Locally Adaptive Constrained Energy Minimization for AVIRIS Images

By Dmitry Frolov and Randall B. Smith

Introduction

The Constrained Energy Minimization (CEM) technique (Harsanyi, 1993) has become quite popular in recent years as a means for constructing a linear operator to perform matched filtering of hyperspectral images. The CEM algorithm tries to maximize the response of the target spectral signature while suppressing the response of the unknown background signatures. This technique has its roots in signal processing theory and has long been used in electrical engineering applications. The CEM technique constructs an optimal linear operator that, when applied to a hyperspectral image, produces an abundance map for the every image pixel vector. The algorithm uses an estimate of the sample correlation matrix as a basis for determining the unknown background signatures and is computationally efficient. Application of the algorithm to a hyperspectral image usually provides strong suppression of the undesired signatures and enhances the contrast between target spectra and the background. These properties make the CEM algorithm a very popular processing technique as well as a common component of more complicated algorithms. However, the CEM approach has a number of significant shortcomings. It doesn't perform well in the presence of low probability background signatures, and an abundance image may contain a significant amount of noise and false alarms that make detection of small targets difficult. The standard CEM algorithm uses a single overall correlation matrix as a description of the undesired background signatures. Consequently, the performance of the algorithm is directly related to the complexity and diversity of the hyperspectral scene. Some of the existing modifications of the CEM technique try to improve the performance of the method by introducing spectral variability to the traditional single spectrum design (Farrand and Harsanyi, 1997; Boardman, 1998).

Approach

We propose a modification of the CEM algorithm that improves its ability to suppress low-probability undesired signatures and reduces the amount of noise and false alarms in the abundance image. The Locally Adaptive Constrained Energy Minimization (LA-CEM) technique tries to improve the performance of the CEM algorithm by introducing a spatial variability factor. The algorithm estimates a sample correlation matrix for every pixel vector in the hyperspectral image and uses it to compute a linear operator using the standard CEM formula. This approach allows the algorithm to adapt properties of the CEM operator to the image content around the pixel that is being processed.

The CEM optimal linear operator (Harsanyi, 1993) is given by:

$$W = R^{-1}d\,(d^{T}R^{-1}d)^{-1} \qquad (1)$$

where R^{-1} is the inverse of the sample correlation matrix of the observation pixel vectors and d is a target signature. A shortcoming of the standard CEM technique is that the linear operator W is the same for every pixel in the image. We introduce an extension to this technique that will allow the algorithm to improve its performance based on the local statistical properties of the image.

Consider establishing a grid of non-overlapping or overlapping sampling windows of size N x M pixels over the entire image. We can compute an estimate of the sample correlation matrix for every sampling window as well as its generalized pseudoinverse. Based on that information we can estimate a correlation matrix for every pixel in the image using a matrix interpolation scheme. The simplest interpolation scheme for two-dimensional data on a regular grid is a bilinear interpolation. The bilinear interpolation ensures that interpolated function values change continuously, but the gradient of the function changes discontinuously at the boundaries of each sampling window. However, this approximation is good enough for most applications considering its computational simplicity.

We can't interpolate the correlation matrix itself for a particular pixel, because we are actually after the inverse of the correlation matrix. Computing a generalized pseudoinverse for the every pixel in the image makes the algorithm completely impractical because of the extremely high computational expenses associated with it. The matrix inversion is considered an N^3 algorithm, so the performance of the algorithm will quickly go down as the number of bands in the image increases. Increasing the number of bands from 10 to 100 slows down the algorithm approximately 1000 times. We also can't apply a Sherman-Morrison formula for quick update of the matrix inverse, because it is applicable only to cases in which just a few matrix elements are changed. In order to solve this problem, we can consider interpolating not the correlation matrices themselves, but actual inverses of them. This approach might be viewed as an approximate method, because the matrix interpolated in the inverse space loses some of the properties associated with a correlation matrix. The corresponding matrix will still be symmetrical, but diagonal elements will be not equal to 1. Although such an interpolation is not physically correct, the error associated with it is not significant enough to offset the useful properties of the resulting CEM operator. However, the problem of estimating the actual amount of distortion introduced by this approach is probably worth a special investigation that is beyond the scope of the present study.

Another potential problem is that the scale of values produced by any particular CEM linear operator is based on the correlation matrix. The only output value that is fixed by the standard CEM operator results from the unity constraint that ensures an abundance value of 1 when the operator is applied to the target signature itself. The locally adaptive modification of the CEM uses a separate correlation matrix for every pixel in the image, so the raw output abundance values are not strictly comparable to each other. We have to introduce a common scaling or normalization procedure that brings different abundance values into a comparable range. The unity constraint places an upper limit on the output values, so the natural choice for the scaling procedure is to fix an output value in the lower range of abundance values. A simple scaling procedure that satisfies these requirements is given by:

$$A = 1 + (w * s - 1) / (2 - w * I) \qquad (2)$$

where A is an output abundance value, s is a current pixel vector, w is a linear CEM operator and I is an identity vector. This scaling procedure preserves a unity constraint and also ensures an abundance value of zero when the operator is applied to a pixel signature that is an "inverted" version of the target signature:

$$S_0 = I - d \qquad (3)$$

where S_0 is a signature that produces a zero abundance value and I is a unity vector.

Results

We applied the LA-CEM algorithm to the spectral subset 2.0 – 2.4 mm of the high-altitude AVIRIS image of the Cuprite area, Nevada. The target spectrum was obtained from the same image and closely resembled typical alunite spectra. The same spectrum was used to produce an abundance image using the traditional CEM approach. The abundance images are compared in Figure 1, and profiles across the images are shown in Figure 2. These results illustrate that the LA-CEM algorithm produces a much better contrast between target and background and

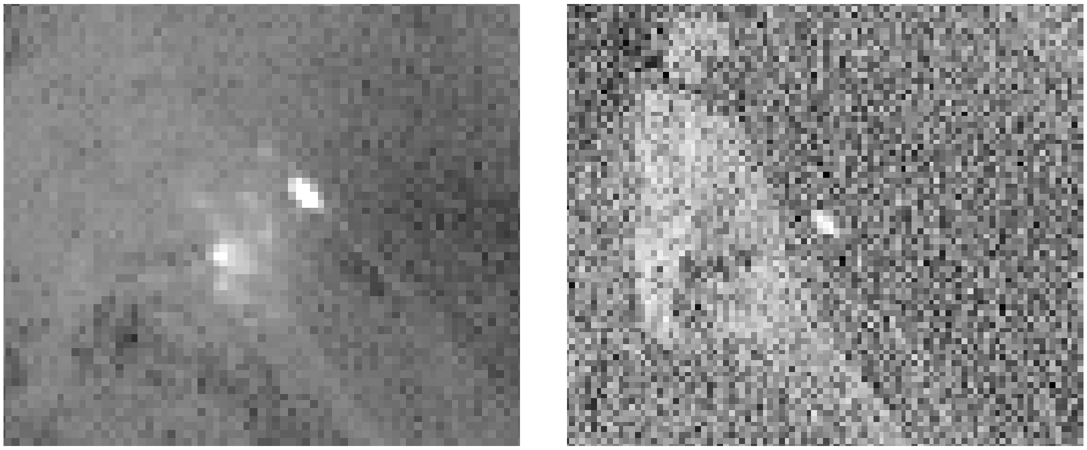

Figure 1. Comparison of LA-CEM (left) and traditional CEM (right) abundance images for the same area of an AVIRIS image of Cuprite, NV. Lighter tones indicate greater relative abundance of the target spectrum (alunite). Contrast of the images is enhanced for clarity. The LA-CEM image has less noise, and clearly resolves two small areas of high alunite abundance. The smaller alunite area is not evident in the traditional CEM abundance image.

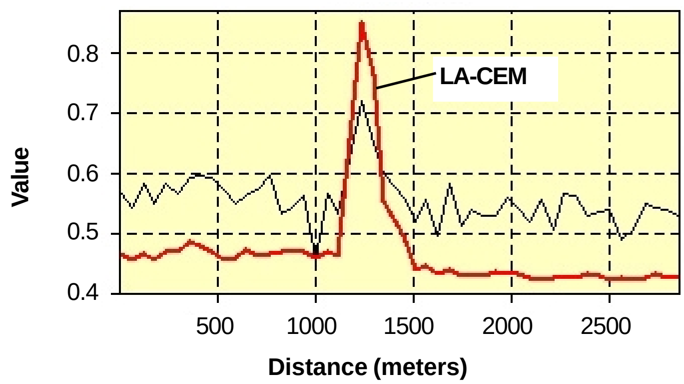

Figure 2. Profiles across LA-CEM and traditional CEM abundance images. The LA-CEM profile (thicker line) shows higher contrast and less noise in the background.

significantly reduces the amount of noise in the output. One small occurrence of alunite that is positively identifiable in the LA-CEM output is completely obscured by noise in the standard CEM output image, probably due to low-probability background signatures.

Discussion

The LA-CEM approach shows remarkable performance in detecting small (pixel and sub-pixel) targets obscured by complex backgrounds, but it doesn't perform well in image areas where the target spectrum becomes dominant or is present in significant amounts. This is completely understandable considering that the algorithm depends on local image statistics. The algorithm doesn't produce a good contrast between target and background in such areas, but the output still contains less noise than a traditional CEM abundance image. The average abundance values computed by LA-CEM for dominant signatures are significantly overestimated compared to a traditional CEM. There is a partial solution to this problem. The performance of the algorithm might be tuned by adjusting the size of the sampling window. The sampling window size should be proportional to the spatial dimension of the targets that the algorithm is looking for. As an approximation, the size of the window must be at least 5-10 times bigger than the typical size of the objects in order to provide good target differentiation. Another way to improve LA-CEM results is to use overlapping sampling windows in order to provide smooth transitions in the estimation of the CEM operator across the image. This also makes the algorithm less sensitive to the way that the sampling grid is superimposed on the ground features. The use of overlapping sampling windows is crucial in scenes with abrupt changes in spectral properties within the image, such as agricultural fields or coastal zones. The non-overlapping scheme might produce noticeable artifacts in the case where the division line between two windows falls near the border between different agricultural fields or near the coastline.

Another significant drawback of the CEM technique is that it is more sensitive to signal energy than to signal shape. The CEM operator produces different abundance values when applied to spectra of the same shape but with different constant components. As a result, the algorithm is vulnerable to the influence of topographic illumination effects and additive noise. In order to correct this problem, the pixel spectra might be converted into a different form that is less sensitive to the signal energy and is able to emphasize the shape of the spectra. Use of first derivatives might improve the results and produce less noisy abundance images (Woods, 1984).

Conclusions

The Locally Adaptive Constrained Energy Minimization (LA-CEM) algorithm has been successfully applied to a number of high altitude AVIRIS images and produces superior results in comparison with the traditional CEM method in detecting small targets obscured by complex backgrounds. The algorithm is suitable for detecting multi- and sub-pixel occurrences of non-dominant materials with a significantly reduced number of false alarm pixels and high contrast between target and background. The abundance images produced by the LA-CEM method are less plagued by noise and provide more spatially accurate mapping of the target materials than the traditional technique.

References

Boardman, J.W., 1998, "Leveraging the High Dimensionality of AVIRIS Data for Improved Sub-Pixel Target Unmixing and Rejection of False Positives: Mixture Tuned Matched Filtering," Seventh JPL Airborne Earth Science Workshop, January 12-16, 1998, pp. 55.

Farrand, W.H., and J.C. Harsanyi, 1997, "Use of a Modified Constrained Energy Minimization Technique to Map Ferruginous Sediments Along The Alamosa River, Colorado," Twelfth International Conference and Workshops on Applied Geologic Remote Sensing, Denver, Colorado, 17-19 November 1997, Vol.2, pp. 385-392.

Harsanyi, J.C., 1993, "Detection and Classification of Subpixel Spectral Signatures in Hyperspectral Image Sequences," Ph.D. Dissertation, University of Maryland, Baltimore County, 116 pp.

Woods, J.W., 1984, "Image Detection and Estimation" in "Digital Image Processing Techniques," edited by Ekstrom, P. M., Academic Press, Inc., Orlando, Florida, pp.80-81.

To be published in: Summaries of the Eighth JPL Airborne Earth Science Workshop, February 8-11, 1999, Volume 1. AVIRIS Workshop.

About the first author:
Dmitry Frolov is a software engineer with MicroImages, responsible for processes such as mosaic, surface modeling, and hyperspectral analysis. He graduated in 1988 with a diploma in Remote Sensing Engineering from the Department of Applied Astronautics of the Moscow Institute of Geodesy, Aerial Survey, and Cartography.