

14. Query, Measurement, and Transformation

Geographic Information Systems and Science

SECOND EDITION

Paul A. Longley, Michael F. Goodchild, David J. Maguire, David W. Rhind

© 2005 John Wiley and Sons, Ltd



Outline

- What is spatial analysis?
- Queries and reasoning
- Measurements
- Transformations



Spatial analysis

- Turns raw data into useful information
 - ▣ by adding greater informative content and value
- Reveals patterns, trends, and anomalies that might otherwise be missed
- Provides a check on human intuition
 - ▣ by helping in situations where the eye might deceive

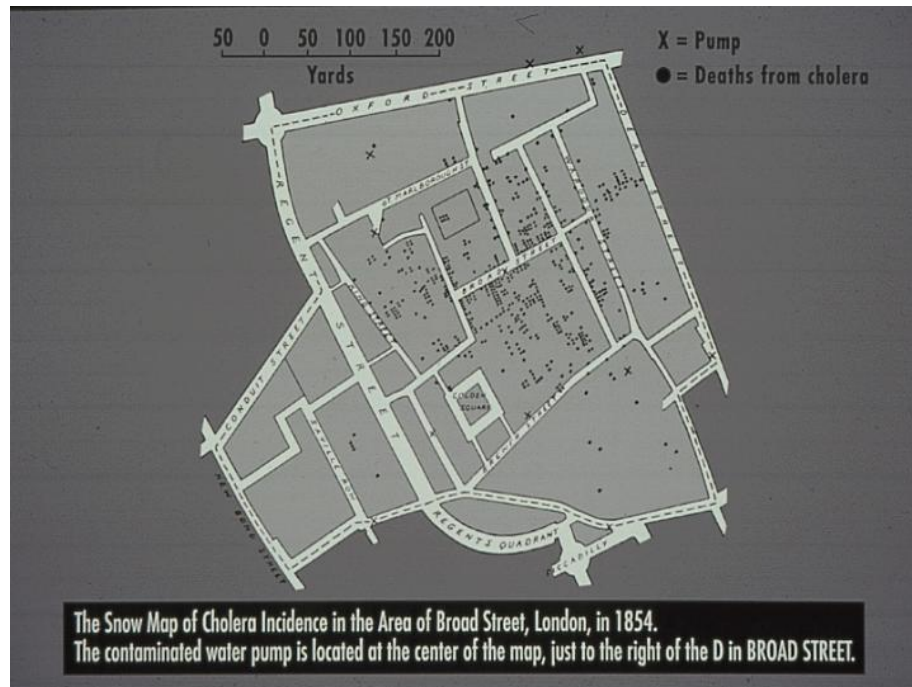


Definitions

- A method of analysis is spatial if the results depend on the locations of the objects being analyzed
 - ▣ move the objects and the results change
 - ▣ results are *not invariant* under relocation
- Spatial analysis requires both attributes and locations of objects
 - ▣ a GIS has been designed to store both



The Snow map



Shows the locations of cholera cases in the Soho area of London during an outbreak in the 1850s. The map shows that the outbreak was centered on a pump in Broad Street, and provided evidence in support of Dr John Snow's hypothesis that contaminated drinking water was causing the outbreak



The Snow map

- Provides a classic example of the use of location to draw inferences
- But the same pattern could arise from contagion
 - ❑ if the original carrier lived in the center of the outbreak
 - ❑ contagion was the hypothesis Snow was trying to refute
 - ❑ today, a GIS could be used to show a sequence of maps as the outbreak developed
 - ❑ contagion would produce a concentric sequence, drinking water a random sequence



Types of spatial analysis

- There are literally thousands of techniques
- Six categories are used in this course, each having a distinct conceptual basis:
 - ▣ Queries and reasoning
 - ▣ Measurements
 - ▣ Transformations
 - ▣ Descriptive summaries
 - ▣ Optimization
 - ▣ Hypothesis testing

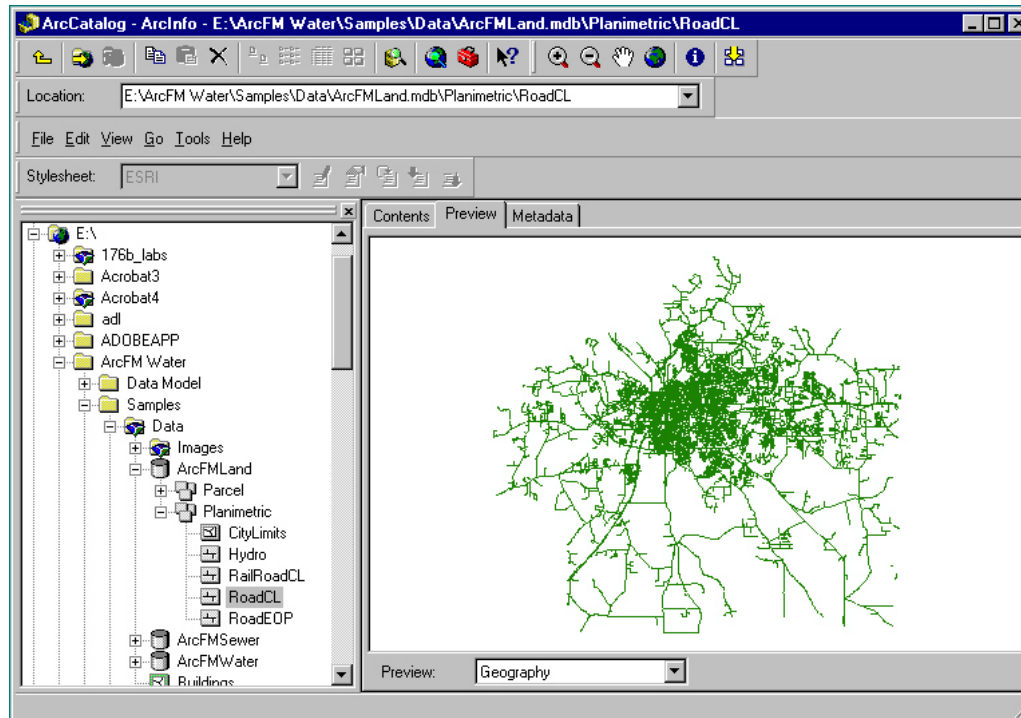


Queries and reasoning

- A GIS can respond to queries by presenting data in appropriate *views*
 - and allowing the user to interact with each view
- It is often useful to be able to display two or more views at once
 - and to link them together
 - linking views is one important technique of *exploratory spatial data analysis* (ESDA)



The catalog view



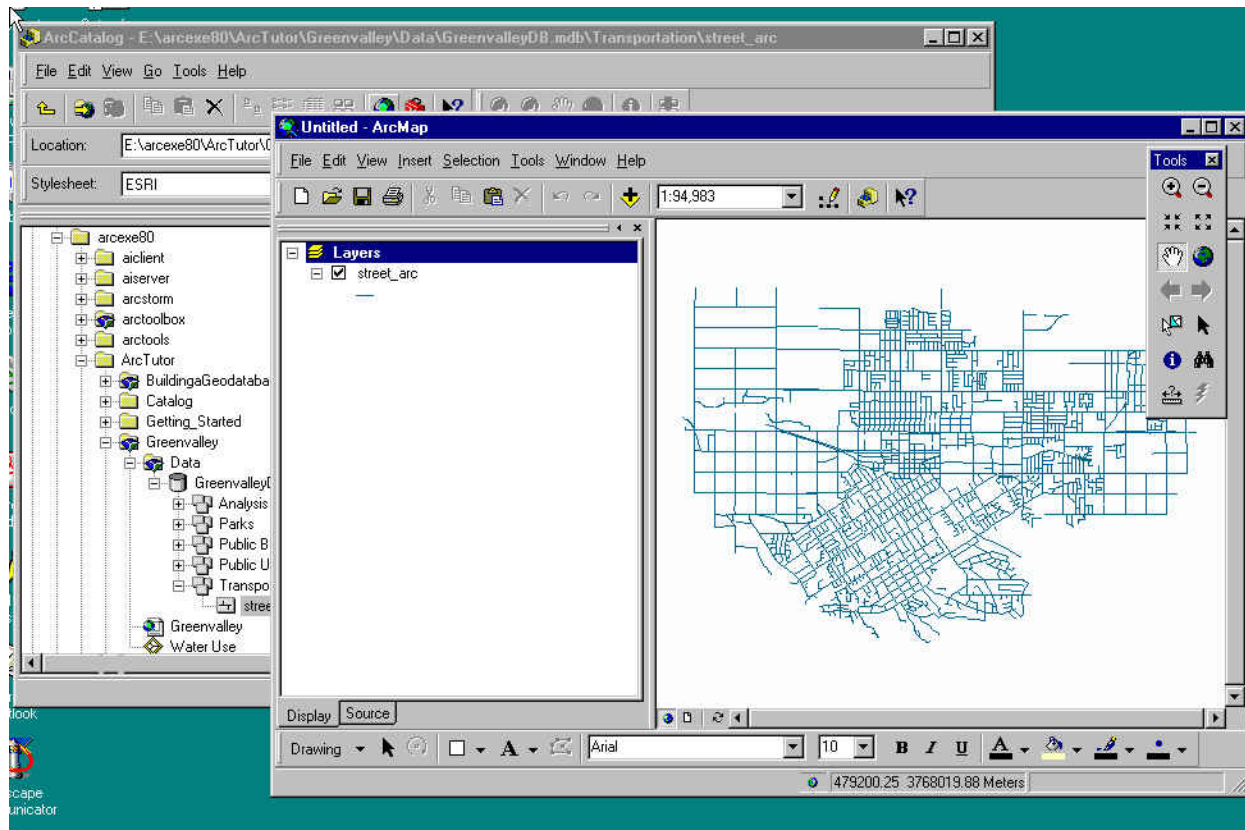
Shows folders, databases, and files on the left, and a preview of the contents of a selected data set on the right. The preview can be used to query the data set's metadata, or to look at a thumbnail map, or at a table of attributes. This example shows ESRI's ArcCatalog.



SECOND EDITION

The map view

A user can interact with a map view to identify objects and query their attributes, to search for objects meeting specified criteria, or to find the coordinates of objects. This illustration uses ESRI's ArcMap.





The table view

FID_1	Shape	NAME	Shape_Length	Type
1	Polyline		136.500014433901	4
2	Polyline		267.503634104294	4
3	Polyline		403.25717951764	4
4	Polyline		806.22123346338	4
5	Polyline		403.006998608611	3
6	Polyline		799.93784833492	4
7	Polyline		800.741232233492	5
8	Polyline		53.400023458049	5
9	Polyline		398.754707416395	5
10	Polyline		123.104856726918	5
11	Polyline		95.6591882595931	5
12	Polyline		40.0053830742638	5
13	Polyline		135.607670712223	5
14	Polyline		90.6600280511209	5
15	Polyline		30.5051829119346	5
16	Polyline		91.3370072599947	5
17	Polyline		158.455728217768	5
18	Polyline		108.500000101048	5
19	Polyline		266.572276546268	5
20	Polyline		269.274931124029	5
21	Polyline		88.7598024628139	5
22	Polyline		134.601737966072	5
23	Polyline		82.888220664282	5

Here attributes are displayed in the form of a table, linked to a map view. When objects are selected in the table, they are automatically highlighted in the map view, and vice versa. The table view can be used to answer simple queries about objects and their attributes.



Measurements

- Many tasks require measurement from maps
 - ▣ measurement of distance between two points
 - ▣ measurement of area, e.g. the area of a parcel of land
- Such measurements are tedious and inaccurate if made by hand
 - ▣ measurement using GIS tools and digital databases is fast, reliable, and accurate




Measurement of length

- A metric is a rule for determining distance from coordinates
- The Pythagorean metric gives the straight-line distance between two points on a flat plane
- The Great Circle metric gives the shortest distance between two points on a spherical globe
 - given their latitudes and longitudes



Issues with length measurement


- The length of a true curve is almost always longer than the length of its polyline or polygon representation





Issues with length measurement

- Measurements in GIS are often made on horizontal projections of objects
 - ▣ length and area may be substantially lower than on a true three-dimensional surface

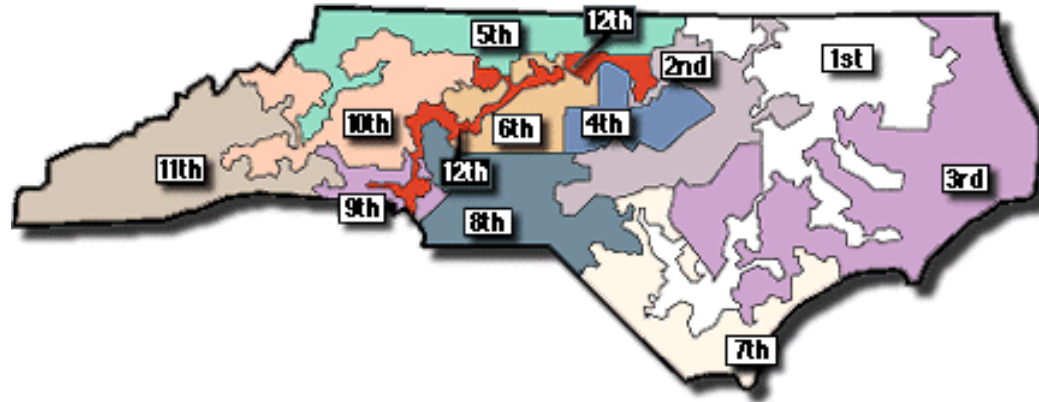




Measurement of shape

- Shape measures capture the degree of contortedness of areas, relative to the most compact circular shape
 - ❑ by comparing perimeter to the square root of area
 - ❑ normalized so that the shape of a circle is 1
 - ❑ the more contorted the area, the higher the shape measure

Shape as an indicator of gerrymandering in elections



The 12th Congressional District of North Carolina was drawn in 1992 using a GIS, and designed to be a *majority-minority* district: with a majority of African American voters, it could be expected to return an African American to Congress. This objective was achieved at the cost of a very contorted shape. The U.S. Supreme Court eventually rejected the design.

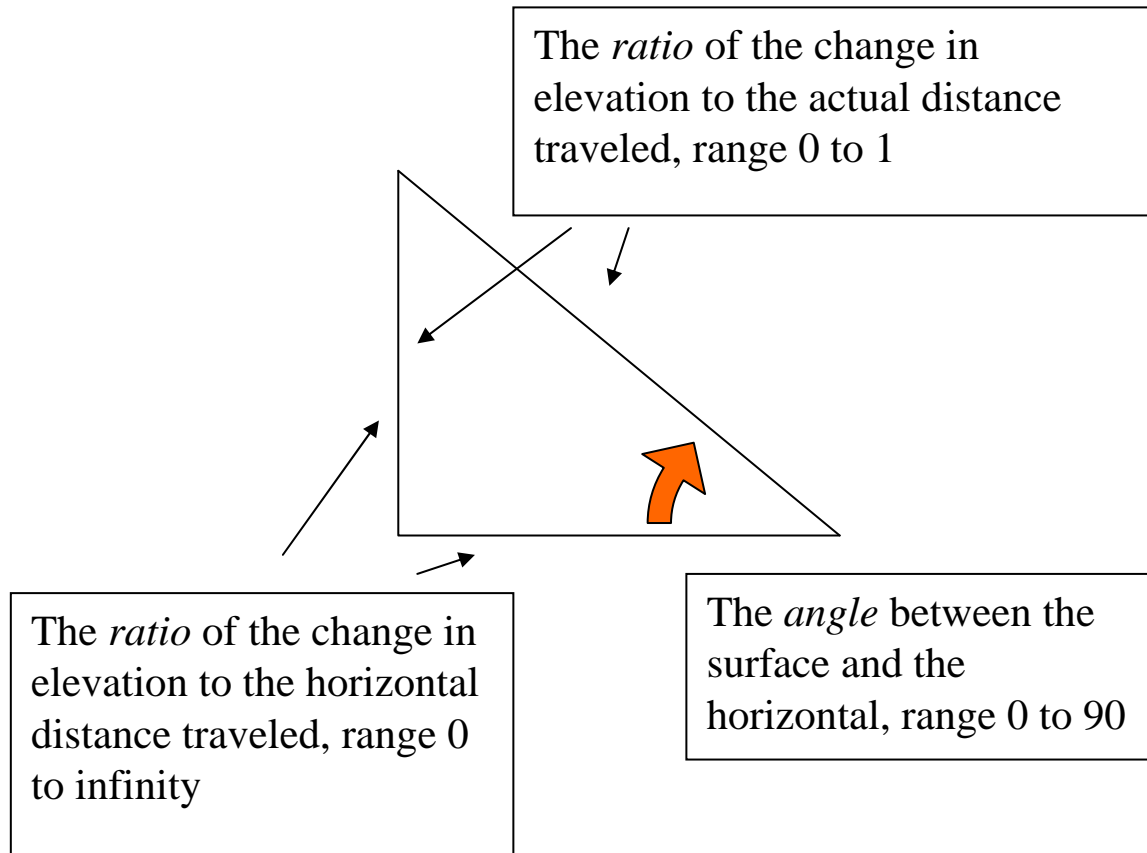


Slope and aspect

- Calculated from a grid of elevations (a digital elevation model)
- Slope and aspect are calculated at each point in the grid, by comparing the point's elevation to that of its neighbors
 - usually its eight neighbors
 - but the exact method varies
 - in a scientific study, it is important to know exactly what method is used when calculating slope, and exactly how slope is defined



Alternative definitions of slope





Transformations

- Create new objects and attributes, based on simple rules
 - ▣ involving geometric construction or calculation
 - ▣ may also create new fields, from existing fields or from discrete objects



Buffering (dilation)



- Create a new object consisting of areas within a user-defined distance of an existing object
 - e.g., to determine areas impacted by a proposed highway
 - e.g., to determine the service area of a proposed hospital
- Feasible in either raster or vector mode




Buffering

Polyline

Point




Polygon





Raster buffering generalized

- Vary the distance buffered according to values in a *friction* layer



- City limits
- Areas reachable in 5 minutes
- Areas reachable in 10 minutes
- Other areas




Point-in-polygon transformation

- Determine whether a point lies inside or outside a polygon
 - ❑ generalization: assign many points to containing polygons
 - ❑ used to assign crimes to police precincts, voters to voting districts, accidents to reporting counties



The point in polygon algorithm



Draw a line from the point to infinity in any direction, and count the number of intersections between this line and each polygon's boundary. The polygon with an odd number of intersections is the containing polygon: all other polygons have an even number of intersections.

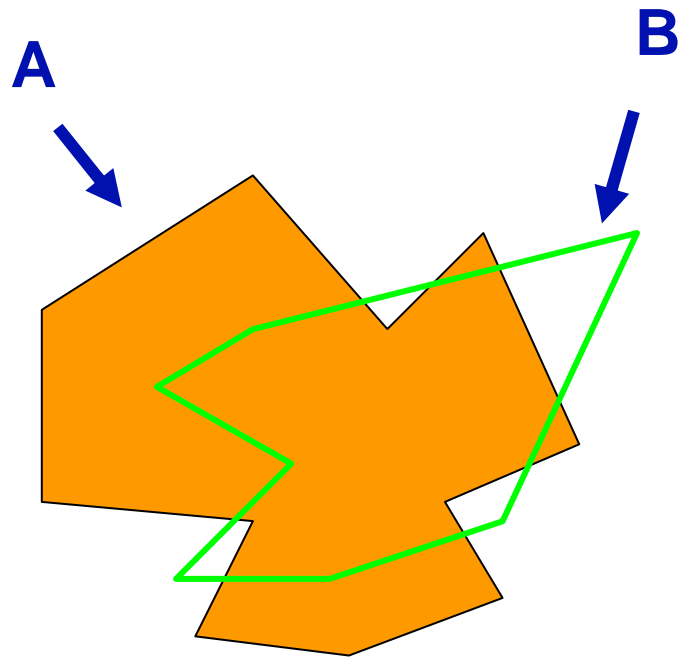


Polygon overlay

- Two case: for discrete objects and for fields
- Discrete object case: find the polygons formed by the intersection of two polygons. There are many related questions, e.g.:
 - do two polygons intersect?
 - where are areas in Polygon A but not in Polygon B?



Polygon overlay, discrete object case




In this example, two polygons are intersected to form 9 new polygons. One is formed from both input polygons; four are formed by Polygon A and not Polygon B; and four are formed by Polygon B and not Polygon A.



Polygon overlay, field case

- Two complete layers of polygons are input, representing two classifications of the same area
 - e.g., soil type and land ownership
- The layers are overlaid, and all intersections are computed creating a new layer
 - each polygon in the new layer has both a soil type and a land ownership
 - the attributes are said to be concatenated
- The task is often performed in raster

Polygon overlay, field case



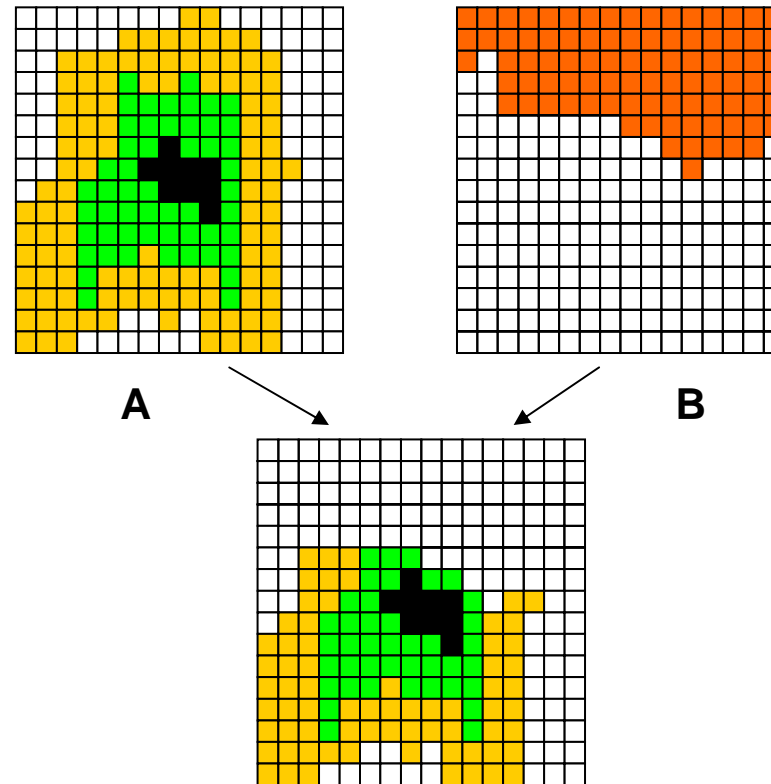
A layer representing a field of land ownership (colors) is overlaid on a layer of soil type (layers offset for emphasis). The result after overlay will be a single layer with 5 polygons, each with a land ownership value and a soil type.



Spurious or sliver polygons

- In any two such layers there will almost certainly be boundaries that are common to both layers
 - e.g. following rivers
- The two versions of such boundaries will not be coincident
- As a result large numbers of small sliver polygons will be created
 - these must somehow be removed
 - this is normally done using a user-defined tolerance

Overlay of fields represented as rasters



The two input data sets are maps of (A) travel time from the urban area shown in black, and (B) county (red indicates County X, white indicates County Y). The output map identifies travel time to areas in County Y only, and might be used to compute average travel time to points in that county in a subsequent step.



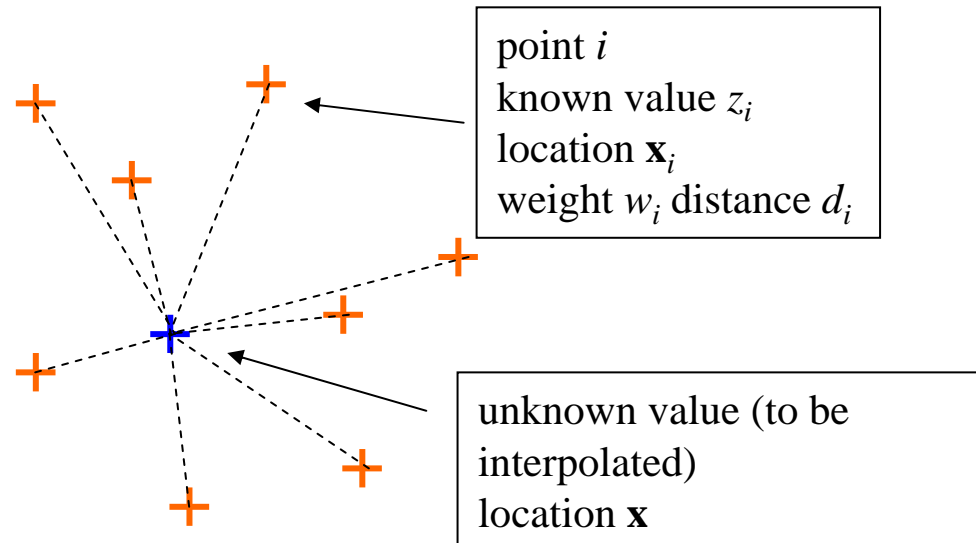
Spatial interpolation

- Values of a field have been measured at a number of sample points
- There is a need to estimate the complete field
 - to estimate values at points where the field was not measured
 - to create a contour map by drawing isolines between the data points
- Methods of spatial interpolation are designed to solve this problem



Inverse distance weighting (IDW)

- The unknown value of a field at a point is estimated by taking an average over the known values
 - ▣ weighting each known value by its distance from the point, giving greatest weight to the nearest points
 - ▣ an implementation of Tobler's Law



$$z(\mathbf{x}) = \frac{\sum_i w_i z_i}{\sum_i w_i}$$

The estimate is a weighted average

$$w_i = 1/d_i^2$$

Weights decline with distance

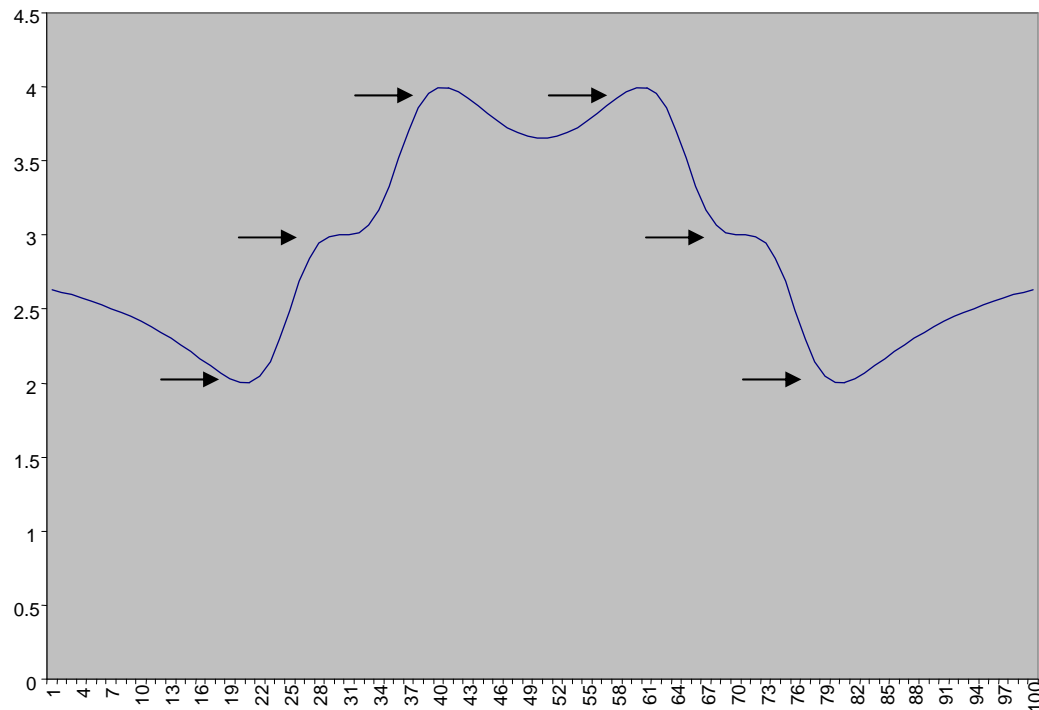


Issues with IDW

- The range of interpolated values cannot exceed the range of observed values
 - ▣ it is important to position sample points to include the extremes of the field
 - ▣ this can be very difficult



A potentially undesirable characteristic of IDW interpolation




This set of six data points clearly suggests a hill profile. But in areas where there is little or no data the interpolator will move towards the overall mean. Blue line shows the profile interpolated by IDW



Kriging

- A technique of spatial interpolation firmly grounded in geostatistical theory
- The *semivariogram* reflects Tobler's Law
 - ▣ differences within a small neighborhood are likely to be small
 - ▣ differences rise with distance

*One half the mean squared
difference (semivariance)*



A semivariogram. Each cross represents a pair of points. The solid circles are obtained by averaging within the ranges or *bins* of the distance axis. The solid line represents the best fit to these five points, using one of a small number of standard mathematical functions.



Stages of Kriging

- Analyze observed data to estimate a semivariogram
- Estimate values at unknown points as weighted averages
 - obtaining weights based on the semivariogram
 - the interpolated surface replicates statistical properties of the semivariogram



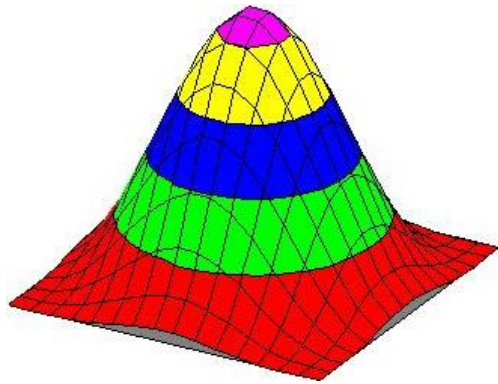
Density estimation and potential

- Spatial interpolation is used to fill the gaps in a field
- Density estimation creates a field from discrete objects
 - the field's value at any point is an estimate of the density of discrete objects at that point
 - e.g., estimating a map of population density (a field) from a map of individual people (discrete objects)




The kernel function

- Each discrete object is replaced by a mathematical function known as a kernel
- Kernels are summed to obtain a composite surface of density
- The smoothness of the resulting field depends on the width of the kernel
 - narrow kernels produce bumpy surfaces
 - wide kernels produce smooth surfaces



A typical kernel function

The result of applying a 150km-wide kernel to points distributed over California



When the kernel width is too small (in this case 16km, using only the S California part of the database) the surface is too rugged, and each point generates its own peak.

