

GIS Error Models and Visualization Techniques for Spatial Variability of Soils

Michael F. Goodchild. *National Center for Geographic Information and Analysis, and Department of Geography, University of California, Santa Barbara, CA 93106-4060, USA.*

Abstract. The paper begins by describing the major data sources, applications, and software characteristics of GIS, and then compares them within a consistent terminological framework that emphasizes the digital representation of continuously varying spatial data. Because the spatial continuum can be discretized in many different ways, and because each necessarily approximates the truth, GIS data are subject to error and uncertainty. The paper reviews the models of error that have been developed in recent years for spatial data and examines their use in GIS. Choice of representation is also driven to some extent by the needs of analysis and modeling, although these have emerged relatively recently. The paper looks at the functional requirements of modeling, and includes discussion of error propagation. The results of GIS are often presented in visual form, and the paper concludes with a review of recent research in visualization of spatial data.

Introduction. A variety of technologies for collecting, handling, and analyzing spatial data have emerged in the past three decades. Of these, GIS in particular is increasingly used in mapping the spatial variability of soils. A geographic information system (GIS) is a somewhat broad and nebulous concept - a system for the input, storage, manipulation, and output of geographically referenced data (for general overviews see Star and Estes, 1991; Maguire, Goodchild, and Rhind, 1991; Burrough, 1986). Although there is wide acceptance of this general definition, the actual specification of a GIS can vary widely depending on the assumptions that its designers are willing to make about its uses. For example, a GIS designed to process only data derived by remote sensing would be very different from one designed to process only data derived from maps.

The purpose of this paper is to develop a general conceptual framework for geographical data, and soil mapping in particular, and to embed GIS within it. Because the technology creates approximate representations of geographical distributions, issues of error and uncertainty are critical to its successful use. The GIS field has developed several apparently incompatible approaches to error description, and the paper presents a general framework for overcoming this significant impediment. Finally, spatial data technologies are inherently visual, and the paper concludes with some observations on the role of visualization in understanding error and uncertainty and their implications.

The Nature of Geographical Data. At the most fundamental level, geographical data can be defined as a collection of facts about places. The atom of spatial data is the tuple $\langle x, y, z \rangle$, where x and y define a place, and z some fact about the place, such as soil class. A convenient way to define the difference between spatial and geographical data is to insist that for geographical data, the coordinates (x, y) be defined by some system of measurement on the Earth, such as latitude/longitude, or the UTM (Universal Transverse Mercator) coordinate system.

Current measurement technologies place fairly coarse limits on one's ability to determine geographical position accurately. A hand-held GPS receiver, for example, can determine absolute position on the Earth to perhaps 1000 sq m, or approximately the area of one arc-second of latitude and longitude, outside the polar regions. The problems of approximating the complex shape of the Earth with simple mathematical functions also place limits on absolute positioning. On the other hand, it is possible to establish relative positions with millimeter accuracy, even over distances of hundreds of kilometers, using modern surveying techniques.

For many types of geographical facts, it is appropriate to conceive of a measurement z being made at a point, or infinitely small area, even though the accuracy of that point's location may be limited. For example, ground elevation can be measured at a point, as can air temperature, rainfall, or atmospheric pressure. Although a rain gauge will average rainfall over a discrete radius, that radius will be much less than the positional error in the location of the rain gauge, and irrelevant in any likely analysis.

For other types of geographical facts, the definition of the value at a point requires an observation over some significant radius around the point. Soil classes, for example, can only be defined over discrete areas that substantially exceed the area covered by any one soil sample, and depend on the specific level in the taxonomic hierarchy. This spatial resolution of the fact should not be confused with the positional accuracy of the point, or with measures of the size of the study area. In a typical example, the positional accuracy of the observation might be 1000 sq m using GPS, while the spatial resolution of soil class might be 1 hectare, implying that it is necessary to observe an area of 1 hectare to determine the soil class recorded at its central point. Unfortunately, neither measure is explicitly defined in much soil mapping practice.

In any discussion of positional accuracy it is common to mix units - some parameters may be expressed as areas, some as radii, and some as diameters. To avoid confusion, area measures will be used consistently throughout the paper. Thus the size of a pixel or raster cell, minimum area of a polygon, and positional accuracy of a point will all be expressed in area measure. In the case of positional accuracy of a point, the measure will express the area of a circle centered on the point within which the true location is assumed to lie with some determined probability.

In this conceptual framework, it is possible to determine the value of any fact z at any location (x,y) . For some types of geographical data, such as ground elevation, the value at any location is potentially unique, and thus an infinite number of tuples would be required to capture geographical variation over a finite area completely. In practice, measurement systems determine values only at a discrete number of locations, or adopt other strategies to limit the volume of data that must be handled.

Fields and Entities. Two distinct mental models underlie the various strategies that have been devised for the representation of geographical variation (Goodchild, 1992). In the field view, geographical variation can be conceived as a set of distinct variables $\{z_1, z_2, \dots, z_n\}$, each with a precise definition that allows the variable to be measured at any point (x,y) .

Mathematically, geography is a multivariate or vector field. Some variables may be measured on continuous (interval or ratio) scales, while others, such as soil class, may be discrete (nominal or ordinal). Because all variables show some degree of spatial dependence or autocorrelation, continuous variables may be differentiable, although ridge lines, cliffs, and faults are obvious sources of exceptions. Similarly, the values of discrete variables may be constant over substantial contiguous areas.

While the field view may be appropriate for rigorous definition and precise measurement, it is generally inadequate for interpretation and many forms of modeling, where it is replaced by the entity view (the terms field and entity have many precise and imprecise synonyms in the geographical data handling field). In this mental model, geography is conceived as an empty space littered with various kinds of objects, with associated characteristics. For example, interpretation of a raw field of ground elevations would yield such topographic objects as mountain, valley, watershed, ridge, pass, or saddle. Interpretation of a raw field of atmospheric pressure might yield high, low, trough, or front. Many models address the behavior of these discrete entities, rather than the continuous fields from which they may have been derived.

In some cases, the derivation of entities may be well-defined. However, not all maxima in a continuous field of ground elevation are significant enough to be called summits - the rule base of interpretation is likely to be complex. Definitions of entities often fail to give them precise footprints (the summit of a mountain is often better defined than its geographic limits; there are no precisely defined limits to the Atlantic Ocean) and entities of the same class may overlap (the Gulf of Mexico may be considered part of the Atlantic Ocean) or have hierarchical relationships with each other.

The field and entity views are mental models of geographical variation that allow the mind to comprehend what is potentially an infinitely complex world. In everyday life the entity view is dominant, and reflected in the construction of topographic maps, or the giving of navigational directions. The State of Minnesota is described as having 10,000 lakes, not a continuous field of "lakeness" that can be evaluated objectively at any point on the surface of the state. In science, on the other hand, both views are in common use; atmospheric pressure is captured and recorded by sampling a field, and the field view is reflected in the finite difference simulation models used to predict future conditions; but weather forecasts provide the public with predictions of the behavior of highs, lows, and fronts as discrete entities moving in an otherwise empty space.

Digital Representation. Although fields and entities are useful mental models of real geography, they both contain a potentially infinite amount of information. The digital representation of entities is relatively straightforward. A finite set of point entities can be represented with no information loss as a series of tuples $\langle x, y, z \rangle$, and the numerical precision of a digital representation normally exceeds both the positional accuracy of the coordinates and the measurement accuracy of z . A line is conventionally represented as a polyline, an ordered set of points assumed to be connected by straight line segments. Although accuracy can be improved by increasing the density of points, the polyline convention has two significant disadvantages: the length of a polyline is normally less than

the curve it approximates; and it is difficult to estimate tangents or perpendiculars from a polyline representation. Area entities are conventionally represented as polygons, with similar problems regarding the measurement of perimeters, perpendiculars, and tangents.

The representation of fields offers more alternatives, and at the same time raises more problems. Six methods are in common use (Goodchild, 1992):

- randomly located sample points;
- sample points in a regular rectangular array;
- a regular rectangular array of cells, with some aggregate value of the field recorded for each cell;
- a polyline representation of selected isolines (contours);
- an irregular array of triangles, with the value of the field recorded at every vertex, and the field assumed to vary linearly within each triangle (the TIN, or triangulated irregular network model);
- an irregular array of polygons, with the value of the field assumed constant within each polygon.

Although many other representations are possible, such as arrays of hexagonal cells, they have not been exploited to any great extent in geographical data modeling to date.

Each of the six field models creates sets of points, polylines, or polygons, and thus at this superficial level the distinction between field and entity becomes moot within a digital representation. The term "vector" is used to refer ambiguously to all three types of entity models, and all but two of the field models, while the term "raster" refers to the second and third field models. Moreover, any set of non-overlapping entities can be represented as a field whose value is the entity identifier at locations inside an entity, and zero elsewhere.

However, the distinction between entities and fields is critical when the behavior or processing of points, polylines, or polygons is considered. The polygons of a field representation must collectively exhaust the space, and cannot overlap; one polygon cannot be dragged over another during editing, and common boundaries can be moved only as long as they do not intersect other common boundaries. The polylines of a contour representation cannot cross, and adjacent contours must have adjacent values. At a more sophisticated level, the common process of interpolation between a sample of points representing a field, exemplified by the contouring operation, makes no sense whatever if the points are conceived as entities littering an empty space.

Each of the six field models has its own characteristics, and some of these are important in soil mapping. First, only two of the field models are appropriate for discrete (categorical, multinomial) fields: the regular array of cells, termed the "raster" model for the purposes

of this paper, and the polygon model. In both cases the value assigned to each object in the representation, respectively cell and polygon, is assumed to be the most common or modal value of the field within the limits of the object.

Second, only three of the field models include a representation of the value of the field at every point in the plane: the regular array of cells, the polygon model, and the TIN. In the remaining cases (irregular points, regular array of points, and digitized contours) the model must be coupled with some set of rules about interpolation if values are to be determinable at any point in the plane. From an object oriented perspective, it would be desirable to encapsulate the method of interpolation with the data model, but in practice this is rarely done; instead, selection of an appropriate rule or rules is left to the user.

The next section of the paper discusses the traditions of GIS within this conceptual framework.

Geographic Information Systems. GIS uses a broad range of data models, and some currently available GIS support all of the entity and field models. In this paper the focus is on soil mapping, conceptualized as a discrete or multinomial field. In GIS, such data may have been obtained by remote sensing or aerial photography, or alternatively by digitizing or scanning an existing map. Thus the field's digital representation may use either the raster or polygon model. The choice may depend on the original source of data (raster is more likely if the data came from a classified photograph, or by scanning a map, and polygon is more likely if the data were created by digitizing a map); the processing software used (some GIS require the raster model for analysis and modeling, and some require the polygon model); the analysis and modeling to be performed; or the preferences of the user.

Thus several options for the representation of a soil map can be found in GIS practice. A raster model may have been obtained by classifying a photograph, or by scanning a map; a polygon model may have been obtained by digitizing a map, or by vectorizing a raster representation. In the raster case the error inherent in the model is attributable to the replacement of the true value at a point by the dominant value in one or a small array (e.g., 3 by 3) of pixels; in the polygon case, the model imposes an error through the replacement of the true value at a point by the dominant value within a polygon of arbitrary size and shape.

Thus both approaches involve information loss. In the raster case, this is imposed by the pixel size, which is likely a characteristic beyond the user's control, and by the perceived need to remove "salt and pepper". In the polygon case, the reasons behind the lumping of space into polygons are more complex, and are reviewed here at some length. In each case, the discussion includes an interpretation of the minimum mapping unit, the area of the smallest polygon in the model.

Cartographic: The traditions of map making have evolved under the constraints imposed by available technology. It is comparatively easy to draw a map of homogeneous areas separated by sharp boundaries - the boundaries can be drawn by pen, and the areas can be filled using a variety of forms of shading or by applying prepared textures. It is much more

difficult to portray fuzzy boundaries, or small inclusions within larger areas, or spatially continuous change of any kind. The minimum mapping unit may be the smallest area that can be conveniently drawn and labeled at the planned scale of the map. In addition, area boundaries may be smoothed to create a cleaner appearance.

Management: Maps of soils are often made for well-defined purposes, such as agriculture, and the areas shown on them may be managed as agricultural fields. While soils may change continuously, management practice must work within well-defined boundaries, and must be applied uniformly over each well-defined area. The minimum mapping unit may be the smallest area that can be conveniently managed.

Cognitive: Much mapping tradition is concerned not so much with the accurate measurement of conditions on the ground as with the communication of an impression of geographic variation; in this interpretation, the cartographer plays a distinct role. It was observed earlier that the entity models have more in common with everyday human spatial cognition than the field models; the polygon model, which blocks space into entities of uniform character, can be seen as a response to human patterns of thinking about geographic variation. The minimum mapping unit may reflect a human need to impose a certain level of uniformity on the landscape.

For some mix of the above reasons, it is common for soil data obtained in the field or by interpretation of aerial photographs to be aggregated into polygons with area greater than some minimum mapping unit. Note that when this strategy is used there is an implicit assumption that the spatial resolution is much less than the minimum mapping unit; if it is not, it is in principle impossible to define the class to which a polygon belongs. Aggregation produces polygons whose value becomes the commonest class observed within the polygon. Other classes which may be present as inclusions contribute to the error inherent in the model, unless they are part of the definition of the class.

Error models. For the reasons discussed in the preceding section, it is impossible to create an exact representation of the spatial distribution of a phenomenon as complex as soil within a digital database. At the same time, a digital representation is precise, and capable of giving unambiguous answers to simple queries. This discussion will focus on two particular generic queries, on the assumption that more complex queries can be reduced to them. The two queries of interest are:

- the value of the soil class field at some specified point (x,y); and
- the area of a specified soil class.

In order not to be misled by the precision of answers to these queries, it is desirable that the database respond with some appropriate measure of uncertainty in both cases. In the case of a point query, the response should include not only the class of the containing pixel or polygon, but also some measure of the variation known to exist within the spatial object, whatever its source. In the case of an area estimate, the response should include some expression of confidence limits around the value of the estimate.

Databases containing soil maps are being widely used at this time for purposes such as land use control, and environmental management. If uncertainty is known to exist in the responses to these two generic queries, it is crucial that that information be conveyed as succinctly as possible to the user or decision-maker. It would also be helpful if map displays could also convey some impression of uncertainty, as discussed later in the paper.

For these reasons, there has been increasing interest in accuracy issues in recent years, particularly in the research community. This research follows three distinct threads, which will be discussed here under the headings of classification, cartographic, and spatial statistical. Each offers its own answers to questions of how accuracy should be described and measured, and propagated into estimates of the uncertainty associated with responses to queries. In a much quoted paper, Anderson *et al.* (1976) proposed that 85% be an appropriate accuracy standard for land cover mapping. The following sections review the various meanings that have been ascribed to that statement, and associated methods of measurement.

Classification. In this context, accuracy is viewed as a problem of misclassification. A pixel or polygon is said to be misclassified if its true class, as determined by ground check or from a source of higher accuracy, is j , but its assignment by the database is i (in this discussion the terms true or truth should always be taken to imply ground check or a source of higher accuracy; the general term reference data is often used). Pixels or polygons assigned to a class other than their true class are termed errors of omission, or false negatives; pixels or polygons assigned to class i that are not truly of that class are termed errors of commission, or false positives. Note that any one misclassified object can be regarded as either an error of omission or an error of commission, depending on the perspective taken. Story and Congalton (1986) refer to omission errors as producer's accuracy because the producer's concern is presumably to avoid them; commission errors define user's accuracy because a user is interested in knowing how many objects that appear to be of class i are actually of class j .

Consider the idealized case discussed earlier, an image of an agricultural area, divided into fields growing different crops. Because of the inherent limitations of data gathering and spatial databases, it is inevitable that some misclassification or confusion of classes will occur, where pixels or polygons falling in fields growing crop j are mistakenly assigned to class i . If the data have been obtained by remote sensing, the likelihood that this will happen for a given pixel depends largely on the difference in spectral response of the two crops, and also on the technical efficiency of the classification method used. It is helpful to think of misclassification in the form of a matrix, in which the cell in row i , column j gives the number of pixels that are truly class j but have been assigned to class i . This matrix has also been called the error matrix or confusion matrix. The term error matrix will be used in this paper.

If the pixels that have been subject to accuracy assessment can be regarded as a random sample of the entire population of pixels, then the entries in row i can be divided by the row total to give an estimate of the probability that a pixel of class i is actually of class j , $p(j|i)$. Alternatively, the entries in column j can be divided by the column total to give an

estimate of the probability that a pixel that is actually in class j has been assigned to class i , $p(i|j)$. Here, as elsewhere in this approach, it is assumed that the probability of misclassification is constant for a given class over the entire image. In practice, some fields may be more likely to be misclassified than others, because of differences in the growing stage of the crop, and numerous other factors. Several studies (e.g., Campbell, 1981; Congalton, 1988a) have looked at the extent to which errors tend to cluster in space, indicating that the probabilities of errors are not constant.

In order to estimate the error matrix, it is common to conduct a random sample stratified by class. A random sample of pixels classified as j is selected and checked against the truth. The proportion that are found to be of true class i is then used to estimate the probability $p(i|j)$. Stratification is used because some classes tend to be more abundant than others; in order to obtain a reasonably sound estimate of the accuracy in each class, it is necessary to assign a greater number of samples to the rarer classes than would occur if sample points were located randomly. There have been several studies comparing the effectiveness of alternative sampling designs, particularly concerning the number of samples to be allocated to each class in the stratified approach, and the location of samples within areas of each class (e.g., Rosenfield, Fitzpatrick-Lins, and Ling, 1982; Hay, 1979; Fitzpatrick-Lins, 1981; Congalton, 1988b).

Much research effort has gone into devising suitable overall measures to summarize the contents of an error matrix. Let x_{ij} denote the number of cases recorded in row i , column j of the matrix, that is, the number of cases where the true class was j and the assigned class was i . The percent correctly classified compares the number of correctly classified pixels (those appearing on the diagonal of the matrix) to the total number of pixels. In order for this statistic to be meaningful, as an estimate of the probability that a randomly chosen pixel will have been assigned the correct class, it is necessary that the proportion of test pixels in each column (the ratio of the column total to the grand total) be the same as the proportion of total area that is truly of that class. Since this is generally unknown, the proportion of total area assigned to that class is used instead. Thus the row totals, rather than the column totals, are in the same proportions as the observed areas. When a stratified sampling scheme has been used, the totals must be adjusted by weighting each case by the inverse of that class's sampling density.

Unfortunately, and although it is the most obvious interpretation of Anderson *et al.*'s (1976) recommended 85% accuracy threshold, the percent correctly classified can be a misleading statistic, because a certain number of correctly classified cases are expected to occur by chance, even in the most confused classification. Thus it is often replaced by a statistic which allows for chance, and ranges from 0 in the case of the most confused classification to 1 in the case of the most accurate. Various known as Cohen's Kappa and KHAT, it is defined as follows:

where a subscript replaced by a dot indicates summation over that subscript, and N is the total sample size (Congalton, Oderwald, and Mead, 1983; Hudson and Ramm, 1987). Rosenfield and Fitzpatrick-Lins (1986) describe a variant on the Kappa statistic that can be calculated separately for each class; this is especially useful when it is desirable to know the

$$\hat{k} = \frac{\sum_i x_{ii} - \sum_i \frac{x_{i.} x_{.i}}{N}}{N - \sum_i \frac{x_{i.} x_{.i}}{N}}$$

classification accuracy of each class.

In some applications it may be useful to compare error matrices. For example, a study may compare the effectiveness of two image classifiers applied to the same scene; it is then important to know whether one classifier has outperformed the other. It may also be interesting to compare the effectiveness of a classification to a random assignment of classes, to see if the classifier has done better than chance. For many purposes, a simple comparison of Kappa values is sufficient. But since any error matrix or Kappa statistic is based on a limited sample, one might want to know whether the apparent difference between two Kappa statistics could have occurred by chance, implying that if other samples had been taken, the answer might have been different.

Congalton, Oderwald, and Mead (1983; see Hudson and Ramm, 1987, for corrected equations) describe methods for conducting inferential tests of Kappa, using the methods for testing discrete contingency tables described by Bishop, Fienberg, and Holland (1975). In their simplest form, such tests assume that the entries in the error matrices are strictly the number of times a condition was observed. Weighting, or other numerical manipulation of the table, will invalidate this assumption, and must be dealt with using special variants of the test.

The approach just described relies for its success on the validity of its underlying conceptual model, of areas of homogeneous class divided by sharp discontinuities. This model may apply well to agricultural fields, but in soil mapping it may be confused by continuous variation within more or less homogeneous areas, and slow transitions across boundaries. Both of these lead to a condition described statistically as non-stationarity, in which the error matrix captures only an average over the entire study area; locally, and for a variety of reasons, the probabilities of misclassification may vary markedly. To deal effectively with non-stationarity, it is necessary to partition the study area into regions of more or less constant error probabilities, and to sample them with as many samples as would have been used to characterize the entire study area under stationarity. Moreover, the basis of regionalization is unlikely to be known in advance, although Congalton, Oderwald, and Mead (1983) describe discrete multivariate methods that can potentially be used to identify the factors that cause spatial variation in rates of error.

In addition to the problems of misclassification of spectral response, some pixels will fall on the boundaries between homogeneous patches, and will therefore contain more than one class. Such mixed pixels are not misclassifications in the same sense, and it is common to avoid pixels that straddle boundaries in accuracy assessment and the development of error matrices. Recently, there has been much interest in improved image classification methods that deal explicitly with the mixed pixel problem.

Cartographic. To a cartographer, the problem of map accuracy assessment is perceived very differently. Whether its source was aerial photography or ground mapping, a map of soils shows a set of non-overlapping, space-exhausting areas with homogeneous characteristics, bounded by lines of constant width. In one of the final stages of the map-making process, the boundaries will likely have been smoothed, or splined, to create a pleasing, orderly appearance.

The accuracy of such maps is often seen in terms of two questions:

- are the boundaries in the correct locations; and
- have the areas been assigned to the correct classes?

Hord and Brooner (1976) use this approach in their analysis of land-use map accuracy. The new spatial data transfer standard, now known as Federal Information Processing Standard (FIPS) 173 (see Morrison, 1992, and other papers in that journal issue), refers to these as positional accuracy and attribute accuracy respectively. Positional accuracy can be measured by the Perkal epsilon band, a band of width epsilon centered on the observed location of the boundary, within which the true position of the boundary is assumed to lie (Blakemore, 1984), and Mark and Csillag (1989) have described an alternative probabilistic approach. Attribute accuracy can be described by a form of error matrix, but note that the matrix now refers to the classification of each area, rather than to each pixel. In this approach, the percent of areas correctly classified is compared to Anderson *et al.*'s (1976) threshold of 85%.

This form of accuracy assessment asks only whether each area or polygon has been assigned to the correct class, and thus ignores the almost inevitable heterogeneity that results from the approximation of continuously varying land characteristics. If an area is known to be heterogeneous, the "correct" class is normally assumed to be the commonest class, but no assessment is made of the occurrence of other classes within the polygon. This approach to accuracy assessment is often referred to as per polygon accuracy assessment, to distinguish it from the per pixel or per point approach of the classification school. Unfortunately, this means that an accuracy assessment based on the cartographic approach cannot be transformed into one based on the classification approach, or vice versa.

The cartographic approach has clear advantages if it is reasonable to assume that areas are homogeneous, and that boundaries are sharp and distinct (crisp). On the other hand the classification approach is clearly more appropriate if areas are heterogeneous, with substantial inclusions of unmapped classes, or if boundaries are indistinct (fuzzy). Chrisman (1989) shows how it is possible using the cartographic approach to distinguish between different sources of error, by observing whether they result in small shifts in boundaries, or in reclassification of entire areas. In reality, some boundaries on soil maps are more distinct than others, and some areas are more homogeneous than others; the distinction is rarely clearcut.

Spatial statistical. From a statistical perspective, error is variation that cannot be explained.

Most forms of measurement are subject to error, and statistics has developed methods for describing the amount of error present, and for analyzing data despite the presence of error.

In the measurement of a simple quantity like temperature, it is conventional to conceive of a true value, which has been distorted to an unknown degree by error. Repeated measurements using different observers, or different measuring instruments, would produce a range of values characterized by a mean and a standard deviation, and often with a histogram that follows the normal or Gaussian distribution closely. A range of well-developed methods allow such data to be analyzed.

The spatial equivalent of the Gaussian distribution is a range of maps, each showing the same general characteristics but with individual patterns of distortion, and each being equally likely to have been observed. The range might represent the work of different interpreters, or the use of different training sites for image classification. An error model is defined as a statistical process capable of generating such a range of maps given appropriate parameters (the Gaussian distribution is the most commonly used error model for simple measurements). Any one such map is termed a realization of the model. Simple measures of error, such as the percent correctly classified, are termed error descriptors, and may be related to the parameters of an error model. For example, the standard deviation, a commonly used measure of measurement uncertainty, is a parameter of the Gaussian distribution.

The use of an error model has numerous advantages. The effects of uncertainty can be visualized by using the error model to simulate a sample of possible maps. Errors can be propagated by performing analysis on several realizations of the error model, and computing the variation in results across them. The parameters of an error model can be calibrated by adjusting them so that the range of outcomes under the error model matches the range observed in reality.

One of the more straightforward applications of soil mapping is in the estimation of areas having particular characteristics. For example, one might want to know the total area of sandy loam in Mendocino County. Knowing that the database includes misclassifications and errors, it would be useful to know the standard error associated with the estimate of area. Unfortunately, neither the classification approach, with its error matrix and Kappa statistics, nor the cartographic approach with its epsilon bands and polygon error matrix, are capable of providing such an estimate. However, given an error model, an estimate of standard error can be made by generating a sample of realizations, calculating area on each, and computing the variation between them.

Several recent papers have described error models for maps, and their use in predicting the effects of uncertainty. Of most relevance to this paper are those that deal with maps of classified land, generally known as area class maps. Goodchild, Sun, and Yang (1992) propose a general error model, and show how it can be used to obtain estimates of uncertainty in such GIS products as area estimates. In their model, the uncertainty associated with the class at any point is represented by a probability vector in a raster representation. Each pixel ij is associated with a vector of probabilities $\{p_{ij1}, p_{ij2}, \dots, p_{ijn}\}$,

giving the probability that the pixel truly belongs to each class 1 through n. These probabilities can be interpreted as the consequence of mixed pixels, of uncertainty of class definition, or in remote sensing as the effect of confusion of spectral signatures - in this spatial statistical approach, the origins of error are not of immediate concern. The classification approach described earlier assumes that the contents of the vector are determined by the class to which the pixel appears to belong.

In this model, the classes allocated to the pixels in the database represent one realization of a stochastic process defined by these vectors of probabilities. For example, if there are two classes and the probabilities for a given pixel are $\{0.5, 0.5\}$, then one interpretation might be that in the maps made by two independent observers, one map would assign the pixel to class 1 and the other to class 2. Over a large number of realizations, the proportion of times the pixel is assigned to each class will converge on each class's probability.

In addition, Goodchild, Sun, and Yang (1992) propose that within any one realization, the outcomes in neighboring pixels be correlated, and the model also includes parameters describing the level of spatial dependence. For example, if a large area of many pixels is suspected of being sandy loam, but with a small probability of being clay, a low level of spatial dependence would imply that clay occurs in small inclusions; a high level of spatial dependence corresponds to large inclusions, or to the possibility that the entire patch is clay rather than sandy loam.

The relationship between this model and the error matrix discussed earlier is straightforward: the error matrix assumes that probability vectors are a function only of the class assigned to each pixel in an observed realization, and leaves spatial dependence undefined. The relationship to the cartographic approach requires more detailed explanation.

Consider a fuzzy boundary line between two soil classes A and B. In probabilistic terms, the fuzziness can be represented as a rate of change of probability as the boundary is traversed. Well to one side of the boundary, the probability that a pixel belongs to class A is 1, and the probability of class B is 0. As the boundary is approached, the probability of class A falls, at a rate depending on the degree of fuzziness. Each realization of the stochastic process will assign each pixel to one class, allowing the boundary to be inferred in a position that will vary from one realization to another within the general area of the fuzzy boundary. Crisp boundaries are represented by sharp changes in pixel probabilities as the boundary is crossed; every realization will place the inferred boundary in the same position.

A less general model that omits spatial autocorrelation between outcomes in neighboring cells has been described by Fisher (1991) for soil maps. A general model for raster data has been developed by Haining and Arbia (1993), and Fisher (1992) has used a model of spatially autocorrelated errors in digital elevation models.

In summary, while the classification and cartographic approaches are useful to monitor errors in image classification and map making respectively, they are mutually incompatible,

and only the spatial statistical approach offers sufficient generality to deal comprehensively with the problem of error propagation into the products of GIS analysis, and to provide a bridge between the distinct traditions of remote sensing and GIS.

Visualization. The first section of this paper discussed the issues of data modeling, as they affect the representation of fields in general, and soil maps in particular. GIS was seen as making use of two specific representations of multinomial fields: the raster and polygon models. Various forms of approximation and generalization are involved in each of these methods. In the second section, the concept of error modeling was introduced as an approach to allow representation of the uncertainty inherent in each data model, and its incorporation into the reported responses to two basic forms of query.

GIS is a visual technology, and derives much of its value from its ability to present a visual representation of geographical variation. In turn, the human eye and brain are extraordinarily effective processors of visual information, and able to detect patterns, particularly lineations, with great speed and efficiency.

In the classification approach described above, it is conventional to present spatial variation in the form of a display of each pixel's most likely class. Knowledge of error and uncertainty is captured in aspatial (or spatially stationary) form in the misclassification matrix.

The cartographic approach offers somewhat greater opportunity to communicate knowledge of uncertainty visually. If an area of a given class is known to contain a mixture of some other class, cartographers will sometimes cross-hatch the area with alternating stripes of colors, with widths appropriate to the proportions of each class in the mixture. The regularity of the stripes should be sufficient to discourage any notion in the mind of the map reader that they are real features of the landscape. Other techniques, such as blurring or widening, can be used to express uncertainty about boundary positions.

In general, however, few techniques are available for communicating knowledge of uncertainty. In 1991, the U.S. National Center for Geographic Information and Analysis instituted a research initiative in the Visualization of Spatial Data Quality (Beard, Battenfield, and Clapham, 1991) to try to respond to a perceived need for improved techniques.

Leung, Goodchild, and Lin (1992) identify three approaches that can be taken to visualization within the spatial statistical approach discussed earlier. First, uncertainty can be hidden by displaying only the most likely class at any point in the plane. Second, a series of displays can be used to show the probabilities of each class at each point. While this approach conveys some information on uncertainty, it fails to deal with spatial dependence, and Leung, Goodchild, and Lin (1992) show that it can confuse the reader by showing a pattern of spatial variation that cannot exist in reality (see also Englund, 1993). Finally, uncertainty can be communicated through the display of a sample of realizations of the error model. Each realization is a possible map, and the reader gains an impression of uncertainty through the variation observed between realizations.

Conclusions. Mapping of soils has a long and largely effective history. However, the issue of error and uncertainty in soil maps has led to two contrasting and potentially conflicting approaches, termed here the classification and cartographic. The two approaches are particularly distinct in the communication of information on map uncertainty to the user.

This paper has proposed a reconciliation of these two distinct traditions, through what has been termed the spatial statistical approach. A soil map is viewed as a multinomial field, a single realization of a multinomial error model. Uncertainty is described through the parameters of the model, and expresses itself in the variation between realizations of the model. The representation of the infinite information present in a field in the discrete space of a digital computer is one significant source of error, together with a variety of other forms introduced at various stages in the data collection and interpretation processes. Uncertainty can be visualized through the parameters of the model, but the most direct and unambiguous way to visualize it is by generating a series of realizations of the model.

Recent interest in conservation, and in environmental change, has led to a renewal of efforts to devise better methods of land surface characterization, of which soil mapping is a significant element. Such mapping will necessarily always fall short of being a precise exercise in scientific measurement, despite the availability of remote sensing. However, the unification of classification and cartographic approaches under a single framework may help to remove some of the problems that have plagued the field in the past.

Acknowledgment. The National Center for Geographic Information and Analysis is supported by the National Science Foundation, grant SES 88-10917.

References

- Anderson, J.R., E.E. Hardy, J.T. Roach, and R.E. Witmer, 1976, A land use and land cover classification system for use with remote sensor data. *Professional Paper 964*. Washington, DC: U.S. Geological Survey.
- Beard, M.K., B.P. Battenfield, and S.B. Clapham, 1991, Scientific report of the specialist meeting: NCGIA Research Initiative 7: Visualization of Spatial Data Quality. *Report 91-26*. Santa Barbara, CA: National Center for Geographic Information and Analysis.
- Bishop, Y.M.M., S.E. Fienberg, and P.W. Holland, 1975, *Discrete Multivariate Analysis: Theory and Practice*. Cambridge, MA: MIT Press.
- Blakemore, M., 1984, Generalization and error in spatial data bases. *Cartographica* 21: 131-139.
- Burrough, P.A., 1986, *Principles of Geographical Information Systems for Land Resources Assessment*. Oxford: Clarendon.

- Campbell, J.B., 1981, Spatial correlation effects upon accuracy of supervised classification of land cover. *Photogrammetric Engineering and Remote Sensing* 47(3): 355-363.
- Chrisman, N.R., 1989, Modeling error in overlaid categorical maps. In M.F. Goodchild and S. Gopal, editors, *Accuracy of Spatial Databases*. London: Taylor and Francis, pp. 21-34.
- Congalton, R.G., 1988a, Using spatial autocorrelation analysis to explore the errors in maps generated from remotely sensed data. *Photogrammetric Engineering and Remote Sensing* 54(5): 587-592.
- Congalton, R.G., 1988b, A comparison of sampling schemes used in generating error matrices for assessing the accuracy of maps generated from remotely sensed data. *Photogrammetric Engineering and Remote Sensing* 54(5): 593-600.
- Congalton, R.G., R.G. Oderwald, and R.A. Mead, 1983, Assessing Landsat classification accuracy using discrete multivariate analysis statistical techniques. *Photogrammetric Engineering and Remote Sensing* 49(12): 1671-1678.
- Englund, E., 1992, Spatial simulation: environmental applications. In M.F. Goodchild, B.O. Parks, and L.T. Steyaert, editors, *Environmental Modeling with GIS*. New York: Oxford University Press, pp. 432-437.
- Fisher, P.F., 1991, Modeling soil map-unit inclusions by Monte Carlo simulation. *International Journal of Geographical Information Systems* 5(2): 193-208.
- Fisher, P.F., 1992, First experiments in viewshed uncertainty - simulating fuzzy viewsheds. *Photogrammetric Engineering and Remote Sensing* 58(3): 345-352.
- Fitzpatrick-Lins, K., 1981, Comparison of sampling procedures and data analysis for a land-use and land-cover map. *Photogrammetric Engineering and Remote Sensing* 47(3): 343-351.
- Goodchild, M.F., 1992, Geographical data modeling. *Computers and Geosciences* 18(4): 401-408.
- Goodchild, M.F., G.Q. Sun, and S. Yang, 1992, Development and test of an error model for categorical data. *International Journal of Geographical Information Systems* 6(2): 87-104.
- Haining, R.P., and G. Arbia, 1993, Error propagation through map operations. *Technometrics* 35(3): 293-305.
- Hay, A.M., 1979, Sampling designs to test land-use map accuracy. *Photogrammetric Engineering and Remote Sensing* 45(4): 529-533.

- Hord, R.M., and W. Brooner, 1976, Land use map accuracy criteria. *Photogrammetric Engineering and Remote Sensing* 42(5): 671-677.
- Hudson, W.D., and C.W. Ramm, 1987, Correct formulation of the Kappa coefficient of agreement. *Photogrammetric Engineering and Remote Sensing* 53(4): 421-422.
- Leung, Y., M.F. Goodchild, and C.-C. Lin, 1992, Visualization of fuzzy scenes and probability fields. *Proceedings, Fifth International Symposium on Spatial Data Handling*, 2: 480-490.
- Maguire, D.J., M.F. Goodchild, and D.W. Rhind, 1991, *Geographical Information Systems: Principles and Applications*. London: Longman Scientific and Technical.
- Mark, D.M., and F. Csillag, 1989, The nature of boundaries on 'area-class' maps. *Cartographica* 26(1): 65-78.
- Morrison, J.L., 1992, Implementing the spatial data transfer standard - introduction. *Cartography and Geographic Information Systems* 19(5): 277.
- Rosenfield, G.H., and K. Fitzpatrick-Lins, 1986, A coefficient of agreement as a measure of thematic classification accuracy. *Photogrammetric Engineering and Remote Sensing* 52(2): 223-227.
- Rosenfield, G.H., K. Fitzpatrick-Lins, and H.S. Ling, 1982, Sampling for thematic map accuracy testing. *Photogrammetric Engineering and Remote Sensing* 48(1): 131-137.
- Star, J.L., and J.E. Estes, 1991, *Geographic Information Systems: An Introduction*. Englewood Cliffs, NJ: Prentice Hall.
- Story, M., and R.G. Congalton, 1986, Accuracy assessment: a user's perspective. *Photogrammetric Engineering and Remote Sensing* 52(3): 397-399.