Structural Lines for Triangulations of Terrain

James J. Little and Ping Shi
little@cs.ubc.ca and pshi@cs.ubc.ca
Department of Computer Science
University of British Columbia
Vancouver, BC, Canada, V6T 1Z4

Abstract

To build triangulated approximations to terrain surfaces from dense elevation models, we find structural lines for the initial skeleton of the triangulation. We describe the surfaces as regions; the spines of these regions are the structural lines on the surface. We use these lines as a skeleton of the surface, points and edges that initialize the triangulation. Simple local curvature analysis of the terrain results in many lines whose significance is only local. We show how the curvature analysis can be extended to find local regions around the structural lines. We use various properties of the regions to assign a significance to the lines, to rank them for inclusion in the skeleton. We then use the lines to approximate the surface with a compact triangulation.

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).

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[5, 11]. From this initial triangulation, the surface is improved by adding points. One particular method [3] finds in each triangle the point that is most poorly fit by the current triangulation, and adds that point to the Delaunay triangulation of the points. Iteratively following this process produces triangulations that eventually fit the surface well, but with many fewer points than the dense source data.

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. The literature [5] abounds with counterexamples. To avoid these failings, [3] first identifies *ridges* and *channels*, surface lines determined by the flow of water away from them (ridges) and into them (channels), by simple local geometric operations. These *structural lines* are then fit by a polygonal approximation 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[8].

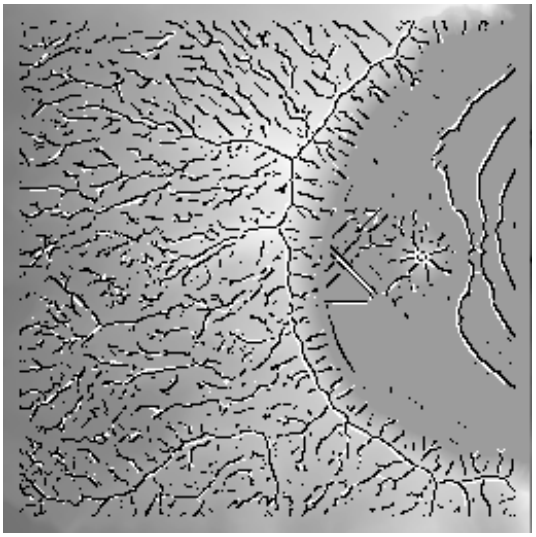

Figure 1: Crater Lake terrain with ridges (black:σ = 1.0, and white:σ = 2.5).

However, any errors in the initial points/lines force the triangulation method to introduce further points, reducing the savings.

Schmitt and Chen[2] have updated this 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. [7, 4] also inserts "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). Little and Shi[9] showed how to extract structural lines based on local curvature descriptions and use them as the basis for a constrained Delaunay triangulation of the surface. The structural lines for a section of the Crater Lake DEM are shown in white in Fig. 1. Their method of deriving these lines will be explained in Section 3.

However, surface curvature is a purely local property; it measures only the local change in the direction of the surface normal. At small scales, i.e., when σ of the Gaussian smoothing is small, there are many local

"creases" in the surface (black lines in Fig. 1). Including all the lines contradicts the goal of producing a compact triangulation.

Increasingly triangulated approximations are used in surface visualization and interaction, where the size of the triangulation determines its rendering time. In this paper we will argue that we can treat each structural line as the spine or skeleton of a region around it. We then rank each line by a measure of the local region. Only lines with high measures are explicitly included in the triangulation. Section 2 describes the triangulation process and how lines are used. Section 3 shows how we find the structural lines by local curvature analysis of the surface and introduces various measures of the surface regions. Section 4 describes experiments that show how the significance measures affect the triangulation results.

2 Triangulation

In the original work in this area[3], 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 include in each triangle the point (the "worst" point) with most vertical error in the current approximation. [3] proceeded by inserting every "worst" point if it exceeded a desired error tolerance, continuing until all points were within this tolerance. [5] 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.

The second innovation of [3] is not often used since 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 Curvature Descriptions

Like [4, 10], we find determine the local surface curvature description, which yields a local frame where the two curvatures are k_1 and k_2 with k_1 the curvature of maximum absolute value, and t_1 and t_2 are vectors in the local tangent plane pointing the direction of maximal and minimal curvature. To compute curvature, we first locally determine the surface derivatives, which requires smoothing. A curvature description depends on the scale, σ, of the Gaussian smoothing applied before computing derivatives. At each point, we determine whether a the maximum curvature k_1 at the point is locally maximal, in the associated principal direction. We use non-maximum suppression to identify these points, looking in the direction of maximal curvature, t_1.

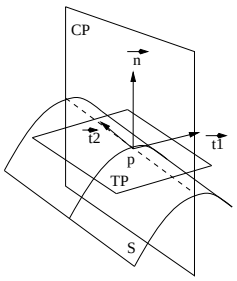

Figure 2: Principal curvatures and directions.

We define *p-lines*, lines connecting points of locally maximal curvature, where that curvature is positive. An *n-line* 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, tracking in the direction of local minimal curvature.

3.1 Snakes: Deforming Large Scale Lines

Because lines at a fine scale ($\sigma = 1.0$) are too numerous, we use the lines at a coarse scale ($\sigma = 2.5$), but these may have been displaced by the smoothing process. The crest on an asymmetrical hill, 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 undesirable "corrective" points near the p-line to model the actual location of the crest. To move the coarse-level line to the location of the fine-level line, we use the "snake" method[6, 1]. 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 curvature field computed at the fine scale. See [9] for more details on how structural lines are corrected.

3.2 Significance of Structural Lines

There still remain lines at the coarse scale that may not contribute to the approximation. We would like to use a measure along the line or of the surrounding region to assess the significance of the line before it is included in the triangulation. We might use the local vertical, z, difference along a line as the significance of the line, but this is only local. Alternatively, since curvature is our important feature, we can explore using the Gaussian curvature of the surface. The sign of Gaussian curvature divides the surface into hyperbolic (negative), elliptic (positive) and parabolic (zero) regions. Typically the lines of zero Gaussian curvature (parabolic lines) partition the surface, but are too numerous and do not correspond to natural regions around the features.

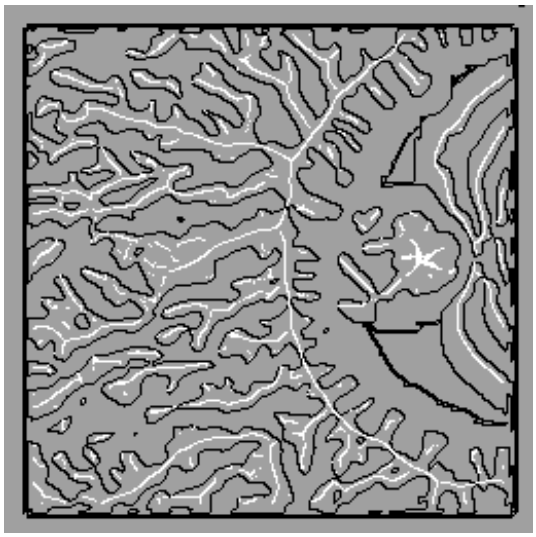

Figure 3: P-lines (white) and zero crossings in the maximal direction, $\sigma = 3.0$

Each p-line is locally convex transverse to its tangent direction. Following the transverse lines of curvature (direction t_1), we encounter loci where the normal curvature in this direction passes through zero. We can track curves from each point on a structural line along lines of curvature. This creates a domain around the p-line where the maximal curvature is positive. Similarly, n-lines have regions of negative maximal curvature surrounding them. The boundaries of these regions form closed loops on the surface. Since the p-lines are elongated, this, in effect, creates a local cylindrical approximation to the p-line or n-line, with a non-circular cross-section. Figure 2 shows a surface with a p-line at the top; the transverse lines are in the direction of maximal curvature locally (t_1).

Along the p-lines, the Gaussian curvature can pass from positive to negative and back, generating parabolic lines, but we ignore these fluctuations. Only zero-crossings along the curvature lines transverse to the direction of the p-lines are counted. Figure 3 shows the p-lines found at a coarse scale plotted over the zeroes of maximal curvature.

Before tracking curvature regions, we segregate p-lines into extrema-bounded subsets by stepping along the linked p-lines and breaking the line wherever there is a local maximum (a peak) or a local minimum (a pass). Tracking transverse to p-lines produces the regions shown in Fig. 4.

We can associate with each region some measure of the surface as its significance value. There are several options, each depending on a local differential property:

- sum of maximal curvature
- sum of magnitude of the gradient
- sum of plan area of the region
- sum of the local curvature

The first option, summing the total maximal curvature, measures the amount of maximal curvature $|k_1|$

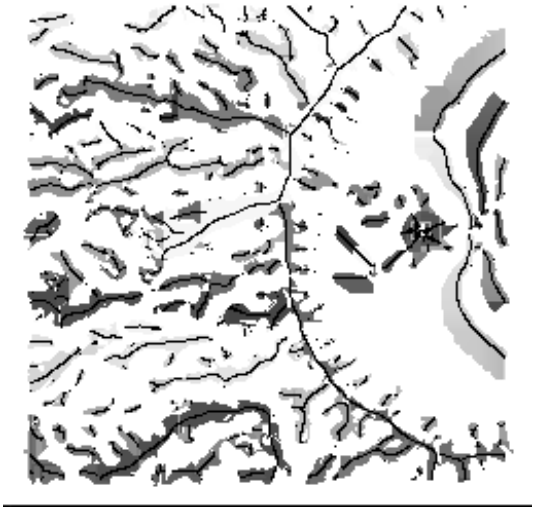

Figure 4: Regions of positive normal curvature surrounding p-lines. Each region is labeled with the number of the p-line it surrounds. p-lines are shown in black.

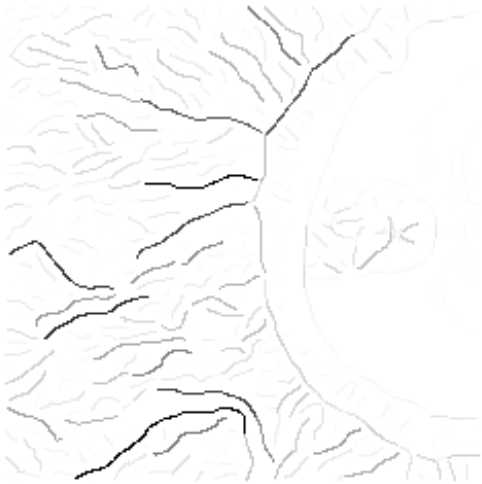

Figure 5: Total maximal curvature: darkness is proportional to the sum of the maximal curvature over the surrounding region.

in the total area; a more "creased" region should have a higher value. Figure 5 shows this scalar measure on both p-lines and n-lines; darkness is proportional to total maximal curvature across the entire region around the line. The second, magnitude of the gradient, estimates surface area by assuming that the value of the surface area is proportional to the magnitude of the gradient: $\sqrt{1 + z_x^2 + z_y^2}$ (Fig. 6). The third measure uses the projected "plan" area of the region (Fig. 7). The fourth estimates the "creasiness" of the line; it ascribes to the line the sum of maximal curvature only

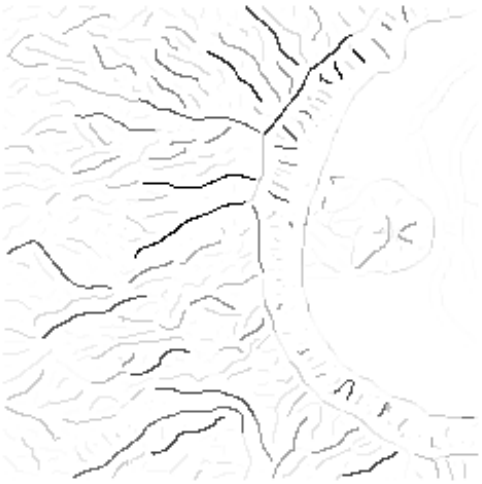

Figure 6: Total surface area: darkness is proportional to the sum of the magnitude of the gradient over the surrounding region.

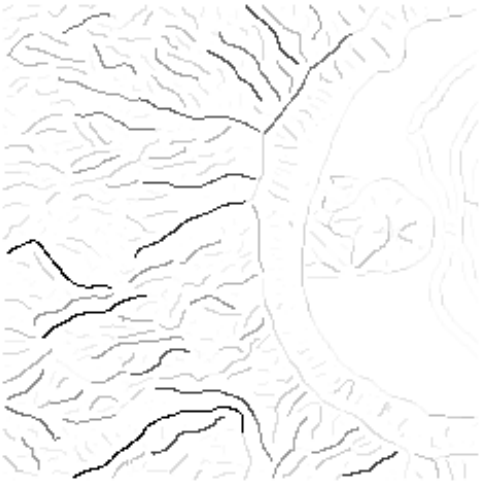

Figure 7: Total projected area: darkness is proportional to the plan area of the region around the line.

along the line, as shown in Fig. 8.

We ascribe the measure of a line to all points in the line and then link lines using hysteresis thresholding, eliminating small lines. Most of these measures reject the small p-lines found in the (flat?) basin of the Crater Lake, appearing as long curves in the right of Figure 1. These lines do not improve the approximation of the triangulation. However, visual comparison of the measures reveals differences, but their effect on the surface approximation can only be determined by experiment.

4 Experiments

To determine whether including structural lines can improve the resulting triangulation, we compare tri-

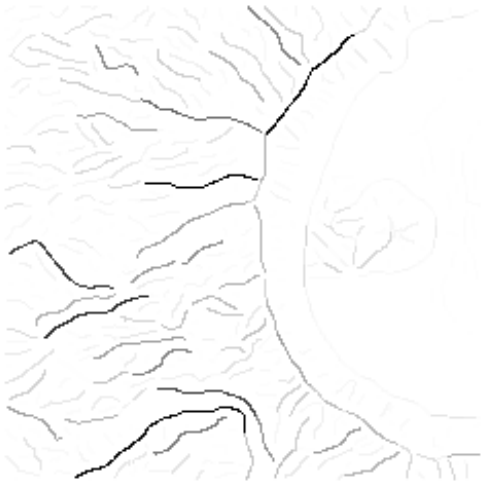

Figure 8: The sum of local maximal curvature along the structural line.

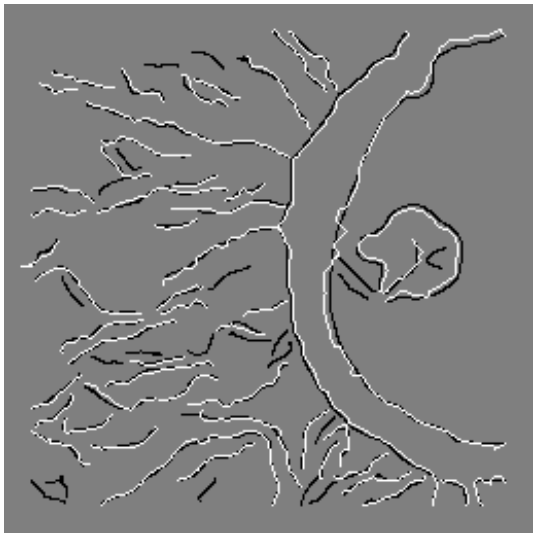

Figure 9: Lines at $\sigma = 2.5$ in black with the corrected result in white for $\sigma = 1.0$

angulations produced by pure "greedy" triangulation, with no lines (called `no lines`), with "greedy" constrained Delaunay triangulation (CDT) with a variety of structural lines. Our experiments included lines at fine scale (noted by `fine`) and lines derived at a coarse scale and then deformed ("snaked") to a fine scale. Unless described as `fine`, structural lines have been derived at a coarse scale and then snaked. The various feature lines are:

- maximal curvature, along the line (`local sum`), as in Fig. 8
- maximal curvature, along the line, at fine scale (`local sum fine`)
- maximal curvature, pointwise, (`local`)
- maximal curvature, at fine scale (`local fine`)

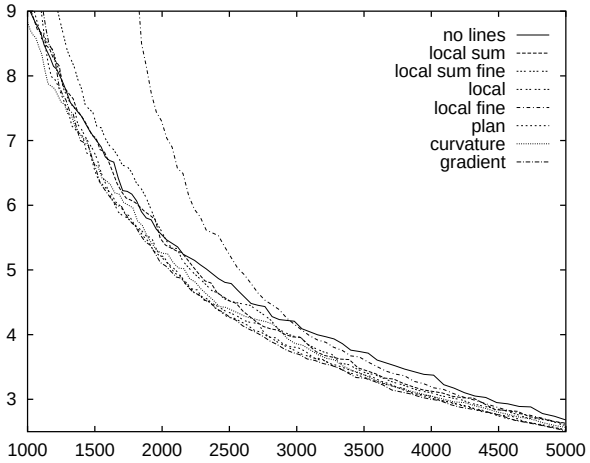

Figure 10: RMS error vs. points.

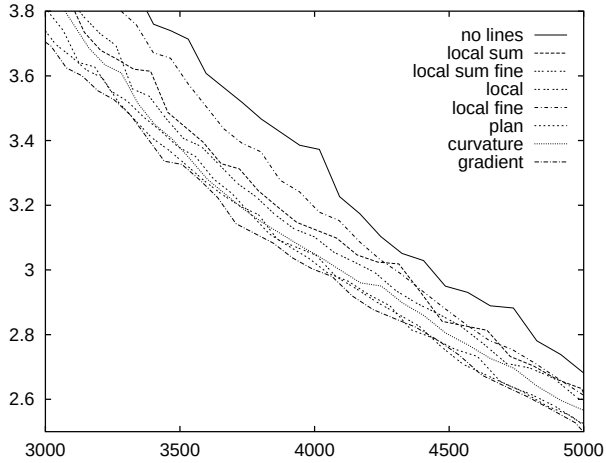

Figure 11: RMS error vs. points (close-up).

- plan area of the region (`plan`), as in Fig. 7
- maximal curvature, summed over the region (`curvature`), as in Fig. 5
- gradient estimate of surface area of the region (`gradient`), as in Fig. 6

The coarse scale features for `local` are shown in black in Fig. 9; the white curves are corrected to fine scale using snakes to the final curves used in `local`. In both of these, only the pointwise maximal curvature is used in linking, in contrast to `local sum`, where the points on the line get the *sum* along the line.

Our scalar measures provide metric criteria 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.

Tests were run on the Crater Lake DEM, running the batched greedy triangulation (subject to constraints provided by structural lines). 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 (RMS) error versus the number of points in Fig. 10, at greater magnification in Fig. 11, and very large in Fig. 12.

The figures showing structural lines in this paper are of a central 257x257 region from the Crater Lake DEM. The results shown here are from a larger area containing this region, with additional gradually sloping terrain. We consider two means of comparison. First, for a particular target RMS, we can find the lines that use the smallest number of points. All triangulations using initial skeletons from the curvature data improve upon the triangulation with no lines, expect for large RMS error. Among the structural lines, the `gradient` measure worked best, needing only 85-88% as many points for comparable RMS error.

Second, for a given number of points, we find the type of structural line that yields the smallest RMS error. Almost all triangulations using initial skeletons from the curvature data improve upon the triangulation with no lines.

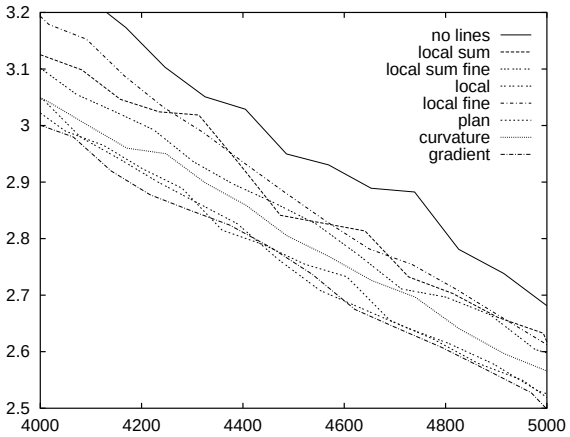

Figure 12: RMS error vs. points (very close-up).

The results can be presented in a tabular form, as shown in Tables 1 and 2. Table 2 shows the ratio of the RMS error for each type of line over the RMS error using no lines. The values are not measured at a point, but represent the *average* ratio of the RMS error from the given value up to 5000. The entry for 1000 is the average ratio from 1000 to 5000. The `gradient` produces an RMS 92% of `no lines`, on average.

Table 1 shows the ratio of the number of points for each type of line to the number for no lines. An entry under x is not just a point sample of the ratio at RMS=x, but shows the average ratio for the range x to $x + 1.0$. The `gradient` uses only 88% as many points as `no lines`, on average.

Table 3 shows the ratio of the number of points for each type of line to number for no lines, using the mean absolute error, which reduces the effect of large errors compared with RMS error. The `gradient` uses only 84% as many points as `no lines`, on average.

5 Discussion

Using structural lines as a skeleton for the triangulation reduces the number of points needed to achieve a

RMS	5.1	4.5	3.9	3.3	2.7
local sum	0.982	0.961	0.933	0.921	0.947
plan	0.908	0.889	0.875	0.878	0.898
curvature	0.935	0.914	0.912	0.902	0.909
gradient	0.903	0.885	0.864	0.858	0.887
local fine	1.175	1.104	1.230	0.968	0.959
local	0.910	0.895	0.881	0.875	0.895
local sum fine	0.996	0.952	0.940	0.920	0.925

Table 1: The ratio of the number of points in the triangulation produced using each type of structural line to the number of points in a triangulation with no lines, for various RMS values.

Points	1000	2000	3000	4000
local sum	0.964	0.951	0.950	0.963
plan	0.924	0.911	0.915	0.927
curvature	0.936	0.929	0.927	0.940
gradient	0.922	0.903	0.906	0.921
local fine	-	1.019	0.969	0.968
local	0.932	0.913	0.913	0.928
local sum fine	0.988	0.947	0.940	0.952

Table 2: The ratio of the RMS of the triangulation of each type of structural line to that using no lines; each entry is the average RMS from the given number up to 5000. The "local fine" starts with 1743 points.

given error. It is clear, also, that the `gradient` measure selects lines that yield the most compact triangulations, both for attaining a specific RMS error with minimal points, and for minimizing the RMS error for a given number of points.

Further work will address the issue of how the areal and volumetric measures of the terrain approximation relate to the linear scale used in the curve approximation for the structural lines. Moreover, we still must determine how to understand the perceptual fidelity of the approximation: how structure is preserved.

We have advanced a method of assigning a significance to structural lines by, first, finding a region of local positive maximal curvature surrounding structural lines, and, second, aggregating various measures of the region for ranking the structural line. We have built triangulations using them, and shown how they can improve the selection of initial surface structure, leading to more compact triangulations.

MAE	5.0	4.2	3.6	3.3	2.7
local sum	1.010	0.986	0.974	0.936	0.949
plan	0.848	0.843	0.850	0.840	0.882
curvature	0.888	0.903	0.910	0.899	0.914
gradient	0.863	0.828	0.825	0.821	0.863
local fine	0.943	1.144	1.024	0.921	0.896
local	0.863	0.844	0.835	0.828	0.866
local sum fine	1.003	0.988	0.957	0.914	0.917"

Table 3: The ratio of the number of points in the triangulation produced using each type of structural line to the number of points in a triangulation with no lines, for various mean absolute errors (MAE).

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] Xin Chen and Francis Schmitt. Surface modelling of range data by constrained triangulation. *Computer-Aided Design*, 26(8):632–645, 1994.

[3] 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.

[4] 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.

[5] 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.

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

[7] 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.

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

[9] James J. Little and Ping Shi. Structural lines, TINs, and DEMs. In *International Symposium on Spatial Data Handling '98*, pages 627–636, July 1998.

[10] 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.

[11] 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.