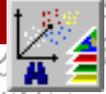
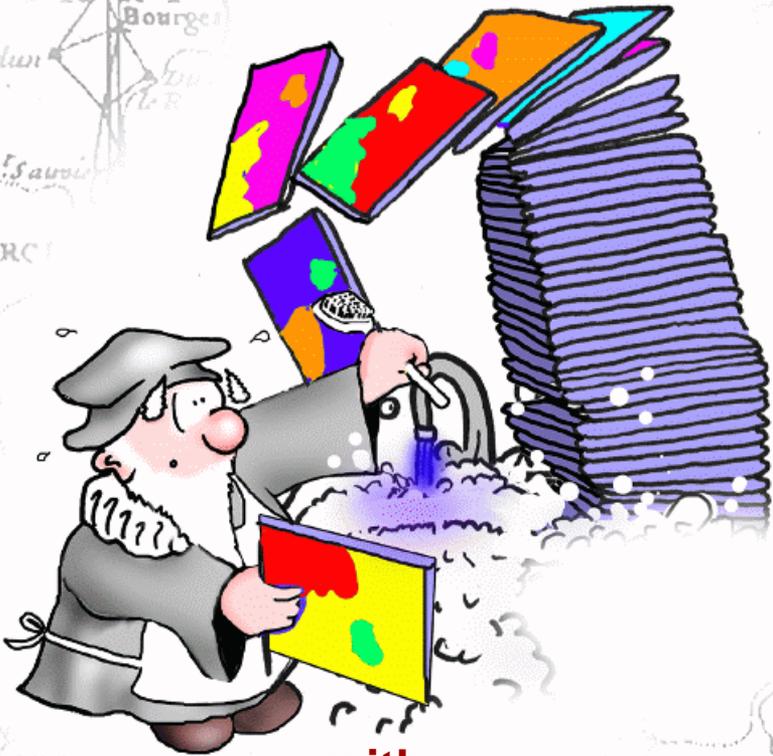


Tutorial



Analyzing Hyperspectral Images



with
TNTmips[®]

The hyperspectral analysis techniques described in this tutorial can be applied **FREE** by anyone with AVIRIS or other hyperspectral imagery using TNTmips **FREE**

Before Getting Started

This booklet introduces the Hyperspectral Analysis process in TNTmips®, which provides specialized methods for visualizing, processing, and analyzing images produced by the new generation of hyperspectral remote sensors. These exercises will familiarize you with basic procedures in Hyperspectral Analysis, including acquiring and saving image spectra, working with spectral libraries, calibrating images to reflectance, and spectral matching and classification.

Prerequisite Skills This booklet assumes that you have completed the exercises in the tutorial booklets entitled *Displaying Geospatial Data* and *TNT Product Concepts*. Those exercises introduce essential skills and basic techniques that are not covered again here. Please consult those booklets for any review you need. Additional background on hyperspectral analysis can be found in the booklet *Introduction to Hyperspectral Imaging*.

Sample Data The exercises presented in this booklet use files in the CUP97 data collection. Because of the large size of this dataset, it is not distributed with the TNT products. You can download the CUP97 dataset from microimages.com or contact MicroImages to obtain a free copy of these Project Files on CD.

More Documentation This booklet is intended only as an introduction to the Hyperspectral Analysis process. Background information on hyperspectral images and their analysis can be found in the booklet *Introduction to Hyperspectral Imaging*. Further information can be found in a variety of tutorial booklets, Technical Guides, and Quick Guides that are available from microimages.com.

TNTmips® Pro and TNTmips Free TNTmips (the Map and Image Processing System) comes in three versions: the professional version of TNTmips (TNTmips Pro), the low-cost TNTmips Basic version, and the TNTmips Free version. All versions run exactly the same code from the TNT products DVD and have nearly the same features. If you did not purchase the professional version (which requires a software license key) or TNTmips Basic, then TNTmips operates in TNTmips Free mode.

All the exercises can be completed in TNTmips Free using the sample geodata available from MicroImages.

Randall B. Smith, Ph.D., 23 August 2013

©MicroImages, Inc., 2000-2013

You can print or read this booklet in color from MicroImages' Web site. The Web site is also your source for the newest tutorial booklets on other topics. You can download an installation guide, sample data, and the latest version of TNTmips.

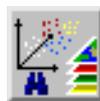
<http://www.microimages.com>

Welcome to Hyperspectral Analysis

Hyperspectral remote sensors collect image data in many narrow, contiguous spectral bands. The resulting datasets contain numerous image bands, each depicting the scene as viewed within a narrow wavelength range. Although this multiband conceptualization is a natural outgrowth of our experience with multispectral images, an alternative approach is more useful for analyzing and interpreting hyperspectral data. The hyperspectral scene can be thought of instead as a single image with a spectrum of brightness values stored for each image cell. These image spectra can be compared with field or laboratory spectra in order to recognize and map the spectral signatures of ground materials such as green vegetation or particular minerals.

The prototype Hyperspectral Analysis process in TNTmips provides specialized interactive analysis tools that allow you to fully exploit the spectral range and spectral resolution provided by hyperspectral datasets. The process allows you to view and save image spectra and to compare image spectra with laboratory spectra stored in spectral libraries. Several methods of spectral classification are also provided.

To illustrate the use of the Hyperspectral Analysis process, we will work with a sample scene from NASA's Airborne Visible/Infrared Imaging Spectrometer (AVIRIS). This sensor collects data in 224 contiguous spectral bands with a bandwidth of 0.10 μm . Each 20-m square cell in the scene has a continuous spectrum over the range from 0.4 to 2.5 μm . The scene covers the Cuprite mining district in western Nevada, USA. Rocks in several areas of the scene have been altered and mineralized by hot water solutions (hydrothermal alteration), creating concentrations of alteration minerals that provide good targets for spectral analysis.



STEPS

- launch TNTmips
- choose Image / Hyperspectral from the TNTmips menu

The exercises on pages 4-10 show you how to open a hyperspectral image, examine image bands and band information, and use the Hyperspectral Explorer to find a useful RGB combination of bands for visual reference during analysis. On pages 11-14 you learn how to acquire, view, and save image spectra. Classification of the hyperspectral image using image spectra is introduced on pages 15-21, and pages 22-23 introduce dimensional reduction methods. Procedures for finding spectral end-members for analyzing images are covered on pages 24-29. Pages 30-36 introduce methods of correcting the raw image values to reflectance and discuss the use of laboratory reflectance spectra and spectral matching procedures. Pages 37-39 discuss continuum-removal and some additional image processing options.

Open a Hyperspectral Image

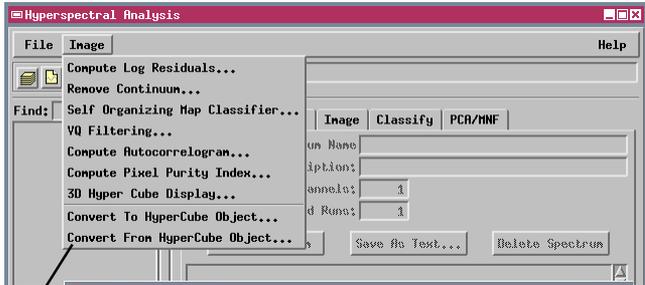
NOTE: The sample files used in these exercises are not distributed on the TNTmips CD. You can download these files from the Microlimages web site (www.microimages.com) or contact Microlimages to obtain a free copy of the files on CD-R.

STEPS

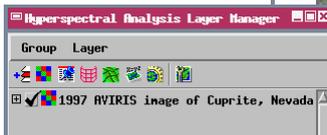
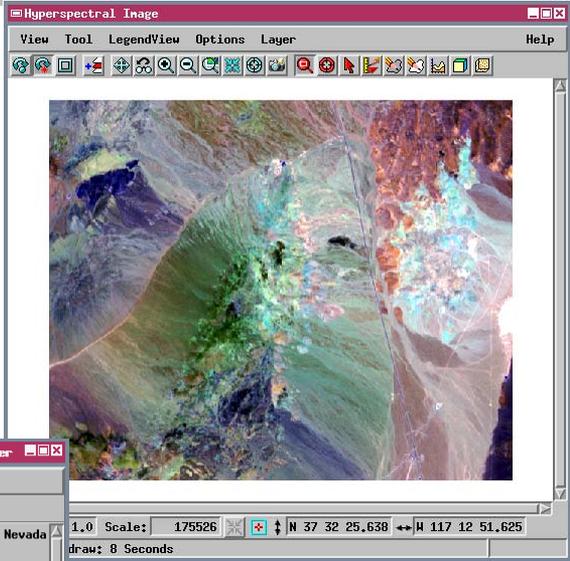
- click the Open icon button in the Hyperspectral Analysis window and select HyperCube Object from the dropdown menu
- use the standard Select Object dialog to choose the CUPRITE97 object in the CUP97CUB Project File



A hyperspectral image can be stored in a TNTmips Project File in one of two ways: as a set of separate raster objects, with one object for each wavelength band, or as a single hyperspectral object. A hyperspectral (or hypercube) object includes raster value and wavelength information for all of the bands in a hyperspectral image. The hyperspectral object incorporates lossless image compression, and so offers a significant reduction in file size in comparison to the multiraster format. Image display in the Hyperspectral Analysis process is also much faster when using a hypercube object.



When you import a hyperspectral image you can choose either the multiraster or hypercube format. You can also convert an image from one format to another within the Hyperspectral Analysis process using options on the Image menu.



3D Hypercube Display

The 3D Hypercube Display emphasizes the high spectral content of the hyperspectral image while allowing you to quickly examine any of the individual wavelength bands. The 3D Hypercube Display window portrays the hyperspectral image as a three-dimensional “image cube”. The top and right panels of the cube show the corresponding edge cells of each wavelength band, with wavelength increasing toward the back of the cube. The shortest wavelength band is displayed by default on the front face of the cube. You can select any of the wavelength bands for display on the front of the cube by using the Band Number slider or the arrow icon buttons.

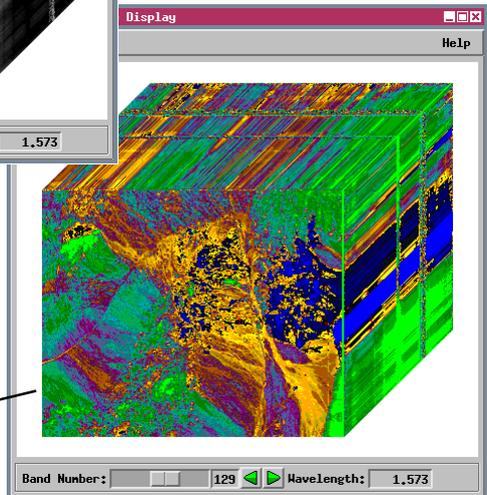
STEPS

- choose 3D Hyper Cube Display from the Image menu in the Hyperspectral Analysis window
- move the Band Number slider at the bottom of the 3D Hyper Cube Display window to view different wavelength bands on the front face of the cube
- choose Close from the File menu in the 3D Hyper Cube Display window



You can also apply a color map to the image cube. Select Edit Color Map from the File menu to open a standard Color Palette Editor window. You can use the Palette menu to choose one of the many predefined color palettes, or design your own color palette. The illustration to the right shows the Cuprite image cube with the predefined Beach color palette.

The “fuzzy” zones on the edges of the hypercube indicate two groups of very noisy image bands. Each of the two groups spans a wavelength region in which atmospheric water vapor absorbs almost all incoming and reflected solar radiation. As a result the bands in these wavelength regions (near 1.4 and 1.9 micrometers) have little image content and are dominated by sensor noise.



Select Bands for RGB Display

STEPS

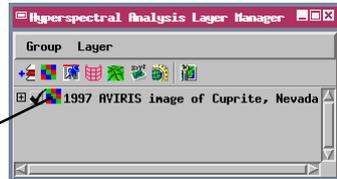
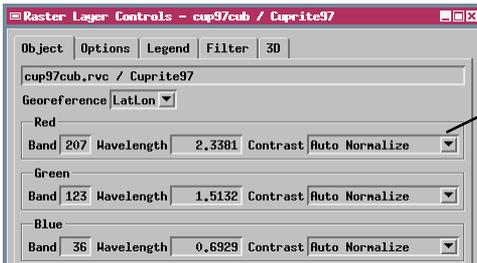
- in the toolbar in the Hyperspectral Analysis window, click on the Select Bands icon button 
- in the Select Bands window, use the Red, Green, and Blue menus to change the bands selected for display and press [OK]



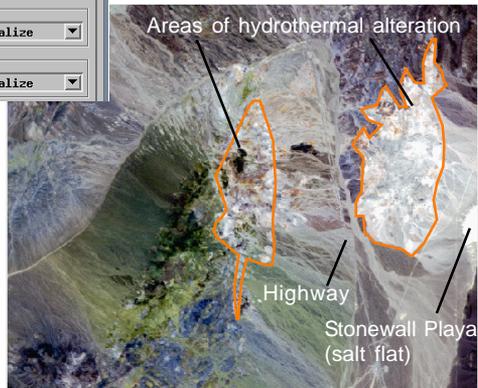
- in the Layer Manager window, click on the icon for the Cuprite hypercube layer 
- note the Band number and wavelength reported for Red, Green, and Blue, then click [OK]

The Hyperspectral Analysis process automatically selects three bands for RGB display in the Hyperspectral Image window. The bands are initially chosen to span most of the spectral range of the data while avoiding bands with excessive noise or no data. The selection of bands for display has no impact on analysis processes, so you can use the Select Band window to manually change the combination of bands you want to use as a visual reference. The Raster Layer Controls window shows the band numbers that have been selected for Red, Green and Blue color components along with their assigned wavelengths. You can also use this window to select any of the standard automatic contrast-enhancement methods for each color channel. Your band selections and contrast settings are automatically saved as part of the display parameters for the hypercube object.

The Hyperspectral Explorer procedure (described on the following pages) provides a rapid means for evaluating many RGB band combinations.



Cuprite image with bands selected to approximate natural color view. The Cuprite district is a desert area with very little vegetation. It includes several hills surrounded by alluvial fans, and is bordered on the east by a dry lake bed (Stonewall Playa). Hydrothermally altered rocks are exposed in the hills on either side of the highway.



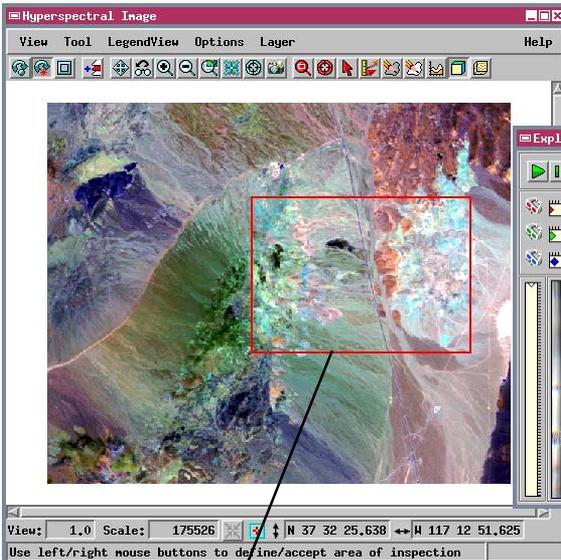
Start the Hyperspectral Explorer

The Hyperspectral Explorer is a unique, automated tool for visualizing the hyperspectral image. It allows you to automatically create a large number of potential RGB band combinations for a sample area. You can then preview the combinations as an animated sequence and select one to use as a reference image during further analysis.

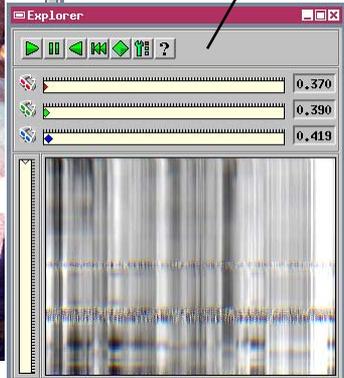
After you select the Explorer tool, use the mouse to drag open a rectangular inspection box in the Hyperspectral Image window. When you right-click to accept the box, the portion of each band included within it is loaded in memory so that different RGB combinations can be displayed rapidly. (The size of the inspection box cannot exceed the product of 256 by 256 raster cells.) Other controls for the procedure are accessed from the Explorer window, which is explained on the following pages.

STEPS

- press the Hyperspectral Explorer icon button on the Hyperspectral Image window 
- hold the left mouse button as you drag the cursor over the image to draw a rectangular inspection area; release the button to cease drawing
- resize or move the rectangle to match the illustration
- click the right mouse button to accept the inspection area and load the image bands (this may take several minutes)



Explorer window after sample images are loaded.



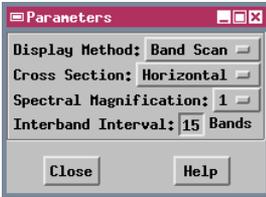
Hyperspectral Image window with completed Explorer inspection box. (If the box reaches the Explorer size limit as you are drawing, the initial box disappears, and a new box begins to draw at the current cursor location.)

This exercise continues on the next page.

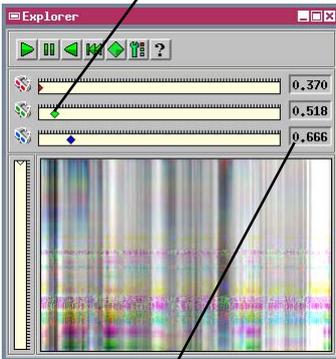
Explorer Window Display

STEPS

- press the Parameters button  on the Explorer window
- change the Interband Interval value in the Parameters window to 15, and press [Close]



You can also drag the colored sliders for the Red, Green and Blue components to set an initial RGB combination. This sets both the wavelength order and interband intervals for all RGB combinations.



These fields show the wavelength (in micrometers) of each component in the current RGB set.

This exercise continues on the next page.

The Hyperspectral Explorer creates a suite of RGB band combinations using constant Red-Green and Green-Blue intervals (in number of bands). You can set a single value for both intervals using the Interband Interval field in the Parameters window. When you use this option, the shortest-wavelength band in each combination is assigned to the Red display channel and the longest-wavelength band to the Blue channel. For example, an Interband Interval of 15 creates RGB combinations using bands 1, 16, and 31; 2, 17, and 32; 3, 18, and 33; and so on until each band has been used in at least one RGB image.

To understand the lower part of the Explorer window, imagine that all of the prepared RGB images of the inspection box are stacked on top of each other, with the lower wavelength combinations at the top. Then make a vertical slice down through the stack along a line through the middle of the inspection box (shown as the dashed cross section line in the view window). The Explorer window displays this vertical slice. Each horizontal line of pixels in the window pane shows the colors produced by one of the RGB combinations along the cross section line. The display gives an overview of the range of colors produced by each of the band combinations, allowing you to focus on useful ranges of combinations.

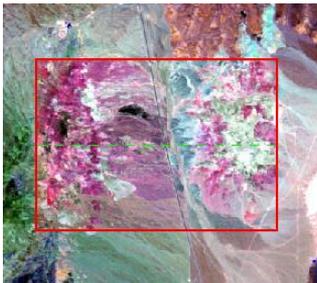
Because AVIRIS bands are narrow and have no wavelength gaps between them, adjacent bands appear very similar to each other. When a small value (such as 3, the initial default) is used for the Interband Interval, each RGB set consists of three very similar bands. Consequently, each RGB image cell has nearly equal contributions of Red, Green, and Blue. The resulting color images (and the Explorer window view) are dominated by gray tones (see the illustration on the previous page). Increasing the Interband Interval reduces the similarity between component rasters, producing more colorful RGB combinations.

Run the Hyperspectral Explorer

When you press the Play button on the Explorer window, the process cycles through the prepared set of RGB combinations in increasing wavelength order and displays them at one second intervals in the inspection box. You can use the other icon buttons to pause, reverse, or restart the sequence, and the vertical slider to change position in the sequence. When you find a combination that highlights features of interest, press the Full View icon button to open this combination to the full image.

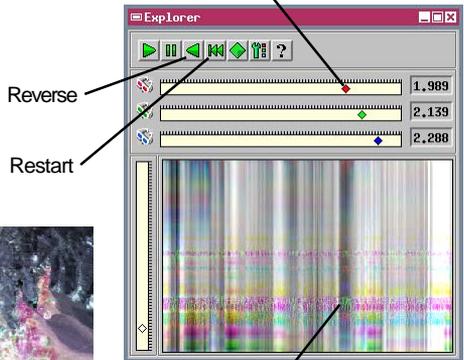
STEPS

- click the Play icon button on the Explorer window 
- press the Pause icon button when a color combination similar to the one in the illustration is displayed 
- press the Full View icon button 



Portion of the original RGB image with a new RGB combination displayed in the inspection box.

When you play the sequence of Explorer images, the colored sliders move to indicate the band wavelengths used in the current combination, and the vertical slider shows the current position in the stack of combinations depicted in the lower part of the window.



Bands 172 (1.989 μm), 187 (2.139 μm), and 202 (2.288 μm) displayed respectively as R, G, and B. This band combination highlights the hydrothermally altered rocks in the Cuprite scene and distinguishes different types of alteration.

The “fuzzy” horizontal stripes in the Explorer window represent groups of RGB sets that include one or more bands with significant image noise. You should avoid including very noisy bands in analysis operations, since they may not provide useful information.

- press the Zoom Box icon button in the Hyperspectral Image window to hide the Explorer window 

Examine Band Characteristics

Hyperspectral analysis procedures require spectral information for each band in the hyperspectral image. Controls for viewing and editing this information are found on the Image tabbed panel in the Hyperspectral Analysis window.

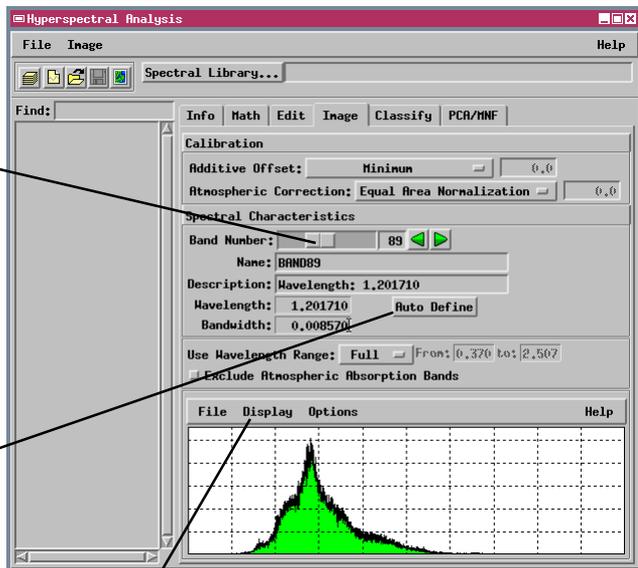
STEPS

- ☑ click on the Image tab on the Hyperspectral Analysis window
- ☑ move the Band Number slider to examine the Wavelength and Bandwidth values and the histogram for various bands

The Band Number slider on the Image panel allows you to quickly select any one of the component bands and review its spectral characteristics. A histogram of the current band is shown at the bottom of the panel. The Wavelength field shows the wavelength in micrometers of the band center, while the Bandwidth field shows the width of the spectral band. For AVIRIS images, the import process automatically incorporates this information with the spectral bands. If you import hyperspectral data from another source not yet supported by TNTmips, you will need to enter the spectral information for each band. If the wavelength value (in micrometers) is included in the object description, you can use the Auto Define utility to assign the appropriate wavelength to each band; approximate bandwidths are then calculated when needed during processing.

Use the Band Number slider or the arrow icon buttons to select the desired band.

Use the Auto Define option only when spectral information has not yet been associated with the spectral bands.



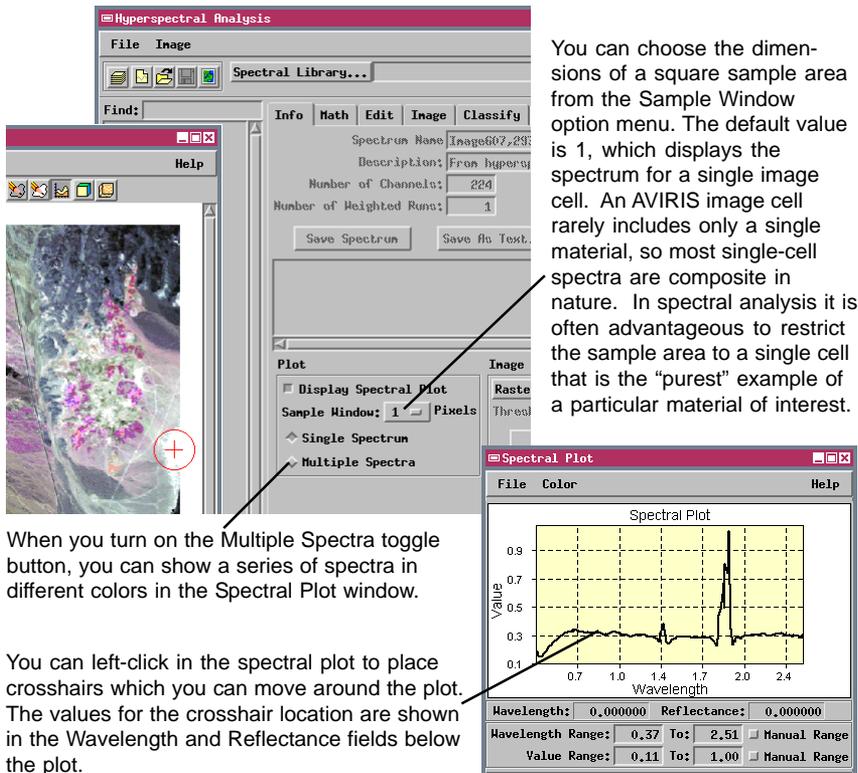
The Display menu provides several modes for viewing histograms, including outline (shown), bars, and 3D strips.

View Image Spectra

The Spectral Profile tool allows you to select an image cell and display its spectrum in the Spectral Plot window. You can also specify a small area and display its average spectrum. A spectral plot is a graph of spectral brightness value (vertical axis) versus wavelength in micrometers (horizontal axis). The Spectral Plot window can display one or more spectra drawn from the hyperspectral image, or spectra stored in a spectral library. The value and wavelength ranges and grid intervals for the plot are computed automatically from the selected spectrum or spectra. You can also choose to set the range for either plot axis manually using the controls at the bottom of the Spectral Plot window.

STEPS

- click the Info tab on the Hyperspectral Analysis window
- turn on the Display Spectral Plot toggle button
- select the Spectral Profile tool from the Hyperspectral Image window toolbar 
- move the cursor to the middle of Stonewall Playa (see illustration) and left-click to place the selection cursor
- right-click to read and display the spectral profile



You can choose the dimensions of a square sample area from the Sample Window option menu. The default value is 1, which displays the spectrum for a single image cell. An AVIRIS image cell rarely includes only a single material, so most single-cell spectra are composite in nature. In spectral analysis it is often advantageous to restrict the sample area to a single cell that is the "purest" example of a particular material of interest.

When you turn on the Multiple Spectra toggle button, you can show a series of spectra in different colors in the Spectral Plot window.

You can left-click in the spectral plot to place crosshairs which you can move around the plot. The values for the crosshair location are shown in the Wavelength and Reflectance fields below the plot.

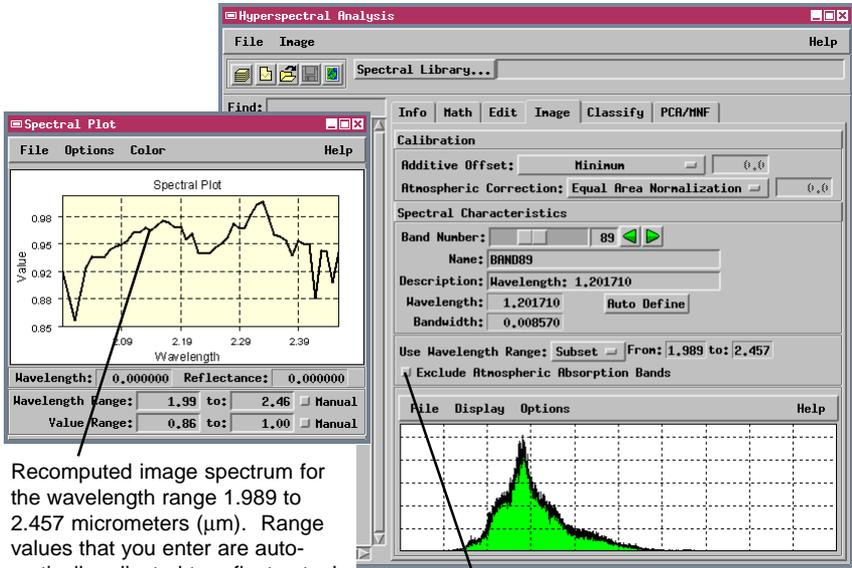
Wavelength:	0.000000	Reflectance:	0.000000
Wavelength Range:	0.37 To: 2.51	<input type="checkbox"/> Manual Range	
Value Range:	0.11 To: 1.00	<input type="checkbox"/> Manual Range	

Set a Wavelength Range

STEPS

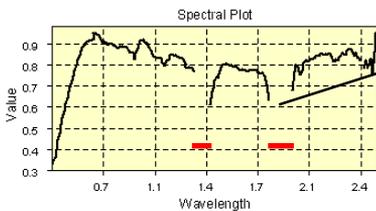
- ☑ click on the Image tab in the Hyperspectral Analysis window
- ☑ select Subset from the Use Wavelength Range option menu
- ☑ enter 1.989 in the From text field and 2.457 in the To text field
- ☑ move the cursor into the Hyperspectral Image window and click the right mouse button to recompute the current image spectrum

The characteristic absorption features of many of the alteration minerals in the Cuprite scene lie in the wavelength range from 2.0 to 2.46 μm . The Subset option on the Use Wavelength Range menu allows you to restrict processing to the hyperspectral image bands within a specified wavelength range. Once you designate the subset range, image spectra and all spectral analysis functions are limited to that group of image bands. Specifying a wavelength subset can speed processing and increase the quality of spectral classification results. You can change the subset range at any time during an analysis session.



Recomputed image spectrum for the wavelength range 1.989 to 2.457 micrometers (μm). Range values that you enter are automatically adjusted to reflect actual image band wavelengths.

The Exclude Atmospheric Absorption Bands toggle excludes from processing two preset wavelength ranges: 1.33 to 1.41 μm , and 1.81 to 1.95 μm . Atmospheric water vapor absorbs most of the energy in these spectral regions, so the included image bands contain much random noise and little information about surface materials. These ranges are shown as gaps in spectral plots, as in the plot to the left (highlighted here by the red bars).



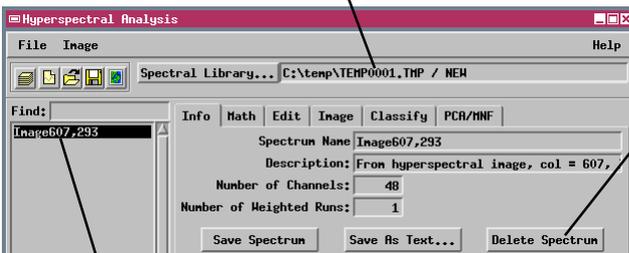
Save Image Spectra in a Spectral Library

In some instances you may want to acquire a number of image spectra for use in analyzing and classifying the hyperspectral image. To do so, you need to save the spectra in a spectral library object. Choosing the New Spectral Library option opens a temporary library. After one or more spectra are added to the library, you can save it as an object in a Project File.

If you save an image spectrum acquired from a subset of the image bands, only the portion of the spectrum corresponding to the selected wavelength range is saved in the library. Hyperspectral analysis operations that compare spectra automatically compensate for differences in spectral range and spectral resolution. Before comparison functions are applied, overlapping spectra are truncated to their common spectral range, and spectra with differing band widths are resampled to the coarsest input spectral resolution.

When you extract an image spectrum with the Spectral Profile tool, a default spectrum name is created with the form ImageXX,YY; XX is the raster column number and YY is the raster line number. You can rename the spectrum if you like by editing the Spectrum Name text field before saving the spectrum in the library.

The name of the file containing the currently open spectral library is shown in the text field.



The names of spectra in the open library are listed in the pane on the left side of the Hyperspectral Analysis window.

STEPS

- click the New Spectral Library icon button on the Hyperspectral Analysis window
- move the cursor into the Hyperspectral Image window and right-click to reacquire the image spectrum
- click the Info tab and note the default name in the Spectrum Name field
- click [Save Spectrum]
- click the Save As icon button
- use the standard Select Object procedure to name a new Project File and library object

The spectral values in an image spectrum are adjusted automatically before plotting or processing according to the options selected on the Calibration subpanel of the Image tabbed panel. These options are used to adjust radiance values measured by a sensor to values approximating the spectral reflectance of the ground materials; they are discussed in a later exercise.

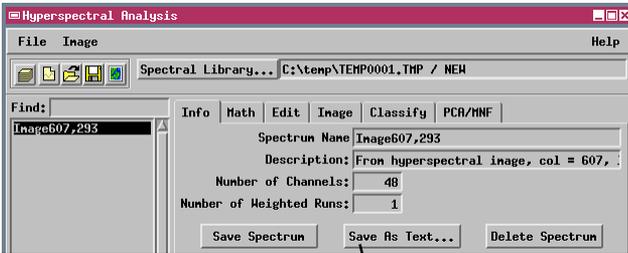
The spectrum currently highlighted in the library listing can be deleted using the Delete Spectrum button.

Save a Spectrum as Text

STEPS

- ☑ click the Save as Text button on the Info panel
- ☑ use the standard Select File dialog to name an output text file for the spectrum
- ☑ open the OUTPUT .CSV file in a text editor to examine its format
- ☑ exit the text editor when you are finished examining the file

You may wish to export spectra for use in other software for further analysis and plotting. The Save As Text button saves the data for a selected spectrum in a text file with the file extension .csv (comma-separated values). The format of the file is illustrated by the sample below. The file has one line for each spectral band, with the data lines preceded by a header line. There are three columns of data that are separated by commas: wavelength, reflectance (or uncalibrated brightness value), and standard deviation (bandwidth). Files in this format can be opened directly by most spreadsheet programs. If you want to export more than one spectrum, each one must be saved separately to its own file.



Use the Save As Text button to save the selected spectrum in a CSV file.

wavelength	reflectance	standard deviation (bandwidth)
1.98867,	0.91535,	0.010990
1.99869,	0.88486,	0.010970
2.00870,	0.85609,	0.010950
2.01872,	0.89134,	0.010940
2.02873,	0.92012,	0.010920
2.03874,	0.93310,	0.010900
2.04875,	0.93234,	0.010880
2.05876,	0.93215,	0.010870
2.06876,	0.94171,	0.010850
2.07877,	0.94504,	0.010830
2.08877,	0.94841,	0.010810
2.09877,	0.95240,	0.010800
2.10876,	0.96267,	0.010780...

Use Image Spectra for Classification

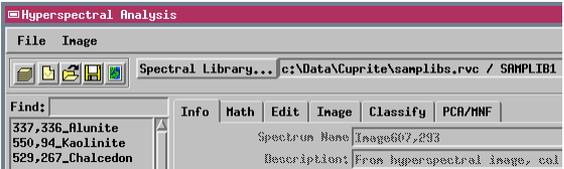
If you know the locations of key surface materials in portions of the hyperspectral scene, you can use image spectra from these locations to identify other occurrences in the scene. This procedure is in some ways similar to supervised classification, and several methods are available on the Classification tabbed panel.

The SAMPLIB1 spectral library contains image spectra from relatively pure occurrences of three of the common alteration minerals in the Cuprite scene.

However, most cells in the image probably include more than one material or mineral, each of which has its own characteristic spectrum. The spectrum of such a cell has a composite shape that depends on the spectral shapes and relative abundances of the included materials. The hyperspectral classification methods produce a raster with cell values that measure the similarity between each image spectrum and the target spectrum. Be cautious when interpreting the values in this raster, because either a high abundance of the target material or the presence of another material with a similar spectrum can produce equivalent “similarity” values.

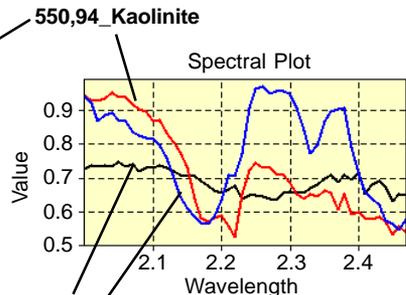
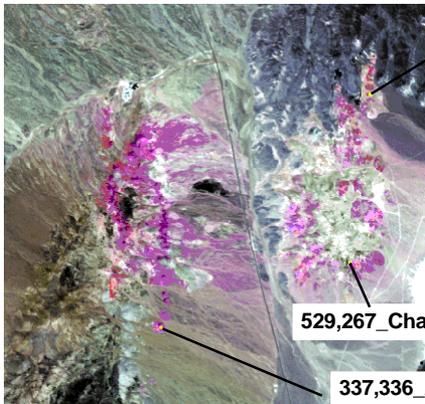
STEPS

- click [Spectral Library...]
- select object SAMPLIB1 from the SAMPLIBS Project File
- on the Info panel, turn on the Multiple Spectra toggle button for spectral plots



- select each of the three spectra listed in the pane on the left side of the Hyperspectral Analysis window to add them to the Spectral Plot window
- click the Classify tab

This exercise continues on the next page.



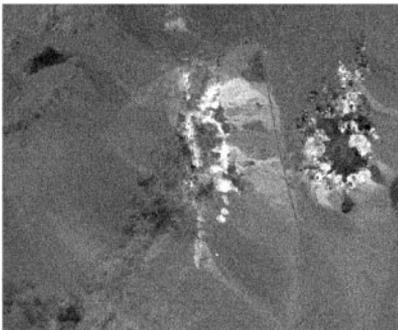
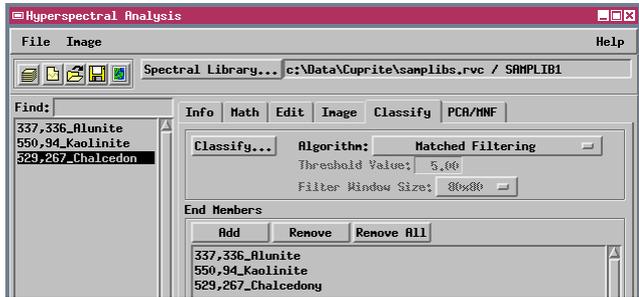
Locations and spectral plots of image spectra included in the SAMPLIB1 spectral library.

Matched Filtering

STEPS

- select Matched Filtering from the Algorithm option menu
- select spectrum 337,336_Alunite
- press [Add] to add this spectrum to the End Member list
- repeat the last two steps for the remaining two spectra in the library
- press [Classify]
- use the standard Select Objects dialog to name each output raster and direct them to a new Project File (processing may take several minutes or more)
- use the Layer Manager for the Hyperspectral Image window to view the output rasters

The Matched Filtering algorithm uses a Constrained Energy Minimization technique to assess the spectral composition of each image cell. Each image spectrum is assumed to be a linear mixture of a target spectrum and multiple unknown spectral signatures. The process identifies what proportion (if any) of each composite image spectrum could be produced by the target spectrum (end member). A perfect match yields an output value of 1, while poor matches yield low positive or even negative values. The process creates a separate floating-point output raster for each target spectrum (end member).



Matched Filtering result for spectrum 337,336_Alunite with linear contrast enhancement. Brightness indicates degree of match to the target spectrum; only the brightest areas indicate a confident identification of the target material.

Two variants of the standard matched filtering procedure are also available on the menu: Derivative Matched Filtering and Locally Adaptive Matched Filtering. The Derivative method uses the first derivative (slope) of the spectral curves rather than the actual spectral values, providing a similarity check that is more sensitive to curve shape and less sensitive to overall brightness differences. The Locally Adaptive method determines a local unknown background signature for each image cell, and is best suited for identifying occurrences of materials that are relatively rare in the scene.

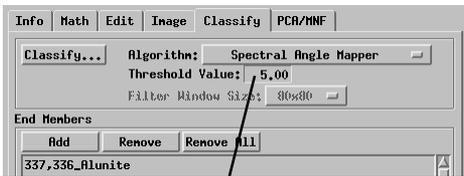
Spectral Angle Mapper

The Spectral Angle Mapper treats target spectra and image spectra as vectors in n-dimensional spectral space. Each spectrum defines a point in spectral space, and this point can also be treated as the end of a vector that begins at the origin of the coordinate system. The angle between a pair of vectors is a measure of the similarity of the spectra; smaller spectral angles indicate greater similarity. This method is insensitive to differences in average brightness between spectra that may be due to topographic or sensor gain effects, because these factors change the length of the spectral vector, but not its orientation.

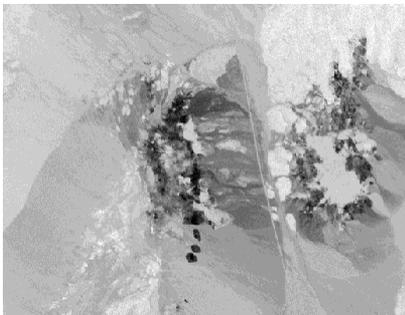
This process produces two output rasters: a Spectral Angle raster that contains values of the spectral angle for each image cell, and a Class raster in which cells are assigned to end-member classes based on the angle value that you set for the Threshold Value parameter. If you select multiple end-member spectra, the spectral angle is calculated between an image spectrum and each end member spectrum, and the minimum value is shown in the Spectral Angle raster. However, this raster is more readily interpreted if only a single end member is used.

STEPS

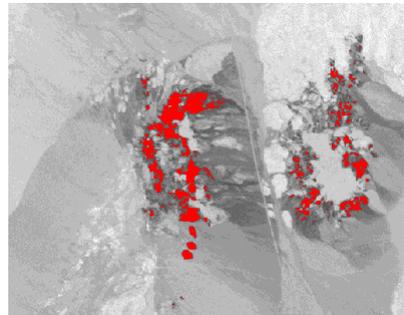
- select the Kaolinite image spectrum in the End Member list and press [Remove]
- repeat for the Chalcodony image spectrum
- select Spectral Angle Mapper from the Algorithm option menu
- press [Classify]
- name each output raster and direct them to your classification Project File



Cells with a spectral angle value less than the Threshold Value are assigned to the endmember class in the Class raster. You will need to experiment to find the appropriate value for a particular target spectrum.



Portion of the spectral angle raster for the target Alunite spectrum. Low values (darkest tones) indicate areas with the best spectral match.



Class raster displayed over the Spectral Angle raster, with cells assigned to the end member class shown in red.

Cross-Correlation

STEPS

- remove the Alunite image spectrum from the End Member list
- add the Kaolinite image spectrum to the End Member list
- select Cross Correlation from the Algorithm option menu
- enter 0.87 in the Threshold Value parameter field
- press [Classify]
- name each output raster and direct them to your classification Project File

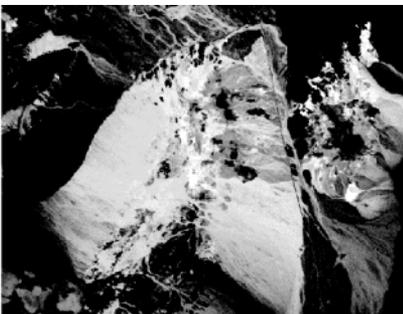
The Cross-Correlation matching method calculates the linear correlation coefficient between each image spectrum and the target spectrum for different match positions (spectral band shifts). The correlation values in the best matching position for each image cell are shown in the Correlation raster (values between 0 and 1, with 1 indicating a perfect match). Cells in the Class raster are assigned to an end-member class using the correlation value that you specify for the Threshold Value parameter. This method is relatively insensitive to differences in average brightness.

Comparing the results of these classification exercises, note that the output rasters produced by the

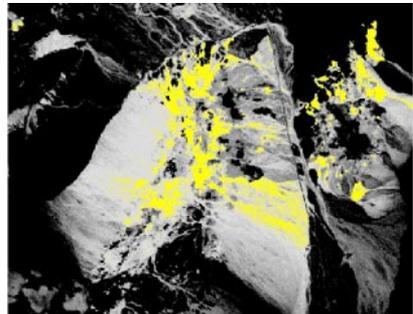
Cross-Correlation and Spectral Angle Mapper methods look like a coherent image, while much of the Matched Filtering (MF) image appears “blurry”. The MF results are actually the best, because most of



the scene has low match values for a given target spectrum, and there is little spatial correlation between these low values. The MF method does a better job of recognizing the target spectrum within various composite (mixed material) image spectra.



Correlation raster for the target kaolinite image spectrum with auto-linear contrast enhancement. High values (bright tones) indicate areas with the best spectral match.



Class raster (threshold value = 0.87) displayed over the Correlation raster, with cells assigned to the end-member class shown in yellow.

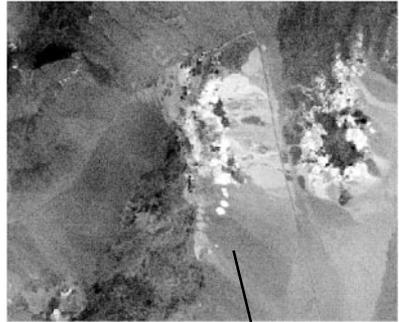
Linear Unmixing

The Linear Unmixing method uses a suite of end-member spectra to assess the spectral composition of each image cell. Cells are assumed to contain varying proportions of the selected end-member materials, and the resulting image spectra are assumed to be linear combinations of the end-member spectra. The process produces a raster for each end-member that shows the fractional abundance of that end-member material for each image cell. The process also creates an Error raster.

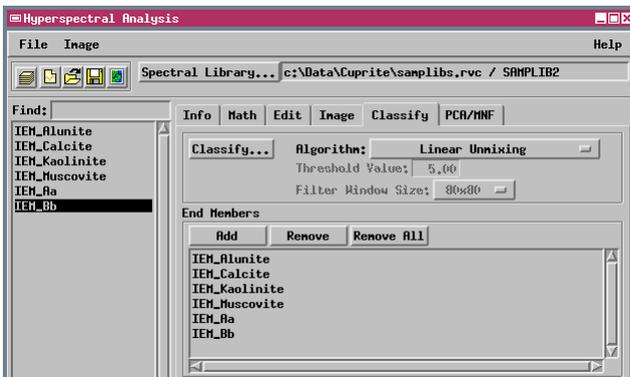
Selection of appropriate end-member spectra is important to the success of this process. The selected endmembers should represent “pure” (single-material) cells, if they exist in the image. Procedures for identifying such image cells are outlined on pages 21-27. These procedures were used to identify the spectra included in spectral library SAMPLIB2. On the basis of spectral shape and published information about the Cuprite area, four of these spectra can be attributed to single minerals (alunite, calcite, kaolinite, and muscovite). The other two spectra in the library probably represent spatial mixtures of two or more minerals that do not occur in pure form throughout any single image cell in the scene.

STEPS

- remove the Kaolinite image spectrum from the End Member list
- select Linear Unmixing from the Algorithm option menu
- click [Spectral Library...]
- select object SAMPLIB2 from the SAMPLIBS Project File
- add all six spectra to the End Member list
- press [Classify]
- name each output raster and direct them to your classification Project File

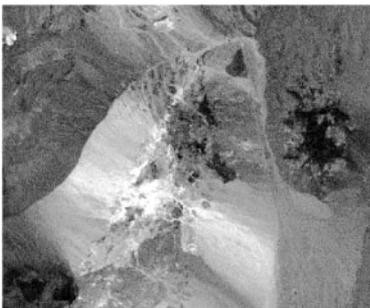
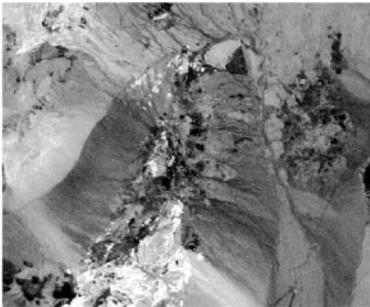
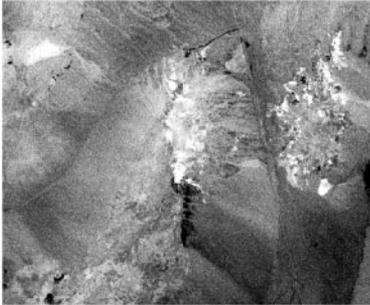


Fraction image for the IEM-Alunite endmember displayed with autonormalized contrast enhancement.



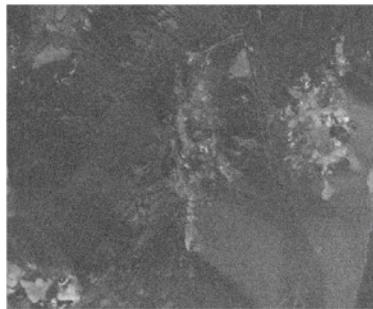
Assessing Unmixing Results

Fraction images for the IEM_Kaolinite (top), IEM_Calcite (middle), and IEM_Muscovite (bottom), endmembers, all displayed with autonormalized contrast enhancement.



In the ideal linear unmixing case, fraction values for individual endmembers should range from 0 to 1, and the sum of the fractional abundances for each image cell should equal 1. But if you supply too few endmembers, or inappropriate endmembers, the fractional abundances for each cell may not sum to the desired value of 1. If a selected endmember is a mixed cell, and “pure” cells representing its components are present in the image, the pure cells will be assigned negative abundance values or values greater than one for some end members. If there are a large number of such “out of range” values in the fraction images, you may need to add or replace endmember spectra.

The linear unmixing equations include an error term that is minimized to identify the best-fit set of endmember abundance values. The computed error value for each image cell is incorporated in the Error raster. Bright cells in the Error raster (high error values) indicate areas that were not adequately modeled by the selected set of endmembers, so you may need to concentrate your search for additional endmembers in those areas.



Error raster for the Cuprite linear unmixing run using the endmember spectra in `SAMPLIB2` (displayed with autoexponential contrast enhancement). Bright areas probably contain materials not included in the endmember set.

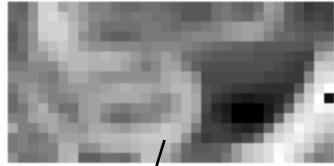
Self-Organizing Map Classifier

The Self-Organizing Map Classifier performs an empirical (unsupervised) classification of the hyperspectral image, so you don't have to specify target spectra. Like the Spectral Angle Mapper, image spectra are treated as points in n-dimensional space. The process uses neural network techniques to find a best-fit set of 512 class-centers, then assigns all image spectra to one of these classes using spectral angle as the measure of similarity.

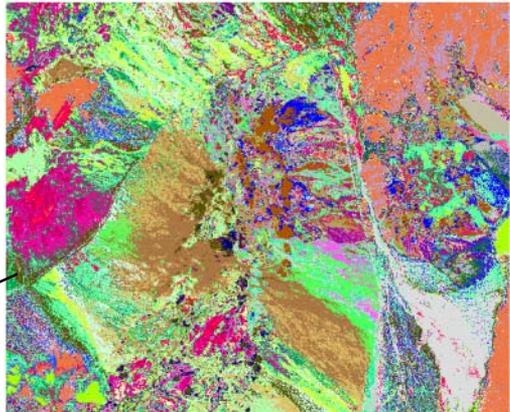
The neural network used to determine the class-center vectors is a square 16×32 array of nodes, each representing one class. Initially the spectral values for each node are assigned randomly (producing randomly-oriented class vectors). A large sample of image spectra are compared one by one to the full set of nodes. For each test spectrum, the node with the most similar spectrum has its spectral values adjusted to improve the match. In addition, nearby nodes in the array are also adjusted to a lesser extent. After many sample spectra have been processed, the node values converge to a set of class-center vectors that approximate the distribution of all image spectra in N-dimensional space. Similar class vectors also lie close together in the node array, and the greater the area that similar materials cover in the hyperspectral image, the more classes are used to represent them. This property of the classifier ensures adequate spectral discrimination of different varieties of common, widespread materials.

STEPS

- select Self Organizing Map Classifier from the Image menu
- use the standard Select Objects dialog to name each output raster and direct them to your classification Project File



The Distance raster is a graphical representation of the final state of the Self-Organizing Map, the 16×32 array of nodes used to determine class centers. Each cell represents one class, and the cell value is the average spectral distance to neighboring classes. Dark areas represent tightly grouped class centers.



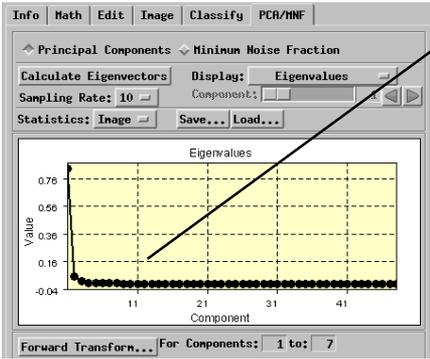
Class raster with 512 classes created by the Self-Organizing Map classifier, shown with a random color map.

Principal Components Transform

STEPS

- click the PCA/MNF tab
- click [Calculate Eigenvectors] (processing will take several minutes)
- choose each of the four options on the Display option menu and study the resulting plots
- press [Forward Transform] to create a set of PC component rasters
- use the standard Select Objects procedure to direct the output rasters to a new Project File
- use the Layer Manager for the Hyperspectral Image window to view the PCA rasters

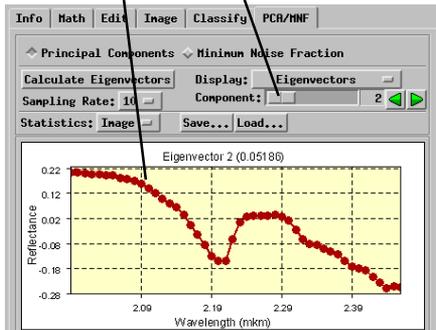
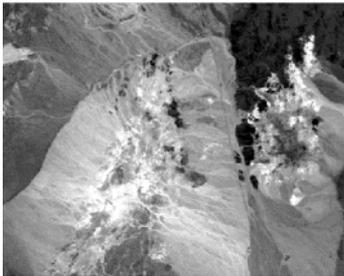
Adjacent hyperspectral image bands are visually and numerically similar, and therefore contain much redundant information. The Principal Components transform is a standard method for deriving a new set of images with reduced spectral redundancy. The process is a linear transformation that projects each image cell's spectrum to a new set of orthogonal coordinate axes. These axes are chosen so that the output images are uncorrelated and ordered by decreasing variance, with the first principal component axis corresponding to the direction of maximum variance in spectral space. Important image information is generally concentrated in the low-order components, while noise increases with increasing component number. Use of low-order PC rasters in place of original image bands can speed visual analysis and classification of the hyperspectral image.



The *eigenvalue* plot shows the image variance for each component axis. The sum of the eigenvalues equals the sum of the band variances in the original image. You can also display plots of Component Variance and Band Variance (in percentage of total variance).

The *eigenvector* plot for each principal component shows the set of weighting coefficients (one for each image band) used in the coordinate transformation for that component. Use the slider or arrow buttons to choose the desired component.

Second principal component raster



Minimum Noise Fraction Transform

If bands in a hyperspectral image have differing amounts of noise, standard principal components derived from them may not show the usual trend of steadily increasing noise with increasing component number. The Minimum Noise Fraction transform (MNF) is a modified version of the Principal Components transform that orders the output components by decreasing signal to noise ratio.

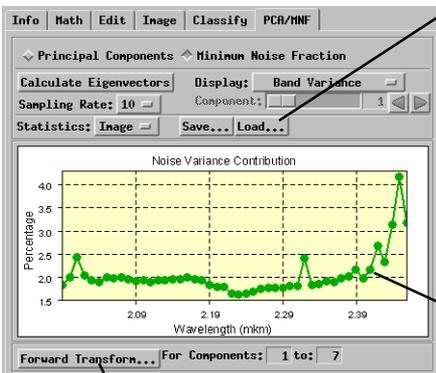
The MNF procedure first estimates the noise in each image band using the spatial variations in brightness values. It then applies two successive principal component transforms. The first uses the noise estimates to transform the dataset to a coordinate system in which the noise is uncorrelated and is equal in each component. Then a standard principal components transform is applied to the noise-adjusted data, with output components ordered by decreasing variance. This procedure produces a component set in which noise levels increase uniformly with increasing component number. The low-order components should contain most of the image information and little image noise.

STEPS

- click the Minimum Noise Fraction radio button at the top of the PCA/MNF panel
- click [Calculate Eigenvectors]
- display the various plots



RGBI display of the first 4 MNF components (3 = Red, 2 = Green, 4 = Blue, 1 = Intensity).



You can use the Save and Load buttons to save and reload the statistics calculated by the Principal Components and MNF transforms. For example, you might calculate the transform in one session, and reload the statistics later to create component rasters. For the MNF transform, you can save the image statistics and noise statistics using the options on the Statistics button.

The plots produced by the MNF transform show the noise statistics calculated from the original input bands.

Creating MNF component rasters using the Forward Transform option can require several hours or more, depending on the number of input components. You can find the first 5 MNF component rasters for the Cuprite scene in the MNF_CUP Project File.

Pixel Purity Index

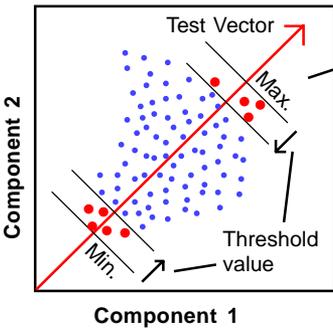
STEPS

- click the Open icon button and select RVC Image List from the dropdown menu
- select all raster objects in the MNF_CUP Project File
- on the Image panel, select None from the option menus for both Additive Offset and Atmospheric Correction calibration methods
- select Compute Pixel Purity Index from the Image menu
- in the Pixel Purity Index window, select 200 from the Test Directions Per Pass option menu, and change the value in the Threshold Value field to 0.10
- click [Run] and direct the output PPI raster to a new Project File



The set of low-order MNF components provides a “distilled” version of the hyperspectral image that can be used to rapidly identify relatively “pure” image spectra for use as spectral endmembers in the Linear Unmixing operation. The Pixel Purity Index (PPI) operation is the first step in identifying these endmember spectra.

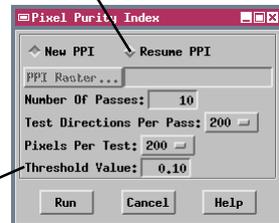
When image spectra are plotted as points in n-dimensional space, endmember spectra should lie at cusps along the margins of the data cloud. (This is true regardless of whether we use the original spectral bands or the transformed MNF components). The PPI creates a large number of randomly-oriented test vectors emanating from the origin of the coordinate space. The spectral points are projected onto each test vector and spectra within a threshold distance of the minimum and maximum projected values are flagged as extreme. As directions are tested, the process tallies the number of times an image cell is found to be extreme. Cells with high values in the resulting PPI raster should correspond primarily to the locations of “edge” spectra in the image.



A two-dimensional illustration of how the Pixel Purity Index identifies potential extreme image spectra (large red dots).

You can select both the number of passes that the PPI operation makes through the full set of image spectra and the number of test directions examined per pass. You can use the Resume PPI toggle to rerun and add to an existing PPI raster, continuing until the number of new cells identified per pass falls to a value close to 0.

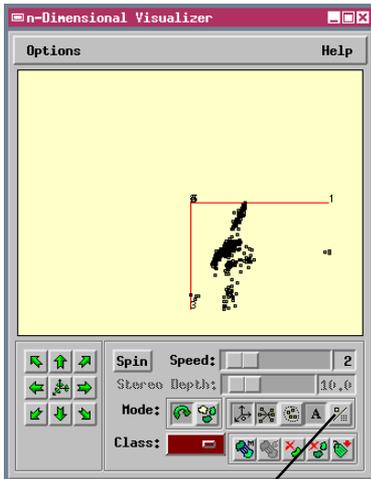
When the PPI operation is run on raster sets with integer values, the default threshold value of 1.00 corresponds to a distance of 1 digital number. When the PPI is run on floating-point rasters with non-integer values, like MNF component sets, the default value is interpreted as 1 Standard Deviation. A threshold close to 0.10 is more appropriate for MNF component sets.



Start the n-Dimensional Visualizer

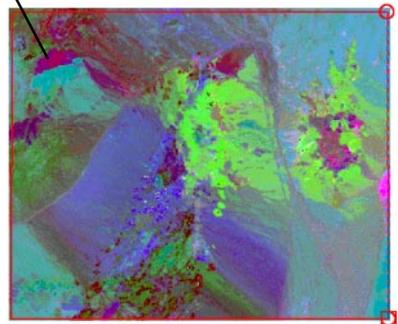
To assess the results of the PPI operation we will use the n-Dimensional Visualizer tool. This tool allows you to create and view an n-dimensional scatterplot of spectra for an area in the hyperspectral image or MNF component set. You can view the scatterplot from different viewpoints and rotate the plot manually or automatically in real time. You designate the selected area by drawing a polygon in the Hyperspectral Image window. You can use this tool to investigate the spectral properties of different materials in the image and to search for image cells with extreme spectral values that might represent pure endmembers.

When you use the PPI raster as a mask, only cells with PPI values above the specified threshold value (and within the selection polygon) are displayed in the scatterplot. The threshold eliminates cells that were only flagged as extreme in a few spectral directions. When no mask is used, you should use a smaller selection polygon to obtain a manageable set of spectral points.



The Style icon button toggles the point display style between two modes: single pixel (button in) and box (shown here).

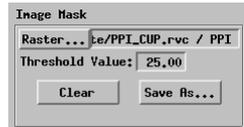
The grayscale PPI mask raster is displayed over the MNF image with 0-value (null) cells transparent.



This exercise continues on the next page.

STEPS

- in the Image Mask section of the Info panel, click [Raster...]
- select the PPI raster you made in the previous exercise, or the one provided in the PPI_CUP Project File
- change the Threshold Value field for the image mask to 25.00



- click the n-Dimensional Visualizer icon button in the Hyperspectral Image window 
- use the Line / Polygon Edit Controls to draw a polygon that includes most of the image, as shown below
- click [Apply] (or press the right mouse button)
- make sure the Style icon button on the Visualizer window is toggled out 
- turn on the Label Axes icon button 

Rotate the n-Dimensional Scatterplot

STEPS

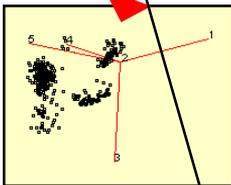
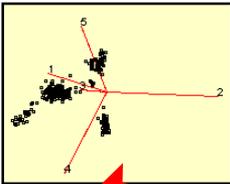
- ☑ move the Speed slider on the Visualizer window to speed 2
- ☑ press [Spin]; the button changes to read Stop
- ☑ after watching the plot automatically rotate for a while, press [Stop]
- ☑ use the direction arrow icon buttons to rotate the plot by increments
- ☑ press the Change Viewpoint icon  button several times
- ☑ place the mouse cursor in the view pane of the Visualizer window, hold down the left mouse button, and drag in the desired direction to rotate the plot manually

This exercise continues on the next page.

You can use the n-Dimensional Visualizer to find and mark the most extreme cells among those flagged by the PPI operation. A typical n-dimensional scatterplot shows better separation of point clusters when viewed from certain spectral directions than from others. Usually you need to rotate the plot in order to find these directions and to understand the distribution of points. The rotation options in the n-Dimensional Visualizer provide a great deal of flexibility, including manual, incremental, and automatic rotation modes.

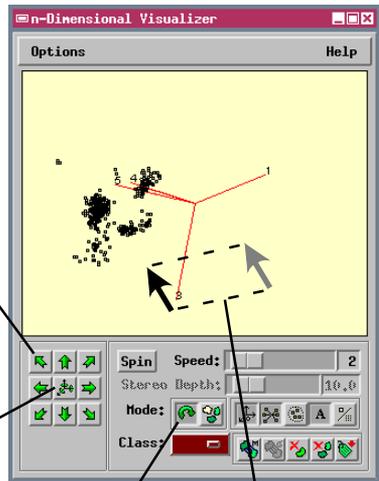
The Spin button starts an animated rotation of the scatterplot in successive, randomly-chosen directions. You can use the Speed slider to vary the rotation speed. The rotation speed is also dependent on the number of spectral points and the number of dimensions. The animation works best with a small number of dimensions (less than 15). You can also rotate the plot by small increments in a specific direction using the arrow icon buttons. With the window mode set to Rotate, you can rotate the plot

manually in any direction by dragging the cursor in the Visualizer window.



The arrow icon buttons rotate the plot by a small, constant amount in the specified direction.

The Change Viewpoint icon button alters the viewing direction by 90 degrees. Clicking it several times allows you to rapidly view the plot from a number of different directions.



With Rotate mode turned on you can rotate the plot in any direction by dragging the cursor in the Visualizer window.

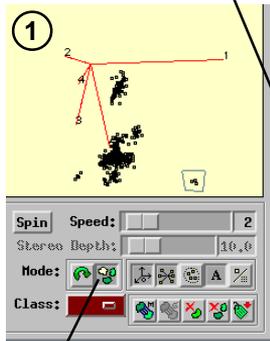
Mark Endmember Classes

As you rotate the n-dimensional scatterplot, look for clusters of points that define the edges of the distribution in various directions. In some cases these are isolated small clumps of points, in others the ends of elongated clusters of points that define a mixing line between two endmembers. Stop rotating when you can distinguish several edge clusters. The n-Dimensional Visualizer allows you to select each such cluster of spectral points and group them into a class. The points you assign to a class are shown in the window in the selected color. You can repeat the grouping operation for other points to assign them to the same class, or choose a new color to create additional classes. Class colors are maintained when you rotate the plot so that you can better visualize the distribution of the spectral points. The Show Classes option enhances visibility of classes when only a few spectral points are assigned to each. It uses solid lines to connect the spectral points in the class to the class mean.

STEPS

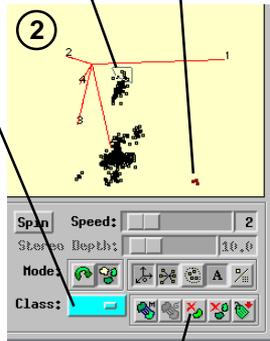
- when you have identified several edge clusters, press the Classify icon button to change the cursor mode 
- make sure that the Show Classes icon button is toggled in 
- left-click in the plot area to place vertices defining a polygon that surrounds a point cluster (as in illustration 1 below)
- right-click to accept the polygon (illustration 2)
- click on the colored button labeled Class and choose another color such as cyan
- draw a polygon around another cluster of points (2) and right click (3)
- press [Spin] and observe how the marked clusters relate to the remaining points
- press [Stop]

Use the Class color option menu to select a color for marking a class.

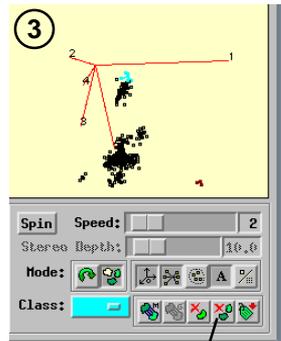


In Classify mode you can use the mouse cursor to draw a polygon to select a cluster of points to assign to a class.

Example first class
Polygon for second class



Press the Clear Class icon button to release all points assigned to the currently selected class color.



Press the Clear All Classes icon button to release all classified points from their respective classes.

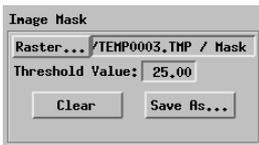
This exercise continues on the next page.

Save Endmember Classes as a Mask

STEPS

- continue rotating the plot and grouping points until you have marked six or seven classes
- when you are satisfied with the set of classes, press the Classes to Source icon  button on the Visualizer window
- click [Yes] in the first Verify window that opens (to create a new mask raster containing the class cells)
- click [No] in the second Verify window that asks whether to save the current mask raster (no changes have been made that require saving)

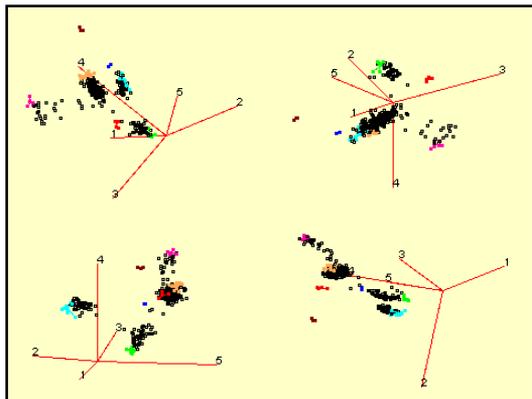
A temporary mask raster showing the locations of cells in the marked classes has now been substituted for the PPI raster you have been using as a mask.



- to save the temporary class mask, click [Save As]
- save the class mask in the same Project File as the PPI raster you created earlier

Continue changing viewpoints and rotating the plot to find additional endmember clusters. As you view the plot, resist the natural impulse to visualize it as a rotating three-dimensional volume. At any moment the plot is showing a two-dimensional projection of an n-dimensional hyperspace (5 dimensions in this example). The spatial relationships of point clusters can only be determined by examining projections from many viewpoints. For example, a point cluster that always appears to lie on a straight line between the same two endmember clusters is probably a mixture of the two, rather than an additional endmember.

The Classes to Source option allows you transfer the class cell locations to the current mask raster, or to a new mask raster. Each class is assigned a unique raster value, and a Colormap subobject containing the class colors is automatically created. A new mask created by this procedure is a temporary object that is automatically substituted for the current mask. To reuse this object you must save it using the Save As button on the Image Mask section of the Info panel.



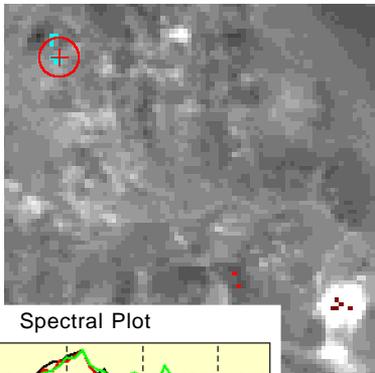
Example endmember classes for the Cuprite MNF component set shown from different viewpoints in the n-Dimensional Visualizer. Your classes may be larger or smaller or use different colors. Note that a given endmember cluster lies on the edge only when seen from certain viewpoints.

Select Endmember Spectra

In the final step in defining image endmembers, we can use the endmember class raster as a mask with the original Cuprite hyperspectral image to select and save endmember spectra. The colored cells in the mask identify the locations of candidate image spectra. Other cells in the mask are transparent, revealing the underlying image. To enhance visibility of the colored class cells, you can add one of the grayscale MNF component rasters as a reference layer. You may also want to use the controls in the Layer Controls window to hide temporarily all layers except the mask in order to locate all of the class cells.

For each endmember class, use the Spectral Profile tool to extract and display the image spectra for the individual class cells. (Use the Multiple Spectra option for the Spectral Plot window.) Choose the most representative spectrum and save it in a spectral library. Repeat until you have saved a spectrum for each endmember class.

Zoomed view of part of the northeast quadrant of the Cuprite scene showing three spatial clusters of class cells representing different endmember classes.



Plot of several individual spectra for the endmember class represented by the color cyan.

STEPS

- repeat the steps on page 4 to reopen the CUPRITE97 hypercube object 
- on the Image panel, select Subset from the Use Wavelength Range option menu and set the range to 1.98 to 2.46
- In the Layer Manager window, click the Add Raster icon  and select Single from the dropdown menu
- select MNF_1 from the MNF_CUP Project File
- in the Visualizer Image Mask section of the Info panel, click [Raster...]
- select the Classes mask raster you saved in the previous exercise, or the one provided in the PPI_CUP Project File
- use the tools in the Hyperspectral Image window to zoom in on the northeast quadrant of the scene and locate one of the collections of colored mask cells that mark the members of an endmember class
- click on the Spectral Profile icon button on the Hyperspectral Image window 
- on the Info panel, turn on the Display Spectral Plot toggle button
- move the cursor to the individual class cells and right-click to view their spectra

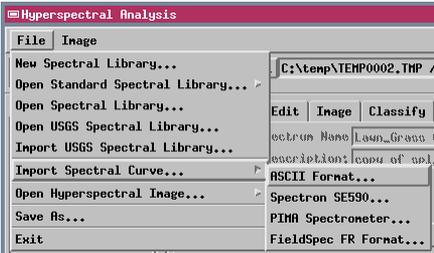
Import a Field or Laboratory Spectrum

STEPS

- in the Layer Manager window, right-click on the entry for the MNF reference layer and select Remove Layer from the dropdown menu
- in the Image Mask section of the Info panel, click [Clear..]
- on the Image panel, select Full from the Wavelength Range option menu
- click the New Spectral Library icon button



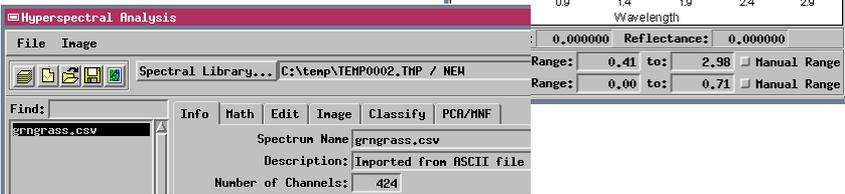
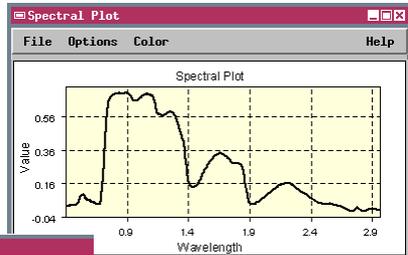
Spectral reflectance is the ratio of radiant energy reflected by a material at a particular wavelength to the energy incident on the material at that wavelength. Variations in spectral reflectance with wavelength constitute a material's **reflectance spectrum**, and arise from its chemical make-up and physical structure. Instruments called spectrometers are used to measure reflectance spectra of materials such as soil, rock, and vegetation. Reflectance spectra can be measured in the field under natural illumination, or in the laboratory under more controlled conditions. Field and laboratory spectra provide known spectral signatures that can be used to identify different materials in a hyperspectral image.



- from the File menu on the Hyperspectral Analysis window select Import Spectral Curve, then ASCII Format
- use the standard Select Files window to navigate to and select file GRNGRASS.CSV
- turn on the Display Spectral Plot toggle button and select the new library spectrum

You can import spectra from external sources into a TNTmips spectral library by using the Import Spectral Curve procedure in the Hyperspectral Analysis process. You can choose from several common spectrometer file formats, or select the generic ASCII format. ASCII files must be in the .csv

format described in the earlier exercise entitled "Save an Image Spectrum as Text". Each spectrum must be in a separate external file.



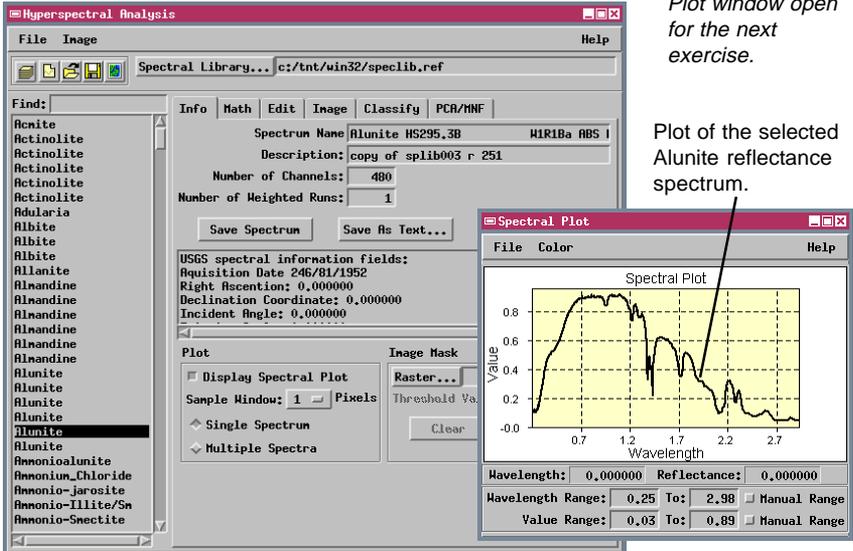
Open the USGS Spectral Library

The United States Geological Survey's Spectroscopy Lab has measured the spectral reflectance of about 500 minerals (basic components of rocks and soil) under standard laboratory conditions. They have compiled the measurements in the USGS Spectral Library, which is freely available to the public. MicroImages has imported the USGS Spectral Library and incorporated it within the Hyperspectral Analysis process as a Standard Spectral Library. Other public spectral libraries may be added to the Standard library list in the future.

The reflectance spectra in the USGS Spectral Library cover a wavelength range from 0.2 to 3.0 micrometers, with spectral bandwidth less than 0.01 micrometers (comparable to the spectral resolution of an AVIRIS image spectrum). In addition to mineral spectra, the USGS spectral library also includes a few spectra acquired from different types of vegetation foliage. The vegetation spectra are grouped together following the alphabetical listing of mineral spectra.

- ☑ click the Open Standard Spectral Library icon button on the Hyperspectral Analysis window 
- ☑ open the Spectral Plot window if it is not currently open, and turn on the Single Spectrum toggle
- ☑ in the spectrum list, select the next-to-last Alunite spectrum

You can find additional information about the USGS Imaging Spectroscopy program and the USGS Spectral Library on the World Wide Web at: <http://speclab.cr.usgs.gov>



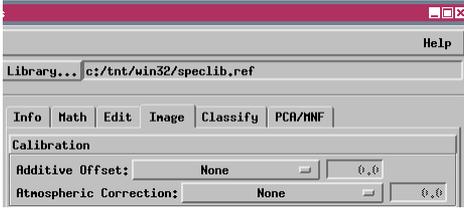
Keep the Spectral Plot window open for the next exercise.

Plot of the selected Alunite reflectance spectrum.

Correction to Reflectance

STEPS

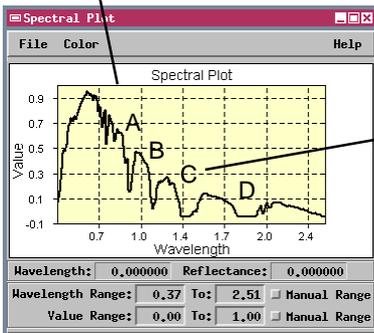
- on the Image panel, choose None from both the Additive Offset and the Atmospheric Correction option menus



- select the Spectral Profile tool from the Hyperspectral Image window toolbar
- obtain an image spectrum from the Stonewall Playa area near the east edge of the image (see page 4)



Raw radiance spectrum for Stonewall Playa. The overall shape of the spectrum, with maximum brightness in the visible light range, replicates the solar irradiance curve and the dropoff in atmospheric scattering with increasing wavelength.



An airborne or satellite-mounted hyperspectral imager measures **spectral radiance**: the upwelling radiant energy at the sensor for many narrow wavelength bands. These radiance values are only partly determined by the spectral reflectance

properties of surface materials in the scene. A number of other factors contribute to the measured radiance, with effects that vary with wavelength. For example, the amount of solar energy illuminating the surface varies with wavelength

(described as the solar irradiance curve) and with the sun's height in the sky when the image was acquired (solar zenith angle). Other factors include wavelength-selective absorption and scattering by atmospheric gases and particulates, and illumination differences due to topography (direction and angle of slope, and shadowing).

In order to make meaningful comparisons between image spectra and laboratory reflectance spectra, the image radiance values must be corrected (calibrated) to reflectance by removing these other contributing effects. The Hyperspectral Analysis process in TNTmips offers several commonly-used empirical correction methods. These methods are applied to the raw hyperspectral image on-the-fly, so you do not need to recompute and store different versions of the image. The steps in this exercise disable these calibrations so that you can examine raw radiance spectra.

Several major lows in the radiance spectrum (labeled A through D) result from strong atmospheric absorption at specific wavelengths. If the atmosphere was uniform over the scene and elevation differences are not extreme, raw spectra from all image cells should include about the same effects from atmospheric scattering and absorption, as well as the same solar irradiance spectrum.

Flat Field Correction

The Flat Field correction method reduces atmospheric and solar irradiance effects using the mean spectrum of an area you designate with the Flat Field Calibration tool. The area you select should be topographically and spectrally flat: it should have a uniform spectral reflectance at all wavelengths, without significant absorption features. The average spectrum from such a “flat field” will be an almost pure signature of solar irradiance plus superimposed atmospheric scattering and absorption effects. The Flat Field method divides each image spectrum by the flat field spectrum. With an appropriate flat field spectrum, this procedure will largely remove atmospheric and solar irradiance effects. (The selected area should be bright in order to reduce the effects of random image noise on the correction.) If the scene includes large variations in elevation or the atmosphere was not uniform over the scene, some residual topographic and atmospheric effects will remain in the calibrated image spectra.

Choosing an appropriate flat field is obviously a critical step in this correction method. No natural material has a completely flat reflectance spectrum. Bright man-made materials such as concrete have a fairly flat spectral response and thus can provide appropriate targets in hyperspectral scenes of urban areas. In the Cuprite scene, the bright, salt-encrusted surface of Stonewall Playa has no significant mineral absorption features, so it provides an approximate flat field spectrum.

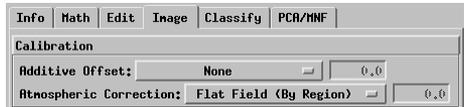
The Maximum correction method assembles an approximate flat field spectrum automatically using the maximum value in each spectral band. Image spectra are divided by the maximum spectrum to produce the corrected image spectra. This method may be appropriate if the scene includes a variety of materials of differing overall brightness.

STEPS

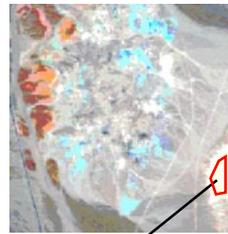
- use the standard View window tools to zoom in on Stonewall Playa
- select the Flat Field Calibration  tool from the Hyperspectral Image window toolbar
- use the standard Line/Polygon Edit Controls to draw a polygon covering the brightest part of Stonewall Playa (see illustration)



- choose Flat Field (By Region) from the Atmospheric Correction option menu



- reacquire an image spectrum from Stonewall Playa



Draw a polygon with the Flat Field Calibration tool to indicate the area to be used to determine the flat field spectrum.

Equal Area Normalization

STEPS

- select Equal Area Normalization from the Atmospheric Correction option menu
- select Minimum from the Additive Offset Calibration option menu
- reacquire an image spectrum from Stonewall Playa



The methods on the Additive Offset Calibration option menu correct for atmospheric backscatter and instrument artifacts that raise the measured radiance values by a constant amount for each cell in an image band. The default Minimum method determines the minimum value in each band and subtracts it from all cells in the band. The Dark Field Calibration subtracts from each cell the average determined from an area that you outline using the Dark Field Calibration tool on the Hyperspectral Image window. The area you select should be dark in all bands; a deep water body or heavily shadowed area would be the best target.

Equal Area Normalization is the default method of atmospheric correction in Hyperspectral Analysis; it is enabled automatically when you start the process. This calibration method does not require any knowledge of surface materials in the scene. The radiance values in each image spectrum are first scaled so that the sum of the values for each image cell is constant over the entire scene. This procedure shifts all image spectra to nearly the same relative brightness, removing differences in overall brightness between materials as well as illumination differences caused by topography. An average spectrum for the entire scene is then calculated from the normalized spectra. For a scene that includes many different land cover types, averaging may largely remove spectral features attributed to the surface materials. The average spectrum should then resemble the flat field spectrum described in the previous exercise. Finally, each normalized image spectrum is divided by the average spectrum. The resulting spectral values represent reflectance relative to the average spectrum, and in ideal cases should be comparable to true reflectance spectra.

Large elevation differences or variations in the atmosphere across the scene may cause atmospheric effects on the image spectra that are not removed by Equal Area Normalization. In addition, if there are a few widespread materials in the scene, the average spectrum may include absorption features related to these surface materials instead of representing only atmospheric and solar irradiance effects. In these cases the calibrated spectra produced by Equal Area Normalization may differ significantly from true reflectance spectra.

NOTE: If the hyperspectral image has already been processed to reflectance, you can choose the Manual options on the two Calibration menus and set numeric values for offset and scaling. For example, sample AVIRIS scenes calibrated to reflectance by NASA use integer values from 0 to 10,000 to correspond to the reflectance range 0 to 1.0. In this case set the Additive Offset value to 0 and the Atmospheric Correction value to 10,000.

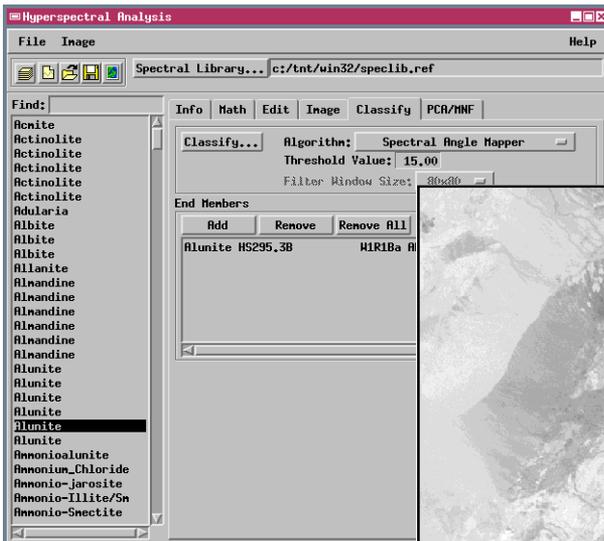
Classification with Reflectance Spectra

Once the hyperspectral image has been calibrated to reflectance, you can use reflectance spectra of the materials in the USGS Spectral Library (or any other library of reflectance spectra) to locate and classify image cells with similar spectra. You use the same classification methods and follow the same procedures outlined in the previous exercises in this booklet on spectral classification, but select library reflectance spectra instead of archived image spectra as end members. The spectral classifiers allow you to produce a series of material maps (class rasters) showing the distribution of image cells matching the target spectrum within the threshold you have set.

With the empirical calibration methods currently implemented, the Spectral Angle Mapper and Cross Correlation methods produce the best results using reflectance spectra. For the Cuprite scene, the default Equal Area Normalization does not appear to achieve a sufficiently accurate reflectance calibration to produce good results with the Matched Filtering method using reflectance spectra.

STEPS

- press the Full icon button on the Hyperspectral Image window 
- select Subset from the Use Wavelength Range option menu on the Image panel, and set the wavelength range from 1.98 to 2.46
- in the spectrum list for the spectrum list for the USGS Spectral Library, reselect the next-to-last Alunite spectrum
- press [Add] on the Classify tabbed panel to add the Alunite spectrum to the End Member list
- use the default Spectral Angle Mapper algorithm, and set the Threshold Value to 15
- press [Classify] and direct the output rasters to your classification Project File



Class raster displayed over the Spectral Angle raster, with cells assigned to the Alunite class shown in red.

Spectral Matching

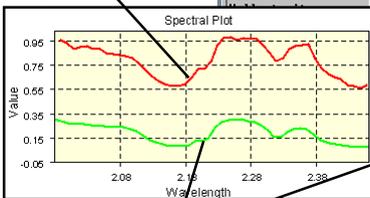
STEPS

- ☑ click [Spectral Library...] on the Hyperspectral Analysis window
- ☑ select object SAMPLIB3 from the SAMPLIBS Project File
- ☑ in the Plot controls on the Info panel, turn on the Multiple Spectra radio button
- ☑ in the Spectral Plot window, turn on the upper Manual Range toggle button and set the Wavelength Range from 1.98 to 2.46
- ☑ scroll to the bottom of the spectrum list and select 337,336_Alunite
- ☑ click the Math tab, and select Band Mapping from the Algorithm option menu in the Spectral Matching control panel
- ☑ press [Find Match]
- ☑ select the first spectrum in the list of matching spectra

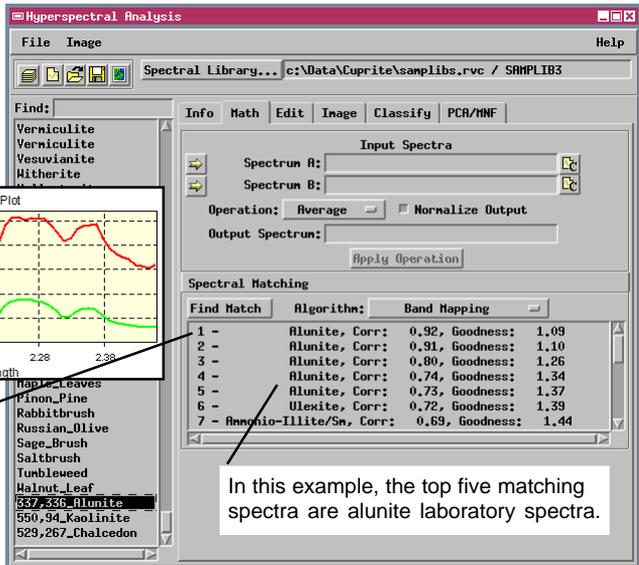
The spectral matching procedure is useful for identifying specific unknown materials in a hyperspectral image. You can save an image spectrum for the target material into a copy of a reference spectral library, then search the library for the reflectance spectra that best match the unknown spectrum. (In this exercise you use a library containing the reflectance spectra in the USGS Spectral Library along with three Cuprite image spectra.) The matching process finds the twenty spectra that best match the target spectrum and ranks them by closeness of fit.

Two spectral matching methods are available. The Spectral Angle Mapper method is identical to the classification method described previously. The Band Mapping method compares spectra on the basis of the positions and relative depths of absorption features (local minima in the spectral curves). Both methods are relatively insensitive to differences in average brightness (note the large vertical separation between the target spectrum and best matching spectrum in the spectral plot).

337,336_Alunite image spectrum.



Best matching spectrum from the USGS Spectral Library: an alunite reflectance spectrum.



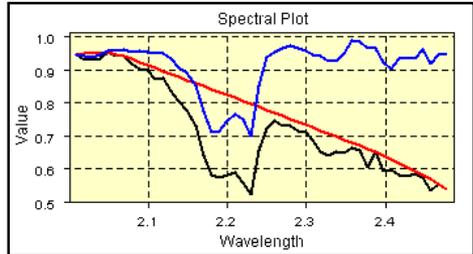
In this example, the top five matching spectra are alunite laboratory spectra.

Remove Continuum from Image

Many minerals can be recognized by the position and shape of one or a few narrow absorption features in a particular part of the spectrum. These features (commonly called “bands”) appear as valleys in the spectrum. In a composite spectrum from a hyperspectral image, the shape of the spectral curve surrounding an absorption band may vary. This broader curve shape, or *continuum*, is the aggregate response of all the materials that contribute to the spectrum. A sloping continuum causes an apparent shift in the location of an absorption band’s reflectance minimum. Spectral matching of individual absorption bands (such as in the Band Mapping method) is more accurate if the continuum is removed from the spectra prior to matching. This involves dividing the spectral value in each channel by the continuum value.

You can use the Remove Continuum option on the Image menu to create a continuum-removed version of the hyperspectral image (for the current wavelength subset). The operation computes the continuum for each image cell, then creates and saves the corresponding continuum-removed spectrum for the cell. Spectra you extract from this processed image can then be directly compared to continuum-removed library spectra. The next exercise discusses how to remove the continuum from library spectra.

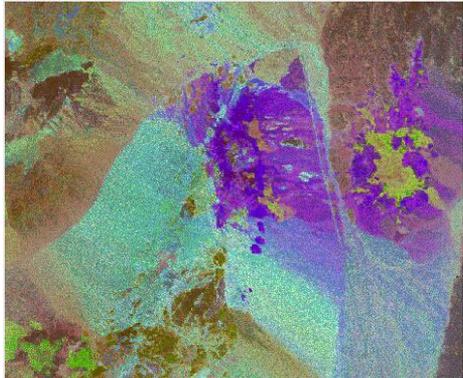
NOTE: the continuum-removed image is created with a separate raster object for each band rather than as a single hyperspectral object.



Spectral plot showing the kaolinite image spectrum (bottom) from the SAMPLIB3 spectral library, the continuum computed for it (smooth line), and the continuum-removed kaolinite spectrum (top).

STEPS

- select Remove Continuum from the Image menu on the Hyperspectral Analysis window
- use the Select File procedure to name a new Project File to contain the continuum-removed image bands



RGB display of continuum-removed bands 172, 187, and 202. To examine the image without interrupting the continuity of these exercises, launch another instance of the Hyperspectral Analysis process and open the continuum-removed image there.

Remove Continuum from Library Spectra

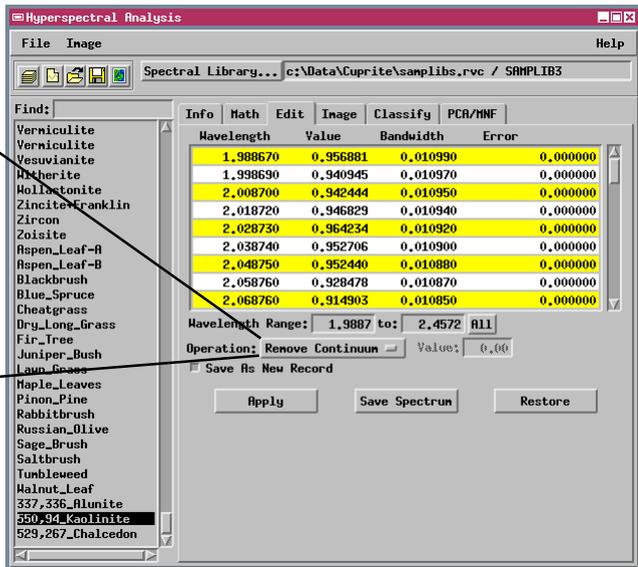
STEPS

- select the 550,94_Kaolinite spectrum from the SAMPLIB3 spectral library
- in the Plot portion of the Info panel, turn on the Single Spectrum radio button, then the Multiple Spectra button
- in the Spectral Plot window, turn on the lower Manual Range toggle button and set the Value Range from 0.5 to 1.0
- click on the Edit tab
- from the Operation menu, select Continuum, then press [Apply]
- press [Restore]
- select Remove Continuum from the Operation menu, then press [Apply]
- press [Save Spectrum] to save the continuum-removed spectrum in the library

You can remove the continuum from individual spectra in a library by using the controls on the Edit tabbed panel. The spectrum you select from the library is shown in tabular form in this panel. Select Continuum from the Operation menu to merely compute the continuum for display in a spectral plot, or Remove Continuum to compute a continuum-removed spectrum from the original library spectrum.

When you apply any of the mathematical transformations selected from the Operation menu, the spectrum table on the Edit panel is updated and the result is stored temporarily in memory. Use the Restore button to remove the transformation or the Save Spectrum button to save the result as a new spectrum in the library. This spectrum is assigned the same name as the parent spectrum. You can use the controls on the Info panel to edit this name and resave the spectrum.

- Set
- Add
- Subtract
- Multiply
- Interpolate
- Smooth
- Normalize
- Continuum
- Remove Continuum**
- Derivative
- Binimize
- Parabolic Fit

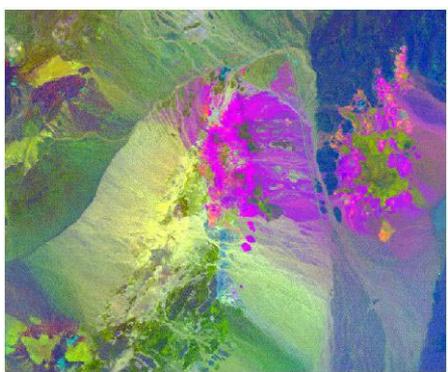


The steps in this exercise create a spectral plot similar to the one illustrated on the preceding page.

Other Image Processing Options

The Image menu on the Hyperspectral Analysis window offers several additional processing options which are summarized briefly below.

Log Residuals is a calibration technique that models the observed radiance value for each image cell as the product of the surface reflectance, a topographic factor, and an illumination factor. The topographic factor (slope direction and shadowing) is constant for all wavelength bands for any image cell, but varies spatially. The illumination factor (solar irradiance) is constant over the image (assuming a cloud-free sky), but is unique for each wavelength. In order to determine reflectance from the measured radiance value, we need to remove the effects of the topographic and illumination factors, neither of which is usually known. The Log Residuals algorithm mathematically eliminates these factors by performing a series of divisions for each image cell using the measured band values, the geometric mean over all wavelength bands for each cell, and the spatial geometric mean for each image band. The geometric means are computed from the arithmetic means of the logarithms of the measured values. These calculations are too involved to be done on-the-fly, so the Log Residuals process produces a transformed set of image bands. The resulting log residual image spectra may lack some of the original image information. If any material spectrum is present in all image spectra, that spectral component will be removed by the process. Log residuals images also tend to accentuate image noise.



Log residuals for bands 172, 187, and 202 displayed respectively as R, G, and B. (The original RGB combination is shown on page 8).

VQ Filtering is an image transformation procedure that is related conceptually to the Self-Organizing Map (SOM) classifier discussed earlier in this booklet. The procedure applies the SOM classifier to the image to derive 256 spectral classes, and assigns each image cell to the class with the closest mean in spectral space. Then the mean spectrum for the assigned class is subtracted from each observed spectrum. The resulting residual (filtered) image of each input band shows how well (or how poorly) each cell fits its assigned class at that wavelength. You might use the results to pinpoint the locations of unusual (but potentially interesting) materials in the scene.

Advanced Software for Geospatial Analysis

MicroImages, Inc. publishes a complete line of professional software for advanced geospatial data visualization, analysis, and publishing. Contact us or visit our web site for detailed product information.

- TNTmips Pro** TNTmips Pro is a professional system for fully integrated GIS, image analysis, CAD, TIN, desktop cartography, and geospatial database management.
- TNTmips Basic** TNTmips Basic is a low-cost version of TNTmips for small projects.
- TNTmips Free** TNTmips Free is a free version of TNTmips for students and professionals with small projects. You can download TNTmips Free from MicroImages' web site.
- TNTedit** TNTedit provides interactive tools to create, georeference, and edit vector, image, CAD, TIN, and relational database project materials in a wide variety of formats.
- TNTview** TNTview has the same powerful display features as TNTmips and is perfect for those who do not need the technical processing and preparation features of TNTmips.
- TNTatlas** TNTatlas lets you publish and distribute your spatial project materials on CD or DVD at low cost. TNTatlas CDs/DVDs can be used on any popular computing platform.

Index

Auto Define wavelength.....	10	matched filtering.....	16,18
Band mapping (spectral matching).....	36	minimum noise fraction transform.....	23
continuum, removal.....	37-38	n-Dimensional Visualizer.....	25-28
classification		pixel purity index.....	24
of image.....	15-21	principal components transform.....	22
in n-Dimensional Visualizer.....	27	self-organizing map classifier.....	21
using reflectance spectra.....	35	spectral angle mapper.....	17,18,36
continuum.....	37,38	spectral libraries.....	13,15,29,31
cross-correlation.....	18	USGS spectral library.....	31
equal area normalization.....	34	spectral matching.....	36
flat field correction to reflectance.....	33	Spectral Profile tool.....	11,29
Hyperspectral Explorer.....	7-9	spectral radiance.....	32
hyperspectral (hypercube) object.....	4	spectral reflectance.....	30,31
3D display of.....	5	correcting radiance to.....	32-34
RGB display of.....	4,6-9	spectrum (spectra)	
illumination effects.....	32	endmember.....	19,20,24-29
imaging spectrometer.....	4,11,18	image.....	3,11,13-20
import, spectrum.....	30	reflectance.....	32-36
irradiance, solar.....	32,33	saving in library.....	13
linear unmixing.....	19,20	saving as text.....	14
log residuals calibration method.....	39	viewing.....	11



MicroImages, Inc.

11th Floor - Sharp Tower
206 South 13th Street
Lincoln, Nebraska 68508-2010 USA

Voice: (402) 477-9554
FAX: (402) 477-9559

email: info@microimages.com
internet: www.microimages.com