A Cluster-Based Approach for Detecting Man-Made Objects and Changes in Imagery

Mark J. Carlotto, *Senior Member, IEEE*

Abstract—A new unified approach to object and change detection is presented that involves clustering and analyzing the distribution of pixel values within clusters over one or more images. Cluster-based anomaly detection (CBAD) can detect man-made objects that are: 1) present in a single multiband image; 2) appear or disappear between two images acquired at different times; or 3) manifest themselves as spectral differences between two sets of bands acquired at the same time. Based on a Gaussian mixture model, CBAD offers an alternative to compute-intensive, sliding-window algorithms like Reed and Yu's RX-algorithm for single-image object detection. It assumes that background pixel values within clusters can be modeled as Gaussian distributions about mean values that vary cluster-to-cluster and that anomalies (man-made objects) have values that deviate significantly from the distribution of the cluster. This model is valid in situations where the frequency of occurrence of man-made objects is low compared to the background so that they do not form distinct clusters, but are instead split up among multiple background clusters. CBAD estimates background statistics over clusters, not sliding windows, and so can detect objects of any size or shape. This provides the flexibility of filtering detections at the object level. Examples show the ability to detect small compact objects such as vehicles as well as large, spatially extended features (e.g., built-up and bomb-damaged areas). Unlike previous approaches to change detection, which compare pixels, vectors, features, or objects, cluster-based change detection involves no direct comparison of images. In fact, it is identical to the object detection algorithm, different only in the way it is applied. Preliminary results show cluster-based change detection is less sensitive to image misregistration errors than global change detection. The same cluster-based algorithm can also be used for cross-spectral anomaly detection. An example showing the detection of thermal anomalies in Landsat Thematic Mapper imagery is provided.

Index Terms—Change detection, Gaussian processes, image analysis, image representations, image sequence analysis, object detection, pattern clustering methods, vector quantization.

I. INTRODUCTION

THE DETECTION of man-made objects and changes is important across a broad range of mapping and reconnaissance applications. Although a variety of change detection algorithms exist, most cannot effectively discriminate between man-made and natural change. This is a problem in that changes in the natural clutter (e.g., differences in vegetation state, soil moisture, water quality, etc.) often exceed man-made changes in magnitude and number. This results in a large number of false alarms that reduce system performance. Performance is further reduced by false alarms induced by differences in lighting and imaging geometry between images. In single-image object detection, man-made objects are usually modeled as compact regions (e.g., vehicles, buildings, etc.) that are spectrally (or texturally) different from the background. When the characteristics of clutter are similar to the objects of interest, or when the objects of interest are not compact but spatially extended (e.g., built up areas), the false-alarm rate will be high, resulting in poor performance.

The goal of the present work is to develop techniques for detecting man-made objects and changes while maintaining a low false-alarm (FA) rate in the presence of significant natural background clutter. Historically, a considerable amount of cross fertilization has occurred between object and change detection research. Hunt and Cannon [15] first showed that images could be modeled as Gaussian-distributed intensity fluctuations around a nonstationary ensemble mean. Motivated by this observation, Chen and Reed [11] and Reed and Yu [23] developed object detection techniques based on a local Gaussian model (RX-algorithm). Margalit *et al.* [36] applied the model to change detection; Hoff *et al.* [14] and Yu *et al.* [33] further extended it to cross-spectral anomaly detection.

Therrien *et al.* [29] proposed an object detection approach based on linear filtering similar to Chen and Reed's. Tom developed a linear prediction algorithm for change detection [16] based on Therrien's image filtering approach. The linear prediction algorithm can also be applied globally to compute the linear transformation that best predicts (in the minimum mean square error sense) a new image from a reference image. This is the basis of the hyperspectral change detection technique described by Stein *et al.* [27]. Instead of using a linear estimator, Carlotto [3], [5], [6] computes a nonlinear minimum mean square error estimate of the background from multiple reference images for change detection.

Early change detection techniques were based on statistical features such as cross correlation and entropy [17]. Mandelbrot [19] showed that fractals are good models for a variety of natural phenomena that are self-similar in structure over scale or resolution. That man-made objects tend to have structure that varies over scale led Stein [28] to develop an object detection technique that fits a fractal model to the image within a sliding window, and detects pixels within the window whose features deviate from a fractal model. Carlotto and Stein [8] used differences between fractal features to detect man-made changes.

A variety of change vector techniques have been developed for multispectral change detection. Byrne *et al.* [1] perform a principal components analysis of Landsat data and show that gross differences in overall radiation and atmospheric changes appear in the major component images and changes in land

cover appear in the minor component images. Others compare physically significant band combinations such as tasseled-cap features [9], or change vectors [35].

Most change detection techniques require the images to be physically registered. Symbolic techniques [22] detect changes by segmenting and comparing regions in terms of their size, shape, spectral properties, and spatial relations. Object-level change detection [12], [30] is an outgrowth of the symbolic approach.

From this brief overview, change detection seems inexorably tied to the comparison of images, e.g., of pixels, vectors, features, or objects extracted from two or more images. Differences in sensor geometry, environmental conditions, and other factors have limited the application of change detection techniques to images taken under relatively similar conditions. Methods have been developed to estimate and account for global atmospheric and illumination differences [13], [27]. However, it is more difficult to model and correct for local (but correlated) changes in the background (e.g., changing patterns of shading, shadows, layover, vegetative state, etc.), which are often a major source of false alarms.

A new unified approach for detecting man-made objects and changes is described that involves clustering and analyzing the distribution of pixel values in background clusters over one or more multispectral image. Known as the cluster-based anomaly detector (CBAD), this approach assumes the image background can be partitioned into a set of clusters, where the pixel values in a cluster are modeled by a multivariate Gaussian distribution. In single-image object detection, when the frequency of man-made object pixel values is small compared to background pixel values, man-made objects do not form their own, distinct clusters. Instead they are each assigned the nearest background cluster, becoming an outlier (anomaly) in that cluster (i.e., their value lies farther from the mean than background pixel values in the cluster). For change detection, a reference image (against which change is measured) is divided into clusters. Each cluster represents a homogenous population of pixels in the reference image. Over the set of pixel locations in a reference image cluster, a different set of pixel values are observed in a second test image. If there is no change over this cluster, pixel values in the test image will be clustered around the mean. In clusters affected by change (i.e., clusters that contain changed pixels), new pixel values will be introduced, which will tend to lie farther from the mean than background pixel values in those clusters.

CBAD detects man-made objects and changes by finding spatially connected groupings of outlier pixels in clusters. Based on a Gaussian mixture model [27], CBAD can be used to detect man-made objects that are as follows:

- present in a single multiband image;
- appear or disappear between two images acquired at different times;
- manifest themselves as differences between two sets of spectral bands acquired at the same time.

The organization of the paper is as follows. Section II describes the CBAD algorithm. In the same way RX is predicated on Hunt and Cannon's observation that an image can be modeled as Gaussian fluctuations around a local mean, it is shown that an image can also be described as Gaussian fluctuations around a mean that varies cluster to cluster. Although CBAD does not require any particular clustering algorithm, a vector quantization (VQ) approach was used for its computational efficiency. Unlike RX whose window size is related to the size of the objects one wishes to detect, CBAD makes no assumptions about an object's size or shape. Examples are presented illustrating the detection of small compact vehicles as well as large spatially extended features. Section III describes the application of CBAD approach to change detection. Unlike previous methods for change detection that involve a comparison of pixels, vectors, features, or objects, the cluster-based method involves no direct comparison between images. Preliminary results show it is less sensitive to image registration errors than pixel-based global change detection. Section IV shows how CBAD can be used to detect man-made objects that manifest themselves in a wavelength-dependent manner between two sets of bands. Section V addresses similarities between CBAD and other methods, discusses its shortfalls, and outlines future work. The Appendix illustrates how size, shape, and other filtering techniques can be used to reduce false alarms at the object level for both object and change detection.

II. CLUSTER-BASED ANOMALY DETECTION

Hunt and Cannon [15] proposed that images could be modeled as Gaussian-distributed intensity fluctuations around a nonstationary ensemble mean. They demonstrated this behavior in optical imagery by blurring the image, which provides a local estimate of the mean, and subtracting the blurred (local mean) image from the original image. The statistics of the resulting difference image closely approximate white Gaussian noise. Reed and Yu [23] applied this model for detecting objects of known shape, but unknown spectral characteristics. Following their definitions, the shape of the object (template) is specified by an N pixel (lexicographically ordered) vector $\mathbf{s}$; $\mathbf{X}$ is an M band by N pixel matrix of measurements over the region under test. They derive the generalized-likelihood ratio test (GLRT), also known as the RX-algorithm for detecting instances of objects of known shape whose spectral characteristics differ from those of the background:

$$r(\mathbf{X}) = \frac{(\mathbf{Xs}^T)^T (\mathbf{XX}^T)^{-1} (\mathbf{Xs}^T)}{\mathbf{ss}^T} \begin{array}{l} \geq r_0, \text{ then object} \\ < r_0, \text{ then background.} \end{array} \tag{1}$$

Stein *et al.* [27] show that for a single-pixel object, denoted by the M band vector $\mathbf{x}$, as the window size $N \to \infty$, the test statistic is

$$r(\mathbf{x}) = (\mathbf{x} - \mathbf{m})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{m}) \tag{2}$$

where $\mathbf{m}$ and $\mathbf{C}$ are the M band sample mean vector and $M \times M$ band covariance matrix computed over the sliding window. For the hypothesis under test (the null hypothesis), the test statistic has a Chi-square distribution with M degrees of freedom, and so has a constant false-alarm rate (CFAR) independent of the background clutter statistics [11]
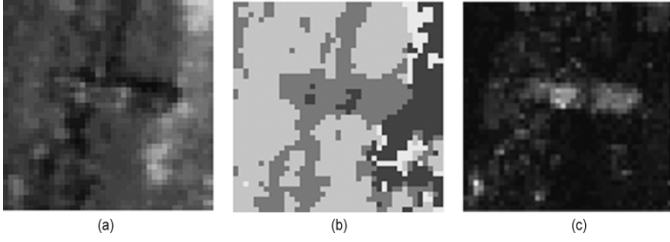
Fig. 1. Cluster-based anomaly detection example. Portion of multispectral image containing a man-made object (a). Cluster image (b). Assignment of grayscales to clusters is arbitrary. Note that object pixels do not form their own cluster but are assigned multiple background clusters. Mahalanobis distance image (c) shows pixels that do not fit their assigned background cluster model.

A. Gaussian Mixture Model

Instead of a local Gaussian model, CBAD uses a Gaussian mixture model [27]. Image statistics are computed over clusters as opposed to a sliding window. CBAD assumes that background pixel values within clusters can be modeled as Gaussian distributions about mean values that vary cluster-to-cluster, and that anomalies (man-made objects) have pixel values that deviate significantly from the distribution of the cluster. It is conjectured that this model is valid in situations where the frequency of occurrence of man-made objects is low compared to the background and so do not form distinct clusters. Instead, man-made object pixels are split up among multiple background clusters as shown in Fig. 1. CBAD exploits this effect by using the Mahalanobis distance [31], which is the same as the test statistic (2), as a means for identifying pixels that do fit well in their assigned cluster.

Multiband data $\mathbf{x}(i,j)$ are decomposed into a set of Gaussian clusters $X_r = \{\mathbf{m}_{x,r}, \mathbf{C}_{x,r}\}$

$$\mathbf{m}_{x,r} = \frac{1}{N_r} \sum_{(i,j) \in X_r} \mathbf{x}(i,j)$$

$$\mathbf{C}_{x,r} = \frac{1}{N_r - 1} \sum_{(i,j) \in X_r} [\mathbf{x}(i,j) - \mathbf{m}_{x,r}][\mathbf{x}(i,j) - \mathbf{m}_{x,r}]^T \quad (3)$$

where N_r is the number of pixels in cluster r. The cluster map [Fig. 1(b)] assigns an index to each pixel location based on its cluster membership

$$r(i,j) = r | \mathbf{x}(i,j) \subset X_r. \quad (4)$$

The Mahalanobis distance of a pixel relative to its cluster [Fig. 1(c)] is

$$d_{x,r}(i,j) = [\mathbf{x}(i,j) - \mathbf{m}_{x,r}(r(i,j))]^T$$
$$\times \mathbf{C}_{x,r}^{-1}(r(i,j))[\mathbf{x}(i,j) - \mathbf{m}_{x,r}(r(i,j))] \quad (5)$$

where $\mathbf{m}_{x,r}(r(i,j))$ and $\mathbf{C}_{x,r}(r(i,j))$ denote the mean and covariance for the cluster at location (i,j).

The Mahalanobis distance (5) is the same test statistic as the RX-algorithm for single pixels under test against the local background (2), and so has the same CFAR property independent of the background clutter statistics. In CBAD, statistics are computed over clusters that are spatially distributed over the image,

in contrast to the those computed within a sliding window in RX. The size of the window in the RX-algorithm is related to the size of the objects of interest. CBAD estimates background statistics over clusters, not sliding windows, and so can detect anomalies of any size (or shape) as shall be demonstrated below. This provides the flexibility of filtering detections by size and shape at the object-level. The Appendix discusses detection and object-level (postdetection) processing in CBAD.

B. Clustering by Vector Quantization

CBAD is potentially more compute-efficient than the RX-algorithm, since it operates over clusters. If L is the number of pixels in an image, the complexity of RX is $O(NL)$, where N is the window size. CBAD's complexity is $O(N_1 + N_2 + \ldots N_R) = O(L)$ plus clustering. This factor of N savings in computation can only be realized if an efficient clustering algorithm is chosen. Iterative algorithms like K-means [31], expectation-maximization [37], and others can require as many as $O(LR)$ operations per iteration, and require many iterations to converge. Although the cluster-based approach is not predicated on any particular clustering algorithm, a non-iterative, vector quantization (VQ) technique, based on ideas summarized in Pratt [21], was chosen for its computational efficiency.

In one dimension, the VQ algorithm assigns image pixel values one of R levels

$$x \to x_r | x_r \leq x < x_{r+1}. \quad (6)$$

The levels $\{x_1, x_2, \ldots x_R\}$ are chosen so as to minimize the mean square quantization error

$$\text{MSQE} = \sum_{r=0}^{R-1} \left[\int_{x_r}^{x_{r+1}} (x - x_r)^2 p(x) dx \right] \quad (7)$$

where $p(x)$ is the probability density of the image. For a uniform density, the levels are equally spaced between the minimum and maximum image values. In general, a companding approach [21] can be used that histogram equalizes the image to force it to have a uniform density

$$y = c[x]. \quad (8)$$

The cumulative distribution function (cdf) of the histogram-equalized image is divided into R equally spaced probability intervals

$$y_r = \frac{r}{R}(y_{\max} - y_{\min}) + y_{\min}. \quad (9)$$

These intervals are mapped back via the inverse transform c^{-1} to obtain the variably spaced quantization levels of the original image

$$x_r = c^{-1}[y_r]. \quad (10)$$

Pratt [21] discusses several VQ strategies for multiband data. If the data are Gaussian, one approach is to spectrally
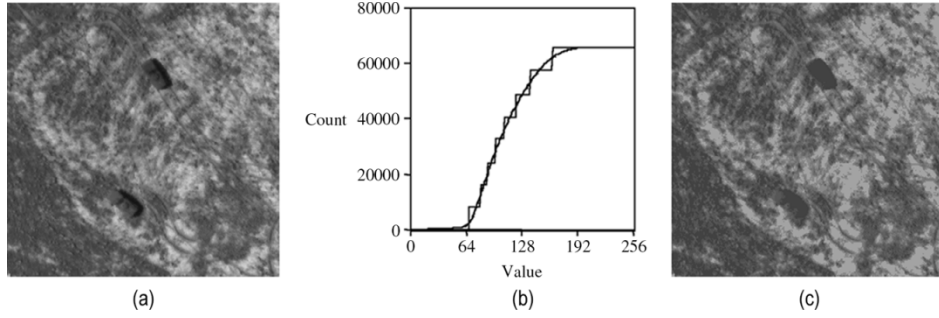
Fig. 2.   (a) Example image, (b) cdf, and quantization levels, and (c) cluster mean image.

decorrelate (rotate) the data using a principal components (PC) transformation

$$\mathbf{z} = \mathbf{W}\mathbf{x} = \begin{bmatrix} \mathbf{e}_1^T \\ \vdots \\ \mathbf{e}_N^T \end{bmatrix} \mathbf{x} \tag{11}$$

where rows of the spectral whitening matrix are the eigenvectors are the covariance matrix

$$\mathbf{C} = E[\mathbf{x}\mathbf{x}^T] - [E(\mathbf{x})E(\mathbf{x})^T] \tag{12}$$

and the eigenvectors $\{\mathbf{e}_n\}$ and eigenvalues $\{\lambda_n\}$ of the covariance matrix satisfy

$$\mathbf{C} = \sum_n \lambda_n \mathbf{e}_n \mathbf{e}_n^T. \tag{13}$$

Each PC image can then be quantized individually and the results combined. For Gaussian data, the total entropy

$$H = -\int p(\mathbf{z})\ln p(\mathbf{z})d\mathbf{z} = \sum_{n=1}^{N} -\int p(z_n)\ln p(z_n)dz_n = \sum_{n=1}^{N} H_n \tag{14}$$

where $p(z_n)$ is the density of the nth PC image. The entropy in the nth PC image is [24]

$$H_n = \ln(\sigma_n \sqrt{2\pi e}) \tag{15}$$

where $\sigma_n^2 = \lambda_n$.

It is noted that (14) is only an approximation if the data are not Gaussian. Provided the Gaussian assumption is valid, bits can be allocated to PC images based on their relative information content. If R is the desired number of clusters, $B = \log_2 R$, where B is an integer (i.e., $R = 1, 2, 4, 8, \ldots$), and $B \leq H$,

$$B_n = \frac{BH_n}{H} \tag{16}$$

is the number of bits assigned to the nth PC image. A block quantization scheme is used to allocate an integer number of bits to each PC based on their information content. The resulting cluster decision regions are the intersection of individual PC decision intervals.

C. Cluster-Based Whitening

RX assumes that images can be modeled as Gaussian-distributed intensity fluctuations around a nonstationary local mean. CBAD assumes that images can be modeled as Gaussian distributions about a mean value that varies cluster to cluster. To assess the empirical validity of the cluster-based model, consider the image in Fig. 2(a) containing two vehicles in a complex natural background. The cdf of the image is plotted along with $R = 8$ (plus zero) quantization levels [Fig. 2(b)]. Fig. 2(c) is the quantized image where each pixel value in the original image is replaced by the average of its respective upper and lower quantization levels. This is an image of ensemble means computed over clusters (quantization intervals) as opposed to local means computed in a sliding window.

It is now demonstrated that the statistics of the difference between the cluster-based mean image and the original image approximates white Gaussian noise. Fig. 3(a) is the upper left corner of the image in Fig. 2(a) containing only background clutter. Fig. 3(b) is its histogram; Fig. 3(c) is the autocorrelation function. Fig. 3(d) is the upper left corner of the difference between the image in Fig. 2(a) and the image of cluster means in Fig. 2(c). Its histogram [Fig. 3(e)] is much more Gaussian than that of the original image [Fig. 3(b)], and its autocorrelation [Fig. 3(f)] closely resembles that of white noise.

Fig. 4(a) is the difference between the original image in Fig. 2(a) and the image of ensemble means in Fig. 2(c). Log plots of the cdf of difference images computed for $R = 2$, 8, and 64 levels are shown in Fig. 4(b). For $R = 8$ levels, essentially Gaussian behavior is exhibited over a large portion of the distribution. Too many $(R = 64)$ or too few $(R = 2)$ clusters result in difference images that are not Gaussian, analogous to those resulting from smoothing windows that are either too small or too large [15].

The image of Mahalanobis distances [Fig. 4(c)] reveals those pixels that do not cluster well. They can be visualized in the isometric plot [Fig. 5(a)] which displays the Mahalanobis distance as a function of image value and cluster. Clusters at the low and high ends contain the tails of the distribution [Fig. 5(b)]. Shadows and bright materials are farthest from the mean in these clusters and so have the highest Mahalanobis distance.

D. Object Detection Examples

Fig. 6(a) shows a slightly larger portion of the image of the two vehicles shown earlier in Fig. 2 (image size 400×400 pixels). Fig. 6(b) is the output from single-pixel RX-algorithm (2) and Fig. 6(c) is the Mahalanobis distance (5). Even though the vehicles are similar in tone to the background, both RX and CBAD are to detect them from their cast shadows. Fig. 6(d) and (e) plot the receiver operating characteristic (ROC) curves for
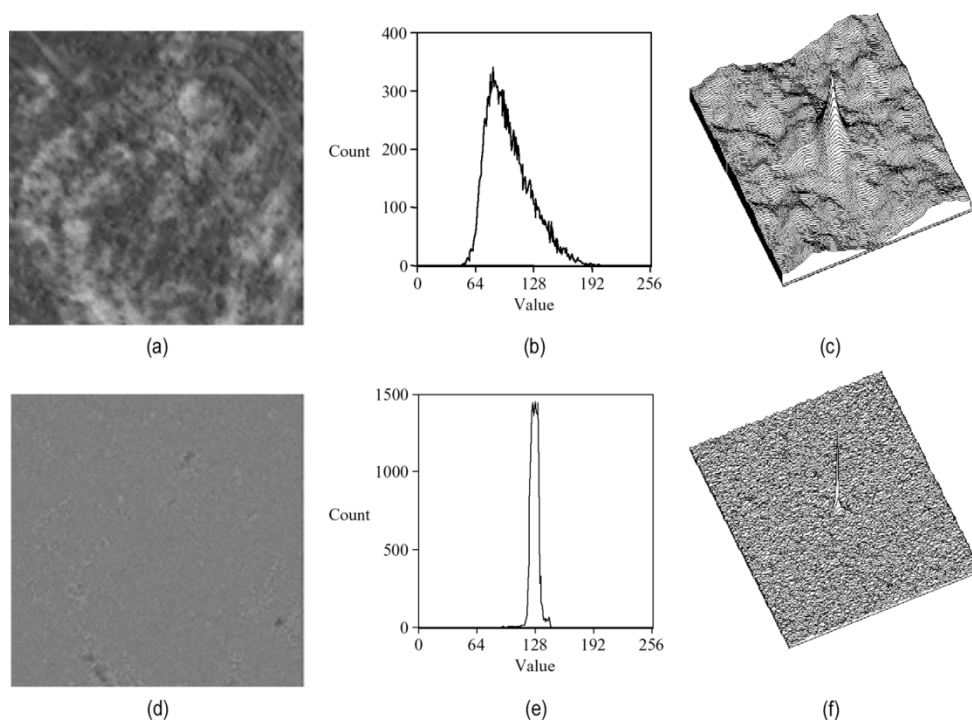
Fig. 3.   Subimage containing background (a) only, (b) histogram, and (c) autocorrelation plot. (d) Difference image, (e) histogram, and (f) autocorrelation plot.
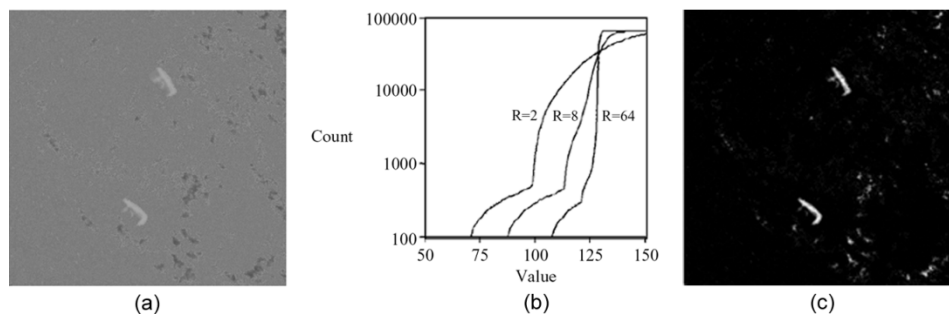

Fig. 4.   (a) Difference image, (b) cdfs of difference images, and (c) Mahalanobis distance image. CDFs (of pixel counts) in (b) are plotted on a log scale. Linear portions of plot are indicative of Gaussian behavior.
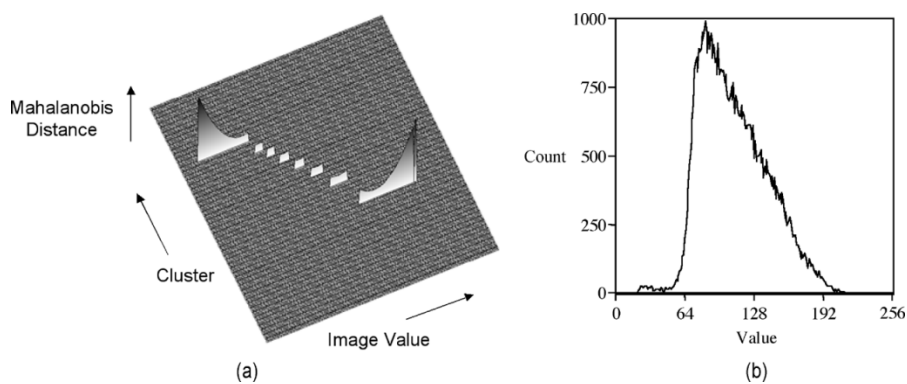

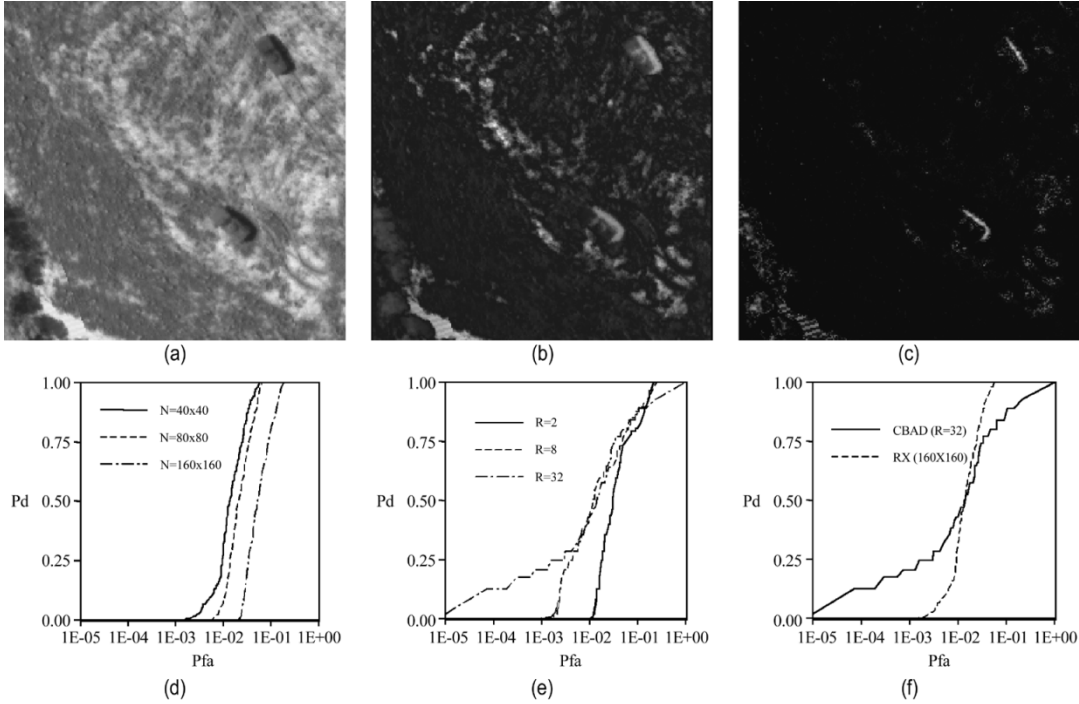Fig. 5.   (a) Isometric plot and (b) histogram of image in Fig. 2(a).

Fig. 6. Comparison of CBAD and RX performance for single-band image. (a) Image containing two vehicles. (b) RX output ($N = 160 \times 160$). (c) CBAD output ($R = 32$). (d) Pd versus Pfa (log scale) for RX. (e) Pd versus Pfa (log scale) for CBAD. (f) comparison of CBAD and RX performance.

RX and CBAD. The probability of detection (Pd) and probability of false alarm (Pfa) are based on the fractions of object and background pixels detected relative to a truth mask as a function of detection threshold. The single-pixel version of RX was applied using window sizes of 40×40, 80×80, and 160×160 pixels. CBAD was run at $R = 2, 8$ and 32 clusters (quantization levels). A comparison of curves [Fig. 6(f)] indicates that CBAD has better a Pd at low Pfa, and RX has a better Pd at high Pfa in this particular example.

Figs. 7 and 8 present two multispectral object detection examples illustrating the use of CBAD in detecting large spatially extended features and small compact objects. A portion of an Ikonos image[1] over Dawrah, Iraq [Fig. 7(a)] shows a built up area (top) and bomb-damage (bottom). The first three Ikonos bands were divided into $R = 16$ clusters ($B = 4$ bits). Fig. 7(b) is the cluster map, and Fig. 7(c) plots cluster versus Ikonos band 1 and band 3 values. The rectangular decision regions are result of block quantization. The global Gaussian assumption (14) appears to be reasonably valid for the data in this example. Fig. 7(d) plots the Mahalanobis distance versus Ikonos bands 1 and 3. Pixels that lie in the tails of the global distribution have high Mahalanobis distance relative to their assigned cluster [Fig. 7(e)]. Being large and spatially extended, these regions would be difficult to detect using the RX-algorithm. The result of size filtering after detection (CFAR = 0.1) for regions larger than 500 pixels in area (arbitrary threshold) is shown in Fig. 7(d). (The Appendix discusses detection, region labeling, and filtering.)

The second example involves the detection of small compact objects (ground vehicles) in multispectral imagery consisting

of three visible and one near-infrared (NIR) band [Fig. 8(a)]. Fig. 8(b) is the cluster map, and Fig. 8(c) plots cluster versus band 4 and band 3 values. The Gaussian assumption does not appear as valid as in the previous example, yet the algorithm still performs well. Fig. 8(d) plots the Mahalanobis distance versus bands 4 and 3. The vehicles (and some trees), which lie at the periphery of the global distribution [Fig. 8(d)], have a high Mahalanobis distance with respect to their assigned clusters [Fig. 8(e)]. Like ground vehicles, certain tree spectra have a low frequency of occurrence in the limited area processed and are a source of false alarms. Size filtering after detection (CFAR = 0.01) eliminates trees leaving vehicle-sized detections [Fig. 8(f)].

III. CLUSTER-BASED CHANGE DETECTION

Consider two coregistered multiband images: a reference image $\mathbf{x}(i,j)$ and a test image $\mathbf{y}(i,j)$ acquired at a latter time containing a man-made change. The test image contains two components

$$\mathbf{y}(i,j) = \mathbf{y}_b(i,j) + \mathbf{z}(i,j) \tag{17}$$

where $\mathbf{y}_b(i,j)$ is the background, and $\mathbf{z}(i,j)$ is change. If the two images are acquired under identical conditions, $\mathbf{y}_b(i,j) = \mathbf{x}(i,j)$. Image subtraction yields

$$\mathbf{y}(i,j) - \mathbf{x}(i,j) = \mathbf{y}(i,j) - \mathbf{y}_b(i,j) = \mathbf{z}(i,j). \tag{18}$$

In practice, atmospheric and sensing conditions also change, and so the appearance of the background is different in the two images. The key challenge in detecting man-made change is the elimination of nonsignificant changes in the background.
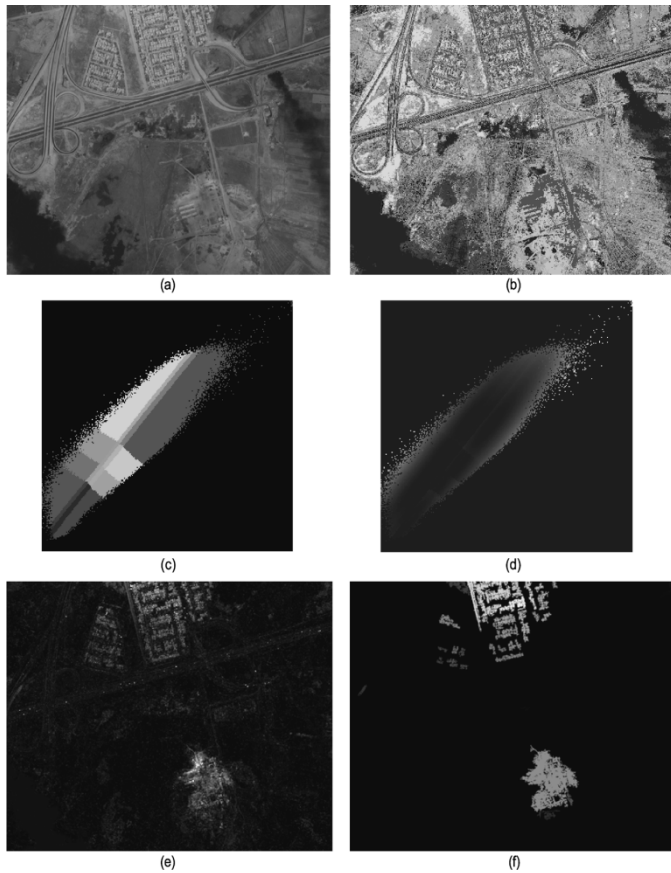
[1]http://www.spaceimaging.com/gallery/default.htm.

Fig. 7. Extraction of large spatially extended features using CBAD. (a) Ikonos bands 1–3 over Dawrah, Iraq. Band centers at 0.48, 0.55, and 0.66 μm. (b) False-color (in grayscale) rendition of cluster map ($R = 16$ clusters). (c) Cluster map displayed as a function of the band 1 (horizontal axis) and band 3 (vertical axis) value. (d) Mahalanobis distance (brightness) plotted as a function of the band 1 (horizontal axis) and band 3 (vertical axis) value. (e) Mahalanobis distance image. Large values occur over built up area (top) and bomb damage (below). (f) Detected regions greater than 500 pixels in area ranked by their average Mahalanobis distance (Pfa = 0.1).
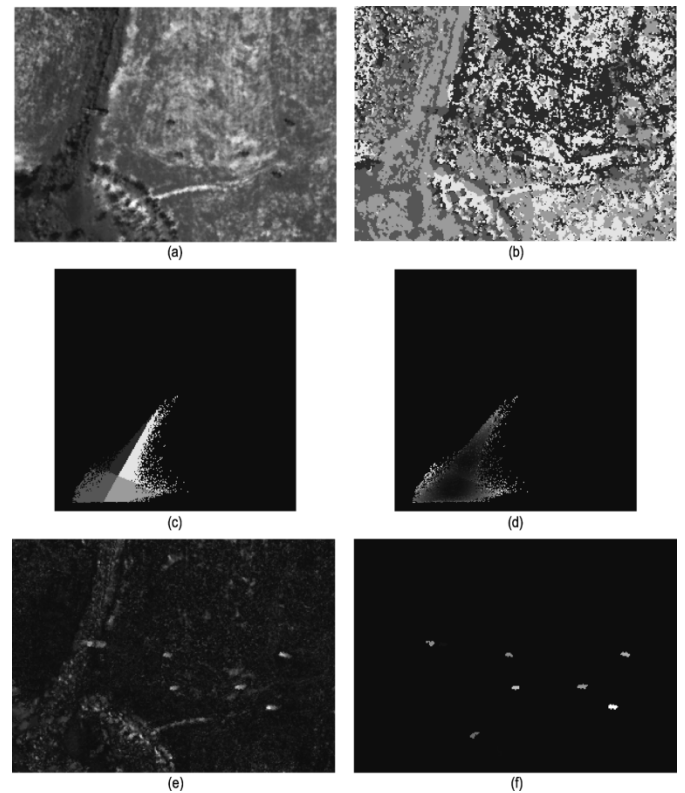

Fig. 8. Extraction of compact objects using CBAD. (a) False-color (in grayscale) multispectral image (bands 4, 3, and 2 in medium, light, and dark areas, corresponding to red, green, and blue). (b) False-color (in grayscale) rendition of cluster map ($R = 4$ clusters). (c) Cluster map displayed as a function of the band 4 (horizontal axis) and band 3 (vertical axis) value. (d) Mahalanobis distance (brightness) plotted as a function of the band 4 (horizontal axis) and band 3 (vertical axis) value. (e) Mahalanobis distances image. Large values occur over vehicles and trees. (f) Detected regions within the size range of vehicles ranked by their average Mahalanobis distance and compactness. (Pfa = 0.01).

A change detection technique described by Stein *et al.* [27] exploits the fact that under certain conditions spectral measurements of the same material viewed at different times are linearly related. This leads to the linear transform

$$\mathbf{y}_b(i,j) = f\left[\mathbf{x}(i,j)\right] = \mathbf{A}\mathbf{x}(i,j) + \mathbf{b} \qquad (19)$$

where the parameters are determined by linear regression. Often images have systematic background differences, e.g., poorly drained fields are wetter in one image, or bodies of water have a higher turbidity, etc. These correlated differences in the background are not considered significant, yet they are a major source of false alarms.

Rather than use a single (global) function [Fig. 9(a)], one could cluster the two images, match clusters in the reference image to those in the test image, and use a different function for the pixels in each cluster. For example, a linear transform could be computed over each cluster

$$\mathbf{y}_b(i,j) = f_r\left[\mathbf{x}(i,j)\right] = \mathbf{A}\left[r(i,j)\right]\mathbf{x}(i,j) + \mathbf{b}\left[r(i,j)\right] \quad (20)$$

where the functions vary by cluster [Fig. 9(b)]. The would provide the change detector additional degrees of freedom to adapt to background differences.

Instead of actually estimating these functions over each cluster and performing change detection by subtraction (18), an alternative method is proposed that is based on analyzing the distribution of pixel values within clusters. In cluster-based change detection (CBCD), the reference image is divided into clusters as in CBAD. Each cluster represents a homogenous population of pixels in the reference image. Over the set of pixel locations in a reference image cluster, a different set of pixel values are observed in the test image. If there are no man-made changes affecting this cluster, the pixel values in the test image will be related to those in the first. Instead of trying to determine an explicit relation to account for any background differences over the cluster, pixel statistics in the second image are examined to determine whether or not man-made changes have occurred. If there are correlated changes in the background over the cluster, the test image mean and covariance may be different from the reference image, but the pixels in the cluster will remain clustered around the mean. But if the cluster is affected by man-made change, new (uncorrelated) values will be introduced tending to produce clusters with a mixed population of pixels, having a wider spread of values. As in CBAD these man-made change pixels will have a higher Mahalanobis distance relative to the background.
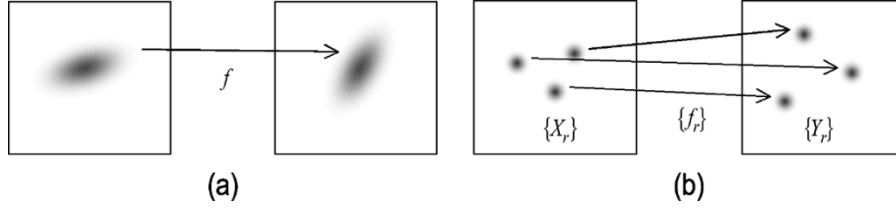
Fig. 9.   (a) Global change detection uses a single function to relate reference and test image spaces. (b) Cluster-based change detection effectively assumes different functions for each cluster.
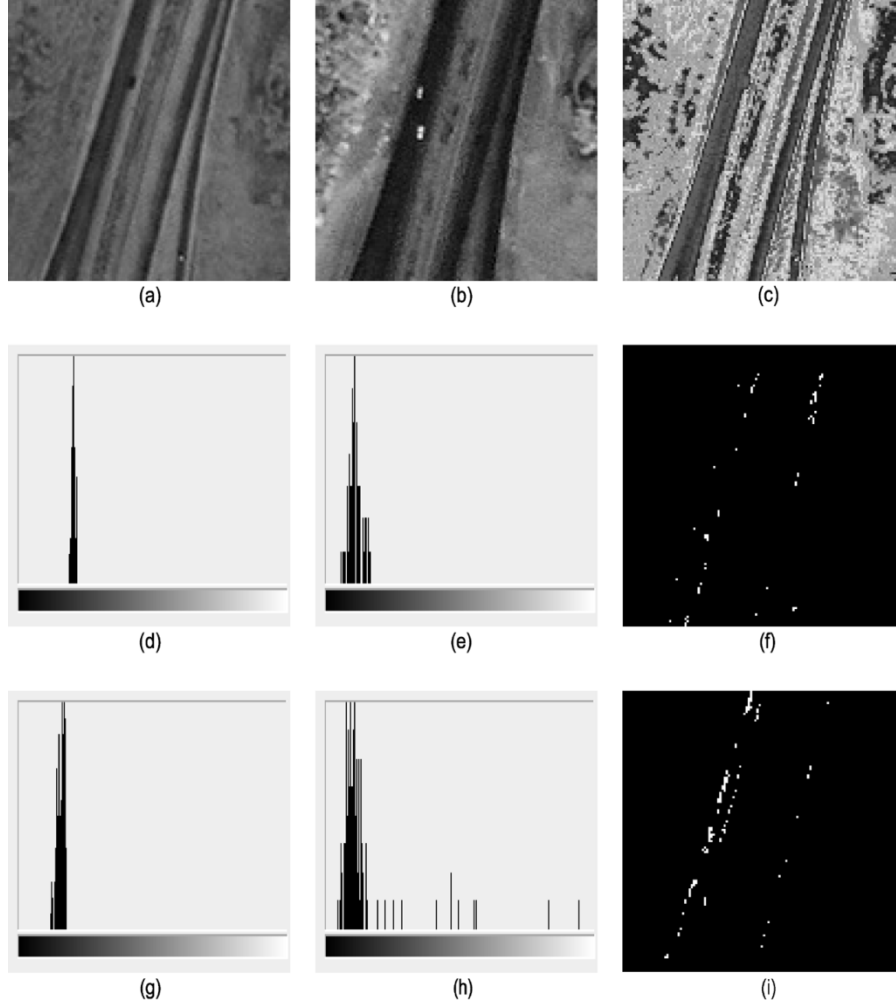

Fig. 10.   Illustration of cluster-based change detection concept. (a) Reference image (blue band), (b) test image (blue band). (c) False-color (in grayscale) rendition of reference image cluster map. (d) Blue band reference image histogram over no change cluster. (e) Blue band test image histogram over no change cluster. (f) Pixel locations of no-change cluster. (g) Blue band reference image histogram over cluster with a change. (h) Blue band test image histogram over cluster with a change—note presence of second population. (i) Pixel locations of cluster containing change.

Assume the reference image has been divided into a set of clusters $\{X_r\}$ where each cluster is represented by a set of parameters (3). Let $\{Y_r\}$ be the set of cluster parameters computed from the test image $\mathbf{y}(i,j)$ over the spatial extent of the clusters derived from $\mathbf{x}(i,j)$, i.e., $Y_r = \{\mathbf{m}_{y,r}, \mathbf{C}_{y,r}\}$ where

$$\mathbf{m}_{y,r} = \frac{1}{N_r} \sum_{(i,j)\in X_r} \mathbf{y}(i,j)$$

$$\mathbf{C}_{y,r} = \frac{1}{N_r-1} \sum_{(i,j)\in X_r} [\mathbf{y}(i,j)-\mathbf{m}_{y,r}][\mathbf{y}(i,j)-\mathbf{m}_{y,r}]^T . \quad (21)$$

The reference image statistics $\{X_r\}$ model random background fluctuations within a cluster as a multivariate Gaussian distribution. For each pixel in the reference image, its Mahalanobis distance is given by (5). Now, for each pixel in the test image, its Mahalanobis distance image relative to $\{Y_r\}$ is

$$d_{y,r}(i,j) = [\mathbf{y}(i,j) - \mathbf{m}_{y,r}\,(r(i,j))]^T$$
$$\times \mathbf{C}_{y,r}^{-1}\,(r(i,j))\,[\mathbf{y}(i,j) - \mathbf{m}_{y,r}\,(r(i,j))]. \quad (22)$$

Within-cluster differences in the test image are caused by background fluctuations and changes in the test image. It is con-
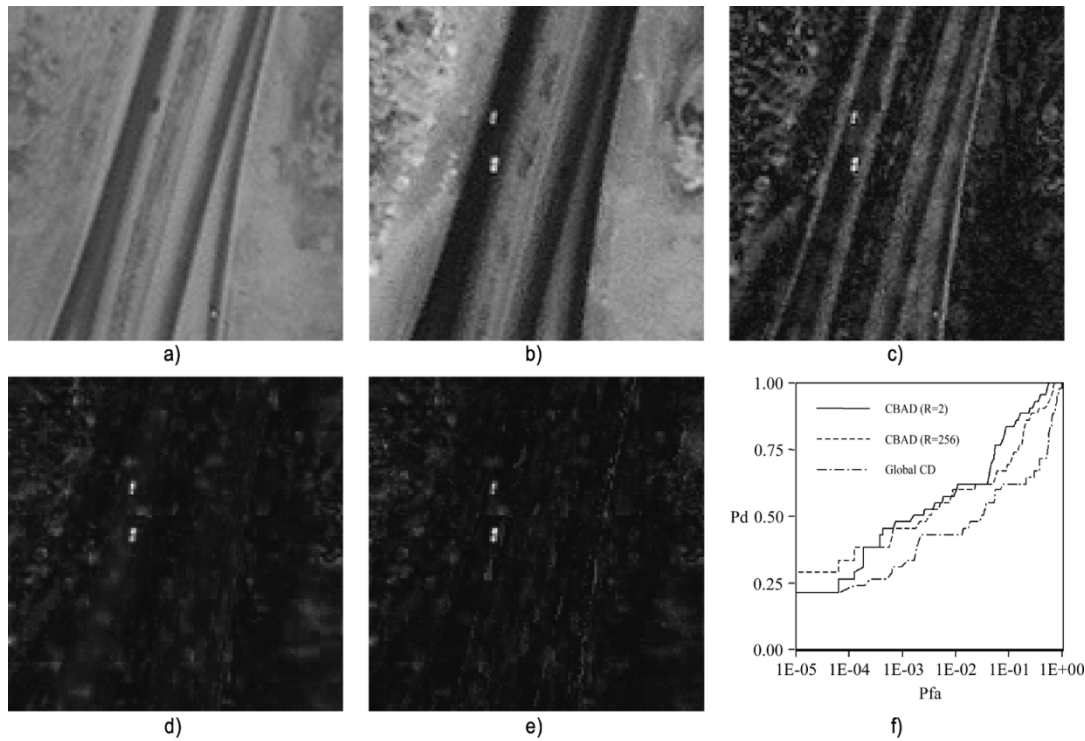
Fig. 11.   Comparison of CBCD and global change detection algorithm results. (a) Reference image. (b) Test image. (c) Global change detection difference image. d) CBCD Mahalanobis distance image ($R = 2$ clusters). (e) CBCD Mahalanobis distance image ($R = 256$ clusters). (f) ROC curves plotting Pd versus Pfa in (percent object and background pixels detected as a function of operating point).

jectured that changed pixels in a cluster will have a higher Mahalanobis distance than background (no change) pixels.

A. Change Detection Example

Fig. 10 illustrates the change detection concept. A pair of Ikonos images (bands 1–3) were registered. Consider a region containing vehicle changes on a road [Fig. 10(a) and (b)]. A cluster map ($R = 256$ clusters) was computed using the VQ algorithm over this region [Fig. 10(c)]. Let us examine two clusters: one over a region of no change [Fig. 10(f)], the other containing a vehicle change in the test image [Fig. 10(i)]. The VQ algorithm produces clusters in the reference image whose distributions are approximately Gaussian [Fig. 10(d)]. Assuming a linear relationship between the test and references image pixels within this cluster, the distribution of pixel values in the test image [Fig. 10(e)] should also be approximately Gaussian. Changes in the test image generally affect multiple clusters, e.g., the appearance of a vehicle replaces background pixel values in some clusters by those of the vehicle. The value of changed pixel(s) will, in general, differ from those of the background leading to a mixed population. This is seen in one such cluster [Fig. 10(i)], where the reference image distribution is approximately Gaussian [Fig. 10(g)], while the test image distribution is mixed [Fig. 10(h)]. If the frequency of change within a cluster is small, the change pixels do not significantly affect the statistics of the cluster, which is dominated by the background. However, change pixels, which deviate from the background, produce larger Mahalanobis distances than do background pixels.

B. Comparison of CBCD and Global Change Detection

Fig. 11 compares CBCD with a global change detection algorithm for the two images in the previous example. The global change detection algorithm uses the reference image to estimate the background in the test image by assuming a linear relation (19). Significant differences in atmospheric state, illumination, sensor geometry, and ground cover result in large background differences between the two images [Fig. 11(a) and (b)]. In this example, two vehicles appear, and one vehicle disappears. The appearance of the two vehicles is clearly evident in the Mahalanobis distance image [Fig. 11(d) and (e)]. By comparison these changes are much less obvious in the global change detection result [Fig. 11(c)], in which background differences dominate. (Objects that disappear are detected by running the change detector in the reverse direction.) Fig. 11(f) plots ROC curves for the three results [Fig. (11c)–(e)]. At a given Pd, CBCD has about an order of magnitude fewer false-alarm pixels than the global change detection algorithm. Changing the number of clusters does not greatly affect CBCD results, since the changes are spectrally distinct from the background in this example.

C. Effect of Misregistration on Detection Performance

Misregistration is typically a significant factor limiting change detection performance [10]. In order to assess its effect on CBCD, the reference and test images were misregistered by shifting one relative to the other. Fig. 12 plots the performance of CBCD and global change detection for shifts of 0, 1, and 4 pixels in the horizontal direction. The performance of global change detection decreases significantly as misregistration
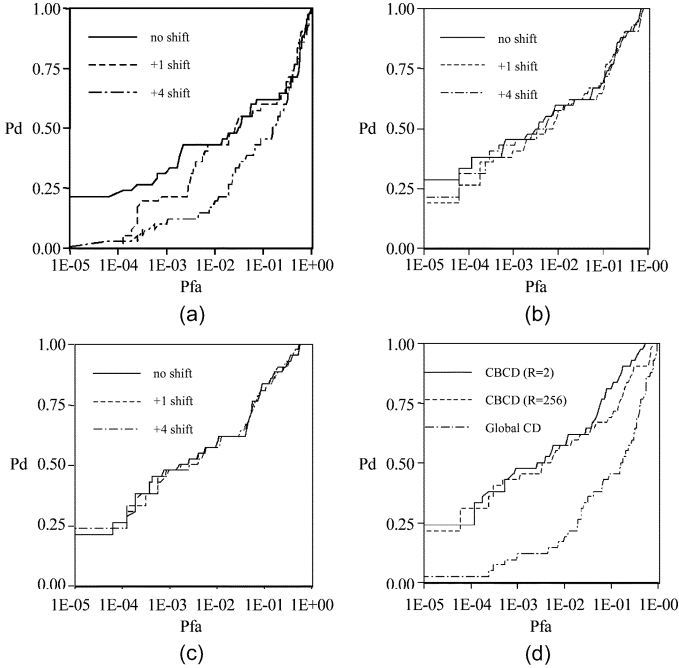
Fig. 12. Effect of misregistration on CBCD and global change detection performance (Pd versus Pfa). (a) Global change detection performance (Pd versus Pfa). (b) CBCD performance ($R = 256$ clusters). (c) CBCD performance ($R = 2$ clusters). (d) Comparison of global change detection and CBCD (+4 shift).
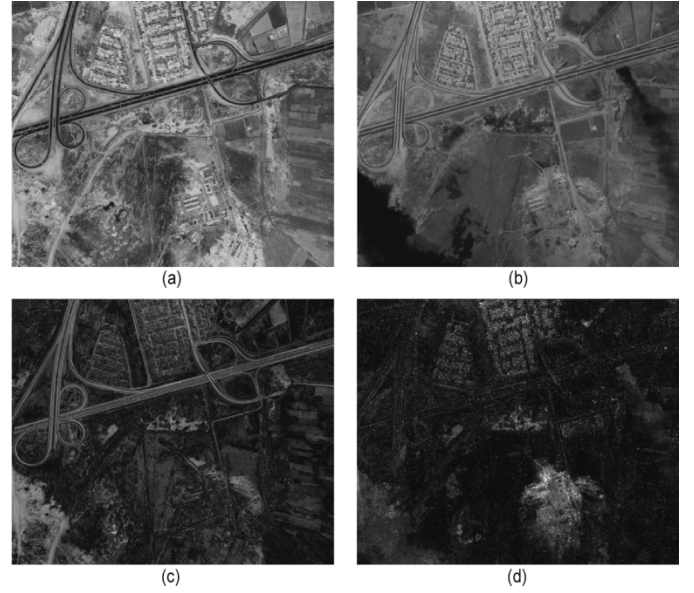

Fig. 13. Comparison of global change detection and CBCD results over a larger image. (a) Reference (old) image. (b) Test (new) image. (c) Global change detection difference image. (d) CBCD Mahalanobis distances image ($R = 256$ clusters).

increases [Fig. 12(a)]. In comparison, CBCD performance ($R = 256$ clusters) decreases only slightly [Fig. 12(b)]. For $R = 2$ clusters, performance is relatively constant [Fig. 12(c)]. CBCD's stability in performance is attributed to the fact that the background is estimated over a set of pixels that are distributed over the entire image. Spatial structure in images causes neighboring pixels to be statistically correlated. The correlation distance depends on the scale of the structure. In parts of the image containing larger structures, the distance will be greater than in parts containing smaller structures. CBCD operates over spatially distributed, nonconnected pixel sets and so can take advantage of spatial correlations where and when they occur. At Pd = 0.5 CBCD has almost two orders of magnitude fewer FA pixels than global change detection [Fig. 12(d)].

D. Results for Larger Image Pair

Results for a larger portion of the previous image pair are now presented. Fig. 13(a) and (b) shows a registered pair of Ikonos images (bands 1–3) acquired before and after the start of the Iraqi War. Global change detection and CBCD results are shown in Fig. 13(c) and (d). The global change detection algorithm responds largely to background differences, including changes in the appearance of roads due to shading and shadowing differences, and vegetation changes. CBCD, on the other hand, responds strongly to the bomb damage, and to a lesser extent, to changes in the built-area near the top of the image. The latter are false-alarms caused by changing patterns of building layover and shadows between the two dates. Discussions of these effects and methods to reduce their impact on change detection performance are contained in Section V and the Appendix.

IV. CROSS-SPECTRAL ANOMALY DETECTION

Instead of operating on two images over time, one set of image bands can be used to estimate another set of bands (from the same or a different coregistered sensor) imaged at the same time. For example, a technique for removing space-varying, wavelength-dependent haze (and smoke) in multispectral imagery uses reflective infrared (IR) bands, which are less affected by haze, to estimate the visible bands [4]. The technique effectively removes the space-varying haze in the visible bands [26]. Haze and smoke can be interpreted as spectral anomalies—features in one set of bands that are not predictable from another set of bands. Hoff *et al.* [14] describe a generalized spectral difference algorithm for detecting weak targets in multiband imagery. It assumes that natural backgrounds behave as a gray body (emitting more or less uniformly over wavelength), while man-made objects emit radiation more strongly at particular wavelengths. Stated another way, natural backgrounds tend to be highly correlated while man-made objects exhibit less correlation (are not predictable) band to band.

The cluster-based approach can be applied to cross-spectral anomaly detection in the same way as change detection. Two sets of bands are defined: reference bands in which the objects of interest cannot be distinguished from the background, and test bands which respond weakly to the objects of interest. The reference bands are clustered (3), statistics of the test bands are estimated over the spatial extent of the clusters derived from the reference bands (21), and the Mahalanobis distance is computed for each test pixel relative to its corresponding cluster (22). Within-cluster differences in the test bands are caused by background and object fluctuations. The cluster-based model assumes spectrally anomalous object pixels will be farther from the mean (have a higher Mahalanobis distance) than background pixels in a cluster.

Fig. 14 illustrates cluster-based cross-spectral anomaly detection on Landsat Thematic Mapper (TM) imagery. Landsat
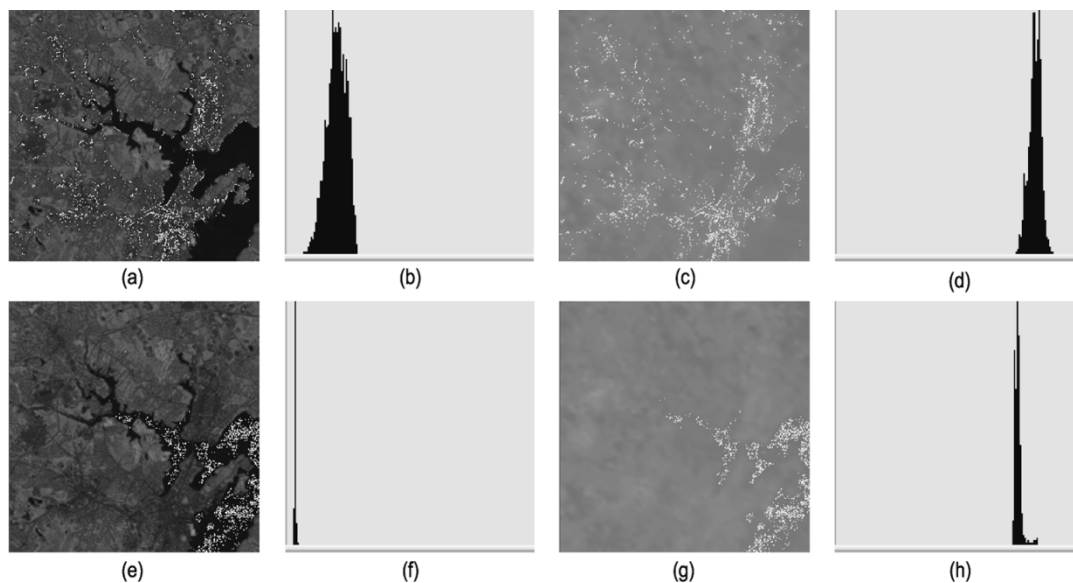
Fig. 14. Illustration of cluster-based cross-spectral anomaly detection concept. (a) One of several mixed urban clusters over band 4. (b) Band 4 histogram. (c) Same mixed urban cluster over band 6. (d) Band 6 histogram. (e) One of several water clusters over band 4. (f) Band 4 histogram. (g) Same water cluster over band 6. (h) Band 6 histogram showing evidence of second component due to warmer pixels from power plant.

TM contains three visible bands (bands 1, 2, and 3) and three reflective IR bands (bands 4, 5, and 7) at 30 m/pixel, and one thermal band (band 6) at 120 m/pixel. Landsat bands 1–5, and 7 (reference bands) were clustered over a region north of Boston. Band 6 (the test band) responds to the objects of interest (thermally emissive features). Fig. 14(a) and (c) highlights the pixel locations of a cluster (mixed urban) over bands 4 and 6. Fig. 14(b) and (d) shows the histograms of bands 4 and 6 over that cluster. This cluster does not contain any thermally anomalous features, i.e., the thermal band is correlated with the visible/relective IR bands. Fig. 14(e) and (g) plots the pixel locations of a cluster (water) containing a spectral anomaly (thermal discharge from a power plant) over the same bands; Fig. 14(f) and (h) shows the histograms over the cluster. The thermal discharge from the power plant raises the temperature of the water which produces a mixed population of cooler (background) and warmer (object) pixels in the band 6 histogram [Fig. 14(h)].

Fig. 15 plots bands 4 and 6 values, and the Mahalanobis distance along a transect over water through the power plant's thermal plume. While the IR reflectivity (band 4) over water remains relatively constant, the thermal emissivity (band 6) increases slightly in the plume. The mixed population of band 6 values over clusters containing the plume (one of which is shown in Fig. 14) produces large within-cluster differences (large values of the Mahalanobis distance). Stated another way, the thermal discharge is not detectable in the visible and reflective IR bands, and so cannot be predicted from these bands. This leads to large prediction errors in the weighted spectral difference algorithm [14, eq. (13)].

A. Thermal Anomaly Detection Example

Fig. 16(a) is a false-color image of Landsat TM bands 4, 3, and 2 (same area as shown in Fig. 14). The thermal band image is shown in Fig. 16(b). Two features of interest are identified
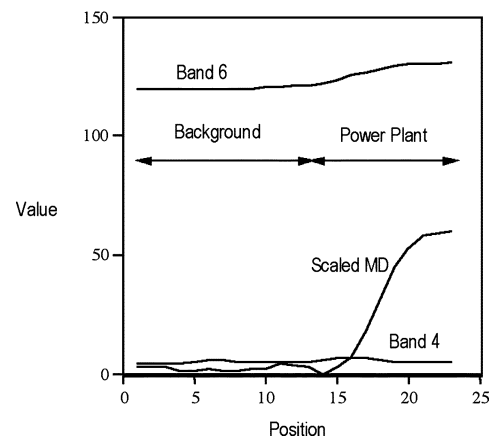

Fig. 15. Data along transect through power plant's thermal plume. Bands 4 and 6 are raw (uncalibrated) pixel values. Mahalanobis distance increases as one moves through the plume.

with arrows: a shopping center (upper left), and a coal-burning power plant (lower right). Neither feature stands out from the background in the visible and NIR (reference) bands. In the test band (thermal band), the shopping center is warmer than the surrounding area, but the power plant's thermal discharge into the cold water is not noticeably warmer than nearby regions [Fig. 16(b)]. Mahalanobis distance images were computed for different numbers of clusters. Fig. 16(c) is the result for $R = 256$ clusters, which clearly reveals the power plant, shopping center, and several other smaller thermal anomalies. Fig. 16(d) plots ROC curves for the thermal (band 6) image, and for CBAD with $R = 32$ and 256 clusters. As in previous examples the Pd and Pfa are based on the fractions of object and background pixels detected relative to a truth mask as a function of detection threshold. For $R = 256$ clusters, at Pd $= 0.5$, CBAD has roughly two orders of magnitude fewer false-alarm pixels than simple thresholding of band 6.
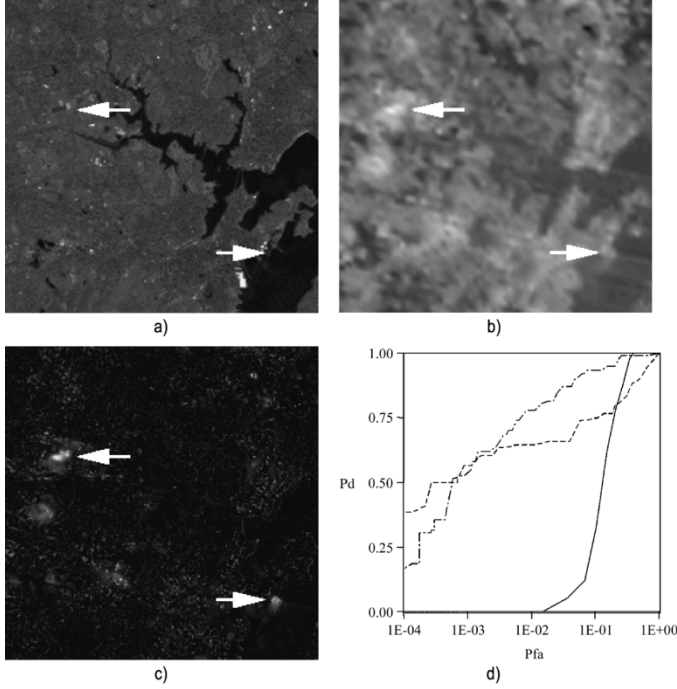
Fig. 16. Thermal anomaly detection example. The object of interest is the Salem Harbor power plant near the lower right corner of the image. (a) Landsat TM bands 4-3-2. (b) Thermal band (TM 6). Arrows show thermally anomalous features. (c) Mahalanobis distance image ($R = 256$ clusters). (d) ROC curves plotting detection performance around the power plant. (Solid line) Band 6. (Dashed line) CBAD ($R = 32$). (Dashed–dotted line) CBAD ($R = 256$).

V. DISCUSSION

CBAD provides a unified approach to object and change detection based on a Gaussian mixture model. It has been shown to be useful in detecting man-made objects that are: 1) present in a single multiband image; 2) appear or disappear between two images acquired at different times; or 3) manifest themselves as spectral differences between two sets of bands acquired at the same time.

The cluster-based method was motivated by previous work in nonlinear mean square image estimation. Nonlinear image estimation algorithms have been applied to change detection [5], elevation data sharpening [6], and haze equalization [4]. In each application, a nonlinear function is assumed to exist between input and output subspaces. For change detection, it is between two images acquired at different times, for haze equalization it is between two sets of spectral bands acquired at the same time, and for elevation sharpening it is between an image and gradient information derived from lower resolution terrain elevation data. A hash table implementation [2] stores the nonlinear estimate $\hat{\mathbf{y}} = E[\mathbf{y}|\mathbf{x}_k]$ for each unique combination of input pixel values $\mathbf{x}_k$ (pixel sets) computed from the output image:

$$E[\mathbf{y}|\mathbf{x}_k] = \frac{1}{N_k} \sum_{(i,j)|\mathbf{x}(i,j)=\mathbf{x}_k} \mathbf{y}(i,j) \qquad (23)$$

where N_k are the number of pixels with value $\mathbf{x}_k$.

Instead of computing estimates over pixel sets, the cluster-based approach forms its estimates over clusters. It can be viewed as a generalization of the nonlinear image estimation technique. In nonlinear estimation, averaging is over

pixels with the same value. In cluster-based anomaly detection averaging is over the pixels in a cluster

$$E[\mathbf{y}|X_r] = \frac{1}{N_r} \sum_{(i,j)|x(i,j)\subset X_r} \mathbf{y}(i,j) \qquad (24)$$

where N_r are the number of pixels in cluster r. The cluster-based estimator thus reduces to the nonlinear image estimator when clusters consist of unique pixel-value combinations.

The number of clusters is currently a user-specified parameter that depends on the spectral diversity of the scene. For large spectrally diverse scenes, hundreds of clusters may be needed to provide a sufficient number of degrees of freedom to adapt to the different background types present. As a result noniterative clustering algorithms like VQ, capable of efficiently partitioning large images into large numbers of clusters, are desirable. Although a VQ clustering algorithm was used here for efficiency reasons, there is no reason why other techniques could not be substituted. One area of future work is to evaluate alternative algorithms, e.g., stochastic EM clustering algorithm [18], and others.

The key question is how to determine the optimal number of clusters. CBAD performance (Pd/Pfa) increases as the number of clusters increases to a point, beyond which it decreases. Pfa decreases as the background is subdivided into smaller, more homogeneous clusters; Pd decreases when the number of clusters become so large that the objects of interest form their own clusters. The minimum description length (MDL) method [25] can be used to pick the number of clusters that results in the greatest compression of the data. McKenzie and Adler [20] find MDL to be more efficient than Akaike information criterion for Gaussian mixtures. Yeung et al. [34] describe a model-based approach using a Bayesian information criterion.

CBCD provides a fundamentally new approach to change detection, one which does not involve a comparison of pixel values (or linear functions of pixel values) between images. In contrast, the Chronochrome algorithm estimates the background in the test image as linear function of the reference image, and detects changes on the whitened difference image [27]. Since CBCD detects deviations in the background distribution over clusters, changes that affect an entire cluster (correlated changes) are not detected by CBCD. A large fraction of the background differences in the examples in Figs. 11 and 13 are correlated changes (e.g., shading and shadow changes). They are a major source of false change in the global change detection algorithm, but are ignored by CBCD. Unfortunately, certain types of man-made change (e.g., large scale forest clearing, cultivation, or urban development) may appear as correlated changes and be missed by CBCD.

CBCD, like most change detection techniques, requires the images to be physically registered to each other. Preliminary results indicate that CBCD is more tolerant of registration errors than global change detection. The point where performance breaks down needs to be explored further to determine its relation to the size of the changes, their signal to noise ratio, the number of clusters, and other factors. Differences in sensor and lighting geometry produce false alarms due to layover and shadow changes. CBAD currently uses an object-level strategy to reduce false alarms as discussed in the Appendix. Another
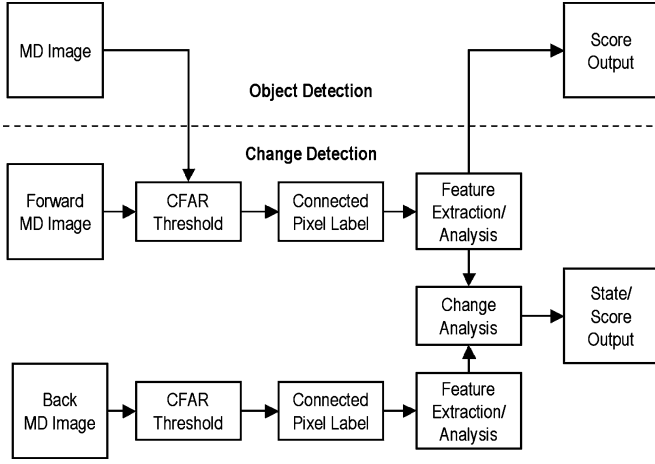
Fig. 17.   CBAD object-level processing flow for object and change detection.
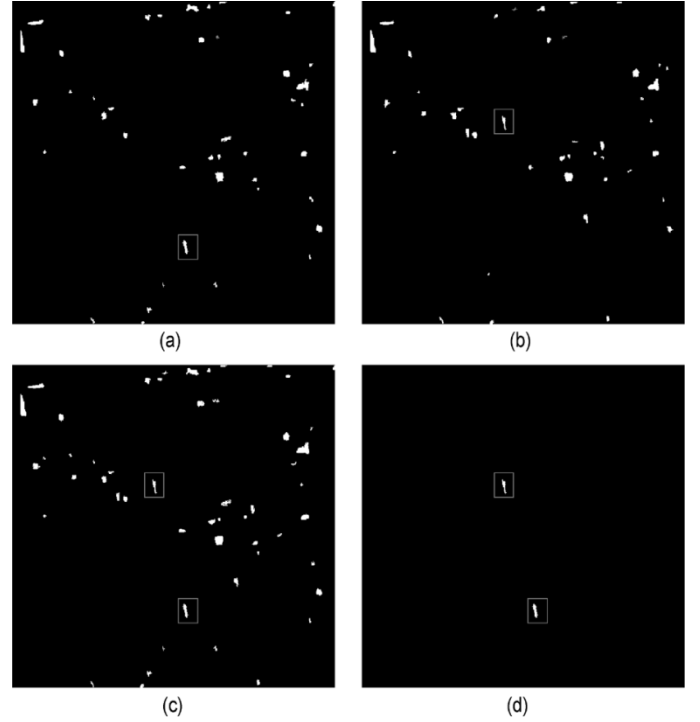

Fig. 18.   Change analysis example illustrating the remove of state changes. (a) Forward change image. (b) Backward change image. (c) All changes. (d) Object appearance and disappearance changes.

method, based on minimizing the Mahalanobis distance over a local search window, is also being explored.

CBAD can be applied to cross-spectral (cross-sensor) anomaly detection in the same way as it is used for two-image change detection. Like CBCD, cross-spectral anomaly detection does not involve a comparison of pixel values (or linear functions of pixel values) between bands. It differs from the generalized spectral difference algorithm [14], which is a whitened difference between linearly transformed reference bands and test bands. The generalized spectral difference algorithm operates between bands in much the same way as the Chronochrome algorithm operates between images (compare Hoff *et al.* [14, eq. (15)] with Stein *et al.* [27, eq. (24)]).

Algorithm fusion by combining local, reference image/band, and cluster based models can offer significant performance gains [7], [27] and is an important area of future work.

APPENDIX
OBJECT-LEVEL PROCESSING

Unlike RX, CBAD does not assume knowledge of an object's size or shape at the pixel level. Instead, it defers the use of size/shape information until after detection, at the object level. This allows different objects to be detected without having to run multiple anomaly detectors.

Object and change detection systems often use a combination of pixel- and object-level processing techniques [7], [12]. Fig. 17 shows the object-level processing flow for object and change detection. Objects (top path) are detected by thresholding the Mahalanobis distance (MD) image from CBAD. Candidate man-made object pixels are those whose MD values (5) exceed a threshold τ

$$\Delta(i,j) = \begin{cases} 1, & d_{x,r}(i,j) > \tau \\ 0, & \text{otherwise} \end{cases} \quad \text{(A1)}$$

where the threshold depends on the desired false-alarm rate. Detected pixels $\Delta(i,j)$ are processed as spatially connected regions [32]. This involves connected pixel labeling, which builds a label map that represents each region (object) by a unique value in the map. Next, a set of features are computed for each connected region such as area, perimeter, centroid, length, width, pose, average Mahalanobis distance within the region, and others. Regions outside a given size range can be

eliminated at this stage of processing, and remaining regions scored based on desired object size/shape characteristics.

The two examples in Figs. 7 and 8 illustrate size filtering and scoring. In the first example, the objective was to extract large, spatially extended regions affected by bomb-damage. Regions smaller than 500 pixels in area were eliminated. Those that remained were scored by their average MD value. In Fig. 8, the goal was to detect vehicle-size objects. Regions outside the range of 15–30 pixels were eliminated. Remaining regions were scored by their average MD value and compactness (area divided by perimeter squared). (The values used in these examples were empirically derived and are intended only to demonstrate the ability to do filtering at the object-level.)

The same process is used for processing forward CBCD changes. Forward changes are objects not present in a reference image that appear in a new image. Candidate forward change (appearance) object pixels are those whose MD values (22) exceed a threshold τ_y

$$\Delta_f(i,j) = \begin{cases} 1, & d_{y,r}(i,j) > \tau_y \\ 0, & \text{otherwise.} \end{cases} \quad \text{(A2)}$$

The change detection examples addressed in Figs. 10–13 involved forward change detection only.

Backward changes are objects present in the reference image that disappear in the new image. Candidate backward change (disappearance) object pixels are those whose MD values exceed a threshold τ_x

$$\Delta_b(i,j) = \begin{cases} 1, & d_{x,s}(i,j) > \tau_x \\ 0, & \text{otherwise} \end{cases} \quad \text{(A3)}$$

where

$$d_{x,s}(i,j) = [\mathbf{x}(i,j) - \mathbf{m}_{x,s}(s(i,j))]^T$$
$$\times \mathbf{C}_{x,s}^{-1}(s(i,j))[\mathbf{x}(i,j) - \mathbf{m}_{x,s}(s(i,j))] \quad \text{(A4)}$$

$s(i,j)$ is the new image cluster map, and

$$\mathbf{m}_{x,s} = \frac{1}{N_s} \sum_{(i,j)\in Y_s} \mathbf{x}(i,j)$$

$$\mathbf{C}_{x,s} = \frac{1}{N_s-1} \sum_{(i,j)\in Y_s} [\mathbf{x}(i,j) - \mathbf{m}_{x,s}][\mathbf{x}(i,j) - \mathbf{m}_{x,s}]^T. \quad (A5)$$

Detected changes can be further processed at the object level to eliminate false alarms caused by differences in shadowing, layover, object pose, etc. between images. This can be accomplished by removing spatially overlapping changes that occur in both directions, which is illustrated in Fig. 18. Fig. 18(a) are forward changes, where the box identifies a vehicle of interest. Fig. 18(b) are backward changes, where the box indicates the same vehicle at a different location at a later time. Fig. 18(c) are all forward and backward changes. Fig. 18(d) are only those changes greater than 50 pixels in area that do not overlap any opposite change.

ACKNOWLEDGMENT

REFERENCES

[1] G. F. Byrne, P. F. Crapper, and K. K. Mayo, "Monitoring land-cover change by principal component analysis of multitemporal Landsat data," *Remote Sens. Environ.*, vol. 10, pp. 175–184, 1980.

[2] M. J. Carlotto, "Non-linear mean-square estimation with applications in remote sensing," *Proc. SPIE*, vol. 2758, pp. 206–217, 1996.

[3] ——, "Detection and analysis of change in remotely-sensed imagery with application to wide area surveillance," *IEEE Trans. Image Process.*, vol. 6, no. 1, pp. 189–202, Jan. 1997.

[4] ——, "Reducing the effects of space-varying wavelength-dependent scattering in multispectral imagery," *Int. J. Remote Sens.*, vol. 20, no. 17, pp. 3333–3344, 1999.

[5] ——, "Nonlinear background estimation and change detection for wide area search," *Opt. Eng.*, vol. 39, no. 5, pp. 1223–1229, May 2000.

[6] ——, "Spatial enhancement of elevation data using a single multispectral image," *Opt. Eng.*, vol. 39, no. 2, pp. 430–437, Feb. 2000.

[7] M. J. Carlotto, M. A. Nebrich, and D. Crary, "Order of battle change detection, fusion, and tracking," *Proc. SPIE*, vol. 4729, Apr. 2002.

[8] M. J. Carlotto and M. C. Stein, "Detecting man-made changes in imagery," *Proc. SPIE*, pp. 38–45, Nov. 1988.

[9] E. P. Crist and R. C. Cicone, "A physically-based transformation of thematic mapper data—The TM tasseled cap," *IEEE Trans. Geosci. Remote Sens.*, vol. GE-22, no. 3, May 1984.

[10] X. Dai and S. Khorram, "The effects of image misregistration on the accuracy of remotely sensed change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 36, no. 5, pp. 1566–1577, Sep. 1998.

[11] J. Y. Chen and I. S. Reed, "A detection algorithm for optical targets in clutter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-23, no. 1, Jan. 1987.

[12] G. G. Hazel, "Object-level change detection in spectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 3, pp. 553–561, Mar. 2001.

[13] G. Healey and D. Slater, "Models and methods for automated material identification in hyperspectral imagery acquired under unknown illumination and atmospheric conditions," *IEEE Trans. Geosci. Remote Sens.*, vol. 37, no. 6, pp. 2706–2717, Nov. 1999.

[14] L. E. Hoff, A. M. Chen, X. Yu, and E. M. Winter, "Generalized weighted spectral difference algorithm for weak target detection in multiband imagery," *Proc. SPIE*, vol. 2561, pp. 141–152, 1995.

[15] B. R. Hunt and T. M. Cannon, "Nonstationary assumptions for Gaussian models of images," *IEEE Trans. Syst., Man Cybernet.*, vol. SMC-6, no. 6, pp. 876–882, Nov. 1976.

[16] B. G. Lee, V. T. Tom, and M. J. Carlotto, "A signal-symbol approach to change detection," *Proc. Nat. Conf. Artificial Intelligence*, Aug. 11–15, 1986.

[17] J. G. Kawamura, "Automatic recognition of changes in urban development from aerial photographs," *IEEE Trans. Syst., Man Cybernet.*, vol. SMC-1, no. 3, Jul. 1971.

[18] P. Masson and W. Pieczynski, "SEM algorithm and unsupervised statistical segmentation of satellite images," *IEEE Trans. Geosci. Remote Sens.*, vol. 31, no. 3, pp. 618–633, May 1993.

[19] B. B. Mandelbrot, *The Fractal Geometry of Nature*. New York: W. H. Freeman, 1983.

[20] P. McKenzie and M. Alder, "Selecting the optimal number of components for a Gaussian mixture model," in *Proc. IEEE Int. Symp. Information Theory*, 1994, p. 393.

[21] W. Pratt, *Digital Image Processing*. New York: Wiley, 1991.

[22] K. Price and R. Reddy, "Change detection and analysis in multispectral images," in *Proc. 5th Int. Joint Conf. Artificial Intelligence*, 1977, pp. 619–625.

[23] R. Reed and X. Yu, "Adaptive multi-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 3, pp. 293–305, Mar. 1990.

[24] F. Reza, *An Introduction to Information Theory*. New York: Dover, 1994.

[25] J. Rissanen, "Universal coding, information, prediction, and estimation," *IEEE Trans. Inform. Theory*, vol. IT-30, no. 4, pp. 629–636, Jul. 1984.

[26] J. Rogan, J. Franklin, D. Stow, L. Levien, and C. Fischer, "Toward operational monitoring of forest cover change in California using multitemporal remote sensing data," in *Proc. IGARRS*, Sydney, Australia, Jul. 9–13, 2001.

[27] D. Stein, S. Beaven, L. Hoff, E. Winter, A. Schaum, and A. Stocker, "Anomaly detection from hyperspectral imagery," *IEEE Signal Process. Mag.*, no. 1, pp. 58–69, Jan. 2002.

[28] M. C. Stein, "Fractal image models and object detection," *Proc. SPIE*, vol. 845, pp. 293–300, 1987.

[29] C. W. Therrien, T. F. Quatieri, and D. D. Dudgeon, "Statistical model-based algorithms for image analysis," *Proc. IEEE*, vol. 74, no. 4, pp. 532–551, Apr. 1986.

[30] V. T. Tom, H. F. Webb, and S. K. Bento, "Context-based SAR object level change detection," in *Proc. 1999 IRIS Nat. Symp. Sensor and Data Fusion*, Laurel, MD, May 24–27, 1999.

[31] J. T. Tou and R. C. Gonzalez, *Pattern Recognition Principles*. Reading, MA: Addison-Wesley, 1974.

[32] P. Winston and B. Horn, *Lisp*. Reading, MA: Addison-Wesley, 1981.

[33] X. Yu, L. E. Hoff, I. S. Reed, A. M. Chen, and L. B. Stotts, "Automatic target detection and recognition in multiband imagery: A unified ML detection and estimation approach," *IEEE Trans. Image Process.*, vol. 6, no. 1, pp. 143–156, Jan. 1997.

[34] K. Y. Yeung, C. Fraley, A. Murua, A. E. Rafery, and W. L. Ruzzo, "Model-based clustering and data transformations for gene expression data," *Bioinformatics*, vol. 17, no. 10, pp. 977–987, 2001.

[35] X. Zhan, C. Huang, J. Townshend, R. DeFries, M. Hansen, C. Dimiceli, R. Sohlberg, J. Hewson-Scardelletti, and A. Tompkins, "Land cover change detection with change vector in the red and near-infrared reflectance space," in *Proc. IGARSS*, vol. 2, 1998, pp. 859–861.

[36] A. Margalit, I. S. Reed, and R. M. Gagliardi, "Adaptive optical target detection using correlated images," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AE-21, no. 5, pp. 394–405, May 1985.

[37] R. Redner and H. Walker, "Mixture densities, maximum likelihood and the EM algorithm," *SIAM Rev.*, vol. 26, no. 2, 1984.

Mark J. Carlotto (S78–M'81–SM'90) was born on May 16, 1954, in New Haven CT. He received the B.S., M.S., and Ph.D. degrees in electrical engineering from Carnegie-Mellon University, Pittsburgh, PA, in 1977, 1979, and 1981, respectively.

He is currently a Senior Staff Scientist in the Advanced Information Systems Division, General Dynamics (GD), Rosslyn, VA. He has held a variety of positions in academia and industry. From 1981 to 1993, he was a member of the technical staff at The Analytic Sciences Corporation, Reading MA. He was also an Assistant Adjunct Professor in the College of Engineering, Boston University from 1981 to 1983. In 1993, he joined Pacific-Sierra Research (PSR), Arlington VA, and remained with PSR until they were purchased by Veridian Systems in 2001. In 2003, Veridian was acquired by GD. His research interests involve the development of automated image processing/understanding techniques for reconnaissance and remote sensing applications.