The Stability of Geometric Inference in Location Determination

Karen T. Sutherland

A revised version of the thesis submitted to the faculty of the Graduate School of the University of Minnesota in partial fulfillment of the requirements for the degree of Doctor of Philosophy

UUCS-94-021

Department of Computer Science
University of Utah
Salt Lake City, UT 84112 USA

July 8, 1994

Abstract

Geometric inference is widely used in computer vision, but very little attention has been given to the question of how geometric properties affect the resulting errors in the inferences made. This thesis addresses the problem of the stability of geometric inference in determining locations with a goal of being able to predict type and magnitude of the errors which occur and to determine on what basis to make geometric inferences which will minimize error. It is shown that the amount of the error occurring in a localization process using angular measurements to features depends heavily on which features are used, that the amount of the error occurring in such a localization process is not a function of the number of features used, that it is possible to develop simple heuristic functions for choosing features for localization which will significantly decrease error in that localization, that it is possible to decrease localization error in a particular direction, and that, if features have been identified but knowledge of left to right order in the view is unknown, simple steps can be taken to aid in determining that ordering. This knowledge is applied in the domain of robot navigation in outdoor, unstructured environments.

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

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Biological vision systems are constantly faced with tasks requiring geometric inference. Computational vision systems have been developed to perform many of those tasks. Both biological and computational systems must deal with errors in the inferences they make. This thesis addresses the problem of the stability of the vision-based geometric inferences made in location determination. Location determination can be intuitively described as the process of establishing a location using information which is often incomplete or inexact. The stability of an inference is determined by the amount of error in the resulting estimate of location with the most stable inferences producing the least error.

We humans use configurations of landmarks in the environment to determine our own location. We make decisions as to the particular landmarks to use, whether or not we should move and take another look before we estimate where we are, and, if we do move, in what direction(s) and how far. We make such decisions when recognizing objects: How do we identify a three dimensional object given a noisy, partially occluded two dimensional image?

The above decisions require that inferences be made regarding angular and/or distance measures. Certain configurations of features or properties of movement will allow for less error in inference than others. We say that these are more stable or less sensitive to noise. The geometric inferences which will make use of the stable configurations or properties are said to be more stable than those which do not. When we speak of the stability of geometric inference in location determination, we are referring to the above types of decisions.

This thesis addresses the problem of the stability of geometric inference in determining locations, (i.e., *localization*), with the double goal of being able to predict the type and magnitude of errors which will occur as well as to know on what basis to make inferences which will minimize error. Chapter 2 summarizes the research which has been done on the geometric inferences made by insects and early human navigators. Chapter 3 analyzes the sensitivity of a configuration of features and shows how some configurations are more stable (lead to less error in localization) than others. Chapter 4 applies this knowledge to several problems in the domain of robot navigation in outdoor, unstructured environments. Chapter 5 summarizes the contributions of this work and discusses possible future extensions of what has been done here.

The following restrictions were placed on the analysis and applications:

- It was assumed that features were point features, that some sort of map was available and that features in the view were matched to that map. The problems of setting up view to map correspondences and dealing with any major correspondence errors were not addressed. [Thompson *et al.*, 1993] addresses those problems. A small set of correspondences are established, used to infer a viewpoint, which is then used to establish more correspondences. A constraint satisfaction approach is taken such that ambiguous correspondences generate multiple hypotheses. Any hypotheses which lead to incompatible constraints are candidates for correspondence errors and discarded. Minor correspondence errors, such as misidentifying a subpeak as a peak, could in most cases be handled as measurement errors, whereas major correspondence errors, such as those which would reverse the direction of view, were found to be better handled using constraint satisfaction.

- Distance estimates, because distance is so often impossible to estimate, were not used.

- Due to the fact that in many practical applications, conditions on the earth's surface render navigational instruments useless, such instruments, including a compass, were assumed not available.

- An effort was made to keep any heuristics used to minimize error as simple as possible. Even when maps were assumed available, an attempt was made to use only simple geometric information from those maps. This need for simplicity was emphasized by W. Grey Walter [Walter, 1953] in his description of the behavior of *Machina speculatrix*. The electromechanical *M. speculatrix* was built to exemplify the behaviors of a living creature. Three of those behaviors were "parsimony", "discernment" and "optima". Parsimony referred to the economy of structure and function in living animals while discernment concerned the ability to distinguish between effective and ineffective behaviors. Optima was "the tendency to seek conditions with moderate and most favourable properties, rather than the maxima". As Walter so eloquently states: "Little virtue or interest would lie in achieving life like effects with a multiplication of mechanisms greater than life would tolerate. Creatures with superfluous organs do not survive; the true measure of reality is a minimum."

Previous work in localization, when it addressed error at all, did so after the error occurred. What has been done here demonstrates that steps over and above improving sensor accuracy can be taken to decrease it before it occurs. The specific contributions of this work, to be further summarized in Chapter 5, are as follows:

- The amount of the error occurring in a localization process using angular measurements to features depends heavily on which features are used.

- The amount of the error occurring in such a localization process is not a function of the number of features used.

- It is possible to develop simple heuristic functions for choosing features for localization which will significantly decrease error in that localization. This applies not only to choosing features in the view which have been previously matched to a map, but to choosing directions in which to look for new features.

- It is possible to decrease localization error in a particular direction. This is useful for a robot attempting to stay close to a path.

- If features have been identified but knowledge of left to right order in the view is unknown, simple steps can be taken to aid in determining that ordering.

Chapter 2

Background

The next two sections summarize research on the geometric inferences made by two different biological systems, as notable in their differences as they are in their similarities. Section 2.1 discusses how geometric inferences are made by insects. Section 2.2 looks at some of the primitive human navigational techniques used to deal with errors in geometric inferences.

2.1 Geometric inferences made by insect navigators

Although animal behavior has been studied for centuries, the fairly new field of ethology[1], which involves observations of animals in their natural environments rather than in a laboratory setting has contributed much to the understanding of how insects as well as other animals can successfully find their way home after foraging large distances from a nest site [Huntingford, 1984, Slater, 1985].

Several well known computational navigation systems have been modeled after insect-type navigation [Brooks, 1986, Maes and Brooks, 1990, Lucarini *et al.*, 1993]. However, these systems are based on low level reactive behaviors or self-organization of multiple agents. Many insect-like robots working at a very low level accomplish a possibly complex task. Studies of actual insect navigators [Wehner and Raber, 1979, Cartwright and Collett, 1982, Collett *et al.*, 1992] have shown that some insects use higher level geometric inferences when navigating. It is these navigational techniques which relate to this work.

In order to use geometric inferences in navigation, some sort of representation of geometric relationships must exist. A cognitive map is one such representation. A cognitive map can be defined in different ways. It could be a simple record of geometric relationships in the environment which can be used to plan movements through that environment. It is possible for an animal to have such a record of geometric relationships among points in its environment, but not have a complete map of the surrounding terrain. Gallistel [Gallistel, 1990] claims that the popular belief that the cognitive maps of "lower" animals are weaker than those of humans is not well-founded. It is

[1] Many dictionaries still define this word as "the science of the formation of human character".

important to distinguish between a map having less in it (being "impoverished") and actually being weaker in a geometric sense. He classifies geometric maps into four categories: topology, projective geometry, affine geometry and metric geometry, with each one being weaker than the next (i.e., metric maps are also affine but affine maps are not necessarily metric). As an example, bees can distinguish between square and rectangular configurations around a feeding site. Thus, their maps must be more than affine because a square and a rectangle are equal under affine transformations. Distance and angle measures, fundamental to a metric map, are used to differentiate configurations.

At the other extreme, a cognitive map can also be defined as the mental analogue of a topographical map [Wehner, 1990] where an animal can determine its position relative to any other point within its environment, even if it has been displaced. In this wider sense, it is questionable whether or not insects possess cognitive maps at all. If they do, they are most likely "weak". However, it is possible for humans (see Section 2.2) as well as insects to navigate successfully without a mental topographic map.

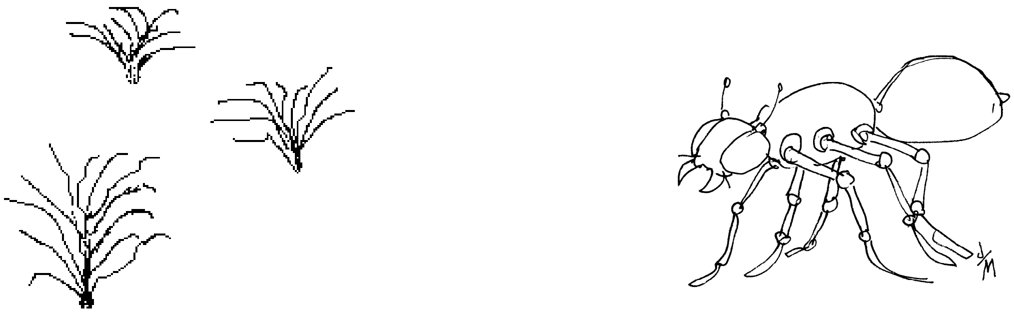

Figure 2.1: Cataglyphis bicolor, the desert ant, navigates using point sightings on mapped terrain. (Ant drawing courtesy of James Moen.)

Experimental evidence substantiates an assumption that animals do not use celestial fixes to determine their own position. This makes sense, considering that an error of 1° in the sun's azimuth will put a navigator off by 100 km. Gallistel believes that animals obtain direction from the heavens but determine their position from an interplay between dead reckoning[2] and piloting with the aid of point sightings on mapped terrain. For example, the desert ant, Cataglyphis bicolor, looks for its nest only when reckoning indicates that it is in the general vicinity. At this point, the dead reckoning stops and piloting takes over. This dead reckoning or *path integration* is not done by performing a vector summation such as a modern human navigator would do, but by using computationally simple approximations [Muller and Wehner, 1988]. A path integration system produces large errors only when the ant makes sharp backward turns, and it very seldom makes this type of turn. The errors in distance and angular measure that it does make tend to cancel each other out. If its reckoning contains a great deal of error, it will either run right past its nest or start piloting too soon and misidentify the landmarks it needs for that piloting.

[2] See Section 2.2 for an explanation of the source of this term.

This combining of the two determination methods conforms with the theory set forth in [McFarland and Houston, 1981] that the survival of any animal depends to a large extent on its optimal use of resources. Returning to its nest with the use of only landmarks would require a large store of noted landmarks and most likely limit the distance it could forage. On the other hand, dead reckoning alone could cause it to miss the nest entirely, with often fatal results due to the fact that its survival time in the desert sun is only a matter of hours. The search process that the ant goes through if it does get lost is also governed by this concept [Wehner and Srinivasan, 1981].

As the desert ant nears the nest site, it switches into its piloting stage, using a small number of prominent landmarks. This helps eliminate the uncertainty present in a cluttered environment as well as the directional errors which accumulated during path integration [Collett *et al.*, 1992, Wehner and Harkness, 1983]. Rather than recalling route details as it approaches the nest site, it simply passes either to the left or right of specific landmarks. There is usually a prominent landmark close to the nest site which plays two roles: it shows the correct general path while the ant is still quite a distance away, and its size and angular position from the nest are stored visual properties which help define the nest entrance as the ant moves closer.

Experiments in changing the size of landmarks placed near a nest entrance support the theory that a simple image matching routine is used at this point. When landmarks are replaced with ones identical except for size, the ants don't stop. It is clear that they economize, not learning the topography of a wide area, but rather basic properties of landmarks which help them stay on the path. They do not seem to possess more general information such as that the nest site is located halfway between two landmarks. Similar recent observations have been made in the case of foraging honey bees [Dyer, 1991]. Rather than a cognitive map in the sense of an internal topographic map, the bees use landmarks associated with routes which they have traveled previously. They will also head toward a panorama of landmarks associated with a food site. In neither case do they need or seem to use the complex geometric relationships available in a topographic map.

2.2 Geometric inferences made by human navigators

According to Hutchins [Hutchins, 1993], human navigators have three basic questions:

- Where Am I? (Localization)
- If I am here, how do I get there? (Path Planning)
- If I am here and move in a certain way for a certain time, where will I be? (Dead Reckoning)

Errors in measurement will particularly affect answers to the first and third questions. Western navigators determine their position on a map by the use of sightings to objects with known coordinates on that map. Lines, arcs and circles of position are established. Lines of position are determined either by lining the viewpoint up with two charted points or by using the compass bearing of one charted point. Arcs of position are determined by measuring the angles between two charted points. Circles of position are determined by the sighting of a known point together with knowledge of distance to that point. Distance is often extremely difficult to estimate, resulting in circles of position not being used very often. A stationary *fix* is obtained from intersecting lines and/or arcs of position. A common fix is called the three-point fix and is the intersection of arcs of position using three points.

The fact that errors which occur in localization are not directly proportional to errors in sensing has been known to navigators for many years. Nathaniel Bowditch (1773–1838) [Bowditch, 1802] said: "A person who understands the nature of errors avoids many pitfalls. Thus, the magnitude of the errors of individual lines of position is not a reliable indication of the size of the error of the fix obtained from them. Two lines of position with small error might produce a fix having a much larger error if the lines cross at a small angle."

Bowditch also explained the source of and confusion over the phrase *dead reckoning*. Early navigators threw a buoyant object overboard to determine the speed of the vessel relative to the object "dead" in the water. The course steered and this estimated speed through the water was then used to determine position, thus the term *dead reckoning*. This was confused with the process of *deduced* or *ded* reckoning in which a ship's position was deduced with allowances made for current and wind. The dead reckoning referred to in Section 2.1 and used throughout this work should be defined in the first, simpler way. The typical nautical error in dead reckoning for western sailors in the last century ranged from 5 to 10%. It is interesting to note that this is also the average error made by cataglyphis bicolor during path integration.

Taking fixes requires the second definition of a cognitive map, (i.e., a mental topographic map). However, it is not the case that all human nautical navigators have used or needed such a map. The navigators of the Central Caroline Islands of Micronesia use ancient techniques which require metric relationships but no maps. In fact, these methods are so different from those used in western cultures that they were not fully understood until recently.

One possible explanation for the inability to comprehend the reasoning behind these techniques is that an assumption was made that both Micronesian and western navigators were processing information in the same way. According to David Marr [Marr, 1982], any information processing task is carried out at three levels: computational theory, representational and implementational. The computational level gives the goal of the computation and the logic by which it can be carried out, while the representational level deals with the actual algorithm needed for accomplishing the task. Hutchins claims [Hutchins, 1993] that early studies of Micronesian navigation assumed that representations used in western navigation were also used by the Micronesian navigators. Due to this blindness to other feasible navigational algorithms, they attributed the navigators' success to luck rather than to skill. The Micronesians, although sharing the same computational theory as western navigators, hold an egocentric point of view. They do not have the bird's eye world view idea of western cultures. They do not use a map of the navigational area [Hutchins, 1983], nor do they use instruments. The navigational methods they do use are much like those of the desert ant, a combination of dead reckoning and piloting.

In the area surrounding the Caroline Islands, approximately $.2\%$ of the earth's surface is land. Navigation using landmarks would be extremely difficult to implement. One concept widely employed in Pacific non-instrumental navigation is that of using a "star path". The positions of the stars relative to each other are fixed from a viewpoint on the earth. The stars appear to move from east to west as the earth moves. A linear constellation (star path) is a set of stars which follow the same path. Each star path is made up of 6–10 stars, fairly evenly spaced. Connecting the stars produces an east to west path, converting the moving stars into a fixed frame of reference. These stellar maps are good examples of pure geometric representations. A star is not recognized by its own intrinsic properties, but by its geometric relationship to other stars. Different star paths are learned by navigators traveling in different areas of the Pacific. This is not an easy task. Not everyone navigates, and those who do so traditionally learn these paths as children.

This is not unlike the methods used by those birds in the northern hemisphere which migrate at night, flying away from the small cluster of stars around Polaris. Regardless of the rotation of individual stars around Polaris itself, flying away from the cluster results in a path heading south [Slater, 1985]. Any constellations used for navigation are learned by these birds as nestlings [Gallistel, 1990]. Birds in different areas of the world learn different constellations.

It is also not unlike the cognitive map as defined by Kuipers [Kuipers, 1978]. If, as Kuipers states, the body of common-sense knowledge which is built up from observations as one travels through the environment is considered to be the cognitive map, it follows that, although humans may have the ability to generate and refer to actual topographical maps, they may in many instances only need and only use a cognitive map that is just as impoverished as that of an insect.

Micronesian navigators conceived of the horizon as being a straight line parallel to the course of the canoe, rather than as a circle around it. Their cognitive maps, just as those of "lower" animals may have had less in them, but they were not any weaker in the geometric sense than the navigational maps referred to by modern navigators. In fact, this lack of cluttered detail most likely contributed to their ability to compute their own location without the use of instruments. Shepard and Hurwitz [Shepard and Hurwitz, 1985] discuss the mental rotations required to follow left and right turns on a map and the confusion resulting when "up" on the map does not coincide with straight ahead in the surrounding environment. They distinguish three different frames of reference used when localizing and moving on two-dimensional terrain: An egocentric frame, defined by up-down, left-right and front-back; an object-centered frame, defined in terms of some other object; and an environmental frame, defined in terms of north-south and east-west from a particular location on the earth's surface. The insects discussed in Section 2.1 as well as the early Micronesian navigators use the simpler egocentric frame of reference rather than the more complicated environmental frame required for modern map and instrument based navigation.

Chapter 3

Sensitivity of feature configuration in viewpoint determination

Viewpoint determination often must be based on the apparent position of features in the environment. One-dimensional constraints on position can be combined to estimate location. The two most commonly used constraints are distance and absolute bearings (angular measure from true north).

A perfect estimate of distance to one feature in three-dimensional space constrains the viewpoint to the surface of a sphere centered at that feature with radius equal to the given distance. An estimate of distance to a second feature constrains the viewpoint to the circle of intersection of the two spheres. When the features can be ordered with respect to the viewpoint position (e.g., A is left of B), the viewpoint is constrained to half that circle. In the special case when the viewpoint lies on the line joining the two features, the spheres are tangent and the viewpoint is uniquely determined, lying at the point of tangency. If, as shown in Figure 3.1, a two-dimensional approximation of the environment is assumed, knowing the distance to one feature constrains the viewpoint to the circumference of a circle of radius equal to that distance. Knowledge of the distance to two features restricts the viewpoint to the intersection of two such circles. In most cases, the circles will intersect in two points. In the special case when the viewpoint lies on the line joining the two features, the circles are tangent and the viewpoint is unique. If the left to right order of the features is known, the viewpoint is also unique. In all other cases, knowing the distance to a third feature is required for uniquely determining the viewpoint.

Absolute angular measure to one feature in both two and three-dimensional space constrains the viewpoint to a line such as that in Figure 3.2a. Knowledge of absolute angular measure to a second feature constrains the viewpoint to the intersection of two lines, as shown in Figure 3.2b. This location is unique unless the lines are antiparallel (i.e., the viewpoint lies on the line joining the two features).

However, it is frequently the case that actual distance to features is unknown. Likewise, absolute bearings are not always available and can be unreliable in many environments. In such situations,

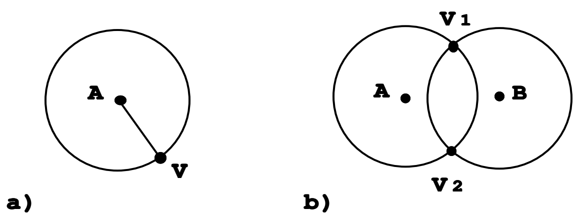

Figure 3.1: Knowledge of the distance to feature A constrains the viewpoint V to the circumference of a circle as shown in a. Knowledge of the distance to both features A and B constrains the viewpoint to the intersection of two circles, V_1 or V_2, as shown in b.

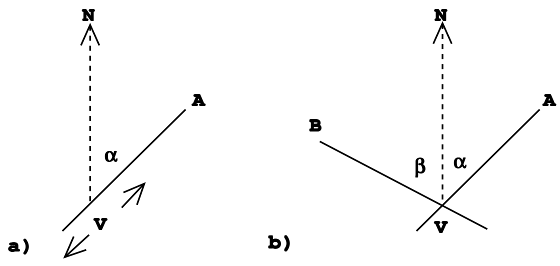

Figure 3.2: Absolute angular measure to feature A in two-dimensional space constrains the viewpoint V to a line as shown in a. Absolute angular measures to features A and B constrain the viewpoint V to the intersection of two lines as shown in b.

alternative methods must be employed. The alternative addressed in this thesis is the use of relative angular measure. This approach not only avoids the above mentioned problems, but may be a more appropriate model for applications involving biological systems. Human cultures, for example, have successfully navigated significant distances without any concept of a distance measure, and did so long before the introduction of the magnetic compass in 1100 A.D.

The *visual angle* from a viewpoint to two point features will be defined as the angle formed by the rays from the viewpoint location to each point feature. It has been shown [Levitt and Lawton, 1990] that a perfect estimate of the visual angle between two points constrains the viewpoint to a surface of revolution somewhat resembling the torus shown in Figure 3.3a. If a two–dimensional approximation of the environment is assumed, as in Figure 3.3b, the viewpoint is constrained to the intersection of the surface of revolution with a plane passing through the two feature points and the viewpoint. When the points can be ordered with respect to the viewpoint position, the viewpoint is restricted to half the surface (if *up* is known) or to one arc of the boundary. In the two-dimensional case, as shown in Figure 3.4, visual angles between three points will constrain the viewpoint to the intersection of three circles, unless all three points and the navigator lie on the same circle [Krotkov, 1989, Sugihara, 1988]. This *single circle* configuration restricts the viewpoint to the arc of the circle boundary from A to C, providing the same localization as would knowledge of the visual angle between only two points.

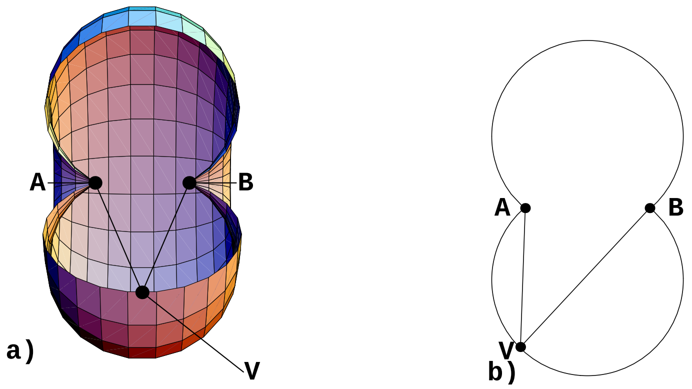

Figure 3.3: Knowledge of the visual angle measure between two points, A and B, constrains the viewpoint V to a) a surface of revolution. If a two-dimensional approximation of the environment is assumed, the viewpoint is constrained to b) the intersection of the surface in a) with a plane through A, B, and V.

The above three types of one-dimensional constraints can be combined to provide unique localization. For example, the distance to one feature and the absolute angular measure to another can be used to determine the viewpoint. However, any measurements taken will never be completely accurate. As a result, it is important to have an understanding of the errors which can develop when using such methods. Due to the previously mentioned problems with the use of distance and absolute angular measures, this analysis deals with the errors involved when using relative angular measurements only. *A posteriori* analysis can predict the precision of an estimated viewpoint given expectations about the errors associated with the determination of visual bearings. Perhaps even more importantly, *a priori* analysis can be used to choose feature configurations which are least sensitive to error, thus yielding the most reliable localizations.

In Section 3.1 the errors associated with this form of localization when the three ordered features are included within a visual angle of $< 180°$ are analyzed. A method for determining the area around the actual viewpoint in which an estimated location will lie given a particular configuration of features and specific error bounds in visual angle estimates to those features is developed. Section 3.2 covers the situation when the viewpoint lies within the feature configuration (i.e., the visual angle is not restricted to less than $180°$). The effects on the estimated location when a two-dimensional environment is modified so that features are at different elevations (e.g., landmark features on uneven terrain) are shown in Section 3.3. Section 3.4 analyzes the distribution of error in the area around the actual viewpoint. Localization using different configurations of features when identical errors in the estimate of visual angles are made is discussed in Section 3.5. This *sensitivity* to error in visual angle estimate will vary considerably among configurations, with the result that feature configuration can dramatically affect the precision with which localization is accomplished. Section 3.6 addresses the question of how visual angle measures can aid in ordering features.

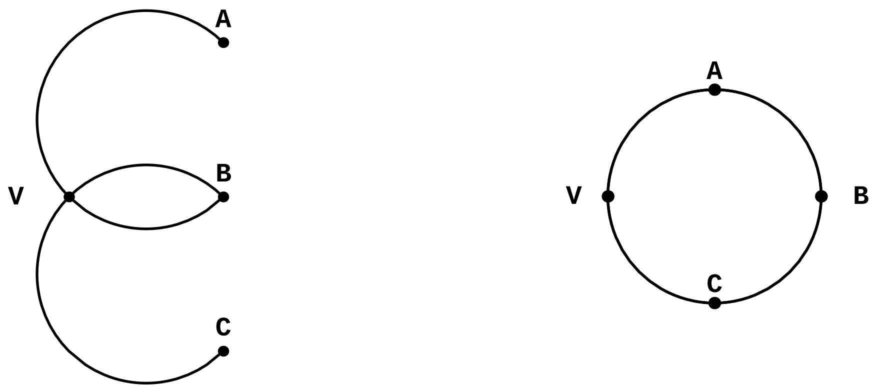

Figure 3.4: Knowledge of the visual angle measures from the viewpoint to three points A, B and C in two-dimensional space will uniquely determine the viewpoint V, as shown on the left, unless all three points and the viewpoint lie on the same circle, as shown on the right.

3.1 Area of uncertainty - viewpoint outside configuration

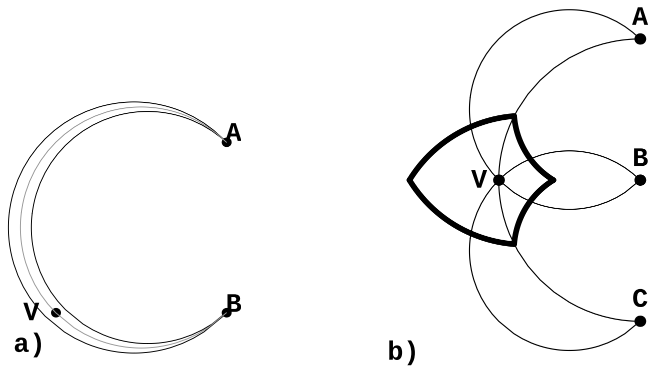

Figure 3.5: The error in visual angle estimate to two points constrains the viewpoint V to a) a thickened ring. When three points are used, the viewpoint is constrained to b) the intersection of two such thickened rings.

When a two-dimensional approximation of the environment is assumed, an error in the visual angle estimate to two features, as shown in Figure 3.5a, will constrain the viewpoint to a thickened ring, the thickness of the ring determined by the amount of error [Krotkov, 1989, Levitt and Lawton, 1990]. When three features are used, any given error in estimate constrains the viewpoint to the intersection of two such rings [Kuipers and Levitt, 1988, Sutherland, 1992].[1] See Appendix A for

[1] A third ring passing through the two features lying at greatest distance from each other can be computed, but it does not affect area size.

details of the computation. In Figure 3.5b, let α be the visual angle from the viewpoint subtended by chord AB and β be the visual angle from the viewpoint subtended by chord BC. The intersection of the thickened rings, termed the *area of uncertainty*, is surrounded by the dark lines. In this particular example, the features are in a straight line, the distance between features equals the distance from the viewpoint to the center feature and both α and β measure $45°$. The area represents an error less than or equal to $\pm 13.5°$ or $\pm 30\%$ in both α and β.

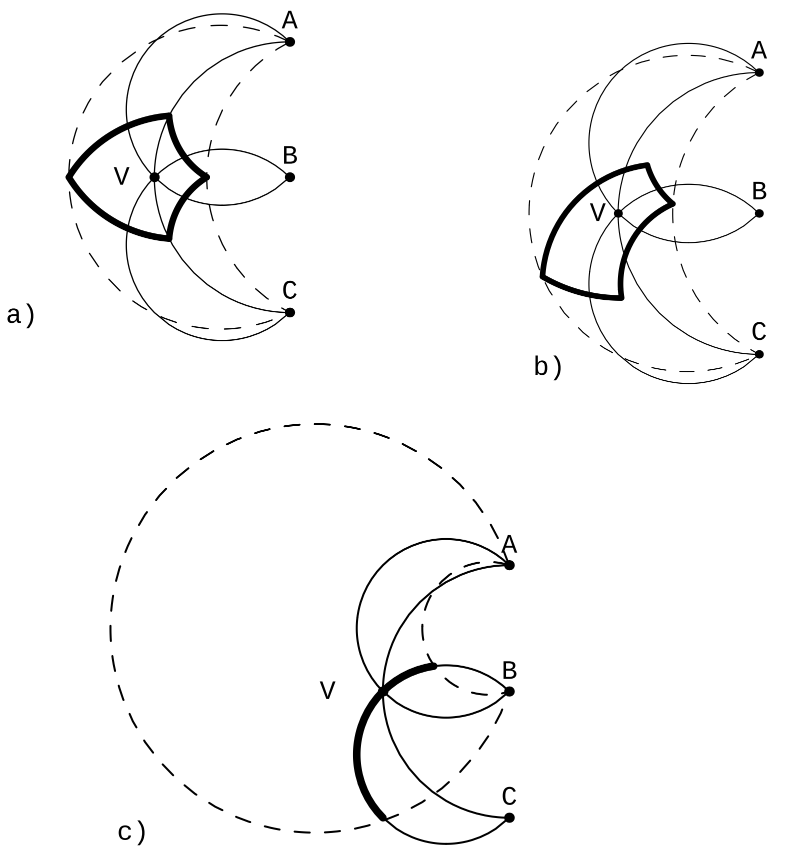

Figure 3.6: a) Dark lines surround the area of uncertainty for a $\pm 30\%$ error bound with the same configuration as in Figure 3.5b. Dashed lines surround the error area for angle γ. b) Dark lines surround the area of uncertainty for a $\pm 40\%$ error bound in α and a $\pm 20\%$ error bound in β. c) The error in the estimate of angle α is bounded by $\pm 60\%$. Angle β is estimated perfectly. The result is a $\pm 30\%$ error bound in estimate of γ. Dashed lines surround the error area for angle α.

Figure 3.5b is somewhat simplified. The features will not always be in a straight line, angles α and β will not always be identical and the same error in estimate will not always be made for each angle. An example of how different errors in estimate of angles α and β will affect the size and shape of the area of uncertainty is shown in Figure 3.6. Figure 3.6a shows the area of uncertainty for an error bound in both α and β of $\pm 30\%$. Figure 3.6b shows the area of uncertainty for an error bound in α of $\pm 40\%$ and an error bound in β of $\pm 20\%$. Figure 3.6c shows the area of uncertainty for a $\pm 60\%$ error bound in α and a perfect estimate of β. Note that, in this case, the area is reduced to an arc of the circle passing through B and C.

In general, if the error is additive with γ_e the error in γ, for any given γ_e, $\gamma_e = \alpha_e + \beta_e$ implies that $\alpha_e = \gamma_e - \beta_e$ for all β_e such that $0 \leq \beta_e \leq \gamma_e$. If the error is multiplicative, $\gamma_e = \alpha/\gamma * \alpha_e + \beta/\gamma * \beta_e$ implies that $\alpha_e = \gamma/\alpha * \gamma_e - \beta/\alpha * \beta_e$ for all β_e such that $0 \leq \beta_e \leq \gamma/\beta * \gamma_e$. In all cases, the resulting area of uncertainty equals the intersection of the two (possibly thickened) rings corresponding to the error in estimate of angles α and β. This intersection will always lie within the thickened ring corresponding to the error in γ, with the relationship of γ_e to α_e and β_e as given in the above equations.[2] For all examples in this analysis, a choice of the type of error introduced, either multiplicative or additive, was made, unless error type affected results. In these cases, both types were used.

3.1.1 Size of the area of uncertainty

Three factors affect the size of the area of uncertainty: the amount of the error in angle estimate, the relative distance of the viewpoint to the configuration and the shape of the configuration.

The area size will change according to the amount of error. The larger the error, the thicker the ring produced by the error bounds on one visual angle and the larger the area resulting from the intersection of two such rings. As an example, Figure 3.7 shows two areas of uncertainty. The outer area, with an error bound of $\pm 30\%$ of visual angle measure, is the same as shown in Figure 3.5. The inner area is for an error bound of $\pm 10\%$ of visual angle measure.

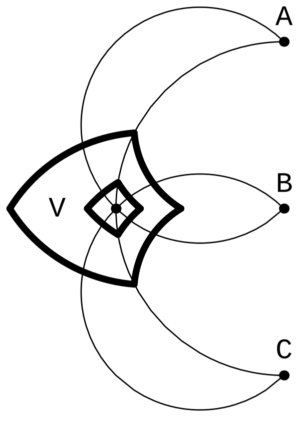

Figure 3.7: The size of the area of uncertainty is a function of the error in the visual angle measure.

It has been shown [Kuipers and Levitt, 1988, Levitt and Lawton, 1990] that the lines joining pairs of features divide space into distinguishable areas (orientation regions). [Levitt and Lawton,

[2] Although all graphs in this example show a straight line configuration of features, the described conditions also hold for nonlinear configurations.

1990] called these lines LPB's (landmark pair boundaries). When all three features (landmarks) lie on a straight line, two orientation regions are created and the viewpoint is constrained to a half plane. When the features are not in a straight line, the three LPB's create seven orientation regions, as shown in Figure 3.8. The viewpoint is constrained to one of those regions. Crossing into a different orientation region changes the ordering of the features. Since feature order has been determined, it can be assumed that an LPB will not be crossed. Thus, the LPB puts a bound on the size of the area in the direction of overestimation of the visual angle. However, for 6 of the 7 orientation regions formed by a non-linear configuration of three features, there is no bound on how far back the viewpoint can be located, causing the total possible area of uncertainty due to the amount of error in the angle estimate to be unbounded.

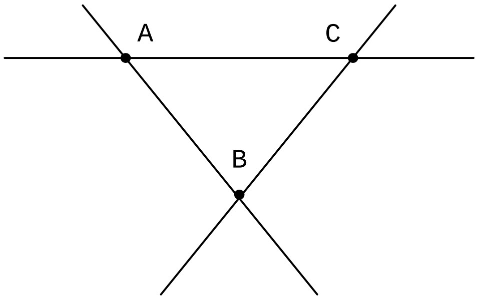

Figure 3.8: The LPB's in a non–linear configuration restrict the viewpoint to one of 7 orientation regions.

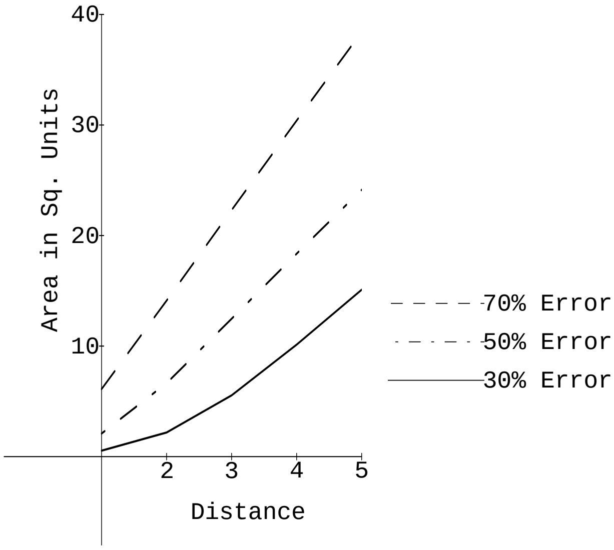

Figure 3.9: The size of the area of uncertainty as the viewpoint moves away from the configuration.

The relative distance of the viewpoint to the configuration will also affect the area size.[3]

[3]Because the visual angles alone are used as a measure, distance is relative (e.g. distance of the viewpoint of 1000

Figure 3.9 shows the change in size of the area of uncertainty for a straight line configuration with features one unit apart and the viewpoint lying on the perpendicular bisector of the line joining the features.

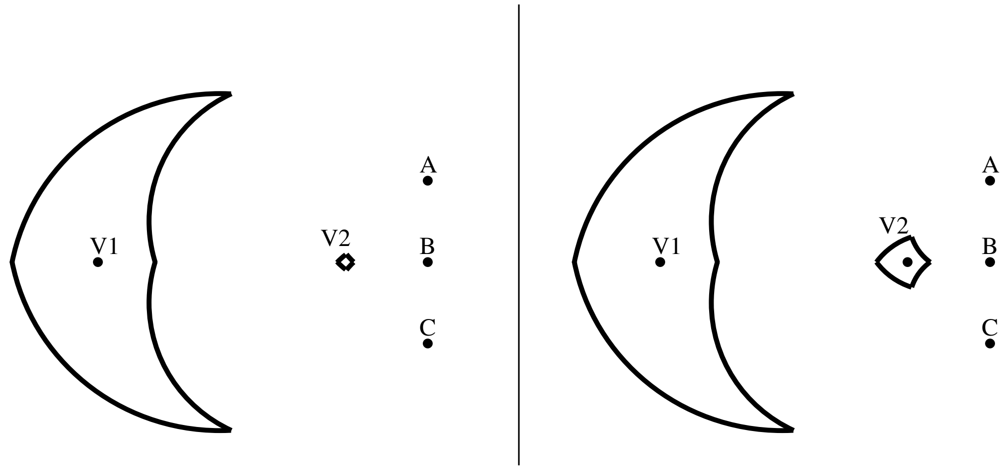

Figure 3.10: An additive error bound of $\pm 2.8°$ was introduced in the figure on the left while a multiplicative error bound of $\pm 20\%$ was introduced in the figure on the right. The change in size of the area of uncertainty is due to the change in distance of the viewpoint from the feature configuration.

Figure 3.10 shows two examples of the change in the area size as the viewpoint distance to the configuration changes. Due to the fact that whether the error is multiplicative or additive will affect results, both types of error were introduced. The figure on the left uses additive error bounds. The figure on the right uses multiplicative error bounds. The larger areas are for an error bound of $\pm 20\%$ or $\pm 2.8°$ in the visual angle measure. The smaller areas are for the corresponding amount of either multiplicative or additive error. The point representing V2 in the additive example is not shown because the area is so small.

The third parameter to consider is the shape of the configuration. Thus far the configuration has been held constant with all three features in a line and equally spaced. The comparative distance between straight line features and the angular relationship between non–linear features will both affect area size.

In Figure 3.11, dark lines surround the area of uncertainty resulting from a $\pm 30\%$ error bound in the visual angle measure for the straight line configuration with features A, B and C one unit apart and the viewpoint five units from the configuration. The resulting areas with feature C′ moved away from C are surrounded by dashed lines. C′ is 2 units from B on the left and 4 units from B on the right. In both figures, the dark and dashed lines coincide where the boundary is determined by the error ring for the circle through A and B. A skewness develops on the boundary determined by

feet to a straight line configuration with features located 1000 feet apart is considered to be the same as a distance of 4000 feet with the features 4000 feet apart).

the error ring for the circle through B and C. The area of uncertainty becomes smaller as C′ moves away from C.

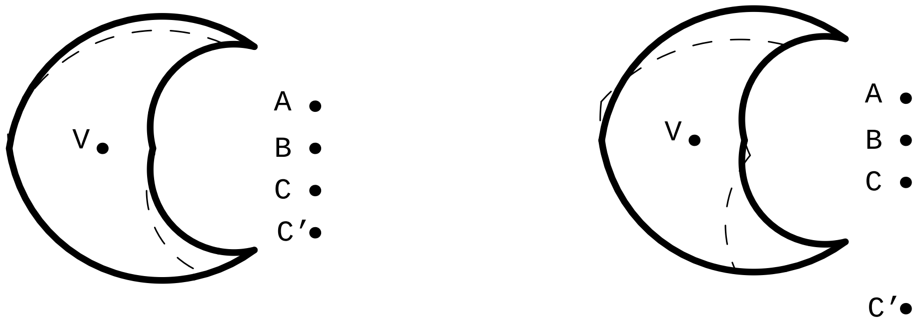

Figure 3.11: Both graphs show an error bound of $\pm 30\%$ in visual angle estimate with the viewpoint 5 units from configuration. The dark lines surround the area of uncertainty for the ABC configuration. The dashed lines surround the area for the ABC′ configuration.

The first step in analyzing how a change in the angular relationship between features affects the size of the area of uncertainty is to consider the LPB's and the resulting orientation regions. Note that regardless of the error size, the viewpoint cannot move out of the orientation region. The area of uncertainty will always be within that region.

If the locations of features A and C are fixed and feature B is moved along a line equidistant from A and C, the largest area of uncertainty for any give error amount will occur when B lies on the same circle as A, C and the viewpoint. Feature positions alone will not reveal whether or not this single circle condition, shown in Figure 3.4, exists. It is also necessary to know the viewpoint position. However, it can, for example, be ruled out if no part of the circle through the three features lies in the viewpoint's orientation region, which holds if the center feature is closer to the viewpoint than the other two.

Figure 3.12 shows the area of uncertainty for a $\pm 30\%$ error bound and the same basic configuration as in Figure 3.11 ($\alpha = \beta = 11.3°$). Feature B′ is one unit closer to the viewpoint in Figure 3.12a and one unit further away in Figure 3.12b. Although both show a decrease in area, the resulting area in Figure 3.12a is significantly smaller. Thus, the nonlinear configuration in Figure 3.12a is the least sensitive to error and would produce the most precise localization. Note that the length of the chord lying on the axis of symmetry does not change as feature B is moved. For a viewpoint facing the configuration, change in area is lateral only.

To summarize, the area of uncertainty corresponding to a given visual angle and error in that visual angle varies greatly for different configurations of features [Sutherland, 1992, Sutherland and Thompson, 1993]. Figure 3.13 shows a comparison of different areas with visual angle error range of $\pm 13.5°$ or $\pm 30\%$ in both α and β. Visual angles are identical for all configurations.

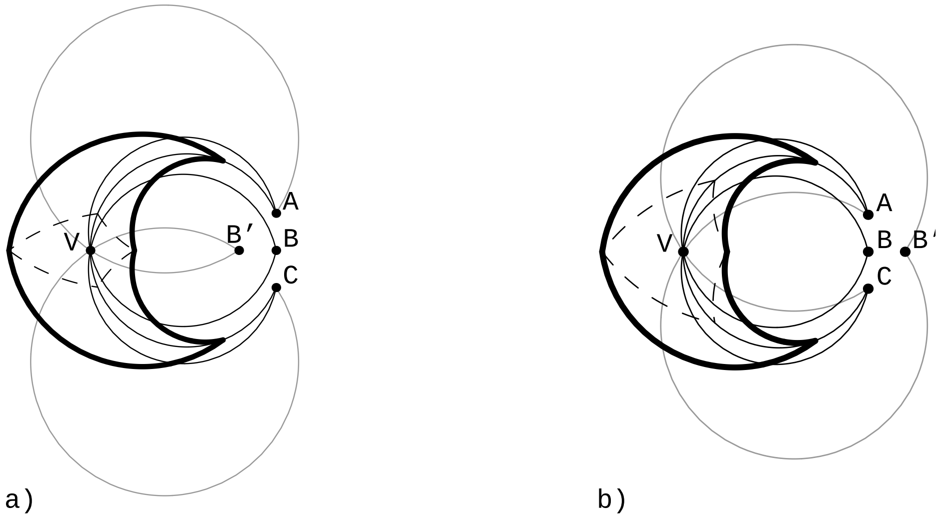

Figure 3.12: The heavy dark lines surround the area of uncertainty with a $\pm 30\%$ error bound in estimate of the visual angle, features A, B and C one unit apart and the viewpoint 5 units away. The dashed lines in a) surround the area of uncertainty resulting when B′ is one unit closer to the observer than B. The dashed lines in b) surround the area of uncertainty resulting when B′ is one unit further away. The error circles are black for the linear configuration and grey for the non–linear configuration.

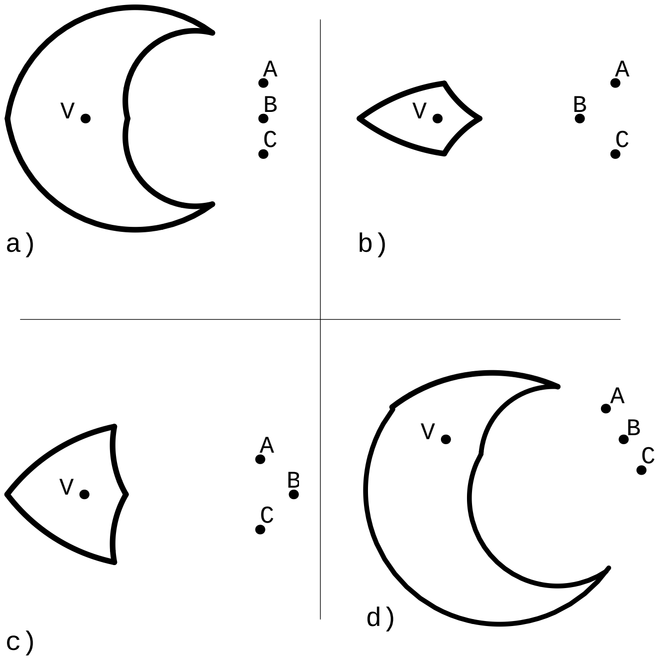

Figure 3.13: The area of uncertainty will vary with feature configuration: a) features in a straight line b) center feature pushed forward c) center feature pushed backward d) viewpoint off the line of symmetry of the configuration. All visual angles and error bounds are the same.

3.1.2 Shape of the area of uncertainty

Two shape properties must be considered. The first is symmetry. If the configuration is symmetric, the viewpoint is located on a line of symmetry of that configuration and the bound on the error in visual angle estimate is the same for both angle α and angle β, then the area of uncertainty will also be symmetric. If any of those properties do not hold, the area will not be symmetric.

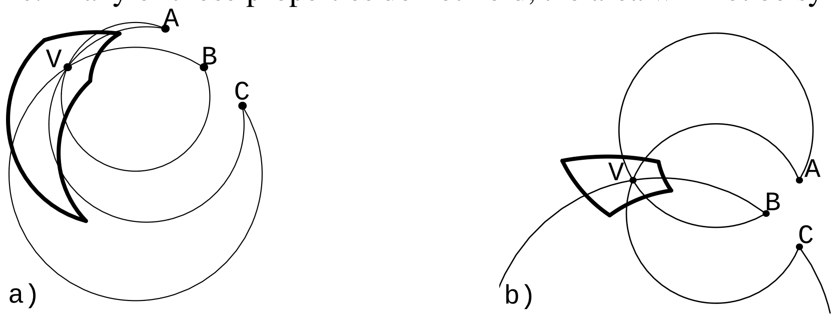

Figure 3.14: The shape of the area of uncertainty is skewed when the viewpoint is moved off the line of symmetry of the configuration.

Figure 3.6 shows how unequal error bounds in the estimates of angles α and β affect shape. In Figure 3.11, as feature C was pulled away from the configuration, the area decrease resulted in an asymmetric shape. Figure 3.14 shows how the area shape is skewed when the viewpoint is moved off the line of symmetry of the configuration.

The second property is eccentricity, a measure of the "roundness" of an area of uncertainty. The same three factors which affect the size of the area of uncertainty (the amount of error in the visual angle estimate, the relative distance of the viewpoint to the configuration and the configuration shape) also affect its eccentricity. However, the last factor has much more effect on the eccentricity than do the first two. The ratio W/L where L = the diameter of the area of uncertainty in the direction of feature B and W = the diameter of the area of uncertainty perpendicular to the direction of feature B can be used as a measure of eccentricity. Figure 3.7 showed an example of how the area of uncertainty changed as a function of error in visual angle measure. Note that the eccentricity of the outer area is slightly less than that of the inner area, but the basic shape remains the same. How the relative distance of the viewpoint to the configuration affects eccentricity was demonstrated in Figure 3.10. In both cases, the W/L measure is greater for the area surrounding V1. However, the small change in roundness is overshadowed by the large change in the size of the area of uncertainty. Figures 3.12 and 3.13 show the significant difference that configuration shape has on eccentricity.

3.2 Area of uncertainty - viewpoint inside configuration

It was shown in Section 3.1 that when all three point features, ordered A,B,C, lie within an angle of $180°$ from the viewpoint, the area of uncertainty consists of the intersection of two thickened rings passing through the AB and BC pairs. This area lies within the thickened ring passing through A and C. The sum of the visual angles to AB and BC equals the visual angle to AC. When the viewpoint lies within the configuration, the situation is modified. All three thickened rings contribute to the formation of the area of uncertainty. The sum of the visual angles to AB, BC and AC is $360°$. Figure 3.15 shows the modification in area of uncertainty shape caused when the viewpoint moves inside the configuration. The new area is formed by six circular arc segments. If, as in Figure 3.15, the viewpoint lies in the center of equally spaced features, this area is symmetric. If it is not located at the center and/or if the features are not equally spaced, the shape of the area changes.

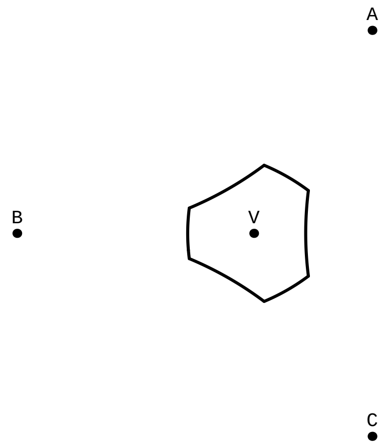

Figure 3.15: When the viewpoint is inside the configuration of features, the area of uncertainty is formed by six circular arc segments.

Whereas when the viewpoint was outside the configuration, the area of uncertainty was only bounded in size of underestimate of the visual angle by an estimate of $0°$ and, thus, unbounded in how far away from the configuration an estimate of the viewpoint could lie, the viewpoint being in the center of the configuration implies that it is surrounded by LPB's, putting a bound on how far away from any feature pair the viewpoint can be. Figure 3.16 shows one example of such a situation. Angle α measures $152°$. An underestimate of only $10°$ would put the viewpoint on the LPB joining features B and C. If the LPB is not crossed, the amount of error in underestimate is bounded by that amount.

Figure 3.17 shows a comparison of the areas of uncertainty with two different configurations of features, the viewpoints inside and outside the configurations and both additive and multiplicative error bounds. An additive error bound of $\pm 12°$ was introduced on the left while a multiplicative

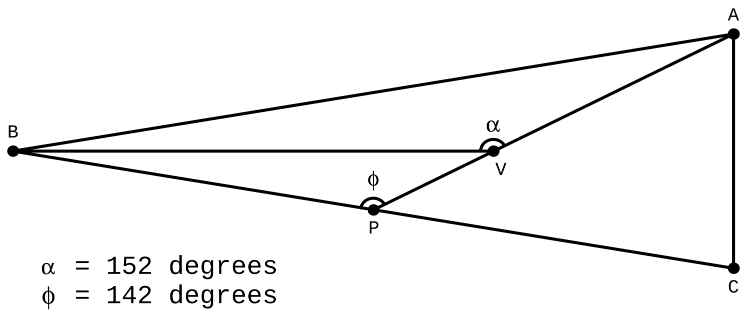

Figure 3.16: The LPB's can put a bound on the error in underestimation of the visual angle when the viewpoint lies inside the configuration of features.

error bound of $\pm 20\%$ was introduced on the right. The dashed lines surround the areas of uncertainty when the viewpoint is outside the configuration of features. The visual angles from the viewpoint to AB′ and B′C are both $60°$, so an error of $12°$ equals an error of 20%, causing the areas surrounded by dashed lines to be the same in corresponding rows. Since features A, B, and C do not move, the areas surrounded by solid lines are the same in corresponding columns. It is clear from this example that a comparison of area size cannot be made without a prior assumption on error type.

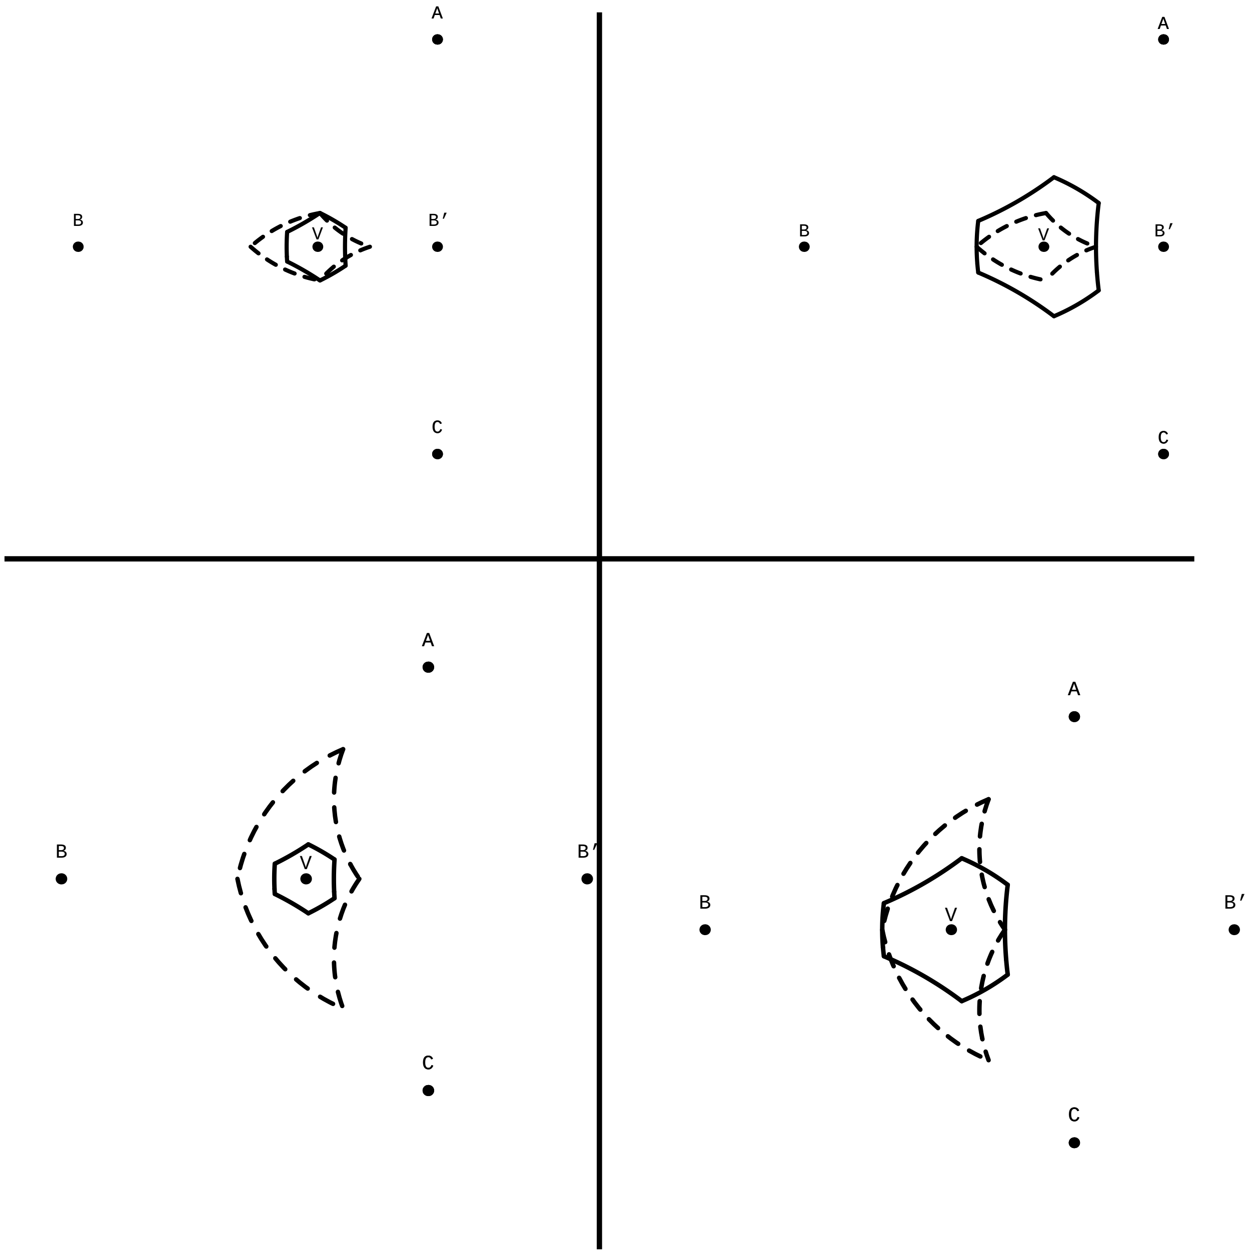

Figure 3.17: A comparison of the areas of uncertainty with the viewpoint inside and outside the configuration of features: the areas on the left result from an additive error bound of $\pm 12^\circ$ while those on the right result from a multiplicative error bound of $\pm 20\%$.

3.3 Area of uncertainty - 2.5 dimensions

When the environment is three rather than two dimensional, perfect knowledge of visual angles between three features does not necessarily produce a unique viewpoint. One of the following will hold:

Case I: All three features and the viewpoint lie in the same plane. This is identical to assuming a two-dimensional approximation of the environment. The only time that exact localization cannot be determined is when the three features and the viewpoint lie on a single circle. (See Figure 3.4.)
Case II: All three features and the viewpoint do not lie in a single plane. This implies that they do not lie on a single circle, excluding the exception in Case I. Let A, B, C be the features and V the viewpoint. Without loss of generality, it can be assumed that A is at the origin and B at point $(d_1, 0, 0)$ where d_1 is the distance between A and B. C is located at (Cx, Cy, Cz) such that d_2 (the distance between B and C) = $\sqrt{Cx^2 + Cy^2 + Cz^2}$.

The parametric equation of the surface of revolution through A, B and V is given by:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}^T = \begin{bmatrix} \frac{d_1}{2\sin\alpha} \sin u_1 + d_1/2 \\ \frac{d_1}{2\sin\alpha} \cos t_1 (\cos\alpha - \cos u_1) \\ \frac{d_1}{2\sin\alpha} \sin t_1 (\cos\alpha - \cos u_1) \end{bmatrix}^T$$

where $\pi/2 \leq t_1 \leq 3\pi/2$ [4], $\alpha \leq u_1 \leq 2\pi - \alpha$ and α is the visual angle.

The parametric equation of the surface of revolution through B, C and V is given by:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}^T = \begin{bmatrix} \frac{d_2}{2\sin\beta} \sin u_2 + d_2/2 \\ \frac{d_2}{2\sin\beta} \cos t_2 (\cos\beta - \cos u_2) \\ \frac{d_2}{2\sin\beta} \sin t_2 (\cos\beta - \cos u_2) \end{bmatrix}^T * R + \begin{bmatrix} d_1 \\ 0 \\ 0 \end{bmatrix}^T$$

where

$$R = \begin{bmatrix} \frac{Cx}{d_2} & \frac{Cy}{d_2} & \frac{Cz}{d_2} \\ \frac{-Cy}{\sqrt{Cx^2+Cy^2}} & \frac{Cx}{\sqrt{Cx^2+Cy^2}} & 0 \\ \frac{-Cx*Cz}{d_2*\sqrt{Cx^2+Cy^2}} & \frac{-Cy*Cz}{d_2*\sqrt{Cx^2+Cy^2}} & \frac{\sqrt{Cx^2+Cy^2}}{d_2} \end{bmatrix}$$

is the rotation matrix, $\pi/2 \leq t_2 \leq 3\pi/2$, $\beta \leq u_2 \leq 2\pi - \beta$ with β the visual angle.

Solving these two equations simultaneously gives the parametric equation of a curve in three-space. Although a general formulation of the equation of intersection of the surfaces is rather

[4] Because feature order is known, only half the surface of revolution is needed.

complicated, it is not difficult to express three of the variables in terms of the fourth (e.g., t_1, t_2, u_1 in terms of u_2) for any specific configuration. Thus, requiring the navigator to lie on both surfaces constrains the viewpoint to a single curve.

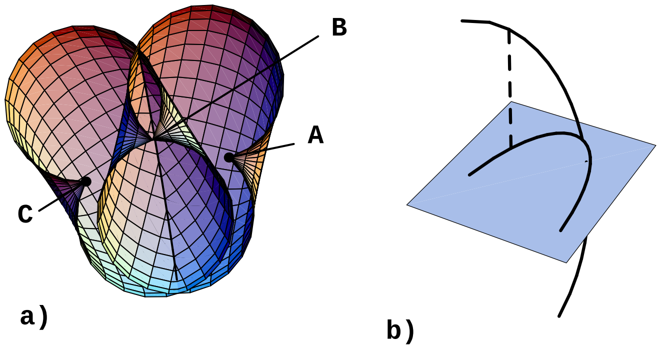

Figure 3.18: Perfect measurement of visual angles to three features in three-space constrains the viewpoint to a) the intersection of two surfaces. This intersection is b) projected onto the map to determine the viewpoint.

In Figure 3.18a, the scene is being viewed from the side opposite the viewpoint. The viewpoint lies somewhere on the black line running along the intersection of the two semi-surfaces.[5] As in the two-dimensional approximation of the environment, the third surface, formed by the visual angle to points A and C, will intersect the other two surfaces along the same line, so is of no help in providing a unique location.

Now, consider a 2.5 dimensional environment: The curve on which the viewpoint must lie is projected onto a surface, as shown in Figure 3.18b. Feasible viewpoint locations are at the points (x,y,z) where the z coordinate of the curve matches the elevation of point (x,y) on the surface. It is possible, but highly unlikely, that more than one point on the curve will intersect with the surface. The assumption can therefore be made that the location is unique. It follows that a perfect measurement of the visual angles to three features in 2.5 dimensional space will provide exact localization.

With imperfect visual angle measurements, the area of uncertainty is affected not only by the error amount and the two-dimensional configuration shape, but also by the elevation of the features measured from a horizontal plane passing through the viewpoint. Although the actual area size on a given plane will depend on local features around any individual viewpoint, the area of uncertainty projected onto a horizontal plane through the viewpoint provides a bound on lateral distance from that viewpoint. Figure 3.19 shows an example of the difference that elevated features can make in the area of uncertainty. The visual angles to feature points A, B and C are both $45°$. The angles from the viewpoint to the projection of points A, B and C on the plane are $52°$. An error bound of $\pm 10°$ on the actual visual angles produces an error of $+11°$ to $-17°$ for the projected angles. The smaller area on the plane is what the area of uncertainty would be for angles of $45°$ and an error bound of $\pm 10°$ if the feature points were at the same elevation as the viewpoint. The larger area is

[5] "semi" due to the feature ordering restrictions.

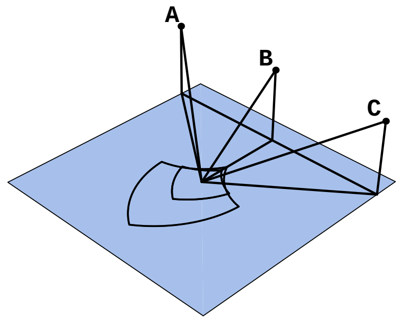

Figure 3.19: Visual angles to A, B, and C of $45°$ and an error bound of $\pm 10°$ produce the larger area of uncertainty. The smaller area is what would result with the same visual angles and error amounts if the features were in the same plane as the viewpoint.

the actual area of uncertainty for this configuration. It is not always the case that areas are larger if features are elevated. Elevation differences in features as well as the configuration of features affect area size.

3.4 Error distribution within the area of uncertainty

In addition to the differences in the area of uncertainty size and shape for identical visual angle measures, the distribution of points within an area is not solely 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 Figure 3.20, the area is divided into sections, each representing a combination of errors in the estimations of the visual angles α, from the viewpoint to features A and B, β, from the viewpoint to features B and C, and γ, from the viewpoint to features A and C such that $\gamma = \alpha + \beta$. Table 3.1 shows which conditions hold in each section. For any feature 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 feature 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 Figure 3.20, can be used to determine the probability of each of the numbered sections. The numbers in the diagram correspond to the numbers of the sections in the figure. If er_i is the

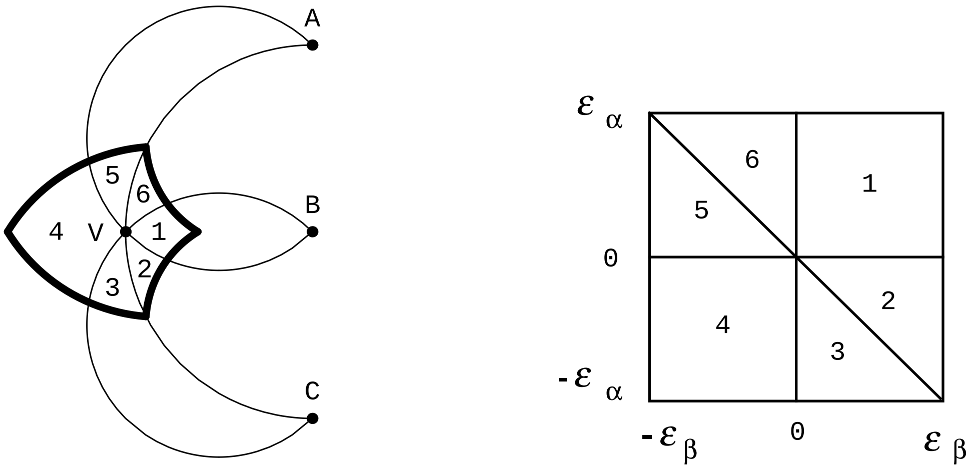

Figure 3.20: Uncertainty due to error in visual angle measure: On the left, arcs of the circles divide the area of uncertainty into 6 sections. The features are at A, B and C. The viewpoint is at V. α is the visual angle from the viewpoint to A and B. β is the visual angle from the viewpoint to B and 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.

error in angle i, the top half of the square represents $\mathrm{er}_\alpha > 0$, the right half of the square represents $\mathrm{er}_\beta > 0$, and the diagonal line represents the location where $\mathrm{er}_\gamma = \mathrm{er}_\alpha + \mathrm{er}_\beta = 0$. Note that the diagram has been drawn to match this particular figure, so the error bounds are the same for both α and β. When equality does not hold, the square becomes a rectangle. The probabilities can be easily computed as in the following examples:

$$\mathrm{Prob}(\alpha, \beta, \gamma \text{ overestimate}) = \frac{\text{Area 1}}{\text{Total Area}} = \frac{\epsilon_\beta * \epsilon_\alpha}{4 * \epsilon_\beta * \epsilon_\alpha} = \frac{1}{4}$$

or:

$$\mathrm{Prob}(\beta, \gamma \text{ under, } \alpha \text{ over}) = \frac{\text{Area 5}}{\text{Total Area}} = \frac{\frac{\epsilon_\beta * \epsilon_\alpha}{2}}{4 * \epsilon_\beta * \epsilon_\alpha} = \frac{1}{8}$$

The result of running 10,000 iterations, introducing a random, uniformly distributed error for the same feature configuration and error measure used in Figure 3.20 is shown in Figure 3.21. The true viewpoint is at the large black point. In all 10,000 iterations, the highest multiplicity of a single estimated viewpoint location was four. Sections 1 and 4, although quite different in size, each contained approximately one quarter of the points, as predicted, leading to the less dense distribution on the left of the scatter plot in Figure 3.21.[6] This can be predicted from the diagram on the right of Figure 3.20. The areas in the diagram corresponding to Sections 1 and 4 are each equal in size to one quarter of the total area of the square.

[6] Error amounts here and in subsequent sections were generated using an implementation of the Wichmann-Hill algorithm [Wichmann and Hill, 1982].

Area	Overestimate	Underestimate
1	α, β, γ	—
2	β, γ	α
3	β	α, γ
4	—	α, β, γ
5	α	β, γ
6	α, γ	β

Table 3.1: The visual angles are either under or overestimated in each of the 6 areas of Figure 3.20. 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.

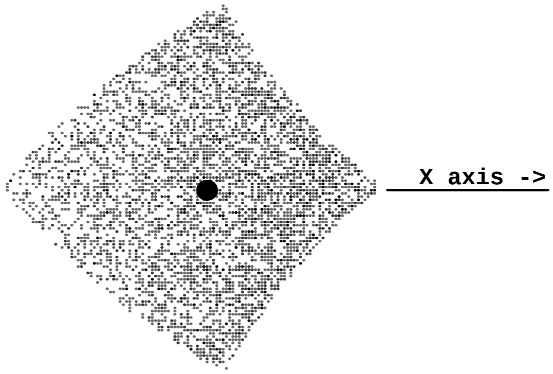

Figure 3.21: Assuming a uniform distribution of the error in visual angle measure, the distribution of points within the area of uncertainty is affected by the shape of that area. Each small point represents 1 of 10,000 iterations. The actual viewpoint is at the large black point. The example shown is for the area of uncertainty diagrammed in Figure 3.20.

A similar situation exists with the distribution of error by amount within any one section of the figure on the left of Figure 3.20. As an example, if Area 4 was divided according to error ranges of $0 - 15\%$ underestimate and $15 - 30\%$ underestimate, it would contain six subareas as shown in Figure 3.22. The conditions which hold in each area are given in Table 3.2. This property prevents the distribution within any one area from having anything close to a normal distribution in spite of the wedge shaped areas suggesting a denser distribution near the actual viewpoint.

If a normal distribution of error in visual angle measure is assumed, the distribution around the true viewpoint exhibits a definite central tendency with the same skewing as shown for the uniform distribution. In any case, if the presumed bounds on under and overestimate in the angle measure are equal and the area in which an overestimate lies is significantly smaller than is that for a corresponding underestimate, it may be better to overestimate.

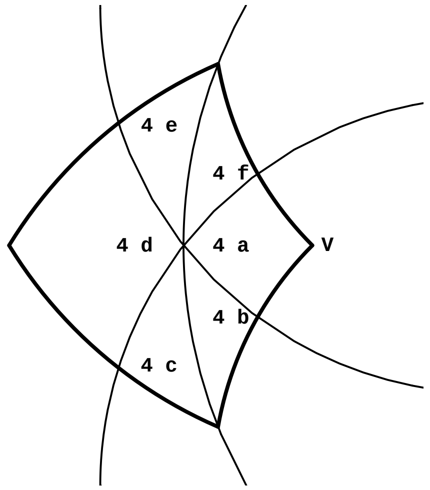

Figure 3.22: Conditions which hold in the 6 subareas of Area 4 are given in Table 3.2.

Subarea	Underestimate of $0 - 15\%$	Underestimate of $15 - 30\%$
4a	α, β, γ	—
4b	β, γ	α
4c	β	α, γ
4d	—	α, β, γ
4e	α	β, γ
4f	α, γ	β

Table 3.2: The visual angles are underestimated by differing amounts in each of the 6 subareas of Area 4 in Figure 3.20. 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.

3.5 Conditions affecting sensitivity

Two basic conditions affect the sensitivity of a configuration of features. The first is the rate of change of each visual angle measure as the viewpoint moves in the environment. The second is the rate of change of the ratio of angles α and β as the viewpoint moves.

As an example of the first condition, consider the viewpoints in Figure 3.10. Much less movement from the viewpoint is required to exit the area of uncertainty when the viewpoint is closer to the configuration because visual angles change at a greater rate with movement in that area than when the viewpoint is a large distance from the configuration. The total size of an area of uncertainty is based on this first condition. If the rate of change of the visual angle with respect to the viewpoint is high at a given location, a moving viewpoint will leave the area quickly. If the rate of change is low, a significant distance could be traveled before the area boundary is reached.

Figure 3.23 shows two three-dimensional graphs. The visual angle α to the features A and B is

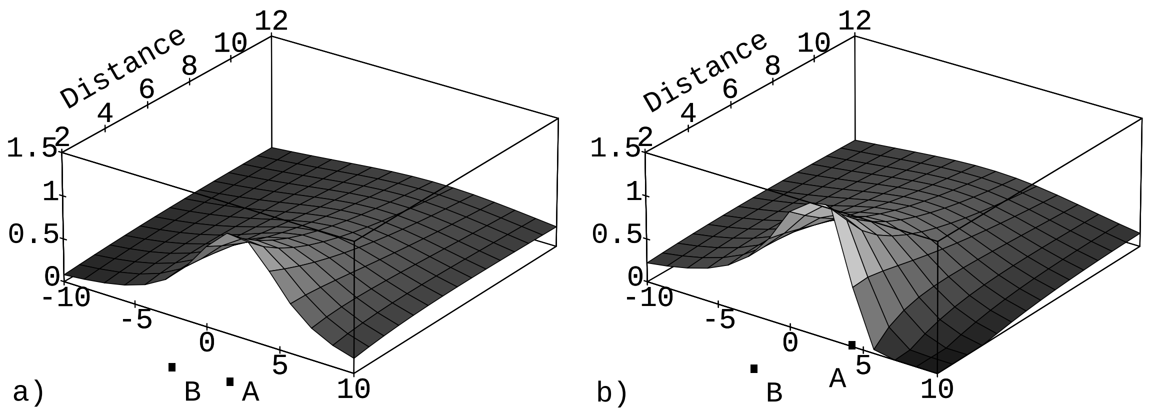

Figure 3.23: The visual angle α is represented by the height of the surface. Surface height is 0 at the LPB and outside of the orientation region. Distance is units to the furthest feature. Features are 4 units apart.

represented by the height of the surface. The rate of change of α is greatest where the surface slope is steepest. Less error in localization should be made when the viewpoint is situated at a point of steep slope than than when it is at a point where the slope is shallow. A viewpoint located at (0,4) on the XY plane of Figure 3.23a with Y being the axis labeled as Distance would correspond to V2 in Figure 3.10 and is at a point where the rate of change of visual angle α is much greater than it is at a viewpoint located at (0,20) on the same XY plane (off the graph of Figure 3.23a), which would correspond to V1 in Figure 3.10. Moving feature A closer to the graph of possible viewpoint locations, as shown in Figure 3.23b, causes an increase in the height of the surface at x = 0 and a decrease in the height of the surface at x = 5. The visual angle to features A and B from a viewpoint at (0,4) on the XY plane will change more rapidly as that viewpoint moves, for example, toward the point (4,4) than would the visual angle for a viewpoint moving between the same two locations in Figure 3.23a. This results in the bounds on the area of uncertainty created by the error in a given angle being close to the true viewpoint when that viewpoint is located at a point of steep slope on the graph. The closer to the viewpoint those bounds are, the smaller the area of uncertainty.

The second condition, the comparison of angle measure, also affects sensitivity. The rate of change of this ratio will be large when α is increasing rapidly and β is decreasing rapidly or vice-versa. This affects the eccentricity of the area of uncertainty. As an example of this second condition, consider a viewpoint facing the configuration in Figure 3.12a. Very little lateral movement off the viewpoint is required to exit the area of uncertainty in the non-linear configuration because the ratio of the visual angles is changing at a much greater rate than it does with the linear feature configuration. It follows that Areas 2, 3 5 and 6 of the area of uncertainty shown in

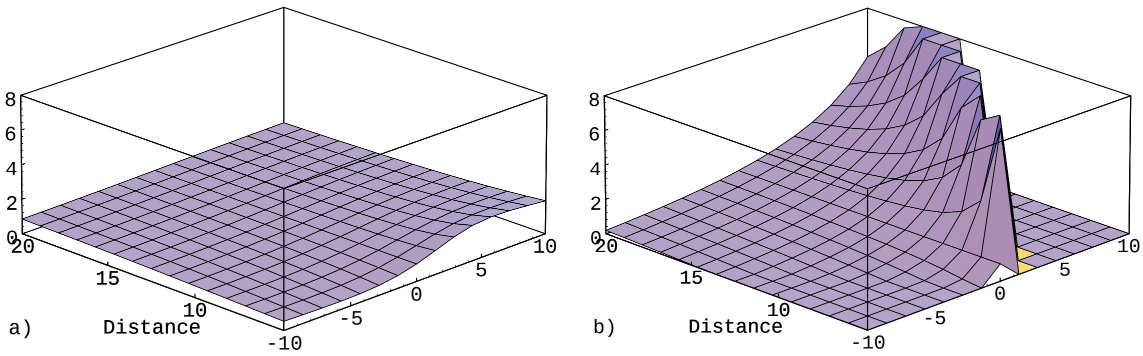

Figure 3.24: The ratio of the visual angles α/β is represented by the height of the surface. The height is 0 outside of the orientation region. The distance shown is in linear units to the furthest feature.

Figure 3.20 will be small. In the graphs of Figure 3.24, the ratio α/β is represented by the height of the surface. The features A, B, and C (not shown) are located at (0,4), (0,0) and (0,-4) in Figure 3.24a and at (0,4), (4,4) and (0,-4) in Figure 3.24b. This corresponds to the feature locations in Figure 3.12a. The maximum rate of change occurs at the points of steepest slope. The visual angle ratio α/β from a viewpoint located at (0,20) on the XY plane of Figure 3.24b with the Y-axis labeled Distance, corresponding to the viewpoint in Figure 3.12a, would change more rapidly as the viewpoint moves, for example, toward (4,20) than would the visual angle ratio in Figure 3.24a for the same viewpoint and the same move.

To summarize, the rates of change of the visual angles and their ratio in an area around the actual viewpoint location depend on configuration shape and the location of the viewpoint with respect to the configuration. Large rates of change result in less sensitivity to error.

3.6 Ordering features

The work described thus far in this chapter assumed that features had been identified and ordered. An interesting question to pose is whether or not visual angle measures can aid in ordering features. The answer is mixed.

3.6.1 Angle measures from one viewpoint

Assume that for a two-dimensional approximation of the environment, the visual angle measure is known exactly. When the configuration of three figures is linear, there are two possible viewpoint locations, one on each side of the single LPB passing through the features. It is easily shown that both of these locations are physically possible, regardless of spacing between features: Without loss of generality, it can be assumed that features are at $(0,0)$, $(1,0)$ and $(X_0,0)$ where $X_0 > 0$.

The viewpoint could then be either at V_1 or V_2 in the figure on the left of Figure 3.25.

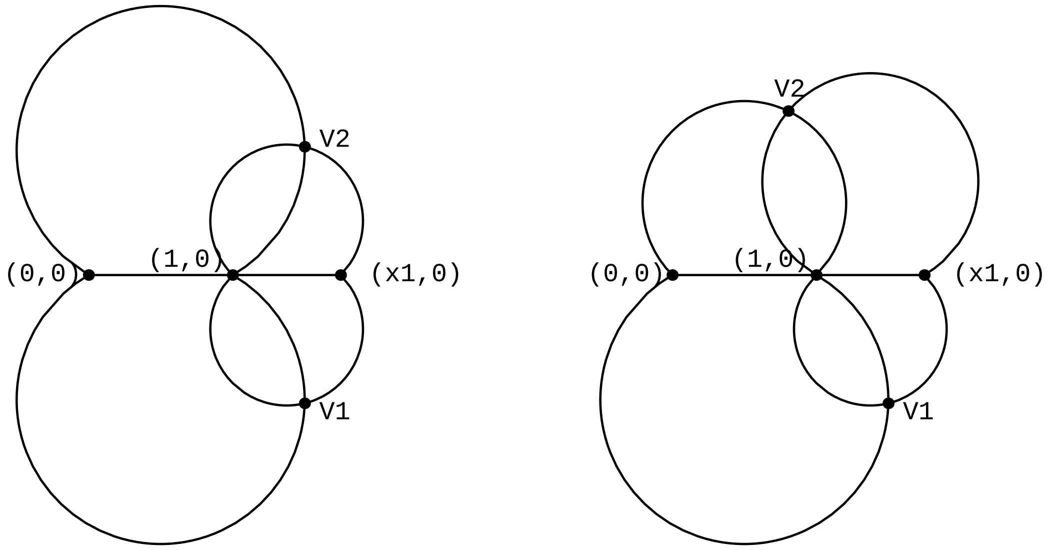

Figure 3.25: The viewpoint can be either at V_1 or V_2 when features lie in a straight line. In the figure on the left, the viewpoint rotates around the feature line. In the figure on the right, it is assumed that the viewer is standing on a surface on either side of the feature line.

If the viewpoint coordinates are given by (x, y), eliminating a_1 and r_1 in the system of equations representing the circle with radius r_1 and center (a_1, b_1) through $(0, 0)$, $(1, 0)$ and (x, y):

$$a_1^2 + b_1^2 = r_1^2 \tag{3.1}$$

$$(1 - a_1)^2 + b_1^2 = r_1^2 \tag{3.2}$$

$$(x - a_1)^2 + (y - b_1)^2 = r_1^2 \tag{3.3}$$

gives the resulting equation:

$$x^2 - x + y^2 - 2b_1 y = 0 \tag{3.4}$$

Eliminating a_2 in the corresponding system of equations representing the circle through $(1, 0)$, $(X_0, 0)$ and (x, y):

$$(1 - a_2)^2 + b_2^2 = r_2^2 \tag{3.5}$$

$$(X_0 - a_2)^2 + b_2^2 = r_2^2 \tag{3.6}$$

$$(x - a_2)^2 + (y - b_2)^2 = r_2^2 \tag{3.7}$$

gives the resulting equation:

$$x^2 - (1 - X_0)x + y^2 - 2b_2 y + X_0 = 0 \tag{3.8}$$

Solving equations 3.4 and 3.8 simultaneously gives the point $(1, 0)$, which is the center feature point and:

$$x = \frac{4(b_1 - b_2)^2 + 4b_1 X_0 (b_1 - b_2)}{4(b_1 - b_2)^2 + X_0^2} \tag{3.9}$$

$$y = \frac{2X_0(b_2 + b_1(X_0 - 1))}{4(b_1 - b_2)^2 + X_0^2} \tag{3.10}$$

This second solution provides the coordinates for possible viewpoints. Due to the fact that $X_0 > 1 > 0$, the denominators will always be non-zero, with the result that both x and y are always defined. When $b_1, b_2 > 0$, y is greater than 0 and when $b_1, b_2 < 0$, y is less than 0, providing two possible viewpoint locations. The viewpoint can, of course, be determined exactly if either the leftmost or rightmost feature is known.

The implicit assumption was made in the above argument that the viewpoint would rotate around the feature line, i.e., a viewer standing at V_1 would have to be standing on its head at V_2 in order for the angle measures to be in the same left to right order. If, as in many realistic applications, the viewer is standing on a surface, the viewpoint will still be restricted to two possible locations, but as shown in the figure on the right of Figure 3.25, these locations are not directly across the feature line from each other.

When the feature configuration is non-linear and the viewpoint is outside the configuration, a partial ordering combined with constraints on angle measure can limit the viewpoint location. Table 3.3 shows the possible orientation regions for the viewpoint location in Figure 3.26 if the position of only one feature is known. Furthermore, if it is known that B is in the center and if the visual angle to the outer two features is greater than angle ABC, then the viewpoint must lie in orientation region 2 because the angle to the two outer features from any point in orientation region 5 is less than angle ABC. It follows that if it is known that A is to the left and the visual angle to AB is greater than the angle ACB, the viewpoint must lie in orientation region 2. Likewise, if it is known that C is to the right and the visual angle to BC is greater than the angle BAC, the viewpoint must also be in orientation region 2. In general, for static localization, knowing the position of one feature point restricts the viewpoint to one of two orientation regions, one with the center feature closer to the viewpoint than the line joining the two outer features and one with the center feature further from the viewpoint than that line. If the visual angle is greater than the angle subtended by the three features, the orientation region, and thus total ordering can be uniquely determined. If there is no knowledge about feature order, the feature positions can still minimally constrain the viewpoint. For example, angle ABC in Figure 3.26 measures $30°$. Angles BAC and BCA measure $75°$ each. Any visual angle measure between the two outer features greater than $30°$ eliminates orientation region 5. Any visual angle measure between the two outer features greater than $75°$ eliminates orientation regions 1, 3 and 5.

It follows that when the assumption is made that any given error in visual angle measure lies within certain bounds and the above hold for those bounds, the same restrictions can be made. As an example, if it is known that B is in the center and that the error in visual angle measure is bounded by $\pm 20\%$, then an angle measure between the two outer features of $38°$ or greater would force the viewpoint to be in orientation region 2.

Knowledge	Visual Angle	Orientation Region
B in center	$\gamma > ABC$	2
	$\gamma < ABC$	2 or 5
A to left	$\gamma > ACB$	2
	$\gamma < ACB$	2 or 3
C to right	$\gamma > BAC$	2
	$\gamma < BAC$	1 or 2

Table 3.3: Possible orientation regions for the viewpoint location if the position of only one feature is known. The numbers refer to the numbered regions in Figure 3.26. Angle γ is the visual angle between the outer two features.

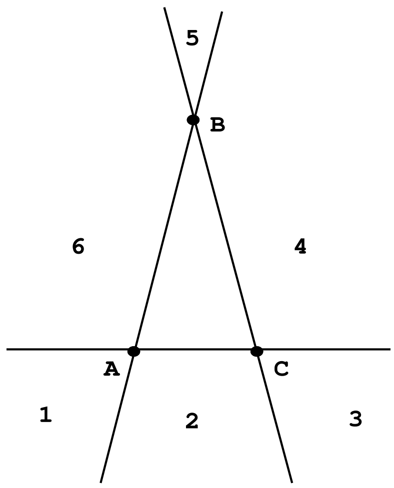

Figure 3.26: A partial ordering of non-linear features limits the orientation regions in which the viewpoint can lie.

When a 2.5-dimensional environment such as the earth's surface is assumed, feature altitude could perhaps help determine order, but only in conjunction with other visual cues. As an example, consider the sequence of images of three point features in Figure 3.27. One of the features is higher than the other two. However, as the viewpoint position changes from frame to frame, one is not able to distinguish the difference in height from the points alone. In the sequence of images in Figure 3.28, the boxes and termination points of the rectangular solids provide clues which the points alone did not, making a determination of the tallest feature fairly simple.

3.6.2 Angle measures from multiple viewpoints

A second question to pose is whether or not taking angle measures from more than one viewpoint would help in resolving the ambiguity resulting from no or partial knowledge of order. First, assume a two-dimensional approximation of the environment, exact visual angle measure and an

Figure 3.27: Four different views of three point features, one of which is taller than the other two.

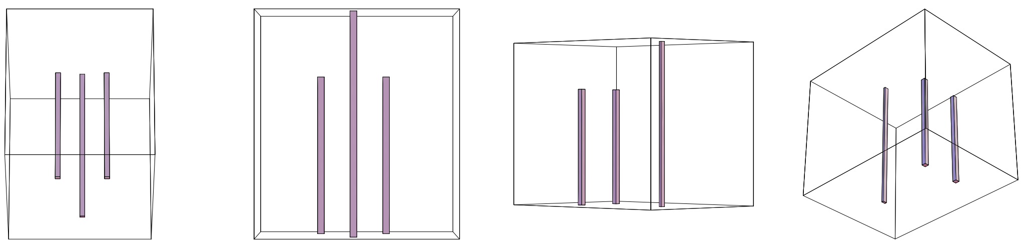

Figure 3.28: The same sequence of features as in Figure 3.27 with extended altitudes and 3-d boxes.

equally spaced, linear configuration of features. It is clear that for any movement on one side of the single LPB, a comparable one is possible on the other side with the same change in visual angle measures. Moving and taking a second measure will not provide any additional information.

If the configuration of features is straight line asymmetric, as in Figure 3.25, or non-linear, moving and taking a second measure will provide information which will help in determining feature order. For example, a move parallel to the line passing through the outer two features followed by a second error-free angular reading will provide sufficient information to determine in which orientation region the viewpoint lies. However, knowledge of the visual angle measures alone are not enough to determine in which direction to head so that movement is parallel to the line passing through the outer two features.

One possible strategy is to make a small move toward a feature. This can easily be accomplished by keeping the feature centered in the image plane. Consider the configuration on the right of Figure 3.25. As was previously discussed, the viewpoint can be at either V_1 or V_2. Movement toward the leftmost feature from each of the two viewpoints would result in different visual angle measures after the move. A series of experiments were run in simulation, using different positions of the center feature, different length moves and different bounds on the errors occurring in the visual angle measures. One thousand runs were made with each situation.

The scenario is as follows: The viewer computes the visual angle measure and determines the two possible viewpoints. Due to the fact that the measurements contain error, these locations are probably not correct. It then computes what visual angle measures should be after a move of a given distance toward the leftmost feature from each of the two possible viewpoints. Finally, it makes that move and computes the actual visual angle measures. This last computation can also contain error. It compares the last measurements to the two computations of what the angles should measure and chooses its location determined by which computation is closer to the actual measurements.

When the measurements were error free, the only situation which did not produce the correct result was when the center feature was equally spaced between the outer two. When error was introduced, the bounds on that error, the length of the move and the distance from the configuration all affected the results. As the center feature was moved closer to the midpoint between the two outer features, performance deteriorated. The distance of the move was given as a percent of the distance between the outer two features (feature spread) rather than as an absolute distance. The initial viewpoints were randomly picked in an area extending from the feature line a distance of 1.5 time the distance between the outer two features. This area was divided into three equally sized subareas with Area 1 being the furthest from the configuration. This division of space, as diagrammed in Figure 3.29, demonstrates how the distance from the feature line affects results.

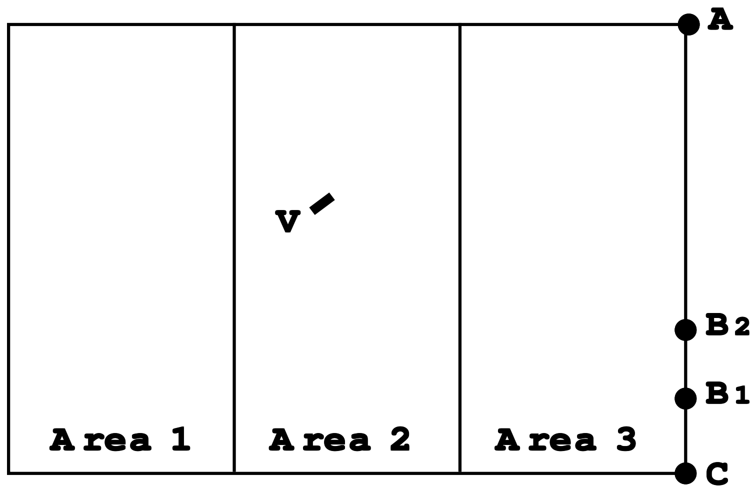

Figure 3.29: Initial viewpoints, such as V, were picked in the three labeled subareas. A movement of a percent of the feature spread was made toward feature A.

For the results summarized in Table 3.4, the actual move distances were 400 and 800 meters with the spread between features A and C equal to 20,000 meters and the random viewpoints extending 30,000 meters back from the feature line. Feature B was located 2760 meters from feature C. For the results summarized in Table 3.5, the actual move distances were again 400 and 800 meters with the spread between features A and C equal to 20,000 meters and the random viewpoints extending 30,000 meters back from the feature line. However, feature B was located 5760 meters from feature C. Figure 3.29 shows a diagram of one such move with B_1 the position of the center landmark for the first set of results, B_2 the position of the center landmark for the second set of results, and the line from V showing the approximate length of a move of 4% of feature spread. The results show

how performance is improved when the middle feature is moved away from a position equally spaced between the two outer features.

	Movement as % of Feature Spread					
Measurement	Area 1		Area 2		Area 3	
Error	2%	4%	2%	4%	2%	4%
5%	.91	1.00	.89	1.00	.87	.96
10%	.77	.97	.78	.93	.80	.88
15%	.69	.94	.72	.86	.75	.83

Table 3.4: Percent of correct orientation region guesses, given as a decimal, for a linear asymmetric configuration of features after one local move and a second angle measure. The center feature was 2760 meters from the feature on the right. The total spread between features was 20,000 meters.

	Movement as % of Feature Spread					
Measurement	Area 1		Area 2		Area 3	
Error	2%	4%	2%	4%	2%	4%
5%	.62	.74	.59	.69	.65	.73
10%	.56	.62	.53	.59	.62	.68
15%	.55	.56	.51	.54	.62	.64

Table 3.5: Percent of correct orientation region guesses, given as a decimal, for a linear asymmetric configuration of features after one local move and a second angle measure. Center feature was 5760 meters from the feature on the right. The total spread between features was 20,000 meters.

A similar set of experiments was done with non-linear feature configurations. Random viewpoints were chosen in the same three areas used in the straight line trials but the constraint of requiring a second position in the "opposite" orientation region with the same visual angle measures to features had to be imposed. There was still, of course, a need for partial ordering knowledge such as "B is in the center". Results for the three configurations shown in Figure 3.30 are summarized in Tables 3.6, 3.7 and 3.8. The question marks lie in the areas which were being chosen between for possible viewpoint location.

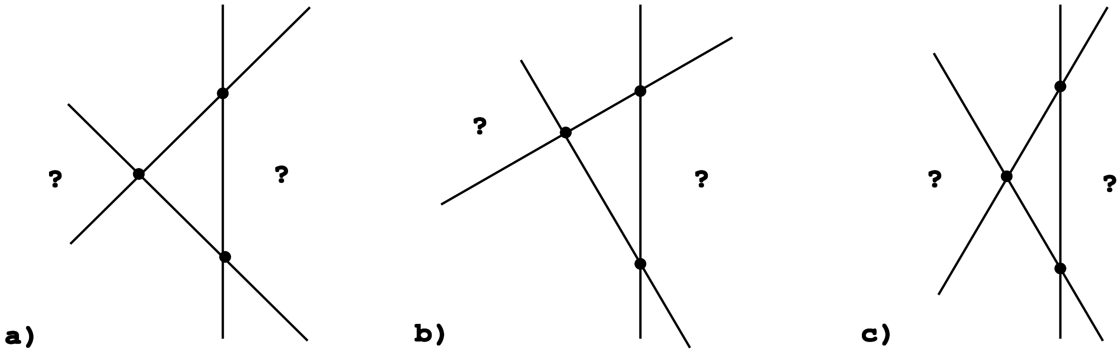

Figure 3.30: The three configurations of features for which the localization data is given. The question marks lie in the areas which were being chosen between for possible viewpoint location.

	Movement as % of Feature Spread					
Measurement	Area 1		Area 2		Area 3	
Error	2%	4%	2%	4%	2%	4%
5%	.87	.89	.85	.87	.91	.91
10%	.83	.84	.81	.83	.86	.84
15%	.78	.78	.78	.80	.81	.78

Table 3.6: Percent of correct orientation region guesses, given as a decimal, for the non-linear configuration of features shown on the left of Figure 3.30 after one local move and a second angle measure.

	Movement as % of Feature Spread					
Measurement	Area 1		Area 2		Area 3	
Error	2%	4%	2%	4%	2%	4%
5%	.99	.98	.97	.97	.94	.92
10%	.95	.95	.94	.94	.89	.87
15%	.93	.94	.93	.94	.83	.83

Table 3.7: Percent of correct orientation region guesses, given as a decimal, for the non-linear configuration of features in the center of Figure 3.30 after one local move and a second angle measure.

	Movement as % of Feature Spread					
Measurement	Area 1		Area 2		Area 3	
Error	2%	4%	2%	4%	2%	4%
5%	.82	.84	.86	.87	.95	.96
10%	.78	.80	.85	.86	.92	.93
15%	.76	.76	.82	.83	.89	.90

Table 3.8: Percent of correct orientation region guesses, given as a decimal, for the non-linear configuration of features on the right of Figure 3.30 after one local move and a second angle measure.

As discussed in Section 3.1, the area of uncertainty increases in size as the viewpoint approaches a single circle configuration with the features. There is a large area in the right section of the rightmost configuration of Figure 3.30 for which this property can hold. The result is that movement toward the leftmost feature changes visual angle measure very little, accounting for the improved performance shown in Table 3.8, particularly in the area closest to the configuration.

These simulations show that small movements can aid in reducing the ambiguity resulting from only partial knowledge of order, especially when the three feature points are non-linear. Does this process fall into the category of *active perception* or *active vision* as in [Bajcsy, 1988]? Bajcsy states: "perceptual activity is exploratory, probing, searching;" In that general sense, this process could be considered a type of active perception. Something is being actively done to aid in the perceptual process. However, most research in active machine perception has dealt with ideas such as the use of retina-like lenses [Sandini and Tagliasco, 1980] or gaze control [Ballard, 1989] in which small, often continuous movement of the imaging device implements low level visual goals. The activity here involves moving the viewpoint a distance which, although small compared to the overall size of the environment, is large compared to that involved in movements required for gaze control, and compared to those tightly controlled movements is rather loose and full of error. It is more closely tied to the concept of *purposive vision* as described in [Maver and Bajcsy, 1993].

Animals will use motion parallax to judge distance [Gallistel, 1990] but this again usually consists of either small movements of the head from side to side or a series of vertical head bobs. The zigzag flight motions employed by bees and wasps to measure distance as they back away from a nest or feeding site [Wehner, 1981] contain large movements. However, the decision being made here (i.e., what side of the configuration are we on?) is at a higher level of cognition than is that of estimating distance, regardless of the length of the movement made in the estimating process.

Chapter 4

Robot navigation in unstructured environments

4.1 Introduction

Determining locations in an unstructured environment is a major task for an autonomous mobile robot navigating with a map. It must, as in Figure 4.1, 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.

The navigation problem of self-localization in the world is an absolute orientation problem. *Absolute orientation* [Haralick *et al.*, 1989] is defined as the recovery of the relationship between two 3D coordinate systems. To find this relationship, one must determine scaling, translation and rotation. The early algorithms for determining absolute orientation were developed by researchers in the field of photogrammetry. In 1958, E. H. Thompson [Thompson, 1958] matched model coordinates with geodetic coordinates. He projected the sphere onto the complex plane and used a system of linear equations to solve the 3D $\rightarrow$ 3D absolute orientation problem for rotation only. In 1973, Sanso [Sanso, 1973] extended this solution to include scaling and translation. Sanso's solution used the quaternion algebra, a method which has recently regained popularity, particularly in the area of computer vision.

Most recent work in autonomous robot navigation has been done in structured indoor environments. A sampling of this work can be found in [Crowley, 1985, Tsubouchi and Yuta, 1987, Kriegman *et al.*, 1989, Kosaka and Kak, 1992, Fennema *et al.*, 1990, Ferrari *et al.*, 1990, Atiya and Hager, 1993, D'Orazio *et al.*, 1992]. Of the work in outdoor navigation, the following is associated most closely with this research.

Dickmanns' experimental five ton van "VaMoRs" navigated in outdoor environments, but remained (understandably) on paved roads, using road edges for tracking [Dickmanns and Graefe, 1988b, Dickmanns and Graefe, 1988a, Dickmanns, 1992]. Thus, the environment was large, but

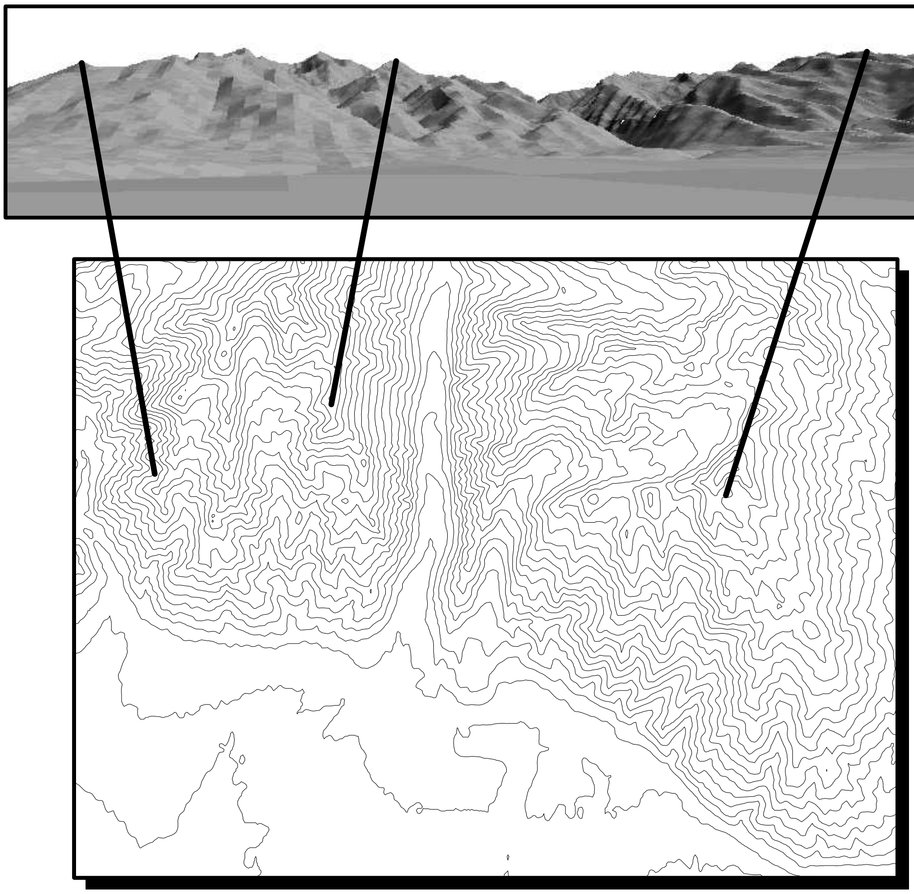

Figure 4.1: An established match between landmarks in the environment and features on a map can be used by a robot navigator to determine its own location.

structured, and the edges used for determining locations were close to the vehicle. Work in both the DARPA ALV program and the Carnegie Mellon Navlab also was carried out in a structured outdoor environment with roads, their intersections and lanemarkings the predominant terrain features [Davis *et al.*, 1987, Thorpe *et al.*, 1987].

Talluri and Aggarwal, [Talluri and Aggarwal, 1992], used a constrained search paradigm, deriving a significantly pruned search space in which to determine robot location by making terrain elevation assumptions based on altimeter and compass readings. Although their assumptions hold when sensing errors are small, they failed to simulate the frequently occurring conditions in a real outdoor environment which cause altimeter readings to err in the magnitude of hundreds rather than their stated tens of feet and compass readings to be rendered completely useless.

Yacoob and Davis, [Yacoob and Davis, 1992], used a single point range finder, an altimeter, a compass and an inclinometer along with Digital Elevation Map (DEM) data to prune the search space. They also made assumptions such as a rather small error bound on both the range finder and altimeter measurements of $\pm 2\%$. The work of [Levitt *et al.*, 1987, Levitt *et al.*, 1988, Levitt and Lawton, 1990, Kuipers and Levitt, 1988] referenced in Chapter 3 addressed the problem of navigating in large-scale space, defining it as space whose structure cannot be observed from a single viewpoint. They built up a representation of such space based upon observations of visual events. Due to the fact that their environment was rich in landmarks and that their "qualitative navigation" approach was intended to function with only approximate knowledge, types and magnitudes of errors were not an issue. [Dai and Lawton, 1993] expanded on this approach, developing algorithms which deal with landmarks for which there are no range estimates. Again, the number of available landmarks in their environment made it possible to navigate by moving from one orientation region to another without straying too far from the desired path.

Autonomous underwater vehicles (AUV's) often navigate in their completely unstructured environments by either adding structures for reference or navigating only over short local distances. Methods proposed include the use of a network of bottom-laid acoustic transponders combined with a high accuracy dead reckoning system [Babb, 1990], electronic still cameras with high performance, large format CCD imagers [Harris *et al.*, 1987], and the exploitation of surface texture properties [Negahdaripour *et al.*, 1990, Negahdaripour and Yu, 1990]. However, even when several methods are combined, the error involved is significant.

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 instrumental delays [Wu and Melbourne, 1993]. Selective availability encryption degrades the positioning accuracy for any user without access to the encryption [Wu *et al.*, 1992]. These users make up a large group: the entire civilian community. Averaging techniques such as the Extended Kalman filter are used to add differential corrections to the GPS measurements and improve accuracy [Negast and Paschall, 1992]. When using the Kalman filter in this type of application, the dynamic and observation models which are assumed are often incorrect due to causes such as cycle slips, leading to significant errors or possible nonconvergence in the filter results [Lu and Lachapelle, 1992].

DGPS (Differential Global Positioning System)/radiobeacon broadcast networks have been developed to provide position fixing for maritime applications which require more than an accuracy within 100 meters [Enge, 1993]. A local reference station with a high quality GPS receiver and an antenna at a known location estimates the error in the GPS signal and transmits it as a correction to users within the range of the station. This is necessary because GPS is not accurate enough for such maritime applications as navigating a harbor entrance or offshore surveying.

Atmospheric conditions also affect GPS. The signal strength is a function of the thickness of the air mass which the signal passes through [Cohen *et al.*, 1993]. Experiments have been conducted [Tranquilla and Al-Rizzo, 1993] 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. The elevation angle of the signal also affects the result. The path length through the clouds usually increases as the elevation angle decreases. The longer the path length through the clouds, the larger the error. Note that 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 [Mattos, 1992]. Duerr, [Duerr, 1992], has studied the effect of terrain masking on the precision of GPS readings. Most published results on GPS accuracy have assumed an unobstructed view of the sky. Duerr showed that a GPS receiver located in a valley can have a significant loss of accuracy for long periods of time. The accuracy is a function of time of day and latitude. These results not only affect land vehicles but also aircraft flying at low altitudes and some missile systems.

In most of the above cases, assumptions such as existence of a road, nearby features or special instruments were made. In the particular case of GPS, the assumption is that measurements will be precise. Without making these assumptions, an outdoor, unstructured environment with its rugged terrain and few distinguishable landmarks presents unique challenges to a robot navigator:

- Errors in distance traveled can be significant and unpredictable, compounding as the robot moves. Whether the robot is wheeled [Talluri and Aggarwal, 1992] or legged [Hoffman and Krotkov, 1993], terrain surface conditions exacerbate these errors for land vehicles. Existing tools for determining position and translational motion in underwater vehicles (e.g., linear accelerometers or Doppler techniques) are insensitive to slow positional drift [Negahdaripour *et al.*, 1990]. Kosaka and Kak, in their very thorough treatment of indoor navigation techniques [Kosaka and Kak, 1992], 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.

- 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 [Talluri and Aggarwal, 1992] 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 which requires absolute bearings to two or more landmarks, cannot be used.

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

- 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 the discretization error. In addition, conditions such as fog surrounding a land navigator or particles suspended in the water around an undersea system can cause image blurring and back-scattered noise, both of which affect sensor accuracy [Stewart, 1991].

- 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 (DEMs) [Thompson *et al.*, 1993]. 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 after the fact. 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 above mentioned Extended Kalman filter (EKF) assumes white Gaussian noise [Kalman, 1960, Smith and Cheeseman, 1986, Matthies and Shafer, 1987, Ayache and Faugeras, 1989]. 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 [Sorenson, 1970]. Due to the problems mentioned above, 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 [Cui *et al.*, 1990]. 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, a variety of accurate sensors were available, the environment surrounding the navigator was conducive to signal transmission and objects sensed were bar codes rather than flora covered terrain, outdoor navigation would be much easier. Unfortunately, that is not the case. A mobile robot, navigating in real time and limited in the number of measurements it can take, must make the best use of the sensory data it has. In short, a successful navigator must use *inexact navigation* (i.e., produce good localization with only approximate information)[Thompson and Kearney, 1986]. This chapter approaches the task of navigating in an unstructured world by critically analyzing the types of errors which will occur and, utilizing only those tools and techniques which are available in such environments, exploiting the geometric properties involved in the localization process so that steps can be taken to decrease the possibility of error before it occurs. The result is less error and greater success for whichever filter is used to handle the error that does exist.

4.2 Localizing to a point

Drawing on the results of Chapter 3, if the landmarks are point features[1], and left to right order of those landmarks is known, it follows that, for a navigator located at the viewpoint, the configuration

[1] The mountain peaks in Figure 4.1 are one example of such features.

of landmarks used for localization will significantly affect how successful that navigator is.

As shown in Section 3.3, it can be safely assumed that, for a navigator traveling on terrain (as opposed to being in space), exact knowledge of the visual angles between three points constrains the viewpoint to the intersection of three surfaces and the terrain.

When the visual angle measure is not exact but within a given range, location is constrained to an area on the terrain. The *area of uncertainty* for a terrestrial navigator is defined to be the area on a plane passing through the true viewpoint in which the navigator may self-locate for any given error range in visual angle measure. Section 3.1 described how areas of uncertainty are determined for any configuration of three landmarks and any given error in visual angle estimate to those landmarks when the viewpoint was outside the configuration. This information will be used to show how an algorithm incorporating a wise choice of landmarks on which to base localization will lead to a significant decrease in the resulting area of uncertainty. The performance of this algorithm will then be compared to indiscriminate choice of landmarks used for localization.

4.2.1 Choosing good configurations

In order to make the best use of available information, the successful navigator must choose landmarks which will give the least localization error regardless of the amount of error in the visual angle measure. As shown in Section 3.1.1, the area of uncertainty corresponding to a given visual angle measure and given error in that visual angle measure varies greatly for different configurations of landmarks. The areas of uncertainty calculated in Chapter 3 were computed knowing the viewpoint location on the map. The navigator will only know landmark map location and landmark viewing order. This information, together with knowing that the viewpoint is located somewhere on the map, can be used to choose configurations of landmarks which will lead to good localization.

Identifying landmarks in large-scale space is difficult and time consuming [Thompson and Pick, 1992]. Rather than depend on a large pool of identified landmarks, one can start with a small "basis" set. Additional landmarks can be identified relative to those in the basis set only if they are needed. It makes no sense to spend time identifying landmarks which will not help in localization.

An algorithm has been developed which can be used to choose good configurations for localization. As discussed in Section 3.1.1, since the computation of the area of uncertainty is dependent on landmark order, the area will always be bounded by the orientation region formed by those three landmarks used for localization. As shown in Figure 4.2, the algorithm begins by picking the triple of landmarks from the basis set which produce the smallest orientation region on the map. An ongoing rule of thumb is to avoid anything close to a single circle configuration (i.e., all landmarks and viewpoint on one circle). Although it is not known exactly where the viewpoint is, any configuration which will result in a circle which passes through all three landmarks also passing near the estimated viewpoint can be avoided. If, as in Figure 4.3, all triples produce the same orientation region, the most widely spaced landmarks should be chosen.

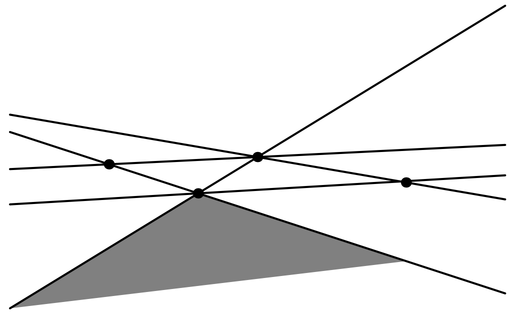

Figure 4.2: Lines joining the landmark points divide space into orientation regions such as the shaded area in the foreground.

The outer area of Figure 4.3 is the area of uncertainty with angles α and β of $30°$ and an error bound of $\pm 10°$ in each. The landmarks are at A, B and C. The inner area is for visual angles of $60°$ with the same error bound. The landmarks are at A′, B and C′. The further apart these landmarks are, the better the localization.

The Algorithm:

Step I: Choose the first set of landmarks as those best satisfying the above qualifications. An estimate of the viewpoint using this set will locate the navigator at point V_0.

Step II: Choose the "best" configuration based on V_0, and estimate a new viewpoint, V_1. In preliminary work [Sutherland and Thompson, 1993], the best configurations in two dissimilar directions were used to estimate viewpoint. The weighted average of these two values was taken as the estimated viewpoint. As explained in Section 4.2.2, results were not as good as when only one configuration was used. If necessary (and possible), identify needed new landmarks.

Incorporating the following constraints, based on the error analysis of Chapter 3, a "goodness" function was developed 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 Figure 4.3.

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 configuration is a straight line.

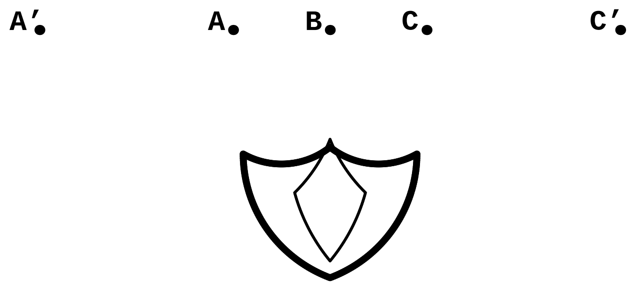

Figure 4.3: The outer lines surround the area of uncertainty with angles α and β of $30°$ and an error bound of $\pm 10°$ in each. The landmarks are at A, B and C. The inner area is for visual angles of $60°$ with the same error bound. The landmarks are at A′, B and C′.

3. The closer a configuration is to single circle (i.e., all landmarks plus the viewpoint on one circle), the greater the error in localization. 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 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.

This heuristic function uses the locations of landmarks A, B, C and the first estimated viewpoint V_0. The larger the function value, the better the configuration. Although V_0 is not necessarily the true viewpoint, experiments have shown that this function discriminates in such a way that the best configuration to be used for localization can be determined using this estimate. In Figure 4.4, 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.

Then:

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

where

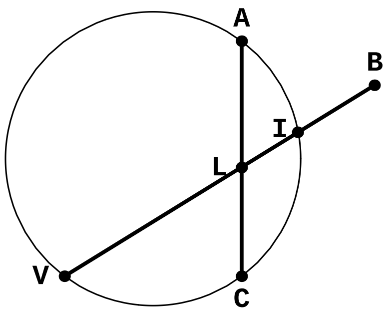

Figure 4.4: Simple geometric relations can be used to rank landmark configurations.

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

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

The 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. An example was given in Figure 3.19 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 the elevation data was in meters, *k* was set to .005 in the experiments described here.

- 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. 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 Figure 4.4 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. The point where that line intersects the AC line is labeled L and the point where it intersects the circle is labeled I.

The line joining V and B is split 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 $d(V, B) - d(V, I)$ gives 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. So, 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 Figure 3.13, 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 which occurred in our trials, the factor $\frac{2}{\pi}$ does a good job of equally rating two configurations which produce similar sized areas of uncertainty where one has B positioned toward V and the other has B outside of the circle.

The function rates the configurations a) through d) in Figure 3.13 as 1, 2.73, 2.15, and 1.74, respectively, for V_0 at actual viewpoint. Varying the position of V_0 in the orientation region will change the actual measures, but, in most cases, not their ordering (i.e., the second configuration remains the best). The important point to be made here is not that a particular heuristic function is superior to any other for choosing configurations, but rather that it is possible to use a simple function to choose landmark triples which will decrease the resulting error in localization regardless of error in visual angle measure.

4.2.2 Why use only one landmark triple?

The question arises as to whether or not better localization could be attained with more than the minimal three landmarks. Levitt and Lawton [Levitt and Lawton, 1990] 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. Figures 4.3 and 4.5 show two examples of situations where the benefit of adding more landmarks is questionable. In Figure 4.3, the innermost area, which is

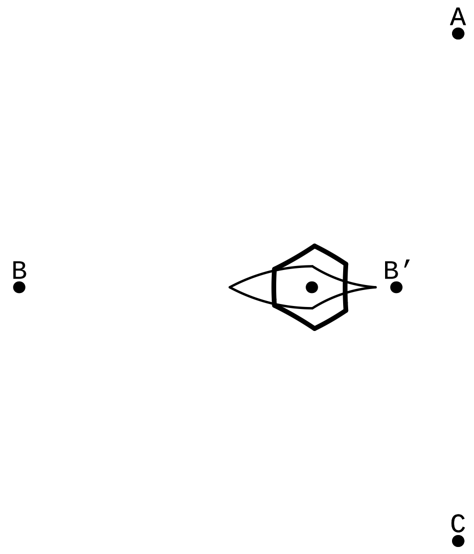

Figure 4.5: It is not necessarily the case that the more landmarks used for localization, the better. The area surrounded by 4 arcs resulted from using landmarks A, B', and C.

completely inside the outer 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 bound was $\pm 30\%$ of the visual angle measure for both configurations. A multiplicative error was used 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 a $\pm 9°$ error bound was used, the smaller area, produced by A', B and C', would decrease. If a $\pm 16.9°$ error bound was used, the larger area, produced by A, B and C, would increase. In this situation, if landmarks A', B and C' had already been used for localization, identifying and adding A and C (or any other features on line segment $A'C'$) would not improve the size of the area of uncertainty.

It is important to keep in mind the fact that the true viewpoint is not known. 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 in Figure 4.3, one area was nested inside another and viewpoints were estimated using both triples, an estimated viewpoint using the configuration which produced the smaller area would have to be somehow combined with the viewpoint estimated using the configuration that produced the larger area. A general rule for combining two estimated viewpoints, (i.e., without knowing the true viewpoint or the size and shape of the areas of uncertainty), will not, in all cases, produce a better result than using only one of the estimates, particularly if that estimate was made using the configuration which produced the smallest area of uncertainty.

In Figure 4.5, 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 A, B', and C. The error bounds were $\pm 12°$ of visual angle measure for both configurations. For this example, an additive error model was used. A multiplicative error would put the $AB'C$ area completely inside the ABC area. 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. 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 actual incurred error is additive, multiplicative or (most likely) a combination of the two, it is clear that some configurations will not improve localization and that when processing sensor feedback in real time and having to possibly search for and identify new landmarks, the assumption of the more landmarks used the better is not necessarily true.

A sequence of experiments was run comparing 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 outcomes for a typical set of runs are shown in Tables 4.1-4.4. Triples of landmarks were picked from a group of five, providing 10 different triples from which to choose. A random error with uniform distribution within the given bounds was added to the angle measures for 100 iterations. The tables show the number of iterations (as a percent) in which the given combination of triples produced an estimated viewpoint closest to the actual one. If more than one triple was used, the estimate was determined by linearly combining the x and y coordinates of each viewpoint. This does not, of course, account for the cost of processing the additional sensor readings. In the run summarized in Table 4.1, 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. Figure 4.6 shows the view and map with landmarks labeled which were used for this set of runs.[2] The navigator is located at the star in the foreground. The landmarks creating the "good"

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

triple are marked by triangles.

Error in Angle	"Good" Triple	2 Triples	5 Triples	10 Triples
10%	71%	20%	1%	8%
20%	72%	22%	2%	4%
30%	67%	21%	1%	11%

Table 4.1: Percentage of 100 iterations for which the given combinations of landmark triples produced an estimated viewpoint closest to the actual viewpoint. In all cases, the 2 Triple and 5 Triple columns included the "Good" Triple.

The results of ordering all triples using the heuristic and combining them in the order: the best, the two best, the five best, etc. are shown in Table 4.2. The results of combining them in the order: the worst, the two worst, the five worst, etc. are shown in Table 4.3. The results of combining them in a random order are shown in Table 4.4.

Error in Angle	"Best" Triple	2 "Best"	5 "Best"	10 Triples
10%	38%	17%	45%	0%
20%	40%	23%	37%	0%
30%	34%	25%	41%	0%

Table 4.2: Percentage of 100 iterations for which the given combinations of the best landmark triples produced an estimated viewpoint closest to the actual viewpoint.

Error in Angle	"Worst" Triple	2 "Worst"	5 "Worst"	10 Triples
10%	8%	1%	0%	91%
20%	10%	7%	0%	83%
30%	22%	4%	0%	74%

Table 4.3: Percentage of 100 iterations for which the given combinations of the worst landmark triples produced an estimated viewpoint closest to the actual viewpoint.

What is clear from these results is that, when landmarks are sparse, if any combinations of landmark triples are made, the triples used should be chosen with care. However, it also should be noted that the heuristic used for ranking the triples is simply that: a heuristic. The "goodness" ratings for the 10 triples are given in Table 4.5. It cannot be assumed that the heuristic is, for example, able to clearly differentiate between configuration ABC with a rating of .61 and configuration ABE with a rating of .58. In addition, if some of the triples available for localization rank very poorly, it would be best not to use them at all. Furthermore, although the combination of the triples with the five best ratings gave results comparable to using the triple with the highest rating alone, using those five triples required the identification of all five landmarks and the making of five estimates. In all of these trials, the combination of triples which resulted in the estimate

Error in Angle	Random Triple	2 Random Triples	5 Random Triples	10 Triples
10%	46%	33%	13%	8%
20%	47%	36%	15%	2%
30%	42%	34%	21%	3%

Table 4.4: Percentage of 100 iterations for which the given combinations of random landmark triples produced an estimated viewpoint closest to the actual viewpoint.

closest to the true viewpoint received the vote. How much closer it was than the second best estimate was not considered. Two interesting extensions of this analysis would be to measure the differences in distances and to consider the variance as well as the mean of each set of estimates.

In summary, exploiting the geometric properties of a few identified landmarks rather than assuming that numerous landmarks have been identified will minimize cost, particularly since the number of significant landmarks in outdoor environments is usually quite small.

4.2.3 Experimental results for static localization

The performance of the algorithm described in Section 4.3.4 was compared with using indiscriminate choice of landmarks for localization. The size of the basis set in all cases was five, providing ten different combinations of ordered landmarks. Experiments were run in simulation using real topographic map data. A uniform distribution of error of $\pm 10°$ in visual angle measure was assumed.

The statistical distribution of sensor errors and whether they affect measurements in an additive, multiplicative, or non-linear manner are heavily dependent on the specific sensor technologies used. In navigation, these can range from very wide field image sensors to sensors which mechanically move to scan for landmarks. An assumption of a uniform distribution of error does not necessarily mean that error is 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. [Hager and Mintz, 1991, McKendall and Mintz, 1990] for indoor use. However, many of the errors a navigator 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, whatever error does occur is likely to be a subset of that which was assumed.

Whether or not assumed error should be multiplicative or additive is also a point of debate. Sources of sensor measurement errors for the visual angle between two landmarks include positional quantization, lens distortion and camera positioning errors. The latter two will almost always dominate, particularly if visual angles are large. The magnitude of these errors is multiplicative with respect to the visual angle. Another important source of error in visual angle measurement results from the fact that landmarks are seldom truly point features, resulting in some uncertainty in the bearing that should be assigned to an individual landmark. This error is independent of

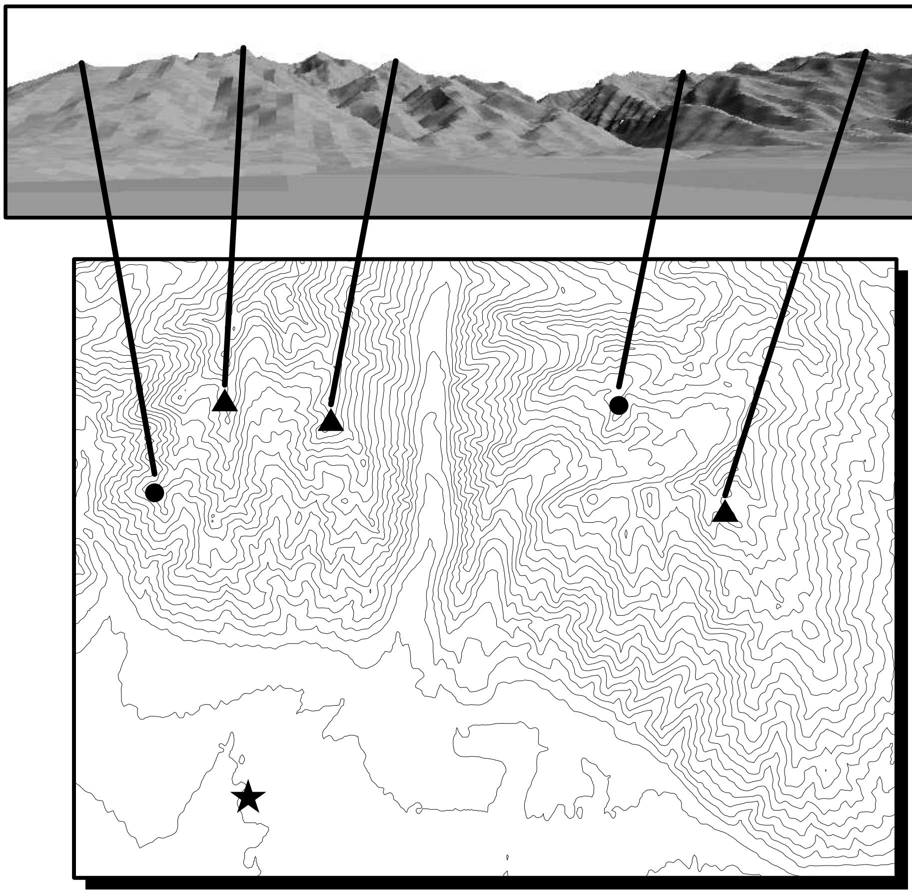

Figure 4.6: The five marked mountain peaks in the view correspond to points on the map. The navigator is at the star in the foreground. The good triple is marked by triangles.

the actual visual angle measure so contributes in a more additive manner to the visual angle error. However, it is not necessarily completely additive. As an example, an error of $20°$ could be made in a large visual angle due to a mismatch of the true "point" for one feature. The assumption that landmark ordering is known precludes this type of error from occurring with visual angles of measurement, in this example, less than $20°$. The error cannot be larger than the angle without changing the order of the landmarks. The assumption of additive or multiplicative error was varied in the simulations run here.

The identified landmarks for each experiment were chosen based on their locations relative to the viewpoint. All landmarks could theoretically be seen from the viewpoint location. Figure 4.7 shows the same map as in Figure 4.1 with the basis set of landmarks identified with large dots and three of the ten LPB's drawn on the map.[3] The landmark labels run across the top of the map directly above the corresponding landmarks.

A navigator located at the star near the left-middle of the map views the landmarks in the order

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

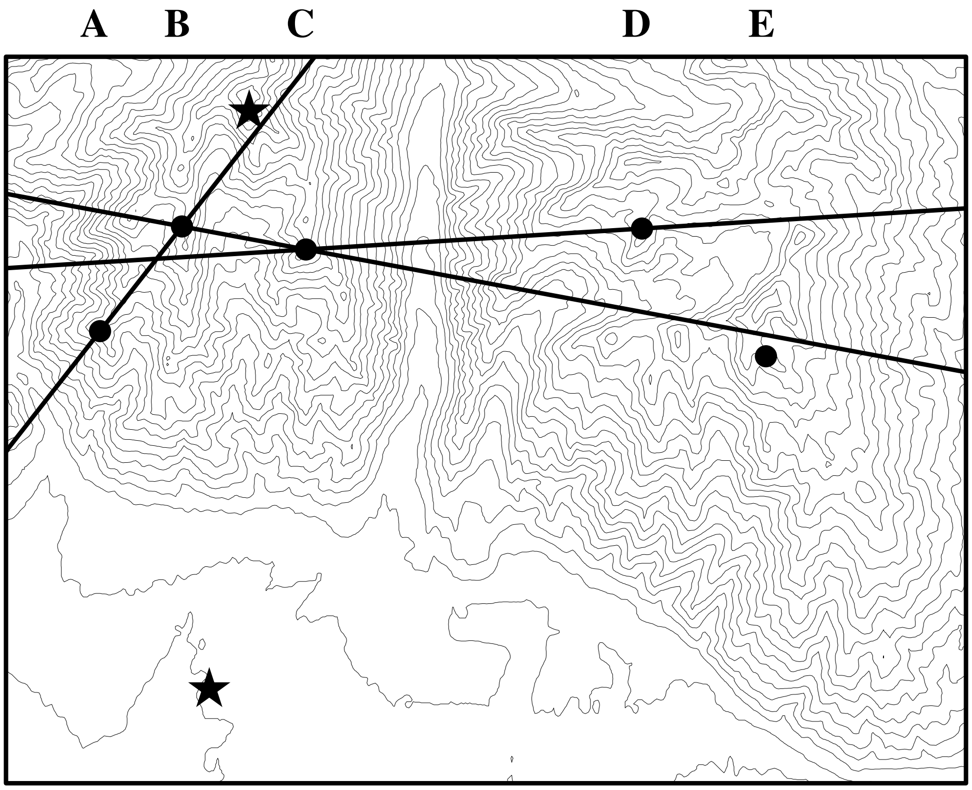

Figure 4.7: The algorithm begins by picking landmarks which produce the smallest orientation region on the map.

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 V_0. 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 in the foreground then estimates viewpoint V_0 using these landmarks. The estimated V_0 is used in the goodness function. Table 4.5 shows the results. The navigator will estimate location using configuration BCE. This is the configuration with the highest rating used for the runs summarized in Table 4.1.

This heuristic function is not guaranteed to always provide the best localization. However, it is easy to implement, intuitive in its method of rating, computationally inexpensive and, in general, does a good job.

Multiple runs of 100 iterations each were made at different locations and with different groups

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

Table 4.5: Based on the above results, the navigator in the foreground of the map in Figure 4.7 will estimate location using configuration BCE.

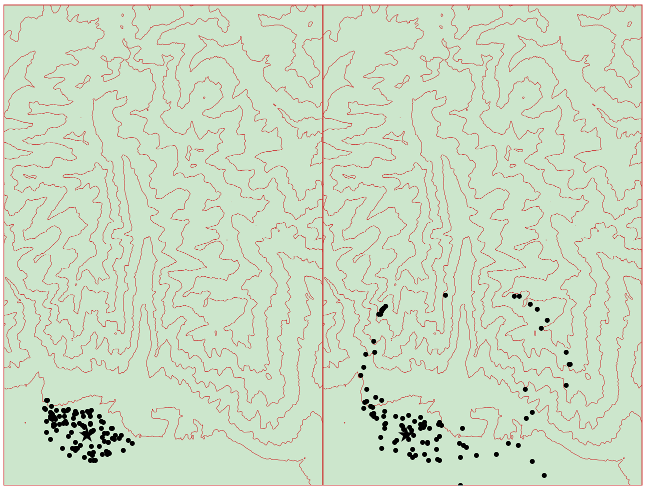

Figure 4.8: Scatter plots of 100 viewpoint estimates for the viewpoint shown in the foreground of Figure 4.7. The plot on the left is the result of localizing using a good landmark configuration. The plot on the right is the result of choosing configurations at random. The map boundaries were 427890E, 4482180N by 452160E, 4498320N UTM coordinates. The maps are shown with east toward the top.

of five landmarks. A random error with a uniform distribution bounded by $\pm 20\%$ of visual angle was added to all angle measures. Figure 4.8 shows scatter plots for the run which was done using the configuration shown in Figure 4.7 with 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 4.6 shows mean distance and standard deviation in meters from actual viewpoint for a sampling of runs done on the map in Figure 4.9, including the run plotted in Figure 4.8. The wise choice of landmarks consistently performed better than a random choice.

Although these errors may seem large, two things should be noted:

- More accurate sensing will result in tighter error bounds with the wise choice continuing to perform better than a random choice of configurations.

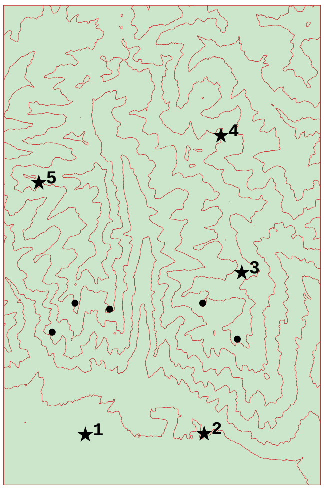

Figure 4.9: The labeled stars are the locations from which the localization runs summarized in Table 4.6 were made.

		Good Landmark Choice		Random Landmark Choice	
Label	Location	Mean Dist	Std Deviation	Mean Dist	Std Deviation
1	430560E,4494240N	1264.89	513.12	3237.74	2726.13
2	430620E,4488210N	1142.44	670.81	3136.97	2520.50
3	438840E,4486290N	402.73	202.35	688.74	478.83
4	445830E,4487340N	1368.04	1136.07	2876.42	3288.25
5	443460E,4496760N	987.45	358.96	1185.76	819.96

Table 4.6: Results of a sampling of localization runs using the map shown in Figure 4.9. The labels refer to the labeled points in Figure 4.9. Each run was of 100 iterations. The map boundaries were 427890E, 4482180N by 452160E, 4498320N UTM coordinates.

- Distances *are* large in large-scale space. The visual angle measure to landmarks is a global constraint. After the viewpoint has been estimated to a general area, local constraints can provide the necessary information for further localization. The smaller that area is, the more successful those local constraints will be.

4.2.4 Experimental results for localization while moving

A moving navigator is faced with a double challenge: not only is there error in visual angle estimate, but also error in movement. As in the static localization, it is assumed that landmarks are point features and can be ordered, that the navigator is traveling on terrain (as opposed to being in space), and that perfect measurement of the visual angles to three landmarks will provide exact localization.

The performance of the algorithm was again compared with randomly choosing landmarks to be used for localization. All experiments were run in simulation using real topographic data. It was assumed that the navigator had a map of the area and knew map locations of points which defined both the path and the landmarks as well as the order of landmarks with respect to initial navigator location. Results for one example are shown in Figure 4.10 and in the sequence of frames in Figure 4.11. 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. All landmarks are mountain peaks which are visible from the given path.[4] As in the static localization experiments, a small set of landmarks was used, with additional landmarks being identified when needed. The eight landmarks used for these trials provided 56 different combinations of ordered landmark triples.

Consider two navigators moving along a path toward a goal. They have identified visible landmarks on a map and know the left to right order of those landmarks. Both begin by using their knowledge of landmark order to determine the smallest orientation region in which they are located. They use the landmarks which form that region to estimate their initial location. Those three landmarks are shown as triangles in Figure 4.10. 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 Figure 4.11 show each step as the navigators progress toward the goal. A configuration of three landmarks to use for localization (triangles) is chosen. The viewpoint (the empty square) is estimated and a move is made toward the next path point (the line ending in a solid square). The sequence on the left shows a wise choice of landmarks. Landmarks are chosen randomly in the sequence on the right.

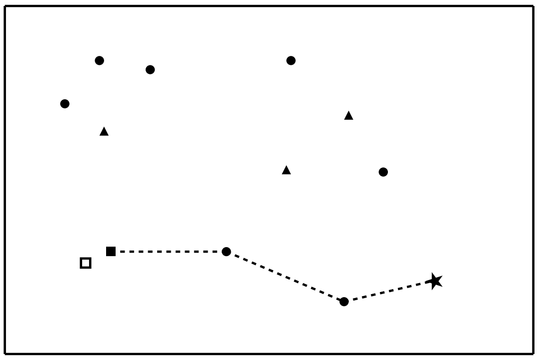

Figure 4.10: The eight points at the top of the figure represent the eight landmarks used for localization. Both navigators start at the solid square on the lower left. The viewpoint is estimated (the empty square) using the three landmarks (triangles) which produce the smallest orientation region. The desired path is shown as a dotted line. The goal is marked by a star.

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)

[4]Landmark locations and elevations were taken from USGS 30m DEM data.

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.[5]

Error Bounds:		$\pm 20\%$ Angle 0 Move	$\pm 30\%$ Angle 0 Move	$\pm 20\%$ Angle $\pm 20\%$ Move
Mean Extra Distance	Wise Landmark Choice	344	2883	474
	Random Landmark Choice	4273	18657	4576
Mean Distance to Path	Wise Landmark Choice	452	513	387
	Random Landmark Choice	1106	1227	861
Mean Distance to Goal	Wise Landmark Choice	711	1166	769
	Random Landmark Choice	3239	4781	3290

Table 4.7: Results after 100 trials. The total path length was 11352 meters. All distances have been rounded to the nearest meter.

In order to illustrate this approach, simulations were run using a multiplicative error, uniformly distributed over a fixed range. The three pairs of frames in Figure 4.12 show navigator positions for 50 trials, assuming a uniform distribution of error within $\pm 20\%$ in visual angle measure and no error in movement, error within $\pm 30\%$ in visual angle measure and no error in movement, and error within $\pm 20\%$ in both visual angle and direction and distance of move.[6] The clustering around the path points is quite marked on the left, the result of using the algorithm to choose landmark configurations.

Table 4.7 gives the results for all three cases after 100 trials each. The distances have been rounded to the nearest meter. "Mean Extra Distance" is the average number of meters $\pm$ total path length that each navigator traveled. Due to the fact that paths in unstructured environments are seldom straight, total distance traveled does not necessarily reflect how well the navigator stayed on the desired path. For that reason, distance of each path segment of the desired path to the corresponding path taken was also recorded. The perpendicular distance of the midpoint of the desired path segment to the 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 which used the 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

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

[6]A point was picked from a uniform distribution within a circle of radius 20% of path segment length around the desired path point.

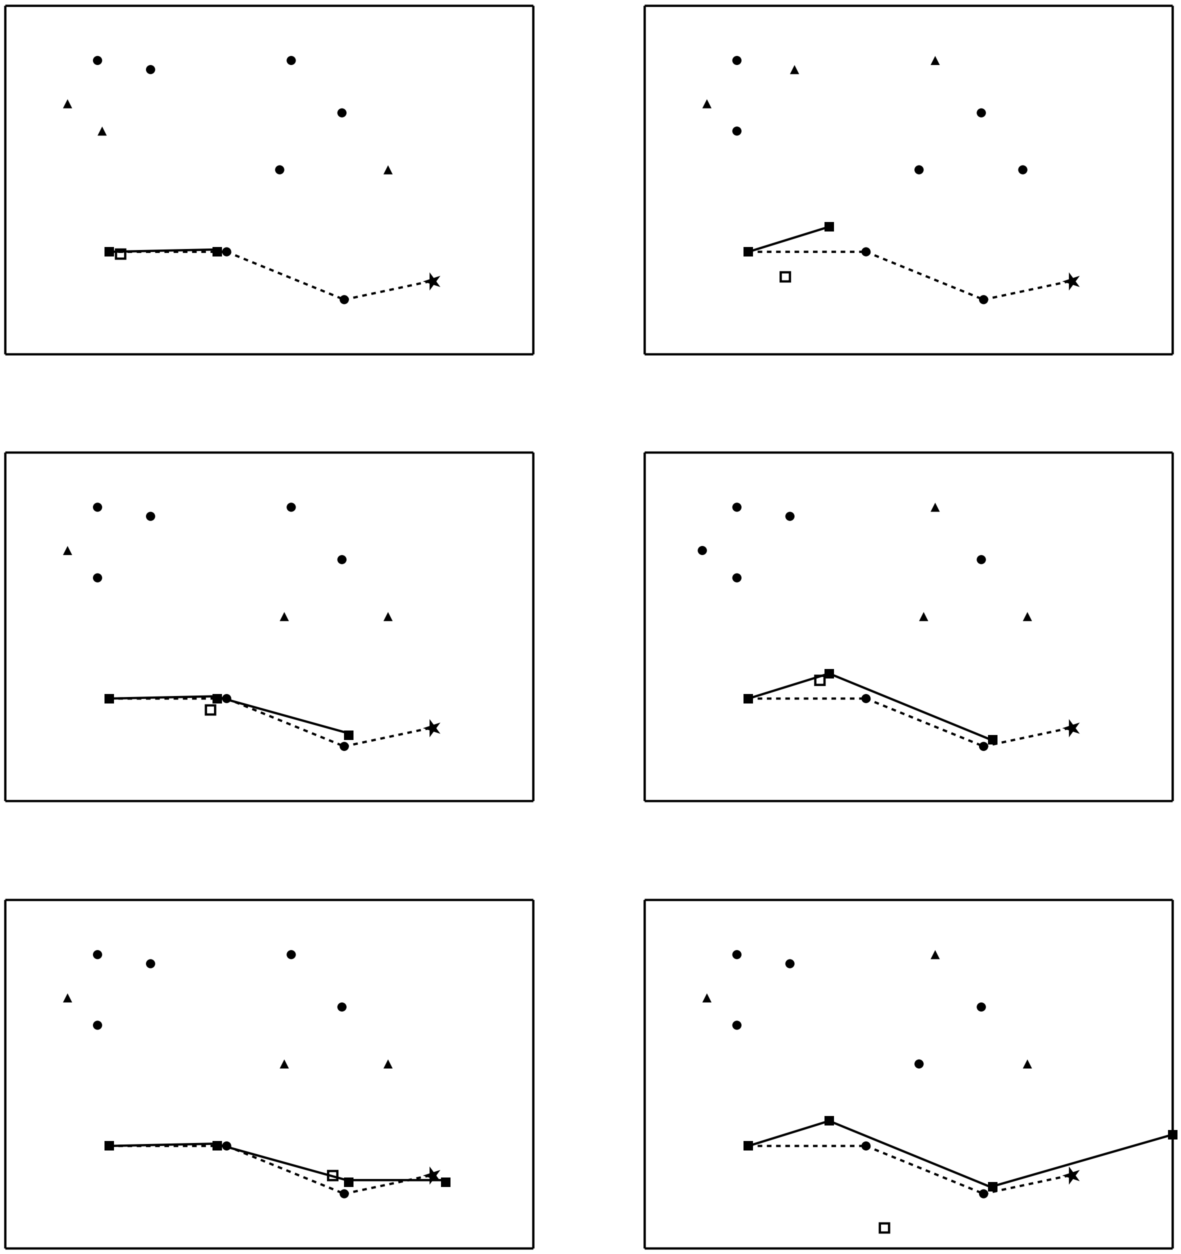

Figure 4.11: The sequence on the left shows the path taken by the navigator using the algorithm. The sequence on the right shows the path taken when landmarks used for localization are chosen randomly. The landmarks used for localization are shown as triangles. The desired path is a dotted line. The path taken is a solid line. The viewpoint is estimated at the empty square, and the navigator moves to the next path point (the end of the solid line furthest to the right).

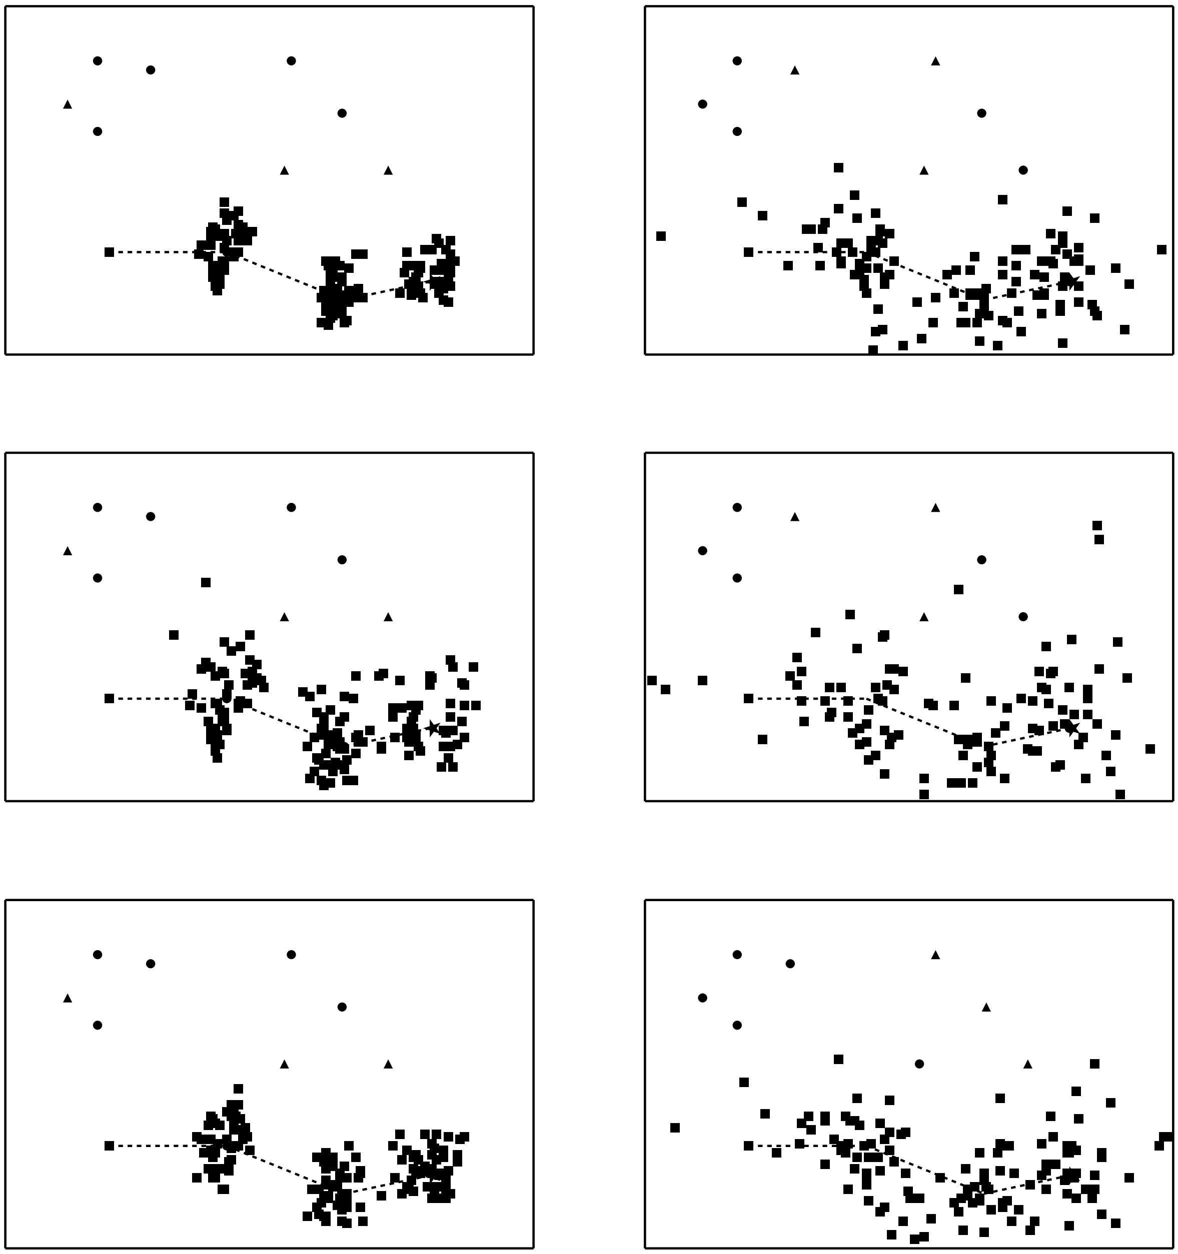

Figure 4.12: After fifty trials, clustering on the left shows how better localization results when landmarks are chosen wisely. The error bounds were $\pm 20\%$ in visual angle for the top pair of frames, $\pm 30\%$ in visual angle for the second pair of frames, and $\pm 20\%$ in both visual angle and direction and distance of move for the third set of frames.

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 is close enough 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.

4.2.5 Adding new landmarks

For the example in Section 4.2.4, all landmarks were initially identified and matched to the map. The landmarks were also all on the navigator's left as it moved. It is often the case that new features to be used for localization must be added as movement continues. This is almost always true when distances are long or movement is through an area such as a canyon with landmarks on both sides of the path. It would be to the navigator's advantage to look for new landmarks which would provide the best localization. This would also be an appropriate strategy to be used in the placement of synthetic landmarks such as beacons. Thus, whether placing beacons to be used for localization or identifying new natural landmarks, the cost can be reduced significantly if locations which will provide good localization are chosen. Turning again 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 [Collett *et al.*, 1992]. The results of the analysis in Chapter 3 can be used to choose areas on the terrain map in which to look for landmarks.

Consider two navigators following the path shown in Figure 4.13.[7] 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 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, 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 the 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 Figure 4.13, the navigator on the left traveled less distance, remained closer to the path and ended closer to the next path point. It also

[7] They are using the same map as in Figure 4.7 and are moving through the canyon which runs up the center of the map.

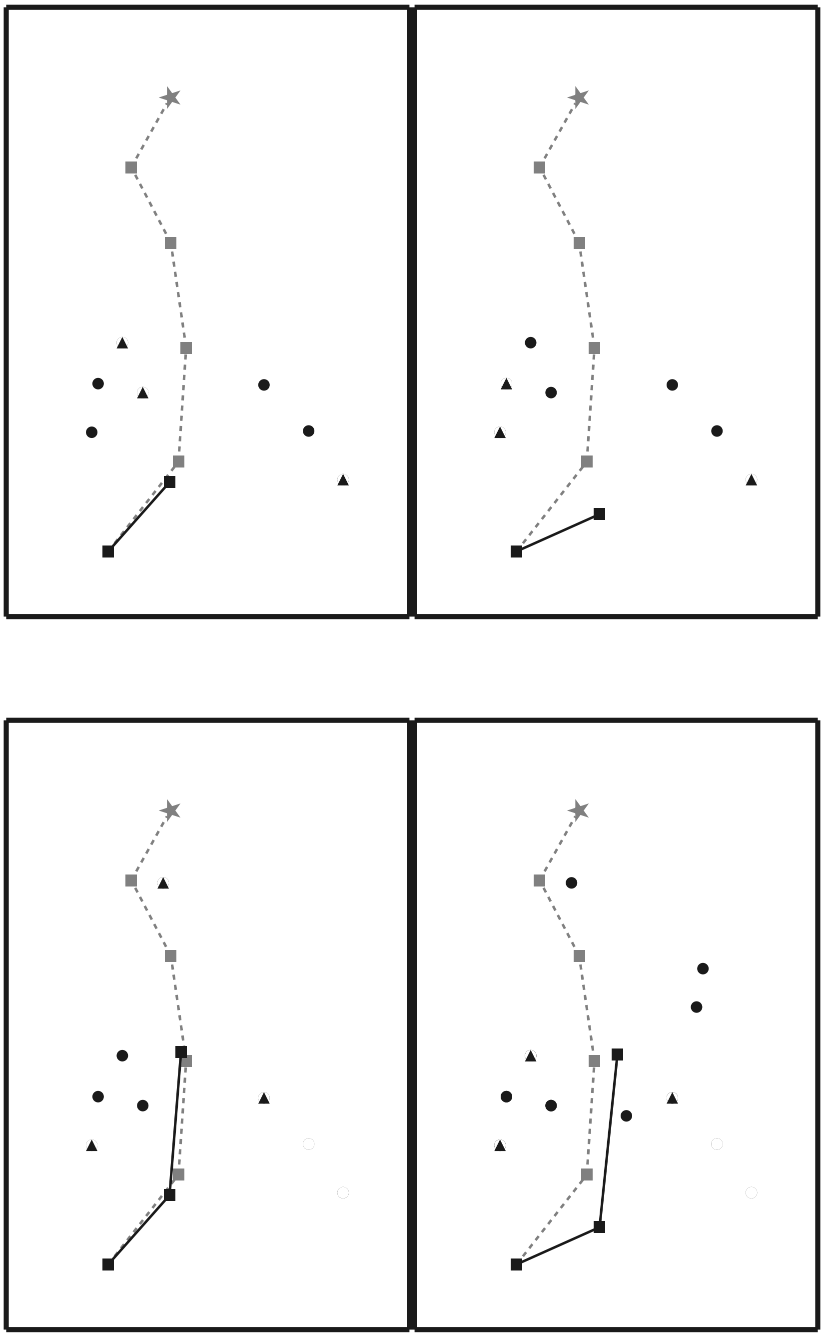

Figure 4.13: An example of two navigators traveling toward a goal. They begin by picking 3 landmarks with which to localize, shown by the triangles in the top set of frames, and move along the dark line to the second dark box. They then take a new image. The navigator on the left looks for an additional 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.

spent minimal time identifying new landmarks because it knew where the best place was to look for them.

Figures 4.14-4.16 show the results after 50 trials for the same multiplicative error amounts used in the example shown in Figure 4.12, a uniform distribution of error within $\pm 20\%$ in visual angle measure and no error in movement, error within $\pm 30\%$ in visual angle measure and no error in movement, and error within $\pm 20\%$ in both visual angle and direction and distance of move. Note that no other heuristics, such as knowledge of orientation region, were used. Thus, outliers are present on both the left and the right. The visual angles were, in general, larger than those in Figure 4.12, producing very large absolute errors for the given multiplicative error amounts. In spite of that, better clustering is apparent on the left, where the heuristic was used. Table 4.8 gives the results for all three cases after 100 trials each. The distances have been rounded to the nearest meter. The table is labeled in exactly the same way as Table 4.7, which gave the results for the example in Section 4.2.4.

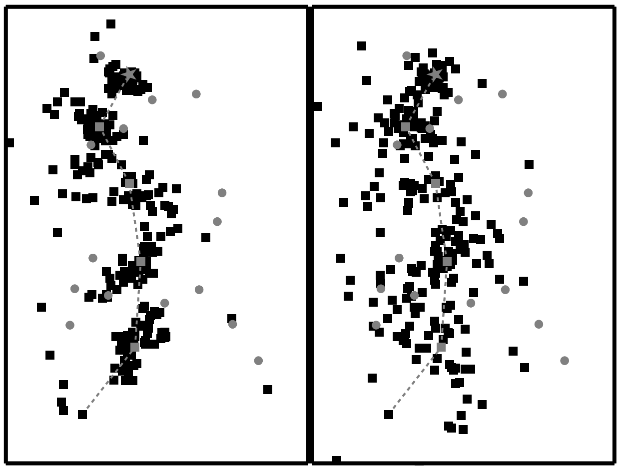

Figure 4.14: Results after 50 trials with error bounds of $\pm 20\%$ in visual angle measure. The heuristic is used for choosing landmarks in the frame on the left.

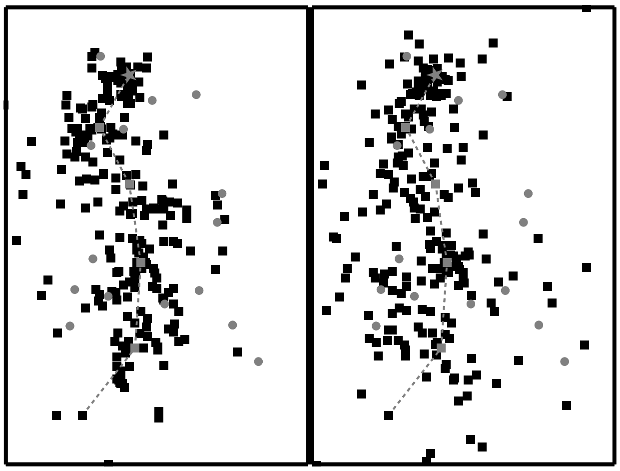

Figure 4.15: Results after 50 trials with error bounds of $\pm 30\%$ in visual angle measure. The heuristic is used for choosing landmarks in the frame on the left.

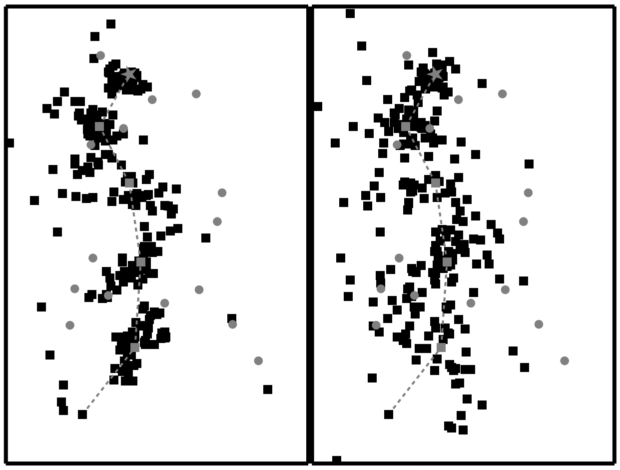

Figure 4.16: Results after 50 trials with error bounds of $\pm 20\%$ in both visual angle measure and the direction and distance of the move. The heuristic is used for choosing landmarks in the frame on the left.

Error Bounds:		$\pm 20\%$ Angle 0 Move	$\pm 30\%$ Angle 0 Move	$\pm 20\%$ Angle $\pm 20\%$ Move
Mean Extra Distance	Wise Landmark Choice	2969	4522	4764
	Random Landmark Choice	5430	8070	5379
Mean Distance to Path	Wise Landmark Choice	687	964	710
	Random Landmark Choice	871	1302	998
Mean Distance to Goal	Wise Landmark Choice	761	1438	900
	Random Landmark Choice	1177	1838	1215

Table 4.8: Results after 100 trials with the navigator traveling through the area containing the landmarks. The total path length was 19490 meters. All distances have been rounded to the nearest meter.

4.3 Localizing to a path

4.3.1 Pursuing projections

Section 4.2.1 described how heuristics can be used to choose landmarks for localization which reduce the error in location estimate. 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.

Techniques frequently applied to statistical pattern recognition problems can be used to develop heuristics for choosing landmark configurations which will produce good path localization. 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 [Diaconis and Freedman, 1984, Friedman, 1987, Intrator, 1992, Intrator, 1993]. 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 [Diaconis and Freedman, 1984] 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 will be pursued herein.

The error in location estimate can be divided into two components: error in lateral distance from the path and error along the path. As shown in Figure 4.17, 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

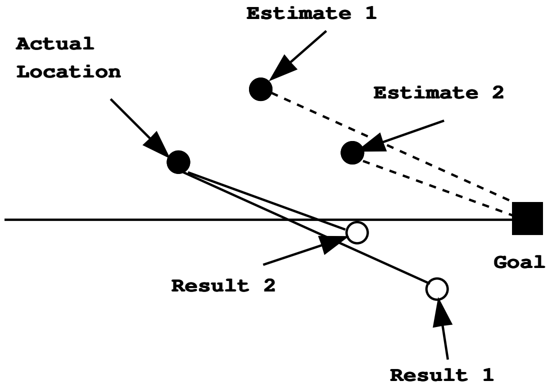

Figure 4.17: When the navigator is trying to follow the path toward the goal, Estimate 1 is closer to actual location but Estimate 2 is closer to the path. The dashed lines show computed paths from the two estimated locations. The solid lines show actual paths. The resulting locations are shown by empty circles. If the goal is to stay as close to the path as possible, Estimate 2 is preferable to Estimate 1.

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.

Consider the area of uncertainty shown in Figure 3.21. Assuming uniform distribution of error in 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 the two-dimensional distribution, the distribution of points along both the x and y axes, as shown in the graphs at the top of Figure 4.18 show 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 at the bottom of Figure 4.18, 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.

A second example in Figure 4.19 shows an elongated area of uncertainty. The landmarks used for localization are at A, B, and C. The error bound was $\pm 10\%$ of the 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.

The goal: to find an area of uncertainty which has a Gaussian distribution with small variance

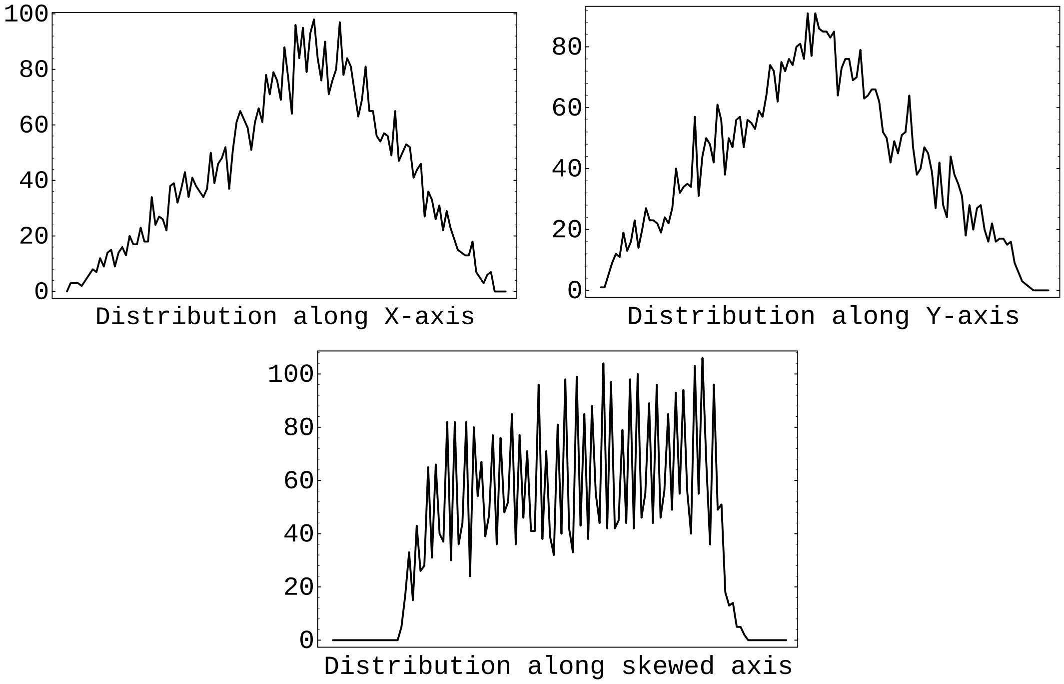

Figure 4.18: Distribution of points along x-axis, y-axis and along an axis rotated $45°$ counterclockwise from the x-axis.

when projected onto an axis perpendicular to the 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 the landmark configuration and the path location.
- The ratio W/L where L = the diameter[8] 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 4.19 is "thin". In most cases, the ratio W/L can be estimated with knowledge of landmark configuration, the path location and the orientation region in which the viewpoint lies. This ratio was used in Chapter 3 to measure the eccentricity of an area of uncertainty.

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

[8]"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.

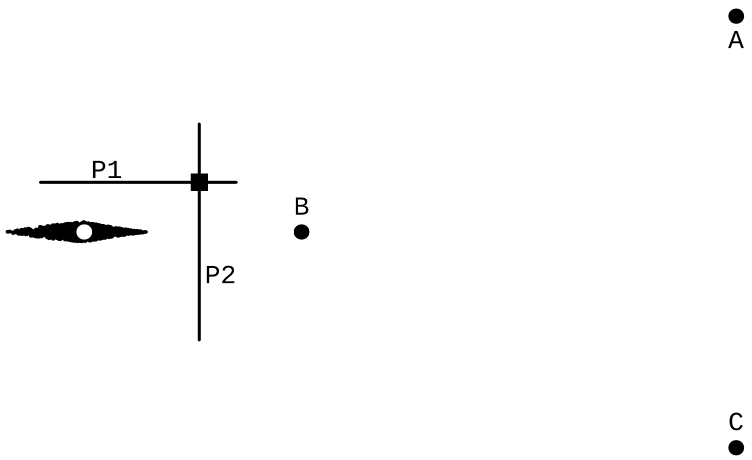

Figure 4.19: The landmarks used for localization are at A, B, and C. The error bound was $\pm 10\%$ of the visual angle measure. The actual viewpoint is at the white dot surrounded by the area of uncertainty shown by the black scatter plot of possible points. When the navigator moves to the black square, final location will be within an area matching the area of uncertainty centered at the square.

4.3.2 Effect of path-axis angle

To show how the angle between the path and the 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 4.20. In each set of simulations, trials of 1000 runs each were done. The visual angles to all configurations were the same. The location was estimated and a move was made toward the goal.

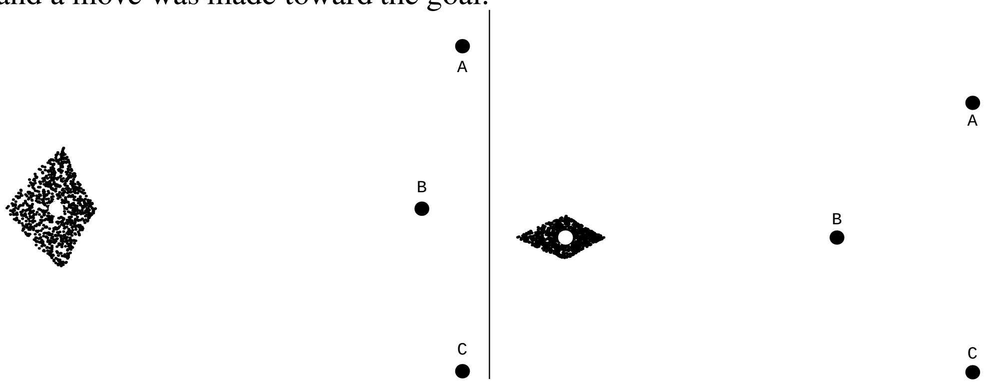

Figure 4.20: Configurations used to show how orientation to path affects path localization. Landmarks are at A, B and C. Viewpoint is at the white dot. Three orientations are shown for the left configuration in Figure 4.21 and for the right configuration in Figure 4.22.

The only parameter which varied was the angular relationship of the configuration to the path direction. A uniform error bound of 10%, 20% and 30% of the visual angle measure was introduced.

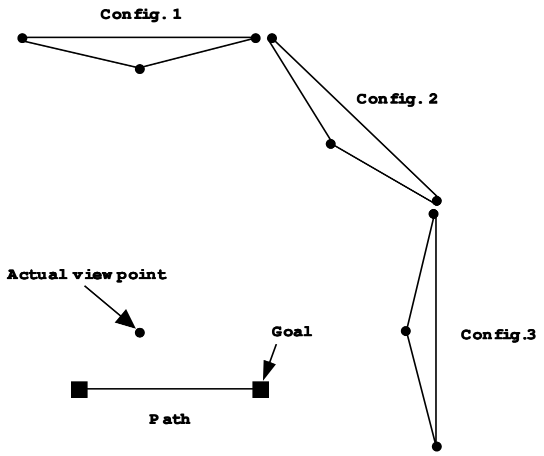

Figure 4.21: The center landmark was 225 linear units and the line joining the outer landmarks was 250 linear units from the actual viewpoint for all three configurations. The angle from the viewpoint to all outer landmark pairs was $43.6°$. As shown in Table 4.9, path localization is best for Configuration 1.

The results for the first set of simulations are shown in Table 4.9. The positions of the configurations are shown in Figure 4.21. Although, as expected, the mean distance to the 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 4.10. The configuration positions are shown in Figure 4.22. In this case, Configuration 3 provides better path localization than does 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 4.9: Results of simulated trials using the three configurations shown in Figure 4.21. The distances are given in linear units.

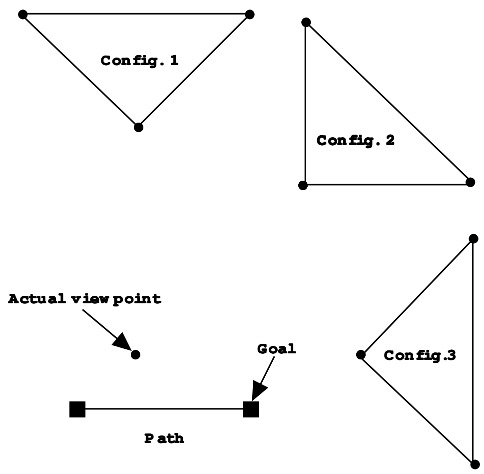

Figure 4.22: The center landmark was 200 linear units and the line joining the outer landmarks was 300 linear units from the actual viewpoint for all three configurations. The angle from the viewpoint to all outer landmark pairs was $36.6°$. As shown in Table 4.10, path localization is best for Configuration 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 4.10: Results of simulated trials using the three configurations shown in Figure 4.22. The distances are given in linear units.

4.3.3 Effect of W/L ratio

Changing the angle between the path and the axis of the orientation region changes the path localization but how it is changed depends on the shape of the area of uncertainty, which can be estimated by the eccentricity ratio W/L, which was described in Section 3.1.2. 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 [Pedoe, 1970] 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. The landmark pair boundaries, as described in Section 3.1, joining AB and BC (and intersecting at B) will be used 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 4.23a, 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.

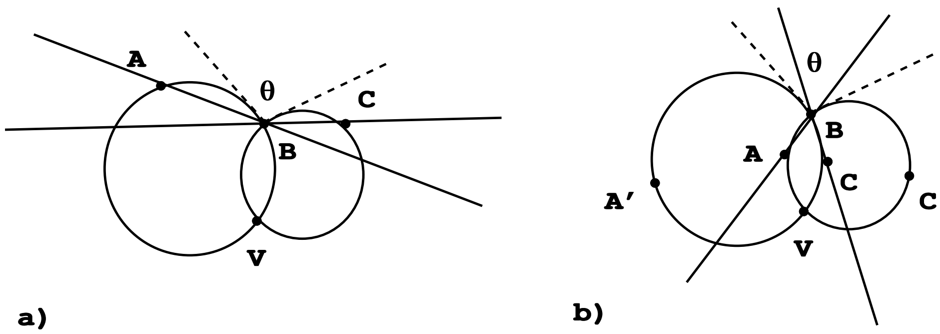

Figure 4.23: In a), the angle θ of intersection of the circles cannot be greater than the angle of intersection of the LPB's. In b), the angle θ of intersection of the circles cannot be less than the angle of intersection of the LPB's.

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 4.23b, but they could also lie on the outer circular arcs, as do A' and C' in Figure 4.23b. 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 4.24. 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). Since it is assumed that anything close to a single circle configuration has already been eliminated, it follows that the estimate of the visual angle to A and C will either be significantly less 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 4.24, 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.

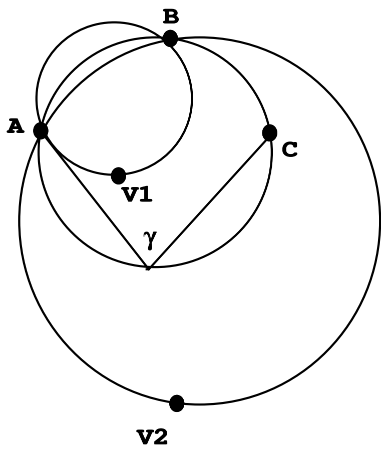

Figure 4.24: Visual angle to A and C is larger than γ at V_1 and smaller than γ at V_2.

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 error (See Section 3.1). As shown in Figure 4.25, 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 4.26 shows the shapes of different areas of uncertainty within given orientation regions, including those formed by the configurations in Figure 4.20. Although the areas change shape somewhat according to the viewpoint location, the general shape is the same within each area.

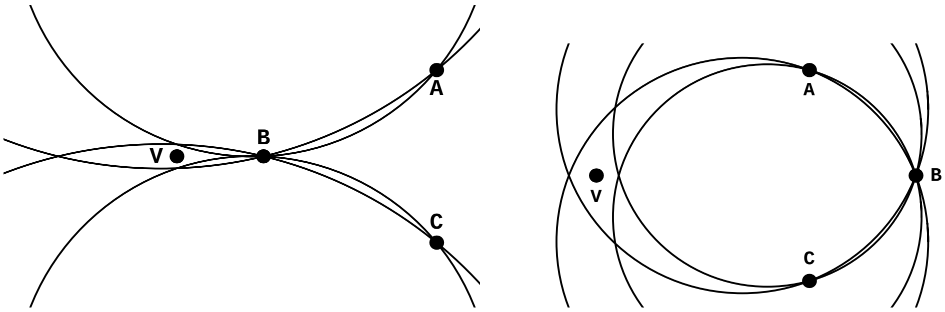

Figure 4.25: The angle of intersection of the circles affects the shape of the area of uncertainty.

To summarize, it can be determined, 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.

4.3.4 Choosing good configurations

The function described in Section 4.2.1 was first used to rank configurations for general goodness, based on the size of the area of uncertainty.

The original function was then augmented 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:

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

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.

The first factor of the upper bound function, *k*, is a weighting factor which was set to 1 for the experiments described here. The third factor, $\frac{1}{1-\frac{1}{\frac{\pi^2}{2}+1}}$, is the multiplicative inverse of the second factor when it reaches its maximum, i.e., when either angle α or angle β are 0. The sole purpose of this factor is to force the range of the function to be $[0, 1]$. The second factor, which is responsible

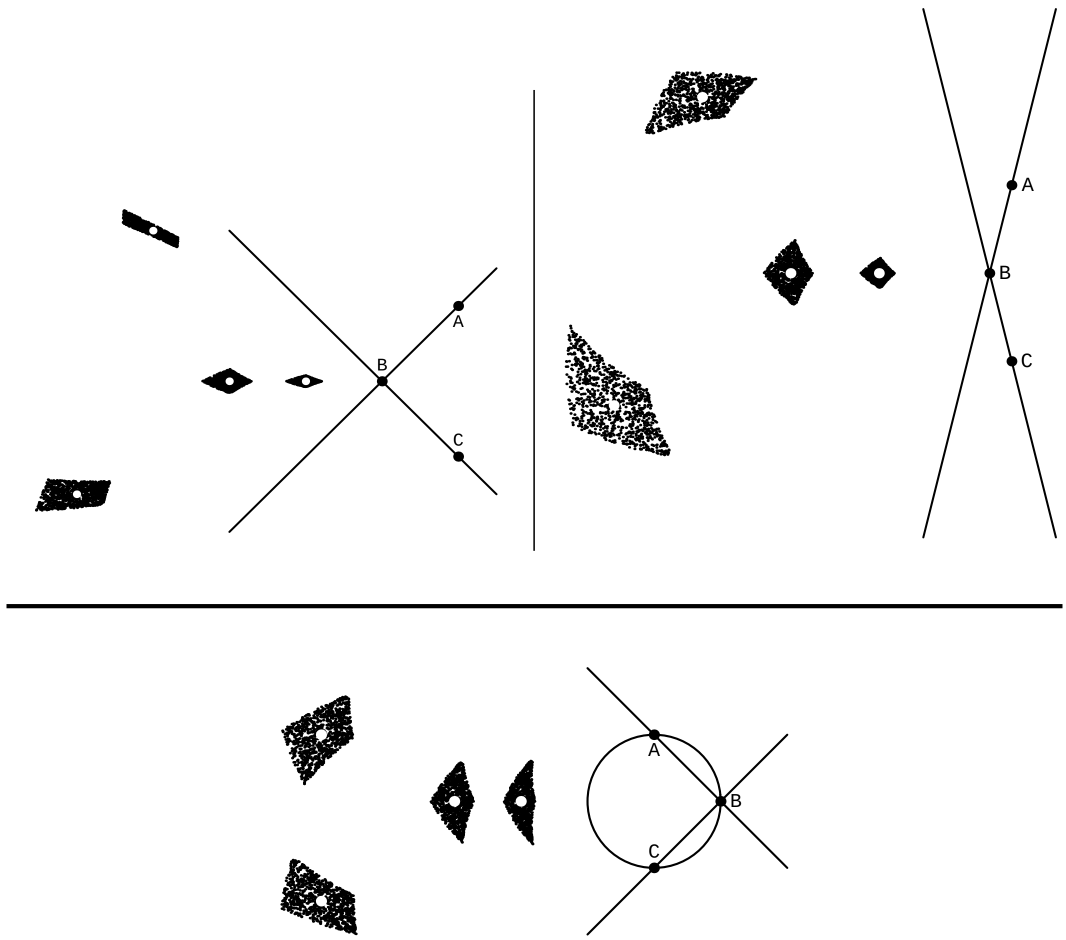

Figure 4.26: In these three examples, four different viewpoints are chosen with each configuration. The LPB's provide an upper limit on the angle of intersection of the circles in both top examples. They provide a lower limit in the example at the bottom, where the circle through A, B and C has been added to show how the four viewpoints lie outside the circle.

for the rating, equals 0 when angle $\beta = \frac{\pi}{2}$ and angle $\alpha = \pi$. In such a case, the upper bound exists, but it is large. In addition, the path is perpendicular to the axis of the orientation region. When such conditions hold, this factor should not, and does not, add to the goodness measure of the configuration. On the other hand, when angle α is small, providing a small upper bound, and/or angle β is small, the function should affect the goodness measure of the configuration. A three-dimensional plot of this function is shown in Figure 4.27.

The rationale for the lower bound function is that a large lower bound is best for path localization when the path is perpendicular to the axis of the orientation region. When angle $\beta = 0$, $p = 0$. As angle β approaches $\frac{\pi}{2}$, p approaches $k * \frac{\alpha}{\pi}$. At the same time, angle α is ranging from a theoretically infinitesimal lower bound when landmark B is pushed back an infinite distance from the viewpoint to π for a straight line configuration of landmarks. This function also ranges in value from 0 to 1,

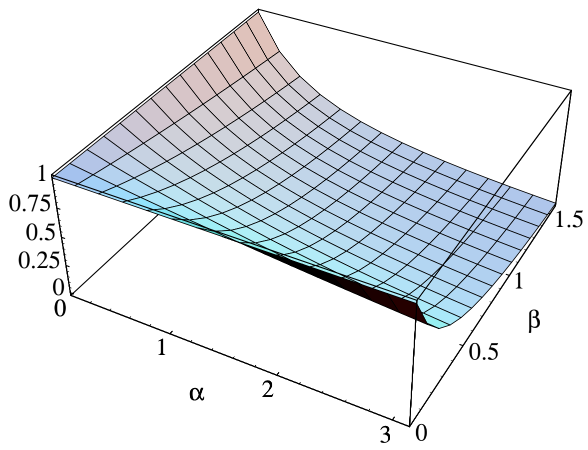

Figure 4.27: When an upper bound exists and the weighting factor $k = 0$, the p function ranges in value from 0 to 1, depending on the measure of α, the angle between landmarks A and C with vertex B and β, the angle between the axis of the orientation region and the path.

with 0 being poor for path localization and 1 being good. The weighting factor k was set to 1, as for the upper bound function. In either case, if path localization is not desired, $k = 0$.

4.3.5 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 4.28 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. The UTM coordinates of the southwest corner are 426180E, 4511040N. The path, shown on the map, runs along the banks of the creek through City Creek Canyon. The start position and each point at which a new image was taken are shown by squares. The goal is marked by a star. The landmarks are marked by filled black circles. The panorama in Figure 4.29 is the terrain viewed from the nearest high point to the robot start position. The canyon is visible in the center frame.

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 the path to the 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

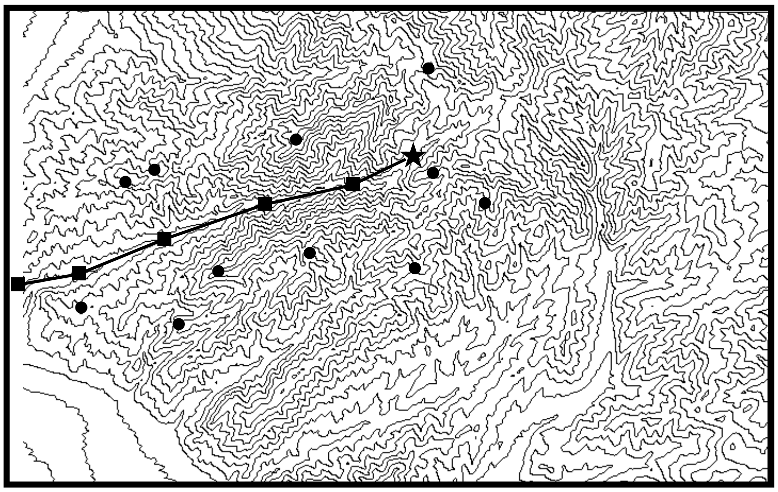

Figure 4.28: A contour map of the area in which simulations were run. The path is shown by a heavy black line. The start position is at the left. The goal is at the star.

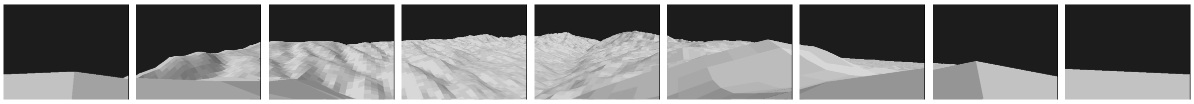

Figure 4.29: Terrain view from nearest high point to start position. The entrance to City Creek Canyon is visible in the center frame.

configurations which had a goodness measure greater than 1 were considered for the additional path weighting. Although this threshold worked well in the 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 of the simulations, its value could also be chosen based on the range of goodness measures for point localization. The results are in Table 4.11.

Error in	Heuristic for Choosing Landmark Triple		
Angle	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 4.11: Results of runs through City Creek Canyon. The mean distance to the path (in meters) is given for three different bounds in angular measure error. Fifty trips were recorded with each error amount.

Adding the heuristic for path localization did not change total distance traveled. The final distance from the goal averaged 10% better when the heuristic was used. As the error bounds on the angular measure increase, the areas of uncertainty increase in size, resulting in the possible estimates being further from the true location and, thus, further from the path. When errors are large, areas of uncertainty are not long and narrow, so few areas are good for path localization. For that reason, the increased error not only causes an increase in the mean distance from the path but also less improvement from using a path heuristic.

In summary, an inversion of the classic statistical projection pursuit problem shows 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. Experiments using USGS 30m 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.

Chapter 5

Conclusion

5.1 Contributions

The results of this work have shown that:

- The amount of the error occurring in a localization process using angular measurements to features depends heavily on which features are used. An example of how different features with identical angular measurements to those features and identical error bounds on visual angle measure can produce significantly different areas of uncertainty was shown in Figure 3.13.

- The amount of the error occurring in such a localization process is not a function of the number of features used. The key point here is that the viewpoint is not known, so the area of uncertainty is not known. The problem is not one of intersecting areas of uncertainty, but of averaging multiple viewpoint estimates. Unless numerous estimates are made, averaging can worsen the final estimate. This problem was discussed in Section 4.2.2.

- It is possible to develop simple heuristic functions for choosing features for localization which will significantly decrease error in that localization. This applies not only to choosing features in the view which have been previously matched to a map, but to choosing directions in which to look for new features. One such function was developed and tested in simulation on real terrain data, using mountain peaks as features. These experiments in localizing to a point were described in Section 4.2.

- It is possible to decrease localization error in a particular direction. A landmark triple which produces a distribution with a small variance when projected into one-dimensional space in a particular direction is "pursued". For a robot attempting to stay close to a path, the most favorable direction for such a projection would be perpendicular to the direction of movement. Experiments in localizing to a path were discussed in Section 4.3.

- If features have been identified but knowledge of left to right order in the view is unknown, simple steps can be taken to aid in determining that ordering. Trials were run introducing varying amounts of error and different feature configurations. This work was summarized in Section 3.6.

5.2 Future work

The three main areas summarized below hold possibilities for future work. The first, outlined in Section 5.2.1, is a continuation of the navigation work described in this thesis. The second is an expansion of the preliminary analysis which has already been done in the area of object recognition. This analysis is described in Section 5.2.2. The third, outlined in Section 5.2.3, is a natural extension of these same ideas in the area of camera calibration.

5.2.1 Navigation

In the area of navigation, there are five immediate extensions of this work which should prove interesting:

1. Experiments should be run with a robot navigating in an actual unstructured outdoor environment both on terrain and underwater.

2. Experiments should be run in an indoor environment which is structured but large. Robot runs were made indoors, but were done so in such a confined space that localization error was not a major problem. The time to rate and choose landmark triples is not well spent if all triples produce good localization.

3. The symmetry properties described in Section 3.1.2 have not been exploited. It would be very interesting to explore methods of using the symmetry properties of an area of uncertainty to aid in predicting location.

4. The navigating skills of cataglyphis bicolor should be analyzed carefully with a focus on discovering which of them would be well suited for use by an autonomous robot.

5. The extraordinary navigation skills of the Micronesians were described in Chapter 2. Several land-based ancient cultures are also known for their ability to navigate without sophisticated instruments. It might be just as interesting to investigate these navigational techniques as was the work of [Pick *et al.*, 1993] with modern, experienced land navigators.

5.2.2 Object recognition

Several extensions to what has been done here exist in the area of pose estimation in model–based object recognition. In most cases, the goal in object recognition is not necessarily to reconstruct a 3D scene, but to identify objects in it. The image is searched for matches to models in memory in an attempt to identify those objects. After image features (e.g., points, edges) have been found, combinations of the features are matched against transformations of the stored models. A common domain used in model–based object recognition is that of two–dimensional rigid objects in a three–dimensional scene. The features to match are points and a bounded amount of sensing error is assumed. Three non–collinear model points are transformed, usually by orthographic projection plus scale (equivalently, by an affine transform [Huttenlocher, 1988]) then matched against triples of non–collinear image points to determine a model pose. This pose is then used to find the location of additional model points in the image [Fischler and Bolles, 1981, Huttenlocher and Ullman, 1990]. Lamdan, Schwartz and Wolfson have shown that a set of three noncollinear model points m_1, m_2, m_3 can be used as an affine basis such that all other model points m_i are expressed as:

$$m_i = \zeta_i(m_2 - m_1) + \eta_i(m_3 - m_1) + m_1$$

where (ζ_i, η_i) are the *affine coordinates* of m_i. The affine coordinates are invariant under affine transforms but they are dependent on the three points which were chosen for the basis [Lamdan *et al.*, 1988].

As in visual navigation, there is sensing error in the presumed location of image points. Jacobs [Jacobs, 1991] assumed a sensing error in image point locations up to ϵ in any direction, such that the correct model pose would align each basis point within a circle of radius ϵ around the corresponding image point. After using model points m_1, m_2, m_3 and image points i_1, i_2, i_3 to determine an affine transform, he applied that transform to a fourth model point m_4. He then showed that the potential location of the fourth image point, i_4, will lie within a circle of radius $(|1 - \zeta_4 - \eta_4| + |\zeta_4| + |\eta_4| + 1)\epsilon$ around the transformed m_4. The size of this circle does not depend on the transformations which have been made, but only on the affine coordinates of m_4.

However, just as the configuration of three landmarks can have a dramatic affect on the size of the area of uncertainty, the three points chosen for the basis will affect the size of the *circle of uncertainty*. Since the affine coordinates of a point are invariant under affine transforms and circle radius depends only on the affine coordinates, a wise choice of model basis points can significantly limit the size of the area in which to search for a match. Limiting the size of this area has two advantages: less time is required to search a smaller area and it is less likely that an incorrect match will be made.

As an example, consider the object in Figure 5.1. Let $m_1 = (-2, 6)$, $m_2 = (0, 4)$, $m_3 = (2, 6)$, $m_2' = (-1, 2)$ and $m_4 = (0, 0)$. Without loss of generality, we can assume that no transformation was required to line up the model with the image. If m_1, m_2, m_3 are used as a basis, then

$$m_4 = (0, 0) = \zeta_4(2, -2) + \eta_4(4, 0) + (-2, 6)$$

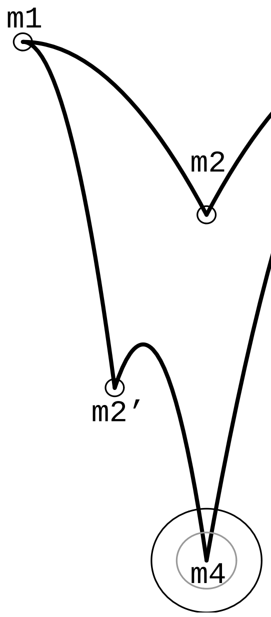

Figure 5.1: Assuming that sensing error is within a circle of radius ϵ within each basis point, the dark circle around point m_4 surrounds the search area when m_1, m_2, m_3 are used as a basis. The light circle surrounds the area when m_2' replaces m_2. The area is significantly reduced in the second case.

$$\Rightarrow \zeta_4 = 3, \eta_4 = -1$$
$$\Rightarrow \text{radius} = (|1 - 3 + 1| + 3 + 1 + 1)\epsilon = 6\epsilon$$

If m_1, m_2', m_3 are used as a basis, then

$$m_4 = (0,0) = \zeta_4(1,-4) + \eta_4(4,0) + (-2,6)$$
$$\Rightarrow \zeta_4 = 1.5, \eta_4 = .125$$
$$\Rightarrow \text{radius} = (|1 - 1.5 - .125| + 1.5 + .125 + 1)\epsilon = 3.25\epsilon$$

Thus, the size of the *circle of uncertainty* can be significantly affected by the choice of basis points.

It is interesting to note that the size of this circle can easily be determined based on the area in the affine plane in which the affine coordinates of the point lie. Consider the numbered areas in the plane of Figure 5.2. The radius of the circle for each of the areas is:

1. 2ϵ
2. $2\epsilon(1-\xi)$
3. $2\epsilon\eta$
4. $2\epsilon(\xi+\eta)$
5. $2\epsilon\xi$
6. $2\epsilon(1-\eta)$
7. $2\epsilon(1-\xi-\eta)$

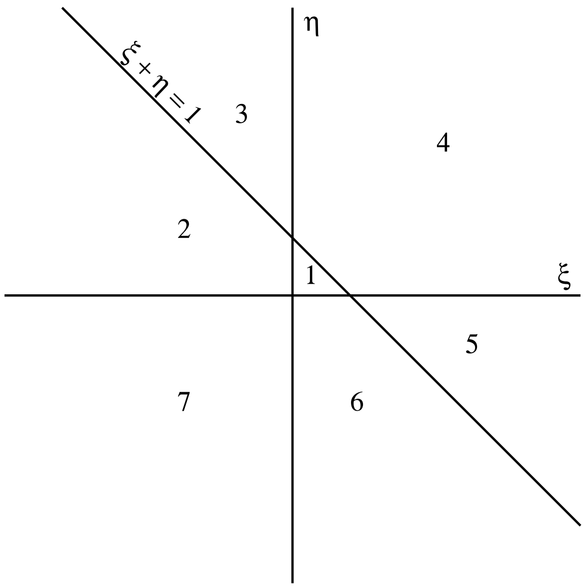

Figure 5.2: The size of the *circle of uncertainty* is determined by the area in the affine plane in which the transformed point lies.

These values can be easily computed from the previously stated general formula [Grimson *et al.*, 1991]. What is interesting about these values is that areas 4 and 7 are the only ones in which the size of the circle depends on both affine coordinates of the point. For example, the circle would be the same size for any point in area 5 for a constant value of ξ. Likewise, any point within area 1 will have a circle of uncertainty of radius 2ϵ.

Another interesting question is: what would be the result if the assumed error in each basis point was *not* a circle of radius ϵ, but an oval or a crescent or some other shape? Would it, for example, be possible to choose basis points which would reduce the search area for an edge to a long narrow strip?

5.2.3 Camera calibration

Camera calibration in stereo vision provides many of the same problems as navigation and object recognition. If the camera is calibrated by using 6 known points and their respective images on the two image planes, the question of how the location of those points affects the error in calibration parallels that of selecting basis points in pose estimation. An error in estimate of an extrinsic camera parameter such as pan or tilt will give a result similar to that of a visual angle error in navigation. [Blostein and Huang, 1987] have derived the probability that position estimate due to quantization error in triangulation is within a specified error tolerance. They assume uniform distribution of the possible locations of the 3D point and analyze the effect on the pixel area. The parallel of this problem to what this research addressed can be seen by comparing Figure 5.3 and Table 5.1 with Figure 3.20 and Table 3.1 in Chapter 3.

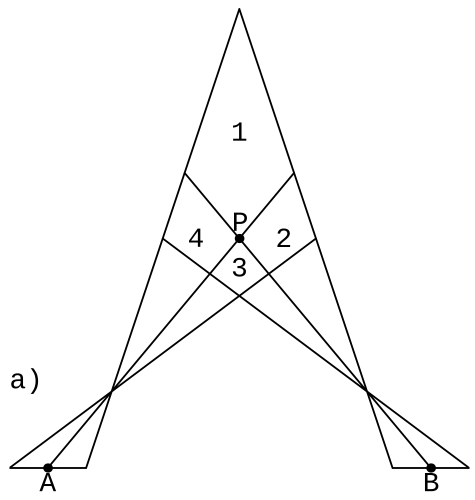

Figure 5.3: Uncertainty due to quantization: A and B are the midpoints of the pixels on which the world point is projected. Actual point projection can be anywhere within the intervals. If A and B were exact, they would project from point P.

Area	Overestimate	Underestimate
1	B	A
2	A, B	—
3	A	B
4	—	A, B

Table 5.1: Given that the world point is in one of the 4 areas shown in Figure 5.3, whether A and/or B are under or overestimates of the actual projected image point is shown.

An analysis of the geometric constraints involved in the use of a third camera should be done. [Dhond and Aggarwal, 1991] in their cost-benefit analysis of adding a third camera to the stereo process, mention that the pan and tilt angles of the second and third cameras. respectively, are usually set equal in magnitude and equal to $\arctan \frac{b}{Z_{max}}$ where b is the baseline distance and $\mathrm{Z}_{\max}$ the maximum expected depth in the scene. The baselines are also all equal. This situation is full of possibilities for testing the configurations used in the 2–camera analysis as well as searching for other geometric properties which could affect stability. The positioning of the three cameras seems to have been done as a matter of convenience only. How results are affected by that positioning is a problem which has not yet been addressed.

Appendix A

Computing the area of uncertainty

The following computation is for the area of uncertainty for an error bound of $\pm 30\%$ in the visual angle measure and a straight line configuration with features one unit apart and the viewpoint one unit from feature B, as shown in Figure 3.5b. It is assumed that V is at (0,0), A at (1,1), B at (1,0) and C at (1,-1).

1. Determine equations of the four error circles:
   The circle passing through V, A and B has equation $(x-.5)^2+(y-.5)^2=(.707)^2$ with center P at (.5,.5). Due to the fact that the inscribed angle at V equals half the central angle, the isosceles triangle ABP with base AB has altitude d and angle APD equal to α where α is the visual angle from V to A and B (See Figure A.1). To find the equation of the error circle through A and B for a 30% underestimate, a new center P′ must be found such that angle AP′D = .7*α. P′ will lie on an extension of altitude PD because distance(P′,A)= distance(P′,B). Thus, $\tan(.7*\alpha)$=distance(A,D)/$(d+\Delta d)$ where Δd is the distance from P to P′.

   For this particular example, α=45°, distance(A,D)= .5 and $d=.5$. Therefore,
   $\tan(31.5°)=.5/(.5+\Delta d)$
   $\Rightarrow .613=.5/(.5+\Delta d)$
   $\Rightarrow \Delta d=.316$

   Thus, the error circle for a 30% underestimate of visual angle passes through A(1,1) and B(1,0) and is centered at (.184,.5). By solving the equation $(1-.184)^2+(1-.5)^2=r^2$, the radius $r=.957$ is found.

   The error circle for a 30% overestimate of α is found as above, using the equation $\tan(1.3*\alpha)=.5/(.5-\Delta d)$. To find the two error circles for angle β, repeat the entire process using B and C.

   The equations for the four error circles for this example are:

(a) $(x - .184)^2 + (y - .5)^2 = (.957)^2$
(α underestimate)

(b) $(x - .694)^2 + (y - .5)^2 = (.586)^2$
(α overestimate)

(c) $(x - .184)^2 + (y + .5)^2 = (.957)^2$
(β underestimate)

(d) $(x - .694)^2 + (y + .5)^2 = (.586)^2$
(β overestimate)

(The symmetry in the values is due to the symmetric configuration.)

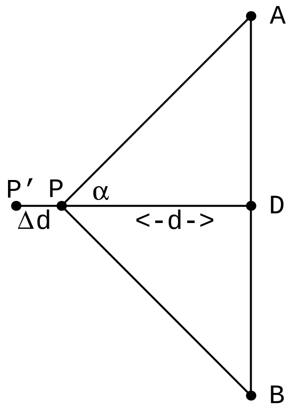

Figure A.1: The center P′ of the error circle for a 30% underestimate of visual angle α lies on an extension of the altitude of triangle APB.

2. Determine the intersection of the error circles:
The important point here is to solve for the correct intersections. Referring to Figure 3.5b and beginning with the intersection point to the left of the viewpoint and proceeding clockwise, the equation pairs to solve are (a) and (c), (c) and (b), (b) and (d), (d) and (a). Note that each pair of circles will intersect in two points. However, one of those points is a feature point. To find the intersection point to the left of the viewpoint, the system :
$(x - .184)^2 + (y - .5)^2 = (.957)^2$ and
$(x - .184)^2 + (y + .5)^2 = (.957)^2$
is solved to get the points (0,1) and (-.632,0). The second point is the one desired. The other three intersection points are (.109,.454), (.387,0) and (.109,-.454).

3. Graph the arcs:
If the circles are defined parametrically, only a single function is needed for the graphing. The following calculations are for the arc of circle (a). The equation $(x - .184)^2 + (y -$

$.5)^2 = (.957)^2$ which must be graphed from (-.632,0) to (.109,-.454) is given parametrically by $(.184 + .957 * \cos\theta, .5 + .957 * \sin\theta)$ with θ ranging from $\pi + \arctan(.5/.816)$ to $\pi + \arctan(.954/.075)$ or 3.69 to 4.63. This process must be repeated for each arc.

Bibliography

[Atiya and Hager, 1993] Sami Atiya and Greg Hager. Real-time vision-based robot localization. *IEEE Transactions on Robotics and Automation*, 9(6):785–800, December 1993.

[Ayache and Faugeras, 1989] Nicholas Ayache and Oliver D. Faugeras. Maintaining representations of the environment of a mobile robot. *IEEE Transactions on Robotics and Automation*, 5(6):804–819, December 1989.

[Babb, 1990] Richard J. Babb. Navigation of unmanned underwater vehicles for scientific surveys. In *Proceedings of the IEEE Symposium on Autonomous Underwater Vehicle Technology*, pages 194–198. IEEE, 1990.

[Bajcsy, 1988] Ruzena Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):996–1005, August 1988.

[Ballard, 1989] Dana H. Ballard. Reference frames for animate vision. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1635–1641. Morgan Kaufmann, 1989.

[Blostein and Huang, 1987] Steven D. Blostein and Thomas S. Huang. Error analysis in stereo determination of 3-d point positions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(6):752–765, November 1987.

[Bowditch, 1802] Nathaniel Bowditch. *The American Practical Navigator (an epitome of navigation)*. Blunt, 1802.

[Brooks, 1986] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, March 1986.

[Cartwright and Collett, 1982] B. A. Cartwright and T. S. Collett. How honey bees use landmarks to guide their return to a food source. *Nature*, 295:560–564, February 1982.

[Cohen *et al.*, 1993] 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*, 18:72–81, 1993.

[Collett *et al.*, 1992] T. S. Collett, E. Dillmann, A. Giger, and R. Wehner. Visual landmarks and route following in desert ants. *Journal of Comparative Physiology A*, 170:435–442, 1992.

[Crowley, 1985] James L. Crowley. Navigation for an intelligent mobile robot. *IEEE Transactions on Robotics and Automation*, RA-1(1):31–41, March 1985.

[Cui *et al.*, 1990] Ning Cui, Juyang Weng, and Paul Cohen. Extended structure and motion analysis from monocular image sequences. In *Proceedings of the International Conference on Computer Vision*, pages 222–229. IEEE, 1990.

[Dai and Lawton, 1993] David Dai and Daryl Lawton. Range-free qualitative navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1, pages 783–790. IEEE, May 1993.

[Davis *et al.*, 1987] L. S. Davis, D. Dementhon, R. Gajulapalli, T. R. Kushner, J. LeMoigne, and P. Veatch. Vision-based navigation: A status report. In *Proc. DARPA Image Understanding Workshop*, pages 153–169, San Mateo, CA, February 1987. Morgan Kaufmann.

[Dhond and Aggarwal, 1991] Umesh R. Dhond and J. K. Aggarwal. A cost-benefit analysis of a third camera for stereo correspondence. *International Journal of Computer Vision*, 6(1):39–58, April 1991.

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

[Dickmanns and Graefe, 1988a] Ernst Dickmanns and Volker Graefe. Applications of dynamic monocular machine vision. *Machine Vision and Applications*, 1:241–261, 1988.

[Dickmanns and Graefe, 1988b] Ernst Dickmanns and Volker Graefe. Dynamic monocular machine vision. *Machine Vision and Applications*, 1:223–240, 1988.

[Dickmanns, 1992] Ernst D. Dickmanns. Expectation based dynamic scene understanding. In A. Blake and A. L. Yuille, editors, *Active Vision*, pages 285–318. MIT Press, 1992.

[D'Orazio *et al.*, 1992] T. D'Orazio, M. Ianigro, E. Stella, and A. Distante. Self localization of a mobile robot using visual landmarks. In *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1869–1874. IEEE, July 1992.

[Duerr, 1992] T. E. Duerr. Effect of terrain masking on GPS position dilution of precision. *Navigation*, 39(3):317–323, 1992.

[Dyer, 1991] Fred C. Dyer. Bees acquire route-based memories but not cognitive maps in a familiar landscape. *Animal Behavior*, 41(2):239–246, 1991.

[Enge, 1993] Per K. Enge. Forward error correction for radiobeacon broadcast of differential GPS data. *IEEE Transactions on Aerospace and Electronic Systems*, 29(1):223–232, January 1993.

[Fennema *et al.*, 1990] C. Fennema, A. Hansen, E. Riseman, J. R. Beveridge, and R. Kumar. Model-directed mobile robot navigation. *IEEE Transactions on Systems, Man and Cybernetics*, 20(6):1352–1369, November/December 1990.

[Ferrari *et al.*, 1990] F. Ferrari, E. Grosso, G. Sandini, and M. Magrassi. A stereo vision system for real time obstacle avoidance in unknown environment. In *Proceedings of the 1990 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 703–708. IEEE, July 1990.

[Fischler and Bolles, 1981] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, June 1981.

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

[Gallistel, 1990] Charles R. Gallistel. *The Organization of Learning*. MIT Press, 1st edition, 1990.

[Grimson *et al.*, 1991] W. Eric Grimson, Daniel P. Huttenlocher, and David W. Jacobs. Affine matching with bounded sensor error: A study of geometric hashing and alignment. Technical Report AI-Memo-1250, Massachusetts Institute of Technology, August 1991.

[Hager and Mintz, 1991] Greg Hager and Max Mintz. Computational methods for task-directed sensor data fusion and sensor planning. *The International Journal of Robotics Research*, 10(4):285–313, August 1991.

[Haralick *et al.*, 1989] Robert M. Haralick, Hyonam Joo, Chung-Nam Lee, Xinhua Zhuang, Vinay G. Vaidya, and Man Bae Kim. Pose estimation from corresponding point data. *IEEE Transactions on Systems, Man and Cybernetics*, 19(6):1426–1446, November 1989.

[Harris *et al.*, 1987] Stewart E. Harris, Robert H. Squires, and Emile M. Bergeron. Underwater imagery using an electronic still camera. In *Oceans '87*, pages 1242–1245. IEEE, 1987.

[Hoffman and Krotkov, 1993] Regis Hoffman and Eric Krotkov. Terrain mapping for outdoor robots: Robust perception for walking in the grass. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1, pages 529–533. IEEE, May 1993.

[Huntingford, 1984] Felicity Huntingford. *The Study of Animal Behavior*. Chapman and Hall, 1984.

[Hutchins, 1983] Edwin Hutchins. Understanding micronesian navigation. In Dedre Gentner and Albert L. Stevens, editors, *Mental Models*, pages 191–225. Lawrence Erlbaum Associates, 1983.

[Hutchins, 1993] Edwin Hutchins. *Cognition in the Wild*. MIT Press - Bradford Books, 1993.

[Huttenlocher and Ullman, 1990] Daniel P. Huttenlocher and Shimon Ullman. Recognizing solid objects by alignment with an image. *International Journal of Computer Vision*, 4:195–212, 1990.

[Huttenlocher, 1988] Daniel P. Huttenlocher. *Three-Dimensional Recognition of Solid Objects from a Two-Dimensional Image*. Tech Report 1045, MIT, 1988.

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

[Intrator, 1993] 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.

[Jacobs, 1991] David W. Jacobs. Optimal matching of planar models in 3D. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 269–274. IEEE Computer Society Press, 1991.

[Kalman, 1960] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, pages 35–45, 1960.

[Kosaka and Kak, 1992] Akio Kosaka and Avinash C. Kak. Fast vision-guided mobile robot navigation using model- based reasoning and prediction of uncertainties. *CVGIP: Image Understanding*, 56(3):271–329, November 1992.

[Kriegman *et al.*, 1989] D. J. Kriegman, E. Triendl, and T. O. Binford. Stereo vision and navigation in buildings for mobile robots. *IEEE Transactions on Robotics and Automation*, 5(6):792–803, October 1989.

[Krotkov, 1989] 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.

[Kuipers and Levitt, 1988] Benjamin J. Kuipers and Tod S. Levitt. Navigation and mapping in large-scale space. *AI Magazine*, pages 25–43, Summer 1988.

[Kuipers, 1978] Benjamin Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978.

[Lamdan *et al.*, 1988] Yehezkel Lamdan, Jacob T. Schwartz, and Haim J. Wolfson. Object recognition by affine invariant matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 335–344. IEEE, June 1988.

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

[Levitt *et al.*, 1987] T. S. Levitt, D. T. Lawton, D. M Chelberg, and P. C. Nelson. Qualitative navigation. In *Proc. DARPA Image Understanding Workshop*, pages 447–465, Los Altos, CA, February 1987. Morgan Kaufmann.

[Levitt *et al.*, 1988] T. S. Levitt, D. T. Lawton, D. M Chelberg, K. V. Koitzsch, and John W. Dye. Qualitative navigation II. In *Proc. DARPA Image Understanding Workshop*, pages 319–326, Los Altos, CA, April 1988. Morgan Kaufmann.

[Lu and Lachapelle, 1992] G. Lu and G. Lachapelle. Statistical quality control for kinematic GPS positioning. *Manuscripta Geodaetica*, 17(5):270–281, 1992.

[Lucarini *et al.*, 1993] G. Lucarini, M. Varoli, R. Cerutti, and G. Sandini. Cellular robotics: Simulation and HW implementation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 846–852. IEEE, May 1993.

[Maes and Brooks, 1990] Pattie Maes and Rodney A. Brooks. Learning to coordinate behaviors. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 796–802. AAAI Press, July 1990.

[Marr, 1982] David Marr. *Vision*. W. H. Freeman and Company, 1982.

[Matthies and Shafer, 1987] Larry Matthies and Steven A. Shafer. Error modeling in stereo navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):239–248, June 1987.

[Mattos, 1992] Philip Mattos. GPS. *Electronics and Wireless World*, 98:982–987, 1992.

[Maver and Bajcsy, 1993] Jasna Maver and Ruzena Bajcsy. Occlusions as a guide for planning the next view. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(5):417–433, May 1993.

[McFarland and Houston, 1981] David McFarland and Alasdair Houston. *Quantitative Ethology: The State Space Approach*. Pitman Advanced Publishing Program, 1981.

[McKendall and Mintz, 1990] Raymond McKendall and Max Mintz. Sensor-fusion with statistical decision theory: A prospectus of research in the GRASP lab. Technical Report MS-CIS-90-68, University of Pennsylvania, September 1990.

[Muller and Wehner, 1988] Martin Muller and Rudiger Wehner. Path integration in desert ants, cataglyphis fortis. *Proceedings of the National Academy of Science*, 85:5287–5290, July 1988.

[Negahdaripour and Yu, 1990] Shariar Negahdaripour and C. H. Yu. Passive optical sensing for near-bottom stationkeeping. In *Oceans '90*, 1990.

[Negahdaripour *et al.*, 1990] Shahriar Negahdaripour, Chih Ho Yu, and Amir Shokrollahi. Recovering shape and motion from undersea images. *IEEE Journal of Oceanic Engineering*, 15(3):189–198, July 1990.

[Negast and Paschall, 1992] 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 Conference*, pages 356–362, May 1992.

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

[Pick *et al.*, 1993] 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.

[Sandini and Tagliasco, 1980] Giulio Sandini and Vincenzo Tagliasco. An anthropomorphic retina-like structure for scene analysis. *Computer Graphics and Image Processing*, 14:365–372, 1980.

[Sanso, 1973] Fernando Sanso. An exact solution of the roto-translation problem. *Photogrammetria*, 29:203–216, 1973.

[Shepard and Hurwitz, 1985] Roger N. Shepard and Shelley Hurwitz. Upward direction, mental rotation and discrimination of left and right. In Steven Pinker, editor, *Visual Cognition*, pages 161–194. MIT Press, 1985.

[Slater, 1985] P. J. B. Slater. *An Introduction to Ethology*. Cambridge University Press, 1985.

[Smith and Cheeseman, 1986] Randall C. Smith and Peter Cheeseman. On the representation and estimation of spatial uncertainty. *The International Journal of Robotics Research*, 5(4):56–68, Winter 1986.

[Sorenson, 1970] Harold W. Sorenson. Least-squares estimation: from Gauss to Kalman. *IEEE Spectrum*, pages 63–68, July 1970.

[Stewart, 1991] W. Kenneth Stewart. Remote-sensing issues for intelligent underwater systems. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 230–235. IEEE, June 1991.

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

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

[Sutherland and Thompson, 1994] Karen T. Sutherland and William B. Thompson. Pursuing projections: Keeping a robot on path. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, pages 3355–3361. IEEE, May 1994.

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

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

[Talluri and Aggarwal, 1992] Raj Talluri and J. K. Aggarwal. Position estimation for an autonomous mobile robot in an outdoor environment. *IEEE Transactions on Robotics and Automation*, 8(5):573–584, October 1992.

[Thompson and Kearney, 1986] W. B. Thompson and J. K. Kearney. Inexact vision. In *Proc. Workshop on Motion: Representation and Analysis*, pages 15–21. IEEE, May 1986.

[Thompson and Pick, 1992] William B. Thompson and H. L. Pick, Jr. Vision-based navigation. In *Proc. DARPA Image Understanding Workshop*, pages 149–152, San Mateo, CA, January 1992. Morgan Kaufmann.

[Thompson *et al.*, 1993] 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.

[Thompson, 1958] E. H. Thompson. An exact linear solution of the problem of absolute orientation. *Photogrammetria*, 13(4):163–178, 1958.

[Thorpe *et al.*, 1987] Charles Thorpe, Steven Shafer, and Takeo Kanade. Vision and navigation for the Carnegie Mellon Navlab. In *Proc. DARPA Image Understanding Workshop*, pages 143–152, San Mateo, CA, February 1987. Morgan Kaufmann.

[Tranquilla and Al-Rizzo, 1993] J. M. Tranquilla and H. M. Al-Rizzo. Investigation of GPS precise relative static positioning during periods of ice clouds and snowfall precipitation. *IEEE Transactions on Geoscience and Remote Sensing*, 31(1):295–299, January 1993.

[Tsubouchi and Yuta, 1987] T. Tsubouchi and S. Yuta. Map assisted vision system of mobile robots for reckoning in a building environment. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1978–1984. IEEE, May 1987.

[Walter, 1953] W. Grey Walter. *The Living Brain*. W. W. Norton, New York, 1953.

[Wehner and Harkness, 1983] Rudiger Wehner and Robert Harkness. *Foraging Strategies in Individually Searching Ants: Cataglyphis Bicolor*. Akademie der Wissenschaften und der Literatur, 1983.

[Wehner and Raber, 1979] Rudiger Wehner and F. Raber. Visual spatial memory in desert ants, cartaglyphis bicolor. *Experientia*, 35:1569–1571, 1979.

[Wehner and Srinivasan, 1981] Rudiger Wehner and Mandyam Srinivasan. Searching behavior of desert ants, genus cataglyphis. *Journal of Comparative Physiology*, 142:315–338, 1981.

[Wehner, 1981] Rudiger Wehner. Spatial vision in anthropods. In H. Autrum, editor, *Physiology and Evolution of Vision in Invertebrates*, pages 287–617. Springer, 1981.

[Wehner, 1990] Rudiger Wehner. Do insects have cognitive maps? *Annual Review of Neuroscience*, 13:403–414, 1990.

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

[Wu and Melbourne, 1993] Sien-Chong Wu and William G. Melbourne. An optimal GPS data processing technique for precise positioning. *IEEE Transactions on Geoscience and Remote Sensing*, 31:146–152, January 1993.

[Wu *et al.*, 1992] Sien-Chong Wu, W. I. Bertiger, and J. T. Wu. Minimizing selective availability error on satellite and positioning. *Journal of Guidance, Control and Dynamics*, 15:1306–1309, September/October 1992.

[Yacoob and Davis, 1992] Yaser Yacoob and Larry Davis. Computational ground and airborne localization over rough terrain. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 781–783. IEEE, June 1992.