ROBUST HYPERSPECTRAL DETECTION WITH ALGORITHM FUSION

Steven Adler-Golden and Patrick Conforti

Spectral Sciences, Inc., Burlington MA 01803-3304

ABSTRACT

Two simple methods are described for fusing the outputs of hyperspectral rare target detection algorithms to achieve more consistent results across a variety of images and objects of interest. The methods are demonstrated with atmospherically corrected (spectral reflectance) visible/near-infrared/shortwave-infrared and long-wavelength infrared hyperspectral imagery using five different detection algorithms that output generalized likelihood ratio decision statistics. Results are presented for nine test cases.

***Index Terms*—**hyperspectral, fusion, detection, reflectance

1. INTRODUCTION

Detection algorithms for hyperspectral imagery can display a wide variation in effectiveness, depending on the scene and target of interest, even for a single sensor or class of problems. Often no one algorithm is a clear "winner" or "loser" across a variety of test cases. Development of robust, effective, and automated detection systems would greatly benefit from methods that assess the performance of candidate algorithms and select a "best" one, or that fuse algorithms to provide equivalent, or even better, performance.

This paper demonstrates a simple but effective general approach for fusing "detection algorithms for "rare" targets in an image that have known "ground truth" spectra. It is applied to a family of algorithms that includes the spectral angle mapper (SAM), ACE (adaptive covariance estimator) detector, and linear unmixing [1, 2]. The algorithms outputs are written in a Generalized Likelihood Ratio form that contains the reciprocal of a target-fitting residual, so that a large positive value corresponds to a high likelihood of target presence. Writing the output, or response, image for the *i*th algorithm as a column vector $\mathbf{r}_i$ of length p pixels, the responses for n different algorithms are stacked to form an $n \times p$ array given by

$$\mathbf{R} = [\mathbf{r}_1, \mathbf{r}_2, \ldots \mathbf{r}_n]^T \qquad (1)$$

The essence of the fusion approach is to apply standard multiband rare-target detection algorithms to this response stack. One method uses the matched filter technique [1] with an appropriately defined target "spectrum". Another uses a simple modification of the RX anomaly detection method [3]. The RX method has been previously used for fusing multimodal data, such as hyperspectral plus synthetic-aperture radar [4], and time series of multispectral images [5].

The utility of the fusion methods is illustrated with results from nine test cases based on three different examples of hyperspectral data. The first dataset consists of a long-wavelength infrared (LWIR) SEBASS hyperspectral image that has been processed to spectral reflectance [6]. The second dataset is from the Rochester Institute of Technology "self-test" [7], which is based on a HyMap visible/near-infrared/short-wave-infrared (visible/NIR/SWIR) reflectance image. The third data example is an EO-1 Hyperion image of the Galveston, TX area, also converted to reflectance. Five individual detection algorithms and the two fusion methods were run on a total of nine target-image combinations and the detection results compared.

The two fusion methods are found to provide more consistent detection performance across the range of test cases than any single algorithm, and for a given test case they usually perform as well as the best or next-best algorithm. As the methods are not inherently specific to a particular wavelength range, type of scene, or target of interest, they may be useful in a wide variety of applications. The sensitivity of the results to the mathematical form of the algorithm response remains to be investigated. For an individual algorithm, responses that are monotonically related yield equivalent receiver-operator characteristic (ROC) curves, but this is not the case with the fusion methods in their current form.

2. FUSION METHODS

The first fusion method is based on linearly combining the algorithm responses $\mathbf{r}_i$ to generate an image that maximizes target signal-to-noise. This requires defining an n-band signal vector $\mathbf{t}$ that represents the "spectrum" of expected responses to the target. The ground truth target spectrum turns out not to be useful for calculating $\mathbf{t}$, as it fails to account for unknowns such as spectrum variability, sensor artifacts and background contributions. Instead we derive $\mathbf{t}$ from the response stack itself, assuming that the algorithms are sufficiently sensitive and selective that their responses to an actual target are close to their maximum responses in the scene. This assumption is most sensible when there is at least one target present. Accordingly, the target vector is taken as

$$\mathbf{t} = [\max(\mathbf{r}_1), \max(\mathbf{r}_2), \ldots \max(\mathbf{r}_n)]^T \qquad (2)$$

Under the assumption of Gaussian noise the signal-to-noise optimization solution is given by the output of a matched filter for $\mathbf{t}$. This may be written to within a normalization factor as

$$\mathbf{f}_{MF} = (\mathbf{R}-\mathbf{m})^T [\mathbf{K}^{-1}(\mathbf{t}-\mathbf{m})] \qquad (3)$$

where $\mathbf{m}$ is the $\mathbf{R}$ mean column vector, $\mathbf{K}$ is the covariance matrix for $\mathbf{R}$, and the quantity in brackets is the matched filter vector.

The second fusion method makes an opposite assumption: it presumes that no *a priori* information is available about the algorithm responses to the target. This turns the problem into one of anomaly detection in $\mathbf{R}$ space, which can be thought of as

treating the algorithm response images as measurements of an unknown target taken with different bandpasses or sensors. The standard RX anomaly detector is given by

$$\mathbf{f}_{RX} = (\mathbf{R}-\mathbf{m})^T \mathbf{K}^{-1}(\mathbf{R}-\mathbf{m}) \quad (4)$$

which represents the squared Mahalanobis (whitened) amplitude of the pixel in **R** space.

Eq. (4) by itself does not take advantage of the fact that the algorithm responses to the target are, by design, expected to be both large and positive. A very small probability exists that a non-target pixel might have multiple large *negative* values of $\mathbf{r}_i$-$\mathbf{m}_i$, yielding a large Mahalanobis amplitude and hence a false positive with Eq. (4). These pixels can however be filtered out using a simple criterion such as non-negativity of the total response, which may be applied via

$$\mathbf{f}_{RX} = 0 \text{ for } \Sigma_i(\mathbf{r}_i-\mathbf{m}_i) < 0. \quad (5)$$

3. HYPERSPECTRAL TEST CASES

To evaluate the fusion methods we have assembled nine different hyperspectral test cases, spanning visible through long-wavelength infrared (LWIR) wavelengths.

Test cases 1-4 are from a LWIR hyperspectral image taken by the SEBASS sensor [8] at the Department of Energy's Lamont, OK Atmospheric Radiation Monitoring Site in June, 1997. The site has various ground covers (grass, water, soil, and gravel), buildings, and an array of calibrated emissivity panels. The panel array includes two moderately reflective targets, labeled E1 and E4, which each occupy around 40 pixels. Another panel, E4U, located elsewhere in the image has the identical spectrum. Ground truth panel pixel locations and emissivities were supplied. The E1 spectrum is shown in Ingram and Muse [9]; the E4 spectrum is less reflective but otherwise similar. The SEBASS data were converted from spectral radiance to spectral emissivity estimates using an algorithm described by Adler-Golden *et al.* [6], which follows the basic method of Borel [10]. The results were processed in the form of reflectance = 1-emissivity. Test cases 1 and 2, in which the E1 and E4 plus E4U panels are the targets, utilize the full 700-line image. Test cases 3 and 4 were generated by cropping the image to a 128-line square containing the E1 and E4 panels.

Test cases 5-8 are from the Rochester Institute of Technology (RIT) public hyperspectral self-test and blind test web site (http://dirs.cis.rit.edu/blindtest/) [7]. The targets are fabrics and vehicles of measured reflectance located in a visible/NIR/SWIR image taken in July, 2006 at Cooke City, MT by the HyMap hyperspectral sensor. The data are in reflectance units. We selected four self-test cases, fabrics 1, 2 and 4 and vehicle 1. Target pixel locations are provided as separate region of interest (ROI) maps for subpixel and whole-pixel fills.

Test case 9 is from a February, 2008 visible/near-IR/short-wave IR image of the Galveston, TX area taken by the Hyperion hyperspectral sensor on the EO-1 satellite. We processed the data to spectral reflectance using the FLAASH code [11] in the ITTVIS ENVI software version 4.6.1. The original 3242-line image was cropped to 600 lines to accommodate limited computer RAM. The target to be detected is a large open pile of elemental sulfur, covering around 50 image pixels, located across from Pelican Island. A whole-pixel target ROI was constructed by inspection. The sulfur ground truth spectrum is taken from a US Geological Survey spectral library in ENVI.

For evaluating detection algorithm performance, the whole-pixel ROIs were taken as the "truth" locations, and one-pixel-wide buffers around the ROIs were taken as "ignore" locations (not counted as either detections or false positives). The buffer regions correspond to the subpixel ROIs in the RIT test cases.

4. DETECTION ALGORITHMS

The fusion methods have been exercised with five different hyperspectral detection algorithms that are based on a Generalized Likelihood Ratio (GLR) decision statistic. For the general case in which the pixels contain unknown amounts of target, background, and superimposed Gaussian white noise fluctuations, as described by Bajorski *et al.*[12], the statistic is taken as

$$\mathbf{r} = \frac{\text{RMS background fit residual}}{\text{RMS target and background fit residual}} \quad (6)$$

$$= (\mathbf{x}^T\mathbf{P}_B\mathbf{x} \,/\, \mathbf{x}^T\mathbf{P}_Z\mathbf{x})^{1/2} \quad (7)$$

where **x** is the pixel spectrum and $\mathbf{P}_B$ and $\mathbf{P}_Z$ are, respectively, matrices that project onto the space orthogonal to the background (B) and target-plus-background (Z) subspaces. This **r** statistic is the GLR in Bajorski *et al.* [12] raised to the $1/p$ power, where p is the number of bands, and is therefore monotonic in a Matched Subspace Detector [13]. In the present examples of reflectance imagery, the target subspace is one-dimensional, defined by a scaling of the ground truth reflectance spectrum **s**.

4.1. Unmixing Algorithm

This algorithm models the background subspace with a basis set of endmember pixels. The Eq. (7) projection matrices may be computed from Eq. (9) in Bajorski *et al.* [12].

Automated routines are available for finding suitable scene endmembers or other background basis spectra. With any routine, the appropriate number of basis spectra needs to be made. Too many spectra could accurately fit target as well as background signatures, making it difficult to distinguish them. Too few may result in poor background fits that allow substantial residual reduction when the target is included in the basis set, leading to large Eq. (7) ratios even for non-target pixels. A further complication is the possibility that the selected basis spectra will be contaminated with the target spectrum. With SVD eigenvector basis spectra, used by Thai and Healey [14] and Bajorski *et al.* [12], a high-contrast target, particularly one occupying a significant fraction of the image (*i.e.*, a non-rare target), can contaminate the eigenvectors. With endmember pixel basis spectra, a target present in the image may show up as one of the endmembers, even if the target is rare.

We have recently developed an automated method for selecting background endmember basis vectors that largely overcomes these problems. The procedure is as follows:

1. The target spectrum or subspace is projected out of the data.

2. Endmember pixels of the reduced-dimension result are found. In the MaxD algorithm [12] the first endmember is the brightest pixel and the second endmember is the dimmest. Before running MaxD we add a zero spectrum to the image, forcing the second endmember to be zero. The total number of endmembers acquired should be more than what will eventually be needed; around 25 usually suffice. The locations of the endmember pixels are tabulated.
3. The original, full-dimension spectra of the tabulated pixels form an initial background endmember set. The image pixels are then sequentially least-squares fit with these endmembers via projections until the fit residual is reduced to a specified value, thereby defining the final, selected endmember set.

In the present case of a single-spectrum target, we determined the number of background endmembers by requiring that the scene RMS residual in step 3 be no smaller than a value representing the estimated RMS level of sensor noise and artifacts in the data. This value, which is sensor-dependent, is the only adjustable parameter in all of the algorithms in this paper. It has been set to 0.002 for the HyMap data, 0.007 for the Hyperion data, and 0.01 for the SEBASS data; the values are in reflectance units.

4.2. Whole Pixel Algorithm

If only whole-pixel targets are to be detected, and detector white noise is a dominant source of fluctuations in the target-containing pixels, it may be sensible to ignore the background subspace, in which case Eq. (7) becomes

$$\mathbf{r} = [\ \mathbf{x}^T\mathbf{x} / (\mathbf{x}^T\mathbf{x} - \mathbf{x}^T\mathbf{s}(\mathbf{s}^T\mathbf{s})^{-1}\mathbf{s}^T\mathbf{x})\ ]^{1/2} = \csc(\theta)\ , \quad (8)$$

where θ is the vector angle between the target spectrum **s** and the pixel spectrum **x**. This **r** statistic corresponds to the spectral angle mapper (SAM) algorithm [15], whose output is θ.

4.3. Three Stochastic Algorithms

The whole-pixel algorithm is readily generalized to account for non-white-noise fluctuations in the data (*i.e.*, having a non-unit covariance or correlation matrix) by applying a whitening transform to the spectra. If the transform is based on the scene second-order statistics, the algorithm functions as either a subpixel or whole-pixel detector. Subtracting the mean spectrum **m** from all the data and using the scene covariance matrix **Γ** for whitening, the algorithm is described by Eq. (8) with the replacements

$$\mathbf{x} \rightarrow \mathbf{\Gamma}^{-1/2}(\mathbf{x}-\mathbf{m}),\ \mathbf{s} \rightarrow \mathbf{\Gamma}^{-1/2}(\mathbf{s}-\mathbf{m}) \quad (9a)$$

The resulting **r** statistic corresponds to the Adaptive Covariance Estimator (ACE) detector [16]. Alternatively, the scene correlation matrix **C** may be used to whiten non-mean-subtracted data, *i.e.,*

$$\mathbf{x} \rightarrow \mathbf{C}^{-1/2}\mathbf{x},\ \mathbf{s} \rightarrow \mathbf{C}^{-1/2}\mathbf{s} \quad (9b)$$

The use of non-mean-subtracted data in Eq. (9b) is a characteristic shared with algorithms such as SAM, Constrained Energy Minimization [17], and "invariant" subspace algorithms that utilize SVD basis spectra [14]. We call the Eq. (9b) algorithm Whitened Angle Mapper (WAM).

Finally, to avoid the potential problem of contamination of the scene second-order statistics by the presence of a large or high-contrast target we may use the background endmembers from the unmixing algorithm to build a "target-free" version of the data, from which the covariance or correlation matrix is then calculated. We thus derive a target-excluded variant of WAM, which we call TWAM.

5. RESULTS

For the nine test cases, ROC curves, plotting detected target pixel fraction versus false positive fraction, were constructed for each of the five detection algorithms and two fusion algorithms. The results are summarized in Fig. 1 in the form of the false positive fraction when detecting 50% or more of the target pixels, a useful number for characterizing ROC curve performance. For the single-whole-pixel targets this corresponds to detection of the whole target. The displayed quantity is $-\log_{10}$(false positives per pixel+10^{-7}), which indicates better detection with larger values and allows many orders of magnitude of false positives, including zero, to be shown on the same scale.

The Fig. 1 results show the superiority of the matched filter fusion (MFF) and RX fusion (RXF) methods (solid lines) over the individual algorithms. The closely related stochastic algorithms ACE and WAM (dashed lines) are the best single algorithms overall. However, they perform poorly in the cropped SEBASS test cases 3 and 4, where the targets are not rare and therefore contaminate the covariance and correlation matrices. Here the TWAM algorithm, which seeks to isolate the background subspace from the target, works the best. The two fusion algorithms are of very similar quality, replicating the best detector results in nearly every case. The RXF algorithm (triangles) does not quite match the effectiveness of ACE and WAM in case 7, but outperforms all other algorithms in case 9, where it yields zero false positives. The Unmixing and SAM algorithms perform especially poorly in the current reflectance test cases; however, we have had reasonable success with analogous target subspace versions of these algorithms [6].

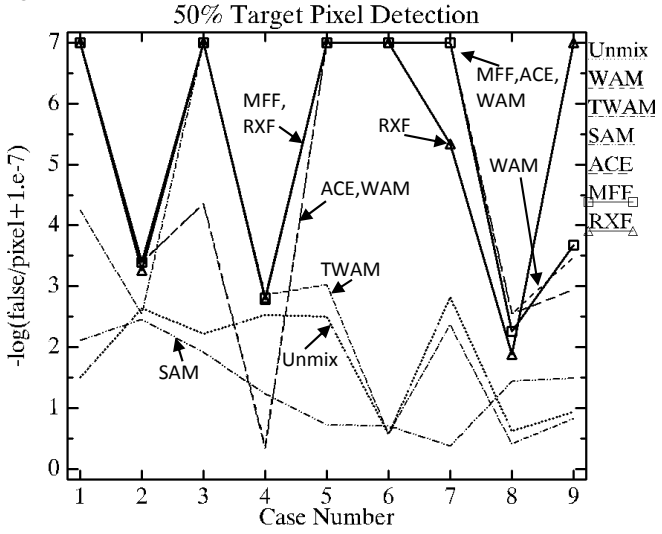

Fig. 1. Detection performance of the five individual algorithms (Unmixing, WAM, TWAM, SAM and ACE; dashed and dotted lines) and two fusion algorithms (MFF and RXF; solid lines) in the hyperspectral test cases.

6. SUMMARY AND CONCLUSIONS

Application of basic matched filtering or anomaly detection techniques to an output stack of GLR decision statistics has been shown in nine different test cases to provide effective means for fusing results from multiple rare target hyperspectral detection algorithms, providing more consistent performance than can be obtained from any one algorithm. The fusion approach generally replicates the results of the best algorithms for a given test case and is insensitive to the presence of poor performers. This enables typically mediocre algorithms that are occasional standouts to be included in the mix. One such algorithm, TWAM, uses endmember reconstruction to suppress contamination of scene second-order statistics by the target presence. Although the results are not included here, we have had similar success in fusing target subspace algorithms, which are very useful for LWIR radiance imagery [6].

Several issues remain unresolved. The fusion results depend on the precise form of the algorithm decision statistic; this dependence should be investigated to better understand the favorable properties of the RMS residual ratio used here, and perhaps to develop a superior statistic. The matched filter and RX fusion methods follow from an assumption of a multidimensional Gaussian distribution of the stacked statistics. The actual distribution is skewed, and can suffer from target contamination problems analogous to those in stochastic detection algorithms. We therefore expect that further work will yield improved fusion methods.

7. ACKNOWLEDGEMENTS

The authors are grateful to Joao Romano (US Army Picatinny Arsenal) for programmatic guidance, Dr. Alan Stocker (Space Computer Corp.) for sharing results of previous research, Pat Cappelaere (Vightel Corp.) for providing Hyperion data, and John Gruninger (Spectral Sciences, Inc.) for technical review. This work was funded by the US Army under Contract No. W15QKN-08-C-0017.

8. REFERENCES

[1] Manolakis, D. and G. Shaw, "Detection Algorithms for Hyperspectral Imaging Applications," IEEE Signal Processing Magazine, 19, Issue 1, pp. 29-43 (2002).

[2] Keshava, N., "A Survey of Spectral Unmixing Algorithms," MIT Lincoln Laboratory, www.ll.mit.edu/news/journal/pdf/vol14_no1/14_1suvery.pdf (2006).

[3] Reed, I.S. and X. Yu, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," IEEE Trans. Acoustics, Speech, and Signal Processing, 38, pp. 1760-1770 (1990).

[4] Nasrabadi, N.M., "A nonlinear kernel-based joint fusion/detection of anomalies using Hyperspectral and SAR imagery," 15th IEEE International Conference on Image Processing, 12-15 Oct. 2008, pp. 1864 – 1867 (2008).

[5] Theiler, J. and S. Adler-Golden, "Detection of Ephemeral Changes in Sequences of Images," 2008 International Conference on Artificial Intelligence and Pattern Recognition, Orlando, FL, 7-10 July (2008).

[6] Adler-Golden, S., J. Gruninger and R. Sundberg, "Hyperspectral Detection and Identification with Constrained Target Subspaces," Proc. 2008 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), vol. 2, pp. 915-918, Boston, MA (2008).

[7] Snyder, D., J. Kerekes, I. Fairweather, R. Crabtree, J. Shive, and S. Hager, "Development of a Web-based Application to Evaluate Target Finding Algorithms," Proc. 2008 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), vol. 2, pp. 915-918, Boston, MA (2008).

[8] Hackwell, J.A., D.W. Warren, R.P. Bongiovi, S.J. Hansel, T.L. Hayhurst, M.G. Sivjee, and J.W. Skinner, "LWIR/MWIR imaging hyperspectral sensor for airborne and ground-based remote sensing," SPIE Proceedings, Imaging Spectrometry, Vol. 2819, pp. 102-107 (1996).

[9] Ingram, P.M. and A. H. Muse, "Sensitivity of Iterative Spectrally Smooth Temperature/Emissivity Separation to Algorithmic Assumptions and Measurement Noise," IEEE Trans. Geosci. Remote Sens. 39, pp. 2158-2167 (2001).

[10] Borel, C.C., "ARTEMISS – an Algorithm to Retrieve Temperature and Emissivity from Hyper-Spectral Thermal Image Data," 28th Annual GOMACTech Conference, Hyperspectral Imaging Session, Tampa, FL, Los Alamos National Lab. Rpt. No. LA-UR-027907 (2003).

[11] Matthew, M.W., S.M. Adler-Golden, A. Berk, G. Felde, G.P. Anderson, D. Gorodetzky, S. Paswaters and M. Shippert, "Atmospheric Correction of Spectral Imagery: Evaluation of the FLAASH Algorithm with AVIRIS Data," SPIE Proc. Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery IX (2003).

[12] Bajorski, P., E.J. Ientilucci and J.R. Schott, "Comparison of Basis-Vector Selection Methods for Target and Background Subspaces as Applied to Subpixel Target Detection," Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery X, Sylvia Shen and Paul Lewis, eds., Proc. SPIE 5425, 97-108 (2004).

[13] Scharf, L.L. and B. Friedlander, "Matched Subspace Detectors," IEEE Trans. Signal Processing, 42, pp. 2146-2157 (August 1994).

[14] Thai, B. and G. Healey, "Invariant Subpixel Material Detection in Hyperspectral Imagery," IEEE Trans. Geosci. Remote Sensing, 40, 599-608 (2002).

[15] Kruse, F.A. *et al.*, "The Spectral Image Processing System (SIPS)--Interactive Visualization and Analysis of Imaging Spectrometer Data," Remote Sens. Environ., 44, pp. 145-163 (1993).

[16] Kraut, S.S., L.L. Scharf, R.W. Butler, "The Adaptive Coherence Estimator: A Uniformly Most-Powerful Invariant Adaptive Detection Statistic," IEEE Trans. Signal Processing, Vol. 53, No. 2 (2005).

[17] Harsanyi, J.C., W. Farrand and C. -I. Chang, "Detection of subpixel spectral signatures in hyperspectral imaging sequences," in Annual Meeting, Proc. American Society of Photogrammetry and Remote Sensing, Reno, NV, pp. 236-247 (1994).