Spatial Cognition and Computation **1:** 291–321, 1999.

Geometric reasoning under uncertainty for map-based localization

WILLIAM B. THOMPSON[1], CAROLYN M. VALIQUETTE[1], BONNIE H. BENNETT[2] and KAREN T. SUTHERLAND[3]
[1]*Department of Computer Science, University of Utah, U.S.A.;* [2]*Graduate Programs in Software, University of St. Thomas, St. Paul, MN, U.S.A.;* [3]*Department of Computer Science, Augsburg College, Minneapolis, MN, U.S.A.*

Abstract. Map-based navigation in outdoor terrain lacking man-made structures or other highly distinctive landmarks can produce severe localization problems. This paper presents an approach to navigation which implements high level geometric reasoning and matching strategies based on those used by skilled human navigators. This approach, which is demonstrated on a real example involving imagery of mountainous terrain obtained with a video camera and USGS map data, is designed to avoid many of the pitfalls occurring when an attempt is made to navigate by modeling the environment mathematically. It exploits feature attributes which cannot be easily expressed quantitatively but are central to the successful human navigation process.

Key words: localization, maps, navigation

1. Introduction

An essential aspect of map-based navigation is the determination of an agent's current location based on sensed data from the environment. Formally, this amounts to specifying the current viewpoint in some world model coordinate system. This *localization* process has two distinct components: one involving the establishment of correspondences between aspects of the sensed data and the map or model, and the other involving derivation of constraints on the viewpoint based on the correspondences that have been determined.[1]

Correspondences can be established at the signal or feature level. Signal-level matching correlates sensed data with predictions of how the sensed data should appear. It works best when the uncertainty in the viewpoint is small and when it is relatively easy to accurately generate expected sensor data. For example, in the TERCOM and SITAN cruise missile guidance systems, a digital elevation model is matched against a downward looking, radar sensed elevation profile (Andreas et al. 1978; Baird and Abramson 1984).

Several researchers have addressed the more difficult problem of signal-based localization at or near ground level using horizontally oriented imaging systems and passive sensing. Ernst and Flinchbaugh (1989) determine deviations between expected and observed views using curve matching algorithms. Yacoob and Davis (1991) and Talluri and Aggarwal (1992) determine viewpoint under the assumption that viewpoint elevation is known with high precision in the reference frame of the map, a situation which dramatically reduces complexity but is unfortunately not likely to hold in practice. Stein and Medioni (1992) proposed an alternate method for determining viewpoint based on the observed horizon line which is similar to the characteristic view approach in object recognition. Cozman and Krotkov (1997) use a statistical approach to related landmarks in views and on a map. In all of these methods, actual viewpoint determination is done using the same types of methods involved in photogrammetry – which solves the same problem (Sanso 1973; Thompson 1958), in alignment approaches to object recognition (Huttenlocher and Ullman 1987; Grimson 1990), or using an exhaustive search hypothesize-and-test approach.

For a vision-based system in such environments, there is limited *a priori* knowledge about the viewpoint due to travel through indistinct terrain, temporary occlusion of landmark features, or errors in position updating processes. The view of the world at or near ground level is difficult to generate from map data with sufficient fidelity to allow signal-level matching. Sensed data can contain substantial geometric aberrations not easily described using simple error models. Finally, available digital cartographic data sets contain inaccuracies that can cause serious problems for correlation-based analysis. For example, in one of the USGS DEMs that make up our test data, the location of the high point of a significant peak is off by over 200 m. It is not surprising that most of the published work on vision-based localization from a ground-level perspective has been demonstrated only on synthetic data, where these problems do not occur.

Feature-based approaches to vision-based localization hold the potential for avoiding many of these problems. As shown in Figure 1, features are extracted independently from sensed data (*view features*) and maps (*map features*) and then matched symbolically if they are likely to correspond to the same physical landmarks (*terrain features*). As a result, there is no longer a need to be able to synthesize an accurate rendition of expected sensed data. The symbolic nature of matching and viewpoint inference allows the introduction of sophisticated problem solving methods which are able to deal with issues such as ambiguity and complex error models.

In the remainder of this paper, we describe one possible approach to feature-based localization in outdoor terrain. Our emphasis is on matching

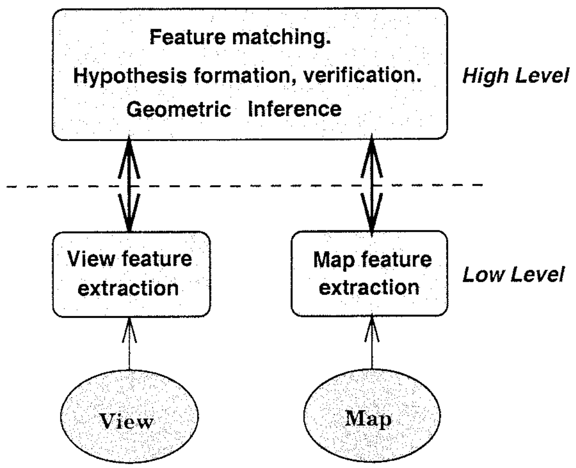

Figure 1. Symbolic matching of landmark features extracted from map and view.

strategies that can accommodate ambiguity due to correspondence errors and on qualitative geometric reasoning procedures for determining viewpoint while maintaining an explicit representation of the uncertainty associated with that determination. Rather than the traditional approach of applying mathematical equations which often break down when faced with the above described errors, this approach is modeled after the skills used by human navigators, whose expertise is based on high level matching schemes and reasoning instead of the solving of equations. The approach is demonstrated on a real example involving imagery obtained with a video camera and map data provided by the USGS.

2. Geometric inference about viewpoint

In photogrammetry and pose estimation, a set of view/model correspondence is used to solve for the relative orientation between viewpoint and model. A follow-up confirmation step is then often performed to verify that additional model features appear where expected in the view. Some methods also yield a simple statistical measure of error. The difficulty of uniquely identifying landmarks in outdoor terrain makes this sort of approach problematic for navigation due to the combinatorics of possible correspondences. In addition, the nature of positional uncertainty cannot effectively be represented using low order models of statistical deviation (Sutherland and Thompson 1994).

The methods of photogrammetry and pose estimation work by finding a view transformation that optimally accounts for a set of view/model correspondences. In landmark-based navigation, the high likelihood of error in at least some of the presumed correspondences plus the complex nature of the errors associated with those correspondences which are in fact correct argues for a separate analysis of small sets of correspondences, rather than a single batch solution to the localization problem. This allows a separation between viewpoint inference based on selected landmarks and the combining of information about viewpoint based on an analysis of many landmarks.

2.1. *Inference methods*

Accurate measurements of the range to distant landmarks are seldom possible without the use of specialized active sensors. As a result, we limit our analysis to inference methods that depend primarily or exclusively on bearing and visual angle in the view. We have identified four classes of geometric inference relevant to vision-based navigation using bearing or visual angle. A fifth inference class that does not necessarily require visual processing also appears to be quite important and will be briefly discussed in Section 3.1.

- *Absolute bearing*: This is the standard way to solve localization problems. The viewpoint is on a line from the landmark along a reciprocal bearing to that of the viewed landmark (Bowditch 1958). To use this inference method, an accurate compass registered to the map coordinate system is required.
- *Ordinal view*: The ordinal position of two features (e.g., "**A** is *left-of* **B**") can be used to constrain the viewpoint to lie on one side of a line through the positions of **A** and **B** (Kuipers and Levitt 1988; Levitt and Lawton 1990; Schlieder 1995) (see Figure 2).
- *Relative bearing*: Relative bearings between three or more image features with known map positions lead to a classical pose estimation problem. Well established numerical techniques exist for solving such problems. Levitt and Lawton describe an alternate method in which only two features are considered at a time (Levitt and Lawton 1990). The visual angle between the two features with known map positions constrains the viewpoint to lie on a particular circle on the map (see Figure 3).
- *Alignment*: If two features line up along a line of sight, then the viewpoint is constrained to lie on a line connecting the two features. In almost all circumstances encountered in outdoor navigation, it is possible to determine which of the two features is more distant and as a result the viewpoint can be constrained to a half-line. This inference method appears to be critical when human map users are solving difficult localization problems (Pick 1996).

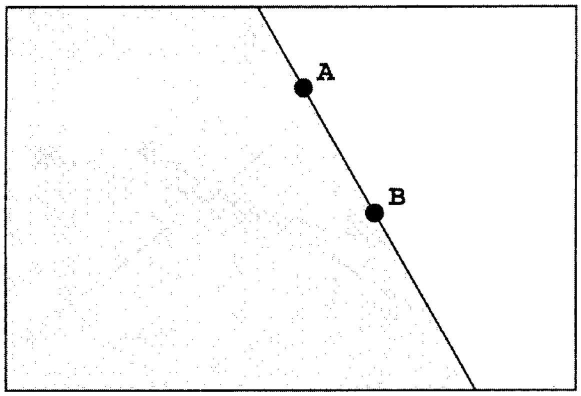

Figure 2. Viewpoint is in shaded area if landmark **A** is seen to the left of landmark **B**.

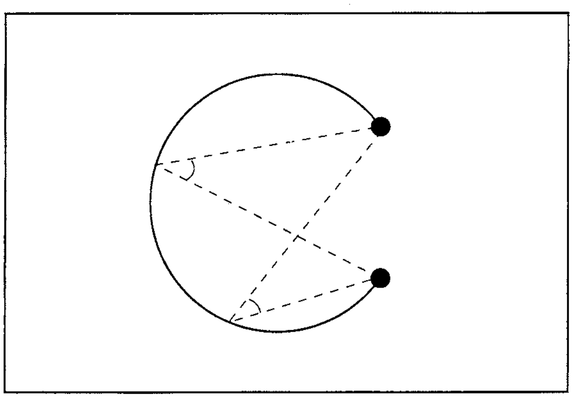

Figure 3. Relative bearing to two landmarks constrains viewpoint to a circle.

2.2. *Combining uncertain constraints on viewpoint*

Three of the four inference methods above can be used to constrain the viewpoint to a line or arc on the map. Two non-degenerate line or arc constraints can be used to uniquely determine the viewpoint, although the arcs resulting from the relative bearing methods may yield multiple solutions which require additional constraints in order to select the true location. Uncertainty in bearing estimates and map locations makes it desirable to use additional constraints, though the resulting over-constrained system must then be solved using some sort of optimization technique.

Figure 4 illustrates the common situation in which the absolute bearing to more than two landmarks is measured. The classical approach to determining a viewpoint estimate in this situation would be to use a linear least-squares

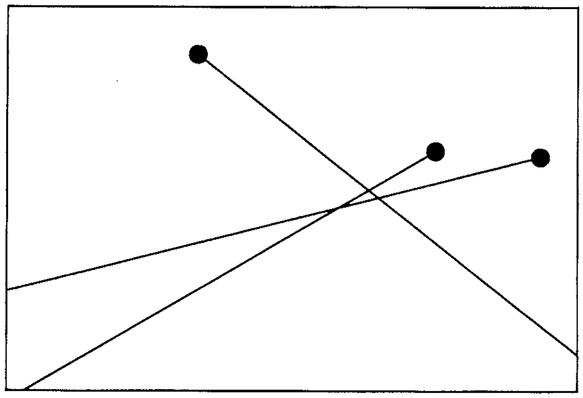

Figure 4. Absolute bearing to three landmarks results in an over constrained system for determining viewpoint.

method to find a viewpoint location "close" to all of the constraint lines. For some configurations of landmarks, this is in fact a reasonable thing to do. For other configurations, however, the use of least squares methods to find a single estimate masks important aspects of the uncertainty in the estimates. This becomes even more important with the non-linear cases resulting from the use of relative bearing (Sutherland and Thompson 1994).

Two very different sorts of errors affect viewpoint estimates in feature-based localization. Features are seldom sufficiently distinctive to be uniquely identifiable in either the view or map. As a result, correspondence errors are frequent. The effect is to associate a view feature with an incorrect map location. Our approach to dealing with correspondence errors is discussed in Section 4. Even when correspondences are correctly determined, there will be uncertainty in the estimated positions of features in the map and view. These feature localization errors, which can cause significant uncertainty in viewpoint determinations, can be classified as either *sensing errors* or *position specification errors*.

Most of the research on uncertainty in localization has been limited to indoor applications for which 3–D sensors with well described error models are used (e.g., Brooks 1985; Moravec 1988). The nature of localization uncertainty in large-scale outdoor environments is much different. As previously mentioned, accurate range to distant landmarks is seldom available. Measurements of bearing or relative visual angle suffer from *sensing errors* due to distortion, calibration inaccuracies, and sampling errors which can often be severe, particularly over larger viewing angles. In addition, there is uncertainty in the exact location of landmark features in both view and map, since the features used as landmarks in outdoor terrain have substantial but ill-

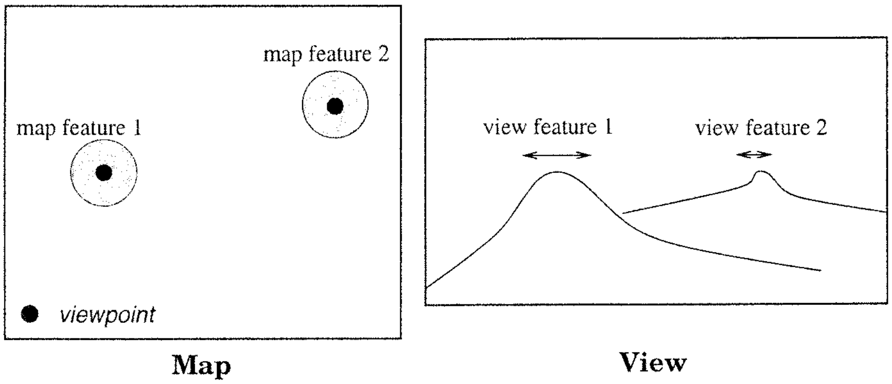

Figure 5. View localization uncertainty modeled as map uncertainty.

defined physical extent. These *position specification errors* are compounded by viewpoint dependencies. For example, except when the line of sight is horizontal, the apparent high point of a peak in the view will usually not correspond to the actual high point.

It is important to distinguish between sensing and position specification error when implementing algorithms for geometric inferences about viewpoint and correspondences. Sensing errors manifest themselves as uncertainties in either relative or absolute visual angle. The actual error model needed is a function of the sensors involved. Our geometric inference techniques presume the availability of a bound on maximum angular error. While a statistical evaluation of possible viewpoints is feasible, assuming a uniform distribution of sensor error across an interval is likely to be more believable than those probabilistic models which require unrealistic assumptions about the distribution of error.

Position specification errors for map features can easily be modeled by introducing a region of uncertainty around the presumed map location of the feature. View position specification errors would appear to be more difficult to deal with, since they are viewpoint dependent: the larger the spatial extent of the feature in the view the larger the uncertainty in feature location is likely to be. It turns out to be possible to transfer this uncertainty to the map and treat it in the same manner as for map features.

Figure 5 shows the reasoning. The uncertainty in locating a view feature is a function of the extent of the feature in the view. This in turn is directly related to the physical extent of the feature in the world and the viewing distance. Thus, position specification errors can be accounted for by assuming that there is an uncertainty in the location of map features sufficient to account for errors in extracting both map and view feature locations. The magnitude

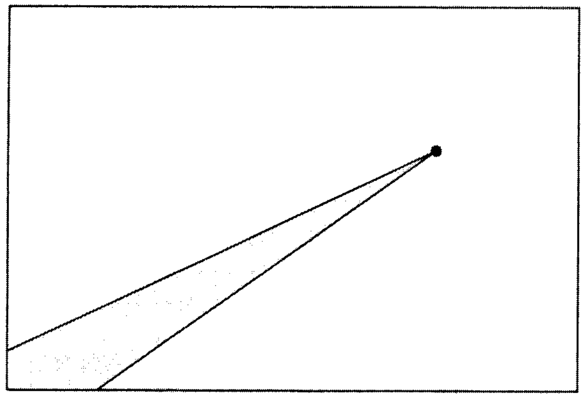

Figure 6. The viewpoint is constrained to the gray area when there is uncertainty in absolute bearing.

of this uncertainty is dependent on the nature of the local topography: Glaciated terrain such as used in the example shown in Section 5 is characterized by a predominance of sharp features that are relatively easy to localize, while topographic features in rolling terrain have much greater positional uncertainty.

We have found it useful to employ interval models of uncertainty in sensor values and map accuracy, and use these to determine feasible viewpoint regions. These regions are chosen to be compatible with the expected variability of measured bearings. Multiple constraints are evaluated using intersection to determine the region or regions compatible with each of the constraints individually (A discussion of one way to do a probabilistic analysis of this problem can be found in Sutherland (1994)). Figure 6 shows the possible viewpoint region associated with a single absolute bearing constraint, when there is uncertainty in the bearing. Figure 7 shows the resulting viewpoint region in dark gray after intersecting three uncertain absolute bearing constraints. Figure 8 illustrates the results of a complete error analysis for three intersecting absolute bearing constraints. Each landmark is represented with a circle of uncertainty that captures imprecision in the actual landmark location, plus likely difficulties in precisely locating the landmark in the view. The wedge spreading out from these circles accounts for errors in the measured bearings to each landmark. The viewpoint region is shown in dark gray.

The gray area in Figure 9 is the area in which the viewpoint could lie given an uncertain relative bearing measurement to two landmarks, assuming sensing error within a given bound but no position specification error. The landmarks are represented by the small black disks. An extensive analysis

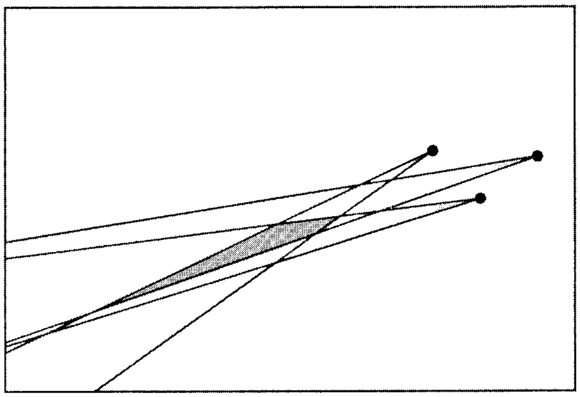

Figure 7. Intersecting multiple uncertain absolute bearing constraints. The viewpoint is constrained to the dark gray region.

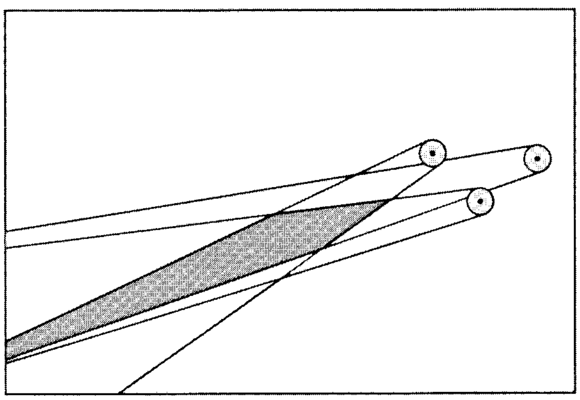

Figure 8. Intersecting absolute bearing constraints with both sensing and position specification errors. The viewpoint lies in the dark gray region.

of the results in intersecting two or more such regions can be found in Sutherland and Thompson (1994). Figure 10 shows an equivalent viewpoint region with the uncertain relative bearing due to position specification error but no sensing error. The landmark points are now surrounded by dark gray circles of uncertainty. In Figure 11, the errors shown in the previous figures are combined, resulting in a larger area in which the viewpoint might be located. Unlike in the method of absolute bearing, the size of this area is not only a function of the bound on sensing error and the radii of the circles of uncertainty surrounding the landmark points, but also of the true visual angle measure. For any given bounds on sensing and position specification errors,

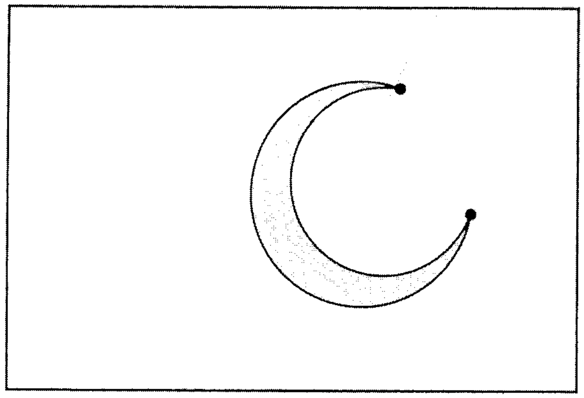

Figure 9. Viewpoint constraint based on relative visual angle with sensing errors.

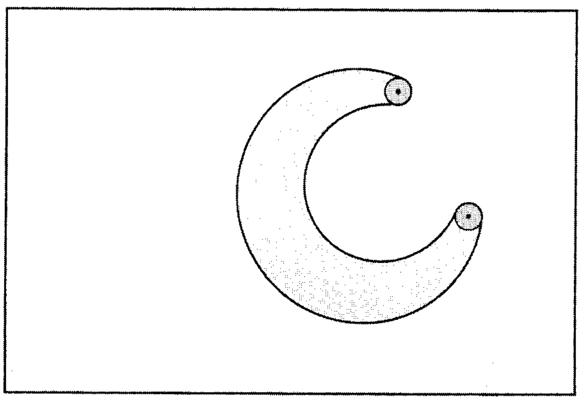

Figure 10. Viewpoint constraint based on relative visual angle with position specification errors.

the area will increase in size as the visual angle decreases or, equivalently, as the viewpoint moves further from the landmark pair.

Figure 12 shows the effect of sensing errors on viewpoint localization using feature alignment. Rather clearly, it is important that the distance from the actual viewpoint to the nearer feature not be too much larger than the distance between the features. A similar effect occurs with ordinal position, as shown in Figure 13. This serves to limit the usefulness of the ordinal position inference methods, except for widely spaced landmarks.

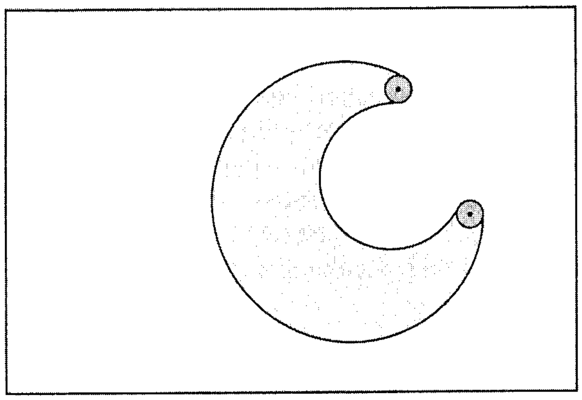

Figure 11. Viewpoint constraint based on relative visual angle with both sensing and position specification errors.

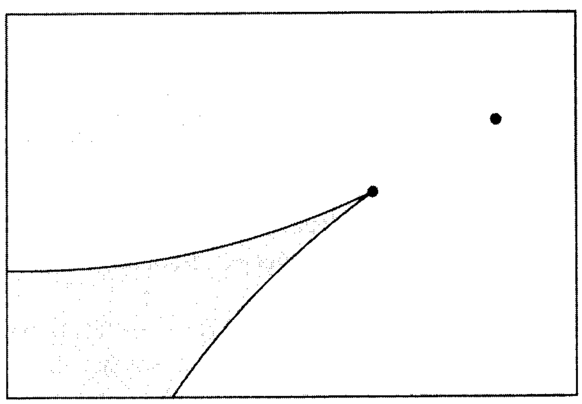

Figure 12. Viewpoint constraint based on feature alignment with sensing errors.

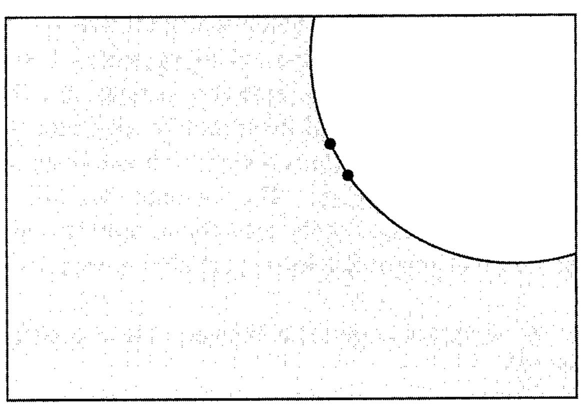

Figure 13. Viewpoint constraint based on ordinal position with sensing errors.

3. Constraint-based analysis

Ambiguity in landmark recognition and matching is a central problem in outdoor navigation. Landmarks are seldom so distinctive that they can be unambiguously recognized as unique entities. Dead reckoning errors and errors in prior localization determinations hamper methods based solely on predict-verify operations. These problems are particularly troublesome in outdoor environments with few cultural features.

3.1. *Interacting constraints*

One way to reduce the impact of these effects is to use a constraint-based approach which interleaves the establishment of correspondences with the estimation of viewpoint: Easily determined correspondences are used to obtain an initial estimate of viewpoint location which facilitates establishment of additional correspondences. This in turn allows the estimated viewpoint location to be further refined. Uncertainty is represented explicitly by allowing for alternate hypotheses about possible landmark correspondences, with each hypothesis specifying the resulting regions within which the viewpoint must lie if the hypothesis is correct. Constraint satisfaction is used to discard hypothesized correspondences that lead to implausible predictions and to refine the possible viewpoint regions associated with the remaining hypothesized correspondences.

The geometric inference methods described in Section 2.1 are most commonly used to generate *viewpoint constraints* specifying where the viewpoint must lie to be consistent with a particular set of corresponding map and view features. The validation of viewpoint hypotheses involves the generation of *distant constraints* which specify additional correspondences for landmarks distant from the proposed viewpoint that should hold if the hypothesis is valid. These constraints can involve correspondences between map and view features as well as relationships between features in either the map or view. Information about the nature of the terrain in the immediate vicinity of the viewpoint leads to *local constraints* which limit the viewpoint to compatible terrain features on the map (E.g., "I'm standing on a hill. Therefore, the viewpoint must be located at one of the hills found represented on the map.") These appear to be quite important for expert human map users (Pick et al. 1996).

Three types of geometric inference between these constraint types are possible (Figure 14):

– *Distant constraints* $\Rightarrow$ *constraints on viewpoint*: Map/view feature correspondences for sets of distant features can be used to determine

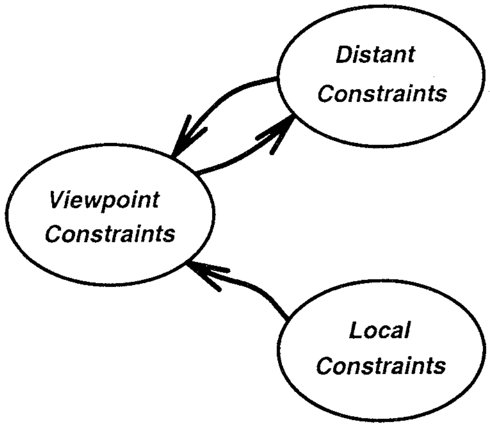

Figure 14. Reasoning about viewpoint involves interacting constraints.

constraints on the viewpoint using various forms of trigonometric analysis.

- *Constraints on viewpoint ⇒ expectations about distant constraints*: Hypotheses about viewpoint can be evaluated by examining distant features. A possible viewpoint, together with one or more view features, can be used to predict the location of the corresponding map feature(s). Likewise, a map feature can be used to predict the nature and location of view features. If these expectations fail to be met, the hypothesized viewpoint is likely in error.
- *Local constraints ⇒ constraints on viewpoint*: Local constraints allow for the enumeration of possible viewpoints. Such an enumeration can be intersected with the constraint regions that usually arise from consideration of distant features.

3.2. *Evaluating viewpoint hypotheses*

In addition to generating new hypotheses about the viewpoint, the inference methods described in Section 2.1 can be used to to evaluate existing viewpoint hypotheses as well. Two sorts of evaluation can occur. Features in the view not used in forming the original viewpoint hypothesis can be used to generate expectations about compatible terrain features that should appear within particular regions of the map if the hypothesis is correct. Likewise, features on the map can be used to generate expectations about features that should appear in the view.

Outdoor terrain contains many similar appearing landmark features. As a result, the discovery of a particular landmark where one is expected is only weak evidence confirming a viewpoint hypothesis. On the other hand, the failure to find a feature where one is expected is strong evidence that the hypothesis on which the expectation is based is incorrect. This *disconfirmation strategy* plays an important role in human map usage (Heinrichs et al. 1992).

3.2.1. *Confirming view expectations on the map*

Given an hypothesized viewpoint on the map, the absolute bearing to view features can easily be used to specify a map region where the corresponding map feature should appear. This map region is then searched for viable matches to the view feature. This process is very important since it either generates new feature correspondences, or, when no matching map features are found in the region, provides sufficient disconfirming evidence to reject the hypothesis.

Since hypothesized viewpoints typically specify regions on the map and not just a single location, the area in which the feature is expected to lie is also a region, not just a line specified by the bearing. The region is found by taking the union of the half lines specified by the measured bearing from all points within the supposed viewpoint region. Uncertainty due to sensing errors can be accounted for by recognizing that a viewpoint location plus a bearing specify a wedge within which the feature should lie. Figure 15 illustrates the method. The shaded area indicates a search region on the map in which a feature appearing in the view at a particular bearing is expected to be. This search region includes both the dark gray viewpoint region itself and the light gray wedge emanating from it. As before, feature position specification errors can be modeled by adding a circle of uncertainty around map features. The expectation generated from the view is confirmed on the map if the search region includes any part of the uncertainty circle from a compatible map feature. Note that the search region can be constructed based only on an analysis of the external contour of the viewpoint region. This is also true of the other inference techniques described below.

Figure 16 shows the situation when the ordinal position of two features in the view is known and one of those features has been localized on the map. From the same dark gray viewpoint region as shown in Figure 15, any feature seen to the "left" of **F** will be located either in region **C1** or in one of the darker gray areas shown in the figure. Any feature seen to the "right" of **F** will be located either in region **C2** or in one of the darker gray areas.

If we have one view feature that corresponds to a map feature, and another view feature with a particular visual angle with respect to the first view feature, then the map region in which we expect to find a match for the second view feature is shown in Figure 17, with the darkly shaded region

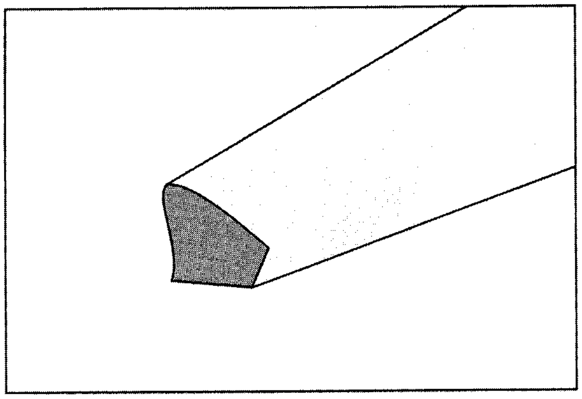

Figure 15. Given a a viewpoint specified by the dark gray region and an uncertain absolute bearing to a view feature, the corresponding map feature lies somewhere in either the viewpoint region or the light gray wedge emanating from it.

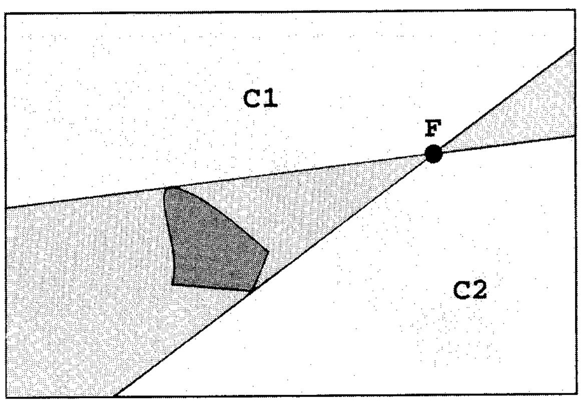

Figure 16. The region in which a map feature is located can be determined using ordinal position with respect to a second feature, F, previously localized on the map.

being the viewpoint region. Figure 18 is a special case of Figure 17 showing the possible locations of a second map feature that appears in the view to align with the marked feature, given a bound on the angular uncertainty with which the features can be localized in the view.

The geometric constructions needed to determine these regions are complex. Computational implementations can often benefit from simplifications. It is important, however, to use a conservative approach. Since constraints are intersected to find features and relationships that satisfy all relevant constraints, it is important that the actual search regions used cover all of the theoretically possible locations.

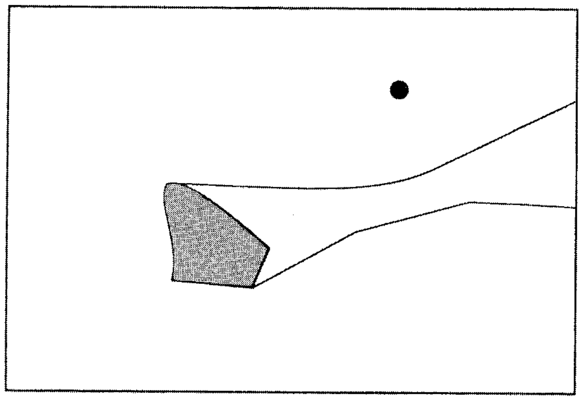

Figure 17. Given a viewpoint specified by the dark gray region and an uncertain relative bearing between one feature indicated by the black dot and another view feature, the corresponding map feature lies somewhere in either the viewpoint region or the light gray region emanating from it.

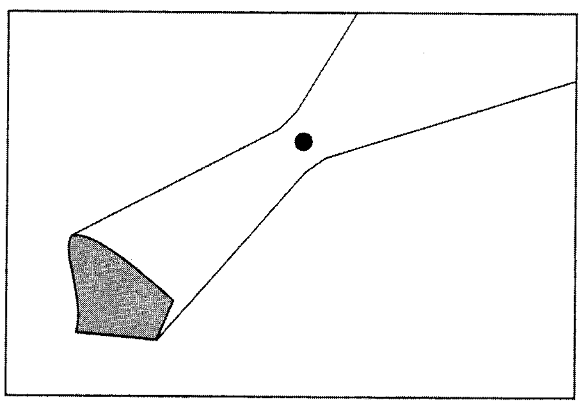

Figure 18. Possible map locations for a feature, given that it aligns with another feature in the view.

3.2.2. *Confirming map expectations in the view*

Given landmark locations on the map and a map region which is known to contain the viewpoint, it is possible to determine constraints on landmarks in the view. In general, we assume that all view features are present on the map, but the converse is not necessarily true. The map may well cover more terrain than is visible in the view. Occlusion also results in missing view features. Thus, the lack of a view feature where the map suggests one is not sufficient cause to reject a hypothesis, making this a less useful process than the view driven searches discussed in the previous section. For the same reasons, human map users tend to concentrate on view features first (Heinrichs et

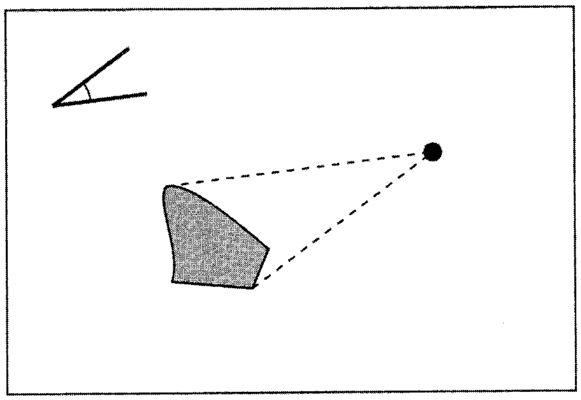

Figure 19. Using absolute bearing to determine where a map feature should appear in the view.

al. 1992). That is not to say that this process is not useful at all. It can be used to verify an hypothesized viewpoint or possibly to establish additional correspondences between features in the view and map.

If map and view orientations are registered using a compass or other means, then absolute bearing can easily be used to determine possible viewing directions from a viewpoint region to any particular map feature not in the region. (If the feature is within the region, any viewing direction is possible.) Figure 19 illustrates the situation. The region boundary is traversed and the range of possible viewing directions is determined, as shown in the upper left of the figure. Though not indicated here, a search for the feature in the view should take into account any uncertainty in determining view bearings that might exist.

Ordinal position can be used to predict that some view feature will be found to the right of (or left of) some other view feature. In Figure 16, any map feature in region **C1** will show up in the view to the left of feature **F** from anywhere within the viewpoint region. Likewise, any map feature in region **C2** will show up in the view to the right of **F**. Any feature located in the dark gray areas might appear to either side of **F**, depending on where the actual viewpoint is located. Such features are thus not useful for validating a viewpoint hypothesis.

Figure 18 can be used to illustrate the way in which it is possible to determine if two map features might appear aligned in the view. Any landmark feature located in one of the gray areas will appear aligned with the marked feature from at least one location within the region of possible viewpoints. Alignment will *not* in general occur at other possible viewpoints,

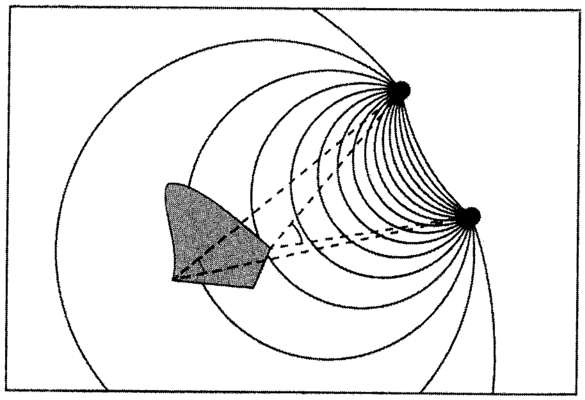

Figure 20. Determining possible visual angles between two landmarks.

however. As a result, this form of analysis is of little value in confirming a particular viewpoint hypothesis unless the viewpoint region is quite small.

As with absolute bearing, the range of relative visual angles between two landmarks outside of the possible viewpoint region can be found by traversing the region boundary, keeping track of the maximum and minimum angles that are found. Figure 20 shows the two extrema angles, along with constant visual angle contours in 10° increments in order to provide a sense of scale and sensitivity.

3.3. *Orienting a view*

In the discussion thus far, the use of absolute bearing has presumed that some sort of compass was available in order to transform viewing directions into map orientations. In fact, it is possible to orient the view with respect to the map in the absence of a compass once one or more of the other inference methods have been used to narrow down the possible viewpoint region. If, in the situation shown in Figure 19, we have no initial means to register map and view orientations, but have determined *both* the view and map locations of the feature by an inference method other than absolute bearing, we can then determine the bearing to the view feature with respect to the map coordinate system. This registration of the two coordinate systems is accurate to within an uncertainty range that is a function of the size and shape of the viewpoint region and the distance to the feature. It follows that north in the view, corresponding to north on the map, can then be hypothesized. For a well calibrated camera system, equivalent information is then provided on all of the other view features as well. This is a potent tool, since it lets us bring into play absolute bearing inferences even if we start out not knowing the orientation of the view and have no direct way of determining that orientation. As demonstrated in the example in Section 5, as more view features are matched to features on the map, this hypothesis of *view north* can be refined. There is evidence that human map users will try to orient the view with respect to the map before generating detailed hypotheses about the viewpoint if they do not have a compass or other indication of direction in the world (Pick et al. 1996).

4. Matching

Although constraint-based analysis using multiple hypotheses and a disconfirmation strategy provides a powerful tool for dealing with the potential of incorrect correspondences between map and view features, it is critical that feature matching be done in a manner that reduces as much as possible the chances for false matches. It is particularly important to restrict the possible associations between map and view features (i.e., the distant constraints), when initial hypotheses about the viewpoint are being formed. At this point in the problem solving process the distant constraints will have little effect in limiting matches. The likelihood that a view feature and a map feature correspond to the same terrain landmark can be estimated based on two aspects of the features: features must be of compatible types and should have compatible geometric properties.

While landmarks are seldom uniquely identifiable in outdoor terrain where distinctive man-made features are absent, they can be grouped into categories based on feature type with some reliability (Brändli 1996). View features, for example, can be divided into the classes of gaps, ridges, saddles, valleys, inclines and peaks. Map features can be organized into a richer structure, since more information is available about their actual geometric shape. This much richer structure is due in large part to the fact that while 3-D structure is directly available from the map, it is quite difficult to recover from the view for both people and machines using passive range sensors. Figure 21 shows a taxonomy of features in mountainous terrain that has proven useful both to account for human performance and in computational simulations (Bennett 1992).[2]

To deal with viewpoint dependencies, it is possible to specify the *a priori* likelihood that a particular view feature type and a particular map feature type are compatible with a single terrain feature. Figure 22 shows a possible empirically determined assignment of such likelihoods (Bennett 1992).[3] Likelihood values range from a value of 5, indicating a strong likelihood of compatibility to a value of 0, indicating no likelihood of compatibility between map and view features. As an example, consider the map feature "Ridges" in Figure

Peak properties.	Ridge properties.
height: average vertical distance from top to adjacent saddles.	*flatness, steepness:* functions of overall ridge slope.
elevation: height of peak.	*extent:* length of ridge line.
sharpness, roundness: functions of the slope of the sides of a peak.	*average elevation:* mean elevation computed over ridge line.
flatness, steepness: opposite of sharpness and roundness.	*elevation variation:* change in elevation over ridge line.
	elevation consistency: opposite of elevation variation.

Figure 24. Peak and ridge feature properties.

rise angle (the angle of the ridge itself relative to the horizontal) and the *break angles* of each of the faces (the slope of the face measured along its fall line).

For a horizontal viewing direction, the projection process is such that the angle of the ridge in the view is never less than the rise angle. This is apparent in Figure 23, where the angle of the ridge relative to the horizontal is equal to the rise angle in the view on the left and greater in the view on the right. It follows that the projected ridge angle in the view for ridges seen in profile is never more than the break angle of the hidden face. If ridge features have been flagged as having a high likelihood of compatibility, additional weight for matching these features is given if the angle of inclination of the ridge in the view lies between the rise angle of the ridge and the break angle of the face. Since a ridge line is made up of multiple ridges, knowledge of the minimum rise angles and maximum break angles in the ridge line on the map constrains to an interval the ridge in the view. Likewise, if the angle of inclination in the view lies outside of this interval, the match can be eliminated from further consideration. In most realistic situations, the viewing angle is sufficiently close to horizontal for this constraint to be useful.

While the geometric and topographic properties of map and view features can help determine whether or not a particular feature match is valid, it is still necessary to determine which features to attempt to match in the first place and also necessary to resolve ambiguous matches. Geometric similarity can only help in matching when there are few features present with similar shapes. As a result, the key features to match are those with geometric properties that are both prominent and unusual.

In the example presented in Section 5, we have implemented the selection of features in the following way. Each feature type is characterized in terms of a set of scalar *properties*. Figure 24 lists the peak and ridge properties used

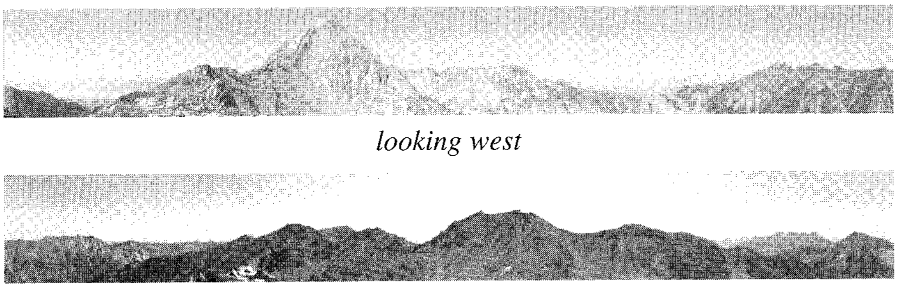

looking west

looking east

Figure 25. Panoramic image of terrain.

in the example. In this case, the same properties are used for view and map features, with the values being actual on the map and apparent in the view. Due to viewpoint dependencies and the lack of reliable 3-D information in the view, these are not necessarily the same. For a given localization problem, each property value has an associated *prominence* value, ranging from 0.0 to 1.0, that indicates the conspicuousness of the feature value within the context of other features of the same type. Prominence alone is not a sufficient criterion on which to select features for matching, however, since the terrain may contain many prominent features, all with about the same shape. As a result, a *distinctiveness* value is computed for each feature type, which has a high value only if there are few features that are prominent with respect to the feature type. Finally, the *saliency* of each feature type for each feature is computed as the product of its prominence and distinctiveness. By ranking each feature in terms of the maximum saliency over that feature's properties and focusing on those features with a saliency over a predefined threshold, attention is concentrated on features most likely to be easily matched.

5. Example of feature-based localization

The geometric reasoning techniques described above have been tested on real terrain data from the Wasatch mountains surrounding Salt Lake City, Utah (Valiquette 1995). The panoramic image shown in Figure 25 provided the view data. The map data was USGS 30 m DEM data covering approximately 21.4 by 28 km.

The view was searched for significant features. Only one feature, the largest, highest peak, has an overall saliency value, as described in Section 4, above a predefined saliency threshold. Two proximity configurations were formed involving the selected peak and each of two nearby ridge lines, one immediately to the left of the peak and one immediately to the right. The physical characteristics of the peak and ridge line features combined are likely to produce a smaller number of potential matches on the map than would be produced if each feature was considered separately. The previously described ordinal view and relative bearing inference methods were used in the formation and refinement of localization hypotheses. Four of the six strategies for localization introduced in Heinrichs et al. (1992) were used: concentrate on the view first, organize features into configurations, pursue multiple hypotheses and evaluate hypotheses using a disconfirmation strategy.

With no *a priori* knowledge of view orientation (see Section 3.3), the map was searched for configurations involving a peak and a nearby ridge line. Nineteen localization hypotheses were generated. An initial map location region was generated for each hypothesis using the map feature positions and the relative bearing between the view features. The view and map coordinate systems were then used to determine for each hypothesis an estimate of absolute bearing which will be referred to as *view north.* The accuracy of the view north estimate determines the angular extent of map search and constraint regions and must be updated as new feature correspondences are formed.

After estimating view north, a search was made for map features to match other highly salient view features. If exactly one map feature was found to match the view feature, another feature correspondence was added to the hypothesis, used to refine view north, and then to refine the viewpoint region. Ambiguous features were noted, and the search for a match was repeated later when improvements in view north and/or reductions in the size of the location region might lead to a smaller search region and subsequently a single match. If no map features were found to match the view feature, this disconfirming evidence was sufficient to reject the hypothesis. The steps of refining view north, refining the map location region, and searching for ambiguous view features were repeated until there was no longer improvement in localization.

Figures 26–28 demonstrate the refinement of the actual correct hypothesis. In the figures, four types of uncertainty regions are overlaid on top of the topographic map for a portion of the area in which the viewpoint lies. Light gray specifies a region within which the viewpoint is hypothesized to lie at a given point in the analysis. Mid-light gray illustrates the search region used in the map to find a feature which appears in the view. When a possible correspondence is found between a map and view feature, absolute bearing can be used to constrain the possible viewpoint, as shown in mid-dark gray. The intersection of the current (light gray) viewpoint hypothesis and the (mid-dark gray) constraints derived from corresponding map and view features often allows a reduction of the hypothesized viewpoint region, as shown in dark gray.

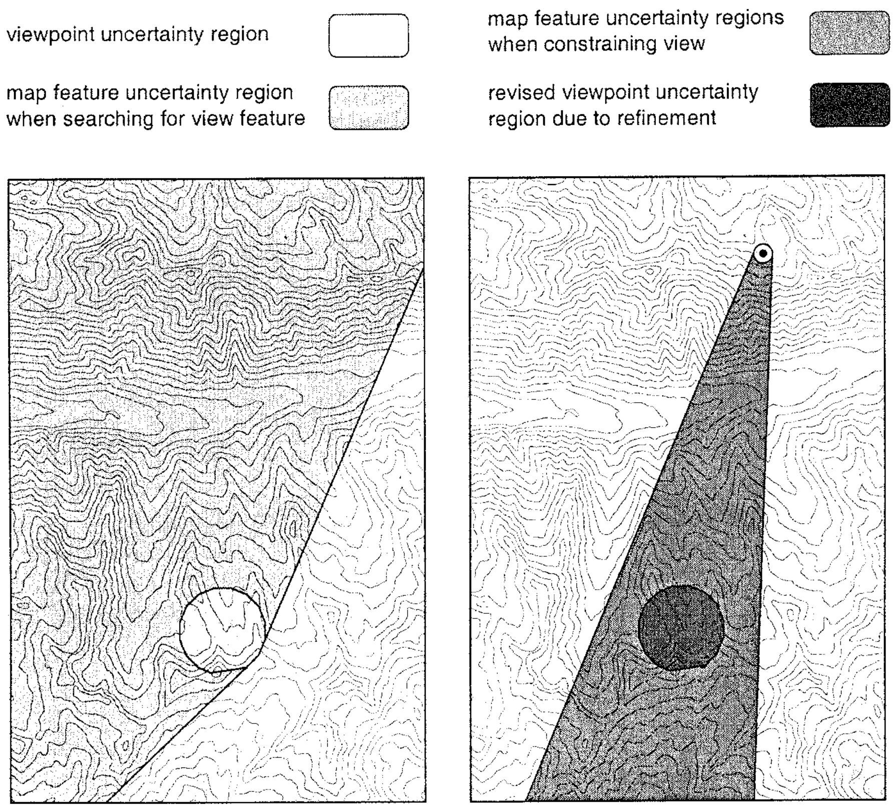

Figure 26. On the left is the search region for a map feature to match the long ridge in the view. On the right is the constraint region based on that match. The circular region in both figures is the initial map location region.

The left panel in Figure 26 shows a viewpoint location and the search region for a map feature to match the long view ridge that wraps from one edge of the image to the other. One map ridge line meeting the description of the view ridge line was found in that region. This map feature and the associated view bearing were then used to refine both view north and the map location region. The map location region and the constraint region associated with the newly matched ridge line are shown on the right. Although the map location region was not reduced by this additional correspondence, it did significantly reduce the view north estimate, evident in the smaller angular extent of the constraint region as compared to that of the search region.

As a result of the newly matched ridgeline, the estimate of view north changed. This affects the size of both search regions and constraint regions, so the refinement based on the previously corresponded features is repeated

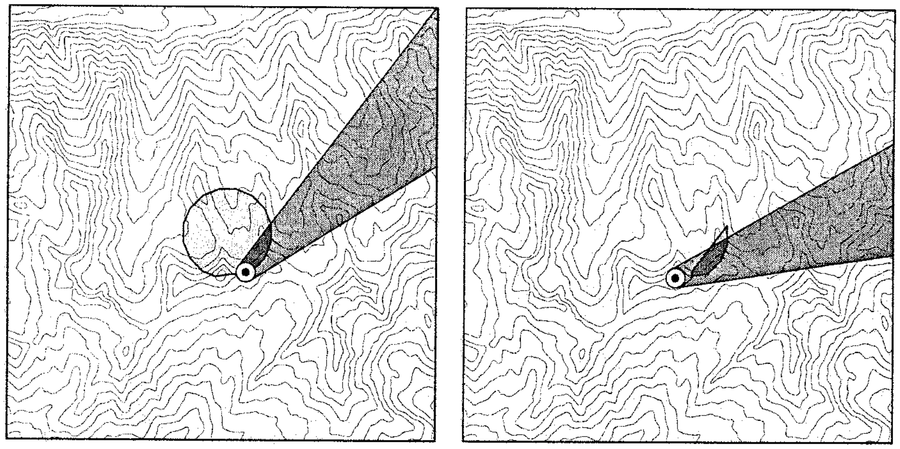

Figure 27. The size of the map location region shown in Figure 26 is reduced by repeating the refinement based on previously corresponded features.

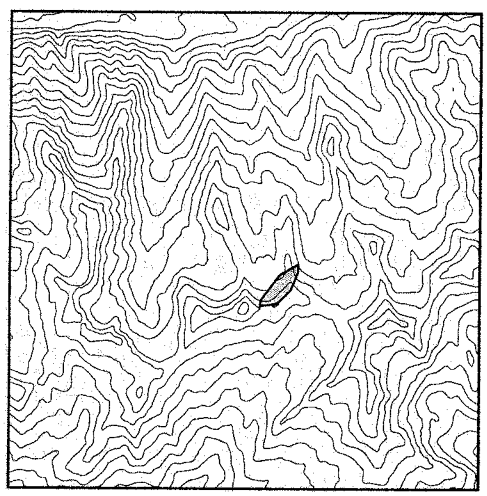

Figure 28. The resulting viewpoint uncertainty region after applying the constraints illustrated in Figure 27.

in an attempt to further reduce the map location region. The peak and nearby ridgeline correspondences, on which the hypothesis was formed, are used in this refinement. The constraint regions associated with each of these feature correspondences are shown in Figure 28. The reduction of 79.1% in the size of the map location region, from 1.486 km^2 to 0.1405 km^2 shows the benefit of using distant features to refine view north followed by nearby features to refine the map location region.

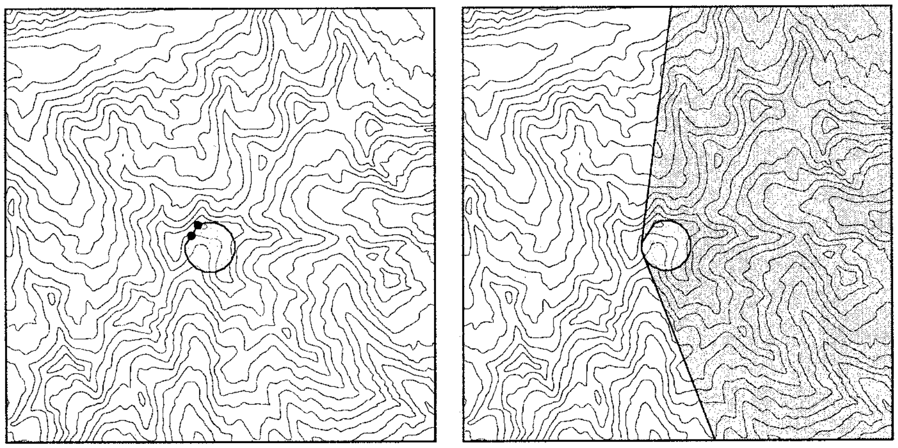

Figure 29. The localization hypothesis with map location region shown on the left is rejected when the search for a map feature to match the long view ridge line within the dark gray area on the right is unsuccessful.

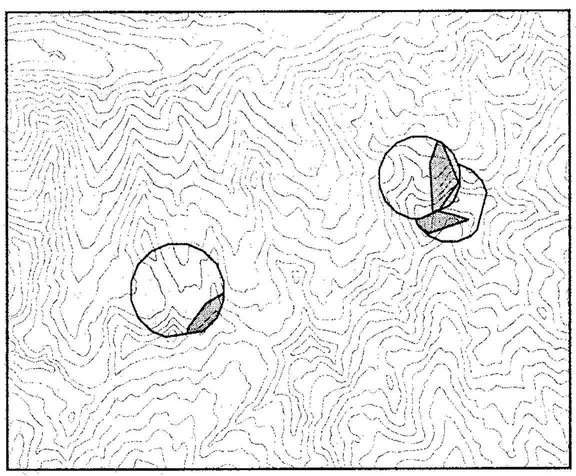

Figure 30. Initial and final location regions for the three non-rejected hypotheses.

Sixteen of the 18 incorrect hypotheses were rejected. One of these had the initial location region shown on the left of Figure 29. After estimating view north, the search for a map ridgeline to match the long view ridge line was unsuccessful. No ridgeline meeting the description of the view ridge was found within the search region, shown on the right of Figure 29. This disconfirming evidence caused the hypothesis to be rejected.

Three hypotheses, including the correct solution, were accepted. Both of the accepted but false hypotheses had incorrect estimates of view orientation. This is a common occurrence among human navigators, who will often look in the wrong direction on the map for a given view feature. The initial and final map location regions for all three non-rejected hypotheses are shown in Figure 30. The set of regions associated with the correct hypothesis is farthest west on the map (with north being "up"). The two incorrect hypotheses were not rejected because an inaccurate estimate of view north led to plausible but incorrect feature matches. This error in estimate, caused by an initial mismatch of features, was very large. Due to the complete lack of *a priori* knowledge of view orientation, the northern direction was assumed to be opposite what it actually was. Indeed, when view north was pre-specified in an analysis of the same data set, only the correct hypothesis was accepted. One goal of future work is to refine the feature correlation process so that disconfirming evidence will force rejection of such an hypothesis, or, preferably, prevent the feature mismatch in the first place. A second goal would be to compare this approach with one which incorporates existing work (Papadias and Egenhofer 1997; Goyal and Egenhofer, in press) on hierarchical reasoning about direction relations and determining cardinal directions between objects into the decision algorithm.

6. Conclusion

We have presented in this paper an approach to feature-based localization in outdoor terrain based on techniques used by experienced human navigators. It implements high level matching strategies which are designed to handle ambiguities such as those due to correspondence errors and gross misidentification of landmarks, and so can be used in situations where mathematical modeling fails. This approach is demonstrated on a real example taken from an outdoor environment lacking man-made features or other clearly distinctive landmarks. Successful application of this type of reasoning in such a problem points out that the exploitation of high level, but often ignored, information can lead to a more robust navigational system.

Our work demonstrates the feasibility of using geometric reasoning strategies which explicitly account for positional uncertainty, combined with high-level problem solving, to perform localization tasks of real-world complexity. Future work needs to address at least three limitations of the current model. While our model captures positional uncertainly quite well, it does not provide an effective mechanism for representing linear and area features with imprecise extent (Burrough 1992; Clementini and Di Felice 1996). The model applies to a single, static viewpoint and doesn't account for either viewpoint motion to confirm hypotheses (Heinrichs et al. 1989) or the viewpoint estimation updating that occurs when an observer moves through outdoor terrain starting from a known location. Finally, while the general principles described here are likely to apply to a wide range of topographies, the detailed analysis of features will require significant extensions in environments different than the high alpine region used in the example.

Acknowledgements

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

Rory Cejka provided many of the support tools used in this effort and assisted in the development of the code for extracting view features. Sean Igo implemented the geometric reasoning subsystem. Lisa Dick developed and implemented the system for automatically extracting terrain features from DEM maps. Thomas Colvin implemented the system for extracting terrain features from digital panoramic images. Thomas Henderson provided much advice, support, and encouragement.

Notes

1 The Navstar Global Positioning System (GPS) provides an alternate method of localization. However, due to frequent occlusion of satellite signals in mountainous terrain, signal distortion due to weather conditions, and other effects (Duerr 1992; Mattos 1992; Cohen et al. 1993; Tranquilla and Al-Rizzo 1993), there is still an important role for perceptually based competence in localization.

2 No single definitive taxonomy of terrain features exists. As a result, we have chosen terms familiar to the experienced map readers that have been involved in our studies of human and machine navigation.

3 The taxonomy in Figure 21 and the compatibilities listed in Figure 22 were derived based on experience with mountainous terrain in Wyoming and Utah. Modifications will likely be required for different types of topography and different geographic areas.

References

Andreas, R.D., Hostetler, L.D. and Beckmann, R.C. (1978). Continuous Kalman Updating of an Inertial Navigation System Using Terrain Measurements. In *Proc. IEEE National Aerospace Electronics Conference* (pp. 1263–1270).

Baird, C.A. and Abramson, M.R. (1984). A Comparison of Several Digital Map-Aided Navigation Techniques. In *Proc. IEEE Position Location and Navigation Symposium* (pp. 294–300).

Bennett, B.H. (1992). A Problem-Solving Approach to the Localization Problem. Ph.D. thesis, University of Minnesota.

Bowditch, N. (1802). *American Practical Navigator*. Blunt.

Brändli, M. (1996). Hierarchical Models for the Definition and Extraction of Terrain Features. In P.A. Burrogh and A.U. Frank (eds.), *Geographic Objects with Indeterminate Boundaries* (pp. 257–270). Taylor & Frank.

Brooks, R.A. (1985). Aspects of Visual Map Making. In *Proc. IEEE International Conference on Robotics and Automation* (pp. 824–829).

Burrough, P.A. (1992). Development of Intelligent Geographic Information Systems, *International Journal of Geographic Information Systems* 6(1): 1–11.

Clementini, E. and DiFelice, P. (1996). An Algegraic Model for Spatial Objects with Indeterminate Boundaries. In P.A. Burrogh and A.U. Frank (eds.), *Geographic Objects with Indeterminate Boundaries* (pp. 155–169). Taylor & Frank.

Cohen, S.C., Chinn, D.S., Torrence, M.H. and Dunn, P.J. (1993). Spaceborne Laser Ranging: Sensitivity to Ephereris and Range Measurement Errors, *Manuscripta Geodaetica*: 72–81.

Cozman, F. and Krotkov, E. (1997). Automatic Mountain Detection and Pose Estimation for Teleoperation of Lunar Rovers. In *Proc. IEEE International Conference on Robotics and Automation* (pp. 2452–2457). Albuquerque, NM.

Duerr, T.E. (1992). Effect of Terrain Masking on GPS Position dilution of Precision, *Navigation* (3): 317–323.

Ernst, M.D. and Flinchbaugh, B.E. (1989). Image/Map Correspondence Using Curve Matching. In *AAAI Symposium on Robot Navigation* (pp. 15–18).

Goyal, R.P. and Egenhofer, M.J. (in press). Cardinal Directions between Extended Spatial Objects. In *IEEE Transactions on Knowledge and Data Engineering*.

Grimson, W.E.L. (1990). *Object Recognition by Computer: The Role of Geometric Constraints*. The MIT Press.

Heinrichs, M.R., Montello, D.R., Nusslé, C.M. and Smith, K. (1989). Localization with Topographic Maps. In *Proceedings of the AAAI Symposium on Robot Navigation* (pp. 29–32).

Heinrichs, M.R., Smith, K., Pick, Jr. H.L., Bennett, B. and Thompson, W. (1992). Strategies for Localization. In *Proc. DARPA Image Understanding Workshop* (pp. 321–325).

Huttenlocher, D.P. and Ullman, S. (1987). Object Recognition Using Alignment. In *Proc. First International Conference on Computer Vision* (pp. 102–111).

Kuipers, B. and Levitt, T.S. (1988). Navigation and Mapping in Large-Scale Space, *AI Magazine* 9(2): 25–43.

Levitt, T.S. and Lawton, D.T. (1990). Qualitative Navigation for Mobile Robots, *Artificial Intelligence* 44(3): 305–4360.

Mattos, P. (1992). GPS, *Electronics and Wireless World*: 982–987.

Moravec, H. (1988). Sensor Fusion in Certainty Grids for Mobile Robots, *AI Magazine* 9(2): 61–74.

Papadias, D. and Egenhofer, M.J. (1997). Algorithms for Hierarchical Spatial Reasoning, *GeoInformatica* 1(3): 251–273.

Pick, Jr., H.L., Heinrichs, M.R., Montello, D.R., Smith, K., Sullivan, C. and Thompson, W. (1996). Topographic Map Reading. In J. Flach, P. Hancock, J. Caird and K. Vicente (eds.), *Ecology of Human-Machine Systems*. Lawrence Erlbaum Associates.

Sanso, F. (1973). An Exact Solution of the Roto-Translation Problem, *Photogrammetria* 29: 203–216.

Schlieder, C. (1995). Reasoning About Ordering. In A.U. Frank and W. Kuhn (eds.), *Spatial Information Theory: A Theoretical Basis for GIS* (pp. 341–349). Springer.

Stein, F. and Medioni, G. (1992). Map-Based Localization using the Panoramic Horizon. In *Proc. IEEE International Conference on Robotics and Automation*. Nice, France.

Sutherland, K.T. (1994). The Stability of Geometric Inference in Location Determination. *Technical Report* UUCS-94-021, University of Utah.

Sutherland, K.T. and Thompson, W.B. (1994). Localizing in Unstructured Environments: Dealing with the Errors. In *IEEE Trans. on Robotics and Automation* (pp. 740–754).

Talluri, R. and Aggarwal, J.K. (1992). Position Estimation for an Autonomous Mobile Robot in an Outdoor Environment. In *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 8 (pp. 573–584).

Thompson, E.H. (1958). An Exact Linear Solution of the Problem of Absolute Orientation, *Photogrammetria* 13(4): 163–178.

Tranqulla, J.M. and Al-Rizzo, H.M. (1993). Investigation of GPS Precise Relative Static Positioning During Periods of Ice Clouds and Snowfall Precipitation. In *IEEE Trans. Geoscience and Remote Sensing* (pp. 295–299).

Valiquette, C.M. (1995). Localization in Outdoor Static Terrain. Master's thesis, University of Utah.

Yacoob, Y. and Davis, L.S. (1991). Computational Ground and Airborne Localization Over Rough Terrain. Technical Report CAR-TR-591, Computer Vision Laboratory, University of Maryland.