/N-19-CR
182469

19P

Spectral Analysis for Automated Exploration and Sample Acquisition

Susan Eberlein
Gigi Yates

May 1, 1992

NASA

Spectral Analysis for Automated Exploration and Sample Acquisition

Susan Eberlein
Gigi Yates

May 1, 1992

NASA

National Aeronautics and
Space Administration

Abstract

Future space exploration missions will rely heavily on the use of complex instrument data for determining the geologic, chemical and elemental character of planetary surfaces. One important instrument is the imaging spectrometer, which collects complete images in multiple discrete wavelengths in the visible and infrared regions of the spectrum. Extensive computational effort is required to extract information from such high-dimensional data. A hierarchical classification scheme allows multispectral data to be analyzed for purposes of mineral classification while limiting the overall computational requirements. The hierarchical classifier exploits the tunability of a new type of imaging spectrometer which is based on an acousto-optic tunable filter. This spectrometer collects a complete image in each wavelength passband without spatial scanning. It may be programmed to scan through a range of wavelengths or to collect only specific bands for data analysis. Spectral classification activities employ artificial neural networks, trained to recognize a number of mineral classes. Analysis of the trained networks has proven useful in determining which subsets of spectral bands should be employed at each step of the hierarchical classifier. The network classifiers are capable of recognizing all mineral types which were included in the training set. In addition, the major components of many mineral mixtures can also be recognized. This capability may prove useful for a system designed to evaluate data in a strange environment where details of the mineral composition are not known in advance.

Table of Contents

I. Background and Motivation

Planetary and Lunar exploration missions envisioned for the future rely heavily on the use of complex instrument data for determining the geologic, chemical and elemental character of the environment. Extensive computational effort is often required before the important information may be extracted from such data. The lengthy communication time and limited bandwidth involved in many exploration missions creates a need for autonomous onboard analysis of data. This need is accentuated in missions involving sample acquisition, where the acquisition process itself will produce a stream of data requiring real-time analysis in order to drive the activity.

In an effort to form a technology base supporting real-time, autonomous instrument data analysis, the Autonomous Exploration task and the Sample Acquisition, Analysis and Preservation (SAAP) task have been examining methods for spectral data analysis. The particular application is the interpretation of multiband imagery spanning the visible and short wave infrared regions of the spectrum to extract information on mineral composition. However, the methods developed for this application may also be revised for use with other types of spectral data, supporting autonomous extraction of chemical and elemental information from a variety of instruments.

The algorithms for spectral analysis have been developed concurrently with the development and laboratory testing of an imaging spectrometer instrument which is based on an acousto-optic tunable filter (AOTF). The potential availability of this instrument motivated consideration of data analysis algorithms that differed from traditional algorithms to exploit the instrument capabilities. The most important characteristic of the AOTF-based spectrometer is its tunability.

A standard imaging spectrometer uses a grating or prism to diffract or refract incoming light, causing spatial separation of the different wavelengths. The separated wavelengths are collected on a linear array of detectors, producing a spectrum for a single pixel encompassing the complete wavelength range of the instrument (Figure 1). To produce an image, the single-pixel focal point may be swept in two dimensions. Alternately, a linear array of gratings/prisms may focus a line of light onto a two dimensional detector array, by means of a sweeping motion in a single spatial dimension. The grating/prism-based spectrometers have the attributes of requiring spatial motion and always collecting a complete spectrum for each pixel.

The AOTF-based spectrometer passes light through a crystal while simultaneously passing an electrically generated acoustic wave through the crystal. The acoustic wave interferes with the light wave such that a narrow passband of light is diffracted while the residual light passes directly through (Figure 2). The light passing into the crystal may be an image rather than a point. The diffracted light maintains the spatial integrity of the image, and is collected on two-dimensional detector arrays. A spectrum is formed by changing the acoustic wave in the crystal, thereby changing the center of the passband for the diffracted light. The spectrum is collected one band at a time, with the complete image collected simultaneously in each band. There is no requirement for spatial motion, and the tunability of the filter allows the user to select only those spectral passbands of importance for the current application.

1

Spectral image processing algorithms have been developed to exploit the AOTF spectrometer tunability. These algorithms selectively acquire data in a few passbands, analyze that data, and use the results to select the next set of passbands for each region in the image. The resulting spectral classifier has a hierarchical structure, shown in Figure 3. The classifier algorithm applied at each step involves a trained neural network [1]. Below we describe the development of these classification approaches and the results obtained applying the approaches to a number of datasets.

II. Experimental Results - Hierarchical Classifier

The initial design and testing of the classification system used a simulated dataset, an image overlain with multispectral data. For each pixel, a single library spectrum or a mixture of two library spectra was assigned, Gaussian noise was added to the spectrum, and the spectrum was scaled in intensity to match the pixel brightness in the visible image. The library spectra and mixtures were chosen to simulate those that would result from the expected mineral composition of the surface of Mars. Thirty-two spectral bands in the wavelength range of 2.0 to 2.5 microns were chosen. There is good variability among different mineral classes in this range, which assists in distinguishing mineral types. This simulated dataset was useful for testing since the desired results of the classification were known for each pixel.

After initial tests on the simulated data proved successful, testing was performed on real datasets, collected by the Airborne Imaging Spectrometer (AIS) and Airborne Visible/Infrared Imaging Spectrometer (AVIRIS). These images consisted of 32 spectral bands ranging from 2.04 microns to 2.34 microns. Because the wavelength range did not match that of the simulated data, a second hierarchy was constructed for the new range, using the same protocol. Table 1 shows results of tests on the simulated dataset and one AIS image.

For each dataset, three classification schemes were compared (Table 1):

(1) A single-step matched filter, placing each input into one of 28 mineral classes based on all 32 spectral bands.

(2) A two-step matched filter, which first placed each input into one of seven classes, then matched within each class.

(3) A multi-step hierarchical classifier that used reduced dimensions to subdivide the input space followed by a matched filter to do final detailed classification.

Approaches 1 and 2 reflect methods that have been considered in the past for spectral analysis, utilizing all dimensions of the data (i.e. all spectral bands) for classification. The third approach utilizes only a few spectral bands at each step, and is the approach which exploits the tunability of the AOTF spectrometer. Table 1 shows the number of operations required to perform each activity. The accuracy of classification for the simulated dataset reflects the number of pixels placed into the correct geological class.

The two-step matched filter scheme places spectra into broad classes by matching against an average spectrum for each class. This approach causes a degradation of performance accuracy, although a significant savings in computation is realized. The poor performance results from the fact that major differences exist between the spectra of different minerals within a geological class, even though certain class specific spectral features are preserved. Using a matched filter classifier with averaged spectral memories causes each spectral band to be given equal weight in the classification. Regions of the spectra which are highly variable within the class contribute noise to the matching procedure, and can outweigh the contributions of the preserved spectral features.

The hierarchy of reduced-dimension classifiers overcomes this problem by using only the dimensions which are conserved within members of a class at each step in the hierarchy. It is this characteristic that allows the resulting classification to improve significantly over the classification produced by a single-step matched filter, while simultaneously reducing the computational requirements.

Accuracy was not determined for the real AIS dataset because the pixel by pixel mineral identity was not known. However, overall mineral composition in the larger regions was known, and the results of all classification methods were consistent at this broad level. Distribution of some mineral classes was compared to the distribution found by an alternate analysis method (SPAM - Spectral Analysis Manager) [2,3]. The alternate method produced the same results as were seen with the simulated dataset. The complete matched filter and pairwise averaged matchers differed very little in classification results. Accuracy degraded when the two-step matched filter was used, and improved considerably with the hierarchy of classifiers.

Table 2 shows a breakdown of the computational requirements for the classification hierarchy used for the simulated Mars dataset. This hierarchy is diagrammed in Figure 3. Since only three to six input dimensions were used for the first broad classification steps, very little computation was needed. This simple hierarchy reduces computational requirements by a factor of more than three compared to a single-step matched filter. In the future, more geological classes and mineral types will be incorporated into the hierarchy, and more steps will be added. While this expansion would dramatically increase computation were a matched filter used, the hierarchical approach should allow computation to be kept to a minimum.

One important benefit of the hierarchy of classifiers occurs in real applications where not all data needs to be classified. In an application where spectral data are being collected during planetary exploration, it is likely that some geological classes will be of great interest and others insignificant. The classifier tree allows one to incorporate a scientific goal into the classification. Only geological classes of current importance are classified in detail. Table 3 shows how specific types of goals will decrease the amount of computation required for the Mars dataset. In this application, finding clays or carbonates may be important because clays can be indicators of past water activity, carbonates indicators of fossil life. Other applications, such as searching for specific mineral deposits on earth or other planets, would invoke other goals.

3

III. Selecting the Reduced-Dimension Space

There is a significant body of work in the area of pattern recognition and classification supporting the contention that hierarchical classifiers are valuable both for reducing computation and for improving accuracy [5,6]. Similarly, it is accepted that careful selection of a subset of dimensions to be used for pattern classification will improve results, whether or not a hierarchical classification scheme is used. However, the selection of reduced dimensions, in this case specific spectral bands for each classification step, is a significant problem.

Segmentation of a vector space, such as the space of mineral spectra, into only two groups can often be achieved with a very small number of dimensions. The difficulty lies in determining the best dimensions to classify the vectors reliably. The following method for determining appropriate dimensions for vector space segmentation is applicable when the vectors are not randomly distributed, but fall into a number of classes. In this case, the classes represent geological groups of minerals (e.g. carbonates, borates).

A software simulation of a fully connected, feed-forward neural network was constructed [4]. This network consists of three layers, an input layer with number of nodes equal to the full dimensional pattern vector, a middle (hidden) layer with fewer nodes, and an output layer with nodes corresponding to the geological groups (Figure 4A). Each node in a layer is connected to each node in the next layer with a given connection weight.

Processing of the pattern vectors occurs by inputting a vector, and setting each hidden node value equal to the weighted sum of all the input nodes. The connection weight between a given input node and a given hidden node differs from those between the same input node and other hidden nodes. Since the input values are multiplied by these connection weights, the weighted sum of the values collected at a given hidden node differs from the sum at each other hidden node. The value at each output node similarly equals the weighted sum of all the hidden node values, where the weighting factor is the connection weight between a hidden node and an output node.

At each level, a sigmoid function is performed on the weighted sum to drive the value towards 0 or 1. When the net is properly designed and trained, a pattern may be input to the network to produce an output value of 1 for the appropriate geological group node, and a value of 0 for all other outputs. Although these networks may be constructed in hardware, they are readily simulated in software as vector-matrix multiplications.

The network weights are set by training the net. A set of known vectors which span the expected test space are input and the weights are altered until the desired output classes are obtained [1]. During the training procedure, the hidden nodes often come to act as feature detectors. They select from the higher dimensional inputs those features (maxima, minima, ratios of values) which are important in distinguishing the desired classes.

4

Analysis of the hidden nodes and their associated weights can assist in choosing the appropriate subset of dimensions that will allow the input vectors to be placed into certain classes [4,7,8]. For example, given ten output classes, a single hidden node may always be on when the input vector belongs to one of classes one through four, off for classes five through ten (Figure 4A). By finding which input nodes have very strong weights connecting to that hidden node, one can determine which dimensions are most important for distinguishing members of classes one through four from classes five through ten. A much smaller network may then be constructed (Figure 4B) which uses only those dimensions of the input vector to make a binary classification decision.

The result of these steps is a tree of classifier networks, each dedicated to making a simple segmentation (in this case always binary), and each working on a very few dimensions. The network architecture allows the classification to be made with computational requirements on the order of (I * H) + (H * O) where I is the number of input dimensions, H the number of hidden nodes, and O the number of output classes (usually two). Using a series of these networks, followed by a matched filter classifier for the final step, total computation may be reduced significantly.

IV. Application of Hierarchical Classification Scheme to AOTF Data

The hierarchical classifier scheme has been applied to a limited set of multiband images obtained from the AOTF-based imaging spectrometer. Unfortunately, the initial datasets have been confined to the visible region of the spectrum where differentiation among minerals and mineral classes is minimal. Several hierarchical classifiers have been developed which segment the images reasonably well, but the meaning of the classes is not clear.

Images are being acquired in the infrared region of the spectrum for use with a hierarchical classifier. The images will be of rocks rather than pure minerals, each rock potentially being composed of mixtures and alterations of several minerals. This dataset will provide an interesting test of classifier algorithms, and will no doubt uncover areas where more work will need to be performed.

Several steps are planned for the analysis of this data. The first step is the simple construction of a hierarchical classifier based on pure-mineral spectra, using the methods outlined above. The second step will incorporate findings (described below) which may improve the spectral "decomposition" of mineral mixtures. The third step will incorporate the use of textural information, should it become evident that differing textural character affects spectral reflectance.

Although the results of previous work with imagery obtained from airborne sensors indicate that neural-network based classification hierarchies are promising for the analysis of real data, much more work needs to be done to identify and deal with problems that do not arise in the classification of pure laboratory spectra. A series of experiments involving the classification of simulated and real mineral-mixture spectra is the first step in this work.

V. Spectral Mineral-Mixture Decomposition

The problem of robust, automated mixture decomposition is one of the most important issues in the area of imaging spectrometer data analysis. Mineral mixture data acquired by remote sensing techniques are generally analyzed using the assumption that mixing is linear. Linear mixing of spectra implies that the component spectra are added together proportionately to the amount of mineral present. This is a reasonable assumption to use for remotely sensed data, where large pixel sizes may encompass several discrete rocks, each of which contributes reflectance proportionately to its surface area within the pixel field of view.

When sensing is performed at very close range, the predominant form of mixing results from a physical intermingling of mineral particles at a scale comparable to or smaller than the wavelength of the light used for sensing. This is referred to as "intimate mixing" and can produce a non-linear spectral response in that one mineral component may contribute disproportionately to the overall spectral reflectance [9-11].

Much of the current effort in mixture decomposition work focuses on developing better algorithms for extracting the relative quantities of different mixture components, given that a human user inputs the component minerals. This approach is designed for non-real-time use by a user. It requires that the user have *a priori* knowledge of the minerals present at the test site, or is able to make an educated guess. Such an approach is not appropriate for use in an automated system used for planetary exploration; for this application, determination of component minerals must be automated as well.

Our current goal is to identify the major components of mineral mixtures without attempting to produce accurate quantitative estimates of mixture composition. These efforts apply equally to large scale and intimate mixtures. In the long term we hope to incorporate models of intimate mixing that will support accurate quantitative decomposition of both types of mixtures.

Trained neural networks have an advantage over many other types of pattern classification techniques for the spectral decomposition of mixture patterns. During the learning phase, the networks begin to extract the critical features from the patterns in each class, disregarding unimportant variations. In spectral data, the critical features are absorption features, where reflectance is minimal. Many mixtures retain significant feature information which can be used to extract major mineral components. Once likely mixture components are identified, it is a much easier task to model the quantitative composition based on either the linear or intimate mixing model.

The first step in the mixture identification process is recognizing that a spectrum may result from a mixture of minerals rather than a pure mineral. The method used to make this determination involves the application of an "entropy" measure [12]. The entropy measure is applied to the outputs of the neural network classifier to determine if a particular spectrum is well classified as a single known mineral. Those spectra which have a high entropy do not match well with any single mineral spectrum used to train the neural network. Such a spectrum may result from an unknown mineral or

from a mixture of minerals.

Using a test set of 180 synthetic two-component linear mixtures and a hierarchy trained by using only pure mineral spectra, 80% of the mixtures were flagged as being poorly classified based on the entropy measure. However, the hierarchy correctly identified at least one component of the mixture in 70% of the cases. In an effort to extract information regarding secondary mixture components, the first identified component was subtracted from the mixture spectrum (after scaling to represent 50% composition), and the remainder spectrum re-evaluated. This resulted in correct identification of both mixture components in 24% of the cases. Although this is a very preliminary result, it suggests that the combined use of neural network classifiers and subtractive analysis schemes may produce valuable results.

A small set of real intimate-mixture spectra was also tested using a single-stage trained classification network. The mineral composition of the materials causing these spectra was known precisely from separate analysis. The network was trained to recognize classes rather than single minerals. In all cases, at least one component of binary and trinary mixtures was recognized as the major class. In this set of tests, subtractive analysis was not performed, but the second-best output class proposed by the neural network was examined. In more than half the cases, the second-best class represented the second component of the mixture.

Although the spectral mixture-decomposition work is still in the early stages, the initial results suggest that neural network based classification systems may be effective for this problem. Future work will examine the relative merits of hierarchical and single-stage classifiers for mixture decomposition, and will consider approaches to extracting multiple components from the mixture spectrum. The goal is to develop a set of neural network based classification tools that will address the real needs of extraterrestrial mineral classification.

VI. Conclusion

A hierarchy of spectral pattern classifiers has proven to be a useful tool for reducing computation in a large pattern matching problem. With careful selection of the reduced dimensions for initial grouping of patterns, final classification accuracy may be increased as well. The hierarchical classification scheme fits well with the data acquisition scheme utilized by the AOTF-based imaging spectrometer, exploiting the spectral tunability of this instrument. Neural network based pattern matching has been utilized successfully within the hierarchical classifier. In addition, neural networks show promise when applied to the problem of spectral mineral-mixture decomposition; the networks are usually able to identify one or more of the major mineral components.

The combination of neural network processing units and hierarchical classification schemes has great potential for automated mineral classification as required by the Sample Acquisition, Analysis and Preservation task. It is quite possible that this classification approach will also be applicable for the analysis of other types of spectral data acquired for chemical and elemental characterization.

VII. References

1. Rumelhart, D., G. Hinton, and R. Williams (1986). Learning internal representations by error propagation. In D. Rumelhart and J. McClelland, Parallel Distributed Processing, pp. 318-362. MIT Press, Cambridge, MA.

2. Mazer, A., M. Martin, M. Lee, and J. Solomon (1987). Image processing software for imaging spectrometry. Proceedings of SPIE 31st Annual International Technical Symposium on Optical and Optoelectronic Applied Science and Engineering, SPIE Vol. 384, pp. 136-139, Ed. Gregg Vane.

3. Mazer, A., M. Martin, M. Lee, and J. Solomon (1988). Image processing software for imaging spectrometry data analysis. Remote Sensing of Environment, Vol. 24, No. 1, pp. 201-210.

4. Eberlein, S., and G. Yates (1990). Neural Network Based System for Autonomous Data Analysis and Control. Progress in Neural Networks, Volume 1, pp. 25-55.

5. Hauska, H. and P.E. Swain (1975). The decision tree classifier: design and potential. Proc. 2nd Symp. on Machine Processing of Remotely Sensed Data.

6. You, K. C. and K. S. Fu (1976). An approach to the design of a linear binary tree classifier. Proc. 3rd Symp. on Machine Processing of Remotely Sensed Data. pp. 3a-10.

7. Gorman, R., and T. Sejnowski (1988). Analysis of hidden units in a layered network trained to classify sonar targets. Neural Networks 1, pp. 75-89.

8. Sanger, D. (1989). Contribution analysis: a technique for assigning responsibilities to hidden units in connectionist networks. University of Colorado at Boulder Technical Report CU-CS-435-89.

9. Hapke, B. (1981). Bidirectional reflectance spectroscopy 1: Theory. J. Geophys. Res. 86, pp. 3039-3054.

10. Hapke, B. (1984). Bidirectional reflectance spectroscopy 3: Correction for macroscopic roughness. Icarus 59, pp. 41-59.

11. Hapke, B. (1986). Bidirectional reflectance spectroscopy 4: The extinction coefficient and the opposition effect. Icarus 67, pp. 264-287.

12. Eberlein, S., G. Yates, and E. Majani (1990). Spectral Image Classification in an Alien Environment. Proceedings of the Fourth Annual Symposium on Parallel Processing, Vol. 2, pp. 630-638. Sponsored by the IEEE Computer Society/California State University, Fullerton. Ed. L. H. Canter.

Figure 1. Traditional Imaging Spectrometer.
A single line of pixels is collected at one time, and each pixel is separated into the full range of wavelengths using a grating or prism. A full image is collected by sweeping in one spatial dimension.
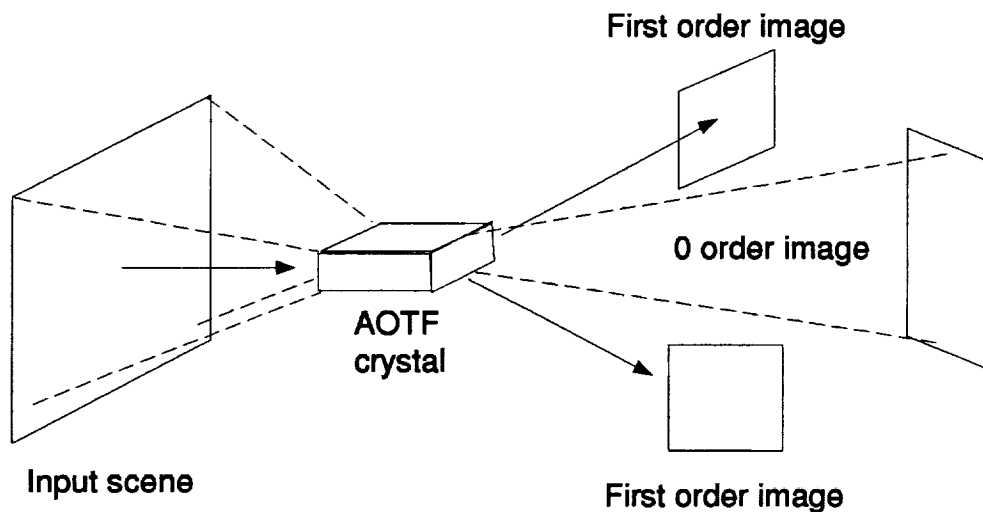

Figure 2. AOTF-Based Imaging Spectrometer.
A complete image scene in one passband is collected from the first order images. The crystal is tuned sequentially to different wavelengths.
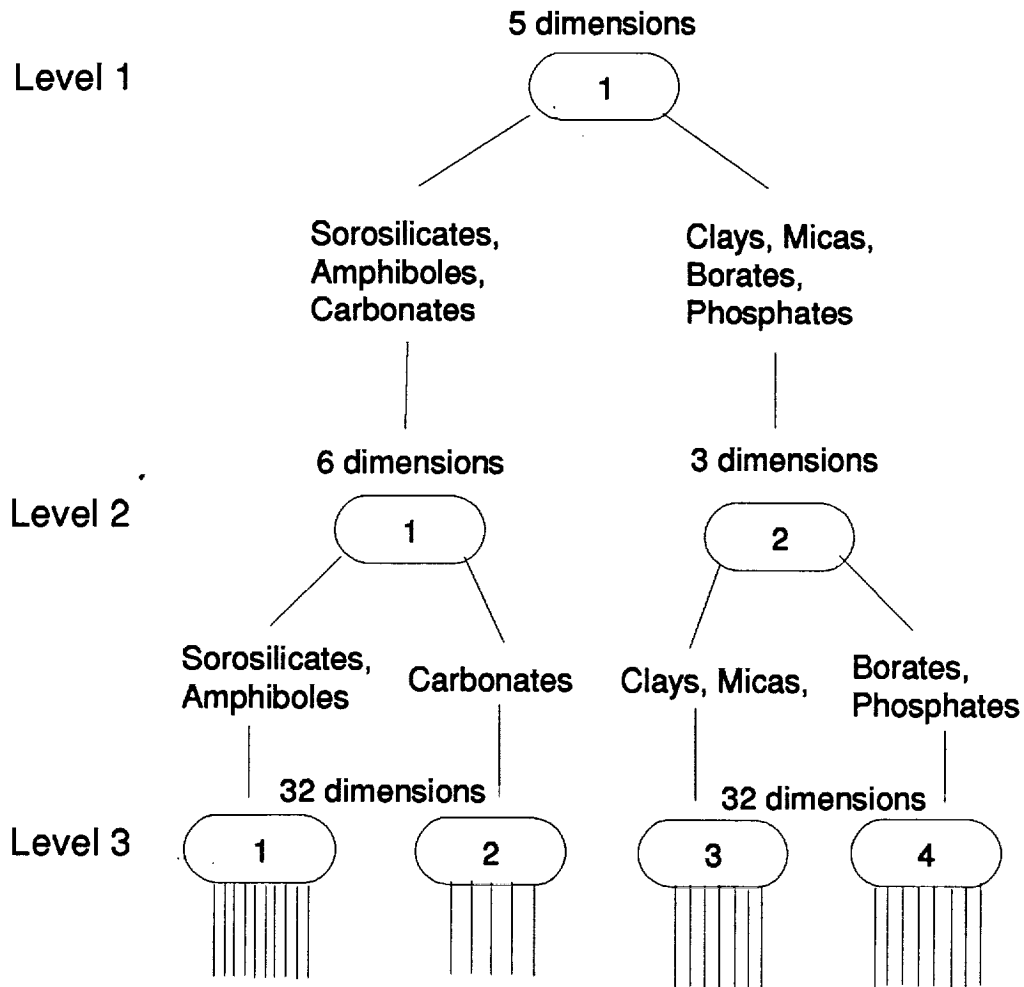
Figure 3. Three-Level Hierarchical Classifier.
Spectra from seven classes of minerals (sorosilicates, amphiboles, carbonates, clays, micas, borates and phosphates) are placed into progressively finer groups. A subset of the total 32 spectral bands (dimensions) is used at preliminary stages of classification.
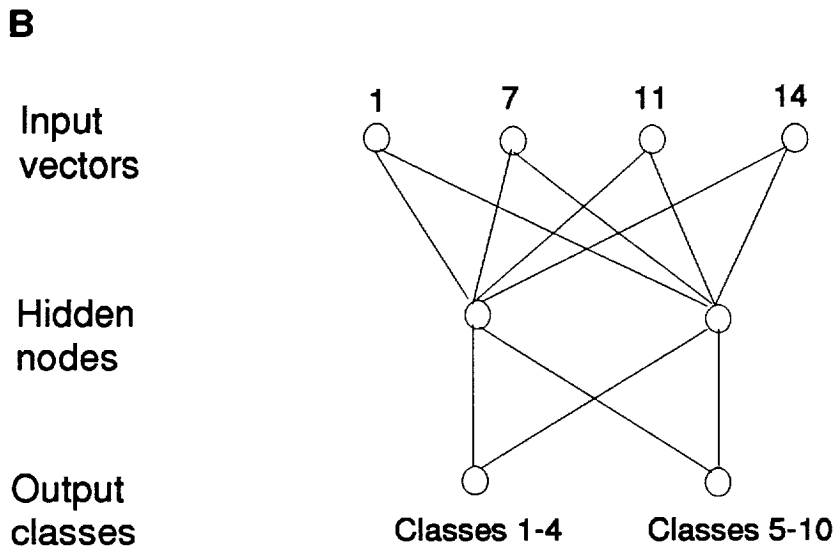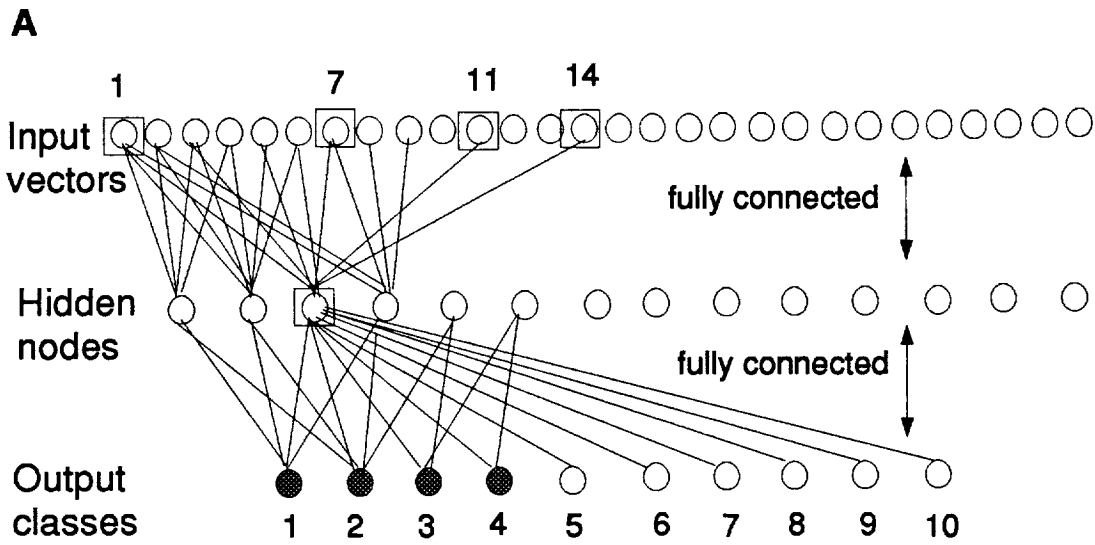
A

B

Figure 4. Neural network for mineral spectra classification.

A. A completely connected, feed-forward network has been trained on all input vector dimensions. Input vectors are placed into one of 10 classes. Analysis of the connection weights shows which input nodes are important for distinguishing output classes 1-4 from classes 5-10.

B. A small network uses only a subset of the total input dimensions to distinguish between two groups of classes.

Table 1. COMPARISON OF 3 CLASSIFICATION METHODS			
METHOD	DATASET	TOTAL OPERATIONS	ACCURACY
Single Matched Filter	Mars	16,226,560	80%
	AISA	5,017,600	
Two Step Matched Filter	Mars	6,374,720	69%
	AISA	1,971,712	
Hierarchy	Mars	4,858,284	89%
	AISA	1,006,099	

Table 2. NUMBER OF OPERATIONS AT EACH HIERARCHY LEVEL (MARS SIMULATED DATA)

LEVEL	NET	# OF SAMPLES	OPERATIONS PER SAMPLE	OPERATIONS PER NET	TOTAL OPERATIONS
1	1	18,110	(5*2 + 2*2) = 14	253,540	253,540
2	1	12,576	(6*3 + 3*2) = 24	301,824	
	2	5,534	(3*2 + 2*2) = 10	55,340	610,704
3	1	1,854	(32 * 8 memories) = 256	474,624	
	2	3,202	(32 * 8 memories) = 256	731,648	
	3	10,916	(32 * 8 memories) = 256	2,794,496	
	4	2,138	(32 * 4 memories) = 128	273,664	4,885,136

Table 3. NUMBER OF OPERATIONS WITH GOAL DRIVEN ANALYSIS (MARS SIMULATED DATA)				
GOAL	LEVELS	NETWORKS	OPERATIONS PER NET	TOTAL FOR GOAL
1) Find Clays	Level 1	1	253,540	
	Level 2	2	55,340	
	Level 3	3	2,794,496	1) 3,103,376
2) Find Carbonates	Level 1	1	253,540	
	Level 2	1	301,824	
	Level 3	2	731,648	2) 1,287,012
3) Classify Silicates	Level 1	1	253,540	
	Level 2	1,2	357,164	
	Level 3	1,3	3,269,120	3) 3,879,824
4) Classify All Areas	Level 1	1	253,540	
	Level 2	1,2	357,164	
	Level 3	1,2,3,4	4,274,432	4) 4,885,136