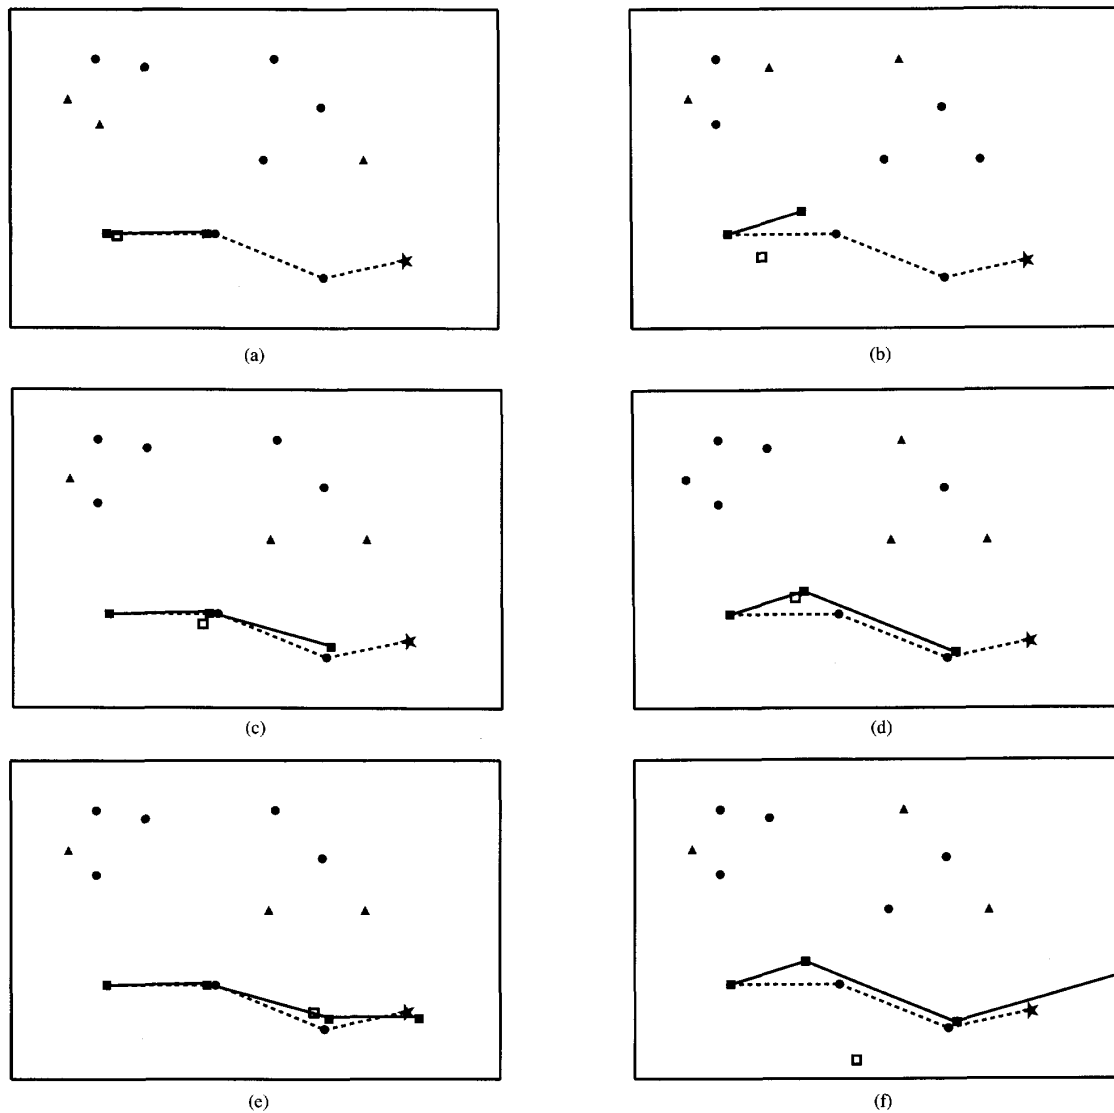
Fig. 16. The sequence on the left shows the path taken by the navigator using our algorithm. The sequence on the right shows the path taken when the landmarks used for localization are chosen randomly. The landmarks used for localization are shown as triangles. The desired path is a dotted line. The path taken is a solid line. The viewpoint is estimated at the empty square, and the navigator moves to the next path point (end of the solid line furthest to right).

It is interesting to note that for this particular example, the w function had no effect on which landmark triples were chosen by the wise navigator. When so few possible choices of landmark triples exist, it is often the case that there are no triples that are equally good for point localization but differ significantly in minimizing the lateral distance from the path.

E. Adding New Landmarks

Whether placing beacons to be used for localization or identifying new natural landmarks, we can significantly reduce our cost if we choose locations that will give us good localization. We again turn to the example of the desert ant. *Cataglyphis* economizes on what it learns about its surroundings. It learns only significant landmarks and no more than are necessary to stay on course. It does not learn the complete topography of a large area [2]. We have used the results of our analysis in Section III to choose areas on the terrain map in which to look for landmarks.

Consider two navigators following the path shown in Fig. 18.[6] Their landmarks are mountain peaks. As they move along the path, some landmarks are no longer visible and new ones are identified. They start at the square shown at the bottom of each frame. Their goal is to reach the star at the top. They have identified the seven landmarks marked in the

[6] They are using the same map as in Fig. 13 and are moving through the canyon that runs up the center of the map.
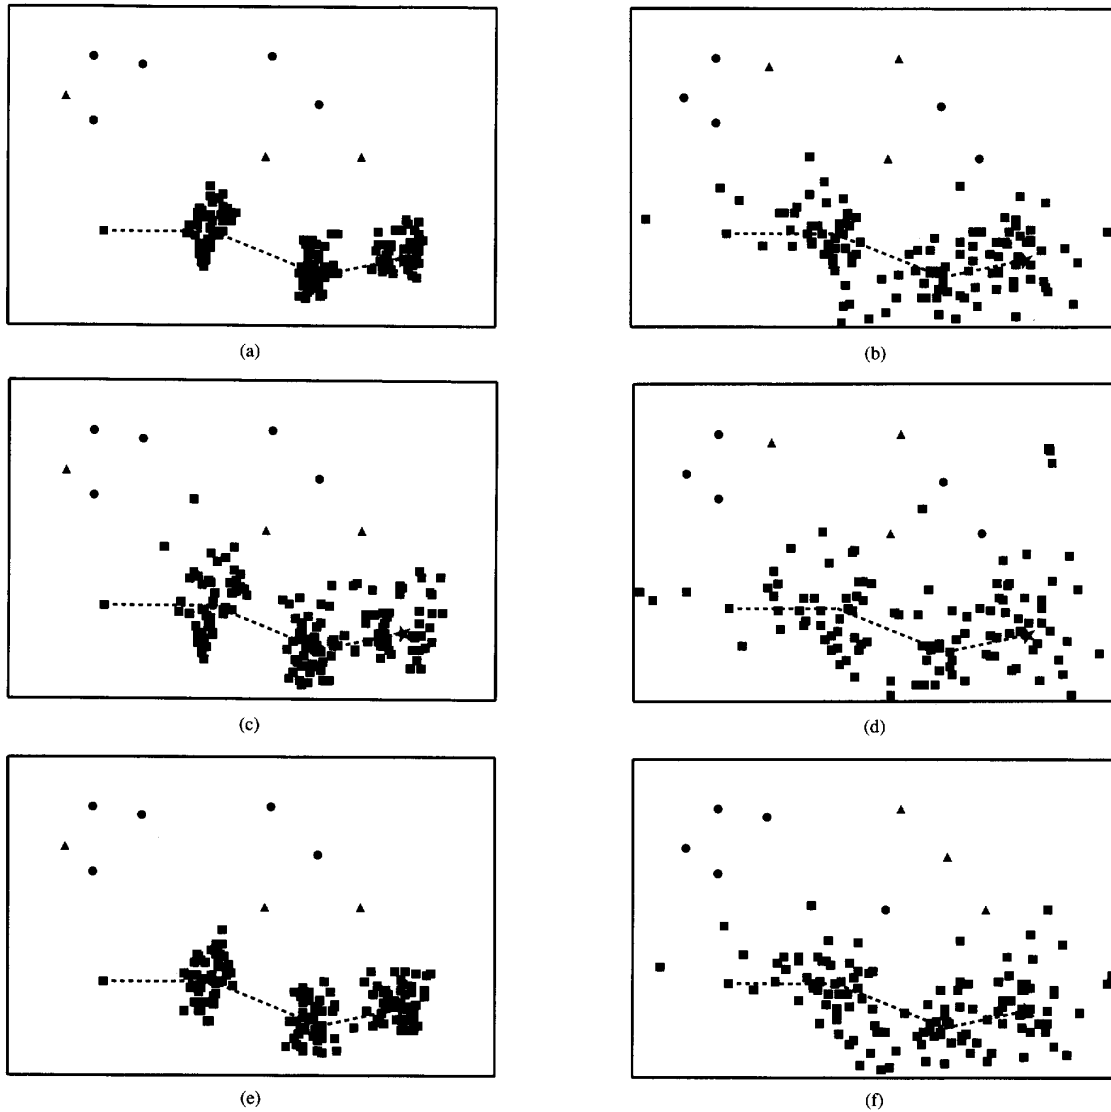
Fig. 17. After fifty trials, clustering on the left shows improved localization when the landmarks are chosen wisely. The error bounds were ±20% in the visual angle measure for the top pair of frames, ±30% in the visual angle measure for the second pair of frames, and ±20% in both the visual angle measure and the direction and distance of the move for the third set of frames.

TABLE V
RESULTS AFTER 100 TRIALS

	Wise landmark choice			Random landmark choice		
Error bounds	±20% angle 0 move	±30% 0 move	±20% angle ±20% move	±20% angle 0 move	±30% angle 0 move	±20% angle ±20% move
Mean extra distance traveled	344	2883	474	4273	18657	4576
Mean distance to path	452	513	387	1106	1227	861
Mean distance to goal	711	1166	769	3239	4781	3290

*The total path length was 11352 m. All distances have been rounded to the nearest meter.

top two frames, chosen the landmarks marked by triangles to estimate location and moved along the solid line to the next dark box. The navigator on the left made a wise choice of

landmarks with which to localize, and stayed closer to the path as a result. At this point, they take another image. The two landmarks on the far right are no longer seen in the image,
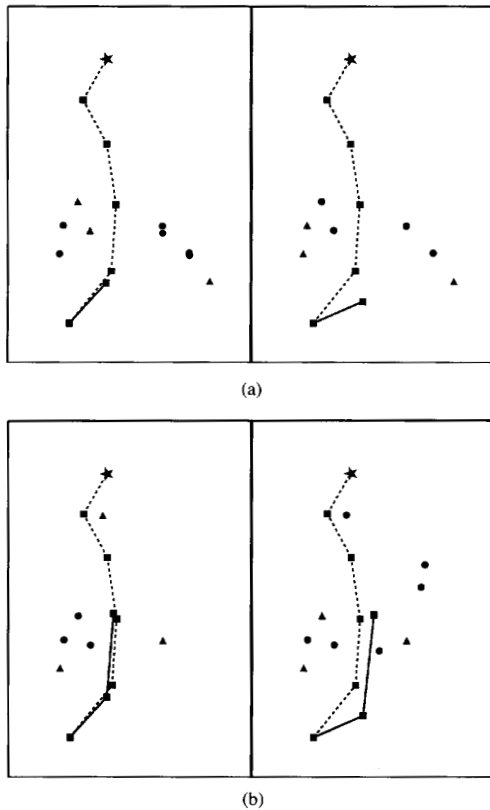
(a)

(b)

Fig. 18. An example of two navigators traveling toward a goal. They begin by picking three landmarks with which to localize, shown by the triangles in the top set of frames, and move along the dark line to the second dark box. They then take a new image. The navigator on the left looks for an additional landmark in a "good" place, and finds the one marked by the center triangle on the left of the second set of frames. The navigator on the right identifies all new landmarks and randomly picks a triple. The result of their next move is shown in the second set of frames.

and new landmarks appear. The navigator on the left looks for a landmark in the same direction as the next path point. It identifies and uses the one marked by the center triangle on the left side in the second set of frames, knowing that it should get good two-dimensional localization with a landmark pushed so far back from the ones it has already identified and good lateral localization with a landmark lying in the same direction as the next point on the path. The navigator on the right identifies all the peaks it sees in the image before it picks a triple and moves. Due to the fact that in our simulated runs, that navigator picks a random triple, it happens to pick a triple from the pool of landmarks it had already identified. It could have not identified any new landmarks before it moved. Whatever the scenario, it was in a position where none of the known landmarks was very good for localization and it did not do anything to improve the situation. As can be seen in the second set of frames of Fig. 18, the navigator on the left traveled less distance, remained closer to the path and ended closer to the next path point. It also spent minimal time identifying new landmarks because it knew where the best place was to look for them.

V. SUMMARY

Errors, their magnitude and unpredictability are one of the major factors that make outdoor navigation so much harder than indoor navigation. It is difficult and time consuming to obtain sensory data in such an environment. The data that is obtained should provide as much correct information as possible. Measurement errors cannot be avoided. They are often large and nonstandard in distribution. Filters and averaging techniques commonly used indoors quickly reach the limits of their robustness when applied to localization problems in this type of environment. It is crucial to limit the effect of the measurement errors as much as possible. In this paper, we have described the types of errors faced by a robot attempting to localize in an unstructured environment. Avoiding dependence on instruments with questionable accuracy in that environment, we have proposed a method of localizing, analyzed the errors that occur in that localization, and described steps that can be taken to help minimize those errors before they do occur. We have made a point for making the best use of the data one has and have shown that such a process may be better from a viewpoint of cost effectiveness than simply accumulating more data. We claim that a simple algorithm can be implemented to exploit the geometric properties of landmarks in the environment in order to decrease error and have shown the result of testing this concept in experiments run in simulation with real USGS 30 m DEM data for a robot statically locating, following a path and identifying new landmarks.

Our current work is addressing the following questions:

1) Can we develop better heuristics than our w function to be used in choosing landmarks for path localization that will produce a distribution with a small variance when the area of uncertainty is projected into one-dimensional space perpendicular to the direction of movement, as shown in Fig. 10? The results of this attempt to "pursue projections" can be found in [22].

2) How can a navigator best move either on the terrain or, if possible, above it, in order to put itself in a better position to self-locate?

3) How, if faced with only a partial ordering of landmarks (e.g., "I know that B is in the center"), can a small movement toward one landmark and a remeasuring of visual angles be used to determine total ordering?

We have also tested our method for choosing features for localization with a robot navigating in an indoor structured environment. Since in most cases there are more features to choose from, the results have not suffered from lack of good landmark triples. On the other hand, the errors that occurred in the confined area in which the robot could wander were quite small. It is debatable whether or not the extra time required to choose triples was well spent.

The problem of establishing the set of view to map correspondences that we assume already exists has been addressed in [24]. A constraint satisfaction approach was used such that hypotheses of view to map correspondences were discarded if they led to incompatible constraints. We have found this type of reasoning to be better suited for dealing with incorrect

correspondences than treating them as problems in geometric uncertainty.

REFERENCES

[1] S. C. Cohen, D. S. Chinn, M. H. Torrence, and P. J. Dunn, "Spaceborne laser ranging: Sensitivity to ephemeris and range measurement errors," *Manuscripta Geodaetica*, vol. 18, pp. 72–81, 1993.
[2] T. S. Collett, E. Dillmann, A. Giger, and R. Wehner, "Visual landmarks and route following in desert ants," *J. Comparative Physiology A*, pp. 435–442, 1992.
[3] T. E. Duerr, "Effect of terrain masking on GPS position dilution of precision," *Navigation*, vol. 39, no. 3, pp. 317–323, 1992.
[4] Charles R. Gallistel, *The Organization of Learning*, 1st ed. Cambridge, MA: MIT Press, 1990.
[5] Greg Hager and Max Mintz, "Computational methods for task-directed sensor data fusion and sensor planning," *Int. J. Robotics Res.*, vol. 10, no. 4, pp. 285–313, 1991.
[6] Regis Hoffman and Eric Krotkov, "Terrain mapping for outdoor robots: Robust perception for walking in the grass," in *Proc. IEEE Int. Conf. Robotics Automat.*, May 1993, pp. 1-529–1-533.
[7] Akio Kosaka and Avinash C. Kak, "Fast vision-guided mobile robot navigation using model-based reasoning and prediction of uncertainties," *CVGIP: Image Understanding*, vol. 56, no. 3, pp. 271–329, 1992.
[8] Eric Krotkov, "Mobile robot localization using a single image," in *Proc. IEEE Int. Conf. Robotics Automat.*, pp. 978–983, 1989.
[9] Benjamin Kuipers, "Modeling spatial knowledge," *Cognitive Science*, vol. 2, pp. 129–153, 1978.
[10] Benjamin J. Kuipers and Tod S. Levitt, "Navigation and mapping in large-scale space," *AI Magazine*, pp. 25–43, Summer 1988.
[11] Tod S. Levitt and Daryl T. Lawton, "Qualitative navigation for mobile robots," *Artificial Intell.*, vol. 44, no. 3, pp. 305–360, 1990.
[12] G. Lu and G. Lachapelle, "Statistical quality control for kinematic GPS positioning," *Manuscripta Geodaetica*, vol. 17, no. 5, pp. 270–281, 1992.
[13] Philip Mattos, "GPS," *Electronics and Wireless World*, vol. 98, pp. 982–987, 1992.
[14] Raymond McKendall and Max Mintz, "Sensor-fusion with statistical decision theory: A prospectus of research in the grasp lab," Tech. Rep. MS-CIS-90-68, University of Pennsylvania, September 1990.
[15] William J. Negast and Randall N. Paschall, "Compensation of selected availability using a GPS/INS extended Kalman filter," in *Proc. IEEE 1992 National Aerospace and Electronics Conf.*, pages 356–362, May 1992.
[16] Herbert L. Pick, Jr., Albert Yonas, Douglas Gentile, Patricia Melendez, Douglas Wagner, and Dominick Wegesin, "Perceptual aspects of navigation," in *Proc. DARPA Image Understanding Workshop*, April 1993.
[17] Harold W. Sorenson, "Least-squares estimation: from Gauss to Kalman," *IEEE Spectrum*, pages 63–68, July 1970.
[18] W. Kenneth Stewart, "Remote-sensing issues for intelligent underwater systems," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 230–235, IEEE, June 1991.
[19] Kokichi Sugihara, "Some localization problems for robot navigation using a single camera," *Computer Vision, Graphics, and Image Processing*, vol. 42, pp. 112–129, 1988.
[20] Karen T. Sutherland, "Sensitivity of feature configuration in viewpoint determination," in *Proc. DARPA Image Understanding Workshop*, pages 315–319, January 1992.
[21] Karen T. Sutherland and William B. Thompson, "Inexact navigation," in *Proc. IEEE Int. Conf. Robotics Automat.*, pages 1-1—1-7. IEEE, May 1993.
[22] Karen T. Sutherland and William B. Thompson, "Pursuing projections: Keeping a robot on path," in *Proc. IEEE Int. Conf. Robotics Automat.*, pp. 4-3355–4-3361, 1994.
[23] Raj Talluri and J. K. Aggarwal, "Position estimation for an autonomous mobile robot in an outdoor environment," *IEEE Trans. Robotics Automat.*, vol. 8, no. 5, pp. 573–584, 1992.
[24] William B. Thompson, Thomas C. Henderson, Thomas L. Colvin, Lisa B. Dick, and Carolyn M. Valiquette, "Vision-based localization," in *Proc. DARPA Image Understanding Workshop*, April 1993.
[25] William B. Thompson and H. L. Pick, Jr., "Vision-based navigation," in *Proc. DARPA Image Understanding Workshop*. San Mateo, CA: Morgan Kaufmann, January 1992, pages 149–152.
[26] J. M. Tranquilla and H. M. Al-Rizzo, "Investigation of GPS precise relative static positioning during periods of ice clouds and snowfall precipitation," *IEEE Trans. Geosci. Remote Sensing*, vol. 31, no. 1, pp. 295–299, January 1993.
[27] Albert J. Wavering, John C. Fiala, Karen J. Roberts, and Ronald Lumia, "Triclops: A high-performance trinocular active vision system," in *Proc. IEEE Int. Conf. Robotics Automat.*, pages 3-410—3-417, IEEE, May 1993.
[28] B. A. Wichmann and I. D. Hill, "An efficient and portable pseudo-random number generator," *Applied Statistics*, vol. 31, pp. 188–190, 1982.
[29] Sien-Chong Wu and William G. Melbourne, "An optimal GPS data processing technique for precise positioning," *IEEE Trans. Geosci. Remote Sensing*, vol. 31, pp. 146–152, 1993.

Karen T. Sutherland received the A.B. degree in mathematics from Augustana College, Rock Island, IL; an M.S. degree in mathematics and an M.S. degree in computer science, and the Ph.D. degree in computer science, all from the University of Minnesota. She is an Assistant Professor in the Department of Computer Science at the University of Wisconsin–La Crosse. Her research interests include robot navigation, object recognition, and filtering algorithms.

Prof. Sutherland is a member of the AMS, the MAA, the ACM, and the AAAI. She is an affiliate of the IEEE Computer Society and the IEEE Society on Robotics and Automation.

William B. Thompson (S'72–M'75–SM'91) received the Sc.B. degree in physics from Brown University, Providence, RI, in 1970, and the M.S. and Ph.D. degrees in computer science from the University of Southern California, Los Angeles, CA, in 1972 and 1975, respectively. He is currently a Professor in the Department of Computer Science at the University of Utah, previously having been a member of the Computer Science Department at the University of Minnesota. His primary research interest is in the area of computer vision, with an emphasis on the perception of spatial organization. As part of his work, he has been active in the exploration of techniques for analyzing visual motion, the development of computational models for vision-based navigation (during which he has been lost in the mountains of four continents), and new approaches to using computer vision in manufacturing.

Prof. Thompson has served a term as an Associate Editor of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE. He was co-chair of the 1989 AAAI Symposium on Robot Navigation.