Extracting Meaningful Slopes from Terrain Contours

Maciej Dakowicz and Christopher Gold

Department of Land Surveying and Geo-Informatics
Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong
Tel: (852) 2766-5955; Fax: (852) 2330-2994
maciej.dakowicz@dtm.prv.pl, christophergold@voronoi.com

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 because, despite modern techniques, contour maps are still the most available form of elevation information. Recent work on the automatic reconstruction of curves from point samples, and the generation of medial axis transforms (skeletons) has greatly helped in expressing the spatial relationships between topographic sets of contours. With these techniques the insertion of skeleton points into a TIN model guarantees the elimination of all "flat triangles" where all three vertices have the same elevation. Additional assumptions about the local uniformity of slopes give us enough information to assign elevation values to these skeleton points. In addition, various interpolation techniques were compared using the enriched contour data. Examination of the quality and consistency of the resulting maps indicates the required properties of the interpolation method in order to produce terrain models with valid slopes. 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.

1. Introduction

This paper concerns the generation of interpolated surfaces from contours. While this topic has been studied for many years [5], [6], [7], [11], [13], [15], the current project 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 recently been discovered [1]. 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 light on the key components of a good weighted-average 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 cannot claim satisfactory results for these properties.

2. 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. The Delaunay triangulation may easily be constructed using its "empty circumcircle" property – this circle is centred at the Voronoi node associated with each triangle.

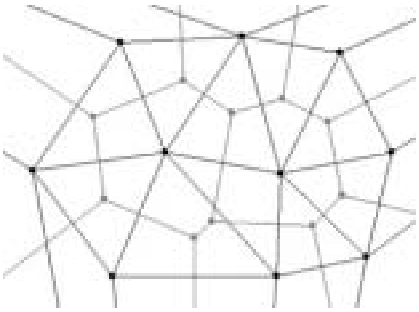

Fig. 1. Delaunay triangulation and Voronoi diagram

The Voronoi diagram and Delaunay triangulation are associated with other geometric structures, known as the crust and the skeleton (or "medial axis transform"), using algorithms introduced by (Amenta *et al.*, 1998). They 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 formed by the remaining Voronoi edges the crust was guaranteed to be correct. Fig. 2 shows the triangulation of several points: the crust is marked with thick black lines and the skeleton with thick grey lines.

Fig. 2. Crust and skeleton of triangulation

(Gold, 1999) and (Gold and Snoeyink, 2001) simplified Amenta's algorithm for the extraction of the crust by showing that, in every Delaunay/Voronoi edge pair, either the Delaunay edge could be assigned to the crust or else the dual Voronoi edge could be assigned to the skeleton. The Delaunay edge belongs to crust when there exists a circle through its two vertices that does not contain either of its associated Voronoi vertices; if not then the corresponding Voronoi edge belongs to the skeleton. In practice their approach boils down to a simple InCircle test applied for each Delaunay/Voronoi edge pair, checking whether the second Voronoi vertex is inside the circle formed by the Delauany edge and the first Vorono vertex, as shown in Fig. 3.

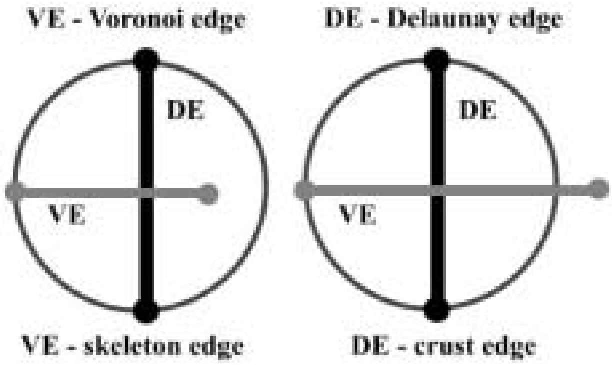

Fig. 3. Results of crust/skeleton test

3. Generation of Ridge and Valley Lines

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 perceived form of the surface at a particular scale – and it would be desirable to preserve this, as derived ridges and valleys.

Fig. 4a shows a small part of a contour map. The contourlines are processed and samples are extracted (Fig. 4b).

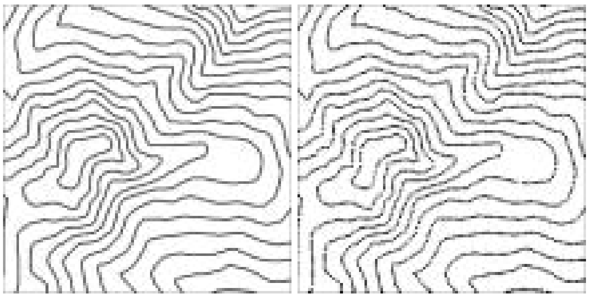

Fig. 4. Contours – a) contourlines; b) data points

Samples are triangulated and a terrain model is generated. Fig. 5a shows a perspective view of terrain model with the contourlines draped over the terrain and Fig. 6 shows a vertical view. In both figures it is readily seen, that in many places, instead of ridges, valleys or summits, flat terraces has been created. These flat features are formed by triangles having all three vertices at the same elevation, these triangles are called flat triangles. The problem of flat triangles is a very well know problem of TIN models.

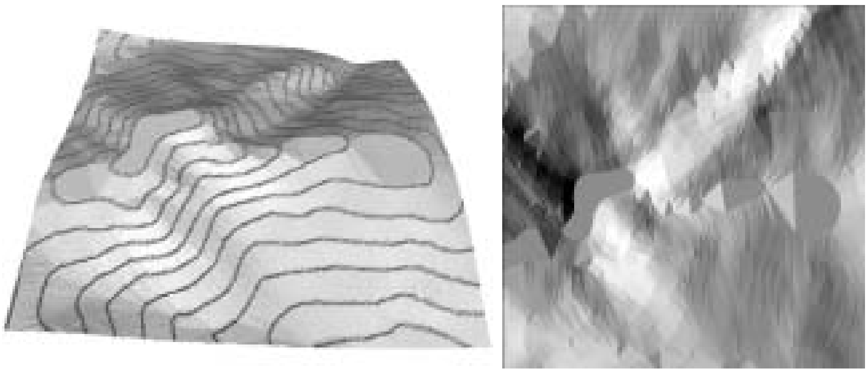

Fig. 5. Terrain model created from simple triangulation.

Figs. 6a and 6b show close-ups of a triangulation created from another data set, the crust and skeletons are drawn and flat triangles are shaded. It is readily seen that skeletons are present in these flat areas and they approximate the shape of a missed ridge (Fig 6a) and a summit (Fig. 6b).

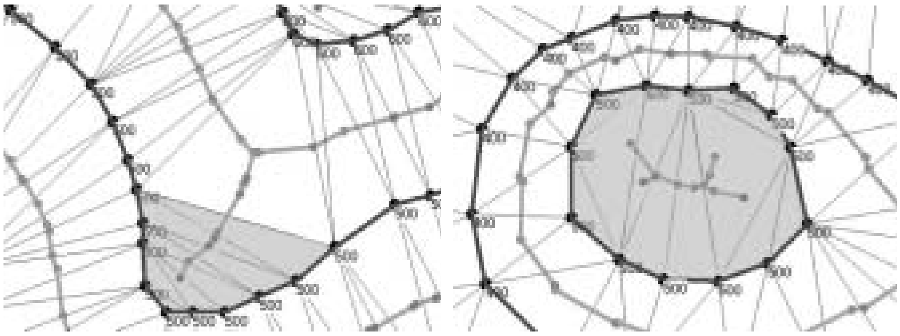

Fig. 6. Skeleton and "flat triangles" – a) ridge; b) summit

To remove flat triangles, further processing of the triangulation is required. To this end the Voronoi diagram is computed from the Delaunay triangulation and the crust and skeleton are extracted - these are shown in Fig. 6a. Aumann *et al.* [3] produced somewhat similar results by raster processing. The extracted crust, due to a sufficient sampling of the contourlines, reconstructs their shape. Skeletons between different contourlines do not provide any interesting information, thus Fig. 6c shows the crust and only skeleton branches. Skeleton vertices of these branches separate points on the same contour and can be used to solve the problem of flat triangles – they can reconstruct ridge and valley lines. Therefore skeleton points can be included in the triangulation, and fig. 6c shows the enriched data set consists of original samples and vertices of skeleton branches.

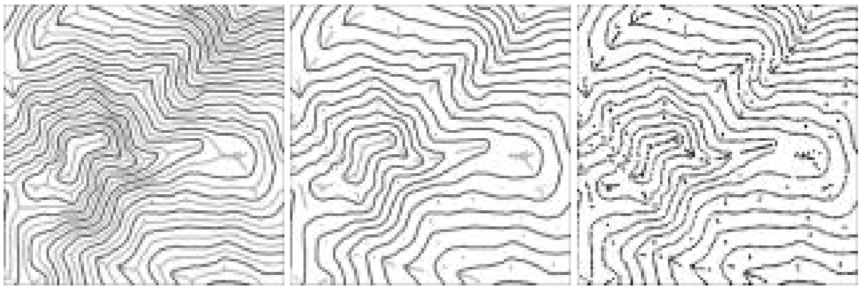

Fig. 7. Crust and skeleton a) crust and skeleton; b) crust and skeleton branches; c) enriched data set

In a Delaunay triangulation all circumcircles must be empty, and the insertion of a Voronoi vertex (circumcentre or skeleton point) will force the deletion of its forming triangle. Thus the insertion of the skeleton point of a flat triangle guarantees that it is replaced by new triangles with the skeleton point as a vertex. The challenge is to assign meaningful elevation values to skeleton points.

Two techniques have been developed for this, each with its own physical interpretation. The first, following Thibault and Gold [14], uses Blum's [4] 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. This is illustrated in Figs. 8a and 8b, where points on a simple closed curve are used to generate the crust and skeleton.

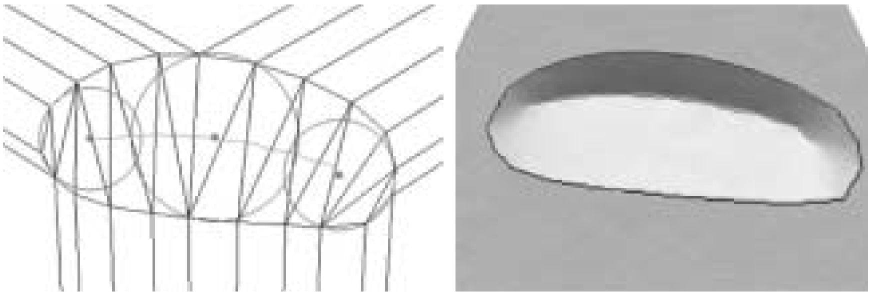

Fig. 8. Triangulation of a summit – a) skeleton and circumcentres; b) elevation model after adding skeleton vertices with assigned height values

In Fig. 8a, the circumcentres of the skeleton points are given a height above the previous contour level equal to the circumradius. The resulting interpolated model is shown in Fig. 8b. 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. 7b, the slope would initially be unknown, and would be estimated using circumradii from the next contour level down – see [14].

In the case of a ridge or valley, the circumradius may also be used, as in Fig. 9a, 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 estimate the local slope. The elevation of the centre of the smaller circle is thus based on the ratio of the two radii. For more details see [14].

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. 9b, 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.

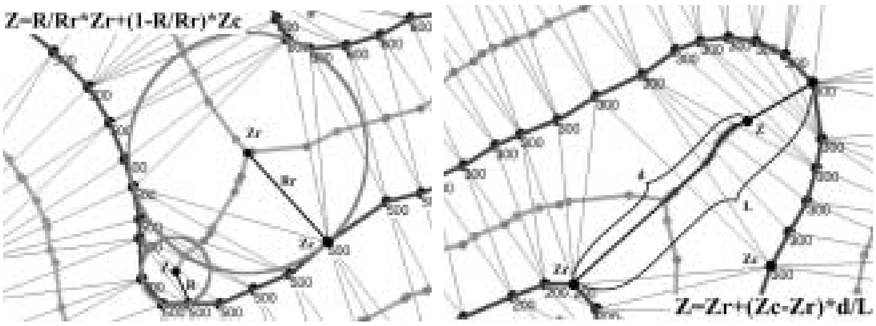

Fig. 9. Estimating skeleton heights – a) from circumradii; b) from valley length

Fig. 10 shows the improved model when estimated skeleton points are added. All flat triangles are removed and ridges, valleys and the summit are reconstructed.

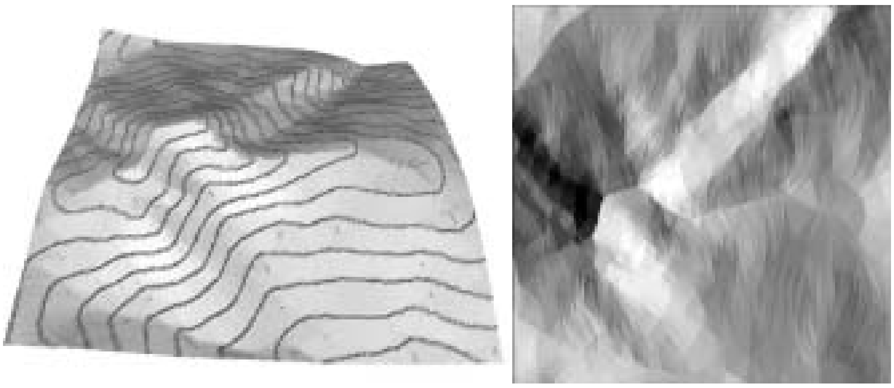

Fig. 10. Enriched terrain model.

4. 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. To obtain perspective views we used Genesis II, available from www.geomantics.com. Vertical views were generated using version 5 of the Manifold GIS, available from www.Manifold.net. We feel that 3D visualization has been under-utilized as a tool for testing terrain modelling algorithms, and the results are often more useful than a purely mathematical or statistical approach.

We have restricted ourselves to an evaluation of several weighted-average methods, 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, and not necessarily better. 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

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.)

One simple weighted-average model is triangle-based interpolation in which a linear interpolation is performed within every triangle T and areas of three triangles formed inside triangle T by interpolated point and two triangle T vertices are used as weights, as shown in Fig. 11. This method was used in the previous paragraph to illustrate a problem of flat triangles (Fig. 5) and to show a result of simple TIN model enrichment.

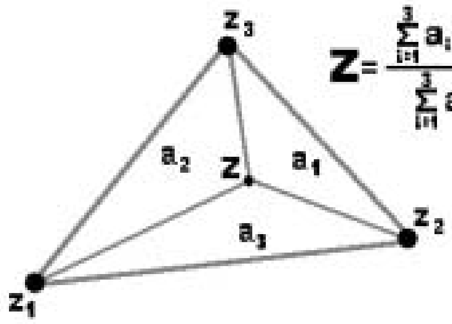

$$Z = \frac{\sum_{i=1}^{3} a_i z_i}{\sum_{i=1}^{3} a_i}$$

Fig. 11. Triangle-based interpolation

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 (([8], [12], [16]).

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. Fig. 12 gives an idea of this method for an interpolated point having four neighbours selected by a circle of specified radius.

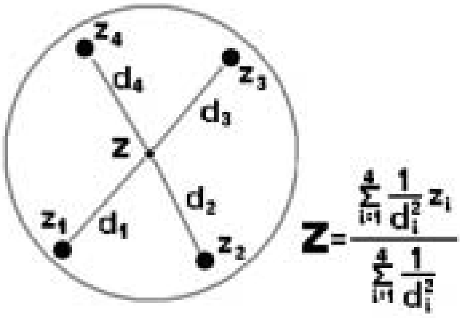

Fig. 12. Gravity interpolation

Fig. 13 shows the resulting surface for a radius of 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.

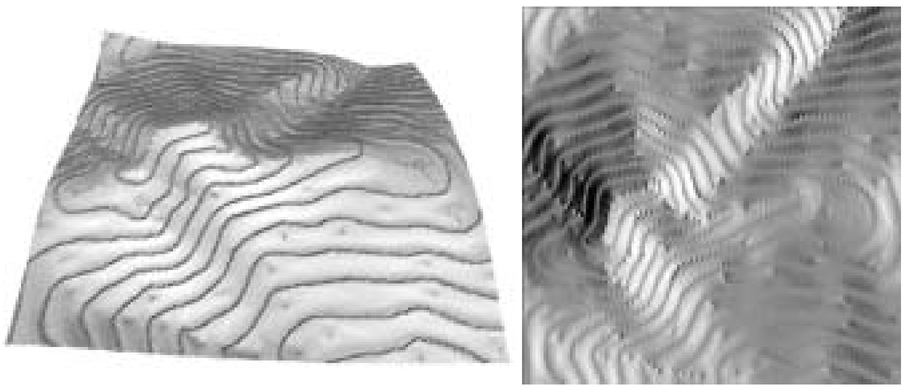

Fig. 13. Interpolation using the gravity model – a) perspective view; b) vertical view

The Sibson method is based on the idea of inserting the point being interpolated 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. In order to obtain the stolen areas an expensive insertion and deletion can be replaced by simulation of the insertion only without modifying the triangulation. Fig. 14 shows this process for a sample data set – the elevation value z is estimated using areas stolen from five neighbouring Voronoi cells.

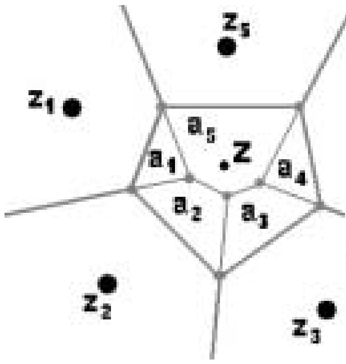

$$Z = \frac{\sum_{i=1}^{5} a_i z_i}{\sum_{i=1}^{5} a_i}$$

Fig. 14. Sibson interpolation

Fig. 15 shows a sample data set and the neigbour selection for the same point in both the gravity and Sibson method. In the Sibson method natural neigbour selection results in a reasonable set of neigbours, but the circle used in the gravity method may not select a sufficient number of neigbours to produce a valid elevation value for the interpolated point. The Sibson method is particularly appropriate for poor data distributions as the number of neighbours used is well defined. In the gravity model, when the data distribution is highly anisotropic, there is considerable difficulty in finding a valid counting circle radius.

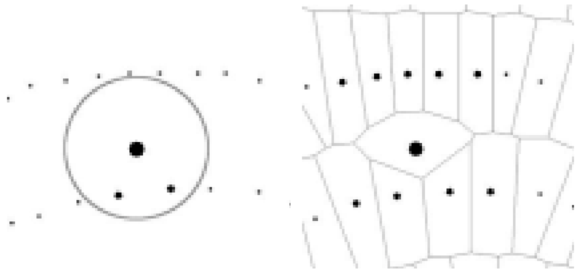

Fig. 15. Neighbour selection a) using a counting circle; b) using Voronoi neighbours

Fig. 16 shows the results of using Sibson interpolation. The surface behaves well - it fits the original data and is smooth in areas between data points, but is angular at ridges and valleys. Indeed, slopes are discontinuous at all data points (Sibson, 1980).

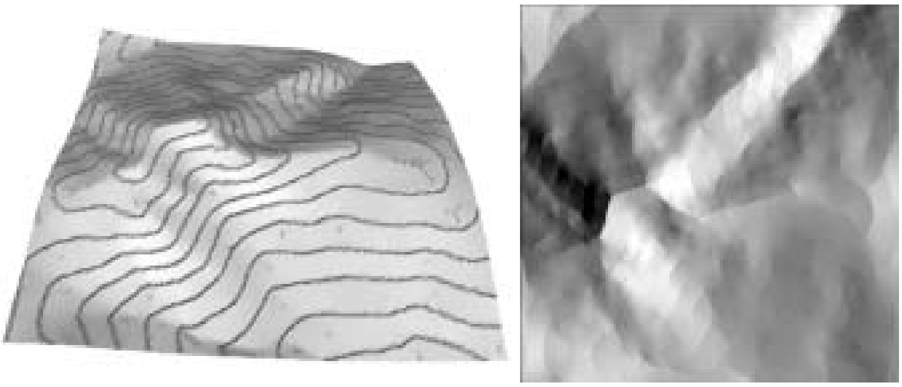

Fig. 16. Sibson interpolation - a) perspective view; b) vertical view

To remove discontinuities at data points, 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). In our case we replace every weight *w* by a new weight *w'* obtained by means of a simple function:

$$w' = 3 * w^2 - 2 * w^3$$

Fig. 17 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.

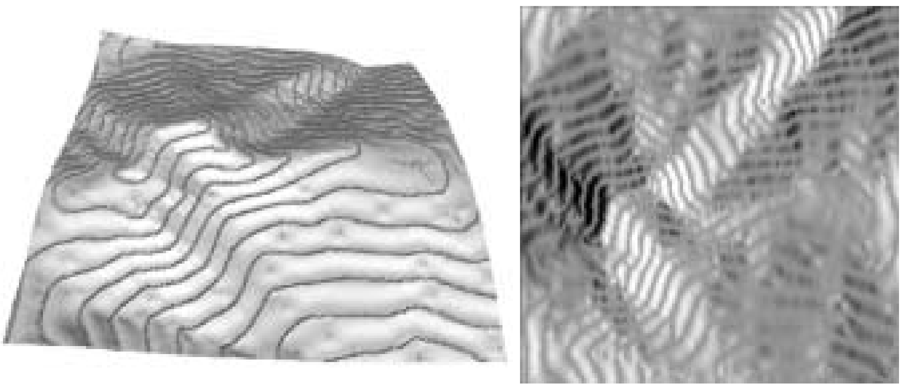

Fig. 17. Adding smoothing to Fig. 16 - a) perspective view; b) vertical view

5. Slopes – The Ignored Factor

This brings us to a subject often ignored in selecting a method for terrain modelling – 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 modelling and erosion. Clearly an assumption of zero slope at each data point, 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. The idea of this process is shown in Fig. 18 - point P at location (x,y) is being estimated and P_i is one of its neighbours. During interpolation the z value of each neigbouring data point P_i is replaced by value z' of the plane (tangent to the surface at P_i) at location (x,y). 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.

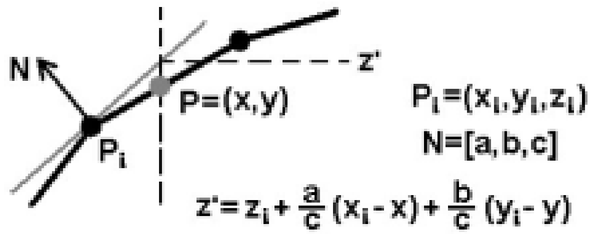

Fig. 18. Plane value at location of interpolated point P

Fig. 19 shows the result of using Sibson interpolation with data point slopes. The form is good, but slight breaks in slope can be seen at contour lines. When using smoothing and slope information together, the surface is smooth, but has unwanted oscillations, see Fig. 20. Clearly an improved smoothing function is desirable to eliminate these side-effects.

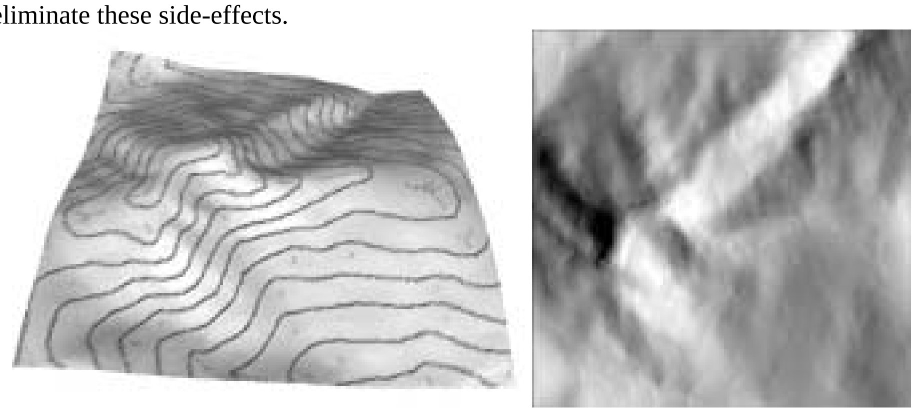

Fig. 19. Sibson interpolation using slopes at data points – a) perspective; b) vertical view

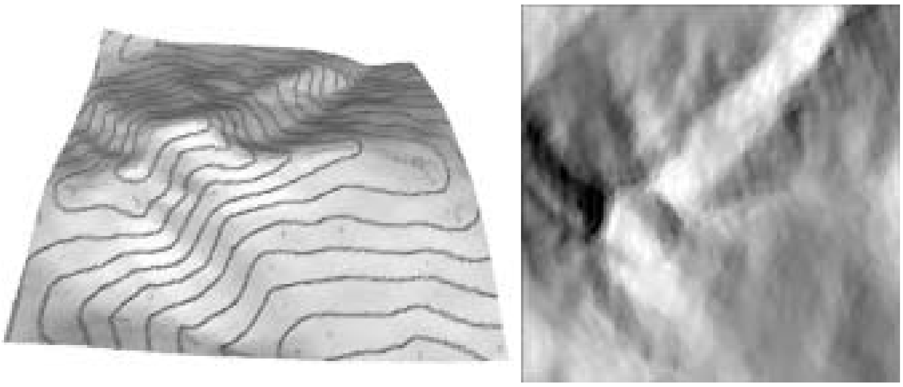

Fig. 20. Sibson interpolation using slopes and smoothing function at data points – a) perspective view; b) vertical view

Adding slopes to the simple TIN model (i.e. using the position in the triangle to provide the weights) produced results (Fig. 21a) that were almost as good as the Sibson method when the sample points were closely spaced along the contours. However, 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 (Fig. 21b), but even here including the data point slope function produces a significant improvement.

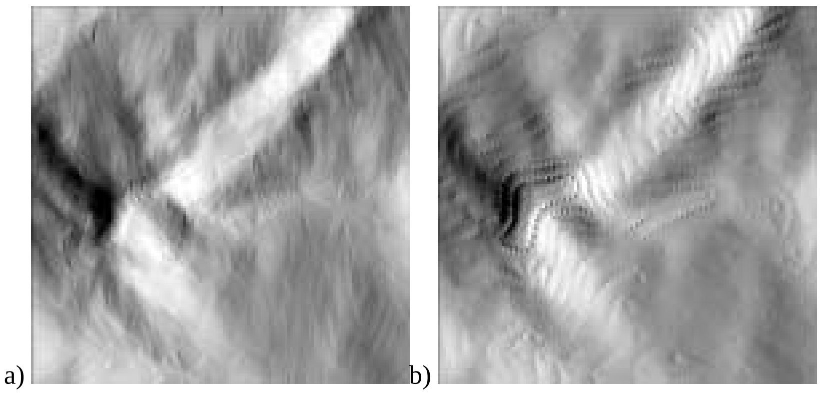

Fig. 21. Adding slopes at data points -a) triangle-based interpolation; b) gravity interpolation

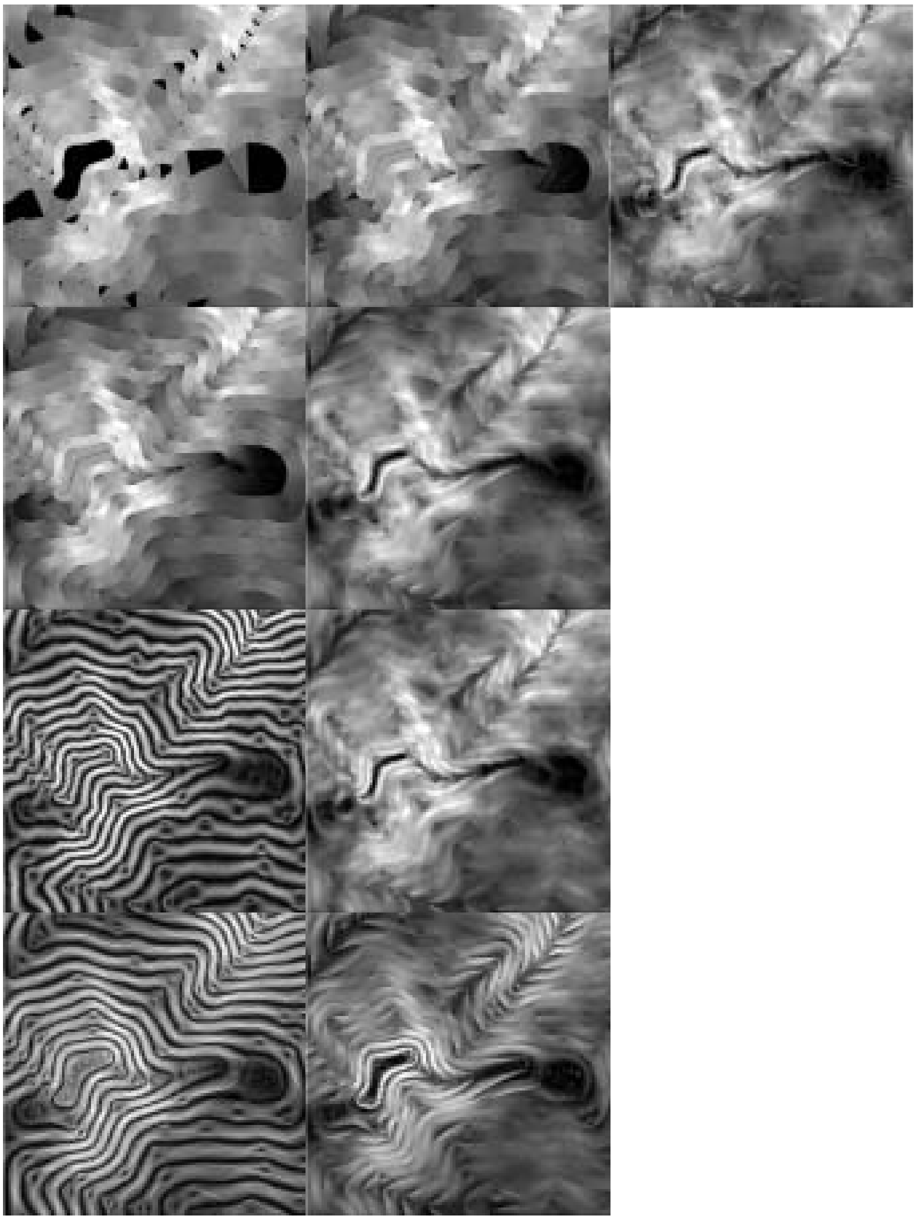

6. Evaluating the quality of terrain models

Fig. 22 shows relative heights distribution in grids created from our data set using various interpolation methods. Each histogram presents a number of grid cells having an elevation value within specified relative heights interval.

Fig. 23 shows a distribution of slope in grids created using various interpolation techniques. Each histogram shows number of cells having a slope value within specified range of degree.

Fig. 24 shows profile curvature distribution in grids created using various interpolation techniques. Each histogram shows number of grid cells having a profile curvature value within a specified range of values.

Fig. 25 shows plan curvature distribution in grids created using various interpolation techniques. Each histogram shows number of grid cells having a plan curvature value within a specified range of values.

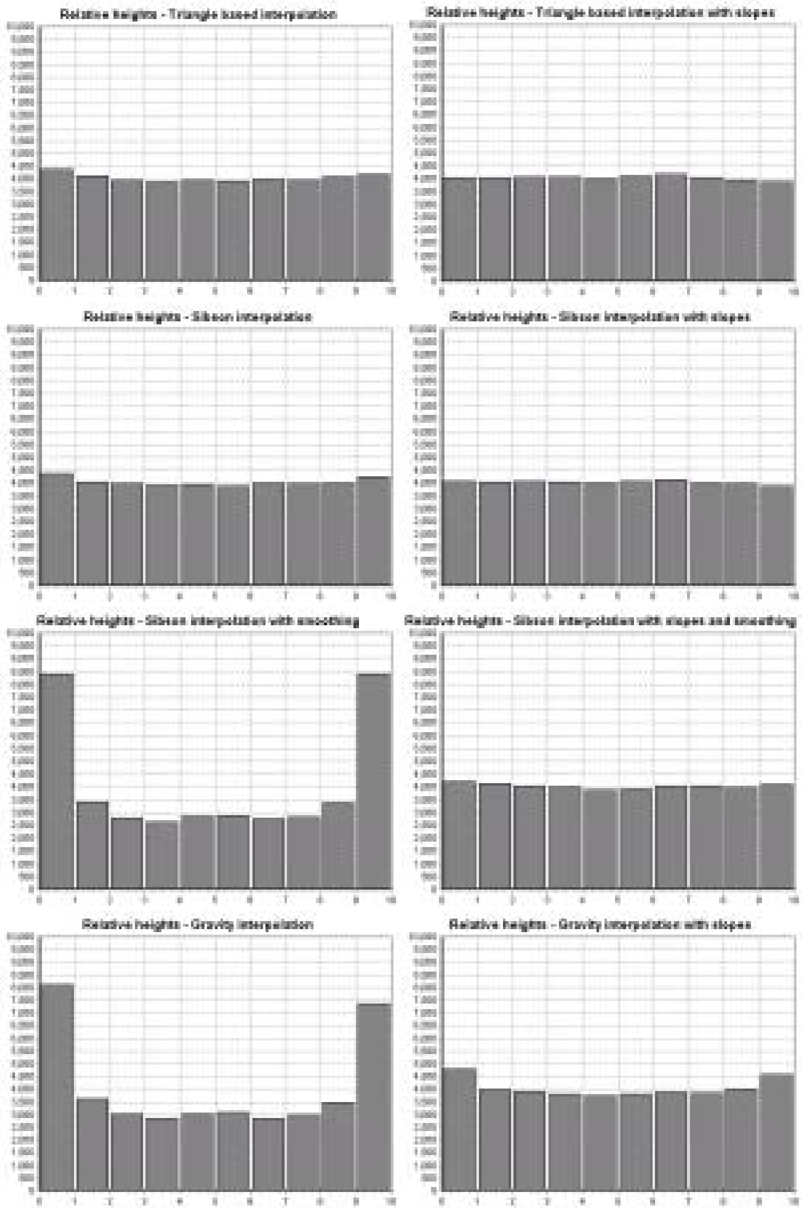

Fig. 22. Relative heights distribution in grids created using various interpolation methods

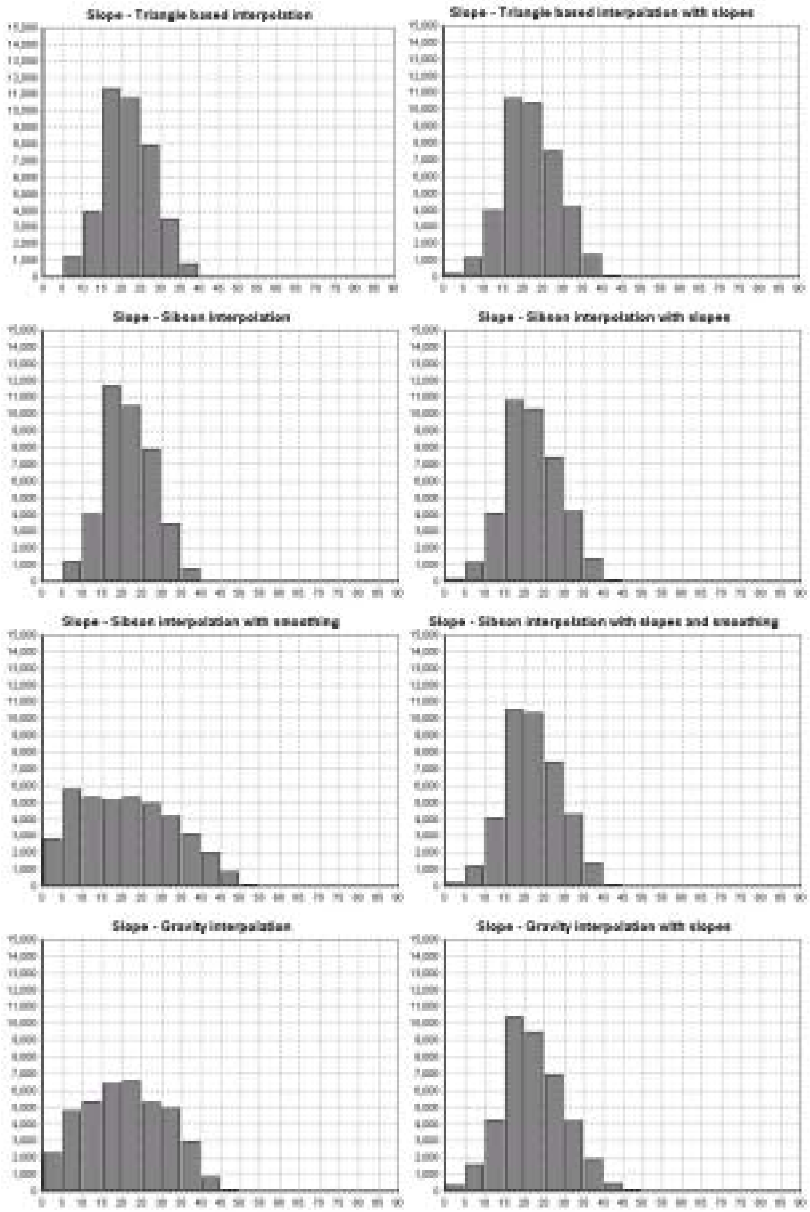

Fig. 23. Slope distribution in grids created using various interpolation methods

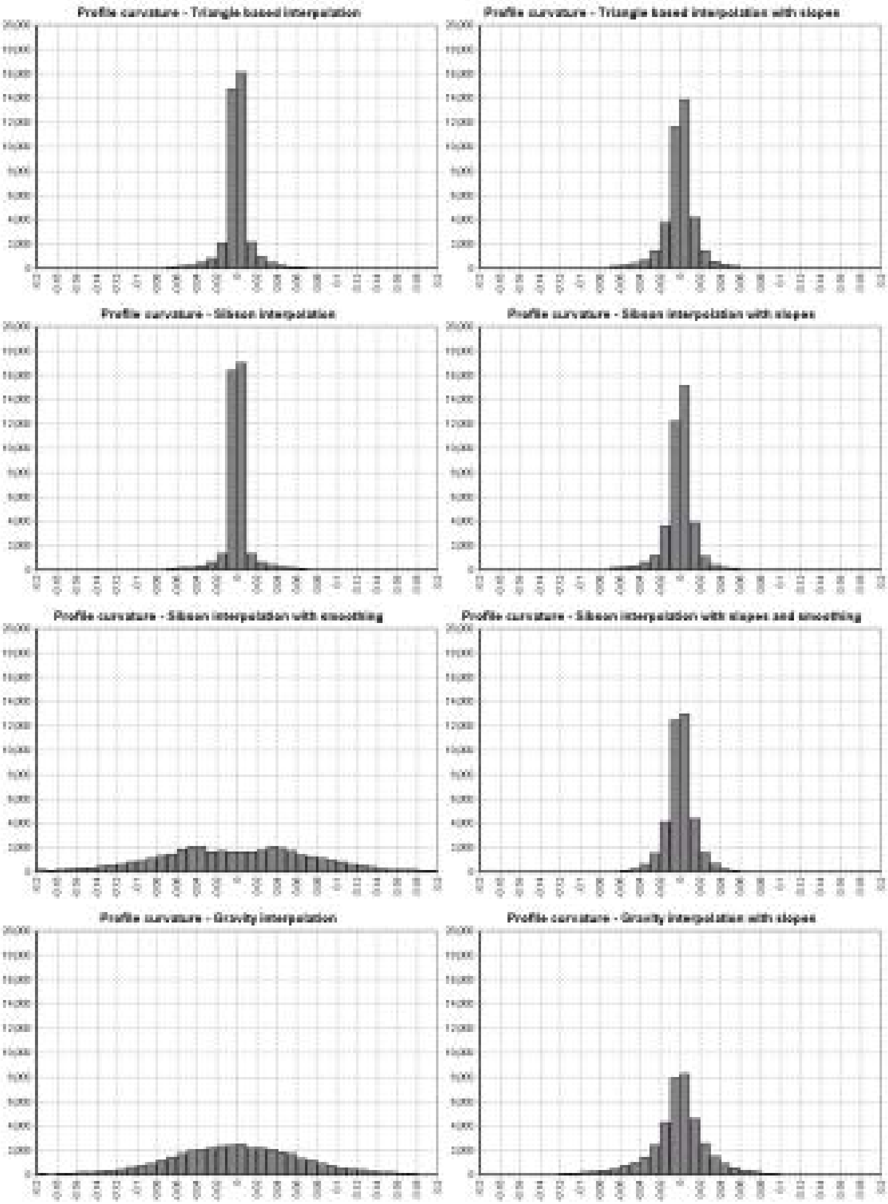

Fig. 24. Profile curvature distribution in grids created using various interpolation methods

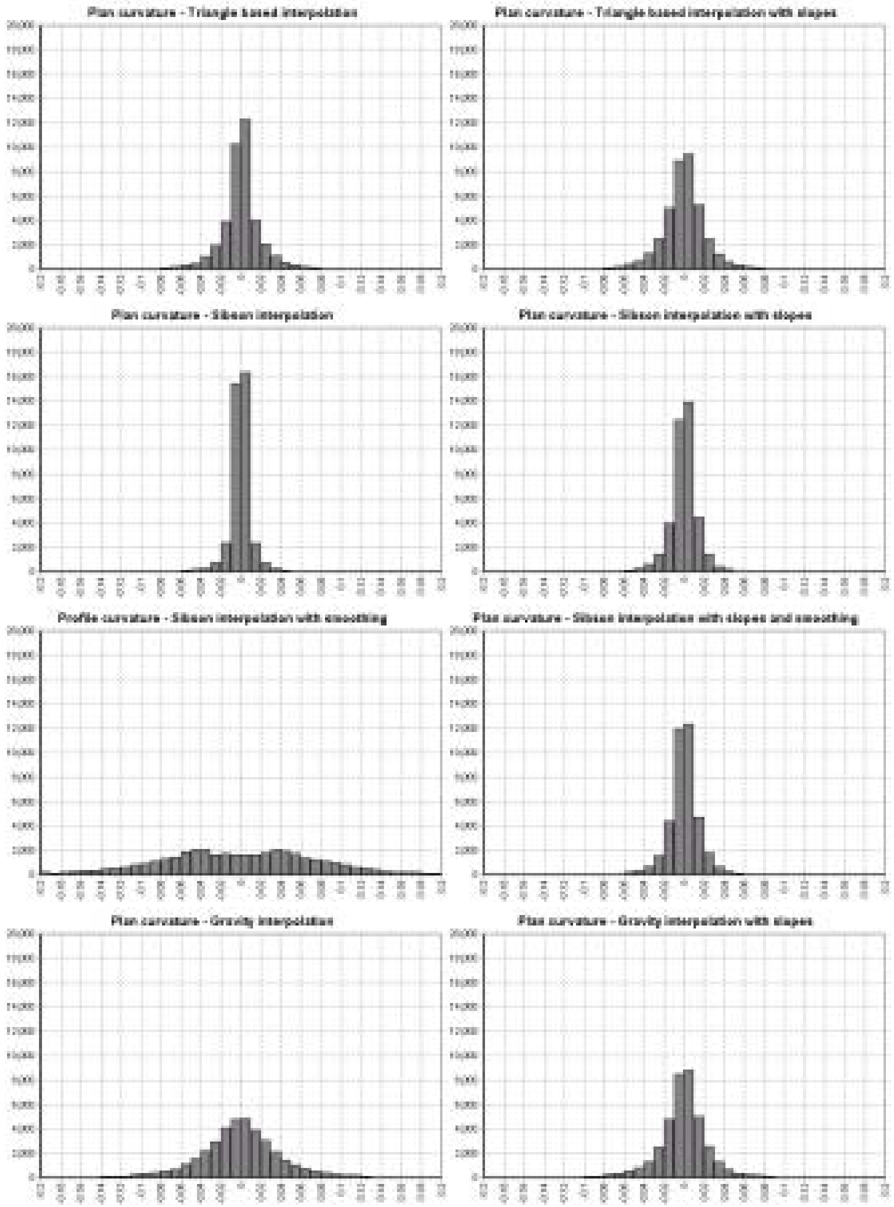

Fig. 25. Plan curvature distribution in grids created using various interpolation methods

7. Proposed methodology and conclusions

For the common problem of deriving surfaces from contours, we propose a general approach:

1. Generate skeleton points along the ridges, valleys, pits, summits and passes by the method of Aumann *et al.* [3] or of Thibault and Gold [14].
2. Assign elevations to these skeleton points by the methods described here, or other suitable techniques.
3. Eliminate flat triangles by the insertion of these skeleton points into the original TIN.
4. Estimate slope information at each data point by any appropriate technique.
5. Perform weighted-average interpolation using the previously estimated slope information. Avoid methods, such as the gravity model, with exponentially large close-range weightings, and avoid neighbour selection techniques which require user-specified parameters, such as counting-circle radius.

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.* [2]. Nevertheless, the moral is clear: both for finding adjacent points and for skeleton extraction, a consistent definition of neighbourhood is essential for effective algorithm development.

We conclude with another imaginary example. Fig. 13a shows four small hills defined by their contours, modelled by a simple triangulation. Fig. 13b 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 landform, 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.

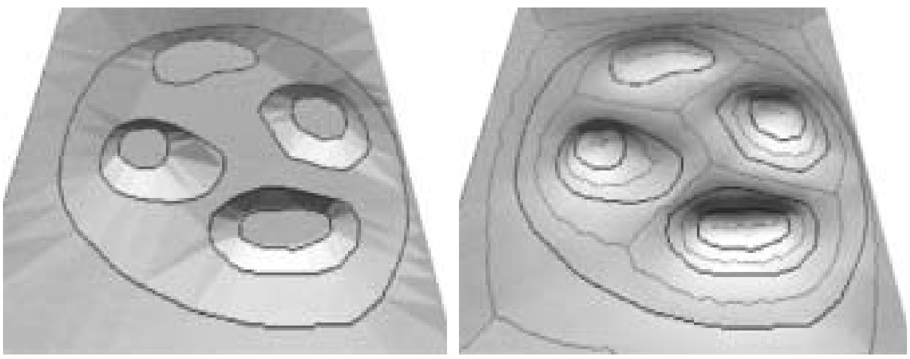

Fig. 13. Triangulation of several small hills – a) triangle based interpolation; b) Sibson interpolation with slopes

Acknowledgments

The work described in this paper was substantially supported by a grant from the Hong Kong Polytechnic University (Project No. PolyU A-PB79).

References

1. Amenta, N. Bern, M. and Eppstein, D. (1998) “The crust and the beta-skeleton: combinatorial curve reconstruction”, Graphical Models and Image Processing, 60, 125-135.
2. 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, 9-12.
3. Aumann, G. Ebner, H. and Tang, L. (1991) “Automatic derivation of skeleton lines from digitized contours”, ISPRS Journal of Photogrammetry and Remote Sensing, 46, 259-268.
4. Blum, H. (1967) “A transformation for extracting new descriptors of shape”, In: Whaten Dunn, W. (eds.), “Models for the Perception of Speech and Visual Form”, 153-171, MIT Press.
5. Crain, I.K. (1970) “Computer interpolation and contouring of two-dimensional data: a review”, Geoexploration, 8, 7x-86.
6. Davis, J. C. (1973) “Statistics and data analysis in geology”, 313, New York, John Wiley and Sons.
7. Dayhoff, M.O. (1963) “A contour map program for X-ray crystallography”, Communications of the Association for Computing Machinery, 6, 620-622.
8. Gold, C.M. (1989) “Chapter 3 - Surface interpolation, spatial adjacency and GIS”, In: Raper, J. (eds.), “Three Dimensional Applications in Geographic Information Systems”, 21-35, Taylor and Francis, Ltd., London.
9. Gold, C.M. (1999) "Crust and anti-crust: a one-step boundary and skeleton extraction algorithm", Proceedings of the ACM Conference on Computational Geometry, Miami, Florida, 189-196.
10. Gold, C. M. and Snoeyink, J. (2001) “A one-step crust and skeleton extraction algorithm”, Algorithmica, 30, 144-163.
11. Peucker, T.K. (1978) “The triangulated irregular network”, Proceedings, Digital Terrain Model Symposium, American Society of Photogrammetry, St. Louis.
12. Sibson, R. (1980) “A Vector Identity for the Dirichlet Tessellation”, Math. Proc. Cambridge Philos. Soc., 87, 151-155.
13. Sibson, R. (1982) “A brief description of natural neighbour interpolation”, In: Bamett, V. (eds), “Interpreting Multivariate Data”, 21-36, John Wiley and Sons, London.
14. Thibault, D. and Gold, C.M. (2000) “Terrain Reconstruction from Contours by Skeleton Construction”, GeoInformatica, 4, 349-373.
15. Walters, R.F. (1969) “Contouring by machine: a users’ guide”, American Association of Petroleum Geologists, Bulletin, 53, 2324-2340.
16. Watson, D.F. and Philip, G.M. (1987) “Neighborhood-based interpolation”, Geobyte, 2, 12-160.