LA-UR- 09-01655

Approved for public release; distribution is unlimited.

Title: A Structural Framework for Anomalous Change Detection and Characterization

Author(s): Lakshman Prasad

James P. Theiler

Intended for: SPIE 2009 Conference Proceedings

Los Alamos
NATIONAL LABORATORY
EST. 1943

A Structural Framework for Anomalous Change Detection and Characterization

Lakshman Prasad*, James Theiler

Space and Remote Sensing Group, Los Alamos National Laboratory, Los Alamos, NM 87545

ABSTRACT

We present a spatially adaptive scheme for automatically searching a pair of images of a scene for unusual and interesting changes. Our motivation is to bring into play structural aspects of image features alongside the spectral attributes used for anomalous change detection (ACD). We leverage a small but informative subset of pixels, namely edge pixels of the images, as anchor points of a Delaunay triangulation to jointly decompose the images into a set of triangular regions, called trixels, which are spectrally uniform. Such decomposition helps in image regularization by simple-function approximation on a feature-adaptive grid. Applying ACD to this trixel grid instead of pixels offers several advantages. It allows: 1) edge-preserving smoothing of images, 2) speed-up of spatial computations by significantly reducing the representation of the images, and 3) the easy recovery of structure of the detected anomalous changes by associating anomalous trixels with polygonal image features. The latter facility further enables the application of shape-theoretic criteria and algorithms to characterize the changes and recognize them as interesting or not. This incorporation of spatial information has the potential to filter out some spurious changes, such as due to parallax, shadows, and misregistration, by identifying and filtering out those that are structurally similar and spatially pervasive. Our framework supports the joint spatial and spectral analysis of images, potentially enabling the design of more robust ACD algorithms.

Keywords: Segmentation, Delaunay triangulation, anomalous, change, feature

1. INTRODUCTION

The problem of Anomalous Change Detection (ACD) in images pertains to the detection of unusual differences between a pair of images of a particular scene, each taken at a different time. Here, one is not interested in pervasive changes such as due to lighting, seasonal variation, camera imaging characteristics, or non-local image editing. Rather, one is looking for changes that are of a rare kind in comparison to the statistics of the collective changes across the two images. Figure 1 illustrates this with two images of the same scene (Figs. 1(a) & 1(b).) The image on the left is blurred and contrast-enhanced compared to the one on the right. Thus it is different from the image on the right in every location. In addition, the image on the right shows an additional 'vehicle' on the road segment shown in inset 1(d) as a real anomalous change.

It is important to emphasize that ACD is different from the problem of anomaly detection wherein one is looking for objects or features that are somehow unusual with respect to the features in the rest of the image. In particular, the occurrence of the same unusual feature in two different images of a scene is not considered unusual in ACD as its peculiarity is of a similar nature in both images.

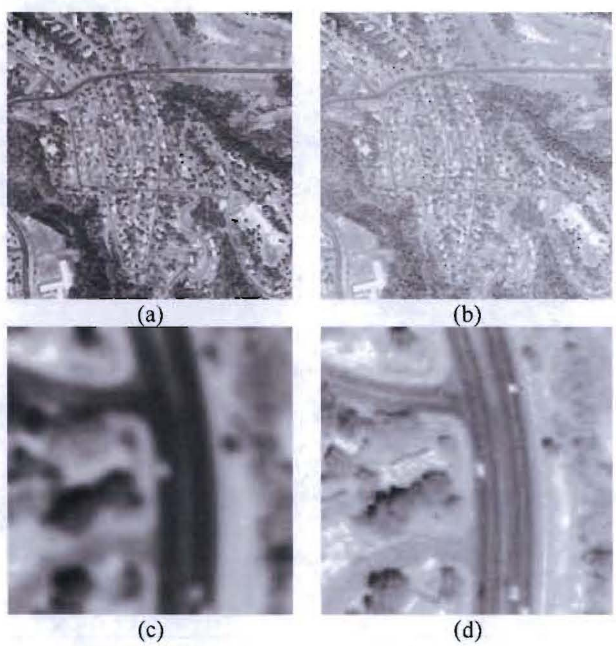

(a) (b) (c) (d)

Figure 1: Pervasive versus anomalous changes

* prasad @lanl.gov, Phone (505) 667-4587

2. FRAMEWORK FOR ANOMALOUS CHANGE DETECTION

Theiler and Perkins[1,2] have proposed a spectral framework for ACD in a pair of images. In this framework they propose an anomalous change detector in terms of level curves of the likelihood ratio

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

Here $P(x,y)$ is the probability distribution of (x,y) with x being the vector of spectral values of a pixel in one image and y being that of the corresponding pixel in the other image. $P(x)$ and $P(y)$ are the corresponding marginal distributions given by

$$P(x) = \int P(x,y)dy, \tag{2}$$

and

$$P(y) = \int P(x,y)dx, \tag{3}$$

In general, the underlying distribution P is not known but a sampling of it is available at the corresponding pixel pairs $\{(x_i,y_i)\}_{i\ 1..N}$ (the normal class) taken from the two images. A resampling approach[2] is employed to estimate the background class, wherein the x and y are treated as if they are independent and the samples $\{(x_i,y_j)\}_{i,j\ 1..N}$ correspond to samples from the distribution $P(x)P(y)$. A binary classification scheme to separate the two classes is then employed to train a discriminant $f(x,y)$ where $f(x,y) >> 0$ represents the anomalous changes. We refer the reader to Theiler and Perkins[1,2] for a detailed discussion of this framework, which we will refer to in the rest of this paper as the framework for anomalous change detection (FACD).

Although the motivating context for FACD is imagery, the generality of the framework extends to non-image data as well. In particular, even in the case of image data, it is not necessary that the elements of picture description be pixels. This suggests that one can perhaps incorporate spatial information in images, along with spectral information, to improve anomalous change detection. For instance, smoothing two images before ACD makes them less variable locally, making real anomalous differences between them stand out. However, smoothing can also affect or even obliterate the very changes sought[3]. Hence, avoiding smoothing across feature boundaries is prudent. Further, apart from knowing there is an unusual change, one would like to know, a) what changed? and b) is it significant, meaningful, or interesting? Consider, for example, two images of an outdoor scene captured an hour apart. Some typical changes that might occur between two such images are: the change in the length of shadows cast by objects in the scene, change in the brightness of certain objects due to change in incident light and reflection, and parallax shift due to a slight change in the camera's position or viewing angle. None of these changes are remarkable from an ACD perspective, but nevertheless may qualify as unusual changes in the statistical landscape of the differences between the images. Worse still, they may dominate in the hierarchy of anomalousness over 'true' anomalous changes such as the appearance, disappearance, or displacement of an object. If one could tease out such pedestrian changes and suppress them, then an ACD algorithm would fare better at detecting true anomalous changes. Feature-adaptive smoothing to mitigate image noise and detection of image-light projection artifacts (such as parallax, shadows, and misregistration,) both require structural knowledge of the images. Structural information in an image is captured to a large extent by the contours that delineate its mostly uniform regions. In this paper we propose the use of feature boundaries in a pair of candidate images for FACD to recast them in terms of spatially *and* spectrally adaptive regional primitives. FACD, originally applied to pixels, will be applied to the spectral attributes of these primitives. Subsequently, the anomalous differences between these elements of the images will be further qualified based on their structural attributes to decide whether they are to be considered as true anomalous changes or, say, merely spatio-temporal artifacts of imaging. We propose the use of an image segmentation scheme that is particularly suited for this task. It first serves as a nonlinear image smoothing mechanism that preserves feature boundaries. Next, it delineates feature elements that can be evaluated based on their structures to determine the nature of the change and characterize it.

3. POLYGONAL IMAGE SEGMENTATION

Image segmentation is a key step in automating image understanding by computers. Segmentation decomposes an image into its constituent salient features defining the semantic content of the image. Image segmentation sets the stage for object detection and recognition by providing a high-level representation of an image in terms of regions with spectral and spatial cohesiveness. Image segmentation may be thought of as analogous to perception in human vision.

Most of the many approaches to segmentation can be broadly classified into one of two categories: 1) methods that seek structure by decomposing an image into spectrally uniform regions, and 2) methods that seek structure by identifying parts of an image that exhibit rapid spectral change, generally assuming that boundaries of objects are sites of such rapid change in intensity. In contrast, our method of segmentation is based on the premise that both region and boundary information are necessary for obtaining meaningful segmentations of images. So, rather than rely on regional uniformity or edge continuity alone for segmentation, we adopt a hybrid region-contour approach that performs segmentation using both aspects of images at the same time.

We now briefly describe the key steps in our segmentation method. We refer the reader to our earlier work[4,5] for a more detailed discussion. Edge pixels of a digital (i.e., pixel-based) image (Fig. 2(a).) are detected by means of an edge detection algorithm such as the well-known Canny edge detector[6] (Fig. 2(b).). Neighboring edge pixels are then linked to obtain contour chains. The space between the contour chains is decomposed into regional units that adapt to the local contour geometry. To achieve this, a constrained Delaunay triangulation (CDT)[7] of contour chains is performed to obtain a complete tessellation of the intercontour regions without any triangle edge crossing a contour chain (Fig. 2(c).)

(a) Input image (b) Canny edge detection (c) Delaunay triangulation

(d) Sample triangles to get trixels (e) Trixel grouping graph (f) Polygonal segmentation

Figure 2. Key stages of edge-based polygonal image segmentation

The reason for using a CDT as against some other kind of triangulation is that each edge of a CDT connects proximate points on the contour chains. Each edge is also transverse to the local axis of the region determined by the two contour elements connected by the edge. Thus the triangles serve as regional elements that capture the local geometry of the image region they represent, with their edges serving to relate contour fragments that bound a common region. The triangles' edges also help interpolate missing edge information to attain completions of object contours. Each triangle in the CDT is sampled for the aggregate spectral value of pixels in its interior (Fig. 2(d)).

At this point, each triangle is a unit comprising the image and is called a *trixel* (TRIangular eXcision ELement). A trixel, like a pixel, is regional unit of uniform color. However, it differs from a pixel in that it is not of fixed size. Its corners lie on image boundaries (i.e, contour chains) and its edges are potential completions of image boundaries in the segmentation process. Thus, the trixel serves as an image adaptive pixel that contains boundary information. As mentioned earlier, the task of image segmentation corresponds to the process of perception in human vision. Accordingly, we model our segmentation method after this process of perception by taking advantage of the optimality properties of CDTs to implement 'perceptual' filters that model certain well-known criteria for perceptual organization employed in human vision. Some examples of such criteria are proximity, transversality (a.k.a amodal completion), closure, and good continuation. Each criterion is implemented as a Boolean filter on the set of all triangle edges. We note that it is possible to restrict the criteria to purely structural and regional aspects of the image's edges, without appealing to the spectral values of the trixels[4]. The criteria jointly determine (via a Boolean expression involving the individual filters) whether the edge between two adjacent triangles should be deleted or retained, or, what is the same, whether two adjacent triangles should be merged into a single polygon or not. Thus, whenever an edge is deleted, a region continuation is achieved. Similarly, whenever an edge is retained a contour continuation is achieved. Thus, our approach to segmentation explicitly accounts for the region-contour duality of inferring form from images. The trixels are represented as the vertices of a graph, with two adjacent trixels connected by an edge if their common edge is not retained as a contour completion element by the perceptual filters (Fig. 2(e)). Connected components of this graph correspond to trixel groupings, yielding a polygonal decomposition of the image with each polygon being assigned the area-weighted average color of the triangles constituting it (Fig. 2(f)).

These polygons are attributed with structural information such as area, aspect ratio, orientation, etc, as needed in a specific application. This representation will enable further high-level processing of structural information to extract and recognize objects, detect features of interest based on generic structure or color properties, and efficiently represent images in terms of constituent features.

The segmentation algorithm outlined above results in polygonal regions that correspond to objects or parts of objects. The triangles belonging to each such excised part can be used to efficiently compute geometric attributes of the part for its characterization and recognition[8,9]. Examples of such attributes are its area, average width, intrinsic aspect ratio, etc. Thus each part is assigned a vector of shape attributes that characterizes it. Ideally, the set of attributes characterizing a shape or its parts should be invariant to rotation and translation. Scale invariance may be desirable when there is no control over the imaging distance of the objects of interest, such as in real-world imagery obtained from a video camera. However, in applications such as remote sensing, industrial machine vision, and radiographic object detection, the imaging distance from objects of interest is typically relatively fixed or easy to ascertain. In such cases, the observed size of an object is an important attribute that is *relatively* invariant because of the fixed/known range of imaging distance. Examples of such attributes are area and average width of a shape. The internal aspect ratio mentioned above is an example of an *absolutely* invariant attribute that is constant under translation, rotation, and scaling. There are many other ways to generate invariant attributes of shapes, such as the computation of invariant shape moments.

4. APPLICATION TO ANOMALOUS CHANGE DETECTION

The edge-constrained triangulation of an image, along with association of the average spectral value of the pixels in each triangle to the triangle, constitutes a local edge-adaptive smoothing operation that is restricted to non-edge pixels. This mitigates image noise by regularizing the variability of pixel values within each triangle. In this sense, as descriptors of images, the trixels are less susceptible to image noise than pixels. However, since we are interested in detecting changes between two images, we have to account for the possibility that there is a difference between the images. In this case the triangulation of one image will not conform to that of the other in regions where change has occurred. As a result, some trixels in one image will straddle some edges in the other. This means that neither triangulation will serve as a spatially adaptive smoothing grid for both images. This calls for a refined grid that adapts to the edges of both images. A natural solution is the triangulation on the union of edges of both images. That is to say, once we obtain the edge pixels of both images, we combine these edge pixels into a single edge image and compute the constrained Delaunay triangulation of this combined edge set. With this 'hyper-triangulation' we can compute the trixel-smoothing of either image without straddling edges. Doing so will yield two spectral values for each trixel; one for each image. Thus distributions $P(x,y)$ and $P(x)P(y)$ can be estimated, as described in section 2, with the additional modification of associating a frequency to each trixel equal to its area relative to the total image. We have now reformulated FACD for images sampled with trixels. This scheme can be expressed schematically as follows:

$$\left.\begin{array}{l} A \rightarrow \partial A \\ B \rightarrow \partial B \end{array}\right\} \partial A \cup \partial B \rightarrow \Delta(\partial A \cup \partial B) \rightarrow \mathrm{FACD}(\Delta(\partial A \cup \partial B), A, B) \tag{4}$$

Here, A and B are two candidate images, ∂ indicates edge detection, Δ indicates constrained Delaunay triangulation, and FACD(T, A, B) is anomalous change detection between A and B with respect to the triangulation T.

We can also independently compute polygons by perceptual grouping of these trixels as described in section 3. As mentioned before, the polygons are associated the area-weighted aggregate spectral value of their constituent trixels. For the purpose of structural FACD, however, the polygons are associated the aggregate anomalousness score of their constituent trixels. In effect, each polygon has a measure of anomalousness associated with it. This has two advantages: First, it puts the anomalousness of a trixel in context by associating it with an object or feature polygon. That is, a trixel's anomalousness is appreciated or devalued based on the collective contribution of its fellow trixels in its parent polygon. Second, the shape of a polygon can be used to characterize the nature of the change. This adds a new dimension to ACD by providing a structural component in deciding the nature and importance of an anomalous change.

One can control the degree of spatial localization of an anomalous change by varying the sensitivity of the edge detector. A more sensitive edge detector would pick up fainter edges and the resulting triangulation would yield more trixels. The polygons resulting from perceptually grouping these trixels would also typically increase in number. An anomalous change encased in a smaller polygon would be stronger than when it is encased in a larger polygon. However, the finer polygonization will likely be structurally less informative because of oversegmentation and fragmentation of features. Thus, one can trade-off between anomaly localization and characterization. This structural tunability is in addition to the spectral tunability of the anomalous change detector where changing the level curve specifying the detector sensitivity will trade-off between probability of detection and false alarm rate.

5. EXPERIMENTAL RESULTS

In this section we illustrate the benefits of applying FACD to polygons instead of pixels. We will do so with two types of data; one set consists of multiband images of an indoor scene of objects on a table imaged with a long wave infrared (LWIR) hyperspectral sensor with 255 channels[10] (we will use the first three principal components of these images in our experiments), and the other consists of aerial grayscale images of a residential neighborhood. The indoor images used here show the introduction of a bottle containing a fluid, and the displacement of the bottle. The outdoor aerial scene images used here show the appearance of a vehicle on a road and its displacement along that road. In the first set, the goal is to detect the introduction and displacement of the bottle as anomalous changes between appropriate images, while ignoring changes in ambient lighting and contrast. In the second set, the goal is to detect the appearance of the car as the anomalous change while ignoring changes in parallax due to camera movement and misregistration of image pairs.

In all figures to follow that show the result of FACD applied to either pixels or polygons, only those pixels/polygons whose anomalousness exceeds the mean value by two standard deviations are shown. In the electronic online color version of this paper, anomalous polygons are outlined in red and are filled with a green color whose intensity is proportional to the anomalousness of the polygon.

Between figures 3(a) and 3(b) the anomalous change is the introduction of the bottle in the right image. Figure 3(c) shows the anomalous change pixels obtained by applying FACD. Figure 3(d) shows the anomalous polygons obtained by applying FACD to trixels and then averaging their values over the polygons to which they belong. In this case, the sensitivity of the Canny edge detector was set at 0.1. The sensitivity specifies the maximum fraction of image pixels allowed to be edge pixels. Typically, the actual number of edge pixels detected with a specified sensitivity will be much less than the specified upper bound fraction of image pixels. For this value of sensitivity (i.e., 0.1,) the bottle and its reflection are detected as anomalous change polygons. On raising the sensitivity to 0.25, the fluid-filled region of the bottle is identified as an anomalous polygon (figure 3(e)). Figure 3(f) shows the polygon as a strong anomalous change. The latter polygon is more specific and localized but perhaps less informative from a shape-theoretic point of view. This ability to obtain the shape of the anomalous change by applying FACD to trixels instead of pixels helps qualify the change by shape analysis.

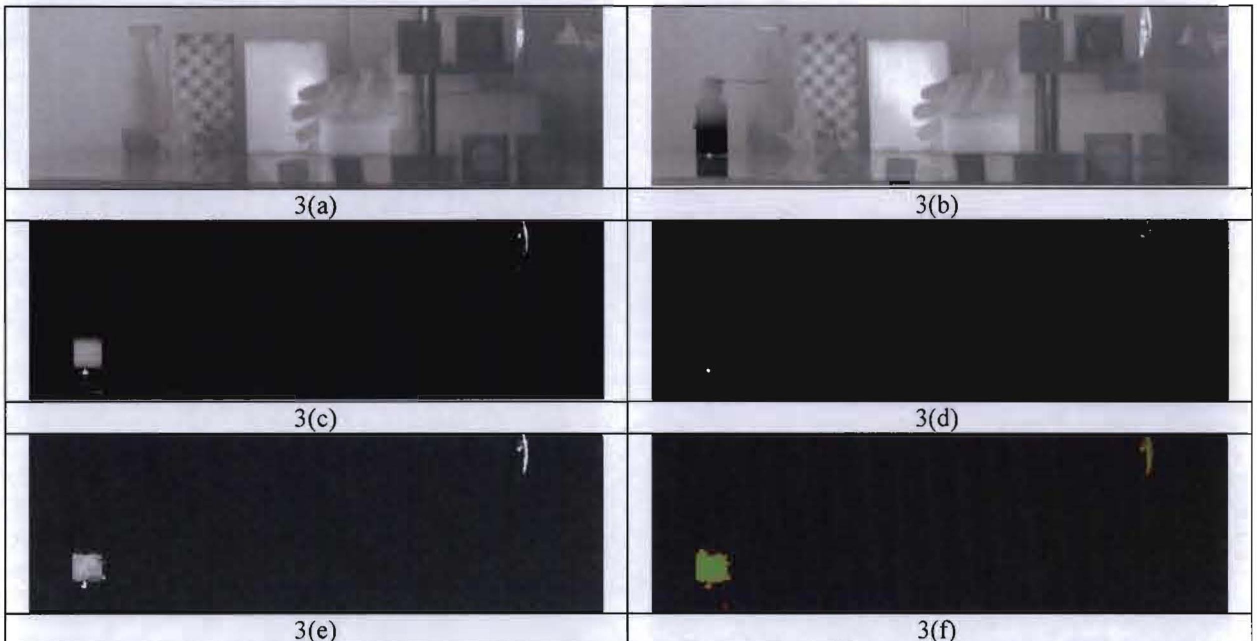

Figure 3: Detecting the shape of an anomalous change.

Between figures 4(a) and 4(b), the anomalous change is that the bottle has been slightly displaced to the right in the image on the right. Figure 4(c) shows the anomalous pixels obtained by pixel-based FACD. Figure 4(d) shows the anomalous polygons obtained by trixel-based FACD. The bottle's displacement is detected as thin anomalous polygons in the joint triangulation of the two images' edges that are absent as features in the polygonal segmentations of the individual images. Here we have a characterization of the nature of the anomalous change, namely movement.

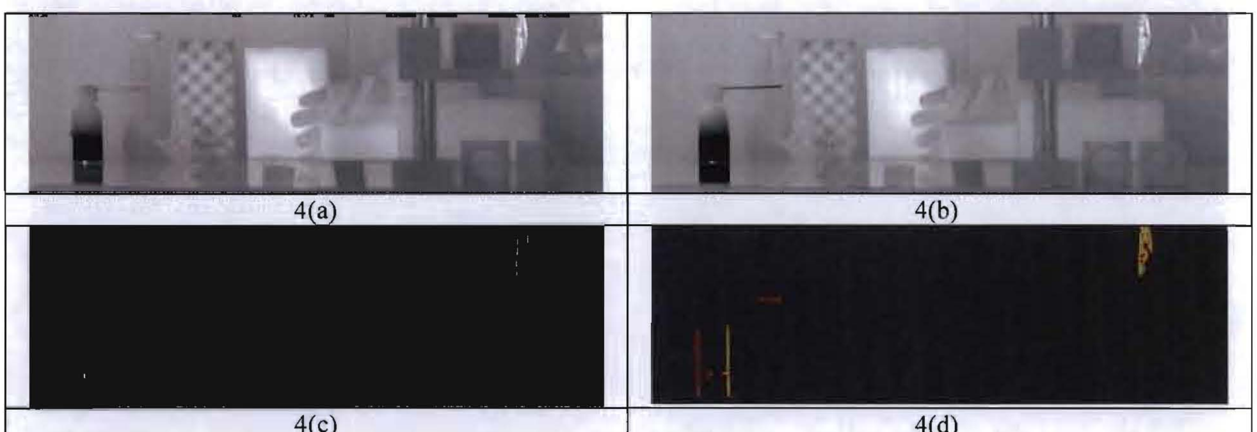

Figure 4: Detecting displacement of an object as an anomalous change

We now turn to the outdoor aerial image data set. In figures 5(a) and 5(b) the anomalous change is the appearance of the car at the left edge of the image on the right. The other changes are caused due to change in the brightness of certain features and a slight parallax shift due to the aircraft's movement. Figure 5(c) shows the pixel-based FACD results and figure 5(d), that of the trixel-based FACD. In the latter case, a geometric filter has been applied to the polygonal FACD image, wherein thin polygons are identified using the geometry of shapes. The criterion for thinness is based on the notion of the intrinsic aspect ratio of a polygon easily computed from the triangles constituting it as mentioned at the end of section 3. By eliminating polygons whose intrinsic aspect ratio is less than a specified threshold, we obtain the reduced set of anomalous polygons shown in 5(d) with the polygon corresponding to the car being the strongest. We check for a pervasive presence (> 50%) of such polygons in the anomalous set before eliminating them as being non-anomalous changes.

5(a)

5(b)

5(c)

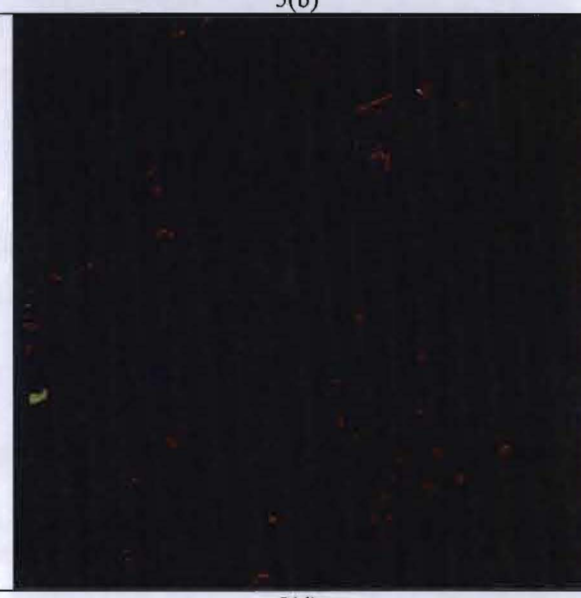

5(d)

Figure 5: Mitigating anomalous changes due to parallax via trixel-based FACD

In figure 6(a) a later image of the car in image 5(b) further down the road is shown. The result of pixel-based FACD on this image along with the image without the car in figure 5(a) is shown in figure 6(b). An increase in parallax artifacts can clearly be seen as thin ridges of anomalous change pixels. Further, it should also be noted that the anomalousness of car is diminished with respect to other spurious anomalous changes. Figure 6(c) and 6(d) show anomalous change polygons as result of trixel-based FACD before and after filtering out thin polygons.

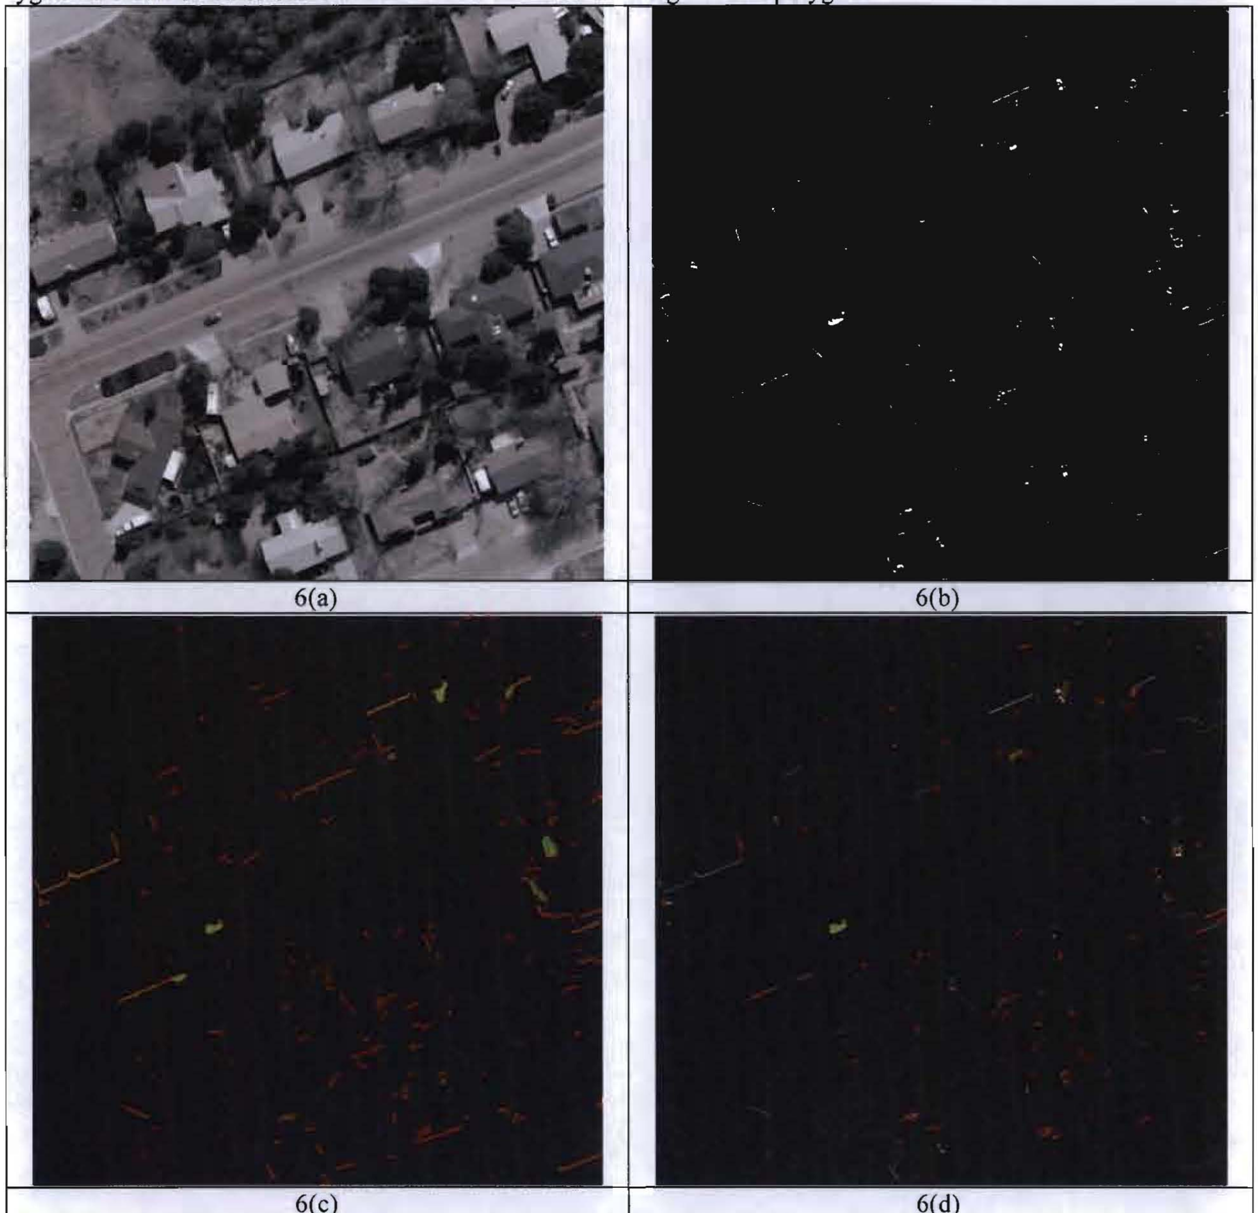

6(a) 6(b) 6(c) 6(d)

Figure 6: Mitigating parallax effects in anomalous change detection via trixel-based FACD.

The abundance of thin polygons in FACD is not only a hallmark of parallax, but also of shadow encroachment with the passage of time. Thus the pervasive presence of thin anomalous polygons corresponding to dark regions in one of the images is indicative of shadow lengthening (or shortening in the other image). Another common source of artifacts that confounds FACD is image misregistration. In the case of misregistration of a few pixels on average across the image, it manifests itself as locally linear shifts and gives rise to thin polygons. Here again, we can use trixel-based FACD to mitigate the effect. Figures 7(c) shows the result of pixel-based FACD after a constant relative shift misregistration of 2 pixels in each direction is applied to the image pair in figures 7(a) and 7(b). We note that the car is significantly attenuated as an anomaly with respect to the misregistration artifact pixels. In figure 7(d) we see a significant mitigation of the misregistration artifacts by elimination of thin polygons from the anomalous polygon population.

In each of the cases illustrated so far, we do not attempt to correct the images to reverse the unwanted effects. Rather, we set the anomaly value of the false positives corresponding to the thin polygons to zero and renormalize the values of the remaining polygons. This reinstates the dominance of true anomalous changes in the hierarchy, bolstering robust detection.

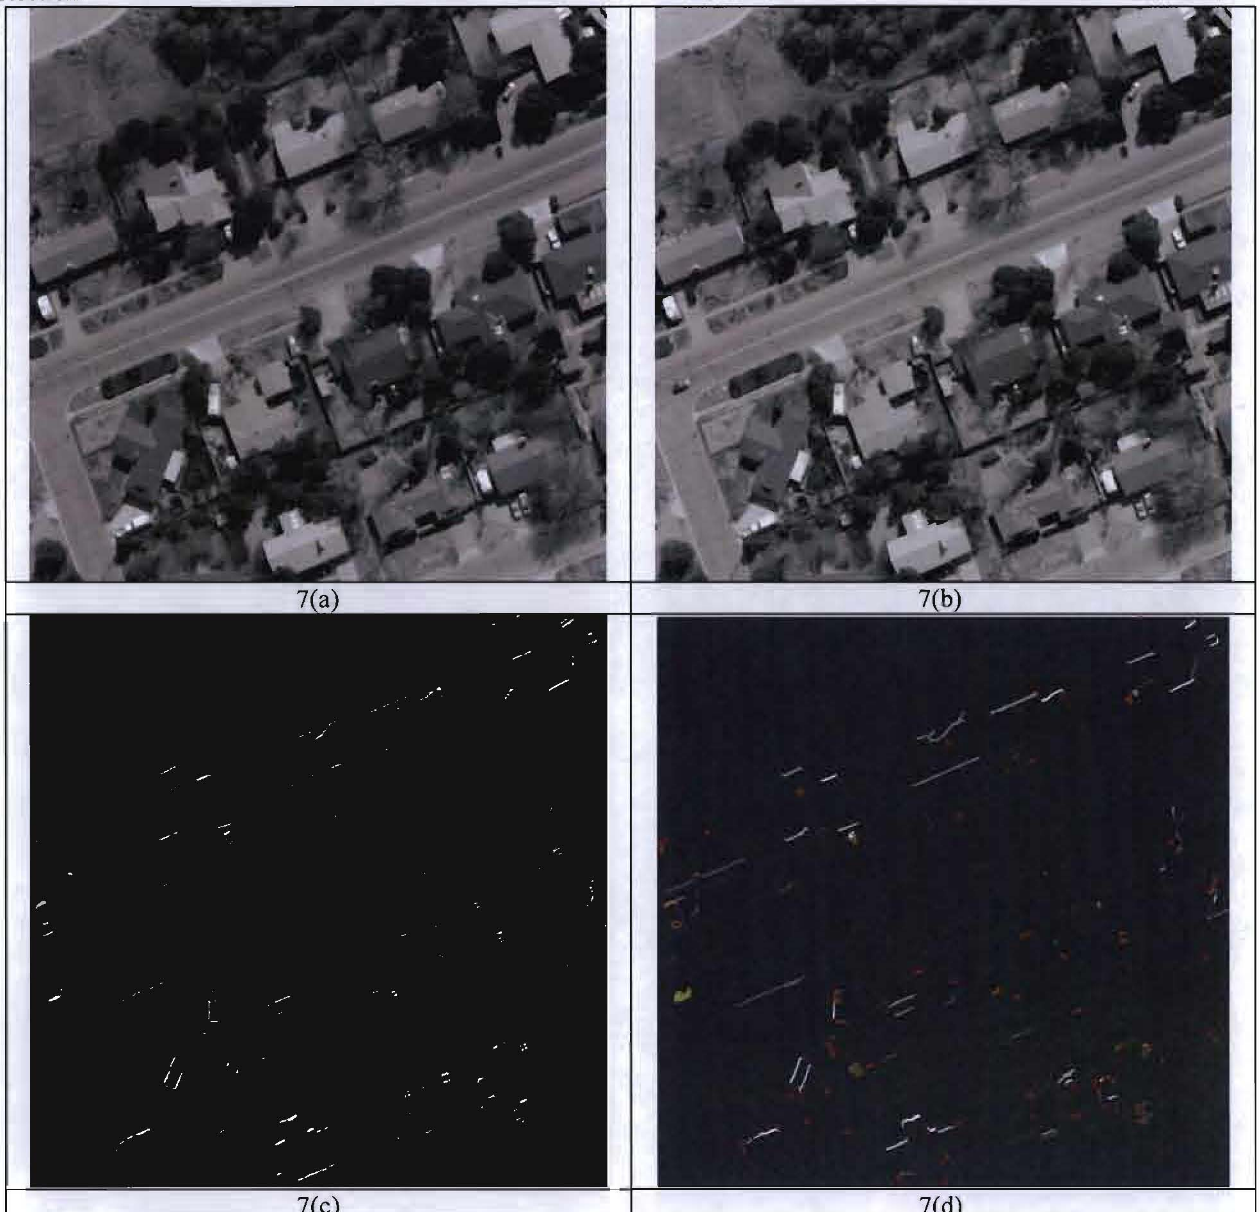

7(a) 7(b) 7(c) 7(d)

Figure 7: Mitigating misregistration artifacts in anomalous change detection via trixel-based FACD

6. SUMMARY AND FUTURE WORK

In this paper we have briefly presented a structural enhancement to the framework of anomalous change detection[1,2]. The motivation for this is the need to mitigate common confounding factors affecting change detection in real-world imagery. Further, we would also like to identify and characterize anomalous changes to understand what changed between two images of a scene. We believe the proposed modification to FACD provides a flexible paradigm for addressing and automating spatial analysis in change detection. In particular, we have identified a generic structural property, namely thinness, as a hallmark of spurious anomalous changes due to parallax, shadows and misregistration that are a common source of errors in ACD. We have demonstrated by example how the proposed structural framework helps identify and mitigate such effects.

Our future work will address the quantitative comparison of pixel-based vs trixel-based FACD by comparing them on a common footing. One way to do this would be to remap anomalous polygons into pixels and use the resulting regularized anomalous values of pixels to perform pixel-based FACD. This will enable the comparison of detector efficiencies via ROC curves on a quantitative basis. Another aspect of spatial analysis that will provide contextual understanding of ACD is to discover spatial relationships between changes. This is important for applications to surveillance and tracking. We will draw upon hierarchical polygonal segmentation[5], which is an extension of image polygonization approach described in this paper to extract context of changes and relate them.
Change detection is an important problem with strong implications to safety, security, and surveillance. While an exhaustive and thorough semantic understanding of imagery at large by computers is still a distant possibility, many of today's imminent image-based detection problems can be addressed with a judicious combination of spectral and spatial analysis. The structural framework for anomalous change detection is a step in that direction.

7. ACKNOWLEDGEMENT

This work described in this paper has been supported by two U. S. DOE LDRD DR grants under contract No. W-7405-ENG-36. We would like to thank Dr. Steven P. Love for kindly providing the ORCAS LWIR imagery used in this paper[10].

REFERENCES

[1] Theiler J. and Perkins S., "Proposed framework for anomalous change detection," ICML Workshop on Machine Learning Algorithms for Surveillance and Event Detection, pp. 7-14, 2006.
[2] Theiler J. and Perkins S., "Resampling approach for anomalous change detection," Proc. SPIE, vol. 6565, pp. 6565-65, 2007
[3] Theiler J., Harvey N. R., Porter R., and Wohlberg B., "Simulation framework for spatio-spectral anomalous change detection," To appear in Proc. SPIE, paper no. 7334-25, 2009
[4] Prasad L. and Skourikhine A. N., "Vectorized image segmentation via trixel agglomeration," Pattern Recognition, vol. 39, no. 4, pp. 501-514, 2006
[5] Prasad L. and Swaminarayan S., "Hierarchical image segmentation by polygon grouping," Proc. IEEE CVPR 2008 Workshop on Perceptual Organization in Computer Vision, Anchorage, AK, June 23, 2008
[6] Canny J., "A computational approach to edge detection," IEEE Transactions on PAMI, vol. 8, no. 6, pp. 679-698, 1986
[7] Goodman J. E. and O'Rourke J. (Eds.), "Handbook of discrete and computational geometry," CRC Press, 1997
[8] Prasad L. and Rao R. L., "A Geometric Transform for Shape Feature Extraction," Proc. SPIE, vol. 4117, Vision Geometry IX 2000.
[9] Prasad L., "Quantifying qualitative features," Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp 360-363, 2006.
[10] Love S. P., Hale T. C., Jolin L. J., Barefield J. E., Atkins W. H., and Tiee J. J., "The ORCAS long-wave infrared hyperspectral sensor," Proc. MSS (Military Sensing Symposium) (2008). LA-CP-08-0070.