EXTRACTION OF MICRO-TERRAIN FEATURES

William B. Thompson
University of Utah

Gregory W. Thoenen
University of Utah

Ronald G. Moore
Evans & Sutherland

Thomas C. Henderson
University of Utah

Abstract

Micro-terrain features are topographic structures relevant to the behavior of a simulation but with a horizontal extent significantly less than the resolution of the base-level terrain data covering the area in which these structures occur. These features are thus not directly extractable from elevation data. While they are often apparent in aerial imagery, micro-terrain features are easily missed or confused with other features, making reliable detection based on imagery alone problematic. This paper describes an automated method for extracting high fidelity models of small-scale ravine features by augmenting a hydrological analysis with computer vision techniques.

Introduction

Sensor technology, limitations of photogrammetry, storage constraints, and requirements for real-time rendering all limit the fidelity with which terrain can be effectively represented in a geospatial database. For certain applications, it is critical that these databases include specific *micro-terrain* features with a lateral extent less than the resolution of base-level terrain description. Ravines and other, similar drainage features are examples of topographic micro-terrain that can critically affect the realism of simulations.

While ravines and dry washes are usually at least partially visible on aerial photographs, accurate detection and localization is difficult. In addition, ravines are easily confused with roads, tracks, and other structures commonly appearing in non-urban environments. Photogrammetry fails to extract many ravine features because of their restricted depth and small width relative to the resolution with which terrain elevation is extracted. The photogrammetric overlaps usually used for non-urban terrain preclude the ability to see into many ravine bottoms or measure the slopes of their sides. As a result, adding such features to terrain databases currently requires substantial manual processing.

The dry wash shown in Figure 1 is located in the live fire range (Range 400) of the USMC Air Ground Combat Center, located at Twentynine Palms, CA. The wash is 1m-2m deep and 2m-3m across. Though *very* small compared to the resolution at which terrain features are usually modeled, such ravines are of critical tactical significance to both dismounted infantry and trafficability, and therefore all units in a combat force. This area is included within the coverage of the Army's Close Combat Tactical Trainer Southwest United States Desert Database (CCTT Primary Two) (Pope et al., 1995). As a result, it provides an effective opportunity to investigate the utility of methods such as presented here compared to current practice.

Figure 1: Shallow wash located within Range 400.

Standard NIMA data product are available for this area, including DTED level 1 and 2 elevation data and ITD feature data. DTED, and to a lesser degree ITD, was never intended to support high resolution simlulation requirements. In practice, however, these and similar resolution data sources will continue to form the basis for terrain skin construction for some time to come. In addition, a 3.3km by 2.2km area is covered by a custom-produced high-quality DEM with a 1m post spacing and a relative vertical accuracy on the order of 0.1 to 0.3 meters with a matching orthoimage at the same resolution (Richbourg et al., 1995; Richbourg and Olson, 1996). This provides an indication of the true topography, which we can compare with estimates extracted from more commonly available data sources.

Approach

The nominal resolution of elevation data on which terrain models are based is commonly on the order of 30m or greater. Due to the smoothing inherent in the manner in which the elevation data is obtained, the effective res-

Presented at the **1998 IMAGE Conference**
Scottsdale, Arizona. 2–7 Auguest 1998.

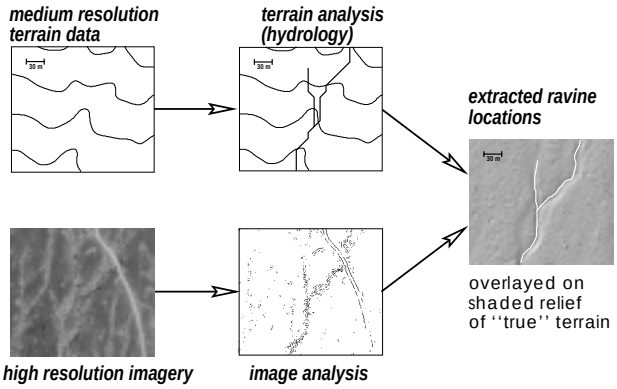

Figure 2: Combining terrain analysis and computer vision for ravine extraction.

olution, measured in terms of the size of distinct features apparent in the data, is much coarser. Thus, even with DEM data finer than a 30m grid, terrain structure such as shown in Figure 1 is likely to go unrepresented. Nevertheless, coarse resolution DEMs can be used to predict likely locations where smaller terrain deformations are to be expected.

Ravines are erosional features generated by large-scale processes, even if the final effect is visible mostly on a fine-scale. As a result, hydrological analysis can be used to predict where such erosion is most likely to occur. Effective algorithms for performing this analysis even on low-resolution, error-full representations of the terrain skin are now a standard function in many geographic information systems (GIS). Such operations can be used to effectively estimate the existence and location of ravine features using DEM data with a post spacing significantly greater than the width of the features of interest.

Hydrological analysis alone is not sufficient to confirm that a ravine actually exists or to accurately determine its location. Our approach therefore uses computer vision techniques applied to higher resolution aerial images to refine the results of the hydrological processing. Figure 2 illustrates the method. The contour maps show elevations in a small patch of terrain, determined based on DTED level 2 data with a post spacing of approximately 30m. Hydrologic features are overlayed on the right-most map. The lower left shows a 1m resolution orthoimage of the same area. To its right is the results of a standard edge detector applied to this image. Neither the hydrological drainage features nor the edge features are enough in isolation to accurately located the ravines. Taken together, however, accurate localization is possible, as shown on the far right of the figure. Here, the results of automated ravine extraction are overlayed on a shaded relief rendering of the area, derived from the high-resolution DEM. Essentially, this indicates the "ground truth" topography.

Background

Hydrological Analysis for Terrain Analysis

The earliest ideas on using DEM data to find ravines were based upon using local surface properties to look for a part of the topographic surface that is locally concave-upward, and mark this position as a valley or ravine, presuming that it is where surface water runoff is likely to be concentrated (e.g., (Peucker and Douglas, 1975; Chorowicz et al., 1989; Tribe, 1992)). Many researchers (e.g., (Mark, 1983; Jenson and Domingue, 1988)) have used a method that is more physically justifiable in nature. In this method, a direction is assigned to each cell of the DEM, corresponding to the direction that water would flow out of that cell. This direction is that of steepest decent (i.e. one of the 8 compass directions that corresponds to the steepest downhill slope from that cell). Given this "direction matrix", the total number of cells of the DEM that contribute drainage through each cell is calculated. Those cells that accumulate drainage above a certain threshold are considered part of the drainage network.

Active Contours for Image Analysis

Active contours (often called *snakes*) are a powerful tool for finding long curving linear structures in images (Kass et al., 1988). Snakes are introduced as energy-minimizing splines, whose energy is a weighted sum of internal and external energies. There are two different internal energies which may be weighted in order to force a snake to act more like a membrane or string, in the sense of it resisting stretching, or more like a thin-plate or rod, in that it resists bending (Leymarie and Levine, 1993). The external energy equation is a function of the image on which it is acting. This equation can be specified to favor various image properties, such as edges and lines. Snakes bring to bear a high-level, global knowledge across the entire curve, instead of relying solely on local, low-level knowledge(Kass et al., 1988; Menet et al., 1990).

In recent years, much research has been directed at various aspects of snakes, including initialization (e.g., (Berger and Mohr, 1990; Neuenschwander et al., 1994)), different underlying representation of snakes (e.g., (Menet et al., 1990; Wang et al., 1996)), formulation of energy functions (e.g., (Radeva et al., 1995; Lai and Chin, 1993)), imposing constraints (e.g., (Neuenschwander et al., 1995; Fua and Brechbuehler, 1995)),and the method of solution (e.g., (Wang et al., 1996; Amini et al., 1990)).

Ravine Extraction

Our method uses a hydrological analysis to infer likely ravine locations. This information is used to initialize the snaking process. The external energy function is based on standard edge detection methods applied to an aerial image covering the same terrain. The refinement of con-

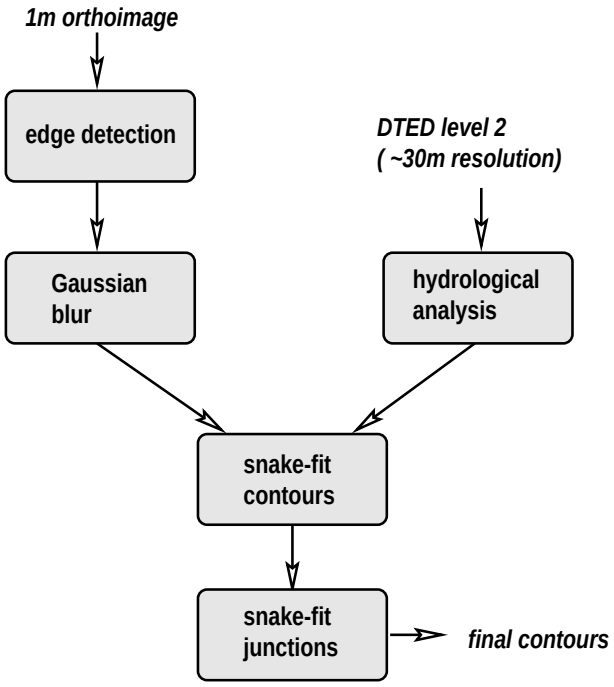

Figure 3: Overview of ravine extraction process.

tour location using the snaking process proceeds in two steps. The first adjusts individual segments of the tree of drainage patterns to best fit the imagery, subject to the contour smoothing criteria. The second step adjusts the junctions connecting multiple contour segments. (See Figure 3.)

Finding Drainage Features

Hydrological analysis was done using Arc/Info functions `flowdirection`, `flowaccumulation`, `con`, and `streamline`, which implement the steepest descent approach. DTED level 2 elevation source data was used. This is represented in a geographic (latitude/longitude) coordinate system, with a post spacing of 1 arc-second in both latitude and longitude. We used the slightly unconventional approach of doing the hydrological analysis in this spherical coordinate system to avoid the interpolation artifacts that would have been generated by reprojecting the geographic coordinates into a planar coordinate system. Instead, reprojection to a UTM coordinate system was done after the drainage patterns were found.

Figure 4 shows the NIMA ITD drainage features for a portion of the Range 400 area, overlayed onto an aerial image of the same region. Only one feature is present in this area and it is mislocated by a significant amount. Figure 5 shows the drainage features for the same area, determined using the hydrological analysis. Many more relevant features are included. Locations are still incorrect, however, which is not surprising given the limited resolution of the source data.

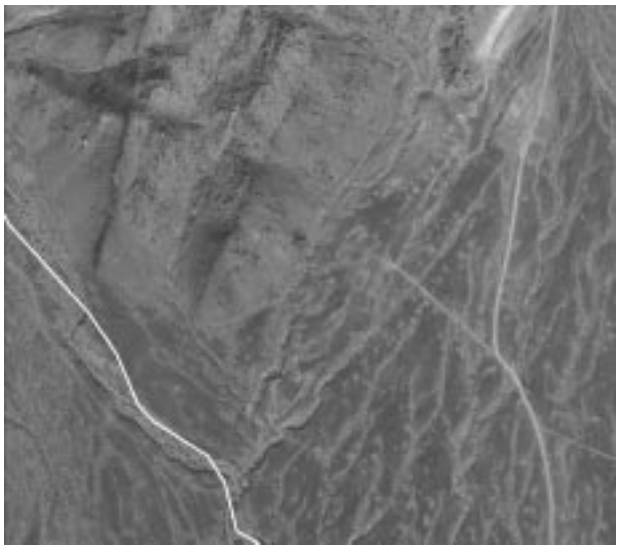

Figure 4: ITD drainage features in a portion of Range 400.

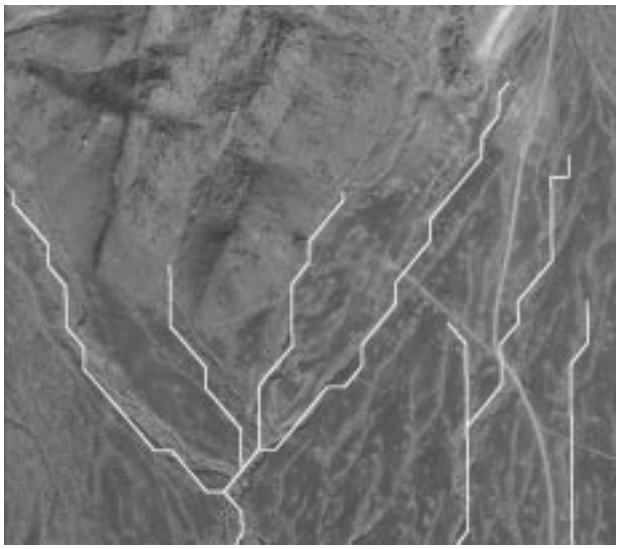

Figure 5: Drainage features for the same regions as shown in Figure 4, determined using hydrological analysis.

Refining Drainage Features Using Aerial Imagery

A single aerial image, georeferenced to the elevation data on which the hydrological analysis is based, can be used to refine the location of drainage patterns to precisions on the order of the resolution of the image. Since high resolution imagery is much more likely to be available for a given area than is high resolution elevation data, this can greatly improve the fidelity of a terrain database.

While we can guess at the locations of ravines in images such as shown in Figure 4, the visual signature of ravines is easily confused with other commonly occurring terrain and cultural features. This ambiguity can be significantly reduced if we use a top-down approach which uses the results of hydrological analysis as a starting point

for searching for ravines in the aerial imagery. Such an approach is easily implemented by using the results of the hydrological analysis as the starting estimate for a snaking computation. Snakes have the computational property that they will modify this estimate only in ways that make it better conform to the imagery.

The detailed appearance of ravines varies enormously, depending on local topography, soil type, and vegetation. As a result, standard contour extraction methods from computer vision won't work. Instead, we need a method for finding linear patterns that vary in some visual property, without knowing the details of the variation. Using edge density has proven to be effective in implementing this criteria. Figures 6–12 show an example. Figure 6 is a 1m orthoimage of the same area shown in Figures 4 and 5. Figure 7 shows the output of a Canny edge detector applied to this orthoimage. Automated methods for choosing the threshold and size parameters of the Canny operator are described in (Thoenen, 1998).

Snaking requires a potential function with a spatial extent at least as big as the likely uncertainty in initial contour location. We do this by applying a Gaussian blur to a binary representation of the edge image, using a standard deviation on the order of the resolution of the elevation image used as the basis of the hydrological analysis (Figure 8). As iterations of the snaking process proceed, the potential function is modified by reducing the amount of blur applied to the edge image (Thoenen, 1998). This has the effect of increasing the spatial localization of the final estimates.

Figure 9 shows the the final results of this process, indicating how much refinement has occurred over an analysis based only on the DTED level 2 elevation data (Figure 5). In Figure 10, these features have been overlayed on top of the original aerial image. The availability of accurate, high-resolution elevation data for Range 400 lets us compare the results with the actual terrain. Figure 11 shows the final estimated ravine locations overlayed on a shaded relief rendering of the terrain, based on the 1m DEM. Over most of the area, ravine centerlines are found to within a few meters of the correct location. It is important to note that the 1m elevation data is used only to validate results. The actual data extraction is based on DTED level 2 elevation data (with approximately 30m post spacing and well above 30m resolvable detail) and a single, 1m resolution aerial image. Figure 12 shows the final ravine features extracted over the whole of the Range 400 area, again overlayed onto a shaded relief rendering of the high-resolution DEM.

Figures 13 and 14 give an indication of the improvements that can be obtained from this approach on real terrain simulation applications. Figure 13 shows a view generated from the CCTT Primary Two database from a vantage point in the Range 400 area. The ravine that is shown was entered into the database using ITD features.

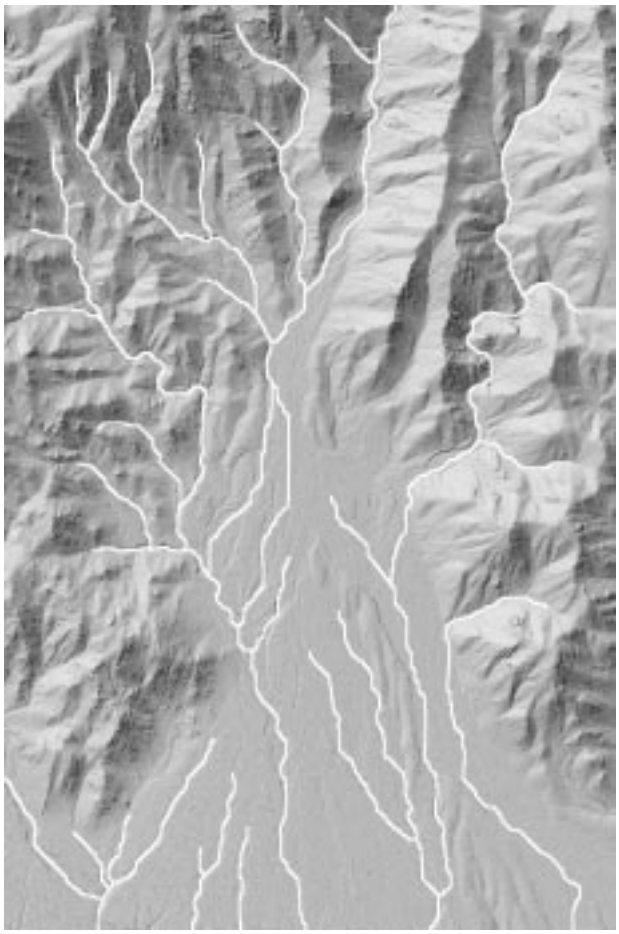

Figure 12: Final features – full Range 400, overlayed onto shaded relief.

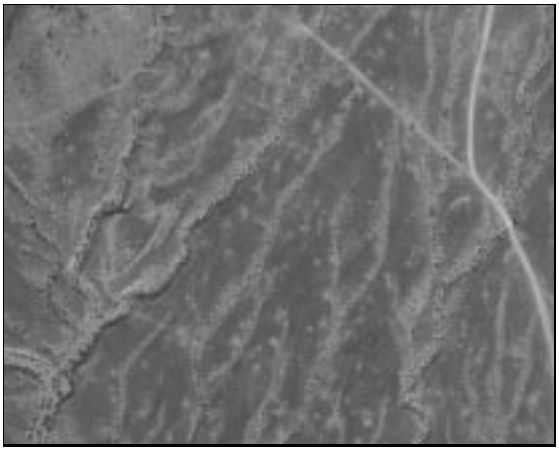

Figure 6: Aerial image.

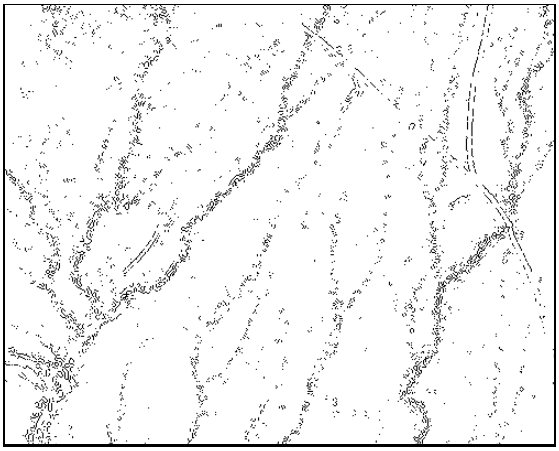

Figure 7: Edges derived from Figure 6.

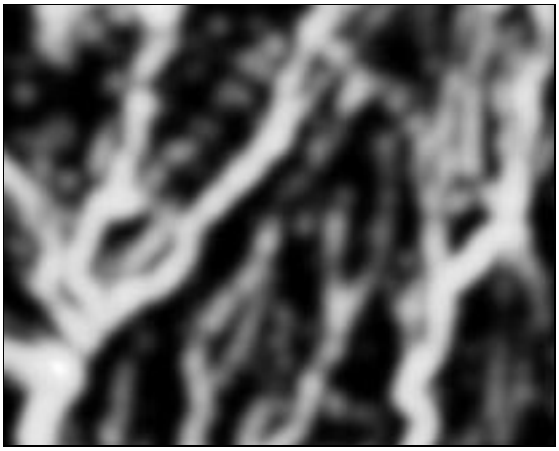

Figure 8: Potential function derived from edges shown in Figure 7.

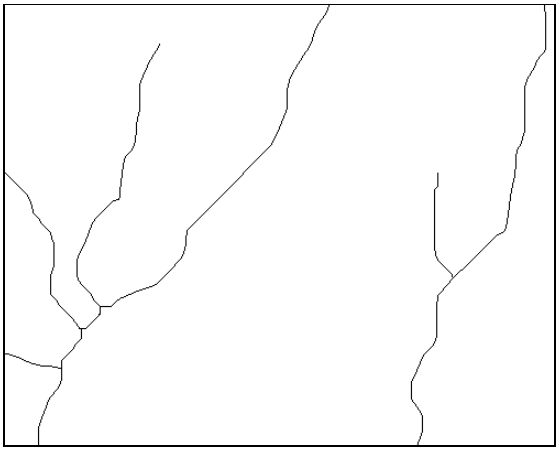

Figure 9: Final features.

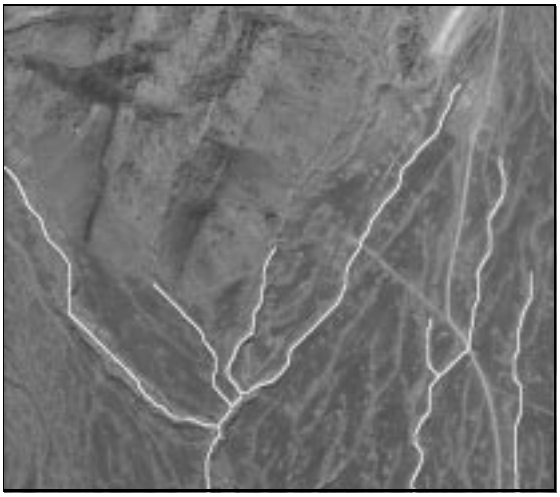

Figure 10: Final features, overlayed onto aerial image.

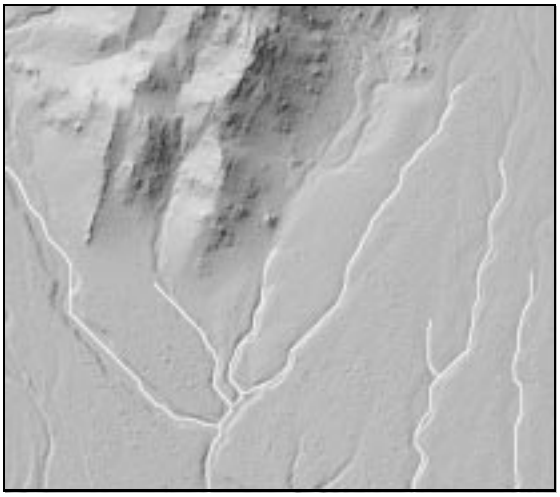

Figure 11: Final features, overlayed onto shaded relief generated from 1m elevation data.

In Figure 14, the database has been modified by removing ITD drainage features and replacing them with the features shown in Figure 12. The improvement in realism is immediately apparent. (The ravines remain standard width in order to satisfy CCTT SAF constraints – see next section.)

Determine the Width of Ravines

The method just described is effective at finding the centerlines of ravines, but it tells us nothing about the cross-section of the ravine. A complete cross-section analysis would require high-resolution 3–D photogrammetry, which in turn depends on the availability of very precisely controlled and calibrated stereo imagery. A straightforward extension of the snaking procedure outlined above, however, can be used to find the location of the two side-walls of a ravine using only a single aerial image.

If ravines had near constant width, *ribbon snakes* (Fua and Leclerc, 1990; Neuenschwander et al., 1994; Davatzikos and Prince, 1996) could be used. Since width is highly variable, a different approach is needed. Our method starts by first finding the location of ravine centerlines. This is used to initialize a second snaking process. Each centerline is split into two contours. These contours are optimized using a second potential function that is the gradient magnitude of the first potential function (which in turn was based on edges in the image). The effect is to create a new potential function with local maxima where the edge density in the image is changing most rapidly. The dot product of the gradient direction of the first potential function and the normal of the initial contours is used to insure that one contour is drawn to each side of the ravine.

Figures 15–20 show the results of applying this approach. Figures 15 and 16 are the same image and initial potential function as shown in Figures 6 and 8. Figure 17 shows the gradient magnitude of Figure 8. Figures 19 and 20 show the final ravine sidewall locations, overlayed onto the original aerial image and onto the shaded relief rendering of the 1m elevation data. Again, the high-resolution elevation data is used only to validate the results, it was not part of the feature extraction process. Figure 21 shows the final ravine sidewall features extracted over the whole of the Range 400 area, overlayed onto a shaded relief rendering of the high-resolution DEM. Figure 22 gives an indication of just how accurate these extracted features are. The figure was produced by draping the aerial image and the ravine sidewall extracted from the lower-resolution elevation data onto the high resolution DEM. It is likely that hand processing would not have produced results more accurate than this automated process.

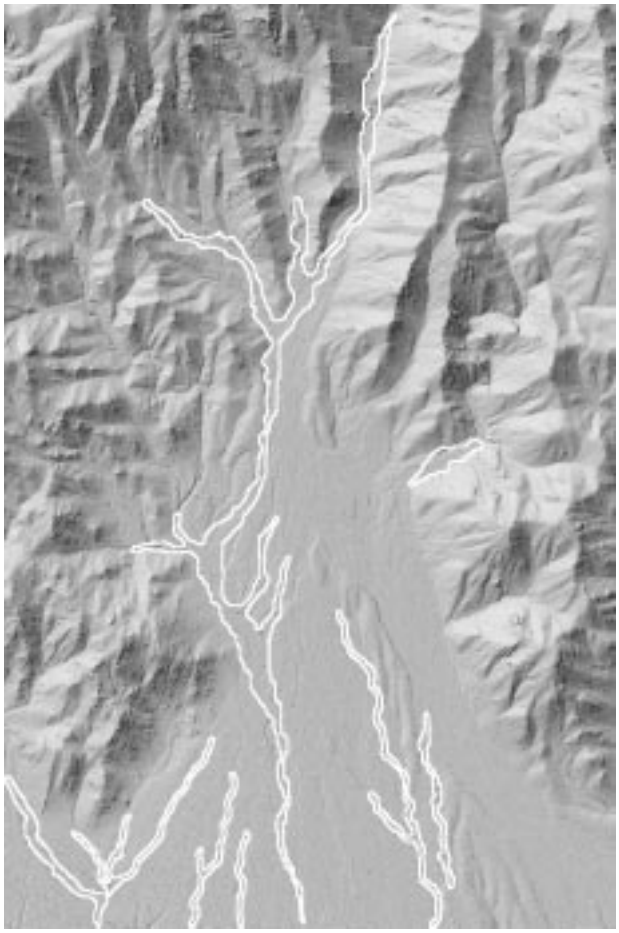

Figure 21: Extracted ravine sides – full Range 400, overlayed onto shaded relief.

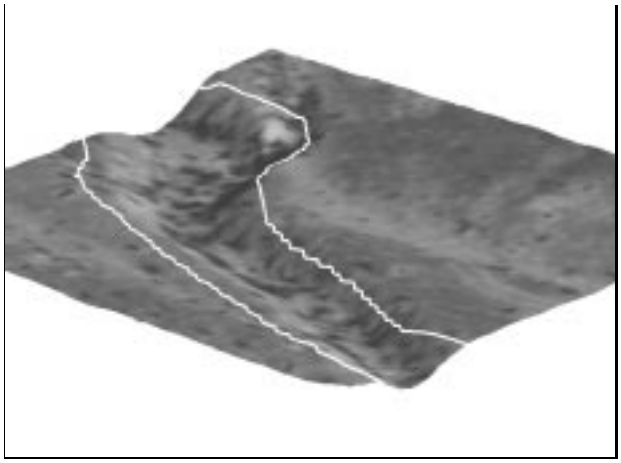

Figure 22: Extracted ravine sides rendered onto high-resolution DEM.

Figure 13: CCTT database with ITD-based drainage features.

Figure 14: CCTT database with improved ravine features.

Using Active Contours to Refine Existing Features

The method described above uses snakes (active contours) to refine the location of linear features first estimated from medium resolution elevation data using a hydrological analysis. The computer vision component of this approach can also be used to improve the accuracy of features already present in a database. Figures 23 and 24 give an example. Figure 23 shows an ITD drainage feature in a canyon section of the Range 400 area, overlayed onto an aerial image of the same region. Clearly, the feature is not accurately located. Figure 24 shows the results of using the ITD feature as an initial estimate in the snaking process and then refining this estimate using the imagery. The result is far better localization of the features.

More generally, tools such as this will become increasingly important as more and more visual simulations use image mapping techniques rather than stereotypical texture mapping. When image mapping is utilized, it is critical that the geometry and functional behavior of the terrain model match the image that is rendered. The same top-down, model-driven image analysis techniques used for ravine extraction can aid in consistent modeling of a wide range of features.

Figure 23: ITD drainage feature.

Figure 24: Refined ITD drainage feature.

Acknowledgment

This work was supported in part by the Defense Advanced Research Projects Agency, under contract DACA76-97-K-000.

References

Amini, A., Weymouth, T., and Jain, R. (1990). Using dynamic programming for solving variational problems in vision. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(9):855–867.

Berger, M.-O. and Mohr, R. (1990). Towards autonomy in active contour models. In *10th International Con-*

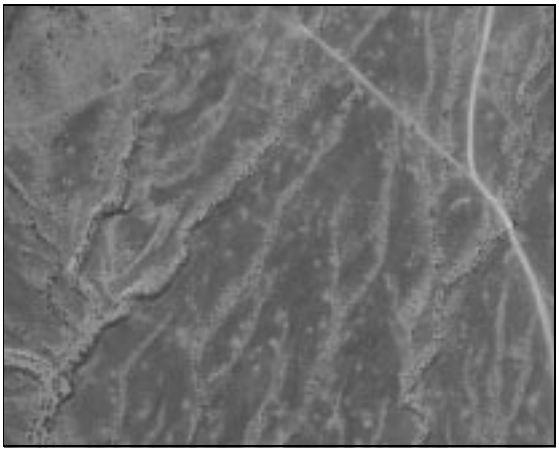

Figure 15: Aerial image.

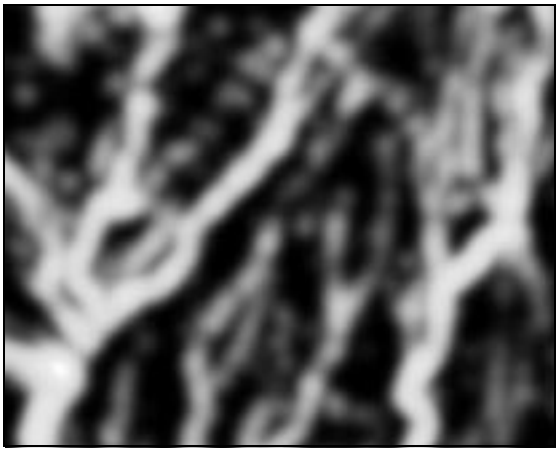

Figure 16: Potential function.

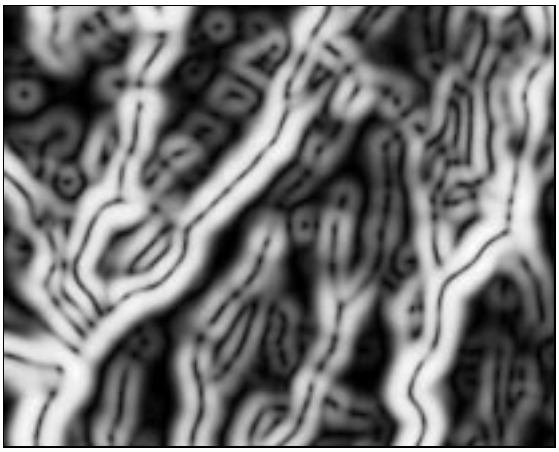

Figure 17: Gradient of potential function.

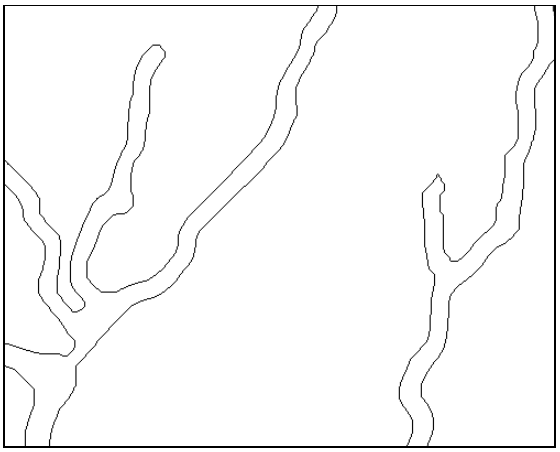

Figure 18: Extracted ravine sides.

Figure 19: Extracted ravine sides, overlayed onto aerial image.

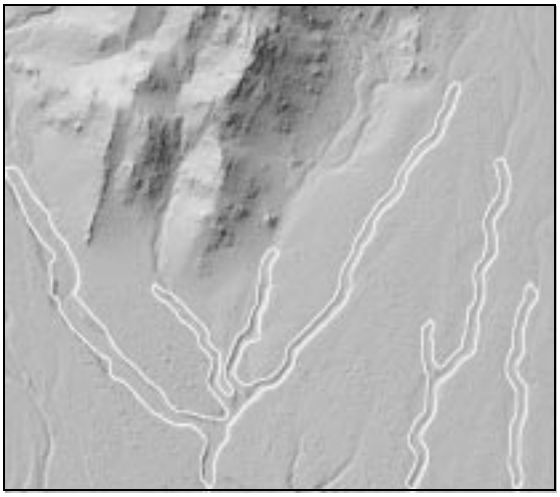

Figure 20: Extracted ravine sides, overlayed onto shaded relief.

ference on Pattern Recognition, pages 847–851, Piscataway, NJ.

Chorowicz, J., Kim, J., Manoussis, S., Rudant, J. P., Foin, P., and Veillet, I. (1989). A new technique for recognition of geological and geomorphological patterns in digital terrain models. *Remote Sensing of Environment*, 29:229–239.

Davatzikos, C. and Prince, J. (1996). Convexity analysis of active contour problems. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 674–679.

Fua, P. and Brechbuehler, C. (1995). Imposing hard constraints on soft snakes. Tech Note 553, Artificial Intelligence Center, SRI International.

Fua, P. and Leclerc, Y. (1990). Model driven edge detection. *Machine Vision and Applications 3*, pages 45–56.

Jenson, S. K. and Domingue, J. O. (1988). Extracting topographic structure from digital elevation data for geographic information system analysis. *Photogrammetric Engineering and Remote Sensing*, 54(11):1593–1600.

Kass, M., Witkin, A., and Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, pages 321–331.

Lai, K. and Chin, R. (1993). On regularization, formulation and initilization of the active contour models (snakes). In *Second Asian Conference on Computer Vision*, pages 542–545, Singapore.

Leymarie, F. and Levine, M. D. (1993). Tracking deformable objects in the plane using an active contour model. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(6):617–634.

Mark, D. M. (1983). Automated detection of drainage networks from digital elevation models. In *Auto-Carto VI: Proceedings Sixth International Symposium on Computer Assisted Cartography*, pages 288–298.

Menet, S., Saint-Marc, P., and Medioni, G. (1990). Active contour models: Overview, implementation, and applications. In *IEEE International Conference on Systems, Man and Cybernetics*, pages 194–199, Los Angeles, CA.

Neuenschwander, W., Fua, P., Szekely, G., and Kubler, O. (1994). Initializing snakes. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 658–663, Seattle.

Neuenschwander, W., Fua, P., Szekely, G., and Kubler, O. (1995). From ziplock snakes to velcro surfaces. In *Automatic Extraction of Man-made Objects from Aerial and Space Images*, pages 105–114. Birkh auser Verlag.

Peucker, T. K. and Douglas, D. H. (1975). Detection of surface-specific points by local parallel processing of discrete terrain elevation data. *Computer Graphics and Image Processing*, 4:375–387.

Pope, C. N., Vuong, M., Moore, R. G., and Cowser, S. S. (1995). A whole new CCTT world. *Military Simulation and Training*, (5).

Radeva, P., Serrat, J., and Marti, E. (1995). Snakes for model-based segmentation. In *Fifth International Conference on Computer Vision*, pages 816–821, Cambridge MA.

Richbourg, R. and Olson, W. K. (1996). A hybrid expert system that combines technologies to address the problem of military terrain analysis. In *Expert Systems and Applications*.

Richbourg, R., Ray, C., and Campbell, L. L. (1995). Terrain analysis from visibility metrics. In *Conference on Integrating photogrammetric techniques with scene analysis and machine vision II, (SPIE Proceeding Volume 2486-2)*, Orlando, FL.

Thoenen, G. W. (1998). Extraction of micro-terrain ravines using hydrological analysis combined with active contours. Master's thesis, University of Utah.

Tribe, A. (1992). Automated recognition of valley lines and drainage networks from grid digital elevation models: A review and a new method. *Journal of Hydrology*, 139:263–293.

Wang, M., Evans, J., Hassebrook, L., and Knapp, C. (1996). A multistage, optimal active contour model. *IEEE Transactions on Image Processing*, 5(11):1586–1591.