

Uncorrelated versus Independent Elliptically-Contoured Distributions for Anomalous Change Detection in Hyperspectral Imagery

James Theiler and Clint Scovel

Los Alamos National Laboratory, Los Alamos, NM, USA

ABSTRACT

The detection of actual changes in a pair of images is confounded by the inadvertent but pervasive differences that inevitably arise whenever two pictures are taken of the same scene, but at different times and under different conditions. These differences include effects due to illumination, calibration, misregistration, etc. If the actual changes are assumed to be rare, then one can “learn” what the pervasive differences are, and can identify the deviations from this pattern as the anomalous changes. A recently proposed framework for anomalous change detection recasts the problem as one of binary classification between pixel pairs in the data and pixel pairs that are independently chosen from the two images. When an elliptically-contoured (EC) distribution is assumed for the data, then analytical expressions can be derived for the measure of anomalousness of change. However, these expressions are only available for a limited class of EC distributions. By replacing independent pixel pairs with uncorrelated pixel pairs, an approximate solution can be found for a much broader class of EC distributions. The performance of this approximation is investigated analytically and empirically, and includes experiments comparing the detection of real changes in real data.

Keywords: change, hyperspectral, elliptically-distributed, covariance

1. INTRODUCTION

Each pixel in a hyperspectral image consists of a radiance (or, with some processing, reflectance) spectrum with typically hundreds of high signal-to-noise-ratio channels, each channel corresponding to a very narrow wavelength band. Such exquisitely detailed datasets provide an opportunity for precise discrimination of constituent materials from remote platforms. The precision also permits the detection of weak signals, as from subpixel targets or gaseous plumes, in broad area surveys.

Although researchers often bemoan the *deluge* of data provided by hyperspectral imagery, the fact that there is so much information *in each pixel* can actually simplify the analysis. In fact, most algorithms for hyperspectral image analysis effectively treat the image as a “bag of pixels” – with each pixel treated as an independent sample. Although the spatial correlations are nontrivial and not inconsiderable, useful analysis can often be obtained even while neglecting these correlations.

Not only are the pixels treated as if they were independent samples from a common distribution, that distribution is usually assumed to be a multivariate Gaussian. This is simplistic, but it does capture some important aspects of the data. The covariance matrix that characterizes the Gaussian encapsulates the (linear) correlations between every pair of image channels. The dynamical range between the largest and smallest eigenvalues of the covariance matrix can span many orders of magnitude.

One important aspect of hyperspectral image data which is *poorly* captured by the Gaussian model is its behavior on the tails of the distribution. It is widely recognized that the tails of most hyperspectral datasets are much fatter than Gaussian. And since detection of rare signals requires comparison out at the tails of the distribution, this is particularly important for anomaly detection applications.

The problem of “white balance” that bedevils amateur (and professional) photographers is particularly onerous in hyperspectral imagery. The problem is that the observed spectrum for a given material (whose “actual color” is fixed) will be different when viewed under even slightly different conditions (of illumination, sensor calibration, atmospheric distortion, *etc.*). For target detection applications, this means that the effective target signatures vary from image to image.¹ And for the change detection problem, it confounds the search for actual changes because under different conditions, the spectrum of every pixel changes.

2. MACHINE LEARNING FRAMEWORK FOR ANOMALOUS CHANGE DETECTION

Given two images, call them the x-image and the y-image, the aim is to find those few pixels for which the x-to-y change is unusual compared to the changes exhibited by the rest of the pixels.

Let $\mathbf{x} \in \mathcal{R}^{d_x}$ denote (the observed radiance spectrum observed at) a pixel in the x-image, and $\mathbf{y} \in \mathcal{R}^{d_y}$ be the corresponding pixel in the y-image. We assume that the images are registered (*i.e.*, that corresponding pixels \mathbf{x} and \mathbf{y} correspond to the same location in the scene), but we acknowledge that this registration is not always precise.^{2,3} Here, d_x and d_y are the number of spectral channels in the x-image and y-image, respectively.

In the machine learning framework introduced in Ref. 4, the data is modeled as random samples from a probability distribution $P(\mathbf{x}, \mathbf{y})$. In this model, straight anomaly detection seeks points \mathbf{x}, \mathbf{y} on the “tail” of the distribution; that is, where $P(\mathbf{x}, \mathbf{y})$ is small. But straight anomaly detection identifies pixels where \mathbf{x} and \mathbf{y} are individually unusual (*e.g.*, they might correspond to particularly dark or particularly bright pixels), as well as pixels where the *relationship* between \mathbf{x} and \mathbf{y} is unusual. If we write $P(\mathbf{x})$ as the distribution just of the pixels in the x-image, then this $P(\mathbf{x})$ will be the marginal distribution of $P(\mathbf{x}, \mathbf{y})$. We can similarly write $P(\mathbf{y})$ as the distribution of pixels in the y-image. Then the product $P(\mathbf{x})P(\mathbf{y})$ describes a distribution of \mathbf{x} and \mathbf{y} values that are independent of each other. When $P(\mathbf{x})P(\mathbf{y})$ is small, that means that either \mathbf{x} or \mathbf{y} (or both) are individually unusual, but does not saying anything about whether the relationship between them is unusual. When the ratio

$$\frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})P(\mathbf{y})}. \quad (1)$$

is small, that means not only that $P(\mathbf{x}, \mathbf{y})$ is small, but that it is small compared to $P(\mathbf{x})P(\mathbf{y})$. That is, the pair \mathbf{x}, \mathbf{y} is more anomalous than would be expected, given the individual anomalousnesses of \mathbf{x} and \mathbf{y} . This enables us to isolate the notion of anomalous change from that of straight anomaly.

For a given ρ , we can define a set of anomalies as

$$A_\rho = \left\{ (\mathbf{x}, \mathbf{y}) \mid \frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})P(\mathbf{y})} < \rho \right\}. \quad (2)$$

In seeking a function $\mathcal{A}(\mathbf{x}, \mathbf{y})$ which quantifies the “anomalousness” of the change that has occurred at this pixel location, we can take a function of this ratio

$$\mathcal{A}(\mathbf{x}, \mathbf{y}) = f \left(\frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})P(\mathbf{y})} \right) \quad (3)$$

where f is a monotonically decreasing function of its argument. When the ratio is small, the anomalousness is large, and in particular, the set A_ρ defined in Eq. (2) is given by those \mathbf{x}, \mathbf{y} for which $\mathcal{A}(\mathbf{x}, \mathbf{y}) > f(\rho)$.

2.1. Gaussian Model

The ratio in Eq. (1) takes a simple form when the distribution is modeled as a multivariate Gaussian. In general, a d -dimensional Gaussian depends on a d -dimensional mean $\boldsymbol{\mu} \in \mathcal{R}^d$ and a covariance matrix $K \in \mathcal{R}^{d \times d}$. We can write

$$\boldsymbol{\mu} = \langle \mathbf{z} \rangle, \text{ and} \quad (4)$$

$$K = \langle (\mathbf{z} - \boldsymbol{\mu})(\mathbf{z} - \boldsymbol{\mu})^T \rangle \quad (5)$$

where the angle brackets denote a mean over the distribution (in practice, these quantities are estimated by taking a sample average over the data), and the superscript T denotes a matrix transpose. The density of the distribution at a point $\mathbf{z} \in \mathcal{R}^d$ is given by (*e.g.*, see Eq. (2.5) in Kay⁵)

$$P(\mathbf{z}) = (2\pi)^{-d/2} |K|^{-1/2} \exp \left[-\frac{1}{2} (\mathbf{z} - \boldsymbol{\mu})^T K^{-1} (\mathbf{z} - \boldsymbol{\mu}) \right]. \quad (6)$$

It is useful to more specifically identify a stacked vector

$$\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \in \mathcal{R}^{d_x+d_y} \quad (7)$$

which denotes a corresponding pixel pair in both images. This leads to

$$\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_x \\ \boldsymbol{\mu}_y \end{bmatrix} \quad (8)$$

where $\boldsymbol{\mu}_x = \langle \mathbf{x} \rangle$ and $\boldsymbol{\mu}_y = \langle \mathbf{y} \rangle$. Further, we can write

$$K = \begin{bmatrix} X & C^T \\ C & Y \end{bmatrix} \quad (9)$$

where

$$X = \langle (\mathbf{x} - \boldsymbol{\mu}_x)(\mathbf{x} - \boldsymbol{\mu}_x)^T \rangle, \quad (10)$$

$$Y = \langle (\mathbf{y} - \boldsymbol{\mu}_y)(\mathbf{y} - \boldsymbol{\mu}_y)^T \rangle, \text{ and} \quad (11)$$

$$C = \langle (\mathbf{y} - \boldsymbol{\mu}_y)(\mathbf{x} - \boldsymbol{\mu}_x)^T \rangle. \quad (12)$$

The marginal distributions $P(\mathbf{x})$ and $P(\mathbf{y})$ are also Gaussian, and are given by

$$P(\mathbf{x}) = (2\pi)^{-d_x/2} |X|^{-1/2} \exp \left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_x)^T X^{-1}(\mathbf{x} - \boldsymbol{\mu}_x) \right], \text{ and} \quad (13)$$

$$P(\mathbf{y}) = (2\pi)^{-d_y/2} |Y|^{-1/2} \exp \left[-\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_y)^T Y^{-1}(\mathbf{y} - \boldsymbol{\mu}_y) \right]. \quad (14)$$

Finally, we can combine Eq. (6) with Eq. (13) and Eq. (14) to express the ratio in Eq. (1)

$$\begin{aligned} \frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})P(\mathbf{y})} &= \frac{(2\pi)^{-(d_x+d_y)/2} |K|^{-1/2} \exp \left[-\frac{1}{2}(\mathbf{z} - \boldsymbol{\mu})^T K^{-1}(\mathbf{z} - \boldsymbol{\mu}) \right]}{(2\pi)^{-(d_x+d_y)/2} |X|^{-1/2} |Y|^{-1/2} \exp \left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_x)^T X^{-1}(\mathbf{x} - \boldsymbol{\mu}_x) - \frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_y)^T Y^{-1}(\mathbf{y} - \boldsymbol{\mu}_y) \right]} \\ &= \left[\frac{|K|}{|X||Y|} \right]^{-1/2} \exp \left[-\frac{1}{2}(\xi_z - \xi_x - \xi_y) \right] \end{aligned} \quad (15)$$

where the three scalars ξ_x , ξ_y and ξ_z are given by

$$\xi_x = (\mathbf{x} - \boldsymbol{\mu}_x)^T X^{-1}(\mathbf{x} - \boldsymbol{\mu}_x), \quad (16)$$

$$\xi_y = (\mathbf{y} - \boldsymbol{\mu}_y)^T Y^{-1}(\mathbf{y} - \boldsymbol{\mu}_y), \text{ and} \quad (17)$$

$$\xi_z = (\mathbf{z} - \boldsymbol{\mu})^T K^{-1}(\mathbf{z} - \boldsymbol{\mu}). \quad (18)$$

Since $|K|$, $|X|$, and $|Y|$ are constants that do not depend on \mathbf{x} or \mathbf{y} , we can derive a simple expression for anomalousness using Eq. (3) with $f(r) = -2\log(r) - \log(|K|/(|X||Y|))$:

$$\mathcal{A}(\mathbf{x}, \mathbf{y}) = -2\log \left[\frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})P(\mathbf{y})} \right] - \log \left[\frac{|K|}{|X||Y|} \right] = \xi_z - \xi_x - \xi_y. \quad (19)$$

Equivalently, $\mathcal{A}(\mathbf{x}, \mathbf{y}) = (\mathbf{z} - \boldsymbol{\mu})^T Q(\mathbf{z} - \boldsymbol{\mu})$, where the quadratic coefficient matrix is given by

$$Q = \begin{bmatrix} X & C^T \\ C & Y \end{bmatrix}^{-1} - \begin{bmatrix} X & 0 \\ 0 & Y \end{bmatrix}^{-1}. \quad (20)$$

2.2. Elliptically-Contoured Distributions

The class of elliptically-contoured (EC) distributions has found utility both for radar⁶ and hyperspectral imagery.^{7,8} These distributions depend on the covariance matrix K and are of the form

$$P(K; \mathbf{z}) = |K|^{-1/2} H(d, \xi_z) \quad (21)$$

where $|K|$ is the determinant of the covariance matrix K , and H is a function which depends on the dimension d of the vector \mathbf{z} , and on \mathbf{z} via the scalar $\xi_z = (\mathbf{z} - \boldsymbol{\mu})^T K^{-1} (\mathbf{z} - \boldsymbol{\mu})$, which is the squared Mahalanobis distance to the centroid of the data. Note that for the Gaussian distribution, $H(d, \xi) = (2\pi)^{-d/2} e^{-\xi/2}$.

2.2.1. Consistent families

A consistent family $H(d, \xi_z)$ has the property that: if $P(\mathbf{z}) = |K|^{-1/2} H(d, \xi_z)$, where \mathbf{z} is the stacked vector in Eq. (7), then $P(\mathbf{x}) = |X|^{-1/2} H(d_x, \xi_x)$ is the marginal distribution associated with the projection of \mathbf{z} onto the $d_x < d$ dimensional subspace corresponding to \mathbf{x} . Given $H(d, \xi)$ for a given d , there exists a consistent family of lower dimensional distributions,⁹ given by

$$H(d', \xi) = c(d', d) \int_0^\infty w^{(d-d')/2-1} H(d, w + \xi) dw, \quad (22)$$

where $c(d', d)$ is a scalar constant that ensures that the distribution is normalized.

For a consistent family, we can write an explicit expression for the ratio in Eq. (1):

$$\frac{P(\mathbf{x}, \mathbf{y})}{P(\mathbf{x})P(\mathbf{y})} = \left[\frac{|K|}{|X||Y|} \right]^{-1/2} \frac{H(d, \xi_z)}{H(d_x, \xi_x)H(d_y, \xi_y)} \quad (23)$$

and from that derive a closed-form expression for anomalousness $\mathcal{A}(\mathbf{x}, \mathbf{y})$.

The Gaussian is an example of a consistent family, and as already seen in Eq. (19), provides a simple anomalous change detector. Another example of a consistent family is the multivariate- t statistic.^{7,9} Here,

$$H(d, \xi) = \frac{\Gamma\left(\frac{d+\nu}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) \pi^{d/2} (\nu-2)^{d/2}} \left(1 + \frac{\xi}{\nu-2}\right)^{-(d+\nu)/2}. \quad (24)$$

This is a fatter tailed distribution than the Gaussian, and it gets fatter as ν gets smaller. In fact, as $\nu \rightarrow 2$, the variance diverges. The limit $\nu \rightarrow \infty$ recovers the Gaussian distribution.

The multivariate- t leads to an anomalousness measure (using Eq. (3) with $f(r) = -2\log(r) + \text{constant}$)

$$\mathcal{A}(\nu; \mathbf{x}, \mathbf{y}) = (d_x + d_y + \nu) \log(\xi_z + \nu - 2) - (d_x + \nu) \log(\xi_x + \nu - 2) - (d_y + \nu) \log(\xi_y + \nu - 2) \quad (25)$$

that may be more effective than Eq. (19) when the data is fatter tailed than Gaussian.

One simplifying limit takes place for $d_x = d_y \gg \nu$ and $\nu \rightarrow 2$. In this limit, Eq. (25) becomes $\mathcal{A}(\mathbf{x}, \mathbf{y}) = 2d \log(\xi_z / \sqrt{\xi_x \xi_y})$, which can be monotonically transformed to:

$$\mathcal{A}(\mathbf{x}, \mathbf{y}) = \frac{\xi_z}{\sqrt{\xi_x \xi_y}}. \quad (26)$$

Although the expression in Eq. (22) ensures the *existence* of a family of distributions, it does not say that the family can be expressed in a tractable closed-form. For instance, a popular choice of EC distribution is given by the generalized Gaussian:

$$H(d, \xi) = c(d, \beta, \gamma) \exp(-\gamma \xi^\beta) \quad (27)$$

with $c(d, \beta, \gamma)$ a scalar constant. Here $\beta = 1$ produces the Gaussian distribution, and $\beta < 1$ is a fatter tailed distribution. However the generalized Gaussian does not satisfy the condition in Eq. (22), and it is *not* a consistent family.

One can, however, take the expression in Eq. (27) for a specific value of d and derive a consistent family of distributions for other values of d , but the corresponding expressions for these other values of d will not have the nice form in Eq. (27).*

3. UNCORRELATION AS AN APPROXIMATION TO INDEPENDENCE

It is for these inconsistent families that we have introduced the concept of uncorrelated versus independent distributions as a denominator in the ratio in Eq. (1). Here, if

$$P(K; \mathbf{x}, \mathbf{y}) = P\left(\begin{bmatrix} X & C^T \\ C & Y \end{bmatrix}; \mathbf{x}, \mathbf{y}\right) \quad (28)$$

is an EC distribution for the stacked vector $\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}$, then we will approximate the product of the marginals $P(\mathbf{x})P(\mathbf{y})$ which corresponds to independence of \mathbf{x} and \mathbf{y} with an EC distribution in which \mathbf{x} and \mathbf{y} are merely uncorrelated:

$$P(\mathbf{x})P(\mathbf{y}) \approx P_u(\mathbf{x}, \mathbf{y}) = P\left(\begin{bmatrix} X & 0 \\ 0 & Y \end{bmatrix}; \mathbf{x}, \mathbf{y}\right). \quad (29)$$

That is,

$$P_u(\mathbf{x}, \mathbf{y}) = \left| \begin{array}{cc} X & 0 \\ 0 & Y \end{array} \right|^{-1/2} H(d, \xi'), \quad (30)$$

where

$$\xi' = (\mathbf{z} - \boldsymbol{\mu})^T \begin{bmatrix} X & 0 \\ 0 & Y \end{bmatrix}^{-1} (\mathbf{z} - \boldsymbol{\mu}) = \xi_x + \xi_y. \quad (31)$$

We remark that the projection of $P_u(\mathbf{x}, \mathbf{y})$ to the \mathbf{x} or \mathbf{y} space produces marginals $P(\mathbf{x})$ and $P(\mathbf{y})$ that are the same as the marginals of the parent distribution $P(\mathbf{x}, \mathbf{y})$. The anomalousness measure then varies inversely with the modified ratio

$$\frac{P(\mathbf{x}, \mathbf{y})}{P_u(\mathbf{x}, \mathbf{y})} = \left[\frac{|K|}{|X||Y|} \right]^{-1/2} \frac{H(d, \xi_z)}{H(d, \xi_x + \xi_y)}. \quad (32)$$

For the Gaussian case, we have that $P_u(\mathbf{x}, \mathbf{y}) = P(\mathbf{x})P(\mathbf{y})$ exactly, and we get the same anomalousness measure as the independent case in Eq. (19). For the multivariate- t in Eq. (24), using $f(r) = \text{constant} \times r^{-2/(d+\nu)}$, we obtain

$$\mathcal{A}(\nu; \mathbf{x}, \mathbf{y}) = \frac{\xi_z + \nu - 2}{\xi_x + \xi_y + \nu - 2}. \quad (33)$$

which is substantially simpler than Eq. (25), and has no dependence on d_x or d_y . The fat tailed $\nu \rightarrow 2$ limit produces

$$\mathcal{A}(\mathbf{x}, \mathbf{y}) = \frac{\xi_z}{\xi_x + \xi_y}, \quad (34)$$

which recalls Eq. (26), except that the harmonic mean is replaced by an arithmetic mean. We also remark that the $\nu \rightarrow \infty$ limit of Eq. (33) leads to the Gaussian anomalousness measure defined in Eq. (19).

The main utility of this approximation is that there are no constraints on the function $H(d, \xi)$ that appears in Eq. (32). Since both numerator and denominator use the same dimension d , there is no need to identify a consistent family. For instance, even though the family of generalized Gaussian distributions in Eq. (27) is inconsistent, we can still write a simple expression for the uncorrelation-based anomalous change detector:

$$\mathcal{A}(\beta; \mathbf{x}, \mathbf{y}) = \xi_z^\beta - (\xi_x + \xi_y)^\beta. \quad (35)$$

Here $\beta = 1$ corresponds to a Gaussian distribution. In the fat-tailed limit, as $\beta \rightarrow 0$, the anomalousness expression becomes $\mathcal{A}(\mathbf{x}, \mathbf{y}) = \xi_z/(\xi_x + \xi_y)$ which agrees with the fat-tailed limit of the multivariate- t in Eq. (34).

An alternative interpretation of this approximation is given in the Appendix.

*We can show that it *is* possible to create a family of distributions which for even d take the form $H(d, \xi) = G(d, \beta, \gamma; \xi) \exp(-\gamma\xi^\beta)$, where $G(d, \beta; \xi)$ is a polynomial in ξ^β . But to find the members of this family for odd d , one must fall back on the less tractable expression in Eq. (22).

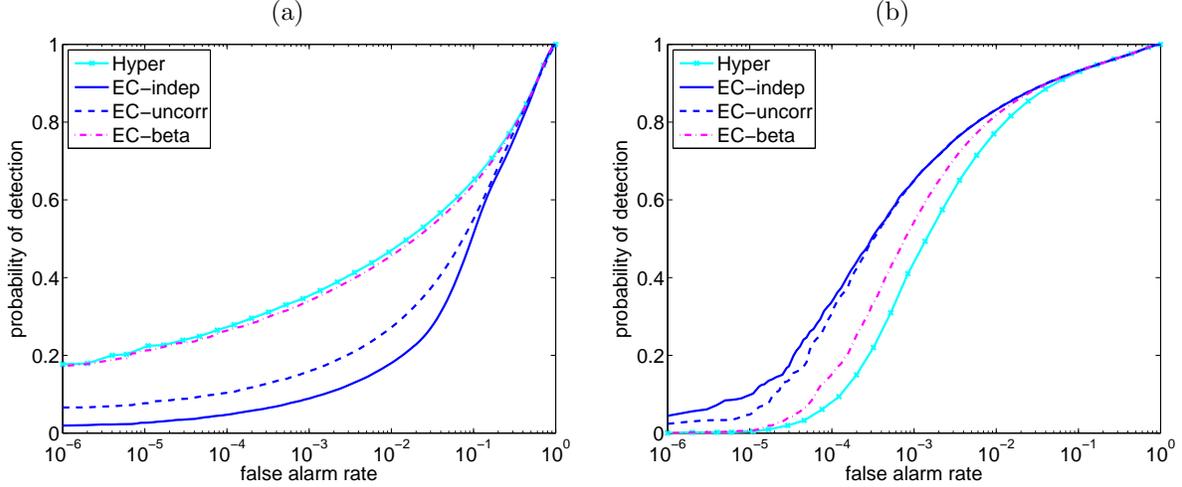


Figure 1. ROC curves for Gaussian and EC detectors applied to simulated data. In both cases, $n = 10^6$ samples are drawn from a single-channel ($d_x = d_y = 1$) distribution specified by covariances X and Y , and cross-covariance C . (a) Gaussian data is generated with $X = 2$, $Y = 1$, and $C = 1.3$. For this data, we see that the Gauss-based anomalous change detector achieves the highest performance. (b) Elliptically-contoured multivariate- t data is generated using $X = 2$, $Y = 1$, and $C = 1.41$, and $\nu = 2.1$; see Eq. (24). Here, we see that the EC detectors outperform the Gaussian detector. We also see that the detector based on the independence formula, in Eq. (25), outperforms (though only slightly) the detector based on the approximate, but simpler, formula in Eq. (33). Both plots also show the performance of the detector in Eq. (35) with $\beta = 0.5$.

4. NUMERICAL EXPERIMENTS

To illustrate the utility of EC distributions for anomalous change detection, and in particular to evaluate the effect of approximating independence with uncorrelation in this context, we provide three sets of numerical experiments: pure simulation, a hybrid simulation in which pervasive differences and anomalous changes are artificially generated in a real image, and a real pair of images which exhibits both pervasive differences in the images and some actual changes in the scene.

4.1. Pure simulation

In this experiment we generate two image pairs, one Gaussian and one elliptically contoured with very fat tails. The data are specified in terms of X , Y , and C defined in Eqs. (10-12). Single-channel images are generated using $X = 2$ and $Y = 1$. The pervasive differences between the x and y images are encapsulated in the cross-covariance term C ; for the Gaussian pair, we take $C = 1.3$ and for the EC pair, we take $C = 1.41$. Note that the closer C is to $\sqrt{2}$, the more nearly deterministic is the relationship between the pixels in the two images. For the EC distribution, we use a multivariate- t distribution with $\nu = 2.1$.

Following the simulation framework described in Ref. [10], we generate an anomalous change at a given pixel location by keeping \mathbf{x} fixed and replacing \mathbf{y} with a random sample drawn from elsewhere in the image. This way, neither \mathbf{x} nor \mathbf{y} are individually unusual, but the relationship between \mathbf{x} and \mathbf{y} breaks the pattern of correlation in the simulated image pairs.

Fig. 1(a) shows that of the four detectors, the “Hyper” detector, which is optimized for Gaussian distributions, performs best. Also shown are the generalized Gaussian (EC-beta) with $\beta = 0.5$, the independence-based multivariate t statistic (EC-indep) with $\nu = 2.1$, and the uncorrelation-based multivariate t (EC-uncorr) again with $\nu = 2.1$. Choosing ν so close to 2 produces an extreme variant of the change detector that is optimized for very fat tails.

When applied to non-Gaussian fat-tailed data, however, the EC detectors outperform Hyper. As shown in Fig. 1(b), the best detector is EC-indep, the detector that is exactly matched to the statistics of the data. But

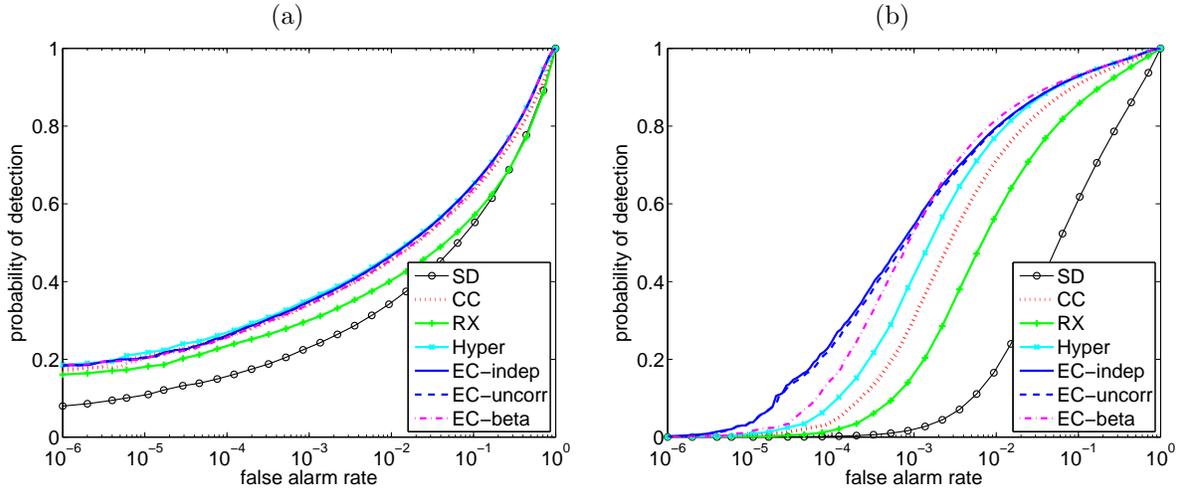


Figure 2. Same data as previous figure, but including other change detection algorithms, and using more moderate value of $\nu = 10$ instead of $\nu = 2.1$ in the EC-indep and EC-uncorr detectors.

even in this extreme case, we see that the approximation provided by EC-uncorr performs nearly as well as EC-indep. The generalized Gaussian with $\beta = 0.5$ and the Gaussian itself (Hyper) perform less well.

In Fig. 2, we revert to a less highly optimized variant of the EC detector, and use a moderate $\nu = 10$. This figure also includes the performance of some previously described detectors; these detectors are surveyed in Ref. [10], and include the simple difference (SD), the chronochrome (CC) of Schaum and Stocker,^{11,12} and a straight anomaly detector (RX) based on Mahalanobis distance from the centroid of the stacked \mathbf{z} space.¹³ For the Gaussian data in Fig. 2(a), the performance of all three of the EC detectors is virtually identical to that of Hyper, which is known to be optimal. For the EC data in Fig. 2(b), even though it is extremely fat tailed, the more moderate EC detectors (EC-indep, EC-uncorr) still perform the best, and are virtually identical in their performance. Comparing the moderate EC detectors in Fig. 2(b) with the extreme EC detectors in Fig. 1(b), we see that the moderate detectors perform nearly as well.

4.2. Simulating pervasive differences and anomalous changes

In this hybrid simulation, we begin with real 224-channel hyperspectral data from the AVIRIS sensor.^{14,15} See Fig. 3. Using the simulation framework outlined in Ref. 10, we generate pervasive differences by applying some operation to the whole scene. For the results shown in Fig. 4, four operations were considered: multiplicative noise, splitting the image into two 112-channels images, smoothing the image with a 3×3 kernel, and (after smoothing) misregistering the image by one pixel. As with the pure simulation, anomalous changes are produced by replacing a pixel with one drawn at random from another part of the image. After this simulation, the first ten principal components are used in place of the full image as a dimension reduction measure. In all four cases, the EC detectors outperformed the Gaussian-based detectors, and in particular EC-uncorr and EC-indep performed essentially identically.

4.3. Real anomalies in real imagery

In a long running experiment, Eismann *et al.*¹⁶ made a series of hyperspectral images of a grassy field with trees in the background (see Fig. 5(a,b)). As well as the grass and trees, four panels were placed in the scene, exhibiting spectra that were unlike that of most of the background, and which might be identified as anomalies in the image. But those panels were kept in place throughout the experiment, so they did not represent anomalous changes. Periodically, a pair of tarps would be placed on the grass, and these were the anomalous changes that algorithms were challenged to find. In particular, two images were taken on October 14th, one without and one with the emplaced tarps. The images are 800×1024 pixels, and have 124 spectral channels. Following the



Figure 3. Broadband image of AVIRIS data over the Florida coastline. Shown here is the first principal component of the 224-channel image, which is 150×500 pixels.

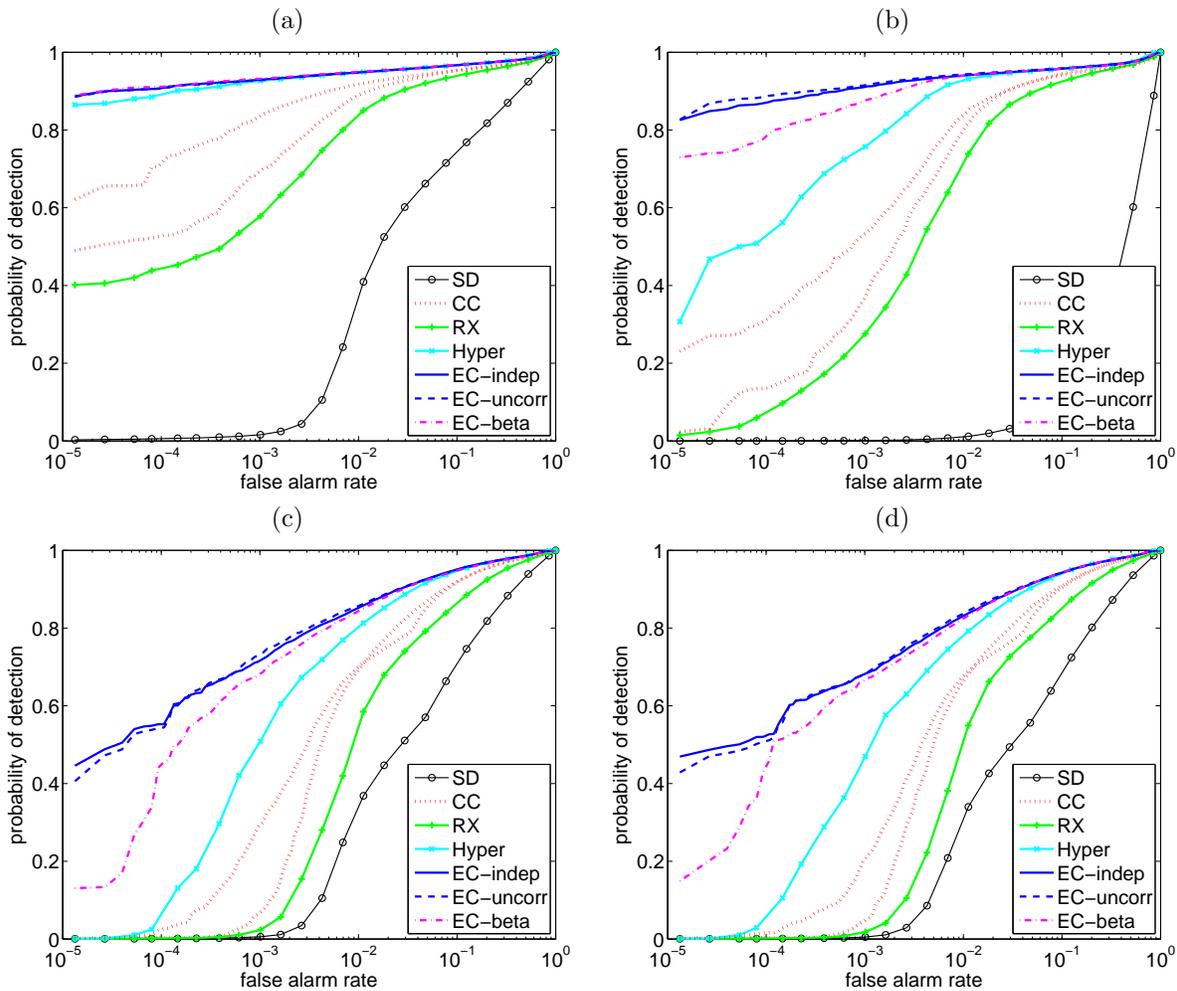


Figure 4. ROC curves produced by various anomalous change detectors, applied to the AVIRIS data shown in Fig. 3 in which pervasive differences have been simulated in four ways: (a) multiplicative noise (each pixel is multiplied by a value randomly chosen from the interval $[1,2]$); (b) the image bands are split into two groups, one with the first 112 bands of the image, and one with the last 112 bands; (c) the image is smoothed with a 3×3 kernel; and (d) the image (after smoothing) is misregistered by one pixel.

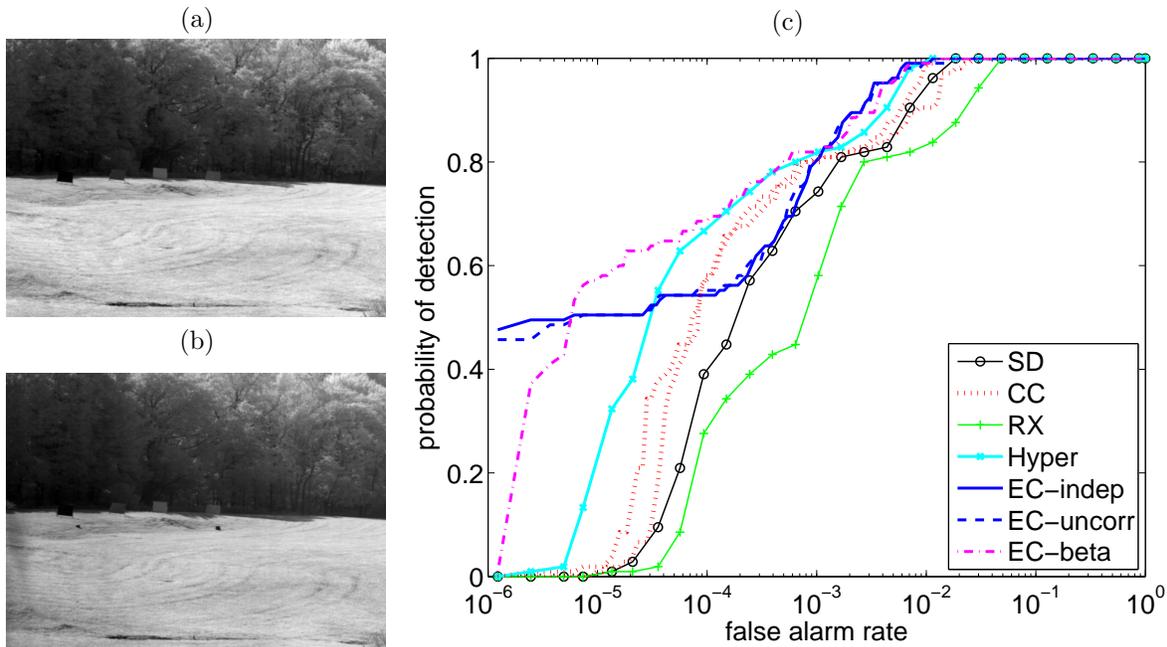


Figure 5. (a,b) First principal component of two hyperspectral images taken of a grassy field with trees in the background. There is a set of four panels on the horizon line that are in both images, and some actual changes which are evident in the second image (b) as two darker spots, halfway up the image, with one near the center and the other to the left. (c) ROC curves for the detection of the actual changes seen in these two images.

approach taken by Meola and Eismann,² we reduce that to ten channels each, using the principal components computed from a third image (taken August 25).

Fig. 5(c) compares the performance of the different detectors applied to this data. Although one should be cautious about drawing conclusions from single examples, we see in this case that the EC detectors outperform the Gaussian (Hyper) detector in the very low false-alarm rate regime, but the multivariate- t based detectors (EC-indep and EC-uncorr) perform more poorly in the intermediate false-alarm-rate regime. The detector based on a generalized Gaussian (EC-beta) outperforms Hyper over the whole range. We recall that EC-beta is an uncorrelation-based detector and that we found that a formally optimal generalized Gaussian detector (*i.e.*, one that used independence of $P(\mathbf{x})P(\mathbf{y})$) was intractable, but that the uncorrelation approximation that is so successful on this data is given by the very simple formula in Eq. (35).

5. VISUALIZATION

The simple form given in Eq. (32) ensures that any EC-uncorr anomalous change detector will be a function only of the two scalars ξ_z and $\xi_x + \xi_y$. This permits a two-dimensional visualization of the data, as shown in Fig. 6. It is evident, in these plots, that the relative performance of different detectors depends on how heavy-tailed the distribution of data is.

The reason EC-uncorr is able to produce a two-dimensional plot is that the expression for “individual anomalousness” depends only on $\xi_x + \xi_y$. For the EC-indep algorithms, on the other hand, the individual anomalousness term depends on more general functions of ξ_x and ξ_y . Fig. 7 compares the shapes of the contours, in $\xi_x - \xi_y$ space, that are obtained with these two different approaches to EC anomalous change detection.

Acknowledgments

We are grateful to Joseph Meola and Michael Eismann for providing hyperspectral data with ground-truth changes. This work was supported by the Laboratory Directed Research and Development (LDRD) program at Los Alamos National Laboratory.

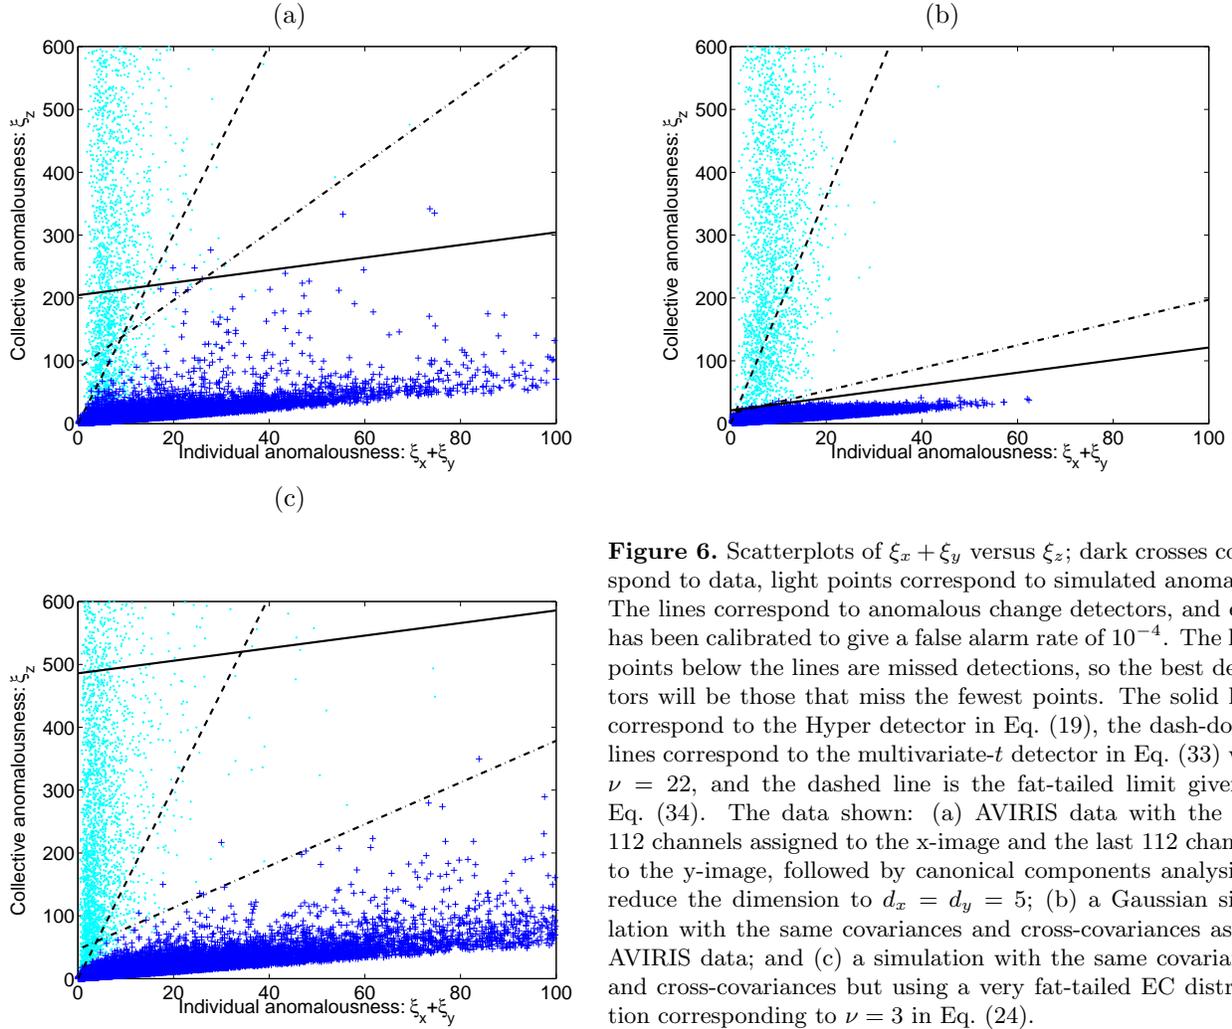


Figure 6. Scatterplots of $\xi_x + \xi_y$ versus ξ_z ; dark crosses correspond to data, light points correspond to simulated anomalies. The lines correspond to anomalous change detectors, and each has been calibrated to give a false alarm rate of 10^{-4} . The light points below the lines are missed detections, so the best detectors will be those that miss the fewest points. The solid lines correspond to the Hyper detector in Eq. (19), the dash-dotted lines correspond to the multivariate- t detector in Eq. (33) with $\nu = 22$, and the dashed line is the fat-tailed limit given in Eq. (34). The data shown: (a) AVIRIS data with the first 112 channels assigned to the x-image and the last 112 channels to the y-image, followed by canonical components analysis to reduce the dimension to $d_x = d_y = 5$; (b) a Gaussian simulation with the same covariances and cross-covariances as the AVIRIS data; and (c) a simulation with the same covariances and cross-covariances but using a very fat-tailed EC distribution corresponding to $\nu = 3$ in Eq. (24).

REFERENCES

1. R. Mayer, F. Bucholtz, and D. Scribner, "Object detection by using "whitening/dewhitening" to transform target signatures in multitemporal hyperspectral and multispectral imagery," *IEEE Trans. Geosci. Remote Sens.* **41**, pp. 1136–1142, 2003.
2. J. Meola and M. T. Eismann, "Image misregistration effects on hyperspectral change detection," *Proc. SPIE* **6966**, p. 69660Y, 2008.
3. J. Theiler, "Sensitivity of anomalous change detection to small misregistration errors," *Proc. SPIE* **6966**, p. 69660X, 2008.
4. J. Theiler and S. Perkins, "Proposed framework for anomalous change detection," *ICML Workshop on Machine Learning Algorithms for Surveillance and Event Detection*, pp. 7–14, 2006.
5. S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, vol. II, Prentice Hall, New Jersey, 1998.
6. E. Conte and M. Longo, "Characterisation of radar clutter as a spherically invariant random process," *IEE Proc. F (Communications, Radar and Signal Processing)* **134**, pp. 191–197, 1987.
7. D. B. Marden and D. Manolakis, "Modeling hyperspectral imaging data," *Proc. SPIE* **5093**, pp. 253–262, 2003.

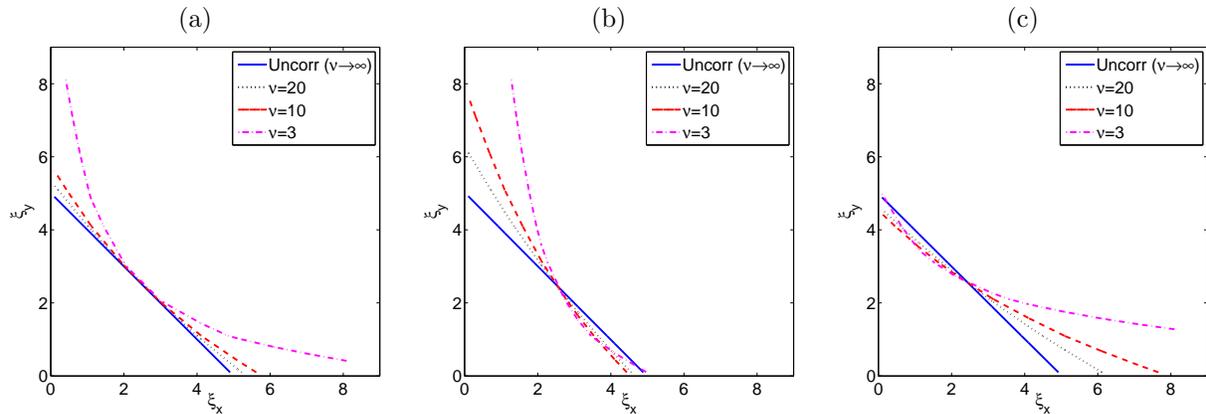


Figure 7. Contours in ξ_x - ξ_y space (for fixed ξ_z) for EC-based ACD algorithms based on the multivariate- t statistic. The three cases shown here correspond to: (a) $d_x = d_y = 10$, (b) $d_x = 10$, $d_y = 3$, and (c) $d_x = 3$, $d_y = 10$. The uncorrelation approximation corresponds to contours in which $\xi_x + \xi_y = \text{constant}$, which agrees with the independence results in the $\nu \rightarrow \infty$ limit. Only a single contour is shown, and in all cases it is the the contour that includes the point $\xi_x = \xi_y = 2.5$. The level that is associated with a particular false alarm rate will depend on the data, but these figures show how the shapes of the contours vary with ν and how they are approximated by the straight line. It also appears that the approximation works best when $d_x = d_y$.

8. A. Schaum, E. Allman, J. Kershenstein, and D. Alexa, "Hyperspectral change detection in high clutter using elliptically contoured distributions," *Proc. SPIE* **6565**, p. 656515, 2007.
9. E. Gomez-Sanchez-Manzano, M. A. Gomez-Villegas, and J. M. Marin, "Sequences of elliptical distributions and mixtures of normal distributions," *J. Multivariate Analysis* **97**, pp. 295–310, 2006.
10. J. Theiler, "Quantitative comparison of quadratic covariance-based anomalous change detectors," *Applied Optics* **47**, pp. F12–F26, 2008.
11. A. Schaum and A. Stocker, "Spectrally selective target detection," *Proc. International Symposium on Spectral Sensing Research*, 1997.
12. A. Schaum and A. Stocker, "Long-interval chronochrome target detection," *Proc. 1997 International Symposium on Spectral Sensing Research*, 1998.
13. I. S. Reed and X. Yu, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoustics, Speech, and Signal Processing* **38**, pp. 1760–1770, 1990.
14. Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), Jet Propulsion Laboratory (JPL), National Aeronautics and Space Administration (NASA) <http://aviris.jpl.nasa.gov/>.
15. AVIRIS Free Standard Data Products, Jet Propulsion Laboratory (JPL), National Aeronautics and Space Administration (NASA). <http://aviris.jpl.nasa.gov/html/aviris.freedata.html>.
16. M. T. Eismann, J. Meola, and R. Hardie, "Hyperspectral change detection in the presence of diurnal and seasonal variations," *IEEE Trans. Geoscience and Remote Sensing* **46**, pp. 237–249, 2008.
17. S. Cambanis, S. Huang, and G. Simons, "On the theory of elliptically contoured distributions," *J. Multivariate Analysis* **11**, pp. 368–385, 1981.
18. K.-T. Fang and Y.-T. Zhang, *Generalized multivariate analysis*, Springer-Verlag, Science Press, Berlin, Beijing, 1980.
19. I. Steinwart, D. Hush, and C. Scovel, "A classification framework for anomaly detection," *J. Machine Learning Research* **6**, pp. 211–232, 2005.
20. C. Scovel, D. Hush, and I. Steinwart, "Learning rates for density level detection," *Analysis and Applications-Special Issue on Machine Learning* **3**(4), pp. 357–371, 2005.

APPENDIX A. GROUP INVARIANT ANOMALY DETECTION

In this appendix, we provide an alternative interpretation of the anomalous change detector that we have proposed in the text. In this interpretation, we do not see the uncorrelated case as an *approximation* to the *correct* independence-based anomalous change detector, but instead consider it a valid anomaly detector in its own right, one of a family of detectors that is consistent with a symmetry property that is derived in terms of group invariants.

According to Cambanis et. al.¹⁷ a measure μ on \mathcal{R}^d is an EC distribution with covariance parameter Σ (corresponding, up to a multiplicative constant, to the covariance of the measure, when the covariance exists[†]) if the characteristic function

$$\phi(\mathbf{t}) := \int_{\mathcal{R}^d} e^{i\mathbf{t}\cdot\mathbf{z}} d\mu(\mathbf{z}) \quad (36)$$

is a function of $\mathbf{t}^T \Sigma \mathbf{t}$. When Σ is invertible, it is easy to show that the function $\mathbf{t} \mapsto \mathbf{t}^T \Sigma \mathbf{t}$ is a maximal invariant for the representation $g \mapsto \Sigma^{1/2} g \Sigma^{-1/2}$, with g a member of the orthogonal group $O(d)$, in the sense (see, e.g., Ex. 1.7.1 in Ref. [18]) that any function which is invariant under this representation is a function of the maximal invariant. Consequently, EC distributions are simply those which are invariant under a representation of the orthogonal group. Similar statements can be made when the matrix Σ is degenerate but for simplicity we restrict to the non-degenerate case. In the anomaly detection framework of Refs. [19, 20], for a given measure μ one must select a reference measure ν such that μ is absolutely continuous with respect to ν and then the anomalies at level ρ are defined to be the set $\{\mathbf{z} : \frac{d\mu}{d\nu}(\mathbf{z}) \leq \rho\}$ where $\frac{d\mu}{d\nu}$ is the Radon-Nykodym derivative.

When μ is an EC distribution it at first appears reasonable to assume that the sets of anomalies at any level should be invariant under the associated orthogonal symmetry group. Unfortunately, taking the reference measure ν to have the same symmetries as μ leads to the conclusion that different choices of reference measure lead only to different parameterization of the level function: namely, that all symmetric anomaly detectors are reparameterizations of Mahalanobis distance. Or said differently, if we look over the set of level parameters ρ there really is only one symmetric anomaly detector – the Mahalanobis distance.

Instead, let us consider anomalous change detection where μ is a measure on $\mathcal{R}^{d_x} \times \mathcal{R}^{d_y}$. If μ is an EC distribution and therefore symmetric, and we wish to detect anomalous changes, then the above argument implies that we must *not* select the reference measure ν to have the same symmetry as μ , for otherwise we would simply be detecting straight anomalies and not anomalous changes. That is, symmetric anomalous change detection is *symmetry breaking*. However there are still symmetries available. In particular, Thm. 2.6.3 in Ref. [18] implies that the marginal distributions μ_x on \mathcal{R}^{d_x} and μ_y on \mathcal{R}^{d_y} are EC with parameters Σ_x and Σ_y respectively and therefore they are orthogonally symmetric with the representations $\Phi_x(g_x) := \Sigma_x^{1/2} g_x \Sigma_x^{-1/2}$ with $g_x \in O(d_x)$ and $\Phi_y(g_y) := \Sigma_y^{1/2} g_y \Sigma_y^{-1/2}$ with $g_y \in O(d_y)$ respectively. Consequently, it is natural to require the reference measure ν to be symmetric with respect to Φ_x in its \mathcal{R}^{d_x} coordinates and symmetric with respect to Φ_y in its \mathcal{R}^{d_y} coordinates. That is, it should be symmetric with respect to the direct product symmetry $\Phi_x \times \Phi_y$. If we take this as a natural assumption for symmetric anomalous change detection and define the maximal invariants $\xi_z = \mathbf{z}^T \Sigma^{-1} \mathbf{z}$, $\xi_x = \mathbf{x}^T \Sigma_x^{-1} \mathbf{x}$ and $\xi_y = \mathbf{y}^T \Sigma_y^{-1} \mathbf{y}$, it therefore follows that all symmetric anomalous change detector have the following form

$$\{\mathbf{z} = (\mathbf{x}, \mathbf{y}) \mid \frac{h(\xi_z)}{g(\xi_x, \xi_y)} \leq \rho\} \quad (37)$$

for a univariate function h and bivariate function g . A particular case of this general formulation is given by $g(\xi, \eta) := h(\xi + \eta)$ and this leads to a family of anomalous change detectors determined by a single univariate function h :

$$\{\mathbf{z} = (\mathbf{x}, \mathbf{y}) \mid \frac{h(\xi_z)}{h(\xi_x + \xi_y)} \leq \rho\}. \quad (38)$$

This is just the family proposed in Eq. (32).

[†]When the distribution is so fat-tailed that the second moment diverges, then the covariance is not defined. But it is still possible to describe such EC distributions with this more general formulation.