Algorithms for the Detection of Sub-Pixel Targets in Multispectral Imagery

Edward A. Ashton and Alan Schaum

Abstract

A new sub-pixel target detection algorithm is developed that integrates a linear mixing model (LMM) with the powerful "RX" anomaly detector of Reed and Yu (1990). RX is applied to mixing model errors instead of to measured radiances, because they are more nearly multivariate Gaussian. The integrated method consistently outperforms spectral anomaly detectors that are based on either RX or an LMM alone. A novel method of image-based endmember selection is also presented, and a simple method of computing the fully constrained LMM residuals is described.

Introduction

It has long been recognized that use of multispectral imagery can enhance the detectability of targets that are poorly resolved spatially (Margalit *et al.*, 1985; Hoff *et al.*, 1992; Schmalz and Ritter, 1994). Reed and Yu have shown that use of multiple spectral bands can dramatically improve detection for targets with known spatial shape and either known (Yu *et al.*, 1993) or unknown (Reed and Yu, 1990) spectral characteristics. In these papers the spatial distribution of targets played an incidental role, and backgrounds were assumed to be spatially uncorrelated. Furthermore, performance modeling assumed a multivarate Gaussian description of background spectral channels. In this paper, we examine the consequences for point-target detection of describing the background with simple phenomenology, in the form of a linear mixing model (LMM) (Gillespie *et al.*, 1990). The errors in this description are modeled as Gaussian distributed noise, suitable for application of the anomaly detector of Reed and Yu (the RX algorithm).

The linear mixing model has frequently been used in conjunction with multispectral remote sensing devices in the automated analysis of scene composition (Gillespie *et al.*, 1990). The approach taken here, however, is to make use of an LMM to model and remove background clutter without regard to its specific composition. We then use the RX algorithm, which in the case of sub-pixel targets reduces to spectral decorrelation filtering followed by energy detection, to detect targets in the resultant residual images. Because we are not overly concerned with obtaining realistic estimates of abundances of particular endmembers, we are able to use the computationally simple inscribed simplex LMM, as described by Craig (1990), with a novel and fully automated technique for endmember selection that makes use of a training sequence of typical background data.

Clutter reduction techniques using the LMM commonly take two forms: *continuum analysis*, in which the target spectrum is known and is included as an endmember in the LMM, and *residual analysis*, in which only the background is modeled, and pixels which are not well modeled are considered to be potential targets (Sabol *et al.*, 1992). Our algorithm makes use of the second (residual) analysis, as described in Schaum and Stocker (1994), Schaum and Stocker (1995), and Michalowicz and Schaum (1994). This type of analysis is more robust than continuum analysis, in that it assumes that the target spectrum can be mimicked by mixtures of the background elements in all but a few bands. Its essence is that the spectral vector at each pixel is replaced by the residual vector from the "best" modeled pixel estimate to the actual pixel. Note that the best estimate may take on different meanings, depending upon the strictness of the constraints that are placed upon the model. This topic will be addressed in detail in the following section. The residual vector is small in magnitude for well modeled pixels (background) and large for poorly modeled pixels (targets). Therefore, target detection can be achieved by simple thresholding of the vector magnitudes. However, we have found that, with certain statistical assumptions and careful consideration of the dimensionality of the residual vector space, it is possible in many cases to use the RX algorithm to increase target detectability in the residuals.

The endmember selection technique described in this paper was tested on both simulated data with known endmembers, and actual multispectral data for which endmembers have been manually selected by independent expert observers. The target detection algorithms were tested on infrared multispectral background images into which targets were artificially (but realistically) implanted, and on visible spectrum image cubes containing rural backgrounds that were processed similarly. Implanted target spectra were taken from actual spectral measurements of military vehicles. As a final validation, the algorithms were also tested on IR imagery, taken under realistic conditions, of regions containing actual military vehicles. The combined algorithm approach was tested against both RX and residual analysis alone, and proved in each case to provide significant improvement in target detectability over either algorithm.

Modeling

Linear Mixing

The core assumption of the linear mixing model is that a scene consists of a small (fewer than the number of spectral bands) number of distinct elements or endmembers, i.e., vegetation, sand, water, soil, etc. Each pixel in the scene must contain some combination of these endmembers, for example, 60 percent soil and 40 percent vegetation. The measured spectrum at that pixel should then be given by a weighted

Code 5621, Naval Research Laboratory, 4555 Overlook Ave. SW, Washington, DC 20375 (ashton@dynasun.nrl.navy.mil).

Photogrammetric Engineering & Remote Sensing, Vol. 64, No. 7, July 1998, pp. 723–731.

0099-1112/98/6407–723$3.00/0

sum of the endmember spectra. Mathematically, if $\mathbf{y}_i$ is the measured spectrum at a given pixel location i, N is the number of endmembers present in the scene, and the vector set $\mathbf{e}_n$ gives the spectra of the endmembers, the best estimate $\mathbf{x}_i$ of pixel $\mathbf{y}_i$ is given by

$$\mathbf{x}_i = \sum_{n=1}^{N} w_{n,i}\, \mathbf{e}_n, \tag{1}$$

where $w_{n,i}$ is the relative fraction of endmember n at pixel location i.

Two useful pieces of information come out of this equation: the weighting fractions $w_{n,i}$, which have frequently been used to determine the relative proportions of different elements in a scene, as in (Gillespie *et al.*, 1990), and the residual vector which can be obtained by subtracting the modeled estimate from the observed spectral data. The weighting fractions can be used for target detection using the continuum analysis technique (Sabol *et al.*, 1992), which assumes that the target is one of the endmembers and can be detected by thresholding the appropriate fraction plane. However, more interesting in our case is the residual vector $\mathbf{r}_i$ which is given by

$$\mathbf{r}_i = \mathbf{y}_i - \mathbf{x}_i. \tag{2}$$

If, as we hope, the LMM has accurately modeled the background, but is not an accurate model for the target, the vector magnitude of $\mathbf{r}$ at a target pixel should be larger than the expected background magnitude. Additionally, the *direction* of $\mathbf{r}$ in vector space may be different for target and background. It is this difference that we hope to exploit using the RX algorithm.

It should be noted at this point that the LMM is applicable only to cases of macroscopic mixing — those in which surface components are large and opaque enough to allow photons to interact with only one component (Adams *et al.*, 1986). Intimate mixing requires nonlinear modeling, which renders the equations given previously invalid.

Calculation of Weighting Fractions

Given a set of endmembers, the only information required to calculate $\mathbf{x}_i$ is the set of weighting fractions $w_{n,i}$. The problem of estimating these fractions given raw data and an endmember set has a unique, closed-form solution which can be calculated by inverting Equation 1 using a least-squares regression, while subjecting $w_{n,i}$ to the constraint,

$$\sum_{n=1}^{N} w_{n,i} = 1. \tag{3}$$

More precisely, if $\mathbf{s}$ is a matrix whose columns are the N endmember vectors $\mathbf{e}_n$, $\mathbf{y}_i$ is a column vector whose values give the measured spectrum of the pixel at location i, and $\mathbf{s}\#$ is the pseudoinverse of $\mathbf{s}$, given by

$$\mathbf{s}\# = (\mathbf{s}^{\mathrm{T}}\mathbf{s})^{-1}\,\mathbf{s}^{\mathrm{T}}, \tag{4}$$

then the weighting fraction $w_{n,i}$ is given by

$$w_{n,i} = (\mathbf{s}\#\ \mathbf{y}_i)_n - \frac{\left[\sum_{n=1}^{N} (\mathbf{s}\ \mathbf{y}_i)_n - 1\right]\left[\sum_{m=1}^{N} (\mathbf{s}^{\mathrm{T}}\mathbf{s})^{-1}_{m,n}\right]}{\sum_{n=1}^{N}\sum_{m=1}^{N} (\mathbf{s}^{\mathrm{T}}\mathbf{s})^{-1}_{m,n}}. \tag{5}$$

The inverse in Equation 4 and, consequently, the pseudoinverse, exists if and only if $\mathbf{e}_n$ form a linearly independent set, a condition that we always assume to hold. The first term in Equation 5 represents a projection of $\mathbf{y}_i$ into the subspace spanned by $\mathbf{e}_n$. The second term further restricts the solution to the affine subspace consistent with Equation 3. Once $\mathbf{w}_i$ is calculated through Equation 5, Equations 1 and 2 can be used to calculate the residual vector $\mathbf{r}_i$.

It is useful to consider the geometric implications of Equation 3. Because the LMM supposes that the spectrum of each pixel in a scene results from a summation of several pure endmember spectra, the pixel spectra must lie within a simplex in spectral space whose vertices are the endmember spectra. For example, three endmembers define a triangular simplex. A physically possible combination of the vertex spectra subject to Equation 3 must lie within that triangle. However, these equations are not sufficient to guarantee a physically possible solution. Equation 5 permits negative fractional weightings. In spectral space this corresponds to points that are outside the simplex defined by the endmembers, but are within their dimensional space, i.e., outside the triangle, but within the *plane* that the three endmembers define. For the remainder of this paper this residual, which provides the vector from the endmember space to the pixel vector and therefore does not fully implement the endmember model, will be referred to as the *first residual*.

In order to fully implement the model, it is necessary to impose a second constraint on the weighting fractions: i.e.,

$$0 < w_{n,i} < 1, \quad \text{for all } n, i. \tag{6}$$

Unfortunately, under this constraint the problem no longer has a closed form solution, and must be solved by a directed search of the solution space. This can be accomplished by first solving Equation 5 for all n, i. If all n fractions in the vector $\mathbf{w}_i$ are positive, then the projection of pixel $\mathbf{y}_i$ into the space defined by the endmembers does fall within the simplex. If there are negative fractions, then the true projection to the surface of the simplex must lie along an edge or at a vertex (assuming our simple three endmember model), i.e., one or more of the endmembers should have a weighting fraction of $w_{n,i} = 0$. This situation is illustrated for the two-band case in Figure 1.

The correct estimate $\mathbf{x}_i$ will be obtained if Equation 5 is re-calculated using a new endmember set which does not contain the endmembers which should have $w_{i,n} = 0$. The difficulty in this case lies in determining how many and

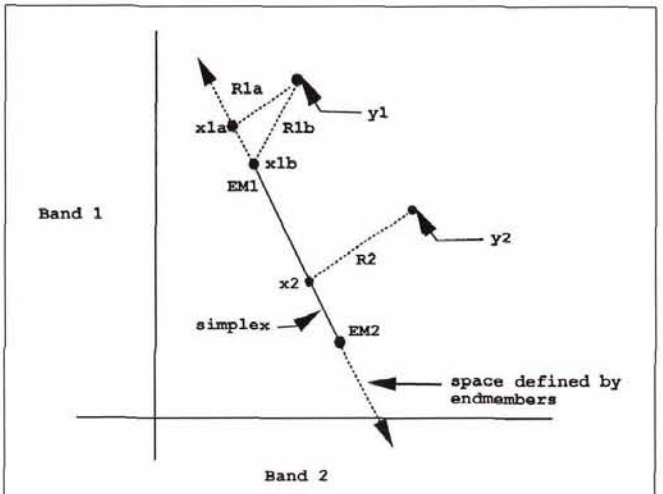

Figure 1. A two-band case with two endmembers EM1 and EM2, and two data pixels y1 and y2. Points x1a and x2 give the estimates of y1 and y2 as calculated by Equation 5. x2 lies within the simplex. x1a does not. A realistic projection, given by x1b, will only be obtained by giving Equation 5 a new endmember model which does not contain EM2.

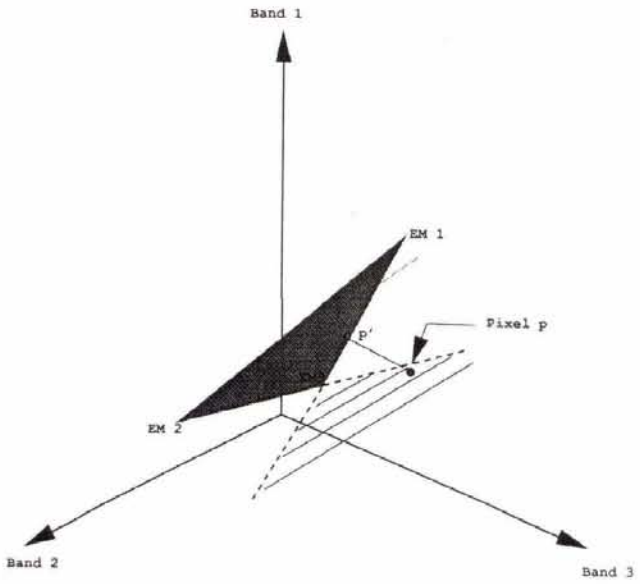

Figure 2. A scatterplot of three-band data in which each pixel is a linear combination of three endmembers. Equation 5 will calculate negative weighting fractions for both EM 1 and EM 2 for any pixel in the cross-hatched region, including p. However, the correct projection of p to the simplex, given by p', should contain EM 1. A negative weighting fraction is not a reliable indicator that an endmember should be eliminated. Exhaustive search methods are therefore necessary.

which endmembers should be null. A simple example, given in Figure 2, shows that it is not necessarily true that the fractions which are negative in the initial calculation of $\mathbf{w}_i$ are the ones that should be zero. One way to ensure optimality is to conduct an exhaustive search of all possible $N_1 = N - 1$ endmembers. If this yields no solutions, then a search must be made of all possible $N_2 = N - 2$ endmbembers, and so on until $N_{N-1} = 1$, at which point, of course, all fractional weightings will be 1. If more than one viable solution is found at a given value of N, then the solution for which $\mathbf{r}_i$ has the lowest vector magnitude is the correct one. This process is computationally expensive, but can be carried out in a reasonable amount of time if the number of endmembers is relatively small. The resulting residual, which provides the vector from the nearest point on the surface of the simplex to the pixel vector, will be referred to in this paper as the *total residual.*

Each of the previously described residuals has both advantages and drawbacks. If true endmembers are known, the total residual will provide optimal detection. However, because we make use of the inscribed simplex (a simplex whose endmembers are taken directly from available data), we are unlikely to find pure endmembers. Still, the total residual will prove superior for cases in which the target is similar in spectral shape to the background. The first residual requires much less time to compute. Also, it frequently provides superior detection when the target is spectrally distinct from the background. Another very large advantage to the first residual is that, for reasons which will be elaborated in a following section, the total residual is not statistically amenable to the RX algorithm. Primarily for this reason, the first residual is our preferred method.

Endmember Selection

The key to our automated endmember-selection algorithm is an understanding of the geometrical implications of the model. That is, the knowledge that an endmember model must result in a data distribution in spectral space that resembles a simplex, not the commonly assumed multivariate Gaussian distribution. An example of this is given in Figure 2, which shows a scatterplot of artificially produced three-band data representing a uniform mixture of three endmembers. Note that the data points fall into a triangular region, with the endmembers at the vertices. The endmembers, therefore, represent spectral extrema. This fact allows their identification. Our algorithm for endmember identification has five basic steps, and is based upon the "pixel purity index" developed by Dr. Joseph Boardman:

(1) Cull data to eliminate near-duplicate spectra.
(2) De-mean surviving spectra.
(3) Project spectra onto many lines through the origin – score 1 point to each extreme projection.
(4) Select the spectra with the most points as endmembers.
(5) Check for representation — eliminate anomalous spectra.

This algorithm exploits the convex property of a simplex; if all points within a simplex are projected onto any line, the extremal points correspond to two of the vertices. For the case shown in Figure 2, where pure endmembers are present in the data and there is no noise, this algorithm is guaranteed to identify the correct endmembers. Unfortunately, neither of these assumptions is generally true. A lack of pure elements in the data, combined with system noise, generally causes the data to take on a rounded shape, like a teardrop in 2D or a round-topped cone in 3D, with some anomalous pixels widely separated from the bulk of the data. These anomalous pixels are eliminated in Step 5, which calculates the $\mathbf{w}_i$ vector set and eliminates any selected endmembers which do not have at least 20 percent representation in at least 10 percent of the available data. These eliminated endmembers are then replaced by new ones, which are again tested until all endmembers have sufficient representation.

This selection algorithm was first tested on simulated data generated by selecting three four-band endmembers and combining them in random proportions to produce an image cube. Each endmember was represented in pure form at least once. The range of pixel values was 0 to 255, and the resultant data cube was corrupted by spatially invariant white Gaussian noise with $\sigma = 0.10\mu$. The endmembers were selected correctly, within the limits of the noise. Data enclosure for this case was 93 percent.

When this algorithm is applied to measured data, which generally is not only noise corrupted, but also rarely contains pure endmembers, it produces endmembers that define the largest possible inscribed simplex within the data cloud. Because only a fraction of the data is enclosed by the simplex — approximately 41 percent for the four endmember case (see Craig, 1990) — this technique is not well suited to scene analysis. However, because most of the excluded data are close to the surface of the simplex, while a target presumably is not, this algorithm is useful for our purposes. Comparison to manually selected endmembers for two data sets which will be described in the following section shows that in both cases our selection algorithm selects endmembers which are very close to the manually selected ones. An analysis of performance in target detection using both pure residual analysis and residual analysis followed by RX showed one case where the two endmember sets had nearly equivalent performance and three cases in which the algorithmically selected endmembers out-performed the manually selected ones. No cases were found in which the manually selected endmembers showed superior target discrimination.

TABLE 1. CHARACTERISTICS OF THE TIMS SENSOR

Band	Center (μm)	Bandwidth (μm)
1	8.4	0.4
2	8.8	0.4
3	9.2	0.4
4	9.9	0.6
5	10.7	0.8
6	11.5	0.4

The RX Algorithm

The RX algorithm is intended for the detection of targets with known shape in the spatial domain but unknown spectral distribution. It combines a spatial matched filter with a spectral anomaly detector, and has been shown to provide superior detection over either technique alone. We are interested in gains achievable through spectral analysis, and consider only targets that are unresolved and backgrounds that have been spatially whitened (Soni *et al.*, 1992) by local mean removal. In this case, the spatial filter kernel reduces to an impulse function. In short, the data vector $\mathbf{y}_i$ is replaced by a scalar z_i, whose magnitude provides a measure of the abnormality of the pixel at location i. This scalar is given by

$$z_i = \mathbf{y}_i^T \mathbf{R}^{-1} \mathbf{y}_i, \tag{7}$$

where $\mathbf{R}$ is the local spectral correlation matrix. The resultant image matrix $\mathbf{z}$ can then be thresholded to provide anomaly detection. Should $\mathbf{y}$ be multivariate Gaussian distributed, this procedure is equivalent to thresholding the background probability density function.

The only significant implementation question concerning this algorithm is related to the adaptive estimation of the correlation matrix $\mathbf{R}$ at each pixel, using only a window of surrounding pixels with a set width W. Our experiments indicate that near-optimal results can be achieved with window sizes ranging from 16 to 64 pixels. Clearly, calculating a new $\mathbf{R}$ at each pixel is a huge computational burden, particularly as the window size becomes larger. We avoid this problem by calculating $\mathbf{R}$ at a grid of points in the image, with grid points spaced every $W/2$ pixels. The intervening values for $\mathbf{R}$ are then calculated through bilinear interpolation, in the manner outlined by Pappas (1992). This produces a significant reduction in the computational burden, while sacrificing little in terms of performance.

Experimental Procedure

The algorithms described above were tested on targets artificially but realistically implanted into data from the Thermal Infrared Multispectral Scanner (TIMS), and on similarly processed data from the Naval Research Lab's Portable Hyperspectral Imaging Low Light Sensor (PHILLS). The target spectra were obtained from spectral measurements of portions of military vehicles. These spectra were then scaled to minimize thermal contrast with the background.

The TIMS data consists of six bands of longwave IR collected from an airborne platform. Two cubes of TIMS data were used, one taken over an airstrip in Kona, Hawaii, and the other over a region of the White Mountains in California. These data were acquired from SETS Technologies under a contract with the Naval Research Lab (NRL). SETS also provided a group of manually selected endmembers with each data set. These served as a point of comparison for our automated endmember-selection algorithm. The spectral characteristics of the TIMS sensor are given in Table 1.

NRL's PHILLS sensor provides spectrally simultaneous imaging of 1024 bands with a spectral range of 250 nm to 1208 nm. The volume of this data was reduced to manageable levels by excluding the noisy spectral extrema and decimating the remainder to 25 distinct spectral bands, then selecting a six-band subset of these 25 that provided acceptable contrast between the target and background. The data were taken from a ground-based platform at the Hogback Overlook near Luray, Virginia, and consisted primarily of forested regions.

Full-pixel targets were implanted into each of the image cubes, and the results were then blended with the original background cubes using linear interpolation to produce new cubes with target pixel-fill factors ranging from 0 percent to 100 percent in 2 percent increments. Each of these cubes was processed using our combined algorithm, the RX algorithm alone, spectral unmixing alone, and a spectral matched filter. The results were then thresholded at the highest level such that all targets were detected, and the resultant number of false alarms was plotted against the pixel-fill factor to obtain a measure of performance. Note that the matched filter was included only as a point of comparison, because it requires knowledge of the target spectra which is not assumed known for the other algorithms.

Our final set of experiments involved data taken using Night Vision Laboratory's linear variable filter (LVF) infrared camera. This sensor obtains 256 spectral bands ranging from 3 to 5 μm. These image cubes were taken from the Hogback and Mount Marshall Overlooks and contained sub-resolved military vehicles. These data were reduced to six bands through decimation and elimination of atmospheric absorption regions, and were processed in the same manner as the

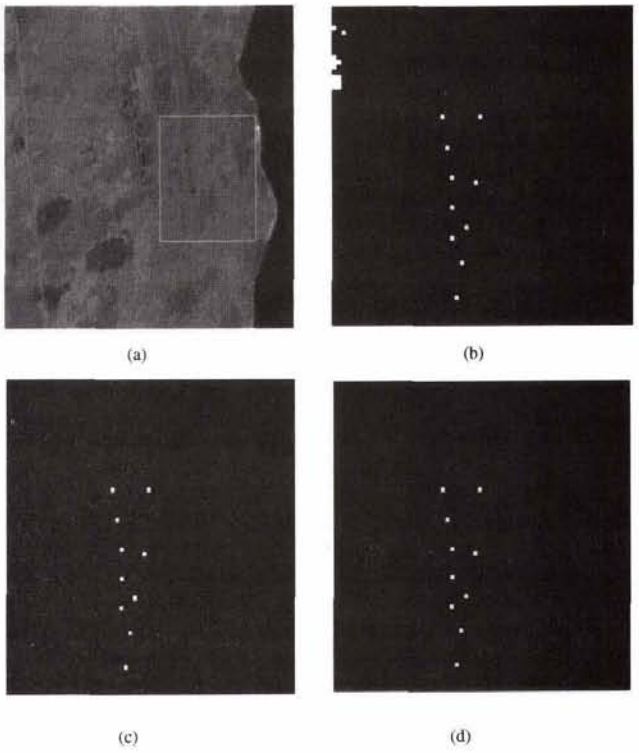

Figure 3. (a) A 200- by 200-pixel section of band 1 from the TIMS Kona data cube. This section contains ten targets at 16 percent pixel fill. (b) The outlined section of (a), processed using RX filtering and thresholded at 100 percent detection. False alarm rate is 0.35 percent. (c) Processed using first residual unmixing followed by RX. This case yields 0 false alarms. (d) Processed using spectral matched filtering. This case yields 0 false alarms.

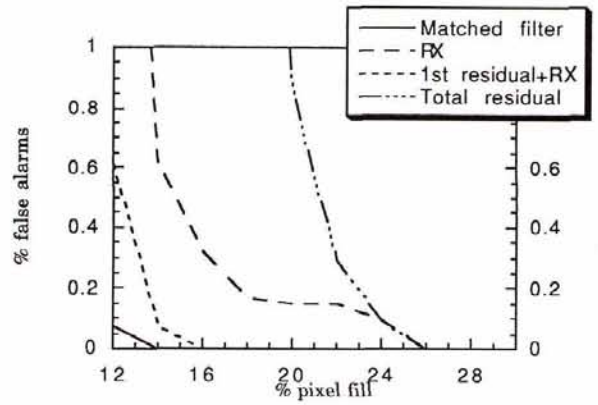

Figure 4. Plot of pixel fill versus false alarm rate at 100 percent target detection for the TIMS Kona2 data cube. Note that the combination of first residual unmixing provides better detection than either RX or total residual unmixing alone.

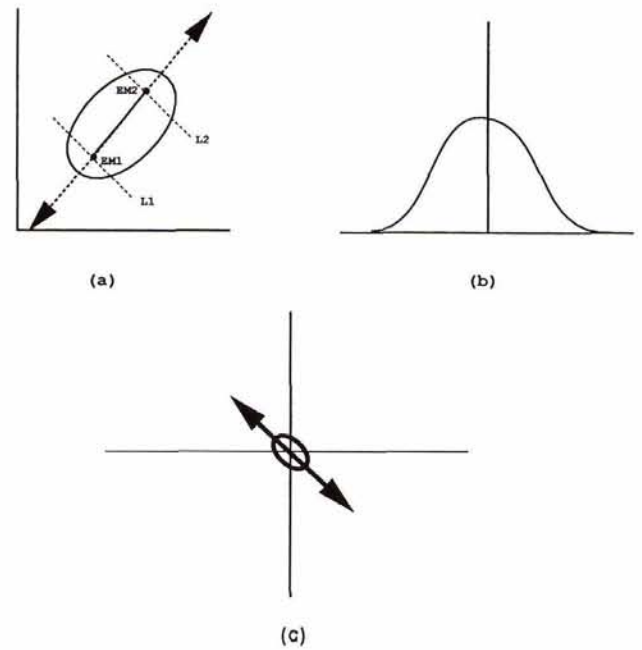

Figure 5. (a) Two-band data distribution, with two endmembers and a 1D simplex. (b) 1D, approximately Gaussian distribution resulting from the first residual. Only the dimension perpendicular to the extended simplex has value. The dimension parallel to the extended simplex will have a null eigenvalue and will be eliminated. (c) 2D distribution resulting from the total residual. All pixels falling between L1 and L2 in (a) will form a 1D distribution similar to (b), in a direction perpendicular to the simplex. The pixels outside L1 and L2 will form a smaller 2D distribution which will be nearly Gaussian. The superposition of these two distributions, given in (c), will be highly non-Gaussian.

TIMS and PHILLS data. In these data, the precise pixel-fill factors of the targets are unknown, but relative false-alarm rates for the various algorithms could be derived. These results are given in the following section.

Results and Discussion

Application to TIMS data

Our first application experiments were carried out on two sets of TIMS data. Figure 3(a) shows band one of a 64- by 64-pixel block of the TIMS Kona data cube. Ten targets are present. Each is a CARC brown metal panel, and fills 16 percent of its pixel. Figure 3(b) gives the thresholded output of this cube processed using the RX algorithm with window size 32. Note that, although the targets are visible, there are significant numbers of false alarms. Figure 3(c) shows the results obtained through background reduction using linear unmixing followed by application of the RX algorithm to the first residual. Figure 3(d) gives the output of the spectral matched filter.

Figure 4 plots false alarm rate versus pixel fill factor for this data cube. As expected, the matched filter provides the best results, reaching zero false alarms at a pixel fill factor of 14 percent. Of the pure anomaly detectors, a combination of first residual unmixing with RX provides the best result by a significant margin, followed by RX alone and unmixing using the total residual without any postprocessing.

Two complications in the application of RX to unmixing residuals bear mentioning. First, the first residual has a dimensionality which is less than that of the original data. In order to apply RX to the first residual, it is necessary to apply the Hotelling transform first, and eliminate those dimensions with which null eigenvalues are associated. The number of null dimensions is one less than the number of endmembers. This situation is illustrated in Figure 5. Second, the total residual is not statistically suited to the RX algorithm. RX as-

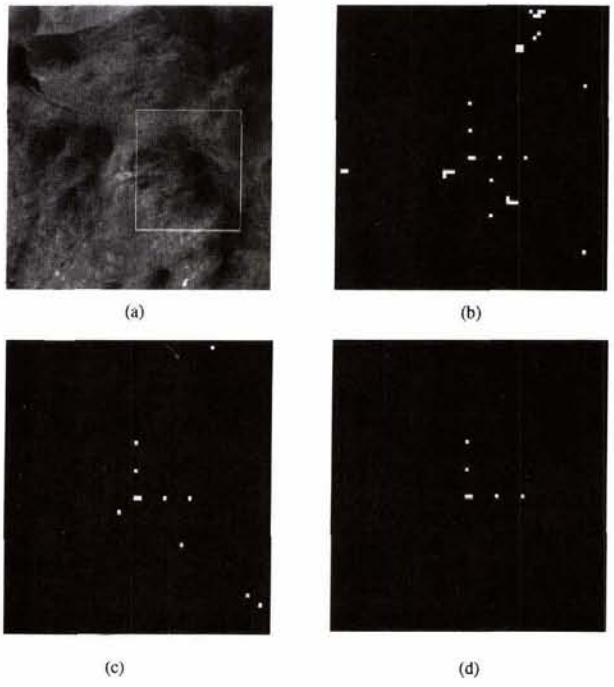

Figure 6. (a) A 200- by 200-pixel section of band 1 from the TIMS White Mountains data cube. This section contains five targets at 70 percent pixel fill. (b) Processed using RX filtering and thresholded at 100 percent detection. FAR in this case is 0.60 percent. (c) Processed using spectral unmixing followed by RX. FAR is 0.15 percent. (d) Processed using spectral matched filtering. FAR is 0.04 percent.

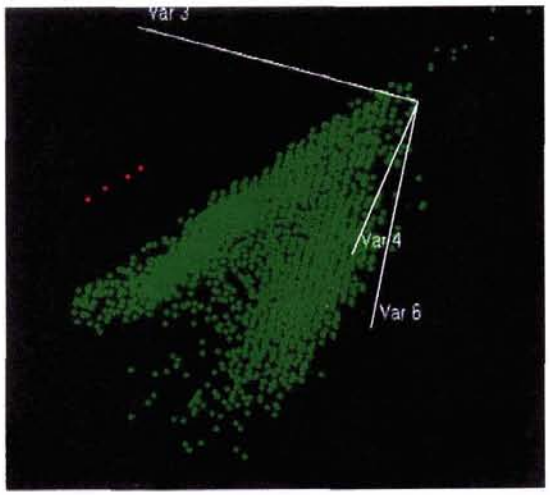

(a)

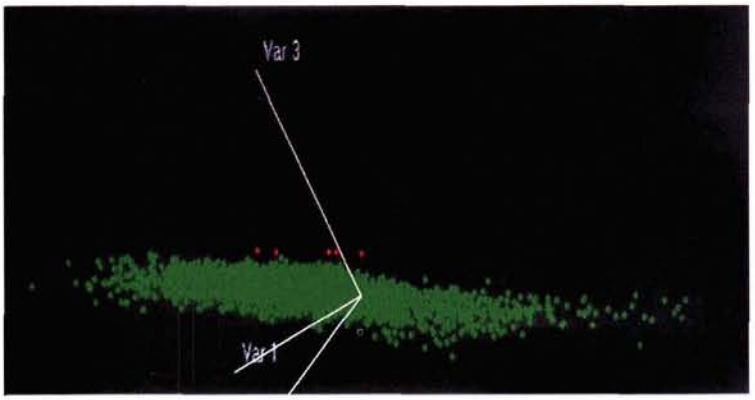

(b)

Plate 1. (a) Scatterplot in three bands of targets (red) and background pixels (green) for the TIMS Kona data cube with 100 percent pixel fill factor. Note first that the background does fall into a well defined simplex, and second that the targets are well separated from the background distribution. This data set is ideally suited for application of a LMM. (b) Scatterplot in three bands of targets and background for the TIMS White Mountains data cube with 100 percent pixel fill. The simplex in this case is less well defined, and the targets are very close to the background in spectral space. This is a much more difficult detection problem.

sumes a multvariate Gaussian distribution of data. Such a distribution can be reasonably expected to result from the first residual. The total residual, however, generally results in a distribution which is the superposition of two or more multivariate Gaussian distributions *of different dimensionality*. Points which project to the interior of the simplex will be reduced in dimensionality by a factor of $N-1$, where N is the number of endmembers. Points which project to an edge or vertex will be reduced in dimensionality by a smaller number in the first case, and not at all in the second. Applying RX to the total residual causes noise boosting in the pixels with null dimensions, producing results which are unpredictable.

Figure 6a shows band one of a 64- by 64-pixel block of the TIMS White Mountains data cube. This cube contains five targets. Again, each is a CARC brown metal panel. In this case, each has a pixel fill factor of 70 percent. Figures 6b, 6c, and 6d show the thresholded output of our processing algorithms, as above. Note that, in this case, the required pixel fill factor for target identification is considerably higher than that observed in the Kona data cube. The reason for this is obvious from examination of Plate 1, which shows scatterplots of background and target pixels for both data sets in three bands. Clearly, the target and background pixels in the White Mountains cube are much closer in spectral space than the target and background pixels in the Kona cube. Furthermore, the Kona background data falls into a well-defined simplex distribution in spectral space. This data set fits the LMM's assumptions. The White Mountains data, however, fall into a more Gaussian distribution. Spectral unmixing should provide less benefit in this case. This observation is confirmed in Figure 7, which shows the relative performances of the detection algorithms on this cube. The results are similar to those obtained for the Kona cube, but the pixel fill factor required for equivalent levels of detection is much higher. Also, the relative performance of total residual unmixing has improved from that seen in the Kona cube, due to the closeness in spectral shape of the target and background, and the total gains achievable through unmixing are reduced. However, the use of RX as a postprocessor still allows the first residual to provide the most effective detection.

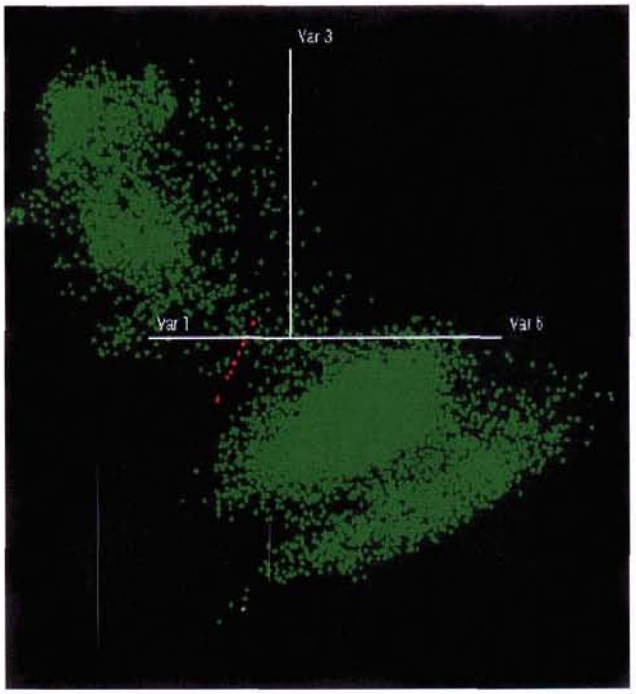

Plate 2. Scatterplot in three bands of targets (red) and background pixels (green) for the PHILLS Hogback image cube with 100 percent pixel fill factor. The distribution of these data cause the enclosing simplex to also enclose the targets in several dimensions. Because of this, the gains from the LMM will be much smaller than those seen in the infrared imagery.

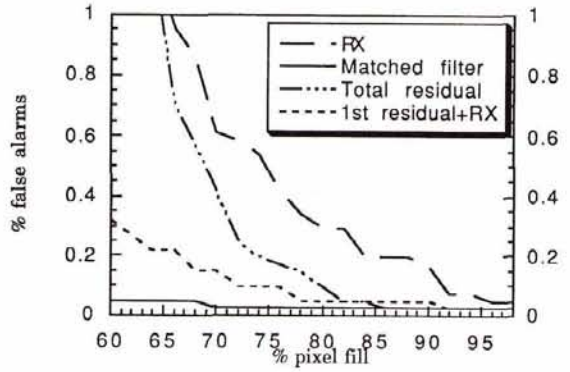

Figure 7. Plot of pixel fill versus false alarm rate at 100 percent target detection for the TIMS White Mountains data cube. Spectral unmixing in combination with RX again provides the best anomaly detection, but all algorithms require a higher pixel fill factor in this case due to the similarity of the background and target spectra.

Application to PHILLS Data

Our second set of application experiments involved data taken using the NRL PHILLS sensor from the Hogback overlook in Shenandoah National Park. This image cube contains primarily forested areas. The target spectrum was taken from the CARC green metal side panel of a military vehicle. The image cube was processed in the same manner as the TIMS data, and detection was attempted using each of the algorithms. Thresholding was carried out at the highest level such that the single implanted target was detected. False alarms were then calculated as a percentage of total pixels, and plotted against pixel fill.

Figure 8a shows band one of a 64- by 64-pixel block of the PHILLS data cube. This cube contains one target, at a pixel fill factor of 30 percent. Figures 8b, 8c, and 8d show the thresholded output of our processing algorithms. Note that the results for this case are consistent with those obtained in the thermal IR experiments described above. Plate 2 shows the distribution of these data in spectral space. The targets in this case appear to be lost in clutter. This data set is not particularly well suited to spectral unmixing, but there is separation in the out-of-plane direction which the LMM is able to exploit. Figure 9 shows the relative performances of the various detection algorithms on this cube. Although the gap in performance between the matched filter and the anomaly detection algorithms is greater in this case than in the previous ones, and unmixing using the total residual provided results which were far worse than the other algorithms, the results were generally consistent with those obtained using TIMS data. Again, unmixing using the first residual in combination with RX proved to be the most effective of the anomaly detection algorithms.

Application to LVF Data

Our final set of validation experiments involved data taken using the LVF sensor. These image cubes were intended to be realistic approximations of possible field applications. The images were of forested rural regions, and contained actual military vehicles at sub-pixel ranges. Figure 10a shows band one of a 100- by 100-pixel block of the LVF Hogback data cube. Figures 10b, 10c, and 10d show the results of processing this cube using our algorithms and thresholding at the highest level that preserves the known target. It is difficult to determine false-alarm rates in these cases, because other man-made objects in the scene should also appear as anomalies, and the number of such objects is unknown. FAR is therefore most useful as a relative measure. The results of processing in Figure 10 demonstrate moderate utility for the unmixing technique, with total residual unmixing and first

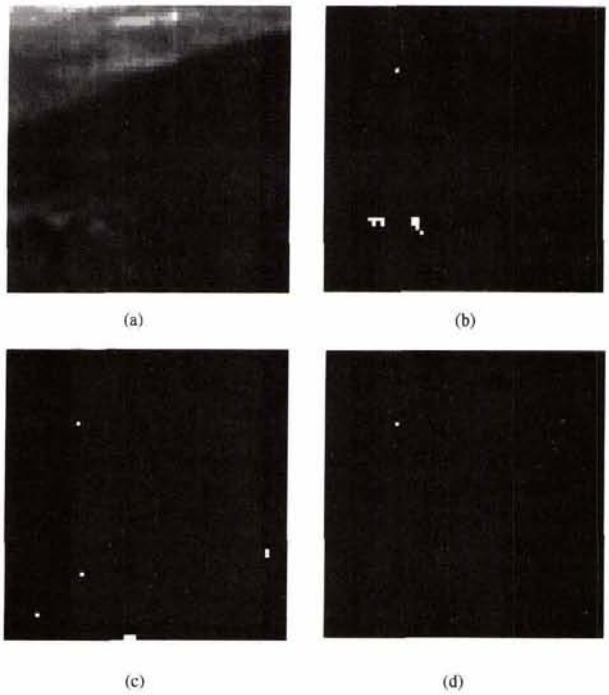

Figure 8. (a) A 64- by 64-pixel section of band 1 from the PHILLS Hogback Overlook data cube. This section contains one target at 30 percent pixel fill. (b) Processed using RX filtering followed by thresholding at 100 percent detection. FAR in this case is 0.29 perecent. (c) Processed using spectral unmixing followed by RX. FAR is 0.17 percent. (d) Processed using spectral matched filtering. This case yields 0 false alarms.

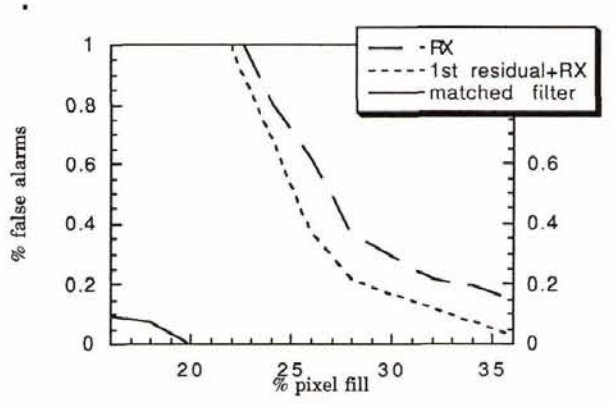

Figure 9. Plot of pixel fill versus false alarm rate at 100 percent target detection for the PHILLS Hogback data cube. Note that these results are consistent with those achieved using TIMS data. Total residual unmixing in this case was of limited utility.

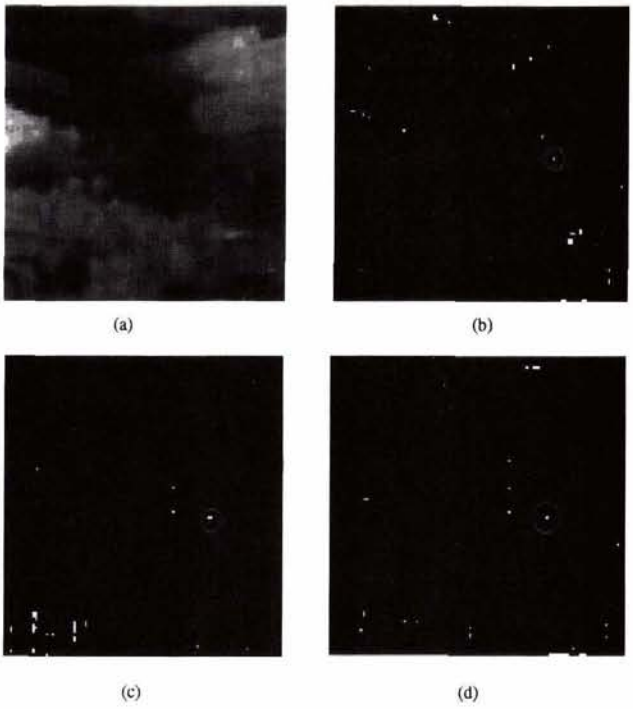

Figure 10. (a) A 100- by 100-pixel section from the LVF Hogback data cube. This section contains one target of sub-pixel extent. (b) Processed using the RX filter. FAR is 0.34 percent. (c) Processed using total residual unmixing. FAR is 0.24 percent. (d) Processed using first residual unmixing followed by RX. FAR is 0.27 percent.

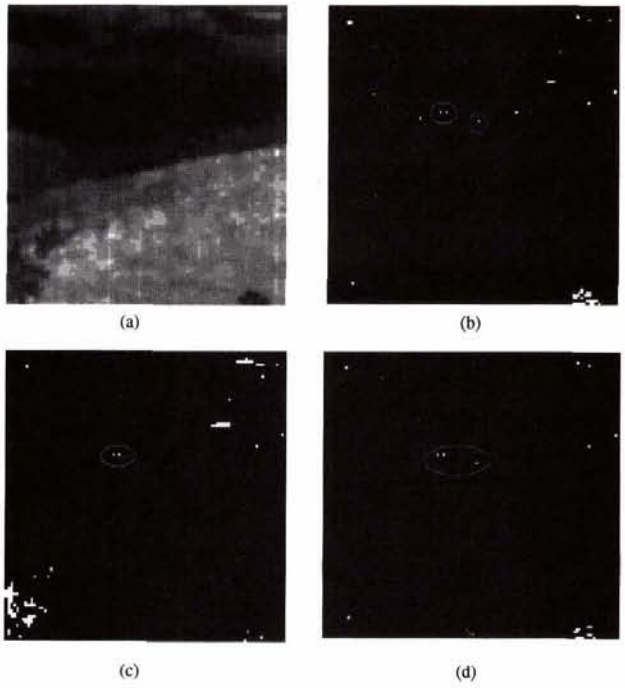

Figure 11. (a) A 100- by 100-pixel section from the LVF Mount Marshall data cube. This section contains three targets of sub-pixel extent. (b) Processed using the RX filter. FAR is 0.40 percent. (c) Processed using total residual unmixing. FAR is 0.76 percent. Note also that one target is lost. (d) Processed using first residual unmixing followed by RX. FAR is 0.18 percent.

residual unmixing followed by RX providing roughly equivalent performance.

More interesting results are seen in Figure 11, which shows a scene taken from the Mount Marshall LVF data cube. In this case, there are three sub-pixel targets. RX and first residual unmixing followed by RX both identify all three targets, with unmixing providing a reduction of more than 50 percent in false alarms. Total residual unmixing, however, not only has a higher false alarm rate, but also fails to identify one of the targets. This result highlights the fact that, although the combination algorithm is sometimes sub-optimal, as in Figure 10, its advantage in robustness compensates to an extent for the possibility of a fall-off in performance.

Conclusions

In this paper we have addressed the most difficult problem related to detection of targets in multispectral imagery: detection of sub-pixel targets with unknown spectral signatures. We have demonstrated the application of a known algorithm, RX, to the residuals produced through an application of the linear mixing model, and shown that the combination algorithm provides consistently superior results to those obtained through either method alone. We have also presented a novel approach to the automated determination of endmembers for the LMM, which is the most challenging problem associated with that model.

Our experiments have indicated that these algorithms are applicable to both the infrared (thermal) and visible (reflective) spectra. Clearly, the benefits of applying this modeling technique are directly related to the extent that the background distribution matches the assumptions of the LMM. Future work will include a more detailed analysis of this phenomenon. In particular, we would like to determine whether it is possible to assess the applicability of these algorithms to a given problem *a priori*. Future work will also include application to a wider range of target/background combinations, more extensive application to mid-wave IR imagery, and further testing of these algorithms on imaged targets in realistic settings.

References

Adams, J., M. Smith, and P. Johnson, 1986. Spectral mixture modeling: A new analysis of rock and soil types at the viking lander 1 site, *J. Geophysical Research*, 91:8098–8122.

Craig, M., 1990. Unsupervised unmixing of remotely sensed images, *Proceedings of the Fifth Australasian Remote Sensing Conference*, Perth, Australia, pp. 4–10.

Gillespie, A., M. Smith, J. Adams, and S. Willis, 1990. Spectral mixture analysis of multispectral thermal infrared images, *Proc. 2nd Thermal IR MUltispectral Scanner (TIMS) Workshop*, JPL Publ. 90–33, pp. 57–74.

Hoff, L., J. Zeidler, and C. Yerkes, 1992. Adaptive multispectral image processing for the detection of targets in terrain clutter, *Signal and Data Processing of Small Targets*, Orlando, Florida, SPIE, 1698:100–114.

Margalit, A., I. Reed, and R. Gagliardi, 1985. Adaptive optical target detection using correlated images, *IEEE Trans. on Aerospace and Electronic Systems*, AES-21:394–405.

Michalowicz, J., and A. Schaum, 1994. Infrared spectral signal processing for small target detection, *Proc. of the Joint Symposium on Multisensors and Sensor Fusion*, NATO AC/243 Panel, Brussels, Belgium.

Pappas, T., 1992. An adaptive clustering algorithm for image segmentation, *IEEE Trans. Signal Processing*, 40:901–914.

Reed, I., and X. Yu, 1990. Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution, *IEEE Trans. Acoustics Speech and Signal Processing*, 38:1760–1770.

Sabol, D., J. Adams, and M. Smith, 1992. Quantitative subpixel spectral detection of targets in multispectral images, *Journal of Geophysical Research*, 97:2659–2672.

Schaum, A., and A. Stocker, 1994. Subpixel detection methods: Spectral unmixing, correlation processing, and when they are appropriate, *International Symposium on Spectral Sensing Research*, San Diego, California.

———, 1995. Chronochrome image processing for infrared target detection, *IRIS Specialty Group on Targets Backgrounds and Discrimination*, Albuquerque, New Mexico.

Schmalz, M., and G. Ritter, 1994. Image-algebraic design of multispectral target recognition algorithms, *Image Algebra and Morphological Image Processing V*, pp. 213–228.

Soni, T., J. Zeidler, and W. H. Ku, 1992. Adaptive whitening filters for small target detection, *Signal and Data Processing of Small Targets*, SPIE, 1698:21–31.

Yu, X., I. Reed, and A. Stocker, 1993. Comparative performance analysis of adaptive multi-spectral detectors, *IEEE Trans. Signal Processing*, 41:2639–2656.

(Received 07 November 1996; accepted 01 December 1997; revised 22 December 1997)