Structural Lines, TINs, and DEMs

James J. Little and Ping Shi
Department of Computer Science
University of British Columbia
Vancouver, BC, Canada, V6T 1Z4
little@cs.ubc.ca 604-822-4830 (Tel) 604-822-5485 (FAX)

Keywords: TIN, DEM, surface-specific lines, constrained Delaunay triangulation, curvature

Abstract

The standard method of building compact triangulated surface approximations to terrain surfaces (TINs) from dense digital elevation models(DEMs) adds points to an initial sparse triangulation or removes points from a dense initial mesh. Instead, we find structural lines to act as the initial skeleton of the triangulation. These lines are based on local curvature of the surface, not on the flow of water. These features use only local information. When linked into curves, we can ascribe a "scale" to the entire curve; ranking them using this scale measure lets us choose the most important curves for inclusion in the surface description. Experiments in building TINs from DEMs with points, structural lines and structural lines ranked by "scale" show that including such lines initially reduces the error in the resulting as well as improving its fidelity to the structure of the terrain.

1 Introduction

We begin with dense terrain data specified on a grid of points, a digital elevation model (DEM), and derive a triangulation, a collection of nonoverlapping planar triangular regions that fit the DEM with minimal error, often called a Triangulated Irregular Network (TIN)[20].

Many techniques for approximating a surface, usually a terrain height field, begin by selecting points that are expected to be critical in the final approximation[8, 22, 9, 15, 10]. From this initial triangulation, the surface is improved by adding points. One particular method, [6], finds in each triangle the point that is most poorly fit by the current triangulation, and adds that point to the Delaunay triangulation[17] of the points. Iteratively following this process produces triangulations that eventually fit the surface well, but with many fewer points than the source dense data.

It is clear that this method is not guaranteed to fit the surface well at surface discontinuities, or at slope discontinuities, both of which occur frequently in terrain, and especially in range maps produced in computer vision. The literature[8] abounds with counterexamples. To avoid these failings, Fowler and Little[6] first identify *ridges* and *channels*, surface lines determined by the flow of water away from them (ridges) and into them (channels), by simple local geometric operations[21]. These *structural lines* are then fit by a polygonal approximation[4] and included in the triangulation, by forcing the triangulation to include these lines. Modern methods allow incremental construction of Delaunay triangulation "constrained" by initial line segments[12]. However, any errors in the initial points/lines force the triangulation method to introduce further points, reducing the savings.

Schmitt and Chen[23] have updated Fowler and Little's method by first identifying surface and slope discontinuities and including these lines in the resulting approximation, which uses their own triangulation criterion. They choose lines based on the local differential structure of the surface, which is independent of the choice of coordinate system, and is not necessarily coincident with the paths determined by the flow of water[14]. [16, 7] also insert "crest" lines into adaptive meshes to improve stereo-driven surface approximation. The resulting lines are not in fact ridges, rather local extrema of curvature (coordinate system independent), but we will knowingly abuse terminology and continue to refer to them as ridges. The ridges for a section of the Crater Lake DEM are shown in white in Fig. 1. The method of deriving these ridges will be explained in Section 2.

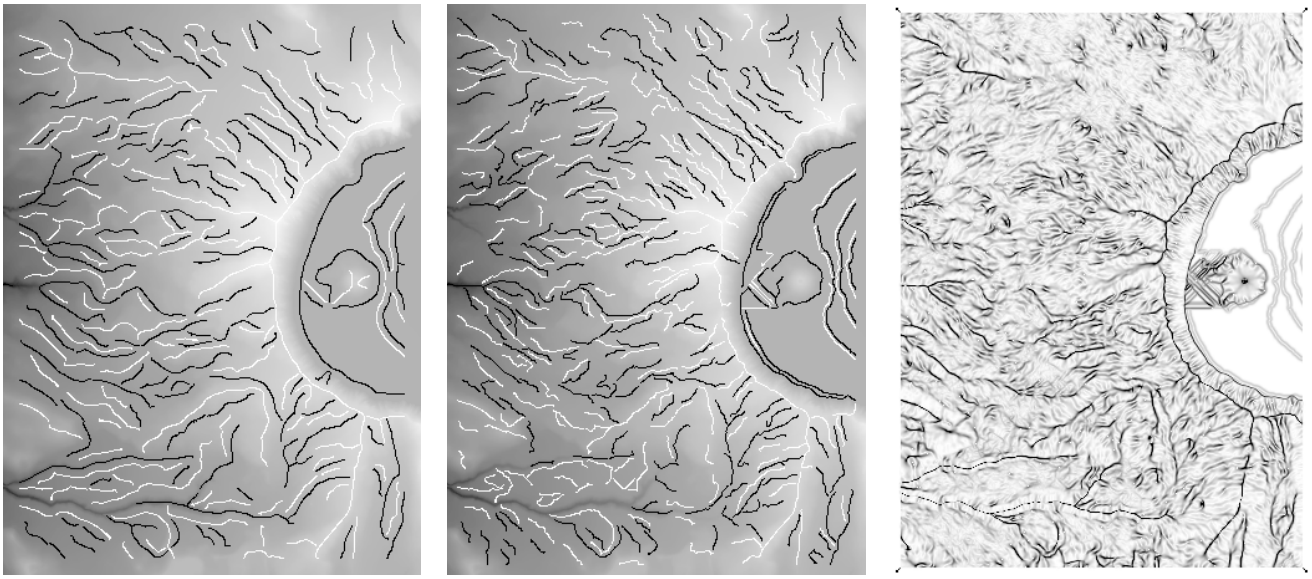

(a) (b) (c)

Figure 1: (a) Curvature-based ridges and channels ($\sigma = 2.5$) overlaid on the Crater Lake DEM; (b) ridges and channels ($\sigma = 1.0$); c) absolute value of maximal curvature ($\sigma = 1.0$): darker is larger.

The significance of surface lines is related to their role in the surface representation. Increasingly these triangulated approximations are used in surface visualization and interaction. The usual criterion for quality of approximation is the root-mean-square of the error, but this does not capture the importance of the "features" such as ridge lines, streams, and slope breaks.

What is the fidelity of an approximation? Typically the root-mean-square (RMS) error of the vertical difference is used, but in many visualization tasks, retaining features such as horizon lines is more important than aggregate errors. Also, in drainage queries[24] and other such applications, preserving the surface-specific lines is critical.

2 Curvature Descriptions

In [6], ridges and channels are found by marking the points to which water would flow[21], leaving unmarked the ridge lines (see Fig. 2). In a symmetric fashion the channels are found. Even in this method there are many small segments identified to which little significance can be given. The figure shows the results of marking ridges, after the DEM has been smoothed by a Gaussian filter with $\sigma = 2.5$ to reduce quantization effects. Despite the smoothing, the ridge lines are interrupted by many gaps.

Instead, we determine the local surface properties independent of the coordinate system. Thus we can select linear features along ridges and channels, as well as other lines where surface slope changes sharply. Like [16, 7, 19], we determine the local surface curvature description. At each point the tangent plane is computed; it is orthogonal to the surface normal $\vec{n}$. In any direction, the surface can be cut by a plane containing $\vec{n}$; the *normal* curvature in that direction is the curvature of the curve formed by the intersection of the surface and the plane containing $\vec{n}$. The *principal directions* are the two directions $\vec{e_1}$ and $\vec{e_2}$ where the value of the normal curvature reaches its maximum and minimum values, k_1 and k_2. We choose subscripts so that k_1 is the curvature of maximum absolute value, and $\vec{e_1}$ and $\vec{e_2}$ are vectors in the local tangent plane pointing the direction of maximal and minimal curvature.

To compute these quantities, we first locally determine the surface derivatives and thence the first and second fundamental forms, and thence the curvature and the principal directions (see [5]). The principal directions on the surface define a net[13].

Since the hills under ridges are elongated, the curvatures create a local cylindrical approximation to the ridge or channel, with an non-circular cross-section. Figure 3 shows a ridge with a ridge line, at the top; the transverse lines are in the direction of maximal curvature locally ($\vec{e_1}$).

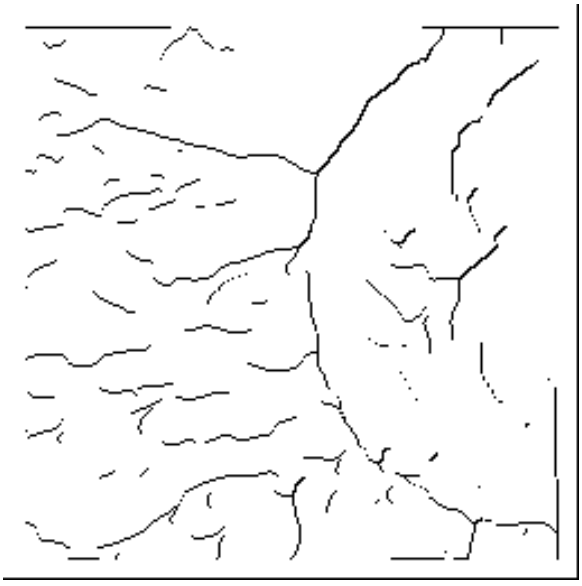

Figure 2: The ridge lines found in the right central section of the Crater Lake DEM by the marking method of Peucker and Douglas; the DEM has been smoothed by a Gaussian filter ($\sigma = 2.5$).

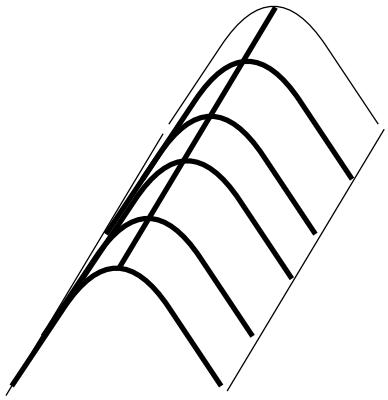

Figure 3: Ridge (line at the top) with transverse curves in direction of maximal curvature.

At each point, we determine whether the maximum curvature k_1 at the point is locally maximal, in the principal direction $\vec{e_1}$. The image of $|k_1|$ appears in Fig. 1(c); higher values appear darker. Unlike [16], we use non-maximum suppression to identify these points, looking in the direction of maximal curvature, $\vec{e_1}$, and marking points that are greater (in absolute value) than the neighboring points along the line of curvature. Thus in this description a "ridge" is a line connecting points of locally maximal curvature, where that curvature is positive. A "channel" or "course" is such a line with negative curvature. To find these lines, we track lines and connect the points, employing hysteresis with thresholding, using the magnitude of the maximal curvature[2]. This tracking, followed by pruning short features, produces the ridge and channel lines shown in Fig. 1.

An important issue in computing curvature-based ridges and channels is the scale of the features, the amount of simplification or smoothing of the surface when curvature is computed. Because of quantization and noise, it is important to filter the surface by smoothing with a Gaussian filter whose scale is described by the parameter σ. Figure 1 shows the effect of smoothing with different scales—at the coarser scale the curves are better connected and some small curves have been eliminated.

3 Triangulation

In the original work in this area[6], two innovations were proposed: incremental "greedy" triangulation of a TIN by inserting points based upon the error in each triangle, and preservation of structural lines found by marking ridges and channels and then generalizing these 3D lines. Incremental improvement is widely used now, together with many variations in criteria for adding points. We have not experimented with this, and include in each triangle the point (the "worst" point) with most error in the current approximation.

Many different strategies can be used for determining the order of insertion. [6] proceeded by inserting every "worst" point if it exceeded a desired error tolerance, continuing until all points were within this tolerance.

[8] introduced the idea of "batching" updates, collecting the worst points in each triangle, and only selecting points whose error exceeded some fraction α of the current maximum error. As α approaches 1.0, the triangulation becomes sequential, inserting one point for each pass over the data. This is extremely expensive, so α is set at 0.8 in our experiments, which is slow, but can, we hope, permit the triangulation process to take advantage of the different initial skeletons provided by various feature lines.

The second innovation of [6] is not often used as the processing to determine structural lines is more complex. Fowler and Little forced the structural lines into the triangulation after inserting points. Since that time, Constrained Delaunay Triangulation (CDT) has become well understood, so the structural lines will be inserted initially as part of the triangulation. We have the adapted the incremental CDT software of Dani Lischinski, available at `http://www.cs.huji.ac.il/ danix/` to insert points based on the error between the current triangulation and the DEM.

3.1 Snakes: Deforming Large Scale Lines

Lines found at the coarse scale ($\sigma = 2.5$) may have been displaced by the smoothing process. The crest on an asymmetrical ridge, where slope on one side is significantly steeper than the other, will be displaced toward the less steep side. When these lines are used as the skeleton, the triangulation process must include "corrective" points near the ridge to model the actual location of the crest. This is undesirable. To move the coarse level line to the location of the fine-level line, we use the "snake" method[11, 1, 3]. The essential idea of snakes is to deform the initial line until it reaches a new position close to the line at the fine scale. The snake method allows the line to deform to minimize the sum of "internal energy", the energy of stretching the line, and "external energy", the attractive force applied by some external source, in this case the proximity to the lines at the finer scale. We used the actual curvature field computed at the finer scale (Fig. 1(c) instead of the proximity.

Deforming the lines is an iterative process; at each step, each point checks neighboring points and finds the best location. The local solutions are combined using dynamic programming[7]; we have used the implementation described in [1].

4 Experiments

To determine whether including structural lines can improve the resulting triangulation, we compare triangulations produced by pure "greedy" triangulation, with no lines (called `no lines`), with "greedy" constrained Delaunay triangulation (CDT) with a variety of structural lines. The various feature lines are:

- curvature-based feature lines at fine scale (`fine`)
- coarse scale feature lines corrected to a fine scale (`snaked`)
- selected lines at fine scale (`selected`)
- selected lines at coarse scale corrected to a fine scale (`sel-snaked`)

The curvature-based feature lines are shown in Fig. 4(a), in white, for the fine scale (used in `fine`). The coarse scale features are shown in black in both parts of the figure; in Fig. 4(b), they appear, in white, as corrected to fine scale using the snake method of Sec. 3.1. These are used in `snaked`.

[18] describes how to determine the region surrounding a ridge (or channel) and then compute a surface measure that can be integrated over the region. Such a scalar measure provides a metric criterion for preferring some feature lines over others. Reducing the number of initial lines should improve the fit of the surface, all other things being equal. We use an estimate of "creasiness" of the ridge, which is computed by summing the absolute value of maximum curvature along the ridge. Lines selected on the basis of "creasiness" appear in Fig. 5(a) in black. Because they have been computed at coarse scale ($\sigma = 2.5$), they can be "corrected" using the snake method to better fit the surface at fine scale (shown in white); these are used in `sel-snaked`. The structural lines in `selected` are shown in Fig. 5(b).

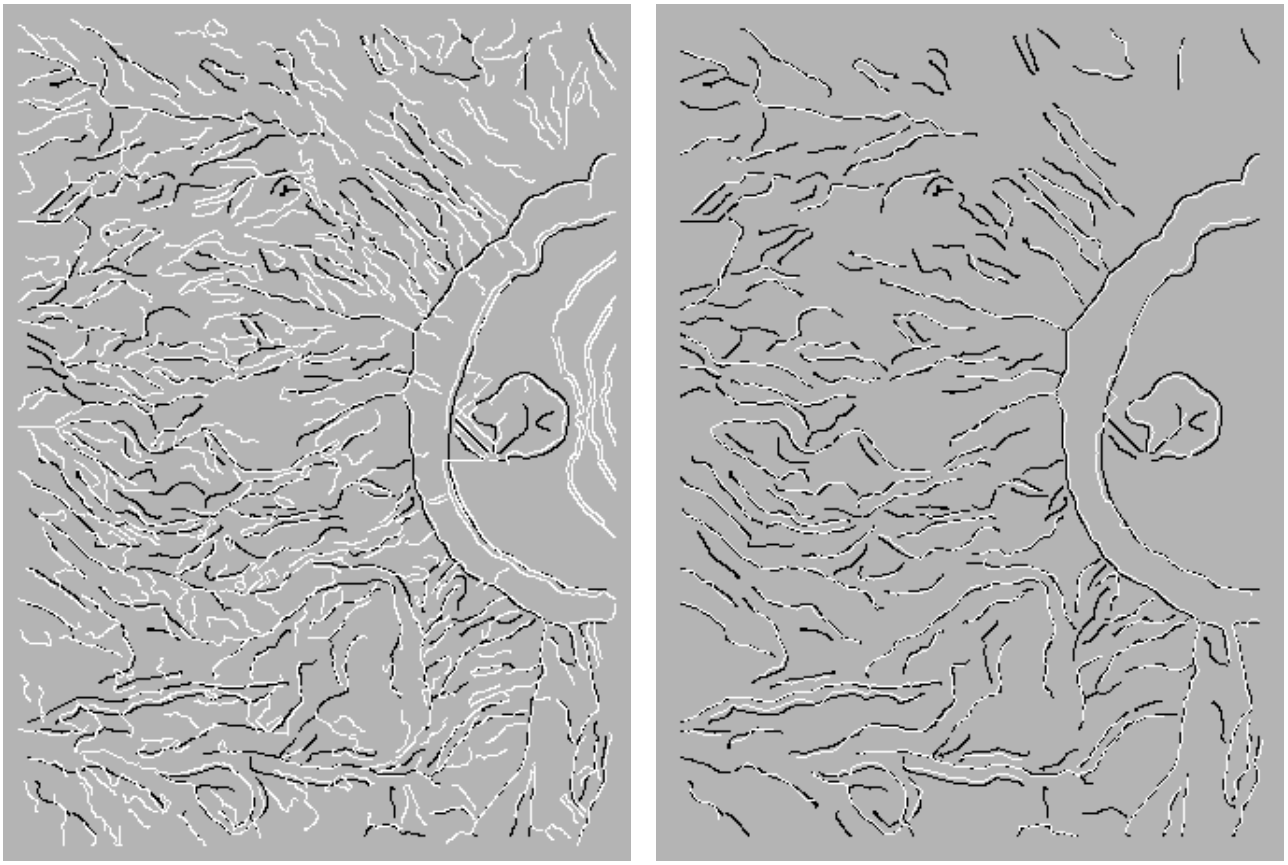

(a) (b)

Figure 4: (a) lines at fine scale ($\sigma = 1.0$) in white, lines at $\sigma = 2.5$ in black; (b) lines at $\sigma = 2.5$ in black with the corrected result in white 1.0

RMS	5.0	4.0	3.0	2.5	2.25	2.0
Ratio of points (`snaked/no lines`)	0.92	0.93	0.86	0.89	0.91	0.96
Points `no lines`	1611	2097	3219	3995	4036	4653
Points `snaked`	1478	1949	2781	3555	3696	4463
Percent of points `no lines`	1.04	1.36	2.08	2.59	2.62	3.02
Percent of points `snaked`	0.96	1.26	1.80	2.30	2.40	2.89

Table 1: The ratio of the number of points in `snaked` to points in `no lines` for various RMS values

The tests were run on the Crater Lake DEM, running the batched greedy triangulation (subject to constraints provided by structural lines), with $\alpha = 0.8$. The triangulation stopped when 5000 points had been included. To understand the effect of using structural lines in the triangulation, we plot the root-mean-square error versus the number of points, for several experiments in Fig. 6 and, at greater magnification, in Fig. 7.

The results can also be seen by considering the number of points required to achieve a particular RMS error, as shown in Fig. 8, which is simply a rotation of Fig. 6. In all these cases the results given by using `snaked`, the lines from a coarse scale corrected to a fine scale by snakes, were superior. Figure 9 shows the triangulations (limited to 2500 points) produced by the two methods; the constrained lines are shown in black and other edges in white.

(a) (b)

Figure 5: (a)Selected lines at $\sigma = 2.5$ in black with the corrected result in white 1.0 (`sel-snaked`); (b) Selected lines at $\sigma = 1.0$ (`selected`)

Points	1000	2000	3000	4000	5000
Ratio of RMS error (`snaked/no lines`)	0.94	0.90	0.85	0.88	0.97
RMS `no lines`	6.93	4.31	3.17	2.42	1.83
RMS `snaked`	6.47	3.88	2.70	2.14	1.77

Table 2: The ratio of the RMS of `snaked` to `no lines` for various numbers of points

The results can be presented in a tabular form, as shown in Tables 1 and 2. In both tables, it is clear that the `snaked` TIN uses fewer points than the TIN produced with points along, as much as 15% fewer.

5 Discussion

How have we advanced beyond the work of Fowler and Little[6]? We have a sounder basis for local surface description, more immune to quantization introduced in the DEM, and not coordinate system dependent, permitting description of surface breaks as well as ridges and channels. We also can compute measures of the feature lines or the regions surrounding them that allow us to rank the significance of the lines.

The inclusion of lines in the construction of TINs reduces the number of points required to achieve a particular RMS error. Alternatively, at a particular size of the TIN, the RMS achieved is less when the structural lines are included.

Finding curvatures, estimating derivatives, requires smoothing or regularization for stability, but this impedes

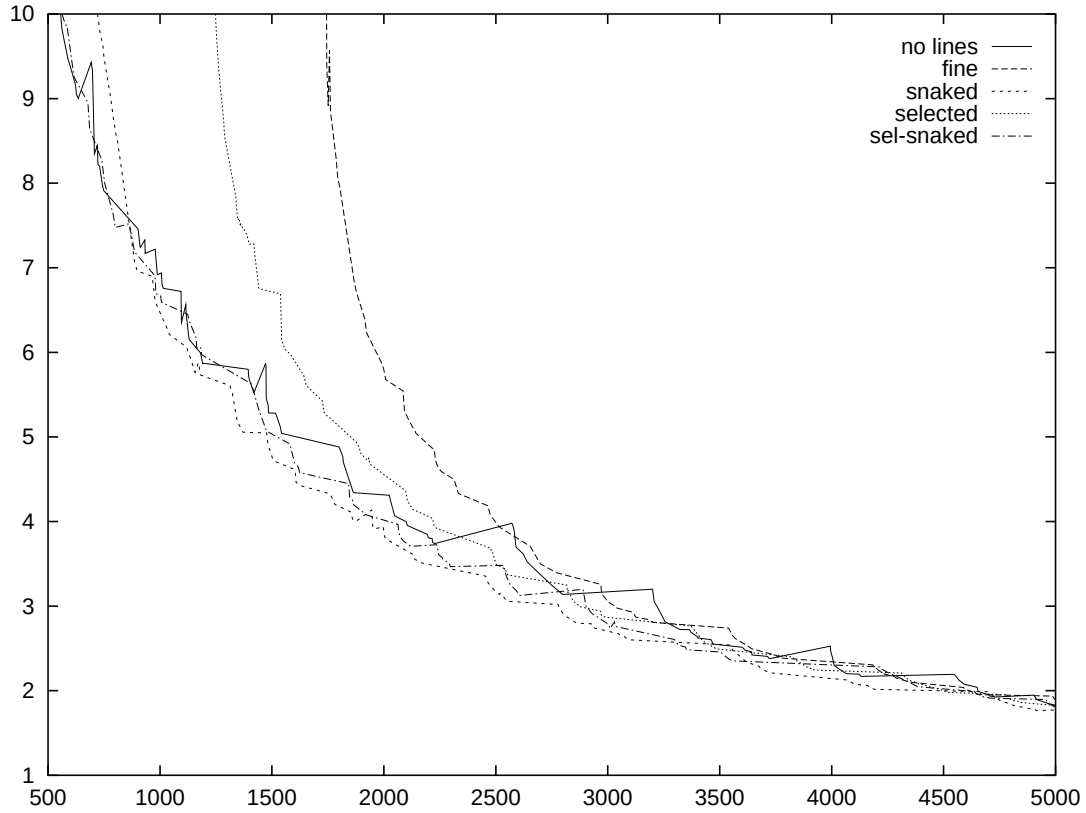

Figure 6: Root-mean-square error versus the number of points, for several experiments.

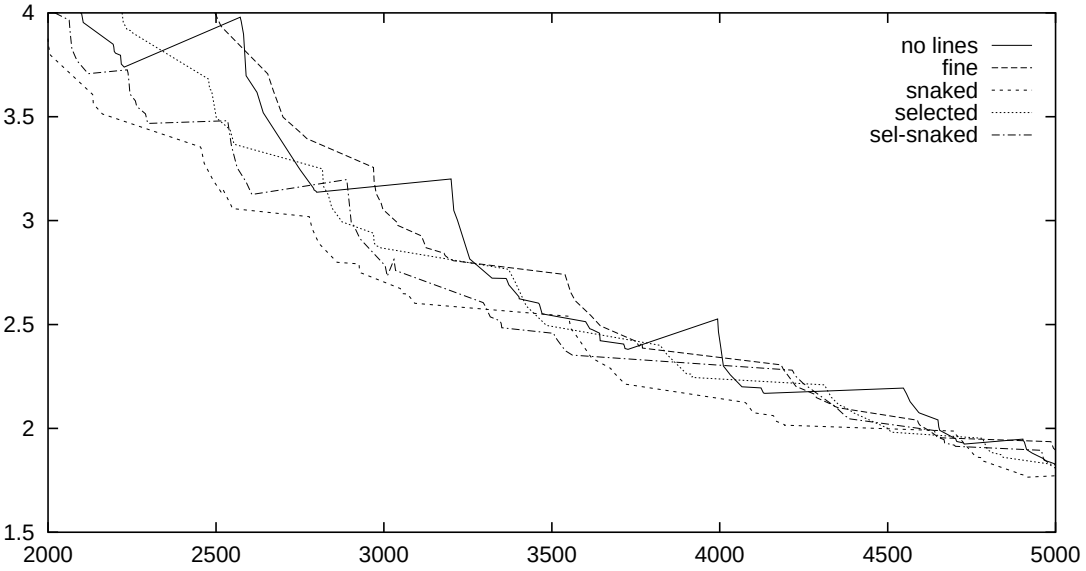

Figure 7: Root-mean-square error versus the number of points, for several experiments; only a portion of the previous graph is shown

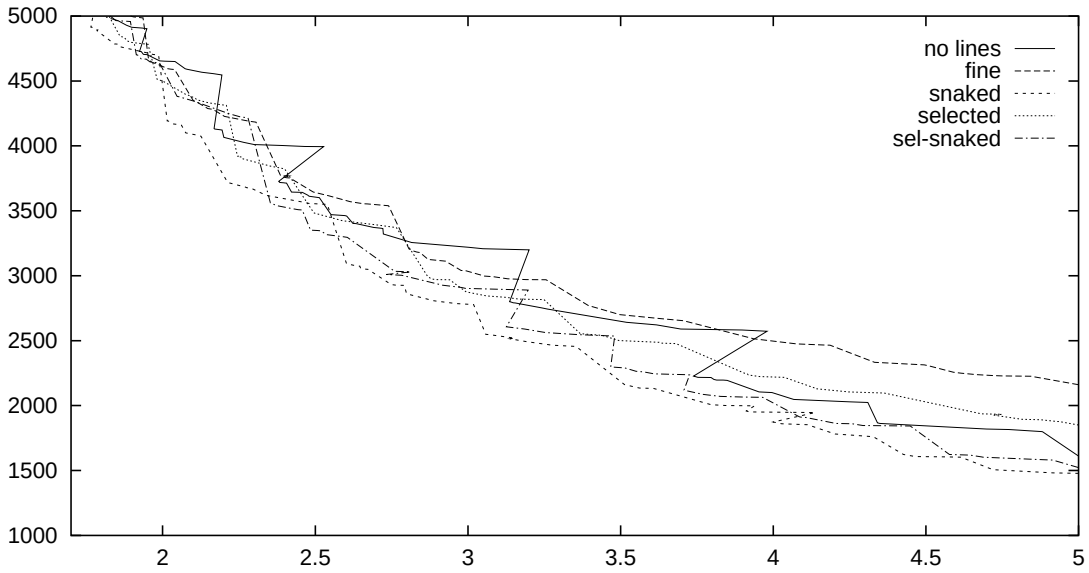

Figure 8: The number of points in the TIN versus the root-mean-square error. The experiments varied the number of points. At times, increasing the number of points increased the RMS error, which accounts for the multiple values for a particular RMS at some locations.

localization. The snake method lets us extract significant lines at a coarse scale and correct them to improve their localization for use in the TIN. Methods that would combine smoothing and localization, such as anisotropic diffusion, may aid determination of structural lines.

Further work will address the issue of how the scalar measures of the feature lines relate to the linear scale used in the curve approximation for the structural lines: how does a measure of approximation on lines translate to a measure of curvature fidelity? There are many alternative methods for choosing the best point for insertion, but we have not explored these here. Some may perform much better when structural lines are used. The normal error (the error measured perpendicular to the surface), not the vertical error, is arguably more useful, but our experiments with the normal error are still preliminary. There remains much to understand on the interaction of smoothing, the scale of features, and incremental improvement strategies for creating TINs from DEMs.

References

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

[2] J. F. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[3] S. Chandran, T. Maejima, and S. Miyazaki. Global minima via dynamic programming: Energy minimizing active contours. In *SPIE Vol 1570, Geometric Methods in Computer Vision*, pages 391–402, 1991.

[4] David H. Douglas and Thomas K. Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *The Canadian Cartographer*, 10(2):112–122, 1973.

[5] O. Faugeras. *Three-Dimensional Computer Vision*. MIT Press, Cambridge, MA, 1993.

[6] Robert J. Fowler and James J. Little. An automatic method for the construction of irregular network digital terrain models. In *Proceedings of SIGGRAPH '79*, pages 199–207, Chicago, Illinois, August 1979.

[7] Pascal Fua. Model-based approach to accurate and consistent 3-D modeling of drainage and surrounding terrain. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1997*, pages 903–908, 1997.

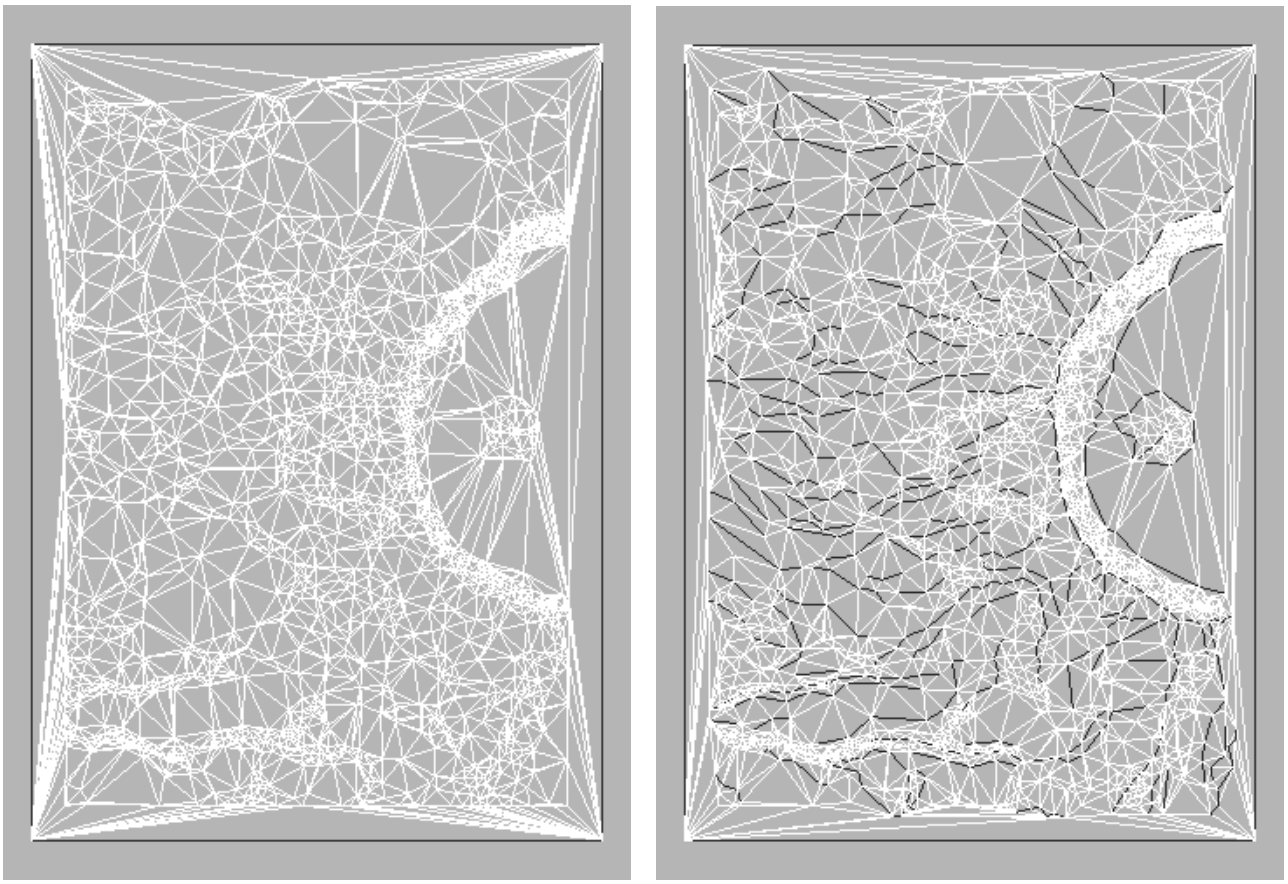

(a) (b)

Figure 9: (a) The triangulation (2500 points) for `no lines`; (b) The triangulation (2500 points) for `snaked`–in both the constrained lines are shown in black and other edges in white.

[8] Michael Garland and Paul S. Heckbert. Fast polygonal approximation of terrains and height fields. Technical Report CMU-CS-95-181, Carnegie Mellon U., September 1995.

[9] Paul S. Heckbert and Michael Garland. Survey of polygonal simplification algorithms. Technical report, CMU, 1997.

[10] Martin Heller. Triangulation algorithms for adaptive terrain modeling. In *Proceedings of the 4th International Symposium of Spatial Data Handling*, pages 163–174, 1990.

[11] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1988.

[12] D. G. Kirkpatrick. Efficient computation of continuous skeletons. In *Proc. of the 20th Annual IEEE Symposium on Foundations of Computer Science*. 1979.

[13] J. J. Koenderink. *Solid Shape*. MIT Press, Cambridge, MA, 1990.

[14] J.J. Koenderink and Andrea van Doorn. Local features of smooth shapes: Ridges and courses. In *SPIE(2031)*, pages 2–13, 1993.

[15] Mark P. Kumler. An intensive comparison of triangulated irregular networks (tins) and digital elevation models (dems). *Cartographica*, 31(2):1–48, 1994.

[16] Richard Lengagne, Oliver Monga, and Pascal Fua. Using differential constraints to reconstruct complex surfaces from stereo. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1997*, pages 1081–1086, 1997.

[17] Dani Lischinski. Incremental Delaunay triangulation. In P. Heckbert, editor, *Graphics Gems IV*. Academic Press, 1994.

[18] James J. Little and Ping Shi. The scale of structural lines. Submitted to CVPR-98, December 1997.

[19] O. Monga and S. Benayoun. Using partial derivatives of 3D images to extract typical surface-features. *Computer Vision Graphics and Image Processing: Image Understanding*, 61(2):171–189, March 1995.

[20] T. K. Peucker, R. J. Fowler, J. J. Little, and D. M. Mark. The triangulated irregular network. In *Proc. of the Digital Terrain Models Symp.*, pages 516–532, St. Louis, MO, 1978.

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

[22] M.F. Polis and D.M. McKeown. Iterative TIN generation from digital elevation models. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1992*, pages 787–790, 1992.

[23] Francis Schmitt and Xin Chen. Vision-based construction of CAD models from range images. In *Proc. 4th International Conference on Computer Vision*, pages 129–136, 1993.

[24] Sidi Yu, Marc van Kreveld, and Jack Snoeyink. Drainage queries in tins: From local to global and back again. In *7th Symp Spatial Data Handling*, pages 13A.1–13A.14, 1996.