

Leveraging structural information for enhanced coherent change detection

Scott Dayton^a, Oliver Milledge^a, Jovana Nikitovic^a, Anne Gelb^a, Dylan Green^a, and Aditya Viswanathan^b

^aDartmouth College, 27 N. Main Street, Hanover, NH, 03755, USA

^bUniversity of Michigan – Dearborn, 4901 Evergreen Road, Dearborn, MI, 48128, USA

ABSTRACT

We consider the two-pass coherent change detection problem for SAR imaging. Inspired by classical maximum likelihood-based coherent change detectors (Jakowatz, 1996)¹ and multi-polarization SAR change detection techniques (Novak, 2005),² we propose a method of incorporating underlying structural image information using specially formulated kernels. In particular, we utilize a class of convolutional edge detection kernels to extract underlying edge information in the scene of interest given noisy and potentially incomplete data. We then adapt existing multi-polarization SAR change detection methods to incorporate such edge information to improve the quality and robustness of resulting change maps. We validate the proposed method using real-world SAR images from the CCD Challenge Problem dataset and demonstrate improved change detection performance using empirical ROC studies.

Keywords: coherent change detection, synthetic aperture radar, edge detection, GOTCHA

1. INTRODUCTION

We consider the two-pass synthetic aperture radar (SAR) change detection problem,^{1,3} where we have access to data (or suitably processed complex-valued SAR imagery) of the same underlying scene at two different times. These are often referred to as the reference (or time-1) and mission (or time-2) images. It is assumed that the data are co-registered; i.e., the pixel locations in the processed SAR imagery of any common (unchanged) objects in the scene are identical between the two times. Changes in the scene may be significant – such as insertions or deletions of objects (such as cars or buildings) between the two passes, or more subtle – such as a disturbance of the ground in some small local region. We are interested in identifying such changes or regions of changes in the scene accurately and efficiently, with particular emphasis on urban environments. We do this by computing a *change map*, or more specifically a change statistic $\alpha \in [0, 1]$ for each pixel location in the scene, such that $\alpha = 0$ denotes a change between time-1 and time-2 while $\alpha = 1$ denotes no change. The continuum of values for α between 0 and 1 indicate varying levels of confidence in identifying changes in the scene.

Fig. 1a and Fig. 1b display representative time-1 and time-2 images (magnitudes of complex-valued SAR recoveries) of such a scene (see (Scarborough et al., 2010)³ for details of the scene and corresponding dataset), with the corresponding change “map” computed using the classical maximum-likelihood-based change detector (see [1, Section 5.5] for details) shown in Fig. 1c. The two passes correspond to a single co-registered underlying scene containing a parking lot with cars (towards the bottom right of the scene), a large region of foliage (towards the top left of the scene), a couple of buildings, some roads, and a running track towards the right-center of the scene. Of note is the large shadow in the immediate vicinity of the building at the bottom right. From time-1 to time-2, there are insertions and deletions of vehicles in the parking lots, especially in the bottom right of the scene, as well as changes due to usage of the running track. The change map computed using the MLE-based change detector identifies many of these changes. However, it is evident that there are significant false changes recorded in regions of foliage and around building shadows. Results using the proposed method incorporating edge information are shown in Fig. 1d and discussed in §3.

Further author information: (Send correspondence to Dylan Green)

Dylan Green: E-mail: dylan.p.green.gr@dartmouth.edu

S. Dayton, O. Milledge, J. Nikitovic: {scott.d.dayton.23,oliver.p.milledge.23,Jovana.Nikitovic.25}@dartmouth.edu

A. Gelb: anneglb@math.dartmouth.edu; A. Viswanathan: adityavv@umich.edu




Figure 1: Representative SAR change detection scene and change maps. The top panel shows the time-1 and time-2 images of the scene (magnitude images plotted on a log scale with 60dB of dynamic range). Fig. 1c shows the change map computed using the classical MLE-based method; Fig. 1d shows equivalent results using the proposed method.

1.1 Problem Setup

The scene is represented with images of size $N_1 \times N_2$. The *complex*-valued pixel reflectivities are respectively denoted by f_{j_1, j_2} and g_{j_1, j_2} at time-1 and time-2, with $(j_1, j_2) \in \{1, \dots, N_1\} \times \{1, \dots, N_2\}$. The *neighborhood* $\mathcal{N}_{j_1, j_2}^M$ of pixel (j_1, j_2) is defined by

$$\mathcal{N}_{j_1, j_2}^M = \{(k_1, k_2), \mid k_1 = j_1 - M, \dots, j_1 + M, k_2 = j_2 - M, \dots, j_2 + M\}, \quad M \in \mathbb{N}, \quad (1)$$

which describes a symmetric $(2M + 1) \times (2M + 1)$ square patch of pixels closest to pixel (j_1, j_2) . The pixel intensities in the scene are further assumed to follow a zero mean* i.i.d. (complex) Gaussian distribution, with

*In practice, we subtract the sample mean (evaluated over the neighborhoods (1)) before computing the change map.

independent (complex) zero mean additive white Gaussian noise so that

$$\begin{aligned} \text{(time-1)} \quad & f_{j_1, j_2} = r_{j_1, j_2} + n_{f_{j_1, j_2}} \\ \text{(time-2)} \quad & g_{j_1, j_2} = \alpha_{j_1, j_2} r_{j_1, j_2} e^{i\phi} + \sqrt{1 - \alpha_{j_1, j_2}^2} z_{j_1, j_2} + n_{g_{j_1, j_2}}. \end{aligned} \quad (2)$$

Here $r_{j_1, j_2} \sim \mathbb{CN}(0, \sigma_r^2)$ and $z_{j_1, j_2} \sim \mathbb{CN}(0, \sigma_r^2)$ are uncorrelated with each other and used to model terrain reflectivities at pixel (j_1, j_2) . The additive noise $n_{f_{j_1, j_2}}, n_{g_{j_1, j_2}} \sim \mathbb{CN}(0, \sigma_n^2)$ is such that the individual noise components are both uncorrelated with each other as well as with the signal components. The *change statistic* is denoted by α_{j_1, j_2} , and finally $e^{i\phi}$ with $\phi \in [0, 2\pi)$ is a phase term included to account for unavoidable phase errors in multi-pass SAR setups. Note that if the phase term is neglected, setting $\alpha = 1$ in (2) corresponds to the “no change” scenario (as desired), while setting $\alpha = 0$ corresponds to change in reflectivity between the two passes. For reference, the classical MLE change detector is obtained by computing the maximum likelihood estimate for α in (2) (while treating ϕ as an unknown nuisance parameter). Assuming independent observations of the reflectivities in a neighborhood of pixel (j_1, j_2) , it can be shown (see [1, Section 5.5] for details) that

$$\alpha_{j_1, j_2}^{\text{MLE}} = \frac{2 \left| \sum_{(k_1, k_2) \in \mathcal{N}_{j_1, j_2}^M} \overline{f_{k_1, k_2}} g_{k_1, k_2} \right|}{\sum_{(k_1, k_2) \in \mathcal{N}_{j_1, j_2}^M} |f_{k_1, k_2}|^2 + \sum_{(k_1, k_2) \in \mathcal{N}_{j_1, j_2}^M} |g_{k_1, k_2}|^2}, \quad (3)$$

where $\overline{f_{k_1, k_2}}$ denotes the complex conjugate of f_{k_1, k_2} and $|\cdot|$ denotes the component-wise absolute value.

1.2 Related Work

An early method proposed in (Touzi et al., 1988)⁴ involved estimating and comparing the mean backscatter powers of the scene in the two passes. The classical MLE change detector described in (3) is derived, for example, in (Jakowatz et al., 1996).¹ This method essentially computes a scaled version of the magnitude of the sample coherence, also called complex cross-correlation. For low clutter-to-noise scenarios, a related method for estimating complex reflectance change in (2) and introduced in (Wahl et al., 2016)⁵ has been shown to provide better results. An alternate Bayesian hypothesis testing-based problem formulation and resulting log-likelihood change statistic is described in (Preiss et al., 2006).⁶ In (Cha et al., 2015),⁷ the authors propose a two-stage technique which combines a non-coherent intensity change statistic – suitable for identifying significant scene changes – with an alternative (Berger’s) coherence-based estimator to identify more subtle scene changes. A framework based on tracking the evolution of local statistics between passes, computed using cumulant-based series expansions and compared using Kullback-Leibler divergence measures, along with the potential for its application in a multi-scale approach, is discussed in (Inglada and Mercier, 2007).⁸

The methods mentioned above are generally utilized with single polarization SAR data. For multi-polarization SAR, a generalized likelihood ratio test (GLRT) for the change detection problem was introduced in (Novak, 2005)² and partly motivates the proposed method in this manuscript. False alarm mitigation techniques for SAR change detection are addressed, for example, in (Phillips, 2011)⁹ and (Newey et al., 2013).¹⁰

1.3 Contributions

We propose a new coherent change detection scheme which incorporates relevant underlying structural information. In particular, we extract descriptors of object and region boundaries in the magnitude image via edge detection and use a multi-polarization change detection framework to incorporate such information into the change statistic. Moreover, we use image connectivity analysis in local patches, as well as entropy filtering-based low-RCS detection techniques to distinguish true changes from false positives. Empirical results on real-world SAR images are presented using the CCD challenge problem dataset.³

The rest of this paper is organized as follows: Section 2 explains the components of the proposed method, with §2.1 detailing the entropy filtering technique to identify low-RCS regions, §2.2 explaining the false positive

mitigation technique based on image connectivity analysis, §2.3 reviewing the multi-polarization change detection framework of (Novak, 2005),² and §2.4 detailing the edge detection methods used in our empirical results. Section 3 explains how these components work together to construct an enhanced change detection algorithm and provides representative empirical results (including ROC analysis) on real-world SAR images. Finally, Section 4 provides some concluding remarks and avenues for future work.

2. COMPONENTS OF THE PROPOSED METHOD

2.1 Low-RCS Masking

It is easy to verify that (3) is undefined when $|f_{k_1,k_2}| = |g_{k_1,k_2}| = 0$ for all pixels in a neighborhood. Moreover, there is a distinct possibility of spurious change detection results in areas of the scene corresponding to low signal or radar cross-section (RCS), as is evident in building shadow regions, and to a lesser extent along roads in Fig. 1c. Identifying such regions and setting the corresponding pixel values in the time-1 and time-2 images to be identical mitigates this issue. Following the convention in (Stojanovic and Novak, 2013)¹¹ we set the pixel values to be identically unity, although any non-zero value will also work. Furthermore, since entropy is a statistical measure of randomness, it is useful in distinguishing pixels containing random noise only – as is likely the case with low radar signal – versus regions of structure. We therefore propose the use of entropy filtering (Gonzalez and Woods, 2017)¹² to identify such regions robustly, with

$$e_{j_1,j_2} = - \sum_{(k_1,k_2) \in \mathcal{N}_{j_1,j_2}^M} p_{k_1,k_2} \log_2(p_{k_1,k_2}). \quad (4)$$

Here e_{j_1,j_2} denotes the computed entropy at pixel (j_1, j_2) , and p_{k_1,k_2} represents normalized histogram counts for the magnitude image. Note that we apply (4) in neighborhoods defined by (1) to be consistent with the change detection scheme. Moreover, we apply the entropy filter to the time-1 and time-2 magnitude images separately before taking the pixel-wise minimum entropy value between the two results. The results of such a computation, with values re-scaled to $[0, 1]$, are shown in Fig. 2a. Entropy values above a user-specified threshold (0.85 in our computations) are used to identify regions of low signal, while the corresponding pixel locations are used to generate a mask identifying such regions. The results of applying the MLE change detector (3) after such masking and signal correction is shown in Fig. 2b. Comparing with Fig. 1c, note the lack of change recorded in the building shadow region toward the bottom right of the scene.




Figure 2: Low-RCS region identification via entropy filtering (Fig. 2a) and change detection (using the MLE method) with low-RCS masking (Fig. 2b). Note the absence of false changes in the building shadow regions at the bottom right of the image when compared to Fig. 1c

2.2 False Positive Mitigation via Image Connectivity

Fig. 1c identifies changes in regions with foliage, the running track, and cars (parking lots). Since we are interested in change detection in urban environments, large scale and intermittent (scattered) changes in foliage regions are likely to be false positives caused by sensitivity of the radar system to the rustling of leaves. We identify such changes from the computed change map and apply a post-processing procedure to minimize false positives. Towards this end, we utilize ideas from *image connectivity*. A *connected component* (Gonzalez and Woods, 2017)¹² in a *binary* image is a set of “adjacent” pixels, with two pixels being consider adjacent if they are both nonzero, and occupy consecutive horizontal, vertical or diagonal pixel locations.




Figure 3: Connected components – illustration and application to the scene. Fig. 3a illustrates the concept of connected components in (binary) images. Fig. 3b applies this concept to the change map image from Fig. 1c in three distinct regions of the scene. Note the significantly larger number of connected components in the foliage region.




Figure 4: Applying image connectivity analysis to reduce false positives in change maps. Fig. 4a plots the number of connected components in local 31×31 neighborhoods across the entire scene; Fig. 4b plots a post-processed (using image connectivity information) change map and shows reduced false positives.

Fig. 3a illustrates this concept by plotting the matrix representing a binary image, the corresponding image,

and the extracted (indicated using different colors) connected components. The results of applying this analysis to three distinct regions (identified using yellow boxes in Fig. 1c) of the scene – foliage, running track, and a parking lot – are shown in Fig. 3b. The top panel in Fig. 3b shows the corresponding regions of the change map from Fig. 1c while the bottom panel shows the connected components. In each case, the grayscale change map is converted to a binary image via thresholding[†] before the connected components are identified using convolutional kernels or graph-theoretic methods (see, for example (Gonzales and Woods, 2017)¹²). We use Matlab's `bwconncomp` implementation to perform such connectivity analysis. While properties of these connected components (area, shape, orientation etc.) may be useful and present an opportunity for further investigation, we restrict our current attention to the *number* of such connected components. In particular, note the significantly larger number of components in the foliage region in comparison to the other regions. Indeed, the results of performing this analysis across the entire scene (using image patches of size 31×31 in a raster scan fashion) are shown in Fig. 4a. Note the large number of connected components in the foliage (and building shadow) regions directly coincide with regions yielding a significant number of false positives. By suitably thresholding this image, for example, here we use 35% of the maximum number of connected components in the scene, we may generate a binary mask identifying regions with likely false positives. Corresponding pixels in the change map may be artificially set to 1 (denoting no change) as a post-processing step. The resulting post-processed change map is shown in Fig. 4b; note the significant reduction in false positives in comparison to the classical MLE-based results in Fig. 1c.

2.3 Multi-Polarization SAR Change Detection

In multi-polarimetric SAR imaging, we acquire measurements corresponding to multiple polarization returns, such as HH, HV and VV configurations. For change detection problems, we further assume that two or more of these returns at two distinct passes are co-registered between themselves as well as between passes, so as to represent the same underlying scene. For simplicity in presentation we vectorize the SAR image so that $j = (j_1, j_2)$, for $j = 1, \dots, N_1 N_2$. This leads to, for example, $\mathbf{f}_j = [f_j^{HH} \ f_j^{HV} \ f_j^{VV}]^T$, where each component represents the reflectivity of the specified area of the scene under the corresponding wave polarization. Assuming the reflectivities can be modeled as (complex; zero mean) Gaussian random variables, $\mathbf{f}_j$ is a three-dimensional complex Gaussian vector with some underlying covariance matrix $C_j^f = \mathbb{E}(\mathbf{f}_j \mathbf{f}_j^*)$. Similarly, we have the time-2 vector $\mathbf{g}_j = [g_j^{HH} \ g_j^{HV} \ g_j^{VV}]^T$, characterized by the corresponding covariance matrix $C_j^g = \mathbb{E}(\mathbf{g}_j \mathbf{g}_j^*)$. If there is no change in the underlying scene between these two passes, we expect (statistically) that $C_j^f = C_j^g$. The alternate hypothesis therefore corresponds to a change, that is $C_j^f \neq C_j^g$. With independent observations of the reflectivities in a neighborhood of pixel $\mathbf{x}_j$, following (Novak, 2005),² we formulate a generalized likelihood ratio test resulting in the change detection statistic

$$\alpha_j^{\text{pol}} = \frac{\det \left(\frac{1}{|\mathcal{N}_j^M|} \sum_{i \in \mathcal{N}_j^M} \mathbf{f}_i \mathbf{f}_i^* \right) \det \left(\frac{1}{|\mathcal{N}_j^M|} \sum_{i \in \mathcal{N}_j^M} \mathbf{g}_i \mathbf{g}_i^* \right)}{\left[\det \left(\frac{1}{2} \left(\frac{1}{|\mathcal{N}_j^M|} \sum_{i \in \mathcal{N}_j^M} \mathbf{f}_i \mathbf{f}_i^* + \frac{1}{|\mathcal{N}_j^M|} \sum_{i \in \mathcal{N}_j^M} \mathbf{g}_i \mathbf{g}_i^* \right) \right) \right]^2}, \quad j = 1, \dots, N_1 N_2. \quad (5)$$

Here, $|\mathcal{N}_j^M|$ denotes the number of pixels in the neighborhood (1) of pixel $\mathbf{x}_j$. As before $\alpha_j^{\text{pol}} \in (0, 1]$, with $\alpha_j^{\text{pol}} \approx 0$ denoting a change. Representative results using this change detection statistic are shown in Fig. 5b with the MLE-based change map from Fig. 1c re-plotted in Fig. 5 for reference and easy comparison. Both methods result in similar change maps, with some minor differences due to the different signal returns from the multi-polar components. For example, the HH, VV and HV polarizations have slightly different sensitivities to vegetation,

[†]Global or adaptive local thresholds may be used and are handled automatically by Matlab's `imbinarize` command.

roads and other such scene features. Changes in such features are therefore more (or less) pronounced when using the multi-polar formulation. Nevertheless, this method provides a simple and effective way to incorporate multi-component data into the change detection framework.




Figure 5: Change map (a) classical MLE-based method (Fig. 1c) and (b) using multi-polarization data and change statistic (5).

Remark: To account for potential disparities in scale with multi-polarization data, we scale the images (across all polarizations and passes) to have the same Frobenius norm.

2.4 Edge Detection

Edges often encapsulate important features in the underlying scene, especially in urban environments where rigid objects are prevalent. Rather than constructing the underlying image, it may even be easier to approximate internal boundaries, especially in cases where the data are corrupted and/or incomplete. Sparse regularization techniques such as (Stefan et al., 2011)¹³ are particularly useful for this purpose. Due to the combination of speckle and additive Gaussian noise components, SAR data is often highly contaminated, however. Moreover, sparse regularization techniques are typically not designed for complex-valued signals.

Under the assumption that edges manifest in the magnitude rather than the phase of a signal, we apply convolutional edge detection to the magnitude component, while also incorporating filtering techniques on both pre- and post-processed data. More specifically, we apply edge preserving anisotropic diffusion-based speckle filtering (Yu and Acton, 2002),¹⁴ followed by a convolutional Sobel kernel edge detector. We then used matched filtering to reduce false positives. We note that our three-step approach preserves edge “heights” since we skip the thresholding step typically seen in conventional edge detectors, and that such information proves useful in subsequent change statistic computations.

A couple of remarks are in order. First, when the underlying scene of interest contains few anisotropic scatterers, SAR data can be interpreted as nonuniform Fourier data (Jakowatz et. al., 1996).¹ Although the data in our experiments come from pre-formed images, it is possible to use the so-called concentration kernels (Gelb and Tadmor, 1999).¹⁵ These convolution kernels are used to identify edges in real-valued signals directly from frequency data, and as such data loss is minimized since an initial image recovery is not needed. The technique can be adapted to recover edges in the square magnitude of complex-valued data. Finally, the use of concentration kernels has been extended to nonuniform Fourier data (see, for example (Gelb and Song, 2020)¹⁶ and (Stefan et al., 2011)¹³), potentially allowing for use with SAR phase history data (PHD). Second, we use edges as a prototypical representative of useful features in the underlying scene. Other structurally informed

features (such as shape and texture) may be used in place of edges; we defer investigations involving such features to future work.

3. PROPOSED METHOD

The following two main ideas are utilized in constructing an enhanced coherent change detection procedure: (i) The construction of a structurally informed multi-component representation of the scene at any given pixel; and (ii) the formulation of a statistical test for identifying change in the *relationship between the components* of this multi-component representation at any given pixel from one pass to another. The results in this manuscript use edges as a surrogate for such structural information. As already noted, other features (topological shape, texture, morphology etc.) may also be used for this purpose. In addition, we propose use of the entropy filtering-based low-RCS and/or signal identification scheme from §2.1 to pre-process the given data and reduce resulting false positives in the change map. Finally, we note that the radar system may be sufficiently sensitive to record undesirable changes in the scene (such as the rustling of leaves). We propose use of the image connectivity-based analysis from §2.2 to post-process any intermediate change maps to suppress such false positives. Our proposed method is summarized in Alg. 1, with representative results provided in the center panel of Fig. 6. For reference and ease of comparison, the classical MLE-based change map is re-plotted in the left panel of Fig. 6. Observe that while there is significant reduction in the number of false positives, the changes in the running track region are no longer clearly identified. There is an inherent trade-off between such changes (which manifest in the phase of the SAR data) and false positives in regions of vegetation, building shadows etc. Further refinements in performing the image connectivity analysis in §2.2 will likely lead to better results. Alternatively we can take the convex combination of two change maps, such as those from Alg. 1 and Fig. 4b, to yield a composite change map with desired properties, as displayed in the right panel of Fig. 6. Freely available Matlab code used to generate these figures can be found at (SI-CCD, 2024).¹⁷




Figure 6: Change map using the proposed method. The left panel re-plots Fig. 1c for reference and comparison; the center panel is the change map obtained using Alg. 1 (proposed method); the right panel is a composite change map obtained by taking a convex combination of the results from the proposed method and Fig. 4b

Since ground truth changes are difficult to determine with the CCD challenge problem dataset, we instead present empirical ROC studies on representative (synthetic) test functions shown in Figs. 7a and 7b. These complex-valued test functions have two large rectangular model objects which do not change from time-1 to time-2. However, the three smaller square and circular objects change position between the two times. Both images are subject to i.i.d. additive complex Gaussian noise. The empirical ROC is computed using Monte Carlo simulations over 1000 trials and the four locations indicated by red x's in Fig. 7 (these include two locations each of change and no change). Figs. 7d and 7e plot representative change maps without and with incorporated edge information respectively. Note the improved performance when incorporating edges in the change detection procedure. In particular, we observe reduced false positives in regions of no change ($\alpha \approx 1$) and stronger correct detection performance in regions of change ($\alpha \approx 0$). This is also reflected in the empirical ROC plot in Fig. 7f.




Figure 7: Empirical ROC studies: (Top panel) (a) and (b) magnitudes of the synthetic test functions; (c) a representative edge detection plot. (Bottom panel) representative change maps (d) without and (e) with edge information. Note the improved performance (reduced false positives in regions of no change ($\alpha \approx 1$) and stronger correct detection in regions with jumps ($\alpha \approx 0$) at potentially the expense of a small loss of resolution. (f) Empirical ROC plot in the bottom right for red $\mathbf{x}$'s in (a), (b), (d), and (e).

Algorithm 1 Structurally informed coherent change detection

Input: Reference (time-1) and mission (time-2) images $f, g \in \mathbb{C}^{N_1 \times N_2}$

Output: Change map $\alpha \in \mathbb{C}^{N_1 \times N_2}$

- 1: Apply the entropy filtering-based low-RCS masking procedure detailed in §2.1.
 - 2: Compute edges in the time-1 and time-2 magnitude images as detailed in §2.4.
 - 3: Apply the multi-polarization change detection statistic of §2.3 - with the input images and computed edge information.
 - 4: Apply the image connectivity-based false positive mitigation scheme from §2.2.
-

4. CONCLUDING REMARKS

This paper introduces a new change detection technique which incorporates *structural information* in the underlying scene. This was accomplished via two primary avenues: *edge detection* to collect descriptors of region and object boundaries, and *image connectivity analysis* to distinguish true changes from false positives. Other descriptors of structure such as texture, topological shape, and morphology could be used in place of or in addition to edges and/or image connectivity, as well as within machine learning frameworks; these would provide potential avenues for future investigation. This work also assumed pre-registered time-1 and time-2 images of the scene.

Extensions of the proposed method to unregistered datasets could provide an appealing challenge. Finally, we have restricted our attention to change detection in urban environments. It would also be instructive to establish the efficacy of the proposed method in other types of datasets in environmental monitoring, agriculture, disaster response and border security.

ACKNOWLEDGMENTS

This work was supported in part by AFOSR grant FA9550-22-1-0411, DMS #1912685, DOE ASCR #DE-ACO5-000R22725, and DOD ONR MURI grant #N00014-20-1-2595. Parts of this work were completed during the final author's sabbatical visits to the Department of Mathematics at Dartmouth College and the School of Mathematical and Statistical Sciences at Arizona State University.

REFERENCES

- [1] Jakowatz, C. V., Wahl, D. E., Eichel, P. H., Ghiglia, D. C., and Thompson, P. A., [*Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*], Springer US, Boston, MA (1996).
- [2] Novak, L. M., "Change detection for multi-polarization multi-pass SAR," in [*Algorithms for Synthetic Aperture Radar Imagery XII*], **5808**, 234–246, SPIE (May 2005).
- [3] Scarborough, S. M., Gorham, L., Minardi, M. J., Majumder, U. K., Judge, M. G., Moore, L., Novak, L., Jaroszewski, S., Spoldi, L., and Pieramico, A., "A challenge problem for SAR change detection and data compression," in [*Algorithms for Synthetic Aperture Radar Imagery XVII*], **7699**, 287–291, SPIE (Apr. 2010).
- [4] Touzi, R., Lopes, A., and Bousquet, P., "A statistical and geometrical edge detector for SAR images," *IEEE Transactions on Geoscience and Remote Sensing* **26**, 764–773 (Nov. 1988).
- [5] Wahl, D. E., Yocky, D. A., Jakowatz, C. V., and Simonson, K. M., "A New Maximum-Likelihood Change Estimator for Two-Pass SAR Coherent Change Detection," *IEEE Transactions on Geoscience and Remote Sensing* **54**, 2460–2469 (Apr. 2016).
- [6] Preiss, M., Gray, D., and Stacy, N., "Detecting scene changes using synthetic aperture Radar interferometry," *IEEE Transactions on Geoscience and Remote Sensing* **44**, 2041–2054 (Aug. 2006).
- [7] Cha, M., Phillips, R. D., Wolfe, P. J., and Richmond, C. D., "Two-Stage Change Detection for Synthetic Aperture Radar," *IEEE Transactions on Geoscience and Remote Sensing* **53**, 6547–6560 (Dec. 2015).
- [8] Inglada, J. and Mercier, G., "A New Statistical Similarity Measure for Change Detection in Multitemporal SAR Images and Its Extension to Multiscale Change Analysis," *IEEE Transactions on Geoscience and Remote Sensing* **45**, 1432–1445 (May 2007).
- [9] Phillips, R. D., "Clean: A false alarm reduction method for SAR CCD," in [*2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*], 1365–1368 (May 2011).
- [10] Newey, M., Barber, J., Benitz, G., and Kogon, S., "False alarm mitigation techniques for SAR CCD," in [*2013 IEEE Radar Conference (RadarCon13)*], 1–6 (Apr. 2013).
- [11] Stojanovic, I. and Novak, L., "Change detection experiments using Gotcha public release SAR data," in [*Algorithms for Synthetic Aperture Radar Imagery XX*], **8746**, 144–153, SPIE (June 2013).
- [12] Gonzalez, R. and Woods, R., [*Digital Image Processing*], Pearson, New York, NY, 4th edition ed. (Mar. 2017).
- [13] Stefan, W., Viswanathan, A., Gelb, A., and Renaut, R., "Sparsity Enforcing Edge Detection Method for Blurred and Noisy Fourier data," *Journal of Scientific Computing* **50**, 536–556 (Sept. 2011).
- [14] Yu, Y. and Acton, S., "Speckle reducing anisotropic diffusion," *IEEE Transactions on Image Processing* **11**, 1260–1270 (Nov. 2002).
- [15] Gelb, A. and Tadmor, E., "Detection of edges in spectral data," *Applied and computational harmonic analysis* **7**(1), 101–135 (1999).
- [16] Song, G., Tucker, G., and Xia, C., "Admissible concentration factors for edge detection from non-uniform Fourier data," *Journal of Scientific Computing* **85**, 1–16 (2020).
- [17] Gelb, A., Green, D., and Viswanathan, A., "SI-CCD: Matlab software for Structurally Informed Coherent Change Detection." <https://github.com/char-md/si-cdd> (2024).