Auto-correlogram approach for estimating the spectral-spatial variability of hyperspectral images

By Dmitry Y. Frolov

Abstract

The amount of information contained in a single hyperspectral image is overwhelming for the human operator. As a result, assessing the spatial and spectral variability of a hyperspectral image is very difficult. The existing techniques mainly rely on different preprocessing algorithms that reduce the high-dimensionality of the hyperspectral data down to a few images that can be visualized using traditional RGB or RGBI combinations. The proposed auto-correlogram approach provides a simple framework for reducing a hyperspectral image cube to a single grayscale image that is easy to interpret and screen for spectral anomalies.

Introduction

The recent advances in hyperspectral data collection platforms pose some important questions about the strategies being used for initial data analysis and interpretation. There is no doubt that computer-based analysis procedures play an important role in processing such highly dimensional data as hyperspectral imagery. The amount of information contained in a single hyperspectral image like AVIRIS is beyond human abilities to understand, visualize, and comprehend. The problem is not just the amount of information but its extreme high-dimensionality and what is more important — high correlation between bands. Nevertheless, the human involvement in hyperspectral image analysis is still very high. The important decisions are still being made by a human operator concerning initial data quality assessment, data processing strategy, algorithms to apply, and features to extract. We are still far away from the moment when computer algorithms will be able to make those decisions in the broad range of applications. With the expected launch of OrbView-4, the first space-based hyperspectral imaging platform, the amount of available hyperspectral information is going to explode. There is a definite need for simple and efficient approaches to visualize hyperspectral data in "human perceptible" ways that will help an operator unambiguously assess the spatial-spectral variability of a scene. The existing common approach usually involves applying a dimensional reduction scheme such as the Minimum Noise Fraction (MNF) transformation and displaying the most informative components as an RGB or RGBI image. The resulting representation of the hyperspectral image cube is colorful, but not easy to interpret. The color assignment varies from scene to scene and depends on a number of factors such as spectral properties of the ground features and their coverage in the image area, amount of noise, etc. In most cases this kind of visualization technique does not give a quantitative representation of the scene and serves merely as an indicator of spatial variability. Another popular approach is a perspective 3D visualization of the image cube. The cube cell values are usually color-mapped to enhance the inter-band differences. This visualization method serves primarily as a data presentation. The visualization method that we are looking for should provide a single image-like representation of the whole hyperspectral data cube, allow quantitative interpretation of the imagery, be scene-independent, and allow a direct comparison of the different images. What we are looking for is a visualization algorithm that will provide a human operator a simple way to assess the spectral-spatial (SS) variability of the image. Visualization of this kind of variability provides an efficient way to quickly screen input imagery for the presence of local anomalies and objects of interest and to estimate the quality and potential value of the data. The visualization algorithm must produce results that might be interpreted by a non-expert and unambiguously map the spectral property of individual pixels to a single value or a color triplet. We can loosely define the spectral-spatial variability of a hyperspectral image as a 2D rectangular array of values that assign a spectral variability factor to the every cell on the image. The spectral variability factor is some arbitrary metric that reflects the difference in spectral properties (some sort of "distance") between a given hyperspectral image cell and some subset of its neighbors. There are no real limitations for selecting a distance measure as long as it provides meaningful interpretation of the results, but it might be desirable to select measures that satisfy triangular inequality.

Approach

The proposed algorithm creates an auto-correlogram, an image-like representation of the spectral proximity between each cell in the image and its neighbors within a user-defined distance. The key part of constructing a useful auto-correlogram is the selection of the appropriate metric for computing the spectral "distance" or similarity between hyperspectral image cells. The auto-correlogram approach is similar to the traditional spectral similarity mapping methods, such as spectral angle mapper (SAM) or matched filtering (MF) algorithms, but instead of computing the spectral similarity between each cell and some target material, it calculates the similarity between each cell and some subset of its neighbors. There is no single universal metric that is superior to others for defining spectral similarity. The spectral curve might be viewed and analyzed from many different viewpoints, so the similarity measure should be chosen with respect to the specific application. A variety of algorithms have been used to define the similarity between spectral curves. They all might be roughly classified into several major groups: statistical metrics, vector geometry, and information theory based.

Probably the most popular geometrical method is the spectral angle mapper, which treats spectra as vectors in n-dimensional space. The spectral angle mapper algorithm uses the angle between two vectors as a spectral similarity metric. The advantage of this approach is its relative insensitivity to variation in the scene illumination and albedo effects. Smaller angles represent a higher degree of match between spectral curves.

The RMS distance between two spectra is an example of the popular statistical distance metric (Clark et al, 1990):

$$RMS = [S\ (R_i - P_i)^2]^{1/2} / N \qquad (1)$$

where N is the number of bands, and R_i and P_i are cell values of the spectra being compared in the band i. The RMS distance is a good statistical indicator of the overall spectral similarity. However, it is more sensitive to albedo variations than to changes in the spectral curve shape. Another vector-based measure, the vector dot product, suffers from the same problem.

The Hamming distance is a common distance metric that is used for computing spectral similarity. This metric is based on binary encoding of the spectra into 1- or 2-bit elements that represent the cell value of a given band with relation to its average value and spectral slope. The binary encoding methods have their roots in digital communications and information theory and are very computationally efficient. They have been successfully used in mineral mapping applications (Kruse et al 1993).

The constrained energy minimization (CEM) technique, also known as matched filtering, is an example of a more complicated statistical measure that exploits not only the original spectra, but also statistical information about a hyperspectral image cube in the form of the sample correlation matrix. The CEM optimal linear operator (Harsanyi, 1993) is given by:

$$W = R^{-1}d\ (d^TR^{-1}d)^{-1} \qquad (2)$$

where R^{-1} is the inverse of the sample correlation matrix of the observation pixel vectors and d is a target signature. The linear operator W is applied to every cell in the image and yields an abundance map. The higher abundance values indicate a closer spectral match. This technique might be easily adopted as a distance metric for the auto-correlogram approach, but it will require some additional computation of the optimal linear operator W for every cell in the image. Equation (2) contains a target spectrum that in the auto-correlogram case is the spectrum of the cell that is being processed. A further extension of the CEM operator is a locally adaptive constrained energy minimization algorithm (LA-CEM) that uses a local estimation of the sample correlation matrix in some floating window around the current cell. This approach requires even more computations, but in some cases it produces better results than a traditional CEM operator.

All of the similarity measures mentioned above use a spectral curve in its original form. In some cases this characteristic becomes a real shortcoming, because it makes most of the algorithms more sensitive to variations in the overall energy of the spectral curve than to its shape. Spectral curve shape analysis is definitely under-exploited in the field of spectral similarity mapping. As a partial remedy, most of the algorithms above can be changed easily to use not an original spectral curve, but a transformed version that enhances the influence of the curve's shape on the output of the algorithm. The simplest approach is to use derivatives of the spectra (Woods, 1984). It has been shown that using the first derivative of the signal can significantly enhance the output of the CEM operator (Frolov and Smith, 1999). The abundance map produced contains less noise and has better contrast characteristics.

We also have to mention the importance of image calibration for mapping a spectral similarity. It is desirable to apply an appropriate calibration to the image in order to standardize the signal amplitude in the different bands. The auto-correlogram approach is also sensitive to the image calibration, because that might result in unequal weighting of the bands in the output similarity map. However, the selection of the particular image calibration algorithm is not as important as for traditional spectral matching applications and might be considered more as a statistical normalization procedure.

So far we have been discussing spectral similarity measures and ignoring another important aspect of the auto-correlogram approach. The spatial aspect of the SS-variability is probably equally significant for achieving good results. We have defined an auto-correlogram operator as a measure of the spectral similarity between an image cell and some set of its neighbors. There is an obvious analogy with conventional spatial image filtering techniques that use a moving window for computing a filter response. The auto-correlogram (AC) approach might be viewed as a combination of image filtering and spectral similarity mapping techniques. One can easily derive a number of useful auto-correlogram operators by combining an appropriate spectral similarity measure with a spatial filter design. The first approach is to construct an auto-correlogram operator that computes a spectral distance between the center cell of the moving filter window and every other cell in it. The spatial filtering portion of the algorithm then applies a filter kernel or function to the set of similarity values and outputs a single parameter to the auto-correlogram image.

The advantage of this approach is that we can use a great variety of existing spatial filters and utilize the statistical properties of the spectral similarity values for the given cell. The alternative, but less attractive approach is to somehow process all cell spectra within the moving window and assemble a single representative spectrum that will be used in the spectral similarity computation. The drawback of this method is that the statistical information about the local spectral variability of the image will be lost for further processing or output. Besides, it will limit the possible range of processing algorithms that might be generated using this approach.

We can start the list of possible spatial filters with the simple filter that does a weighted averaging of spectral similarity values within the window. The output auto-correlogram image then represents an average spectral similarity value between each cell and its neighbors. The result is easy to understand and provides a good estimation of the spectral-spatial variability. The weakness of this filter is that it blurs the image and is sensitive to the presence of boundaries between two spectrally different materials within the window. This is usually not a problem for the scenes that do not contain any human-made structures. The averaging of the SS values within the window is similar to applying a Laplacian edge enhancement omni-directional operator. The common Laplacian mask is composed of an eight in the center pixel location with -1 weighting coefficients in the surrounding locations. The SS-variability factor computed for the cell itself is 0 (if we are using the spectral angle mapper algorithm) and the operator basically computes the average of the SS

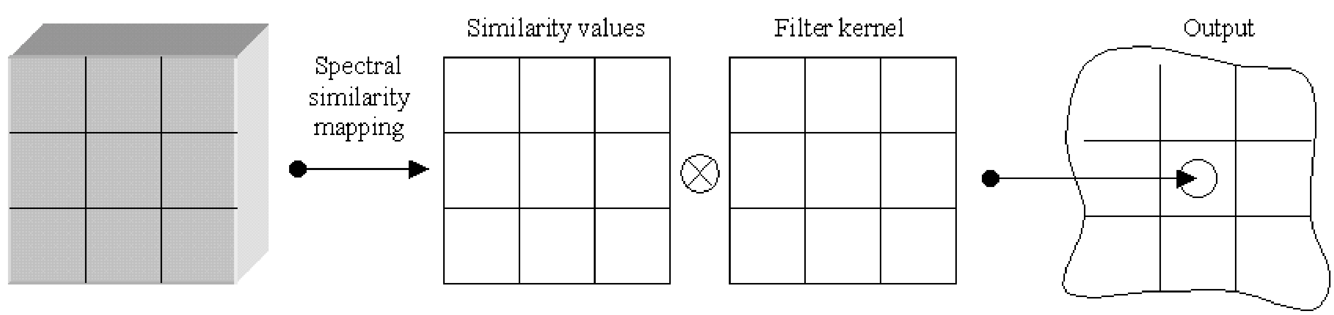

Figure 1. An example of auto-correlogram framework for a 3x3 filter window.

values for the eight immediate neighbors. The output of this kind of auto-correlogram filter is visually very similar to a typical Laplacian operator output.

The nonlinear spatial filters represent another important class of filters that might be used successfully in the auto-correlogram design. This class of filters does not compute a linear summation of elements multiplied by constant weights (filter kernel values). Instead, they use various other methods to process input cell brightnesses and compute filter output. The most popular non-linear filter is definitely a median filter. It has been widely used in noise and bad data removal algorithms. The median filter computes the median value of all cells within a filter window and uses it as an output. The drawback of using these filters is their cumulative nature. These filters might help to visualize inter-band and inter-pixel relationships that occur in some subset of the bands, but they fail to detect a single band occurrence of a sensor problem or some spectral phenomena. The “cumulative” filters work with statistical information and provide a good representation of overall spectral behavior within the scene. Basically, they work as hyperspectral edge detectors. They are not intended to locate subtle anomalies or detect sensor artifacts. For these tasks we have to introduce a different kind of filter that will exhibit spectral and spatial selectivity. One of the simplest approaches is to construct three auto-correlograms for different parts of the spectrum and then display them as an RGB image. However, this does not solve all the problems and results still depend on an arbitrary selection of wavelength ranges.

An interesting set of possible AC-operators might be constructed by applying a filter function that estimates how fast a spectral similarity factor declines with the distance from the center of the window. Another interesting set of operators arises from applying local directional estimators to the filter window. In this case, in addition to the AC output discussed above, the algorithm can generate an image that depicts the direction of the strongest spectral similarity for every cell in the image. All the cells that have no preferred direction of the SS-coefficient and belong to a spectrally isotropic area will get a 0 value, that refers to the absence of a dominant direction. However, image cells that belong to linear features will receive a certain value that will indicate their orientation. This kind of filter might be used as a good quality assurance test for an image, because it allows detection of different kind of resampling and post-processing algorithms applied to an image.

A completely different idea might be used to create an auto-correlogram operator for detecting spectrally localized dissimilarities between bands. In many cases it is desirable to find out how the spectral variability of the ground features changes with the wavelength. A simple implementation of such an operator might search for a wavelength that maximizes the output of some local function F(x,y) for every image cell. A variety of functions might serve as local estimators of image variability: absolute maximum differences, local standard deviation, or some other statistical measures. The wavelength value (or band number) that maximizes the function output for

Figure 2. Example of the application of average spectral angle auto-correlogram operator to an agricultural scene. Wavelength range 0.5 - 1.1 micrometers.

a given pixel is then stored in the output auto-correlogram. The output depicts the spatial distribution of maximum spectral variability of the hyperspectral image. This kind of operator is very sensitive to the calibration procedures applied to the image. Application of such operators to the raw digital numbers of the hyperspectral sensors is not going to produce any meaningful results without some sort of calibration or normalization preprocessing.

Results

We have applied different kinds of auto-correlogram operators to a number of hyperspectral scenes produced by the AVIRIS sensor and other commercial hyperspectral imagery. The goal of our study was to evaluate the potential value of the approach for data visualization, image quality inspection, and detection of small anomalies.

The average spectral angle auto-correlogram operator has been applied to a low altitude AVIRIS image (Figure 2) taken over the state of Maryland. This is a typical agricultural scene containing a variety of bare soil and vegetated fields, forested areas, and a river. It is easy to see that visually this kind of auto-correlogram resembles the output of the Laplacian edge detection filter. For every pixel we computed an average spectral angle between the cell and its eight immediate neighbors. The maximum values (that correspond to maximum spectral differences) are concentrated along edges between distinct ground features such as roads, different kinds of agricultural fields, water, and human made structures. The cumulative nature of the operator provides a good general overview of the spectral-spatial variability within the image and is relatively insensitive to the noise or sensor problems in individual bands. The output auto-correlogram provides a visual clue to the texture of the objects on the ground. Vegetated and bare soil fields are easily distinguished by absolute values as well as by image textures. The highest variability is found in the forested areas, partially because of their irregular texture and a significant shadow effect. The output of the auto-correlogram operator might be analyzed also on a local basis, for example by inspecting spectral variability properties of the individual fields. The auto-correlogram of the agricultural field might be extracted out of the output and displayed independently. The contrast-enhanced auto-correlogram may reveal the existence of patterns that might correspond to the different soil or plant conditions within a field.

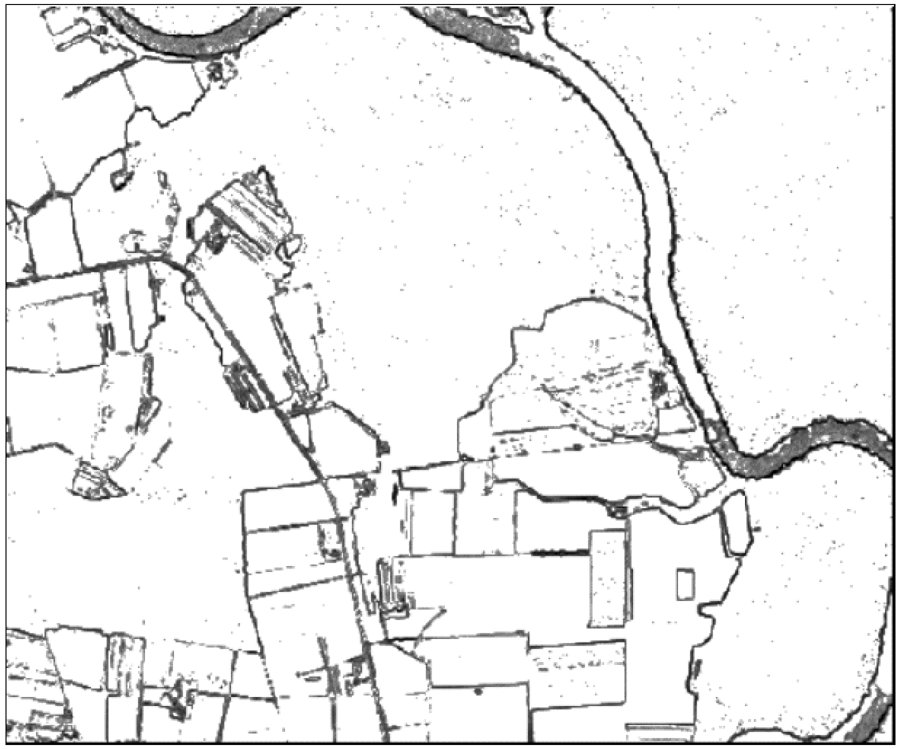

Figure 3. Thresholded version of the average spectral angle auto-correlogram.

Another possible application of the average spectral angle auto-correlogram is an edge detection scheme that works simultaneously on all image bands. The thresholded version (Figure 3) of the auto-correlogram might be used as input to an automatic image vectorization procedure. The advantage of using the whole set of image bands over picking a single band for edge/ border vectorization is that it unambiguously uses all available information and is not dependent on selecting a correct image band.

However, constructing a hyperspectral edge detector is probably not a very practical use of the auto-correlogram. One of the most useful areas where auto-correlograms might really produce interesting results is in sensor defects detection and image quality assurance. Everybody who has worked with the currently available hyperspectral imagery has probably encountered problems associated with the sensor or registration system malfunctions. In many cases such image degrading artifacts occur in some small subset of the bands and appear on the image in some regular geometrical shapes. The typical sensor errors include change in the signal amplitude, excessive noise, or even complete loss of the image in some spectral bands. More interesting artifacts are introduced by registration systems including spatially misplaced portions of images. Some post-processing operation such as geo-rectification is also a possible source of systematic errors that might significantly affect the quality of further hyperspectral image analysis and classification. Sometimes, if the detailed post-processing history of the image is not available, it is important to screen the data for possible systematic errors or the presence of artificially shaped spectral anomalies. Detection of image resampling and the actual resampling pattern might also prove to be valuable.

Detection of spatial resampling applied to the image is a relatively simple task and might be efficiently achieved by constructing a directional auto-correlogram operator that depicts the direction of strongest spectral similarity for every cell in the image. We have constructed an auto-correlogram operator that finds for every image cell a most spectrally close neighbor (within the processing window centered around the current cell) and checks a similarity threshold. If spectral similarity is within a given threshold range this pixel is considered to have a significantly similar neighbor. The relative position of the similar cell is then converted to one of the four or eight possible directions and the output auto-correlogram cell is marked with this value. If the closest spectral similarity value is below a user-specified threshold, then such cell is considered to have no preferred direction and it is appropriately marked in the output. One of the possible ways to define a threshold is to use the local standard deviation of the spectral similarity as a basis for the value. We applied such a directional auto-correlogram operator to the low-altitude AVIRIS scene taken over agricultural fields of the Salinas Valley (Figure 4).

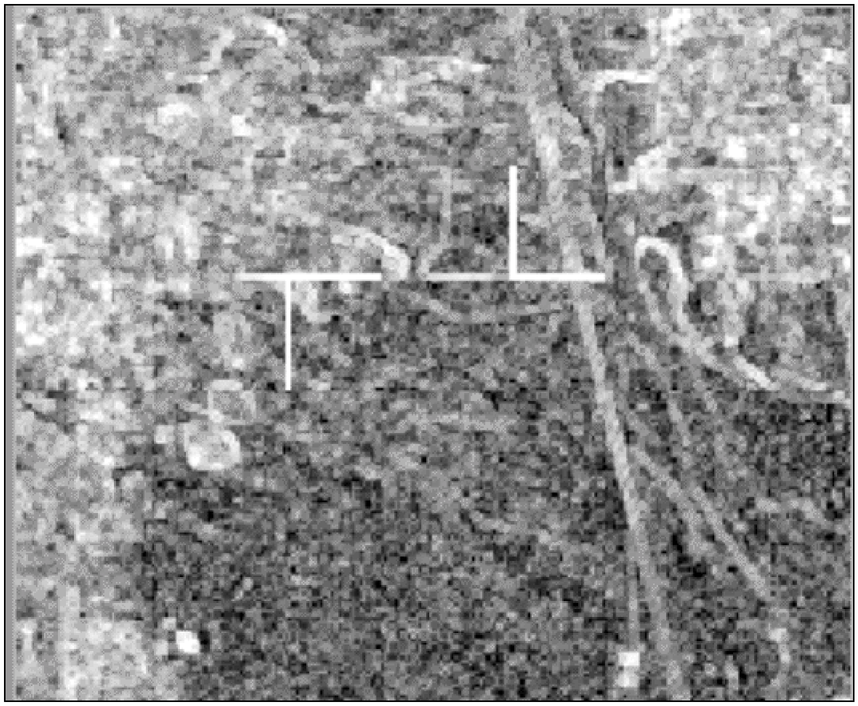

Figure 5. “Maximum spectral difference” auto-correlogram for the portion of AVIRIS image taken over Cuprite site.

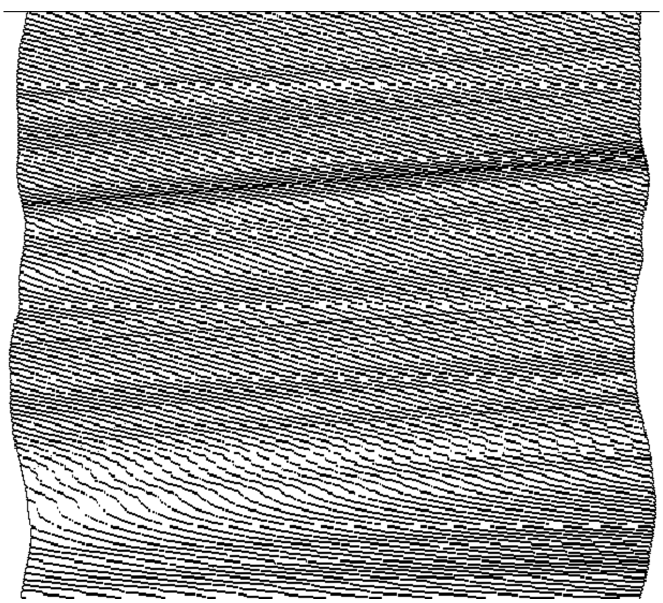

Figure 4. The directional auto-correlogram operator applied to the low-altitude AVIRIS image that has been geo-rectified using a nearest neighbor resampling.

This low-altitude AVIRIS image has been post-processed in order to remove a distortion typically associated with the low-altitude scan-line hyperspectral images. It has been geo-rectified and resampled using a nearest-neighbor resampling technique. A significant amount of new image cells that are just duplicates of their neighbors have been introduced in the image. The set of nearly horizontal and diagonal lines are actually a pattern of the resampling procedure applied to the image. Each of the black cells exhibits unusually high spectral similarity with one of its neighbors (in this particular case of nearest-neighbor resampling, cells are identical to each other). The directional auto-correlogram for the original image does not contain any of those patterns and contains a very small amount of spectrally similar cells. This is a good illustration of how an auto-correlogram might find and reveal statistical and functional relationships that exist in the image. This approach might be used also in the case when it is necessary to restore an original unprocessed image from the resampled one, but the original resampling map is lost or not available. The drawback of this operator is that it does not allow us to determine the exact position of the original cells, because we do not have a basis to distinguish the newly introduced cells from the original ones. Both kinds have high similarity values with each other and no signs of being original. This consideration limits the accuracy of the method to ± 1.

The directional auto-correlogram is effective for detecting patterns that are not localized to a few individual bands but occur in most bands of the image. A significant number of defects in hyperspectral images actually occur within a single band and might call into question all further processing results with that data. Figure 5 shows an example of the application of a maximum difference auto-correlogram operator to a portion of the high-altitude AVIRIS scene taken over Cuprite, Nevada. Some of the hyperspectral image artifacts might be detected by finding for every image cell a wavelength (band) that maximizes differences between the cell and its neighbors. The auto-correlogram output then may contain two possible types of images: absolute maximum difference within the window and a wavelength (band number) where this difference was recorded. When such a simplified operator was applied to the Cuprite scene, it detected some regular geometrical shapes in the both types of auto-correlogram images. The wavelength image allowed us to quickly pinpoint the band number 129 (1.573 micrometers) responsible for these problems. This approach still does not guarantee that every kind of single-band sensor defect might be identified automatically. Besides finding sensor problems, the detection of the wavelengths where image cells exhibit the most variability or some specific behavior is also useful for general scene understanding, because it provides some initial clues about where the image information is and how it is distributed across the bands. The atmospheric absorption bands are a significant source of noise and generally should be excluded from processing. A great variety of auto-correlograms might be constructed by selecting a different variability criterion or function allowing us to adjust the algorithm to different problems and hyperspectral data sources.

Conclusion

A new, easily implemented approach for visualizing hyperspectral images has been demonstrated. The proposed auto-correlogram method can be easily modified to suit a great variety of quantitative processing and analysis needs. The basic idea, elaborated in a collection of different algorithms, is to explore how spectral or other properties of each cell in the image differ from some subset of its neighbors. The output of all the algorithms is always an image that depicts the spatial distribution of the explored phenomenon. A further extension of visualization techniques that use the concept of the auto-correlogram is an RGB or RGBI display of a number of different auto-correlograms where each correlogram type is responsible for either the red, green, blue or intensity channel. The application of auto-correlograms to sensor defect detection and general scene understanding has been demonstrated using publicly available hyperspectral imagery. All of the algorithms were implemented inside the hyperspectral module of TNTmips, an advanced GIS and image processing system.

Acknowledgments

The author would like to thank Dr. Randall Smith for his valuable suggestions and support.

References

Clark, R.N., A.J. Gallagher, and G.A. Swayze, 1990, “Material Absorption Band Depth Mapping of Imaging Spectrometer Data Using a Complete Band Shape Least-Squares Fit with Library Reference Spectra,” Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop. JPL Publication 90-54, pp. 176-186.

Frolov, D.Y., and R.B. Smith, 1999, “Locally Adaptive Contrained Energy Minimization for AVIRIS Images,” Eighth JPL Airborne Earth Science Workshop, edited by Rober O. Green, (in press).

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

Kruse, F.A., A.B. Lekhoff, and J.B. Dietz, 1993, “Expert System-Based Mineral Mapping in Northern Death Valley, California/ Nevada, Using the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS),” Rem. Sens. Envir., vol. 44, pp. 309-335.

Woods, J.W., 1984, “Image Detection and Estimation” in “Digital Image Processing Techniques,” edited by P.M. Ekstrom, Academic Press, Inc., Orlando, Florida, pp. 80-81.