Preprint
UCRL-JC-141883

Automatic Extraction of Closed Pixel Clusters for Target Cueing in Hyperspectral Images

D. W. Paglieroni, D. E. Perkins

This article was submitted to 46th Annual Meeting of the International Symposium Optical Science and Technology
San Diego, CA
July 29 through August 3, 2001

U.S. Department of Energy
Lawrence Livermore National Laboratory

June 5, 2001

Approved for public release; further dissemination unlimited

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

Automatic Extraction of Closed Pixel Clusters for Target Cueing in Hyperspectral Images[1]

David W. Paglieroni and Dwight E. Perkins
Lawrence Livermore National Laboratory
P.O. Box 808, Livermore, CA, 94550

ABSTRACT

Traditional algorithms for automatic target cueing (ATC) in hyperspectral images, such as the RX algorithm, treat anomaly detection as a simple hypothesis testing problem. Each decision threshold gives rise to a different set of anomalous pixels. The clustered RX algorithm generates target cues by grouping anomalous pixels into spatial clusters, and retaining only those clusters that satisfy target specific spatial constraints. It produces one set of target cues for each of several decision thresholds, and conservatively requires $O(K^2)$ operations per pixel, where K is the number of spectral bands (which varies from hundreds to thousands in hyperspectral images).

A novel ATC algorithm, known as *Pixel Cluster Cueing (PCC)*, is discussed. PCC groups pixels into clusters based on spectral similarity and spatial proximity, and then selects only those clusters that satisfy target-specific spatial constraints as target cues. PCC requires only $O(K)$ operations per pixel, and it produces only one set of target cues because it is not an anomaly detection algorithm, i.e., it does not use a decision threshold to classify individual pixels as anomalies. PCC is compared both computationally and statistically to the RX algorithm.

Key Words: automatic target cueing (ATC), anomaly detection, Pixel Cluster Cueing (PCC) algorithm, RX algorithm

1. INTRODUCTION

Each pixel in an overhead hyperspectral image contains a wealth of spectral information. By identifying spectral signatures of targets of interest, it may be possible to detect targets in overhead hyperspectral images (with hundreds or thousands of spectral bands) that cannot be detected in single-band images or even multi-spectral images (with tens of spectral bands). Unfortunately, it is not computationally feasible to identify spectral signatures of interest in each pixel individually. Automatic target cueing (ATC) algorithms are typically much more computationally efficient than spectral identification algorithms, and they can be used to dramatically reduce the number of pixels that must be subjected to spectral signature identification. ATC becomes a factor operationally when the volume of hyperspectral imagery becomes too large for a team of image analysts to process manually.

Traditional algorithms for automatic target cueing in hyperspectral images, such as RX [7-8], stochastic expectation maximization (SEM) [3] and endmember-based algorithms [1-2], are based on anomaly detection. Anomaly detectors classify individual pixels as either of type "anomaly" or "background clutter". They produce one set of anomalies for each of several decision thresholds. Target cues can be formed by grouping anomalous pixels into spatial clusters, and retaining only those clusters that satisfy target specific spatial constraints. SEM is by far, the most computationally expensive of these algorithms (in fact, SEM is not computationally feasible for hyperspectral imagery). The RX algorithm is less expensive, but still requires $O(K^2)$ operations per pixel (by conservative estimates), where K is the number of spectral bands. Endmember-based algorithms can be somewhat more efficient.

A novel algorithm for automatic target cueing in hyperspectral images, known as *Pixel Cluster Cueing (PCC)* is discussed. PCC requires only $O(K)$ operations per pixel. It produces only one set of target cues because it is not an anomaly detection algorithm - it does not use a decision threshold to classify individual pixels as anomalies. Instead, PCC groups pixels into clusters based on spectral similarity and spatial proximity, and then specifies only those clusters that satisfy target-specific spatial constraints as target cues. The statistical and computational performance of PCC is compared to that of the

[1] This work was performed under he auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

benchmark RX algorithm. A brief review of the RX algorithm is given in Section 2. Section 3 describes the two stages of the PCC algorithm. Section 4 gives statistical performance metrics for automatic target cueing and compares PCC to RX by applying each algorithm to an AVIRIS scene.

2. OVERVIEW OF THE RX ALGORITHM

Traditional ATC algorithms treat anomaly detection as a simple hypothesis testing problem. The pixel spectrum $\boldsymbol{x}_{i,j}$ (a $Kx1$ column vector with samples $x_{i,j}(k)$, $k = 0,\ldots,K-1$) at row i and column j of a hyperspectral image with K spectral bands is classified as either of type "background clutter" or "anomaly". The null hypothesis H_0 is the anomaly absent hypothesis and the alternative hypothesis H_1 is the anomaly present hypothesis. Classification is based on a decision rule in which the decision threshold can either be specified manually or selected automatically. One can vary the number of pixels classified as anomalies by varying the decision threshold.

A given anomaly detection algorithm is often assessed by comparing its performance to that of the RX algorithm. The RX algorithm is a traditional industry standard algorithm developed by Reed, Yu and Stocker in the early 1990's for detecting anomalies in hyperspectral images [7-8]. It is motivated by the concept of matched filtering. The decision rule for RX anomaly detection is

$$|y_{i,j}| \stackrel{\Delta}{=} |\tilde{\boldsymbol{x}}_{i,j}^T \boldsymbol{R}^{-1} \tilde{\boldsymbol{x}}_{i,j}| \underset{H_0}{\overset{H_1}{\gtrless}} y_0 \tag{1}$$

where $y_{i,j}$ is the alarm statistic, $\tilde{\boldsymbol{x}}_{i,j}$ is typically $\boldsymbol{x}_{i,j}$ minus an image-average pixel spectrum, $\boldsymbol{R}$ is typically the KxK sample covariance matrix derived from some or all of the hyperspectral image pixels, and y_0 is the decision threshold. Since it requires $O(K^2)$ operations per pixel to compute $\boldsymbol{R}$ plus $O(K^2)$ operations per pixel to compute $y_{i,j}$ plus $O(K^3)$ operations per image to compute $\boldsymbol{R}^{-1}$ from $\boldsymbol{R}$, the computational complexity of the RX algorithm can be conservatively assessed as $O(K^2)$ operations per pixel. This represents an appreciable computational burden because for hyperspectral images, the number of spectral bands (K) typically ranges from several hundred to several thousand.

Other traditional industry standard algorithms include Stochastic Expectation Maximation (SEM) [3] and endmember-based algorithms [1-2]. For SEM, the number of covariance matrix inversions per iteration is equal to the number of spatial regions that the image was segmented into, which can be large. SEM is thus likely to be computationally feasible for multi-spectral images, but not hyperspectral images. Endmember techniques are feasible for anomaly detection in hyperspectral images, but due to scope constraints, will not be discussed further here.

3. PIXEL CLUSTER CUEING (PCC) ALGORITHM

Pixel Cluster Cueing (PCC) is a novel cluster-based algorithm for automatic target cueing in hyperspectral images [6]. Its computational complexity is typically $O(K)$ operations per pixel, so its speed can be expected to vary in direct proportion to the number (as opposed to the square of the number) of spectral bands. PCC is not an anomaly detection algorithm - it does not apply decision rules and decision thresholds to individual pixels in order to classify them as anomalies. Instead, PCC groups pixels into clusters based on spectral similarity and spatial proximity, and then specifies only those clusters that satisfy target-specific spatial constraints as target cues.

3.1 Percentile Spectrum Analysis (PSA)

In order to assign pixels to clusters based on spectral similarity, it is necessary to establish an appropriate criterion for spectral similarity. The *test of spectral homogeneity*

(2) $$|x(k) - x_0(k)| < T(k) \quad \forall k$$

defines two pixels $x(k)$ and $x_0(k)$ to be spectrally similar if the magnitude of their *difference spectrum* is less than some *similarity threshold spectrum* $T(k)$. The *Percentile Spectrum Analysis (PSA)* algorithm estimates $T(k)$ from the statistics of hyperspectral image pixel spectra [5-6]. The first stage of PCC involves estimating the threshold of spectral similarity between pixels. When the image is large (as in broad area search) PSA is only applied to small representative blocks of image pixels. In this case, the computational cost of PSA is small relative to PCC as a whole, and can thus be discounted as an initialization cost.

Let $\dot{x}_{i,j}(k)$ be the spectral difference between $x_{i,j}(k)$ and its predecessor on the same row or column. The image of difference spectra contains spectral bands that are high-pass filtered versions of spectral bands from the original hyperspectral image (see Fig.1). PSA derives $T(k)$ from statistics of $\dot{x}_{i,j}(k)$. PSA is originally formulated on the basis of a theoretical statistical model for $\dot{x}_{i,j}(k)$, and is then extended to practical cases for which the model is not completely valid.

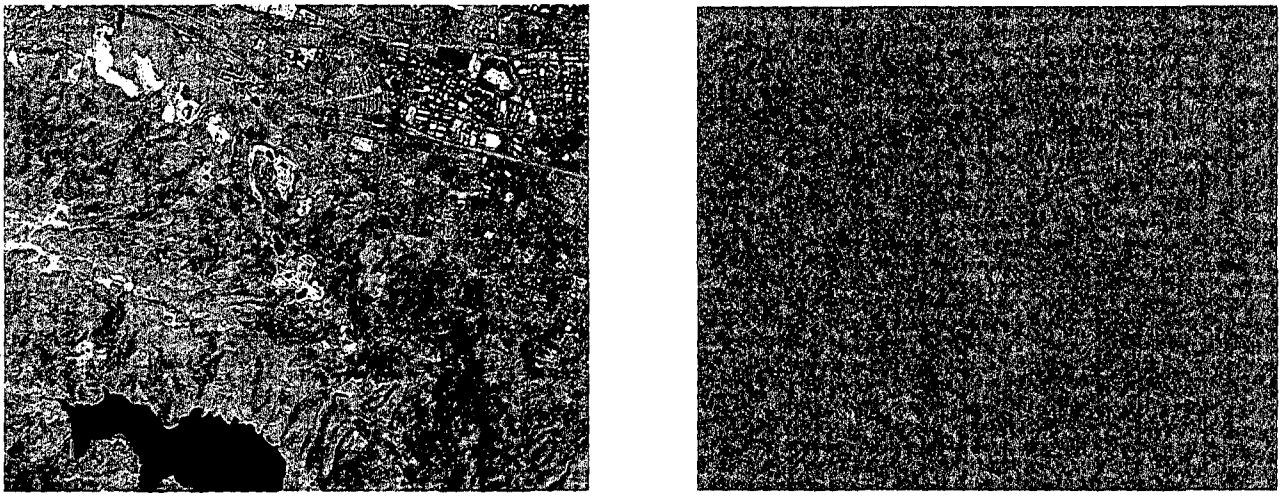

(a) (b)

Fig.1 (a) Broadband image (average of spectral bands) for AVIRIS hyperspectral image with 200 bands, 614 columns and 512 rows. (b) Broadband image for AVIRIS image of difference spectra.

$|\dot{x}_{i,j}(k)|$ is taken to be a sample of some random variable $\dot{X}(k)$. For spatially uniform scenes, almost all pixels are spectrally similar (i.e., there is no clutter), so $\dot{x}_{i,j}(k)$ typically represents deviation in spectral band k between pixels that are spectrally similar. Let us postulate a statistical model for images of difference spectra for spatially uniform scenes in which the random variables associated with the different spectral bands are positively scaled versions of one another. Mathematically, for any two spectral bands k and l, let us postulate that there exists a scale factor $a_{kl} > 0$ such that

(3) $$\dot{X}(k) = a_{kl}\dot{X}(l)$$

In this case, the analytical form of the PDF does not vary between bands, but the ratio of the standard deviation between bands k and l is equal to a_{kl}. Note that the statistical model of equation (3) does *not* assume that the PDF's have any particular analytical form.

Let $\dot{x}_p(k)$ be the *pth* percentile of $\dot{X}(k)$. $\dot{x}_p(k)$ will be referred to as the *pth percentile spectrum* associated with spectral band k (see Fig.2a for examples). It can be shown that

$$\dot{x}_p(k) / \dot{x}_p(l) = a_{kl} > 0 \tag{4}$$

from which it can be seen that the *scale identity for percentile spectra*

$$\dot{x}_p(k) / \dot{x}_q(k) = b_{pq} > 0 \tag{5}$$

is satisfied [5-6]. These equations establish the *scaling property of percentile spectra*:

> *The random variables associated with the different spectral bands in the magnitude of the image of difference spectra for a spatially uniform scene are scaled versions of one another if and only if the percentile spectra associated with the different percentiles are scaled versions of one another.*

Thus, if the percentile spectra measured for an image of magnitudes of difference spectra are close to being positively scaled versions of one another, then the random variables associated with the different spectral bands are also close to being positively scaled versions of one another.

Let $\dot{\boldsymbol{x}}_p \triangleq [\dot{x}_p(0),...,\dot{x}_p(K-1)]^T$ be the column vector associated with percentile spectrum $\dot{x}_p(k)$. The scale factor b_{qp} of best linear fit between $\dot{\boldsymbol{x}}_p$ and $\dot{\boldsymbol{x}}_q$ minimizes $\| \dot{\boldsymbol{x}}_q - b_{qp} \dot{\boldsymbol{x}}_p \|$ and is given by

$$b_{qp} = (\dot{\boldsymbol{x}}_p^T \dot{\boldsymbol{x}}_q) / (\dot{\boldsymbol{x}}_p^T \dot{\boldsymbol{x}}_p) \tag{6}$$

which yields a fractional error of

$$\varepsilon_p \triangleq \frac{\| \dot{\boldsymbol{x}}_q - b_{qp} \dot{\boldsymbol{x}}_p \|}{\| \dot{\boldsymbol{x}}_q \|} = \left[1 - \frac{(\dot{\boldsymbol{x}}_p^T \dot{\boldsymbol{x}}_q)^2}{(\dot{\boldsymbol{x}}_p^T \dot{\boldsymbol{x}}_p)(\dot{\boldsymbol{x}}_q^T \dot{\boldsymbol{x}}_q)} \right]^{1/2} \tag{7}$$

One can generate the *shape deviation profile for percentile spectra*, i.e., the plot of ε_p vs. p for $p > q$ and some relatively low *nominal percentile* q (say $q = 10$). $\varepsilon_p = 0$ for perfect similarity in spectral shape. The postulated statistical model of images of difference spectra for spatially uniform scenes (see equation (3)) is reasonable to the extent that ε_p is small and relatively constant for $p > q$.

Suppose that in addition to a hyperspectral image of a spatially cluttered scene, one also has access to a hyperspectral image of a spatially uniform scene acquired with the same instrument under similar conditions. In this case, one could measure shape deviation profiles ε_p for the spatially cluttered scene and $\varepsilon_{p,0}$ for the spatially uniform scene. Independent of the precise statistical behavior of the difference spectra (i.e., the precise validity of equation (3)), one would expect to be able to estimate a *similarity percentile* $p^* > q$ such that $\varepsilon_p - \varepsilon_{p,0} \approx 0$ for $q < p \leq p^*$ and $\varepsilon_p - \varepsilon_{p,0} > 0$ for $p > p^*$.

Note that $\varepsilon_{p,0} \approx 0$ for $p > q$ if and only if the random variables $\dot{X}(k)$ are close to being scaled versions of one another. p^* is an estimate of the percentage of pairs of adjacent pixel spectra used in computing difference spectra that are appropriately deemed similar for the spatially cluttered hyperspectral image of interest.

Unfortunately, it may often be the case that one has access to a spatially cluttered hyperspectral image, but *not* to a corresponding spatially uniform hyperspectral image. In this case, p^* can be estimated by assuming that independent of the

precise statistical behavior of the difference spectra, $\dot{\varepsilon}_p \triangleq d\varepsilon_p/dp$ can be expected to be relatively small for $q < p \leq p^*$ and larger for $p > p^*$. Moreover, for $p < p^*$, one can expect ε_p to be similar to the shape deviation profile associated with the spatially uniform hyperspectral image that would have been supplied had it been available. The assumption that $\dot{\varepsilon}_p$ will be relatively small for $q < p \leq p^*$ and larger for $p > p^*$ is certainly valid if the random variables $\dot{X}(k)$ corresponding to the spatially uniform hyperspectral image are roughly scaled versions of one another, in which case $\dot{\varepsilon}_p \approx 0$ for $q < p \leq p^*$. More generally, this assumption will be valid to the extent that scene clutter disrupts the scale relationship between percentile spectra to a larger degree than the noise present in the spectral bands (i.e., the spatially cluttered hyperspectral image is not noise limited). See Fig.2b for an example.

The foregoing discussion suggests that for spatially cluttered scenes, one may choose to estimate the similarity percentile p^* by determining the percentile at which ε_p begins to increase at an accelerated rate. The similarity percentile might, for example, correspond to a knee in the plot of ε_p vs. p. One way to estimate p^* is to use the rule

$$p^* = \text{smallest } p\text{: } \dot{\varepsilon}_p > 2\bar{\dot{\varepsilon}}_p \tag{8}$$

where $\bar{\dot{\varepsilon}}_p$ is the mean of $\dot{\varepsilon}_p$ values for percentiles ranging from q to p (see Fig.2b).

Since the shape of the *similarity percentile spectrum* $\dot{x}_{p^*}(k)$ mimics the shape of the noise spectrum, it is reasonable to represent $T(k)$ as

$$T(k) = \alpha \cdot \dot{x}_{p^*}(k) \tag{9}$$

The parameter $\alpha > 0$ will be referred to as the *similarity threshold scale factor*. $\alpha > 0$ may be chosen such that exactly p^* percent of all difference spectra satisfy the relation $|\dot{x}_{i,j}(k)| < T(k)$, for all k of interest. This ensures that exactly p^* percent of all difference spectra pass the test of spectral homogeneity (equation (2)). To accomplish this, for all pixels (i,j) and over all spectral bands k, compute

$$(\varepsilon_{i,j})_{max} = \max_{k=0,\ldots,K-1} \{ \varepsilon_{i,j}(k) \triangleq |\dot{x}_{i,j}(k)| / \dot{x}_{p^*}(k) \geq 0 \} \tag{10}$$

Now let ε^* be the p^* percentile of $(\varepsilon_{i,j})_{max}$ over all pixels. Also, define Ω to be the set consisting of all ordered pairs (i,j) such that $(\varepsilon_{i,j})_{max} < \varepsilon^*$ (this set contains p^* percent of all pixels). Then there exists a similarity threshold scale factor α in equation (9), namely

$$\alpha = \varepsilon^* \tag{11}$$

such that $|\dot{x}_{i,j}(k)| < T(k)$ for all $(i,j) \in \Omega$ and $|\dot{x}_{i,j}(k)| \geq T(k)$ for all $(i,j) \notin \Omega$ (see Fig.2c for an example).

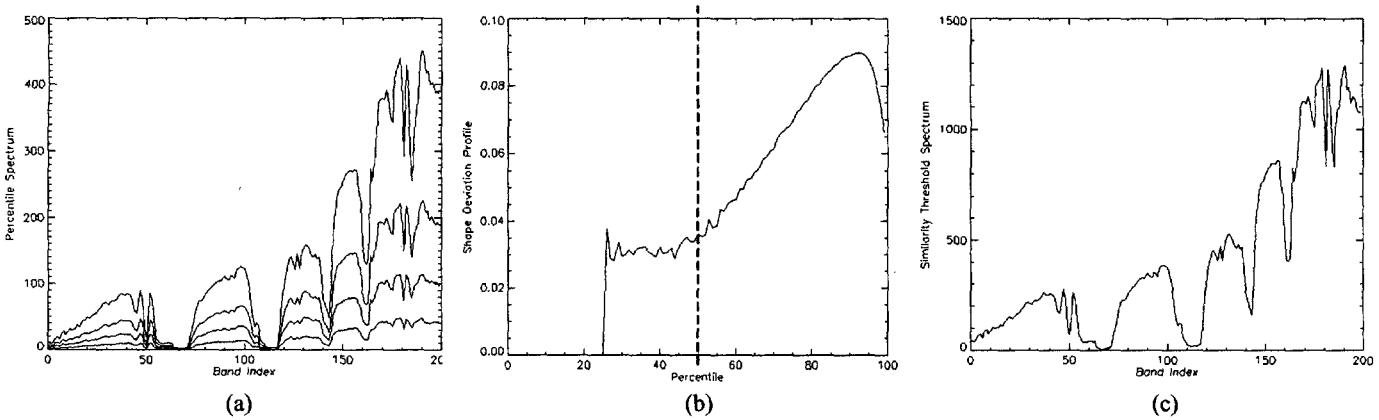

(a) (b) (c)

Fig.2 (a) Percentile spectra (p=20, 40, 60 and 80) and (b) shape deviation profile with p*=50 (dashed line) and (c) similarity threshold spectrum for the image in Fig.1.

3.2 Pixel Clusters

PCC is based on the premise that targets of interest (such as military vehicles scattered across a broad area of interest) can often be modeled as clusters of spatially connected pixels that differ spectrally from background pixels, and that for a given type of target, these clusters have characteristic spatial properties. The second stage of PCC involves pixel clustering. PCC performs clustering in two steps. In the first step, selected *primary clusters* (clusters of individual pixel spectra) are extracted from the original hyperspectral image. In the second step, cued *secondary clusters* (clusters of selected primary clusters) are extracted from the binary image of selected primary clusters. Both types of clusters are obtained by region growing based on local search [4,9]. The local search operation for primary clustering is

$$\Omega_1(i,j) \triangleq \{ (m,n) : \text{ } (m,n) \text{ has not yet been assigned to a cluster and } |x_{m,n}(k) - x_{i,j}(k)| < T(k) \ \forall k \text{ and } |m-i|, |n-j| \le w \} \quad (12)$$

w (typically *1* or *2*) is a *connectivity relaxation parameter* which specifies the maximum allowed spatial separation between nearest neighbors in a cluster. The local search operation for primary clustering requires $O(nK)$ operations per pixel, where the number of spectral comparisons (n) varies from *0* to $(2w+1)^2-1$. For small w, this translates to $O(K)$ operations per pixel. Note that it is not necessary to continue evaluating the spectral difference between two pixels once a spectral band k has been found for which the magnitude of the spectral difference exceeds $T(k)$. The local search operation for secondary clustering is

$$\Omega_2(i,j) \triangleq \{ (m,n): y_{m,n} = 1 \text{ and } |m-i|, |n-j| \le w \} \quad (13)$$

where $y_{i,j}$ is initially the binary image of selected primary cluster pixels (i.e., $y_{i,j} = 1$ for selected primary cluster pixels, and $y_{i,j} \leftarrow 0$ once pixel (i,j) has been added to a secondary cluster). The cost of secondary clustering is negligible relative to primary clustering.

A pixel cluster is said to be *complete*, *full-grown* or *closed* when it is closed under the mathematical operation of local search. Closed clusters are distinct (each pixel belongs to exactly one closed cluster), disjoint (non-overlapping) and unique (for a given image and $T(k)$, there is one and only one set of closed clusters). A pixel cluster is *spectrally homogeneous* if no two cluster pixels are separated by more than $2T(k)$ (i.e., they are all within $T(k)$ of some spectrum). The spectral homogeneity of a pixel cluster can be quickly ascertained as the cluster is grown.

The spatial properties of a pixel cluster can be measured by computing the values of certain spatial features, such as size (the number of pixels in a cluster) and *spatial cohesion*. The spatial cohesion κ of a cluster is a number that varies from *0* to *1* ($\kappa = 1$ for clusters that are as spatially cohesive as possible, i.e. circles). The spatial cohesion of a cluster C of size m

can be computed as the ratio of the spatial variance of the cluster of size m that is as cohesive as possible and the spatial variance of C.

The selected primary clusters are closed and must satisfy primary target specific spatial constraints. Spectral homogeneity can be enforced if so desired. Since selected primary clusters may correspond to complete targets or to pieces of targets, it is important for the primary constraints to be consistent with the spatial properties not only of complete targets, but also pieces of targets. Primary target specific spatial constraints thus tend to be relaxed. In particular, since pieces of targets can be very small, it may often be appropriate to choose a lower bound of 1 on the primary size constraint, even if a complete target covers a much larger number of pixels (the upper bound should still be close to the target size).

The cued secondary clusters are closed and must satisfy secondary target specific spatial constraints. Cued secondary clusters may correspond to complete targets, but not to pieces of targets. Secondary constraints are thus typically more restrictive than primary constraints. For example, both the lower and upper bounds on secondary cluster size should be close to the size of the target. This combination of primary and secondary constraints allows pieces of spectrally non-homogeneous targets to be assimilated into a single cluster, while allowing for the possibility that the target might actually be spectrally homogeneous.

It is important to note that the idea of using clusters to assist in automatic target cueing is not unique to PCC. For example, the *basic RX algorithm* treats target cues as individual anomalous pixels (see section 2). However, the *clustered RX algorithm* forms spatial clusters of anomalous pixels from basic RX, and then selects only those clusters that satisfy the secondary target specific spatial constraints from PCC.

4. PERFORMANCE METRICS AND EXPERIMENTAL RESULTS

Performance metrics for automatic target cueing are designed to reflect how well the ATC algorithm facilitates target detection. They have historically been based on the probability of target detection (P_D) and the number of false alarms (N_{FA}). To estimate these quantities, one must first provide appropriate definitions for "alarm" and "target" in the context of hyperspectral images.

In the context of hyperspectral images, a target is a cluster of pixel spectra that corresponds to the spatial location of an object of interest. An alarm (or target cue) is a cluster of pixel spectra that satisfies certain spectral constraints and target specific spatial constraints. The pixels in an alarm must be aggregated into one or more spectra, which are then subjected to detailed analysis so that the spectral signature of the target of interest can be identified.

One way to measure N_{FA} is to add *false alarm scores* $(S_{FA})_i \in [0,1]$ computed independently for each of N_A alarms $i = 1,\ldots,N_A$, i.e.,

$$N_{FA} = \sum_{i=1}^{N_A} (S_{FA})_i \tag{14}$$

This generalized definition of false alarm rate is not necessarily integer-valued. Each alarm, regardless of the size of the pixel cluster, is weighted the same. The false alarm score is the fraction of alarm pixels that are also target pixels, i.e.,

$$(S_{FA})_i \triangleq \frac{\textit{number of pixels from alarm } i \textit{ that are not pixels from any target}}{\textit{number of pixels from alarm } i} \tag{15}$$

When an alarm is defined to be an individual pixel (as in basic RX), $(S_{FA})_i$ will be 0 if pixel i is a target pixel, and 1 otherwise. In this case, N_{FA} as defined in equation (14) degenerates to the number of alarm pixels that are not target pixels, which is the usual definition.

An analogous way to measure P_D is to average *target detection scores* $(S_D)_i \in [0,1]$ computed independently for each of N_T targets $i = 1,...,N_T$, i.e.,

$$P_D = \frac{1}{N_T} \sum_{i=1}^{N_T} (S_D)_i \tag{16}$$

Each target, regardless of the size of the pixel cluster, is weighted the same. In order to provide an appropriate definition for target detection score, it is helpful to distinguish between alarm clusters and alarm spectra:

1. *Alarm clusters* are target cues. For basic RX, alarm clusters are individual pixels of type "anomaly". For clustered RX, alarm clusters are clusters of pixels of type "anomaly" from basic RX that satisfy target specific spatial constraints. For PCC, alarm clusters are cued secondary clusters. For consistency in computing the false alarm rate, alarm clusters for PCC and clustered RX are chosen to satisfy the same set of target specific spatial constraints.

2. *Alarm spectra* are the spectra subjected to detailed analysis, subsequent to target cueing, for identification of the target spectral signature. For basic RX, alarm spectra are the spectra associated with individual pixels of type "anomaly". For clustered RX, alarm spectra are typically non-uniform aggregations of spectra associated with pixels from the same alarm cluster. For clustered RX, an alarm cluster thus has exactly one alarm spectrum. For PCC, alarm spectra are uniform aggregations of spectra associated with pixels from the same cued primary cluster (i.e., a selected primary cluster contained within a cued secondary cluster). For PCC, an alarm cluster may thus have more than one alarm spectrum.

One could choose to assign a target detection score of *1* to a target only if there is at least one alarm spectrum whose constituent pixels all belong to the target cluster. On the other extreme, one would assign a target detection score of *0* to a target only when there is no overlapping alarm cluster. In general, one could choose the target detection score to be the fraction of alarm pixels that are also target pixels, for the alarm spectrum whose constituent pixels afford the largest fractional overlap with the target cluster:

$$(S_D)_i \triangleq \max_{j=1,...,N_A'} \frac{\textit{number of pixels from alarm spectrum } j \textit{ that are pixels from target } i}{\textit{number of pixels from alarm spectrum } j} \tag{17}$$

where N_A' is the number of alarm spectra. When an alarm is defined to be an individual pixel (as in basic RX), $(S_D)_i$ will be *1* if pixel i is a target pixel, and *0* otherwise. In this case, P_D as defined in equation (16) degenerates to the number of anomaly pixels that are also target pixels divided by the number of target pixels, which is the usual definition.

The statistical performance of basic RX is captured in its *receiver operating characteristic (ROC) curve* (a plot of P_D vs. N_{FA}). P_D and N_{FA} are both non-increasing functions of the RX decision threshold (y_0 in equation (1)). For clustered RX, P_D will not necessarily be a non-decreasing function of N_{FA}, so for clustered RX, (P_D, N_{FA}) vs. y_0 is perhaps more appropriately presented as a ROC scatter plot. For PCC, there is only one ROC point (P_D, N_{FA}) because there is no decision threshold.

A target cueing experiment was conducted using the AVIRIS image in Fig.1a and the target mask in Fig.3. The target mask was created to capture large buildings that are nearly square, ranging in size from 4x4 to 16x16 pixels. PCC was compared to both basic and clustered RX. Computationally, PCC was found to run roughly 80 times faster than RX (8 seconds for PCC vs. 628 seconds for RX on the same 450MHz Pentium II processor). The primary constraints for PCC included a size range from 1 to 256 pixels, and a cohesion range from 0.5 to 1.0. The secondary constraints for PCC included a size range from 16 to 256 pixels, and a cohesion range from 0.8 to 1.0. Fig. 4 shows the alarm mask for PCC. Fig. 5 shows the alarm masks for basic RX and decision thresholds corresponding to signal-to-clutter ratios (SCR's) of 1, 3 and 5 (see equation (2)). Fig. 6 shows the alarm masks for clustered RX obtained by applying the secondary constraints of PCC to the alarm masks for basic RX in Fig. 5. Fig.7 shows plots of the number of targets detected vs. the number of false alarms for

each of the 3 methods. None of the anomaly detectors cued on more than about half of the targets. The detection rate for PCC was comparable to the basic and clustered RX detection rates associated with the most favorable decision thresholds. The false alarm rate for clustered RX was much lower than for basic RX. The false alarm rate for PCC was much lower than for basic RX, but modestly higher than for clustered RX. However, PCC ran almost two orders of magnitude faster than RX, and the alarm mask for PCC is unambiguous (i.e., there is only one PCC alarm mask).

Fig.3 Target mask corresponding to the image in Fig.1a.

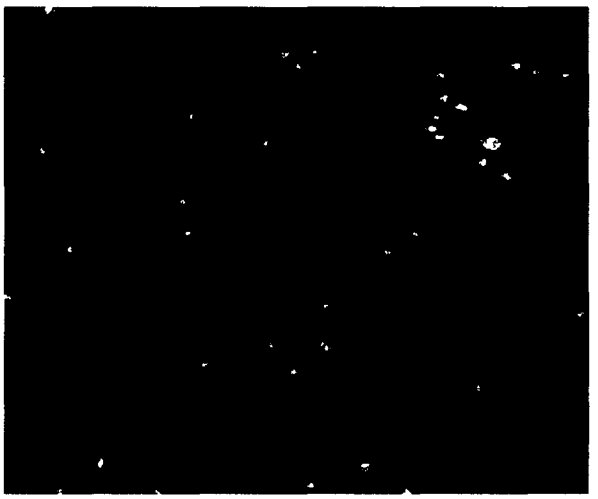

Fig.4 Alarm mask for PCC.

5. CONCLUSIONS AND FUTURE DIRECTIONS

By viewing alarms as clusters of pixels, rather than individual pixels, it is possible to significantly reduce the false alarm rate of an ATC algorithm. In particular, the false alarm rate can be much lower for clustered RX than for basic RX. However, both versions of RX conservatively require $O(K^2)$ operations per pixel, where K is the number of bands in the hyperspectral image (typically ranging from several hundred to several thousand). The statistical performance of PCC can be competitive with that of clustered RX, but PCC typically runs orders of magnitude faster because its computational cost is only $O(K)$ operations per pixel. Furthermore, unlike RX, which produces one alarm mask for each of several decision thresholds, PCC produces only one alarm mask.

One problem with PCC, as currently specified, is that it sometimes misses targets that one would expect to detect. The authors have concluded that this can be partially attributed to two things. First, PCC needs to be slightly modified such that primary clusters satisfying secondary target specific spatial constraints are always selected as target cues, even if they belong to secondary clusters that do not satisfy those secondary constraints. Second, the primary target specific spatial constraints need to be specified more rigorously. For example, for the targets in Fig.3, the cohesion constraint should be made less restrictive as the primary clusters decrease in size from 256 down to 1.

REFERENCES

[1] J. Boardman, "Analysis, Understanding and Visualization of Hyperspectral Data as Convex Sets in n-Space", Proc. SPIE, vol. 2480, 1995, pp.14-22.

[2] J. Bowles, P. Palmadesso, J. Antoniades and M. Baumback, "Use of Filter Vectors in Hyperspectral Data Analysis", Proc. SPIE, vol. 2553, 1995, pp.148-157.

[3] Pascale Masson and Wojciech Pieczynski, "SEM Algorithm and Unsupervised Statistical Segmentation of Satellite Images", IEEE Trans. Geosci. Remote Sens., vol. 31, no. 3, May 1993, pp.618-633.

[4] J. L. Muerle and D. C. Allen, "Experimental Evaluation Techniques for Automatic Segmentation of Objects in a Complex Scene", in Pictorial Pattern Recognition (G. C. Cheng et al., Eds.), Thompson, Washington, 1968.

[5] David W. Paglieroni and Randy S. Roberts, "Optimal Segmentation Strategy for Compact Representation of Hyperspectral Image Cubes", Proc. ASPRS DC 2000 Conf., May 22-27, 2000.

[6] David W. Paglieroni and Dwight E. Perkins, "Spatial Constraints on Segmented Regions for Target Cueing in Hyperspectral Images", Proc. Military Sensing Symposium on Passive Sensors, March 5-7, 2001, Vienna, VA.

[7] Irving S. Reed and Xiaoli Yu "Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution", IEEE Trans. ASSP, vol. 38, no. 10, October 1990, pp.1760-1770.

[8] Alan D. Stocker, Irving S. Reed and Xiaoli Yu, "Multidimensional Signal Processing for Electro-Optical Target Detection", Proc. SPIE, vol. 1305, April 1990, pp.218-231.

[9] S. W. Zucker, "Region Growing: Childhood and Adolescence", CGIP, vol. 5, no. 3, September 1976, pp.382-399.

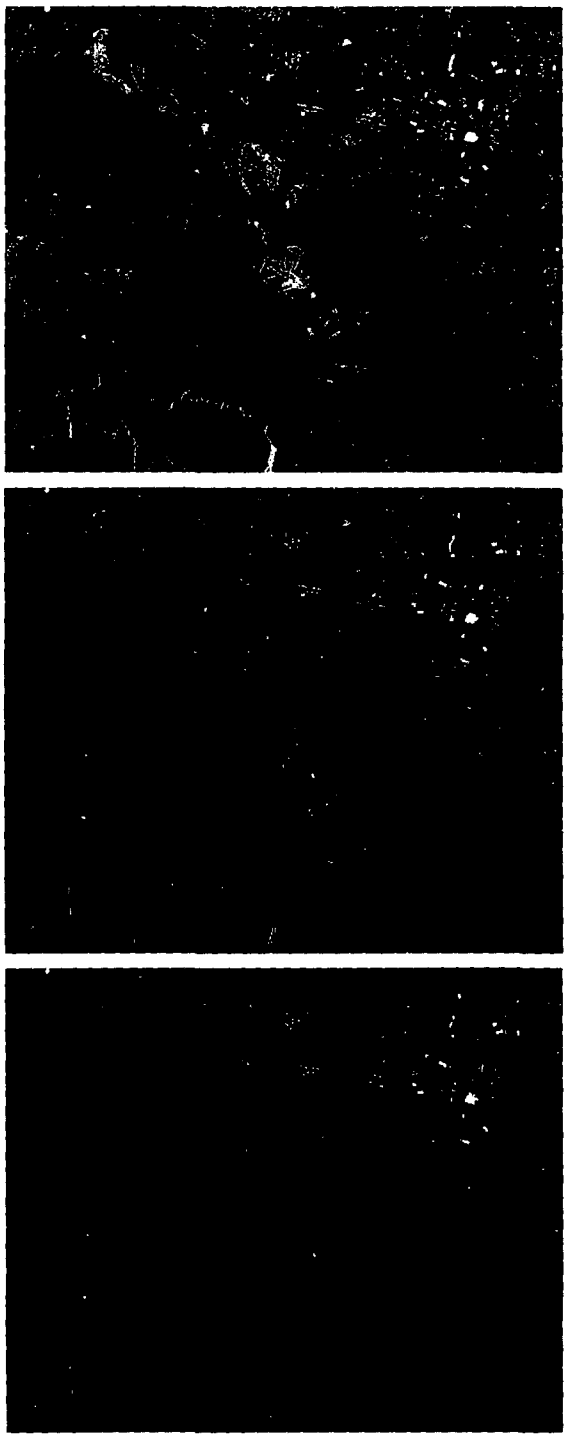

Fig.5 Alarm masks for basic RX corresponding to SCR's of 1, 3 and 5 (top to bottom).

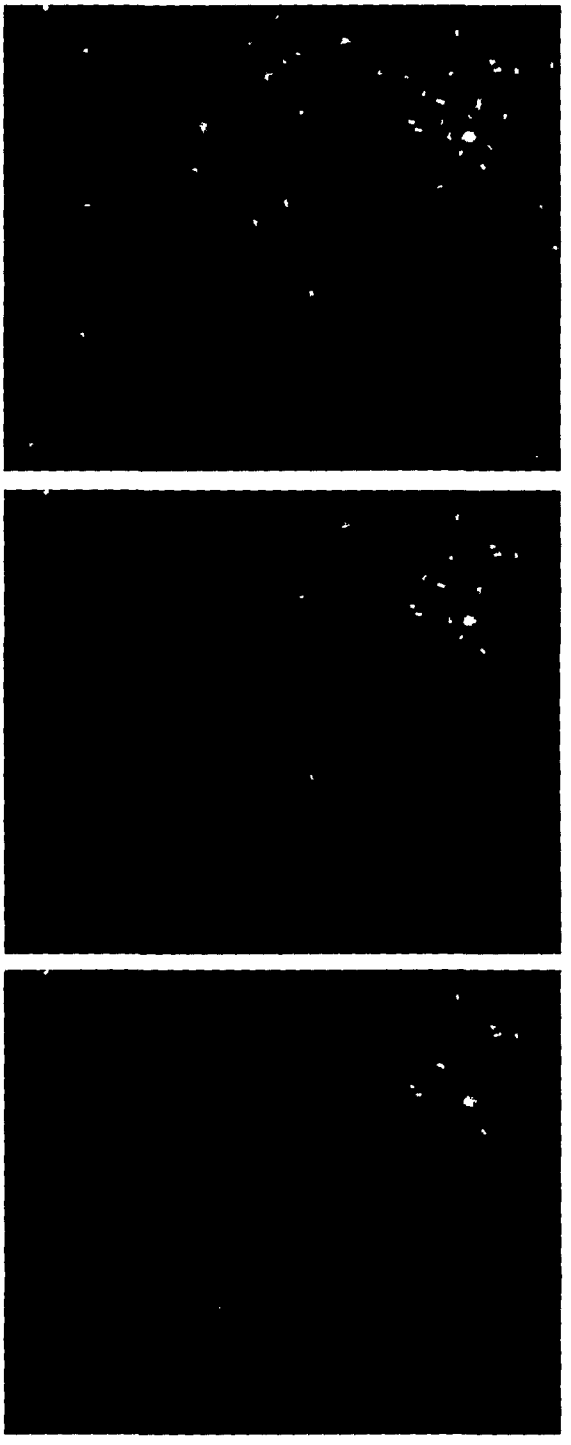

Fig.6 Alarm masks for clustered RX corresponding to the basic RX alarm masks in Fig.5.

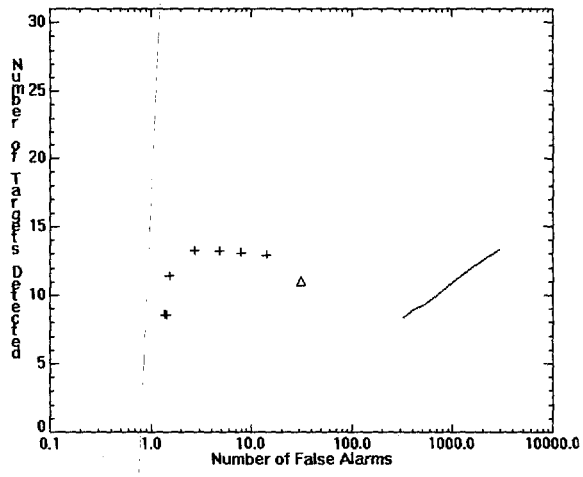

Fig.7 Number of targets detected vs. number of false alarms for basic RX (line), clustered RX (pluses) and PCC (triangle).

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551