

[FC-05-022] Geometric Primitives and Algorithms

Abstract

Geometric primitives are the representations used and computations performed in a GIS that concern the spatial aspects of the data, data objects described by coordinates. In vector geometry, we distinguish in zero-, one-, two-, and three-dimensional objects, better known as points, linear features, areal or planar features, and volumetric features. A GIS stores and performs computations on all of these. Often, planar features form a collective known as a (spatial) subdivision. Computations on geometric objects show up in data simplification, neighborhood analysis, spatial clustering, spatial interpolation, automated text placement, segmentation of trajectories, map matching, and many other tasks. They should be contrasted with computations on attributes or networks.

There are various kinds of vector data models for subdivisions. The classical ones are known as spaghetti and pizza models, but nowadays it is recognized that topological data models are the representation of choice. We overview these models briefly.

Computations range from simple to highly complex: deciding whether a point lies in a rectangle needs four comparisons, whereas performing map overlay on two subdivisions requires advanced knowledge of algorithm design. We introduce map overlay, Voronoi diagrams, and Delaunay triangulations and mention algorithmic approaches to compute them.

Keywords: algorithms and efficiency, basic measures, computational geometry, Delaunay triangulation, map overlay, spatial data structure, topological data structures, topology, triangular irregular network, Voronoi diagrams

Author & citation

Kreveld, M. van (2020). Geometric Primitives and Algorithms. The Geographic Information Science & Technology Body of Knowledge (2nd Quarter 2020 Edition), John P. Wilson (ed.). DOI: [10.22224/gistbok/2020.2.6](https://doi.org/10.22224/gistbok/2020.2.6).

This topic is also available in this earlier edition:

DiBiase, D., DeMers, M., Johnson, A., Kemp, K., Luck, A. T., Plewe, B., and Wentz, E. (2006). Geometric primitives. The Geographic Information Science & Technology Body of Knowledge. Washington, DC: Association of American Geographers.

Explanation

1. [Definitions](#)
2. [Representation of Spatial Data](#)
3. [Basic Computations in Geometry](#)
4. [Geometric Algorithms and Efficiency](#)

1. Definitions



Raster-vector debate: the discussion of advantages and disadvantages of raster-based models and vector-based models when storing spatial data.

Topological data structure: a vector-based model for storing subdivisions where adjacency of zero-, one-, and two-dimensional features are explicitly represented.

Triangular Irregular Network: a type of topological data structure where all bounded regions are triangles. They are important for representing elevation data.

Geometric algorithm: a high-level description of an approach to solve a computational task involving geometry, which is sufficiently precise to allow mathematical efficiency analysis.

Map overlay: the process of combining two subdivisions into one, such that each cell of the output subdivision lies in exactly one cell of each of the input subdivisions.

Voronoi diagram: subdivision of the plane induced by a set S of points where cells correspond to regions where one point of S is closest. Cells are bounded by so-called Voronoi edges and Voronoi vertices.

Delaunay triangulation: subdivision of the plane induced by a set S of points where all cells are triangles, the vertices of the triangles are the points of S , and for every triangle, the circumcircle of its three vertices does not contain any point of S .

2. Representation of Spatial Data

There are two fundamentally different ways to represent space: with a raster representation and with a vector representation. A raster representation is simpler. The whole space of interest is subdivided into small square-shaped cells that form a grid. Each grid cell is an indivisible element that gets a unique value. On a land-use map, a grid cell may receive the value "forest" or "street", for example. On a rainfall map, a cell would get the number of millimeters precipitation. A vector representation allows more detail because it is not bound to a globally chosen cell size. Essentially, the only explicit reference to space in a vector model is via the coordinates of points. Line segments are simply the linear connections between two points. Areas or polygons are the regions bounded by sequences of line segments. Most GIS support both raster and vector structures for their thematic map layers (Duckham and Worboys, 2004; Peuquet, 1984).

This entry focuses on vector-based representation and computation. The advantages and disadvantages of raster and vector models have been discussed extensively in the past in the so-called raster-vector debate. It should be added that raster nor vector models are suitable for all spatial data. For example, for fuzzy data and objects with indeterminate boundaries (Burroughs and Frank, 1996), adaptations are necessary.

2.1 Representation of Subdivisions

One of the most important geometric structures that a GIS must store is the (planar) subdivision. For example, the states of the USA with their boundaries form a subdivision and so do the parcels of land in a land-use map. When we store the boundaries of all



regions, the regions themselves are implicitly the delineated features between those boundaries. The boundaries themselves are polygonal lines (often called lines, which is geometrically incorrect) that can be described by a sequence of points. These points can each be represented by two coordinates.

In the early days of automated cartography and GIS, the common models to represent subdivisions were referred to as spaghetti and pizza (Duckham and Worboys, 2004; Peuquet 1984). The spaghetti model consists of an unstructured collection of polygonal lines, each represented by a sequence of coordinate pairs. Endpoints of polygonal lines correspond to junctions in the subdivision, which are present in three or more polygonal lines. This model makes it difficult to retrieve the boundaries of a single region: all polygonal lines that contribute to the boundary of a region must be determined by a linear scan, which is time-consuming. Then the boundaries must be combined into one polygon boundary.

The pizza model consists of an unstructured collection of region boundaries. For every region in the subdivision there is one (outer) boundary, which is a closed polygonal line. "Closed" means that the polygonal line is in fact a cycle and the begin- and endpoints are the same point. While this representation makes operations on regions easier, it has substantial data duplication because every boundary point is stored (at least) twice. Discovering the regions adjacent to a given region can only be done by a linear scan, which is again inefficient.

The topological data model solves the aforementioned issues at the expense of being more complicated. Here points, edges, and regions should be seen as objects that are interlinked by pointers in a well-defined and structured manner (Duckham and Worboys, 2004). This allows the GIS to traverse the boundaries of regions and access adjacent regions by local operations, giving the model its efficiency. A topological data model can store any embedded planar graph and its faces.

2.2 Representation of Elevation Models

Elevation models as abstract entities are functions from the plane to the real numbers, in a formula: $f(x,y)=z$. They can represent elevation above sea level, annual precipitation, lead concentration in the top soil, and anything else where a location on the Earth's surface has a unique value.

A grid based elevation model, often called a Digital Elevation Model or DEM, is a common and practical way to store such functions. The value in each cell is then the average value over the area that the cell represents.

A vector based elevation model is the alternative. The most common one of these is the TIN, or Triangular Irregular Network. It consists of a collection of points that have three coordinates, (x,y,z) , and the points are connected by line segments so that the bounded region between the points is filled with non-overlapping triangles. By assumption, the elevation model function value at a point on a triangle is the weighted linear interpolation over the three points that are the vertices. The weights are the barycentric coordinates. More visually, we use the plane that passes through the three three-dimensional vertices as the linear interpolator over the triangle.

Tetrahedral meshes are three-dimensional extensions of TINs and can represent volumetric



data. Spatial data can be combined with time to model spatio-temporal data, for example when modeling trajectories of moving entities. These extensions are beyond the scope of this entry.

3. Basic Computations in Geometry

When a road is modeled by a sequence of line segments in some map layer, and a country boundary is represented by a sequence of line segments in a different map layer, we may wish to know where the road exactly leaves the country. This occurs at a point whose coordinates can be computed: it will lie at the intersection point of two line segments. Another basic computation is to compute the center of the circle that passes through three given points. One way to determine this center is to construct two of the three bisecting lines of two of the points and intersecting these. Such basic computations are common in GIS. To most users, they happen behind the scenes.

Similar but slightly less basic operations occur when distances start to play a role. If a river is modeled by a sequence of line segments, then the region within 1 km from a river is modeled by a polygon whose boundary has both straight line segments and circular arcs. So intersection primitives must also deal with non-linear (curved) objects, and a topological data structure must support non-linear boundaries.

4. Geometric Algorithms and Efficiency

Stepping up from basic geometric computations to a higher level, we arrive at computations on data sets that in principle can have any size. Consider a set of points that are sightings of some animal, and we want to know the extent of its habitat. We may then want to compute the two sightings that are furthest apart. When there are 200 points in the set, we can try all 19,900 pairwise distances and report the largest one. When there are 2000 points, there are 1,990,000 pairwise distances and determining them all becomes inefficient. The objective of geometric algorithms research is to find ways to accomplish geometric tasks in ways that scale well in the data size, while guaranteeing the correct output (de Berg et al., 2008). Such algorithms are considered efficient. Indeed, it is possible to compute the two points furthest apart in a point data set with far fewer distance computations than all pairwise distances.

In a more exact formulation, we parametrize the input size (the data set size) by the letter n , and study the function that gives the required number of computations to solve the task as a function of n . The method that considers all pairwise distances will use a number of computations that grows as a quadratic function in n . The method is then called an $O(n^2)$ time method. The best method for computing the two most distant points requires a number of computations that grows proportional to $n \log n$. This method is then called an $O(n \log n)$ time method. Since n^2 as a function grows much faster than $n \log n$, the latter method is a better, more efficient algorithm that scales well with growing data sets. The symbol O in the expressions $O(n^2)$ and $O(n \log n)$ is pronounced order-of; its precise definition is technical and beyond the scope of this entry (Cormen et al., 2008).

4.1 Map Overlay



Thematic map overlay is the most essential GIS analysis task for site planning. Different map layers are combined in order to determine candidate locations for a new factory or an urban extension. Requirements on candidate locations may be that they avoid nature areas, have good accessibility by car and train, et cetera. The conjunction of these requirements that need data from different map layers requires in essence the overlay of a number of planar subdivisions. It yields a new subdivision where the suitable regions are delineated.

Algorithmically, the map overlay operation is done efficiently by a method known as plane sweep. It comes down to a structured way of merging the two input subdivisions into one using efficient dynamic data structures, in particular, binary search trees. With this method it is possible to compute the overlay of two subdivisions consisting of n line segments together, in $O(n \log n)$ time, provided that the output subdivision has no more than $O(n)$ line segments itself (Bentley and Ottmann, 1979; de Berg et al., 2008). In practice this condition is nearly always fulfilled.

4.2 Voronoi Diagrams and Delaunay Triangulations

For a given set of points, a Voronoi diagram is a subdivision that partitions the plane into regions where the same point of the set is the closest one. Voronoi diagrams have very many applications. Besides GIS, they are used in robotics and motion planning, in graphics, and in biology (Aurenhammer et al., 2013).

The mathematical properties of Voronoi diagrams are well understood. Every region is convex, and the locations where three regions meet are center points of circles through three of the input points that have no other input points inside. The Voronoi diagram is the basis of an advanced spatial interpolation method called natural neighbor interpolation (Farin, 1990; Sibson, 1981).

The Voronoi diagram on n points in the plane can be computed efficiently in various ways. The plane sweep approach, the divide-and-conquer approach, and randomized incremental construction all lead to $O(n \log n)$ time algorithms (Aurenhammer et al., 2013; de Berg et al., 2008).

For a given set of points, a Delaunay triangulation is a subdivision of the convex hull of the points into triangles in such a way that for each triangle, the circumcircle of its corners does not contain any other point of the set. It is surprising that this condition actually yields a subdivision with triangular faces. It has the property that over all possible triangulations of the same input points, the Delaunay triangulation maximizes the smallest occurring angle in the triangles. This implies that the triangles have “reasonable shapes”, to the extent possible with these input points (Aurenhammer et al., 2013; Duckham and Worboys, 2004). For this reason the Delaunay triangulation may be the triangulation of choice to convert points with values into a Triangular Irregular Network. It will have better interpolation properties than many other triangulations on the same set of points.

Voronoi diagrams and Delaunay triangulations are intimately connected. For every triple of Voronoi regions that meet in a point, there is a corresponding triangle in the Delaunay triangulation and vice versa. Equivalently, for every two adjacent Voronoi regions there is a line segment between the corresponding points in the Delaunay triangulation and vice versa. It will therefore be no surprise that Delaunay triangulations can also be computed in $O(n \log n)$ time using the aforementioned approaches.






Figure 1. Voronoi Diagrams and Delaunay triangulations. Image source: author.

4.3 Other Geometric Algorithms for GIS

There is a host of other tasks where geometric algorithms are important. These include processing of acquired data for storage, analyzing stored data, and visualizing data. Many of these tasks are overviewed by van Kreveld (2018).

References

- [Aurenhammer, F., Klein, R., and Lee, D.-T. \(2013\). Voronoi Diagrams and Delaunay Triangulations. World Scientific.](#)
- [Bentley, J.L., and Ottmann, T.A. \(1979\). Algorithms for reporting and counting geometric intersections. IEEE Transactions on Computers, \(9\), 643-647.](#)
- [Berg, M. de, Cheong, O., Kreveld, M. van, and Overmars, M. \(2008\). Computational Geometry - Algorithms and Applications, 3rd Edition. Springer.](#)
- [Burrough, P. A. and Frank, A. U. \(eds.\) \(1996\) Geographic objects with indeterminate boundaries. Proceedings, GISDATA Specialist Meeting #2 on Geographical Objects with Undetermined Boundaries. London: Taylor & Francis.](#)
- [Cormen, T., Leiserson, C., Rivest, R., and Stein, C. \(2009\). Introduction to Algorithms, 3rd Edition. Cambridge, MA: the MIT Press.](#)
- [Duckham, M., and Worboys, M. \(2004\). GIS - A Computing Perspective. 2nd Edition. CRC Press.](#)
- [Farin, G. \(1990\). Surfaces over Dirichlet tessellations. Computer Aided Geometric Design, 7\(1-4\), 281-292.](#)
- [Peuquet, D. J. \(1984\). A Conceptual Framework and Comparison of Spatial Data Models.](#)

[Cartographica 21\(4\), 66-113.](#)

[Sibson, R. \(1981\). A Brief Description of Natural Neighbor Interpolation. In: Barnett, V., Ed., Interpreting Multivariate Data, John Wiley & Sons, New York, 21-36.](#)

[van Kreveld, M. \(2018\). Geographic Information Systems. Chapter 59 in: Goodman, J.E., O'Rourke, J., and Toth, C.D. \(editors\). Handbook of Discrete and Computational Geometry, CRC Press, 3rd edition.](#)

