

TERRAIN MODELLING BASED ON CONTOURS AND SLOPES

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@voro noi.com; maciekda@hotmail.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. This paper describes improved methods for extracting good quality terrain models from topographic contour maps, which despite modern techniques 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 the visualization of the relationships between sets of boundaries, and families of curves. The insertion of skeleton points guarantees the elimination of all flat triangles. Additional assumptions about the local uniformity of slopes give us enough information to assign elevation values to these skeleton points. Various interpolation techniques were compared using visualization of 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.

KEYWORDS: Terrain Modelling; Contours; Skeleton; Slopes; Voronoi Diagram.

1. 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. 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 cannot claim satisfactory results for these parameters.

2. GEOMETRIC PRELIMINARIES

The methods discussed here depend on a few fundamental geometrical constructs that are well known – the Voronoi diagram and its dual, the Delaunay triangulation (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, and goes through the three vertices.

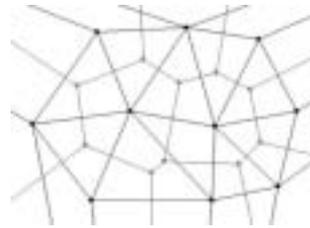


Figure 1. Delaunay triangulation and Voronoi diagram

The Voronoi diagram and Delaunay triangulation are associated with other geometric structures, shown in Fig 2a, 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. 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 the crust when there exists a circle through its two vertices that contain neither of its associated Voronoi vertices; otherwise the corresponding Voronoi edge belongs to the skeleton. A simple InCircle test applied to each Delaunay/Voronoi edge pair (Fig. 2b) distinguishes these cases.

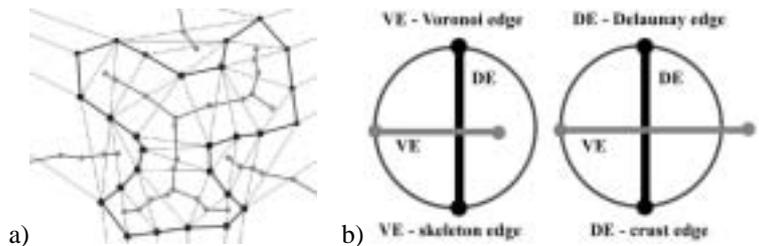


Figure 2. Crust and skeleton: a) crust (thick black lines) and skeleton (thick grey lines); b) crust/skeleton test

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 – and it would be desirable to preserve this, as derived ridges and valleys.



Figure 3. Contours: a) data points; b) crust and skeleton; c) crust and skeleton branches

Fig. 3a shows our raw data set (which is completely imaginary), and Fig. 3b shows the resulting crust, which reconstructs the contour lines and the skeleton. Fig. 3c shows the crust and only those skeleton points that provide unique information – ridge and valley lines that separate points on the same contour (“flat triangles”), rather than merely those points that separate adjacent contours. Aumann et al. (1991) produced somewhat similar results by raster processing.

Fig. 4 shows a close-up of the test data set with shaded flat triangles having all vertices at the same elevation. 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.

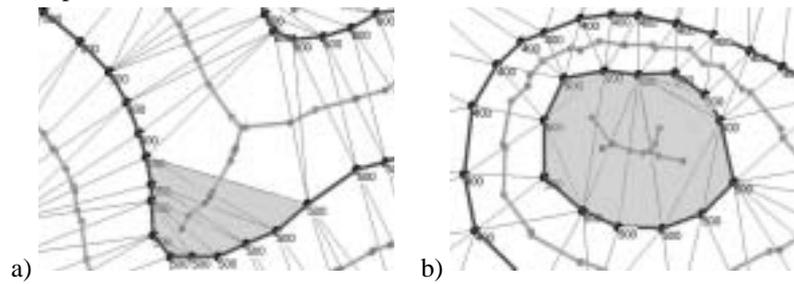


Figure 4. Skeleton and "flat triangles": a) ridge; b) summit

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. This is illustrated in Figs. 5a and 5b, where points on a simple closed curve are used to generate the crust and skeleton. In Fig. 5a the circumcentres of the skeleton points are given an elevation equal to the circumradius. The resulting interpolated model is shown in Fig. 5b. 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. 4b, the slope would initially be unknown, and would be estimated from the circumradii of the next contour level down.

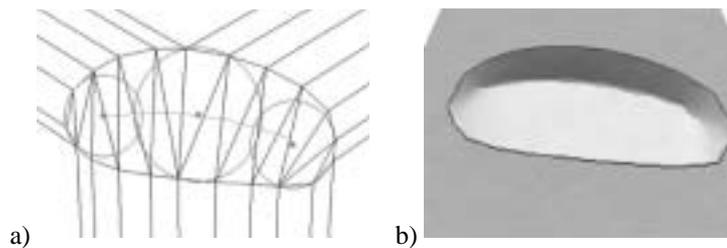


Figure 5. Triangulation of a summit: a) skeleton and circumcentres; b) elevation model after adding skeleton vertices with assigned height values

In the case of a ridge or valley, the circumradius may also be used, as in Fig. 6a, 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 centre of the smaller circle is thus based on the ratio of the two radii; see Thibault and Gold (2000) for details.

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. 6b, 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. This refines and amplifies the work of Thibault and Gold (2000). In particular, instead of enriching the model with all skeleton vertices, only the skeleton vertices of flat triangles are used, and new interpolation techniques are introduced to estimate skeleton point elevations.

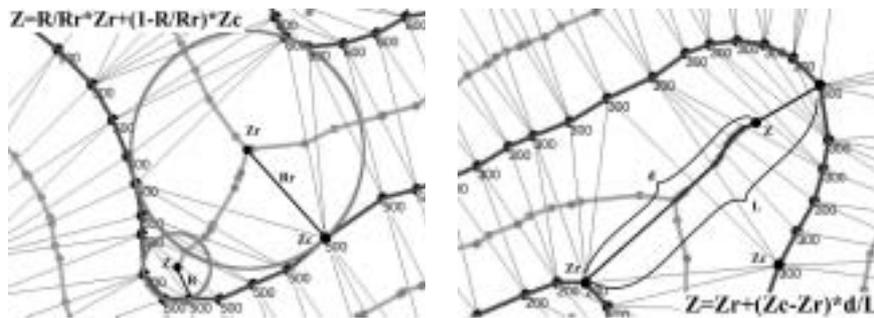


Figure 6. Estimating skeleton heights: a) from circumradii; b) from valley length

4. COMPONENTS OF AN INTERPOLATION MODEL

On the basis of a sufficient set of data points, we then 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 even 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. Nevertheless, many aspects of this study apply to Kriging as well, since it is a weighted-average method with the same problems of neighbour selection and the inclusion of slope values at data points 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: which set of neighbours is 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)

One of the simplest weighted-average models is triangle-based interpolation in which a linear interpolation is performed within each triangle. Fig. 7 shows the result, including the skeleton draped over the flat triangles. Fig. 8 shows the improvement when estimated skeleton points are added, and all flat triangles are automatically removed.

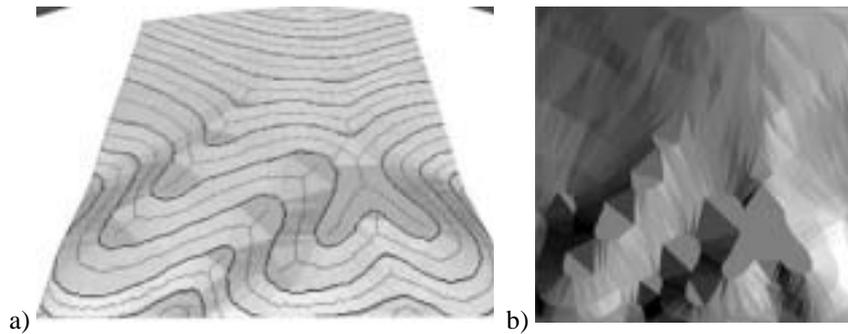


Figure 7. Triangle-based interpolation: a) perspective view; b) vertical view

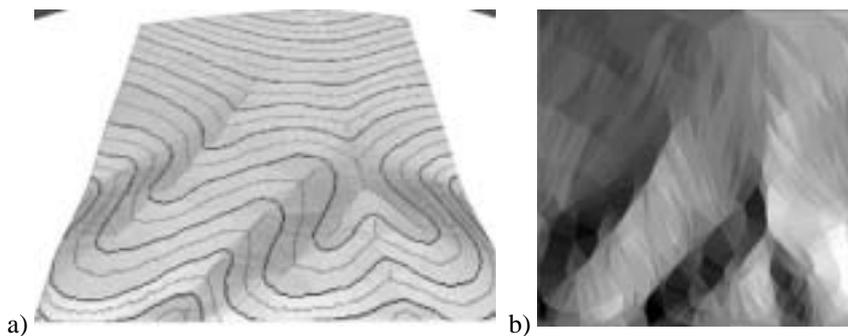


Figure 8. Adding skeleton points to Fig. 7: a) perspective view; b) vertical view

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. Fig. 9 shows the resulting surface for a radius of about a quarter of the map. Data points form bumps or hollows.

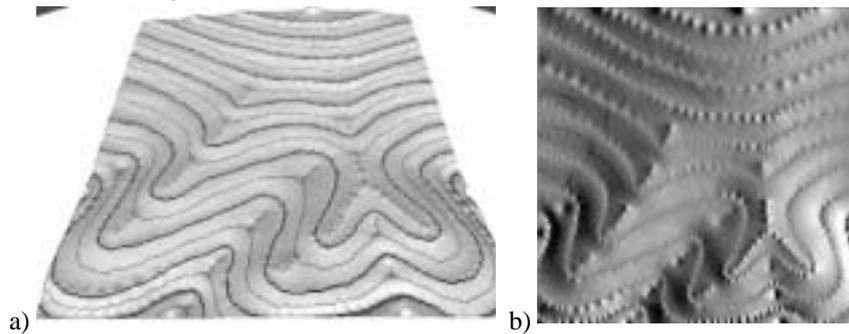


Figure 9. Interpolation using the gravity model with medium radius:
a) perspective view; b) vertical view

If the radius is reduced there may be holes in the surface where no data is found within the circle (Fig. 10a). If the radius is increased the surface becomes somewhat flattened, but the bumps remain (Fig. 10b). The result depends on the radius, and the other selection properties being used. Clearly estimates of slope would be very poor, and very variable.

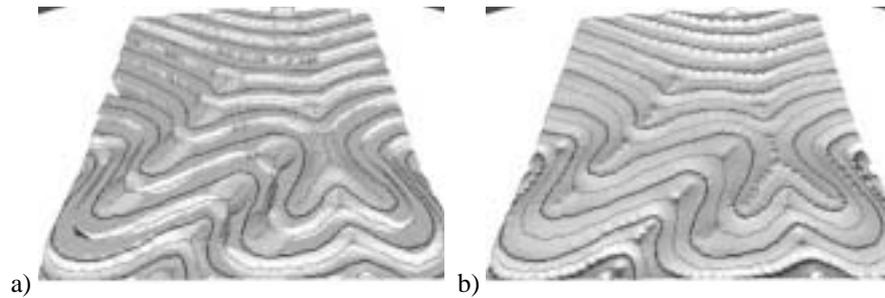


Figure 10. Gravity model interpolation: a) with small radius b) with large radius

The Sibson method, illustrated in Fig. 11a, 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.

Figs. 11b and 11c show a sample data set and the neighbour selection for the same point in both the gravity and Sibson method. In the Sibson method natural neighbour selection results in a reasonable set of neighbours, but the circle used in the gravity method may not select a sufficient number of neighbours 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.

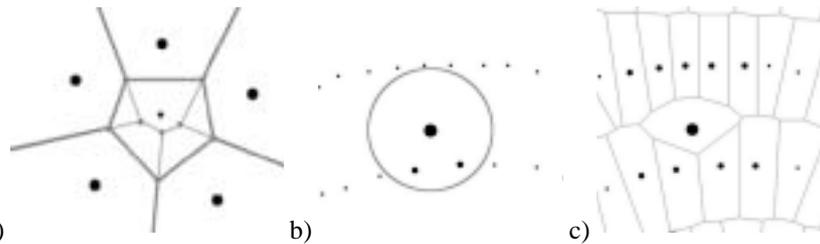


Figure 11. a) Sibson interpolation; b) neighbour selection using a counting circle; c) neighbour selection using Voronoi neighbours

Fig. 12 shows the results of Sibson interpolation. The surface 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). Fig. 13 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.

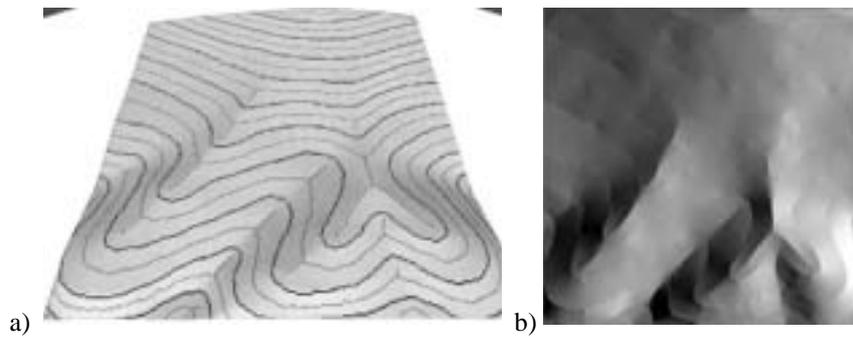


Figure 12. Sibson interpolation: a) perspective view; b) vertical view

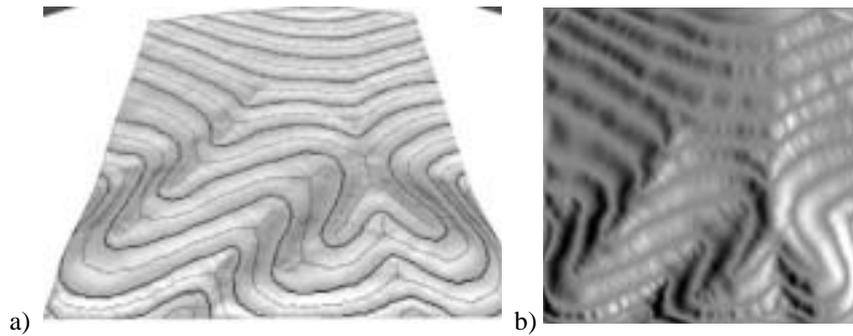


Figure 13. Adding smoothing to Fig. 12: a) perspective view; b) vertical view

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, erosion and insolation. 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. This function may be estimated by a variety of methods. 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 (Gold, 1989).

Fig. 14 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 (Fig. 15). Clearly an improved smoothing function is desirable to eliminate these side-effects.

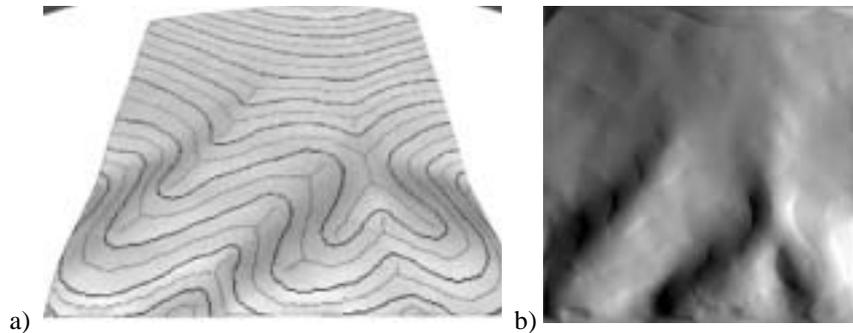


Figure 14. Sibson interpolation using slopes at data points:
a) perspective view; b) vertical view

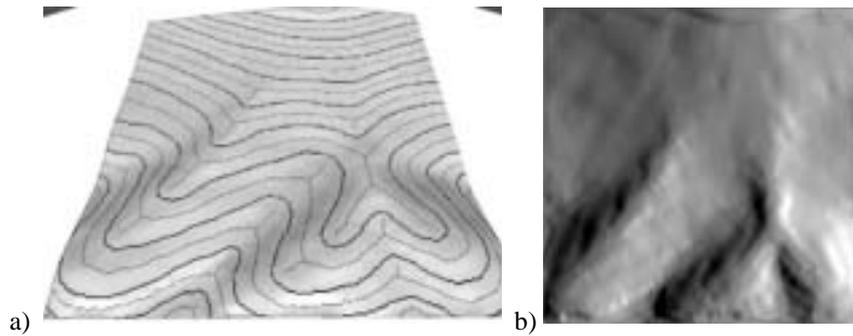


Figure 15. Sibson interpolation using slopes and smoothing function at data points: a) perspective view; b) vertical view

While it is impossible to show the results of all our experiments in this paper, we also used the method of Burrough and McDonnell, (1998) to calculate slopes and profile curvature of grids created from various combinations of our available weighted-average methods. Sibson interpolation with slopes and without smoothing gives more consistent regions of coherent slopes. Adding slopes to the simple TIN model (i.e. using the position in the triangle to provide the weights, as in Fig.16a) produced results 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, but even here including the data point slope function produces a significant improvement, as in Fig. 16b.

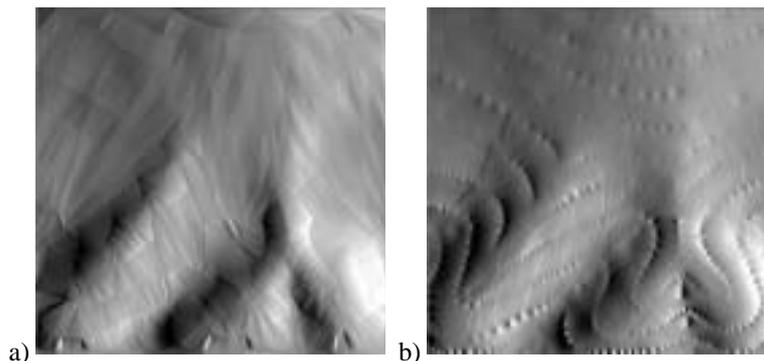


Figure 16. Adding slopes at data points: a) triangle-based interpolation; b) gravity interpolation

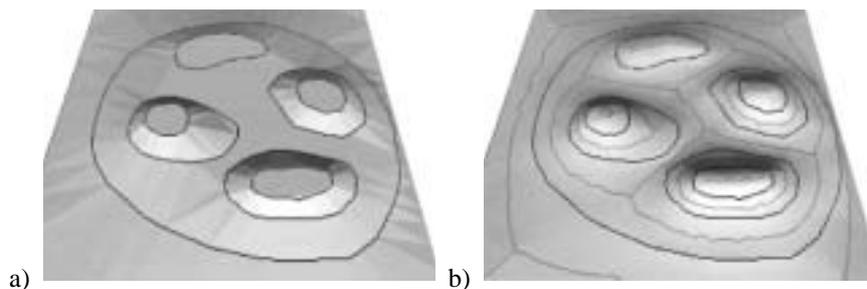
6. SUMMARY AND CONCLUSIONS

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

1. Generate skeleton points by the method of Aumann et al. (1991) or of Thibault and Gold (2000). Ignore skeletons between contours.
2. Assign elevations to these skeleton points by the methods described here based either on longitudinal or lateral slope consistency, 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, which require user specified parameters. Sibson interpolation appears to be the best choice.

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 with estimated elevations, in order to eliminate flat triangles. The second most important contribution is the addition of slope information at the data points, and its use in the weighted average interpolation process - even poor interpolation methods are significantly improved. Also important is the selection of a meaningful set of neighbours around the estimated point. 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.

We conclude with another imaginary example. Fig. 17a shows four small hills defined by their contours, modelled by a simple triangulation. Fig. 17b 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, or generation of DEMs, we believe that our methods are a significant improvement on previous work.



a) b)
Figure 17. Triangulation of several small hills: a) triangle-based interpolation, b) Sibson interpolation with slopes

ACKNOWLEDGEMENTS

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

REFERENCES

- Amenta N, Bern M, Eppstein D (1998) The crust and the beta-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60, 125-135.
- Anton F, Gold CM, 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*, 1998, Kiev, Ukraine, 9-12.
- Aumann G, Ebner H, Tang L (1991) Automatic derivation of skeleton lines from digitized contours. *ISPRS Journal of Photogrammetry and Remote Sensing*, 46, 259-268.
- 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*. MIT Press, 153-171.
- Burrough P, McDonnell RA (1998) *Principles of Geographical Information Systems*. 2nd ed, Oxford University Press, 336.
- Gold CM (1989) Chapter 3 - Surface interpolation, spatial adjacency and GIS. In: Raper J (eds) *Three Dimensional Applications in Geographic Information Systems*. Taylor and Francis, Ltd., London, 21-35.

Gold CM (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.

Gold CM, Snoeyink J (2001) A one-step crust and skeleton extraction algorithm. Algorithmica, 30, 144-163.

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

Thibault D, Gold CM (2000) Terrain Reconstruction from Contours by Skeleton Construction. GeoInformatica, 4, 349-373.

Watson DF, Philip GM (1987) Neighborhood-based interpolation. Geobyte, 2, 12-16.