



Hyperspectral Analytics in ENVI®

Target Detection and Spectral Mapping Methods

Authors: Jason D. Wolfe and Sarah R. Black

Contributor: Joe Boardman

NV5 Geospatial Solutions, Inc.

November 1, 2023 Edition

TABLE OF CONTENTS

INTRODUCTION	4
Knowing Your Data	4
Sensor Characteristics.....	4
Resolution.....	5
Preparing Data for Analysis	6
Removing Bad Bands.....	6
Geometric Correction.....	7
Masks.....	7
Atmospheric Correction.....	7
Spectral Dimensionality Reduction.....	9
Reference Spectra.....	10
Target Detection	13
Selecting Target Spectra.....	13
Selecting Background Spectra (Optional).....	14
Applying MNF Transforms.....	15
Selecting Target Detection Methods.....	16
Setting Thresholds.....	18
Smoothing Results (Optional).....	18
Exporting Results.....	19
Extracting and Mapping Spectral Endmembers	20
Spectral Hourglass Workflow.....	20
Extracting Endmembers.....	21
Spectral Mapping Methods.....	Error! Bookmark not defined.
Summary	30
References	31

INTRODUCTION

The use of imaging spectrometers has gained popularity in the last 20 years, with applications in geology (e.g., Boardman and Kruse 1994; Kosanke and Chen 2017), defense and intelligence (e.g., Probasco 2017; Puckrin et al. 2012), precision agriculture (e.g., Martin and George 2018; Yang 2015), and atmospheric science (e.g., Matsunaga 2015).

Airborne and spaceborne spectrometers contain dozens or hundreds of narrow spectral bands that provide near-continuous reflectance spectra of earth surface features, spanning the visible and near-infrared (VNIR) to shortwave-infrared (SWIR) wavelength range. Examples include AVIRIS (NASA JPL 2018), HYDICE (Mitchell 1995), and CRISM (John Hopkins University APL 2018).

While multispectral imagery can be used to discriminate between different surface materials, imaging spectroscopy—also called *hyperspectral imaging*—provides even more powerful capabilities. These include exploiting unique spectral signatures and absorption features of materials to estimate the sub-pixel abundance of materials or to detect spectral targets of interest.

ENVI® software has traditionally been the software application of choice for analyzing hyperspectral data, whether it involves scientific research or making tactical decisions. With dozens of tools available for analyzing multispectral and hyperspectral data, the choice of which tools to use and what process to follow can be daunting for users who are new to imaging spectroscopy. This paper helps answer questions such as, “How do I get started?” and “What tools do I need?”

While not every possible hyperspectral application can be addressed here, recommended workflows are provided for two common scenarios: (1) locating specific targets of interest, and (2) determining what spectrally unique materials exist in a scene.

This paper does not provide instructional steps for using hyperspectral tools in ENVI; however, it does provide some overall guidance on using them. Screenshots, tools, and workflows are based on ENVI 6.0. That version includes more modern and efficient tools and workflows for processing hyperspectral data. Finally, this paper provides caveats and tips for working with hyperspectral data in ENVI, such as:

- Being aware of limitations and expected results with data
- Knowing when to correct for atmospheric effects and when to apply other preprocessing steps
- Collecting image and reference spectra
- Interpreting results from different workflows and understanding that not all results are absolute, quantitative values
- Learning *why* certain tools are needed and what value they offer, rather than exploring the details of the algorithms

A basic understanding of imaging spectroscopy will be helpful. Consult the references at the end of this paper for background information, or refer to the ENVI Help.

Before analyzing data from an imaging spectrometer, it is important to know the limitations of the data. The next section provides some guidance on things to consider.

KNOWING YOUR DATA

Analysts can easily make the common mistake of relying too much on the application to produce a quick answer without considering the limitations of the data. Often, the results are not what they expect because of incorrect assumptions about the collecting platform and the data. This section provides some important questions to consider before doing any spectral analysis.

Sensor Characteristics

The type of sensor used to collect the scene determines how the digital values are stored and interpreted. With spaceborne and airborne sensors—including unmanned aerial vehicles (UAVs)—it is important to know the scanning direction. There are two main categories: pushbroom and whiskbroom sensors.

A *pushbroom sensor* (also called an *along-track sensor* or *line scanner*) scans incoming radiation along the same direction as the satellite path and collects a line of data at a time for all wavelengths. The radiation is sent to solid detectors to record as digital values (Chuvieco 2016). Examples of pushbroom sensors include HYDICE, CASI, and EO-1 Hyperion.

A *whiskbroom sensor* (also called an *across-track sensor*) uses a rotating mirror to scan incoming radiation in a perpendicular direction to the satellite path. It captures the entire field of view (FOV) at a time in a given wavelength and then repeats the process for a slightly different wavelength. Examples of whiskbroom sensors include AVIRIS and HyMap. Fowler (2014) compares pushbroom and whiskbroom sensors as they relate to imaging spectrometers.




Figure 1: Simplified concepts of pushbroom sensor (left) and whiskbroom sensor (right)

The signal-to-noise characteristics of these two collection methods can be vastly different. In addition, some sensors such as EO-1 Hyperion contain artifacts that can distort images (Yokoya, Miyamura, and Iwasaki 2010).

Analysts should also be familiar with the level of processing that was applied to image data. Data providers typically process the data to remove geometric and radiometric errors associated with the motion of the sensing platform. Also, are the image pixels raw DN values, or have they been calibrated to radiance by the data provider?

Finally, knowing the altitude of the sensor can determine the thickness of the atmosphere to correct for. See the [Atmospheric Correction](#) section on page 7 for a summary of correction tools available in ENVI.

Resolution

Consider the *spatial resolution* of the image, which is the resolution of the smallest object that can be detected. For example, if a pixel is 30 square meters in size (as with EO-1 Hyperion) or even 18 square meters (as with the Mars Reconnaissance Orbiter's CRISM instrument), much of the information in the pixel will be lost. One or two materials (minerals, for example) will dominate the signal in a given pixel, even though far more materials may be within that instantaneous field of view (IFOV). With increased pixel sizes, analysts will never get a true representation of the actual content, compared to what may be identified *in situ*, although they can get an overall sense of what it contains.

Knowing the *spectral resolution* of the data—the number of bands and their ranges/centers—is critical for ensuring that the data captures a specific absorption feature of interest (Figure 2).




Figure 2: A comparison of spectral profiles for various minerals. Arrows point to diagnostic absorption features.

Using a mineralogy example, differentiating between the absorption features of Fe-, Mg-, and Al-OH requires sampling of several spectral bands within the range of 2.1 μm to 2.3 μm . Each of these features produces a characteristic absorption band within that range because of differences in molecular vibration; however, the difference between band centers may be very small (e.g., 2.30 μm for Mg-OH and 2.29 μm for Fe-OH). A benefit of hyperspectral analysis in the VNIR region is that it is sensitive to these subtle differences. When targeting a specific mineral, the band centers of the detectors should align with the characteristic absorption features for that mineral; otherwise, it can be easy to miss the absorption features, which can result in misidentifications.

Also, it is important to ensure that the bands cover the same range of an absorption feature of interest. For example, some materials such as quartz lack diagnostic absorptions in the VNIR range, and instead have diagnostic absorptions in the middle-infrared (roughly 3 to 40 μm). These materials cannot be detected with VNIR bands and require different sensors for detection.

The next section describes common preprocessing steps to prepare data for spectral analysis.

PREPARING DATA FOR ANALYSIS

A fair amount of preprocessing is needed to reduce noise and erroneous data from hyperspectral imagery before meaningful information can be extracted from it. This involves removing bad bands, masking out unwanted features, correcting for atmospheric effects (for spaceborne and airborne sensors), and transforming hyperspectral data into a different image space to reduce its dimensionality.

Geometric correction is also a consideration, along with properly preparing reference spectra. These are also discussed.

Removing Bad Bands

Some bands in imaging spectrometers can produce significant noise, particularly in wavelengths associated with atmospheric absorption and water vapor. These bands need to be removed before subsequent processing.

A recommended approach to identifying bad bands is to display a spectral profile of any pixel. Look for spikes in the plot (Figure 3).

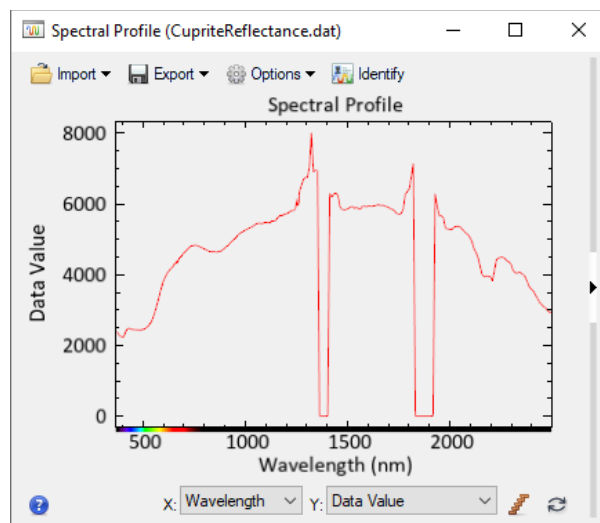


Figure 3: Spectral profile of an AVIRIS image pixel showing spikes and gaps.

Or, use the Band Animation tool (Figure 4) or the Xtreme Viewer to visually inspect each band, and note the bands that appear noisy. To do this, right-click on an image in the Layer Manager and select **Band Animation > Using Raster Series** or **Band Animation > Using Xtreme Viewer**.

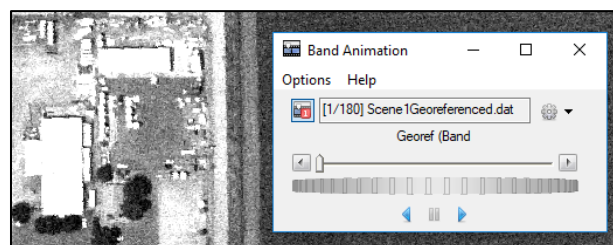


Figure 4: Using the Band Animation tool to animate a hyperspectral image with 180 bands.

Once the bad bands have been identified, there are two ways to remove them from further analysis.

Define Bad Bands in the Image Metadata

Bands can be marked as “bad” to indicate that they should be ignored in subsequent processing. The Bad Bands List in the Edit ENVI Header tool can be used to exclude specific bands. The Edit ENVI Header tool is available in the ENVI Toolbox under the Raster Management category. The bands are not actually removed from the original image; they are just marked as bad. The associated header file (.hdr) that accompanies the hyperspectral image is then updated with the list of bad bands.

Note: The FLAASH atmospheric correction tool determines bad bands based on the strength of the reflectance signal, and it automatically updates the ENVI header file that accompanies the output reflectance image.

With EO-1 Hyperion data, bands 1-7, 58-76, and 221-242 are automatically set to values of 0 by the data provider (Barry 2001). Also, bands 121-126 and 167-180 have severe noise that correspond to strong water vapor absorption and should be removed from processing (Datt et al. 2003).

Define a Spectral Subset

Another way to exclude bad bands from processing is to define a spectral subset of the good bands from the original image, then save that to a new image.

Geometric Correction

When creating maps of material distribution or spectral targets, imagery should be georeferenced to a standard map projection. Some data providers distribute map-ready image products; for example, EO-1 Hyperion L1G, L1Gst, and L1T products. Some UAV images have no spatial reference but are delivered with separate geographic lookup table (GLT) images or input geometry (IGM) images that can be applied to create georeferenced data. An example is imagery from the Corning microHSI imaging spectrometer.

The Geometric Correction folder of the ENVI Toolbox contains a wide range of correction tools, from building/applying GLTs and IGMs to rigorous orthorectification. Refer to the [Georectify Imagery topic in ENVI Help](#).

Masks

Some image pixels in a scene can compromise the spectral integrity of the data. These include image borders, stripes and other bad data, clouds, cloud shadows, and water. These should be masked out prior to analysis. Refer to the [Masks topic in ENVI Help](#) for guidance on applying masks to images.

Atmospheric Correction

The effects of atmospheric absorption and scattering in the VNIR wavelengths are pronounced in data acquired from spaceborne and airborne imaging spectrometers. The composition and thickness of the atmosphere varies spatially and over time, which distorts the perceived reflectance and absorption signals of materials detected by the sensors (Bedell and Coolbaugh 2009). These effects must be accounted for and corrected before doing any spectral analysis.

Raw pixel values (also called *digital numbers* or *DN values*) should be calibrated into physically meaningful units. The three most common radiometric corrections are radiance, top-of-atmosphere (TOA) reflectance, and apparent surface reflectance. The type of calibration to apply depends on the intended application.

Radiance or uncalibrated data can be used as input to spectral analysis tools that do not require any external reference spectra and where spectral endmembers can be derived from the image alone. Examples include Minimum Noise Fraction (MNF), Pixel Purity Index (PPI), Matched Filtering (MF) and Mixture-Tuned Matched Filtering (MTMF).

If library spectra will be used associate image endmembers to known materials, then the image data must be in units of reflectance and must be scaled to match the range of the library spectra. See the [Reference Spectra](#) section on page 10 for tips. In most cases, pixel values should range from 0 to 1, which represents 0 to 100% reflectance. The Spectral Analyst, Spectral Angle Mapper (SAM) classification, and Spectral Feature Fitting tools expect the data to be in units of reflectance, not radiance. Calibrating imagery to surface reflectance also ensures consistency when constructing a time series of data or fusing data from different sensors. Atmospheric correction tools can be used to calibrate data to apparent surface reflectance.

Minu and Shetty (2015) compare and contrast the most commonly used methods for atmospheric correction of hyperspectral imagery, including some that are available in ENVI.

Next is a summary of the atmospheric correction tools that ENVI provides.

Quick Atmospheric Correction (QUAC®)

QUAC is a scene-based empirical method that converts radiance values to apparent surface reflectance. *Scene-based* means that the atmospheric correction parameters are derived strictly from the pixel spectra within the scene and not from any ancillary data. A separate license for the ENVI Atmospheric Correction Module is required to use QUAC.

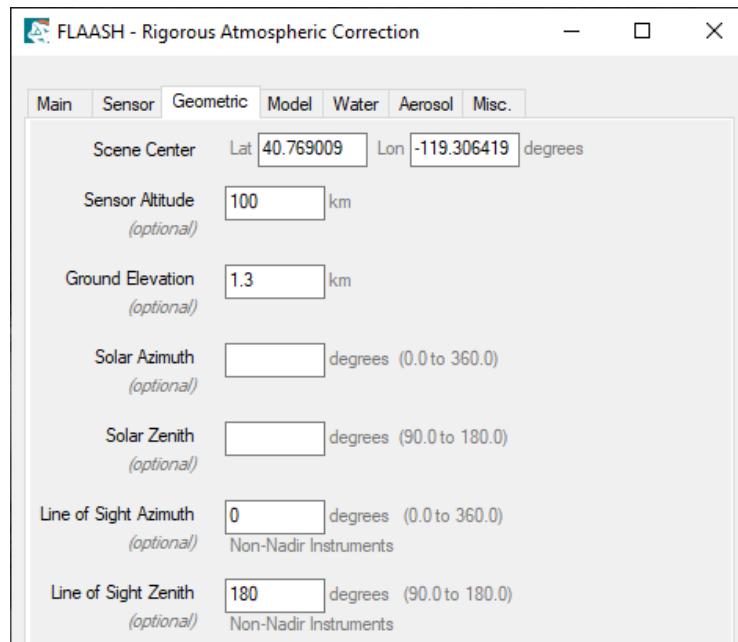
QUAC is the simplest atmospheric correction tool to use in ENVI. It accommodates a wide range of wavelengths (VNIR to SWIR, approximately 0.4 to 2.5 μm) and sensors. The input scene should contain a variety of spectrally diverse materials—at least 10—such as water, soil, vegetation, and man-made structures (Bernstein et al. 2012). It performs best when the imagery is uniformly illuminated, such as in clear-sky conditions or when airborne sensors fly under complete cloud cover.

Where to Find It

- ENVI Toolbox → Radiometric Correction → Atmospheric Correction Module → Quick Atmospheric Correction (QUAC)
- ENVI Toolbox → Workflows → Preprocessing Workflow

Fast Line-of-sight Atmosphere Analysis of Spectral Hypercubes (FLAASH®)

FLAASH is a physics-based method that incorporates MODTRAN® radiative transfer code to model atmospheric water vapor and aerosols. A separate license for the ENVI Atmospheric Correction Module is required to use FLAASH. It is the most rigorous atmospheric correction method available in ENVI and produces accurate surface reflectance data. The associated ENVI header file of the input image must have wavelengths defined. Beginning with ENVI 5.7, FLAASH reads the metadata from an image and pre-populates the FLAASH dialog with as many values as it can determine; for example: acquisition date and time, geographic coordinates, sensor, satellite altitude, and others. FLAASH is also available as a task to use with the ENVI Modeler and programming API for creating custom workflows.



When FLAASH converts radiance data to reflectance, it can introduce artifacts into the spectra. These artifacts can result from any of the following factors:

- Mismatches in the spectral calibration of the hyperspectral dataset and the spectral radiative transfer calculations
- Errors in the absolute radiometric calibration
- Errors in the radiative transfer calculations

Where to Find It

- ENVI Toolbox → Radiometric Correction → Atmospheric Correction Module → FLAASH Atmospheric Correction
- ENVI Toolbox → Workflows → Preprocessing Workflow

Empirical Methods

The following methods can be used as alternatives to removing the effects of the atmosphere. However, they only provide crude approximations and do not calibrate image data to apparent surface reflectance.

Empirical Line

The Empirical Line correction method requires an analyst to collect field or laboratory reflectance spectra, and also to identify dark and bright regions in the image. Calibration targets can be helpful in this case. The image spectra are forced to match the field spectra. This is equivalent to removing solar irradiance and atmospheric path radiance.

Where to Find It

- ENVI Toolbox → Radiometric Correction → Empirical Line Compute Factors and Correct
- ENVI Toolbox → Radiometric Correction → Empirical Line Correct Using Existing Factors
- Included in the Target Detection Workflow

Internal Average Relative (IAR) Reflectance

The IAR Reflectance empirical correction method computes the mean spectrum of the entire scene, then divides it into the radiance value for each band of every pixel. The result is an image of relative reflectance. Its accuracy can vary considerably because it is tied to the most abundant material in a scene (Bernstein et al. 2012). IARR produces relatively good reflectance measurements in cases where the most abundant material is spectrally flat, such as arid scenes with little vegetation (Kruse, Raines, and Watson 1985).

Where to Find It

- ENVI Toolbox → Radiometric Correction → IAR Reflectance Correction
- ENVI Toolbox → Workflows → Preprocessing Workflow

Flat Field

The Flat Field empirical correction method requires an analyst to define a ROI around a spectrally flat area in a scene, preferably one with a high albedo. The mean spectrum of the flat field ROI is divided by the spectrum of each pixel in the scene. This normalizes the entire scene, yielding relative reflectance values (Roberts, Yamaguchi, and Lyon 1986). A drawback to this method is that most spectrally flat areas still contain a lot of spectral variation even though they visually appear to be the same. Experimentation with different display stretches can help to locate the best representative sample of a flat field. Also, because it does not include any baseline subtraction (as with Dark Subtraction), it can magnify errors resulting from dark pixels (Bernstein et al. 2012).

Where to Find It

- ENVI Toolbox → Radiometric Correction → Flat Field Correction
- ENVI Toolbox → Workflows → Preprocessing Workflow

Dark Subtraction

The Dark Subtraction method attempts to remove the effects of atmospheric scattering from a scene (particularly in the blue wavelength region) by subtracting a specific value from every pixel. This value represents a background signature. It can be the band minimum, a mean value based on a region of interest (ROI), or a user-specified value.

Where to Find It

- ENVI Toolbox → Radiometric Correction → Dark Subtraction
- ENVI Toolbox → Workflows → Preprocessing Workflow
- Included in the Target Detection Workflow

Log Residuals

The Log Residuals correction method uses in-scene statistics to produce a pseudo-reflectance image that is useful for analyzing mineral-related absorption features (Green and Craig 1985). It is rarely used.

Where to Find It

- ENVI Toolbox → Radiometric Correction → Log Residuals Correction
- ENVI Toolbox → Workflows → Preprocessing Workflow

Spectral Dimensionality Reduction

A common issue in remote sensing is that adjacent spectral bands often contain redundant information. This is because the bands occupy similar spectral regions, or because some materials have similar radiance values across spectral regions (Chuvienco 2016). Statistically speaking, adjacent bands are often highly correlated.

This issue is more pronounced with hyperspectral data because of its oversampled nature. With hundreds of bands of data, the *data dimensionality* increases significantly. Each contiguous, narrow band of data contains information that is not unique. Only a small portion of each band contributes to the overall signal, with the remainder attributed to noise. Bands that correspond to spectral regions of atmospheric water absorption contain severe noise and are useless as they have no correlation with adjacent bands.

Data transforms are typically used to reduce the dimensionality of datasets and to separate noise from signals. Transforming data into a different space can often reveal spectral features that otherwise would not be found. Methods such as Principal Components Analysis (PCA) and Minimum Noise Fraction (MNF) can be used in ENVI to determine the intrinsic, or *inherent*, dimensionality of a dataset. This refers to the smallest number of dimensions or variables necessary to model the data without incurring loss (Kirby 2001). Once the inherent dimensionality has been determined, the original, high-dimensional dataset is then replaced with the lower-dimensional dataset.

When using hyperspectral data to map the distribution of materials, analysts should ensure that the reduced bands will adequately characterize the materials of interest. For example, if a transformed image with only 10 bands is used to identify 20 different minerals, the mineral identification map can contain many false positives and false negatives.

Machine learning applications typically do not need spectral dimensionality reduction because they automatically discard irrelevant data.

PCA is a commonly used method, but the resulting PC bands contain noise and data. This is less of an issue when performing classification or feature extraction with hyperspectral data, but it should not be used as a preprocessing step for material identification. Ren et al. (2014) describe how PCA does not effectively produce the intended results with large hyperspectral datasets.

Minimum Noise Fraction (MNF)

MNF is the preferred method for reducing spectral dimensionality in hyperspectral data because it has the added benefit of separating noise from data. This is critical for applications like material identification where the goal is to identify the purest spectral endmembers in a scene. MNF reduces a dataset to its inherent dimensionality by retaining only the coherent bands and discarding the remaining noisy bands.

Where to Find It

- ENVI Toolbox → Transform → MNF Rotation → Forward MNF Transform
- Included in the ENVI Spectral Hourglass Workflow
- Included in the Target Detection Workflow

The ENVI implementation of MNF is based on the method described in Green et al. (1988) with some modifications. First, it estimates the noise in an image. In most cases, noise statistics are calculated directly from the data itself based on a *shift difference* method that uses local pixel variance. Analysts can optionally select a spatial subset of an image, or specify a group of pixels in a spectrally flat area, to improve noise estimates. The number of selected pixels must be greater than the number of bands in the image. The result of this step is a noise covariance matrix.

An initial PC transform uses the estimated noise covariance matrix to decorrelate and rescale the noise such the transformed data has unit variance and no band-to-band correlation. A second PC transform is then applied to that transformed data. The result is a set of MNF-transformed bands.

The key process with MNF is determining the threshold at which the bands transition from signal to noise. This is referred to as *spatial coherence*. The lower bands are expected to have spatial structure and will contain most of the information. These are called *coherent images*. Higher MNF bands are expected to have little spatial structure and will contain most of the noise. The bands are ranked with the largest amount of variance in the first few bands and decreasing amounts of variance in the remaining bands.

An Explained Variance plot shows the eigenvalue for each MNF-transformed band (eigenvalue number). Figure 5 shows an example after performing an MNF transform on hundreds of bands of NEON hyperspectral imagery.

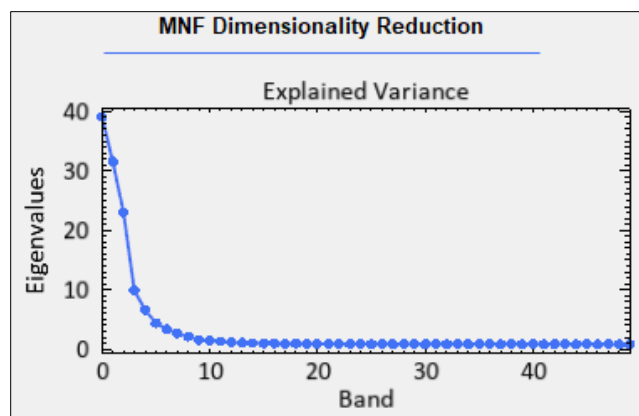


Figure 5: Example MNF eigenvalue plot.

Larger eigenvalues (along the Y-axis) indicate higher data variance in the transformed band. When the eigenvalues approach 1, only noise is left in the transformed band. Using Figure 5 as an example, this occurs near MNF Band 10.

Knowing how many MNF bands to keep is a subjective process. It is complicated by nature and is dependent on the particular scene. It is better to overestimate the dimensionality by using all of the MNF bands that have reasonable image quality and/or eigenvalues that are above unity.

Reference Spectra

This section describes the use of spectral libraries or field/laboratory spectra as ground truth, or *reference*, data. Before continuing, a distinction should be made between the two primary scenarios when working with hyperspectral data.

In the first case, analysts have a scene whose composition and spectral endmembers are unknown. They want to determine the endmembers that are present in the scene. Tools such as the ENVI Spectral Hourglass Workflow and Linear Spectral Unmixing are

used to derive endmember spectra directly from the image to determine the composition of the scene, without relying on ancillary or reference spectra. These tools are discussed in more detail later in this paper.

In the second case, analysts have a scene whose composition and spectral endmembers are unknown, but they are not concerned with the overall composition. They only want to know the locations of a specific material of interest (a target). Reference spectra from a spectral library are used to locate targets and to separate them from non-target, or background, pixels. Reference spectra can even include a group of pixels from the scene that have a high certainty of representing a known material. See the [Target Detection](#) section on page 12.

This discussion pertains to cases when reference spectra are used. The most common type of reference spectra are spectral libraries.

Spectral Libraries

When researchers use spectrometers in the field or laboratory to collect complete spectral signatures of materials, they can build libraries of multiple spectra (Figure 6). Many spectral libraries are available for public access, such as these well-known libraries:

- U.S. Geological Survey (<https://speclab.cr.usgs.gov/spectral-lib.html>)
- ECOSTRESS, which includes data from John Hopkins University, NASA Jet Propulsion Laboratory, and the U.S. Geological Survey (<https://speclib.jpl.nasa.gov/>)
- SPECCHIO Spectral Information System (<http://www.specchio.ch/>)
- CRISM Spectral Library (http://pds-geosciences.wustl.edu/missions/mro/spectral_library.htm)

These libraries do not contain spectra for every possible material, so analysts may need to track down or create additional reference spectra for the exact materials they are interested in. In summary, analysts should have a general sense of what materials may be in a scene and ensure that all of the relevant library spectra are available for reference.

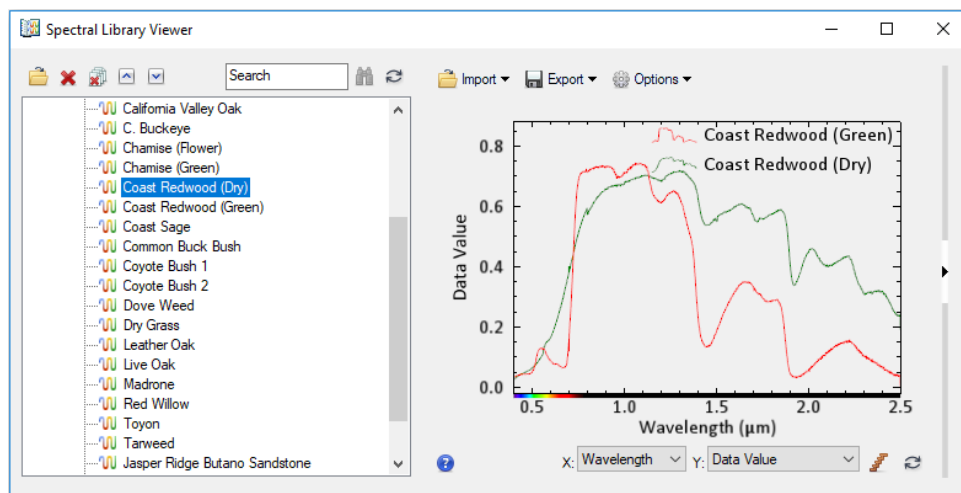


Figure 6: Spectral Library Viewer showing two vegetation spectra.

Where to Find Spectral Library Tools

- ENVI main menu bar → Display → Spectral Library Viewer
- ENVI Toolbox → Spectral → Spectral Library Builder. Use this tool to import spectra from external libraries.
- ENVI Toolbox → Spectral → Spectral Library Resampling

Other Reference Spectra

In addition to spectral libraries, reference spectra can also come from plots, text files, statistics files, binary ASD spectrometer files, and regions of interest (ROIs). An ROI is a group of pixels (in the form of points, polylines, or polygons) that an analyst identifies with a high degree of certainty as containing a material of interest. With polygon ROIs, the mean value is used as the reference spectrum.

Ensuring Consistency with Image Spectra

If spectral libraries or other reference spectra are used in hyperspectral image analysis, they should be consistent with the image spectra. Consider the following:

- The units of the spectral library measurements: reflectance, downwelling irradiance, or others
- Any scaling that has been applied to the library spectra

- The wavelength range covered by the library spectra

For example, a spectral library that contains irradiance values cannot be used directly with an image whose pixels represent reflectance values. Likewise, a material of interest may not be identified in an image if the bands do not cover the same wavelength range as the reference spectra. Figure 7 shows a plot that compares spectrometer data collected in a laboratory for a blue vinyl tarp, and a hyperspectral image pixel of a blue vinyl tarp. The units are the same for both spectra (micrometers), and they cover the same wavelength range (0.48 to 2.5 μm).




Figure 7: Spectral profiles of reference spectra (red) and image spectra (grey).

Reference spectra may need to be resampled to match image spectra. The Spectral Library Resampling tool can resample spectral libraries to match the response of a known instrument such as AVIRIS, an ASCII wavelength file, or the wavelengths of an image file.

The ENVI Band Math tool can be used to scale image pixel values to the range of spectral library data values. For example, some atmospheric correction methods scale pixel values by 10,000 so that they range from 0 to 10,000. However, most spectral libraries have data that range from 0 to 1, as in Figure 7 above. In this case, the image pixel values should be scaled by a value of 0.0001 to match the data range of the spectral libraries.

Importing Reference Spectra Into a Workflow

The Spectral Collection Classification dialog can be used to collect reference spectra from multiple sources. Figure 8 shows an example. In the ENVI Toolbox, expand the **Classification** folder and double-click **Endmember Collection** to display the Spectral Collection Classification dialog.

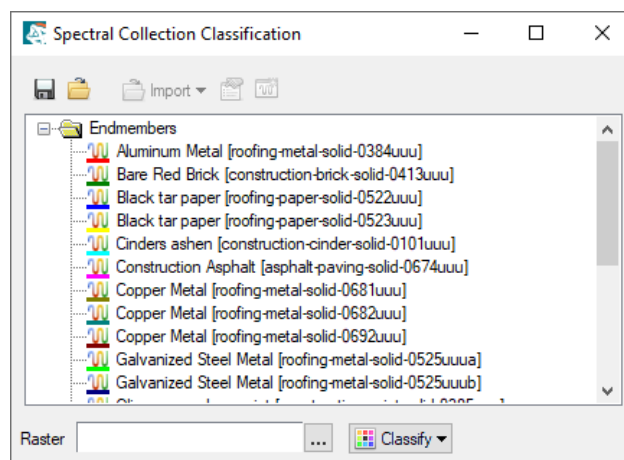


Figure 8: Spectral Collection Classification dialog listing spectra of man-made materials from a spectral library.

TARGET DETECTION

Another common use of hyperspectral imagery is locating a material of interest (a *target*), or discriminating between multiple targets based on their spectral characteristics (Manolakis, Lockwood, and Cooley 2016). Examples include man-made structures, vehicles, and minerals.

Target detection methods look at the spectrum of each pixel to determine whether a target is present or not, based on a known spectral signature. Typically, the target only represents a small fraction of the pixels in the entire scene. The remaining pixels represent the background. Thus, the goal of target detection is to identify known spectral targets in an unknown background. The ability to resolve a target depends on the spatial resolution of the sensor. Some target detection methods incorporate an element of *spectral unmixing* (discussed later in this paper) to make a decision on whether or not a pixel represents a given target when the pixel contains mixed materials. The Mixture-Tuned Matched Filtering (MTMF) method is an example.

Figure 9 shows the recommended workflow for target detection in ENVI when using hyperspectral data from an airborne or spaceborne sensor. Dashed lines indicate optional steps.

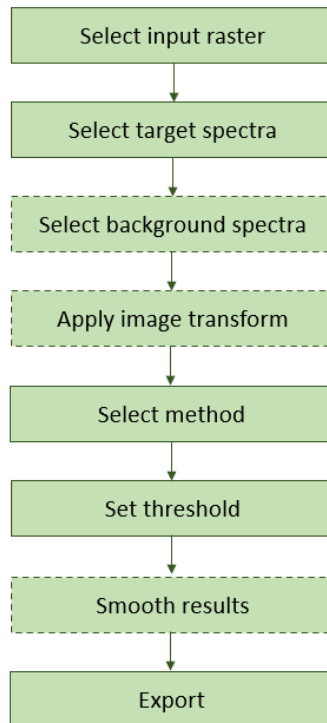


Figure 9: Target Detection workflow

The Target Detection Workflow follows this process; it is available in the ENVI Toolbox under the **Workflows** folder.

Selecting Target Spectra

The next step after selecting an input raster is specifying the spectra of the targets of interest. The target spectra typically come from spectral libraries. See [Reference Spectra](#) on page 10 for tips on collecting spectra. Figure 10 shows a spectral library of roofing materials that will be used as reference spectra.

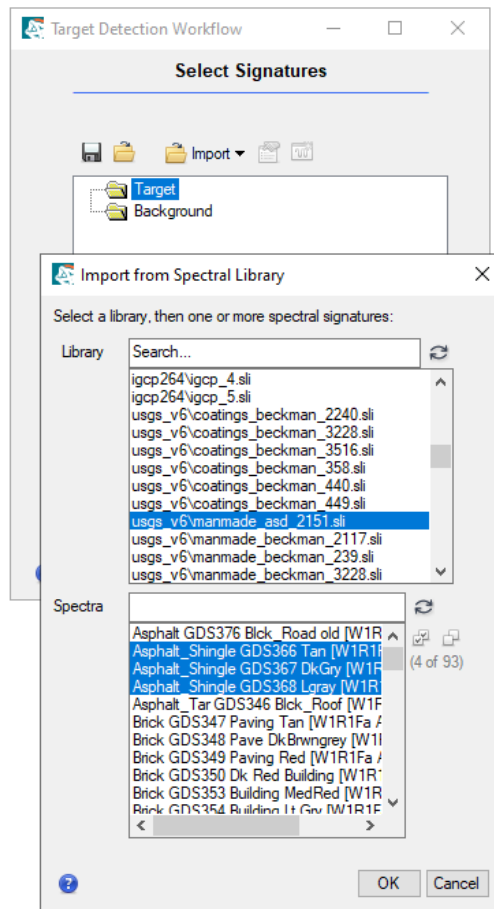


Figure 10: Example of selecting target spectra from a spectral library.

Some target detection methods require more than one target spectrum. See [Selecting Target Detection Methods](#) on page 16 for details.

Selecting Background Spectra (Optional)

Selecting non-target (background) spectra can improve the target detection result when using the Orthogonal Subspace Projection (OSP), Target-Constrained Interference-Minimized Filter (TCIMF), and Mixture Tuned Target-Constrained Interference-Minimized Filter (MTCIMF) methods. These methods are discussed later. Background spectra are unnecessary for other methods. The background spectra can come from spectral libraries, individual spectral plots, text files, ROIs, or statistics files. See [Reference Spectra](#) on page 10 for tips on collecting spectra. Figure 11 shows how ROIs were drawn around vegetation, shadows, and soils with mixed materials, to exclude them from consideration when looking for matches of specific roofing materials. The ROI means are used as the non-target spectra (Figure 12).

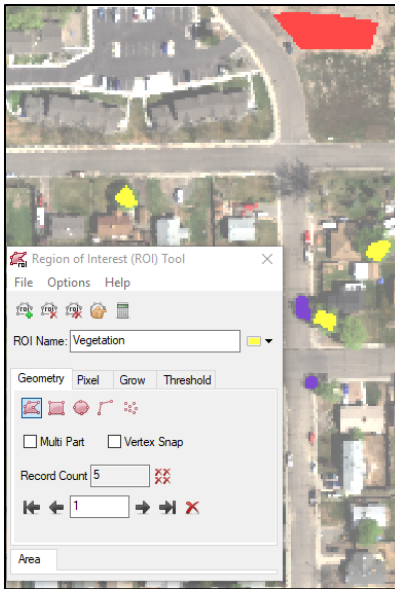


Figure 11: Using ROIs to identify non-target pixels.

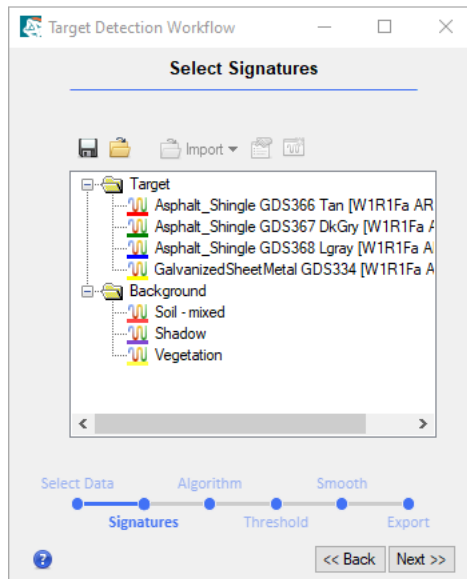


Figure 12: Using ROI means for background spectra.

Applying Image Transforms

Using MNF to reduce spectral dimensionality is generally unnecessary for target detection; however, it depends on the application. Redundant information from adjacent wavelengths can actually help locate better matches to targets. As long as computer memory and speed are not an issue, using the original (non-transformed) data can produce the best results with target detection. Jin, Paswaters, and Cline (2009) recommend preserving the full dimensionality of the data. Since the target is usually a small fraction of all the pixels, it does not contribute much to the covariance matrix and may be hidden in some of the higher MNF bands.

Note, however, that the MTTCIMF and MTMF methods require the data to be in MNF space, as these methods are also used for determining the relative sub-pixel abundance of materials. See [Minimum Noise Fraction \(MNF\)](#) on page 10 for more information about creating MNF-transformed data.

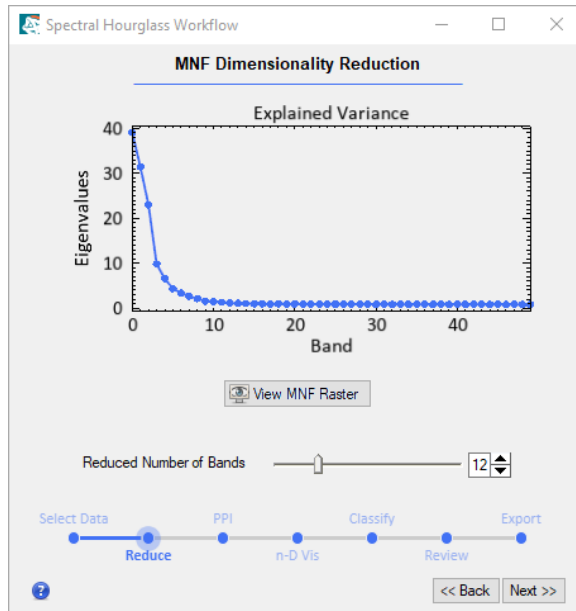


Figure 13: Image Transform step of the Target Detection Workflow.

Selecting Target Detection Methods

ENVI provides a variety of methods for target detection. Jin, Paswaters, and Cline (2009) summarize and compare these methods in more detail. Experimenting with different methods may be needed to achieve the intended results for a particular application (Figure 14).

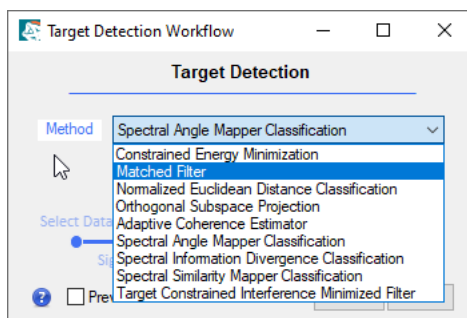


Figure 14: Selecting methods for target detection.

Adaptive Coherence Estimator (ACE)

The ACE method determines if a pixel spectrum possibly consists of a known target signature. It is best used when the background conditions are variable and unknown. It does not vary based on the relative scaling of input spectra, and it does not require knowledge of all the endmembers within a scene (Kraut, Scharf, and Butler 2005; Manolakis, Marden, and Shaw 2003). It performs well with detecting sub-pixel targets.

Constrained Energy Minimization (CEM)

The CEM method does not require knowledge of all the endmembers within a scene. A correlation or covariance matrix is required in order to model the composite unknown background over the whole scene (Harsanyi 1993; Chang et al. 2000). The CEM method works similarly to the OSP method, but it can do a better job of removing unidentified signals and suppressing noise than the OSP method. However, it does not generalize well because it is sensitive to the knowledge of the desired signature (Du, Ren, and Chang 2003).

Matched Filtering (MF)

The MF method is based on the same theory as the CEM method and was originally designed for target detection applications in signal processing. It detects specific materials based on matches to target spectra and does not require knowledge of all the endmembers within an image. However, it does not properly account for spectral unmixing in hyperspectral data.

Pixels that contain rare materials (unrelated to the target of interest) can give false positive responses, which cannot be distinguished from the actual target responses. Because these rare materials occur in only a few pixels, they do not contribute to the background covariance and are not properly ignored by the MF process. The rare materials may exhibit spectra that are different from the target spectra, yet they can have the same score as the target spectra, thus indicating a perfect match to the target. Also, the MF method does not discriminate well among rare targets with similar spectral signatures (Boardman and Kruse 2011). Consider using the MTMF method instead.

Mixture Tuned Matched Filtering (MTMF)

The MTMF method (Boardman 1998) improves upon the MF method by providing better selectivity of targets. It is useful for detecting and discriminating among multiple rare targets whose spectral signatures are similar to the background. The MTMF method recognizes that targets actually *replace* some of the background signature in a pixel, not add to them. It uses the same statistical method as MF but incorporates elements of a linear mixing model. It can provide accurate mapping of very small sub-pixel targets with a low number of false positives (Boardman and Kruse 2011).

The MTMF method requires an MNF-transformed image for input. It calculates an MF Score along with an *infeasibility* measure, which describes the likelihood of each pixel being a mixture of the known target and background materials. The infeasibility measure allows analysts to identify and reject false positives.

Mixture Tuned Target-Constrained Interference-Minimized Filter (MTCIMF)

The MTCIMF method (Jin, Paswaters, and Cline 2009) combines the MTMF and TCIMF methods. It requires MNF-transformed data for input. If background spectra are provided, the MTCIMF method can potentially reduce the number of false positives as compared to the MTMF method.

Normalized Euclidean Distance (NED)

The NED method calculates the distance between two vectors in the same manner as a Euclidean Distance method, but it normalizes the vectors first by dividing each vector by its mean.

Orthogonal Subspace Projection (OSP)

The OSP method (Harsanyi and Chang 1994; Chang 1998) is related to the MF and CEM methods. It eliminates the response of non-targets, then applies MF to match the desired target from the data. Two or more target spectra must be provided, along with background spectra. These can be derived from the image. The OSP method is efficient and effective when target signatures are distinct. When the target signature and background signature are similar, the attenuation of the target signal is dramatic and the performance of OSP can be poor. Refer to Du, Ren, and Chang (2003) for a comparison of the OSP method with the CEM method.

Spectral Angle Mapper (SAM)

The SAM method (Kruse et al. 1993) determines the similarity between image spectra and reference target spectra by measuring their spectral angle in n -dimensional space (where n is the number of bands), without characterizing the background. Smaller angles represent closer matches to the reference spectrum.

Spectral Information Divergence (SID)

The SID method (Du et al. 2004) uses a divergence measure to match pixels to reference spectra. The smaller the divergence, the more likely the pixels are similar. Pixels with a measurement greater than the specified maximum divergence threshold are not classified.

Spectral Similarity Mapper (SSM)

The SSM method combines elements of SAM and Minimum Distance classification into a single measure.

Target-Constrained Interference-Minimized Filter (TCIMF)

The TCIMF method (Ren and Chang 2000) detects the desired targets, removes the background influence, and minimizes interference in one operation. A correlation or covariance matrix is required in order to model the composite unknown background over the whole scene. Previous studies (Johnson 2003; Chang 2003) show that if the spectral angle between the target and background is significant, the TCIMF method can potentially reduce the number of false positives over CEM results. Two or more target or background spectra must be provided.

In general, it is best to compare results from multiple methods to determine the best match between pixels and the targets of interest. Relying on one specific method does not always provide a definitive identification of targets.

Identifying targets based on their visual appearance is not recommended. Figure 15 shows an example where several buildings appear white in a true-color image, leading to the assumption that they contain the same material; however, a SAM image reveals that only one of them contains galvanized metal roofing.




Figure 15: NEON hyperspectral true-color image (left) and SAM image with threshold applied (right).

Setting Thresholds

The next step in the Target Detection Workflow is to interactively choose a threshold for which pixels are categorized as “target” or not. The workflow chooses a default value as a starting point. A slider can be used to modify its value. If the rule value is larger than the threshold, the pixel is highlighted in the display (Figure 16). Pixels with a response larger than the threshold are accepted as target pixels.

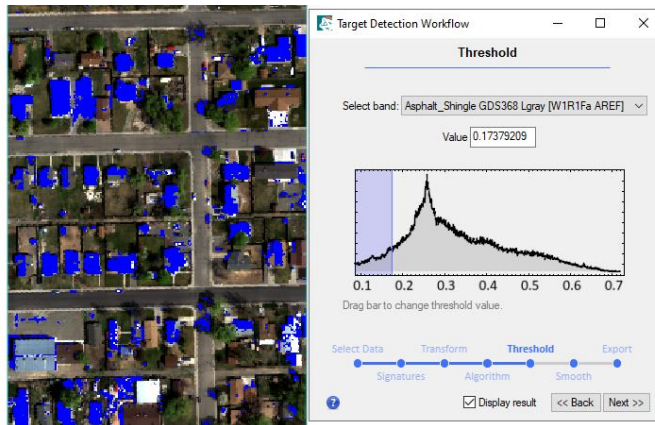


Figure 16: Example of setting a threshold value for a SAM image. Pixels with a closer match to the target spectra are highlighted in blue.

Smoothing Results (Optional)

This step cleans up the classification results by using a moving window (“kernel”) to aggregate adjacent pixels and to remove spurious pixels. See Figure 17.

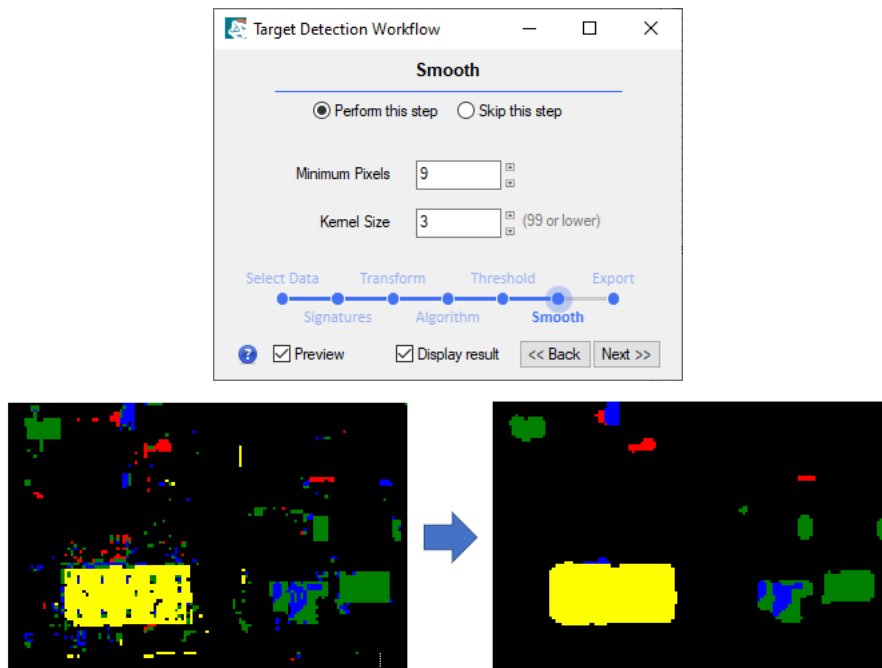


Figure 17: Example of smoothing classes for multiple targets.

Exporting Results

The final step in a target detection workflow is to export pixels that were identified as targets to:

- Classification image
- Shapefile
- ROIs
- GeoJSON
- KML
- Rule images



Figure 18: Shapefile of SAM target detection results for galvanized roofing, displayed over the source image.

While target detection methods can identify materials of interest by comparing the image spectra to library spectra, sometimes researchers do not have specific targets in mind. They do not know the spectral composition of an image and are more interested in finding the purest materials in an image and estimating their relative abundance. The next section describes how this can be accomplished in ENVI.

EXTRACTING AND MAPPING SPECTRAL ENDMEMBERS

In some cases, researchers want to know what spectrally unique materials exist within an image, without focusing on specific targets. An example is a geologist who wants to determine possible minerals that are present in a mining region. The process for doing this involves identifying unique endmembers in an image of the study area, then associating or *mapping* those endmembers to known materials using spectral libraries or other reference spectra. At the same time, they can estimate the relative abundance of each material within the image.

Each image contains a small number of “pure” materials whose spectral properties are constant. These are called *endmembers*. Imaging spectroscopy is frequently used to identify and classify as many of the endmembers as possible within an image. The spectral hourglass workflow described next can be used to derive endmember spectra directly from an image to determine its composition.

Why use image spectra in this case? Library spectra *can* be used to extract endmembers from an image, but consider that library spectra are typically collected under perfect illumination conditions using a laboratory spectrometer. Remote sensing image pixels represent less-than-ideal conditions, but their illumination conditions are relatively similar to each other.

Spectral Hourglass Workflow

ENVI provides a common “hourglass” processing workflow (Figure 19) that guides analysts through a process to:

- Extract endmembers directly from imagery instead of using library spectra.
- Map the distribution and fractional abundances of the materials they are associated with.

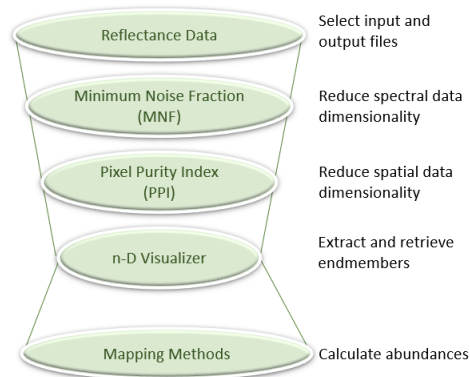


Figure 19: Spectral hourglass workflow.

The Spectral Hourglass Workflow has been used in many scientific applications. Qu et al. (2014) followed the workflow to estimate fractional vegetation abundance with EO-1 Hyperion imagery. Robichaud et al. (2007) used it to map post-fire burn severity with airborne imagery. Pan, Huang, and Wang (2013) used it for spectral feature fitting analysis of rapeseed canopies.

The workflow can be run with individual tools available in the ENVI Toolbox. However, a suggested approach is to use the ENVI Spectral Hourglass Workflow. This is available in the ENVI Toolbox in the Workflows folder. The workflow typically begins with an apparent reflectance image (Figure 20). See [Preparing Data for Analysis](#) on page 6 for guidelines on preparing the input hyperspectral image for analysis and creating an apparent reflectance image.



Figure 20: AVIRIS reflectance image of the Cuprite mining district, Nevada, USA.

An MNF transform determines the inherent dimensionality of the data by separating noise from the signal (Figure 21). See [Minimum Noise Fraction \(MNF\)](#) on page 10 for details.




Figure 21: Color composite of the first three MNF bands of the Cuprite scene.

The Pixel Purity Index (PPI) step identifies pixels that are spectrally pure in the scene. The interactive N-D Visualizer reduces those pixels to a set of endmembers that can be used for spectral mapping. The Material ID tool attempts to identify the materials corresponding to the endmembers. Finally, an appropriate mapping method is selected, based on the objective—whether it is mapping the spatial distribution of materials or mapping their relative sub-pixel abundance.

The next step is to extract spectral endmembers for analysis.

Extracting Endmembers

Every spectrum in an image can be reconstructed as some combination of the image endmember spectra. In theory, the number of endmembers that can be extracted from any image is equal to the number of bands, plus one. The one additional endmember represents the composite background, or all of the other pixels that contain mixtures of the endmembers (Boardman, Kruse, and Green 1995).

The process for extracting endmembers involves two steps: using PPI to identify the purest pixels in the scene, and using the N-D Visualizer to reduce the PPI pixels to a set of endmembers that can be used for spectral mapping methods. The spectral space in which this occurs is multi-dimensional. The next section describes this data space.

Visualizing Data in More Than Two Dimensions

A two-dimensional (2D) scatterplot helps to visualize the relationship between pixel data in two image bands. Figure 22 shows an example of a 2D scatterplot of two MNF bands. A blue color gradient is automatically applied. Brighter areas indicate the most dense concentration of pixels with similar values in both bands.




Figure 22: 2D scatterplot of two MNF bands.

A 3D scatterplot shows the relationship among three bands (Figure 23). This is where the concept of a *data cloud* becomes evident.

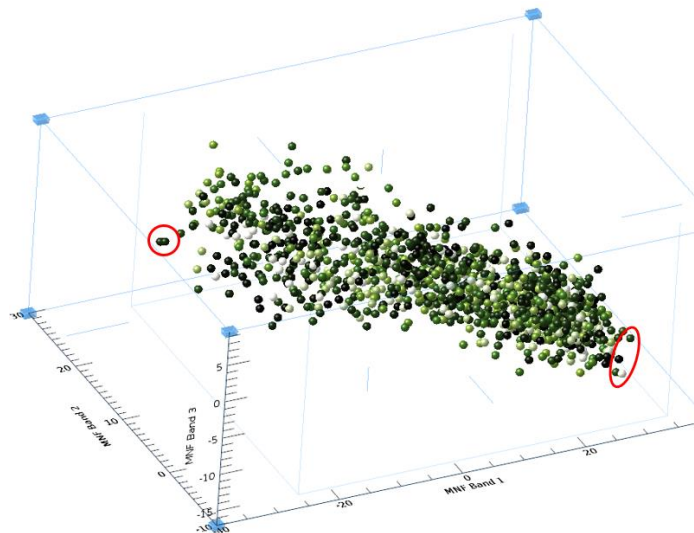


Figure 23: 3D scatterplot of three bands of data. Red circles indicate examples of spectrally "pure" pixels.

With dozens or hundreds of bands of hyperspectral data, however, the complexity of the spectral data distribution requires a different method of visualization. With imaging spectroscopy, the best way to do this is in an N -dimensional Euclidean spectral space, where N is the number of bands in the MNF-transformed image.

For example, if the first nine bands of an MNF-transformed image will be used to extract endmembers, then the spectral space will have nine different dimensions. Each of the nine bands is associated with one axis in spectral space, and the axes are orthogonal. The value of a spectrum in a single band determines its coordinates along the associated axis in spectral space.

Think of the N -dimensional data cloud as an irregularly shaped volume with properties such as shape, position, and density. The most spectrally pure pixels always occur in the corners of the data cloud (see Figures 23 and 26). These pixels form the vertices of the convex hull that surrounds all the other data. Spectrally mixed pixels occur inside of the data cloud. The concept of a convex hull is based on convex geometry, which is beyond the scope of this paper. See Boardman (1993) for details.

Pixel Purity Index (PPI)

The PPI step in the Spectral Hourglass Workflow attempts to separate the most spectrally pure pixels from those that contain a higher mixture of materials in hyperspectral images (Boardman, Kruse, and Green 1995). This reduces the number of pixels to analyze, and it makes separation and identification of endmembers easier. The PPI calculation identifies only the pixels that are the least mixed.

The most common reason for wanting to find the purest pixels is that their spectra are the best candidates for endmember spectra, which are needed for spectral mapping techniques that estimate the abundance of materials in the image.

The result of the PPI process is an image where the value of each pixel corresponds to the number of times it was identified as a pure pixel during a number of iterations. The PPI image can be used to map sites that should be visited for ground truth collection and further spectral measurements.

An MNF-transformed dataset is typically used as input to the PPI process. See [Minimum Noise Fraction \(MNF\)](#) on page 10 for details.

The PPI calculation selects a random vector through the n -dimensional data cloud, passing through its mean value each time. ENVI completes multiple iterations of projecting the image pixels onto the random vector and marks pixels as "pure" if they fall within the extremes of the resulting histogram.

As the PPI is being calculated, ENVI displays and updates a plot of the number of pixels that are determined as pure with each iteration. When the cumulative number of pure pixels begins to level off in the plot, it means that each subsequent iteration is no longer finding new pixels. In other words, those same pure pixels are being identified over and over again. Figure 24 shows an example of a PPI image. Brighter pixels are nearer to the corners of the N -dimensional data cloud; these pixels are relatively more pure than pixels with lower values.



Figure 24: Pixel purity index image stretched to show the highest 10% of values.

A typical PPI image has a large number of pixels with values that are 0 or near 0. Only a small fraction of pixels are considered spectrally pure.

Once the purest pixels have been identified, the next step is to separate them into their respective endmembers.

N-Dimensional Visualizer

The N-Dimensional Visualizer is an interactive tool for finding endmembers by locating and clustering the purest pixels (from the PPI step) in N -dimensional space, where N is the number of MNF bands. It helps visualize the shape of the data cloud that results from plotting image data in spectral space, using the MNF bands as plot axes (Figure 25).

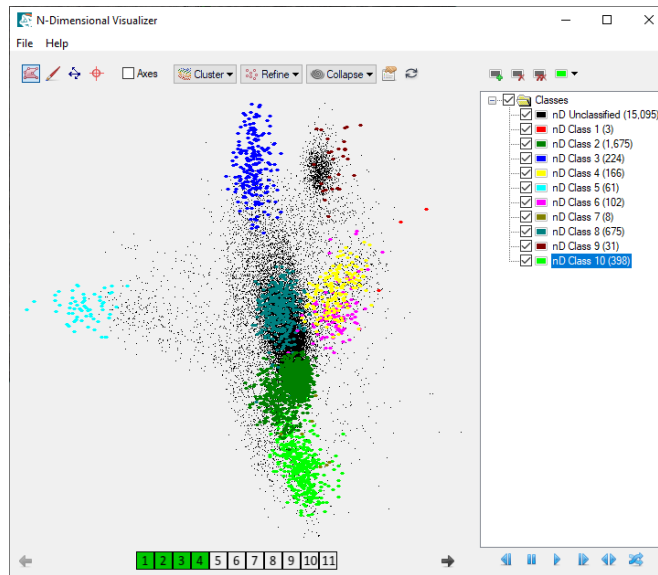


Figure 24: N-D Visualizer scatterplot and controls.

The points in the scatterplot collectively represent the data cloud. The numbers at the bottom of the dialog indicate the number of MNF bands to animate. The data cloud rotates along the axes of all of the selected dimensions (MNF bands). Spectral endmembers are located at the convex corners of the data cloud. Clusters of pixels at various corners of the data cloud can be identified and represented by distinct classes and colors. Figure 26 shows an example.




Figure 25: Data cloud with 10 endmember classes. Three classes are circled for illustration.

Refer to the ENVI Help for details on using the *N*-Dimensional Visualizer and selecting endmembers.

Once the endmembers have been extracted from the MNF-transformed data, they can be used in spectral mapping methods.

Spectral Mapping Methods

Various mapping methods are used in ENVI to project spectral endmembers from *N*-dimensional space back to their locations in the imagery and to their spectral signatures. They can be categorized as follows:

- **Spectral similarity methods:** Separate materials into spectrally similar groups via image classification.
- **Spectral matching methods:** Use specific absorption features in reflectance spectra to identify materials.
- **Unmixing methods:** Estimate the sub-pixel abundances of materials in an image using spectral unmixing techniques.

Spectral Similarity

Spectral Angle Mapper (SAM) is the exclusive spectral similarity measure used in ENVI. It is a whole-pixel technique that is primarily used for target detection; however, it can also be used to create spectral classification maps (Kruse et al. 1993). The input image must be converted to apparent reflectance so that the data units are the same as the library units.

SAM uses an *N*-dimensional angle (in radians) to match pixels to reference spectra. Smaller angles represent closer matches to the reference spectrum. Pixels farther away than the specified maximum angle threshold are not classified.

The endmember spectra used by SAM can come from spectral libraries or directly from an image using ROIs that were exported from the *N*-D Visualizer.

Where to Find It

- Included in the Spectral Hourglass Workflow
- ENVI Toolbox → Classification → Supervised Classification → Spectral Angle Mapper Classification

The result of SAM classification is a set of rule images, one for each selected material. Darker pixels in the rule image represent smaller spectral angles and thus image spectra that are more similar to the reference spectra. A classification image is also created that assigns each pixel to a class that represents one of the materials from the selected library spectrum. The classification chooses the library spectrum that has the smallest spectral angle with each pixel in the input image. Thus, the classification image shows the best match to a given material for each pixel. Figure 27 shows an example.

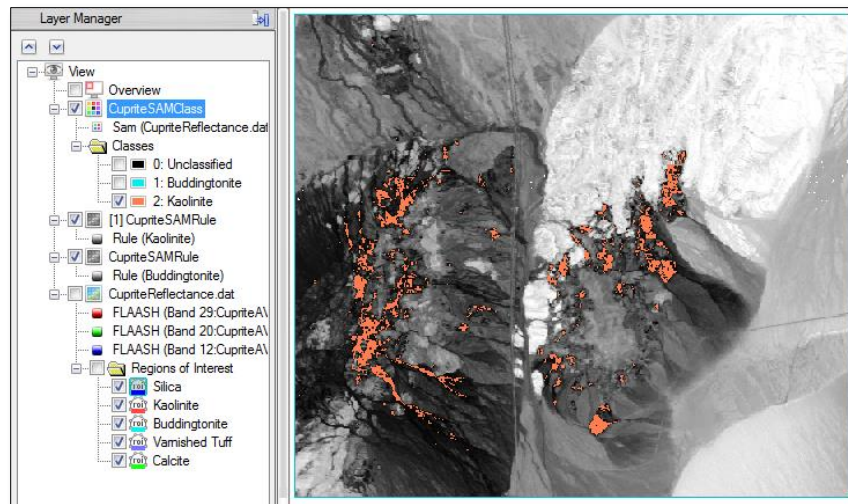


Figure 26: SAM rule image for the mineral kaolinite in AVIRIS imagery, overlaid with the kaolinite class (colored pink).

Spectral Matching

Whole-pixel spectral matching methods such as Spectral Feature Fitting (SFF) and Multi-Range SFF compare specific absorption features from the reflectance spectra of image pixels to those of a known reference spectrum. The basic premise is that pixels whose absorption features match well in width and depth are likely to have a higher abundance of the material of interest, although this requires some caution.

SFF works best with materials that have unique and detailed absorption features such as minerals and some man-made features. In contrast, it does not work well with general categories of materials such as vegetation and water. With some minerals such as hematite, the width and depth of absorption features are affected by crystallinity (Figure 28). So it is important to use the most distinct reference spectrum as possible or include reference spectra that have varying crystallinities.




Figure 27: Plot showing the apparent wavelength centers for absorption features in the mineral hematite.

SFF involves the removal of a *continuum* from the image and reference spectra. This normalizes each spectrum in the input image by comparing it to a continuum curve, thus defining a common baseline from which to measure absorption feature depth and position. A continuum curve is a mathematical function formed by fitting a convex hull to the spectrum (Figure 29). The resulting reflectance curve (with continuum removed) gives a better picture of band centers and depths.

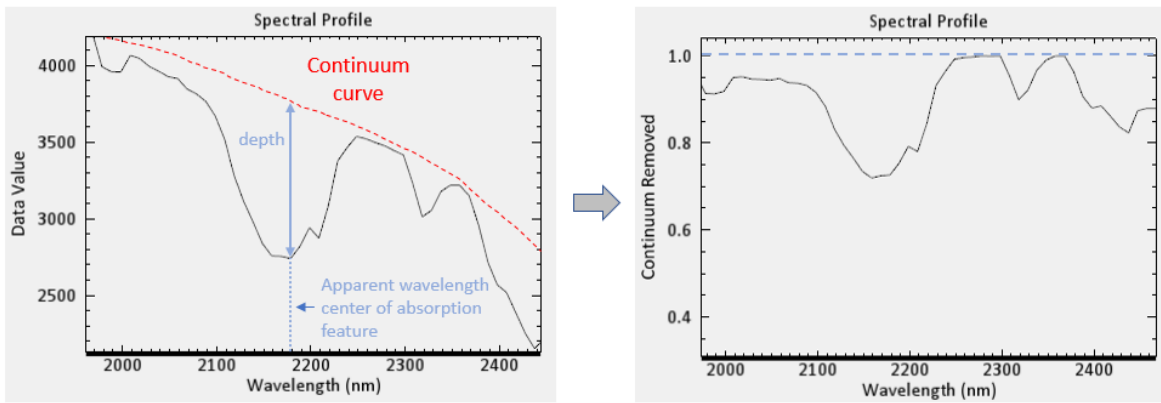


Figure 28: Reflectance curve fitted with a continuum curve (left) and reflectance curve after removing the continuum (right).

Where to Find It

- ENVI Toolbox → Spectral → Mapping Methods → Continuum Removal
- ENVI Spectral Profile plots → Y-axis options → Continuum Removed
- ENVI Toolbox → Spectral → Mapping Methods → Spectral Feature Fitting
- ENVI Toolbox → Spectral → Mapping Methods → MultiRange Spectral Feature Fitting (multiple tools)

The result of SFF is a set of two images for each endmember contained in the spectral library (Figure 30). The first is a Scale image that indicates pixels whose absorption features are similar in depth and width to those of the selected library spectrum. The pixel values indicate the probability (ranging from 0 to 1) that a particular material occurs in that pixel. A value of 1 indicates a perfect match (Kruse 1994). The second image that is created is an RMS error image where dark areas indicate low RMS errors. Comparing the two images can help locate pixels that have a good spectral fit for the selected material absorption feature.




Figure 29: SFF results for the mineral alunitic in AVIRIS imagery. The Scale image is on the left, and the RMS Error image is on the right.

Unmixing Methods

The spectrum from a pixel is a measurement of reflectance coming from all materials within that pixel. If the area within the pixel contains only one material (as with a large body of water or a large roof), that pixel's spectrum represents a single material. (Figure 31, rooftop). However, if the area within the pixel contains more than one material, then it is a mixed pixel (Figure 31, grass + bare soil). Any material in the mixed pixel is at a sub-pixel quantity (HGS 2014).

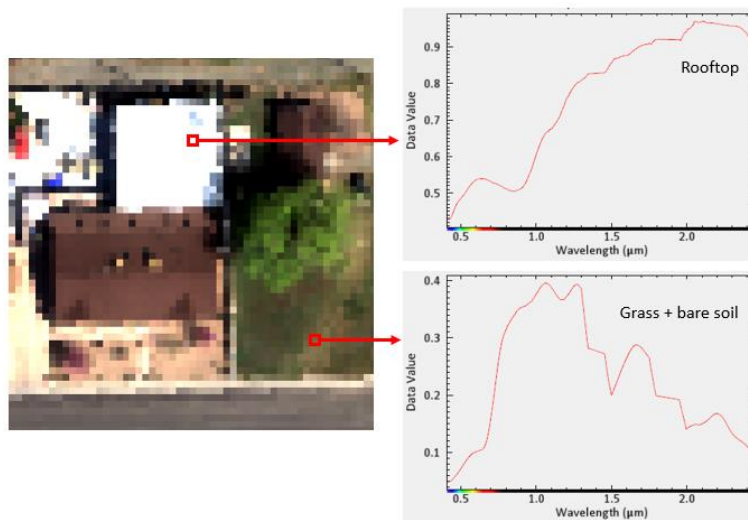


Figure 30: Spectra of a pure pixel (rooftop) and mixed pixel (grass + bare soil).

The reflectance spectrum of a mixed pixel is assumed to be a linear combination of the spectra of each material. In a linear mixture model, a reflectance spectrum from a single pixel is a weighted average of reflectance spectra from each material in the pixel. Figure 32 shows a diagram of a mixed pixel consisting of three materials and how the spectrum for that pixel is measured.

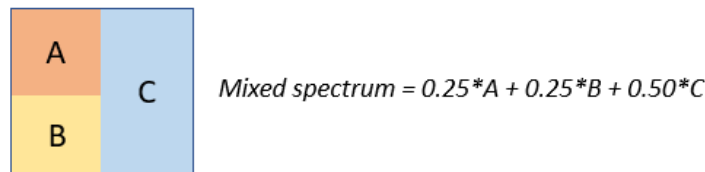


Figure 31: Mixed pixel consisting of three materials.

When mixing occurs, it is not possible to determine what materials are present in the pixel by just measuring the spectrum of that pixel. Spectral unmixing is used to decompose the measured spectrum of a mixed pixel into its constituent spectra (endmembers) as well as the fractional abundances that indicate the proportion of each endmember in the pixel (Sidike et al. 2012). The fractional abundances should be non-negative. For example, in Figure 32:

- The fractional abundance of material A is 25%.
- The fractional abundance of material B is 25%.
- The fractional abundance of material C is 50%.

The most basic form of unmixing—referred to as *linear unmixing*—assumes that the apparent reflectance of each pixel in an image is a linear combination of the apparent reflectance of each material present in the pixel.

Sidike et al. (2012) discuss the concepts of spectral unmixing in more detail, along with comparisons of the various methods used.

The Spectral Hourglass Workflow provides an option for Linear Spectral Unmixing. This is a *complete* linear unmixing method that requires all of the endmembers to be identified in a scene. However, the high dimensionality of hyperspectral data and the complexity of the spectral endmembers means that a complete linear unmixing is not always possible or even desired (Boardman, Kruse, and Green 1995). A complete linear unmixing model only produces valid results if all of the image endmembers have been correctly identified. Additionally, many materials (such as minerals) do not mix linearly in the real world.

Where to Find It

- Included in the Spectral Hourglass Workflow
- ENVI Toolbox → Spectral → Spectral Unmixing → Linear Spectral Unmixing

The result of Linear Spectral Unmixing is a set of abundance images and one RMS error image. Each abundance image corresponds to an individual endmember. The pixel values in the abundance images indicate the percentage of the pixel consisting of that endmember. The RMS error image uses the results of the abundance image to determine the overall error of all of the endmember abundance values for each pixel. RMS error images should appear as noise.

Linear Spectral Unmixing was one of the earliest spectral tools added to ENVI to address the issue of mixed pixels in coarse-resolution imagery such as AVIRIS. A better option for assessing the sub-pixel abundance of materials is the Mixture-Tuned Matched Filtering (MTMF) target detection method.

Matched Filtering Methods

Matched Filtering (MF) and MTF do not require accurate information about all of the endmembers in the scene. They provide abundance estimates of individual endmembers without relying on information from the other endmembers. Pixel values in the resulting rule images are directly proportional to the fractional abundance of the target materials. See [Selecting Target Detection Methods](#) on page 16. MTF also creates an additional “infeasibility” image for each endmember.

Figure 33 shows an example of a MTF abundance image for one endmember derived from an AVIRIS image of the Cuprite mining district in Nevada, USA. The endmember corresponds to a unique class identified in the N-Dimensional Visualizer. Buddingtonite is the presumed mineral type associated with this each endmember, based on previous studies of the region (Swayze et al. 2014). The image was stretched to show the highest 2% of abundance values, and a color table was applied. This removes most of the background signal.

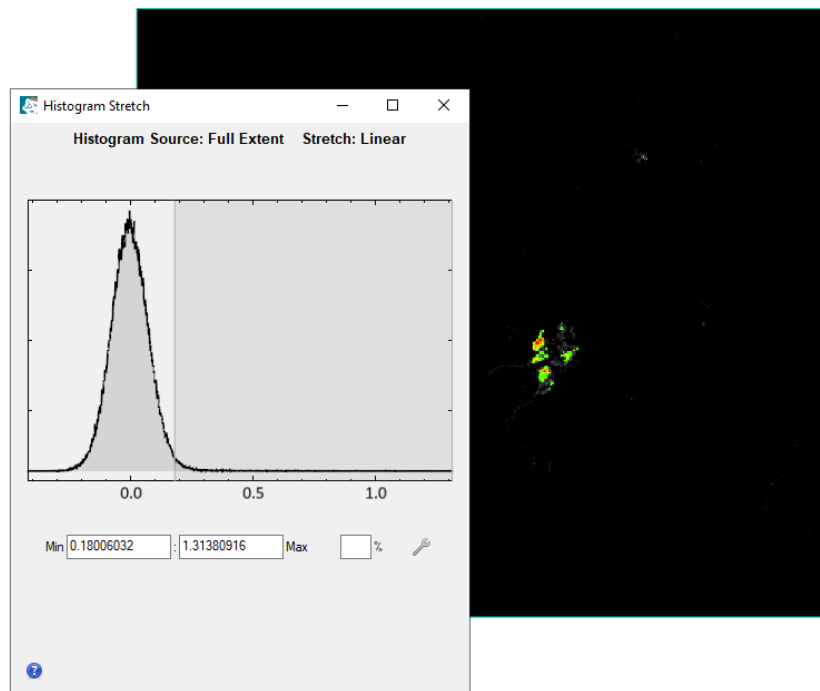


Figure 32: MTF abundance image for buddingtonite.

Another way to visualize a MTF result for a single endmember is to display a scatterplot with the infeasibility band along the Y-axis and the abundance image (called “MF Score”) along the X-axis. Draw a polygon around data points with low infeasibility and high MF Score values. The corresponding pixels are highlighted in the abundance image.

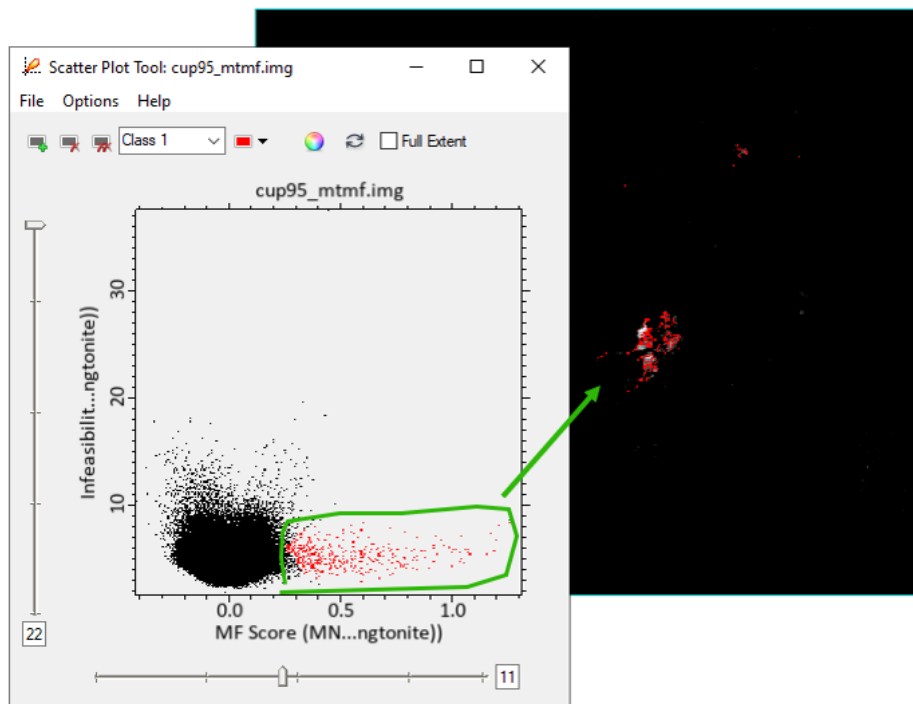


Figure 33: RGB color composite of MTMF abundance images.

Automated methods can provide a starting point for estimating the materials associated with each endmember. The ENVI Material Identification tool was introduced in ENVI 6.0 for this purpose. After plotting a spectral profile of an image pixel, users can further click an **identify** button to estimate that materials associated with the spectrum. The Material Identification dialog shows a ranking of likely materials (Figure 35) based on the comparison of the pixel spectrum with a known library spectrum. Again, this does not provide a definitive answer for what materials comprise a given pixel; however, it provides a starting point for further investigation.

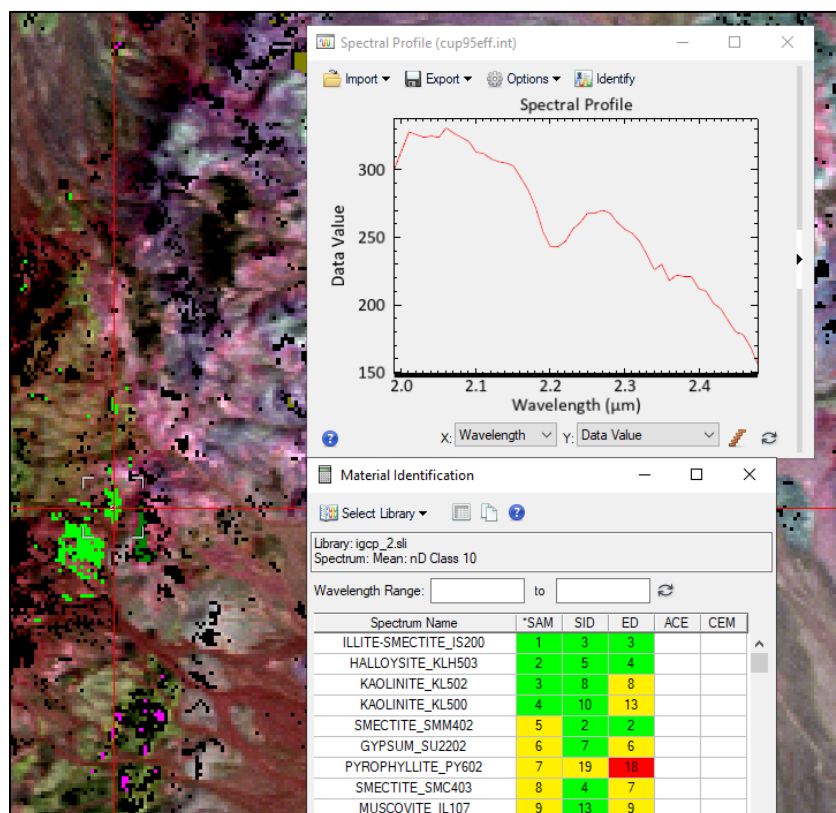


Figure 34: Using the Material Identification tool to estimate the materials associated with an individual pixel spectrum.

SUMMARY

The workflows described in this paper reveal the true power of imaging spectroscopy. Multispectral remote sensing can be used to determine if a given material is present in an image and where it is located. Imaging spectroscopy can provide even further insight by estimating *how much* of the material exists within each pixel. Data from imaging spectrometers are oversampled by nature, which means that they have more spectral channels than inherent data dimensions. This oversampling allows researchers to see inside of pixels to estimate the relative abundance of materials that comprise them.

Imaging spectroscopy requires a unique set of tools and methods, compared to traditional multispectral remote sensing. It is important to know the properties and capabilities of the sensor that acquired the imagery and whether or not its spectral resolution and range will sufficiently identify materials of interest. Requirements for preprocessing vary by application; atmospheric correction and dimensionality reduction techniques such as MNF may or may not be needed depending on the intended result. A particularly important step is preparing image and reference spectra so that they can be directly compared. The success of target detection and spectral mapping methods depends on the quality of the input and reference spectra.

The wide variety of hyperspectral tools that ENVI provides, along with the suggested workflows in this paper, can facilitate the analysis of high-fidelity spectral data.

REFERENCES

- Barry, P. *EO-1/Hyperion Science Data User's Guide*. TRW Space, Defense & Information Systems (2001).
- Bedell, R., and M. Coolbaugh. "Appendix 2: Atmospheric Corrections." In *Remote Sensing and Spectral Geology: Reviews in Economic Geology* 16 (2009), edited by R. Bedell, A. Crosta, and E. Grunsky, pp. 257-263.
- Bernstein, L., X. Jin, B. Gregor, and S. Adler-Golden. "Quick Atmospheric Correction Code: Algorithm Description and Recent Upgrades." *Optical Engineering* 51, No. 11 (2012): 111719-1 to 111719-11.
- Bioucas-Dias, J., A. Plaza, G. Camps-Valls, P. Scheunders, N. Nasrabadi, and J. Chanussot. "Hyperspectral Remote Sensing Data Analysis and Future Challenges." *IEEE Geoscience and Remote Sensing Magazine* (June 2013): 6-36.
- Boardman, J. "Automated Spectral Unmixing of AVIRIS Data Using Convex Geometry Concepts." In *Summaries of the 4th JPL Airborne Earth Science Workshop*, JPL Publication 93-26, Volume 1 (1993): 11-14.
- Boardman, J. "Leveraging the High Dimensionality of AVIRIS Data for Improved Sub-Pixel Target Unmixing and Rejection of False Positives: Mixture Tuned Matched Filtering." In *JPL Proceedings of the 7th Airborne Geoscience Workshop* (1998): 55-56.
- Boardman, J., and F. Kruse. "Automated Spectral Analysis. A Geologic Example Using AVIRIS data, north Grapevine Mountains, Nevada." In *Proceedings of the 10th Thematic Conference on Geologic Remote Sensing* (1994): I-407 to I-418.
- Boardman, J., F. Kruse, and R. Green. "Mapping Target Signatures via Partial Unmixing of AVIRIS Data." In *Summaries of the 5th JPL Airborne Earth Science Workshop*, JPL Publication 95-1, Volume 1 (1995): 23-26.
- Chang, C.-I. "Further Results on Relationship Between Spectral Unmixing and Subspace Projection." *IEEE Transactions on Geosciences and Remote Sensing* 36 (1998): 1030-1032.
- Chang, C.-I. *Hyperspectral Imaging: Techniques for Spectral Detection and Classification*. Kluwer Academic Publishers, 2003.
- Chang, C.-I., J.-M. Liu, B.-C. Chieu, C.-M. Wang, C. Lo, P.-C. Chung, H. Ren, C.-W. Yang, and D.-J. Ma. "A Generalized Constrained Energy Minimization Approach to Subpixel Target Detection for Multispectral Imagery." *Optical Engineering* 39, No. 5 (2000): 1275-1281.
- Chuvieco, E. *Fundamentals of Satellite Remote Sensing: An Environmental Approach*. CRC Press, 2016.
- Clark, R., T. King, M. Klejwa, G. Swayze, and N. Vergo. "High Spectral Resolution Reflectance Spectroscopy of Minerals." *Journal of Geophysical Research* 95, Issue B8 (1990): 12653-12680.
- Datt, B., T. McVicar, T. Van Niel, and D. Jupp. "Preprocessing EO-1 Hyperion Hyperspectral Data to Support the Application of Agricultural Indexes." *IEEE Transactions on Geoscience and Remote Sensing* 41, No. 6 (2003): 1246-1259.
- Du, Q., H. Ren, and C.-I. Chang. "A Comparative Study for Orthogonal Subspace Projection and Constrained Energy Minimization." *IEEE Transactions on Geoscience and Remote Sensing* 41, No. 6 (2003): 1525-1529.
- Du, Q., C.-I. Chang, H. Ren, F. M. D'Amico, and J. Jensen. "New Hyperspectral Discrimination Measure for Spectral Characterization." *Optical Engineering* 43, No. 8 (2004): 1777-1786.
- Ehrlich, M., B. Justice, and B. Baldrige. *Harris Deep Learning Technology*, Harris Geospatial Solutions whitepaper (2016).
- Fowler, J. "Comprehensive Pushbroom and Whiskbroom Sensing for Hyperspectral Remote-Sensing Imaging." In *Proceedings of the 2014 IEEE International Conference on Image Processing* (2014): 684-688.
- Green, A., and M. Craig. "Analysis of Aircraft Spectrometer Data, with Logarithmic Residuals." In *JPL Proceedings of the Airborne Imaging Spectrometer Data Analysis Workshop* (1985): 111-119.
- Green, A., M. Berman, P. Switzer, and M. Craig. "A Transformation for Ordering Multispectral Data in Terms of Image Quality with Implications for Noise Removal." *IEEE Transactions on Geoscience and Remote Sensing* 26, No. 1 (1988): 65-74.
- Harris Geospatial Solutions. *Analysis Fundamentals for Spectral Processing, Exploitation, and Dissemination* (2014).
- Harsanyi, J. *Detection and Classification of Subpixel Spectral Signatures in Hyperspectral Image Sequences*. Ph.D. thesis, University of Maryland (1993).
- Harsanyi, J., and C.-I. Chang. "Hyperspectral Image Classification and Dimensionality Reduction: An Orthogonal Subspace Projection Approach." *IEEE Transactions on Geoscience and Remote Sensing* 32, No. 4 (1994): 779-785.
- Ibarrola-Ulzurrun, E., J. Marcello, and C. Gonzalo-Martin. "Evaluation of Dimensionality Reduction Techniques in Hyperspectral Imagery and Their Application for the Classification of Terrestrial Ecosystems." In *Proceedings of SPIE 10427, Image and Signal Processing for Remote Sensing XXIII* (2017): 104270G-1 to 104270G-10.
- Jin, X., S. Paswaters, and H. Cline. "A Comparative Study of Target Detection Algorithms for Hyperspectral Imagery." In *Proceedings of SPIE 7334, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XV* (2009): 73341W-1 to 73341W-12.
- John Hopkins University (JHU) Applied Physics Laboratory (APL). "CRISM: Compact Reconnaissance Imaging Spectrometer for Mars." <http://crism.jhuapl.edu/index.php>. Accessed August 2018.

- Johnson, S. "Constrained Energy Minimization and the Target-Constrained Interference-Minimized Filter." *Optical Engineering* 42, No. 6 (2003): 1850-1854.
- Kirby, M. *Geometric Data Analysis: An Empirical Approach to Dimensionality Reduction and the Study of Patterns*. John Wiley & Sons, 2001.
- Kosanke, T., and J. Chen. "Hyperspectral Imaging: Geological and Petrophysical Applications to Reservoir Characterization." Unconventional Resources Technology Conference, Austin, TX (2017): 1198-1203.
- Kraut, S., L. Scharf, and R. Butler. "The Adaptive Coherence Estimator: A Uniformly Most-Powerful-Invariant Adaptive Detection Statistic." *IEEE Transactions on Signal Processing* 53, No. 2 (2005): 427-438.
- Kruse, F. "Imaging Spectrometer Data Analysis – A Tutorial." In *Proceedings of the International Symposium on Spectral Sensing Research* (1994): 44-54.
- Kruse, F., G. Raines, and K. Watson. "Analytical Techniques for Extracting Geologic Information From Multichannel Airborne Spectroradiometer and Airborne Imaging Spectrometer Data." In *Proceedings of the 4th Thematic Conference on Remote Sensing for Exploration Geology* (1985): 309-324.
- Kruse, F., A. Lefkoff, J. Boardman, K. Heidebrecht, A. Shapiro, P. Barloon, and A. Goetz. "The Spectral Image Processing System (SIPS) – Interactive Visualization and Analysis of Imaging Spectrometer Data." *Remote Sensing of Environment* 44 (1993): 145-163.
- Manolakis, D., R. Lockwood, and T. Cooley. "Hyperspectral Data Exploitation." In *Hyperspectral Imaging Remote Sensing: Physics, Sensors, and Algorithms*. Cambridge University Press, 2016.
- Manolakis, D., and G. Shaw. "Detection Algorithms in Hyperspectral Imaging Systems: An Overview of Practical Algorithms." *IEEE Signal Processing Magazine* 31, No. 1 (2014): 24-33.
- Manolakis, D., D. Marden, and G. Shaw. "Hyperspectral Image Processing for Automatic Target Detection Applications." *Lincoln Laboratory Journal* 14 (2003): 79-116.
- Martin, S., and T. George. "Applications of Hyperspectral Image Analysis for Precision Agriculture." In *Proceedings of SPIE 10639, Micro- and Nanotechnology Sensors, Systems, and Applications X* (2018): 1063916-1 to 1063916-8.
- Matsunaga, T., S. Yamamoto, and T. Tachikawa. "Detection of Large Point Sources of Carbon Dioxide by a Satellite Hyperspectral Camera." In *7th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing* (2015): 1-4.
- Minu, S., and A. Shetty. "Atmospheric Correction Algorithms for Hyperspectral Imageries: A Review." *International Research Journal of Earth Sciences* 3, No. 5 (2015): 14-18.
- Mitchell, P. "Hyperspectral Digital Imagery Collection Experiment (HYDICE)." In *Proceedings of SPIE 2587, Geographic Information Systems, Photogrammetry, and Geological/Geophysical Remote Sensing* (1995): 70-95.
- NASA Jet Propulsion Laboratory. "AVIRIS Next Generation." <https://aviris-ng.jpl.nasa.gov/index.html>. Accessed August 2018.
- Nascimento, and J. Dias. "Does Independent Component Analysis Play a Role in Unmixing Hyperspectral Data?" *IEEE Transactions on Geoscience and Remote Sensing* 43, No. 1 (2005): 175-187.
- Nasrabadi, N. "Hyperspectral Target Detection: An Overview of Current and Future Challenges." *IEEE Signal Processing Magazine* 31, No. 1 (2014): 34-44.
- National Ecological Observatory Network, 2014. Files were accessed in August 2018, available online at <http://data.neonscience.org> from the National Ecological Observatory Network (NEON), Boulder, CO, USA.
- Pan, Z., J. Huang, and F. Wang. "Multi Range Spectral Feature Fitting for Hyperspectral Imagery in Extracting Oilseed Rape Planting Area." *International Journal of Applied Earth Observation and Geoinformation* 25 (2013): 21-29.
- Probasco, Kevin. "Hyperspectral Imaging: Defense Technology Transfers into Commercial Applications." Photonics for a Better World. <https://photonicsforabetterworld.blogspot.com/2017/09/hyperspectral-imaging-defense.html>. Accessed August 2018.
- Puckrin, E., C. Turcotte, M.-A. Gagnon, J. Bastedo, V. Farley, and M. Chamberland. "Airborne Infrared Hyperspectral Imager for Intelligence, Surveillance and Reconnaissance Applications." In *Proceedings of SPIE 8360, Airborne Intelligence, Surveillance, Reconnaissance (ISR) Systems and Applications IX* (2012): 836004-1 to 836004-10.
- Qu, L., W. Han, H. Lin, Y. Zhu, and L. Zhang. "Estimating Vegetation Fraction Using Hyperspectral Pixel Unmixing Method: A Case Study of a Karst Area in China." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 7, No. 11 (2014): 4559-4565.
- Ren, H., and C.-I. Chang. "Target-Constrained Interference-Minimized Approach to Subpixel Target Detection for Hyperspectral Imagery." *Optical Engineering* 39, No. 12 (2000): 3138-3145.
- Ren, J., J. Zabalza, S. Marshall, and J. Zheng. "Effective Feature Extraction and Data Reduction in Remote Sensing Using Hyperspectral Imaging." *IEEE Signal Processing Magazine* 31, No. 4 (2014): 149-154.
- Roberts, D., Y. Yamaguchi, and R. Lyon. "Comparison of Various Techniques for Calibration of AIS Data." In *JPL Proceedings of the 2nd Airborne Imaging Spectrometer Data Analysis Workshop* (1986): 21-30.

- Robichaud, P., S. Lewis, D. Laes, A. Hudak, R. Kokaly, and J. Zamudio. "Postfire Soil Burn Severity Mapping with Hyperspectral Image Unmixing." *Remote Sensing of Environment* 108, No. 4 (2007): 467-480.
- Sidike, P., J. Khan, M. Alam, and S. Bhuiyan. "Spectral Unmixing of Hyperspectral Data for Oil Spill Detection." In Proceedings of SPIE 8498, *Optics and Photonics for Information Processing VI* (2012): 84981B-1 to 84981B-10.
- Somdatta, C., and S. Chakrabarti. "Pre-processing of Hyperspectral Data: A Case Study of Henry and Lothian Islands in Sunderban Region, West Bengal, India." *International Journal of Geomatics and Geosciences* 2, No. 2 (2001): 490-501.
- Swayze, G., R. Clark, A. Goetz, K Livo, G. Breit, F. Kruse, S. Stutley, L. Snee, H. Lowers, J. Post, R. Stoffregen, and R. Ashley. "Mapping Advanced Argillic Alteration at Cuprite, Nevada Using Imaging Spectroscopy." *Economic Geology* 109, No. 5 (2014): 1179-1221.
- Villa, A., J. Chanussot, C. Jutten, J. Benedicksson, and S. Moussaoui. "On the Use of ICA for Hyperspectral Image Analysis." In *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium* (2008), IV-97 to IV-100.
- Yang, C. "Hyperspectral Imagery for Mapping Crop Yield for Precision Agriculture." In *Hyperspectral Imaging Technology in Food and Agriculture*. Springer, 2015.
- Yokoya, N., N. Miyamura, and A. Iwasaki. "Preprocessing of Hyperspectral Imagery with Consideration of Smile and Keystone Properties." In *Proceedings of SPIE 7857, Multispectral, Hyperspectral, and Ultraspectral Remote Sensing Technology, Techniques, and Applications III* (2010): 78570B-1 to 78570B-9.

