Gary Shaw and Dimitris Manolakis, *Massachusettes Institute of Technology, Lincoln Laboratory*

Signal Processing for Hyperspectral Image Exploitation

Electro-optical remote sensing involves the acquisition of information about an object or scene without coming into physical contact with it. This is achieved by exploiting the fact that the materials comprising the various objects in a scene reflect, absorb, and emit electromagnetic radiation in ways characteristic of their molecular composition and shape. If the radiation arriving at the sensor is measured at each wavelength over a sufficiently broad spectral band, the resulting spectral signature, or simply spectrum, can be used (in principle) to uniquely characterize and identify any given material.

The field of spectroscopy is concerned with the measurement, analysis, and interpretation of such spectra. Combining spectroscopy with methods to acquire spectral in-

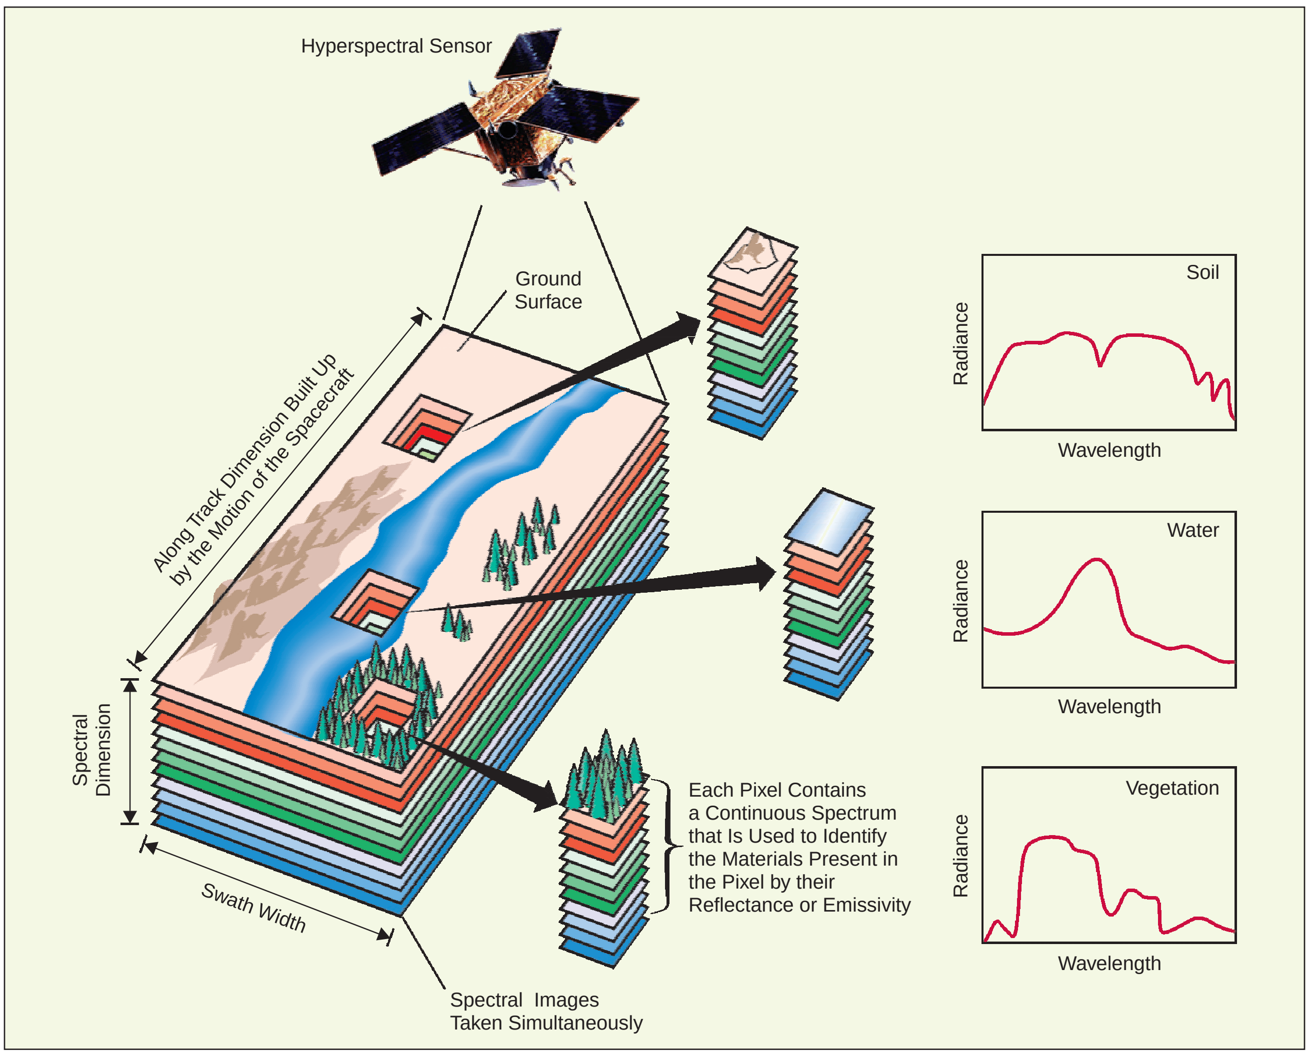

▲ 1. Principle of imaging spectroscopy.

1053-5888/02/$17.00©2002IEEE

formation over large areas is known as imaging spectroscopy. The principles involved in the application of imaging spectroscopy to perform satellite remote sensing are illustrated in Fig. 1. Hyperspectral sensors are a class of imaging spectroscopy sensors in which the waveband of interest is divided into hundreds of contiguous narrow bands for the purpose of signature analysis.

There are four sampling operations involved in the collection of hyperspectral data: spatial, spectral, radiometric, and temporal. The spatial sampling resolution, or ground pixel size, varies from meters to tens of meters and is a function of the sensor aperture and platform altitude that, in turn, depend upon the kind of platform (e.g., spaceborne versus airborne). The spectral sampling can be accomplished by a variety of means such as a prism or interferometer. An A/D converter samples the radiance measured in each spectral channel producing the digital data, at the prescribed radiometric resolution, that comprises a hyperspectral cube. Temporal sampling corresponds to the collection of multiple hyperspectral images of the same scene over time and is an important mechanism for studying environmental change.

Hyperspectral sensors have been developed to sample the reflective portion of the electromagnetic spectrum that extends from the visible region (0.4-0.7 μm) through the near-infrared (about 2.4 μm) in hundreds of narrow contiguous bands about 10 nm wide. Other types of hyperspectral sensors exploit the emissive properties of objects by collecting data in the mid-wave and long-wave infrared (MWIR and LWIR) regions of the spectrum. Hyperspectral sensors represent an evolution in technology from earlier multispectral sensors, which typically collect spectral information in only a few discrete, noncontiguous bands. As indicated in Fig. 2, multispectral imaging (MSI) from space dates back to the launch of Landsat-1 in 1972, whereas the first experimental hyperspectral sensor in space was launched late in 2000 on the NASA EO-1 satellite. The primary sensor on EO-1 is the advanced landsat imager (ALI), a lightweight multispectral imager. However, EO-1 also carries Hyperion, a visible/NIR hyper spectral sensor with 30-m spatial resolution.

The high spectral resolution characteristic of hyperspectral sensors preserves important aspects of the spectrum (e.g., shape of narrow absorption bands) and makes differentiation of different materials on the ground possible. The spatially and spectrally sampled information can be described as a data cube, whose face is a function of the spatial coordinates and depth is a function of spectral band (or wavelength). The data in each band corresponds to a narrowband image of the surface covered by the field of view of the sensor, whereas along the wavelength dimension, each image pixel provides a spectrum characterizing the materials within the pixel. The nature and organization of the collected data is illustrated in Fig. 3.

The spectrum that is observed at each pixel in a scene is determined by a collection of processes as illustrated in Fig. 4. The objective is to identify and segregate materials based on their unique reflective (and for IR wavelengths emissive) properties when ob-

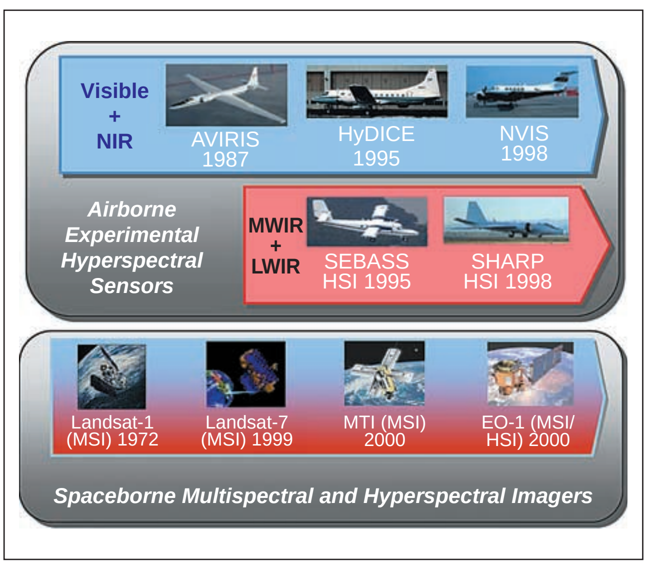

▲ *2. An abbreviated chronology of hyperspectral and multispectral sensor evolution.*

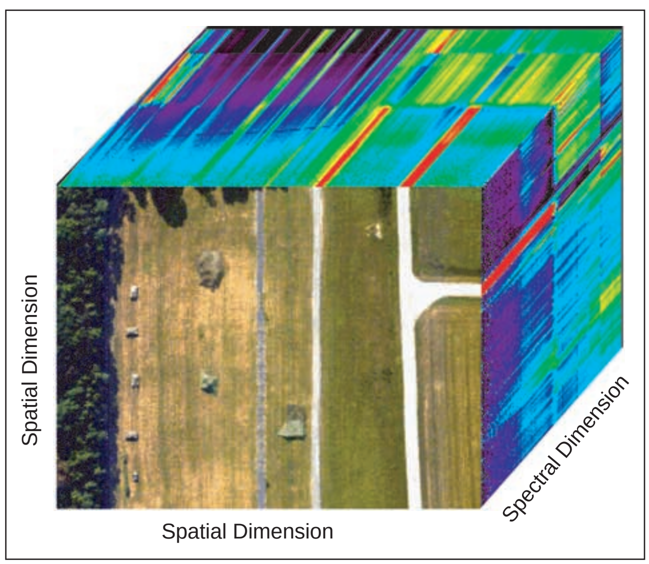

▲ *3. Imaging spectrometry data cube illustrating the 3-D spatial and spectral character of the data.*

served over a wide range of wavelengths. The quantity that can be observed and digitized by a sensor is the radiant flux, or radiance, entering the aperture of the sensor. For a given ground pixel, the radiance observed at any particular wavelength is determined, to first order, by the solar illumination at that wavelength and the reflectivity (reflectance) of the material at that wavelength, which establishes how much of the solar illumination is reflected into the sensor.

There are many important additional effects, however, that may need to be accounted for to uniquely identify materials in the scene. As illustrated in Fig. 4, these effects include

- the angle of the sun;
- the viewing angle of the sensor;
- the upwelling solar radiance from atmospheric scattering;
- the secondary illumination of the material by light reflected from adjacent objects in the scene;
- shadowing;
- the scattering and absorption of the reflected radiance by the atmosphere;
- spatial and spectral aberrations in the sensor.

Characterizing and compensating for these environmental and atmospheric modulations of the radiance spectrum is a key preprocessing step in the exploitation of hyperspectral imagery. In the case of cooperative data collections, the simplest method to compensate for the environmental and atmospheric effects is to place a calibration panel, with known reflectance in the scene, in an open area, and use the observed radiance spectrum from the panel to develop gain and offset corrections for each waveband of interest. Other more sophisticated methods have been developed for atmospheric compensation and sensor calibration. However, to maintain the focus of this special issue, atmospheric compensation and sensor calibration methods are relegated to a preprocessing step with no further elaboration.

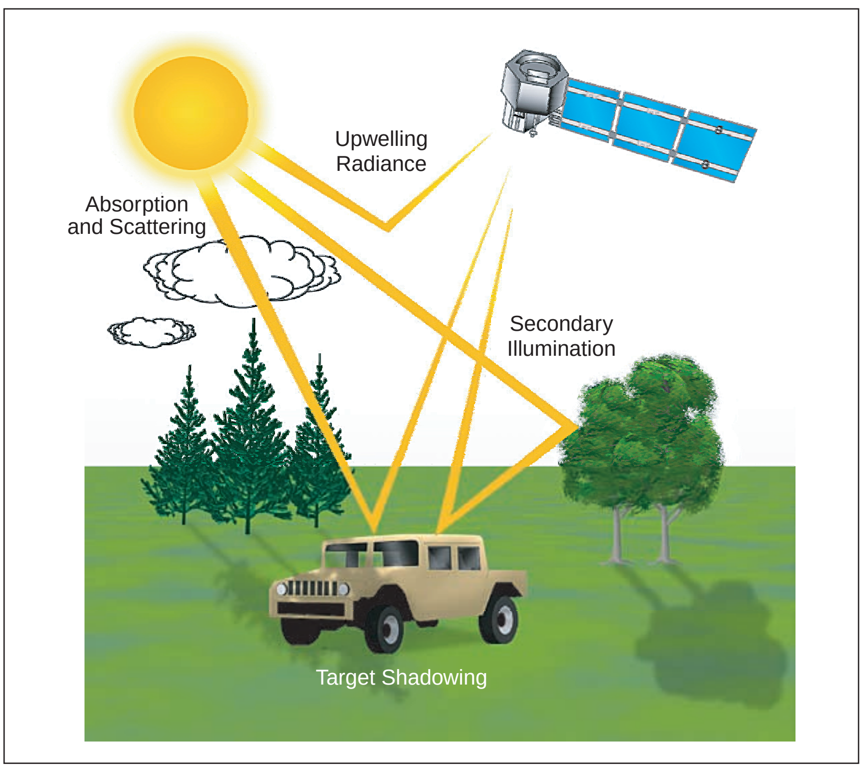

4. *Solar illumination and atmospheric path absorption and scattering modulate the direct path radiance signal observed at the sensor.*

Processing Algorithm Taxonomy

The number and variety of potential civilian and military applications for hyperspectral remote sensing is enormous. However, the majority of algorithms used in these applications can be organized according to the following primitive-application-specific tasks (see Fig. 5).

- a) searching the pixels of a hyperspectral data cube for "rare" (either known or unknown) spectral signatures [target detection (for the purpose of this article a target is defined as any object or material being sought in a hypersepctral data cube)];
- b) finding the "significant" (i.e., important to the user) changes between two hyperspectral scenes of the same geographic region (change detection);
- c) assigning a label (class) to each pixel of a hyperspectral data cube (classification);
- d) estimating the fraction of the pixel area covered by each material present in the scene.

Note that from a signal processing perspective c) is a classification problem whereas d) is an estimation problem.

Dimensionality Reduction

In most cases, hyperspectral sensors oversample the spectral signal to ensure that that any narrow features are adequately represented. In some cases, sensors may oversample the spatial signal as well. An important function of hyperspectral signal processing is to eliminate the redundancy in the spectral and spatial sample data while preserving the high-quality features needed for detection, discrimination, and classification. This dimensionality reduction is implemented in a scene-dependent (adaptive) manner and may be implemented as a distinct step in the processing or as an integral part of the overall algorithm. Regardless of how it is implemented, the dimensionality reduction algorithm must be designed to preserve the information of interest to downstream detection, classification, or spectral unmixing algorithms. Since dimensionality reduction with small loss of information is always possible, the reduction can be achieved without significantly degrading detection performance or decreasing the separability among the different classes (classification performance).

Dimensionality reduction leads to significant reductions in computational complexity and also reduces the number of pixels required to obtain statistical estimates of a given accuracy. (The number of samples (pixels) required to obtain a statistical estimate with a given accuracy increases drastically with the dimensionality of the data.) The most widely used algorithm for dimensionality reduction is principal component analysis (PCA) or, equivalently, Karhünen-Loéve transformation. Fig. 6 is a simplified block diagram of the processing chain showing the common elements of atmospheric compensation and dimensionality reduction. Subsequent processing is specialized depending upon the intended application. Fig. 6 illustrates two of many possible applications, unmixing and detection, each of which is discussed briefly in the following sections and developed more fully in this issue's articles.

Classification versus Target Detection

Formally, classification is the process of assigning a label to an observation (usually a vector of numerical values), whereas detection is the process of identifying the existence or occurrence of a condition. In this sense, detection can be considered as a two-class classification problem: target exists or target does not exist. Traditional classifiers assign one and only one label to each pixel (hard classification) producing what is known as a thematic map. However, the need to deal more effectively with pixels containing a mixture of different materials leads to the concept of soft classification of pixels. A soft classifier can assign to each pixel multiple labels, with each label accompanied by a number that can be interpreted as the likelihood of being correct or, more generally, as the proportion of the material within the pixel.

In terms of data products, the goal of target detection algorithms is to generate target maps at a constant false alarm rate (CFAR). This CFAR property is a highly desirable feature of target detection algorithms. Change-detection algorithms produce a map of significant changes, that, for reliable operation, depend upon the existence of a CFAR change-detection threshold.

At first glance, detection and classification applications may look deceptively similar, if not identical. However, the rarity of the target class, the final product (target detec-

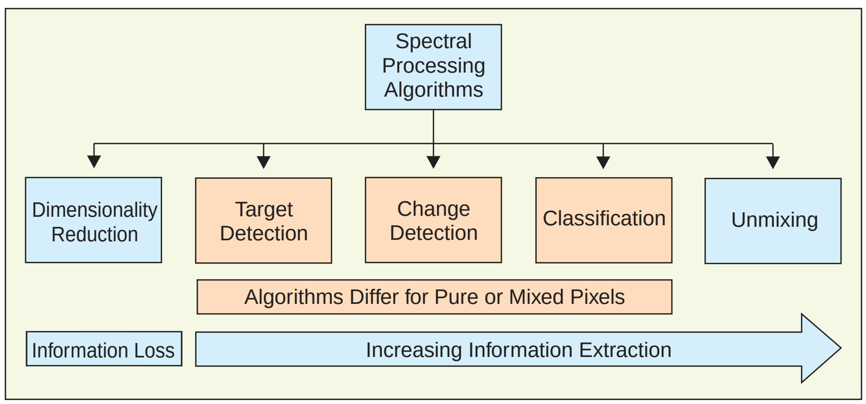

▲ *5. Basic taxonomy of fundamental processing.*

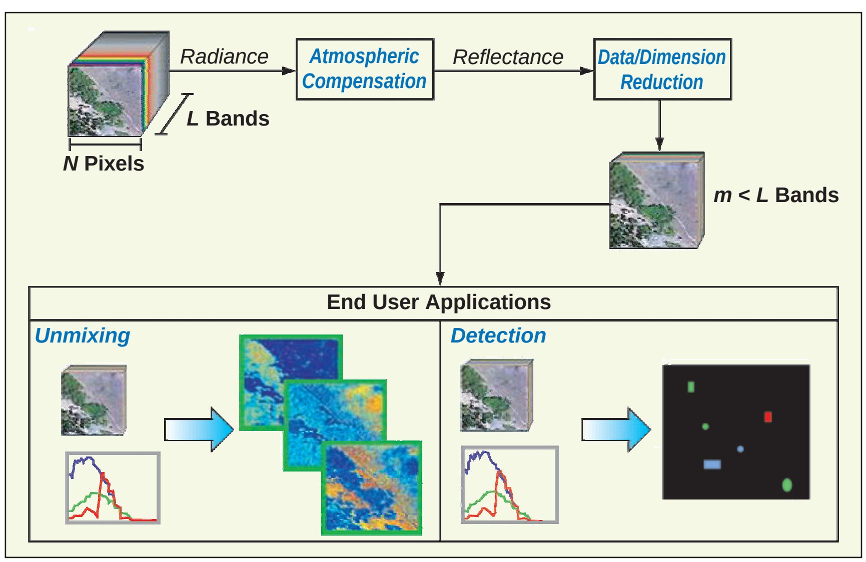

▲ *6. Representative algorithm processing chain for hyperspectral image exploitation.*

tion maps versus thematic maps), and the different cost functions (misclassifying a few pixels in a thematic map is not as critical as missing a target or overloading the tracker with a large number of false alarms), lead to some fundamental theoretical and practical differences between detection and classification applications.

Unmixing

Although the concept of mixed pixels was introduced decades ago, the development of algorithms to extract the constituent spectra comprising a pixel, termed unmixing, has been aggressively pursued only during the last decade. A fundamental question in unmixing is whether the mixture of spectral signatures is formed by linear or nonlinear processes. The answer depends upon a number of factors and conditions in the scene. Nevertheless, the linear mixing model assumption provides data products that usually can be related to the abundances of materials present in the ground pixel footprint and has proven useful in many applications. In contrast to detection and classification, unmixing is an estimation problem. Hence, it is a more involved process and extracts more information from the data.

Issue Overview

The first article, contributed by Prof. David Landgrebe, provides an overview of classification concepts, with an emphasis on some of the challenges that arise due to the high-dimensional nature of hyperspectral data. Prof. Landgrebe has devoted many years to research and algorithm development for multispectral and hyperspectral imagery. His article includes information on accessing a public-domain tool, Multispec, that readers may download and use to gain hands-on experience with visualization and processing of multispectral and hyperspectral data.

The second article, "Detection Algorithms for Hyperspectral Imaging Applications," by the guest editors, examines target detection, in particular adaptive subspace detection algorithms. This article traces many of the algorithm concepts and derivations to earlier work in radar and multispectral systems. Theoretical performance of the algorithms is presented and compared to observed performance on actual data. Departures from theoretical performance are identified and related to departures of the statistics of the data from the usual assumptions of normality.

The third article, "Spectral Unmixing," by Dr. Nirmal Keshava and Prof. John Mustard, introduces the intriguing topic of decomposing or unmixing the composite spectra observed within a pixel. This article highlights some of the open issues and research challenges associated with developing improved unmixing methods and a better understanding the underlying phenomenology. An example of the power of these methods is given using multispectral data collected by the Clementine spacecraft during a lunar flyby.

The final article, by Dr. David Stein and colleagues, also addresses target detection, but for the case where information regarding the target spectrum is either unavailable or unreliable (for example due to deleterious atmospheric effects). This approach to target detection, termed anomaly detection, can also be viewed as an example of change detection in the spatial domain, that is, detection of pixels that appear different from the (local) background pixels in the image. Detection of targets with unknown spectral signature is especially difficult, and this article introduces a novel technique for combining the detection statistics of several different detection algorithms to enhance overall detection performance.

As guest editors, our objective in organizing this special issue is to introduce this fascinating remote sensing application to the broad readership of *IEEE Signal Processing Magazine*. We believe many opportunities remain to develop innovative signal processing algorithms and implementations for hyperspectral image exploitation. We hope you will find these articles interesting, and perhaps be stimulated to contribute to the development of this emerging technology.

Acknowledgment

This work was sponsored by the Department of the Defense Air Force contract F19628-00-C-0002.