Symmetrized local co-registration optimization for anomalous change detection

Brendt Wohlberg and James Theiler

Los Alamos National Laboratory, Los Alamos, NM 87545

ABSTRACT

The goal of anomalous change detection (ACD) is to identify what unusual changes have occurred in a scene, based on two images of the scene taken at different times and under different conditions. The actual anomalous changes need to be distinguished from the incidental differences that occur throughout the imagery, and one of the most common and confounding of these incidental differences is due to the misregistration of the images, due to limitations of the registration pre-processing applied to the image pair.

We propose a general method to compensate for residual misregistration in any ACD algorithm which constructs an estimate of the degree of "anomalousness" for every pixel in the image pair. The method computes a modified misregistration-insensitive anomalousness by making local re-registration adjustments to minimize the local anomalousness. In this paper we describe a symmetrized version of our initial algorithm, and find significant performance improvements in the anomalous change detection ROC curves for a number of real and synthetic data sets.

Keywords: Anomalous change detection, Registration, Multispectral imagery, Hyperspectral imagery

1. INTRODUCTION

Two images of the same scene, taken at different times and under different conditions, will typically exhibit a wide variety of differences. But those differences are not necessarily very interesting. In fact, what an image analyst often seeks is the small change that is hidden among the clutter of pervasive differences. This is the task of *anomalous change detection* (ACD). By distinguishing those changes that are unusual from the differences that occur pervasively throughout the image, the ACD algorithm can help identify potentially interesting changes in the scene. Whether those changes truly are interesting or meaningful is beyond the scope of the problem statement – as an operational issue, that judgement is generally left for a human analyst to decide. What automated ACD offers is a way to cull through a massive archive of imagery, and narrow down the changes that the analyst might want to examine. (See Ref. [1] for an overview.)

One of the most confounding sources of "uninteresting" differences is misregistration of the images. While it is important to align the images as precisely as possible in the first place, so that corresponding pixels in the two images correspond to the same position in the scene, one has to assume that some residual misregistration will inevitably remain. Since the effects of misregistration are pervasive over the whole scene, ACD already provides some robustness to misregistration, at least in principle. But more active compensation is possible, and in an earlier paper, we proposed a misregistration compensation algorithm.[2] In this paper, we extend those efforts, and demonstrate the ability of this extension to further reduce the rate of false alarms caused by the residual misregistration that is inevitable between pairs of images.

Algorithms that have been proposed for ACD include the chronochrome,[3] neural net prediction,[4] covariance equalization,[5] multivariate alteration detection,[6] and a machine learning framework[7] which as led to a number of variations that optimize for different situations, such as subpixel anomalies[8] or fat-tailed data distributions.[9, 10] In all of these pixel-based algorithms, a scalar "anomalousness" value is assigned to every pixel in the image, and those pixels with the highest anomalousness value are the top candidates for the locations of anomalous change. Our approach for misregistration compensation can be applied to any of these pixel-based ACD algorithms, but we will concentrate on the hyperbolic anomalous change detector.[7]

Authors' emails: {brendt,jt}@lanl.gov

1.1. Hyperbolic anomalous change detection

Let $\mathbf{x} \in \mathbb{R}^{d_x}$ be a pixel value in the first image, $\boldsymbol{\chi}$, and $\mathbf{y} \in \mathbb{R}^{d_y}$ correspond to the associated pixel value in the second image, $\boldsymbol{\gamma}$. In general, we will write $\mathcal{A}(\mathbf{x}, \mathbf{y})$ as as our measure of anomalousness.

The Hyperbolic Anomalous Change Detector (HACD) is motivated by modelling the underlying probability distribution $P(\mathbf{x}, \mathbf{y})$ for values $\mathbf{x}$ and $\mathbf{y}$ associated with corresponding pixels in an image. Write $P_x(\mathbf{x}) = \int P(\mathbf{x}, \mathbf{y})\, \mathrm{d}\mathbf{y}$ as the projection of $P(\mathbf{x}, \mathbf{y})$ onto the $\mathbf{x}$ subspace; this is the distribution of pixel values in $\boldsymbol{\chi}$ alone. One can similarly write $P_y(\mathbf{y}) = \int P(\mathbf{x}, \mathbf{y})\, \mathrm{d}\mathbf{x}$. Following the framework introduced in Ref. [7], we can characterize the anomalous changes as those with high values of mutual information. That is,

$$\mathcal{A}'(\mathbf{x}, \mathbf{y}) = \log P_x(\mathbf{x}) + \log P_y(\mathbf{y}) - \log P(\mathbf{x}, \mathbf{y}). \tag{1}$$

When the data distribution is Gaussian, these probability densities can be described in terms of the covariance and cross-covariance matrices of the data. Subtract the mean from both images, so that $\langle \mathbf{x} \rangle = 0$ and $\langle \mathbf{y} \rangle = 0$; then write

$$X = \left\langle \mathbf{x}\mathbf{x}^T \right\rangle \quad Y = \left\langle \mathbf{y}\mathbf{y}^T \right\rangle \quad C = \left\langle \mathbf{y}\mathbf{x}^T \right\rangle .$$

Up to unimportant additive and multiplicative constants, the anomalousness in Eq. (1), in this Gaussian case, becomes a quadratic expression:

$$\mathcal{A}(\mathbf{x}, \mathbf{y}) = \left[\begin{array}{cc} \mathbf{x}^T & \mathbf{y}^T \end{array} \right] Q \left[\begin{array}{c} \mathbf{x} \\ \mathbf{y} \end{array} \right], \tag{2}$$

where the coefficient matrix Q is given by

$$Q = \left[\begin{array}{cc} X & C^T \\ C & Y \end{array} \right]^{-1} - \left[\begin{array}{cc} X & 0 \\ 0 & Y \end{array} \right]^{-1} . \tag{3}$$

For HACD, the matrix Q has negative as well as positive eigenvalues, and the boundaries of constant $\mathcal{A}(\mathbf{x}, \mathbf{y})$ are hyperbolas in $(\mathbf{x}, \mathbf{y})$ space. Another consequence of these negative eigenvalues is that, in contrast most other pixel-based ACD algorithms, the anomalousness measure $\mathcal{A}(\mathbf{x}, \mathbf{y})$ for HACD can be positive *or* negative. The largest (*i.e.,* most positive) anomalousness values correspond to the most anomalous changes.

2. MINIMUM ANOMALOUSNESS REGISTRATION

Since they are based on the statistics of corresponding pixels, the class of ACD algorithms described above depends critically on accurate image registration, which cannot, in practice, be expected to be perfect, given the limitations of registation algorithms. The following simple scenario provides a motivation for our algorithm for reducing the sensitivity of ACD algorithms to registation errors.

In an image pair, consider a pixel $\mathbf{x}$, in the first image, and a small window, containing pixels $\mathbf{y}_m$, about the corresponding pixel, $\mathbf{y}_0$, in the second image. If $\mathbf{x}$ is a true anomalous change consisting of an object not present in the corresponding position in the second image, then all joint vectors $[\mathbf{x}^T \mathbf{y}_m^T]^T$ are likely to have a large anomalousness measure (see Fig. 1). Conversely, if $\mathbf{x}$ does not represent a true anomalous change, but $[\mathbf{x}^T \mathbf{y}_0^T]^T$ has a large anomalousness measure due to misregistration, we can expect that some joint vectors $[\mathbf{x}^T \mathbf{y}_m^T]^T$ will have a low anomalousness measure if the window is large enough to encompass the misregistration (see Fig. 2).

2.1. Asymmetric Algorithm

Motivated by this argument, we previously proposed[2] the following misregistration compensation scheme: For each pixel in $\boldsymbol{\chi}$, consider a window about the corresponding pixel in $\boldsymbol{\gamma}$, and find the pixel within this window that gives the lowest anomalousness when paired with the pixel in $\boldsymbol{\chi}$ (see Fig. 3). Take that pixel as the misregistration compensated pixel value.

We define the window by offset vectors $\mathbf{w}_h$ and $\mathbf{w}_v$, the simplest example being a 3×3 window about the central pixel

$$\left[\begin{array}{c} \mathbf{w}_h \\ \mathbf{w}_v \end{array} \right] = \left[\begin{array}{ccccccccc} 0 & 0 & 0 & -1 & -1 & -1 & 1 & 1 & 1 \\ 0 & -1 & 1 & -1 & 0 & 1 & -1 & 0 & 1 \end{array} \right] ,$$

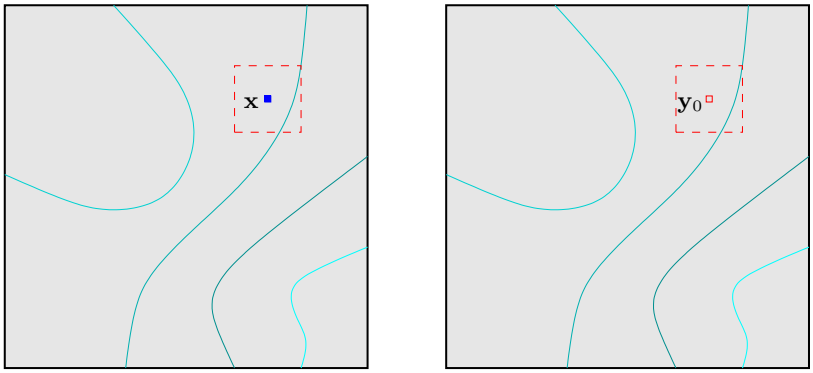

Figure 1. Real anomalies: pixel $\mathbf{x}$ in image $\boldsymbol{\chi}$ has content not present in the window about the registered pixel, $\mathbf{y}_0$ in image $\boldsymbol{\gamma}$, and all pairs $[\mathbf{x}^T \mathbf{y}_m^T]^T$ have high anomalousness.

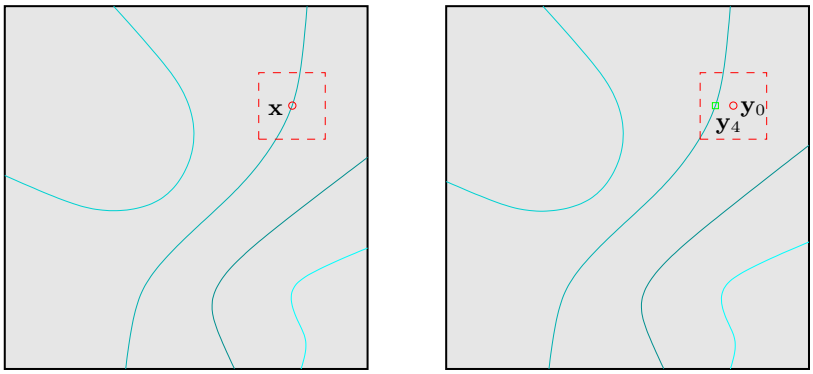

Figure 2. Misregistration anomalies: pixel $\mathbf{x}$ in image $\boldsymbol{\chi}$ has different content from registered pixel, $\mathbf{y}_0$ in image $\boldsymbol{\gamma}$, but pixel $\mathbf{y}_4$ in the window about $\mathbf{y}_0$ has similar content to $\mathbf{x}$, so that the pair $[\mathbf{x}^T \mathbf{y}_4^T]^T$ has low anomalousness.

Algorithm 1 Asymmetric Local Co-Registration Adjustment algorithm

Compute Q for the image pair $\boldsymbol{\chi}$ and $\boldsymbol{\gamma}$
for all pixel indices k, l **do**
  for all window vector indices m **do**
    Set $k' = k + \mathbf{w}_{h,m}$ and $l' = l + \mathbf{w}_{v,m}$
    Set $\mathcal{A}_{k,l,m} = \begin{bmatrix} \boldsymbol{\chi}_{k,l}^T & \boldsymbol{\gamma}_{k',l'}^T \end{bmatrix} Q \begin{bmatrix} \boldsymbol{\chi}_{k,l} \\ \boldsymbol{\gamma}_{k',l'} \end{bmatrix}$
  end for
  Set $\mathcal{A}_{k,l} = \min_m \mathcal{A}_{k,l,m}$
end for

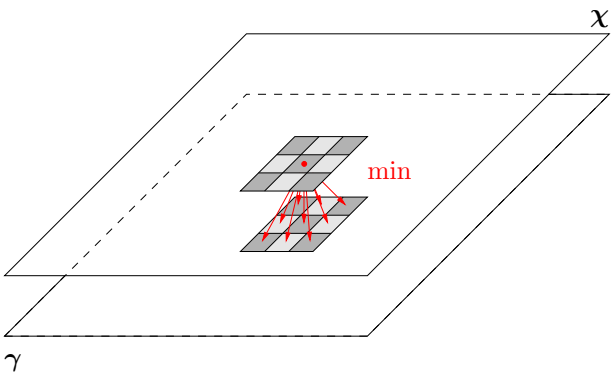

Figure 3. Construction of a set of joint vectors from a single pixel in image $\boldsymbol{\chi}$ and all pixels within a window about the corresponding pixel in image $\boldsymbol{\gamma}$.

which we shall refer to as an $r = 1$ (i.e. radius 1) window, with $r = 2$ denoting a 4×4 window, and so forth. The local co-registration adjustment (LCRA) procedure is defined in Algorithm 1.

An equivalent but more efficient implementation is to apply a shift to $\boldsymbol{\gamma}$ for each relative position in the chosen window and then to compute an anomalousness map for this image pair (see Fig. 4), as described in Algorithm 2.

Algorithm 2 Asymmetric Local Co-Registration Adjustment algorithm (more efficient implementation)

Compute Q for the image pair $\boldsymbol{\chi}$ and $\boldsymbol{\gamma}$
for all window vector indices m **do**
  Construct $\boldsymbol{\gamma}_m$ by applying shift $(\mathbf{w}_{h,m}, \mathbf{w}_{v,m})$ to $\boldsymbol{\gamma}$
  for all pixel indices k, l **do**
    Set $\mathcal{A}_{m,k,l} = \begin{bmatrix} \boldsymbol{\chi}_{k,l}^T & \boldsymbol{\gamma}_{m,k,l}^T \end{bmatrix} Q \begin{bmatrix} \boldsymbol{\chi}_{k,l} \\ \boldsymbol{\gamma}_{m,k,l} \end{bmatrix}$
  end for
end for
for all pixel indices k, l **do**
  Set $\mathcal{A}_{k,l} = \min_m \mathcal{A}_{m,k,l}$
end for

The offset $\mathbf{w}_m$ that minimizes $\mathcal{A}_{k,l,m}$ is naturally interpreted as the misregistration at the point k, l in the image. We do not, however, treat it as an accurate estimator of misregistration *per se*; instead we interpret it more loosely as a way to compensate for the misregistration. Nevertheless, given the mutual information interpretation of the HACD anomalousness measure, as in (1), it is interesting to note the strong similarity between the approach proposed here for reducing the sensitivity of ACD algorithms to registration errors, and the well known mutual information based registration algorithms.[11–13]

In the above algorithms covariance Q is computed once and then applied for every shifted image. A plausible alternative is to recompute Q for every $\boldsymbol{\chi}$ and shifted $\boldsymbol{\gamma}$ image pair, but this turns out to be a bad idea: (i) it effectively (and incorrectly) assumes that the same fixed shift is applied over the entire image, (ii) it results in differences of normalization between the resulting anomalousness maps, and (iii) it works poorly in practice.

2.2. Symmetric Algorithm

The motivation for the LCRA algorithm is based on a scenario in which the anomalous change occurs in the first image, $\boldsymbol{\chi}$, of the pair, and this asymmetry carries through to the resulting algorithm. In particular, when the anomalous change is a pixel in the second image, $\boldsymbol{\gamma}$, it may elude detection. This is because it is always part

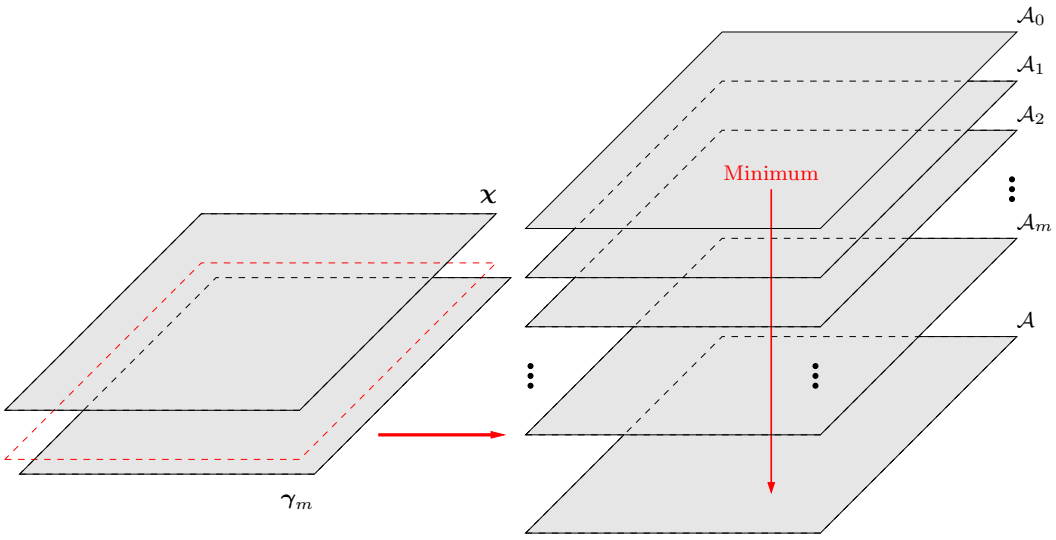

Figure 4. Minimization over the stack of anomalousness maps for each pair of $\boldsymbol{\chi}$ and shifted $\boldsymbol{\gamma}$ images.

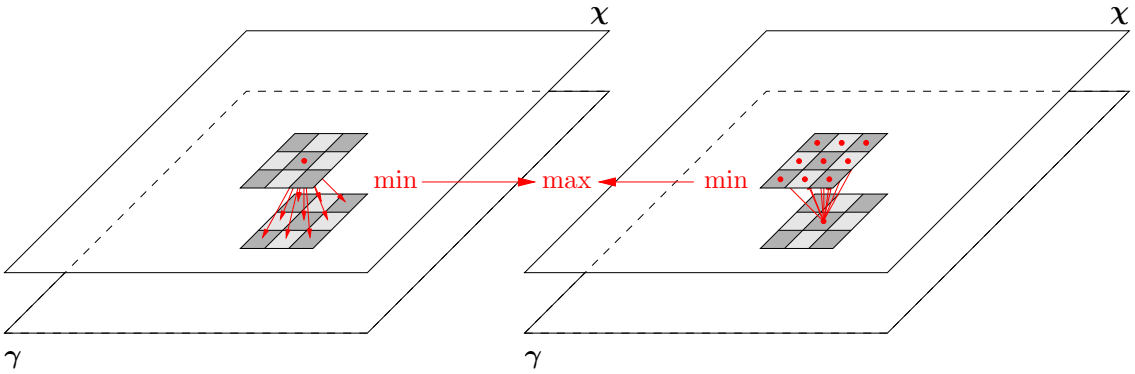

Figure 5. Symmetric version of the LCRA algorithm illustrated in Fig. 3. At each pixel, the minimum anomalousness pair is selected in each direction (i.e. image $\boldsymbol{\chi}$ to $\boldsymbol{\gamma}$, and image $\boldsymbol{\gamma}$ to $\boldsymbol{\chi}$) and the final anomalousness value is the maximum of these two minima.

of a window of pixels, and its neighbors will be chosen as the registration-adjusted matches for the pixels in $\boldsymbol{\chi}$. Note, however, that while a true anomalous change will give either a high or low anomalousness depending on the direction of minimization (i.e. $\boldsymbol{\chi}$ to $\boldsymbol{\gamma}$, or vice versa), a spurious anomalous change due to misregistration is symmetric, in that similar minimum anomalousness should be obtained independent of the direction of minimization. This observation suggests the Symmetric LCRA algorithm, in which minimization is performed in both directions, and the maximum of the resulting minima is selected as the final result, as illustrated in Fig. 5.

The symmetric algorithm, like the original asymmetric version, is more efficiently implemented in terms of shifts applied to whole images, as in Algorithm 3.

2.3. Simulation Framework

Because anomalies are by definition rare, evaluating the utility of anomaly detection algorithms can be problematic; anecdotal evidence has some value, but quantitative comparisons require an adequate supply of anomalies. In the simulation framework proposed in Ref. [14], one can start with two images which are presumed to contain pervasive differences but no anomalous changes.* We call these the *base* image and the *normal* change image

*Of course it is possible to generate these pervasive differences by simulation as well; where the anomalous change is generated for a single pixel at a time, the pervasive difference is applied to the whole image at once.

Algorithm 3 Symmetric Local Co-Registration Adjustment algorithm (more efficient implementation)

Compute Q for the image pair $\boldsymbol{\chi}$ and $\boldsymbol{\gamma}$
for all window vector indices m **do**
  Construct $\boldsymbol{\chi}_m$ by applying shift $(\mathbf{w}_{h,m}, \mathbf{w}_{v,m})$ to $\boldsymbol{\chi}$
  Construct $\boldsymbol{\gamma}_m$ by applying shift $(\mathbf{w}_{h,m}, \mathbf{w}_{v,m})$ to $\boldsymbol{\gamma}$
  for all pixel indices k, l **do**
    Set $\mathcal{A}_{0,m,k,l} = \begin{bmatrix} \boldsymbol{\chi}_{m,k,l}^T & \boldsymbol{\gamma}_{k,l}^T \end{bmatrix} Q \begin{bmatrix} \boldsymbol{\chi}_{m,k,l} \\ \boldsymbol{\gamma}_{k,l} \end{bmatrix}$
    Set $\mathcal{A}_{1,m,k,l} = \begin{bmatrix} \boldsymbol{\chi}_{k,l}^T & \boldsymbol{\gamma}_{m,k,l}^T \end{bmatrix} Q \begin{bmatrix} \boldsymbol{\chi}_{k,l} \\ \boldsymbol{\gamma}_{m,k,l} \end{bmatrix}$
  end for
end for
for all pixel indices k, l **do**
  Set $\mathcal{A}_{k,l} = \max\{\min_m \mathcal{A}_{0,m,k,l} \ , \ \min_m \mathcal{A}_{1,m,k,l}\}$
end for

(and they correspond to the $\boldsymbol{\gamma}$ and $\boldsymbol{\chi}$ images, respectively). A third image, the *anomalous* change image, is simulated by choosing a pixel in the *normal* change image, and replacing it with a randomly chosen pixel from somewhere in the rest of the normal change image. The idea is that the pixel is not itself unusual, but in the context of the corresponding pixel in the *base* image, it exhibits an unusual change. When purely spectral ACD algorithms are employed, one can take a shortcut and produce an *anomalous* change image in which every pixel constitutes an anomalous change – in this case, the *anomalous* change image is obtained from the *normal* change image simply by scrambling its pixels. The anomalous change detection algorithm is "trained" (which is to say that the quadratic coefficient matrix Q is computed) using the *base-normal* pair. Applying the algorithm, at a given threshold, to the *base-normal* pair provides an estimate of the false alarm rate. Applying the same algorithm at the same threshold to the *base-anomalous* pair, one can estimate the detection rate. By varying the threshold, a receiver operator characteristic (ROC) can be generated: this provides detection rate as a function of false alarm rate.

2.3.1. Evaluating LCRA algorithm with simulation framework

When spatial pre-processing is built into the ACD algorithm, the simulation framework requires additional complexity.[15] Because the LCRA algorithm incorporates spatial context, this additional complexity is necessary here. It is possible to avoid this complexity for the specific case where the algorithm is asymmetrical, *and* the anomalies are known to be on a particular image; that is the case we investigated in our earlier work.[2] But that is a special case, and for the results presented here, we employed a more sophisticated (and more robust) simulation framework.

This framework incorporates the same ideas used in other work on spatial processing for anomalous change detection.[15] We begin by introducing another image, the *target* mask, which is a binary image with spatially isolated 1's surrounded by 0's. Anomalous changes are introduced only at locations in the image where the *target* mask is 1. That is, the *anomalous* change image differs from the *normal* image only at those locations where the *target* mask is 1. Those anomalous pixels are chosen at random from the rest of the image.

As before, we apply the algorithm at a given threshold to the *base-normal* pair to provide an estimate of the false alarm rate.[†] We apply the same algorithm at the same threshold to the *base-anomalous* pair, but only consider the pixels where the *target* mask is nonzero. This provides our estimate the detection rate.

Because the anomalous changes are spatially isolated, we can perform spatial processing without having the individual anomalies interfere with each other. It is of course required that the distance between the anomalies be larger than the diameter of the spatial processing window.

[†]We avoid pixels near the edge of the image when we make this estimate, since the shifting introduces artifacts there.

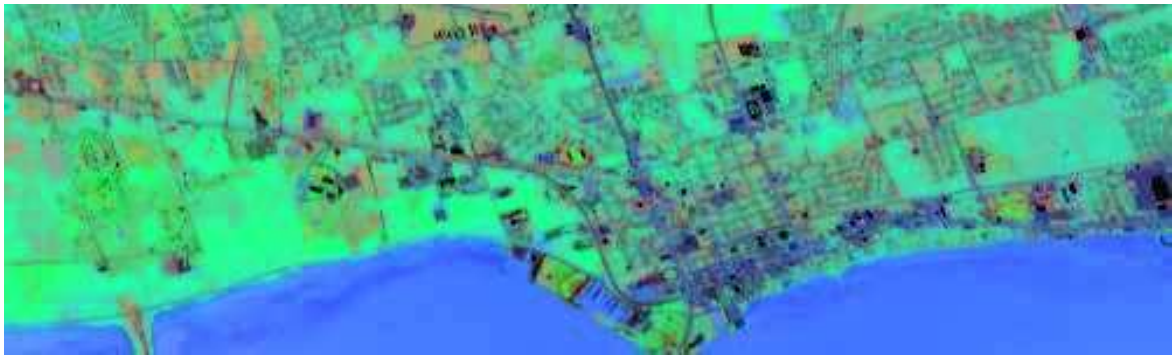

Figure 6. False color AVIRIS[16] hyperspectral image of the Florida coastline, from dataset f960323t01p02_r04_sc01.

For the asymmmetic LCRA algorithm, there are two cases. One in which the $\boldsymbol{\gamma}$ image is shifted (as shown in Fig. 3), and one in which the $\boldsymbol{\chi}$ image is shifted. If we know that the anomalies, if they are to appear, will appear in the $\boldsymbol{\chi}$ image, the the approach illustrated in Fig. 3, and described in the asymmetric LCRA pseudocode (see Section 2.1), is appropriate. But if a single-pixel anomaly were to appear in the $\boldsymbol{\gamma}$ image, then the effect of the `min` operator in Fig. 3 would be to effectively ignore the anomaly.

In our earlier work, we assumed that the changes were introduced in the $\boldsymbol{\chi}$ image and the algorithm was evaluated with simulation framework that made the same assumption. But in practice, this means that the user has to guess correctly which image harbors the anomaly. If the user guesses incorrectly, the asymmetric LCRA can fail catastrophically. In the results presented here, we use the above simulation framework, but employ two different asymmetric LCRA algorithms: one that guesses correctly which image has the anomalies (fwd), and one that guesses incorrectly (rev). For the symmetric LCRA, the forward and reverse algorithms are identical.

3. RESULTS

We present ACD results for both simulated and real anomalous changes. In each case we evaluate the performance of the LCRA algorithm in correcting (i) a simple shift of the entire image in the horizontal direction, and (ii) a misregistration consisting of a random (non-integer) offset at each pixel, in both the horizontal and vertical directions, with the offset vector field smoothed to reduce abrupt changes between adjacent pixels. All LCRA and Symmetric LCRA (SLCRA) results are computed using HACD as the base ACD algorithm.

3.1. Hyperspectral data with simulated anomalous changes

A false color rendition of the AVIRIS (Airborne Visible/InfraRed Imaging Spectrometer[16]) hyperspectral test data[‡] is displayed in Fig. 6. For these results, we distinguish between LCRA in the forward direction ("LCRA (fwd)"), using the correct choice of image in which changes occur, and LCRA in the reverse direction ("LCRA (rev)"), with the incorrect choice of image in which changes occur. Fig. 7 shows detection results for this data with pervasive differences consisting of a uniform horizontal offset. In Fig. 7(a), "LCRA (fwd)" gives the best performance, followed closely by Symmetric LCRA, and "LCRA (rev)" gives worse performance than unmodified HACD. In Fig. 7(b), there is a large performance gap between SLCRA with $r = 1$, and SLCRA with $r = 2$ and $r = 3$; this is consistent with the two-pixel misregistration there. Fig. 8 shows detection results for this data with pervasive differences consisting of a smooth random misregistration. In Fig. 8(a), for which the maximum misregistration radius was 1 pixel, LCRA in the reverse direction substantially reduces performance, while the Symmetric LCRA has very similar performance to the LCRA in the forward direction. Good performance from asymmetric LCRA requires that one know which image the anomalous change will appear; the symmetric LCRA does not require this knowledge. In Fig. 8(b) there is a small performance difference for the three different window radii, but the best performance is obtained with $r = 2$, corresponding to the magnitude of the actual random misregistration applied.

[‡]AVIRIS data is available from the Jet Propulsion Laboratory (JPL) and National Aeronautics and Space Administration (NASA) website: `http://aviris.jpl.nasa.gov/html/aviris.freedata.html`

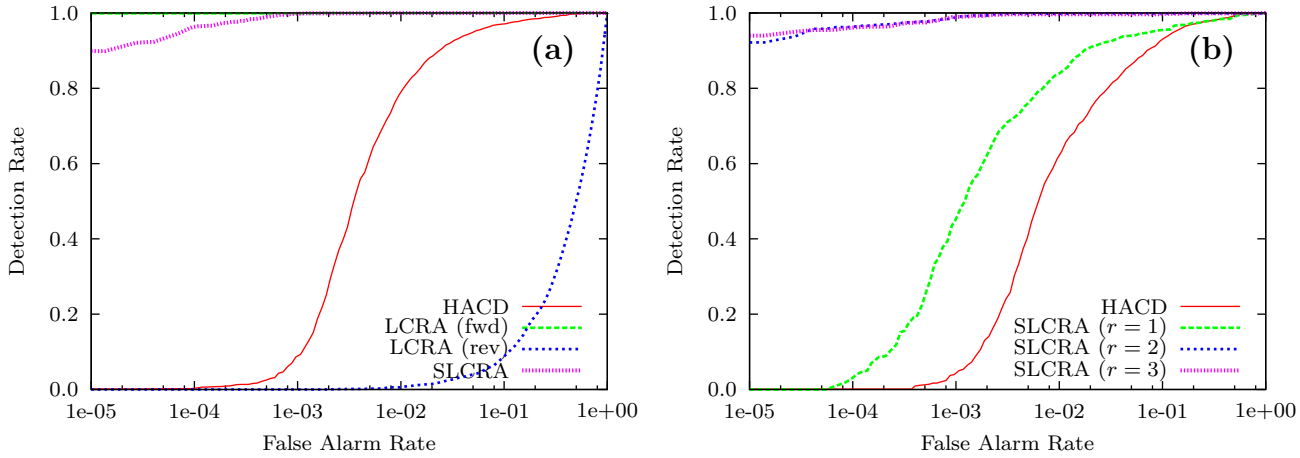

Figure 7. Results using AVIRIS data with simulated anomalous changes and uniform horizontal misregistration: (a) one pixel misregistration, comparing LCRA in both directions with SLCRA (all with $r = 1$), and (b) two pixel misregistration, comparing SLCRA for $r \in \{1, 2, 3\}$. In (a), the "LCRA (fwd)" curve is indistinguisable from the upper horizontal axis.

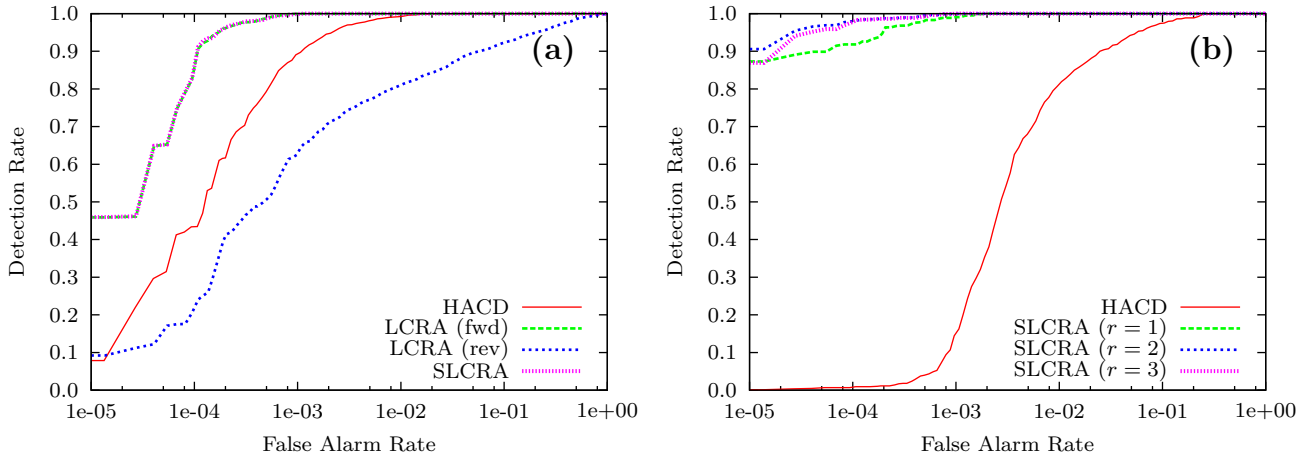

Figure 8. Results using AVIRIS data with simulated anomalous changes and random misregistration: (a) radius one misregistration, comparing LCRA in both directions with SLCRA (all with $r = 1$), and (b) radius two misregistration, comparing SLCRA for $r \in \{1, 2, 3\}$. In (a) the curves for "LCRA (fwd)" and "SLCRA" are indistinguishable.

3.2. Multispectral data with real anomalous changes

A pair of images of desktop clutter is used to compute performance results for real data. These images are displayed in Fig. 9, with a circle surrounding the anomalous change, consisting of a sunflower seed which is rotated in the second image. The pervasive changes consist of different lighting conditions, and a misregistration. In Figs. 10 and 11 the "HACD (aligned)" curve provides reference performance for HACD with no misregistration. For these results, since there is no clear notion of computing LCRA in the correct or incorrect direction, we denote the two different directions by "LCRA($\boldsymbol{\chi}, \boldsymbol{\gamma}$)" and "LCRA($\boldsymbol{\gamma}, \boldsymbol{\chi}$)". With the exception of Fig. 11(a), for which the detection performance of all of the methods is similar, there are significant differences in performance between "LCRA($\boldsymbol{\chi}, \boldsymbol{\gamma}$)" and "LCRA($\boldsymbol{\gamma}, \boldsymbol{\chi}$)", and SLCRA exhibits superior or equal performance to the better of the two.

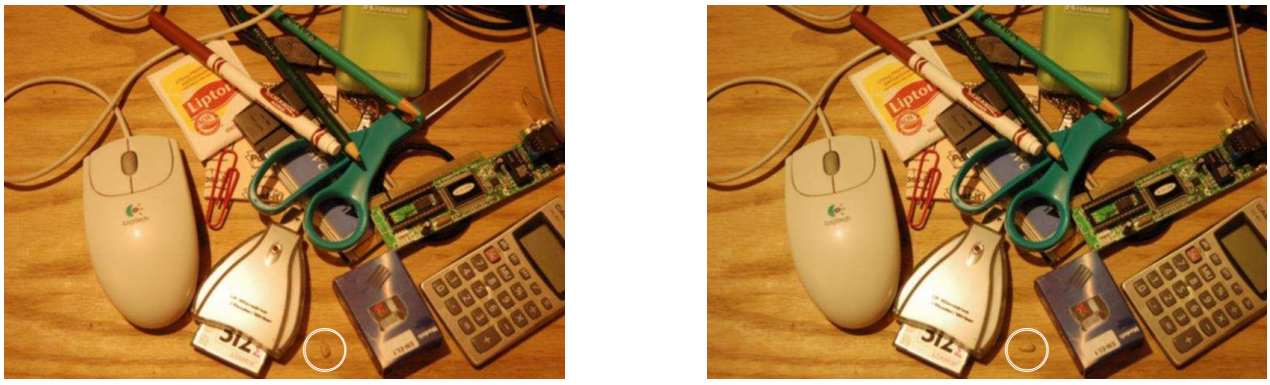

Figure 9. Desktop clutter test image pair. Location of anomalous change indicated by white circle.

4. CONCLUSIONS

The LCRA algorithm is observed to significantly improve detection performance in the low false-alarm regime (sometimes at the expense of performance in the high false-alarm regime) for misregistered data. These per-

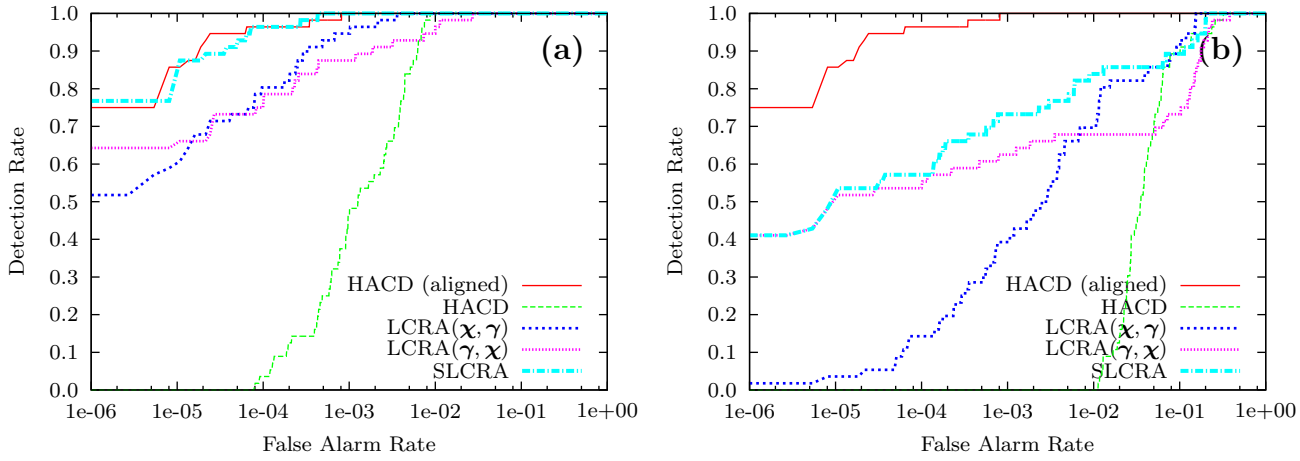

Figure 10. Desktop clutter detection results for uniform horizontal misregistration of (a) 1 pixel, and (b) 2 pixels. The "HACD (aligned)" curve corresponds to no misregistration, and provides an upper bound on misregistration compensation performance. LCRA and SLCRA curves were computed using a window size appropriate for the misregistration magnitude (i.e. $r = 1$ for a 1 pixel offset and $r = 2$ for a 2 pixel offset).

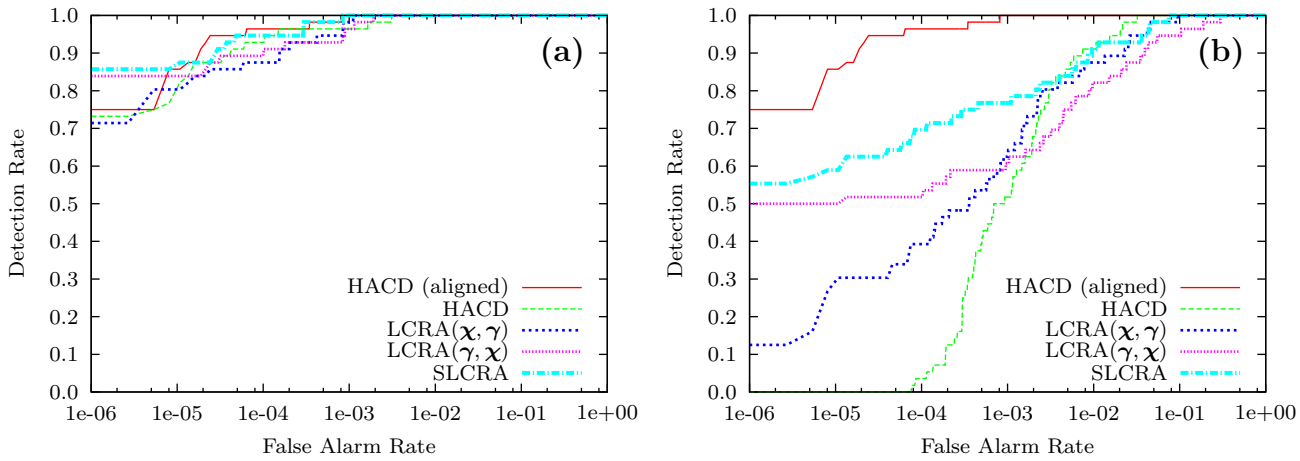

Figure 11. Desktop clutter detection results for random misregistration of (a) radius 1, and (b) radius 2. The "HACD (aligned)" curve corresponds to no misregistration, and provides an upper bound on misregistration compensation performance. LCRA and SLCRA curves were computed using a window size appropriate for the misregistration magnitude (i.e. $r = 1$ for a radius 1 random offset and $r = 2$ for a radius 2 random offset).

formance improvements have been observed for both the simulation framework and real data. The Symmetric LCRA algorithm gives additional performance improvements, over a wider range of false-alarm rates.

Future research will address a number of promising extensions to this approach, including: (i) iteratively applying the estimated registration adjustment, re-computing covariances, and re-computing the local co-registration adjustment; (ii) generalizing the registration window to allow explicit adjustment for sub-pixel misregistrations; and (iii) incorporating appropriate prior knowledge, such as smoothness, on the form of misregistration.

REFERENCES

1. M. T. Eismann, J. Meola, A. D. Stocker, S. G. Beaven, and A. P. Schaum, "Airborne hyperspectral detection of small changes," *Applied Optics* **47**, pp. F27–F45, 2008.
2. B. Wohlberg and J. Theiler, "Improved change detection with local co-registration adjustments," in *Proceedings of IEEE Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing*, (Grenoble, France), Aug. 2009.
3. A. Schaum and A. Stocker, "Long-interval chronochrome target detection," in *Proc. Intl. Symposium on Spectral Sensing Research*, 1998.
4. C. Clifton, "Change detection in overhead imagery using neural networks," *Applied Intelligence* **18**, pp. 215–234, 2003.
5. A. Schaum and A. Stocker, "Linear chromodynamics models for hyperspectral target detection," *Proc. IEEE Aerospace Conference* , pp. 1879–1885, 2003.
6. A. A. Nielsen, K. Conradsen, and J. J. Simpson, "Multivariate alteration detection (MAD) and MAF postprocessing in multispectral bi-temporal image data: new approaches to change detection studies," *Remote Sensing of the Environment* **64**, pp. 1–19, 1998.
7. 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.
8. J. Theiler, "Subpixel anomalous change detection in remote sensing imagery," *Proc. IEEE Southwest Symposium on Image Analysis and Interpretation* , pp. 165–168, 2008.
9. J. Theiler and C. Scovel, "Uncorrelated versus independent elliptically-contoured distributions for anomalous change detection in hyperspectral imagery," *Proc. SPIE* **7246**, p. 72460T, 2009.

10. J. Theiler, C. Scovel, B. Wohlberg, and B. R. Foy, "Elliptically contoured distributions for anomalous change detection in hyperspectral imagery," *IEEE Geoscience and Remote Sensing Letters* . In press. doi:10.1109/LGRS.2009.2032565.
11. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multi-modality image registration by maximization of mutual information," in *Mathematical Methods in Biomedical Image Analysis, 1996., Proceedings of the Workshop on*, pp. 14–22, June 1996.
12. W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Medical Image Analysis* **1**(1), pp. 35–51, 1996.
13. H.-M. Chen, P. K. Varshney, and M. K. Arora, "Performance of mutual information similarity measure for registration of multitemporal remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing* **41**, pp. 2445–2454, November 2003.
14. J. Theiler, "Quantitative comparison of quadratic covariance-based anomalous change detectors," *Applied Optics* **47**, pp. F12–F26, 2008.
15. J. Theiler, N. R. Harvey, R. Porter, and B. Wohlberg, "Simulation framework for spatio-spectral anomalous change detection," *Proc. SPIE* **7334**, 2009.
16. G. Vane, R. O. Green, T. G. Chrien, H. T. Enmark, E. G. Hansen, and W. M. Porter, "The Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sensing of the Environment* **44**, pp. 127–143, 1993.