6-1-2005

Matched Filter Stochastic Background Characterization for Hyperspectral Target Detection

Jason E. West
U.S. Air Force

David W. Messinger
Rochester Institute of Technology

Emmett J. Ientilucci
Rochester Institute of Technology

John P. Kerekes
Rochester Institute of Technology

John R. Schott
Rochester Institute of Technology

Follow this and additional works at: <https://scholarworks.rit.edu/other>

Recommended Citation

Jason E. West, David W. Messinger, Emmett J. Ientilucci, John P. Kerekes, John R. Schott, "Matched filter stochastic background characterization for hyperspectral target detection", Proc. SPIE 5806, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI, (1 June 2005); doi: 10.1117/12.605727; <https://doi.org/10.1117/12.605727>

This Conference Paper is brought to you for free and open access by the Faculty & Staff Scholarship at RIT Scholar Works. It has been accepted for inclusion in Presentations and other scholarship by an authorized administrator of RIT Scholar Works. For more information, please contact ritscholarworks@rit.edu.

Matched filter stochastic background characterization for hyperspectral target detection

Jason E. West^{a*}, David W. Messinger^b, Emmett J. Ientilucci^b, John P. Kerekes^b, and John R. Schott^b

^a U.S. Air Force & Rochester Institute of Technology;

^b Digital Imaging and Remote Sensing Lab, Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, 54 Lomb Memorial Drive, Rochester, NY 14623-5604

ABSTRACT

Algorithms exploiting hyperspectral imagery for target detection have continually evolved to provide improved detection results. Adaptive matched filters can be used to locate spectral targets by modeling scene background as either structured (geometric) with a set of endmembers (basis vectors) or as unstructured (stochastic) with a covariance or correlation matrix. These matrices are often calculated using all available pixels in a data set. In unstructured background research, various techniques for improving upon scene-wide methods have been developed, each involving either the removal of target signatures from the background model or the segmentation of image data into spatial or spectral subsets. Each of these methods increase the detection signal-to-background ratio (SBR) and the multivariate normality (MVN) of the data from which background statistics are calculated, thus increasing separation between target and non-target species in the detection statistic and ultimately improving thresholded target detection results. Such techniques for improved background characterization are widely practiced but not well documented or compared. This paper provides a review and comparison of methods in target exclusion, spatial subsetting and spectral pre-clustering, and introduces a new technique which combines these methods. The analysis provides insight into the merit of employing unstructured background characterization techniques, as well as limitations for their practical application.

Keywords: Hyperspectral, target detection, background characterization, matched filters, covariance

1. INTRODUCTION

Hyperspectral imagery (HSI) may be defined as imagery taken over many (usually more than one hundred) spectrally contiguous and spatially co-registered bands. Target detection exploitation of HSI attempts to locate pixels containing a target material of known spectral composition. The adaptive matched filter is a type of detector that models and suppresses an unstructured background characterized by first and second order statistics (mean and covariance matrix) of a set of background pixels. Much of the work to improve matched filter performance has focused on improved scaling of detector results in order to increase separation between target detects and false alarms. Another and perhaps more fruitful approach to improved detection results is to improve the background model, thus increasing the suppression of unwanted signal and creating greater separation between target detects and false alarms *before scaling*. Based on this concept, several methods for calculating background statistics have been developed to improve detection results [1] [2] [3] [4]. Each of these methods takes cues from the imagery to establish a rationale for selecting which data is to be included in the estimation of background statistics. Establishing this rationale to model scene background and selecting a mathematical technique to formulate background statistics are together called the characterization of background. There are three factors to consider when attempting to improve matched filter background characterization. First and foremost is identifying the source of the signal interfering with the target. This serves two purposes: to suppress spectra mixing with subpixel targets and to suppress the returns on non-target species falsely identified as targets. The second consideration is to ensure that no target signal is present in the background model, as even small amounts of target will be shown to dramatically alter background statistics. Third is to attempt to seed the second moment with data that will adhere to the assumption of multivariate normality inherent in the traditional matched filter detector. This paper will discuss several methods of spatial and spectral subsetting and target exclusion that attempt to raise the detection signal-to-background ratio (SBR) and will introduce a new method combining these concepts.

* email: jew7776@cis.rit.edu, phone: 1-585-475-2453

2. METHODOLOGY

2.1 Matched filter detection

A variety of matched filters have been developed which use the Mahalanobis (or statistical) distance between a known target spectrum and a scene pixel as the primary measure of target presence. This simple matrix multiplication can be expressed

$$T = (d - \mu)^T \Sigma^{-1} (x - \mu) \quad (1)$$

where d is the target spectrum, μ and Σ are the mean and covariance matrix for the background distribution, x is the scene pixel being tested (the test pixel), and superscript T denotes the matrix transpose. Noting that the covariance matrix is inverted, we see that equation 1 is a spectral matched filter measure of signal divided by the statistical model of background. Thus the detection statistic T may also be considered a measure of signal to background ratio (SBR) [1]. In order to increase distance between target and non-target returns in the detection statistic, matched filters have been derived either by minimizing the total energy of the filter output by the Constrained Energy Minimization (CEM) technique [5], or by maximizing the cost function. From the cost function derivation, a family of matched filter detectors has been developed to include the Generalized Likelihood Ratio Test (GLRT), the Adaptive Coherence Estimator (ACE), and the Adaptive Matched Filter (AMF), which all scale the statistical distance differently to achieve improved results [2]. Only the GLRT detector, which can be expressed as

$$T_{GLRT} = \frac{(d^T \Sigma^{-1} x)^2}{(d^T \Sigma^{-1} d)(1 + x^T \Sigma^{-1} x)} \quad (2)$$

for mean centered data, will be used for this study to allow for level comparison of background characterization techniques. When using any of these detectors, it is important to note their underlying assumptions. First, the background is homogeneous and exhibits multivariate Gaussian behavior. Second, the background spectrum interfering with the target signature has the same covariance as the background training pixels. Third, the spectra of the target and background must combine in an additive manner [2].

2.2 Spatial subsetting

The first and perhaps most obvious method of improving on scene-wide background characterization with statistics is to manually select a spatial subset of the scene to represent the background. If the user has some knowledge of general target location, the background subset may be selected from a region of the scene imaged prior to the sensor reaching the target area. This "target approach" method can be used both to exclude targets and to achieve greater multivariate normality (MVN). The drawback of the target approach method is that it is difficult to decide which background best represents the signal mixing with subpixel targets, or that which may cause the most false alarms.

Another method of spatial subsetting, which is aimed at detecting the signal interfering with subpixel targets as well as increasing MVN of the background, is the RX algorithm [3]. The original algorithm implemented a combination of spatial and spectral matched filters, defining the spectral matched filter background with a sliding window to recalculate local background statistics for each pixel. According to the construction of the algorithm, the mean of the background can be calculated within a windowed subset of the data and the more slowly varying covariance can be calculated from a larger target-free subset and approximated by a diagonal matrix. The third moment (approaching zero) was used to approximate the normality of the spatial subset. Minimizing skew in this way was assumed to "create a distribution which is as close to Gaussian as possible" [3]. For this study, we considered only the spectral matched filter portion of the RX algorithm, using a single window to calculate both mean and covariance. Also unlike the original implementation, a fully populated covariance matrix was applied and skew was not explicitly calculated.

Limiting target influence on the sliding window characterization of background is an important consideration for RX because the smaller the window size, the greater the impact a few target pixels will have on background statistics.

We will examine two ways in which targets may be excluded from the sliding window background, which can be applied alone or in combination. First, the background window can be constructed with a central exclusion region spanning a distance twice the approximate size of the target, as depicted in Figure 1(a). Second, a prescreening step employing a loosely thresholded background-independent algorithm like Spectral Angle Mapper (SAM) [2] can exclude possible target species before RX detection, as shown in Figure 1(b). This prescreening draws a spectral distinction between species similar to and different from the target. The same concept can be extended to pre-cluster the image and spectrally subset the data to improve background characterization, as discussed in the next section.

Figure 1. (a) RX background window (in gray) demonstrating the exclusion of the central vehicle target, and (b) the decision map for SAM target exclusion prescreening where white pixels are excluded from the sliding gray background.

2.3 Spectral subsetting

Many new target detection algorithms have sought to exploit methods of HSI classification to gain advantage in the matched filter detection problem. The premise for improved detection results through pre-clustering (spectral subsetting) is achieving greater SBR by comparing the target to pixels in only a single class. The distribution of pixels within a class is intuitively more normal than intra-class mixtures, and a single class may be able to suppress the source of prominent false alarms in a scene. However, selecting a method of combining detection results using pre-clustered data can be complicated for hard target detection. The inclusion of target pixels in the background is also a problematic issue in pre-clustering. While the class covariance that maximizes SBR may not contain any target, at least one class spectrally close to the target will include target pixels. In order to use the class containing target species as background, the target pixels must be excluded from the class. Some techniques aimed at accomplishing pre-clustering target exclusion will be discussed later. First, we will examine a couple of classification algorithms that may be used for pre-clustering.

K-Means is a simple and commonly used unsupervised image classification method which requires the user only to specify the number of classes. Random class means are generated and each image pixel is assigned to a class by calculating a minimum distance to the mean. New class means are calculated based on the pixel assignments and the process is repeated until the means stop changing by a thresholded amount [6]. The matched filter can be adapted to incorporate results from a K-Means clustering, a method that has been shown to increase the performance of several forms of the adaptive matched filter when applied to HSI for plume detection [1]. For the GLRT detector, the adaptation takes the form

$$T_{GLRT}^k = \frac{(d^T \Sigma_k^{-1} x)^2}{(d^T \Sigma_k^{-1} d)(1 + x^T \Sigma_k^{-1} x)} \quad (3)$$

where k is the class number and values are mean subtracted using class means. The K-Means pre-clustering technique has also been used to improve anomaly detection results in multispectral IR imagery [7]. In both of these studies, it was postulated that improvements in the method of classification would lead to improved detection results.

Stochastic Expectation Maximization (SEM) is another classification technique which considers not only first order, but also second order statistics when assigning pixels to a class [8]. SEM is an unsupervised form of the supervised Gaussian Maximum Likelihood (GML) classification technique. The process is initialized by random assignment of pixels or from the results of a previous classification. Classes are modeled as normal distributions, combining additively to form the image by the expression

$$p(x) = \sum_{i=1}^K p(i) \left[\frac{1}{(2\pi)^{-\ell/2} |\Sigma_i|^{1/2}} e^{-\frac{1}{2}[(x-\mu_i)^T \Sigma_i^{-1} (x-\mu_i)]} \right] \quad (4)$$

where K is the number of classes in the image, $p(i)$ is the probability of each class existing, and ℓ is the number of spectral bands. The class statistics may be estimated from the distributions in $p(x)$, which exist in the data. The probability of a class existing in the next iteration can then be calculated by a ratio of the pixels in class i for iteration n to the total number of pixels in the scene, given by

$$p^{n+1}(i) = \frac{N_i^n}{N} \quad (5)$$

which must exceed the user defined threshold for minimum class size. The probability of each class existing $p(x|i)$ in the next iteration can then be estimated from the results, leading to an estimated *a posteriori* probability of each class, which can be expressed

$$p^{n+1}(i | x) = \frac{p^{n+1}(i) p^{n+1}(x | i)}{\sum_{i=1}^M p^{n+1}(i) p^{n+1}(x | i)} \quad (6)$$

where M is the user defined maximum number of classes [8]. Pixels are reclassified based on these probabilities and the new mean, covariance, and $p(i)$ values are calculated for each class. These values then seed the next iteration and the process continues until a convergence criterion has been satisfied. Incorporating second order statistics into the classification process allows the algorithm to consider spectral shape of a class, leading to more accurate classification and more normal class statistics within the classified image. SEM has been used in combination with RX to improve detection results through fusion [9]. Fusion rules for combining results, including model selection, AND/OR, and joint-density, were employed to create decision boundaries encompassing thresholds for each algorithm and for a normality metric, resulting in the reduction of false alarms by between 0.25 and 2.0 orders of magnitude.

2.4 Combined spectral and spatial subsetting with Adaptive RX

A new method of background characterization has been developed during this research which uses the results of classification and incorporates the concept of the RX sliding window. Like RX, the Adaptive RX (ARX) algorithm assumes that the best background for a target comes from the surrounding pixels. A limitation of the RX sliding window is that the size of the window must be small enough to capture the statistics of only the local neighbors of the detection pixel, but large enough to allow for sample sizes needed for well determined statistics. If the number of pixels in a window is less than the number of bands, the covariance matrix will be singular and cannot be traditionally inverted. Additionally, if the number of pixels is only a few times the number of bands, the matrix is usually highly elliptical and errors resulting from its use can be significant [10]. Classification can be used to overcome the problems with underdetermined statistics while preserving the ability of RX to capture the immediate surroundings of the detection pixel. By polling the pixels neighboring the test pixel and outside of an exclusion region, we can construct fractions of the test pixel surroundings that originate from each class. We can then form a linear mixture of the k class statistics to form a single class mean and covariance matrix which characterize the local background with well determined statistics. Figure 2 depicts an ARX window for a single test pixel and equation 7 gives the expression for the linear statistical mixtures, where α_i is the fraction of class i contained in the pixels neighboring the test pixel. Pixels from the class map are polled to establish the α fractions and the resulting class statistics are used in the matched filter applied to the central test pixel. The background for each test pixel is unique and matched to its surrounding, but runtime is greatly reduced compared to RX.

The algorithm also employs a technique to improve class statistics, which are usually derived directly from the class map. Understanding that each pixel in the scene has been forced into a distribution to which it may not belong, improvements may be made by excluding the outliers of each distribution and recalculating class statistics. If the target happens to reside in the tail of the distribution, it will be excluded; if not, targets can be excluded using a loosely

Figure 2. Adaptive RX representation where numbers in the torus represent class identifications.

$$\Sigma = \sum_{i=1}^k \alpha_i \Sigma_i \quad \mu = \sum_{i=1}^k \alpha_i \mu_i \quad (7)$$

thresholded SAM. In either case, statistical distance exclusion (SDE) eliminates the extremity of each class and essentially cuts off part of the heavy tail of the distribution, thereby increasing the MVN of the data seeding the statistics for matched filter detection. Figure 3 shows the distribution of a single SEM-generated class, first with no exclusion 3(a), then with target pixels, marked x, excluded by SAM prescreening 3(b), and finally with the heavy tail reduced by SDE 3(c), where excluded points are omitted. Once these pixels were excluded from the distribution, class statistics were developed from the remaining samples. The χ^2 goodness of fit (GoF) metric in Figure 3, which is described in the next section, shows an improvement in MVN for the final distribution.

Figure 3. Chi-Squared MVN plots showing (a) the original distribution, (b) pixels excluded by SAM prescreening (marked x), and (c) the tail of the distribution reduced by statistical distance exclusion (points omitted).

2.5 Evaluation metrics

Detection results are commonly represented by a receiver operator characteristic (ROC) curve comparing probability of detect to probability of false alarm. For targets with very few pixels, the error in calculating these probabilities can be preventative. We will therefore construct ROC curves with a rate of detection and false alarm, generated by simply counting the false alarm occurrences as each target pixel is detected. This will allow for a comparison of detection results without calculating probabilities while still making a general statement of algorithm performance. For more direct comparison, we can reduce the ROC curve to a single numerical value by taking an average false alarm rate (AFAR) [11]. This representation of the area above the ROC curve is a good indication of detection performance in finding all target pixels. To capture detection performance at low false alarm rates, a partial AFAR can be calculated as the area above the curve to a certain level of detection.

There are a number of available metrics to test a data subset for adherence to the multivariate normality assumption. MVN tests can generally be separated into four categories: graphical examination with correlation coefficients, tests for goodness-of-fit, skewness and kurtosis tests, and consistent procedures based on the empirical characteristic function [12]. For this data, we will use a goodness-of-fit (GoF) metric with the Chi-Squared MVN test.

For this test the rank ordered distribution of the statistical distance from the mean (d^2) is plotted versus the expected value given by the chi-squared distribution quantiles (χ^2). This chi-squared GoF is measured by

$$\chi_{GoF}^2 = \frac{1}{N} \sum_{i=1}^N \frac{(d_i^2 - \chi_i^2)^2}{\chi_i^2} \quad (8)$$

where N is the number of pixels in the subset [13]. While observation of the chi-squared plot gives insight into the shape of the distribution, the GoF measure gives a single numerical value allowing for a rough comparison between distributions.

3. DATA

To illustrate the performance of different background characterization techniques, we will use HYDICE data cubes from the Forest Radiance I and Desert Radiance II collects. Several data cubes from this experiment were analyzed by MIT Lincoln Laboratory for the purpose of exploitation algorithm development. Figure 4 shows a gray scale representation of each of the images with the target locations and target approach background regions indicated.

Figure 4. (a) Forest and (b) desert scene with targets (broken boxes) and background subsets (solid boxes) indicated.

The ground truth spectra and target location maps used to perform this analysis were provided in the data set. For simplicity, we will look at a single target that appears in both the forest and desert scenes. There were 45 full pixels, 28 subpixels, and 55 shadowed target pixels for a total of 128 target pixels in the forest scene, and 115 full pixels, 47 subpixels, and 73 shadowed target pixels for a total of 235 target pixels in the desert scene. This number and variety of target pixels allows for a range of detection performance for each background characterization technique.

4. RESULTS

4.1 Target approach results

Four target approach regions, outlined in Figure 4, were selected from the data sets in a manner similar to previous studies [2]. Each background in the forest scene contained 18K pixels while backgrounds in the desert scene contained 10K pixels. Using general knowledge of target location, the backgrounds were confidently presumed to be target free. Detection results using the GLRT detector with each of these backgrounds as well as the scene-wide backgrounds are shown in Figure 5. The AFAR reported for this and all subsequent figures is a 95% partial AFAR, excluding the last 5% of pixels detected in order to capture detection performance at lower false alarm rates. It is apparent from the chi-squared GoF charts, Figure 5(e) and (f), that each of the subsets exhibits far greater MVN than the scene-wide background. However, comparing Figure 5(c) and (e) it is clear that MVN is not necessarily correlated to performance. Without *a priori* knowledge of the subpixel interference signal or the location of false alarms in the scene, selection of the best spatial subset is difficult.

As demonstrated from Figure 5(b) and (d), the mixed background performed well in the desert image even though the subset was the least multivariate normal of the target approach regions. This illustrates the balance between maximizing SBR and exhibiting MVN, and indicates that the full scene background performed poorly, not because the type of spectra it contained were not suited to maximize SBR, but because the collection of pixels in the full scene strayed far enough from MVN to confound the results. Additionally, noting that the desert road proved to be a worse background than the full scene, we see that it is possible to select a background more normal than the full scene that fails to match (and therefore suppress) the interfering signal. Given these results, it is clearly difficult to optimize background characterization with the target approach method.

Figure 5. (a)-(d) GLRT detection results and (e)-(f) MVN GoF results using target approach and full scene backgrounds.

4.2 Target influence demonstration

The target approach subsets lend themselves well to a simple study of the influence a target species has on the background covariance. By including target pixels from elsewhere in the scene into the target free backgrounds, we can observe the changes in background statistics that would occur had we mistakenly included them. To quantify this change we can decompose the covariance matrix by the generalized eigenvalue problem

$$\Sigma = U^T \Lambda U \quad (9)$$

where U is an orthogonal matrix with columns containing the eigenvectors of Σ and Λ is a diagonal matrix containing the eigenvalues along the diagonal. While addition of a few target pixels in a large background may not change the overall variability in the covariance matrix, the shape of the matrix (which is of great importance for spectral matched filtering) will change significantly with only a few spectrally distinct target pixels. Changes in the shape of the covariance matrix manifest through changes in the ordered eigenvectors. Figure 6(a) shows the spectral angle in radians between the first ten eigenvectors from a target-free covariance and the first ten eigenvectors from a covariance contaminated by up to six target pixels. The significance of change in this angle is a function of the eigenvector number, so small changes in the first few eigenvectors may be more significant than larger changes in higher numbered eigenvectors. Note that there are small changes in the first and second eigenvector with the inclusion of just one or two target pixels and significant changes in the sixth eigenvector with the addition of just three target pixels. To see how this impacts detection, Figure 6(b) shows the GLRT return from a single target pixel compared to the maximum background return (first false alarming pixel) for up to ten target pixels included in the background. At zero target pixels included the target return is above the maximum background, but at six or seven pixels included the target return falls below the maximum background pixel. So, the introduction of just seven target pixels in this background of 18K pixels was sufficient to introduce false alarms for a target that could be detected without false alarms in the uncontaminated case.

Figure 6. (a) Spectral angle between eigenvectors and (b) GLRT return for various levels of target inclusion in background.

4.3 RX sliding window results

Another example of the importance of target exclusion can be found when selecting the parameters of the RX algorithm. The target vehicles in the forest scene are approximately ten pixels on a side and spaced thirty pixels apart. It is important to select correct sizes for the sliding window to prevent both the inclusion of pixels from neighboring targets (cross-contamination) and from the target itself (self-contamination). Figure 7 shows the setup and results for the RX sliding window with correctly selected background and exclusion windows, as well as the cross- and self-contaminated cases. The results are labeled $RX(i,j)$ by the pixel size of the exterior (i) and the interior (j) windows respectively. From Figure 7(a) and (b), it is clear that the target-free background, $RX(33,19)$, outperforms the background which sporadically picks up target pixels from nearby target vehicles, $RX(47,35)$. The properly sized sliding window also outperforms the background which includes pixels from the vehicle being detected, $RX(31,5)$.

Also shown are the results of detection with the contaminated window sizes prescreened with a loosely thresholded SAM. The SAM detection statistic was thresholded to exclude 300 pixels, roughly twice the number of assumed target pixels in the scene. The SAM prescreened target exclusion RX (TERX) elevates results of the poorly selected window sizes nearly to the level of the correctly selected case.

Freedom to select reasonably sized RX windows without concern for target contamination is desirable, recalling that the size of the RX window will also influence the stability of the background statistics. For these results an attempt was made to keep background size ($i^2 - j^2$) consistent. Future studies will include analysis of the influence of total background size on detection performance.

Figure 7. (a)-(b) Detection results and (c) set-up for RX sliding window detection for uncontaminated, cross-contaminated, and self-contaminated cases, as well as target excluded results for the poorly formed cases.

4.4 Pre-clustering results

The next step is to compare the results of pre-clustering for different types of classification algorithms. Five classes in the desert scene and six classes in the forest scene were identified by K-Means and SEM classification and the results were used to develop background statistics to seed the GLRT detector. Figure 8 shows the forest scene detection results for each class background for both classification algorithms. By comparing Figure 8(e) to (f), we see that MVN of the classes was improved by the SEM algorithm, as expected. The detection results in Figure 8(a) through (d) demonstrate that SEM outperformed K-Means in characterizing backgrounds for most classes. It is interesting to compare these results with those in Figure 5(a) and (c), noting that the hand-selected region of trees performed almost as well as the K-Means trees class as a background. Also of note is that the grass was divided by classification into two types (light and dark) with differing performance as backgrounds. The hand-selected grass background performance in Figure 5(c) fell in between the two SEM grass classes in Figure 8(c), indicating a mixture of light and dark grass in the target approach region. The best performing background was the soil class, pieced together from spatially scattered pixels which are impractical if not impossible to gather by hand. This class outperformed the best target approach region, demonstrating the utility of pre-clustering as a background characterization technique.

Figure 9 shows the desert scene detection results, which reveal greater improvement of SEM over K-Means, and demonstrate the detrimental effect of target presence in a cluster. Detection improvements are clear comparing Figure 9(c) to (d), with the SEM medium ground class finding every target pixel within a 0.07 false alarm rate. As for target influence, the K-Means algorithm classified the targets as part of the brush class (nearest in Euclidean distance), while SEM included them in the light ground class (nearest in statistical distance). Detection results using these backgrounds were heavily impacted by target presence, a fact which leads to a prime difficulty in the practical implementation of this method; it is difficult to blindly select a single class to best represent background. In order to use the closest class as a background, target pixels need to be excluded by some method. It was confirmed in processing that most false alarms occur within the class closest in statistical distance to the target, and when that class was used as background the false alarms were abundant and evenly distributed throughout the image. In this case, the closest target-free class may prove to be the best background, confining false alarms to the target-like classes and minimizing the false alarms elsewhere.

Figure 8. (a)-(d) Pre-clustering detection results and (e)-(f) MVN GoF results for the forest scene.

Figure 9. (a)-(d) Pre-clustering detection results and (e)-(f) MVN GoF results for the desert scene.

Intuitively, the best approach would be to exclude target pixels from the cluster in which they reside, preventing them from being suppressed as background and causing them to stand above false alarms in all other classes. Detecting and excluding target pixels in order to develop the background statistics required to detect the target seems to enter a circular argument, but iterations along these lines may be the key to the best automated pre-clustering background characterization technique.

4.5 Adaptive RX results

Combining the RX sliding window and pre-clustering techniques described above, the Adaptive RX algorithm eliminates the need to select an ideal window size or choose from a set of classes. Figure 10 shows the results of ARX for the forest scene, compared with results from the best and worst performing target approach, RX, and pre-clustering background characterization methods. These preliminary results indicate that ARX may outperform other background characterization techniques if poor assumptions have been made during processing, but does not match the intelligent analyst in detection capability. A plausible cause for this lies in the assumption that the best background for a target pixel can be derived from its immediate surroundings. This assumption may be the case for subpixel targets, where the background should model the source of interference mixing with the target signal within the pixel. However, for fully resolved pixels, the best background may be that which suppresses the source of false alarms from elsewhere in the image.

Figure 10. (a)-(b) Adaptive RX results compared to best performance of other techniques.

5. CONCLUSIONS

We have reviewed several techniques for background characterization involving the spatial and spectral subsetting of data in an attempt to improve the separation of target from background in spectral matched filter detection returns. The utility of the target approach method was demonstrated by showing that detection results can be improved by hand-selecting a small background versus using every pixel in the scene. This technique leaves the decision to the analyst as to which pixels best represent background for a given target. We demonstrated the importance of target exclusion by observing changes in the shape of covariance matrices with the addition of only a few target pixels. Spatial and spectral techniques for target exclusion were investigated using the sliding window background of the RX algorithm, revealing

the importance of carefully selecting window parameters. Comparing the results of two pre-clustering techniques, K-Means and SEM, we showed that SEM classification improved the multivariate normality for each class and the detection results for most background clusters. The presence of target species in pre-clustered backgrounds is problematic, and selection of the best background class is easily confounded by contamination. We introduced a new method, called Adaptive RX (ARX) which incorporates pre-clustering statistics into the sliding window background characterization concept. Preliminary results for the unsupervised ARX method indicate that it may improve detection results if less than ideal assumption were made during target approach, RX, or pre-clustering background characterization, but does not outperform correctly characterized backgrounds for these methods.

DISCLAIMER

The views expressed in this article are those of the author and do not reflect the official policy or position of the United States Air Force, Departments of Defense, or the U.S. Government.

ACKNOWLEDGEMENTS

This work was funded under the Office of Naval Research Multi-disciplinary University Research Initiative "Model-based Hyperspectral Exploitation Algorithm Development" #N00014-01-1-0867. The author would also like to thank his family for their support, without which this work would not be possible.

REFERENCES

1. C.C. Funk, J. Theiler, D.A. Roberts and C.C. Borel, "Clustering to Improve Matched Filter Detection of Weak Gas Plumes in Hyperspectral Thermal Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, **39**, pp. 1410-1420, July 2001.
2. D. Manolakis and G. Shaw, "Detection Algorithms for Hyperspectral Imaging Applications," *IEEE Signal Processing Magazine*, **19**, pp. 29-43, January 2002.
3. I.S. Reed and X. Yu, "Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution", *IEEE Transactions on Acoustics, Speech and Signal Processing*, **38-10**, pp. 1760-1770, October 1990.
4. Y. Li, A. Vodecek, R. Kremens, and A. Ononye, "A New Algorithm for Global Forest Fire Detection Using Multispectral Images," *Targets and Background IX, Characterization and Representation*, Proceedings of SPIE, **5075**, pp. 367-377, April 2003.
5. W.H. Ferrand and J.C. Harsanyi, "Mapping the distribution of mine tailings in the Coeur d'Alene River Valley, Idaho through the use of a constrained energy minimization technique," *Remote Sensing of Environment*, **59**, pp. 64-76, 1997.
6. J. Schott, *Remote Sensing: The Image Chain Approach*, Oxford University Press, New York, NY, 1997.
7. E.A. Ashton, "Detection of Subpixel Anomalies in Multispectral Infrared Imagery Using an Adaptive Bayesian Classifier," *IEEE Transactions on Geoscience and Remote Sensing*, **36-2**, pp. 506-517, March 1998.
8. P. Masson and W. Pieczynski, "SEM Algorithm and Unsupervised Statistical Segmentation of Satellite Images," *IEEE Transactions on Geoscience and Remote Sensing*, **33-3**, pp. 618-633, May 1993.
9. D.W.J. Stein, A. Stocker, S. Beaven, "The Fusion of Quadratic Detection Statistics Applied to Hyperspectral Imagery," SPAWAR Systems Center San Diego, Sponsored by CECOM RDEC Night Vision and Electronic Sensors Directorate, January 2001.
10. B-C Kuo and D.A. Landgrebe, "A Covariance Estimator for Small Sample Size Classification Problems and Its Application to Feature Extraction", *IEEE Transactions on Geoscience and Remote Sensing*, **40-4**, pp. 814-819, April 2002.
11. P. Bajorski, E.J. Ientilucci, and J.R. Schott, "Comparison of Basis-Vector Selection Methods for Target and Background Subspaces as Applied to Subpixel Target Detection", *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery*, Proceedings of SPIE, **5425**, pp. 97-108, April 2004.
12. C.J. Mecklin and D.J. Mendfrom, "An Appraisal and Bibliography of Tests for Multivariate Normality", *International Statistical Review / Revue Internationale de Statistique*, **72**, no. 1, pp. 123-138, April 2004.
13. G.W. Snedecor and W.G. Cochran, *Statistical Methods, Eighth Edition*, Iowa State University Press, 1989.