Methods for Detecting Anomalies in AVIRIS Imagery

Brian S. Penn
Mark W. Wolboldt
MASINT Resource Center
Boeing
1330 Inverness Drive, Ste. 330
Colorado Springs, Colorado 80910
Tel. (719) 572-8142
Email: Brian.S.Penn@Boeing.com

Introduction

At a very small scale the Earth's surface is heterogeneous, but the spatial resolution of current satellite and airborne sensors gives the appearance of homogeneity. This is an artifact of digital sensors that integrate the signal over well-defined instantaneous angular fields of view and spectral ranges. Because of limitations in spatial resolution, cursory examination of remotely-sensed imagery data leads to the conclusion there is a significant amount of redundant information, i.e., there are many similar pixels (picture elements) representing homogenous materials. Because so many pixels have similar spectra, finding anomalous pixels is a non-trivial problem.

Significant effort has been spent on classifying pixels in imagery. The goal of image classification is to group spectrally similar pixels together to the exclusion of anomalous pixels, i.e., separate the many from the few. Most classification methodologies do not easily lend themselves to the extraction of anomalous pixels. The reason for this is that these general classification algorithms address the relatively simple problem of analyzing the majority of the pixels. Nonetheless, it is useful to examine general classification approaches in the context of the method detailed here for anomaly detection.

The goal of anomaly detection is to find pixels that are significantly different from the majority of the pixels in an image, i.e., separate the few from the many. The definition of anomalous is rather arbitrary, but it can be expressed statistically, i.e., 2σ, or standard deviations, from the mean of a population with an assumed distribution, or merely as a proportion of the whole dataset, i.e., representing only $<<1$ % of the whole, etc.

Local anomaly detectors generally consider only the surrounding pixels to determine the nature of a pixel and are susceptible to noise, e.g., RX anomaly detection (Yu, Reed, and Stocker, 1993; Yu, et al., 1997). Global anomaly detectors evaluate each pixel in comparison to the entire image. The global approach generally alleviates problems associated with local variations in data values.

The purpose of this paper is to explore methods for finding anomalous pixels in AVIRIS imagery. Toward this end several approaches are described in detail beginning with the simplest and progressing to more complicated techniques. For completeness, a brief overview of finding anomalies using conventional supervised classification is also presented. Histograms are introduced first, followed by simple statistical measures, and concluding with the use of Self-Organizing Maps.

Location of Study Area

The Jet Propulsion Laboratory's Airborne Visible Infrared Imaging Spectrometer (AVIRIS) was flown over the Copper Flat (CF) porphyry copper deposit in the summer of 1998. The CF mine is 8 km NNE of the hamlet of Hillsboro in south-central New Mexico (Figure 1). The CF mine is in the Animas Hills and is part of the Hillsboro mining district of Western Sierra County, New Mexico. CF is one of the older Laramide porphyry-copper deposits in the Arizona-Sonora-New Mexico porphyry-copper belt (McLemore, et al., 1999). The Animas Hills consist of a horst block just west of the axis of the Rio Grande rift underlain by a circular body, nearly four miles in diameter, of andesite. Andesite is a fine-grained volcanic igneous rock extruded on the Earth's surface. The thickness of the andesite and circular shape suggest that the andesite is a remnant of a Cretaceous (144-65 million years ago) caldera or

collapsed volcano. A relatively small body of quartz monzonite intruded the andesite forming the entire CF deposit. Monzonite is a coarse-grained intrusive igneous rock. CF is predominately a low-grade hypogene (hydrothermal fluids carrying minerals up from below) deposit with thin veneer of supergene (meteoric water percolating from above) enrichment at the surface. The numerous latite dikes radiating from the CF quartz monzonite also show significant mineralization. The CF mine produced large quantities of gold, silver, and copper since the late 1800s.

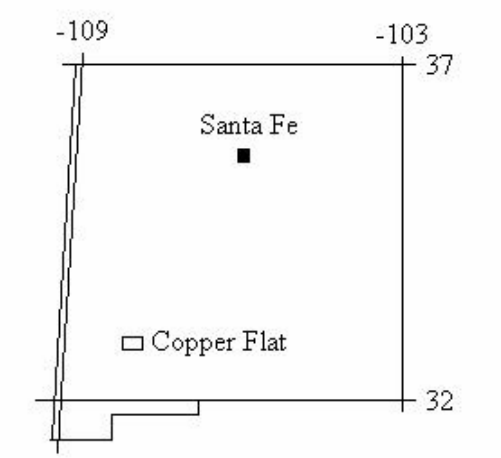

Figure 1 – Location of Copper Flat porphyry copper deposit in New Mexico.
The deposit is located 15 km west of the town of Truth or Consequences.

Gypsum ($CaSO_4 * 2(H_2O)$) is an anomalous mineral occurring only in close proximity to the pit filled with water at the center of the mine and along some of the streambeds in the Hillsboro district. Gypsum occurs as coatings of precipitate on the rocks surrounding the pit and comes from the pit waters. As such, gypsum was selected as the target to test the anomaly detection algorithms.

Finding Anomalies Using Supervised Classification Techniques

In general, supervised image classification is the process by which a classifier is presented with pixels of "known" contents that are used as a guide(s) for classifying other pixels in an image. The classifier function is "seeded" with *a priori* information about particular material types. The classification algorithm then determines which pixels are most like the "seed" pixel based on a similarity condition. Examples of supervised classification techniques include minimum distance, parallelepiped, mahalanobis distance and maximum likelihood. Set membership for these supervised classification techniques is based on Euclidian distance, standard deviation values, etc. A major assumption for these types of classifiers is that the data is linearly separable, i.e., the data is arranged such that a line (or plane or hyperplane) can be situated between different data classes.

Using *a priori* information for target identification is a useful but awkward method of finding anomalies because it is not based on set membership, but the logical negation of set membership. The premise of this approach is to classify every known spectrum and tag as anomalous any pixel that is not recognized.

Using a supervised classification approach to identify anomalies requires possessing a spectral library. Assuming the spectral library is well-populated, any pixel's spectral absorption features not found in the library are classified as anomalous. The U.S.G.S.'s Tetracorder operates in this manner. Tetracorder contains an extensive library of mineral spectra. If Tetracorder is unable to classify a pixel's spectrum, the classification is left blank. This pixel is then considered anomalous.

Unfortunately, this approach is has shortcomings. For example: if the number of material types within an image are not correctly identified prior to processing, then there can be a large numbers of pixels omitted because a class was not properly identified. The result is significant uncertainty in

determining whether a pixel is anomalous or merely incorrectly selected as a member of a larger group of pixels in the image. It is for this reason that unsupervised classification is more appropriate for anomaly detection.

Unsupervised Anomaly Detection

Global Histogram Technique

Histograms are graphs describing the frequency of occurrence for members of a dataset. In the context of imagery, a histogram graphs the digital number on the abscissa and the frequency on the ordinate axis. Histograms are a central concept in remote sensing because they enable image analysts to compress the information in an image into a simple format. One significant use for histograms is to enhance digital images through the use of various contrast stretches, i.e., linear, 2%, Gaussian, square root, etc. Often, as is the case with the linear and 2% stretches, the digital numbers (DNs) with the highest and lowest values are linearly translated to the largest and smallest radiometric values, i.e., for 8-bit data the lowest value is set to 0 and the highest value is translated to 255 with the remaining values linearly mapped between these two values.

For the purpose of anomaly detection, histograms are used in a slightly different manner. The two tails of the histogram are the extremes of the imagery DN population with the lower end representing the darkest pixels and the upper tail representing the brightest pixels. Pixels that repeatedly occur at the tails of histograms represent anomalies. A simple procedure is used to identify these pixels.

First a threshold is arbitrarily selected to mark a pixel as anomalous (< 2%). Next, a histogram for each band is generated and each pixel falling above or below a population threshold is recorded. The result is a two-dimensional array corresponding to the image with counts for high and low ends of the histogram. These images are then further constrained to show the highest number of counts for each tail of the histogram and a color-coded image is generated (Figure 2). In the CF imagery, there are known anomalous areas of gypsum near the mine pit. Figure 3 shows results of applying the histogram anomaly detection for the pit and surrounding pixels with an accompanying average spectra for surface materials adjacent to pit lake.

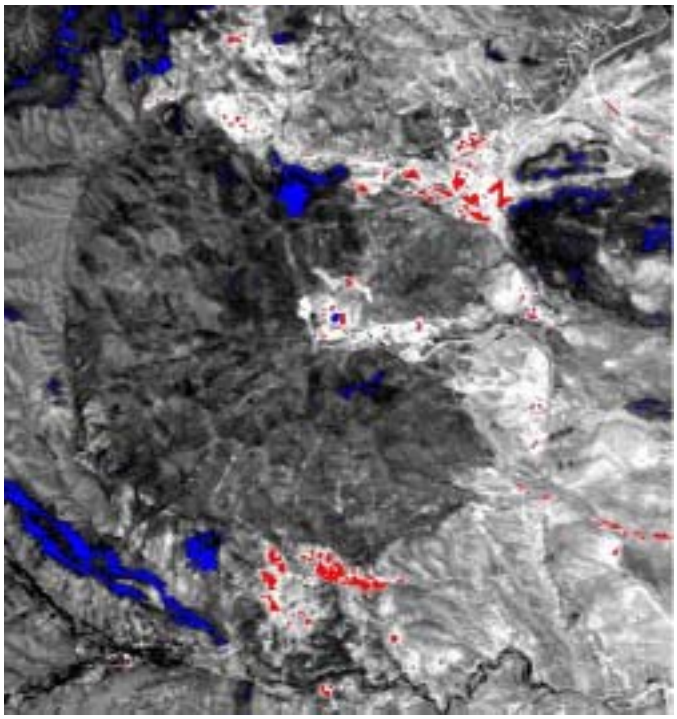

Figure 2 – Results of applying global histogram anomaly detector to AVIRIS imagery data. Red pixels are anomalously bright, blue pixels are anomalously dark. Water in the pit is blue, gypsum is found in pixels adjacent to pit lake.

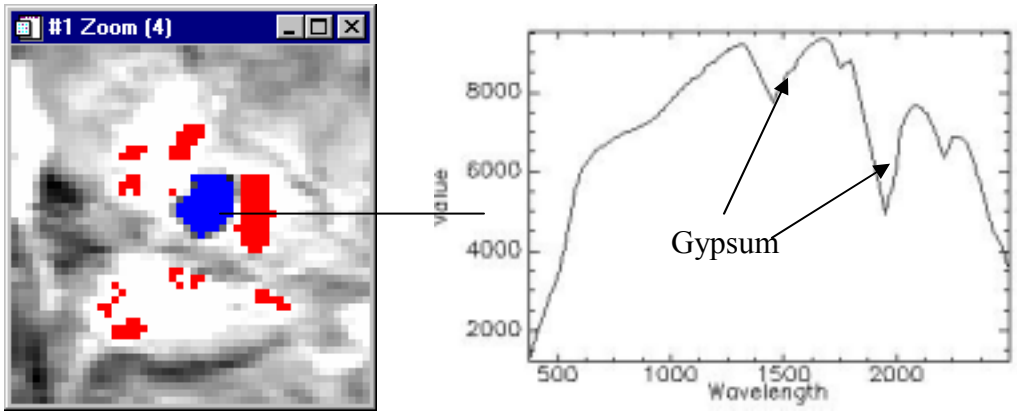

Figure 3 – Image of mine pit with average spectra for area of red pixels to the right of lake in middle. The absorption feature at 2208 nm results from overtones of the sulfate ion restrahlen bands at longer wavelengths. Other red pixels are anomalously bright, but contain no gypsum. Red pixels are anomalously bright, blue pixels are anomalously dark.

Simple Statistical Techniques

A common technique for finding anomalies in a dataset is to examine the data from the perspective of simple statistics. Most familiar statistical models are based on Gaussian or normal distribution, i.e., a bell-shaped curve. The assumptions for a normal distribution are: 1) most samples for any particular dataset are evenly distributed about the mean, or average, for a dataset with proportionately smaller numbers of members occurring progressively farther from the mean; and 2) the mode (greatest number of samples) and the mean for a dataset are the same. This is the basis for variance and standard deviation (σ) calculations. If the assumption of a normal distribution is not true, then basic statistical assumptions may not be valid. Nevertheless, the assumption is usually made that the data have a normal distribution and simple statistics such as variance and standard deviation are applicable. This appears to be a reasonable assumption for satellite and airborne images of the earth.

When working with standard deviation calculations it is necessary to know how much of the area under the curve (percentage of the sample population) is encompassed by various standard deviations. Values occurring $\pm 1\sigma$ from the mean include 68.3% of all the samples; $\pm 2\sigma$ and $\pm 3\sigma$ include 95.4 % and 99.7%; respectively. If an anomalous pixel is described as 1% of a population, then it should have a value that deviates more than $\pm 2\sigma$ from the mean.

Using this approach, the difference between pixels with values less than 2σ below the mean and those values greater than 2σ above the mean are separated. The result of applying this approach to imagery is shown in Figures 4 and 5.

Local Anomaly Detectors

Both the histogram and stand deviation techniques can be applied to an entire image. These approaches distinguish global anomalies, but some features can be globally typical yet locally anomalous. The way to address these types of anomalies is to constrain comparisons between each pixel and its surrounding neighborhood. The process of comparison is achieved by using the standard deviation criteria outlined above to the local environment around the pixel of concern. The results of this approach are shown in Figure 6. In general the larger anomalies were found by differing kernel sizes, but as expected anomalies are scale dependent. Also, the larger the kernel size, the greater the resulting anomalies are smoothed.

One of the reasons for color-coding the results in the previous examples is that there are anomalous bright and dark pixels. Another approach for comparing pixels is to treat spectra as vectors in *n*-space and use the Cauchy-Schwarz Inequality (1) to determine the angular distance between spectra. This approach has the additional benefit of ignoring differences in illumination

$$\cos\theta = \frac{\langle u, v \rangle}{\|u\| \, \|v\|} \tag{1}$$

Once again kernels of differing sizes are used to adjust compare to the center pixel. Figure 6 contains the results of this processing. Figures 6 and 7 show a significant decrease in the number of pixels considered anomalous with increasing kernel size.

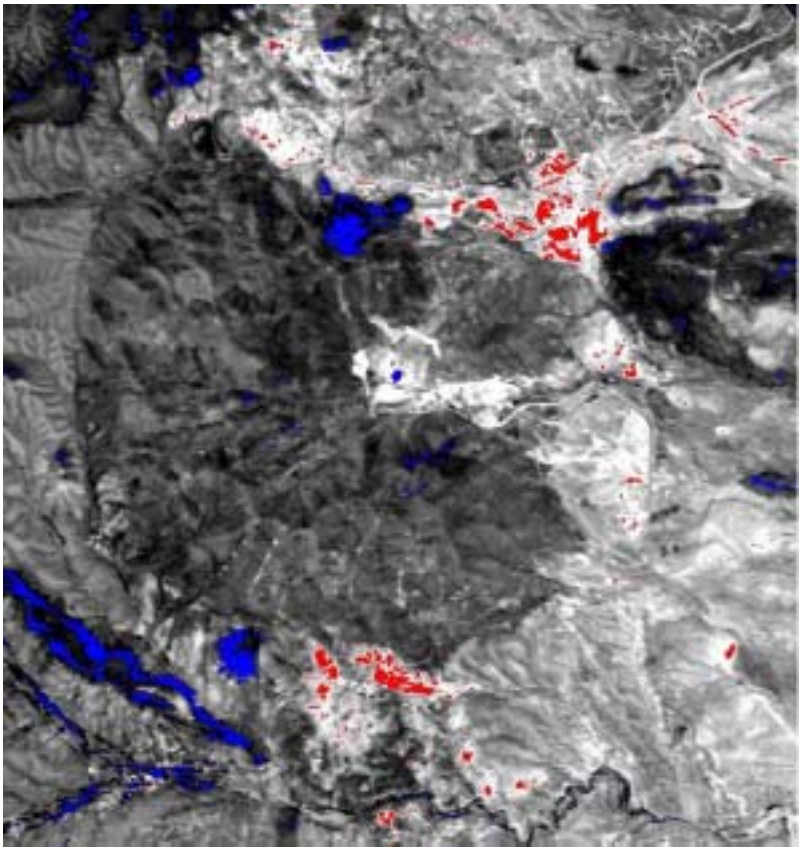

Figure 4 – Output results from standard deviation with threshold at 2.5 σ. Red pixels are those with values > 2.5 σ; blue pixels have values < -2.5 σ.

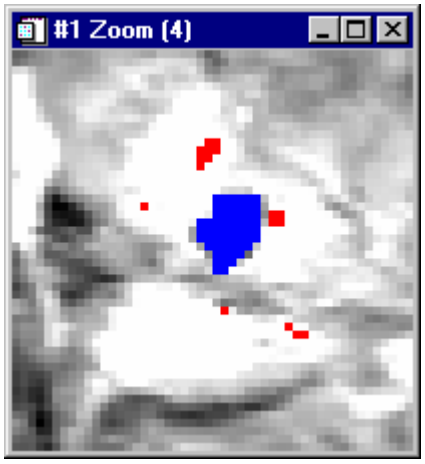

Figure 5 – Magnified image of mine pit lake (blue pixels) with surrounding red pixels.

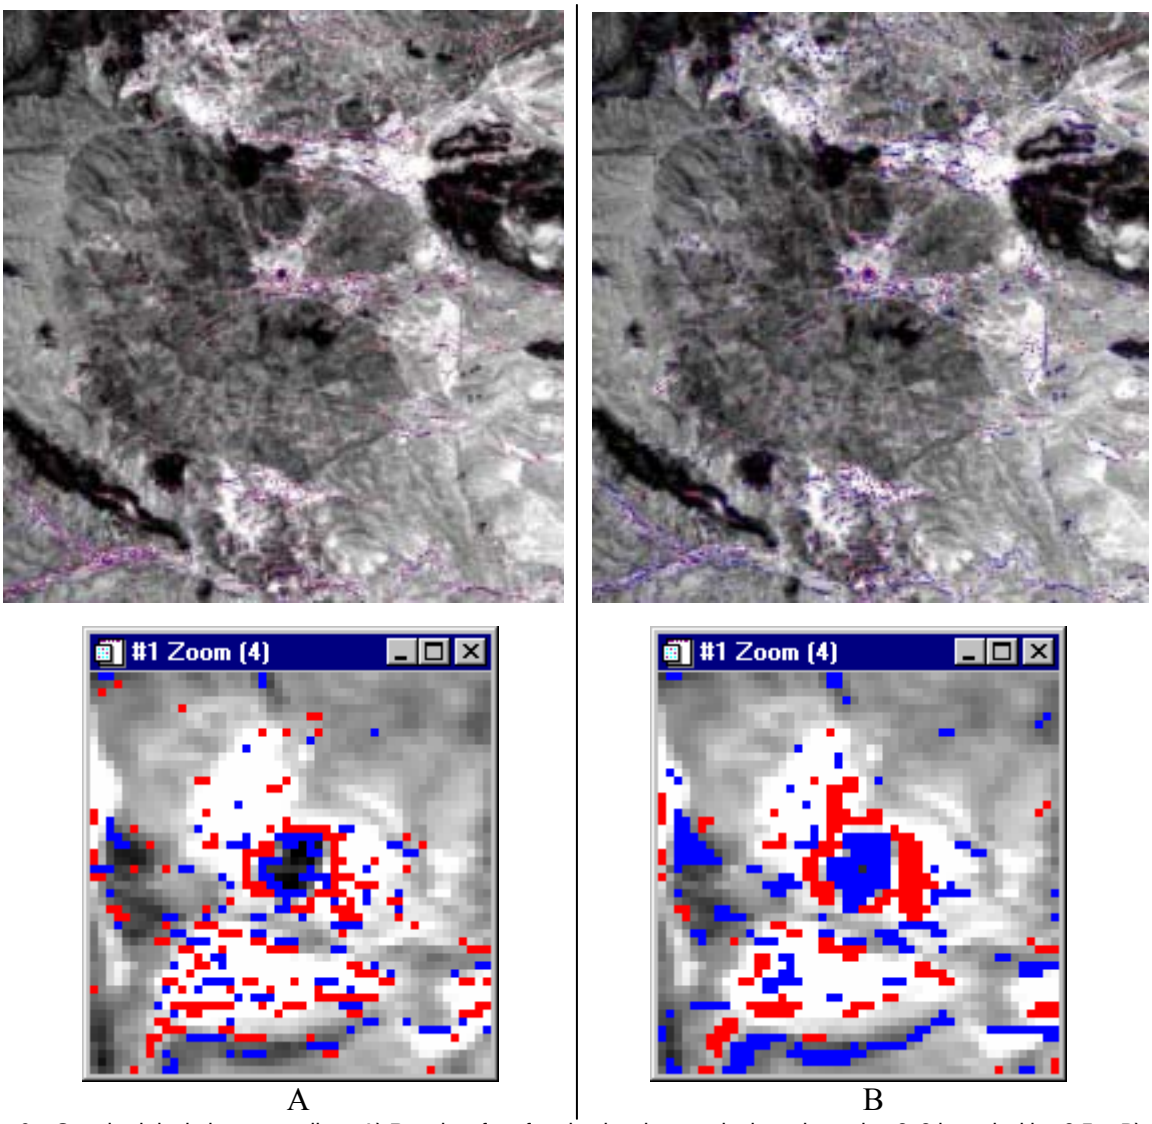

Figure 6 – Standard deviation anomalies. A) Results of performing local anomaly detection using 3x3 kernel with ±2.5σ; B) Results of performing local anomaly detection using 7x7 kernel with ±2.5σ. This approach involves looking at each band and accumulating standard deviation information. Another approach to local anomaly detection uses differences in spectral angle to compare adjacent pixels.

Clustering Techniques

The most common form of unsupervised classification utilizes clustering techniques. Clustering is a non-statistical method for grouping similar data. Clustering methods complement supervised classification. Most unsupervised classification techniques evaluate the entire image data set and assign each pixel to a cluster. There are generally no predetermined limits on the number of clusters. A benefit of this approach is the reduced likelihood of missing a coherent group in the classification process.

Two of the most common unsupervised classification methods are the Iterative Self-Organizing Data Analysis Technique (ISODATA) (Tou and Gonzalez, 1974) and K-Means. ISODATA iteratively evaluates imagery data based on spectral distance from predefined nodes. These nodes are treated as cluster centers and pixels are included in the clusters based on user specified thresholds defined primarily by Euclidian distance or standard deviation values. After each iteration, a new cluster mean is computed based on the actual spectral locations of pixels. These new means are used as the basis for the subsequent iterations. The process continues until change between cluster means falls below a set threshold value.

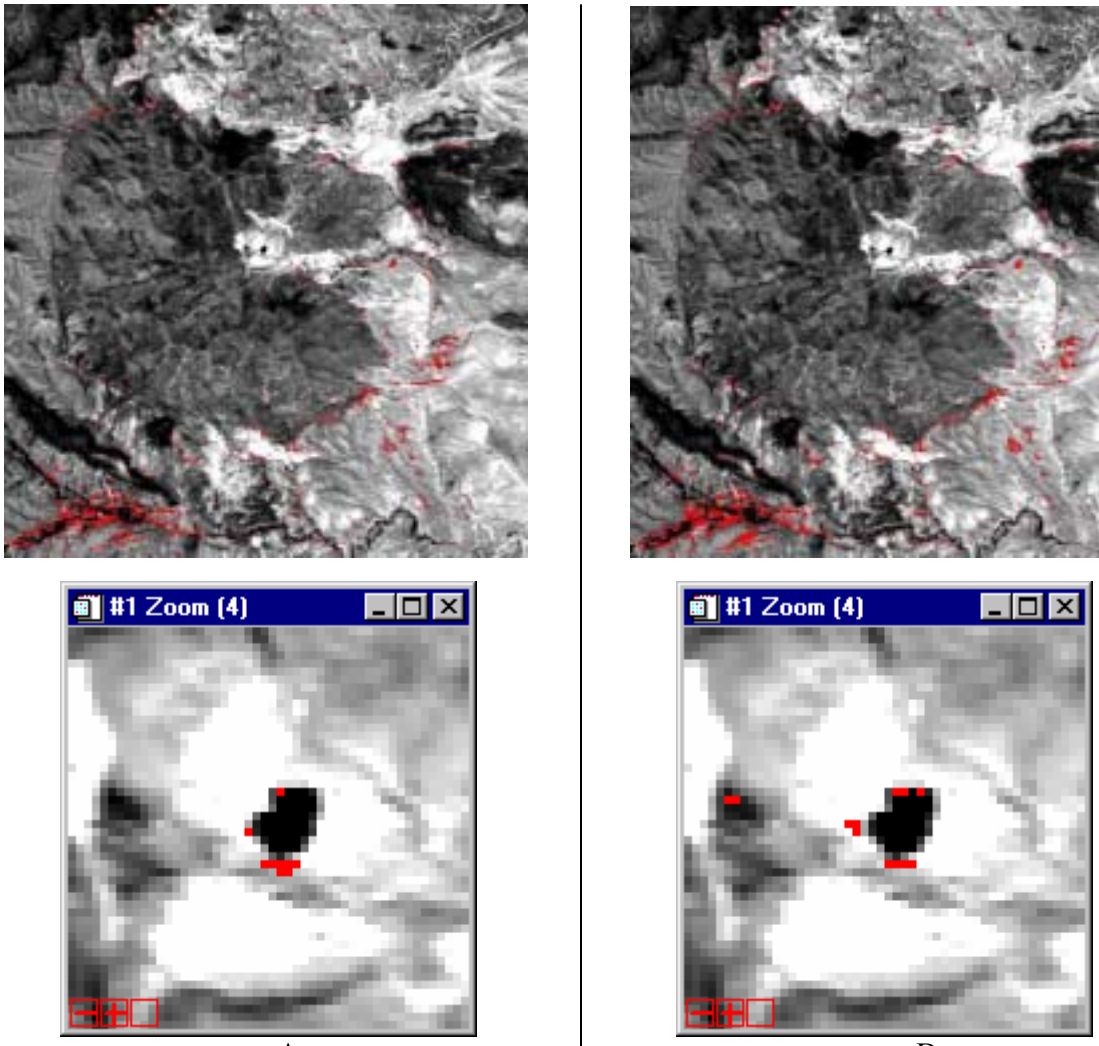

Figure 7 – Spectral angles used to compare pixel with surrounding neighbors. A) 3x3 kernel thresholded to 1%; B) 7x7 kernel thresholded to 1%. Pixels at the bottom of zoom images correspond to gypsum precipitates.

K-Means calculates initial class means evenly distributed throughout the data space. The process then assigns each pixel to the closest cluster. Once all the assignments are made, new cluster means are calculated, followed by the assignment of all pixels to the nearest cluster mean. This process is repeated until the number of user-defined changes fails to occur.

Another clustering technique based on neural networks is Kohonen Self-Organizing Maps (Kohonen, 1986). A Self-Organizing Map (SOM) is an array (1 or more dimensions) of nodes. Each node is composed of a unit vector pointing in a random direction in *n*-dimensional space. After in-band normalization, multi-dimensional data are presented to each of the individual nodes (Figure 8).

Using a "winner-take-all" learning strategy, the node whose vector most closely matches the input data is found. This winning vector incorporates, or adjusts, its vector coefficients (weights) to match the input data. As part of the learning strategy, vectors in the nodes nearest the winning node are modified to look less like the input vector. In this manner, each node in the SOM internally develops the ability to recognize vectors similar to itself. This technique is a form of self-organization, i.e., no external information is supplied to lead to a classification.

Two useful features of a SOM are its topology preserving capability and the automatic generation of probabilities for a dataset. Topology preserving means that the original relationships between the data

points remain intact after processing. This is the desired result when working with hyperspectral data because it maintains the spectral relationships between pixels. Secondly, as a SOM evaluates the data it spontaneously builds a statistical model, or probability distribution, of the dataset. SOMs perform this statistical modeling, even in cases when no closed-form analytic expression can describe the distribution (Caudill, 1988). The SOM approach is similar to ISODATA, but it operates in a much more flexible manner. Cluster centers move around dynamically to account for topological relations in the dataset. Output from the SOM is a series of spectra containing the mean spectra for each cluster group.

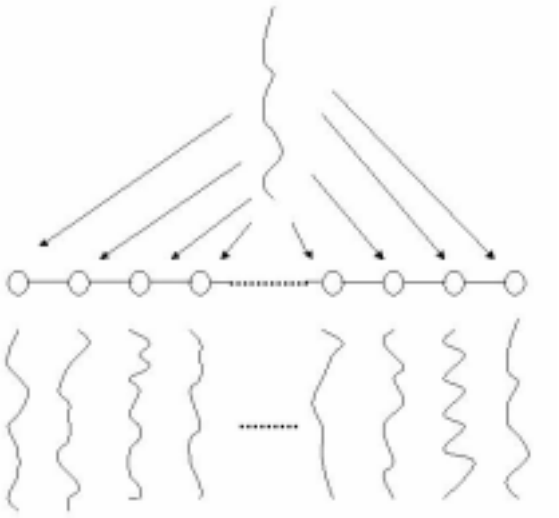

Figure 8 – Kohonen SOM presented with single spectrum input and the resulting spectra developed while processing various input spectra.

During SOM classification the cluster centers constantly adapt themselves to accommodate the spectra presented to them. The end result is a cluster center point representing an average spectrum for the spectra presented to the SOM for classification. The cluster centers are discretely located in *n*-space such that they can be projected into 2-d space (Figure 9). The actual data points lie on the *n*-plane at a distance from the cluster centers forming a cloud or convex hull (or zonatope enclosure) around the cluster centers. Those spectra closest to the cluster center are those that are most abundant while those located farthest away are the least abundant.

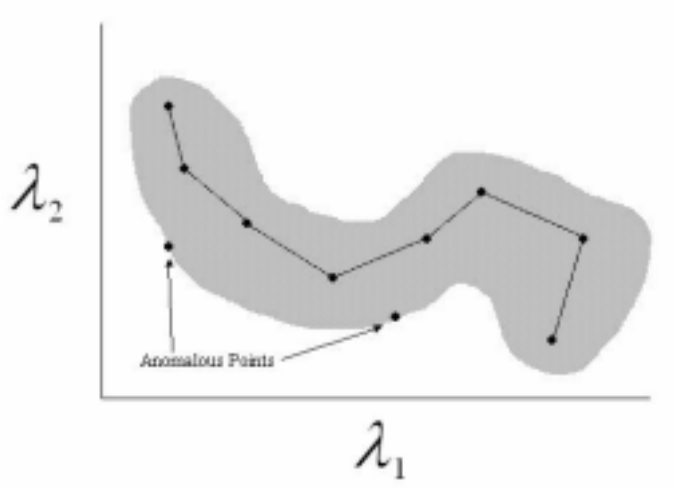

Figure 9 – SOM output for unknown dataset projected into wavelength axes λ_1 and λ_2. Data is arranged around the cluster points in a cloud. Anomalous points lie on the edge of the gray-shaded area indicating least amount of similarity to cluster points. Lines between cluster centers added for emphasis.

After the SOM is built, the spectra are again presented to the SOM for classification. A *2-d* map is generated with the distance from the nearest cluster center stored in the location for each pixel. A density slice of the *2-d* map is then created and those points farthest from a cluster center are identified. Those pixels with the largest values, i.e., greatest distance from a cluster center, are the anomalous pixels.

Data Compression via Minimum Noise Fraction

AVIRIS imagery contains a huge amount of data. A typical AVIRIS data cube contains 614 pixels, 512 lines, and 224 channels totaling slightly more than 170 Megabytes of data. This amount of data presents a significant computational challenge even to the fastest computer. It is possible to compress this amount of data into a manageable form by employing a Minimum Noise Fraction (MNF) transform (Boardman and Kruse, 1994). An MNF transform is a cascade of manipulations including a translation, rotation, scaling and another rotation (essentially an affine transform) of the data.

Stated another way, an MNF is a series of two cascading principal component (PC) transformations resulting in a change of basis for the vector space as defined by the number of bands for each pixel in an image. The high band-to-band correlation inherent in the original data is what makes the MNF so efficient. The net effect of the MNF transform is to compress the variation within the dataset from 224 bands down to 10-20 bands. Reducing the size of number of bands presented to a SOM for classification significantly increases its sped of operation. The results of applying a SOM to an MNF image to find anomalies is shown in Figure 10.

Alternate Hyperspectral Approaches for Anomaly Detection

The most widely known anomaly detection algorithm for hyperspectral data is RX (Reed and Yu, 1990; Yu, Reed, and Stocker, 1993; Yu et al., 1997). The RX algorithm is a variant of the algorithm for generating the Mahalanobis distance and is considered a local anomaly detector. Another way to view the RX algorithm is as the complement of a Principle Components Analysis (PCA), i.e., instead of performing a basis change to maximize the variance into major components emphasis is placed on the finding targets in the minor components. Since SOMs are applied to Minimum Noise Fraction (MNF) bands to generate statistical models, there are significant similarities between these techniques.

Discussion

For hyperspectral imagery, variations of unsupervised classification are possible because methods exist for extracting relevant end-members automatically from imagery based on spectral unmixing (Smith et al., 1985; Barnard and Casasent, 1989; Smith et al., 1990; Roberts et al., 1993; Settle and Drake, 1993; Foody and Cox, 1994; Van der Meer and de Jong, 2000; Penn, 2002). The basic assumption of spectral unmixing is the existence of "pure" pixels (end-members) from which all other pixels in an image can be derived by linearly combining the end-member pixels. While this statement is accurate to a first approximation, the uniqueness of end-member pixels is a function of the method of deriving end-members. Inasmuch as the uniqueness assumption is valid then, based on the method of identifying end-members, there is no guarantee that end-members are homogenous pixels for reasons cited previously. In all probability each end-member pixel is a combination of materials producing a relatively unique integrated signature at the particular resolution of the sensor.

Several tools exist for automatically finding end-members in hyperspectral imagery datasets including N-FINDR (Winter, 1999), ORASIS (Bowles et al., 1995), and Aspect-Sentinel and OSP (Harsanyi, 1993; Harsanyi and Chang, 1994). These programs generally transform the data into various *n*-spaces and fit a convex hull around the dataset. The apices of the hull are treated as the end-members from which all other spectra in the data volume are derived. The logical method for detecting anomalies using this approach is to determine which pixels appear not to fit within the hypervolume generated by the algorithms. These approaches, while automatic, can generate numerous false-alarms especially with the presence of anomalous pixels.

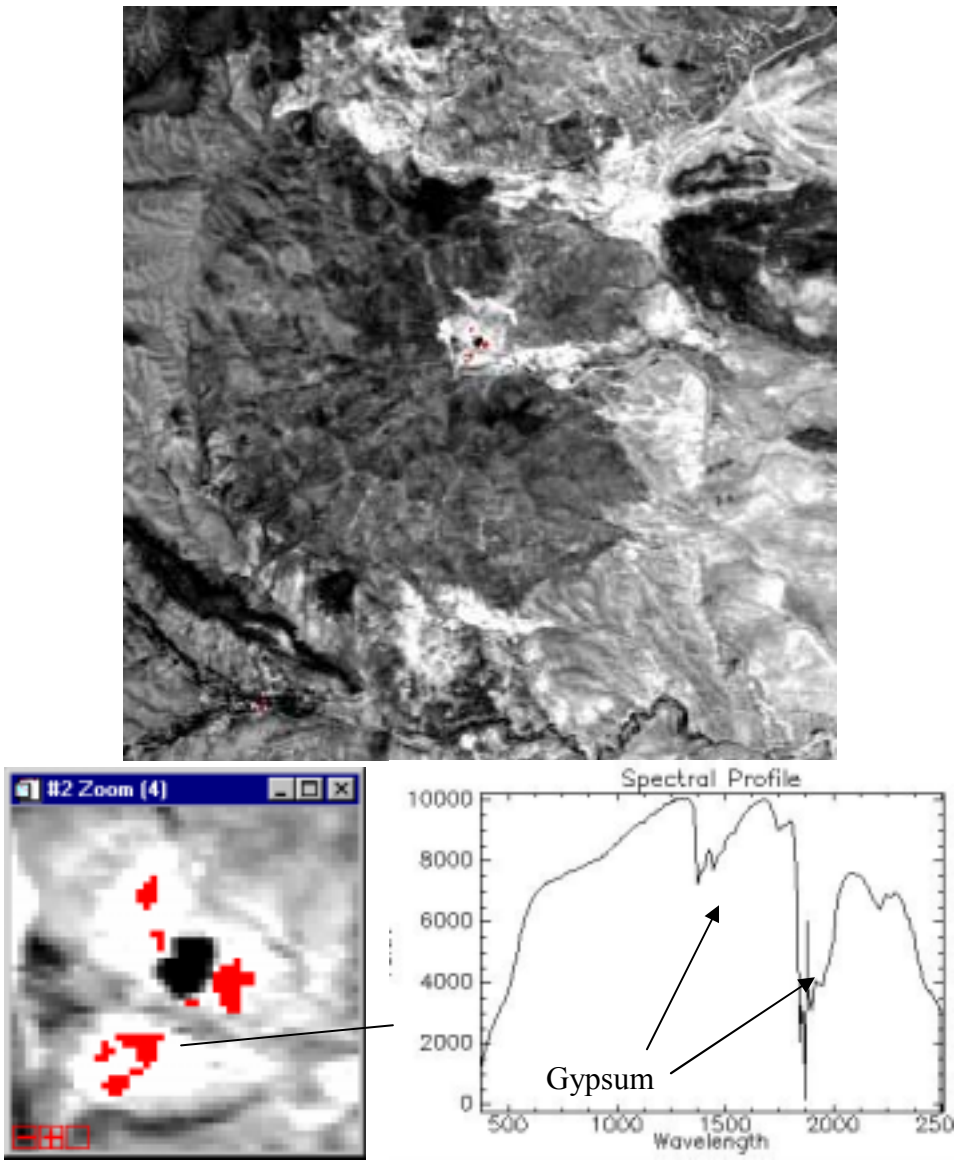

Figure 10 – Results of applying SOM to Copper Flat Imagery. The threshold is set to 0.01% of the output. The most anomalous pixels are found adjacent to the pit lake.

Validation of Anomalies

Once anomalies are detected their uniqueness can be validated using supervised classification methods. Two such methods include Constrained Energy Minimization (CEM) and Spectral Angle Mapper (SAM). This verification process is accomplished by providing proposed anomalous spectra to both CEM and SAM. This is necessary because both techniques require *a priori* information to find targets. An average spectrum from the identified anomalous pixels was provided to both algorithms. CEM and SAM then analyze the image to determine the presence of particular mineral spectra. Both of these approaches corroborate the effectiveness of using SOMs to find anomalous materials (Figure 11).

Conclusions

Several methods for anomaly detection including using Self-Organizing Maps (SOMs) were applied to the AVIRIS imagery of the Copper Flat porphyry copper deposit to extract anomalies. Copper Flat

was selected because it has been mapped extensively and there are several known anomalous occurrences of gypsum ($CaSO_4*2(H_2O)$) near the pit lake in the center of the mine.

The methods employed include both global and local methods of anomaly detection. These methods included using histograms and standard statistical approaches for global analysis. Local anomaly detection methods included standard deviation and spectral angle calculations to ascertain whether a pixel was anomalous compared to neighboring pixels. Each of these approaches was applied to atmospherically corrected data.

The final method used to look for anomalies involved a Self-Organizing Map (SOM). The results from the SOM were slightly more accurate with less clutter as compared to the other approaches. The improved accuracy may have resulted from applying the SOM to data that had a Minimum Noise Transform applied before processing.

The practical use of SOMs to detect anomalies is clearly demonstrated and has been verified by fieldwork. This technique is relatively simple and yields good results. The one drawback is that SOMs are not able to separate out different anomalies. A number of anomalous pixels found using SOMs did not contain any measurable gypsum instead they were dominated by other minerals. The best practical approach is to iteratively use SOMs to locate potential anomalies and use the resulting spectra to classify the image. In this manner the desired anomalies could be found and less desirable anomalies could be eliminated from the processing.

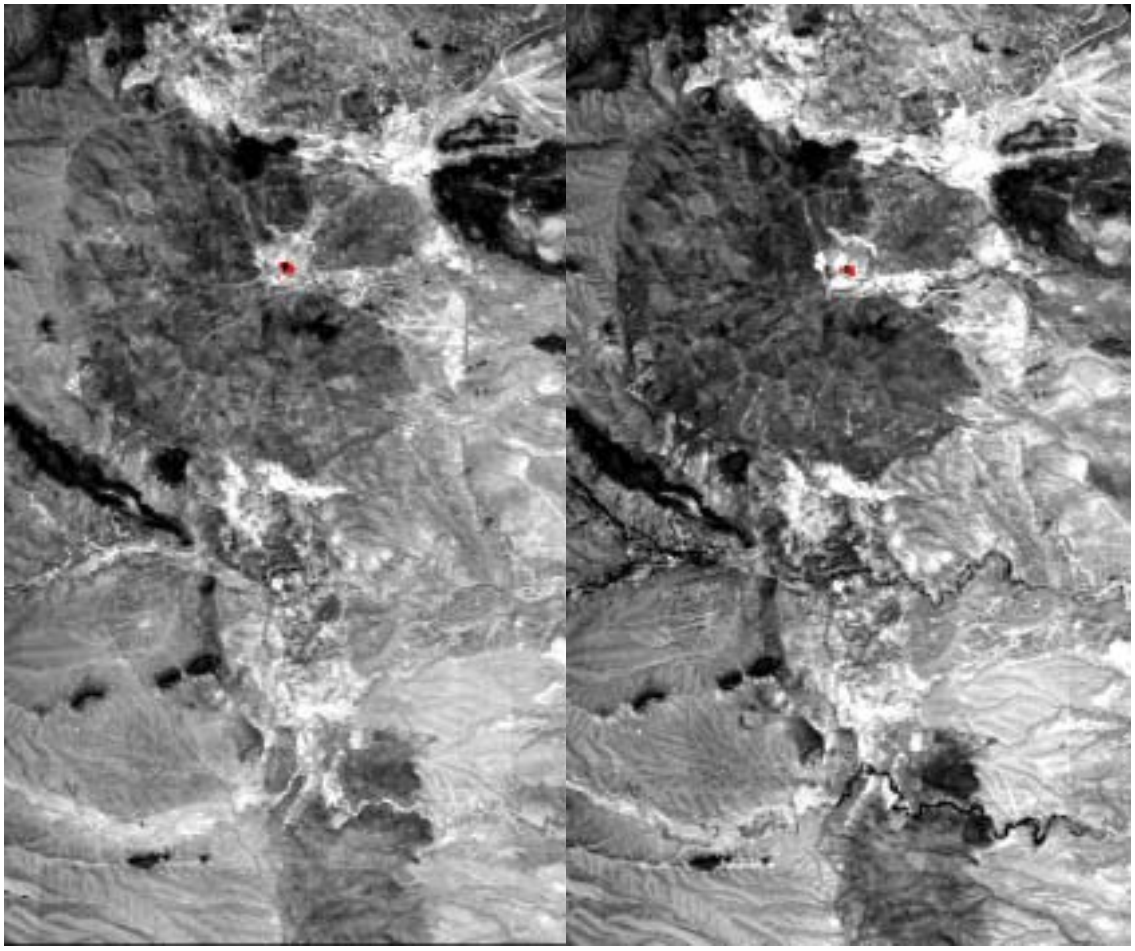

a b

Figure 11 – Results of applying spectral signatures of anomalous pixels using:
a) Constrained Energy Minimization (CEM) and b) Spectral Angle Mapper (SAM).

References

Barnard, E., and Casasent, D. P., 1989. Optical Neural Network for classifying imaging spectrometer data, Applied Optics 28(15), 3129-3133.

Boardman, J.W. and Kruse, F. A., 1994. Automated spectral analysis: a geological example using AVIRIS data, north Grapevine Mountains, Nevada inProceedings ERIM 10th Thematic Conference on Geologic Remote Sensing, Environmental Research Institute of Michigan, p. I407-I418.

Bowles, J. Palmadesso, P., Antoniades, J., Baumbeck, M., and Rickard, L.J. 1995. Useof Filter Vectors in Hyperspectral Data Analysis, Infrared Spaceborne RemoteSensing III, Scholl, M.S., and Anderson, B.F. (eds.), SPIE 2553, 148-157.

Caudill, M. 1988. Neural Networks Primer, Part IV, AI Expert, vol. 3(8), August 1988,61-67.

Foody, G.M., and Cox, D. P., 1994. Sub-pixel land cover composition estimation using a linear mixture model and fuzzy membership functions. International Journal of Remote Sensing 15(3), 619-631.

Harsanyi, J.C. (1993) Detection and Classification of Subpixel Spectral Signatures in Hyperspectral Image Sequences, Ph.D. Dissertation, University of Maryland, Baltimore County, 116 p.

Harsanyi, J.C. and Chang, C. 1994. Hyperspectral Image Classification and Dimensionality Reduction: An Orthogonal Subspace Projection Approach, IEEETransactions on Geoscience and Remote Sensing, vol. 32(4), 779-785.

Kohonen, T. 1986. Representation of Sensory Information in Self-Organizing Feature Maps and Relation of these Maps to Distributed Memory Networks, Proceedings of SPIE Advanced Institute on Hybrid and Optical Computing, vol. 634.

McLemore, V.T., Munroe, E.A., Heizler, M.T., and C. McKee. 1999. Geochemistry of the Copper Flat porphyry and associated deposits in the Hillsboro mining district, Sierra County, New Mexico, USA, Journal of Geochemical Exploration, vol. 67, 167-189.

Penn, B.S. 2002. Using Simulated Annealing to Obtain Optimal Linear End-member Mixtures of Hyperspectral Data," Computers & Geosciences, vol. 28(7), 809-817.

Reed, I.S. and Yu, X. 1990. Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution, IEEE Transactions on Acoustics, Speech and Signal Processing, vol. 38, 1760-17770.

Roberts, D.A., Smith, M.O., Adams, J.B., 1993. Green Vegetation, Nonphotosynthetic Vegetation, and Soils in AVIRIS Data. Remote Sensing of the Environment 44, 255-269.

Settle, J.J., and Drake, N.A., 1993. Linear mixing and the estimation of ground cover proportions, International Journal of Remote Sensing 14, 1159-1177.

Smith, M. O., Johnston, P.E., and Adams, J.B., 1985. Quantitative determination of mineral types and abundances from reflectance spectra using principal component analysis, Journal of Geophysical Research (90), 797-804.

Smith, M.O., Ustin, S.L., Adams, J.B., Gillespie, A.R., 1990. Vegetation in Deserts I: A Regional Measure of Abundance from Multispectral Images. Remote Sensing of the Environment 31, 1-26.

Tou, J.T., and Gonzalez, R.C. 1974. Pattern Recognition Principles, Reading, Massachusetts, Addison-Wesley Publishing Company, New York.

Van Der Meer, F., and De Jong, S.M., 2000. Improving the results of spectral unmixing of Landsat Thematic Mapper imagery by enhancing the orthogonality of end-members, International Journal of Remote Sensing (21:15), 2781-2797.

Winter, M. E. 1999. Fast Autonomous Spectral End-member Determination in Hyperspectral Data, 13th International Conference on Applied Geologic Remote Sensing, Vancouver, British Columbia, March 1-3, 1999, II-337-344.

Yu, X., Reed, I.S., and Stocker, A.D. 1993. Comparative performance analysis of adaptive multispectral detectors, IEEE Transactions on Signal Processing, vol. 41, 2639-2656.

Yu, X., Hoff, L.E., Reed, I.S., Chen, A.M., and Stotts, L.B. 1997. Automatic Target Detection and recognition in multispectral imagery: A unified ML detection and estimation approach, IEEE Transactions on Image Processing, vol. 6, 143-156.