Volume xx (2009), Number 2, pp. 1-4

Visualization of Polygon-based Data as a Continuous Surface

Wolf-Dieter Rase

Federal Office for Building and
Regional Planning, Bonn, Germany

Abstract

Most information for spatial analysis and research in economic and social geography is aggregated into totals associated with units which are polygons. To facilitate comparisons between the units the totals are normalized by a reference variable, such as the area of the polygon. Such kind of data is usually visualized as 2D or 3D choropleth maps, but an alternative is a continuous surface. The main property of the 3D choropleth map, the volume above the polygons, must be preserved, requiring a volume-preserving interpolation method. The average of the heights in a polygon are kept constant by changing heights of the points within the polygon. Obtaining a smooth surface requires additional points to be inserted into the polygon arranged on a regular square or triangular grid. But a regular grid may not be the best solution, especially if the polygons have irregular shapes and wide range of sizes, or if the input geometry of the polygon boundaries must be maintained in the map. In this case an irregular mesh of triangles can be used which must meet certain criteria. The triangles should not have very small or very large interior angles, to avoid arithmetic and visual problems, and the maximum area of any triangle may be limited. Good meshes with customized properties can be constructed using the Triangle program. Only minor modifications in the interpolation algorithm are required for an irregular mesh.

Categories and Subject Descriptors: G.1.1 [Interpolation, smoothing], I.2.1 [Cartography], I.3.8 [Applications]

1. Cartographic visualization of polygon-based data

The results of a census, or a poll, or other administrative data are usually published as totals for geographical regions which are closed polygons, such as counties. This reduces a large number of individual cases suitable for use by organizations such as local government. The accumulated data also preserves the privacy of individuals or the trade secrets of companies. To compare data across reference units it is normalized, for example by dividing the aggregate values by the area of the polygons, the number of inhabitants, or other appropriate variables.

1.1. Choropleth maps

The choropleth map is the most frequently used for the cartographic visualization of normalized variables associated with polygonal units. The range of the values is divided into a small number of classes each of which is represented by a unique color or pattern which is used to depict units in that class. The colors are explained in the map legend.

Strictly two-dimensional choropleth maps have the disadvantage that only class membership can be derived from the map. Actual values are not available and reference units can only be compared by their class membership. This problem is addressed by quasi three-dimensional choropleth maps in which each polygon becomes a prism of a height proportional to the value of the corresponding data (Fig. 1, left).

Different values are now clearly visible as the differences in height. Interactive mapping environments allow the map to be viewed from any angle which eliminates the problem of tall prisms obscuring shorter ones. More advanced systems use virtual reality techniques (VR), including stereograms, or physical models (Fig. 8), allowing even better appreciation of a spatial distribution.

1.2. Continuous surfaces

The polygon-based data can also be interpolated and displayed as a continuous surface (Fig.1, right). This presentation seems to facilitate the mental comprehension of demographic or economic data more effectively than the two-dimensional choropleth map. One reason for this might be that the human information processing system is more inclined towards smooth shapes. For instance, circles seem to be preferred over squares in proportional symbol maps. The other

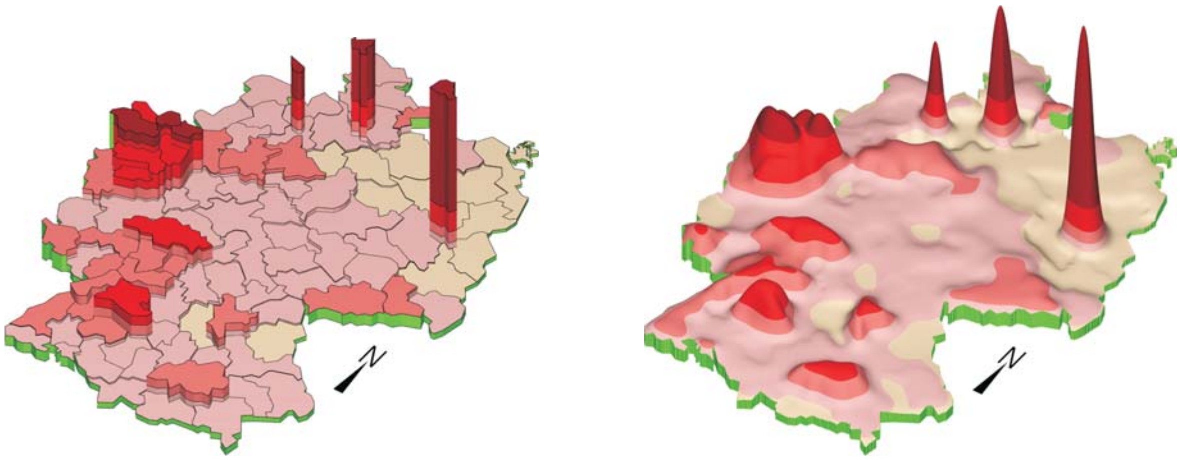

a) 3D choropleth map b) Continuous surface

Figure 1: *3D choropleth map (left) and interpolated smooth surface (right). In both cases the number of inhabitants is represented by volume and the population density by height.*

reason may be the extended range of brightness and color shades that occur in an image of the smooth surface, whereas the shading of a prismatic choropleth map is relatively uniform due to the identical orientation of the prism caps to the main light source.

In the case of most environmental variables the abrupt change in height at the boundary of a polygon on a choropleth map does not conform to reality. The transition of natural phenomena is usually gradual, and natural boundaries seldom coincide with administrative boundaries. The change in height at the administrative boundary is mostly caused by the sampling process or the calculation algorithm, and does not represent the real value of the phenomenon at the boundary.

An interpolated surface also possesses some degree of visual uncertainty or fuzziness. This property may actually be desirable in some applications, for instance, in visualizing concepts in spatial development or planning at the supranational (European) level where blurring of existing political and administrative boundaries can usefully counteract the natural inclination of national and regional authorities to focus on their own territories.

Under certain conditions an interpolated surface can also be considered as a probability surface which can be used for fuzzy reasoning in map algebra. Surfaces interpolated using volume-preserving methods can also be subjected to areal disaggregation, in which data is converted from one system of areal units to another.

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. The height of the prism is calculated by dividing the volume by the area of the reference polygon. The volume of the prisms is proportional to the number of inhabitants, the height (volume divided by the polygon area) is proportional to population density. In the graphic display isolines, isopleths, or simply the color of the prism tops facilitate the comprehension and create a correlation to the two-dimensional choropleth map.

The interpolation options in most GIS packages generate a smooth surface based on point data. Usually the centroids of the polygons are used as geometric proxies to construct the surface. This is a rather crude approach, because interpolation through sample points cannot preserve an important property of the prismatic choropleth map: The volume of the prism above a polygon should remain constant before and after the interpolation of the smooth surface. To preserve the volume for each polygonal unit the interpolation algorithm should use the polygon outline as the geometric reference. For normalized data represented by the height a "pseudo" volume can be calculated by multiplying the height value by the area of the polygon.

2. The pycnophylactic interpolation method

Tobler achieved volume-preserving interpolation of polygon-based data using a method which he dubbed *pycnophylactic interpolation* [Tob1979]. Pycnophylactic is derived from Greek *pyknos* = mass, density and *phylax* = guard, and means *volume-preserving*. Additional points inserted into the polygons allow a variation of heights within a polygon as a prerequisite for a smooth surface, whereas the average of heights within a polygon must be kept constant to ensure the volume-preserving property. The original implementation used a square grid as the data model for the surface. This

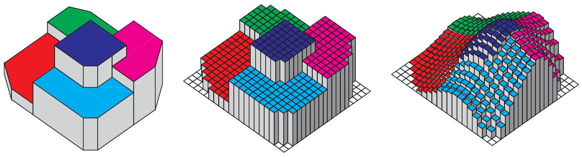

a) Prisms b) Gridding c) Smoothing

Figure 2: *Insertion of additional points (regular grid) and smoothing*

makes it relatively easy to implement as most programming languages provide support for two-dimensional arrays.

The resulting surface is known as a 2.5D surface. Any point in the *xy* plane has only one *z* value. This is valid for either the points of a regular grid or the nodes of an irregular mesh (see later). The term 2.5D (or 2½D) is inaccurate, because there are only integer number of dimensions, but in common use as a shorthand symbol.

The algorithm for the pycnophylactic interpolation consists of two main steps:

1. The polygons (Fig. 2a) are approximated by a regular grid, and the height values associated with the polygons are assigned to the grid points extended to squares (Fig. 2b).
2. The heights of the grid points are individually adjusted to approximate a smooth surface, while the volume-preserving condition for each prism or polygon is simultaneously enforced (Fig. 2c), using an iterative process.

Tobler's original implementation using a square grid is more or less straightforward. Each cycle in the iteration process is divided in two sub-steps:

2a. The surface is smoothed by averaging,

2b. The volume-preserving property is enforced by redistributing the difference between the actual height and the original height for each polygon.

Possible termination criteria for the iteration include the maximum number of cycles, a minimum value of the smoothness measure or of the sum of residuals.

2.1. Smoothing step

During the smoothing step an average value is computed for each point from the heights of the neighboring grid points. Tobler applied two kinds of filters for averaging. The first filter, derived from the Laplacian equation, uses the immediate four neighbors of a grid point (Fig. 3, 1a).

The second filter is a relative of the biharmonic equation. The average is computed from three rings of neighbors using different weights for each ring (Fig. 3, 1b). Special treatment is required for points on the edge of the grid and near the outer boundary. The points on the first and last grid lines and their missing neighbors have to be treated correctly. Other filters and smoothing procedures may also be used [DWW1979].

The pycnophylactic algorithm uses a square grid. It can be modified for a regular mesh of equilateral triangles (Fig. 3, 2). The triangular grid generally has some theoretical advantages. It can be implemented with minor modifications for the geometry and the averaging. For example six immediate neighbors are used for the smoothing step instead of four (Fig.3, 2).

2.2. Enforcing the volume-preserving property

In the second step the difference between the original volume of each polygon and the new volume produced by the smoothing step is calculated. The difference between the actual height average and the original height is redistributed by correcting the height of the points for each polygon. A measure of the roughness of the surface (the complement of the smoothness) is calculated at the end of each cycle.

At each grid point the difference between its previous value and the average value derived from the heights at neighboring points is computed:

$$d^*_{ij} = znew_{ij} - zold_{ij}$$

znew new z value, average from neighboring points
zold old z value from previous step

The difference is divided by 4 in a relaxation step designed to avoid oscillations around the value:

$$d_{ij} = d^*_{ij} / 4$$

For each polygon k the sum s^* of the relaxed values is then computed:

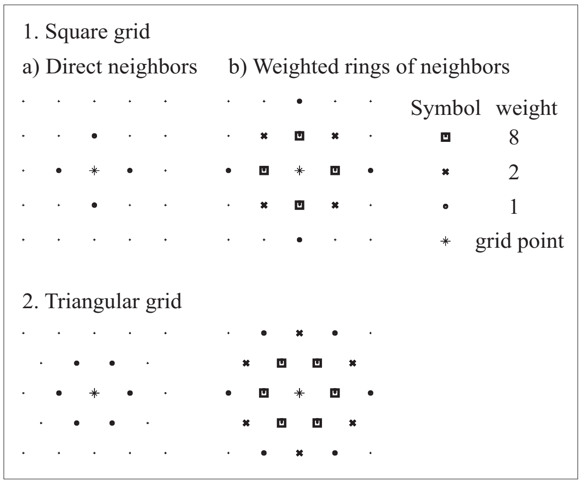

Figure 3: *Filters used for smoothing over a square and a triangular grid*

$$s_k^* = \sum d_{ij}$$

The corrections above are applied to each grid point. The sum of the heights for each polygon k is computed:

$$z_{ij} = z_{ij} + d_{ij} + s_k$$
$$zd_k = \sum_k z_{ij}$$

The difference between the new and previous values is then found:

$$l_k = \left(V_k - zd_k\right) / A_k$$

V_k original volume of polygon k

The difference is added to every grid point.

$$z_{ij} = z_{ij} + l_k$$

Figure 4 shows the first seven iteration steps in a simple test example. R is a measure of the roughness of the resulting surface. Although the roughness measure continues to decrease during the whole iteration, changes in the height of the grid points are barely visible after the first few cycles.

2.4. End of the iteration

The iteration ends when the maximum number of cycles is reached. Alternative termination criteria include smoothness and the sum of the differences between the original and the iterated point heights.

The roughness of the surface is measured by the average of the quadratic deviations of the heights at all points in a polygon between two iteration cycles [HaM1968]:

$$p = 100 * \left[\frac{\sum (znew_{ij} - zold_{ij})^2}{\sum \left(znew_{ij} - \overline{z}\right)^2} \right]$$

$znew_{ij}$ new z values of the points in a polygon
$zold_{ij}$ old z values from the previous step
n number of points in a polygon
$\overline{z}$ average of all $znew_{ij}$

The values for roughness and the changes in volume and height become smaller as the resolution of the grid is increased. Even if the total volume is preserved, there may be large changes in the individual polygons. It is therefore advisable to check the results of the smoothing for all reference units, especially when there is known to be a large difference between the heights of some neighboring polygons, which there is between Berlin and the surrounding areal units in the map of Germany (see section 3.1).

2.4. Sand beach and rock coast

The boundary of the region of interest can either modeled as a “sand beach” or a “rock coast”. The height of a sand beach gradually increases from the outer boundary to the inside of the surface. At the boundary the gradient of the surface is zero, and the height can be preset by the user: the usual value is the data minimum of all polygons.

A “rock coast” has vertical walls and an infinite gradient at the outer boundary. The height values are equal to the original heights of the polygons at the boundary. Tobler put the two coast forms in the context of the solution of elliptic partial differential equations, where the sand beach is known as the Dirichlet condition, and the rock coast as the Neumann condition [TOB1979].

If an island contains only one reference unit it is not necessary to include it in the interpolation process. Solitary polygons appear as prisms, as in a choropleth map, with vertical walls. A sand beach could be created for these polygons, but it is usually unnecessary.

The decision of coastal form is often a matter of personal preference. Sometimes the sand beach can cause peculiar surface shapes. The vertical walls of the rock coast seem to be the better choice in most cases, but the outer boundary is coarsened to match the resolution of the regular grid.

2.5. Shortcomings of the original algorithm

The original version of pycnophylactic interpolation using a square grid has some shortcomings:

- If polygons are small or have a very irregular shape they will only contain a small number of grid points or, in the worst case, no points at all. The intuitive solution is to use a grid with a higher resolution, but this requires more computing resources for the interpolation and the graphic output.
- The fineness of the grid can influence the shape of the surface. If the grid is very fine, the central parts of large

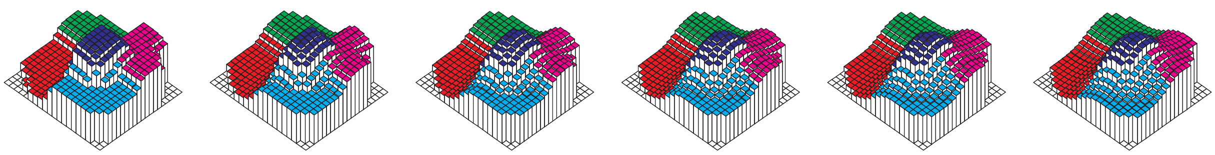

step 1, R=6.31% step 2, R=1.49% step 3, R=0.73% step 4, R=0.53% step 5, R=0.44% step 7, R=0.36%

Figure 4: *Iterative smoothing a continuous surface. R = measure of roughness*

polygons tend to stay flat because the effect of smoothing, which is based on neighborhoods, decreases with distance from the polygon boundary.

- If the difference in height between neighboring units is large it is sometimes impossible to achieve a smooth surface without considerable local deviations in volume preserving. Sometimes it helps to relax the constraints within the process, for instance by allowing negative heights in parts of a polygon. A negative population density, for example, is not realistic, but can be accepted because the average height above the polygon remains constant. Again a finer grid may improve the situation in some cases.
- The exact geometry of boundaries and other linear features cannot be preserved if they are reproduced by points on a grid. A workaround is to combine the grid points with the original boundary vertices during the output step. The height of the boundary must be estimated from the neighboring grid points, and sometimes this produces unwanted visual effects.

An alternative model for the surface must be used if it is essential that the original areas of the polygonal units and the geometric properties of the boundary lines should remain unaltered.

3. Meshes of irregular triangles

A mesh of irregular triangles, also known as triangular irregular network (TIN) is another way of modelling 2.5D surfaces [PFL1978]. A triangular model can be constructed from a map of polygonal units by connecting the points on the boundary lines with additional points within the polygons (Steiner points) by edges which form a triangular mesh. Such mesh approximates a *conforming constrained Delaunay triangulation*. A Delaunay triangulation has the property that no vertex lies inside the circumscribing circle of any triangle.

A *constrained Delaunay triangulation* has an outer boundary defining the area of analysis [HJD2006]. This boundary can consist of several simple polygons which may be of any shape and may contain holes. A constrained Delaunay triangulation is not a Delaunay triangulation in all cases, because some of its triangles might not be Delaunay.

A *conforming Delaunay triangulation* is a triangulation in which every triangle meets the Delaunay criterion. When updating such a triangulation new vertices are frequently necessary to maintain the Delaunay property while ensuring that every line segment is represented. New vertices may be inserted, and input line segments may be subdivided to improve the quality of the triangles [SHE1997].

The resolution of an irregular mesh of triangles can be locally adapted to suit the requirements of the boundary geometry and the local minima and maxima in both frequency and amplitude. In regions with high relief energy, in which significant changes of height occur over small distances, the mesh can be increased in resolution. In flat areas a lower local density, that means larger triangles, is sufficient to represent the surface.

The line segments of the boundaries become edges of the network. Additional Steiner points within the polygons may be required to achieve a sufficiently smooth surface, as well as to make the triangulation as close as possible to a constrained Delaunay triangulation.

3.1. Quality meshes

As well as meeting the Delaunay criterion, the mesh should be able to conform to certain properties appropriate for specific problems and data configurations [SHE2002]:

- The maximum or average area of the triangles should be adjustable. This determines the overall resolution of the mesh and hence the surface.
- The interior angles of the triangles should not be too small or too large. Skinny or fat triangles may cause arithmetic problems, are visually unappealing, and may also influence the shape of the interpolated surface.
- A limit on the total number of points may be required. A mesh with many small, well-shaped triangles, that also conforms strictly to the Delaunay criterion may simply be too large.

The Triangle program [SHE1996, SHE1997] provides the necessary features to construct a quality mesh. Triangle can be used as a stand-alone program, but our implementation of the pycnophylactic interpolation method calls Triangle as a procedure, with data structures as parameters.

Triangle has many options, but constraints on the maximum area of triangles, the minimum interior angle, and the maximum number of Steiner points produce an adequate interpolation in most cases. If boundary lines are very detailed, their subdivision can be prohibited. But is better to simplify the boundaries before the triangulation step to the level of detail appropriate for the scale of the map to avoid problems during interpolation.

Very tiny small segments in the boundary lines can be caused by errors during geometric data capture, untidy weeding operations, or aggregation of the boundary network from smaller polygonal units. The microscopic segments usually remain unnoticed in maps, but they may cause the formation of very small triangles during the triangulation. The small triangles can be the reason for errors in the smoothing step. Very long spikes in the surface, for instance, are caused by floating-point errors due to very small distances. When the tiny line segments are removed, the errors disappear.

Most GIS packages provide appropriate tools for line simplification. The much more detailed original boundary lines in Fig. 5 were simplified using the Douglas-Peucker algorithm [DOP 1973] available in the program ArcView by ESRI, followed by manual deletion of a few extra vertices.

The Triangle program returns the original and the additional points, the triangle definitions and the divided line

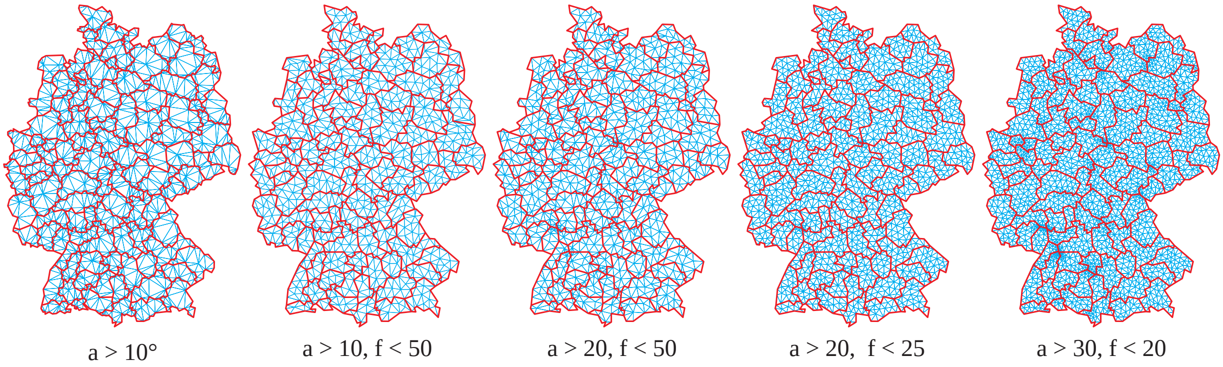

Figure 5: *Examples of quality meshes built from the boundary network (Federal Planning Regions); a = minimum interior angles, f = maximum area of triangles (mm^2 in the original scale, now not at same scale)*

segments in the boundaries. Triangle applies *exact arithmetic* to avoid accuracy problems that are inherent to floating-point numbers [SHE1994]. The sources of Triangle can be downloaded from the Triangle website (http://www.cs.cmu.edu/~quake/triangle.html).

3.2. Quality meshes of the boundary network

Figure 5 shows the effect of different combinations of minimum interior angle and maximum triangle area (not to scale) for the resulting quality mesh. The mesh becomes denser when the minimum interior angle is increased. In theory the angle can be up to 32°, but this requires a large number of triangles, some with a very tiny size.

Reducing the maximum area of a triangle increases the overall resolution of the mesh, as we would expect. Different combinations of the area and angle constraints change the pattern in which the mesh points and triangles are distributed. The Triangle program is very efficient, so an experimental search for the appropriate parameter values is not a significant cost factor.

3.3. Modifications for the triangular mesh

A triangular mesh requires some changes to the original volume-preserving interpolation. For example the distances implied by the equidistant grid lines cannot not be used for averaging, and most vertices on the boundary now belong to more than one polygon.

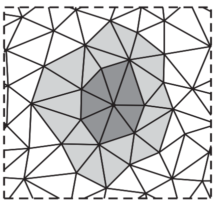

Figure 6: *Neighboring points used for averaging. Dark grey: immediate neighbors; light grey: neighbors of neighbors.*

3.3.1. Smoothing

The two filters implemented in the original version of the interpolation algorithm assume that the data points lie on a uniform square grid. We adapted these filters to the irregular mesh of triangles by modifying the averaging process to apply weights to the heights of the neighboring points to take into account the varying distance. The immediate neighbors of a point to be averaged are the points connected to it by a single triangle edge (dark grey area in fig 6). The second ring consists of the neighbors of those neighbors (light grey area).

Inverse distance is a commonly used weighting factor in simple interpolation algorithms, as Tobler put it in his famous First Geographical Law: “Everything is related to everything else, but near things are more related than distant things” [TOB1970]. In most cases squared distance is used to give a quasi-gravitational relationship, but Tobler proposes an exponent of 1 [TOB2000].

Additional smoothing methods are possible, such as more elaborate inverse distance weighting algorithms, or local polynomials. But the results of the experiments with other local interpolation methods suggest that the programming and computing costs of more complicated approaches are not justified by the results.

3.3.2. Enforcing volume preservation

The adjustment of heights to achieve a constant volume in each polygon basically the same for the irregular mesh. However, extra care must be taken with the vertices that are on polygon boundaries. Except on the outer boundary, these points are members of two or more polygons, and must be averaged accordingly.

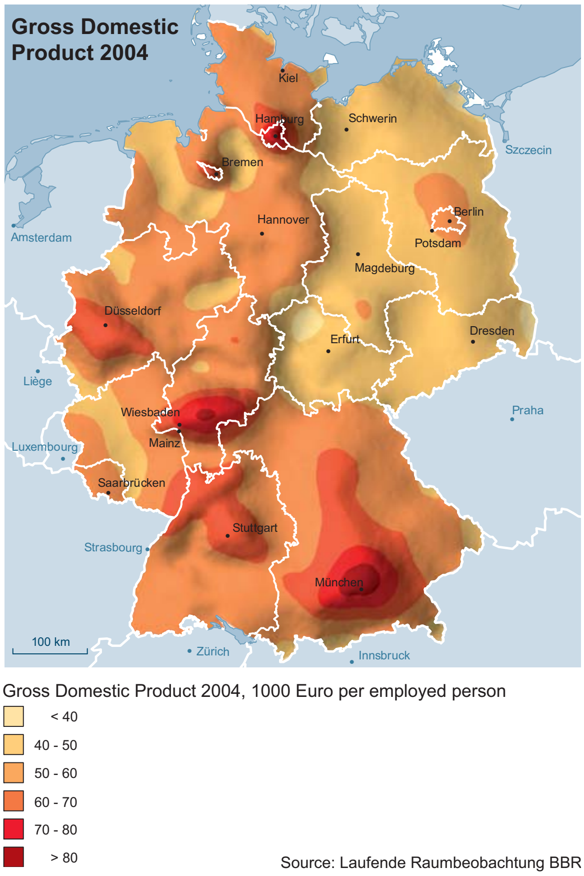

Figure 7: *Surface interpolated from the gross domestic product (GDP) per capita in Germany 2004; for regions used for spatial planning on the federal level (Planungsregionen).*

4. Examples

Pycnophylactic interpolation has been used to display spatial phenomena in many different ways. Numbers of inhabitants and population density (Fig. 1) make good demonstrations of pycnophylactic interpolation, because these variables are widely used as the base for spatial analysis, and the relationship between the volume of a polygon and its height is understood intuitively.

If the variable representing the volume is not normalized by the area of the polygon but by another factor, such as the number of inhabitants, the meaning of the relation between volume and height is not so obvious. Figure 7 shows an interpolated surface with a height corresponding to the gross domestic product (GDP) per capita for 97 regions in Germany used for spatial analysis on the Federal level. On the 3D choropleth map the volume of the prisms is not proportional to the non-normalized value of the GDP in each region.

Rapid prototyping techniques have also been used to realize physical models of surfaces from pycnophylactic interpolation [RAS2006]. Figure 8 is a photograph of a model produced on a 3D color printer (ZCorp Z510). The surface represents the average price of building lots in Euro/m² for German counties in 2003.

5. Conclusions and future work

Pycnophylactic interpolation generates a continuous and smooth surface from polygon-based data, while ensuring that the volume of the prisms above the polygons accurately reflects the data. The remaining error, either caused by coarsening the polygon boundaries or the redistribution of the volume, should be negligible if the parameter values of the algorithm are chosen properly.

5.1. When to use a regular grid and a TIN

The interpolation can be performed with two different models of the 2.5D surfaces, the regular grid of squares or equilateral triangles, or a mesh of irregular triangles. The regular grid has the advantage that the interpolation is usually faster, for a given resolution. The implicit topology of the regular grid enables efficient access to height values and faster smoothing calculations. But this technique coarsens the boundary lines to the resolution of the grid. Using an irregular mesh, the original geometry of the boundary lines is maintained, independent of the resolution specified by the user.

The resolution of an irregular mesh can be adapted to the boundaries and the different sizes and shapes of the polygons. However, for a given resolution, surfaces based on a regular grid appear subjectively smoother than those based

Figure 8: *Physical 3D model produced by 3D printing, representing a smooth surface interpolated from polygon-based data.*

on a irregular mesh, but such observations are not reflected in measurements of roughness. This apparent discrepancy needs further investigation, as well as the arithmetic errors in the averaging step caused by tiny line segments on the boundary.

When the geometric property of the boundary lines must be maintained the triangular mesh will always be preferred. But if approximate boundary lines are allowable the regular grid gives faster interpolation, and a smoother-looking surface.

5.2. Application of tetrahedral meshes

The assessment of triangle quality is based on the 2D data model, and takes no account of the actual shapes of the triangles in three dimensions. But on steep hills the interior angles of the tilted triangles become much smaller than their projection in the plane. These skinny triangles may be a cause for the perceived roughness of the surfaces based on the irregular mesh.

The user can influence the triangle shape, within certain bounds, by supplying a function to Triangle. But this is little more than a workaround in the context of interpolation. If a tetrahedral mesh program equivalent to the Triangle program becomes available it would certainly be worthwhile to see whether it could be used to improve the results of the pycnophylactic interpolation.

The application of tetrahedral meshes for the pycnophylactic interpolation has an additional benefit. The triangular mesh in the 2D plane allows 2.5D surfaces only, where a point in the plane can have only one *z* value. Barriers which inhibit the averaging between neighboring polygons are not possible, because the resulting vertical wall along the boundary contradicts the 2.5D restriction. A workaround used in earlier implementations by generating parallel lines at a very small distance from the boundary line would cause very tiny triangles during triangulation, and subsequently the mentioned arithmetic problems. The tetrahedral mesh overcomes the limitation of the 2.5D surface by allowing more than one *z* value at a point in the plane, and hence vertical walls along boundary lines in case of barriers.

Acknowledgements

I would like to thank Waldo Tobler, the originator of the pycnophylactic interpolation method, for his continuing support and advice, Jonathan Shewchuk for allowing the scientific community to use his Triangle program, and John Woodwark for his valuable advice concerning language and style.

References

[DoP1973] Douglas, D. H., Peucker T. K.: Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. The Canadian Cartographer, Vol. 10, No. 2, December 1973, 112–122

[DWW1979] Dyn, N., Wahba, G., Wong, W. H.: Comment to Waldo Tobler‘s smooth pycnophylactic interpolation for geographical regions. Journal of the American Statistical Association, Vol. 74, No. 357, 1979, 530–535

[HaM1968] Harbaugh, J. W., Merriam, D. F.: Computer applications in stratigraphic analysis. John Wiley & Sons, New York

[HjD2006] Hjelle, Ø., Dæhlen, M.: Triangulations and applications. Springer, Berlin

[PFL1978] Peucker, T. K. , Fowler, R. J., Little, J. J., Mark, D.: Digital representation of three-dimensional surfaces by triangulated irregular networks (TIN). Proceedings Digital Terrain Modeling Symposium, May 1978, ASP, 516–540

[Ras2006] Rase, W.-D.: Physical models for cartographic applications. In: Wiener Schriften zur Geographie und Kartographie, Band 17, Wien 2006, 286–291

[She1994] Shewchuk, J. R.: An introduction to the conjugate gradient method without the agonizing pain, August 1994. http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf

[She1996] Shewchuk, J. R.: Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In: Lin, M. C., Manocha, D. (ed.) Applied Computational Geometry: Towards Geometric Engineering. Lecture Notes in Computer Science, Vol. 1148, 1996, 203–222 http://www.cs.berkeley.edu/~jrs/papers/triangle.pdf

[She1997] Shewchuk, J. R.: Delaunay refinement mesh generation. Ph.D. thesis, Technical Report CMU-CS-97-137, Carnegie Mellon University, Pittsburgh, PA http://www.cs.cmu.edu/~quake-papers/delaunay-refinement.pdf

[She2002] Shewchuk, J. R.: What is a good linear element? Interpolation, conditioning, and quality measures. Proceedings 11th International Meshing Roundtable, Sandia National Laboratories, 2002, 115–126 http://www.cs.berkeley.edu/~jrs/papers/elem.pdf

[Tob1970] Tobler, W. R.: A computer model simulating urban growth in the Detroit region. Economic Geography, vol. 46, no. 2, 1970, 234–240

[Tob1979] Tobler, W. R.: Smooth pycnophylactic interpolation for geographical regions. Journal of the American Statistical Association, Vol. 74, No. 357, (September) 1979, 519–535 http://www.geog.ucsb.edu/~tobler/publications/pdf_docs/geog_analysis/Pycno.pdf

[Tob2000] Tobler, W. R.: personal communication

URLs checked in November 2008