Improved Matched-Filter Detection Techniques

Pierre V. Villeneuve [*a], Herbert A. Fry [b], James Theiler [b], William B. Clodius [b], Barham W. Smith [b], and Alan D. Stocker [a]

[a] Space Computer Corporation, 2800 Olympic Blvd, Santa Monica, CA 90404-4119

[b] Los Alamos National Laboratory, MS C-323, Los Alamos, NM 87545

ABSTRACT

Numerous statistical approaches have been developed for small target detection in cluttered environments. Examples include orthogonal background suppression (OBS) where the initial principal components are suppressed, and the clutter matched filter (CMF) where the principal components are weighted by the inverse of the eigenvalues and the latter principal components are discarded. Our research has shown that improved target detection performance can be obtained by combining certain aspects of both OBS and CMF approaches. This is especially true in the presence of limited scene data (finite number of pixels) or an imperfect reference target spectrum. The basis of this idea is to use weighting by the inverse of the eigenvalues (from CMF) for the initial PCs and then uniform weighting for the later PCs (from OBS). Examples of this new technique and comparisons with OBS and CMF will be shown with model data with realistic clutter containing a chemical plume.

Keywords: matched filter, detection, OBS, eigenvalue, thermal IR

1. INTRODUCTION

The matched filter is an ideal automated detection tool for use with high-dimensional hyperspectral data when searching for small or weak targets of interest. This technique is widely used in statistical signal detection applications (Fukunaga, 1990). In hyperspectral small-target detection, the matched filter may be considered a process by which the target spectrum is modified in such a way as to maximize the signal-to-clutter ratio for a particular realization of clutter (Stocker et al, 1990). In this case clutter is simply signal variance resulting from changing scene spectral properties (i.e. different surface and material spectral signatures or spatially varying atmospheric conditions).

Given the hyperspectral pixel vector $\underline{\mathbf{x}}$ with N channels along with mean spectral covariance matrix **K** (N x N) the matched filter is easily computed and applied. The optimal search vector is calculated as

$$q = K^{-1} \underline{t}, \tag{1}$$

where the vector $\underline{\mathbf{t}}$ is a library spectrum representative of the target of interest. The scalar matched filter output is then simply

$$r = q^{T} \underline{x}, \tag{2}$$

which is normally applied to each pixel vector $\underline{\mathbf{x}}$ of a hyperspectral data cube and the resulting image of matched filter results is interpreted as a mapping of the relative strength of the target $\underline{\mathbf{t}}$ at each pixel.

2. MATCHED FILTERS AND PRINCIPAL COMPONENTS ANALYSIS

Further insight into how the clutter matched filter suppresses clutter is to reconsider the entire procedure in the principal component (PC) domain defined by the covariance matrix decomposed into eigenvectors and eigenvalues:

$$K = UVU. \tag{3}$$

[*] Correspondence: pierre@spacecomputer.com

The matrix K is symmetric and the columns of **U** are the N eigenvectors and the diagonal of **V** is the set of eigenvalues. Quantities expressed in the PC domain are shown with a superscript *. In the PC domain the covariance matrix K^* is simply the diagonal matrix **V** of eigenvalues. In this representation the matched filter in the PC domain is

$$\underline{q}^* = V^{-1}\underline{t}^*. \tag{4}$$

This shows us that the matched filter can be thought of as a version of the original vector $\underline{\mathbf{t}}^*$ whose various principal components are weighted by the inverse of the corresponding eigenvalues. For instance, if the principal components are ordered in descending order by eigenvalue, the projection of $\underline{\mathbf{t}}$ in the direction of the first PC will be suppressed a great deal while the projection into PC number 100 (for example) will be relatively enhanced.

This insight now leads us to consider the question: what is the best estimate for $\mathbf{K}^{-1}$ when it is not known a priori? The quickest way to estimate $\mathbf{K}^{-1}$ is to do so from the data itself:

$$\tilde{K}^{-1} = \left(\frac{1}{N} \sum_{i=1}^{N} (\underline{\mathbf{x}}_i - \underline{\mathbf{m}})(\underline{\mathbf{x}}_i - \underline{\mathbf{m}})^T \right)^{-1}, \tag{5}$$

where $\underline{\mathbf{x}}_i$ is the spectral vector from pixel i and the ~ symbol implies an estimated parameter. There are several possible problems associated with this source of estimating the spectral covariance. One is there may be insufficient data points (scene pixels) to properly sample the covariance matrix. Another is that data processing techniques (i.e. zero padding of interferograms to the next highest power of two) may introduce unusual artifacts to the eigenvalue distribution. A third is that the target itself may be contributing variance at significant levels that would in turn be suppressed by the matched filter. Our observation is that the inverse of the best estimate of the covariance matrix is not necessarily the best estimate of the inverse of the covariance matrix.

We propose that better detection results can be obtained by modifying the eigenvalues of a particular subset of the overall eigenvalue distribution. Partial motivation for modifying the data-derived eigenvalues can be seen when comparing eigenvalue distributions for both model scene data and purely random data.

The model data set was based on real data collected with the Deadelus scanner flying over the Bull Run power plant. This Deadelus data was classified into approximately ten discrete regions with a clustering algorithm. Each class was then assigned a surface temperature and variance as well as a surface material emissivity spectrum at a resolution of 2 cm^{-1}. A MODTRAN atmosphere (transmissivity and path radiance) was superimposed on top of the thermal surface emission to give the top of the atmosphere radiance that a high altitude (aircraft or satellite) sensor would measure. A broadband image of this scene is shown in Figure 1. A Gaussian chemical plume of hot tetrachloroethylene was added to the scene as a target to be detected. Figure 2 shows the tetrachloroethylene spectrum resampled to 2 cm^{-1} and Figure 3 shows an example matched filter output for this chemical. The instrument noise was modeled as white noise at an SNR level of approximately 1000. The scene contains 128 by 128 pixels.

Figure 1. Model broad band image.

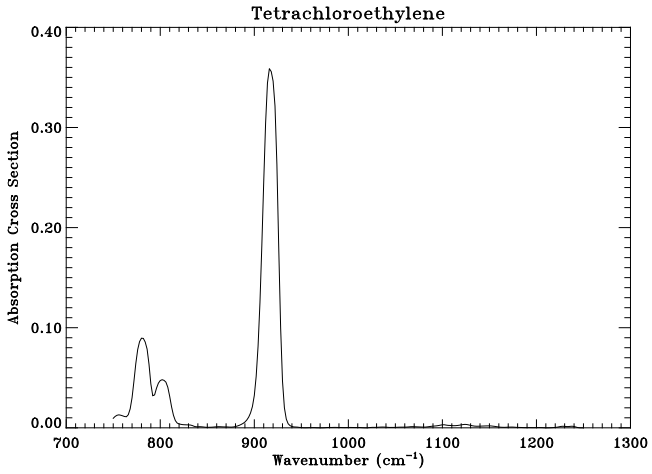

Figure 2. Plume chemical spectrum.

Figure 4 shows the eigenvalues obtained from a simulated hyperspectral data cube (details of the model will be given in the next section). Figure 5 show the eigenvalues for a random dataset of the same dimensions as the model scene data. Note in Figure 5 that the random data results show a non-constant distribution resulting from the particular instance of randomness when the dataset was created. However, since the data are purely random, it is not reasonable to place more value on one PC than another. The distribution seen in Figure 5 is in essence a result of sorting a particular realization of the random data.

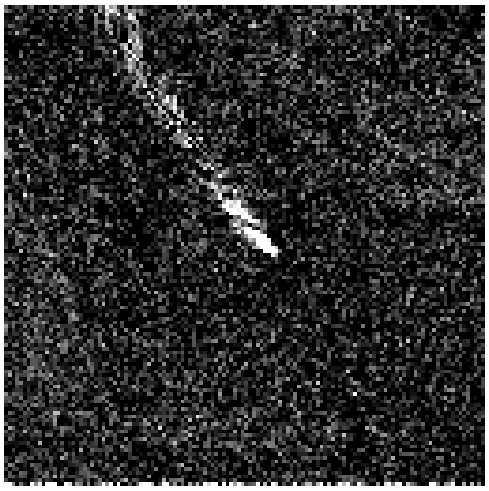

Figure 3. Sample tetrachloroethylene matched filter.

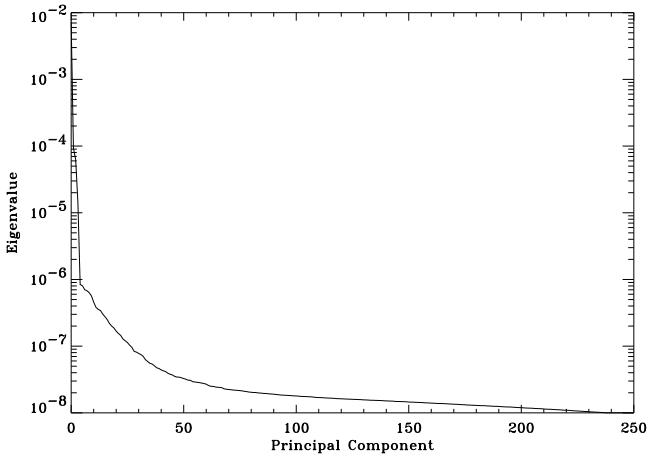

Figure 4. Eigenvalue distribution from model data.

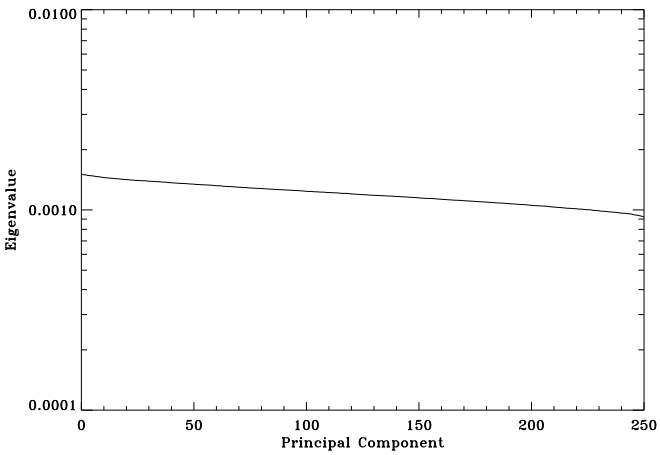

Figure 5. Eigenvalue distribution from random data..

This realization was in part the motivation for modifying the eigenvalue distribution to something other than what was derived from the data. As discussed a few paragraphs above, the inverse of the eigenvalues represent the relative weights in the corresponding hyperspectral directions applied to the library spectrum to create the optimal matched filter. However, if a particular PC has a lower eigenvalue than the noise level, then it is not sensible to allow those eigenvectors to have more influence on the matched filter than other PCs. Thus our advancement over the commonly used clutter matched filter and OBS is to force the PC eigenvalues to remain at a constant level after reaching a certain PC.

Figures 6 through 8 show three different weightings of the eigenvalues in a matched filter implementation. In each case, the dotted line is the inverse of the data-derived eigenvalues, and the solid line is the modified version. Figure 6 illustrates the new proposed distribution where the inverse of the eigenvalues are "saturated" after a certain PC. The purpose of saturating the eigenvalue inverse is to force the later PCs to be weighted equally by the clutter matched filter. This means that the

higher order PC directions will be emphasized identically, which is consistent with how a noise-whitened data set should be treated. This technique will be referred to as CMFsat, where CMF refers to Clutter Matched Filter. Figure 7 and 8 are illustrations of how two other popular techniques fit into this conceptualization.

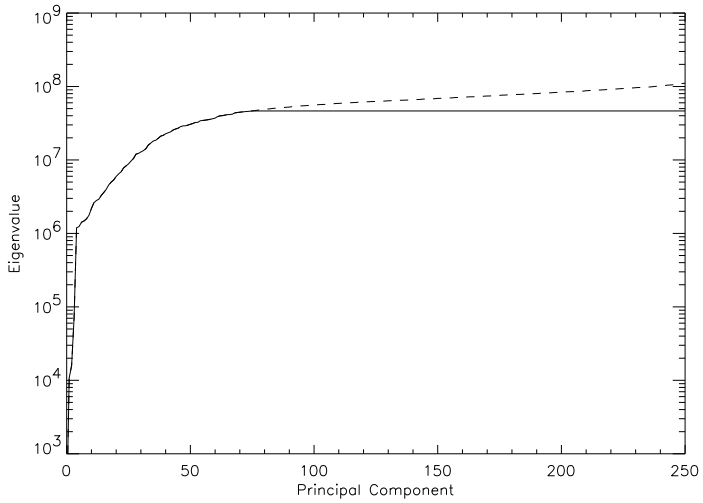

Figure 6. Clutter matched filter with saturated eigenvalue weighting (CMFsat).

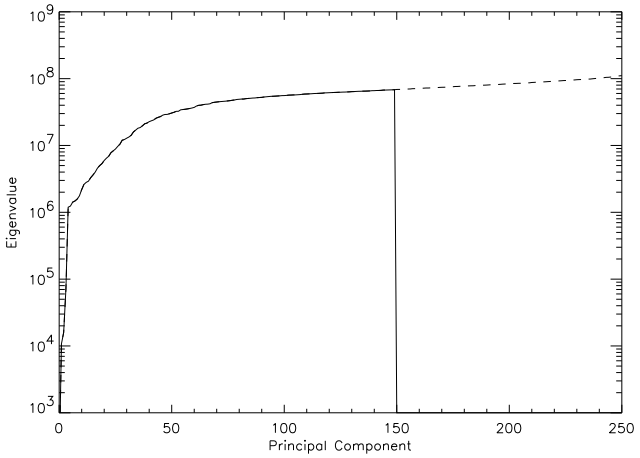

Figure 7. Clutter matched filter with zero weighting of later principal components (CMFcut).

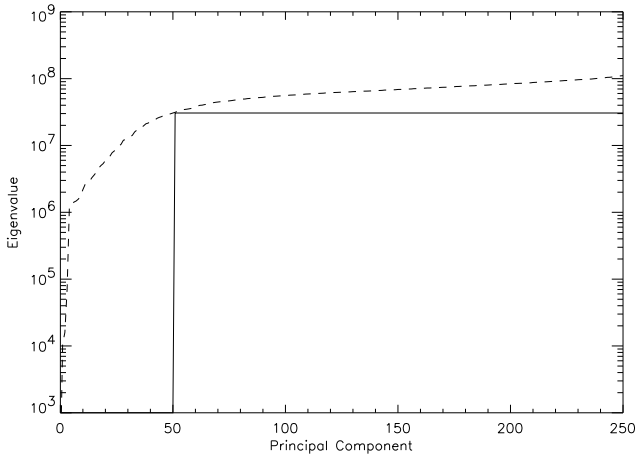

Figure 8. Clutter matched filter with OBS-type weighting of principal components.

Figure 7 describes what is commonly done with the standard clutter matched filter. Typically an analyst will perform a principal component analysis on the dataset and project out the later set of PCs under the assumption that they mostly describe noise. Next the clutter matched filter is performed on the remaining PCs giving optimal weighting by the inverse of the eigenvalues. This is illustrated in Figure 7 by the solid line. Discarding the later set of PCs is equivalent to setting their eigenvalues to zero so that those principal components are completely suppressed from the analysis. This method will be referred to as CMFcut.

Figure 8 shows orthogonal background suppression (OBS). This is a technique where a set of orthogonal endmember vectors (typically obtained from principal component of the scene) is used to represent a set of undesirable signal contributions (i.e., clutter). OBS is simply a mathematical projection of the data set to a subspace orthogonal to the specified orthogonal endmembers. After this step the data pixel spectra are compared to library spectra. The source of these orthogonal vectors is most often the first set of principal component eigenvectors (Harsanyi and Chang (1994) and Tu et al

(1997)). In the frame work described in this report, OBS may be considered as a type of matched filter with the eigenvalue weighting for the first set of PCs set to zero (complete suppression of those components) and uniform weighting for the remaining PCs. In this context it is clearly seen that OBS is really a poor approximation to the optimal weighting by the eigenvalue inverse.

It will be shown in the next sections that the later two techniques (OBS and CMFcut) provide inferior performance when compared to CMFsat because significant numbers of PCs are completely discarded resulting in a decrease in overall signal-to-clutter ratio. This occurs especially with datasets where the target of interest occupies only a small area fraction of the scene, thereby making little if any impact on the estimation of the mean spectral covariance matrix. In this case, the target spectral signature will be spanned almost uniformly by all PCs, including high order so-called "noise PCs".

3. ALGORITHM EVALUATION WITH MODEL DATA

The ideas presented in the previous section were tested with the model dataset described earlier. A broad band image of this scene, calculated as the average of frames spanning the longwave IR window, is shown in Figure 1. An example matched filter output for tetrachloroethylene is shown in Figures 3. The plume originates from a source near the center of the image and propagates to the upper left corner of the scene. The plume is initially quite hot and at high concentration and gradually cools to ambient temperature and diffuses with the air so that by the edge of the scene the plume is hardly detectable.

The various matched filter implementations presented in this report were tested with a pre-defined set of pixels from the model scene. Two specific regions of interest were defined from the matched filter result in Figure 2. An area excluding all plume pixels was defined by simply selecting all pixels (approximately 14000 in total) which did not contain any part of the plume. A second region was found by thresholding the matched filter result (normalized to unit variance) to focus on those pixels with a values greater than 2.0 (approximately 100 pixels). This thresholding selected only those pixels associated with the plume. Algorithm performance results were found by calculating the ratio of the mean matched filter chemical prediction for the ensemble of pixels in the on-plume region to the standard deviation of the algorithm prediction in the off-plume region. This ratio is the mean signal-to-clutter ratio for the particular algorithm when applied to this scene. The two pixel regions defined for this SCR calculation remained constant for all tests in this report giving a common ground for comparison of different ideas.

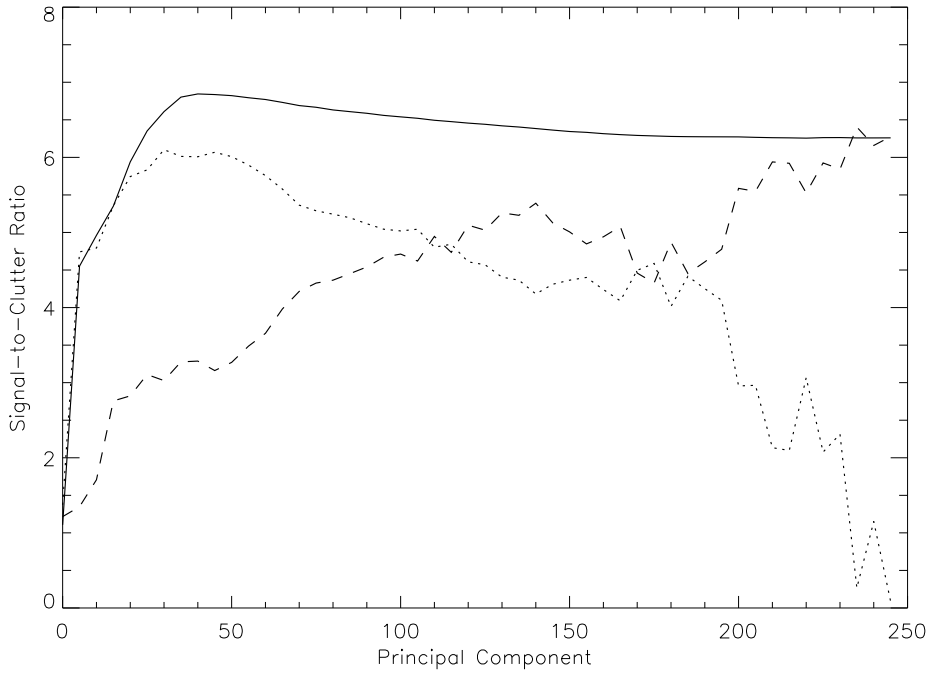

Figure 9. Signal-to-clutter detection ratios for CMFsat (solid), CMFcut (dashed), and OBS (dotted).

The off-plume pixels were used for the calculation of the mean spectral covariance matrix **K**, as opposed to using all pixels including the plume pixels. This is the ideal case where the signal of the target of interest (the tetrachloroethylene plume) does not manifest in any of the eigenvectors. Results of applying the three different approaches to the clutter matched filter (OBS, CMFcut, and CMFsat) are shown in Figure 9 as a function of the PC number at which the particular type of eigenvalue modification is applied.

The dotted line in Figure 9 shows the signal-to-clutter ratio for the OBS method calculated from the mean signal over the on-plume region divided by the standard deviation of the matched filter prediction over the off-plume region. This SCR is plotted as a function of the cutoff PC marking the boundary between where the PCs are completely suppressed and where they are included but with equal weighting. Thus, at low values of the PC cutoff for OBS, the chemical detection SCR is very poor and close to zero since none of the background information is suppressed. However, as the cutoff is moved to higher PCs, the performance improves, quickly at first but then more gradually, as more and more background information is suppressed, allowing the plume chemical signature to be measured with a better SCR with a peak value close to 6.0 near PC number 30. As the cutoff is further increased, more PCs are suppressed resulting in the suppression of significant levels of signal as well as the background clutter. The result is that the overall SCR decreases until finally by PC number 250 the detection SCR is back down to a ratio close to zero indicating that all signals (both target and clutter) are suppressed.

The dashed line in this figure shows the performance of the more standard algorithm CMFcut that suppresses those PCs upward of the cutoff point, as indicated in Figure 7. The results in Figure 9 for CMFcut indicate very poor performance at low PC cutoff values, as expected since at these low cutoff values virtually all target signal is suppressed as all PCs upward of the cutoff point dropped. As the cutoff point is moved to higher PCs, the algorithm performance improves steadily to a maximum SCR of 6.25 at PC number 250.

The solid line in Figure 9 shows the performance of the CMFsat algorithm, where instead of discarding certain portions of the data by suppressing a subset of PCs, all PCs are included with a more appropriate weighting distribution. In this case, all PCs beyond the saturation PC are treated with the same weight as the eigenvalue inverse at the saturation PC. Conceptually this states that the PCs beyond a certain range are mostly describing random instrument noise and that it is more realistic to treats these hyperspectral directions equally. The CMFsat results in Figure 9 support these statements as it can be seen that the CMFsat performance is overall better than or equal to either CMFcut or OBS at any point. Note also that the CMFsat result envelopes the results from each of the other two methods. This indicates that the CMFsat is a better model of how to treat the range of eigenvalues across the entire distribution of PCs. The peak value for CMFsat of 6.9 at PC number 40 is higher than the peak value for CMFcut of 6.25. However, the peak value is not the only measure of best performance. Note that the CMFsat result is higher than both CMFcut and OBS between PCs 50 and 250. This is an important fact since, for any of these three methods, a choice must be made of where to cutoff or saturate the eigenvalue weighting. The fact that CMFsat has a broader PC range of higher performace implies that making a mistake and choosing a PC number slightly away from the peak will still give very strong detection performance when compared to making a similar error with either OBS or CMFcut.

The primary reason the CMFsat algorithm performs better than OBS or CMFcut is that it does not completely discard a subset of the PCs (by setting their eigenvalues equal to zero). As mentioned earlier, the target signature will be virtually equally represented by the entire set of eigenvectors if the target makes only a small fraction (or not at all) of the scene. This is true for many situations and scenarios. Figure 10 illustrates this phenomenon by showing the correlation (normalized dot product) of the tetrachloroethylene spectral signature (Figure 2) with the eigenvectors of each of the principal components. A correlation value of 1.0 would indicate a perfect match between the particular eigenvector and the chemical library spectrum. A value of -1.0 would indicate a perfect anti-correlation. Figure 10 shows that with peak magnitudes near 0.2, not a single eigenvector well represents the target signature. This indicates that the target information is almost equally spanned by all eigenvectors, and that by discarding any of them would result in a loss of information.

4. RANK ESTIMATION

Any of the three methods presented above require a choice to be made as to where to make a change in the eigenvalue distribution, either suppressing or saturating various PCs. Wax and Kailath (1985) and Williams (1994) have shown the use of both the MDL and AIC techniques for estimating the rank of a particular dataset. The two methods are attempts to estimate the total number of uncorrelated signals present in the presence of white noise. This is very similar to the model defined by the CMFsat detection method where the initial principal components are weighted by the eigenvalues and the later

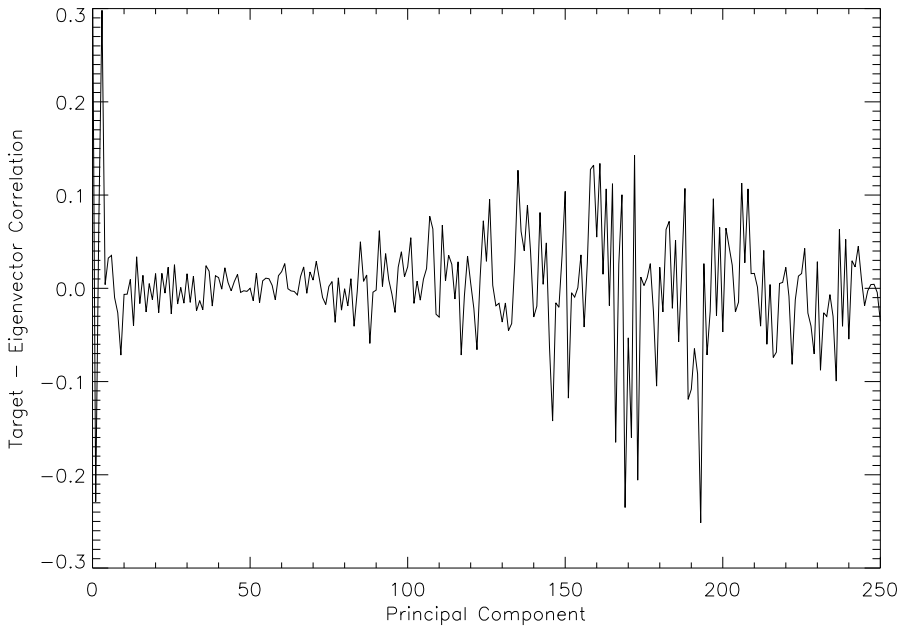

Figure 10. Target-eigenvector correlation for all principal components.

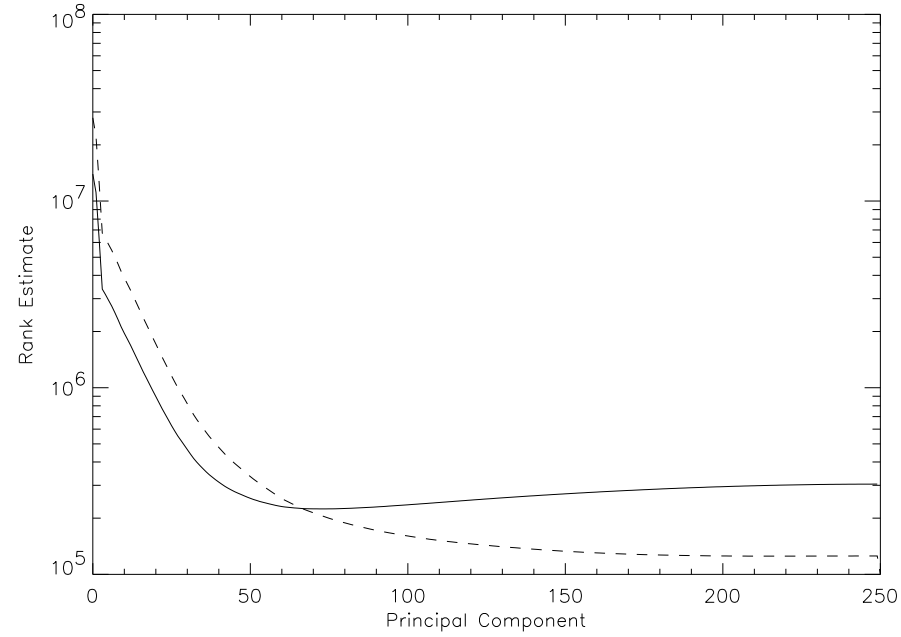

Figure 11. Rank estimates MDL (solid) and AIC (dashed) versus principal component.

PCs are set to uniform (white) weighting. It is thought that MDL or AIC can be used as an estimator for the point at which the eigenvalue distributions are saturated. The MDL and AIC measures are directly calculated from the number of samples (pixels), the number of spectral channels, and the eigenvalue distribution. The resulting plots of either MDL or AIC versus PC number ideally will show a minimum at the PC corresponding to the rank of the data. Examples of these calculations for the model dataset are shown in Figure 11 with the dotted line showing the AIC results and the solid line showing the MDL results. MDL has minimum at PC number 74 and AIC has a minimum at PC number 217. When comparing Figure 11 with Figure 9, it can be seen that the MDL minimum corresponds closely to the point at which the CMFsat method has peak performance.

This point also corresponds to the point where OBS also has maximum performance, which is to be expected since OBS is similar to CMFsat in that the later PCs are forced to a white noise-type uniform weighting. It is not clear why AIC indicates such a high value for the rank of this dataset.

5. CONCLUSIONS

The results presented in this report indicate that improved signal-to-clutter ratio for the detection of small targets or weak signals can be achieved by manipulating the higher-order data eigenvalues. This prevents uncalled-for emphasis of the search vector in the hyperspectral directions associated with the higher order principal components. When comparing the CMFsat results with either CMFcut or OBS, if a choice is to be made on which PCs to include in the analysis, it is clear that best results are obtained when they are all included, but with more appropriate eigenvalue weighting with the CMFsat algorithm. Positive results were shown with estimating the saturation point for CMFsat using the MDL rank estimator.

Further work remains to be done in applying these ideas to real instrument data in both the VIS/NIR and thermal IR regions of the spectrum. The model was based on a white noise sensor model, while real sensors do not behave in this fashion. A noise-whitening step will most likely have to be performed on real data to get maximum performance with the method described in this paper.

6. REFERENCES

Fukunaga, K. (1990), Introduction to Statistical Pattern Recognition, Academic Press.

Harsanyi, J.C., and C.I. Chang (1994), "Hyperspectral Image Classification and Dimensionality Reduction: an Orthogonal Subspace Projection Approach," IEEE Transactions in Geoscience and Remote Sensing, Vol. 32, pp. 779 - 785.

Stocker, A., I. Reed, and X. Yu, "Multi-Dimensional Processing for Electro-Optical Detection," *SPIE Signal and data Processing of Small Targets*, Vol. 1305, 1990.

Tu, T.M., C.H. Chen and C.I. Chang (1997), "A Posteriori Least Squares Orthgonal Subscape Projection Approach to desired Signature Extraction and Detection," IEEE Transactions in Geoscience and Remote Sensing, Vol.. 35, pp. 127 - 139.

Wax, M. and T. Kailath, "Detection of Signals by Information Theoretic Criteria," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-33, pp. 387 – 392, April, 1985.

Williams, B. "Counting of Degrees of Freedom When Using AIC and MDL to Detect Signals," *IEEE Transactions on Signal Processing*, Vol. 42, No. 11, pp. 3282 – 3284, Nov. 1994.