J Geograph Syst (2001) 3: 199–213

Journal of Geographical Systems

Volume-preserving interpolation of a smooth surface from polygon-related data

Wolf-Dieter Rase

Federal Office for Building and Regional Planning (BBR) Am Michaelshof 8, 53177 Bonn, Germany (e-mail: rase@acm.org, rase@bbr.bund.de)

Received: 5 October 2000 / Accepted: 20 December 2000

Abstract. The interpolation of continuous surfaces from discrete points is supported by most GIS software packages. Some packages provide additional options for the interpolation from 3D line objects, for example surface-specific lines, or contour lines digitized from topographic maps. Demographic, social and economic data can also be used to construct and display smooth surfaces. The variables are usually published as sums for polygonal units, such as the number of inhabitants in communities or counties. In the case of point and line objects the geometric properties have to be maintained in the interpolated surface. For polygon-based data the geometric properties of the polygon boundary and the volume should be preserved, avoiding redistribution of parts of the volume to neighboring units during interpolation. The pycnophylactic interpolation method computes a continuous surface from polygon-based data and simultaneously enforces volume preservation in the polygons. The original procedure using a regular grid is extended to surface representations based on an irregular triangular network (TIN).

Key words: Interpolation, surface, TIN, cartography, GIS, regional planning

JEL classification: C63, C88

1 Conceptual surfaces

Surfaces with only one possible height value at a location in the plane (2½D surfaces) are usually generated from discrete geometric elements, such as points or lines with a height information. The height values between the data points or lines are interpolated following rules based on global or local properties of the data, mathematical models, experience, or expert knowledge about the formation of the real surface. The most common example for the model of a surface and its visualization is the surface of the earth as displayed on topographic maps.

I thank Waldo Tobler for his contributions and the constructive discussion during the implementation and test of the TIN-based version of the pycnophylactic interpolation.

The surface of the earth is a physical surface which can be perceived. Surfaces calculated from physical variables, such as temperature or air pressure, are not directly visible. They are nevertheless part of our every-day experience, for example from weather forecasts on TV. Time distances in transportation networks can be visualized as continuous surfaces for analyzing and demonstrating the general accessibility of public infrastructure by computing time distances to, for example, central places or nodes of the transportation network.

Surfaces derived from environmental, demographic, social or economic data are used more and more in spatial analysis. The invisible surfaces are conceptual models which can only be detected by their influence, not by their actual physical presence. They exist only as numbers and files in a computer system. These models are called *conceptual surfaces*. The resolution of the original data and the resolution of the generated surface necessary to obtain a high-quality map differ considerably in most cases. Conceptual surfaces are continuous, with some rare exceptions.

The visualization of conceptual surfaces can be a very useful complement to the usual choropleth maps in spatial analysis. In the case of most environmental variables, for example, the abrupt change in value at the boundary of a polygon on a choropleth map is not appropriate, because the boundaries of natural phenomena coincide rather seldom with administrative boundaries. A continuous surface or its visualization also has a certain amount of uncertainty or fuzziness compared to the strictness of boundary lines. This feature of interpolated surfaces is highly appreciated, for example, to visualize concepts in spatial development or planning on the supranational (European) level to avoid direct references to the existing political and administrative boundaries (BfLR 1995). A conceptual surface can also be considered under certain conditions as a probability surface which can be used for fuzzy reasoning or map algebra (Tomlin 1990).

In contrast to a model of the earth's surface which is calculated from point and line data, conceptual surfaces have to be derived from polygon-based data as well. The data from a poll, a census or from administrative procedures are often published only as totals for spatial units of a certain size to ensure the privacy of individuals or trade secrets of companies. The use of a geometric proxy for the polygon, e.g. the centroid, is common practice, but has some shortcomings. For example the preservation of the volume associated with the polygon cannot be assured in a point-based interpolation method, and the polygon boundary is not present in the surface model.

2 Data structures for continuous surfaces

A data structure is the specification how a model of the real world or parts thereof are represented in a computer system (computer scientists will forgive me the rather coarse definition). For the kind of surfaces discussed here two data structures are used in most applications, the *regular grid* and the *triangular irregular network*, or TIN. Other divisions of the plane, such as hierarchical tesselations (e.g. quadtrees, Samet 1990) seem to be less important in the context of the volume-preserving interpolation. The data structure for representing the surface should not be confused with similar data structures some interpolation methods use to calculate the continuous surface, for example the TIN (Renka 1996b), or a hierarchical tesselation (Mitasova and Mitas 1993).

2.1 Regular grid of rectangles or triangles

The most common representation for a surface is a grid of orthogonal lines with numerical values at the grid intersections. For practical reasons the lines are usually equidistant. Regular rectangular grids can be handled rather comfortably in a computer program because most programming languages provide support for two-dimensional arrays. A grid of equilateral triangles has some advantages over the rectangular grid (Watson 1992), but requires extra provisions in the program code. The immediate neighbors of a point, for example, cannot be accessed by simply incrementing or decrementing the array indices. The regular rectangular grid with numerical height values at the grid intersections is a common model for objects used in display and rendering software, also known as a *heightfield.*

The real world is seldom regular or even rectangular. The data points coincide only by chance with the intersection points of the grid, lines nearly never follow the grid lines. Algorithms and programs have been developed to interpolate irregular points, lines and polygons to a regular grid. To minimize the error induced by the transformation from the irregular data to the regular grid, the size of the grid cells can be made very small. A finer resolution of the grid, however, requires more processing time, memory and I/O bandwidth.

2.2 The triangular irregular network (TIN)

The triangular irregular network (TIN) allows the preservation of the original data points in the model of a surface (Peucker et al. 1978). The data points are the *nodes* of the irregular grid. The nodes are connected by *arcs* or *edges*, forming a mesh of triangles of different size and orientation. In contrast to a regular grid the TIN can be adapted to varying local resolution requirements. In regions of high relief energy (rapid change in numerical value) the mesh can be made denser. In regions with extended areas of small variation the size of the triangles might be larger. The TIN can be *decimated* by eliminating redundant triangles and nodes, for example in flat areas where the nodes have only slight differences in height (Schroeder et al. 1992; Junger and Snoeyink 1998). Decimation is applied to provide surface models with different levels of detail, for example to improve performance in real-time display applications.

The main shortcoming of the TIN model for the implementation as a computer program is the lack of directly applicable language elements such as arrays. Examples of data structures for TINs can be found, for instance, in Renka (1996a), Ruppert (1995), or Shewchuk (1997).

For the assignment of the arcs connecting the nodes several strategies are possible. The Delaunay triangulation is considered to achieve optimal TIN configurations (Ruppert 1995). The more recent implementations of the Delaunay triangulation provide support for *constraint lines*, allowing the definition of a concave outer boundary, or holes and barriers inside the network (*constraint Delaunay triangulation*, CDT). A comprehensive overview of Delauney triangulation algorithms and software including WWW links can be found at Skiena (1998).

To preserve the characteristics of the surface the exact path of the polylines, in our case the boundaries of polygonal units, has to be reproduced in the TIN. The initial Delaunay configuration is modified slightly by forcing certain arcs to coincide with the polylines. The slight deviation from the ideal Delauney triangulation can be tolerated.

3 Volume preservation

Polygon-based data, for example the number of inhabitants in an administrative unit, may be considered as a volume which can be visualized by a prism in the perspective view of a choropleth map (Fig. 1). The height of the prism is calculated by dividing the volume by the area of the reference polygon. The volume of the depicted prisms is proportional to the number of inhabitants, the height is proportional to population density. The color or lightness of the prism tops represent density classes to demonstrate the relationship to a choropleth map in 2D.

During interpolation (the calculation of the smooth surface) parts of the volume should not be redistributed to neighboring units. The volume of the prism above a polygon should remain constant before and after the interpolation of the smooth surface, in our case the number of inhabitants. If the centroid of the polygon is used as a geometric proxy the redistribution cannot be controlled. An interpolation algorithm should use the polygon as the geometric reference, and enforce the preservation of volume for each unit.

Tobler developed a procedure to generate a smooth surface from polygon-based data which he dubbed *pycnophylactic interpolation* (Tobler 1979). Pycnophylactic, from Greek *pyknos* = mass, density and *phylax* = guard, means *volume-preserving*. The procedure consists of two main steps. In the first step the polygons and the

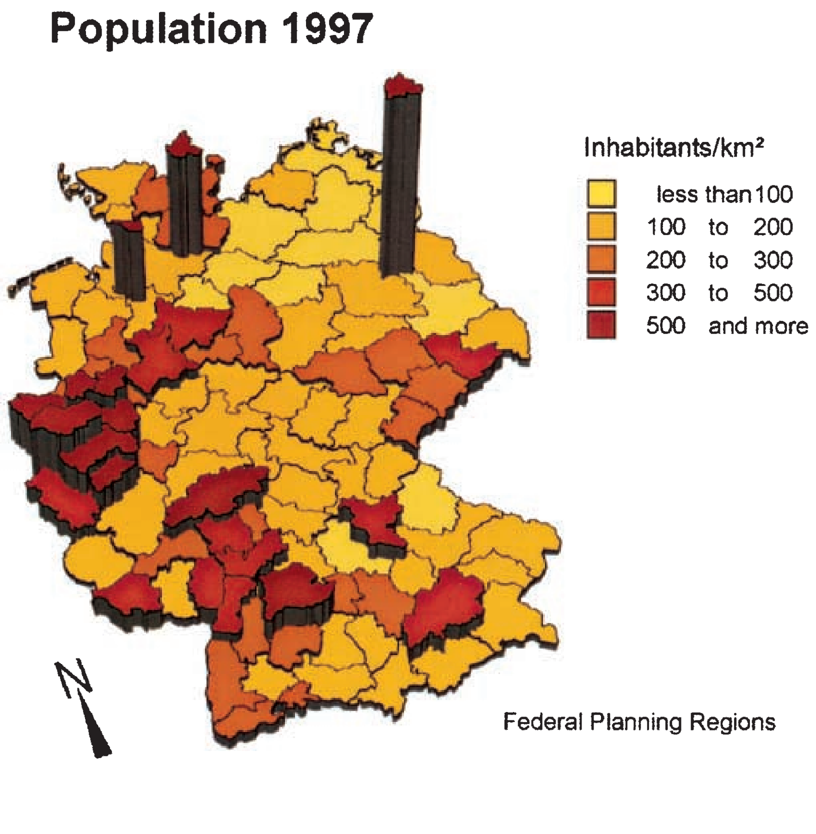

Fig. 1. Perspective view of a choropleth surface

associated volume data are converted to a regular grid with height values assigned to the grid points. In the second step the heights at the grid points are increased or decreased individually to make the surface smooth, while simultaneously enforcing the volume-preserving condition. The second step is repeated until the remaining "roughness", the deviation from the ideal smoothness, has reached a user-defined threshold, or until the maximum number of cycles is reached.

The conversion from the irregular polygons to the regular grid has the disadvantage that the original polygon boundaries and other polyline data are no longer present in the model. The size of the grid cells cannot be reduced below a certain level, for two reasons. Firstly, the doubling of the linear resolution, same number of iterations assumed, leads to a fourfold increase in memory requirement, computing time and display throughput. Secondly, the size of the grid cells influences the form of the surface. A finer grid tends to generate extended flat regions inside the polygons and steeper gradients at the polygon boundaries, an effect not considered to be useful in all cases (Rase 1998).

4 Pycnophylactic interpolation in a triangular irregular network

The volume-preserving properties of Tobler's pycnophylactic interpolation and the advantages of the triangular irregular network for preserving the geometry of lines are combined. The procedure consists of two basic steps:

1. Generation of the TIN from the boundary network by inserting additional points into the polygons.
2. Interpolation of a smooth surface including volume preservation by an iterative procedure.

4.1 Generation of the TIN

The TIN used for interpolation is built from the boundary network of the reference polygons and additional points to be inserted into the polygons. The final TIN should have the following properties:

- The triangles should be of appropriate size in relation to the reference polygons.
- Skinny or obtuse triangles should be avoided, e.g. triangles with one small or one large angle, or one small or one large side.

The number of the additional points in the polygons determines the size of the triangles. The triangles should be small enough to obtain a smooth surface. If the triangles are too small in relation to the area of the polygon the surface tends to stay flat in the center of the polygon. If the triangles are too large a certain amount of "roughness" remains in the surface. Triangles with small angles may cause discontinuities in the surface, or arithmetic problems due to small angles may occur.

Two main approaches for the generation of the TIN were implemented, the insertion of a regular grid into the polygons and the refinement of the initial TIN constructed from the boundary network.

4.1.1 Insertion of a regular grid into the polygons

A regular grid of triangles or quadrangles is inserted into each polygon (Fig. 2a). The boundary and the grid points are triangulated forming a Delaunay triangulation. The

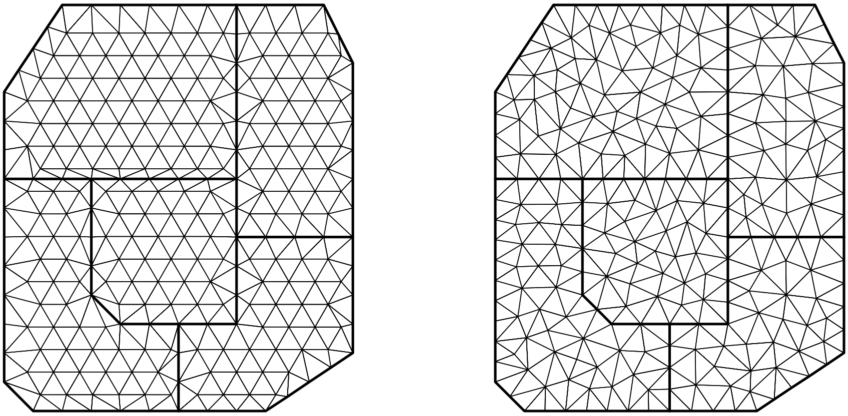

a) Regular triangular grid b) Refinement of TIN

Fig. 2. Two approaches for point generation inside the polygons

path of the boundary polylines is forced to be arcs (or edges) in the network. To minimize the number of obtuse triangles in the neighborhood of the boundaries, only grid points are inserted which have a minimum distance from the boundaries, about half the triangle height. Extra points along the boundary polylines are inserted at a distance approximating the resolution of the grid.

4.1.2 Refinement of the TIN

The second approach tested was the refinement of the initial TIN based on the boundary polylines. The refinement of a TIN is a common procedure in computer-assisted mechanical design, for example to generate a smooth surface for a mechanical part, or to generate a dense grid for simulation of the mechanical stress asserted on the real part before it is actually manufactured (*finite element method*, FEM). The objective is the generation of a *quality mesh* where upper and lower bounds for triangle size, aspect ratio, or angles are assured. Additional points (*Steiner points*) are inserted into the edges and triangles of the initial triangular network, in our case set up from the boundary polylines, in order to split large triangles and to avoid small angles.

Several approaches for the refinement of triangle meshes have been published. The algorithm of Chew (1989) was used to test if mesh refinement achieves an improvement over the more or less heuristic approach of inserting a regular grid of points. The algorithm works as follows. It is assumed that the distances of every point to another point in the initial configuration (the network of polygon boundaries) is between the minimum distance h and $h\sqrt{2}$. Additional points have to be inserted into the polylines, if necessary. A relaxation of the length condition to 2h instead of $h\sqrt{2}$ was found to be tolerable. After triangulation the circumference circle for each triangle is calculated. If the radius of the circle is $> h$, the center of the circle is added as a new node. The triangulation and the check is repeated until no new node is generated.

The procedure results in a TIN with triangles of fairly uniform size. Skinny or obtuse triangles are avoided or at least minimized (Fig. 2b).

The algorithm of Chew was easy to implement using the TRIPACK software by Renka (1996a). The implementation is not optimal but fast enough to test if the approach is useful for the pycnophylactic interpolation.

4.1.3 Results

The best interpolation results were achieved with a grid of equilateral triangles. "Best" means in this context that the resulting surface is smooth, and the amount of "roughness" is small. The smoothness is judged by visual inspection, and the overall roughness is measured by the relative variance (see next section). Obtuse and skinny triangles, however, cannot be excluded completely with the regular grid due to the irregularities of real polygon boundaries.

Other algorithms to generate quality meshes were not tested further. It was concluded from the results that the irregularity of the mesh itself and not the quality of the irregular mesh is responsible for the decline in smoothness (for some suspicions about the reason see next section). A "better" mesh in the sense of the *quality mesh* approach would not lead to a better surface. To test this hypothesis the nodes of the regular grid were displaced slightly by a random amounts with different maximum thresholds. A degradation in smoothness was observed similar to the one with the irregular grids generated by the Chew algorithm.

4.2 Interpolation of the smooth surface

During the construction of the TIN provisional density values are assigned to the nodes. The nodes which are part of the polygon boundaries are preset with the average of the neighboring polygons. Most boundary nodes belong to two adjacent polygons, and receive the average of two values. The nodes where several boundary segments meet belong to several polygons, and receive the average of more than two values. The nodes inside the polygon are preset with the height of the polygon.

The nodes on the outer boundary get a special treatment. For the first option the nodes of the outer boundary are preset with zero or the data minimum. The result is comparable to a coast with a sand beach where the land emerges gradually from the ocean. The second option is similar to a coast with steep rocks. The points on the outer boundary are preset with the height of the neighboring polygon.

The interpolation algorithm is an adaptation of the procedure which Tobler suggested for the regular rectangular grid. A short verbal description of the algorithm should suffice at this point. Readers interested in the mathematical foundations are referred to the original publication (Tobler 1979), or an abbreviated description of the procedure (Rase 1998).

The interpolation procedure consists of two steps:

1. Smoothing: The numerical value of each point is recalculated to achieve a smooth surface.

2. Redistribution of the differences: The deviation in the volume over each polygon is redistributed by increasing or decreasing the node heights for each polygon to enforce the pycnophylactic condition.

The two steps are repeated until the threshold for the overall smoothness measured by the relative variance or the maximum number of iterations are reached.

4.2.1 Smoothing

The numerical value of each node in the TIN is recalculated several times during the iteration to transform the flat tops and steep gradients at the boundaries gradually to a smooth surface. Basically the smoothness is achieved by assigning the arithmetic average of its immediate neighbors to each node in the TIN.

For each node the average of the z-values of the first ring of neighbors (Fig. 3) is calculated with inverse distance weighting (IDW, Shepard 1968).

$$zn_i = \frac{\sum_{j=1}^{m} z_j * d_j^{-p}}{\sum_{j=1}^{m} d_j^{-p}} \tag{1}$$

zn_i new value for point i
z_j value of m nearest neighbors
d_j distance to m nearest neighbors
p exponent of distance

Averaging the immediate neighbors in the TIN is similar to the *Laplace equation* criterion in Tobler's original implementation with the regular grid. The average from the neighbors replaces the previous height value of the node:

$$\partial_i = z_i - zn_i \tag{2}$$

$$z_i = zn_i \tag{3}$$

∂_i deviation at node i

For a node on the outer boundary the boundary condition is checked. In the case of a "sand beach" the node is left at zero or the data minimum.

The relative variance as a measure of the remaining general roughness of the surface (the sum of squared deviations) is calculated for all nodes:

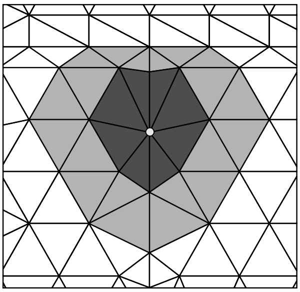

Fig. 3. First and second ring of neighbors of a node in the TIN

$$R = 100 * \frac{\sum_{i=1}^{n} \partial_i^2}{\sum_{i=1}^{n} (z_i - \bar{z})^2} \tag{4}$$

R relative variance
n number of nodes
$\bar{z}$ arithmetic average of all nodes

Inverse distance weighting is supposed to compensate for the irregularity of the TIN (the immediate neighbors have different distances to the node). The value of 2 for the exponent p (inverse squared distance) is used in many applications as an analogy to the gravity model. Tobler (1999) recommends a linear decay function by setting the exponent to 1. A value of 0 for p (unweighted average) can also be used but may cause unwanted irregularities in the surface.

The second smoothness criterion in Tobler's original implementation based on the *biharmonic equation* is approximated by adding the second ring of neighbors (neighbors of neighbors) for the calculation of the average for each node (Fig. 3). The results were similar to the results achieved with the regular grid and the biharmonic criterion (Rase 1998). The gradients at the polygon borders were steeper, and the surfaces had more extended flat areas inside the polygons.

4.2.2 Redistribution of the differences in volume

In the following step the actual volume for each polygon is computed by summing up the volumes of the prisms above the triangles inside each polygon. The difference of the actual volume to the original volume is used to compute a correction factor for the height of the nodes inside the polygon. The height is increased or decreased iteratively by a small amount until the difference between the actual volume and the original volume is smaller than a chosen threshold. A measure to avoid oscillations around the setting point in consecutive iteration steps is underrelaxation, that means that only a percentage of the adjustment value (around 90%) is applied.

Smoothing and redistribution of the differences are repeated until the relative variance is lower than a preset threshold, or the user-chosen maximum number of cycles is reached.

4.3 Example for an interpolation with a small test data set

Figure 4 shows three versions of a continuous surface interpolated from a small test example. Figure 4a displays a perspective view of a choropleth surface representing the volumes associated with polygons. The surface interpolated with a regular triangular grid inserted into the polygons (Fig. 4b) seems to be smoother than the third surface. The surface generated with TIN refinement using the algorithm of Chew (Fig. 4c) has a "bumpier" appearance, which also shows up in the value of the variance. Obviously the averaging step provides the best results if the distances from all neighbors are equal, that means, the neighbors are located on the corners of a regular hexagon. The inverse distance weighting does not compensate enough for the deviations of the irregular grid. Similar observations were made when the

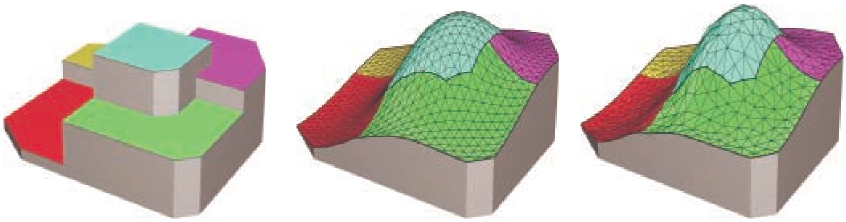

a) Choropleth map b) Regular grid c) Refinement of TIN

Fig. 4. Surfaces interpolated from a small test data set

nodes of the regular triangular grid were relocated by adding random amounts to the coordinates.

4.4 Surface from polygon-based population data

The pycnophylactic interpolation with a TIN was used to interpolate smooth surfaces from several polygon-based variables. The original boundary polylines were simplified using the Douglas-Peucker algorithm (Douglas and Peucker 1973). The data reduction was performed to minimize the number of very small, obtuse and skinny triangles. Figure 5 shows the results of the two basic methods used for the generation of the triangular network.

Both the results of the statistical calculations and the visual appearance of the interpolated surfaces support the observation made with the small data set. The regular triangular grid inserted into the polygons is the preferable solution in comparison to the TIN refinement with the algorithm of Chew.

The distribution of the population in Germany was chosen for the surface displayed in Fig. 6. To avoid false assumptions it has to be made clear that the maps do not try to represent the real distribution of inhabitants within each polygon (Federal Planning Regions, *Gebietseinheiten der Raumordnung* in Germany). For spatial analysis and regional planning on the Federal level the regions and the associated data values (volumes) are sufficiently detailed models of reality. If the planning authorities below the Federal level are interested in the distribution of

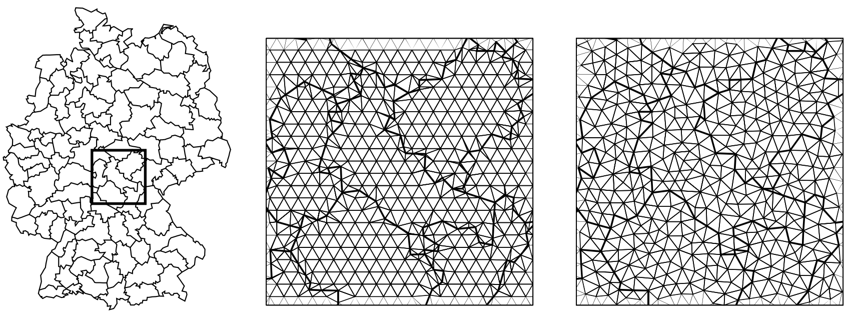

a) Boundary network b) Regular grid c) Refined TIN

Fig. 5. Boundary network and two variants of TINs in enlarged windows

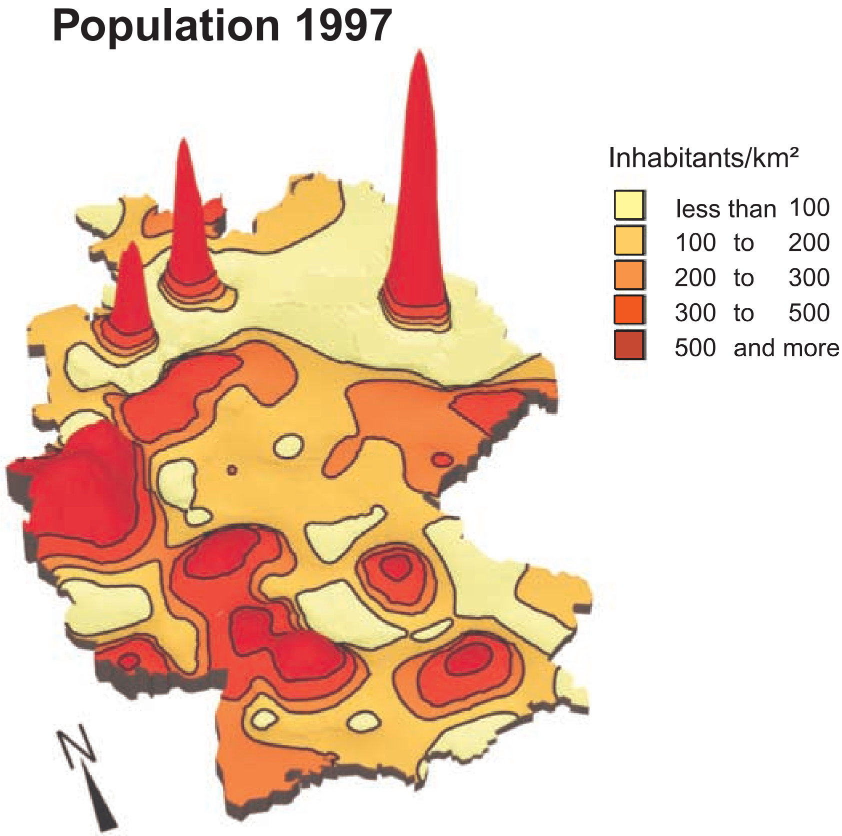

Fig. 6. Continuous surface from population data in Germany

the population within the units, they can use a finer level of regional subdivision, e.g. counties or communities.

4.5 Interpolation of index numbers

The example of the population was used in Figs. 1 and 6 because the variables *inhabitants* and *population density* are well known and easy to explain. It is at least questionable if the continuous surface is the appropriate cartographic method to display population figures (I think it is appropriate). Index numbers composed of several variables, for example variables describing the natural environment, or demographic or economic figures, are inclined to be depicted as continua. The natural environment does not change abruptly at the boundary of an administrative region. The surface expresses a certain amount of fuzziness or probability. In case of the TIN-based pycnophylactic interpolation the original boundary lines are preserved in the surface data structure and could be made visible if required.

The smooth surface in Fig. 7 was interpolated from an index number which is supposed to represent the economic deficiencies in rural areas of Germany (Irmen and Blach 1996). The indicator is composed of ten base variables describing the population density, land use, accessibility, employment, tax revenue, and average income. The base variables were combined by the method of cumulated deficit analysis (Irmen 1995).

The class with the highest values depicts the rural areas with the most severe economic problems, the class with the lowest values are the urban regions (with problems as well, of course, but different ones than the rural areas). The classes between the two extremes represent the gradations for the economic deficiency in rural areas.

In contrast to the variable used in Fig. 6 the index number does not represent a volume but a height value. The “virtual” volume can be calculated by multiplying the height by the area of the polygon.

5 Conclusions

5.1 Regular grid or TIN for the pycnophylactic interpolation?

The pycnophylactic interpolation – or a similar method with volume-preserving properties – is mandatory for the construction of a continuous surface from

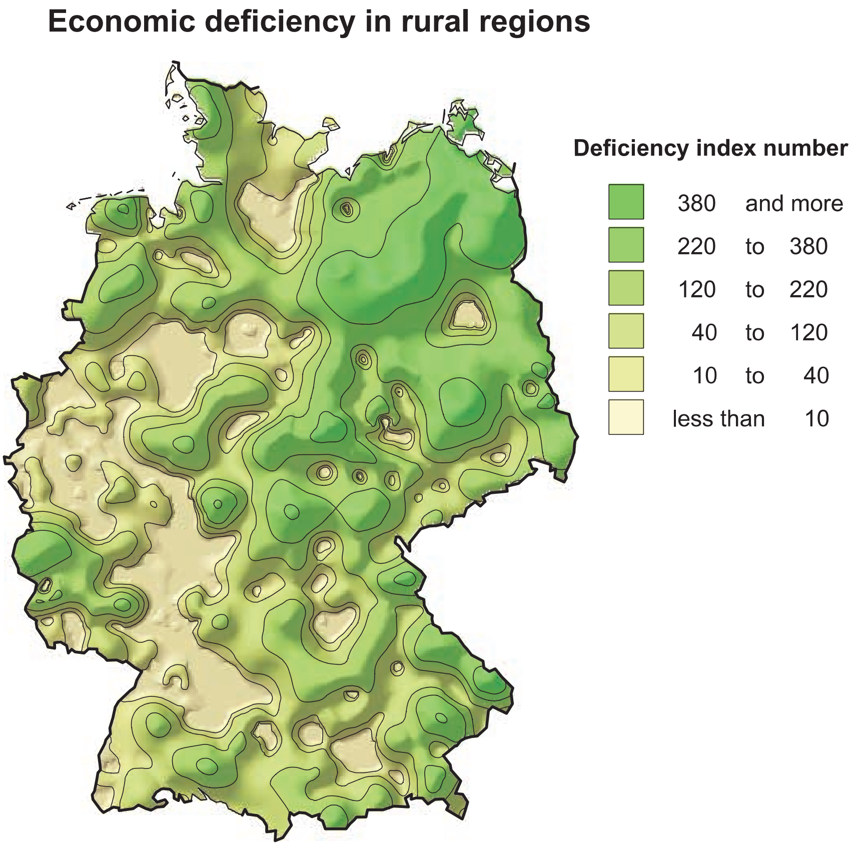

Fig. 7. Surface depicting the economic deficiency in rural regions in Germany

polygon-based data. The use of the centroid as a geometric proxy cannot assure volume preservation during the interpolation. But what are the advantages or disadvantages of using a irregular triangular network (TIN) for the pycnophylactic interpolation in comparison to a regular grid? The following points may help to get an answer which surface model should be used for the specific case.

- The error resulting from converting the polygons to a regular grid in the "classic" version of the pycnophylactic interpolation is avoided, especially for a set of polygons with an extended range in size and resolution. The geometry of the boundary polylines can be preserved to the necessary precision in the data structure of the surface.
- The number of triangles can be kept lower than the number of rectangles in a regular grid, by selection of the resolution, and by decimating the triangular mesh, if appropriate.
- Triangles or triangle meshes are the basic geometric objects for some high-performance graphic devices. The use of triangles can speed up the display process for real-time applications. On the other hand a regular grid can be converted easily into a pixel array on most graphic devices. It depends on the actual device and the application requirements if a regular grid or a TIN is more adequate.
- The TIN model is more difficult to implement than a grid. It requires more effort and provisions for the data and program structures.

5.2 Further work

A regular triangular grid of points inserted into the polygons leads to better interpolation results than a refinement of a TIN based initially on the polygon boundaries. It is suspected that the averaging step with inverse distance weighting (Shepard 1968) is probably too simple to cope with the varying distances to the neighboring nodes. A more advanced algorithm of the same family, for example the algorithm implemented by Renka (1988, 1999), might be a better solution. If smoothing for irregular triangles can be improved other algorithms to obtain quality meshes should be tested.

Although conceptual surfaces are usually considered to be smooth, discontinuities can be present sometimes, for example in case of physical, political, social and cultural boundaries. Indeed the Iron Curtain between the two Germanies until 1990 was such a barrier which will be traceable for many years, visible in the landscape, and invisible in the minds. Barriers have some intricacies concerning the data structure and the interpolation algorithm, for example the steep cliffs which may develop along the barrier lines. The option for barriers was implemented in the pycnophylactic interpolation for TIN-based surfaces to evaluate the necessary modifications to the data structure and the interpolation algorithm, and to demonstrate the use of barriers in general.

The conceptual model for barriers, the proper definition of barrier lines and the logical explanation of the interpolation results, however, seem to be more difficult than the implementation. More examples of smooth surfaces with barriers have to be generated and discussed until conclusions can be drawn about the benefit of barriers in conceptual surfaces.

References

(WWW addresses subject to frequent change, no warranty)

BBR (2000) *Raumordnungsbericht 2000*. Bundesamt für Bauwesen und Raumordnung, Bonn, Germany. http://www.bbr.bund.de

BfLR (1995) *Trendszenarien der Raumentwicklung in Deutschland und Europa*. Beiträge zu einem europäischen Raumentwicklungskonzept. Bundesforschungsanstalt für Landeskunde und Raumordnung, Bonn

Chew LP (1989) *Guaranteed-quality triangular meshes*. Technical Report 89-983, Department of Computer Science, Cornell University, Ithaca, NY. http://cs-tr.cs.cornell.edu/

Irmen E (1995) *Strukturschwäche in ländlichen Räumen – ein Abgrenzungsvorschlag*. Bundesforschungsanstalt für Landeskunde und Raumordnung, Bonn, Arbeitspapiere 15/1995

Irmen E, Blach A (1996) *Typen ländlicher Entwicklung in Deutschland und Europa*. Informationen zur Raumentwicklung, 11/12. pp 713–728

Junger B, Snoeyink J (1998) *Importance measures for TIN simplification by parallel decimation*. Proceedings 8th International Symposium on Spatial Data Handling (SDH98), Vancouver 1998

Mitasova H, Mitas L (1993) Interpolation by regularized spline with tension: I. Theory and implementation. *Mathematical Geology* 25(6): 641–655

Douglas DH, Peucker TK (1973) Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *The Canadian Cartographer* 10(2) December 1973: 112–122

Peucker TK, Fowler RJ, Little JJ, Mark D (1978) *Digital representation of three-dimensional surfaces by triangulated irregular networks* (*TIN*). Proceedings Digital Terrain Modeling Symposium, May 1978, ASP, 516–540

Rase WD (1998) *Modellierung und Darstellung immaterieller Oberflächen* (*Modelling and visualization of conceptual surfaces*). Forschungen des BBR, Band 89, Bundesamt für Bauwesen und Raumordnung, Bonn, Germany

Renka RJ (1988) Algorithm 660: QSHEP2D, Quadratic Shepard method for bivariate interpolation of scattered data. *ACM Transactions on Mathematical Software* 14(2) June 1988: 149–150

Renka RJ (1996a) Algorithm 751: TRIPACK: a constrained two-dimensional Delaunay triangulation package. *ACM Transactions on Mathematical Software* 22(1) March 1996: 1–8

Renka RJ (1996b) Algorithm 752: SRFPACK: Software for scattered data fitting with a constrained surface under tension. *ACM Transactions on Mathematical Software* 22(1) March 1996: 9–17

Renka RJ (1999a) Algorithm 790: CSHEP2D: Cubic Shepard method for bivariate interpolation of scattered data. *ACM Transactions on Mathematical Software* 25(1) March 1999: 70–73

Renka RJ (1999b) Algorithm 791: TSHEP2D: Cosine series Shepard method for bivariate interpolation of scattered data. *ACM Transactions on Mathematical Software* 25(1) March 1999: 74–77

Ruppert J (1995) A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms* 18(3): 548–585

Samet H (1990) *The design and analysis of spatial data structures*. Addison-Wesley, Reading

Schroeder WJ, Zarge JA, Lorensen WE (1992) Decimation of triangle meshes. *Computer Graphics* 26(2) July: pp 65–70

Shepard D (1968) *A two-dimensional interpolation function for irregularly-spaced data*. Proceedings ACM National Conference 1968, pp 517–524

Shewchuk JR (1997) Delaunay refinement mesh generation. PhD Thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh. http://www.cs.cmu.edu/~quake-papers/delauney-refinement.ps.gz

Skiena SS (1998) *The algorithm design manual*. Springer, Berlin Heidelberg New York

Tobler WR (1979) Smooth pycnophylactic interpolation for geographical regions. *Journal of the American Statistical Association* 74(357): 519–535

Tobler WR, Kennedy S (1985) Smooth multidimensional interpolation. *Geographical Analysis* 17(3) July 1985: 251–257

Tobler WR (1999) Personal communication

Tomlin CD (1990) *Geographic Information Systems and cartographic modelling*. Prentice Hall, Englewood Cliff, NJ

Watson DF (1992) Contouring. A guide to the analysis and display of spatial data. Pergamon Press, Oxford