


An Automated Target Detection System for Hyperspectral Imaging Sensors

Marc A. Kolodner



Over the past several years, hyperspectral sensor technology has evolved to the point where real-time processing for operational applications is achievable. Algorithms supporting such sensors must be fully automated and robust. Our approach for target detection applications is to select signatures from a database of target reflectance libraries and project them to the at-sensor and collection-specific radiance domain using the weather forecast or radiosonde data. This enables platform-based detection immediately following data acquisition without the need for further atmospheric compensation. One advantage of this method for reflective hyperspectral sensors is the ability to predict the radiance signatures of targets under multiple illumination conditions. A three-phase approach is implemented, where the library generation and data acquisition phases provide the necessary input for the automated detection phase. In addition to employing the target detector itself, this final phase includes a series of automated filters, adaptive thresholding, and confidence assignments to extract the optimal information from the detection scores for each spectral class. Our prototype software is applied to 50 reflective hyperspectral datacubes to measure detection performance over a range of targets, backgrounds, and environmental conditions.

INTRODUCTION

Timely detection of targets continues to be a top priority for a hyperspectral remote sensing capability. A hyperspectral imaging (HSI) sensor measures the radiance emanating from each pixel in a scene at multiple

wavelengths, creating a datacube of the scene. A “target” for this discussion is any material surface within an imaged scene to be detected. The targets, if present, are generally rare in the image and fall into two

categories, resolved targets (RTs) and subpixel targets (SPTs). Although many novel detection algorithms¹⁻³ have been developed for RTs and SPTs, several current efforts⁴⁻⁶ are focusing on the automation of HSI detection processing for operational use. Examples include search and rescue, emergency response (e.g., recovery of shuttle parts following the Columbia tragedy), and first responder support to natural disasters (e.g., detection of toxic spills following Hurricane Katrina). Sensor platforms that traditionally flew with only an HSI sensor and onboard data storage now include a real-time processor that can produce quick turnaround products from the data. Controllers operating from a ground station task these platforms to collect HSI data over a particular area of interest and provide detection results for specific targets of interest. This article describes a prototype software system to support this emerging capability, including performance of the system for both RT and SPT detection.

A flowchart of our Target Detection System is presented in Fig. 1. The three phases of this system to be discussed here are library generation, data acquisition, and automated detection. The prototype software is controlled by a graphical user interface that simulates the entire process, i.e., planning an operation, selecting targets of interest, tasking the sensor to collect data, calibrating the collected data if necessary, performing the detection processing, and downloading and visualizing the results.

LIBRARY GENERATION PHASE

The objective of the library generation phase is to produce a suite of at-sensor and collection-specific radiance libraries in the planning stage of an operation. These libraries will subsequently be matched against the collected data. Each planned collection will have its own unique set of libraries. The software for this phase can be installed in the platform ground control station to support sensor tasking. This section describes the inputs and method.

Collection and Sensor Specifications

The data collection specifications include the sensor platform location (latitude and longitude), altitude, and collection time. The sensor specifications consist of the number of spectral bands, the wavelengths, and the FWHMs (full widths at half maximum) to provide a custom sensor model, as well as the sensor-to-target area viewing geometry (zenith angle and azimuth angle). Radiance spectral libraries have been produced for several nadir-viewing visible/near-IR/short-wave-IR (VIS/NIR/SWIR) HSI sensors. These include the Advanced Visible and IR Imaging Spectrometer (AVIRIS),⁷ the Hyperspectral Digital Imagery Collection Experiment (HYDICE),⁸ and the Compact Airborne Spectral Sensor (COMPASS).⁹ While both nadir-viewing and off-nadir-viewing sensors are supported,

our modeling currently does not account for adjacency effects or a surface's full bidirectional reflectance distribution function. These effects could be added if needed for an extreme off-nadir-viewing sensor. In addition, the generation of target signatures for longwave-IR (LWIR) sensors such as the Spatially Enhanced Broadband Array Spectrograph System (SEBASS)¹⁰ is also possible. Although our work has focused principally on the VIS/NIR/SWIR spectral regime for material surface detection, augmentation of our modeling approach for LWIR sensors is briefly discussed.

Signature Specifications

The material surfaces to be detected are represented by their spectral reflectance. Spectral data for several materials are collected under controlled conditions in a laboratory. In many cases,

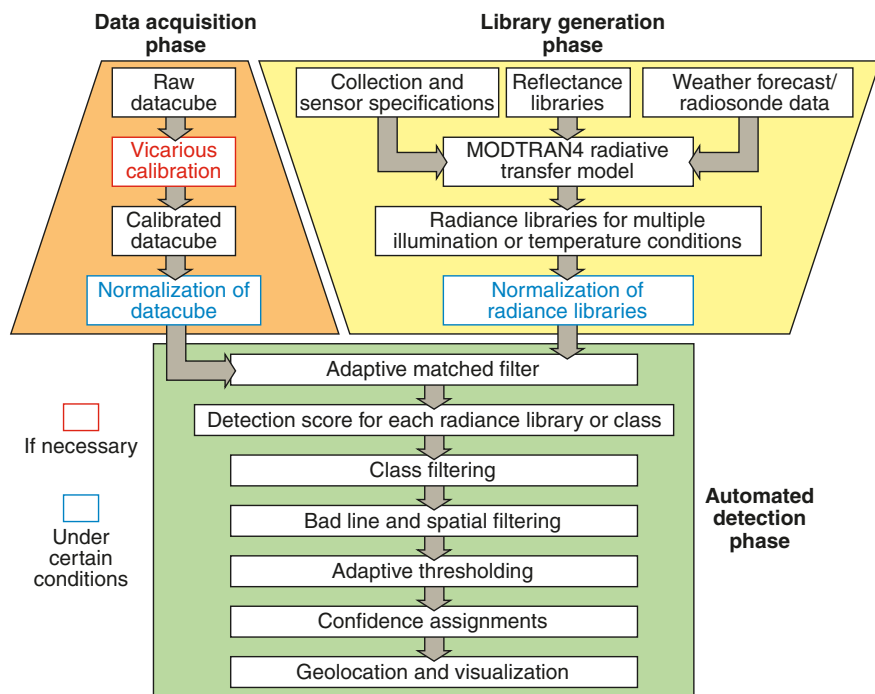


Figure 1. Three-phase HSI Target Detection System.

spectrometers are taken directly into the field to measure surface reflectance. A number of standard reflectance libraries are included with the ENVI (Environment for Visualizing Images) software package.¹¹ They include natural categories such as minerals, rocks, soils, vegetation, water, and snow as well as manmade materials such as galvanized steel, concrete, asphalt, brick, wood, tiles, paints, and metals. Their spectral bands cover 0.4–2.5 μm in the VIS/NIR/SWIR and 6–14 μm in the LWIR.

Environmental Specifications

The intervening atmosphere between the Sun and the target, as well as the target and sensor, is characterized through incorporation of weather forecast data. Over the past 30 years, substantial improvements have been made in both short- and long-term weather forecasting. Weather satellite platforms provide a continuous stream of global weather data for civilian applications. Advances in climatological models and the assimilation of model and satellite data (called analyses) have also enabled more accurate forecasts. There are several weather operations centers throughout the world, such as the Air Force Weather Agency (AFWA), that produce weather forecasts in real time for users. The AFWA uses the fifth-generation Mesoscale Model (MM5)¹² to forecast the state of the atmosphere at a given location for up to 72 h after the time of analysis (or nowcast) in 3-h increments. An updated nowcast is available every 6 h. By contrast, radiosondes are typically launched only twice a day. The AFWA provides MM5 gridded weather data over 15 theater areas around the world, with spatial resolution ranging from 45 to 15 km.

Our prototype software, through a data pipe established with the AFWA, extracts the following surface and atmospheric parameters for each MM5 grid point over an area of interest: surface elevation, surface visibility, surface wind speed, terrain type, precipitation type (if any), geopotential height profile, temperature profile, and relative humidity profile. The weather profiles are provided in pressure levels from the ground up to 5 kPa (20.6 km above sea level). The weather data for a specific grid point are selected graphically by the user. Alternatively, when MM5 data (e.g., historical data collections) are not available, radiosonde data can be read in from the NOAA Forecast Systems Laboratory's Radiosonde Database.¹³ All of the above required data are provided in the radiosonde weather data files except for visibility, terrain type, and precipitation type, which can be supplied by other sources.

Radiative Transfer Modeling

We use the Air Force Research Laboratory's Radiative Transfer Code MODTRAN4¹⁴ to generate the radiance spectra based on collection, sensor, signature, and environmental specifications. MODTRAN4 also provides

climatological parameters such as ozone and aerosol profiles, which are not present in the MM5 or radiosonde weather data. For the radiosondes, MODTRAN4 can also estimate the surface visibility based on surface wind speed for both the sea and barren terrain types. Clouds are present when there is supersaturation, detected when the relative humidity is $\geq 100\%$. The type of cloud (e.g., cumulus, stratus, cirrus) is determined by the cloud height and thickness. MODTRAN4 then provides the density structure for each cloud type.

Shade Characterization

One of the many advantages of this library generation approach is the ability to predict the radiance signatures of targets under multiple illumination conditions. In the VIS/NIR/SWIR, the total at-sensor radiance L_{total} at wavelength λ of a target with diffuse reflectance $r(\lambda)$ can be modeled as a linear combination of three terms, assuming negligible adjacency effects:

$$L_{\text{total}}(\lambda, r) = L_{\text{direct}}(\lambda, r) + L_{\text{indirect}}(\lambda, r) + L_{\text{scattered}}(\lambda), \quad (1)$$

where L_{direct} is the radiance due to direct solar illumination, L_{indirect} is the radiance due to indirect downwelled scattered solar illumination (i.e., the sky shine), and $L_{\text{scattered}}$ is the radiance due to upwelled, scattered solar illumination that is independent of the target reflectance. MODTRAN4 provides the spectra of both L_{total} and L_{direct} for a given target with reflectance r . The radiance spectra of a target under full Sun $L_{\text{full-Sun}}$ equals L_{total} , and the radiance spectra of a target under full shade $L_{\text{full-shade}}$ equals the sum of L_{indirect} and $L_{\text{scattered}}$. Thus, using Eq. 1,

$$\begin{aligned} L_{\text{full-Sun}} &= L_{\text{total}} \\ L_{\text{full-shade}} &= L_{\text{total}} - L_{\text{direct}}. \end{aligned} \quad (2)$$

The radiance spectra of a target under partial shade L_{partial} is modeled as a linear combination of the full Sun and full shade cases:

$$L_{\text{partial}} = \alpha L_{\text{full-Sun}} + (1 - \alpha) L_{\text{full-shade}}, \quad (3)$$

where α is the percentage of full-Sun illumination on a target, ranging from 0 to 1. For each target reflectance, we provide the at-sensor radiance spectra for five illumination conditions, namely, full Sun ($\alpha = 1$), full shade ($\alpha = 0$), and three partial shade cases ($\alpha = 0.25, 0.50$, and 0.75).

As an example of the differences between $\alpha = 1$ and $\alpha = 0$ cases, Fig. 2 plots the VIS/NIR/SWIR radiance spectra of a particular green-painted surface modeled under full Sun and full shade from a nadir-viewing HSI sensor at a 10-km altitude with a solar elevation of 30° .




Figure 2. The radiance spectra from 400 to 2400 nm of a green-painted surface under full Sun ($\alpha = 1$) and full shade ($\alpha = 0$) from a nadir-viewing HSI sensor at a 10-km altitude with a 30° solar elevation.

Note that the paint spectra exhibit the well-known “red-edge” increase in the NIR due to its chlorophyll content. In addition to the reduction in amplitude of the surface’s radiance under full shade, there is a spectral shift toward the blue in the VIS portion of the spectrum due to the dominance of the Raleigh scattered sky shine. The result of this characterization is an illumination invariant set of target radiance libraries. For LWIR sensors, the reflective terms in Eq. 1 are replaced with the appropriate self-emitting thermal terms. Instead of producing a target’s radiance spectra under multiple illumination conditions, the at-sensor spectra are determined for a range of surface temperatures relative to ambient.

Normalization Procedure

Procedures facilitating “spectral shape matching only,” such as normalization, can be extremely valuable under many circumstances. For RT detection, in general, the radiance spectra $L(\lambda)$ are normalized to unity amplitude as follows:

$$L(\lambda)_{\text{norm}} = \frac{L(\lambda)}{\sqrt{L(\lambda)^T L(\lambda)}} = \frac{L(\lambda)}{\|L\|}. \quad (4)$$

This is performed on each of the libraries in the library generation phase and on every pixel in the HSI datacube in the data acquisition phase, as shown in Fig. 1. The normalization procedure is a well-known method to facilitate spectral shape matching only between spectral libraries and HSI data. The common spectral angle mapper classifier performs this inherently. The absolute amplitudes are not considered as they generally have higher uncertainty than their spectral shapes or

band-to-band changes. These uncertainties arise from absolute calibration errors of the sensor and/or errors in the intensities of the reflectance measurements that are propagated during the reflectance-to-radiance conversion. Although the sensor errors can be corrected with “vicarious calibration” (see the next section), the reflectance errors, if present, cannot.

The normalization procedure, though, should not be used for SPT detection where the datacube contains mixed pixels or pixels that are a combination of target and background material. Cross terms from normalization of the mixed pixel spectra cause errors when matching the normalized target radiance libraries to the normalized mixed pixel spectra. These errors are, in general, greater than any errors in

their amplitudes.

Approaches using both “amplitude and spectral shape matching” are also preferred for HSI sensors with a minimal number of high signal to noise ratio (SNR) bands and thus limited spectral shape information. Finally, these approaches are advantageous for HSI scenes exhibiting a significant number of background pixels with amplitudes vastly different than the target library spectra. Thus, a decision point is added regarding whether or not to normalize the spectral data and libraries, as indicated in Fig. 1, based on the target resolution as well as the state of the HSI sensor and scene to be imaged. The resulting target radiance libraries are read into the automated detection processor.

DATA ACQUISITION PHASE

The data acquisition phase simulates the process of the data being acquired by the HSI sensor through selection by the user of a particular datacube with known sensor specifications. The raw datacube can be calibrated in both amplitude and spectral shape, either through onboard calibration techniques or by the process called vicarious calibration. The latter approach is often used on sensors where the calibration is not well understood or is not stable throughout a data collection. It consists of repeating the library generation phase for a reference on the ground where accurate reflectance measurements exist. The reference could be a calibration panel or a very uniform background area such as a dry lake bed or a crop field. The library generator predicts what a calibrated HSI sensor would measure when imaging the reference. By comparing the predicted radiance spectra with the measured spectra over the reference, gain coefficients are derived and used to vicariously calibrate the HSI sensor data.

The measurements of the target reflectances, unlike the well-known reference, may contain uncertainties in their absolute amplitudes that, when converted to spectral radiance, could affect the target's detectability in the HSI data. Thus, as the final step, the normalization procedure is conducted on each pixel in the datacube when appropriate for RT detection as discussed in the previous section. The resulting calibrated data are then read into the automated detection processor.

AUTOMATED DETECTION PHASE

The automated detection phase provides a sequence of steps, first to match the radiance library spectra to the calibrated hyperspectral datacube and then to extract the optimal information from the resulting detection score images. The number of library spectra equals the number of target types, multiplied by the number of illumination (or temperature) conditions. This is referred to as the total number of classes that enter the detection processor. Other input parameters for this phase are read in from an ASCII control file. This section discusses each step of the detection processor.

Adaptive Matched Filter

The process begins with the application of the adaptive matched filter (MF). For all pixels x in the datacube, the MF score is computed as follows:

$$MF_{SCORE} = \frac{(L_T - L_M)^T \Sigma^{-1} (X - L_M)}{(L_T - L_M)^T \Sigma^{-1} (L_T - L_M)}, \quad (5)$$

where

- L_T = the target radiance or class spectra,
- L_M = the mean spectra of the datacube,
- X = the image spectra at pixel x , and
- Σ = the datacube covariance matrix.²

The MF uses the datacube statistics to both suppress the background and enhance the target SNR. The filter assumes that the target spectra in the image are rare, i.e., the datacube statistics (mean plus covariance matrix) are approximately equal to the background statistics. Using Eq. 5, image spectra equal to the mean have MF scores equal to zero, while image spectra identical to the target spectra have MF scores equal to 1.

Although the MF is generally applied to RT detection, it can be used for SPT detection under certain circumstances, namely,

1. The SPT has a large pixel fill factor
2. The amplitude of the SPT's radiance spectrum is large relative to the background
3. Some combination of cases 1 and 2

For SPTs that do not satisfy one of these conditions, the MF can be replaced by a true SPT detector that provides a target abundance estimate. An example of this is the adaptive cosine/coherence estimator or ACE algorithm.¹⁵ Regardless of the detector that is applied, the subsequent steps in the detection processor remain the same as they operate on the resultant detection score images. The number of score images equals the number of classes entered.

Class Filtering

The library generation phase discussed previously allows the target libraries or classes to be selected before data collection occurs. In some cases, particular classes may not be optimal for the imaged scene. The quality of a class is determined by its potential for false alarms (FAs), regardless of whether the associated target is present in the scene. One measure of this is the one sigma standard deviation in the target's detection score histogram curve, which is generally Gaussian out to one sigma and centered around zero (the background mean). If the standard deviation of the histogram curve is large (>0.1), it indicates that the target class closely matches a large number of pixels in the scene, signifying a high probability of FAs. The class filtering step removes such classes from the suite of target detection score images.

Bad Line and Spatial Filtering

A number of preselected classes often fall into the marginal category in terms of quality. Their potential for FAs is low enough to satisfy the class filtering criteria, but they require some "cleaning up" before moving to subsequent steps in the detection processor. These procedures take place on the remaining detection score images at three sigma from the background mean. Bad lines are often artifacts that are inadvertently enhanced by the matched filter because of a few spurious pixels in the cross-track array of a push-broom HSI sensor. Similarly, anomalous spatial features in some images can create contiguous sets of pixels with irregular shapes and/or large sizes relative to the known shapes and sizes of the target suite. These sources of FAs are identified with standard shape and size filters and nullified in the images. Once these cleaning procedures are complete, the classes pass on to the next step in the detection processor.

Adaptive Thresholding

The ability to assign the optimal threshold adaptively to each class is critical to the robustness and automation of the detection processor. One method found in the field of medical image processing¹⁶ is particularly attractive, especially for RT detection, though it can also be applied to SPT detection. The approach is based on the tracking of "blobs," defined as any single pixel or contiguous set of pixels in the detection score images.

The general trend for scenes with no targets is for the number of blobs to decrease as the threshold is raised, starting at three sigma from the background mean. For scenes with RTs, though, the number of blobs increases slightly as the threshold is raised when contiguous sets of target pixels first begin to break apart. This increase is detectable, allowing the threshold between target and background to be drawn just before the increase.

To illustrate this process, Fig. 3 plots the number of blobs versus the threshold for four different detection score images of a scene containing several RTs. Note in some cases that a series of local minimums result in the curves. The optimal threshold occurs at the first local minimum, which is unique for each detection score image as indicated by the vertical dashed lines on the plots. For SPTs, a flattening or leveling off in the blob number curve occurs instead of an increase as the optimal threshold is approached. While not as precise for SPTs, this technique does provide a good estimate to first order. When there are no targets or target-like objects in the detection score image, the number of blobs will continue to decrease and the threshold will eventually exceed the largest detection score value, resulting in no detections.

Confidence Assignments

For each remaining class, all pixels in the score image with magnitudes that meet or exceed their associated thresholds are tagged. The number of tags or hits each

pixel receives from different classes is summed. The frequency of these hits can be used as a measure of detection confidence for both RTs and SPTs for the following reasons. First, for VIS/NIR/SWIR sensors, recall that we compute the radiance spectra of a material surface under five illumination conditions. In many cases, the degree of illumination on a surface may vary across the surface itself based on the position of the Sun and the surface topography. Similarly for LWIR sensors, the temperature across a surface might vary, resulting in numerous hits. In addition, a particular surface may contain different but spectrally similar materials (e.g., various shades of green paint) that are included in our suite of targets. Thus, the natural variability in illumination (or temperature) and material of the surfaces provides the potential for multiple hits on those surfaces and the assignment of confidences.

The latter case, though, does depend on the selection of the targets to be detected when planning an operation. Specifically, the number and variety of selected targets should include all possible target types being sought and capture the expected material variability in the target surfaces. An example of the tag summation process is illustrated in Fig. 4 on a scene containing multiple RTs. As pixels receive more hits across the suite of detection score images, their confidences increase. In this case, confidences of “low,” “medium,” and “high” are assigned to those pixels receiving at least 2, 5, or 10 hits, respectively. Background pixels are labeled with either a zero or 1 tag, indicating “no confidence” in terms of being target pixels.

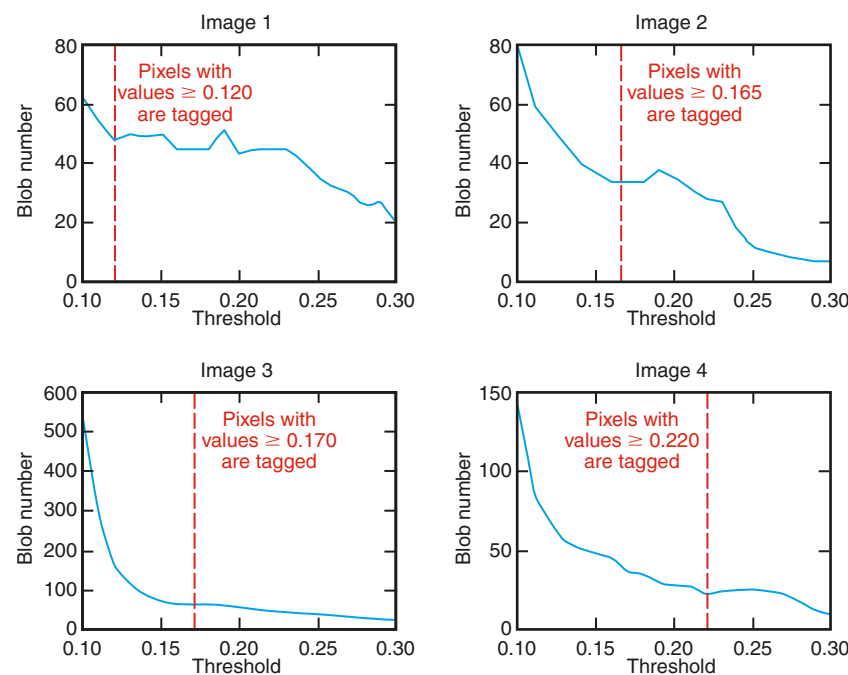


Figure 3. Illustration of adaptive thresholding, with plots of the blob number versus the threshold for four different detection score images of a scene containing several resolved targets. The pixels in each score image are tagged based on an optimal and unique threshold, as indicated by the vertical dashed lines.

Geolocation and Visualization

Pixels with assigned detection confidences are stored as regions of interest and linked to the original datacube for geolocation and visualization purposes. Overlaying the confidence-based color-coded detection tags onto a true color image of the scene is our standard method for visualization of the final results. The software for the automated detection phase can be installed onto a platform-based processor, from which the results can be transmitted via a low-bandwidth link to operators on the ground.

APPLICATIONS

Resolved Target Detection

Extensive testing has been done to ensure that the presented Target




Figure 4. Illustration of the tag summation process for the assignment of confidences on a scene containing multiple resolved targets. In this scene, confidences are allocated to pixels based on the following number of tags: none for 0–1 tag, low for 2–4 tags, medium for 5–9 tags, and high for 10 or more tags.

Detection System performs robustly over a range of targets, backgrounds, and environmental conditions. For RT detection, we use HYDICE test data collected over both desert and rural areas. The vicarious calibration procedure is not required as these data are well calibrated. The shapes and sizes of the targets in the data are vehicle-like in nature, with a variety of painted surfaces. The sensor collected HSI data from a 1.7- to

1.9-km altitude, providing multiple pixels on target depending on size, and from a 3.4- to 3.8-km altitude, producing a small number of pixels on target. The targets in the desert background are deployed in two configurations. In the “open” configuration, the targets are placed in open areas with very little obstruction. In the “closed” configuration, the targets are either partially or almost fully obstructed by natural and manmade canopies in the various desert wash areas. The targets in the rural background are deployed in a “mixed” configuration, with some targets in the open and others either partially or nearly fully obstructed by canopies.

Our Target Detection System is tasked to search for all targets using 8 surface types for the desert areas and 10 surface types for the rural areas. Co-located radiosonde data are used as a proxy for weather forecast data in the generation of the target radiance libraries. The weather profiles were acquired close in time to the HSI data collections. Our RT detection performance over 26 datacubes is summarized in Table 1, with the probability of detection P_d measured by the percentage of targets detected, versus the FA rate, measured by the number of FAs per square kilometer. The performance is provided at three levels of confidence (low, medium, and high) from both the lower- and higher-altitude data for each background and configuration. The total number of pixels processed and targets imaged for each case is also given. These results are plotted in Fig. 5. Unlike typical receiver operating characteristics curves, which vary based on the assigned threshold, these curves vary according to the level of detection confidence. Recall that the adaptive thresholding procedure is part of the automated detection phase. As expected, low confidences result at higher P_d and FA rates, whereas high confidences result at lower P_d and FA rates. Fig. 5 also provides a best-fit curve over all conditions, establishing an average P_d of 77% with 10 FA/km².

Table 1. Resolved target detection performance on HYDICE test data.											
Desert background, open configuration				Desert background, closed configuration				Rural background, mixed configuration			
1.7-km altitude (307 × 7360 pixels over 6 datacubes with 71 targets)		3.8-km altitude (307 × 3200 pixels over 4 datacubes with 68 targets)		1.7-km altitude (307 × 6720 pixels over 6 datacubes with 61 targets)		3.8-km altitude (307 × 1920 pixels over 2 datacubes with 56 targets)		1.9-km altitude (307 × 6400 pixels over 6 datacubes with 174 targets)		3.4-km altitude (307 × 1600 pixels over 2 datacubes with 44 targets)	
P_d (%)	FA/km ²	P_d (%)	FA/km ²	P_d (%)	FA/km ²	P_d (%)	FA/km ²	P_d (%)	FA/km ²	P_d (%)	FA/km ²
Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.
Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.
High conf.	High conf.	High conf.	High conf.	High conf.	High conf.	High conf.	High conf.	High conf.	High conf.	High conf.	High conf.
96	65	94	38	90	67	66	04	82	41	82	08
90	26	90	24	84	30	50	03	79	19	64	03
80	10	62	12	79	08	34	01	75	05	57	00




Figure 5. Resolved target detection performance curves on HYDICE data collections over a range of backgrounds, target configurations, and altitudes. The targets are vehicle-like in shape and size. A best-fit curve over all cases is also provided. (L, M, and H= low, medium, and high confidence, respectively.)

Subpixel Target Detection

For SPT detection, we use COMPASS test data collected over rural areas, with targets placed atop dirt, short grass, and tall grass terrain. The vicarious calibration procedure is conducted for these data using a uniform dirt background as a reference, where detailed *in situ* reflectance measurements were made. The targets considered in these data are mock land mines that vary from 15 to 30 cm across with a variety of painted casings. The sensor collected HSI data from a 1.2-km altitude, imaging SPTs with fill factors between 25 and 50%; from a 0.6-km altitude, imaging partially resolved (i.e., one or two) pixel targets; and from a 0.3-km altitude, imaging RTs with two to four pixels on target. The targets from the 0.6-km altitude are on the border between SPTs and RTs. They are treated as SPTs since the majority of those pixels will contain some mixture of target and background. The normalization procedure is only carried out on the 0.3-km altitude data when the targets become fully resolved.

Our Target Detection System is tasked to search for targets that are dispersed on the dirt, short grass, and tall grass backgrounds using nine surface types. Radiosonde data are again used as a surrogate for weather forecast data

in the generation of the target radiance libraries. In this case, though, the nearest radiosonde station was not co-located at the test site, and the weather profile measurements were not time-coincident with the HSI data collections. Our performance for both SPT and RT detection over 24 datacubes is presented in Table 2 and Fig. 6, using the same format described in the previous section. Table 2 also lists the percentage of targets imaged over each of the three backgrounds at every altitude. Note that the FA rates are much higher than those previously reported in Table 1 and Fig. 5. In reality, the number of FAs per image in the HYDICE and COMPASS datacubes of comparable size is quite similar. The principal difference is that the COMPASS sensor collected HSI data at much lower altitudes than HYDICE because of the smaller sizes of the targets, resulting in less area coverage. In addition, for specific comparisons, note that the COMPASS sensor operated in whisk-broom mode at the 1.2-km altitude, imaging 3 times fewer lines per image along-track than most of the HYDICE datacubes collected in

push-broom mode at the 1.7-km altitude. There is also more uncertainty in the weather profiles used for the COMPASS data, but this has less impact on the accuracy of the target radiance libraries since the COMPASS sensor was imaging through a much shorter atmospheric path.

To assist in assessing detection performance, stoplight criteria have been established for target detection at the COMPASS test site as indicated in Fig. 6. In particular, a P_d of at least 50% at a FA rate of 1000/km² separates very good “purple” from satisfactory “green” performance. We achieve purple status at high confidence at the 0.3-km altitude and at all confidences at the 0.6- and 1.2-km altitudes, with the best performance attained at 0.6 km. Finally, for both the HYDICE and COMPASS test data, the entire automated procedure—from target library selection to visualization of the results—is completed in single-digit minutes on a PC workstation for most datacubes, depending on size.

SUMMARY AND FUTURE WORK

A unique three-phase approach for an automated Target Detection System in support of hyperspectral

Table 2. Subpixel and resolved target detection performance on COMPASS test data.

2–4 pixel resolved targets		1–2 pixel targets		Subpixel targets	
0.3-km altitude (246 × 14,800 pixels over 12 datacubes with 287 targets: 33% on dirt, 32% on short grass, 35% on tall grass)		0.6-km altitude (246 × 5900 pixels over 6 datacubes with 260 targets: 25% on dirt, 40% on short grass, 35% on tall grass)		1.2-km altitude (246 × 2442 pixels over 6 datacubes with 271 targets: 52% on dirt, 35% on short grass, 13% on tall grass)	
P_d (%)	FA/km ²	P_d (%)	FA/km ²	P_d (%)	FA/km ²
Low conf.	Low conf.	Low conf.	Low conf.	Low conf.	Low conf.
Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.	Med. conf.
High conf.	High conf.	High conf.	High conf.	High conf.	High conf.
98.3	1770	90.0	665	65.3	349
98.0	1210	86.9	316	64.9	239
97.0	879	82.7	133	51.3	119

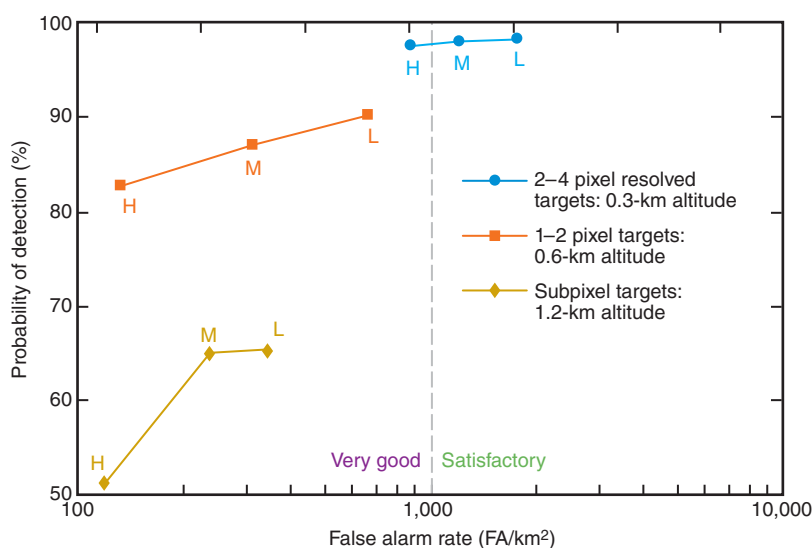


Figure 6. Subpixel and resolved target detection performance curves on COMPASS data collections over dirt, short grass, and tall grass backgrounds at various altitudes. The targets are mock land mines. Stoplight criteria of at least 50% P_d at 1000 FA/km² separate very good from satisfactory performance. (L, M, and H = low, medium, and high confidence, respectively.) Target sizes range from 15 to 30 cm across.

sensor systems has been developed, consisting of a library generation phase, a data acquisition phase, and an automated detection phase. Testing of the system has been conducted on data collected by the HYDICE and COMPASS sensors over both desert and rural area test sites at multiple altitudes. The resulting detection performance approaches the requirements for future operational sensor platforms that warrant quick turnaround products for subsequent action by decision makers or first responders. Future efforts will focus on evaluating the accuracies of target radiance spectral libraries

generated using weather forecast data. Sensitivity of the radiance spectra to uncertainties in temperature, relative humidity, and aerosol loading will be quantified. Such understanding will be valuable in demonstrating this capability over more challenging test sites in littoral areas and even urban areas for the timely detection of spectral targets. (See Ref. 17 for further information on aspects of this work.)

ACKNOWLEDGMENTS. The preparation of this manuscript was supported by a Stuart S. Janney Fellowship courtesy

of the Chair and Committee on Sabbatical Fellows and Professors from The Johns Hopkins University Applied Physics Laboratory.

REFERENCES

- ¹Kraut, S., Scharf, L. L., and McWhorter, L. T., "Adaptive Subspace Detectors," *IEEE Trans. Signal Proc.* **49**, 1–16 (Jan 2001).
- ²Manolakis, D. G., and Shaw, G. A., "Detection Algorithms for Hyperspectral Imaging Applications," *IEEE Signal Proc. Magazine* **19**, 29–43 (Jan 2002).
- ³Thai, B., and Healey, G. E., "Invariant Subpixel Material Detection in Hyperspectral Imagery," *IEEE Trans. Geosci. Remote Sens.* **40**, 599–608 (Mar 2002).
- ⁴Stellman, C. M., Hazel, G. G., Bucholtz, F., Michalowicz, J. V., Stocker, A. D., and Schaff, W. E., "Real Time Hyperspectral Target Detection and Cuing," *Opt. Eng.* **39**, 1928–1935 (Jul 2000).
- ⁵Simi, C. G., Winter, E. M., Schlangen, M. J., and Hill, A. B., "On-Board Processing for COMPASS," *Proc. SPIE* **4381**, 137–142 (Aug 2001).
- ⁶Stevenson, B. P., O'Connor, R., Kendall, W. B., Stocker, A. D., Schaff, W. E., et al., "Design and Performance of the Civil Air Patrol ARCHER Hyperspectral Processing System," *Proc. SPIE* **5806**, 731–742 (Jun 2005).
- ⁷Vane, G., Green, R. O., Chrien, T. G., Enmark, H. T., Hansen, E. G., and Porter, W. M., "The Airborne Visible Infrared Imaging Spectrometer," *Remote Sens. Environ.* **44**, 127–143 (Jun 1993).
- ⁸Basedow, R. W., Carmer, D. C., and Anderson, M. E., "HYDICE System: Implementation and Performance," *Proc. SPIE* **2480**, 258–267 (Jun 1995).
- ⁹Simi, C. G., Winter, E. M., Williams, M. M., and Driscoll, D. C., "Compact Airborne Spectral Sensor (COMPASS)," *Proc. SPIE* **4381**, 129–136 (Aug 2001).
- ¹⁰Hackwell, J. A., Warren, D. W., Bongiovi, R. P., Hansel, S. J., Hayhurst, T. L., et al., "LWIR/MWIR Imaging Hyperspectral Sensor for Airborne and Ground-Based Remote Sensing," *Proc. SPIE* **2819**, 102–107 (Nov 1996).
- ¹¹*Environment for Visualizing Images (ENVI) User's Guide*, ITT Visual Information Solutions, Boulder, CO; <http://www.ittvis.com>.
- ¹²Dudhia, J., Gill, D. O., Manning, K. W., Wang, W., and Bruyere, C., *PSU/NCAR Mesoscale Modeling System Tutorial Class Notes and User's Guide (MM5 Modeling System Version 3)*, University Corporation for Atmospheric Research (UCAR); <http://www.mmm.ucar.edu/mm5/documents/tutorial-v3-notes.html> (Jan 2005).
- ¹³*Radiosonde Database Access*, Forecast Systems Laboratory, National Oceanic and Atmospheric Administration (NOAA), Boulder, CO; <http://raob.fsl.noaa.gov>.
- ¹⁴Anderson, G. P., Berk, A., Acharya, P. K., Matthew, M. W., Bernstein, L. S., et al., "MODTRAN4: Radiative Transfer Modeling for Remote Sensing," *Proc. SPIE* **4049**, 176–183 (Aug 2000).
- ¹⁵Kraut, S., and Sharf, L. L., "The CFAR Adaptive Sub-Space Detector Is a Scale-Invariant GLRT," *IEEE Trans. Signal Proc.* **47**, 2538–2541 (Sep 1999).
- ¹⁶Spisz, T. S., Fang, Y., Reeves, R. H., Seymour, C. K., Bankman, I. N., and Hoh, J. H., "Automated Sizing of DNA Fragments in Atomic Force Microscope Images," *Med. Biol. Eng. Comput.* **36**, 667–672 (Nov 1998).
- ¹⁷Kolodner, M. A., Murphy, P. K., and Hume, E. E. Jr., "Radiance Library Forecasting for Time-Critical Hyperspectral Target Detection Systems," U.S. Patent Number 7,043,369 (May 2006).

The Author

Marc A. Kolodner is a member of the APL Senior Professional Staff in the Space Department. He received his B.S. in physics from Tulane University in 1990 and his M.S. and Ph.D. in physics from the Georgia Institute of Technology in 1994 and 1997, respectively, specializing in microwave remote sensing of planetary atmospheres. Dr. Kolodner joined APL in 1999 and is now part of the Defense Analysis and Applications Group. His work has focused on the development of prototype software for both multi- and hyper-spectral remote sensing applications, including physics-based spectral signature modeling and data exploitation. He was recently awarded a patent for his work. His e-mail address is marc.kolodner@jhuapl.edu.



Marc A. Kolodner