Material detection in hyperspectral imagery in support of nuclear nonproliferation

Amanda Ziemann and James Theiler

Space Data Science and Systems Group, Los Alamos National Laboratory, Los Alamos, NM
ziemann@lanl.gov, jt@lanl.gov

INTRODUCTION

Remote sensing allows for the analysis and understanding of a scene through the use of data collected from satellite or airborne platforms. Hyperspectral imagery (HSI), in particular, is a sensor modality that provides the spectral information necessary to remotely discriminate between materials in the imaged area (*e.g.*, vegetation, soil, buildings, and pavement). Unlike traditional photographic imagery, which is collected at red, green, and blue wavelengths, HSI is collected at hundreds of narrow, contiguous wavelengths. These wavelengths can range from the visible into the near infrared, shortwave infrared, and often midwave and longwave (thermal) infrared, capturing information beyond what the human eye can see. Because different materials have different spectral signatures, analysts can exploit these spectral properties to remotely characterize the material content of a scene in exquisite detail. HSI data supports a variety of materially-focused analyses including classification, change detection, anomaly detection, target detection, and broad area search, with applications ranging from monitoring urban development to detecting camouflaged vehicles. The focus here is target detection, which specifically refers to the algorithms that detect the spectral signatures of targeted materials; within the application space of nuclear security, the ability to remotely detect specific "materials of interest" (MOIs) is of great importance to nonproliferation.

HYPERSPECTRAL TARGET DETECTION

For a hyperspectral image collected over d spectral bands, the pixels can be represented as vectors in $\mathbb{R}^d$. In this *spectral space* representation – which is the predominant data space for hyperspectral image analysis – spatial information is not inherently incorporated. Each vector corresponds to the sensor-reaching radiance spectrum or ground-estimated reflectance spectrum at a particular location on the ground, and that in turn depends on the physical materials at that location. A perennial challenge in hyperspectral analysis is that the spatial extent of a pixel does not necessarily contain a single material; if the pixel corresponds to a region containing multiple materials, then the sensor integrates the radiometric responses of those materials into one spectrum for that pixel. Hyperspectral images may contain "pure" pixels (covering a single material) or "mixed" pixels (with multiple materials in the same pixel location). Even for pure pixels, material discrimination is complicated by uncertainties in illumination, atmospheric distortion, sensor noise, and material composition. The complications are further compounded by spectral variability of the solid materials due to variations in morphology such as grain size, packing density, *etc.*, as demonstrated, for instance, by Myers *et al.* [1].

Target detection algorithms leverage the pixel-level material information in HSI in an attempt to identify all instances of a particular material within the scene. The input signal is a spectrum corresponding to the target material, and that spectrum can be lab-measured, field-measured, or image-derived. The detector then uses a statistical, geometrical, or graph-based measure to assign a target-likelihood score to each pixel in the image, indicating how likely each pixel is to contain the target material in either full or sub-pixel abundances. The most commonly used detectors – which include the Spectral Angle Mapper (SAM) [2], the Adaptive Matched Filter (AMF) [3, 4] and the Adaptive Cosine/Coherence Estimator (ACE) [5, 6] – employ only the spectral information in the image.

As noted in the previous section, there are additional challenges when any of the relevant materials exist in particulate form: traditional detection approaches assume that the target material in the scene "looks like" the input target spectrum, but morphological variations within a single material can lead to highly variable spectra. This is in contrast to gas-phase chemical plumes, for example, which tend to have more consistent spectral features. MOI detection in support of nuclear nonproliferation explicitly refers to materials that are particulates, and the detection of those MOIs can be used to characterize suspected facilities; an illustration of this, along with sample variable spectra, is shown in Fig. 1. This challenge of variable targets has led to the need for algorithms that are more robust to such variability, especially in cluttered backgrounds.

We present two new approaches – a local regression approach and a spectral unmixing approach – that incorporate spatial information in different ways, and compare them to commonly used detection algorithms. Spatial information is not widely used in hyperspectral target detection, and results here will suggest that techniques based on local spatial information can be more robust to target variability. While the motivation here is target detection of MOIs in support of nuclear nonproliferation, we will present an illustrative example using surrogate benign materials. The SHARE 2012 dataset [7] was expressly designed for testing remote sensing analysis schemes, and provides a good example of spectrally-variable targets in cluttered backgrounds.

Regression

To distinguish target signatures in cluttered backgrounds, a key step is the characterization of that background [8]. One very simple background model is the multivariate Gaussian, and it is the basis for classic target detection algorithms such as AMF [3, 4] and ACE [5, 6]. A simple generalization is an elliptically-contoured distribution – like a Gaussian, but with fatter tails – and this has been used with some success in hyperspectral image analysis [9, 10]. These models are highly restrictive, however, and amount to estimating the background at a given pixel by the global mean of all the pixels in the entire image. The use of local estimation for mean and covariance

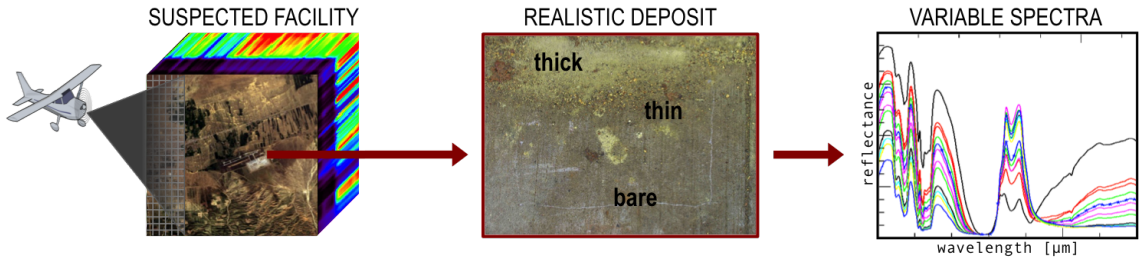

Fig. 1: Illustration of the detection of a material of interest (MOI) deposit for characterizing a facility. The spectra are highly variable, and depend nonlinearly on the morphological properties (particle size, packing density, optical thickness, *etc.*) of the MOI.

was initially suggested by Reed and Yu [11] in the context of anomaly detection [12, 13], but later work considered various schemes for local estimation in target detection [14, 15].

All these local schemes estimate the target-free value of a pixel with the average of pixel values in a spatial annulus surrounding the pixel of interest. Recently, we have been investigating an approach whereby the target-free value of that center pixel is estimated with a general function of those background pixels [16–19]. This corresponds, roughly, to the "filtering" concept advocated by Milanfar [20]. While it is possible to derive this function based on an underlying model (*e.g.*, polynomial) of the background variation [21], or an underlying kernel (*e.g.*, Gaussian) of weights [22], the regression framework is based on the idea of "learning" this function from the image itself.

Unmixing

The second detection approach presented here is CITRUS: Cueing Image Target Regions by Unmixing Spectra [23]. Instead of assigning a target-likelihood score to each *pixel*, as traditional detection algorithms do, CITRUS assigns a score to each local *region*, corresponding to the likelihood that the target material is in that region. CITRUS uses a metric based on the endmembers within each tile, which are the "corners" of the simplex enclosing the data in spectral space; endmember identification is a part of the process known as spectral unmixing. The metric is equal to the Euclidean distance from the locally-mixed input target spectrum to the subspace defined by the endmembers in a tile. If the distance is small, then there is a higher target-likelihood, and if the distance is large, then there is a lower target-likelihood.

Algorithm CITRUS

1: Read in image, target spectrum ($\mathbf{t}$)
2: Segment image into 10×10 pixel spatial tiles
3: Select α
4: **for** each tile **do**
5: Estimate # of endmembers m using Gram matrix [24]
6: Identify m endmembers using MaxD [25, 26]
7: Calculate $\hat{\mathbf{t}} = \alpha\mathbf{t} + (1-\alpha)\boldsymbol{\mu}_{tile}$
8: Compute distance δ from $\hat{\mathbf{t}}$ to the endmember subspace
9: δ = detection score for tile
10: **end for**

RESULTS AND ANALYSIS

The two approaches described here were applied to the SHARE 2012 dataset created by Rochester Institute of Technology's Digital Imaging and Remote Sensing (DIRS) Laboratory in September 2012 [7]. As part of this multimodal experimental campaign, the DIRS lab deployed several red and blue cotton felt panels in both $2m \times 2m$ and $3m \times 3m$ sizes under various states of illumination and occlusion. These were imaged from an aerial platform, resulting in a ground-truthed

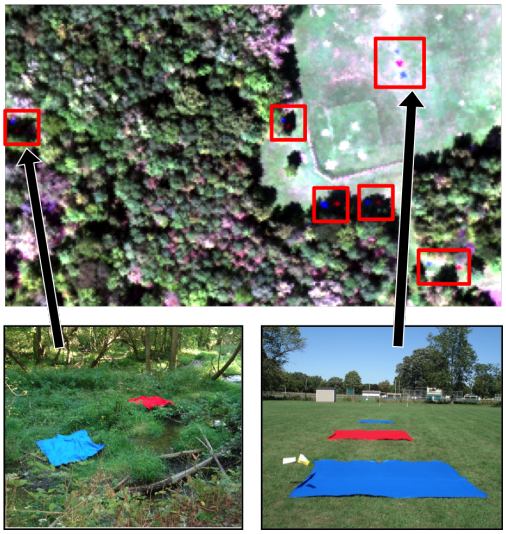

Fig. 2: RGB of SHARE 2012 HSI with known target locations.

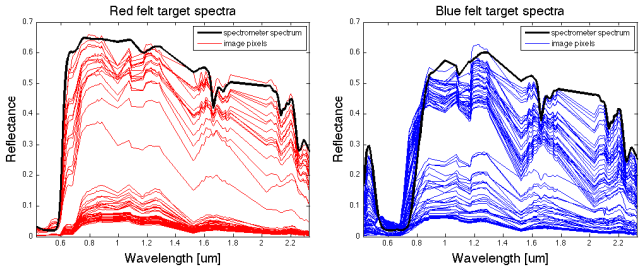

(a) (b)

Fig. 3: Plots demonstrating target variability by comparing the field-measured input spectrum (black) with the spectra of the target-containing pixels in the image for the (a) red and (b) blue felt.

dataset with variable target spectra in both full and subpixel abundances. This imagery has 229 bands (out of 360, after removing atmospherically dominated or low signal to noise bands), over a spectral range of $0.4 - 2.45\mu m$, with a ground sample distance of $\sim 1m$. Fig. 2 shows the SHARE 2012 image with known target locations identified, and Fig. 3 shows the variability of the targets in this dataset. The black lines correspond to the hand-held spectrometer measurement of each material taken at the time of the collect, which are the "input signals" used for the detections in this analysis.

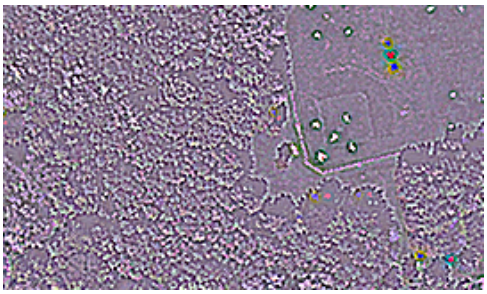

Fig. 4: RGB representation of the residual image.

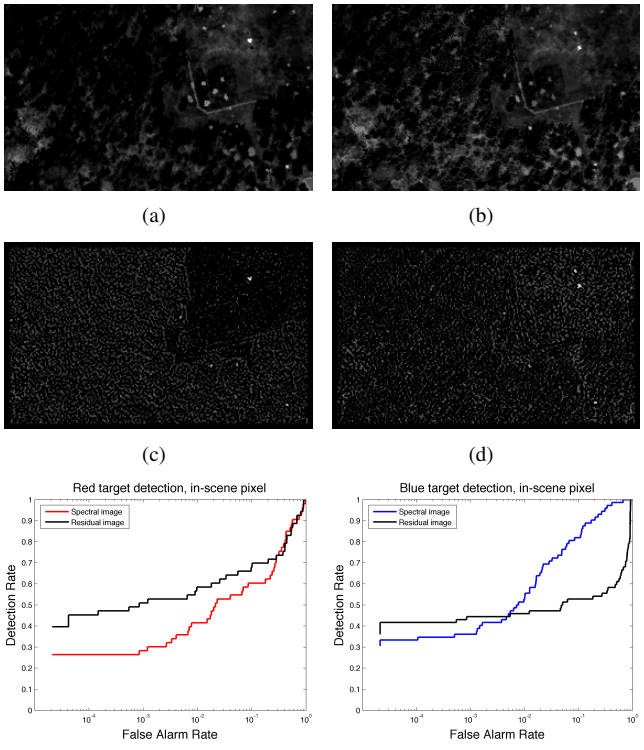

Fig. 5: The left panels (a,c,e) are for the red targets and the right panels (b,d,f) are for the blue targets. The top two panels (a,b) are detection maps against the original spectral imagery, while the middle two (c,d) show detections against the residual image shown in Figure 4. The bottom two panels (e,f) compare ROC curves showing detection performance. The results are not unambiguous, but in the operationally-relevant low false alarm rate regime, the residuals provide a slight advantage.

An RGB representation of the (229-band) residual image is given in Fig. 4, which visually illustrates how the residual image suppresses the background and emphasizes the targets. Fig. 5 shows the results of detection using the Spectral Angle Mapper (SAM) [2] to find targets in both the original scene and in the residual scene built from a local regression estimator.

Fig. 6 shows the results of the ACE detector applied in the spectral space and then compared to CITRUS. To fairly compare the two, we implemented a tiled version of ACE, where the tile score was equal to *max*(ACE scores in tile). The ROC curves were computed using the tiled version of the truth mask. For the red felt targets, CITRUS shows a moderate improvement over ACE at low false alarm rates, and a significant improvement over ACE for the blue felt targets.

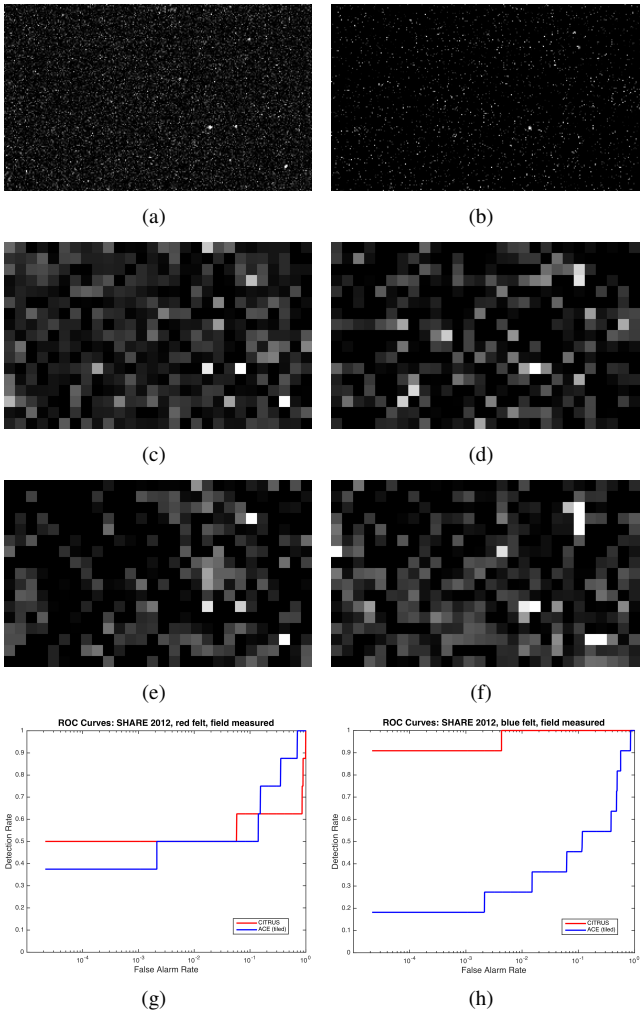

Fig. 6: The left panels (a,c,e,g) are for the red targets and the right panels (b,d,f,h) are for the blue targets. The top two panels (a,b) are per-pixel ACE detection maps against the original spectral imagery, while the second row (c,d) show the tiled version of the ACE detection maps. The third pair of panels (e,f) show the CITRUS detection maps for the red and blue targets, respectively. The bottom two panels (g,h) compare ROC curves showing detection performance. At low false alarm rates, CITRUS performs moderately better than ACE for the red targets, and considerably better than ACE for the blue targets.

CONCLUSIONS

We described two spatio-spectral approaches for improving target detection in hyperspectral imagery, and employed them on an image dataset in which the target signatures varied considerably. The results are preliminary, but may ultimately prove useful against the dangerous proliferation of red and blue felt panels, as well as other spectrally-variable materials of more direct interest to the nonproliferation mission.

ACKNOWLEDGMENTS

We acknowledge support by the United States Department of Energy from the Office on Nonproliferation and Verification Research and Development (NNSA/NA-22).

REFERENCES

1. T. L. MYERS, C. S. BRAUER, Y.-F. SU, T. A. BLAKE, R. G. TONKYN, A. B. ERTEL, T. J. JOHNSON, and R. L. RICHARDSON, "Quantitative reflectance spectra of solid powders as a function of particle size," *Applied Optics*, **54**, 4863–4875 (2015).
2. F. A. KRUSE, A. B. LEFKOFF, J. B. BOARDMAN, K. B. HEIDEBRECHT, A. T. SHAPIRO, P. J. BARLOON, and A. F. H. GOETZ, "The Spectral Image Processing System (SIPS) – Interactive Visualization and Analysis of Imaging spectrometer Data," *Remote Sensing of Environment*, **44**, 145 – 163 (1993).
3. I. S. REED, J. D. MALLETT, and L. E. BRENNAN, "Rapid convergence rate in adaptive arrays," *IEEE Trans. Aerospace and Electronic Systems*, **10**, 853–863 (1974).
4. E. J. KELLY, "Performance of an adaptive detection algorithm: rejection of unwanted signals," *IEEE Trans. Aerospace and Electronic Systems*, **25**, 122–133 (1989).
5. L. T. MCWHORTER, L. L. SCHARF, and L. J. GRIFFITHS, "Adaptive coherence estimation for radar signal processing," in "Conference Record of the Thirtieth Asilomar Conference on Signals, Systems, and Computers.", IEEE (November 1996), vol. 1, pp. 536–540.
6. S. KRAUT, L. L. SCHARF, and R. W. BUTLER, "The adaptive coherence estimator: a uniformly most-powerful-invariant-adaptive detection statistic," *IEEE Transactions on Signal Processing*, **53**, *2*, 427–438 (February 2005).
7. A. GIANNANDREA ET AL., "The SHARE 2012 data campaign," *Proc. SPIE*, **8743**, 87430F (2013).
8. S. MATTEOLI, M. DIANI, and J. THEILER, "An overview of background modeling for detection of targets and anomalies in hyperspectral remotely sensed imagery," *J. Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS)*, **7**, 2317–2336 (2014).
9. D. B. MARDEN and D. MANOLAKIS, "Using elliptically contoured distributions to model hyperspectral imaging data and generate statistically similar synthetic data," *Proc. SPIE*, **5425**, 558–572 (2004).
10. J. THEILER and B. R. FOY, "EC-GLRT: Detecting weak plumes in non-Gaussian hyperspectral clutter using an elliptically-contoured generalized likelihood ratio test," *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, p. I:221 (2008).
11. I. S. REED and X. YU, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoustics, Speech, and Signal Processing*, **38**, 1760–1770 (1990).
12. S. MATTEOLI, M. DIANI, and G. CORSINI, "A tutorial overview of anomaly detection in hyperspectral images," *IEEE A&E Systems Magazine*, **25**, 5–27 (2010).
13. J. THEILER, "By definition undefined: Adventures in anomaly (and anomalous change) detection," *Proc. 6th IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)* (2014).
14. Y. COHEN and S. R. ROTMAN, "Spatial-spectral filtering for the detection of point targets in multi- and hyperspectral data," *Proc. SPIE*, **5806**, 47–55 (2005).
15. Y. COHEN, D. G. BLUMBERG, and S. R. ROTMAN, "Subpixel hyperspectral target detection using local spectral and spatial information," *J. Applied Remote Sensing*, **6**, 063508 (2012).
16. J. THEILER and B. WOHLBERG, "Regression Framework for Background Estimation in Remote Sensing Imagery," *Proc. 5th IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)* (2013).
17. N. HASSON, S. ASULIN, S. R. ROTMAN, and D. BLUMBERG, "Evaluating backgrounds for subpixel target detection: when closer isn't better," *Proc. SPIE*, **9472**, 94720R (2015).
18. J. THEILER, "Symmetrized regression for hyperspectral background estimation," *Proc. SPIE*, **9472**, 94721G (2015).
19. J. THEILER, "Generic target response as a measure of regression accuracy in multispectral background estimation," *Proc. IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI)* (March 2016).
20. P. MILANFAR, "A Tour of Modern Image Filtering: New Insights and Methods, Both Practical and Theoretical," *IEEE Signal Processing Mag.*, **30**, 106–128 (Jan 2013).
21. J. THEILER and J. BLOCH, "Multiple concentric annuli for characterizing spatially nonuniform backgrounds," *The Astrophysical Journal*, **519**, 372–388 (1999).
22. H. TAKEDA, S. FARSIU, and P. MILANFAR, "Kernel regression for image processing and reconstruction," *IEEE Trans. Image Processing*, **16**, 349–366 (2007).
23. A. K. ZIEMANN, "Local spectral unmixing for target detection," in "Proc. Southwest Symposium on Image Analysis and Interpretation (SSIAI)," IEEE (March 2016).
24. D. W. MESSINGER, A. K. ZIEMANN, B. BASENER, and A. SCHLAMM, "Metrics of spectral image complexity with application to large area search," *Optical Engineering*, **51**, *3* (March 2012).
25. K. LEE, *A subpixel scale target detection algorithm for hyperspectral imagery*, Ph.D. thesis, Rochester Institute of Technology (2003).
26. J. SCHOTT, K. LEE, R. RAQUEÑO, G. HOFFMAN, and G. HEALEY, "A Subpixel Target Detection Technique Based on the Invariant Approach," in "Proceedings of the AVIRIS Workshop," NASA JPL (2003).