Image/Map Correspondence Using Curve Matching[†]

Michael D. Ernst and Bruce E. Flinchbaugh

Texas Instruments
Computer Science Center
P.O. Box 655474, M/S 238, Dallas, Texas 75265

[†]**(Presented at: AAAI Robot Navigation Symposium, 28-30 March 1989)**

Abstract

We address the general practical problem of determining correspondences between maps and terrain images, and focus on a static low altitude airborne scenario. For this case we consider the approach of partially matching detected and expected curves in the image plane. Expected curves are generated from a map, using an estimate of the sensor pose in three dimensions, and matched with simulated detected curves in an image. We also outline a method for sensor pose refinement using point correspondences derived from curve matches as input to a relative orientation algorithm.

Image/Map Correspondence

We address the problem of determining correspondences between maps and terrain images in low altitude airborne scenarios. In particular, we consider an aircraft with a wide-angle passive imaging sensor, such as an infrared or TV camera, flying a few hundred feet above tree tops.

Other on-board systems are assumed to provide an estimate of the three-dimensional orientation of the imaging sensor. We also assume availability of an estimate of three-dimensional position, for which there are several sources of information in varying accuracies (e.g. altimeters, inertial navigation, and global positioning systems). The important point is that regardless of how accurate the sensor pose estimate is in principle, it will be inaccurate in detail in practice. Hence we seek effective means of determining correspondences between images and maps, to provide enhanced relative direction and absolute position information.

Finally, we assume availability of a variety of maps for use by the system. The maps must be complete enough to permit prediction of visually prominent aspects of the environment. For this a terrain elevation map is critical but not sufficient. Other maps outlining the locations of areas such as forests, fields, water, roads, and buildings are also needed.

Solutions to the image/map correspondence problem are important for visual support of navigation, and critical in guiding the search for objects in areas anticipated from map information. Image/map correspondence is also important in low altitude airborne surveillance and reconnaissance problems, in which the objective is to refine existing map information on the basis of incoming imagery.

Examples of other work directly addressing various formulations of passive image/map correspondence problems are: (Bolles *et al.* 1979), (Medioni & Nevatia 1984), (Mokhtarian & Mackworth 1986), (McKeown 1987), (Lawton *et al.* 1987), and (Andress & Kak 1988). Our overall approach is conceptually most similar to the approach of Mokhtarian & Mackworth (1986). In contrast to these prior efforts, we emphasize the low altitude formulation of the problem and the roles of partial curve matching and expected curve generation.

The sensor position considered in this paper, at low altitudes above tree tops, poses significantly different problems and considerations than those encountered in land vehicle vision systems. For example, from viewpoints on the ground, forests are significant obstacles to vision and can easily obscure vast amounts of information such as the contour of the horizon. In this case a more map-like scale of visual information is available from the air for comparison with maps because only areas in the immediate vicinity of trees are occluded.

The assumed low altitude sensor position also poses significantly different problems than encountered in high altitude aerial and satellite imagery. In the low altitude case, useful viewing directions tend more toward the horizon than straight down, forcing a three-dimensional formulation of the solution, as opposed to the largely two-dimensional approaches that are viable at high altitudes.

Of course the substantial body of work in model-based object recognition, e.g. as surveyed by Besl & Jain (1985), is also relevant in principle. In fact, the curve matching algorithms we draw upon (Schwartz & Sharir 1987; Wolfson 1987) were motivated primarily by problems with recognizing objects in industrial applications . However, low altitude image/map correspondence problems remain substantially different in almost every detail, as terrain environments are typified by complex irregular geometries, materials, and lighting conditions, and maps are typically sampled at a lower resolution than the visually available information. Partial matching of arbitrary curves is particularly relevant in this context.

Curve Matching Approach

Our curve matching approach is to match curves detected in an image with expected image curves generated from a map. This requires a plausible balance of high speed image processing and high speed computer graphics, while avoiding problems with other matching approaches. For example, matching detected and expected image intensities directly, begs generation of realistic image intensities from maps – realism that is beyond the scope of typical maps. Matching detected and expected three-dimensional surfaces, has potential merit for the passive sensors we consider, but places the heavy burden of solution on robust shape-from-x techniques.

The first three steps of the method, elaborated below, obtain relative image/map correspondence. We have also incorporated the fourth and fifth steps for absolute sensor pose refinement experiments.

Step (1): Detect and extract curves from an image.

Step (2): Project expected curves from a map into the image plane, using the estimated sensor pose (three-dimensional position and orientation).

Step (3): Match detected and expected curves to obtain correspondences between the image and the map.

Step (4): Extract point correspondences at distinctive points along the matched curves.

Step (5): Correct the estimated sensor pose using point correspondences from Step (4) as input to a relative orientation algorithm (Horn 1987) .

Detecting and Extracting Curves

We consider two approaches for the curve detection step, and explain why an approach that seeks only 'visually obvious' curves could be sufficient. Thus far we have simulated the detection of curves from images, by using maps and the expected curve generation technique described below.

One approach for the required curve detection is to attempt to estimate curves of the terrain image by classes such as 'roads', 'water boundaries', and 'vegetation boundaries'. This approach places the primary responsibility for success on devising robust 'map-less' recognition algorithms. Although this approach is potentially feasible, it requires techniques that are still largely beyond the state of the art.

An alternative approach is to use a generic edge/curve extraction algorithm to extract 'obvious' curves without attempting recognition. The expected curve generation component of the system then has a more interesting responsibility of estimating what the obvious curves will be considering the imaging conditions. This approach has considerable potential because graphical generation of curves for expected high contrast image boundaries is more feasible than predicting realistic image intensities *per se*.

Generating Expected Curves

We accomplished expected curve generation for the most part using using straightforward techniques from computer graphics. Map curves were overlaid on a piecewise planar interpolation of the digital elevation map, and projected via a perspective transformation into an image plane using the given estimate of sensor position and orientation. The algorithm we developed maintains contiguous curve representations through the transformation, splitting curves at occluding contours, and removing hidden curve segments. Thus this step obtains a set of idealistic two-dimensional curves in image coordinates, representing the geometrically expected visual locations of the corresponding map information.

Obviously it is not realistic to assume that all of the expected curves will be visible in the image. Although there are many potential criteria to consider for improving expected curve generation, we have implemented this simple approach first. Also, it is not yet clear how much realism is required. As an example, roads roughly parallel to the line of sight tend to be much more visible than roads perpendicular to the line of sight. Expected curve generation criteria could be devised to model such phenomena if the presence of expected, yet invisible, curves turns out to be a major problem. However, in principle the extra expected curves should not matter. They will simply be left unmatched; or only the detected portions will be partially matched, just as a more realistic set of expected curves would be matched.

Matching Curves

The two-dimensional curve matching algorithm we implemented for this step draws largely from algorithms by Schwartz & Sharir (1987) and Wolfson (1987). We selected these algorithms because they are well-suited for arbitrary curves, and because the partial matching component of Wolfson's algorithm shows promise for robust matching of curves derived from low altitude terrain images where occlusions and illumination changes can easily cause fragmented curves, in comparison with the ideal curves generated from maps.

Beyond the prior algorithms, we added a curve connection option and a technique for interpolating matches after the best partial matches have been determined. We also introduced several parameters for the algorithms so that knowledge of constraints on matching can be used to control allowable matches. There are four steps in the curve matching scheme:

Step (3.1) Connecting and Smoothing Curves: We implemented a curve connection criterion to fill gaps between detected curves, using typical separation, orientation, and colinearity tolerances. Two curves with approximately coincident endpoints within the tolerances are connected. Note that the ordinarily unconservative option of using only a separation tolerance, barring intersecting curves, is justified because the partial curve matching that follows should not match across incorrectly connected curves, while correct connections should promote longer matches.

The curve smoothing algorithm we implemented is an algorithm suggested by Schwartz & Sharir (1987). This finds the shortest path within an epsilon neighborhood of the curve. Our use of curve smoothing here is heuristic because the different three-dimensional transformations underlying the detected and expected two-dimensional curve sets do not necessarily satisfy the conditions of a lemma (Schwartz & Sharir 1987) justifying the smoothing operation for other matching problems.

Step (3.2) Finding Plausible Partial Matches: In this step a pool of plausible matches, including partial matches, is constructed. We use an adaptation of the "shape signature string" matching algorithm (Wolfson 1987) which accomplishes partial matching by comparing approximations of curvature as a function of curve length. The resulting match pool contains a large number of pairs of partially matching curves and subcurves. The point is to quickly create a relatively large number of promising matches, discarding those which are obviously wrong, rather than to closely discriminate between correct and incorrect matches.

Step (3.3) Selecting Best Matches: The purpose of this step is to select the best matches from the pool of candidate partial matches. We use a process of elimination, rejecting matches that do not satisfy criteria as follows.

A match is excluded if its translation (the amount by which one curve must be shifted to match the other), rotation, or length (the length of either of the matched curves) are not within specified ranges. These values are computed using the fast curve matching technique of Schwartz & Sharir (1987). The allowed ranges are parameters to allow enforcement of any available knowledge of constraints on the matching.

Another criterion for reducing the population of the match pool requires various measures of global coherence among the matches in two dimensions. For example, it may be known that there is a lateral error in the generated expected curves, although the magnitude of the error is unknown. For this we use histogram pruning on the translation magnitudes, retaining matches near the peak; rotational coherence is also handled this way. However, these particular measures are only reasonable when the underlying three-dimensional transformations differ by a translation perpendicular to the line of sight, and/or by a rotation about the line of sight.

Finally a uniqueness criterion is applied. For this we use a greedy algorithm which adds the best match to the output and removes from the match pool all matches that overlap it. This process is repeated until no matches remain in the pool.

Step (3.4) Interpolating Matches: To complete the curve matching scheme, the set of final partial matches is examined to find pairs of curves between which two or more curve segments have been partially matched. A new partial match is then 'interpolated' between connected partial matches.

This postprocessing of the partial curve matches is intended to help bridge gaps that are natural consequences of map and image resolution, map inaccuracies, and simplifying assumptions made in curve detection. There are many additional strongly model-based ways of proceeding from this point that we have not pursued.

Observations

Image/map correspondence problems and applications are enormous tasks in general, especially considering the processing requirements if correspondences are to be computed in flight. Although our purpose here has been to discuss the basic viability of partial curve matching as a component of such systems, we also believe the approach satisfies real-time implementation requirements because it balances existing and near future capabilities in image processing and computer graphics. The operations posited for curve detection from images are within the capabilities of image processing systems, and there are many computer graphics systems to support the kinds of high-speed expected curve generation operations required in this approach.

We have informally tested the curve matching method for image/map correspondence using digital map data including elevation, vegetation, and roads for a 1920 square meter area. Our initial experiments have focussed on the performance of the curve matching algorithms in the presence of differences between simulated actual and expected views. We believe the curve matching approach has considerable merit, but it is too early commit to particular approaches. Considering requirements for robust operation in the presence of low resolution images and maps, it seems likely that ultimate solutions will draw upon several diverse approaches in concert. We note a basic conflict between the problem of partial matching and the problem of matching curves of different sizes and distortions due to range differences between detected and expected views. Perhaps a combination of partial matching (Wolfson 1987) and scale space matching (Mokhtarian & Mackworth 1986) is viable. Affine invariant partial curve matching (Lamdan, Schwartz, & Wolfson 1988) is also relevant.

The primary challenge for low altitude image/map correspondence is to prove methods that will work using maps and terrain images with ground truth sensor pose information. A next step for the curve matching approach is to devise curve detection criteria coordinated with map-based expected curve generation, and to evaluate using databases of low altitude images of terrain in mapped regions.

Acknowledgments

We thank Marion Lineberry, Vishal Markandey, Joyce Schroeder, and John Wright for various contributions in this work. We also thank Berthold K.P. Horn (MIT) for providing relative orientation code.

References

[1] K. M. Andress and A. C. Kak. Evidence accumulation and flow of control in a hierarchical spatial reasoning system. *AI Magazine*, 9(2):75–94, Summer 1988.

[2] P. Besl and R. Jain. Three-dimensional object recognition. *ACM Computing Surveys*, 17(1):75–145, March 1985.

[3] R. C. Bolles, L. H. Quam, M. A. Fischler, and H. C. Wolf. Automatic determination of image-to-database correspondences. In *Proc. Sixth IJCAI*, pages 73–78, August 1979.

[4] B. K. P. Horn. Relative orientation. AI Memo 994, Massachusetts Institute of Technology AI Lab, 545 Technology Square, Cambridge, MA, September 1987.

[5] Y. Lamdan, J. T. Schwartz, and H. J. Wolfson. Object recognition by affine invariant matching. In *Proc. of Computer Vision and Pattern Recognition*, pages 335–344, June 5-9 1988.

[6] D. Lawton, T. S. Levitt, C. C. McConnell, P. C. Nelson, and J. Glicksman. Environmental modeling and recognition for an autonomous land vehicle. In *Proceedings: Image Understanding Workshop*, pages 107–121. DARPA, February 1987.

[7] D. M. McKeown. The role of artificial intelligence in the integration of remotely sensed data with geographic information systems. *IEEE Transactions on Geoscience and Remote Sensing*, 25(3):330–348, May 1987.

[8] G. Medioni and R. Nevatia. Matching images using linear features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):675–685, November 1984.

[9] F. Mokhtarian and A. Mackworth. Scale-based description and recognition of planar curves and two-dimensional shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):34–43, January 1986.

[10] J. T. Schwartz and M. Sharir. Identification of partially obscured objects in two and three dimensions by matching noisy characteristic curves. *International Journal of Robotics Research*, 6(2):29–44, Summer 1987.

[11] H. Wolfson. On curve matching. In *Proc. of the IEEE Comp. Soc. Workshop on Computer Vision*, pages 307–310, Nov. 30-Dec. 2 1987.