WATER RESOURCES RESEARCH, VOL. 36, NO. 3, PAGES 663–677, MARCH 2000

Mapping soil landscape as spatial continua: The neural network approach

A.-Xing Zhu
Department of Geography, University of Wisconsin, Madison

Abstract. A neural network approach was developed to populate a soil similarity model that was designed to represent soil landscape as spatial continua for hydroecological modeling at watersheds of mesoscale size. The approach employs multilayer feed forward neural networks. The input to the network was data on a set of soil formative environmental factors; the output from the network was a set of similarity values to a set of prescribed soil classes. The network was trained using a conjugate gradient algorithm in combination with a simulated annealing technique to learn the relationships between a set of prescribed soils and their environmental factors. Once trained, the network was used to compute for every location in an area the similarity values of the soil to the set of prescribed soil classes. The similarity values were then used to produce detailed soil spatial information. The approach also included a Geographic Information System procedure for selecting representative training and testing samples and a process of determining the network internal structure. The approach was applied to soil mapping in a watershed, the Lubrecht Experimental Forest, in western Montana. The case study showed that the soil spatial information derived using the neural network approach reveals much greater spatial detail and has a higher quality than that derived from the conventional soil map. Implications of this detailed soil spatial information for hydroecological modeling at the watershed scale are also discussed.

1. Introduction

Information on spatial variation of soil hydraulic and other properties is required for modeling the spatial variability of hydroecological processes over a watershed [*Beven and Kirkby*, 1979; *Burrough*, 1996; *Corwin et al.*, 1997; *Jury*, 1985]. Currently, soil maps produced through conventional surveys are the major source of soil spatial information for many modeling and management activities. Yet soil spatial information derived from these conventional soil maps has been found to be inadequate for hydroecological modeling at watersheds of mesoscale size because of its incompatibility with other landscape data derived from detailed digital terrain analyses and remote sensing techniques [*Band and Moore*, 1995; *Zhu*, 1997, 1999a]. This inadequacy is largely due to the limitations of the discrete data model and polygon-based mapping practice employed in conventional soil map production (see section 2.1).

To overcome these limitations, *Zhu* [1997] developed a soil similarity model to represent soil landscapes as continua. *Zhu et al.* [1997] found that the soil spatial information represented under this similarity model had higher quality both in terms of spatial details and attribute accuracy. Although the model provides a greater flexibility in representing soil landscapes than conventional soil maps, the challenge is how this similarity model can be populated. *Zhu and Band* [1994] and *Zhu et al.* [1996] employed a knowledge-based approach for determining the soil similarity values. The knowledge-based approach is only suitable for areas where there are experienced local soil scientists from whom the needed knowledge on soil-environmental relationships can be obtained [*Zhu*, 1999b]. Experienced local soil scientists may exist for areas which have undergone initial surveys or subsequent updates. However, some of these experienced soil scientists may have retired, and replacements may still be in the training stage. For areas where soil surveys have never been conducted, there are often no soil scientists very familiar with the local soils. Under these circumstances, what is available are the soil maps produced in previous surveys or field samples collected while conducting the initial survey. This paper presents a neural network approach for determining the soil similarity values under these circumstances. The paper also examines the potential impacts of detailed soil spatial information on landscape characterization for hydroecological modeling.

The section 2 provides a background on the limitations of conventional soil maps and an overview of the soil similarity model. The conceptualization of using neural networks to populate the soil similarity model is presented in section 3. Section 4 describes a case study illustrating the implementation of the neural network approach for soil mapping over the Lubrecht watershed in western Montana. The results of the case study and the implications of the detailed soil spatial information for modeling are discussed in section 5 and section 6, respectively. Summaries are presented in section 7.

2. Background

2.1. Limitations of Conventional Soil Survey Maps

Conventional soil maps were mostly produced through aerial photograph interpretation. During a soil survey, soil scientists first investigate the soils in an area to be mapped and establish a set of soil classes. They also gain an understanding of the relationships between soils and their environment through their field activities. The spatial extents of the soil

Paper number 1999WR900315.
0043-1397/00/1999WR900315$09.00

classes are then delineated as polygons on aerial photos through the interpretation of aerial photography [*Hudson*, 1992; *Zhu*, 1997]. The compilation of these polygons (soil polygons) forms the basis of a conventional soil map.

The level of spatial detail represented in a conventional soil map is largely limited by the spatial discrete model and polygon-based mapping practice employed. Under the discrete spatial data model [*Bregt*, 1992; *Goodchild*, 1992; *Mark and Csillag*, 1989], soil spatial distribution is represented through the delineation of soil polygons, with each polygon depicting the spatial extent of a particular soil type (class) or a group of commonly found classes (mixed-class mapping units). The polygons represent only the distribution of a set of prescribed soil classes (ideal concepts of soils), not individual soils in the field that often differ from the prototypes of the prescribed soil classes. In order to be mapped, individual soils in the field must be classified based on some classification scheme. Each individual soil is often assigned to one and only one class (referred to as Boolean assignment). Once assigned to a class, the local soil is said to have the typical properties of that class. As a result, the varying soil continuum is represented by a few distinct and discrete soil classes. The domain of a soil property is thus approximated by some typical values of soil classes, which are often discrete. This generalization of the entire domain of a soil property into a few typical values is referred to as "generalization of soils in the parameter domain" [*Zhu*, 1997, 1999a]. With the polygon-based mapping practice this generalization creates discrete soil polygon boundaries and forces soil spatial variation to be represented as a step function, which means that within a soil polygon, soils are assumed to be the same and differences between soils only occur at the boundaries of soil polygons. This step function of soil spatial variation can bias the estimation of spatial covariation between soils and other landscape parameters. In addition, the discretization of the soil parameter domain creates attribute incompatibility between soil spatial information from soil maps and data from digital terrain analysis and remote sensing techniques [*Zhu*, 1997]. Attribute incompatibility can also introduce bias into the estimation of spatial covariation between soils and other landscape parameters for distributed hydroecological modeling at the watershed scale [*Zhu*, 1999a].

A second problem with polygon-based mapping practice is that it limits the minimum mapping size of "soil bodies" (scale dependent). "Soil bodies" smaller than the established minimum size are either ignored or merged into the larger enclosing soil bodies. The existence of these smaller "soil bodies" (components) may or may not be reported in the map legend. In any case, the spatial extent of these components is not shown in the map. This filtering of small soil bodies due to the limitation of the polygon-based mapping technique is referred to as "generalization of soils in the spatial domain." This spatial generalization can be very significant, and the soil bodies that are filtered out can range from a few hectares on some large-scale (small area) maps to hundreds of hectares on some small-scale (large area) maps. The generalization of soil in the spatial domain can create spatial incompatibility between conventional soil maps and data derived from digital terrain analysis and remote sensing techniques. Spatial incompatibility refers to the difference in spatial resolution between two data sets. Spatial resolution here refers to the level of spatial detail at which the spatial variation is represented. Simply converting a soil map into a raster layer and reducing the pixel size do not increase the spatial resolution of the data layer at all since the process adds no details to the spatial variation of the soils being portrayed. Therefore the spatial resolution of a soil map is the minimum mapping size, which can range from a few hectares to hundreds of hectares. In contrast, data derived from digital terrain analysis and remote sensing techniques often have a spatial resolution of 30 m by 30 m or even finer. Spatially small but potentially important environmental niches can be described using these data sets. Soil information on these niches is often not available from conventional soil maps but could be critical to many detailed hydroecological modeling activities, such as studies on preferential flow paths through which water and chemicals move much more quickly than in the soil as a whole [*Jury et al.*, 1991].

2.2. Overview of Soil Similarity Model

The two generalizations imposed by the discrete spatial data model and the polygon-based mapping technique are the main causes of the inadequacy of conventional soil maps for detailed hydroecological modeling. To reduce the generalization of soils in the spatial domain, the soil similarity model employs a raster Geographic Information System (GIS) representation concept. Under the raster GIS representation an area can be represented by many small squares (pixels). The pixel size can be very small; it is often 30 m on each side, although much finer pixel sizes are possible. The spatial resolution of soil information is then the size of a pixel, not the minimum mapping size defined at a given map scale. With the raster representation, spatial generalization in producing soil spatial information can be largely reduced, and spatial details of soil variation can be represented at a resolution compatible with the detailed terrain and remotely sensed data. This resolves the spatial incompatibility between soil spatial information and other detailed environmental data.

To overcome the generalization of soils in the parameter domain, a similarity representation based on fuzzy logic is developed [*Zhu*, 1997]. With fuzzy logic the soil pedon at a given pixel can be assigned to more than one soil class with varying degrees of class assignment [*Burrough et al.*, 1992, 1997; *McBratney and De Gruijter*, 1992; *McBratney and Odeh*, 1997; *Odeh et al.*, 1992; *Zhu*, 1997]. These degrees of class assignment are referred to as fuzzy memberships. Each can be regarded as a similarity measure between the local soil and the typical case of the given class. All of these fuzzy memberships are retained in this similarity representation, which forms an n-element vector (soil similarity vector or fuzzy membership vector), S_{ij} $(S_{ij}^1, S_{ij}^2, \cdots, S_{ij}^k, \cdots, S_{ij}^n)$, where n is the number of prescribed classes and the kth element S_{ij}^k in the vector represents the similarity value between the soil at pixel (i, j) and soil class k [*Zhu*, 1997]. With this similarity representation a local entity such as a soil pedon at a given pixel is no longer necessarily approximated by the typical case of a particular class but can be represented as an intergrade to the set of prescribed classes. This mode of representation, which allows the local soil to take property values intermediate to the typical values of the prescribed classes, overcomes the attribute incompatibility between the soil spatial information represented under the similarity model and the other detailed environmental data. By coupling this similarity representation with a raster GIS data model, soils in an area are represented as an array of pixels with soil at each pixel being represented as a soil similarity vector (Figure 1). In this way, soil spatial variation can be represented as a continuum in both the spatial and parameter domains.

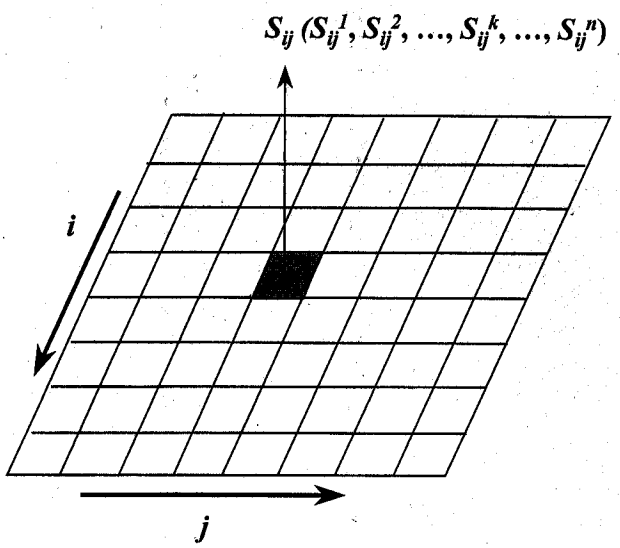

Figure 1. The soil similarity model. Soils are represented as pixels in spatial domain and as soil similarity vectors in parameter domain.

3. Populating the Similarity Model: The Neural Network Approach

The similarity model provides only the flexibility of representing soil spatial variation. The degree of success in using this model depends on how the model is populated or how the soil similarity values in the vector at each pixel are determined. This section presents a basic background on artificial neural networks (ANNs) and the conceptualization of using neural networks to populate the similarity model.

3.1. Basics of ANNs

Artificial neural networks are a form of computing motivated by the functioning of biological neural networks. An ANN solves a problem by first developing a memory associating a large number of input patterns with a set of resulting outputs through training on examples and then by applying this association to produce an output when given an input pattern. A detailed discussion of ANNs and their applications is beyond the scope of this paper. Readers are referred to the following review articles and books for discussion on the structure, functioning, and operation of ANNs [*Aleksander and Morton*, 1990; *Caudill and Butler*, 1990; *Dayhoff*, 1990; *Haykin*, 1994; *Hecht-Nielsen*, 1990; *Hertz et al.*, 1991; *Lippman*, 1987; *Maren et al.*, 1990; *Masters*, 1993; *McClelland et al.*, 1986; *NeuralWare, Inc.*, 1991; *Rumelhart et al.*, 1986; *Simpson*, 1990; *Vemuri*, 1988]. A number of studies on the application of ANNs in water resource research areas have been reported [e.g., *Clair and Ehrman*, 1998; *French et al.*, 1992; *Hsu et al.*, 1995; *Maier and Dandy*, 1996; *Morshed and Kaluarachchi*, 1998; *Pachepsky et al.*, 1996; *Ranjithan et al.*, 1993; *Rizzo and Dougherty*, 1994; *Rogers and Dowla*, 1994; *Schaap and Bouten*, 1996; *Schaap et al.*, 1998; *Shamseldin et al.*, 1997; *Tamari et al.*, 1996; *Wen and Lee*, 1998]. These references provide further information on the application of ANNs. Sections 3.1.1 and 3.1.2 provide a background on multilayer feed forward networks (MFN), the type of ANN used in this study, and a brief discussion on the training algorithms employed.

3.1.1. Structure. A MFN is made of many processing elements (neurons). These neurons are usually arranged in layers: an input layer, an output layer, and one or more layers in between called hidden layers (Figure 2). The neurons in one layer are connected to the neurons in the next layer with different strengths of connection. The strength of connection is referred to as a weight.

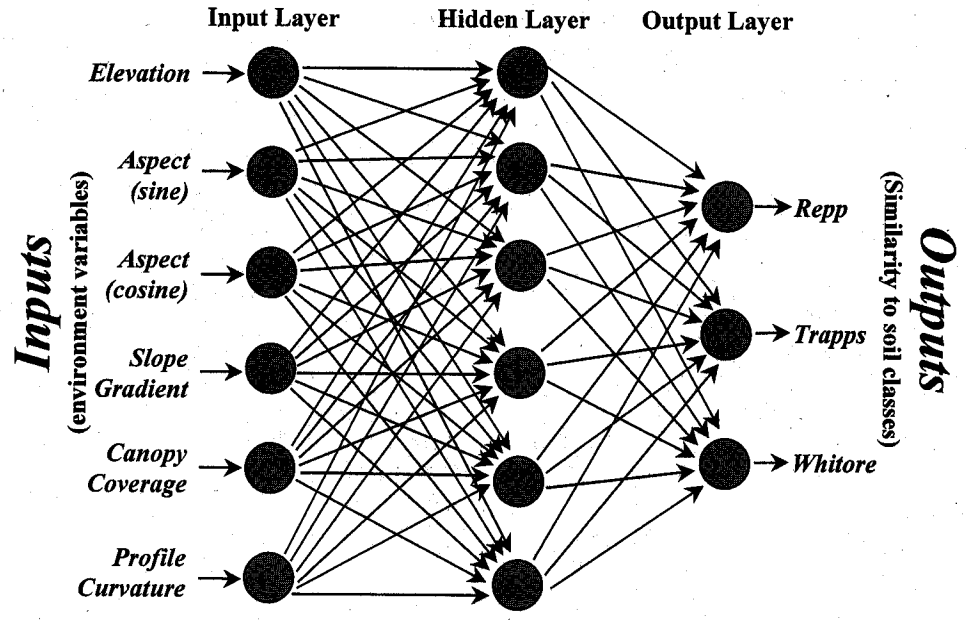

Figure 2. The structure of the neural network used in this study for the limestone area, with environmental variables as inputs, six hidden neurons, and three output neurons representing the three soil series in the area (section 3.2).

The structure of a MFN (the numbers of layers and the number of neurons in each layer) is problem specific. *Masters* [1993, pp. 85–88] reports that for most problems a three-layer model should be sufficient unless the problem domain is highly discontinuous. The number of input neurons should be the number of inputs, and the number of output neurons should be the number of attributes defining the output. The number of hidden neurons depends on the complexity of the problem [*Masters*, 1993, pp. 176–177]. The simpler the problem is, the fewer the number of hidden neurons are needed. The exact number of hidden neurons for a given problem is often determined by trial and error [*Masters*, 1993, p. 178].

3.1.2. Training. Network training (network learning) is a process of determining a set of weights that will produce the best possible input/output mapping. Most network training employs a supervised approach under which the network is presented with a set of input patterns and a set of corresponding desired outputs (together referred to as training data). The training process starts by initializing all weights to small nonzero values. Then, training samples are presented to the network one at a time to produce corresponding results. A measure of error between the network outputs and the desired outputs is computed, and weights are updated to reduce error. Many iterations or epochs (from presenting training samples to measuring error and to updating weights) may be required before a network reaches a given level of accuracy.

Apparently, the magnitude of error is determined by the combination of weights for a given network. There exists an error surface that can be described as a function of weights. The objective of network training is to find a set of weights that will minimize the error function. There are a number of network learning algorithms for minimizing the network error function [*Masters*, 1993, pp. 94–110; *Rogers and Dowla*, 1994]. The conjugate gradient method used in this study is briefly described here.

The conjugate gradient method is an improvement over the conventional back propagation method by using adaptive approaches to the determination of learning rate (μ) and momentum (η) which control the convergence and the speed of network training [*Masters*, 1993, pp. 105–111]. The learning rate determines the rate at which the weights should be modified during each epoch. The momentum controls the direction to search for a minimum. The conjugate gradient method first employs the Polak-Ribiere algorithm [*Polak*, 1971] to determine the best direction to search for the minimum. It then uses the directional minimizing process to determine the location of the minimum. *Johansson et al.* [1992] compared a conjugate gradient method with several other methods including the conventional back propagation and found that the conjugate gradient method speeded convergence by an order of magnitude relative to the conventional back propagation. In addition, they found that the conjugate gradient algorithm is less vulnerable to instabilities and oscillatory difficulties than the conventional back propagation method.

The error surface with respect to weights for a neural network can contain a large number of local minima as a result of a large number of permutations of weights producing similar input/output mapping. These local minima can make network training even more complicated. Avoiding false minima consists of two steps. The first is to avoid initiating weights in the vicinity of these minima, and the second is to determine if the found minimum is local. If it is, try to escape from this local minimum. *Masters* [1993] describes two techniques, simulated annealing and genetic algorithm, for both initially avoiding and later escaping local minima. Interested readers are referred to *Masters* [1993] for further reading on these two techniques. The simulated annealing was employed in this study.

3.2. Computing Similarity Using Neural Networks

The ANN approach for populating the similarity model is based on the soil-landscape relationship concept outlined by V. V. Dokuchaeiv and E. W. Hilgard [see *Hudson*, 1992] and further developed by *Jenny* [1941, 1980]. This concept contends that soil is the result of the interaction of its formative environmental factors over time as described in (1):

$$S = \int f(E)\, dt, \tag{1}$$

where S is soil, f is the relationship of soil development to the soil formative environment (E), and t is time. Since it is difficult, if not impossible, to explicitly describe the t factor at every location across landscape and sometimes information on t is implicitly expressed in other formative environmental factors such as topographic positions, (1) is simplified to

$$S = f(E). \tag{2}$$

Under the soil similarity model, S represents soil similarity vectors. Soil formative environmental conditions E generally include variables describing climate (Cl), topography (Tp), parent materials (Pm), and vegetation factors (Og) (Figure 3). It is understood that current environmental conditions may not reflect the conditions under which a soil was formed. Because of the difficulty of obtaining spatial data on soil paleoenvironment conditions, current environment conditions of soils are used as surrogates to E and can be characterized using GIS and remote sensing techniques.

The exact relationships f between soil development and formative environmental conditions are often very complicated and, in many cases, are unknown for a specific area. These relationships may not be sufficiently characterized by general regression models because of their nonlinear and nonstationary nature. Artificial neural networks have been shown to be capable of approximating virtually any function as long as sufficient representation is available [*Hornik et al.*, 1989]. It is possible then to use neural networks to approximate f for a given study area. The input to the network is a set of environmental conditions characterizing the soil formative environment, with each input neuron representing an environmental variable. The output represents a set of prescribed soil types with each output neuron representing a soil type (Figure 2). The activation level of each output neuron is perceived as the degree to which the input environmental conditions at point (i, j) are similar to the typical environmental conditions of the soil type that neuron represents; thus it is used to approximate the similarity of the soil at (i, j) to the soil type of the output neuron.

A training sample consists of a set of environmental conditions at a sample site and a set of desired membership values expressing the similarity of the local soil to the list of prescribed soil types. Each similarity value is used as the activation level of an output neuron representing the soil type. The neural network can be trained on a set of representative samples to learn the mapping between the input environmental conditions and the set of prescribed soil types for that area by

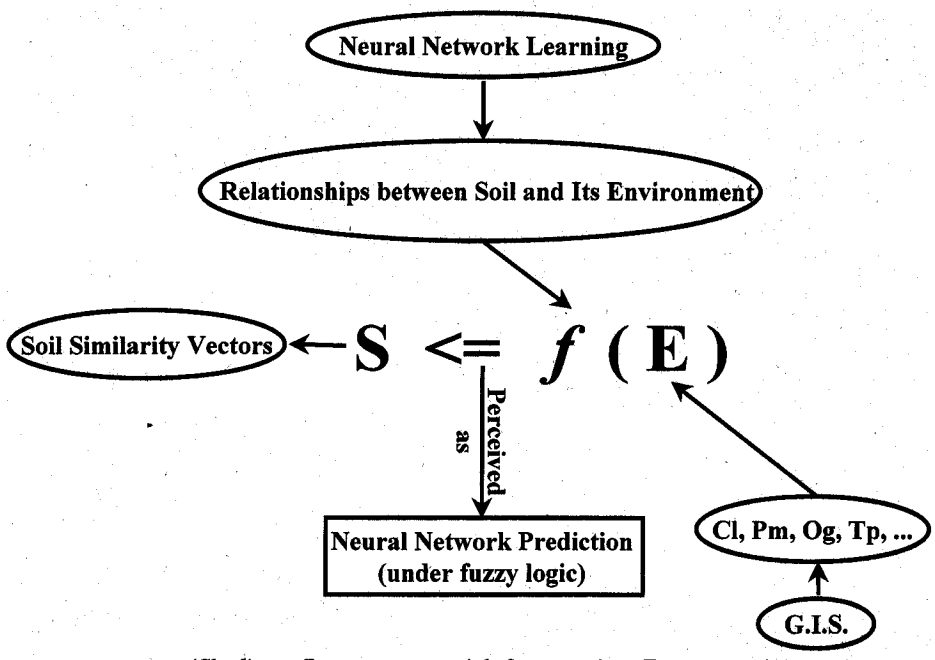

Figure 3. Soil inference based on the neural network approach.

minimizing the difference (sum of squared residuals) between the predicted and desired activation levels. This mapping can then be used to approximate f for the area. Once trained, the network can be used to produce a set of activation levels with respect to a set of environmental conditions for a given location (pixel). This set of activation levels can be used to approximate the similarity values and used to populate the similarity vector for that location (Figure 3).

4. Case Study

4.1. Study Site and Environmental Variables

A case study was conducted to evaluate the usefulness of the neural network approach for populating the similarity model. The study area is the Lubrecht Experimental Forest, which was established in 1937 to foster research on natural resources [*Nimlos*, 1986]. The area is about 50 km northeast of Missoula and is in the mountainous terrain of western Montana with a moderate to strong relief. High elevations are in the east and southwest. Low elevations run from southeast through northwest. The climate is considered to be semiarid to semihumid. There is a strong contrast in terms of moisture conditions between north and south facing slopes and between low and high elevations. Slopes facing south at low elevations have poor moisture conditions, while the moisture conditions on north facing slopes at high elevations are better. Soils on slopes with poor moisture conditions have a shallower soil profile than soils on slopes with better moisture conditions.

Twelve soil series have been identified in the area belonging to three soil orders: Entisol, Inceptisol, and Alfisol [*Nimlos*, 1986]. Entisols are weakly developed soils with very little organic matter accumulation and no illuvial clay or sesquioxides. In Lubrecht, Entisols are usually found on ridge crests. Inceptisols are young soils with little to no illuvial clays but brown subsoil horizons that indicate some translocation of sesquioxides. About 90% of the soils (in terms of areal extent) in the study area are Inceptisols. Alfisols are soils with leached, gray surface horizons (E) and subsurface horizons with accumulations of illuvial clay (Bt). Alfisols are found on high-elevation mountain slopes in the study area.

There are three major bedrock types in the area: Belt rocks, granite, and limestone. Each bedrock type is contiguous with Belt rocks in the north, granite materials in the south, and limestone in the middle (Figure 4). Soils on these three materials are formed from a mantle of colluvium. Belt rocks are the oldest rock in the area and were deposited during the Precambrian period about one billion years ago. The sediments from which they were formed were deposited in a shallow sea, subsequently buried, and then metamorphosed into quartzites, argillites, and siltites [*Nimlos*, 1986]. Soils are similar within each of the three bedrock types, but they are significantly different among the different bedrock types [*Nimlos*, 1986]. There are five different soil series on the Belt rock parent materials (Evaro, Sharrot, Tevis, Winkler, and Winkler Cool), four soil series on the granite materials (Ambrant, Elkner, Ovando, and Rochester), and three on the limestone materials (Repp, Trapps, and Whitore).

A total of six environmental variables (elevation, slope aspect, slope gradient, tree canopy coverage, profile curvature, and bedrock geology) were used in this study to characterize the soil formative environment. Elevation, slope aspect, slope gradient, and profile curvature were used to approximate topography. Tree canopy coverage was used to estimate vegetation conditions, and bedrock geology was used as a surrogate for parent materials. This list does not contain variables that directly measure climatic factors. Although the study area is a small drainage basin (about 5 km by 8 km), great differences in terms of microclimate do exist within the area. However, these differences in microclimate are well expressed by variations in elevation and slope aspect. The list also does not contain variables relating to the time factor. This is because it is difficult, if not impossible, to derive information on the time factor at

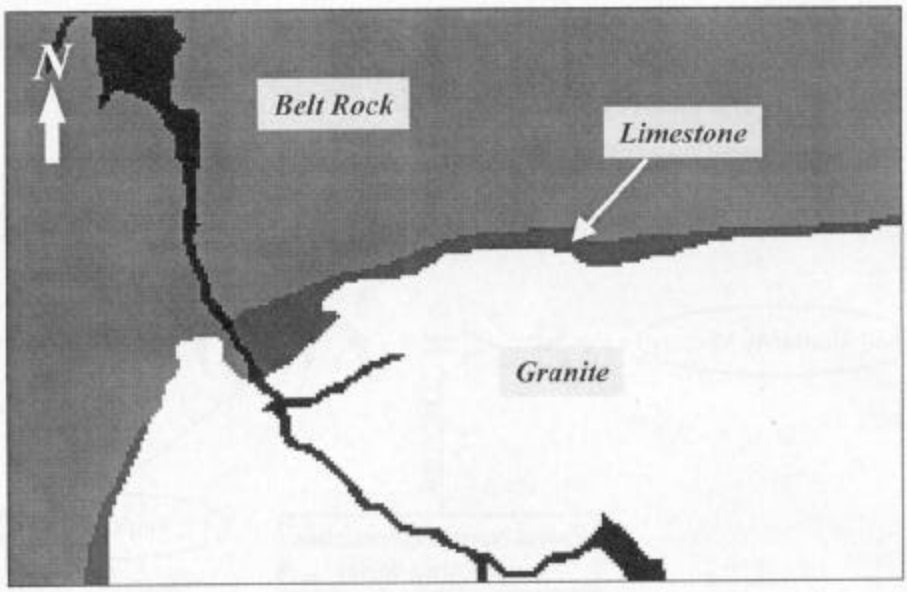

Figure 4. Distribution of bedrock materials in the study area. Solid areas were not included in this study.

every location across the landscape. This variable list is by no mean considered complete in terms of soil formation factors, but it does contain the major variables for which GIS data can be derived to characterize the soil formative environment. The selection of these six variables was also based on the consultation with the local soil scientists who believed these variables captured the major driving force of soil formation in the area [*Zhu et al.*, 1997; *Zhu*, 1999b]. The other reason for using these six variables was that these variables were originally used in a knowledge-based approach [*Zhu et al.*, 1997]. In order to compare this study with the knowledge-based approach, the author believes it is better to use the same set of variables.

Information on elevation, slope aspect, slope gradient, and profile curvature were derived from a digital elevation model (DEM) of the study area. The accuracy of the DEM is comparable with that of level 1 U.S. Geological Survey 7.5-min DEMs. The tree canopy coverage data were approximated with an index derived from the thematic mapper data of the area [*Nemani et al.*, 1993]. Information on parent materials was obtained from a bedrock geology map of the area [*Brenner*, 1968]. The reason for using the bedrock geology map to approximate the soil parent materials is that the bedrock map is the only geological information available for the area, yet the soils exhibit a great dependence on the parent materials in terms of physical and chemical properties.

Data on these variables (except the geology data layer) were preprocessed to facilitate the training process and simplify the network structure. The slope aspect was converted into two variables (a cosine and a sine portion of aspect) to resolve the discontinuity at 360°–0°. Although there are no theoretical limits on the inputs of neural networks, stability of neural networks is usually improved by using comparable limits among the input variables [*Masters*, 1993, p. 262]. All environmental data were thus scaled to the range of 0.0–10.0 for the numerical stability of neural networks. The reason for choosing 10 as the upper limit rather than the commonly used value of 1 is that the values for elevation are in the thousands and scaling them into the range of 0 and 1 could cause a significant amount of rounding. By using an upper bound of 10, this rounding could be reduced.

Soil properties are very different among the different bedrock areas, and sharp boundaries in terms of soil properties do exist between the different bedrock areas in the region [*Nimlos*, 1986]. It is assumed that the pedogenesis and the soil-environment relationships are different among the different bedrock areas but are similar within each of them. On the basis of this assumption the study area was divided into three subregions based on bedrock type: Belt rock subregion in the north, granite in the south, and limestone area in the middle (Figure 4). The soil-environment relationship in each subregion was then approximated by a neural network. For the entire study area a total of three different networks were designed. Since the relationship within each subregion is expected to be more similar, the structure of the neural network for a given subregion is expected to be less complex than the structure of a neural network for the entire study area. A simpler network structure would reduce the complexity of the weight space and thus would facilitate network training.

4.2. Collecting Training and Validation Sets

Two sample sets (one for training and one for validation) were collected for each bedrock subregion. The training set was collected independently of the validation set. The number of samples in a training set was related to the number of soil series in that bedrock area. The assumption is that the more soil series there are, the more complex the input/output mapping will be and, in turn, the more samples will be needed to train the network. For this study the training sample size was set to be at least 30 times the number of soil series (Table 1). The size of the validation set was at least 50% of the size of the respective training set.

For collecting a sample set a GIS-based procedure was developed. The procedure was designed to produce a sample set in which samples were mostly randomly distributed over space, and the attribute frequency distributions approximated those within the given bedrock area. The objective was to gather a

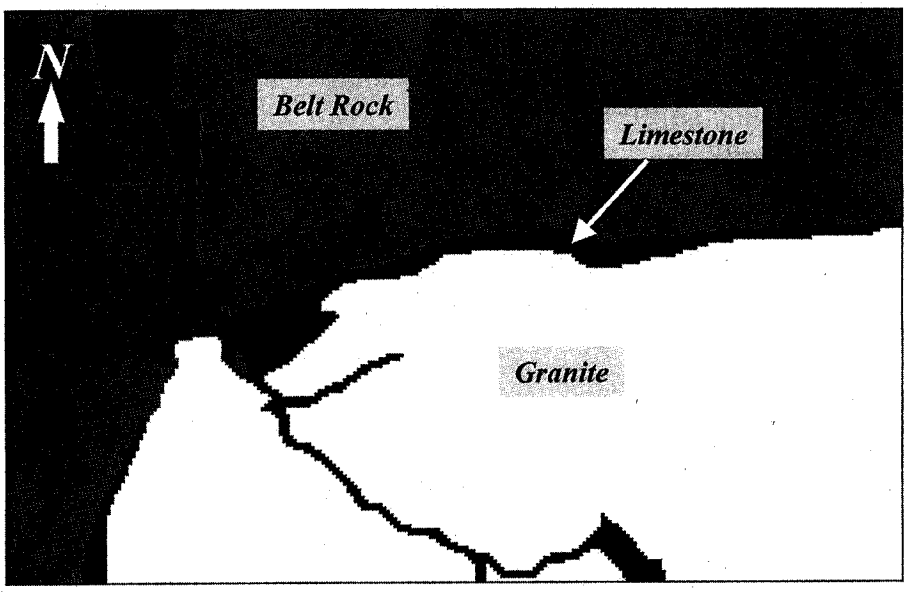

Figure 4. Distribution of bedrock materials in the study area. Solid areas were not included in this study.

every location across the landscape. This variable list is by no mean considered complete in terms of soil formation factors, but it does contain the major variables for which GIS data can be derived to characterize the soil formative environment. The selection of these six variables was also based on the consultation with the local soil scientists who believed these variables captured the major driving force of soil formation in the area [*Zhu et al.*, 1997; *Zhu*, 1999b]. The other reason for using these six variables was that these variables were originally used in a knowledge-based approach [*Zhu et al.*, 1997]. In order to compare this study with the knowledge-based approach, the author believes it is better to use the same set of variables.

Information on elevation, slope aspect, slope gradient, and profile curvature were derived from a digital elevation model (DEM) of the study area. The accuracy of the DEM is comparable with that of level 1 U.S. Geological Survey 7.5-min DEMs. The tree canopy coverage data were approximated with an index derived from the thematic mapper data of the area [*Nemani et al.*, 1993]. Information on parent materials was obtained from a bedrock geology map of the area [*Brenner*, 1968]. The reason for using the bedrock geology map to approximate the soil parent materials is that the bedrock map is the only geological information available for the area, yet the soils exhibit a great dependence on the parent materials in terms of physical and chemical properties.

Data on these variables (except the geology data layer) were preprocessed to facilitate the training process and simplify the network structure. The slope aspect was converted into two variables (a cosine and a sine portion of aspect) to resolve the discontinuity at 360°–0°. Although there are no theoretical limits on the inputs of neural networks, stability of neural networks is usually improved by using comparable limits among the input variables [*Masters*, 1993, p. 262]. All environmental data were thus scaled to the range of 0.0–10.0 for the numerical stability of neural networks. The reason for choosing 10 as the upper limit rather than the commonly used value of 1 is that the values for elevation are in the thousands and scaling them into the range of 0 and 1 could cause a significant amount of rounding. By using an upper bound of 10, this rounding could be reduced.

Soil properties are very different among the different bedrock areas, and sharp boundaries in terms of soil properties do exist between the different bedrock areas in the region [*Nimlos*, 1986]. It is assumed that the pedogenesis and the soil-environment relationships are different among the different bedrock areas but are similar within each of them. On the basis of this assumption the study area was divided into three subregions based on bedrock type: Belt rock subregion in the north, granite in the south, and limestone area in the middle (Figure 4). The soil-environment relationship in each subregion was then approximated by a neural network. For the entire study area a total of three different networks were designed. Since the relationship within each subregion is expected to be more similar, the structure of the neural network for a given subregion is expected to be less complex than the structure of a neural network for the entire study area. A simpler network structure would reduce the complexity of the weight space and thus would facilitate network training.

4.2. Collecting Training and Validation Sets

Two sample sets (one for training and one for validation) were collected for each bedrock subregion. The training set was collected independently of the validation set. The number of samples in a training set was related to the number of soil series in that bedrock area. The assumption is that the more soil series there are, the more complex the input/output mapping will be and, in turn, the more samples will be needed to train the network. For this study the training sample size was set to be at least 30 times the number of soil series (Table 1). The size of the validation set was at least 50% of the size of the respective training set.

For collecting a sample set a GIS-based procedure was developed. The procedure was designed to produce a sample set in which samples were mostly randomly distributed over space, and the attribute frequency distributions approximated those within the given bedrock area. The objective was to gather a

Table 1. Sample Size for Each Bedrock Area

Bedrock Type	Soil Series Number	Training Set Size	Validation Set Size
Belt	5	150	90
Granite	4	140	80
Limestone	3	90	60

sample set that represented as much as possible the characteristics of the population in terms of environmental conditions. Figure 5 displays the elevation and aspect frequency curves of the entire population, the training set, and the validation set for the granite bedrock area. Figure 5 shows that both the training and validation sets approximate the attribute frequency distribution of the whole population well. Each of these sample sets was thus considered here to be a good representative of the population in its respective bedrock area. At each sample site the soil type was determined from an existing soil map, and the environmental conditions described by the six variables were derived from a GIS database.

To train neural networks under the similarity model, the local soil at each sample site must be represented as a similarity vector with each membership being treated as the activation level for an output neuron. It is difficult, if not impossible, to determine the soil similarity vector without introducing a great deal of uncertainty at each sample site. Thus the soil at each sample site was treated just as a typical case of the soil type being assigned to the site and was represented by a vector that contains a membership of 0.9 in the soil type to which the local

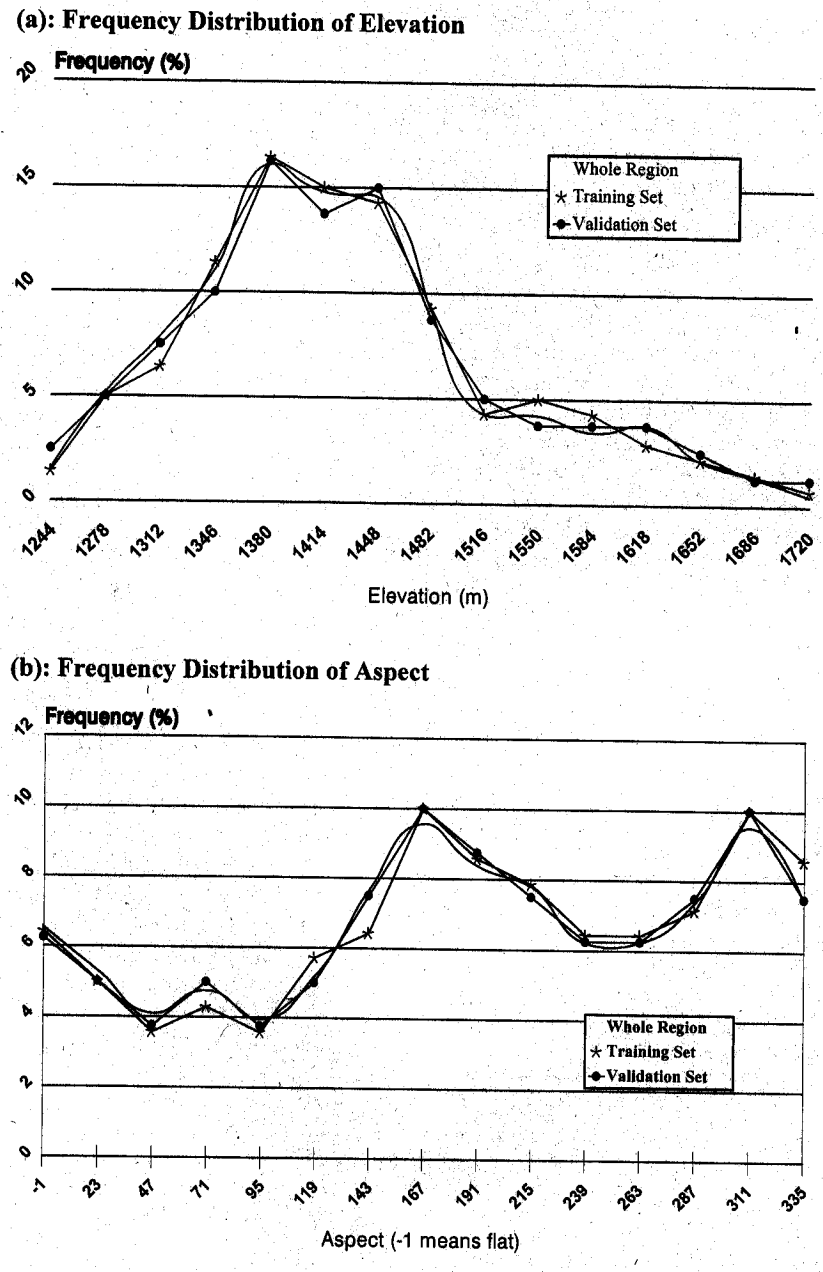

Figure 5. Attribute frequency distributions of the whole population and the training and validation sets for the granite area.

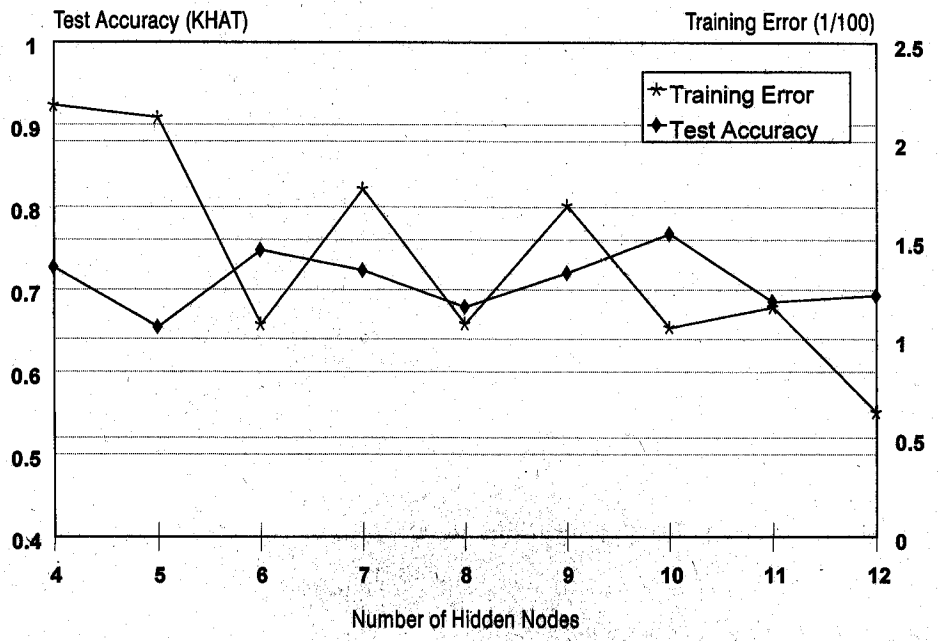

Figure 6. Network performance versus number of hidden neurons for the Belt rock area. Test accuracy is based on the validation set, while training error is based on the training set.

soil was labeled. Membership values in other soil types were set to 0.1. The reason for choosing 0.9 and 0.1 as the membership limits is that these values are the upper and lower bounds of the output from the neural network program used in this study. These limits are fairly typical for neural network outputs.

Treating every training sample as a typical case of a given soil type could introduce errors into the final results. There are two scenarios under which this can happen. The first is that when a training sample is at a transitional area and bears high memberships in two or more soil categories. Assigning 0.9 in one category and 0.1 in others would certainly exaggerate the membership in the first category and bias the training of the network. In order to avoid some of these points, the randomly generated sample sites were inspected to remove points at or close to boundaries of soil polygons where transitional soils were expected to exist. The second scenario is that the soil of a sample is an inclusion in the soil polygon. By assigning the local soil 0.9 membership in the soil category which the soil polygon is labeled will cause the sample to be mislabeled. Since the existing soil map is the only source for obtaining samples in this study and the map legend does not indicate the spatial distribution of inclusions, there is not an easy way to identify these cases. However, the soil map units in the study area are mostly single class units, and inclusions in these units are small in terms of area coverage. It was assumed that these mislabeled samples would not be too many and would not dramatically alter the general patterns of soil-environment relationships extracted by the networks.

4.3. Network Structure

All of the three networks were three-layer feed forward networks. The number of input neurons equaled the number of environmental variables (except the bedrock type variable that was used to divide the area into three subregions); thus all three networks had six input neurons (the aspect variable was split into two). The number of output neurons was the number of soil series in each of the bedrock areas. There were five output neurons (five soil series) for the Belt bedrock area, four for the granite area, and three for the limestone area.

The number of hidden neurons on the hidden layer differed from one bedrock area to another and was determined through a trial-error approach modified from the method described by *Masters* [1993, p. 178]. First, the network training error and test accuracy were used as criteria to assess network performance. The training error (sum of squared residuals) measures the difference between the levels of activation and the prescribed similarity values for the training set. The test accuracy is measured using the kappa (KHAT) coefficient [*Bishop et al.*, 1975; *Cohen*, 1960], which was estimated from the confusion matrix (error matrix) [*Congalton*, 1991] constructed from the results of the validation set. It measures how well the results from the network match the observed for the validation set. The lower the training error is and the higher the test accuracy is, the better the network performance will be. Starting with four hidden neurons which was considered too small a number for the network, the network was trained and tested, and its performance (training error and test accuracy) was recorded. The number of hidden neurons was then increased by 1, and the network was trained and tested again. This process was repeated until the number of hidden neurons was 3 times the number of output neurons or until the number of weights in the network approached the total number of training samples. Figure 6 shows the network performance versus the number of hidden neurons for the Belt rock area.

The number of hidden neurons should be the number at which the network performed well (with a high accuracy and low training error). There are several possible numbers of hidden neurons as shown in Figure 6. The number of hidden neurons could be 6 or 10. The final decision among these possible numbers was made based on the test accuracy for individual soil categories over the validation set. This is justified because the networks were trained to learn from the training data set. It is possible that with some structures (defined by

Table 2. Test Accuracy for the Belt Area With Six and Ten Hidden Neurons

Hidden Neurons	Evaro	Sharrot	Tevis	Winkler	Winkler Cool
6	0.75	0.67	0.81	0.00	0.92
10	0.75	0.80	0.81	0.50	0.95

the number of hidden neurons in this case) the network could be tuned to the training data set too much (overfitting). As a result, the network would not generalize well to a different data set. Thus the selection of the number of hidden neurons should be based more heavily on the test accuracy from the validation data set while still considering the training error. Table 2 shows the test accuracy for each soil series over the Belt area. On the basis of Table 2 the number of hidden neurons of the network for the Belt area was set to be 10 since the test accuracy for soil series Winkler is much higher with 10 hidden neurons than with six. Similarly, the number of hidden neurons was set to seven for the granite area and six for the limestone area. The structure of the network used for the limestone area is shown in Figure 2.

4.4. Network Learning

The above networks were implemented in the NEURAL program developed by *Masters* [1993, p. 404]. Each of the networks was first initialized using the simulated annealing approach [*Masters*, 1993, pp. 118–134]. After initialization the conjugate gradient algorithm was used to locate a nearest minimum (Figure 7). When a minimum was found, simulated annealing was used to attempt to break out of what might be a local minimum. If annealing succeeded in reducing the error (escaping from the local minimum), the conjugate gradient method was used again to find a minimum in this new part of the error surface. This alternation between conjugate gradient minimization and annealing escaping continued until several iterations in a row produced only trivial improvement or no improvement at all, and the minimum was then marked to be a candidate for the global minimum. Annealing was again used to find an entirely new set of starting weights to look for another candidate. This process was repeated 10 times [*Masters*, 1993, p. 409]. All of the candidates were then compared, and the best was chosen to be the global minimum.

5. Results and Discussion

Once a network for a given bedrock area was trained, it was used to populate the soil similarity vectors over that bedrock area. To compute the soil similarity vector for a given pixel, the GIS data on the environmental conditions for that pixel were fed to the trained network as inputs. The network then computed a corresponding output vector containing the activation levels at the output neurons. This vector of activation levels was taken as the similarity vector for the given pixel. Once all pixels had been visited by the network, the similarity representation of soil over the bedrock area was produced.

Soil similarity vectors derived from the ANN approach for a few selected sites on the Belt area are listed in Table 3 to illustrate the concept and the importance of the similarity representation. For example, soil at Lub03_02 bears similarity to all soil series found in the Belt area with the highest membership in Tevis and second highest membership in Winkler.

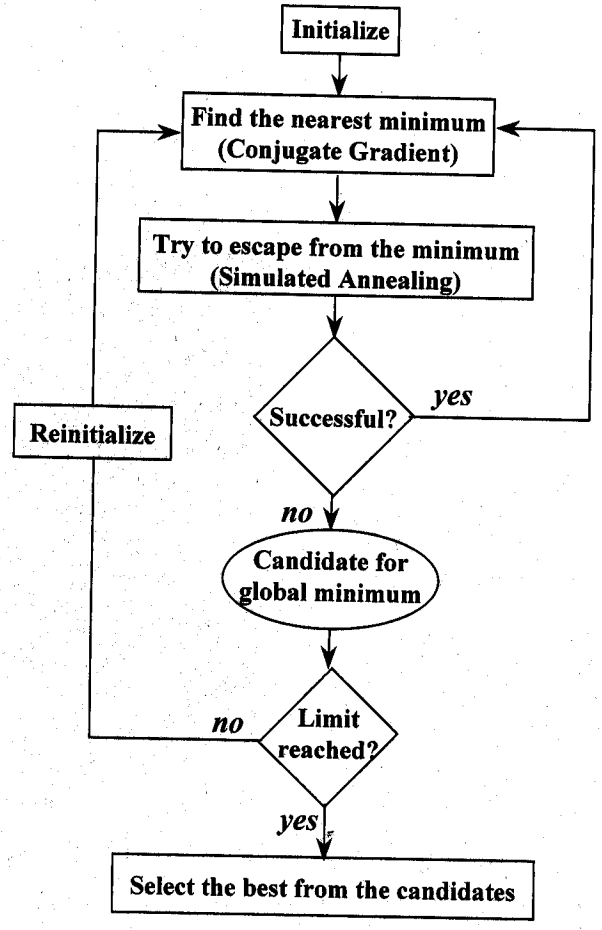

Figure 7. Network training using conjugate gradient and simulated annealing.

Soil at Lub04_01 also has the highest membership in Tevis, but it has a very different membership distribution in the vector. It bears a strong similarity to Evaro. Under the Boolean classification both soils would be assigned to soil series Tevis. Once so assigned, the soils would be assumed to be no different from each other and no different from the typical type of Tevis. Under the similarity representation it is the distribution of membership in the entire vector that is important, not just the largest membership value. It is this membership vector that allows us to record the subtlety of the soil at a given location and its difference from its neighbors. It is this entire membership vector that allows us to derive soil property values intermediate to the typical values of the prescribed soil categories.

Zhu [1997] illustrated the uses of soil similarity values for

Table 3. Soil Similarity Vectors for a Few Selected Sites on the Belt Materials

Point	Evaro	Sharrot	Tevis	Winkler	Winkler Cool
Lub03_02	4*	15	82	30	9
Lub04_01	36	12	61	5	11
T2_07	11	9	48	5	53
T1_18	10	8	8	10	88

*The membership value in a specific category ranges from 0 to 100 (stretched from the output range of 0.1–0.9), but the sum of the membership values in a vector need not to be 100 [*Zhu*, 1997].

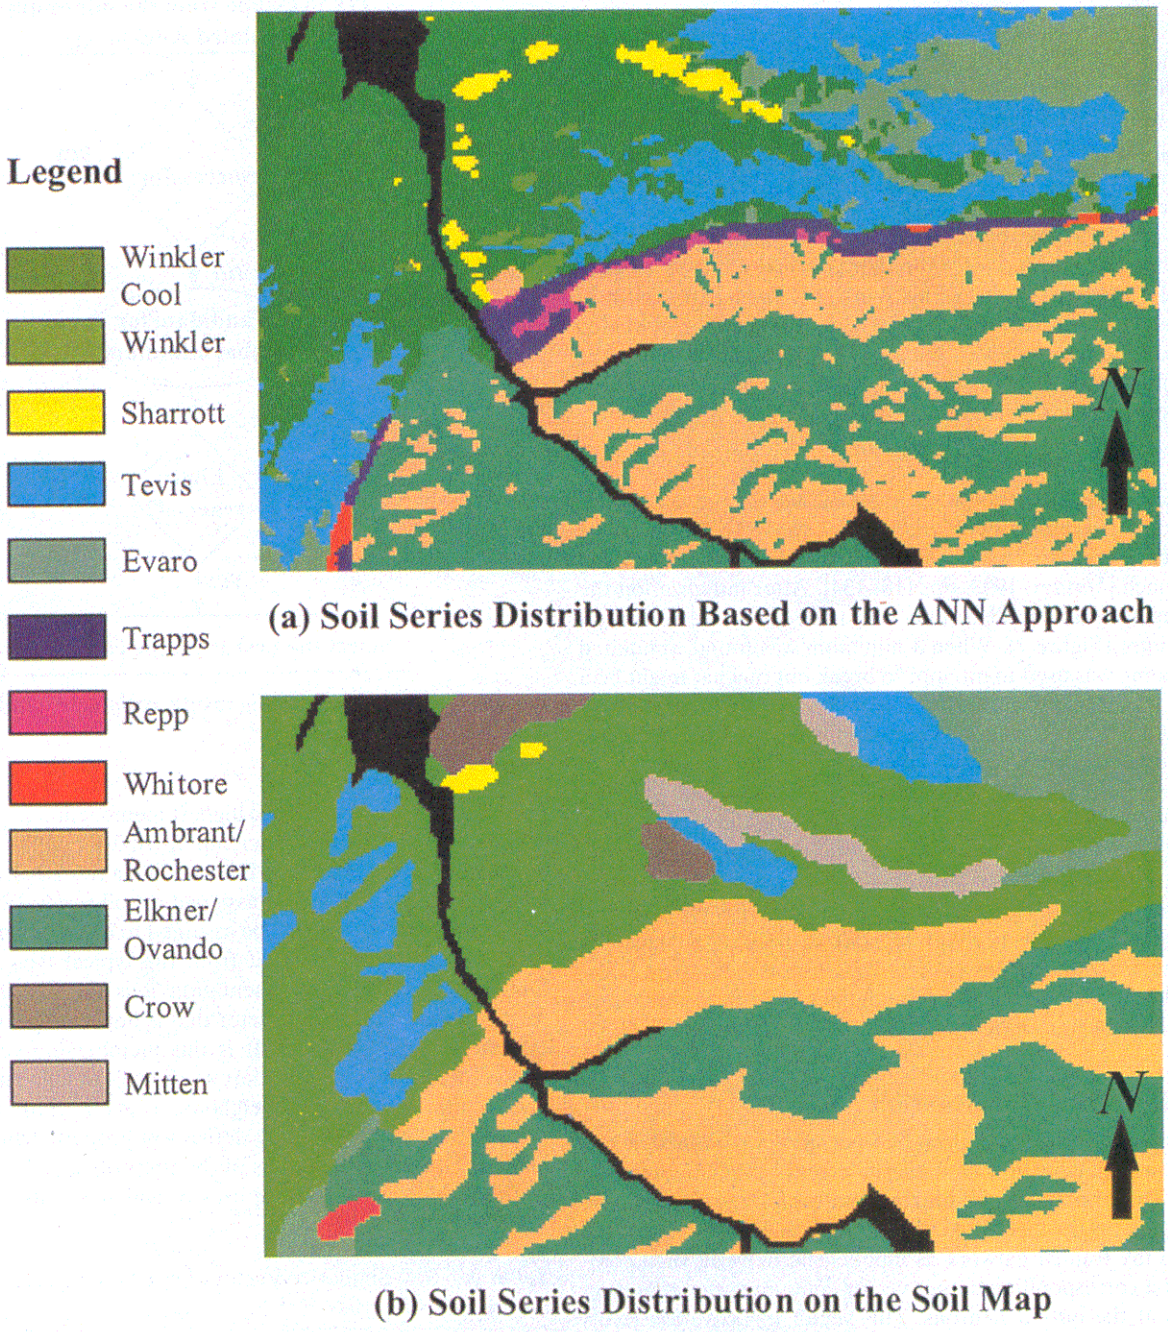

Plate 1. Comparison of the detailed soil map from the (a) artificial neural network (ANN) approach with the (b) conventional soil map. Solid areas were not included in this study.

Table 4. Comparison of Inferred, Mapped, and Observed Soil Series at 64 Sites

	Correct	Total Samples	Percentage
	Overall		
NN approach	49	64	77
KB approach*	52	64	81
Soil map	39	64	61
	Mismatches		
NN approach	13	29†	45
Soil map	4	29†	14
KB approach	17	24‡	71
Soil map	4	24‡	17

*Knowledge-based (KB) approach was conducted in another study [*Zhu*, 1997].

†These are mismatches between the neural network (NN) inferred soil series and those from the conventional soil map.

‡These are mismatches between the KB inferred soil series and those from the conventional soil map.

producing detailed soil spatial information. Two of these uses are employed here to examine the efficiency of this neural network approach to the derivation of soil similarity values. The first is the generation of a soil map for comparing with the conventional soil map to determine if the NN approach produced meaningful similarity values. The soil map derived from the similarity representation is referred to as the inferred soil map, which was produced through the conversion of the similarity representation to a Boolean representation. The conversion was accomplished by assigning the local soil to the soil series with the highest membership value [*Zhu*, 1997]. As shown in Plate 1, the inferred map (Plate 1a) contains much greater spatial detail than the conventional soil map (Plate 1b). In this semiarid to semihumid area of Montana, soil development is heavily influenced by slope moisture condition [*Nimlos*, 1986]. Moisture conditions on the side slopes of small draws (gullies) along a major south facing slope are significantly better than on the area facing directly south. Therefore the soils on these side slopes are expected to be better developed than and thus different from those on the areas facing directly south. These spatial details of soil spatial distribution do not show on the conventional soil map. However, the inferred soil map captures these spatial details.

To validate these spatial details, 64 sites were visited in the field, and the soil series at these sites were identified over the summers of 1991 through 1993. The inferred soil series and mapped soil series at these sites were then obtained from the inferred map and the conventional soil map, respectively. Table 4 summarizes the comparison among the inferred soil series and mapped soil series for field observations at these sites. Overall, the neural network (NN) approach inferred soil series correctly at 49 sites, which accounts for 77%, while the soil series from soil map matched the field observed soil series at only 39 sites. The knowledge-based (KB) approach conducted in another study inferred the soil series correctly at 52 sites [*Zhu*, 1997]. For a given site the NN inferred soil series may differ from that of the soil map. This is referred to as a mismatch. Whenever there was a mismatch, the field-observed soil series was used to determine which (ANN approach or the soil map) is correct. There were mismatches at 29 sites. The ANN approach inferred soil series correctly at 13 of these sites (45%); while at only 4 of these 29 sites, the soil series from the soil map were correct. With the KB approach, there were 24 mismatches between the KB inferred soil series and the soil series from the conventional soil map for these 64 field sites. The KB approach inferred correctly at 17 of these 24 mismatches (71%). These comparisons (both overall and through mismatches) suggest that the NN inferred soil map possesses much greater spatial detail than the conventional soil map, although it is not as good as the KB inferred soil map.

The other use of the soil similarity vector is for the derivation of detailed soil property maps [*Zhu et al.*, 1997]. Figure 8a is the soil A horizon depth image (NN inferred depth image) derived from the soil similarity representation using a linearly additive function described by *Zhu et al.* [1997]. The soil A horizon depth image derived from the KB approach is shown in Figure 8b. Figure 9c depicts the soil A horizon depth image derived from the conventional soil map. The differences between Figures 8a and 8b are much less than the difference between Figures 8a and 8c. As mentioned in section 4.1, slope moisture condition plays an important role in soil development in this semiarid to semihumid region. Soils at high elevations and on north facing slopes are better developed and have deeper A horizons than those at low elevation and south facing slopes. While all three images capture this general pattern of A horizon depth, the images (Figures 8a and 8b) from both the NN and KB approach portray the change of depth as a gradual spatial variation and at much greater spatial detail. Although abrupt changes of A horizon depth in this area are possible along the boundaries of the bedrock types, a more gradual change over space is more realistic, particularly within each bedrock area. Also, the changes of A horizon depth along the small draws on the major south facing slopes are clearly visible on both Figures 8a and 8b but are not shown in Figure 8c. This suggests that the depth image produced from the NN approach is comparable to that from the KB approach but that it has much greater spatial detail than what was portrayed in the soil map.

To examine the validity of the A horizon depth image, the A horizon depths at 33 field sites were measured in the summer of 1993. Figure 9 shows the scatterplots of the observed depths versus the NN inferred depths and the observed depths versus the depths from the conventional soil map at these sites. The distribution of dots in Figure 9a (the NN approach) is very similar to that in Figure 9b (the KB approach); the distribution aligns more closely along the 45° diagonal line than the distribution of the dots in Figure 9c. The correlation between the NN inferred and observed depths at these sites is similar to that for the KB approach but is much stronger than that between the depths from the soil map and the observed depths. In addition, the depths derived from the soil map are stratified along some values (Figure 9c), which are actually the typical values of the prescribed soil classes. Conversely, the inferred depths often have values intermediate to these typical values. From these comparisons it can be concluded that the NN inferred depths are as good as these from the KB approach and that they approximate reality better than the depths derived from the soil map.

The quality of soil information from the NN approach could suffer from the potential errors in the selection of training and validation samples under the two scenarios discussed in section 4.2. Since some caution was exercised to avoid selecting samples from transitional areas, it was expected that the impact on results due to atypical samples (samples in the transitional areas) is small. The small percentages of inclusions in single

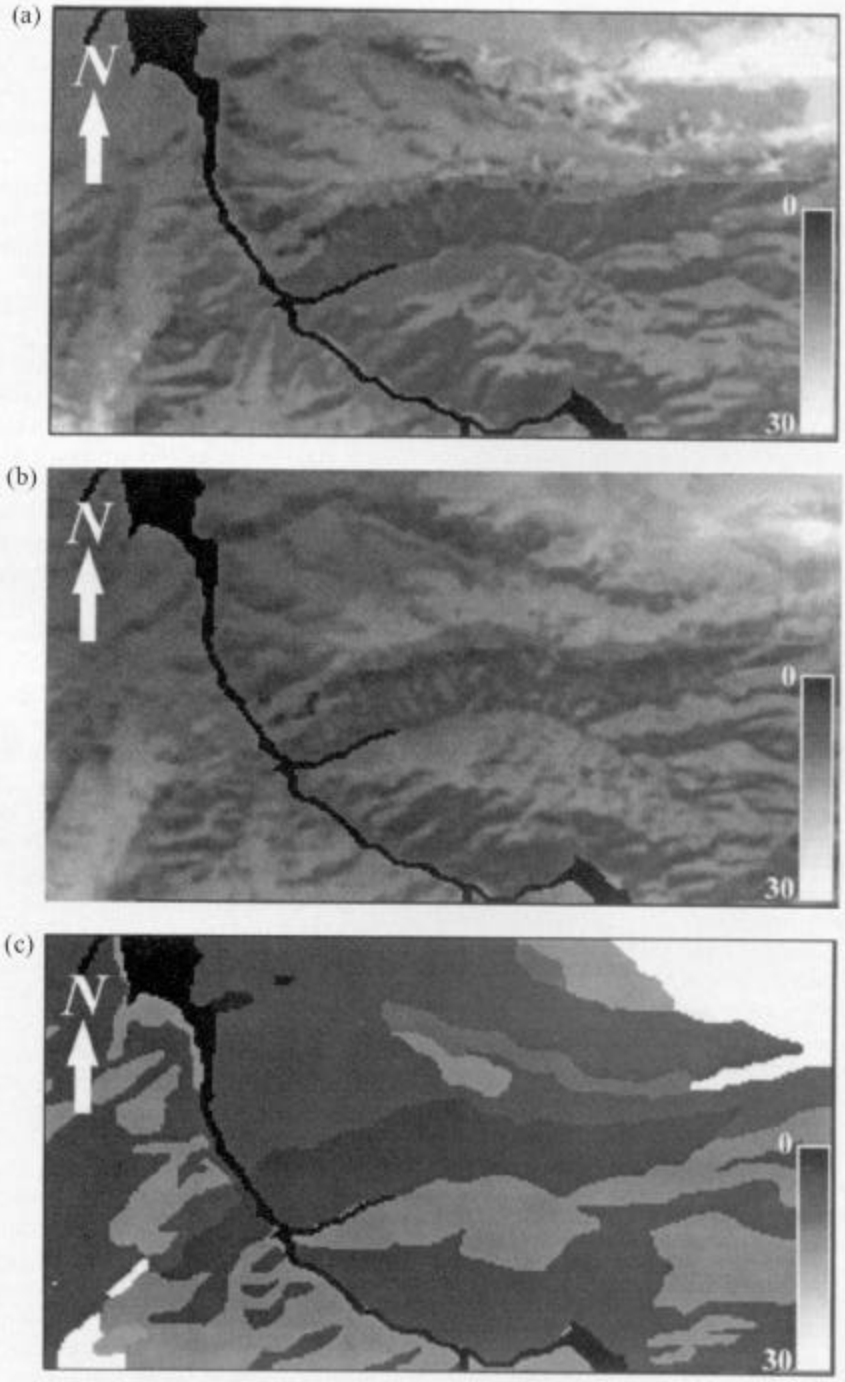

Figure 8. Comparison of A horizon depth images based on (a) the neural network (NN), (b) the knowledge-based (KB) approach, and (c) the conventional soil map. Solid areas were not included in this study.

soil class units could downplay the chance that a sizeable portion of the samples was selected from these inclusions. If a very small portion of samples was, indeed, selected from these inclusions, they will have a negative impact on the results, but they were not expected to hinder the ability of the network to learn the overall pattern of soil-environment relationships from the samples. Nevertheless, refining training and validation samples will improve the quality of results. The first approach to refining samples is to inspect the locations of samples and eliminate samples with the possibility of being in transitional areas and/or in inclusions. This approach would be useful when soil maps are the only source for providing samples. However, identifying possible areas of inclusions is not a simple task if soil maps are the only source, and research in this area is needed. The second approach would be the combination of samples from existing soil maps and samples from the field. The samples from the soil maps are only those with a high degree of typicality. The third approach is to use field samples only. This approach would be very costly, but the quality of results would be the highest.

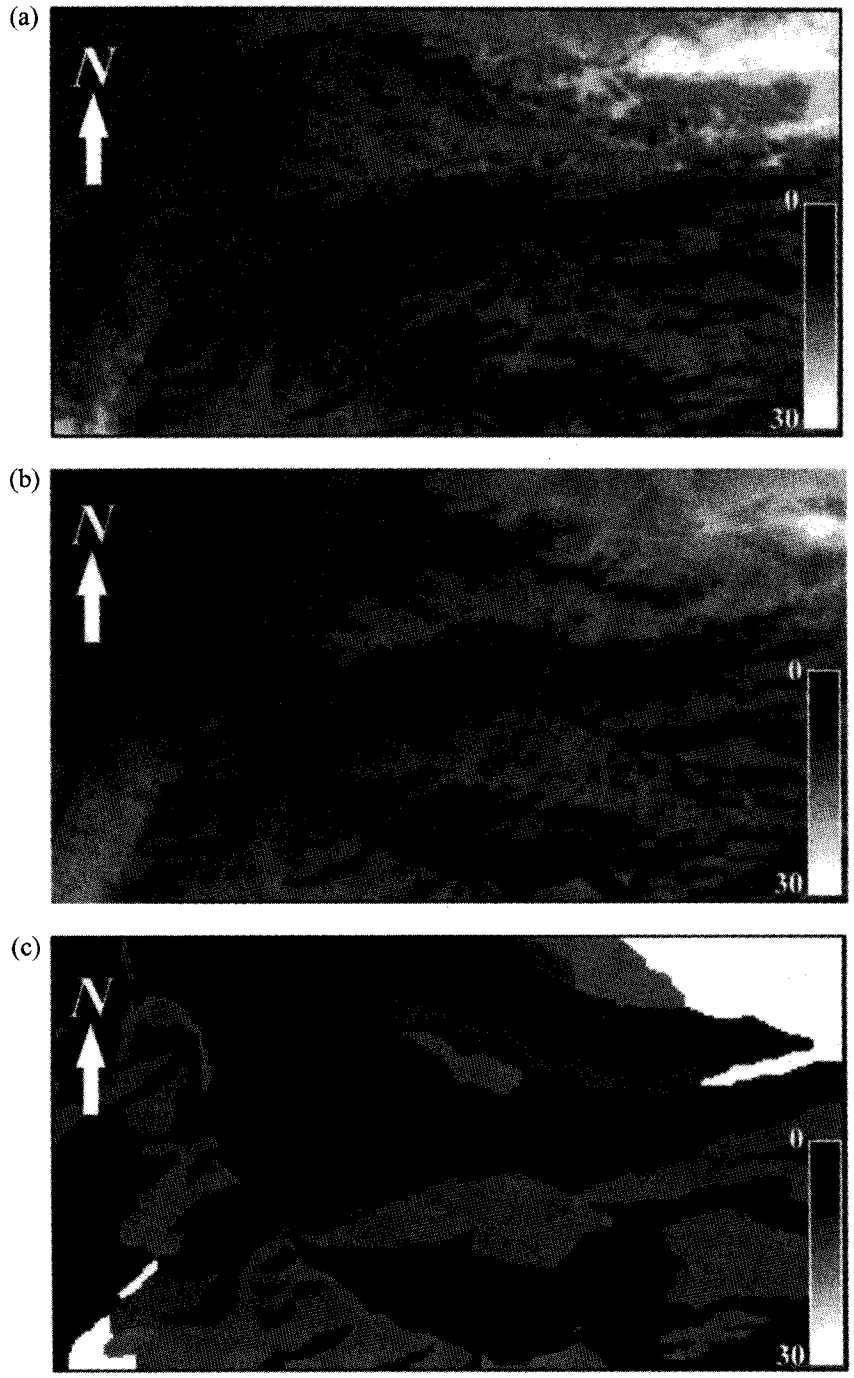

Figure 8. Comparison of A horizon depth images based on (a) the neural network (NN), (b) the knowledge-based (KB) approach, and (c) the conventional soil map. Solid areas were not included in this study.

soil class units could downplay the chance that a sizeable portion of the samples was selected from these inclusions. If a very small portion of samples was, indeed, selected from these inclusions, they will have a negative impact on the results, but they were not expected to hinder the ability of the network to learn the overall pattern of soil-environment relationships from the samples. Nevertheless, refining training and validation samples will improve the quality of results. The first approach to refining samples is to inspect the locations of samples and eliminate samples with the possibility of being in transitional areas and/or in inclusions. This approach would be useful when soil maps are the only source for providing samples. However, identifying possible areas of inclusions is not a simple task if soil maps are the only source, and research in this area is needed. The second approach would be the combination of samples from existing soil maps and samples from the field. The samples from the soil maps are only those with a high degree of typicality. The third approach is to use field samples only. This approach would be very costly, but the quality of results would be the highest.

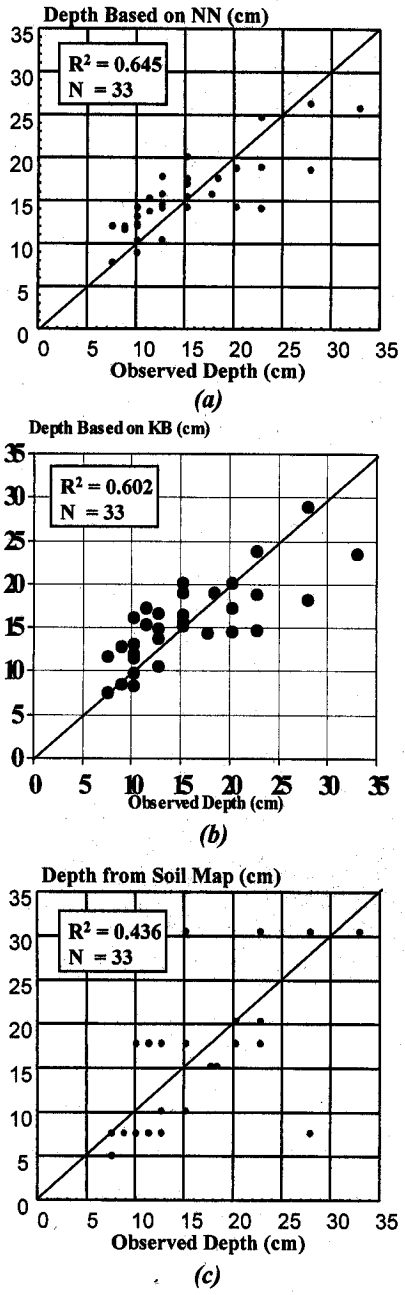

Figure 9. Scatterplots of (a and b) observed depths versus inferred depths and (c) observed depths versus depths from the conventional soil map.

6. Implications for Watershed-Based Hydroecological Modeling

The sensitivity of hydroecological modeling to this detailed soil spatial information is currently under investigation. In this section the potential impacts of this detailed soil spatial information on landscape parameter characterization for watershed-based modeling are discussed in the context of deterministic models. Two general groups of deterministic models [*Chow et al.*, 1988], distributed parameter and lumped parameter models, will be used to illustrate the potential impacts.

A distributed model is one in which the modeled process or phenomenon is explicitly calculated as a function of location in space [*Maidment*, 1993], and the model output at one location will be one of the inputs for a neighboring location. It is not only the overall outcome of the modeled process over the entire watershed that is of concern but, more importantly, the spatial variation of the modeled process within the watershed. Therefore distributed models require model parameters to be estimated at every location across the landscape, particularly the covariation of model parameters over space.

Suppose that a nonpoint source pollution model uses the depth of top soil and other landscape parameters (such as slope gradient and vegetation information) to model the process of pollutant generation and transport. Let us assume that the soil A horizon depth is used to approximate the depth of top soil. The covariation of A horizon depth with other model parameters (such as slope gradient and vegetation coverage) will control the model behavior. The covariation of A horizon depth with slope gradient along a transect in the Lubrecht area is shown in Figure 10. On the basis of the depth from the soil map a general pattern of A horizon depth along the transect is revealed: The greatest depth is on the north facing slope (area A), the shallowest depth is on the major south facing slope, and medium depth is on the plateau (area C). However, the depth does not seem to relate to the slope gradient. With the inferred depth map, not only the above general pattern of spatial variation of depth is depicted, but also an important covariation of depth with slope gradient is revealed. On the north facing slope the depth seems to be negatively related to the slope gradient. This is easy to understand since the moisture condition would be very similar over the north facing slope and the A horizon depth would depend on the erosion process. Slope gradient is often the indicator of the strength of erosion if other factors are the same. The greater the slope gradient is, the stronger the erosion would be and the thinner the soil A horizon would be. On the south facing slope the covariation of depth and slope gradient also follows this pattern at low elevation where both poor moisture condition and steep gradient prevent the development of a deep A horizon. However, this pattern of covariation is no longer valid at high elevation (area B) where the improved moisture condition permits the development of deeper A horizon, even though the slope gradient is high. On the plateau (area C) the covariation takes a very different form. It seems that the depth is positively related to the slope gradient. It is understandable that on a high-elevation plateau the local drainage condition can play a very important role in soil development. It is expected that soils in areas with better drainage (high slope gradient) would develop better than soils in poorly drained areas. This author argues that the characterization of these different patterns of spatial covariation of model parameters would be important to the outputs of distributed models at the watershed scale.

A deterministic lumped parameter model is an abstract representation of spatial features in which properties are averaged over a watershed or stream segment or slope facet [*Maidment*, 1993]. The process over an area is modeled as a single point in space without dimensions (or spatially averaged) [*Chow et al.*, 1988]. The parameters for these models are often the spatial averages (means) for the area units. It is then important to estimate as accurately as possible these means.

To illustrate the impacts on the characterization of means of model parameters, the Lubrecht watershed was partitioned into 42 area units (only 39 reported here, the other three fall outside of the study area) using a method of landscape partitioning from digital elevation data described by *Band* [1989]. For each partition (area unit), two mean A horizon depths

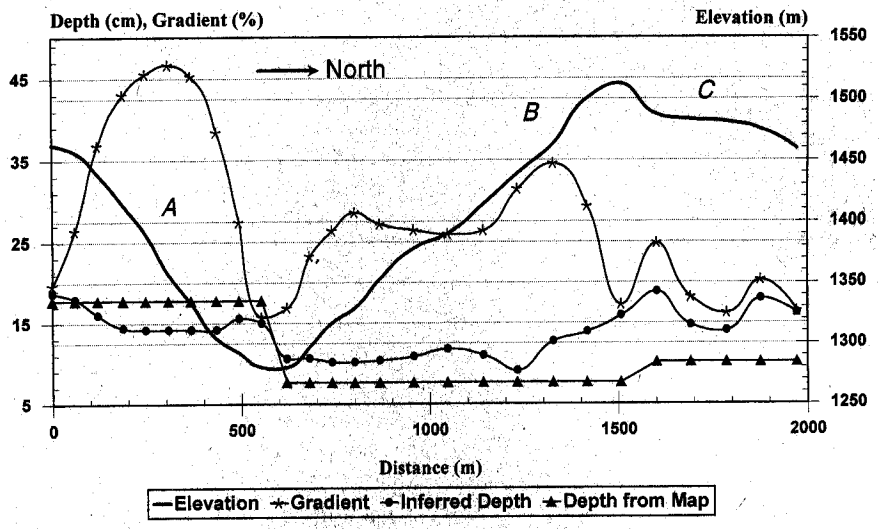

Figure 10. Covariation of A horizon depth with slope gradient and elevation along a transect.

were calculated: one using the inferred depth and the other using the depth from the soil map. The difference between these two means was then computed. The differences for 17 of 39 partitions are greater than 30% of the mean computed from the soil map. The largest difference reaches 110%. On average, the difference between the two sets of means is about 26%. A t test was performed to see whether the difference between the two sets of means is statistically significant. The calculated t is 8.72, and the critical t at 95% confidence with a degree of freedom of 38 is 1.69. The calculated t is much larger than the critical t. It can be concluded that the two groups of means are statistically different from each other. Therefore the output from a lumped parameter model using the means from the inferred depths is expected to be different from that using the means from the soil map.

7. Summaries

This paper presents a neural network approach to populating a soil similarity model for providing information on the detailed spatial variation of soil properties for hydroecological modeling at the watershed scale. Three multilayer feed forward networks were constructed; the inputs to the networks were data on soil formative environmental conditions, and the outputs were a set of prescribed soil categories. Representative samples were selected from a GIS database for training these networks to learn the relationships between soils and their formative environment over three different bedrock areas. The approach employed conjugate gradient and simulated annealing algorithms for training these networks. GIS data on soil formative environmental conditions for every location in each of the bedrock areas were then fed to the trained network to produce a similarity representation of soils over the bedrock area. The advantage of using the ANN approach over the knowledge-based approach [*Zhu et al.*, 1996] is that the ANN approach is not limited by the availability of experienced soil scientists in a local area.

A case study in the Lubrecht watershed in western Montana was conducted to examine the neural network approach. The case study shows that the soil map derived from the similarity representation contains much greater spatial detail and is more accurate than the conventional soil map. A soil A horizon depth image was derived from the similarity representation, and it was found that the spatial gradation of A horizon depth was more realistically portrayed in the inferred depth image than in that derived from the conventional soil map. Also, the inferred depths more closely match the observed depths at 33 field sites than the depths derived from the conventional soil map do.

Implications of this detailed soil spatial information for hydroecological modeling at the watershed scale were discussed. For distributed parameter models it was found that the detailed soil spatial information facilitates the characterization of detailed patterns of spatial covariation of model parameters. The detailed soil spatial information also impacted the characterization of mean conditions of model parameters for lumped parameter models. The sensitivity of hydroecological modeling to this detailed soil spatial information is yet to be studied.

Acknowledgment. The author is grateful to the Graduate School, University of Wisconsin-Madison, for the support that made the study reported in this paper possible.

References

Aleksander, I., and H. Morton, *An Introduction to Neural Computing*, 240 pp., Chapman and Hall, New York, 1990.

Band, L. E., Spatial aggregation of complex terrain, *Geogr. Anal.*, *21*, 279–293, 1989.

Band, L. E., and I. D. Moore, Scale: Landscape attributes and geographical information systems, *Hydrol. Processes*, *9*, 401–422, 1995.

Beven, K. J., and M. J. Kirkby, A physically based, variable contributing model of basin hydrology, *Hydrol. Sci. Bull.*, *24*, 43–69, 1979.

Bishop, Y. M., S. E. Fienber, and P. W. Holland, *Discrete Multivariate Analysis: Theory and Practice*, 557 pp., MIT Press, Cambridge, Mass., 1975.

Bregt, A. K., *Processing of Soil Survey Data*, 167 pp., Wageningen Agric. Univ., Wageningen, Netherlands, 1992.

Brenner, R. L., The Geology of Lubrecht Experimental Forest, *Lubrecht Ser. One*, 71 pp., Mont. For. Conserv. Exp. Stn., Sch. of For., Univ. of Mont., Missoula, 1968.

Burrough, P. A., Opportunities and limitations of GIS-based modeling of solute transport at the regional scale, in *Applications of GIS to the Modeling of Non-point Source Pollutants in the Vadose Zone*, edited by D. L. Corwin and K. Loague, *SSSA Spec. Publ.*, *48*, 19–38, 1996.

Burrough, P. A., R. A. MacMillan, and W. Van Deursen, Fuzzy classification methods for determining land suitability from soil profile observations, *J. Soil Sci.*, *43*, 193–210, 1992.

Burrough, P. A., P. Van Gaans, and R. Hootsmans, Continuous classification in soil survey: Spatial correlation, confusion and boundaries, *Geoderma*, *77*, 115–135, 1997.

Caudill, M., and C. Butler, *Naturally Intelligent Systems*, 304 pp., MIT Press, Cambridge, Mass., 1990.

Chow, V. T., D. R. Maidment, and L. W. Mays, *Applied Hydrology*, 572 pp., McGraw-Hill, New York, 1988.

Clair, T. A., and J. M. Ehrman, Using neural networks to assess the influence of changing seasonal climates in modifying discharge, dissolved organic carbon, and nitrogen export in eastern Canadian rivers, *Water Resour. Res.*, *34*(3), 447–455, 1998.

Cohen, J., A coefficient of agreement for nominal scales, *Educ. Psychol. Meas.*, *20*(1), 37–40, 1960.

Congalton, R. G., A review of assessing the accuracy of classification of remotely sensing data, *Remote Sens. Environ.*, *37*(1), 35–46, 1991.

Corwin, D. L., P. J. Vaughan, and K. Loague, Modeling nonpoint source pollutants in the vadose zone with GIS, *Environ. Sci. Technol.*, *31*(8), 2157–2175, 1997.

Dayhoff, J. E., *Neural Network Architectures: An Introduction*, Van Nostrand Reinhold, New York, 1990.

French, M. N., W. F. Krajewski, and R. R. Cuykendall, Rainfall forecasting in space and time using a neural network, *J. Hydrol.*, *137*, 1–31, 1992.

Goodchild, M. F., Geographical data modeling, *Comput. Geosci.*, *18*, 401–408, 1992.

Haykin, S., *Neural Networks, A Comprehensive Foundation*, 696 pp., Macmillan, Indianapolis, Indiana, 1994.

Hecht-Nielsen, R., *Neurocomputing*, 433 pp., Addison-Wesley-Longman, Reading, Mass., 1990.

Hertz, J., A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computing*, Addison-Wesley-Longman, Reading, Mass., 1991.

Hornik, K., M. Stinchcombe, and H. White, Multilayer feedforward networks are universal approximators, *Neural Networks*, *2*(5), 359–366, 1989.

Hsu, K., H. V. Gupta, and S. Sorooshian, Artificial neural network modeling of the rainfall-runoff process, *Water Resour. Res.*, *31*(10), 2517–2530, 1995.

Hudson, B. D., The soil survey as paradigm-based science, *Soil Sci. Soc. Am. J.*, *56*, 836–841, 1992.

Jenny, H., *Factors of Soil Formation: A System of Quantitative Pedology*, 281 pp., McGraw-Hill, New York, 1941.

Jenny, H., *The Soil Resource: Origin and Behaviour*, 377 pp., Springer-Verlag, New York, 1980.

Johansson, E. M., F. U. Dowla, and D. M. Goodman, Backpropagation learning for multi-layer feed-forward neural networks using the conjugate gradient method, *Int. J. Neural Syst.*, *2*(4), 291–301, 1992.

Jury, W. A., Spatial variability of soil properties, in *Vadose Zone Modeling of Organic Pollutants*, edited by S. C. Hern and S. M. Melancon, pp. 245–269, Lewis, Chelsea, Mich., 1985.

Jury, W. A., W. R. Gardner, and W. H. Gardner, *Soil Physics*, 328 pp., John Wiley, New York, 1991.

Lippman, R. P., An introduction to computing with neural nets, *IEEE Acoust. Speech Signal Process Mag.*, *4*(2), 4–22, 1987.

Maidment, D. R., GIS and hydrologic modeling, in *Environmental Modeling with GIS*, edited by M. F. Goodchild, B. O. Parks, and L. T. Steyaert, pp. 147–167, Oxford Univ. Press, New York, 1993.

Maier, H. R., and G. C. Dandy, The use of artificial neural networks for the prediction of water quality parameters, *Water Resour. Res.*, *32*(4), 1013–1022, 1996.

Maren, A., C. Harston, and R. Pap, *Handbook of Neural Computing Applications*, 448 pp., Academic, San Diego, Calif., 1990.

Mark, D. M., and F. Csillag, The nature of boundaries on 'area-class' maps, *Cartographica*, *27*, 56–78, 1989.

Masters, T., *Practical Neural Network Recipes in C++*, 493 pp., Academic, San Diego, Calif., 1993.

McBratney, A. B., and J. J. De Gruijter, A continuum approach to soil classification by modified fuzzy *k*-means with extragrades, *J. Soil Sci.*, *43*, 159–175, 1992.

McBratney, A. B., and I. O. A. Odeh, Application of fuzzy sets in soil science: Fuzzy logic, fuzzy measurements and fuzzy decisions, *Geoderma*, *77*, 85–113, 1997.

McClelland, J. L., D. E. Rumelhart, and the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 2, MIT Press, Cambridge, Mass., 1986.

Morshed, J., and J. J. Kaluarachchi, Parameter estimation using artificial neural network and genetic algorithm for free-product migration and recovery, *Water Resour. Res.*, *34*(5), 1101–1113, 1998.

Nemani, R., L. Pierce, S. Running, and L. Band, Forest ecosystem processes at the watershed scale: Sensitivity to remotely sensed leaf area index estimates, *Int. J. Remote Sens.*, *14*, 2519–2534, 1993.

NeuralWare, Inc., *Neural Computing, NeuralWorks Professional II/Plus and NeuralWorks Explorer*, 360 pp., Pittsburgh, Pa., 1991.

Nimlos, J. T., *Soils of Lubrecht Experimental Forest, Misc. Publ.*, *44*, 36 pp., Mont. For. and Conserv. Exp. Stn., Sch. of For., Univ. of Mont., Missoula, 1986.

Odeh, I. O. A., A. B. McBratney, and D. J. Chittleborough, Soil pattern recognition with fuzzy-*c*-means: Application to classification and soil-landform interrelationships, *Soil Sci. Soc. Am. J.*, *56*, 505–516, 1992.

Pachepsky, Y. A., D. Timlin, and G. Varallyay, Artificial neural networks to estimate soil water retention from easily measurable data, *Geoderma*, *53*, 237–253, 1996.

Polak, E., *Computational Methods in Optimization*, 329 pp., Academic, San Diego, Calif., 1971.

Ranjithan, S., J. W. Eheart, and J. H. Garrett Jr., Neural network-based screening for groundwater reclamation under uncertainty, *Water Resour. Res.*, *29*(3), 563–574, 1993.

Rizzo, D. M., and D. E. Dougherty, Characterization of aquifer properties using artificial neural networks: Neural kriging, *Water Resour. Res.*, *30*(2), 483–497, 1994.

Rogers, L. L., and F. U. Dowla, Optimization of groundwater remediation using artificial neural networks with parallel solute transport modeling, *Water Resour. Res.*, *30*(2), 457–481, 1994.

Rumelhart, D. E., J. L. McClelland, and the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1, MIT Press, Cambridge, Mass., 1986.

Schaap, M. G., and W. Bouten, Modeling water retention curves of sandy soils using neural networks, *Water Resour. Res.*, *32*(10), 3033–3040, 1996.

Schaap, M. G., F. J. Leij, and M. T. van Genuchten, Neural network analysis for hierarchical prediction of soil hydraulic properties, *Soil Sci. Soc. Am. J.*, *62*, 847–855, 1998.

Shamseldin, A. Y., K. M. O'Connor, and G. C. Liang, Methods for combining the outputs of different rainfall-runoff models, *J. Hydrol.*, *197*, 203–229, 1997.

Simpson, P. K., *Artificial Neural Systems: Foundations, Paradigms, Applications, and Implementations*, 207 pp., Pergamon, Tarrytown, N. Y., 1990.

Tamari, S., J. H. M. Wosten, and J. C. Ruizsuarez, Testing an artificial neural network for predicting soil hydraulic conductivity, *Soil Sci. Soc. Am. J.*, *60*, 1732–1741, 1996.

Vemuri, V. (Ed.), *Artificial Neural Networks: Theoretical Concepts*, Comput. Soc. Press, New York, 1988.

Wen, C. G., and C. S. Lee, A neural network approach to multiobjective optimization for water quality management in a river basin, *Water Resour. Res.*, *34*(3), 427–436, 1998.

Zhu, A. X., A similarity model for representing soil spatial information, *Geoderma*, *77*, 217–242, 1997.

Zhu, A. X., Fuzzy inference of soil patterns: Implications for watershed modeling, in *Advanced Information Technologies for Assessing Nonpoint Source Pollution in the Vadose Zone, Geophys. Monogr. Ser.*, vol. 108, edited by D. L. Corwin, K. Loague, and T. R. Ellsworth, pp. 135–149, AGU, Washington, D. C., 1999a.

Zhu, A. X., A personal construct-based knowledge acquisition process for natural resource mapping, *Int. J. Geogr. Inf. Sci.*, *13*(2), 119–141, 1999b.

Zhu, A. X., and L. E. Band, A knowledge-based approach to data integration for soil mapping, *Can. J. Remote Sens.*, *20*(4), 408–418, 1994.

Zhu, A. X., L. E. Band, B. Dutton, T. J. Nimlos, Automated soil inference under fuzzy logic, *Ecol. Modell.*, *90*, 123–145, 1996.

Zhu, A. X., L. E. Band, R. Vertessy, and B. Dutton, Derivation of soil properties using a soil land inference model (SoLIM), *Soil Sci. Soc. Am. J.*, *61*, 523–533, 1997.

A.-X. Zhu, Department of Geography, University of Wisconsin-Madison, 384 Science Hall, 550 North Park Street, Madison, WI 53706-1491. (axing@geography.wisc.edu)

(Received December 14, 1998; revised October 4, 1999; accepted October 22, 1999.)