A Novel Activity Detector Applied to Sentinel-1 for Surveillance

Axel Davy, Max Dunitz

A NOVEL ACTIVITY DETECTOR APPLIED TO SENTINEL-1 FOR SURVEILLANCE

Axel Davy, Max Dunitz

CMLA, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France

ABSTRACT

Change detection is the challenging process of identifying meaningful changes on an image sequence. In this work, we propose a new change detector for SAR images. We show that by controlling the a priori number of false alarms, one can detect events and filter out very small detections without first applying an anti-speckle filter. A detection is made in the absolute log-ratio image when enough pixels in a neighborhood exceed a statistically pre-determined threshold. This flexible approach allows several neighborhood shapes and sizes to be tested while controlling the number of false alarms. A global bound on the desired expected number of false alarms automatically determines the detection thresholds for each tested configuration.

***Index Terms*—** Change detection, small target detection, *a contrario* statistical model.

1. INTRODUCTION

SAR images have been successfully used to monitor agriculture [1], human activities, and the impact of disasters such as floods [1], forest fires [2], and earthquakes [3, 4]. SAR images have the advantage of being reliable even with adverse atmospheric conditions. Several techniques have been proposed to perform a variety of change detection tasks in SAR images. These include detecting complex coherence or inter-channel correlation changes [3, 4], looking at backscatter changes with a CFAR test [5], and catching textural changes [6, 1] by using techniques such as wavelets and PCA. These methods usually detect changes between two successive images using the log-ratio, which is recommended over a mere difference [7]. Clustering techniques may be used to locate a change [8] if one is known to have occurred.

In our work, we are interested in following human activities of small size compared to the image resolution, though very small and isolated activities of fewer than a predetermined number of pixels should be ignored. The detection thresholds should be adapted to ensure sensitivity to the targets of the intended size. To this end, we introduce a new detection technique. Similar to CFAR, its aim is to control the number of false alarms under a null hypothesis, which in our application is the assumption that no change occurs. Our detection method naturally avoids triggering small detections due to speckle or small targets below the intended size. Plus it can detect directly in noise, eliminating the need for preprocessing with a potentially distortive anti-speckle filter. We will illustrate the performance of the detector on Sentinel-1 sequences. Sentinel-1 is a constellation of satellites run by the ESA with global coverage and free data access. The constellation is currently composed of two satellites with a repeat cycle of 6 days.

Acknowledgement: work partially financed by ONR grant N00014-17-1-2552, DGA Astrid ANR-17-ASTR-0013-01, MENRT, FMJH and Kayrros, Inc.

2. OUR STATISTICAL ANOMALY DETECTOR

2.1. A short introduction to *a contrario* detection

Anomaly detection by an unsupervised method can be performed using the *a contrario* detection theory [9], which is a probabilistic formalization of the *non-accidentalness* principle [10]. The *a contrario* framework has produced impressive results in many different detection or estimation computer vision tasks, such as segment detection [11], ellipse detection [12], spots detection [13], and vanishing points detection [14], among others. The fundamental property of the *a contrario* theory is that it provides a way to automatically compute detection thresholds that yield a bound on the overall number of false alarms (NFA). It can be seen as an extension of the classic CFAR method.

The *a contrario* framework is a general methodology to set a detection threshold in terms of hypothesis testing. This is done by computing a number of false alarms (NFA) rather than just the probability of false alarm (PFA) used in hypothesis testing. It relies on the following definition.

Definition 1 *[13] Given a set of random variables $(X_i)_{i\in[1,N]}$ with known distribution under a null-hypothesis $(\mathcal{H}_0)$, a test function f is called an NFA if it guarantees a bound on the expectation of its number of false alarms under $(\mathcal{H}_0)$, namely:*

$$\forall \varepsilon > 0, \mathbb{E}[\#\{i, f(i, X_i) \leq \varepsilon\}] \leq \varepsilon.$$

To put it in words, raising a detection every time the test function is below ε should give under $(\mathcal{H}_0)$ fewer than ε false alarms in expectation. An observation $\mathbf{x}_i$ is said to be "ε-meaningful" if it satisfies $f(i, \mathbf{x}_i) \leq \varepsilon$, where ε is the predefined target for the expected number of false alarms.

While the definition of the background model $(\mathcal{H}_0)$ doesn't contain any *a priori* information on what should be detected, the design of the test function f reflects expectations on what is an anomaly. In short, applying the *a contrario* framework requires only a stochastic background model $(\mathcal{H}_0)$ of the random variables X_i, and a test function f. We refer the reader to [15] for several examples of NFA tests for image anomaly detection.

2.2. A novel detector for anomalous areas

For our detection framework, we assume the image on which we wish to perform anomaly detection follows the given model $(\mathcal{H}_0)$: the (X_i), which represent the pixel intensities, are independent random variables each following a known continuous distribution of survival functions $\mathrm{S}_i(x) := \mathbb{P}(X_i > x)$.

Our proposed NFA test is the following function:

$$\mathrm{NFA}_k(i) := NC_{n,k} \left(\{\mathrm{S}_j(x_j) : j \in B(i)\}_{(k)}\right)^k$$

where k is an integer, $C_{n,k}$ denotes a binomial coefficient and $\{\cdot\}_{(k)}$ denotes the k-th smallest value over a set, which in this case is the set of $\mathrm{S}_j(x_j)$ for pixels j in $B(i)$ a neighborhood of pixel i. It can be shown that $\mathrm{NFA}_k(i) \leq \varepsilon \Leftrightarrow \{\mathrm{S}_j(x_j) : j \in B(i)\}_{(k)} \leq \alpha_{N,\varepsilon,k} \Leftrightarrow \mathrm{D}_{N,\varepsilon,k}(i) = 1$ where

$$\mathrm{D}_{N,\varepsilon,k}(i) := \mathbb{1}_{\left(\sum_{j\in B(i)} \mathbb{1}_{(\mathrm{S}_j(x_j)\leq \alpha_{N,\varepsilon,k})}\right)\geq k}$$

and

- N is the number of tests, i.e. a multiple of the number of pixels in the image (ignoring border pixels).
- $B(i)$ is the neighborhood (such as a square or rectangular patch) of pixel i and is composed of n pixels.
- $\alpha_{N,\varepsilon,k} := \left(\frac{\varepsilon}{NC_{n,k}}\right)^{1/k}$, with ε the target number of false alarms per image.

We present here a proof that this test indeed raises fewer than the target ε false alarms under the null hypothesis:

$$\begin{aligned}
\mathbb{E}\left[\sum_i \mathrm{D}_{N,\varepsilon,k}(i)\right] &= N\mathbb{P}\left(\left(\sum_{j\in B(i)} \mathbb{1}_{(\mathrm{S}_j(X_j)\leq\alpha_{N,\varepsilon,k})}\right) \geq k\right) \\
&= N\sum_{l=k}^{n} \mathbb{P}\left(\left(\sum_{j\in B(i)} \mathbb{1}_{(\mathrm{S}_j(X_j)\leq\alpha_{N,\varepsilon,k})}\right) = l\right) \\
&= N\sum_{l=k}^{n} C_{n,l}(\alpha_{N,\varepsilon,k})^l(1-\alpha_{N,\varepsilon,k})^{n-l} \\
&\leq NC_{n,k}(\alpha_{N,\varepsilon,k})^k = \varepsilon
\end{aligned}$$

The passage from the third to fourth line in the demonstration above is effected by the following result:

$$\begin{aligned}
\sum_{l=k}^{n} C_{n,l}p^l(1-p)^{n-l} &= p^k n! \sum_{l=k}^{n} \frac{1}{l!(n-l)!} p^{l-k}(1-p)^{n-l} \\
&= p^k C_{n,k} \sum_{l=0}^{n-k} \frac{k!(n-k)!}{(l+k)!(n-k-l)!} p^l(1-p)^{n-k-l} \\
&= p^k C_{n,k} \sum_{l=0}^{n-k} \frac{k!l!}{(l+k)!} C_{n-k,l}p^l(1-p)^{n-k-l} \\
&\leq p^k C_{n,k} \sum_{l=0}^{n-k} C_{n-k,l}p^l(1-p)^{n-k-l} \\
&\leq p^k C_{n,k}(p+1-p)^{n-k} = p^k C_{n,k}
\end{aligned}$$

Since $\mathbb{E}\left[\sum_i \mathrm{D}_{N,\varepsilon,k}(i)\right] \leq \varepsilon$, this test is an NFA. Thus, $\mathrm{D}_{N,\varepsilon,k}(i)$ gives us a simple, statistically justified way to raise a detection if the values of (X_i) are anomalous (with respect to the background model) on a tested neighborhood $B(i)$.

To put the test in words, an anomaly is declared when a pixel i has a neighborhood $B(i)$ with at least k pixels $j \in B(i)$ more intense than $\mathrm{S}_j^{-1}(\alpha_{N,\varepsilon,k})$, where $\alpha_{N,\varepsilon,k}$ is computed to limit to ε the number of false alarms under the null hypothesis.

The larger the number k of pixels required to exceed their intensity thresholds, the lower those thresholds. Hence, a faint anomaly can be detected with sufficiently large k; conversely, an anomaly that affects fewer pixels will require higher thresholds. Since the intensity thresholds are derived from ε, multiple tests involving different choices of neighborhood B or anomaly size k may be flexibly combined. As the number of tests N increases, the detection thresholds adapt accordingly. Thus, the set of test windows must be complete enough to fit any targeted anomaly, while remaining reasonably limited in number to avoid under-detection.

2.3. Model for Sentinel-1

For our experiments, we took series of registered Sentinel-1 SLC images acquired from the same relative orbit on the same sub-swath, hence guaranteeing an almost identical incidence angle. The image size was 1000×1000 pixels. A radiometric correction was applied to ensure comparable intensities along the series. No speckle-removal filter was applied to the images to prevent targets from being hidden or distorted by denoising. In addition, we wanted to verify that our detector can perform well even on noisy images. We shall show the detection results on the log-ratio between consecutive images and on the log-ratio between an image and the geometric mean of the ten previous images. Under $(\mathcal{H}_0)$, no changes occurred, and the absolute log-ratio images are just noise with identical distribution per pixel.

The distribution of the log-ratio between consecutive images can be reasonably estimated with a Generalized Gaus-

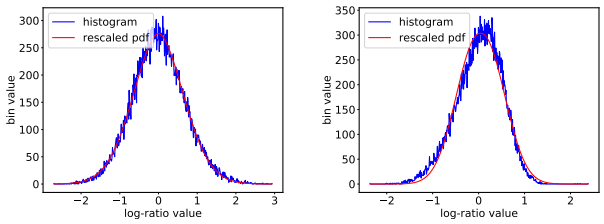

Fig. 1. Histogram of the log-ratio between two consecutive images, and the pdf of the estimated distribution [Left]. Histogram of the log-ratio between one image and the geometric average of the ten previous images, and the pdf of the estimated distribution [Right]

sian distribution, as shown in Figure 1. The distribution's parameters were estimated on a segment centered on the head of the distribution and of length four times the empirical standard deviation, to avoid being affected by potential targets in the histogram tail. We used the same model for the distribution of the log-ratio between an image and the average of the previous images, even though the model is less accurate, as shown on Figure 1. Once the log-ratio distribution is modelled, the absolute log-ratio survival function comes directly.

3. EXPERIMENTAL RESULTS

We shall compare the detection results on two scenarios. (A) corresponds to the log-ratio on two consecutive images, which is typically used in the context of change detection. (B) corresponds to the log-ratio of an image with the geometric average of the ten previous images. (B) has thus a lower noise level, but its practical application is more limited as it assumes no change occurred on the ten previous images. (B) can be used in areas with few changes to enhance detection. The log-ratios and the associated detections for (A) and (B) are shown on Figure 2, on the third and fourth columns, and the fifth and sixth columns, respectively.

Our results were obtained with the proposed detector, using several sets of parameters (the framework allows combinations of tests). We use two neighborhoods B: A square patch of size 2×2 and a square patch of size 3×3. For each B, we tested several k (i.e. the minimum number of pixels above the threshold to trigger a detection): 3 and 4 for the first neighborhood, and 7, 8 and 9 for the second.

To compare our results, we show in the second column of Figure 2 the detections made using a standard 7×7 Lee filter followed by a log-ratio of consecutive images and a manual thresholding (of the same value for all examples). This is most comparable to (A), which uses the same images for its log-ratio. Note, however, that the log-ratios displayed in Figure 2 do not have any Lee filter applied before computation. Our method detects directly on this log-ratio, unlike the compared log-ratio thresholding method, which requires a first denoising pass to get acceptable performance. Not requiring denoising enables better detection contours.

The target number of false alarms per image for our detector was set to $\varepsilon = 10^{-2}$. White pixels correspond to pixels that were above the tested thresholds in at least one successful NFA test.

In the first row of the example Figure 2, a boat and the water waves surrounding it are detected. In the second row, the disappearance or appearance of a few vessels and several ground structure changes are detected. In the third row, a novel human construction is detected. The detection results for (A) and (B) differ slightly. In the case of the second and third rows, the discrepancies are due to actual differences in the log-ratios. In the case of the first row, though, the lower noise level due to the averaging improves the contour of the detection. Overall, our detector shows a more accurate detection contour on these examples than the standard log-ratio test.

These examples highlight the specificity of the detector. Indeed, no very small detections of one or two pixels are present. Furthermore, targets are detected even in the presence of strong speckle noise. The contours of the detections are faithful to the changes detected.

4. CONCLUSION

This paper presented a novel detection test, based on the *a contrario* methodology. The performance of the detector was shown on Sentinel-1 images, on the log-ratio between consecutive images, and on the log-ratio between an image and the geometric average of the ten previous images. The log-ratio images' distributions were estimated with Generalized Gaussian distributions, and the proposed framework gave automatic detection thresholds with false alarms control. With these thresholds, our detector, in the presence of speckle noise, identified the contours of the targeted changes more accurately than a classical method relying on the use of a distortive speckle filter, and without raising smaller detections.

While the detector has been demonstrated here on log-ratio Sentinel-1 images using small square patches as neighborhoods, it need not be limited in this way. Detection problems involving periodic change, subtle but widespread or persistent change, perturbations in patterns of change, or change on different sorts of images could profit from different design choices and parameter combinations.

5. REFERENCES

[1] M. Jia and L. Wang, "Novel class-relativity non-local means with principal component analysis for multitemporal sar image change detection," *International Journal of Remote Sensing, Taylor & Francis*, vol. 39, no. 4, pp. 1068–1091, 2018.

[2] F. Bovolo and L. Bruzzone, "An adaptive multiscale approach to unsupervised change detection in multitemporal sar images," in *International Conference on Image Processing*. IEEE, 2005.

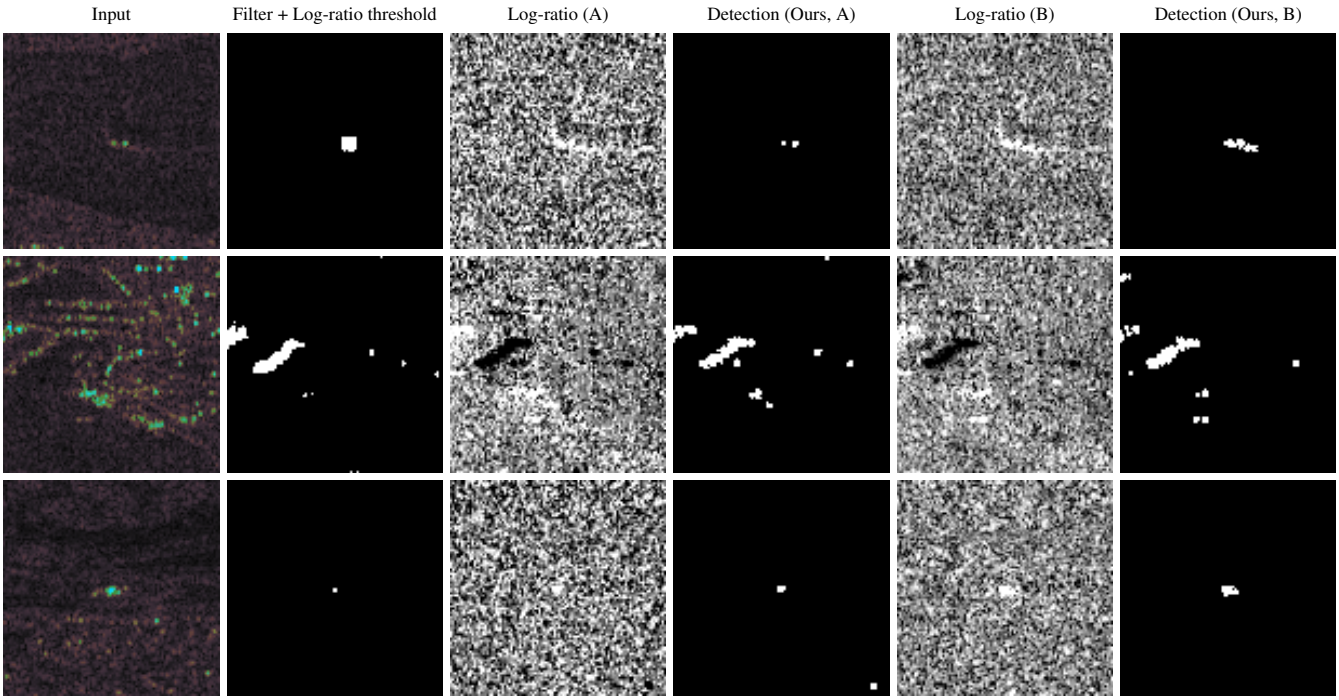

Fig. 2. Detection results on Sentinel-1 satellite imagery taken of a port in Houston, TX, on June 25, 2017 and January 3, 2018, and of a golf course in Kingswood, TX, on January 8, 2018 (rows 1-3, respectively). The reference image is shown on the first column in false color (light green: high backscatter, dark red: low backscatter). The second column shows the manually thresholded absolute log-ratio image, computed after application of a 7×7 Lee filter to the images. Columns 3 and 5 give the log-ratio images and columns 4 and 6 the detections our method produced on the corresponding absolute log-ratio images (no Lee filter was applied). Columns 2-4 used the previous image in the sequence in computing the log-ratio; columns 5-6, the geometric mean of the ten previous images.

[3] E. Ferrentino, A. Marino, F. Nunziata, and M. Migliaccio, "Multi-polarization methods to detect damages related to earthquakes," in *International Geoscience and Remote Sensing Symposium*. IEEE, 2018, pp. 1938–1941.

[4] P. Washaya and T. Balz, "Sar coherence change detection of urban areas affected by disasters using sentinel-1 imagery," *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pp. 1857–1861, 2018.

[5] K. El-Darymli, P. McGuire, D. Power, and C. Moloney, "Target detection in synthetic aperture radar imagery: a state-of-the-art survey," *Journal of Applied Remote Sensing*, vol. 7, no. 1, pp. 071598, 2013.

[6] S. Cui and M. Datcu, "Statistical wavelet subband modeling for multi-temporal sar change detection," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 4, pp. 1095, 2012.

[7] E. Rignot and J. Van Zyl, "Change detection techniques for ers-1 sar data," *IEEE Transactions on Geoscience and Remote sensing*, vol. 31, no. 4, pp. 896–906, 1993.

[8] T. Celik, "Change detection in satellite images using a genetic algorithm approach," *IEEE Geoscience and Remote Sensing Letters*, vol. 7, no. 2, pp. 386–390, 2010.

[9] A. Desolneux, L. Moisan, and J.-M. Morel, *From gestalt theory to image analysis: a probabilistic approach*, vol. 34, Springer Science & Business Media, 2007.

[10] D. Lowe, *Perceptual organization and visual recognition*, Kluwer Academic Publishers, 1985.

[11] C. Liu, R. Abergel, Y. Gousseau, and F. Tupin, "A line segment detector for sar images with controlled false alarm rate," in *International Geoscience and Remote Sensing Symposium*. IEEE, 2018.

[12] V. Patraucean, P. Gurdjos, and R. Grompone, "A parameterless ellipse and line segment detector with enhanced ellipse fitting," in *European Conference on Computer Vision*. Springer, 2012.

[13] B. Grosjean and L. Moisan, "A-contrario detectability of spots in textured backgrounds," *Journal of Mathematical Imaging and Vision, Springer*, vol. 33, no. 3, pp. 313–337, 2009.

[14] J. Lezama, R. Grompone von Gioi, G. Randall, and J.-M. Morel, "Finding vanishing points via point alignments in image primal and dual domains," in *Computer Vision and Pattern Recognition*. IEEE, 2014.

[15] T. Ehret, A. Davy, J.-M. Morel, and M. Delbracio, "Image anomalies: A review and synthesis of detection methods," *Journal of Mathematical Imaging and Vision, Springer*, Apr 2019.