Chapter 1
Introduction

1-1. General

a. This Engineer Technical Letter (ETL) addresses the use of geostatistics at hazardous, toxic, and radioactive waste (HTRW) sites. One very fundamental aspect of perhaps all HTRW site investigations that deal with environmental contamination is the need to characterize the extent and spatial distribution of contamination. Such a characterization typically would include describing, using a variety of statistical or analytical tools, spatial trends and variability. A principal difficulty in doing this is the fact that measurements may be few, or may be sparsely scattered over large regions. A question that arises naturally in this situation is how one might interpolate in order to make predictions (or estimates) at points where measurements of contaminant concentration are not available. Such interpolation will be referred to as point, or punctual, estimation in this ETL. Additionally, an investigator may need to determine a single representative value for an area that is represented by several measured or estimated values or both; this will be referred to in this ETL as block estimation. Geostatistics is a set of statistical procedures designed to accomplish these ends. Geostatistics may be applied to many problems, other than contamination, that occur at HTRW sites. Even though this document addresses only twodimensional applications, geostatistics can be used in three dimensions as well. Indeed, there are many cases in which the third dimension, usually stratification, is desirable to address.

b. Kriging is the principal geostatistical methodology described in this ETL. For introductory purposes kriging can be defined as a technique for determining the optimal weighting of measurements at sampled locations for obtaining predictions, or estimates, at unsampled locations; additional definition of kriging is provided throughout this document. Kriging is well-suited for making point and block estimates. However, much of the advantage of using geostatistical procedures, such as kriging, lies not just in the point and block estimates they provide, but in the information they provide concerning uncertainty associated with these estimates. The uncertainty information is usually quantified as either the standard deviation (or variance) associated with kriging estimates and is referred to as kriging standard deviation (or kriging variance) in this ETL.

c. Original geostatistical work involved making estimates for the areal extent and concentrations of economic mineral deposits, in relation to mining. Today (1996), geostatistical techniques continue to have a function in mining. However, a well-developed methodology that is capable of interpolating a given set of measured values at discrete locations into estimates for new locations or developing an individual estimate for an area including many locations, or both, has attracted users from many disciplines, and there is a trend toward incorporating geostatistics as standard curriculum for most geo-science educational programs. The use of geostatistical techniques as part of HTRW site investigations is becoming common because of the almost routine need for data interpolation as part of these investigations.

d. Once investigators have established that their data are adequate as to quality and quantity, geostatistics can provide powerful analytical tools that result in quantitative characterization of areas of special interest within the study area or the entire study area. These characterizations may address spatial variation; for example, it may be determined where values for concentrations of contaminants in soils are relatively high or low, are less than or greater than a specified value, or even have a high or low probability of exceeding a certain value.

1-2. Scope

a. The scope of this ETL will be limited principally to discussions and examples of two-dimensional point and block estimations using a geostatistical method known as kriging. The ETL will present the technical aspects of geostatistics through discussion of the assumptions behind and the mechanics of several types of kriging, including ordinary kriging, which is applicable when the mean for the variable of interest is constant over the region of interest, and universal kriging, which is applicable when the mean for the variable of interest changes gradually over the region. The discussion also will address a specialized form of kriging known as indicator kriging and the use of information concerning uncertainty associated with kriging estimates. The fundamental concepts of geostatistical kriging theory will be provided in this ETL; however, references will be provided for additional and more detailed information.

b. The practical aspects of kriging will be discussed through categorical examples of HTRW site investigations. The phrase "HTRW site investigations," will refer to planning, analysis, and remediation implementation phases of HTRW projects.

c. Additional topics included in this ETL such as review of applications and of some of the newer geostatistical techniques will be limited. The intent will be to familiarize the reader with these topics and not to provide how-to knowledge.

1-3. Organization

a. This ETL is organized into seven chapters. Chapter 1 is introductory and includes an overview of the technical aspects of spatial prediction in general and certain geostatistical concepts. Chapter 2 provides a detailed discussion of assumptions and theory behind kriging, including equations and concepts that will be useful to investigators who wish to gain a better understanding of the technical aspects, or mathematics, of kriging interpolation. As indicated, many of the concepts developed in Chapter 2 are discussed in very general terms in Chapter 1, so those readers desiring only an overview of kriging concepts may wish to read only Chapter 1 and bypass Chapter 2 altogether.

b. Chapter 3 provides a review of texts that contain much more detailed information regarding kriging theory than material included in Chapter 2. Chapter 3 also provides a brief generic discussion of kriging software.

c. Chapter 4 provides a detailed step-by-step discussion of variogram construction and demonstrates some pitfalls and solutions to this crucial process. Chapter 4 also discusses methodologies which investigators may use to evaluate their variograms.

d. Chapter 5 provides a discussion of practical aspects of geostatistics in a presentation of several example kriging applications with data from the HTRW field. The examples are intended to illustrate a few of the many different ways kriging can be used in HTRW site investigations and are not presented with the same level of detail used in Chapter 4.

e. Chapter 6 provides additional detail on some crucial aspects of kriging applications and includes considerations investigators may use to help determine if kriging is feasible for the application they have in mind, or reviewers can determine if the application of geostatistics was appropriate.

f. Chapter 7 provides an introduction to other methods for spatial modeling. This section also includes discussion of advanced stochastic methods such as simulation.

1-4. An Overview of the Use of Geostatistics in Hazardous, Toxic, and Radioactive Waste Site Investigations

a. *General*.

(1) HTRW site investigations involve complex administrative, scientific, and engineering functions and are truly interdisciplinary. Scientists and engineers, for instance, may be confronted with administrative findings or directives, associated with fiscal, managerial, or regulatory input, that may either guide or constrain their work. In a

likewise fashion, scientific findings may define the scope of administrative effort.

(2) Scientists and engineers involved in HTRW site investigations have found that they have an implicit need for many disciplines to fulfill the objectives of each particular investigation. Frequently, an HTRW site investigation will benefit from input from earth-science disciplines such as geology, hydrogeology, and chemistry, among others. Some HTRW site investigations are large enough to use several individuals from each of these disciplines, as well as many others, for the duration of multi-year investigations. Most disciplines associated with HTRW site investigations will benefit from knowledge or input from specialized and/or interdisciplinary branches; the geologist, for example, will occasionally benefit from knowledge of geophysics. Naturally, interdisciplinary input also can be very helpful, especially in geostatistics, where earth-science disciplines rely on assistance from statisticians.

(3) In this ETL and for its purposes, a complete HTRW site investigation is described concerning three relatively broad sequential activities or phases. These phases are referred to as initial planning, analysis, and implementation of remediation plans. Another very important HTRW site investigation activity, monitoring, is less discrete and is a part of all three phases. Monitoring represents the basis for analysis, is often modified as a result of analysis, and may be newly implemented as part of remediation.

(4) Kriging techniques can and have been used in any of the three phases. Only a few very basic applications of kriging techniques are described in this ETL. The intent of this ETL is to describe basic concepts so that more elaborate applications can be done based on a fundamental understanding of the procedures involved.

(5) For examples of more elaborate applications, the reader can refer to the material cited in Chapter 3. However, the best applications are developed by readers who have a clear understanding of the goals associated with each particular HTRW site investigation and also have a good basic understanding of the fundamental geostatistical techniques. As alluded to here and elsewhere in this ETL, there are many techniques available for gridding data; kriging has an added advantage of generating kriging standard deviations that can be used as a measure of uncertainty.

b. Initial planning.

(1) Initial planning may involve several aspects associated with implementing or operating a monitoring network; it also may involve reconnaissance evaluation of an existing network. Additionally, because monitoring is present in all phases of HTRW site investigations, the same opportunities for geostatistical applications associated with network analysis that occur in the initial planning stages may occur, perhaps often, throughout the investigation. The information available from kriging standard deviations can add much to sampling or monitoring network analysis.

(2) For application of geostatistical techniques, the most likely aspects of network implementation and operation to be addressed certainly include network design, evaluation, and modification. Geostatistics offer the investigator opportunities to:

(a) Locate areas where existing sampling or monitoring networks may provide strong or weak estimates.

(b) Quantify the effect of increasing or decreasing the sampling or monitoring network density.

(c) Evaluate the effect of removing or relocating certain monitoring locations or adding new locations to the sampling or monitoring network.

c. Analysis.

(1) Although aspects of network design can be quite important during analysis, the investigator is likely to be concerned principally with using information from monitoring networks to evaluate

environmental conditions throughout the specified study area. The evaluations may require either point or block estimates. Often, design factors are addressed in the analytical phase as well.

(2) A common application for kriging techniques in HTRW site investigations is estimating real means. More common, however, is estimating the extent of areal contamination. Usually these estimates involve chemicals in air, water, and soil; however, if sufficient information is available, such estimates could include a wide range of environmental factors that involve many issues other than contaminants. Perhaps the most common examples concern geologic and hydrologic factors, such as depth to bedrock and groundwater-level elevations. Investigators need to realize that almost any set of measurements can be distributed using kriging techniques, providing there is a sufficient amount and distribution of measured information.

(3) The investigator also needs to realize that the resultant kriging estimates can be gridded. This gridding affords investigators opportunities to perform mathematical or logical operations, or both, on the kriging estimates, provided that investigators are comfortable with kriging estimates. Saturated thickness could, for example, be calculated from kriging estimates for groundwater elevations and base of aquifer elevations.

(4) Often, after preparing estimates for areal properties, the investigator may appreciate the opportunity afforded by kriging techniques to evaluate the confidence associated with the estimates. Maps of kriging standard deviations can provide the investigator with information concerning the confidence associated with the kriging estimates. Although the areas of lowest confidence may be well-known intuitively, maps of the kriging standard deviation are an important step toward quantification. More often than not, even the most experienced investigator will benefit from careful study of maps of kriging standard deviations.

d. Implementation of remediation.

(1) One of the most common applications for kriging techniques in the final phases of HTRW site investigations is evaluating compliance. For instance, a question such as "Is the mean concentration of constituent x within compliance limits?" is ubiquitous to HTRW site investigations. Making determinations concerning compliance is very similar to estimating areal extent as part of the analysis. Investigators and managers have much to gain from the confidence information available from kriging techniques as to the reliability of estimates as well as in optimizing monitoring networks.

(2) Kriging can also be very useful if managers are interested in making decisions based on the probability of certain conditions existing. If a condition can be defined by the manager, then, providing there are adequate data, indicator kriging can provide an estimate for the probability of existence. A common example of this kind of application is making areal determinations for probabilities that concentrations for a constituent do or do not exceed, for example, an action level.

(3) There are many operational remediation issues that kriging techniques may address as well. Remedial activities at HTRW sites often need estimates for amounts in general. For instance, there could be a need for information regarding volumes of contaminants to be treated, volumes of soil to be excavated, volumes of soil to be stored, and so on. By combining estimates for different geologic, hydrologic, and chemical factors, estimates for these volumes can be obtained from kriging techniques in much the same way as saturated thicknesses can be calculated.

1-5. An Overview of Some Technical Aspects of Geostatistics

The purpose of this section is to provide an overview of some of the procedures and concepts to be treated in detail in this ETL. Some of the technical ideas and terminology will be introduced in very general terms, with the goal of orienting the reader who may not be familiar with the area of geostatistics.

a. General considerations in spatial prediction.

(1) The principal technical issue considered in this ETL is spatial prediction or modeling values of a spatial process; in particular it is considered how best to make use of measurements of a variable (such as pollutant concentration) at sampled locations to make inferences (or predictions) about that variable at unsampled locations or about values of the variable for the region as a whole.

(2) A spatial process can be viewed as having a large-scale or regional component and a smaller scale or local component; both of these components need to be accounted for when modeling a spatial process. The large-scale component is referred to as the mean field and is most often modeled by a spatial trend which may or may not be constant over the region. The smaller scale component is a random fluctuation which is mathematically combined with the trend to make up the sample at a point. The random component is usually assumed to be zero on the average but can be either positive or negative in individual samples. The separation of the trend from the random components is problem- and scale-dependent and requires some judgment to determine. There can be several “solutions” to the problem of separating the trend and random components that may be useful for various geostatistical purposes when using a single set of data.

(3) Local-scale fluctuation of the variable of interest (e.g. water levels or contaminant concentrations) at a sample point, although random, can show some association (i.e. correlation) with the random fluctuations at nearby points. This is referred to as spatial correlation. Positive spatial correlation between measurements means that the random components at both points tend to have the same sign, whereas negative correlation means the random components tend to have opposite signs. Both the “large-scale” trend and the positive spatial correlation of the “local-scale” fluctuations contribute to measurements taken at locations close together being more closely related than measurements taken farther apart.

(4) The most obvious way one might proceed for spatial prediction at unsampled locations is simply to take an average of the sample values that one does have and assume that this value gives a reasonable prediction at all locations in the region of interest. This may work adequately in some cases, but one can also see the pitfalls in doing this. Using a single value for an entire region makes an implicit assumption of spatial homogeneity. It ignores any spatial trends that might exist in the data and it also ignores spatial continuity. If it is known that the variable of interest does have the tendency to be spatially correlated, then it would make sense to use a weighted average rather than a simple average in making a spatial prediction, with measurements at sampled locations that are nearer to the unsampled location being given more weight. This then is the motivation for the geostatistical methods discussed in this ETL. The method known as kriging, which is the principal subject to be considered here, is a technique for determining in an optimal manner the weighting of measurements at sampled locations for obtaining predictions at unsampled locations. These optimal weights depend on spatial trends and correlations that may be present.

(5) There are a number of ways to go about performing spatial prediction. The geostatistical method of kriging covered in this ETL belongs to a class of methods known as stochastic methods. In these methods, it is assumed that the measurements, both actual and potential, constitute a single realization of a random (or stochastic) process. One advantage of assuming the existence of such a random process is that measures of uncertainty, such as the variance used in kriging, can be defined. These measures of uncertainty permit objective assessment of the performance of a spatial prediction technique on the basis of how small such measures are. Once a measure of uncertainty has been selected, the weights to be used in spatial prediction may be determined so as to explicitly minimize the measure of uncertainty.

In short, the use of stochastic techniques provides the investigator with a way of objectively quantifying errors and determining weights. In practice, spatial predictions obtained using kriging are almost always accompanied by a measure of the associated error. Most kriging practitioners consider such an error evaluation to be an integral part of the analysis, and point to error analysis as one of the principal advantages of using kriging (or stochastic techniques in general) over other procedures.

(6) Nonstochastic techniques, on the other hand, are typically applied strictly empirically, with no assumptions concerning the existence of an underlying random process and with no theoretical framework with which to evaluate statistically the performance or optimality of the techniques. When they are applied in such a manner, it is not possible to evaluate in advance whether such a procedure would be expected to yield results that are satisfactory. Two techniques that are commonly applied in a nonstochastic setting are simple averaging, mentioned above, and trend analysis, which is a least- squares method for fitting a smooth surface to the data. Even though these techniques are usually applied nonstochastically, it is still possible to assess their performance if a stochastic setting is assumed. Loosely speaking (these ideas are discussed more precisely in Chapter 7), simple averaging would perform well if there is no trend and no spatial correlation, and trend analysis would perform well if there is a trend that can be modeled, but no spatial correlation. Lack of correlation in the observations is one assumption that is made in ordinary statistical regression analysis, and in fact trend analysis, if it is placed in a stochastic setting, is actually one special type of regression. The stochastic method of kriging explicitly incorporates the spatial correlations which are ignored in trend analysis. In Chapter 7, a few other common techniques that are usually applied in a nonstochastic setting will be discussed briefly. Most of these techniques are designed to incorporate the notion of spatial continuity, but the way it is incorporated may be subjective. Kriging provides an objective means of incorporating the presence of spatial correlation and makes explicit the background assumptions that are being made.

b. Important geostatistical concepts. Below are some of the key ideas in geostatistics that will be given detailed attention in this ETL. They are introduced in much the same order that they are discussed in Chapter 2, where more detail is presented.

(1) Variograms.

(a) A central idea in geostatistics is the use of spatial correlation to improve spatial predictions, or interpolations. The variogram is the principal tool used to characterize the degree of spatial correlation present in the data and is fundamental to kriging. The correlation between measurements at two points is usually assumed, as described above, to depend on the separation between the two points. Values for all possible pairings of sample points can be examined by squaring the difference between the values in each pair. The squared differences are then categorized according to the distance separating the pair. For small separations, or lags, the squared differences are usually small and increase as the lag increases. A plot of the squared differences per sample pair as a function of lag is referred to as the sample variogram.

(b) The general behavior of the sample variogram points relates to the spatial correlation between sample sites and can provide investigators with qualitative information about the spatial process, but in order to use this information in a mathematically rigorous manner as a basis for interpolation, a function with specific properties must be fit to the sample variogram points. The fit, as with all curve-fitting procedures, takes the scattered points and passes a smooth curve through the points. The curve, which can be represented by a mathematical expression or function, is called a model. Several named models with characteristic features introduced in Chapter 2 are commonly used in geostatistics. The resultant variogram model is used to determine kriging weights for use in interpolation.

(2) Directional variogram and anisotropy. It is often the case that spatial correlation depends not only on distance between points, but also on direction. For example, measurements at pairs of points 100 m apart with the line between them oriented in a north-south direction may have a different correlation than measurements at points the same distance apart but with the line joining them oriented in an east-west direction. The spatial process is said to exhibit anisotropy, and what is known as a directional variogram must be used for the geostatistical analysis.

(3) Kriging and kriging variance.

(a) Kriging yields optimal spatial estimates at points where no measurements exist in terms of the values at points where one does have data. As discussed above, placing the problem in a stochastic framework permits precision-defining optimality. In kriging, the restriction is first imposed that the predicted value at any point is a linear combination of the measured values; that is, the kriging estimate is a linear predictor. Given this restriction, the values of the coefficients in this linear function are chosen so as to force the predictor to be optimal.

(b) The first criterion imposed is that the estimate be unbiased, or that in an average sense the difference between the predicted value and actual value is zero. The second optimality criterion is that the prediction variance be minimized. This variance is a statistical error measure defined to be the average squared difference between predicted and actual values. Because the kriging estimate minimizes this variance, it is known as the best (minimum variance) unbiased linear predictor. This minimization is performed algebraically and results in a set of equations known as the kriging equations, which give an explicit representation of the optimal coefficients (weights) in terms of the variogram. The form of these equations is presented in Chapter 2.

(c) Also given in Chapter 2 is an expression for the kriging variance. This variance depends on geometry of the data sites, with the variance at locations near points with measurements tending to be smaller. One can then associate with any spatial prediction a variance, which gives an indication of the uncertainty in that predicted value. As mentioned before, this measure of uncertainty gives kriging one of its principal advantages over many other techniques.

(4) Trends and universal kriging. Special attention must be given in kriging to the question of whether there are spatial trends in the data. A trend in this case is usually any detectable tendency for the measurements to change as a function of the coordinate variables but can also be a function of other explanatory variables. For example, aside from random fluctuations, measurements of groundwater elevations may exhibit a tendency to increase in a consistent manner the farther one proceeds in a certain direction. A kriging analysis in which there is no spatial trend is known as ordinary kriging; when a trend does exist, universal kriging should be considered. In universal kriging, one attempts to account for the trends present. For example, it might be assumed that the trend can be represented as a linear function of coordinate variables. The form of the trend model is then incorporated into the universal kriging equations to obtain the optimal weights.

(5) Block kriging. What has been discussed in the preceding paragraphs is usually known as point, or punctual, kriging. In point kriging, the goal is to predict the value of a variable at discrete locations. By contrast, in block kriging the goal is to predict the average value, over a specified region, of a variable. As in point kriging, the optimal predictor is a linear combination of the measured data values, and degree of uncertainty is indicated by a block kriging variance. Block kriging variances tend to be smaller than point kriging variances because averages tend to be less variable than individual values.

(6) Prediction intervals and normality.

(a) A standard kriging analysis will give two values for any location: the optimal kriging estimate and the kriging variance. The variance provides a measure of uncertainty for the prediction. In some cases, it may be desirable to go even further in specifying the nature of the uncertainty than simply giving the variance. One way to proceed is to try to obtain what is known as a prediction interval. Here one seeks an interval such that there is a certain probability, typically 95 percent, that the actual value lies in this interval.

(b) Finding such an interval often hinges on having knowledge of the probability distribution of the variables being sampled. One ideal situation is when the variable of interest, e.g., contaminant concentration, can be assumed to have a normal distribution. In this case, given the set of measured values, a potential value at an unsampled location has a normal distribution with mean given by the kriging estimate and variance given by the kriging variance. It is thus, using classical statistics, straightforward to use this normal distribution to obtain a 95 percent prediction interval for concentration at the unsampled location.

(7) Transformations. Having a prediction interval will generally be much more informative than simply having the kriging estimate and kriging variance, which explains why investigators often ask whether normality assumptions can be made for their data. When a normality assumption cannot be made, it is sometimes possible to find a transformation that will make the data normal, or nearly so. For example, a transformation that is often tried is the logarithmic transformation. That is, one simply takes the logarithm of all data values (assuming they are > 0) and performs the geostatistical analysis on these transformed values rather than on the original data. Prediction intervals obtained using transformed values can be readily converted to corresponding intervals on untransformed variables. There are, however, subtleties that must be considered in back-transforming the kriging estimate and the kriging variance; these are discussed in more detail in Chapter 2.

(8) Indicator kriging.

(a) In indicator kriging, analysis is performed using what are known as indicator variables rather than the measured data themselves. An indicator variable is thus a special kind of transform of the measured data and can have only two possible values: 0 or 1. To obtain the indicator variables to be analyzed, first specify a threshold value, say *c*, which may represent, for example, a contaminant concentration level which is of particular importance. At each measurement location, the indicator variable is then assigned a value of 1 if the measured value is less than or equal to *c*, and is assigned a value of 0 if the measured value is greater than *c*. This kind of transform will allow censored data, or data reported as less than some reporting limit, to be included in the analysis if the reporting limit is less than or equal to the cutoff value of *c*. After the indicator transform has been performed, the kriging analysis is performed using these indicator variables in the same manner discussed above; first a variogram is obtained, and the kriging equations yield the optimal linear predictor and the kriging variance for the indicators.

(b) Whereas the indicator kriging analysis is done using only 0's and 1's, the interpolated estimates are not restricted to these two values. In most cases the estimates are between 0 and 1, which is interpreted to be the probability that the actual value is less than or equal to the threshold *c*. Performing this analysis for a number of different threshold values, *c*, can give the investigator information about the probability distribution of contaminant values at a location, which may in turn be used to obtain prediction intervals. As discussed above, such intervals may even be more valuable than having only the optimal predictor and variance provided by the usual kriging analysis, particularly if behavior of extremes may be of interest to the investigator. The advantage of using indicator kriging to obtain prediction intervals is that it is not necessary to assume a distribution for the data, as in the discussion of normality above.