Spatial interpolation

Spatial interpolation is a pervasive operation in GIS. Although it is often used explicitly in analysis, it is also used implicitly, in various operations such as the preparation of a contour map display, where spatial interpolation is invoked without the user's direct involvement. Spatial interpolation is a process of intelligent guesswork, in which the investigator (and the GIS) attempt to make a reasonable estimate of the value of a continuous field at places where the field has not actually been measured.

Spatial interpolation is an operation that makes sense only from the continuous-field perspective. The principles of spatial interpolation are discussed in Section 4.5; here the emphasis is on practical applications of the technique and commonly used implementations of the principles.

Spatial interpolation finds applications in many areas:

■ In estimating rainfall, temperature, and other attributes at places that are not weather stations and where no direct measurements of these variables are available.

■ In estimating the elevation of the surface between the measured locations of a DEM.

■ In resampling rasters, the operation that must take place whenever raster data must be transformed to another grid.

■ In contouring, when it is necessary to guess where to place contours between measured locations.

In all of these instances, spatial interpolation calls for intelligent guesswork, and the one principle that underlies all spatial interpolation is the Tobler Law (Section 3.1) – 'all places are related but nearby places are more related than distant places'. In other words, the best guess as to the value of a field at some point is the value measured at the closest observation points – the rainfall here is likely to be more similar to the rainfall recorded at the nearest weather stations than to the rainfall recorded at more distant weather stations. A corollary of this same principle is that in the absence of better information, it is reasonable to assume that any continuous field exhibits relatively smooth variation – fields tend to vary slowly and to exhibit strong positive spatial autocorrelation, a property of geographic data discussed in Section 4.6.

Spatial interpolation is the GIS version of intelligent guesswork.

three methods of spatial interpolation are discussed: Thiessen polygons; inverse-distance weighting

(IDW), which is the simplest commonly used method; and Kriging, a popular statistical method that is grounded in the theory of regionalized variables and falls within the field of geostatistics.

Thiessen polygons

Thiessen polygons were suggested by Thiessen as a way of interpolating rainfall estimates from a few rain gauges to obtain estimates at other locations where rainfall had not been measured. The method is very simple: to estimate rainfall at any point take the rainfall measured at the closest gauge. This leads to a map in which rainfall is constant within polygons surrounding each gauge, and changes sharply as polygon boundaries are crossed. Although many GIS users associate polygons defined in this way with Thiessen, they are also known as Voronoi and Dirichlet polygons (Box 14.5). They have many other
uses besides spatial interpolation:

■ Thiessen polygons can be used to estimate the trade areas of each of a set of retail stores or shopping centers.

■ They are used internally in the GIS as a means of speeding up certain geometric operations, such as
search for nearest neighbor.

■ They are the basis of some of the more powerful methods for generalizing vector databases.

As a method of spatial interpolation they leave something to be desired, however, because the sharp
change in interpolated values at polygon boundaries is often implausible.
Figure 14.23 shows a typical set of Thiessen polygons. If each pair of points that share a Thiessen polygon boundary is connected, the result is a network of irregular triangles. These are named after Delaunay, and frequently used as the basis for the triangles of a TIN representation of terrain

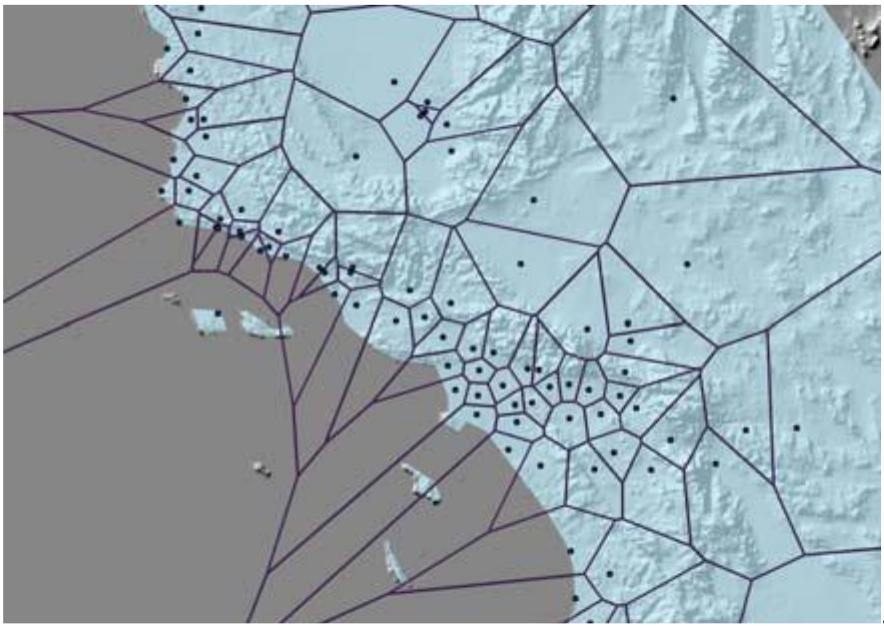
.

Density estimation

Density estimation is in many ways the logical twin of spatial interpolation – it begins with points and ends with a surface.

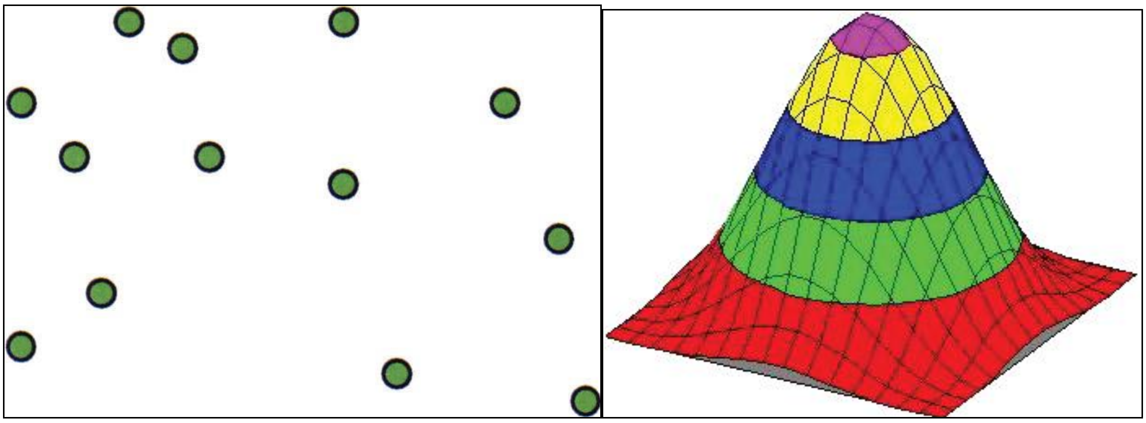

But conceptually the two approaches could not be more different, because one seeks to estimate

the missing parts of a continuous field from samples of the field taken at data points, while the other creates a continuous field from discrete objects. Figure 14.28 illustrates this difference. The dataset can

be interpreted in two sharply different ways. In t he first, itis interpreted as sample measurements from a continuous field, and in the second as a collection of discrete objects. In the discrete-object view there is nothing between the objects but empty space – no missing field to be filled in through spatial interpolation. It would make no sense at all to apply spatial interpolation to a collection of discrete

Density estimation makes sense only fro m the discrete-object perspective, and spatial

interpolation only from the field perspective.

Although density estimation could be applied to any type of discrete object, it is most often applied to the estimation of point density, and that is the focus here. The most obvious example is the estimation of population density, and that example is used in this discussion, but it could be equally well applied to the density of different kinds of diseases, or animals, or any other set of welldefined points.

Consider the mainland of Australia. One way of defining its population density is to take the entire population and divide by the total area – on this basis the 1996 population density was roughly 2.38 per sq km. But we know that Australia’s settlement pattern is very non-uniform, with most of the population concentrated in five coastal cities (Brisbane, Sydney, Melbourne, Adelaide, and Perth). So if we looked at the landscape in smaller pieces, such as circles 10 km in radius, and computed population density by dividing the number of people in each circle by the circle’s area, we would get very different results depending on where the circle was centered. So, in general, population density at a location, and at a spatial resolution of d, might be defined by centering a circle at the location and dividing the total population within the circle by its area. Using this definition there are an infinite number of possible population density maps of Australia, depending on the value selected for d. And it follows that there is no such thing as population density, only population density at a spatial resolution of d. Note the similarity between this idea and the previous discussion of slope – in general, many geographic themes can only be defined rigorously if spatial resolution is made explicit, and much confusionresults in GIS because of our willingness to talk about themes without at the same time specifying spatial resolution.

Density estimation with a kernel allows the spatial resolution of a field of population density to be made explicit.

The theory of density estimation formalizes these ideas. Consider a collection of point objects, such as those shown in Figure 14.29. The surface shown in the figure is an example of a kernel function, the central idea in density estimation. Any kernel function has an associated length measure, and in the case of the function shown, which is a Gaussian distribution, the length measure is a parameter of the distribution – we can generate Gaussian distributions with any value of this parameter and they become flatter and wider as the value increases. In density estimation, each point is replaced by its kernel function and the various kernel functions are added to obtain an aggregate surface, or continuous field of density. If one thinks of each kernel as a pile of sand, then each pile has the same total weight of one unit. The total weight of all piles of sand is equal to the number of points and the total weight of sand within a given area, such as the are ashown in the figure, is an estimate of the total population in that area. Mathematically, if the population density is represented by a field ρ(x, y), then the total population

within area A is the integral of the field function over tha t area, that is:

$$P = \int_A \rho \, \mathrm{d}A$$

A variety of kernel functions are used in density estimation, but the form shown in Figure 14.29 is perhaps the commonest. This is the traditional bell curve or Gaussian distribution of statistics and is encountered elsewhere in this book in connection with errors in the measurement of position in two dimensions (Section 6.3.2.2). By adjusting the width of the bell it is possible to produce a range of density surfaces of different amounts of smoothness.