Hyperspectral target detection using manifold learning and multiple target spectra

Amanda K. Ziemann and James Theiler
Intelligence and Space Research Division
Los Alamos National Laboratory
Los Alamos, NM 87545
Correspondence: ziemann@lanl.gov

David W. Messinger
Carlson Center for Imaging Science
Rochester Institute of Technology
Rochester, NY 14623

Abstract—**Imagery collected from satellites and airborne platforms provides an important tool for remotely analyzing the content of a scene. In particular, the ability to remotely detect a specific material within a scene is of critical importance in nonproliferation and other applications. The sensor systems that process hyperspectral images collect the high-dimensional spectral information necessary to perform these detection analyses. For a $d-$dimensional hyperspectral image, however, where d is the number of spectral bands, it is common for the data to inherently occupy an $m-$dimensional space with $m \ll d$. In the remote sensing community, this has led to recent interest in the use of manifold learning, which seeks to characterize the embedded lower-dimensional, nonlinear manifold that the data discretely approximate. The research presented here focuses on a graph theory and manifold learning approach to target detection, using an adaptive version of locally linear embedding that is biased to separate target pixels from background pixels. This approach incorporates multiple target signatures for a particular material, accounting for the spectral variability that is often present within a solid material of interest.**

I. INTRODUCTION

Unlike traditional color digital images, which are collected at three visible wavelengths (red, green, and blue), hyperspectral imagery (HSI) is collected at hundreds of narrow, contiguous spectral bands. These wavelengths range anywhere from the visible into the near infrared (NIR), short-wave infrared (SWIR), mid-wave infrared (MWIR), and long-wave infrared (LWIR). The advantage of having additional spectral information in HSI is that it allows for greater material separability; materials that look visibly similar can appear very different when examined spectrally. This ability to leverage material separation and identification in remotely sensed hyperspectral scenes is useful to both civilian and military applications, *e.g.*, monitoring urban development, tracking crop growth, and the nonproliferation mission. Broader categories of HSI analysis that utilize this material information include change detection, anomaly detection, target detection, classification, and large area search. Here, we specifically focus on the target detection problem, which – given a spectrum or spectra corresponding to a material of interest – seeks to identify all occurrences of that material in a remotely-sensed scene. Of course, target detection has its own set of challenges, which will be detailed in the next section.

This paper is organized in the following manner. In Section II, we describe the target detection problem, and also explain basic graph theory and manifold learning terminology. Section III presents our graph-based methodology for target detection with spectrally variable targets. Section IV shows the data used in this analysis, and Section V presents the results of the experiment. We conclude with a summary and discussion of future work in Section VI.

II. BACKGROUND

A. Target Detection

One of the challenges in hyperspectral image analysis is within-material variability; unlike gas-phase chemical plumes [1], [2], solid materials cannot be defined be a single, unique, and deterministic spectrum. In other words, there is no single spectrum that describes all grass, and no single spectrum that describes all quartz. This within-material spectral variability is further compounded when the material of interest exists in particle form, due to morphological variations such as particle size and packing density [3], [4]. In target detection, this describes the problem of solid target variability. Target detection also faces the challenge of sub-pixel targets, *i.e.*, when the target material only comprises a fraction of the scene content for a given pixel [5]. When the extent of a pixel corresponds to multiple materials, the sensor integrates the radiometric response of those materials into a single spectral response for that pixel, resulting in a *mixed pixel*. Consequently, the target material often exists in the scene in mixed pixels, or *sub-pixel targets*.

The traditional formulation of the target detection problem uses a spectrum corresponding to a material of interest and, given a hyperspectral image, the detector attempts to identify all occurrences of that target within the scene. This can be done through both statistical [6]–[8] and geometric means [9], [10], and results in a grayscale detection map where each pixel is assigned a detection score $\Delta(\mathbf{x})$. The mathematical framework that is typically used to evaluate these detection scores is hypothesis testing, where the spectra are viewed as random vectors [5]. Given the spectrum of an observed pixel

x, there are two competing hypotheses:

$$H_0 : \text{target absent}$$
$$H_1 : \text{target present.}$$

If $\Delta(\mathbf{x})$ exceeds some threshold η, then the *target present* hypothesis is taken to be true. This detection score is effectively a measure of how "target-like" the pixel is, and can also be viewed as a measure of confidence that the pixel actually contains the target material. However, the "confidence" in these detection scores faces an implicit challenge: as detailed above, it can be difficult to describe a solid material class using a single spectrum. This leads to another formulation of the target detection problem in which multiple target spectra are used. The updated formulation employs a *composite* hypothesis test, so-named because the *target present* hypothesis has multiple components. These multiple target spectra detection approaches are primarily statistical [11]–[14].

It is well-documented that for these detectors, the statistical approaches typically outperform the geometric approaches [15]. However, in complex, materially-cluttered scenes, we expect the structure of the image data to be both nonlinear and non-gaussian [16]. The current models do not exploit the full complexities of these structures, leading to an opportunity to improve upon these models. This motivated the work presented here, which attempts to approach the target detection problem using a completely different mathematical formulation. Specifically, we use a graph-based model (as opposed to the more traditional gaussian or linear mixture models), and then implement a biased nonlinear dimensionality reduction transformation in order to better separate the targets from the background in their new, lower-dimensional space. This is built on previous work, where we addressed the traditional target detection problem using a single target spectrum [17], [18]. Here, we expand on that approach in order to account for solid target variability and multiple target spectra. An overview of graph theory and manifold learning is given in the following subsections, and the full methodology is presented in Section III.

B. Nonlinear Dimensionality Reduction

When working with high-dimensional data, it is often the case that the data can be accurately modeled in lower dimensions. This is particularly true for hyperspectral data. For a $d-$dimensional image (with d spectral bands), it is typical for the data to inherently occupy an $m-$dimensional space, with $m \ll d$. One way to address this is through the use of dimensionality reduction, and these approaches can be both linear and nonlinear. The most well-known linear approach is Principal Components Analysis (PCA) [19]; however, as described in the previous section, hyperspectral data are often highly *nonlinear* in cluttered scenes. As such, our focus here is on nonlinear dimensionality reduction. Another term for dimensionality reduction is *manifold learning*, which assumes that the data discretely approximate a lower-dimensional manifold (*i.e.*, curve or surface) embedded in the original higher-dimensional space containing the data [20].

Popular nonlinear manifold learning algorithms include Kernel PCA [21], ISOMAP [22], Locally Linear Embedding (LLE) [23], [24], and Laplacian Eigenmaps [25], which have been applied to hyperspectral data for a variety of applications [26]–[30]. While manifold learning algorithms do not all initially require a graph-based model, the LLE algorithm - the manifold learning approach used here - *does* use a graph model.

1) Graph Theory: A graph $G = \{V, E\}$ is composed of two finite sets: a vertex set V, an edge set E, and a third optional set of weights ω [31]. In practical applications, the vertices represent the objects, and the edges are the relationships between those objects. A common example of a graph is one in which the vertices are given by user profiles on a social media platform, and the edges represent connections or friendships between those profiles. For an unweighted graph, the relationships are binary, and the edges all have a numerical value of 1. For a weighted graph, ω assigns a positive, numerical value to those edges, or relationships. In the context of hyperspectral data, the spectral vectors (pixels) in the image comprise the vertex set

$$V = \{\mathbf{x}_i\}_{i=1}^N, \tag{1}$$

where $\mathbf{x}_i \in \mathbb{R}^d$ for an image with N pixels and d spectral bands. Although the vertex set in HSI applications is already given, the edge set must still be created. In general applications, the edge set E is typically populated by applying a k-nearest neighbor (kNN) model that creates an edge between each vertex and its k closest neighboring vertices. A user-selected distance metric is used to define "closeness" in this context (*e.g.*, euclidean distance, angular distance).

When two vertices are connected by an edge, they constitute *adjacent* vertices. The adjacency structure of a graph is represented by an adjacency matrix $\mathbf{A}$, which can be defined for both unweighted and weighted graphs. This is an $N \times N$ symmetric matrix whose rows and columns are indexed by the N vertices in V, and the elements are $[\mathbf{A}]_{ij} = [\mathbf{A}]_{ji} = \omega_{ij}$, where ω_{ij} is the weight of the symmetric edge between vertex $\mathbf{x}_i$ and vertex $\mathbf{x}_j$. If the two vertices are *not* connected by an edge, then the corresponding element in $\mathbf{A}$ is assigned a value of zero, $[\mathbf{A}]_{ij} = 0$. The graphs discussed in this paper do not allow for repeated edges and self-loops, which means that each element of the matrix $\mathbf{A}$ has only one value, and the main diagonal entries of the matrix are 0. These basic graph theory definitions play an important role in many manifold learning techniques.

2) Locally Linear Embedding: A manifold may be either linear or nonlinear, and is the higher dimensional analog of 2-D and 3-D curves and surfaces. When examined locally, a manifold behaves like a Euclidean space and appears relatively flat and featureless; globally, a manifold may have far more intricate features. Manifold learning refers to the algorithms that seek to recover these inherently lower-dimensional manifolds that are embedded in the original higher-dimensional space containing data.

The manifold learning algorithm that we utilize here is Locally Linear Embedding (LLE), and was developed by
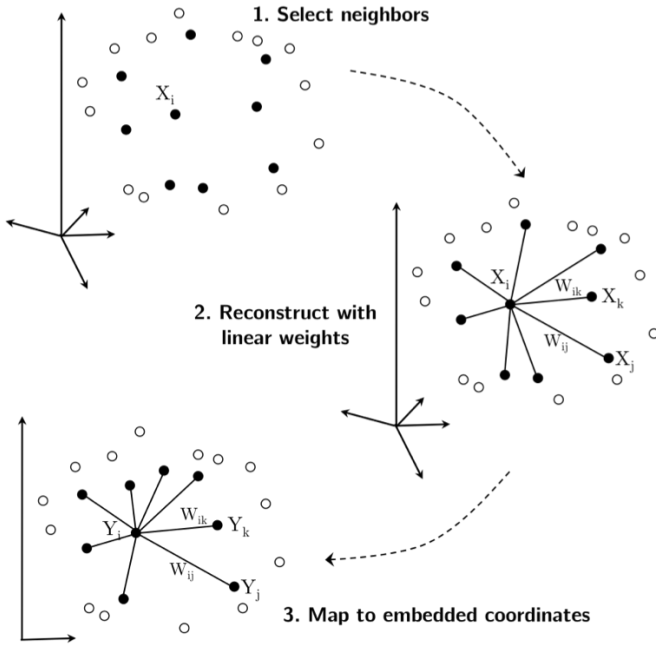
Fig. 1. Chart demonstrating the LLE algorithm for nonlinear manifold learning [23].

Saul and Roweis [23], [24]. LLE has three main steps: (1) nearest neighbor search, (2) constrained least-squares local reconstruction, and (3) spectral embedding. These steps are summarized below, and illustrated in Figure 1. Notationally, $\mathbf{X}$ is the $d \times N$ array of input pixels as columns such that $\vec{\mathbf{X}}_i = \mathbf{x}_i$, and $\mathbf{Y}$ is the $m \times N$ array of output (dimensionality-reduced) pixels as columns such that $\vec{\mathbf{Y}}_i = \mathbf{y}_i$ (where $\mathbf{y}_i$ is the lower-dimensional complement to $\mathbf{x}_i$). Additionally, W is the $N \times N$ matrix of weights associated with each optimized local reconstruction as described in *Step 2*.

1. Nearest neighbor search. The traditional implementation of LLE suggests a kNN approach for building the edge set on the graph, although any initial graph structure may be chosen by the user. In using a kNN graph, the only user-defined free parameter is k, the number of neighbors in each local neighborhood. The kNN approach works well when the data are generally evenly distributed, but for data that have varying neighborhood sizes, the global choice of k can dramatically under-define some neighborhood sizes and dramatically over-define other neighborhood sizes. Depending on the application, the choice of initial graph structure should be carefully considered by the user.

2. Constrained least-squares local reconstruction. After the k respective neighbors of each pixel are identified, the next step is to optimize the unmixing of each pixel with respect to its local neighborhood. The linear reconstruction of each pixel $\vec{\mathbf{X}}_i$ is given by

$$\hat{\mathbf{X}}_i = \sum_{j=1}^{N} W_{ij} \vec{\mathbf{X}}_j, \tag{2}$$

where W_{ij} is the contribution of the j^{th} pixel to the recon-

struction of the i^{th} pixel, and $i \neq j$. The scalar weights W_{ij} are constrained to satisfy $\sum_j W_{ij} = 1$. If $\vec{\mathbf{X}}_j$ does not belong to the neighborhood of $\vec{\mathbf{X}}_i$, then the weighted contribution is $W_{ij} = 0$. Even though the summation is across all N pixels, only the k neighborhood pixels have non-zero values for W_{ij}, and thus only the k neighborhood pixels contribute to the reconstruction. The optimal reconstruction weights in W are computed by minimizing the cost function within each local neighborhood:

$$\mathcal{E}(W) = \sum_{i=1}^{N} \left| \vec{\mathbf{X}}_i - \sum_{j=1}^{N} W_{ij} \vec{\mathbf{X}}_j \right|^2. \tag{3}$$

Equation 3 sums the squared distances between each original pixel and their respective neighborhood-driven reconstructions. The optimized matrix of weight values is denoted $\hat{W}$, given by the minimization of $\mathcal{E}(W)$.

3. Spectral embedding. After $\hat{W}$ is computed in *Step 2*, those weights are then fixed and used to globally optimize the embedded coordinates $\vec{\mathbf{Y}}_i$. The transformation in LLE preserves local linearity (*i.e.*, invariance to scale, translation, and rotation), and so the reconstruction weights within each neighborhood in the spectral space coordinates are preserved within those same neighborhoods in the lower-dimensional transformed coordinates. The cost function that optimizes the global embedding $\mathbf{Y}$ given the fixed neighborhood weights $\hat{W}$ is:

$$\Phi(\mathbf{Y}) = \sum_{i=1}^{N} \left| \vec{\mathbf{Y}}_i - \sum_{j=1}^{N} \hat{W}_{ij} \vec{\mathbf{Y}}_j \right|^2. \tag{4}$$

Roweis and Saul (2000, 2003) show that the optimization of Equation 4 can be computed by solving an $N \times N$ sparse eigenvalue problem. Given the matrix $M = (I - W)^{\mathrm{T}}(I - W)$, the bottom $m + 1$ non-zero eigenvectors are the ordered set of m LLE-derived embedding coordinates (after the single bottom eigenvector of all 1s, and corresponding eigenvalue of 0, is discarded). These m embedding coordinates with N elements each define the $m \times N$ matrix of embedding coordinates $\mathbf{Y}$.

III. METHODOLOGY

The target detection methodology presented here uses a biased nonlinear dimensionality reduction transformation that is designed to separate out the target material from the background data in the image, after which a simple Spectral Angle Mapper detector is applied in the lower-dimensional space. Section II-B2 outlines the steps in the traditional implementation of LLE, where the first step involves building a kNN graph. With hyperspectral data, kNN graphs do not capture the variety of neighborhood sizes, presence of anomalous pixels, etc. As such, it is not appropriate to apply to traditional LLE, with a kNN graph, to a hyperspectral image. In order to more appropriately apply LLE to hyperspectral data, we use our previously developed Adaptive Nearest Neighbors (ANN) graph-building technique [32]. ANN is a data-driven approach

to adaptively identifying a different neighborhood size for each pixel based on an estimate of local density. Here, we implement *adaptive LLE*, where the first step, nearest neighbor search, uses an ANN graph to find a different k value for each pixel. Then, the second and third steps of adaptive LLE proceed in the same manner as traditional LLE.

If adaptive LLE is applied to a hyperspectral chip, as shown in Figure 3(b), then there is material separation in the dimensionality-reduced image. With target detection, however, the goal is to separate out a specific material. To do that, we have to add additional edges to our ANN graph in order to bias our adaptive LLE transformation for target detection. This is done by incorporating the multiple target spectra into our ANN graph, and enforcing additional connectivity on all pixels that are distance-1 or distance-2 (*i.e.*, one or two edges) away from the target cloud. These steps are illustrated in Figures 2(a) & 2(b). In building the edge set in this way, we are exploiting the properties of LLE that attempt to minimize local and global reconstruction errors. By intentionally overconnecting the target and its neighbors, we are forcing LLE to collapse those pixels away from the background. Adaptive LLE, biased for target detection of red felt material, is shown in Figure 3(c). As shown, the pixels corresponding to the red felt become more separated from the background pixels. It is in this lower-dimensional space that target detection is performed.
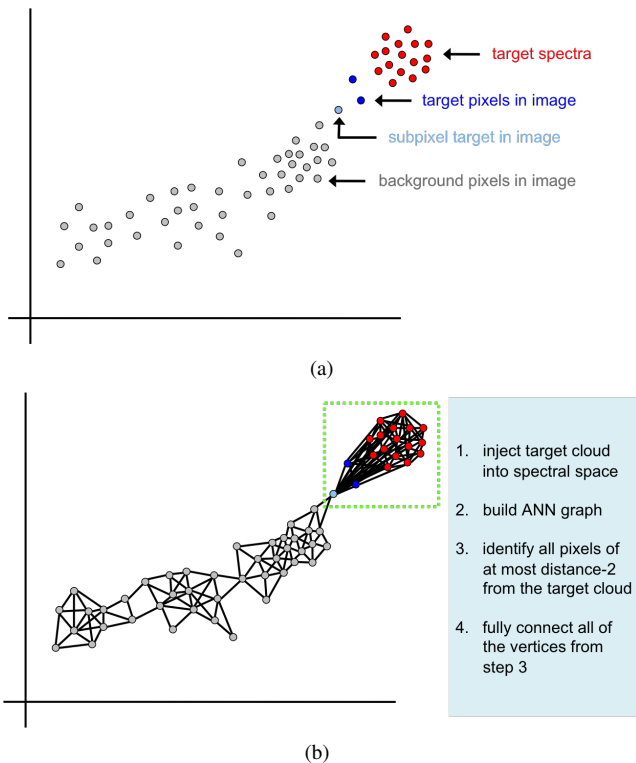

Fig. 2. These two panels illustrate the graph construction used in this methodology in order to bias the LLE transformation to better separate the targets from the background. (a) Pixel labels in the illustration. (b) Steps for building the edge set in the graph.
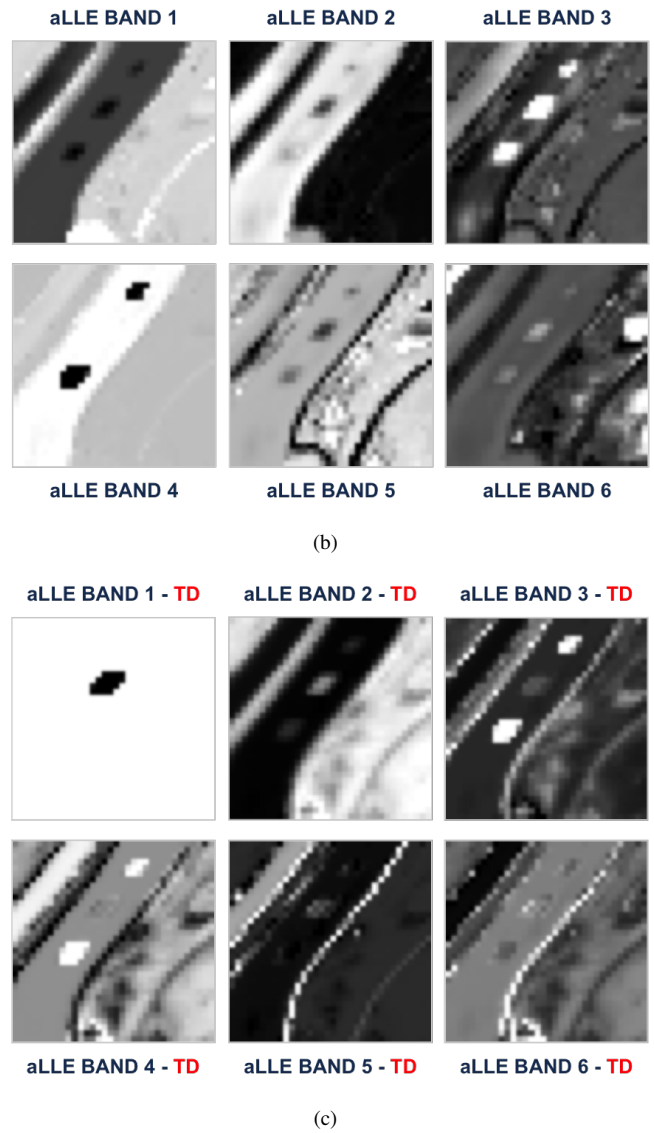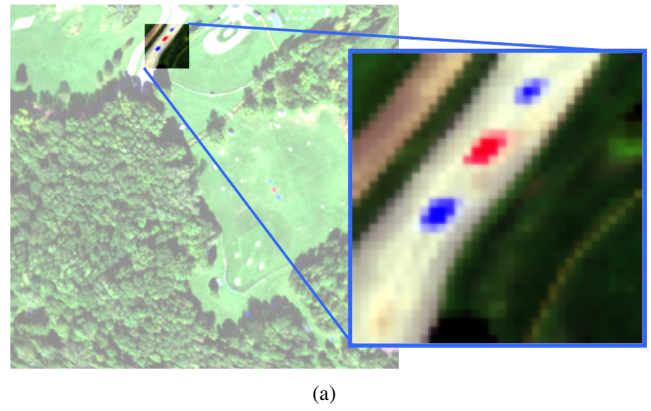

Fig. 3. (a) RGB showing the hyperspectral chip used in this example. This is from RIT's SHARE 2012 experimental campaign, and the specifications for the image are given in Section IV. (b) Image bands when nonlinear dimensionality reduction is performed on the chip using adaptive LLE. (c) Image bands when adaptive LLE is biased for target detection of the red felt.

Due to computational limitations, this is implemented on 20×20 pixel tiles in the image. The steps in this methodology, for each tile, are as follows:

Step 1: Inject target spectra into spectral space with tile data.

Step 2: Unit normalize the tile and targets (so that illumination does not factor into the adaptive nearest neighbors graph).

Step 3: Build ANN graph on tile + targets, and fully connect the target cloud, 1-neighbors, and 2-neighbors.

Step 4: Perform the second and third steps of adaptive LLE to estimate the manifolds.

Step 5: Keep the first m dimensions from adaptive LLE, where m is computed using the Gram Matrix dimensionality estimation approach described in [33], [34].

Step 6: Using the transformed target spectra, perform a rank-based SAM detector on the transformed tile pixels.

Step 7: Convert detection scores to Z-scores, so they can be appropriately compared across the different tiles.

The tile-based approach lends itself well to target detection in HSI because the ability to detect a target in one part of the scene should be independent of the ability to detect it in another part of the scene. For each tile, these steps result in a single-band detection map; they are then stitched together to result in an image-wide detection map.

IV. DATA

The data used here are from the Rochester Institute of Technology's SHARE 2012 experimental campaign, put on by their Digital Imaging and Remote Sensing Laboratory in September, 2012 [35]. The hyperspectral image was collected by the SpecTIR VS sensor, which is a VNIR-SWIR sensor with 360 spectral bands ranging from $0.4\mu m$ - $2.45\mu m$. The sensor had an approximate GSD (ground sample distance) of 1m, and after bad band removal, the image had 229 bands. The 170×280 pixel image is georectified as well as atmospherically compensated to approximate surface reflectance. As part of this experimental campaign, several red and blue felt cotton target panels (in 2m × 2m and 3m × 3m panel sizes) were placed in various illumination and occlusion conditions. The aerial image is shown in Figure 4(a), along with ground photos of some of the target locations. Figure 4(b) shows the two sets of spectra used in this analysis: 10 field-measured spectra for the red felt, and 10 field-measured spectra for the blue felt. These measurements were taken with an ASD spectrometer.

V. EXPERIMENT AND RESULTS

For this experiment, we implemented the methodology outlined in Section III using the 10 red field-measured spectra to detect the red felt targets in the image, and then using the 10 blue field-measured spectra to detect the blue felt targets in the image. We compared our detection results to subspace ACE (ssACE), which is widely recognized in the literature for hyperspectral subspace target detection [15]. The results are shown in Figures 5 & 6.

In these two figures, (a) shows the unthresholded detection map for LLE+SAM, (b) shows the detection map for
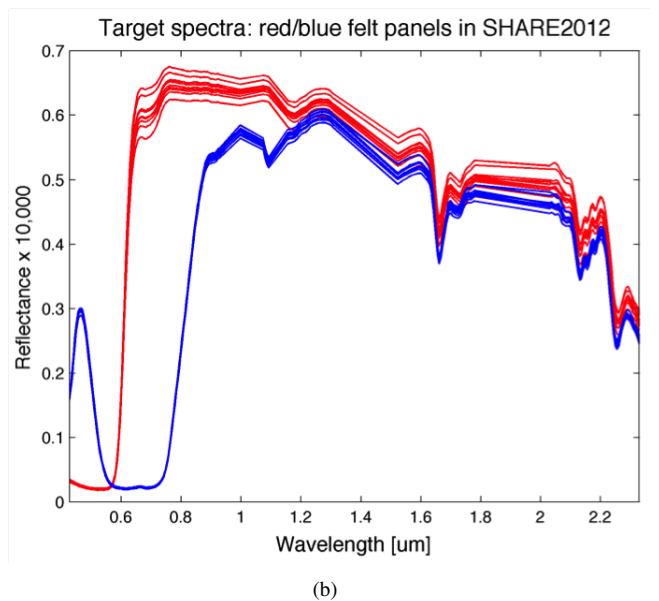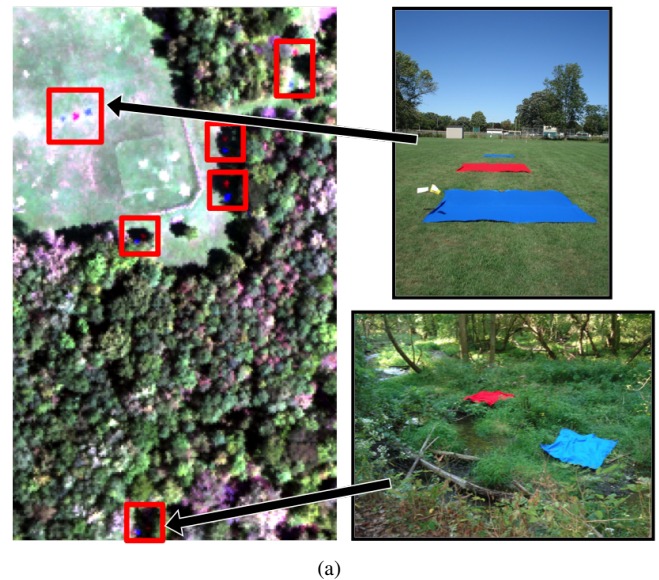

(a)

(b)

Fig. 4. (a) SHARE 2012 image with the red/blue felt target locations identified by the red squares. (b) Multiple spectral measurements, taken in the field, and corresponding to the red and blue felt materials.

LLE+SAM thresholded to the top 1% of scores, (c) shows the unthresholded detection map for ssACE, (d) shows the detection map for ssACE thresholded to the top 1% of scores, and (e) shows the corresponding ROC curves. The boxes in subfigures (b) and (d) show the locations of the targets in the image, and the green boxes indicate target locations that were accurately identified in the top 1% of scores. The red boxes indicate missed target locations within the top 1% of scores. Both approaches, for both target materials, show missed detections within the top 1%, which highlights how challenging some of these targets are (*i.e.*, the ones that are highly subpixel and in hard shadows). For both target materials, the LLE manifold approach with multiple targets outperforms ssACE when considering the ROC curves.
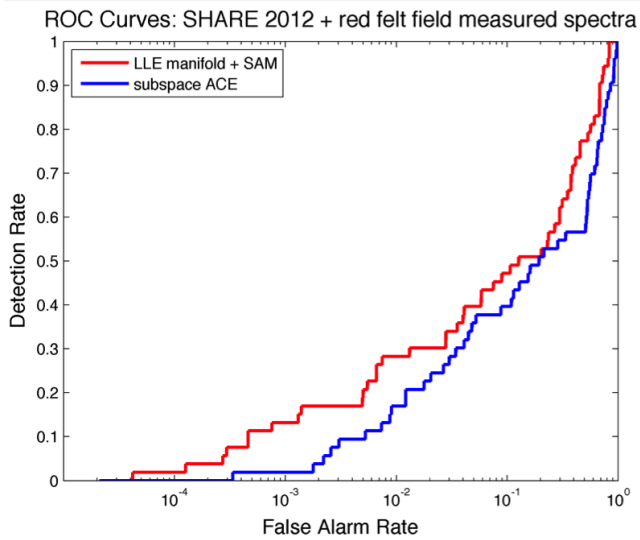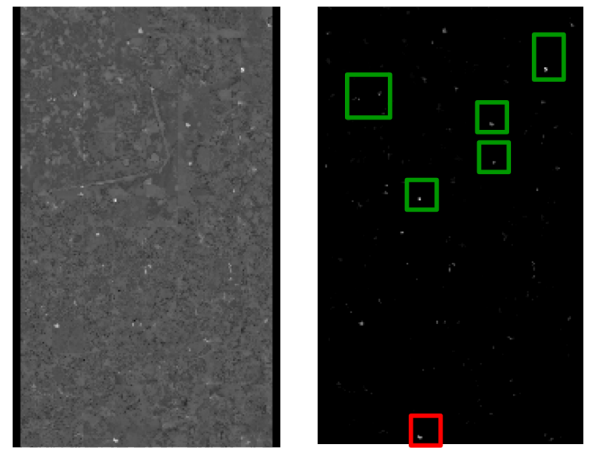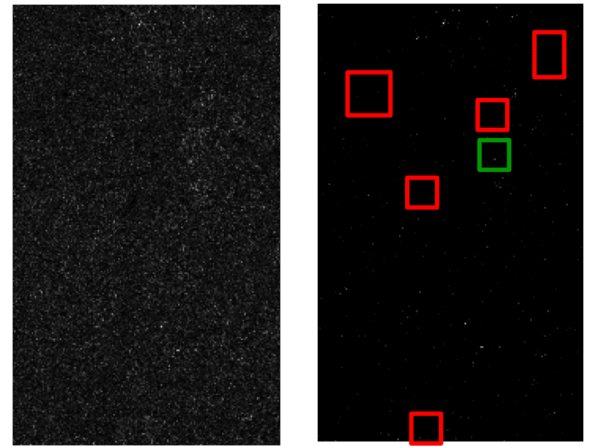
(a) (b)

(c) (d)
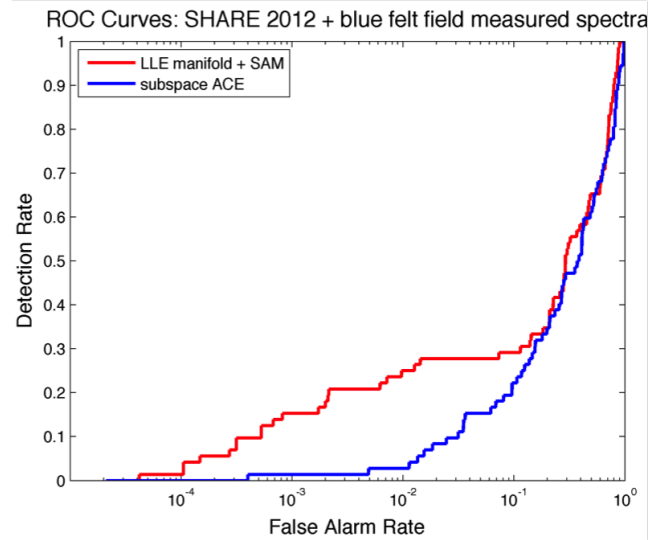
ROC Curves: SHARE 2012 + red felt field measured spectra
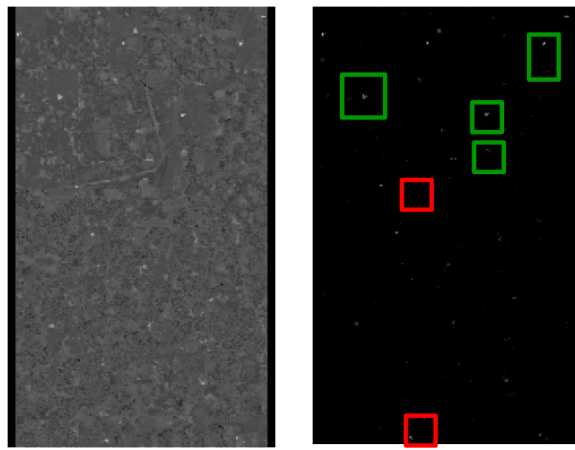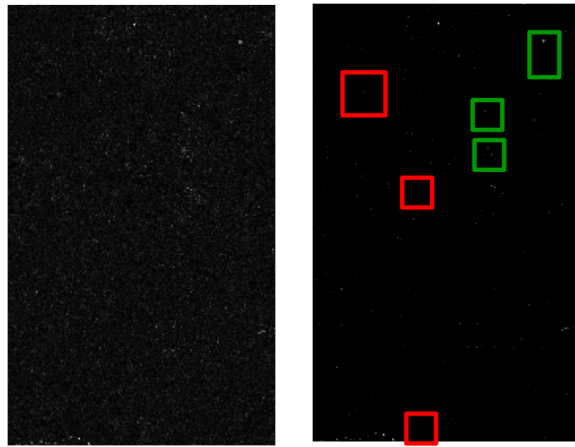
(e)

Fig. 5. Results for the red felt target spectra. (a) Adaptive LLE + SAM detection map, unthresholded. (b) Adaptive LLE + SAM detection map, displaying the top 1% of scores. The boxes show the locations of the targets; green indicates a detection in the top 1%, and red indicates a missed detection in the top 1%. (c) Subspace ACE detection map, unthresholded. (d) Subspace ACE detection map, displaying the top 1% of scores. (e) ROC curves for the two detection maps, plotted an a log scale to emphasize low false alarm rates.

(a) (b)

(c) (d)

ROC Curves: SHARE 2012 + blue felt field measured spectra

(e)

Fig. 6. Results for the blue felt target spectra. (a) Adaptive LLE + SAM detection map, unthresholded. (b) Adaptive LLE + SAM detection map, displaying the top 1% of scores. The boxes show the locations of the targets; green indicates a detection in the top 1%, and red indicates a missed detection in the top 1%. (c) Subspace ACE detection map, unthresholded. (d) Subspace ACE detection map, displaying the top 1% of scores. (e) ROC curves for the two detection maps, plotted an a log scale to emphasize low false alarm rates.

VI. Discussion and Future Work

We have presented a graph-based and manifold learning based approach to target detection in hyperspectral imagery when working with spectrally variable solid targets. This uses a different mathematical formulation than the traditional parametric and linear models used in HSI, and for the data presented here, outperforms the more commonly used ssACE subspace detection algorithm. This research explores alternative models and formulations because in complex, cluttered environments, we expect the structure of the data to be nonlinear and non-gaussian. We implement an adaptive approach to locally linear embedding, and bias the transformation - using multiple target spectra - to separate out the material described by the target spectra. It is in this lower-dimensional space that we perform target detection. Future work for this research would involve exploring the performance of different detectors in the transformed space when considering multiple spectrally variable targets, as well as potentially designing new detectors that are optimized for this target vs. background separation in this transformed space.

References

[1] A. Hayden, E. Niple, and B. Boyce, "Determination of trace-gas amounts in plumes by the use of orthogonal digital filtering of thermal-emission spectra," *Applied Optics*, vol. 35, no. 16, June 1996.

[2] J. Theiler, B. R. Foy, and A. M. Fraser, "Characterizing non-gaussian clutter and detecting weak gaseous plumes in hyperspectral imagery," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI*, vol. 5806. SPIE, June 2005.

[3] 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*, vol. 54, no. 15, pp. 4863–4875, May 2015.

[4] Y.-F. Su, T. L. Myers, C. S. Brauer, T. A. Blake, B. M. Forland, J. E. Szecsody, and T. J. Johnson, "Infrared reflectance spectra: effects of particle size, provenance and preparation," in *SPIE 9253*. October 2014.

[5] D. Manolakis, D. Marden, and G. Shaw, "Hyperspectral image processing for automatic target detection applications," *Lincoln Laboratory Journal*, vol. 14, no. 1, pp. 79 – 116, 2003.

[6] I. Reed, J. Mallett, and L. Brennan, "Rapid convergence rate in adaptive arrays," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-10, no. 6, pp. 853–863, November 1974.

[7] 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*, vol. 1. IEEE, November 1996, pp. 536–540.

[8] 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*, vol. 53, no. 2, Feb. 2005.

[9] F. Kruse, A. Lefkoff, J. Boardman, K. Heidebrecht, A. Shapiro, P. Barloon, and A. Goetz, "The spectral image processing system (SIPS) - interactive visualization and analysis of imaging spectrometer data," *Remote Sensing of Environment*, vol. 44, no. 2, pp. 145–163, 1993.

[10] J. C. Harsanyi and C. Chang, "Hyperspectral image classification and dimensionality reduction: an orthogonal subspace projection approach," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 32, no. 4, pp. 779–785, July 1994.

[11] D. Manolakis, "Taxonomy of detection algorithms for hyperspectral imaging applications," *Optical Engineering*, vol. 44, no. 6, June 2005.

[12] A. P. Schaum and R. Priest, "The affine matched filter," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XV*, vol. 7334. SPIE, April 2009.

[13] A. P. Schaum and B. J. Daniel, "Continuum fusion methods of spectral detection," *Optical Engineering*, vol. 51, no. 11, August 2012.

[14] J. Theiler, "Confusion and clairvoyance: some remarks on the composite hypothesis testing problem," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XVIII*, vol. 8390. SPIE, May 2012.

[15] D. G. Manolakis, R. Lockwood, T. Cooley, and J. Jacobson, "Is there a best hyperspectral detection algorithm?" in *SPIE 7334*. April 2009.

[16] S. Matteoli, M. Diani, and J. Theiler, "An overview of background modeling for detection of targets and anomalies in hyperspectral remotely sensed imagery," *JSTARS*, vol. 7, no. 6, pp. 2317–2336, June 2014.

[17] A. K. Ziemann, "A manifold learning approach to target detection in high-resolution hyperspectral imagery," Ph.D. dissertation, Rochester Institute of Technology, April 2015.

[18] A. K. Ziemann and D. W. Messinger, "An adaptive locally linear embedding manifold learning approach for hyperspectral target detection," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XXI*. SPIE, April 2015.

[19] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosophical Magazine*, vol. 2, no. 11, pp. 559–572, 1901.

[20] Y. Ma and Y. Fu, *Manifold Learning Theory and Applications*. CRC Press, 2011.

[21] B. Schölkopf, A. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Computation*, vol. 10, no. 5, pp. 1299–1319, July 1998.

[22] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, December 2000.

[23] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, December 2000.

[24] L. K. Saul and S. T. Roweis, "Think globally, fit locally: Unsupervised learning of low dimensional manifolds," *Journal of Machine Learning Research*, vol. 4, pp. 119–155, 2003.

[25] M. Belkin and P. Niyogi, "Laplacian eigenmaps and spectral techniques for embedding and clustering," *Advances in Neural Information Processing Systems*, vol. 14, 2001.

[26] C. M. Bachmann, T. L. Ainsworth, and R. A. Fusina, "Improved manifold coordinate representations of large-scale hyperspectral scenes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 10, pp. 2786–2803, October 2006.

[27] Y. Chen, M. M. Crawford, and J. Ghosh, "Applying nonlinear manifold learning to hyperspectral data for land cover classification," in *IGARSS*, vol. 5, July 2005, pp. 4311–4314.

[28] J. Chi and M. M. Crawford, "Selection of landmark points on nonlinear manifolds for spectral unmixing using local homogeneity," *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 4, pp. 711–715, 2013.

[29] D. Gillis, J. Bowles, G. M. Lamela, W. J. Rhea, C. M. Bachmann, M. Montes, and T. Ainsworth, "Manifold learning techniques for the analysis of hyperspectral ocean data," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI*, S. Shen, Ed., vol. 5806. SPIE, 2005.

[30] J. A. Albano, D. W. Messinger, and S. Rotman, "Commute time distance transformation applied to spectral imagery and its utilization in material clustering," *Optical Engineering*, vol. 51, no. 7, July 2012.

[31] D. West, *Introduction to Graph Theory*, 2nd ed. Prentice Hall, 2000.

[32] A. K. Ziemann, D. W. Messinger, and P. S. Wenger, "An adaptive k-nearest neighbor graph building technique with applications to hyperspectral imagery," in *Proceedings of the 2014 IEEE WNY Image Processing Workshop*. Rochester, NY: IEEE, November 2014.

[33] K. Canham, A. Schlamm, A. K. Ziemann, W. Basener, and D. W. Messinger, "Spatially adaptive hyperspectral unmixing," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 11, Nov. 2011.

[34] D. W. Messinger, A. K. Ziemann, B. Basener, and A. Schlamm, "Metrics of spectral image complexity with application to large area search," *Optical Engineering*, vol. 51, no. 3, March 2012.

[35] A. Giannandrea, N. Raqueno, D. W. Messinger, J. Faulring, J. P. Kerekes, J. van Aardt, K. Canham, S. Hagstrom, E. Ontiveros, A. Gerace, J. Kaufman, K. M. Vongsy, H. Griffith, B. D. Bartlett, E. Ientilucci, J. Meola, L. Scarff, and B. Daniels, "The SHARE 2012 data campaign," in *SPIE 8743*. April 2013.