Spatial Interpolation: From Two to Three Dimensions

Hugo Ledoux and Christopher Gold

GIS Research Centre, School of Computing
University of Glamorgan, Pontypridd, CF37 1DL, Wales, UK
hledoux@glam.ac.uk christophergold@voronoi.com

1. Introduction

Interpolation methods are an important part in a geographical information system (GIS) and have been used for years to model elevation data. They are crucial in the visualisation process (generation of contours), for the conversion of data from one format to another, to have a better understanding of a dataset or simply to identify bad samples. The result of interpolation—usually a surface that represents the real terrain—must be as accurate as possible because it often forms the basis for spatial analysis, e.g. runoff modelling or visibility analysis. Although interpolation helps in creating three-dimensional surfaces, it is intrinsically a two-dimensional operation for only the *x*-*y* coordinates of each sample are used and the elevation is considered as an attribute. With the new technologies available to collect information about the Earth, more and more three-dimensional data are collected. A typical dataset in geosciences has samples in 3D space (*x*-*y*-*z* coordinates) to which attributes are attached (e.g. the salinity of the water, or the percentage of gold in the rock). Because of the way they are collected, 3D geoscientific datasets have a highly anisotropic distribution. Geologic and oceanographic data are for example respectively gathered from boreholes and water columns; data are therefore usually abundant vertically but sparse horizontally. While most of the interpolation methods used in GIS intuitively extend to 3D, it is not obvious that they preserve their properties or are appropriate for such datasets.

In this paper, we discuss the extension to 3D of some of the interpolation methods found in GIS or geoscientific modelling packages. We first present some details concerning their generalisation, and then evaluate briefly their properties.

2. What is a Good Interpolation Method?

Given a set of samples to which an attribute *a* is attached, spatial interpolation is the procedure used to estimate the value of the attribute at an unsampled location *x*. To achieve that, it creates a function *f*, called the interpolant, that tries to fit the samples as well as possible, based on some criteria. Interpolation is based on *spatial autocorrelation*, that is the attribute of two points close together in space is more likely to be similar than that of two points far from each other. Watson (1992) lists the essential aspects of an 'ideal' interpolation method: it must be *exact* (the surface must 'pass through' the samples), *continuous* and *smooth* (the surface must have a continuous and finite slope everywhere), *local* (the interpolation function uses only some neighbouring samples), and must also be able to adapt to different densities and distributions of data.

We describe in this paper only weighted-average interpolation methods, i.e. a local method that uses only some sample points, to which weights (an importance) are assigned. The interpolation function *f* of such methods has the following form:

$$f(x) = \frac{\sum_{i=1}^{k} w_i(x) \cdot a_i}{\sum_{i=1}^{k} w_i(x)} \qquad (1)$$

where $w_i(x)$ is the weight of each neighbour p_i (with respect to the interpolation location x) used in the interpolation process, and a_i the attribute of p_i.

3. From 2D to 3D

3.1. Nearest Neighbour

This is not really an interpolation method, but it is nevertheless used. The value of an attribute at location x is simply assumed to be equal to the attribute of the nearest data point. Given a set S of data points, if interpolation is performed at many locations close to each other, the result is the Voronoi diagram (VD) of S (see Figure 1 for 2D and 3D examples), where all the points inside a Voronoi cell have the same value.

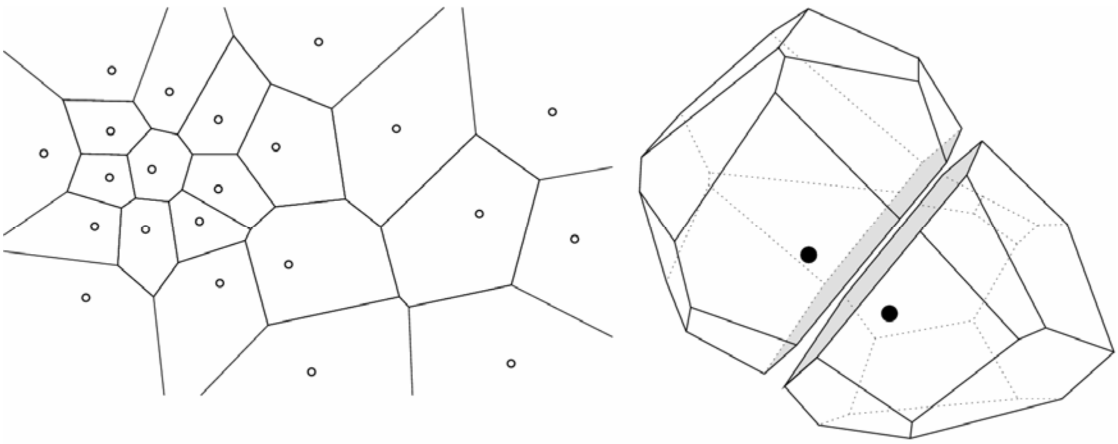

Figure 1. Left: the 2D VD. Right: two 3D Voronoi cells adjacent to each other.

In 2D, each Voronoi cell is a convex polygon, and in 3D it is a convex polyhedron. The VD actually creates a piecewise model, where the interpolation function inside each Voronoi cell is a constant function. Although the method is local, exact and relatively easy to implement (simply search for the closest data point), the interpolation function is discontinuous at the border of cells.

3.2. Distance-based Methods

Distance-based methods are probably the most known methods and they are widely used in many fields. As shown in Figure 2, in 2D they often use a 'searching circle', whose radius is user-defined, to select the data points p_i involved in the interpolation at location x. The weight assigned to each is typically based on the square of the distance from x to p_i. Other weights can also be used, see Watson (1992) for a discussion. The size of the radius of the searching circle influences greatly the result of the interpolation. A very big radius means that the resulting surface will be smooth or 'flattened'. On the other hand, a radius that is too small might have dramatic consequences if for example no data points are inside the circle. This method has many problems when the data distribution is highly anisotropic or varies greatly in one dataset. Figure 2 shows one example with contour lines.

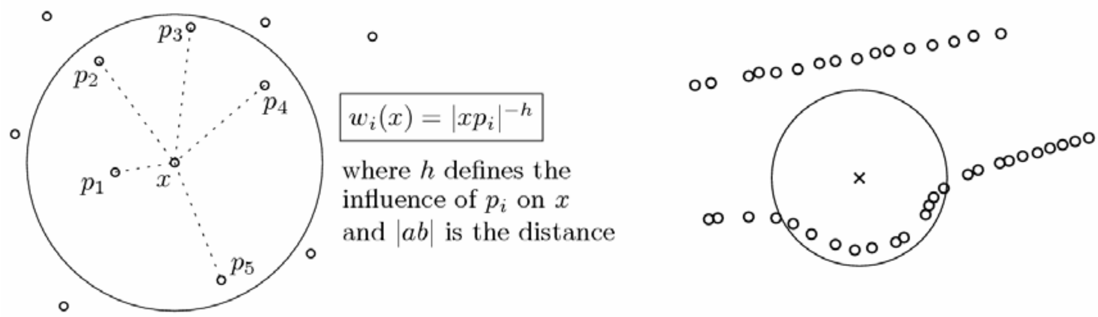

Figure 2. Left: inverse distance to a power interpolation. Right: problems with anisotropic datasets.

The generalisation of this method to 3D is straightforward: a searching sphere with a given radius is used. The one-dimension criterion of the method (the distance) will affect even more the 3D method because the search must be performed in one more dimension.

3.3. Linear Interpolation in Triangles

This method is popular for terrain modelling applications and is based on a triangulation of the data points. As is the case for the VD, a triangulation is a piecewise subdivision, and here a linear function is assigned to each piece (a triangle). Interpolating at location *x* means first finding inside which triangle *x* lies, and then the height is estimated by linear interpolation on the 3D plane defined by the three vertices forming the triangle. The last step can be efficiently performed by using barycentric coordinates, which are local coordinates defined within a triangle (see Figure 3 for details). To obtain satisfactory results, this method is usually used in 2D with a Delaunay triangulation (DT) because it maximizes the minimum angle of each triangle, in other words it creates triangles that are as equilateral as possible (see Figure 3). This ensures that the three vertices used in the interpolation process will most likely be close to the interpolation location.

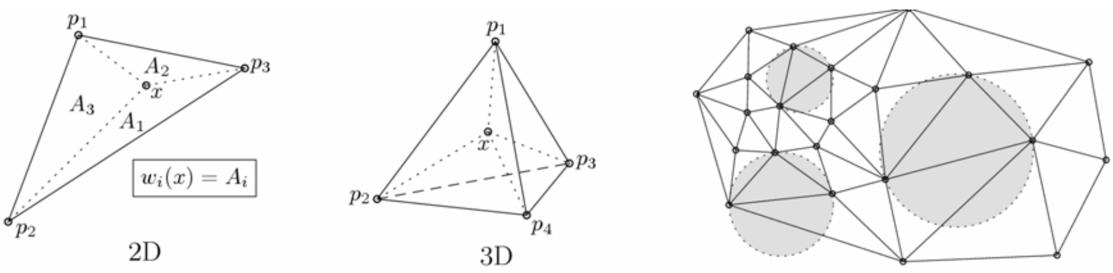

Figure 3. Left: barycentric coordinates. Right: Delaunay triangulation in 2D. The circle circumscribed to a Delaunay triangle is empty of any other data points.

The generalisation of this method to 3D is straightforward: linear interpolation is performed within each tetrahedron of a 3D triangulation, and the barycentric coordinates can also be generalised to use volumes instead of areas. Finding 'good' tetrahedra is however more difficult than finding good triangles because the *max-min* property of Delaunay triangles does not generalise to 3D. A 3D DT can indeed contain some tetrahedra, called *slivers*, whose four vertices are almost coplanar; interpolation within such tetrahedra does obviously not yield good results. Different empirical methods have been developed to improve the shape of certain tetrahedra in a 3D DT (e.g. Cheng et al. (2000)), but they are quite complex to implement.

The method is local and exact, but its major problem is that it is not continuous at the edges or faces of the 3D triangulation.

3.4. Natural Neighbour Interpolation

It has been shown by different researchers (Gold, 1989; Sambridge et al., 1995; Watson, 1992) that the natural neighbour interpolation method (Sibson, 1981) avoids most of the problems the other methods have with anisotropic datasets. This is a method based on the Voronoi diagram for both selecting the data points involved in the process, and to assign them a weight. It uses two VDs: one for the set of data points, and another one where a point x is inserted at the interpolation location. The insertion of x modifies locally a VD: x 'steals' some parts of some Voronoi cells, as shown in Figure 4. The data points p_i used for the interpolation are the ones whose Voronoi cell has been modified by the insertion of x, and the weight of each p_i is proportional to the amount that was stolen (areas in 2D and volumes in 3D). The properties of the method are the same in any dimensions: it is local and exact, and the result is continuous everywhere. The first derivative is also continuous everywhere, except at the data points. It moreover adapts very well to geoscientific data because the neighbours used are based on the VD, which results in data points that both surround and are close to the interpolation location.

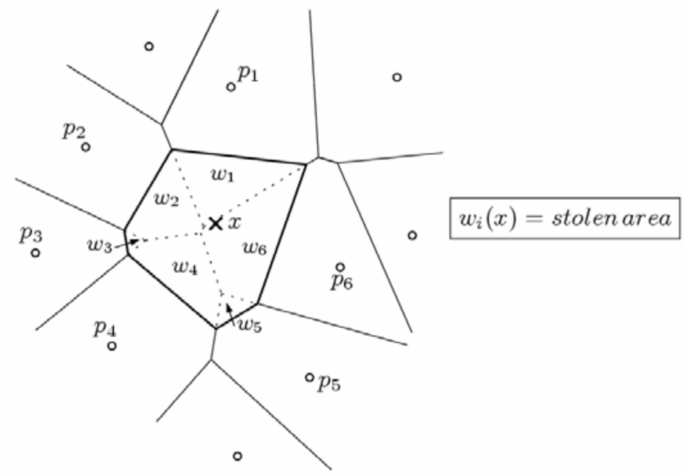

Figure 4. Natural neighbour interpolation in 2D.

Although natural neighbour interpolation yields good results, it has not been widely used, at least among the GIS community. This is probably due to the fact that it is a computationally expensive process, and its implementation is intricate, especially in 3D. It requires the computation of two VDs—just constructing one in 3D is not that simple, see Sugihara and Hiroshi (1995)—and of volumes of cells. We have therefore proposed an algorithm (Ledoux and Gold, 2004), valid in any dimensions, that we believe is easier to implement as it is based on a known and relatively simple to implement algorithm to construct a VD.

4. Discussion

We have shown that most of the weighted-average methods do generalise to 3D, but sometimes new problems appear, for example with triangulations in 3D. The flaws present in some 2D methods will often be amplified in three and higher dimensions. The main aspect not covered in the paper is the continuity of the surface, or interpolant. It is indeed possible to modify the methods described to ensure that the first derivative is continuous everywhere, for example Akima (1978) used higher order functions in each piece of a triangulation. Nevertheless, natural neighbour interpolation, as described, conforms to Watson’s criteria, except for the first derivative continuity at data points, which may be achieved by using additional functions (Gold, 1989).

Acknowledgements

We would like to thank the support from the Hong Kong's Research Grants Council, Hong Kong (Project PolyU 5068/00E and through a research studentship to the first author). The second author also acknowledges the funding of the EU Marie Curie Chair.

References

AKIMA, H., 1978. A method of bivariate interpolation and smooth surface fitting for irregularly distributed data points. *ACM Transactions on Mathematical Software*, 4(2):148–159.

CHENG, S.-W., DEY, T. K., EDELSBRUNNER, H., FACELLO, M. A., and TENG, S.-H., 2000. Sliver exudation. *Journal of the ACM*, 47(5):883–904.

GOLD, C. M., 1989. Surface interpolation, spatial adjacency and GIS. In J. Raper, editor, *Three Dimensional Applications in Geographic Information Systems*, pages 21–35. Taylor & Francis.

LEDOUX, H. and GOLD, C. M., 2004. An efficient natural neighbour interpolation algorithm for geoscientific modelling. In *Proceedings 11th International Symposium on Spatial Data Handling*, pages 97–108. Leicester, UK.

SAMBRIDGE, M., BRAUN, J., and MCQUEEN, H., 1995. Geophysical parameterization and interpolation of irregular data using natural neighbours. *Geophysical Journal International*, 122:837–857.

SIBSON, R., 1981. A brief description of natural neighbour interpolation. In V. Barnett, editor, *Interpreting Multivariate Data*, pages 21–36. Wiley, New York, USA.

SUGIHARA, K. and HIROSHI, I., 1995. Why is the 3D Delaunay triangulation difficult to construct? *Information Processing Letters*, 54:275–280.

WATSON, D. F., 1992. *Contouring: A Guide to the Analysis and Display of Spatial Data*. Pergamon Press, Oxford, UK.

Biography

Hugo Ledoux is a PhD student at the University of Glamorgan, Wales, UK. He holds a BSc in Geomatics from Laval University, Quebec City, Canada. His research focuses on topological data structures, the development of algorithms for three-dimensional modelling, and the use of Voronoi diagram for environmental modelling. He is particularly interested in combining the fields of computational geometry and GIS.