

*Computer Science
Technical Report*

**Colorado
State**
University

Local Search as a Tool for Horizon Line Matching*

J. Ross Beveridge, Christopher Graves and Christopher E. Leshner

December 16, 1995

Technical Report CS-96-109

Computer Science Department
Colorado State University
Fort Collins, CO 80523-1873

Phone: (970) 491-5792 Fax: (970) 491-2466
WWW: <http://www.cs.colostate.edu>

*This work was sponsored by the Advanced Research Projects Agency (ARPA) under grant DAAH04-93-G-422, monitored by the U. S. Army Research Office.

Local Search as a Tool for Horizon Line Matching^{*†}

J. Ross Beveridge
Colorado State University
ross@cs.colostate.edu

Christopher Graves
Colorado State University
gravescc@cs.colostate.edu

Christopher E. Lesher
Colorado State University
lesher@cs.colostate.edu

Abstract

A robust and general local search matching algorithm is used to match fragmentary horizon lines. A horizon line model is extracted from a rendered terrain map and is then matched to features extracted from CCD imagery. Such matching is one means of automating vehicle orientation correction: a problem of practical significance for the UGV program. Currently, final orientation correction relating vehicle coordinates to terrain map coordinates must be performed manually. Results are presented for actual terrain and image data collected at the UGV Demo C test site.

1 Introduction

When a vehicle navigating with GPS and inertial guidance stops, small errors in pointing angle lead to large errors in pixel registration between imagery and stored terrain maps. For the Semi-Autonomous Vehicles (SSVs) developed for the Unmanned Ground Vehicle (UGV) Program, orientation estimates can be off by 1 or more degrees [Ray95b]. The resulting uncertainty precludes terrain guided visual search and target recognition. Thus, a human must hand select registration features before these activities are carried out.

Here, local search matching is tested as a tool for automating registration. Local search matching refers to a body of algorithms we have developed which find, with high probability, the optimal correspondence mapping and geometric transformation between a model and image data [BWR90, Bev93, BHP95]. To match terrain and image data, we anticipate a four step process:

- Render 3D terrain using the estimated vehicle pose.
- Extract matchable features from the rendered terrain and actual imagery.

This work was sponsored by the Advanced Research Projects Agency (ARPA) under grant DAAH04-93-G-422, monitored by the U. S. Army Research Office.

[†]Appears also in the Proceedings of the 1996 ARPA Image Understanding Workshop.

- Optimally match the two sets of features.
- Use matched features in place of hand selected control points to correct the orientation estimate of the vehicle.

Results for steps 1 through 3 are presented on actual data collected at the UGV Demo C test site using the SSV's CCD camera. Step 4, it is assumed, will succeed if matches are comparable to those selected by a person. This assumption is a fair starting point, but future work should address sensitivity of orientation to errors in image space matching.

2 Related Work

In some early terrain feature matching, Clark [CCE⁺80] matched line segments representing dominant image features. Levitt [LLCN87] proposed a way to select salient landmarks from terrain data [LLCN87] for navigation. Stein [Fri92] uses panoramic horizon curve matching for vehicle localization. Thompson and Sutherland [THC⁺93, Tho93, Sut94] have built a sophisticated expert system with a domain specific image feature extraction algorithm for abstracting structural terrain descriptions.

While most of these works address the general problem of vehicle localization anywhere on a map, this paper considers the much simpler problem of orientation correction. Our work is distinguished by: 1) reliance upon the known position to reduce the problem to 2D matching, 2) use of a robust and general matching technique to overcome feature fragmentation, and 3) use of narrow FOV imagery.

3 Step 1: Terrain Rendering

The 5m digital elevation map (DEM) for the Demo C test site was obtained from Lockheed-Martin. This site was selected because test imagery taken directly from the SSV is available along with ground truth indicating vehicle position and pointing angle relative to fixed targets [Ray95a].

A terrain rendering system has been developed using Open-GL which simulates the FOV of the CCD sensor

used on the SSV. A simple lighting model is used and terrain is rendered from positions at which the vehicle actually acquired imagery. The vehicle pointing angle is derived from recorded vehicle and target positions: the targets are other military vehicles. Because target ground truth is being used to derive pointing angles, only images with targets near the image center are used. Figures 1a and 1b show two rendered terrain images for which matching is tested below.

4 Step 2: Extracting Features

The local search algorithm matches one set of line segments to another, and for this problem the model and data segments are extracted from the rendered and actual images respectively. We use our own implementation of the Burns algorithm [BHR86]¹. High frequency texture in these scenes prevents horizon features from being extracted unless the imagery is first smoothed: a 7x7 smoothing kernel has been used here. Even with smoothing, the horizons are still sometimes difficult to extract, and significant fragmentation occurs. Figures 1c and 1d show the images themselves along with the segments extracted by the Burns algorithm.

5 Step 3: Matching

5.1 Review and Overview

A complete explanation of local search matching appears in Beveridge's dissertation [Bev93] and 3D matching results appear in [BR95]. A controlled performance analysis of 2D matching appears in [BRG95]. To briefly review the approach, an iterative generate-and-test strategy moves from a randomly selected initial match to one that is locally optimal. A global least-squares fitting process *always* aligns model and data for *any* correspondence tested. Thus, global geometry implicitly directs search. A match error takes account both of spatial fit and omission: how much model is un-matched.

Search is conducted over a space of correspondence mappings C : C is the powerset of possibly matching features S . Most other algorithms consider one-to-many matches [Gri90] while our C includes many-to-many matches. Without many-to-many mappings, properly matching piecewise approximations to curves with non-coincident breakpoints is impossible. This point is important here because horizon lines involve such non-coincident breakpoints.

While at first the initialization of search from *randomly* chosen matches may seem foolish, it is a strength of the approach. By running multiple trials from independently chosen initial matches, the probability of seeing the best (or near best) at least once may be made arbitrarily high. Past experience has demonstrated 100 trials is adequate to solve most difficult prob-

lems [BR95, BRG95]. Another benefit of multiple trials is the structure and frequency of alternative solutions tells us much about the difficulty of a particular problem.

5.2 Matching Results

Results for two pairs of terrain and image features are presented. Two constraints limit the space of possible matches between horizon model and image line segments. The first assumes that the horizon lies somewhere in a band $1/2$ the height of the image centered about the true position. The second assumes the relative orientation between segments must be less than 17 degrees for them to match. Even with these constraints, the sets of potentially matching features were very large: 1183 for image 1 and 1577 for image 2. The resulting search spaces C contain 2^{1183} and 2^{1577} states respectively.

To explore the space of possible matches, 500 trials of subset-convergent local search were run on each problem. The best match found in each case is shown superimposed in black in Figures 1e and 1f. In both cases, visual inspection shows these to be essentially correct matches. For images 1 and 2, the best matches were found in a single trial with probabilities 0.056 and 0.036 respectively. Based upon this probability, it follows this match may be found with better than 95% confidence running 59 and 90 trials respectively. Being conservative, 100 trials is more than sufficient.


These problems are some of the largest, in terms of search space, for which local search matching has been tested. To find these matches reliably, the current C implementation running on a Sparc 20 requires on the order of 20 minutes for image 1 and an hour for image 2. Clearly, either some domain specific tuning or use of parallel hardware is required to bring run-times down. Both of these are very reasonable options for future work. Use of a better feature extraction algorithm would dramatically simplify the combinatorics and parallel local search is trivial due to the independence of trials.

One general problem with horizon line matching is starting to be evident in Image 2: ambiguous horizon structure leads to ambiguous matches. While the best match shown here is correct, the model for Image 2 is essentially a radial curve. Consequently, it can match the image data at multiple points. Slight changes in the parameterization of the match error uncovered this sensitivity, resulting in shifts and scalings of the true horizon.

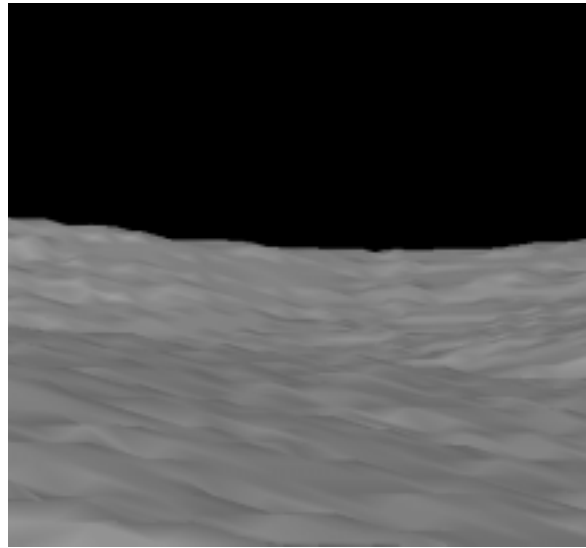
References

- [Bev93] J. Ross Beveridge. *Local Search Algorithms for Geometric Object Recognition: Optimal Correspondence and Pose*. PhD thesis, University of Massachusetts at Amherst, May 1993.
- [BHP95] J. Ross Beveridge, Allen Hanson, and Durga Panda. Model-based fusion of flir, color and

¹This version has a simple single Glyph interface and is publicly available from our ftp site: `ftp.cs.colostate.edu, /pub/vision`




a. Image 1: rendered terrain



b. Image 2: rendered terrain



c. Image 1: extracted features



d. Image 2: extracted features



e. Image 1: horizon match



f. Image 2: horizon match

Figure 1: Results of local search matching on 2 horizon images

- ladar. In Paul S. Schenker and Gerard T. McKeen, editors, *Proceedings: Sensor Fusion and Networked Robotics VIII*, Proc. SPIE 2589, pages 2 – 11, October 1995.
- [BHR86] J. B. Burns, A. R. Hanson, and E. M. Riseman. Extracting straight lines. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-8(4):425 – 456, July 1986.
- [BR95] J. Ross Beveridge and Edward M. Riseman. Optimal Geometric Model Matching Under Full 3D Perspective. *Computer Vision and Image Understanding*, 61(3):351 – 364, 1995. (short version in IEEE Second CAD-Based Vision Workshop).
- [BRG95] J. Ross Beveridge, Edward M. Riseman, and Christopher Graves. Demonstrating polynomial run-time growth for local search matching. In *Proceedings: International Symposium on Computer Vision*, pages 533 – 538, Coral Gables, Florida, November 1995. IEEE PAMI TC, IEEE Computer Society Press.
- [BWR90] J. Ross Beveridge, Rich Weiss, and Edward M. Riseman. Combinatorial Optimization Applied to Variable Scale 2D Model Matching. In *Proceedings of the IEEE International Conference on Pattern Recognition 1990, Atlantic City*, pages 18 – 23. IEEE, June 1990.
- [CCE⁺80] C.S. Clark, D.K. Conti, W.O. Eckhardt, T.A. McCulloh, R. Nevatia, and D.Y. Tseng. Matching of Natural Terrain Scenes. In *ICPR80*, pages 217–222, 1980.
- [Fri92] Fridtjof Stein and Gerard Medioni. Map-based Localization using the Panoramic Horizon. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pages 2631 – 2637, Nice, France, May 1992.
- [Gri90] W. Eric L. Grimson. *Object Recognition by Computer: The Role of Geometric Constraints*. MIT Press, Cambridge, MA, 1990.
- [LLCN87] T.S. Levitt, D.T. Lawton, D.M. Chelberg, and P.C. Nelson. Qualitative landmark-based path planning and following. In *AAAI-87*, pages 689–694, 1987.
- [Ray95a] Ray Rimey. RSTA Sept94 Data Collection Final Report. Technical report, Martin Marietta Astronautics, Denver, CO, January 1995.
- [Ray95b] Ray Rimey and Darrell Hougen. Discussion of SSV Orientation Correction with Lockheed-Martin RSTA Group. Personal Correspondence, 1995.
- [Sut94] Sutherland, K.T. and Thompson, W.B. Localizing in Unstructured Environments: Dealing with the Errors. *Robotics and Automation*, 10:740–754, 1994.
- [THC⁺93] W.B. Thompson, T.C. Henderson, T.L. Colvin, L.B. Dick, and C.M. Valiquette. Vision-Based Localization. In *Proceedings: Image Understanding Workshop*, pages 491–498, Los Altos, CA, 1993. ARPA, Morgan Kaufmann.
- [Tho93] Thompson, W.B. and Pick, Jr., H.L. Vision-Based Navigation. In *Proceedings: Image Understanding Workshop*, pages 127–131, Los Altos, CA, 1993. ARPA, Morgan Kaufmann.