Terrain Modelling from Contours

Christopher Gold and **Maciej Dakowicz**
Department of Land Surveying and Geo-Informatics
Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong
Tel: (852)2766-5955; Fax: (852)2330-2994
christophergold@voronoi.com; lsmaciek@polyu.edu.hk

KEYWORDS: Terrain Modelling; Contours; Skeleton; Interpolation; Delaunay Triangulation; Voronoi Diagram.

ABSTRACT

Good quality terrain models are becoming more and more important, as applications such as runoff modelling are being developed that demand better surface orientation information than is available from traditional interpolation techniques. A consequence is that poor-quality elevation grids must be massaged before they provide useable runoff models.

Rather than using direct data acquisition, this project concentrated on using available contour data, for two reasons. Firstly, despite modern techniques, contour maps are still the most available form of elevation information. Secondly, manual contour tracing has imposed a subjective interpretation of the form of the landscape that is lost with automation, yet which is of considerable value.

The maximum slope of the terrain is perpendicular to the contour, and this permits us to visualize the relationships between pairs of contours. With care this may be modelled by triangulation methods, as the spatial relationships can be preserved, although standard grid interpolation methods based on "n nearest neighbours" often have problems. However, whenever we have relationships between portions of the same elevation contour, such as in peaks, pits or valley heads, our interpretation based on triangle slope is insufficient - we get "flat triangles". In this case we need to re-examine our spatial model.

The usual triangulation approach, the Delaunay triangulation, is effective because it is locally stable - a property based on its dual, the Voronoi diagram. These two spatial structures have been much studied by workers in the field of computational geometry - largely in terms of efficient calculation, but also in terms of their properties. In particular, recent work on the automatic reconstruction of curves from point samples, and the generation of medial axis transforms (skeletons) has greatly helped in the visualization of the relationships between sets of boundaries, and families of curves. This provides us with tools to enrich our original contour data for "flat triangles". The insertion of skeleton points in these cases guarantees the elimination of all flat triangles. Additional assumptions about the local uniformity of slopes, either along or across valleys and other features, give us enough information to assign elevation values to these skeleton points. If required, appropriate interpolation techniques may generate an elevation grid for visualization purposes that preserves reasonable slopes at all points on the model - even at the data points themselves - and that are faithful to the input data. In addition, the algorithms used are only moderately more complex than the underlying Voronoi diagram or Delaunay triangulation. The result provides us with a surprisingly realistic model of the surface - that is, one that conforms well to our subjective interpretation of what a real landscape should look like.

INTRODUCTION

This paper concerns the generation of interpolated surfaces from contours. While this topic has been studied by many people (including the first author) for over 20 years, this work is interesting for a variety of reasons. Firstly, contour data remains the most readily available data source. Secondly, valid theorems for the sampling density along the contour lines have only just been discovered.

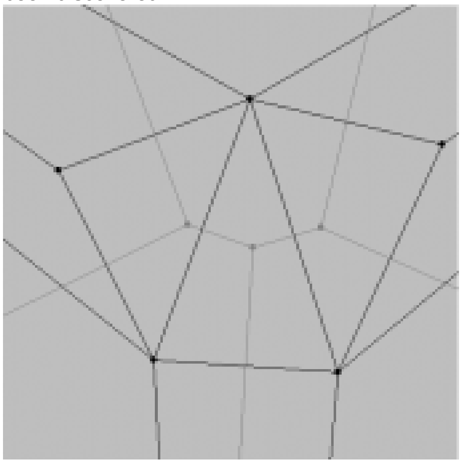

Figure 1: Delaunay triangulation and Voronoi diagram

Thirdly, the same publications provide simple methods for generating the medial axis transform, or skeleton, which definitively solves the "flat triangle" problem which often occurs when triangulating contour data, by inserting additional points from this skeleton. Fourthly, the problem of assigning elevation values to these additional ridge or valley points can be resolved, using the geometric properties of this skeleton, in ways that may be associated with the geomorphological form of the landscape. In addition, comparisons of the methods used in a variety of weighted-average techniques throw a lot of light on the key components of a good interpolation method, using three-dimensional visualization tools to identify what should be "good" results – with particular emphasis being placed on reasonable slope values, and slope continuity. This last is often of more importance than the elevation itself, as many issues of runoff, slope stability and vegetation are dependent on slope and aspect – but unfortunately most interpolation methods can not claim satisfactory results for these parameters.

GEOMETRIC PRELIMINARIES.

The methods discussed here depend on a few fundamental geometrical constructs that are now fairly well known – the Voronoi diagram and its dual, the Delaunay triangulation, as shown in Fig. 1. The first is often used to partition a map into regions closest to each generating point; the second is usually used as the basis for triangulating a set of data points, as it is guaranteed to be locally stable. It may easily be constructed using its "empty circumcircle" property – this circle is centred at the Voronoi node associated with each triangle. As will be seen later, these nodes, and circles, are associated with the skeleton, or "medial axis transform".

GENERATION OF RIDGE AND VALLEY LINES.

Amenta, Bern and Eppstein (1998) examined the case where a set of points sampled from a curve, or polygon boundary, were triangulated, and then attempted to reconstruct the curve. They showed that this "crust" was formed from the triangle edges that did not cross the skeleton, and that if the sampling of the curve was less than 0.25 of the distance to the skeleton the crust was guaranteed to be correct. Their algorithm consisted of inserting all the Voronoi vertices into the diagram. Gold (1999) and Gold and Snoeyink (2001) simplified the approach, showing that, in every Voronoi/Delaunay edge pair, one edge could be assigned to the crust and the other to the skeleton. Fig. 2 shows the results for a simple contour. Thus skeleton points may be inserted into the original diagram, or not, as needed.

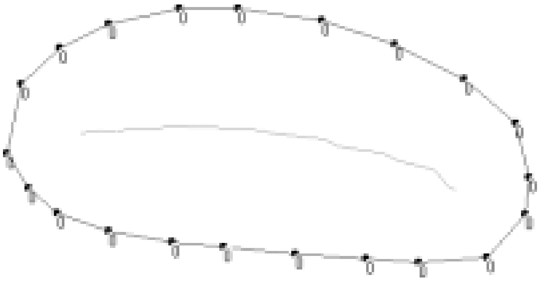

Figure 2: Crust and skeleton of a simple polygon

Figure 3: Contour data points

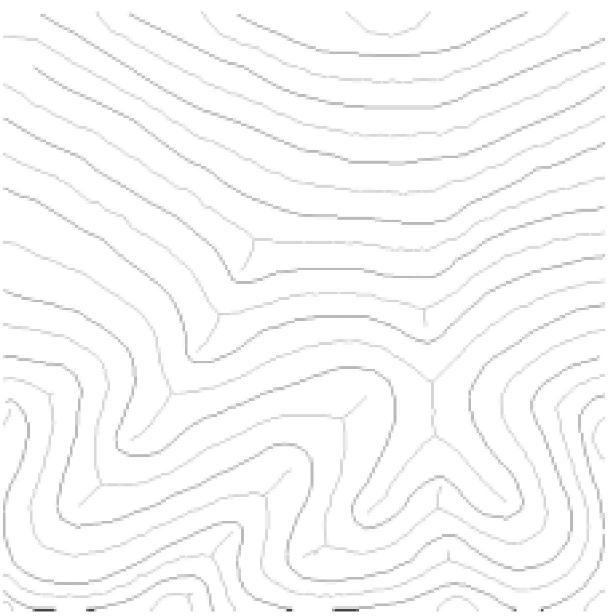

Figure 4: Crust and skeleton of Fig. 3 data

In our particular case, the data is in the form of contour lines, that we assume are sufficiently well sampled – perhaps derived from scanned maps. Despite modern satellite imaging, much of the world's data is still in this form. An additional property is not sufficiently appreciated – they are subjective, the result of human judgement at the time they were drawn. Thus they are clearly intended to convey information about the form of the surface – and it would be desirable to preserve this, as derived ridges and valleys.

Fig. 3 shows our raw data set, and Fig. 4 shows the resulting crust and skeleton. Fig. 5 shows only those skeleton points that provide unique information – ridge and valley lines that separate points on the same contour, rather than merely those points that separate adjacent contours. Aumann et al., (1991) produced somewhat similar results by raster processing.

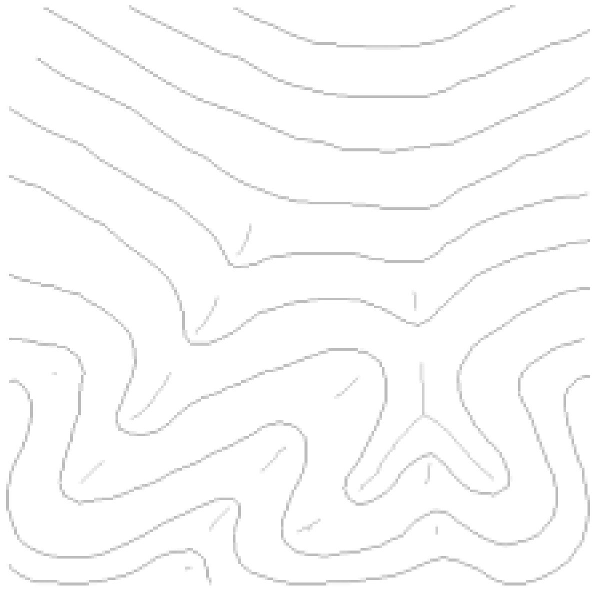

Figure 5: Skeleton branches from Fig. 4

Figure 6 shows a close-up of the test data set, illustrating a key point of Amenta, Bern and Eppstein's work: if crust edges (forming the contour boundary) may not cross the skeleton, then inserting the skeleton points will break up non-crust triangle edges. In particular, if the skeletons between different contours are ignored, then the remaining branch skeleton points will eliminate all "flat" triangles formed by triangles connecting points at the same elevation. Thus ridge and valley lines are readily generated automatically. The challenge is to assign them meaningful elevation values. The same is true in the case of closed summits.

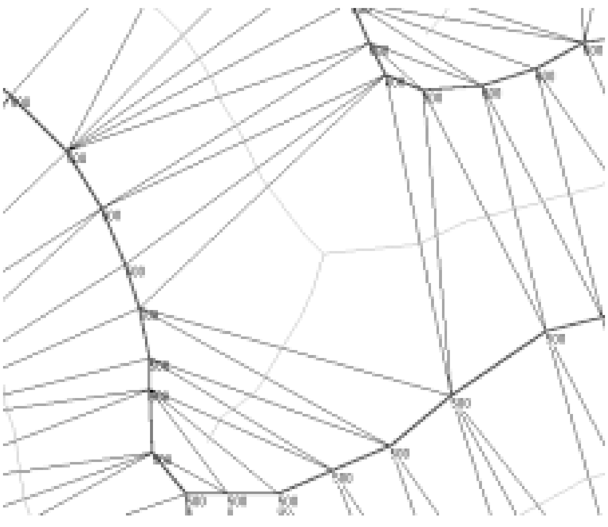

Figure 6: Skeleton and "flat triangles"

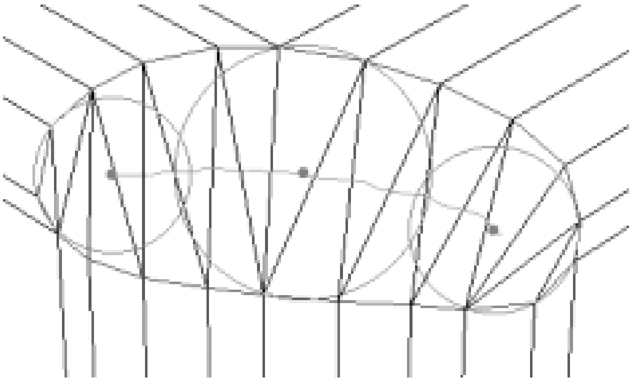

Figure 7: Skeleton and circumcentres

Two techniques have been developed, each with its own physical interpretation. The first, following Thibault and Gold (2000), uses Blum's (1967) concept of height as a function of distance from the curve or polygon boundary, with the highest elevations forming the crest at the skeleton line.

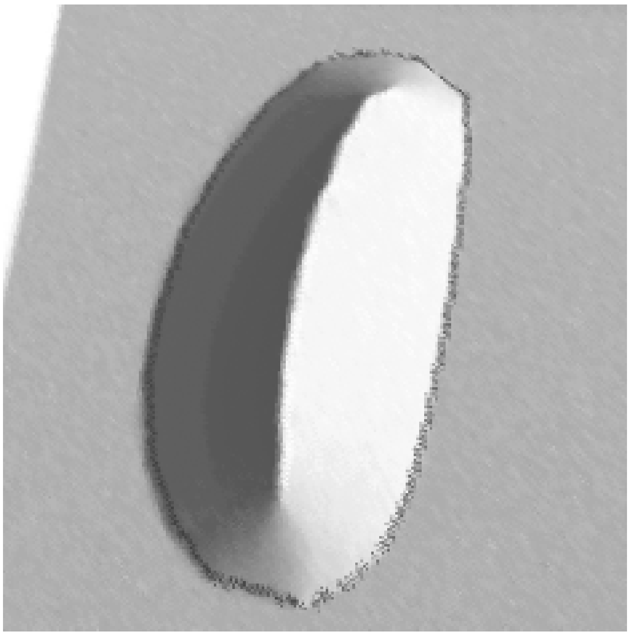

Figure 8: Elevation model of Fig. 7

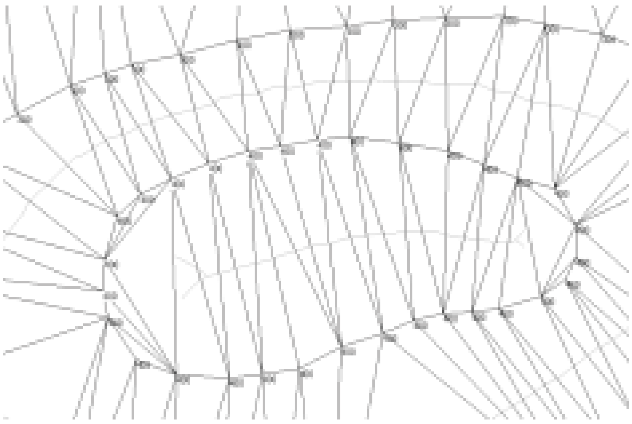

Figure 9: Skeleton of a summit

This is illustrated in Figs. 7 and 8, where points on a simple closed curve are used to generate the crust and skeleton. In Fig. 7, the circumcentres of the skeleton points are given an elevation equal to the circumradius. The resulting interpolated model is shown in Fig. 8. This model is based on the idea that all slopes are identical, and thus the radius is proportional to the height of the skeleton point. Of course, in the case of a real summit as in Fig. 9, the slope would initially be unknown, and would be estimated from the circumradius of the next contour level down.

Figure 10: Estimating skeleton heights from circumradii

In the case of a ridge or valley, the circumradius may also be used, as in Fig. 10, to estimate skeleton heights based on the hypothesis of equal slopes. The larger circle, at the junction of the skeleton branches, has a known elevation – half way between the contours – and may be used to generate the local slope. The elevation of the center of the smaller circle is thus based on the ratio of the two radii. For more details see Thibault and Gold (2000).

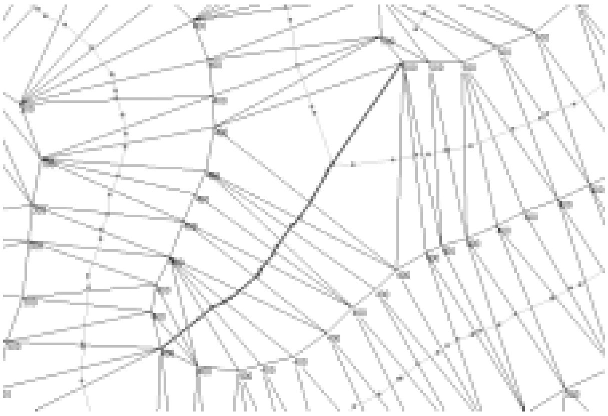

Figure 11: Estimating skeleton heights from ridge or valley lengths

While this method is always available, it is not always the preferred solution where constant slope down the drainage valley, rather than constant valley-side slope, is more appropriate. In a second approach, illustrated in Fig. 11, the line of the valley is determined by searching along the skeleton, and heights are assigned based on their relative distance along this line. This may be complicated where there are several valley branches – in which case the longest branch is used as the reference line. This involves careful programming of the search routines, although the concept is simple. In practice, an automated procedure has been developed, which uses the valley length approach where possible, and the side-slope method when no valley head can be detected, such as at summits and passes.

COMPONENTS OF AN INTERPOLATION MODEL.

On the basis of a sufficient set of data points, we now wanted to generate a terrain model with satisfactory elevations and slopes, as the basis of a valid rainfall runoff model. Our approach was to interpolate a height grid over the test area, and to view this with an appropriate terrain visualization tool – in this case Genesis II, available from www.geomantics.com. We feel that 3D visualization has been under-utilized as a tool for testing terrain modeling algorithms, and the results are often more useful than a purely mathematical, or even statistical, approach.

We have restricted ourselves to an evaluation of several weighted-average techniques, as there are a variety of techniques in common that can be compared. All of the methods were programmed by ourselves – which left out the very popular Kriging approach, as too complicated. Nevertheless, many aspects of this study apply to this method as well, since it is a weighted-average method, with the same problems of neighbour selection, etc., as the methods we attempted.

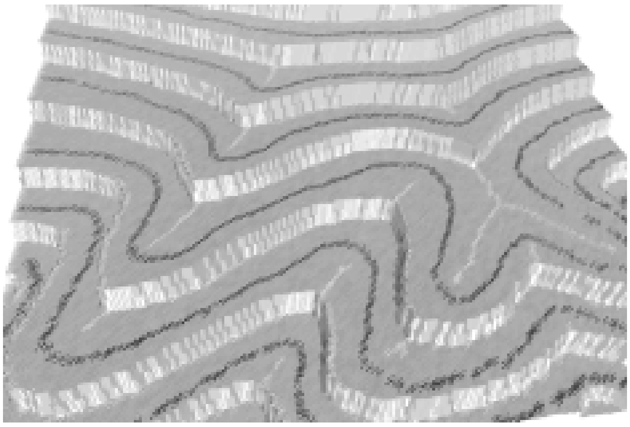

Figure 12: Interpolation from nearest point

In general, we may ask about three components of a weighted-average interpolation method. Firstly: what is the weighting process used? Secondly: what is the set of neighbours used to obtain the average? Thirdly: what is the elevation function being averaged? (Often it is the data point elevation alone, but sometimes it is a plane through the data point incorporating slope information as well.)

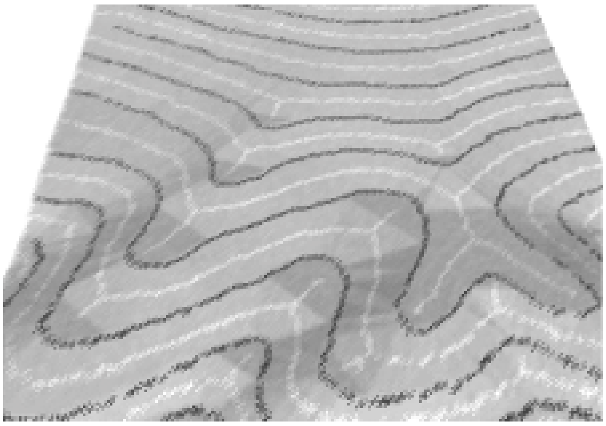

Figure 13: Interpolation from Delaunay triangulation

The simplest possible technique, useful on occasion, is merely to give each grid node (if a grid is being created) the height of the nearest data point. While trivial, it is valuable for a variety of applications, such as image rectification, rainfall estimation, and others. All grid cells falling within the Voronoi cell of a particular data point are assigned its elevation. Fig. 12 shows the result for our contour data set: the skeleton can be seen to separate each plateau around a contour.

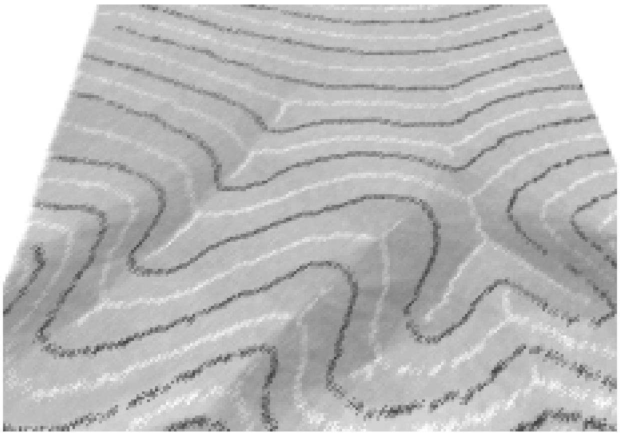

Figure 14: Adding skeleton points to Fig. 14

The next most simple weighted-average model is the triangulation, using the Delaunay triangulation described previously. Fig. 13 shows the result, including the skeleton draped over the flat triangles. Fig. 14 shows the improvement when estimated skeleton points are added.

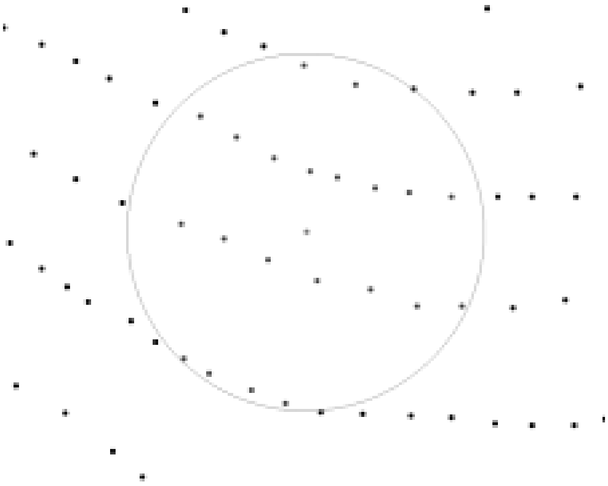

Figure 15: Selecting neighbours using a counting circle

The other weighted average models that were tested were the traditional gravity model, and the more recent “area-stealing” or “natural neighbour” or perhaps more properly “Sibson” interpolation methods (Sibson 1980, Watson and Philip 1987, Gold 1989). Here the number of neighbours used may vary. In the case of the gravity model the weighting of each data point used is inversely proportional to the square of the distance from the data point to the grid node being estimated, although other exponents have been used. There is no obvious set of data points to use, so one of a variety of forms of “counting circle” is used, as in Fig. 15. When the data distribution is highly anisotropic there is considerable difficulty in finding a valid counting circle radius.

Figure 16 shows the resulting surface for a radius of 5 (about a quarter of the map). Data points form bumps or hollows. If the radius is reduced there may be holes in the surface where no data is found within the circle. If the radius is increased the surface becomes somewhat flattened, but the bumps remain.
The result depends on the radius, and other selection properties, being used. Clearly, in addition, estimates of slope would be very poor, and very variable.

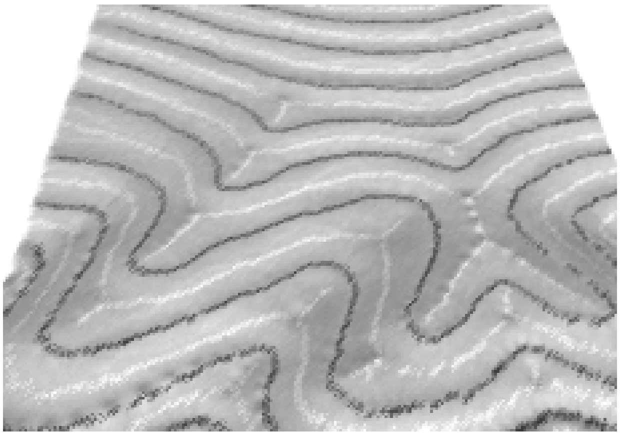

Figure 16: Interpolation using the gravity model

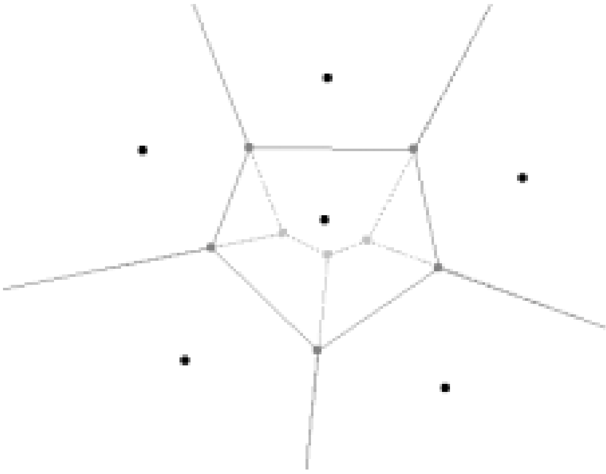

Figure 17: Sibson interpolation

The Sibson method, illustrated in Fig. 17, is based on the idea of inserting each grid point temporarily into the Voronoi diagram of the data points, and measuring the area stolen from each of a well-defined set of neighbours. These stolen areas are the weights used for the weighted average. The method is particularly appropriate for poor data distributions, as illustrated in Fig. 18, as the number of neighbours used is well defined, but dependent on the data distribution.

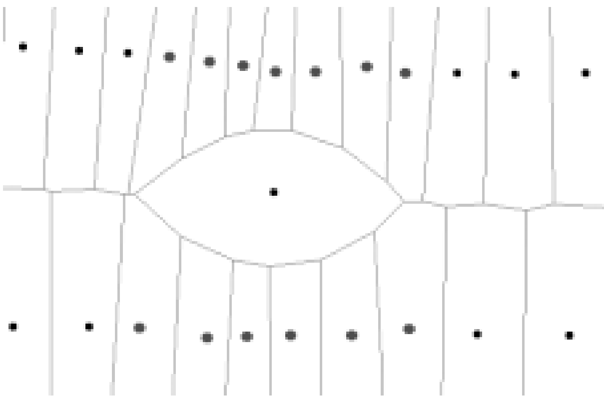

Figure 18: Neighbour selection using Voronoi neighbours

Fig. 19 shows the results. It behaves well, but is angular at ridges and valleys. Indeed, slopes are discontinuous at all data points (Sibson, 1980). One solution is to re-weight the weights, so that the contribution of any one data point not only becomes zero as the grid point approaches it, but the slope of the weighting function approaches zero also (Gold, 1989).

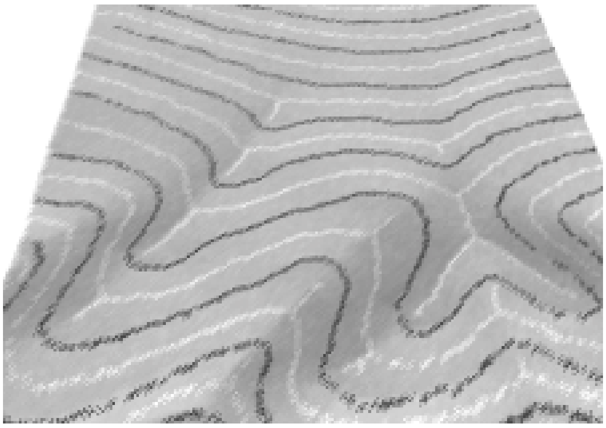

Figure 19: Interplation using Sibson interpolation

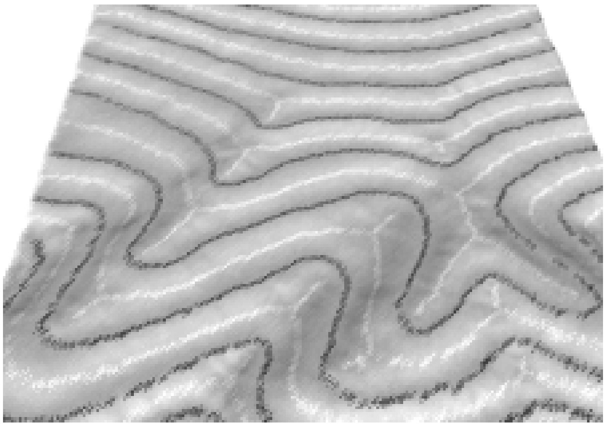

Figure 20: Adding smoothing to Fig. 19

Fig. 20 shows the effect of adding this smoothing function. While the surface is smooth, the surface contains undesirable “waves” – indeed, applying this function gives a surface with zero slope at each data point. This is sometimes called the “wedding-cake effect”, and is also apparent in Fig. 16.

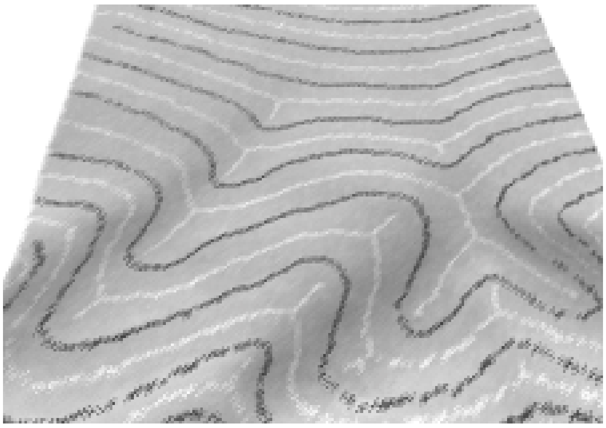

Figure 21: Sibson interpolation using slopes at data points

SLOPES – THE IGNORED FACTOR.

This brings us to a subject often ignored in selecting a method for terrain modeling – the slope of the generated surface. In real applications, however, accuracy of slope is often more important than accuracy of elevation – for example in runoff modeling, erosion, insolation, etc. Clearly an assumption of zero slope, as above, is inappropriate. However, in our weighted average operation we can replace the height of a neighbouring data point by the value of a function defined at that data point – probably a planar function involving the data point height and local slopes. Thus at any grid node location we find the neighbouring points and evaluate their planar functions for the (x, y) of the grid node. These z estimates are then weighted and averaged as before.

Fig. 21 shows the result of using Sibson interpolation with data point slopes. It gives an apparently excellent result – and looks even better than when the smoothing function is added to it.

While it is impossible to show the results of all our experiments in this paper, in order to see what was happening we used the method of (Burrough and Mcdonnell, 1988) to calculate slopes and profile curvature for grids created from various combinations of our available weighted-average methods. Somewhat surprisingly, the version without smoothing gives more consistent regions of coherent slopes, indicating that the smoothing function adds unwanted undulations to the surface. However, examination of the profile curvature map shows that without smoothing there are folds in the surface at the contour lines – as would be expected – although the effects are minor. Adding slopes to the simple TIN model (i.e. using the position in the triangle to provide the weights) produced results that were almost as good as the Sibson method where the sample points were closely spaced along the contours, but the Sibson method is much superior for sparser data, or where the points do not form contour lines. The gravity model does not provide particularly good slope estimates, but even here including the data point slope function produces a significant improvement.

CONCLUSIONS.

From our work, several broad generalizations may be made. To produce good surface models, with reasonable slopes, from contour maps the single most valuable contribution is the addition of skeleton points along the ridges, valleys, pits, summits and passes. These are guaranteed to eliminate flat triangles. Height estimates at these points may be based either on longitudinal or lateral slope consistency, depending on the physical model desired, or the detection of valley-head information.

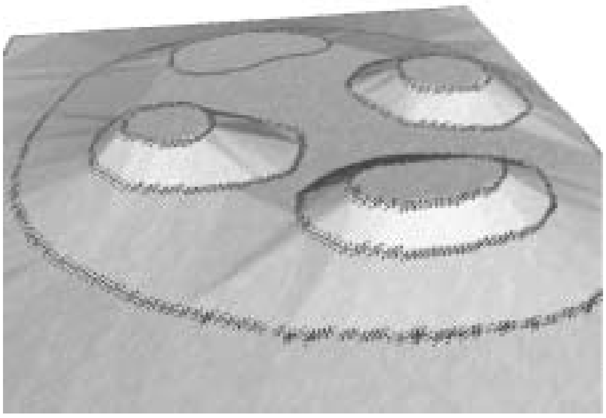

Figure 22: Triangulation of several small hills

The second most important contribution is the addition of slope information at the data points, and its use in the weighted average. Even poor interpolation methods are significantly improved. Also important is the selection of a meaningful set of neighbours around the grid node to be estimated.

Of lesser importance is the particular interpolation method used, although this statement is highly dependent on the data distribution and density. Gravity models in general should be avoided if possible. Surprisingly, mathematically guaranteed slope continuity is not usually critical, although we are continuing to work on an improved smoothing function that guarantees both slope continuity and minimum curvature – probably based on the work of Anton et al. (1998). Nevertheless, the moral is clear: both for finding adjacent points and skeleton extraction, a consistent definition of neighbourhood is essential for effective algorithm development.

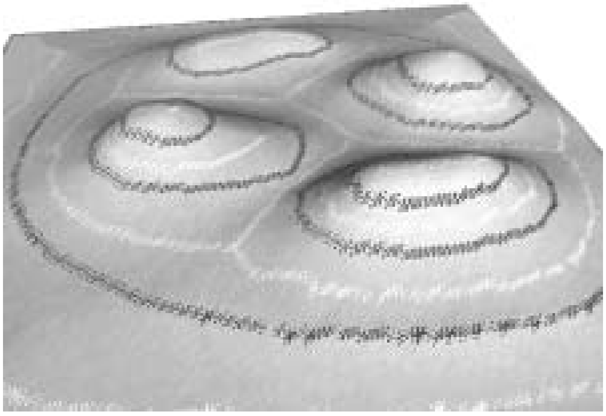

Figure 23: Sibson interpolation with slopes

We conclude with another imaginary example. Fig. 22 shows four small hills defined by their contours, modelled by a simple triangulation. Fig. 23 shows the result using Sibson interpolation, slopes and skeletons. Skeleton heights were obtained using circumcircle ratios, as no valley-heads were detected. While our evaluation was deliberately subjective, we consider that our results in this case, as with the previous imaginary valley, closely follow the perceptual model of the original interpretation. Thus, for the reconstruction of surfaces from contours, we believe that our methods are a significant improvement on previous work.

REFERENCES

Amenta, N., M. Bern and D. Eppstein, 1998. The crust and the beta-skeleton: combinatorial curve reconstruction. Graphical Models and Image Processing, 60/2:2, pp. 125-135.

Anton, F.; Gold, C.M. and Mioc, D., 1998, Local coordinates and interpolation in a Voronoi diagram for a set of points and line segments. Proceedings, 2nd Voronoi Conference on Analytic Number Theory and Space Tillings; Kiev, Ukraine, pp. 9-12.

Aumann, G., H. Ebner, and L. Tang, 1991, Automatic derivation of skeleton lines from digitized contours, ISPRS Journal of Photogrammetry and Remote Sensing, 46 259-268, Elsevier Science Publishers B.V., Amsterdam.

Blum, H., 1967. A transformation for extracting new descriptors of shape. In: W. Whaten Dunn (ed.), Models for the Perception of Speech and Visual Form. MIT Press, pp. 153-171.

Burrough, P. and Mcdonnell, R.A., 1988, Principles Of Geographical Information Systems, 2nd. Ed., 336p. Oxford University Press.

Gold, C.M., 1989. Chapter 3 - Surface interpolation, spatial adjacency and G.I.S. In: Three Dimensional Applications in Geographic Information Systems (J. Raper, ed.), Taylor and Francis, Ltd., London, pp. 21-35.

Gold, C.M., 1999, Crust and anti-crust: a one-step boundary and skeleton extraction algorithm. Proceedings, ACM Conference on Computational Geometry, Miami, Florida, pp. 189-196.

Gold, C.M. and Snoeyink, J., 2001, A one-step crust and skeleton extraction algorithm. Algorithmica, (in press).

Sibson R., 1980, A Vector Identity for the Dirichlet Tessellation, Math. Proc. Cambridge Philos. Soc., 87, pp. 151-155.

Thibault, D. and Gold, C.M. 2000, Terrain Reconstruction from Contours by Skeleton Construction. GeoInfo. v. 4 pp. 349-373.

Watson, D.F. and Philip, G.M., 1987, Neighborhood-based interpolation, Geobyte, 2(2), 12-16.