Global, Local, and Stochastic Background Modeling for Target Detection in Mixed Pixels

Marin S. Halper*
The MITRE Corporation, 7515 Colshire Drive, McLean, VA, USA 22102

ABSTRACT

As hyperspectral sensors and exploitation methods have evolved, the accuracy of conventional background models has become a limiting factor for high confidence and low false alarm detection of mixed pixel targets. Many common target detection algorithms, such as the Adaptive Coherence/Cosine Estimator, implicitly use a global background model that assumes the background can be modeled by a single, multivariate Gaussian random variable with additive independent and identically distributed Gaussian noise. In order to improve the accuracy of the Gaussianity assumptions, a local background model, which models the background as multiple, disjoint Gaussian clusters, has also been widely considered. This paper introduces an improved variant of the local background model, as well as a novel stochastic background model that is free from distributional assumptions and accounts for the spectral variability of the background on a pixel-by-pixel basis. The performance of the global, local, and stochastic background models is evaluated on a controlled data set and the tradeoffs associated with each method are discussed.

Keywords: hyperspectral, background modeling, covariance, clustering, non-Gaussian, stochastic, target detection

1. INTRODUCTION

Hyperspectral sensing is a passive remote sensing technique that measures electromagnetic radiation as a nearly continuous function of wavelength for a given spatial area. This type of measurement results in a spectral signature that represents a physical quantity linked by the underlying physics and chemistry to a particular material or mixture of materials within the given spatial area. By comparing this measured spectral signature to a library of reference signatures derived from laboratory or field measurements, one can uniquely detect and identify a wide range of materials. Although hyperspectral remote sensing is still an evolving technique, it has enjoyed successes across a wide range of domains including geology, forestry and agriculture, environmental monitoring, astronomy, as well as defense and intelligence.

Hyperspectral imaging sensors collect data as a three-dimensional cube. The x and y dimensions of the cube represent spatial pixels, much like a conventional imager, while the z dimension represents spectral information. Therefore, each pixel in the cube represents a single spectral signature. This enables the non-literal, i.e., spectral- and not spatial-based, exploitation of hyperspectral data to derive physical and compositional information. Target detection is one of the most common objectives of the non-literal exploitation of hyperspectral data. Target detection is performed by comparing each spectral signature in the data to a target reference signature in order to estimate the likelihood that each spatial pixel contains the target of interest. If a high abundance of the target, relative to the background, occupies a pixel, it is considered a full pixel or a resolved target. If a low abundance of the target, relative to the background, occupies a pixel, it is considered a mixed pixel, sub-pixel target, or unresolved target. In general, target detection in full pixels is a simpler and more accurate process than target detection in mixed pixels and is often the primary focus for hyperspectral exploitation. However, accurate target detection in mixed pixels represents a critical need within the spectral community.

There exist many scenarios that require accurate target detection in mixed pixels. These include the detection of relatively small targets such as vehicles and aircraft wreckage, naturally mixed pixel targets such as minerals and crop species, layered targets such as gas plumes and residues, as well as targets in data sets with a large ground sampling distance (GSD). While target detection in mixed pixels has become an increasingly ubiquitous challenge area across problem sets and domains, the performance of the associated conventional algorithms has remained limited by incorrect assumptions and a resulting inability to discriminate subtle signatures from noise[1-3]. In particular, the accuracy of conventional background models has become a limiting factor for high confidence and low false alarm detection of

* Electronic address: mhalper@mitre.org

mixed pixel targets[1-3]. Many conventional algorithms, such as the Adaptive Coherence/Cosine Estimator (ACE), implicitly use a global background model that assumes that the background can be modeled by a single, multivariate Gaussian random variable with additive independent and identically distributed Gaussian noise[4-5]. While this model is generally acknowledged to be insufficient[1-3,6-16], there has been limited effort to address the shortcomings. Much of the effort has focused on improving the multivariate Gaussianity and accuracy of the background model by removing the statistical contribution of likely targets and anomalies[2-3,14-15], utilizing spatially local or spectrally local background segmentation[1-3,6], developing Gaussian mixture models[8-10,17], or exploring non-Gaussian distributions within the family of elliptically contoured distributions[9-13].

This paper focuses on two novel background models for accurate target detection in mixed pixels. The first background model, the local background model, is an improved variant within the class of spectrally local background segmentation approaches. The local background model uses specialized spectral clustering and segmentation in order to construct multiple disjoint background clusters and a set of unassigned pixels. Each member of a background cluster uses the associated cluster statistics for background modeling, while the remaining unassigned pixels use a variant of the global statistics. Although the local background model employs Gaussian assumptions, it uses local modeling to improve the multivariate Gaussianity and accuracy of the model. This paper also introduces a second background model, the stochastic background model, which represents the first in a new class of approaches based on stochastic optimization. In this approach, each pixel is treated as an independent stochastic optimization process based on the conventional linear mixing model, but with the target abundance, background abundance, and background signature as free parameters. The stochastic background model leverages more accurate, distribution-free assumptions that explicitly model the background and noise as unknown random processes. In addition, this paper will explore the conventional global background model that models the entire scene as a single Gaussian cluster, as well as an improved variant of the global background model that uses spatial masks to reduce the contribution of likely targets and anomalies to the background statistics.

All four of these background models can be paired, with slight modifications, to existing target detection algorithms in order to evaluate performance. ACE is a target detection algorithm that is particularly well-suited for this pairing, as it is the generalized likelihood ratio test for the linear mixing with replacement model designed for target detection in mixed pixels[16,18]. In addition, ACE is generalized for all elliptically contoured background distributions and is therefore more resistant to non-Gaussian background distributions than many other conventional algorithms[13,16]. For these reasons, the four background models will be discussed and evaluated in the broader context of their implications on ACE target detection performance.

Following this introduction, Section 2 provides more detail on the global, global masked, local, and stochastic ACE algorithms. Section 3 describes the methods used to evaluate the detection performance of the four algorithms, as well as the Rochester Institute of Technology (RIT) Target Detection Blind Test data sets used in this evaluation. The quantitative performance results, as well as relevant qualitative results, are presented in Section 4 and discussed in Section 5. Lastly, the paper concludes with some reflection on the lessons learned and potential next steps in Section 6.

2. ALGORITHMS

2.1 Global ACE

The global ACE algorithm uses the global scene mean and covariance to model the background as a single multivariate Gaussian distribution. The associated detection formula is given by

$$D_{ACE}(x) = \frac{Sign((s-\mu_b)^{\mathrm{T}}\Sigma_{\mathrm{b}}{}^{-1}(x-\mu_b))\left[(s-\mu_b)^{\mathrm{T}}\Sigma_{\mathrm{b}}{}^{-1}(x-\mu_b)\right]^2}{\left[(s-\mu_b)^{\mathrm{T}}\Sigma_{\mathrm{b}}{}^{-1}(s-\mu_b)\right]\left[(x-\mu_b)^{\mathrm{T}}\Sigma_{\mathrm{b}}{}^{-1}(x-\mu_b)\right]}, \quad (1)$$

where s is the target reference spectral signature, x is the test pixel signature, μ_b is the background mean signature, and Σ_b is the background covariance. Implicit in this formula is the transformation to a feature space defined by the background mean and covariance. This transformation is referred to as whitening and the associated feature space is referred to as whitened space[4-5]. Whitening uses the background mean and covariance to transform the background to a zero-mean distribution with a covariance matrix equal to the identity matrix. This results in a background model in which each dimension or band in whitened space is independent and identically Gaussian distributed with zero mean and unit variance, i.e., equivalent to white Gaussian noise. The whitening transformation is designed to suppress background

interference and maximize the target-to-background contrast under the Gaussian assumptions. After applying this transformation, ACE is equivalent to the cosine squared of the angle in whitened space between the target reference spectral signature *s* and test pixel signature *x*. Note that the *sign* function is included solely to distinguish target reference and test pixel signatures pointed in the same direction in whitened space from target reference and test pixel signatures pointed in opposite directions.

While Global ACE is relatively simple and fast, it is the least theoretically robust of the four ACE algorithms discussed in this paper. In modeling the background as a single multivariate distribution, global ACE is subject to biases resulting from the inclusion of targets and anomalies into the background statistics[2-3,14-16,19-21], the incorrect modeling of multi-modal background classes[1-3,6,8-13,17], and other sources of non-Gaussianity in the background model. In addition, Global ACE is subject to biases resulting from the lack of independence, homoscedasticity, or Gaussianity in the noise[4,22].

2.2 Global Masked ACE

The global masked ACE algorithm improves the robustness of global ACE by using two spatial masks to remove the contribution of the likely targets and anomalies, respectively, from the background statistics. In order to derive the spatial mask for likely targets, a global ACE detection score is first calculated, using Equation 1, for each pixel in the data. Then, the pixels corresponding to the top *n* percent of scores are masked. In order to derive the spatial mask for likely anomalies, a global RX detection score is first calculated for each pixel in the data[23]:

$$D_{RX}(x) = (x - \mu_b)^{\mathrm{T}} \Sigma_{\mathrm{b}}^{-1} (x - \mu_b) . \quad (2)$$

Then, the pixels corresponding to the top *m* percent of scores are masked.

After applying the two spatial masks, the remaining pixels are used to calculate a new background mean signature μ_m and background covariance Σ_m, where the subscript *m* denotes the use of a mask to calculate the background statistics. These new background statistics are then used in Equation 1 to calculate the global masked ACE score for each test pixel.

The inclusion of target pixels in the background model has been shown to decrease target-to-background contrast and reduce detection performance[2-3,16,19-21]. Similarly, the inclusion of anomalous pixels in the background model has been shown to skew the background statistics in potentially undesirable and non-Gaussian ways, particularly if the anomalies are related to targets, sensor artifacts, or non-background materials[14-15]. The global masked ACE algorithm attempts to mitigate these issues, as well as improve the accuracy and robustness of global ACE, at the expense of a slower runtime and two thresholding parameters.

2.3 Local ACE

The local ACE algorithm uses a novel spectral clustering and segmentation approach in order to develop multiple disjoint background clusters and a set of unassigned pixels. Each member of a background cluster uses the associated cluster statistics for background modeling, while the remaining unassigned pixels use the global masked statistics, μ_m and Σ_m, discussed in Section 2.2.

Like ACE, the clustering algorithm uses the spectral angle between vectors in whitened space as the fundamental distance metric. This enables the algorithm to cluster pixels based on background class, regardless of background abundance, which is particularly important for mixed pixels. The clustering algorithm begins by whitening the data using the global masked statistics and then truncating all bands in whitened space after a user-supplied threshold *t*. Given that the whitening transformation results in whitened bands that are ordered from highest to lowest contribution to scene variance, the truncation step is designed to create a subspace that is likely dominated by background signal[24]. The cluster labeling process begins by labeling the first pixel in whitened space as the first exemplar. If the second pixel is within a user-defined θ degrees from the first exemplar, then it is clustered with the first exemplar, otherwise it is labeled as the second exemplar. This process is repeated for each pixel in the hyperspectral cube until all pixels are clustered with an exemplar. Given this set of cluster assignments, the original exemplars are replaced with the corresponding cluster means and the labeling process is repeated in order to refine the cluster assignments. After this second pass, if any member of a small cluster, defined as a cluster containing fewer members than ten times the number of spectral bands, is within θ degrees of the mean of a large cluster, it is reassigned to that cluster. This reassignment is designed to reduce the number of pixels that belong to clusters that are too small for relatively robust covariance estimation[16,25]. Lastly, the scene is segmented according to the final set of cluster assignments.

The local ACE algorithm uses the clustering results to construct several disjoint background models, one background model for each large cluster. Members of small clusters that were not reassigned to larger clusters during the clustering process are inferred to lack a robust local background model and therefore instead use the global masked background model. As such, the local ACE detection scores for those pixels are equal to the global masked ACE scores. Prior to calculating the background models for the pixels belonging to larger clusters, the likely target mask described in Section 2.2 is applied to the data. Next, the non-target-like members of each cluster are used to create a cluster mean signature μ_i and cluster covariance Σ_i, where the subscript i denotes cluster i. Each pixel is then whitened, without mean subtraction, using its corresponding cluster covariance. In addition, the target reference and cluster mean signatures are also whitened, without mean subtraction, using the same cluster covariance. The background abundance for each pixel is then estimated by performing linear least-squares regression using the first t bands in whitened space, inferred to represent the background subspace, in order to solve the linear mixing equation

$$\Sigma_{\mathrm{i}}^{-1/2} x = \alpha\left(\Sigma_{\mathrm{i}}^{-1/2} s\right) + \beta\left(\Sigma_{\mathrm{i}}^{-1/2} \mu_i\right) + w, \tag{3}$$

which describes a pixel in whitened space as a linear combination of target and background signatures in whitened space, with abundances α and β, respectively, and additive Gaussian white noise w. After estimating those abundances, the local ACE detection score is computed as

$$D_{ACE}(x) = \frac{Sign\left(s^{\mathrm{T}}\Sigma_{\mathrm{i}}^{-1}(x-\beta\mu_i)\right)\left[s^{\mathrm{T}}\Sigma_{\mathrm{i}}^{-1}(x-\beta\mu_i)\right]^2}{\left[s^{\mathrm{T}}\Sigma_{\mathrm{i}}^{-1}(s)\right]\left[(x-\beta\mu_i)^{\mathrm{T}}\Sigma_{\mathrm{i}}^{-1}(x-\beta\mu_i)\right]}. \tag{4}$$

In contrast to conventional ACE algorithms, linear regression in whitened space is used to suppress the contribution of the background signature in the test pixel in order to obtain a residual signature in whitened space. This residual signature can then be directly compared to the target signature in whitened space. In addition, this residual signature can be transformed back to spectral space using the inverse of the whitening transformation. This enables the visualization and analysis of background suppressed residual signatures in spectral space to assess the confidence of detection results[26-27]. Note that the global ACE algorithms do not have an analog of background suppressed residual spectra as they do not model the background abundances.

In general, an individual cluster of homogenous background materials is more likely to obey the multivariate Gaussian background assumptions than the global scene[1-3,6]. By induction, a set of disjoint clusters is more likely to obey those assumptions than the global scene. This logic represents the motivation for local ACE, which builds upon global masked ACE to further improve the accuracy and robustness of the background model. In addition, unlike the global ACE algorithms, local ACE uses the estimated background abundance to further suppress the background signal on a pixel-by-pixel basis. While local ACE generally obeys the underlying assumptions more closely than the global ACE algorithms, it does so at the expense of a slower runtime and the selection of two additional parameters, cluster angle θ and background bands threshold t. The possible sensitivity of local ACE performance to those parameters represents a source of variability that could potentially offset some of the additional accuracy and robustness of the local ACE algorithm.

2.4 Stochastic ACE

Compared to the global and local ACE algorithms, stochastic ACE is fundamentally different in its assumptions and approach to background modeling. In particular, stochastic ACE does not assume that the variability associated with the background and noise can be adequately modeled by multivariate Gaussian distributions. Instead, stochastic ACE uses a distribution-free stochastic optimization technique to model that variability. This approach assumes only that the noise is a zero-mean random process, i.e., that the data do not exhibit a systematic bias, and that the noise is independent of the abundance and background signature. This directly contrasts to the strong assumptions underlying conventional Gaussian-based algorithms, e.g., the assumptions that the background is normally distributed and that noise is normally distributed, equal variance, and uncorrelated across spectral bands[4,22]. As such, stochastic ACE is able to more closely obey the underlying assumptions of the background model.

At the heart of Stochastic ACE is the conventional linear mixing model for binary mixtures $x = \alpha s + \beta\mu_b + w$, where, as in Equation (3), x is the test pixel signature, s is the target reference spectral signature, α and β are the corresponding abundances, and w is an additive noise vector[4-5,22,28]. In the conventional model, x and s are known, α and β are unknown, and the distribution of w is known. Therefore, the linear mixing model can be solved using linear least-squares regression to estimate the abundances of target and background in the text pixel[22,28]. In practice, this linear unmixing is

often performed in whitened space, as shown in Equation (3), in order to reduce the bias associated with the correlation of spectral bands and obtain more accurate abundance estimates[28]. In addition, non-negativity constraints are often incorporated to ensure that the estimated abundances are physically realizable[5,28].

Stochastic ACE frees the linear unmixing process from the strong assumptions, mentioned above, that are required by linear least-squares regression for the unbiased estimation of abundances. Furthermore, stochastic ACE does not assume that the background signature is strictly known, but instead constrained within some adaptive volume that is unique for each pixel. Stochastic ACE uses stochastic optimization, instead of linear least-squares regression, to unmix the test pixel and estimate the abundances[29]. This optimization process is guided by the loss function to be minimized

$$L(b, \alpha, \beta) = E\left[(x - \alpha s - \beta b)^{\mathrm{T}} {\Sigma_{\mathrm{b}}}^{-1} (x - \alpha s - \beta b) + \gamma(\mu_{\mathrm{b}} - b)^{\mathrm{T}} {\Sigma_{\mathrm{b}}}^{-1} (\mu_{\mathrm{b}} - b)\right], \tag{5}$$

where b is the unknown background signature, γ is a given penalty parameter, μ_{b} is the global masked background mean as described in Section 2.2, Σ_{b} is the global masked background covariance, and E denotes the expectation over all sources of randomness embodied in w. Implicit in this loss function is the use of the whitening transformation. While this transformation is not required for stochastic optimization, it is useful for ensuring that the magnitude of the parameters are relatively equal, which can improve convergence[29]. Note that the loss function is the sum of two distinct quantities. The first quantity represents the variance of the residuals associated with the error between the fit of the linear mixing model and the test pixel. This quantity is equivalent to the loss function for conventional linear least-squares regression in whitened space. The second quantity represents the variance of the residuals associated with the difference between the estimated background signature b and the initial background signature μ_{b}, weighted by the penalty parameter γ. Therefore, the minimization of this loss function represents a tradeoff between reducing the variance of the residual error for the unmixing model and increasing the variance of the background signature. This tradeoff enables the background to vary within a volume, centered at the initial background signature and constrained by the penalty parameter, in order to improve the fit between the linear mixing model and the data.

As the loss function $L(b, \alpha, \beta)$ includes x and therefore w, an underlying random process, the gradient cannot be directly computed. However, given the assumption that w is zero-mean and independent of the estimated parameters, one can derive the stochastic gradient and use the stochastic gradient form of stochastic approximation to minimize the loss function[29]. As w is constant throughout the minimization process, the process is directly analogous to gradient descent[29]. For stochastic ACE, the initial conditions for the gradient descent are given by $b = \mu_{\mathrm{b}}$, $\alpha = 0$, and $\beta = 1$. This represents the null hypothesis that the target is absent in the test pixel. A simple line search is used to determine the step size required for descent. In order to ensure physical constraints, if a descent step results in negative values for α or β, the values are replaced with a value of zero. The minimization process terminates after reaching a user-supplied number of iterations. The final values of α and β represent estimates for the target and background abundances, respectively, while the value for b represents the estimated background signature associated with the test pixel. While convergence at the global minimum is not guaranteed, the final value of the loss function is guaranteed to be less than or equal to the initial value corresponding to the null hypothesis.

Using these estimated values of b, α, β, the stochastic ACE detection score is computed as

$$D_{ACE}(x) = \frac{\left[s^{\mathrm{T}} {\Sigma_{\mathrm{i}}}^{-1} (x - \beta b)\right]^2}{\left[s^{\mathrm{T}} {\Sigma_{\mathrm{i}}}^{-1} (s)\right]\left[(x - \beta b)^{\mathrm{T}} {\Sigma_{\mathrm{i}}}^{-1} (x - \beta b)\right]}. \tag{6}$$

As β and b are free parameters, it is possible for the quantity $(x - \beta b)$ to be exactly equal to zero. In this case, the resulting ACE score is also exactly equal to zero. Stochastic ACE uses stochastic optimization to suppress the contribution of the background signature in the test pixel in order to obtain a whitened residual signature. As in local ACE, this residual signature can be directly compared to the target signature in whitened space and transformed back to spectral space using the inverse of the whitening transformation. In addition to using these residual signatures for confidence building, as described in Section 2.3, these signatures can also be used for false alarm mitigation. By default, stochastic ACE flags and masks detections in which the resulting residual signature is not physically meaningful, e.g., not correlated above a user-supplied correlation value R with either the target or the original pixel.

Stochastic ACE is a novel method that represents a departure from the conventional assumptions and methods. As it explicitly models background and noise variability on a pixel-by-pixel basis without distributional assumptions, stochastic ACE is more robust in its assumptions and may theoretically be more accurate than conventional methods. Like local ACE, this benefit comes at the expense of a slower runtime and additional parameters, e.g., the penalty

parameter and number of iterations. While stochastic ACE uses the global masked mean and covariance as part of the initial conditions, the use of a parametric background model is not required nor assumed. Other selections for initial conditions could be explored as part of future research. The penalty parameter directly controls the volume in which the background signature is able to vary. If the value of the parameter is too small, then background variability may suppress the contributions of subtle, correlated targets. If the value of the parameter is too large, then the background will not be able to vary at all and loss function will be equivalent to that of conventional linear least-squares regression. For this reason, the selection and sensitivity of the penalty parameter is an important consideration for stochastic ACE. This issue is discussed further in Section 5. Note that the whitening transformation standardizes the variance of the background distribution, before the penalty parameter is introduced in the loss function. As such, the behavior of penalty parameters may generalize across data sets.

3. METHODS

3.1 Data Sets and Targets

Two reflectance data sets, both provided as part of the RIT Target Detection Blind Test project[30], were used to evaluate the four ACE target detection algorithms. The first data set, referred to as Self Test, is packaged with complete ground truth information. This includes target spectral signatures from laboratory measurements, photographs, and locations. The second data set, referred to as Blind Test, lacks ground truth target locations. The two data sets were collected over Cooke City, Montana and environs in July 2006 by the HyMap[31] visible near infrared and shortwave infrared (VNIR/SWIR) airborne sensor. The spectral range of the HyMap sensor spans from 0.45 to 2.5 micrometers and consists of 126 spectral bands; however, for this analysis, bands 63-65 and 125-126 were removed to improve the signal to noise ratio. For both data sets, the approximate GSD is 3 meters (m).

There are five significantly sub-pixel targets associated with the data sets, two in Self Test and three in Blind Test, without replication. The two sub-pixel targets in Self Test are labeled F3 and F4, corresponding to 1 m^2 blue cotton and red nylon cloth targets, respectively. The three sub-pixel targets in Blind Test are labeled F5, F6, and F7, corresponding to 1 m^2 maroon nylon, gray nylon, and green cotton cloth targets, respectively. Given the approximate GSD of 3 m, the approximate pixel fill factor of each target in its associated mixed pixel is bounded between 2.7% and 11.2%, i.e., an abundance between 0.027 and 0.112. Because of their relatively low abundances, these five sub-pixel targets were selected to evaluate the performance of the four target detection algorithms. Note that, in addition to these five sub-pixel targets, the data set contains five larger, 4 m^2 cloth panels made from the same materials. For the purpose of this analysis, these five panels are referred to as F3-type, F4-type, F5-type, F6-type, and F7-type panels, respectively.

3.2 Algorithm Parameters

The detection algorithms use several user-input parameters to adjust behavior. Global ACE requires no parameters. Global masked ACE masked the top 1% of anomalies and 0.01% of targets. Local and stochastic ACE employed the same masking percentages. In addition, local ACE used an angle of 70 degrees to cluster the data in whitened space and the first 6 bands in whitened space to form the background subspace. Stochastic ACE used a penalty parameter of 6 and performed 100 iterations during the optimization process. In addition, Stochastic ACE replaced with zero all detections in which the associated background suppressed residual spectrum is neither correlated with the original pixel signature nor the target signature above a value of 0.9. These parameters were empirically selected by quick look qualitative evaluation and not rigorous quantitative optimization.

3.3 Scoring and Visualization

In order to assess the relative detection performance, the number of false alarms was measured and recorded for each combination of detection algorithm and target. For example, given a detection map, the number of false alarms for the F3 target in Self Test was calculated by first masking all pixels associated with F3, except for the pixel associated with the 1 m^2 F3 panel, and then tallying the number of remaining pixels with ACE scores above that of the F3 pixel. This process was repeated for the F4 target in Self Test. The process to calculate the number of false alarms for the F5, F6, and F7 targets in Blind Test was slightly different. For example, given a detection map, the number of false alarms for the F5 target in Self Test was calculated by first masking all pixels directly surrounding the highest scoring pixel, inferred and later confirmed to be associated with the 4 m^2 F5-type panel, and then uploading the masked detection map to the Target Detection Blind Test site[32] for scoring. As this masked detection map still included the two pixels related to

F5, the final number of false alarms was obtained by subtracting two from the score reported by the site. This process was repeated for F6 and F7.

The visualization of spectral signatures can be used to build confidence in detections and mitigate potential false alarms. Generally, this process can be complicated or even impossible for targets in mixed pixels, as the contribution of the target signal can be negligible relative to that of the background signal. However, as described in Sections 2.3 and 2.4, local and stochastic ACE are both capable of deriving background suppressed residual spectra in whitened space and then transforming the spectra back to spectral space for visual inspection. These background suppressed residual spectra were qualitatively used to evaluate the confidence in select detection results in order to demonstrate the utility of using spectroscopy to complement statistical detection methods.

4. RESULTS

The quantitative performance of the global, global masked, local, and stochastic ACE algorithms on targets F3, F4, F5, F6, and F7 is reported in Table 1. This performance is measured by the number of false alarms for each combination of algorithm and target, given a probability of detection equal to one. The average number of false alarms and average rank, as well as the corresponding standard deviations, are also reported. From Table 1, it can be seen that local ACE performed the best overall, with stochastic ACE a relatively close second. Global masked ACE algorithm performed worse than local and stochastic ACE, but better than global ACE, which performed the worst overall. It is also important to note that the global ACE algorithm exhibited a substantially larger standard deviation in number of false alarms than the other three algorithms.

In order to evaluate the qualitative performance of the ACE algorithms, a complete set of results for the F4 target is provided in Figures 1, 2, and 3. This set of results is representative across all five targets. The RGB color composite for Self Test, as well as the target and anomaly mask and segmentation images associated with F4, are provided in Figure 1. As a reference, a red square is centered over the location of the F4 target. Note that, as intended, the mask image appears to discriminate anomalous materials from large background classes and the segmentation image appears to segment the scene into multiple disjoint background classes. The four detection maps for F4, which are histogram equalized to have an approximately equal linear stretch, are provided in Figure 2. As in Figure 1, a red square is centered over the location of the F4 target. In particular, note that approximately 96.5% of the pixels in the stochastic ACE detection plane have scores equal to zero and the remaining scores appear qualitatively uncorrelated with the other three detection planes. Lastly, in Figure 3, the suppressed residual spectra for the pixel associated with F4 are plotted against the original in-scene pixel and measured target signatures for F4 in both absolute (a) and normalized reflectance space (b). Normalized reflectance is obtained by dividing each spectral signature by its vector magnitude. In addition, Figure 3 includes an equivalent plot in normalized reflectance space for the pixel associated with the 4 m^2 F4-type panel. This third plot (c), which is based on a panel with higher abundance than F4, is included to facilitate an evaluation of the utility of spectral visualization, as otherwise it may be difficult to distinguish the subtle features in the relatively small plots. The key feature of the plots in Figure 3 is that, while there is relatively little to no increase in reflectance between 580 and 640 nanometers (nm) in the original pixel signatures, there is an increase in that region for the background suppressed residual spectra. The increase in reflectance in this spectral region can be directly correlated with the target signature, which provides some justification and confidence for a detection result in that pixel. In addition, other subtle features in the background suppressed residual spectra can be directly correlated with subtle features in the target signature. This is especially true in the third plot (c), e.g., the features in the local background suppressed residual and target signatures at approximately 2048 and 2300 nm, but significantly less obvious in the second plot (b).

Table 1. Detection performance, as measured by the number of false alarms, for each combination of algorithm and target. The average number of false alarms and average rank across all targets, as well as corresponding standard deviations, σ, are also reported for each algorithm.

Algorithm	*F3*	*F4*	*F5*	*F6*	*F7*	*Average False Alarms (σ)*	*Average Rank (σ)*
Global ACE	739	107	9	3	2062	584.0 (880.9)	3.2 (1.3)
Global Masked ACE	160	179	3	7	52	80.2 (84.0)	3.0 (1.2)
Local ACE	3	15	4	3	5	6.0 (5.1)	1.2 (0.4)
Stochastic ACE	24	27	4	5	32	18.4 (13.0)	2.2 (0.4)

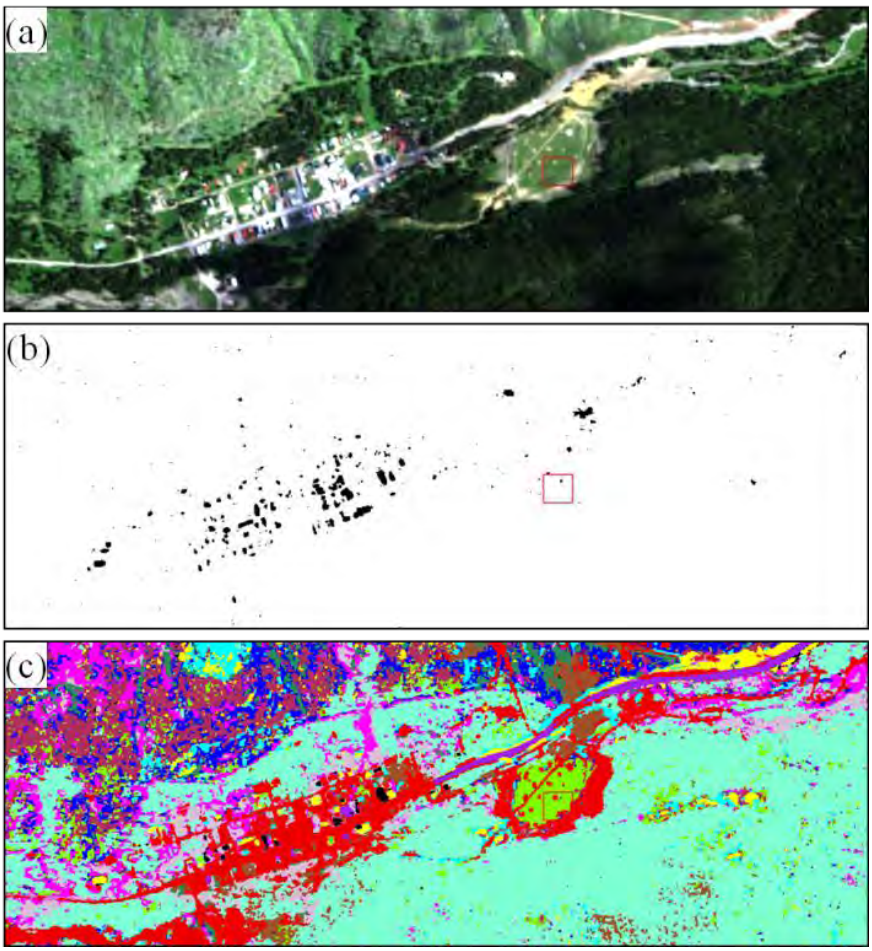

Figure 1. (a) RGB color composite of the Self Test data set, (b) the target and anomaly mask image associated with F4, and (c) the background segmentation image associated with F4. The red square is centered over the location of the F4 target.

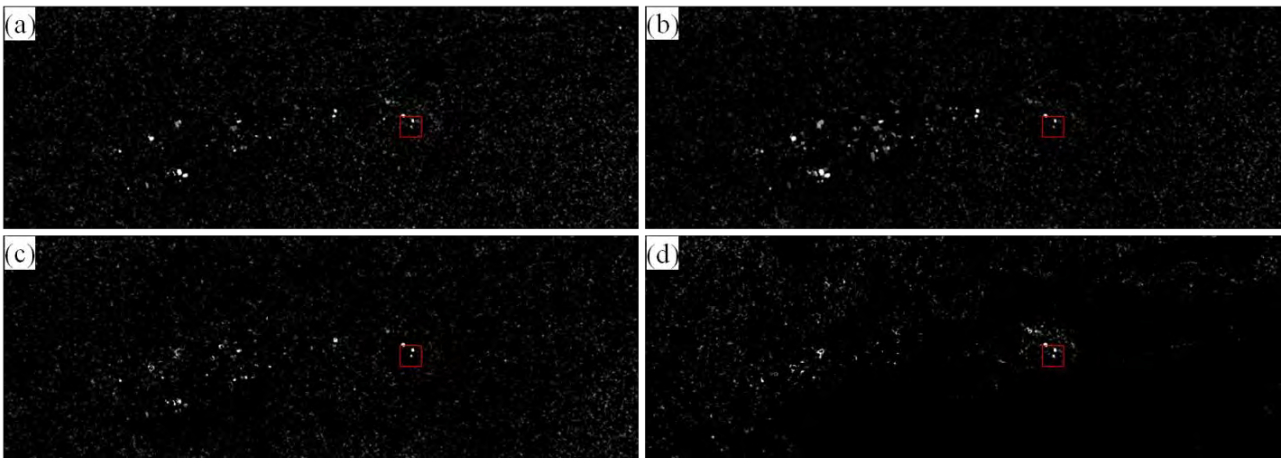

Figure 2. Four detection maps for F4, (a) the detection map for global ACE, (b) global masked ACE, (c) local ACE, and (d) stochastic ACE. The red square is centered over the location of the F4 target. For qualitative comparison, the detection maps have been histogram equalized to have an approximately equal linear stretch. Note that the detection plane shown in (d) appears relatively uncorrelated with the detection planes shown in (a), (b), and (c), respectively.

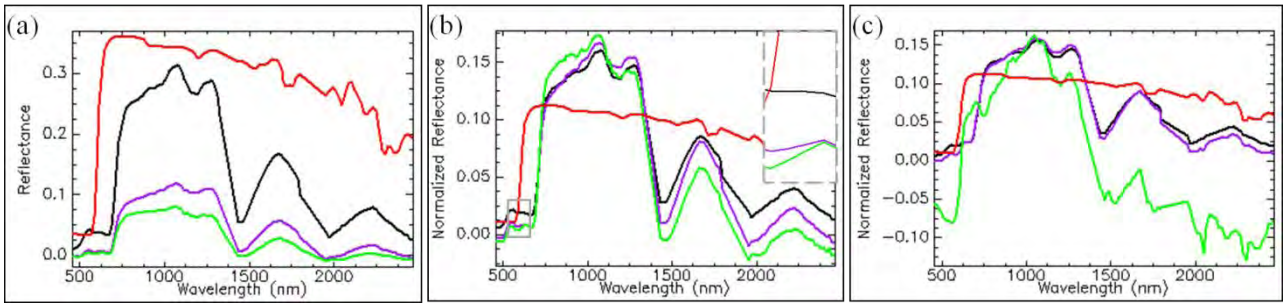

Figure 3. Three spectral plots of the F4 target signature in red, the associated in-scene pixel in black, the background suppressed residual spectrum for local ACE in green, and the background suppressed residual spectrum for stochastic ACE in purple. (a) The first plot shows the signatures in reflectance, (b) the second plot shows the signatures in normalized reflectance with an inset for the region between 580 and 640nm, and (c) the third plot shows the signatures for a 4 m^2 F4-type panel in normalized reflectance.

5. DISCUSSION

The results presented in Table 1 demonstrate that the local and stochastic ACE algorithms performed consistently better than the Global ACE algorithms for these data sets. This is supported by a lower number of average false alarms and average rank across the five targets.

Overall, local ACE performed the best out of the four algorithms. This is not particularly surprising as local ACE is optimized to detect low abundance targets mixed in with large and well-characterized background classes. For example, local ACE modeled the background statistics associated with the F4 target using the moderately large cluster represented by green in the segmentation image provided in Figure 1. Using this background model, local ACE was able to suppress the contribution of the background signal, resulting in the signatures shown in green in Figure 3. These signatures more closely match the target signatures, shown in red, than the in-scene pixel signatures, shown in black, match the target signatures. As a result, local ACE was able to decrease the probability of false alarm for F4 relative to the global ACE algorithms. This example generalizes across the five targets included as part of this analysis and is expected to generalize across all sub-pixel targets mixed in with well-characterized backgrounds.

Stochastic ACE performed the second best out of the four algorithms and was considerably closer to first place than third place. This ordering is more likely to be the result of a particularly good local ACE performance for these targets than a systematic disadvantage for stochastic ACE. In contrast to local ACE, stochastic ACE is able to derive the background model on a pixel-by-pixel basis in order to suppress the contribution of the background signal. In doing so, the algorithm directly models both the variability of the background and noise, as well as indirectly models the variability of the target, for each pixel. This suggests that stochastic ACE may be particularly powerful for finding targets in mixed pixels in which the variance of the underlying signals is unexpectedly large and unaccounted for, e.g., due to background, noise, or target variability. This also suggests that stochastic ACE may be able to detect targets in mixed pixels that cannot be associated with a well-characterized background class, e.g., beyond the binary mixing of target and background. By using a non-Gaussian approach that only requires that the noise has zero mean, stochastic ACE more closely obeys the underlying assumptions. As a result, the algorithm is expected to be more robust for complicated backgrounds. In addition, it complements the conventional detection algorithms by providing a somewhat uncorrelated approach. For example, in Figure 2, it can be seen that approximately 96.5% of the pixels in the detection map have scores identically equal to zero and the remainder of the scores appear to be qualitatively uncorrelated with those in the other detection maps. This is important as approaches that are fundamentally different in their assumptions and uncorrelated in their false alarms can be used to better assess confidence in a detection result.

Global masked ACE performed better overall than the conventional global ACE algorithm, with a substantially lower average number of false alarms and a lower average rank. The underlying difference between the performances of the two algorithms is that, for the seemingly more challenging targets, e.g., F3 and F7, global masked ACE was able to achieve detections with considerably less, e.g., 5x and 40x, false alarms than global ACE. This is evidenced by the relatively large standard deviation in number of false alarms associated with global ACE, which demonstrates a lack of robustness. In contrast, global masked ACE has a relatively modest standard deviation and does not appear to suffer

from the lack of robustness. This suggests that the use of a target and anomaly mask may not only improve detection performance, but also decrease the variance of detection results and improve the overall robustness of the ACE algorithm.

In addition to the improved detection performance demonstrated in Table 1, local and stochastic ACE can enable improved confidence building and false alarm mitigation through the use of background suppressed residual spectra. Coupled with conventional visualization and spectroscopy techniques, these residual spectra can provide insight into the spectral motivation and justification for a high statistical detection source. This is demonstrated in Figure 3. In particular, the residual spectra in the second plot (b) exhibit an increase in spectral reflectance between 580 and 640 nm, which can be directly correlated to a corresponding increase in the target signature, and is not apparent in the original pixel signature. This increase is even more evident in the third plot (c), in which the spectra correspond to a higher abundance F4-type panel. Using this insight, one can assign confidence to those true detections of the F4 target in mixed pixels. The example provided in Figure 3 demonstrates the potential of using residual spectra, a product of improved background modeling, to complement statistical detection results.

Although local and stochastic ACE demonstrated improved detection performance, the algorithms were somewhat slower than the global ACE algorithms. While speed performance was not measured as part of this research, all algorithms completed within a reasonable time, e.g., on the order of one to ten minutes using a Dell Precision T7500 PC workstation and the standard IDL/ENVI programming environment. These benchmarks indicate that the slower speeds processing times are not likely to be a limitation to the practical utilization of these techniques. Furthermore, the additional processing time associated with local and stochastic ACE will be offset by the reduced analysis time required to evaluate fewer false alarms. Note also that stochastic ACE scales linearly with the number of pixels and therefore could be used after triage or cueing to provide higher confidence detection results for a selected set of pixels. In addition, stochastic ACE considers each pixel independently and therefore could be massively parallelized to further improve performance speed.

As discussed in Section 3.2, the local, stochastic, and global masked ACE algorithms require several parameters that are not required by global ACE. Local ACE requires a maximum acceptance angle for clustering and a maximum number of whitened bands for forming the background subspace. Stochastic ACE requires a penalty parameter for constraining the background variability and the maximum number of iterations for optimization. For this research, the parameters were selected by qualitatively inspecting the general structure of the segmentation image and detection map, for local and stochastic ACE, respectively. The parameters were not rigorously optimized. In general, parameter selection offers another dimension of variability, which could potentially decrease the robustness of target detection performance. While sensitivity analysis was not performed on these data sets, parameter selection did not appear to greatly impact the quality of the results. In addition, all of the aforementioned parameters are used in whitened space, which is effectively equivalent across data sets. For this reason, the parameters may generalize across data sets.

There is no best algorithm or "free lunch" across all data sets and targets[16,33]. As such, while local ACE outperformed stochastic ACE and both algorithms outperformed the global ACE algorithms on these targets and data sets, it would be a mistake to choose only one algorithm for target detection in mixed pixels. Instead, it is suggested that these two new algorithms, with very different but improved assumptions, be used to complement each other, as well as the suite conventional algorithms, for improved target detection in mixed pixels.

6. CONCLUSIONS

This paper introduced two new methods for improved background modeling, as well as their associated target detection algorithms: local and stochastic ACE. While local ACE uses spectral clustering to improve upon the validity of the conventional Gaussian background assumptions, stochastic ACE overcomes the limitations of conventional assumptions by employing stochastic optimization to explicitly model background variability. Using two data sets provided by the RIT Target Detection Blind Test project, the detection performance of local and stochastic ACE was compared to the performance of two conventional global ACE algorithms. It was observed that local and stochastic ACE performed better than the global ACE algorithms in detecting five targets in mixed pixels, as well as provided an additional capability for confidence assessment and false alarm mitigation through the visualization of background suppressed residual spectra.

In performing this analysis, several lessons learned were captured and possible next steps identified. These lessons learned and next steps are documented below in the hopes that they will guide and motivate further research efforts towards improving target detection in mixed pixels.

6.1 Lessons Learned

Both the local and stochastic ACE algorithms use several seemingly small, but important modifications of conventional methods. For example, as angle is invariant to abundance, the use of spectral angle as a distance metric improves clustering and detection results for targets in mixed pixels. In addition, angle is invariant to illumination, which suggests that these methods may be readily extensible to exploitation in radiance domain, i.e., the forward modeling of reference signatures to radiance as opposed to the inverse modeling of the data cube to reflectance. Using the first several bands in whitened space to form a background subspace appears to improve both clustering and linear regression results for background segmentation and suppression, respectively. Ensuring that a background cluster has at least ten times the number of members as the number of spectral bands increases the robustness of covariance estimation. While parameter selection can be challenging, it is possible to estimate reasonable parameters through the visual inspection of the target and anomaly mask image, segmentation image, and detection maps. In ACE, replacing conventional mean subtraction with regressed mean subtraction, i.e., subtracting the background abundance multiplied by the background mean, for the test pixel and with no mean subtraction for the target signature appears to improve the performance of target detection in mixed pixels. Background suppressed residual spectra can be correlated with the associated target signature and in-scene pixel in order to ensure that background suppression results in a physically realizable residual signature. In addition, it may be possible to use background suppressed residual spectra to mitigate false alarms. Lastly, the use of non-Gaussian methods, such as stochastic ACE, can complement conventional Gaussian methods by generating detection maps that are relatively uncorrelated in their false alarms, which can be used to further build confidence and mitigate false alarms.

6.2 Next Steps

There are several next steps for future work that were identified during the course of this research. As mentioned in Section 4, a sensitivity analysis could be used to rigorously demonstrate the impact of parameter selection on detection performance. The resulting insights may inform the development of an automated method for parameter selection. In addition, it may be worth exploring the use of additional constraints for stochastic ACE to ensure that the stochastic optimization process models the background with physically realizable and accurate spectra. In parallel, it may be possible to automate false alarm mitigation based upon characteristics of the background suppressed residual spectra. Furthermore, it may be possible to develop novel detection algorithms based on the exploitation of residual spectra. Different forms of stochastic optimization, beyond the stochastic gradient technique, may be used to improve the rate of convergence and identify global optima for stochastic ACE. Similarly, it may be useful to explore the rate and nature of convergence as a function of the number of iterations. The detection performance of stochastic ACE may be improved by allowing the target signature to also vary within a constrained space, much as the background signature is allowed to vary. Stochastic ACE could be paired with local ACE to create a two-step hybrid approach that uses local background modeling to derive cluster statistics and then uses those statistics as the initial conditions for stochastic ACE. Stochastic ACE could also take advantage of novel feature-space transformations, e.g., beyond whitening, to improve the stochastic optimization process and detection performance. Lastly, evaluating the performance of the four target detection algorithms on additional data sets could provide more accurate estimates of detection performance and robustness, reveal new insights, as well as motivate other directions for future research.

ACKNOWLEDGEMENTS

The author gratefully acknowledges help and encouragement from Ron Resmini, Chris Simi, Sherry Olson, James Spall, Scott Tenaglia, and Brian Cavanaugh. The author also would like to acknowledge the Rochester Institute of Technology for providing access to Target Detection Blind Test project. This technical data was produced for the U. S. Government under Contract No. W15P7T-09-C-F600, and is subject to the Rights in Technical Data-Noncommercial Items clause at DFARS 252.227-7013 (NOV 1995).

REFERENCES

[1] Funk, C. C., Theiler, J., Roberts, D. A., and Borel, C. C., "Clustering to improve matched filter detection of weak gas plumes in hyperspectral thermal imagery," IEEE Trans. Geosci. Remote Sens. 39(7), 1410-1420 (2001).

[2] West, J. E., Messinger, D. W., Ientilucci, J. E., Kerekes, J. P., and Schott, J. R., "Matched filter stochastic background characterization for hyperspectral target detection," Proc. SPIE 5806, 1-12 (2005).

[3] West, J., Messinger, D., and Schott, J., "Comparative evaluation of background characterization techniques for hyperspectral unstructured matched filter target detection," J. Appl. Remote Sens. 1, 013520 (2007).

[4] Manolakis, D., and Shaw, G., "Detection algorithms for hyperspectral imaging applications," IEEE Signal Process. Mag., 19(1), 29-43 (2002).

[5] Manolakis, D., Marden, D., and Shaw, G. A., "Hyperspectral image processing for automatic target detection applications", Lincoln Laboratory Journal 14(1), 79-116 (2003).

[6] Messinger, D. W., West, J. E., and Schott, J. R., "Improving background multivariate normality and target detection performance using spatial and spectral segmentation," Proc. IGARSS 2006, 371-374 (2006).

[7] Manolakis, D. G., Marden, D., Kerekes, J. P., and Shaw, G. A., "Statistics of hyperspectral imaging data," Proc. SPIE 4381, 308-316 (2001).

[8] Murat Dundar, M., and Landgrebe, D., "A model-based mixture-supervised classification approach in hyperspectral data analysis," IEEE Trans. Geosci. Remote Sens., 40(12), 2692-2699 (2002).

[9] Theiler, J., Foy, B.R., and Fraser, A. M., "Characterizing non-Gaussian clutter and detecting weak gaseous plumes in hyperspectral imagery," Proc. SPIE 5806, 182-193 (2005).

[10] Kerekes, J.P. and Manolakis, D., "Improved modeling of background distributions in an end to-end spectral imaging system model," Proc. IGARSS 2004 2, 972- 975 (2004).

[11] Marden, D., and Manolakis, D., "Modeling hyperspectral imaging data," Proc. SPIE 5093, 253-262 (2003).

[12] Manolakis, D., and Marden, D., "Non-Gaussian models for hyperspectral algorithm design and assessment," Proc. IGARSS 2002 3, 1664-1666 (2002).

[13] Theiler, J., and Foy, B. R., "EC-GLRT: Detecting weak plumes in non-Gaussian hyperspectral clutter using an elliptically-contoured generalized likelihood ratio test," Proc. IGARSS 2008 1, I-221-I-224 (2008).

[14] Lo, E., and Ingram, J., "Hyperspectral anomaly detection based on minimum generalized variance method," Proc. SPIE 6966, 696603 (2008).

[15] Basener, B., and Messinger, D., "Clutter removal via topology for improved target detection in hyperspectral imagery," IGARSS 2008, (2008).

[16] Manolakis, D., Lockwood, R., Cooley, T., and Jacobson, J., "Is there a best hyperspectral detection algorithm?," Proc. SPIE 7334, 733402 (2009).

[17] Stocker A. D., and Schaum A. P., "Application of stochastic mixing models to hyperspectral detection problems", Proc. SPIE 3071, 47-60 (1997).

[18] Kraut, S., and Scharf, L. L., "The CFAR adaptive subspace detector is a scale-invariant GLRT," IEEE Trans. Signal Process. 47(9), 2538-2541 (1999).

[19] Farrell, M. D., and Mersereau, R. M., "On the impact of covariance contamination for adaptive detection in hyperspectral imaging," IEEE Signal Process. Lett. 12(9), 649-652 (2005).

[20] Foy, B. R., and Fraser, A. M., "Nonlinear signal contamination effects for gaseous plume detection in hyperspectral imagery," Proc. SPIE 6233, 62331U (2006).

[21] Theiler, J., and Foy, B.R., "Effect of signal contamination in matched-filter detection of the signal on a cluttered background," IEEE Geosci. Remote Sens. Lett. 3(1), 98-102 (2006).

[22] Manolakis, D., Siracusa, C., and Shaw, G., "Hyperspectral subpixel target detection using the linear mixing model," IEEE Trans. Geosci. Remote Sens. 39(7), 1392-1409 (2001).

[23] Reed, I. S., and Yu, X., "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution", IEEE Trans. Acoust., Speech, Signal Process 38(10), 1760-1770 (1990).

[24] Schaum, A. P., "Hyperspectral anomaly detection beyond RX", Proc. SPIE 6565, 656502 (2007).

[25] Richards, J.A., and Jia, X, [Remote Sensing Digital Image Analysis: An Introduction], Springer, Berlin, Heidelberg, and New York, 199 (2006).

[26] Hudgins, L. H., Hayashi, J., Blake, P. L., Tran, L. V., and Kilday, E., "Hyperspectral materials detection/identification/quantification using the residual correlation method," Proc. SPIE 4381, 391-408 (2001).

[27] Sweet, J. N., "Dominant component suppression with applications to spectral analysis," Proc. AIPR 2008, 1-7 (2008).

[28] Keshava, N., J. P. Kerekes, D. G. Manolakis, and G. A. Shaw, "Algorithm taxonomy for hyperspectral unmixing," Proc. SPIE 4049, 42-63 (2000).
[29] Spall, J. C., [Introduction to Stochastic Search and Optimization], John Wiley & Sons, Inc., New York, 22-27, 126-147 (2003).
[30] Snyder, D., Kerekes, J., Fairweather, I., Crabtree, R., Shive, J., and Hager, S., "Development of a web-based application to evaluate target finding algorithms," Proc. IGARSS 2008 2, II-915-II-918 (2008).
[31] Cocks, T., R. Jenssen, A. Stewart, I. Wilson, and T. Shields, "The HyMap airborne hyperspectral sensor: the system, calibration and performance," Proc. 1st EARSeL Workshop on Imaging Spectroscopy, 37-43 (1998).
[32] http://dirsapps.cis.rit.edu/blindtest/
[33] Wolpert, D.H., and Macready, W.G., "No free lunch theorems for optimization," IEEE Trans. Evol. Comput. 1(1), 67-82 (1997).