

# Feature-based anomaly detection

Mark J. Carlotto<sup>1</sup>  
General Dynamics, Advanced Information Systems

## Abstract

A feature-based approach for detecting anomalies in spectral, spatial, temporal, and other domains is described. When the frequency of occurrence is small relative to the background, anomalies such as man-made objects in natural image backgrounds do not form their own clusters, but are instead assigned the nearest background cluster, becoming an outlier (statistical anomaly) in that cluster. Our method clusters data, which may be spectral, spatial, or temporal in nature, into one or more background types and computes the Mahalanobis distance between the data and assigned model (background cluster). The detection of a variety of objects and phenomena in panchromatic and multispectral imagery, and video are illustrated.

**Keywords:** Spectral, spatial, and temporal anomnaly detection, Mahalanobis distance, background modeling, wedge-ring filters.

## 1. INTRODUCTION

Anomaly detection is often the first processing step in automated imagery exploitation systems. Stein et al [1] describe a family of anomaly detectors based on the generalized likelihood ratio test. One of these, the Reed Xiaoli (RX) algorithm [2], detects objects of known shape/shape that are spectrally different from the immediate background. Autoregressive [3] and fractal models [4] have been also proposed for detecting manmade objects in textured backgrounds. Many of these algorithms use some form of sliding window to find objects of a given size.

The cluster-based anomaly detector (CBAD) is a pure spectral detector that involves clustering and analyzing the distribution of pixel values in background clusters [5]. CBAD 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. When the frequency of manmade object pixel values is small compared to background pixel values, manmade 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). In this paper we generalize the CBAD approach for detecting anomalies in space, time, as well as spectrum.

## 2. FEATURE-BASED ANOMALY DETECTION

Let  $\mathbf{z} = \{z_n\}$  be a vector of measurements (features). In spectral anomaly detection,  $z_n$  is the n-th spectral measurement (spectral band). To detect anomalies in a time-series,  $\mathbf{z}$  could be a vector of frequency measurements. In a single band (panchromatic) image,  $\mathbf{z}$  can be a vector of local (texture) measurements. One could also envision  $\mathbf{z}$  to be a vector of size/shape features derived from the connected components of an image.

### 2.1 Spectral Anomaly Detection

Assume the normal background in any of these problem domains is multimodal, and can be characterized by a set of  $K$  means  $\mathbf{m}_k$  and covariances  $\mathbf{C}_k$ , where  $k = 1, 2, \dots, K$  enumerates the clusters. If the relative frequency of anomalies is small compared to the natural background, they do not form their own, distinct clusters. Instead clustering assigns

---

<sup>1</sup> [mark.carlotto@gd-ais.com](mailto:mark.carlotto@gd-ais.com)

anomalies to the nearest background cluster, becoming an outlier in that cluster, i.e., their value lies farther from the mean than background pixel values in the cluster. The Mahalanobis distance measures how similar a measurement is from its associated cluster

$$d = (\mathbf{z} - \mathbf{m}_k)^T \mathbf{C}_k^{-1} (\mathbf{z} - \mathbf{m}_k) \quad (1)$$

and provides a convenient test statistic for anomaly detection [2]. Fig. 1 gives an example of spectral anomaly detection from [5]. The image (a) shows a truck in a mixed background containing trees, sparse vegetation, and water. In cluster image (b) each background type has been assigned a unique cluster (color). Notice that object pixels do not form their own cluster but are instead assigned the nearest background cluster. The Mahalanobis distance image (c) shows pixels on the truck, which do not fit their assigned background model.

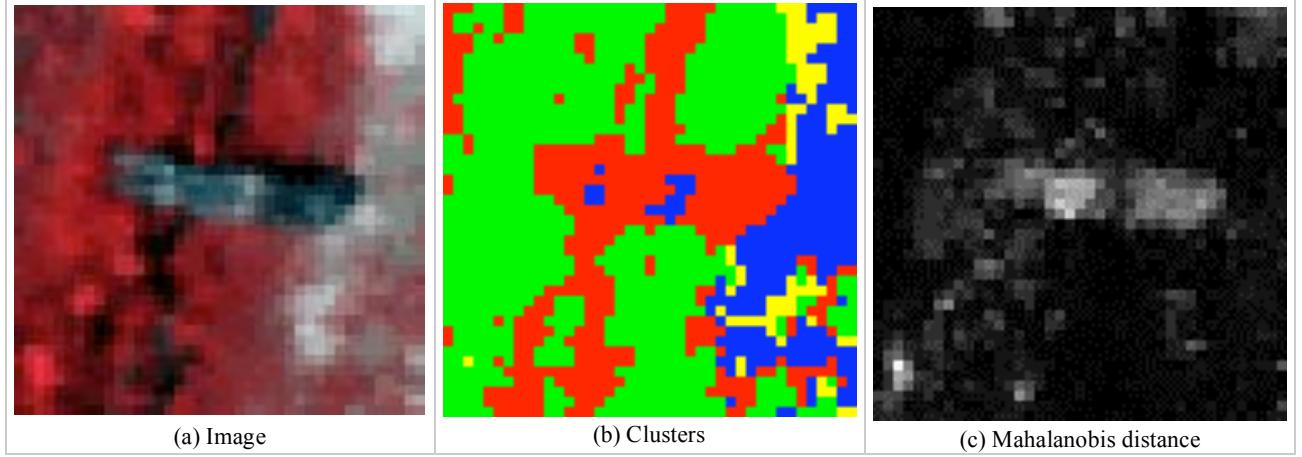


Fig. 1 Cluster-based spectral anomaly detection

## 2.2 Temporal Anomaly Detection

Now we apply this approach to temporal anomaly detection. Let  $x(m\Delta t) = x_m$  be a discrete time series, and  $z_n$  its discrete Fourier transform (DFT)

$$z_n = \sum_{m=0}^{N-1} x_m e^{-j2\pi mn/N} \quad (2)$$

where  $N$  is the length of the sequence. If  $\mathbf{m} = E[\mathbf{z}_{ij}]$  is the average DFT over an ensemble (image) of signals, and  $\mathbf{C} = E[(\mathbf{z}_{ij} - \mathbf{m})(\mathbf{z}_{ij} - \mathbf{m})^T]$  is the covariance, then

$$d_{ij} = (\mathbf{z}_{ij} - \mathbf{m})^T \mathbf{C}^{-1} (\mathbf{z}_{ij} - \mathbf{m}) \quad (3)$$

measures the distance of a signal's spectrum at  $(i, j)$  from the average.

Fig. 2 illustrates the use of a temporal anomaly detector to enhance tracks of moving objects in noisy imagery. Fig. 2a is from a video clip taken by astronauts aboard the Space Shuttle Discovery in 1991 showing a number of moving objects in space near the spacecraft. Although frame averaging (b) reveals the tracks of the objects, they are hard to separate from the background, particularly where they cross the limb of the Earth. Computing the maximum value across frames

(c) works well when the objects have a positive contrast relative to the background but would not work in general. Fig. 2d is the output of the temporal anomaly detector, which clearly shows the tracks.

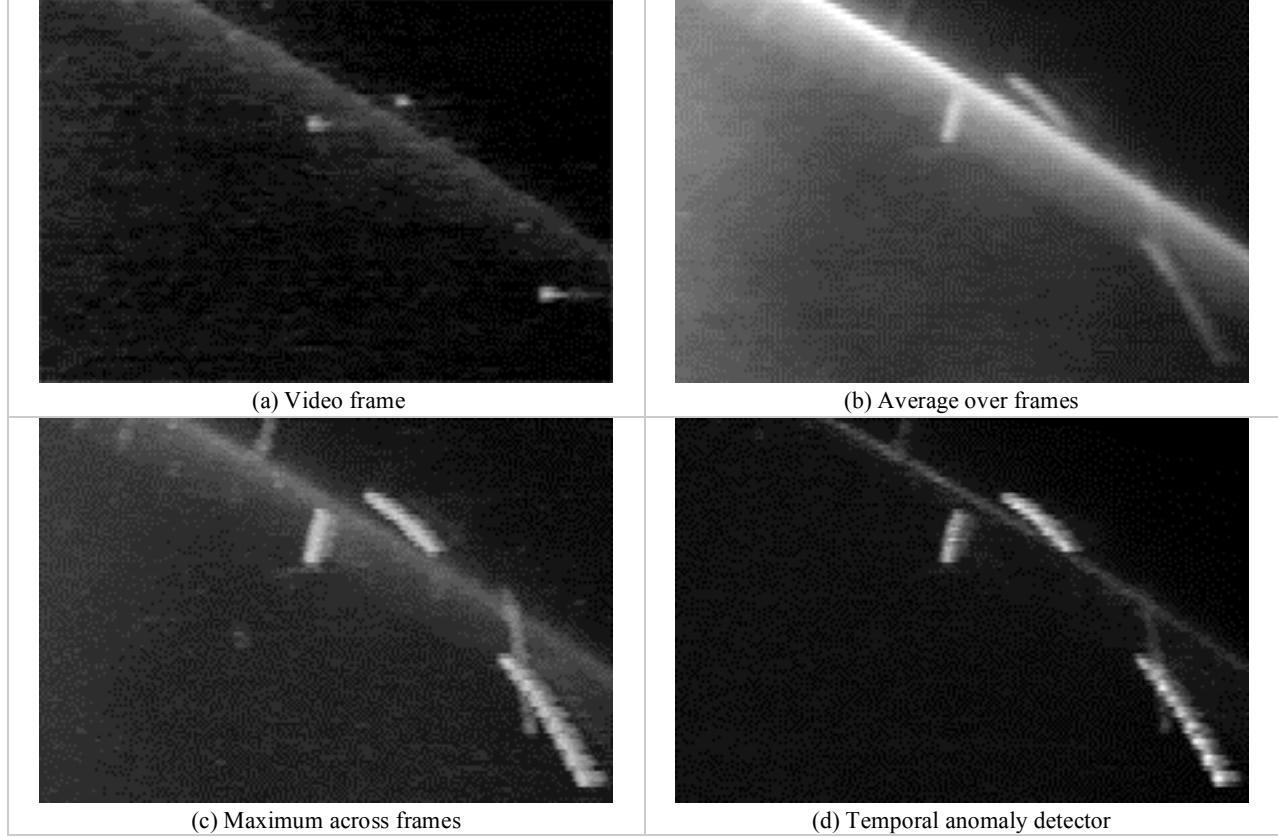


Fig. 2 Temporal anomaly detection in video data (Image courtesy NASA)

### 3. SPATIAL ANOMALY DETECTION

A set of filters applied to an image produces a set of filtered images

$$z_{ij}(n) = x_{ij} * f_{ij}(n) \quad (4)$$

Instead of computing the DFT over a time series, let  $X_{pq} = \mathfrak{F}\{x_{ij}\}$  be the 2-D DFT of  $x_{ij}$ . For convenience we define a set of filters in the frequency domain by wedge/ring regions. The output from the  $n$ -th filter is

$$z_{ij}(n) = \mathfrak{F}^{-1} \left\{ \sum_{I_r} \sum_{I_\theta} X_{p=r \sin \theta, q=r \cos \theta} \right\} \quad (5)$$

where  $r$  is the radial frequency, and  $\theta$  is its direction. Each filter is defined by two intervals – one in angle and the other in frequency. For example,  $I_\theta = [0, 2\pi]$  and  $I_r = [r_n, r_{n+1}]$ ,  $n = 0, 1, \dots$  define a set of concentric ring filters which provide a local estimate of the power spectral density (PSD) of the texture;  $I_\theta = [\theta_n, \theta_{n+1}]$ ,  $n = 0, 1, \dots$  and  $I_r = [1, N/2]$  are a set of directional (wedge) filters, which measure its anisotropy.

With a slight change in notion, letting  $\mathbf{z}_{ij} = [z_{ij}(0), z_{ij}(1), \dots, z_{ij}(N-1)]$ , the Mahalanobis distance (3) now measures the similarity of the image in the vicinity of  $(i, j)$  to the overall average. This provides a means for detecting anomalies in space as deviations from the average background texture. As in CBAD more than one cluster can be defined to adapt to multiple backgrounds.

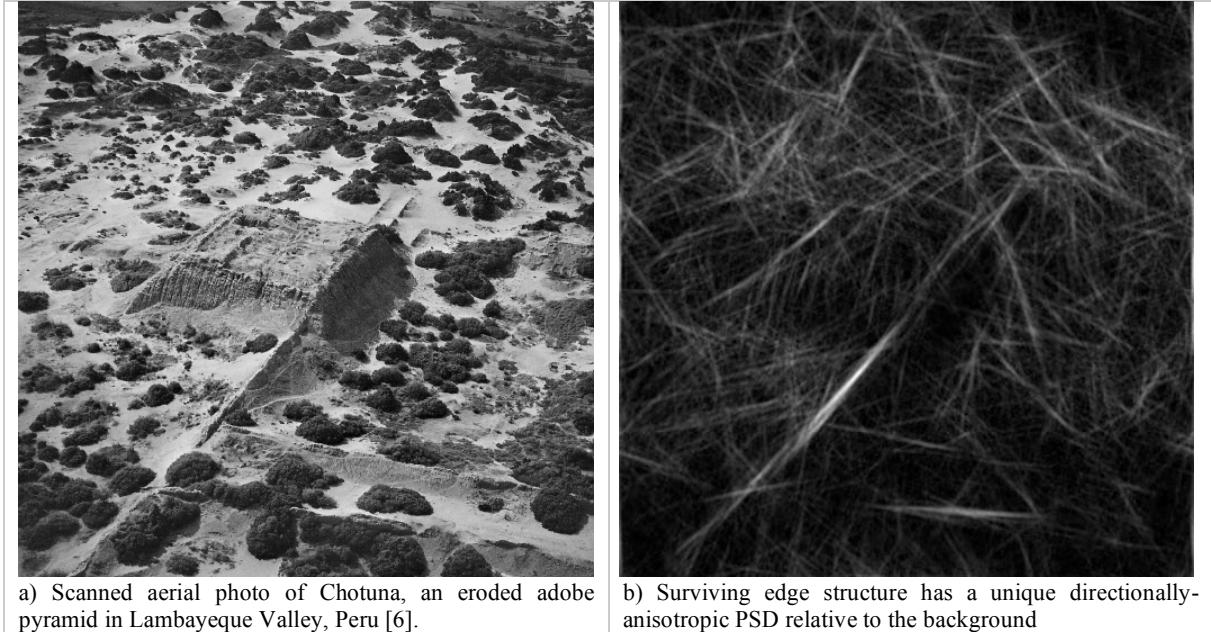


Fig. 3 Linear feature detection in unstructured background (Image courtesy Marillyn Bridges/Aperture Foundation)

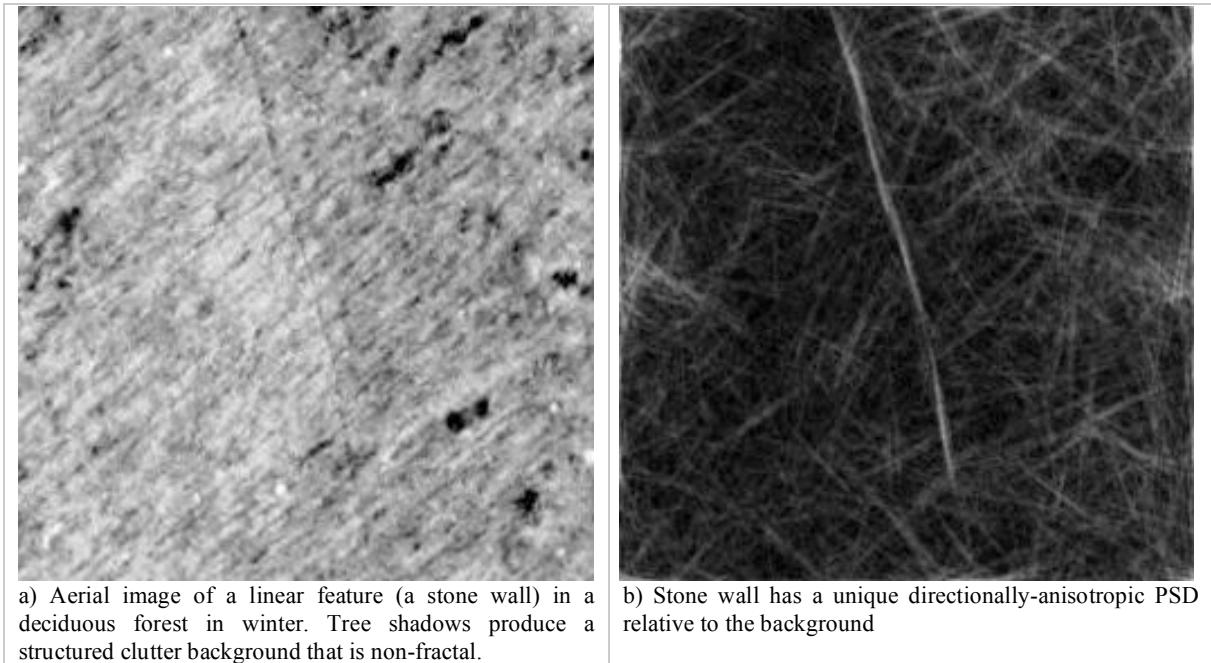


Fig. 4 Linear feature detection in structured background (Aerial image courtesy USGS)

Fig. 3 shows the detection of an extended linear feature within an eroded archaeological site in South America. Here we use a set of wedge (directional) filters to detect areas that are different from the background texture. One might ask how is the result in Fig. 5b any different from simply running a linear feature detector? Fig. 4 is an image of a stone wall that runs through a forested area. With the leaves off in winter, the sun creates tree shadows that produce a highly directional background texture. The stone wall is detected because its directional components are different from those of the background.

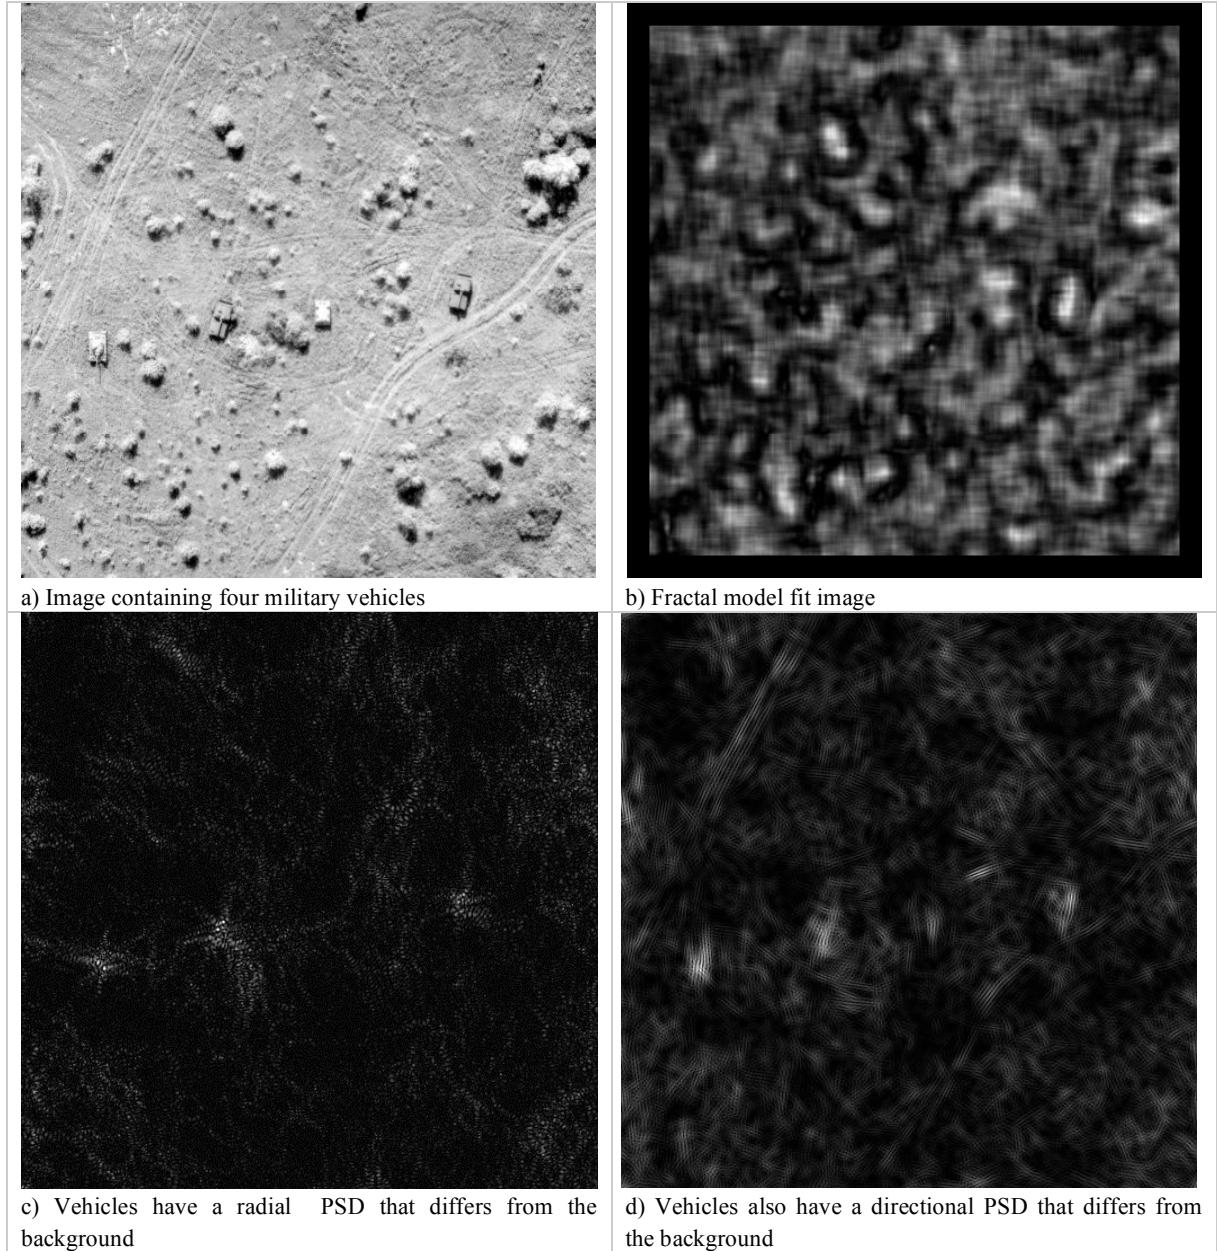


Fig. 5 Manmade object detection in cluttered background

#### 4. COMPARISON WITH FRACTAL OBJECT DETECTION

Fractal object detection assumes natural image backgrounds are fractal, and manmade objects can be detected as non-fractal deviations [4]. Unfortunately the fractal assumption is rarely true, with backgrounds often having multi-fractal or even non-fractal structure. The result of this can be seen in Fig. 5b where the non-fractal character of shadow edges makes it difficult to differentiate between vehicles and bushes.

Our anomaly detector is fundamentally different from the fractal approach. Instead of making explicit assumptions about the background, we measure the local PSD and detect areas that differ from the average. As seen in Fig. 5, the vehicles have a relatively unique radial PSD (c) and directional PSD (d) compared to that of bushes and shadows. Our method is thus effective both in images whose backgrounds are approximately fractal (Fig. 3) as well as those that are not (such as the directionally-textured background in Fig. 4).

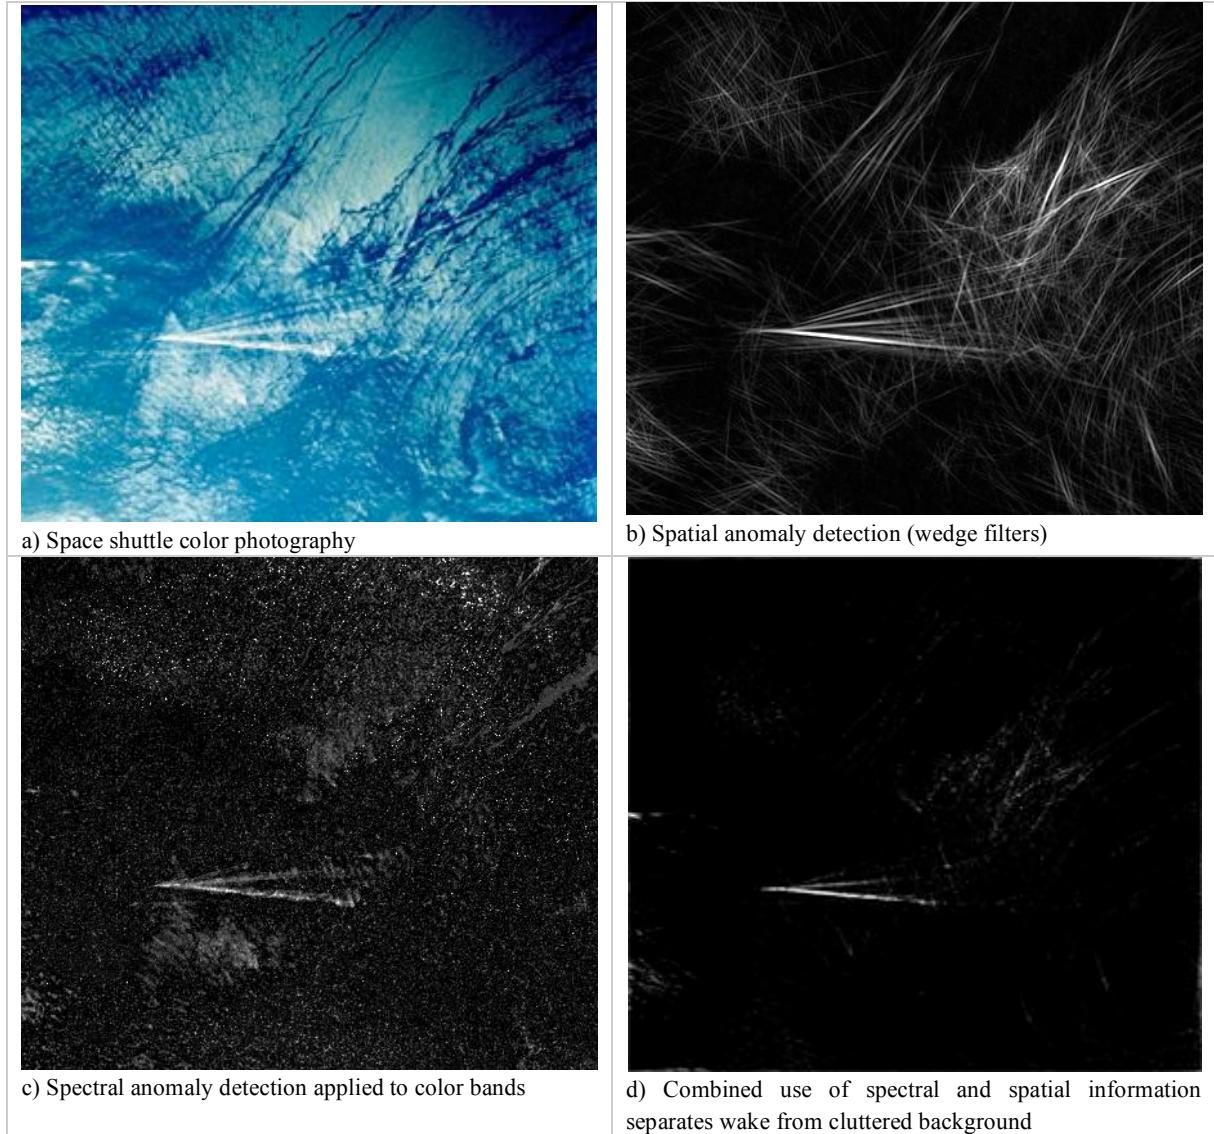


Fig. 6 Example of spatial-spectral anomaly detection (Space Shuttle color photography courtesy NASA)

## 5. COMBINED (SPECTRAL-SPATIAL) ANOMALY DETECTION

It is also possible to detect anomalies in multiple modalities. Fig. 6 is an example of ship-wake detection using spatial and spectral information. The image, which was taken over the Bahamas is highly cluttered by underwater (bathymetric) structure, clouds, eddies, and other features. Individual anomaly detectors contain a significant number of false alarms. The spatial detector responds to the shape of the wake and to lineaments and striated structures in the image. The spectral anomaly detector false alarms on certain clouds. Combining the two using ‘and’-fusion effectively separates the wake from the background.

## 6. SUMMARY

A new feature-based anomaly detection technique was described and examples from several domains presented. It assumes the frequency of anomalies is small relative to the background, which can be represented by a mixture of Gaussian components. The method is an extension of the cluster-based anomaly detector (CBAD), which was originally developed for spectral anomaly detection.

## References

- [1] David Stein, Scott Beaven, Lawrence Hoff, Edwin Winter, Alan Schaum, and Alan Stocker, “Anomaly detection from hyperspectral imagery,” *IEEE Signal Processing Magazine*, pp 58-69, Jan. 2002.
- [2] Jian Yeu Chen and I.S. Reed, “A detection algorithm for optical targets in clutter,” *IEEE Trans. Aerospace and Electronic Systems*, Vol. AES-23, No. 1, Jan. 1987.
- [3] C. W. Therrien, T. F. Quatieri and D. D. Dudgeon, "Statistical model-based algorithms for image analysis," *Proc. IEEE*, Vol. 74, No. 4, April 1986.
- [4] M. C. Stein, "Fractal image models and object detection," *Proc. Society of Photo-optical Instrumentation Engineers*, Vol 845, pp 293-300, 1987.
- [5] Mark J. Carlotto, “A Cluster-based Approach for Detecting Manmade Objects and Changes in Imagery,” *IEEE Trans. Geoscience and Remote Sensing*, Vol. 43, No. 2, Feb. 2005.
- [6] Marilyn Bridges, Planet Peru, An Aerial Journey through a Timeless Land, Kodak/Aperture Books, 1991.