Introduction to

Hyperspectral Imaging

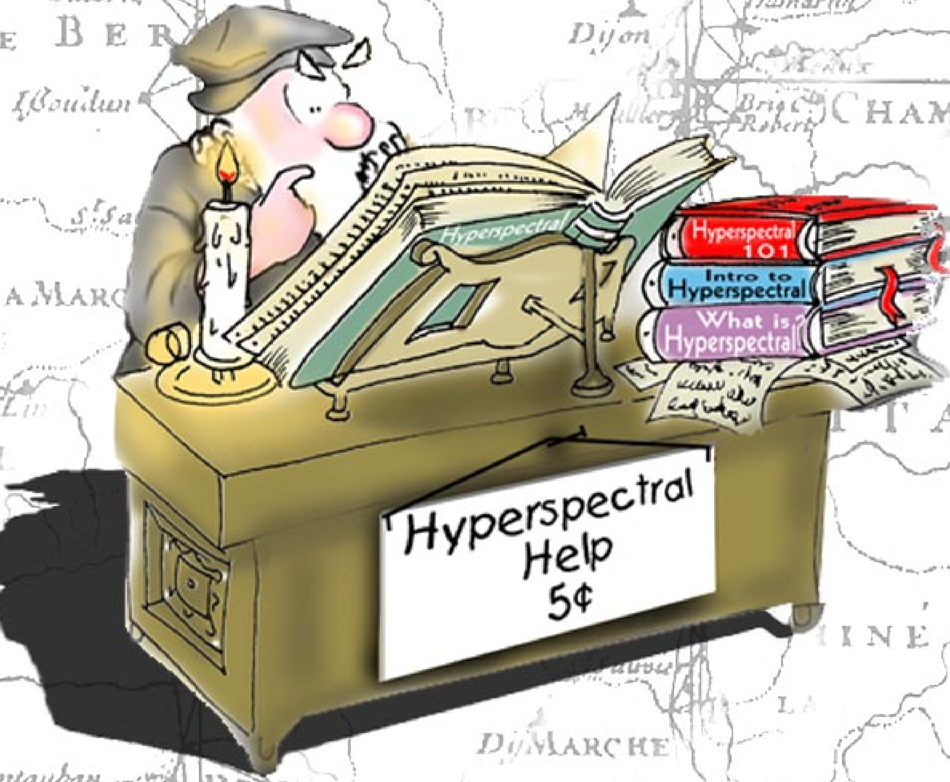

with TNTmips®

Before Getting Started

For much of the past decade, hyperspectral imaging has been an area of active research and development, and hyperspectral images have been available only to researchers. With the recent appearance of commercial airborne hyperspectral imaging systems, hyperspectral imaging is poised to enter the mainstream of remote sensing. Hyperspectral images will find many applications in resource management, agriculture, mineral exploration, and environmental monitoring. But effective use of hyperspectral images requires an understanding of the nature and limitations of the data and of various strategies for processing and interpreting it. This booklet aims to provide an introduction to the fundamental concepts in the field of hyperspectral imaging.

Sample Data Some illustrations in this booklet show analysis results for a hyperspectral scene of Cuprite, Nevada. This scene was acquired using the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), which is operated by the NASA Jet Propulsion Laboratory. The same scene is used in the exercises in the companion tutorial booklet *Analyzing Hyperspectral Images*. You can download this scene in the TNTmips Project File format (along with associated sample data) from the MicroImages web site.

More Documentation This booklet is intended only as a general introduction to hyperspectral imaging. In TNTmips, hyperspectral images can be processed and analyzed using the Hyperspectral Analysis process (choose Image / Hyperspectral Analysis from the TNTmips menu). For an introduction to this process, consult the tutorial booklet entitled *Analyzing Hyperspectral Images*. Additional background information can be found in the booklet *Introduction to Remote Sensing of Environment (RSE).*

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.

Randall B. Smith, Ph.D., 5 January 2012

You can print or read this booklet in color from MicroImages' web site. The web site is also your source for the newest Getting Started 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 Imaging

Multispectral remote sensors such as the Landsat Thematic Mapper and SPOT XS produce images with a few relatively broad wavelength bands. Hyperspectral remote sensors, on the other hand, collect image data simultaneously in dozens or hundreds of narrow, adjacent spectral bands. These measurements make it possible to derive a continuous spectrum for each image cell, as shown in the illustration below. After adjustments for sensor, atmospheric, and terrain effects are applied, these image spectra can be compared with field or laboratory reflectance spectra in order to recognize and map surface materials such as particular types of vegetation or diagnostic minerals associated with ore deposits.

Hyperspectral images contain a wealth of data, but interpreting them requires an understanding of exactly what properties of ground materials we are trying to measure, and how they relate to the measurements actually made by the hyperspectral sensor.

The technological background of hyperspectral sensors is discussed on page 4. Pages 5-10 introduce the concepts of spectral reflectance of natural materials, spectra as vectors in n-dimensional spectral space, and spectral mixing. Factors contributing to the measured radiance values in an image are detailed on pages 11-13, followed by methods for converting from radiance to reflectance on pages 14-15. Strategies for analyzing hyperspectral images are discussed on pages 16 - 21, and a list of literature references is provided on pages 22-23.

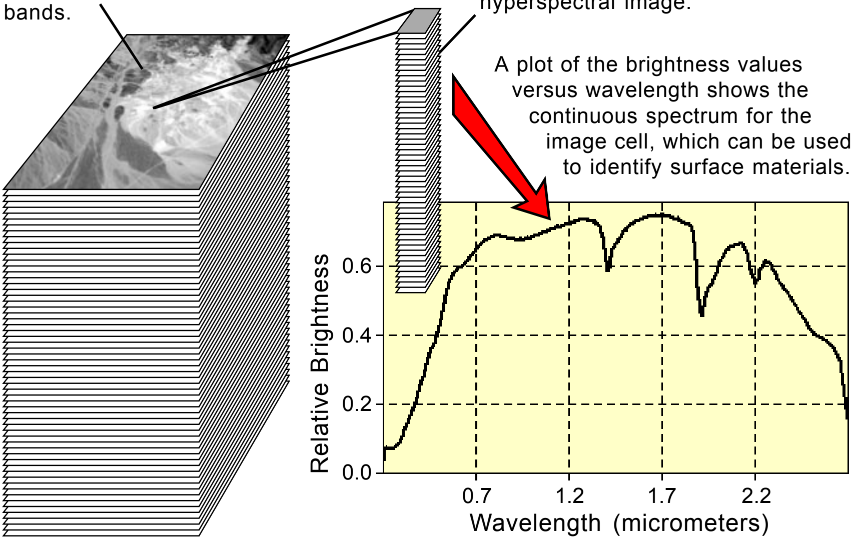

Images acquired simultaneously in many narrow, adjacent wavelength bands.

Set of brightness values for a single raster cell position in the hyperspectral image.

A plot of the brightness values versus wavelength shows the continuous spectrum for the image cell, which can be used to identify surface materials.

The Imaging Spectrometer

Hyperspectral images are produced by instruments called *imaging spectrometers*. The development of these complex sensors has involved the convergence of two related but distinct technologies: *spectroscopy* and the *remote imaging* of Earth and planetary surfaces.

Spectroscopy is the study of light that is emitted by or reflected from materials and its variation in energy with wavelength. As applied to the field of optical remote sensing, spectroscopy deals with the spectrum of sunlight that is diffusely reflected (scattered) by materials at the Earth's surface. Instruments called spectrometers (or spectroradiometers) are used to make ground-based or laboratory measurements of the light reflected from a test material. An optical dispersing element such as a grating or prism in the spectrometer splits this light into many narrow, adjacent wavelength bands and the energy in each band is measured by a separate detector. By using hundreds or even thousands of detectors, spectrometers can make spectral measurements of bands as narrow as 0.01 micrometers over a wide wavelength range, typically at least 0.4 to 2.4 micrometers (visible through middle infrared wavelength ranges).

Remote imagers are designed to focus and measure the light reflected from many adjacent areas on the Earth's surface. In many digital imagers, sequential measurements of small areas are made in a consistent geometric pattern as the sensor platform moves and subsequent processing is required to assemble them into an image. Until recently, imagers were restricted to one or a few relatively broad wavelength bands by limitations of detector designs and the requirements of data storage, transmission, and processing. Recent advances in these areas have allowed the design of imagers that have spectral ranges and resolutions comparable to ground-based spectrometers.

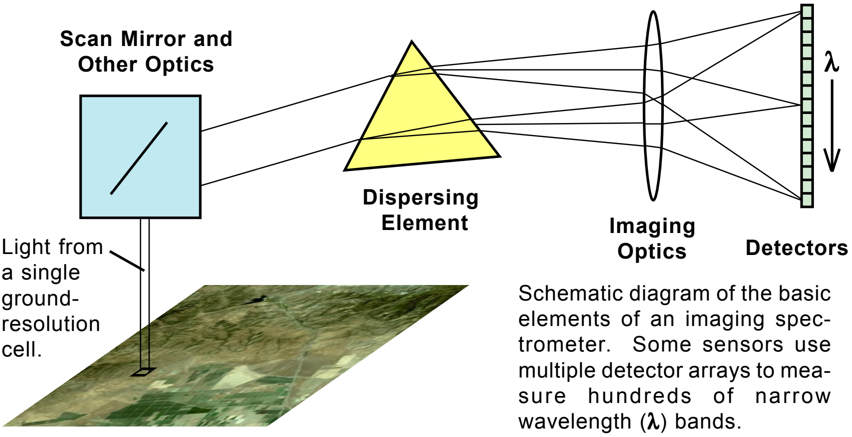

Schematic diagram of the basic elements of an imaging spectrometer. Some sensors use multiple detector arrays to measure hundreds of narrow wavelength (λ) bands.

Spectral Reflectance

In reflected-light spectroscopy the fundamental property that we want to obtain is *spectral reflectance*: the ratio of reflected energy to incident energy as a function of wavelength. Reflectance varies with wavelength for most materials because energy at certain wavelengths is scattered or absorbed to different degrees. These reflectance variations are evident when we compare spectral reflectance curves (plots of reflectance versus wavelength) for different materials, as in the illustration below. Pronounced downward deflections of the spectral curves mark the wavelength ranges for which the material selectively absorbs the incident energy. These features are commonly called *absorption bands* (not to be confused with the separate *image bands* in a multispectral or hyperspectral image). The overall shape of a spectral curve and the position and strength of absorption bands in many cases can be used to identify and discriminate different materials. For example, vegetation has higher reflectance in the near infrared range and lower reflectance of red light than soils.

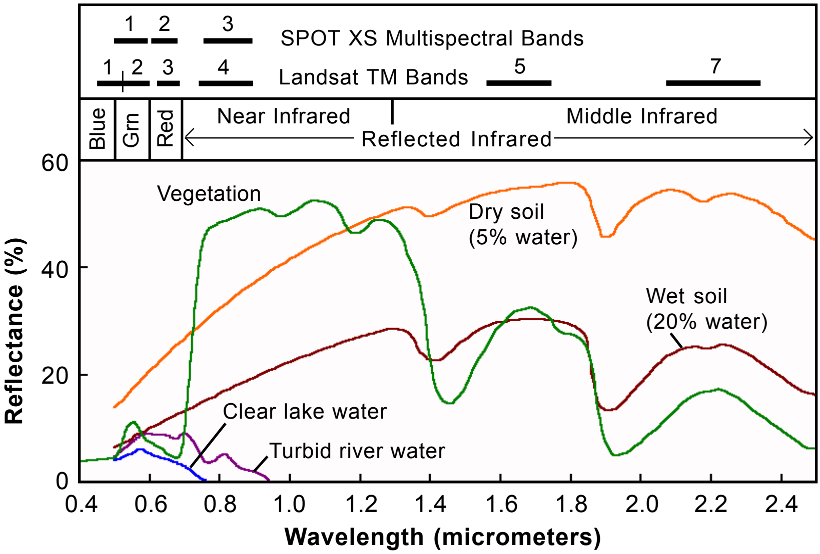

Representative spectral reflectance curves for several common Earth surface materials over the visible light to reflected infrared spectral range. The spectral bands used in several multispectral satellite remote sensors are shown at the top for comparison. Reflectance is a unitless quantity that ranges in value from 0 to 1.0, or it can be expressed as a percentage, as in this graph. When spectral measurements of a test material are made in the field or laboratory, values of incident energy are also required to calculate the material's reflectance. These values are either measured directly or derived from measurements of light reflected (under the same illumination conditions as the test material) from a standard reference material with known spectral reflectance.

Mineral Spectra

In inorganic materials such as minerals, chemical composition and crystalline structure control the shape of the spectral curve and the presence and positions of specific absorption bands. Wavelength-specific absorption may be caused by the presence of particular chemical elements or ions, the ionic charge of certain elements, and the geometry of chemical bonds between elements, which is governed in part by the crystal structure.

The illustration below shows spectra of some common minerals that provide examples of these effects. In the spectrum of hematite (an iron-oxide mineral), the strong absorption in the visible light range is caused by ferric iron (Fe^{+3}). In calcite, the major component of limestone, the carbonate ion (CO_3^{-2}) is responsible for the series of absorption bands between 1.8 and 2.4 micrometers (µm). Kaolinite and montmorillonite are clay minerals that are common in soils. The strong absorption band near 1.4 µm in both spectra, along with the weak 1.9 µm band in kaolinite, are due to hydroxide ions (OH^{-1}), while the stronger 1.9 µm band in montmorillonite is caused by bound water molecules in this hydrous clay. In contrast to these examples, orthoclase feldspar, a dominant mineral in granite, shows almost no significant absorption features in the visible to middle infrared spectral range.

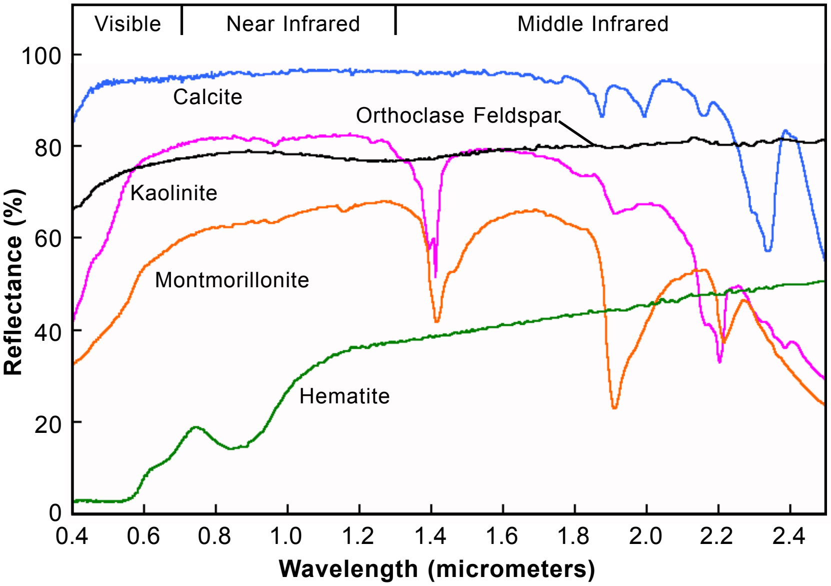

Reflectance spectra of some representative minerals (naturally occurring chemical compounds that are the major components of rocks and soils).

Plant Spectra

The spectral reflectance curves of healthy green plants also have a characteristic shape that is dictated by various plant attributes. In the visible portion of the spectrum, the curve shape is governed by absorption effects from chlorophyll and other leaf pigments. Chlorophyll absorbs visible light very effectively but absorbs blue and red wavelengths more strongly than green, producing a characteristic small reflectance peak within the green wavelength range. As a consequence, healthy plants appear to us as green in color. Reflectance rises sharply across the boundary between red and near infrared wavelengths (sometimes referred to as the *red edge*) to values of around 40 to 50% for most plants. This high near-infrared reflectance is primarily due to interactions with the internal cellular structure of leaves. Most of the remaining energy is transmitted, and can interact with other leaves lower in the canopy. Leaf structure varies significantly between plant species, and can also change as a result of plant stress. Thus species type, plant stress, and canopy state all can affect near infrared reflectance measurements. Beyond 1.3 µm reflectance decreases with increasing wavelength, except for two pronounced water absorption bands near 1.4 and 1.9 µm.

At the end of the growing season leaves lose water and chlorophyll. Near infrared reflectance decreases and red reflectance increases, creating the familiar yellow, brown, and red leaf colors of autumn.

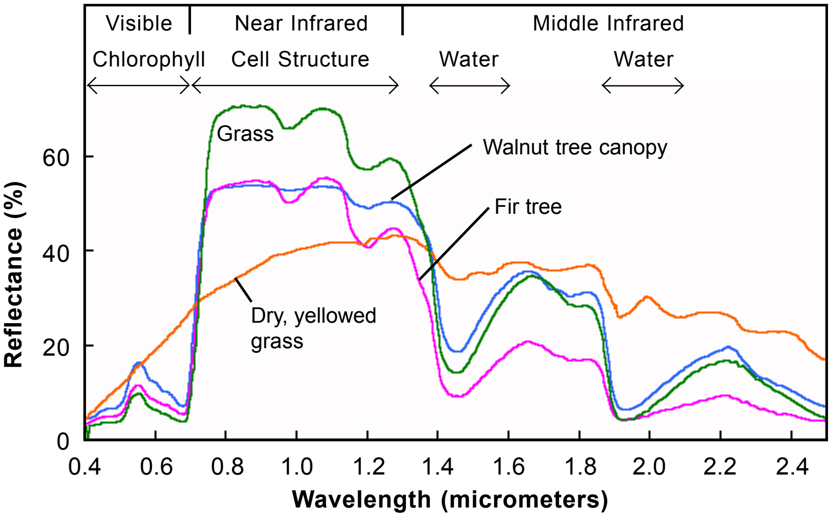

Reflectance spectra of different types of green vegetation compared to a spectral curve for senescent (dry, yellowed) leaves. Different portions of the spectral curves for green vegetation are shaped by different plant components, as shown at the top.

Spectral Libraries

Several libraries of reflectance spectra of natural and man-made materials are available for public use. These libraries provide a source of reference spectra that can aid the interpretation of hyperspectral and multispectral images.

ASTER Spectral Library This library has been made available by NASA as part of the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) imaging instrument program. It includes spectral compilations from NASA's Jet Propulsion Laboratory, Johns Hopkins University, and the United States Geological Survey (Reston). The ASTER spectral library currently contains nearly 2000 spectra, including minerals, rocks, soils, man-made materials, water, and snow. Many of the spectra cover the entire wavelength region from 0.4 to 14 µm. The library is accessible interactively via the Worldwide Web at **http://speclib.jpl.nasa.gov.** You can search for spectra by category, view a spectral plot for any of the retrieved spectra, and download the data for individual spectra as a text file. These spectra can be imported into a TNTmips spectral library. You can also order the ASTER spectral library on CD-ROM at no charge from the above web address.

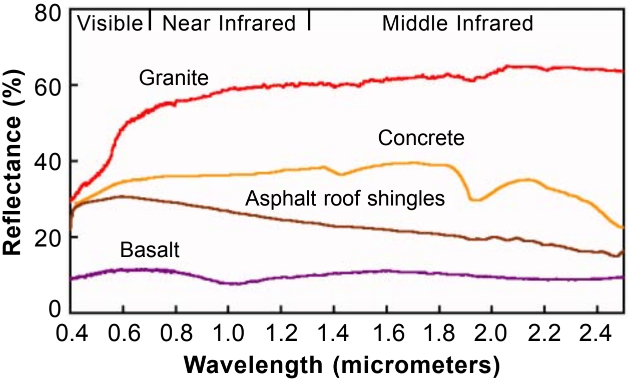

Sample spectra from the ASTER Spectral Library. ASTER will be one of the instruments on the planned EOS AM-1 satellite, and will record image data in 14 channels from the visible through thermal infrared wavelength regions as part of NASA's Earth Science Enterprise program.

USGS Spectral Library The United States Geological Survey Spectroscopy Lab in Denver, Colorado has compiled a library of about 500 reflectance spectra of minerals and a few plants over the wavelength range from 0.2 to 3.0 µm. This library is accessible online at

http://speclab.cr.usgs.gov/spectral.lib04/spectral-lib04.html.

You can browse individual spectra online, or download the entire library. The USGS Spectral library is also included as a standard reference library in the TNTmips Hyperspectral Analysis process.

Plotting Spectra in Spectral Space

The spectral plots on the previous pages provide a convenient way to visualize the differences in spectral properties between different materials, especially when we are comparing only a few spectra. Spectral plots are an important tool to use when you explore a hyperspectral image. But to understand how a computer compares and discriminates among a large number of spectra, it is useful to consider other conceptual ways of representing spectra.

A reflectance spectrum consists of a set of reflectance values, one for each spectral channel (band). Each of these channels can be considered as one dimension in a hypothetical n-dimensional spectral space, where n is the number of spectral channels. If we plot the measured reflectance value for each spectral channel on its respective coordinate axis, we can use these coordinates to specify the location of a point in spectral space that mathematically represents that particular spectrum. A simple two-band example is shown in the illustration. The designated point can also be treated mathematically as the end point of a vector that begins at the origin of the coordinate system. Spectra with the same shape but differing overall reflectance (albedo) plot as vectors with the same orientation but with endpoints at different distances from the origin. Shorter spectral vectors represent darker spectra and longer vectors represent brighter spectra.

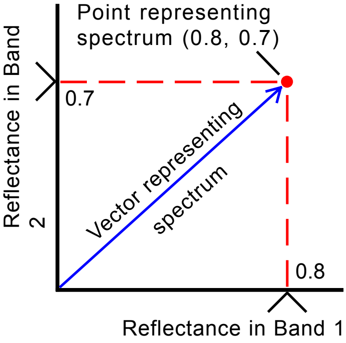

N-dimensional plot of a reflectance spectrum for a hypothetical 2-band case (n = 2).

It may be difficult to visualize such a plot for an image involving more than three wavelength bands, but it is mathematically possible to construct a hyperdimensional spectral space defined by dozens or hundreds of mutually-perpendicular coordinate axes. Each spectrum being considered occupies a position in this n-dimensional spectral space. Similarity between spectra can be judged by the relative closeness of these positions (spectral distance) or by how small the angle is between the spectral vectors.

The spectral reflectance curves shown on the previous pages for various materials represent “averages” or “typical examples”. All natural materials exhibit some variability in composition and structure that results in variability in their reflectance spectra. If we obtain spectra from a number of examples of a material, the resulting spectral points will define a small cloud in n-dimensional spectral space, rather than plotting at one single location.

Spatial Resolution and Mixed Spectra

An imaging spectrometer makes spectral measurements of many small patches of the Earth's surface, each of which is represented as a pixel (raster cell) in the hyperspectral image. The size of the ground area represented by a single set of spectral measurements defines the spatial resolution of the image and depends on the sensor design and the height of the sensor above the surface. NASA's Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), for example, has a spatial resolution of 20 meters when flown at its typical altitude of 20 kilometers, but a 4-meter resolution when flown at an altitude of 4 kilometers.

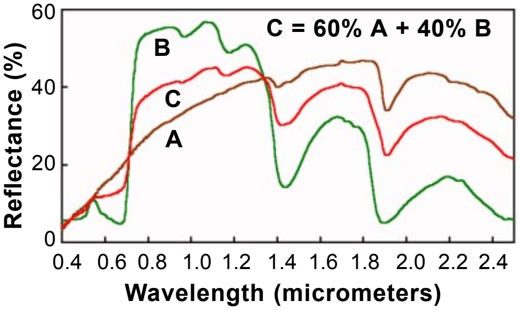

Example of a composite spectrum (C) that is a linear mixture of two spectra: A (dry soil) and B (green vegetation).

When the size of the ground resolution cell is large, it is more likely that more than one material contributes to an individual spectrum measured by the sensor. The result is a *composite* or *mixed* spectrum, and the "pure" spectra that contribute to the mixture are called *endmember* spectra.

Spectral mixtures can be macroscopic or intimate. In a macroscopic mixture each reflected photon interacts with only one surface material. The energy reflected from the materials combines additively, so that each material's contribution to the composite spectrum is directly proportional to its area within the pixel. An example of such a linear mixture is shown in the illustration above, which could represent a patchwork of vegetation and bare soil. In spectral space each endmember spectrum defines the end of a mixing line (for two endmembers) or the corner of a mixing space (for greater numbers of endmembers). Later we will discuss how the endmember fractions can be calculated for each pixel. In an intimate mixture, such as the microscopic mixture of mineral particles found in soils, a single photon interacts with more than one material. Such mixtures are nonlinear in character and therefore more difficult to unravel.

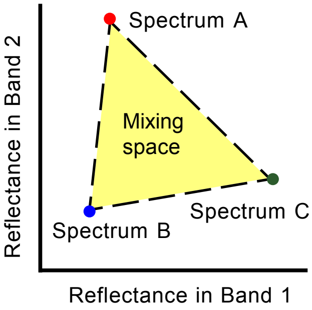

N-dimensional plot of three endmember spectra for a hypothetical 2-band case. All spectra that are mixtures of A, B, and C alone must lie within the mixing space.

Radiance and Reflectance

From the discussions on the preceding pages, it should be clear that spectral reflectance is a property of ground features that we would like to be able to measure precisely and accurately using an airborne or satellite hyperspectral sensor. But look at the brightness spectrum in the illustration below. This is the average of 25 image spectra measured by the AVIRIS sensor over a bright dry lake bed surface in the Cuprite, Nevada scene. The input spectra have been adjusted for sensor effects using on-board calibration data, but no other transformations have been applied.

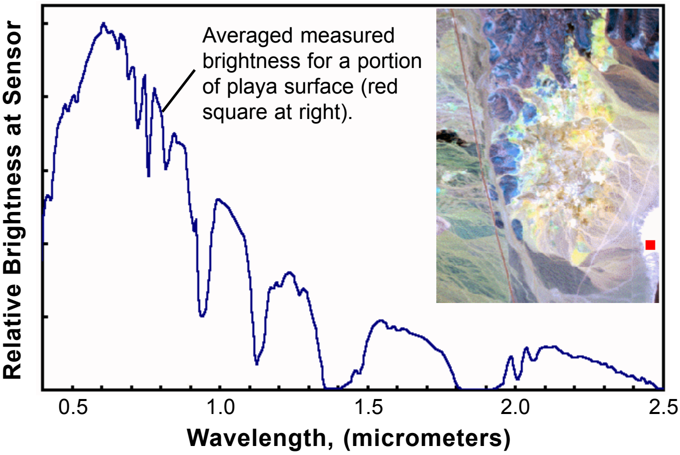

This spectrum does not bear much resemblance to the reflectance spectra illustrated previously. This is because the sensor has simply measured the amount of reflected light reaching it in each wavelength band (*spectral radiance*), in this case from an altitude of 20 kilometers. The spectral reflectance of the surface materials is only one of the factors affecting these measured values. The spectral reflectance curve for the sample area is actually relatively flat and featureless.

In addition to surface reflectance, the spectral radiance measured by a remote sensor depends on the spectrum of the input solar energy, interactions of this energy during its downward and upward passages through the atmosphere, the geometry of illumination for individual areas on the ground, and characteristics of the sensor system. These additional factors not only affect our ability to retrieve accurate spectral reflectance values for ground features, but also introduce additional within-scene variability which hampers comparisons between individual image cells. These factors are discussed in more detail on the next two pages.

Illumination Factors

Source Illumination The figure below shows a typical solar irradiance curve for the top of the Earth's atmosphere. The incoming solar energy varies greatly with wavelength, peaking in the range of visible light. The spectrum of incoming solar energy at the time an image was acquired must be known, assumed, or derived indirectly from other measurements in order to convert image radiance values to reflectance.

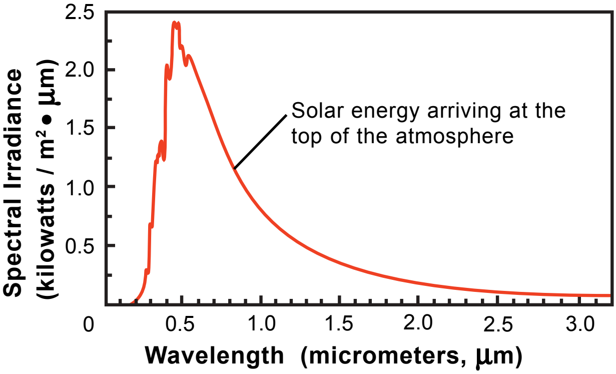

Illumination Geometry The amount of energy reflected by an area on the ground depends on the amount of solar energy illuminating the area, which in turn depends on the *angle of incidence*: the angle between the path of the incoming energy and a line perpendicular to the ground surface. Specifically, the energy received at each wavelength (Eg) varies as the cosine of the angle of incidence (θ): Eg = Eo **x** cos θ, where Eo is the amount of incoming energy. The energy received by any ground area therefore varies as the sun's height changes with time of day and season. If the terrain is not flat, the energy received also varies instantaneously across a scene because of differences in slope angle and direction.

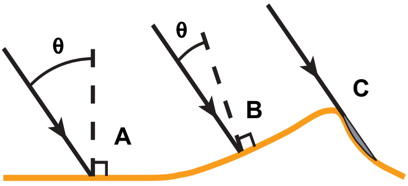

Illumination differences can arise from differing incidence angles (**θ**) as for **A** and **B**, or from shadowing (**C**).

Shadowing The amount of illumination received by an area can also be reduced by shadows. Shadows cast by topographic features or clouds can affect areas including many contiguous image cells. Trees, crop rows, rock outcrops, or other small objects can also cast shadows that are confined to an individual image cell. Both types of shadows have the effect of lowering the measured brightness across all wavelengths for the affected pixels.

Atmospheric and Sensor Effects

Atmospheric Effects Even a relatively clear atmosphere interacts with incoming and reflected solar energy. For certain wavelengths these interactions reduce the amount of incoming energy reaching the ground and further reduce the amount of reflected energy reaching an airborne or satellite sensor. The transmittance of the atmosphere is reduced by absorption by certain gases and by scattering by gas molecules and particulates. These effects combine to produce the transmittance curve illustrated below. The pronounced absorption features near 1.4 and 1.9 µm, caused by water vapor and carbon dioxide, reduce incident and reflected energy almost completely, so little useful information can be obtained from image bands in these regions. Not shown by this curve is the effect of light scattered upward by the atmosphere. This scattered light adds to the radiance measured by the sensor in the visible and near-infrared wavelengths, and is called *path radiance*. Atmospheric effects may also differ between areas in a single scene if atmospheric conditions are spatially variable or if there are significant ground elevation differences that vary the path length of radiation through the atmosphere.

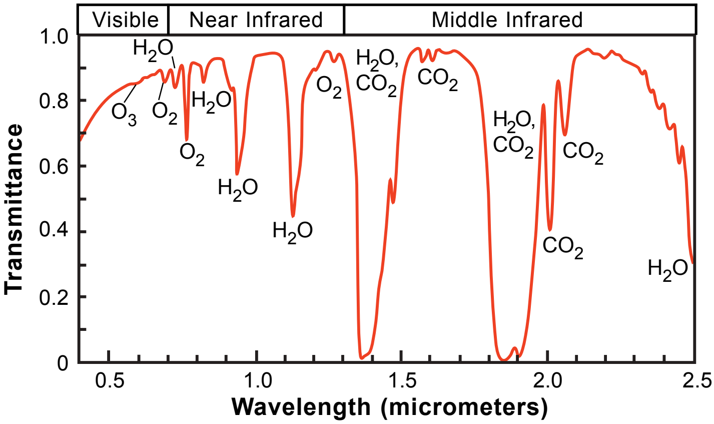

Plot of atmospheric transmittance versus wavelength for typical atmospheric conditions. Transmittance is the proportion of the incident solar energy that reaches the ground surface. Absorption by the labeled gases causes pronounced lows in the curve, while scattering is responsible for the smooth decrease in transmittance with decreasing wavelength in the near infrared through visible wavelength range.

Sensor Effects A sensor converts detected radiance in each wavelength channel to an electric signal which is scaled and quantized into discrete integer values that represent "encoded" radiance values. Variations between detectors within an array, as well as temporal changes in detectors, may require that raw measurements be scaled and/or offset to produce comparable values.

Reflectance Conversion I

In order to directly compare hyperspectral image spectra with reference reflectance spectra, the encoded radiance values in the image must be converted to reflectance. A comprehensive conversion must account for the solar source spectrum, lighting effects due to sun angle and topography, atmospheric transmission, and sensor gain. In mathematical terms, the ground reflectance spectrum is multiplied (on a wavelength per wavelength basis) by these effects to produce the measured radiance spectrum. Two other effects contribute in an additive fashion to the radiance spectrum: sensor offset (internal instrument noise) and path radiance due to atmospheric scattering. Several commonly used reflectance conversion strategies are discussed below and on the following page. Some strategies use only information drawn from the image, while others require varying degrees of knowledge of the surface reflectance properties and the atmospheric conditions at the time the image was acquired.

Flat Field Conversion This image-based method requires that the image include a uniform area that has a relatively flat spectral reflectance curve. The mean spectrum of such an area would be dominated by the combined effects of solar irradiance and atmospheric scattering and absorption The scene is converted to “relative” reflectance by dividing each image spectrum by the flat field mean spectrum. The selected flat field should be bright in order to reduce the effects of image noise on the conversion. Since few if any materials in natural landscapes have a completely flat reflectance spectrum, finding a suitable “flat field” is difficult for most scenes. For desert scenes, salt-encrusted dry lake beds present a relatively flat spectrum, and bright man-made materials such as concrete may serve in urban scenes. Any significant spectral absorption features in the flat field spectrum will give rise to spurious features in the calculated relative reflectance spectra. If there is significant elevation variation within the scene, the converted spectra will also incorporate residual effects of topographic shading and atmospheric path differences.

Average Relative Reflectance Conversion This method also normalizes image spectra by dividing by a mean spectrum, but derives the mean spectrum from the entire image. Before computing the mean spectrum, the radiance values in each image spectrum are scaled so that their sum is constant over the entire image. This adjustment largely removes topographic shading and other overall brightness variations. The method assumes that the scene is heterogeneous enough that spatial variations in spectral reflectance characteristics will cancel out, producing a mean spectrum similar to the flat field spectrum described above. This assumption is not true of all scenes, and when it is not true the method will produce relative reflectance spectra that contain spurious spectral features.

Reflectance Conversion II

The image-based conversion methods discussed on the previous page only account for multiplicative contributions to the image spectra. Most studies that have used these methods have focused on mapping minerals using shortwave infrared spectra (2.0 to 2.5 µm) for which the additive effect of atmospheric path radiance is minimal. If the spectra to be analyzed include the visible and near infrared ranges, however, path radiance effects should not be neglected. If the scene includes dark materials or deep topographic shadows, an approximate correction can be made by determining (for each band) the minimum brightness value (or the average value of a shadowed area) and subtracting it from each pixel in the band.

Empirical Line Method Field researchers using hyperspectral imagery typically use field reflectance measurements from the image area to convert the image data to reflectance. Field reflectance spectra must be acquired from two or more uniform ground target areas. Target areas should have widely different brightness and be large enough to recognize in the image. Using the image radiance and ground reflectance values for the target areas, a linear equation relating radiance to reflectance can be derived for each image band. In a plot of radiance versus reflectance, the slope of the calculated line quantifies the combined effects of the multiplicative radiance factors (gain), while the intercept with the radiance axis represents the additive component (offset). These values are then used to convert each image band to apparent reflectance. The final values should be considered "apparent" reflectance because the conversion does not account for possible effects of topography within the scene (shading and atmospheric path length differences).

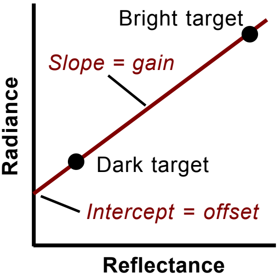

Reflectance conversion parameters for a single image band using known target reflectance values.

Modeling Methods Radiative-transfer computer models start with a simulated solar irradiance spectrum, then compute the scene radiance effects of solar elevation (derived from the day and time of the scene) and atmospheric scattering and absorption. In the absence of measurements of actual atmospheric conditions, the user must estimate some input parameters, such as amount and distribution of scattering agents. Absorption by well-mixed gases (CO_2 and O_2) is assumed to be uniform across a scene but absorption due to water vapor is often variable. Water vapor absorption effects can be estimated and corrected individually for each image pixel using portions of the spectra that include water absorption bands. The final apparent reflectance values may still incorporate the effects of topographic shading, however.

Strategies for Image Analysis

The table below lists some of the imaging spectrometers currently being operated for research or commercial purposes. The hyperspectral images produced by these sensors present a challenge for the analyst. They provide the fine spectral resolution needed to characterize the spectral properties of surface materials but the volume of data in a single scene can seem overwhelming. The difference in spectral information between two adjacent wavelength bands is typically very small and their grayscale images therefore appear nearly identical. Much of the data in a scene therefore would seem to be redundant, but embedded in it is critical information that often can be used to identify the ground surface materials. Finding appropriate tools and approaches for visualizing and analyzing the essential information in a hyperspectral scene remains an area of active research.

Most approaches to analyzing hyperspectral images concentrate on the spectral information in individual image cells, rather than spatial variations within individual bands or groups of bands. The statistical classification (clustering) methods often used with multispectral images can also be applied to hyperspectral images but may need to be adapted to handle their high dimensionality (Landgrebe, in press). More sophisticated methods combine both spectral and spatial analysis. The following pages detail some of the popular methods of analyzing the spectral content of hyperspectral images.

A Sample of Research and Commercial Imaging Spectrometers

Sensor	Organization	Country	Number of Bands	Wavelength Range (µm)
AVIRIS	NASA	United States	224	0.4 - 2.5
AISA	Spectral Imaging Ltd.	Finland	286	0.45 - 0.9
CASI	Itres Research	Canada	288	0.43 - 0.87
DAIS 2115	GER Corp.	United States	211	0.4 - 12.0
HYMAP	Integrated Spectronics Pty Ltd	Australia	128	0.4 - 2.45
PROBE-1	Earth Search Sciences Inc.	United States	128	0.4 - 2.45

Match Each Image Spectrum

One approach to analyzing a hyperspectral image is to attempt to match each image spectrum individually to one of the reference reflectance spectra in a spectral library. This approach requires an accurate conversion of image spectra to reflectance. It works best if the scene includes extensive areas of essentially pure materials that have corresponding reflectance spectra in the reference library. An observed spectrum will typically show varying degrees of match to a number of similar reference spectra. The matching reference spectra must be ranked using some measure of goodness of fit, with the best match designated the "winner."

Spectral matching is complicated by the fact that most hyperspectral scenes include many image pixels that represent spatial mixtures of different materials (see page 10). The resulting composite image spectra may match a variety of "pure" reference spectra to varying degrees, perhaps including some spectra of materials that are not actually present. If the best-matching reference spectrum has a sufficient fit to the image spectrum, then this material is probably the dominant one in the mixture and the pixel is assigned to this material. If no reference spectrum achieves a sufficient match, then no endmember dominates, and the pixel should be left unassigned. The result is a "material map" of the image that portrays the dominant material for most of the image cells, such as the example shown below. Sample mixed spectra can be included in the library to improve the mapping, but it is usually not possible to include all possible mixtures (and all mixture proportions) in the reference library.

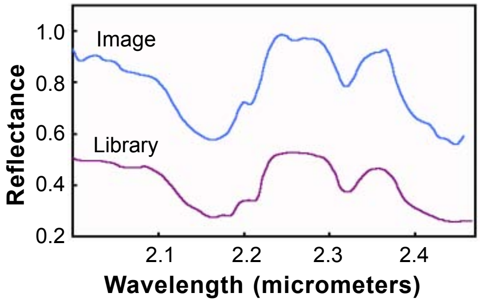

Sample image spectrum and a matched spectrum of the mineral alunite from the USGS Spectral Library (goodness of fit = 0.91).

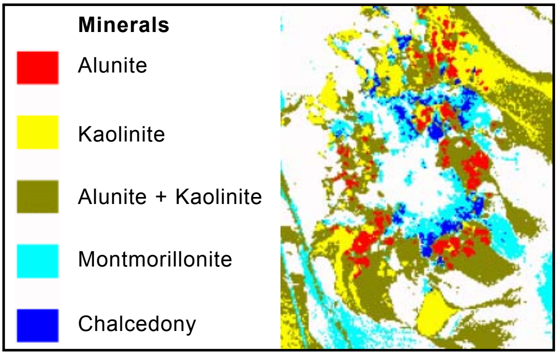

Mineral map for part of the Cuprite AVIRIS scene, created by matching image spectra to mineral spectra in the USGS Spectral Library. White areas did not produce a sufficient match to any of the selected reflectance spectra, and so are left unassigned.

Spectral Matching Methods

The shape of a reflectance spectrum can usually be broken down into two components: broad, smoothly changing regions that define the general shape of the spectrum and narrow, trough-like absorption features. This distinction leads to two different approaches to matching image spectra with reference spectra.

Many pure materials, such as minerals, can be recognized by the position, strength (depth), and shape of their absorption features. One common matching strategy attempts to match only the absorption features in each candidate reference spectrum and ignores other parts of the spectrum. A unique set of wavelength regions is therefore examined for each reference candidate, determined by the locations of its absorption features. The local position and slope of the spectrum can affect the strength and shape of an absorption feature, so these parameters are usually determined relative to the *continuum*: the upper limit of the spectrum's general shape. The continuum is computed for each wavelength subset and removed by dividing the reflectance at each spectral channel by its corresponding continuum value. Absorption features can then be matched using a set of derived values (including depth and the width at half-depth), or by using the complete shape of the feature. These types of procedures have been organized into an expert system by researchers at the U.S. Geological Survey Spectroscopy Lab (Clark and others, 1990).

Many other materials, such as rocks and soils, may lack distinctive absorption features. These spectra must be characterized by their overall shape. Matching procedures utilize full spectra (omitting noisy image bands severely affected by atmospheric absorption) or a uniform wavelength subset for all candidate materials. One approach to matching seeks the spectrum with the minimum difference in reflectance (band per band) from the image spectrum (quantified by the square root of the sum of the squared errors). Another approach treats each spectrum as a vector in spectral space and finds the reference spectrum making the smallest angle with the observed image spectrum.

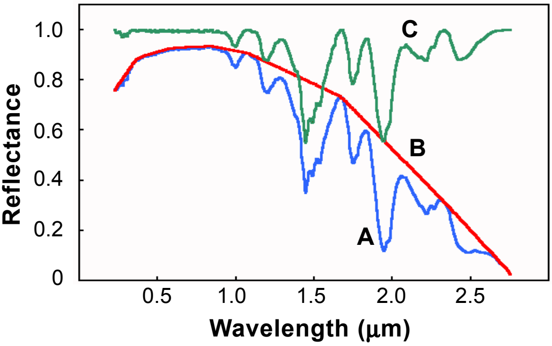

Reflectance spectrum for the mineral gypsum (A) with several absorption features. Curve B shows the continuum for the spectrum, and C the spectrum after removal of the continuum.

Linear Unmixing

Linear unmixing is an alternative approach to simple spectral matching. Its underlying premise is that a scene includes a relatively small number of common materials with more or less constant spectral properties. Furthermore, much of the spectral variability in a scene can be attributed to spatial mixing, in varying proportions, of these common endmember components. If we can identify the endmember spectra, we can mathematically "unmix" each pixel's spectrum to identify the relative abundance of each endmember material.

The unmixing procedure models each image spectrum as the sum of the fractional abundances of the endmember spectra, with the further constraint that the fractions should sum to 1.0. The best-fitting set of fractions is found using the same spectral-matching procedure described on the previous page. A fraction image for each endmember distills the abundance information into a form that is readily interpreted and manipulated. An image showing the residual error for each pixel helps identify parts of the scene that are not adequately modeled by the selected set of endmembers.

The challenge in linear unmixing is to identify a set of spectral endmembers that correspond to actual physical components on the surface. Endmembers can be defined directly from the image using field information or an empirical selection technique such as the one outlined on the next page can be used. Alternatively, endmember reflectance spectra can be selected from a reference library, but this approach requires that the image has been accurately converted to reflectance. Variations in lighting can be included directly in the mixing model by defining a "shade" endmember that can mix with the actual material spectra. A shade spectrum can be obtained directly from a deeply shadowed portion of the image. In the absence of deep shadows, the spectrum of a dark asphalt surface or a deep water body can approximate the shade spectrum, as in the example to the right.

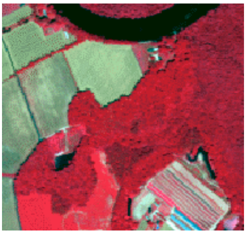

Portion of an AVIRIS scene with forest, bare and vegetated fields, and a river, shown with a color-infrared band combination (vegetation is red). Fraction images from linear unmixing are shown below.

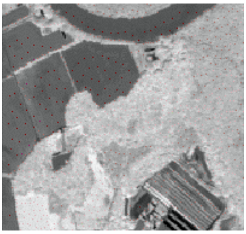

Vegetation fraction

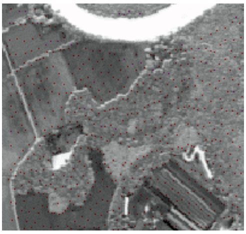

Water / shade fraction

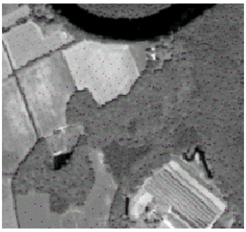

Soil fraction

Defining Image Endmembers

When spectral endmembers are defined from a hyperspectral image, each image endmember should have the maximum abundance of the physical material that it represents. (Ideally, each endmember would be a single pure material, but “pure” pixels of each endmember may not be present in the image). If image spectra are represented as points in an n-dimensional scatterplot, the endmembers should correspond to cusps at the edge of the cloud of spectral points.

One common procedure for isolating candidate image endmembers involves several steps. Because of the high degree of correlation between adjacent spectral bands, the dimensionality of the dataset first can be reduced by applying the Minimum Noise Fraction (MNF) transform and retaining the small number of noise-free components. The MNF transform (Green et al. 1988) is a noise-adjusted principal components transform that estimates and equalizes the amount of noise in each image band to ensure that output components are ordered by their amount of image content. Second, an automated procedure is applied to the MNF components to find the extreme spectra around the margins of the n-dimensional data cloud. One such procedure is the Pixel Purity Index (PPI). It examines a series of randomly-oriented directions radiating outward from the origin of the coordinate space. For each test direction, all spectral points are projected onto the test vector, and the extreme spectra (low and high) are noted. As directions are tested, the process tallies the number of times each image cell is found to be extreme. Pixels with high values in the resulting PPI raster should correspond primarily to the edges of the MNF data cloud. In the third step, the PPI raster is used to select pixels from the MNF dataset for viewing in a rotating n-dimensional scatterplot (using a tool such as the n-Dimensional Visualizer in the TNTmips Hyperspectral Analysis process). By viewing the PPI spectral cloud from various directions, the analyst can identify significant directions, mark the spectral points that are extreme in those directions, and save an image of the marked cells. Finally, the marked cell image is overlaid on the original hyperspectral image and used as a guide to select and examine the image spectra. The best candidate endmember spectra are then saved in a spectral library for use in unmixing the hyperspectral image.

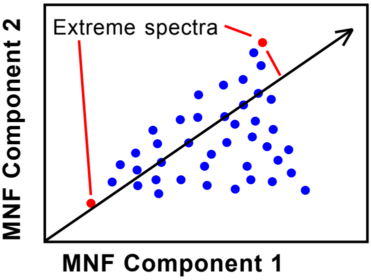

Simple two-component plot showing one of the random vector directions (arrow) tested by the Pixel Purity Index procedure. All spectral points are projected to each test vector, and extreme points are noted.

Partial Unmixing

Some hyperspectral image applications do not require finding the fractional abundance of all endmember components in the scene. Instead the objective may be to detect the presence and abundance of a single target material. In this case a complete spectral unmixing is unnecessary. Each pixel can be treated as a potential mixture of the target spectral signature and a composite signature representing all other materials in the scene. Finding the abundance of the target component is then essentially a partial unmixing problem.

Methods for detecting a target spectrum against a background of unknown spectra are often referred to as *matched filters*, a term borrowed from radio signal processing. Various matched filtering algorithms have been developed, including orthogonal subspace projection and constrained energy minimization (Farrand and Harsanyi, 1994). All of these approaches perform a mathematical transformation of the image spectra to accentuate the contribution of the target spectrum while minimizing the background. In a geometric sense, matched filter methods find a projection of the n-dimensional spectral space that shows the full range of abundance of the target spectrum but “hides” the variability of the background. In most instances the spectra that contribute to the background are unknown, so most matched filters use statistical methods to estimate the composite background signature from the image itself. Some methods only work well when the target material is rare and does not contribute significantly to the background signature. A modified version of matched filtering uses derivatives of the spectra rather than the spectra themselves, which improves the matching of spectra with differing overall brightness.

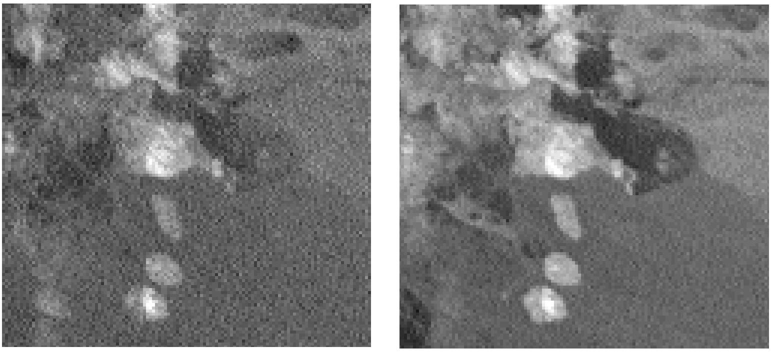

Fraction images produced by Matched Filtering (left) and Derivative Matched Filtering (right) for a portion of the Cuprite AVIRIS scene. The target image spectrum represents the mineral alunite. Brighter tones indicate pixels with higher alunite fractions. The image produced by Derivative Matched Filtering shows less image noise, sharper boundaries, and better contrast between areas with differing alunite fractions.

References

General

Kruse, F.A. (1999). Visible-Infrared Sensors and Case Studies. In Renz, Andrew N. (ed), *Remote Sensing for the Earth Sciences: Manual of Remote Sensing* (3rd ed.), Vol 3. New York: John Wiley & Sons, pp. 567-611.

Landgrebe, David (1999). Information Extraction Principles and Methods for Multispectral and Hyperspectral Image Data. In Chen, C.H. (ed.), *Information Processing for Remote Sensing*. River Edge, NJ: World Scientific Publishing Company, pp. 3-38.

Vane, Gregg, Duval, J.E., and Wellman, J.B. (1993). Imaging Spectroscopy of the Earth and Other Solar System Bodies. In Pieters, Carle M. and Englert, Peter A.J. (eds.), *Remote Geochemical Analysis: Elementatl and Mineralogic Composition*. Cambridge, UK: Cambridge University Press, pp. 121-143.

Vane, Gregg, and Goetz, A.F.H. (1988). Terrestrial Imaging Spectroscopy. *Remote Sensing of Environment*, 24, pp. 1-29.

Spectral Reflectance Signatures

Ben-Dor, E., Irons, J.R., and Epema, G.F. (1999). Soil Reflectance. In Renz, Andrew N. (ed), *Remote Sensing for the Earth Sciences: Manual of Remote Sensing* (3rd ed.), Vol 3. New York: John Wiley & Sons, pp. 111-188.

Clark, Roger N. (1999). Spectroscopy of Rocks and Minerals, and Principles of Spectroscopy. In Renz, Andrew N. (ed), *Remote Sensing for the Earth Sciences: Manual of Remote Sensing* (3rd ed.), Vol 3. New York: John Wiley & Sons, pp. 3-58.

Ustin, S.L., Smith, M.O., Jacquemoud, S., Verstraete, M., and Govaerts, Y. (1999). Geobotany: Vegetation Mapping for Earth Sciences. In Renz, Andrew N. (ed), *Remote Sensing for the Earth Sciences: Manual of Remote Sensing* (3rd ed.), Vol 3. New York: John Wiley & Sons, pp. 189-248.

Reflectance Conversion

Farrand, William H., Singer, R.B., and Merenyi, E., 1994, Retrieval of Apparent Surface Reflectance from AVIRIS Data: A Comparison of Empirical Line, Radiative Transfer, and Spectral Mixture Methods. *Remote Sensing of Environment*, 47, 311-321.

References

Goetz, Alexander F.H., and Boardman, J.W. (1997). Atmospheric Corrections: On Deriving Surface Reflectance from Hyperspectral Imagers. In Descour, Michael R. and Shen, S.S. (eds.), Imaging Spectrometry III: Proceedings of SPIE, 3118, 14-22.

van der Meer, Freek (1994). Calibration of Airborne Visible/Infrared Imaging Spectrometer Data (AVIRIS) to Reflectance and Mineral Mapping in Hydrothermal Alteration Zones: An Example from the "Cuprite Mining District". *Geocarto International*, 3, 23-37.

Hyperspectral Image Analysis

Adams, John B., Smith, M.O., and Gillespie, A.R. (1993). Imaging Spectroscopy: Interpretation Based on Spectral Mixture Analysis. In Pieters, Carle M. and Englert, Peter A.J. (eds.), *Remote Geochemical Analysis: Elementatl and Mineralogic Composition*. Cambridge, UK: Cambridge University Press, pp. 145-166.

Clark, R.N., Gallagher, A.J., and Swayze, G.A. (1990). Material absorption band depth mapping of imaging spectrometer data using a complete band shape least-squares fit with library reference spectra. *Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop*, JPL Publication 90-54, pp. 176-186.

Cloutis, E.A., (1996). Hyperspectral Geological Remote Sensing: Evaluation of Analytical Techniques. *International Journal of Remote Sensing*, 17, 2215-2242.

Farrand, William H., and Harsanyi, J.C. (1994). Mapping Distributed Geological and Botanical Targets through Constrained Energy Minimization. Proceedings of the Tenth Thematic Conference on Geological Remote Sensing, San Antonio, Texas, 9-12 May 1994, pp. I-419 - I-429.

Green, Andrew A., Berman, M., Switzer, P., and Craig, M.D. (1988). A Transformation for Ordering Multispectral Data in Terms of Image Quality with Implications for Noise Removal. *IEEE Transactions on Geoscience and Remote Sensing*, 26, 65-74.

Mustard, John F., and Sunshine, J.M. (1999). Spectral Analysis for Earth Science: Investigations Using Remote Sensing Data. In Renz, Andrew N. (ed), *Remote Sensing for the Earth Sciences: Manual of Remote Sensing* (3rd ed.), Vol 3. New York: John Wiley & Sons, pp. 251-306.

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

absorption
  bands 5-7
  atmospheric 13,18
atmosphere
  absorption by 13,18
  scattering by 13
continuum 18
illumination 11,12
imaging spectrometer 4,10,16
irradiance, solar 12
linear unmixing 19-21
matched filtering 21
matching, spectral 17,18
minimum noise fraction transform 20
pixel purity index 20
resolution, spatial 10
scattering 4,5,13
sensor effects 13
shadowing 12
spectral libraries 8
spectral radiance 11
spectral reflectance 5-11
  converting image to 14-15
  curve *See* spectrum
  defined 5
spectral space 9
spectrometer 4
spectroscopy 4,5
spectrum (spectra)
  endmember 19,20
  image 3,17-20
  in library 8
  mineral 6
  mixed 10
  plant 7
  plotting 9
  reflectance 5-11
  soil 5
  solar 12
  water 5

MicroImages, Inc.
www.microimages.com