In Algorithms for Synthetic Aperture Radar XX, (Proc. SPIE 8746), SPIE Defense and Security Symposium, April 29 - May 3, 2013

Joint imaging and change detection for robust exploitation in interrupted SAR environments

Joshua N. Ash

Ohio State University, Dept. of Electrical and Computer Engineering, Columbus, Ohio, USA

ABSTRACT

Modern radar systems equipped with agile-beam technology support multiple modes of operation, including, for example, tracking, automated target recognition (ATR), and synthetic aperture radar imaging (SAR). In a multi-mode operating environment, the services compete for radar resources and leave gaps in the coherent collection aperture devoted to SAR imaging. Such gapped collections, referred to as interrupted SAR, typically result in significant image distortion and can substantially degrade subsequent exploitation tasks, such as change detection. In this work we present a new form of exploitation that jointly performs imaging and coherent change detection in interrupted environments. We adopt a Bayesian approach that inherently accommodates different interrupt patterns and compensates for missing data via exploitation of 1) a partially coherent model for reference-pass to mission-pass pixel transitions, and 2) the *a priori* notion that changes between passes are generally sparse and spatially clustered. We employ approximate message passing for computationally efficient Bayesian inference and demonstrate performance on measured and synthetic SAR data. The results demonstrate near optimal (ungapped) performance with pulse loss rates up to $\sim 50\%$ and highlight orders of magnitude reduction in false alarm rates compared to traditional methods.

Keywords: Interrupted SAR, change detection, loopy belief propagation, approximate message passing

1. INTRODUCTION

Modern radar systems equipped with agile-beam technology support multiple modes of operation, including, for example, tracking, automated target recognition (ATR), and synthetic aperture radar imaging (SAR). However, in such multi-mode operating environments, the services compete for radar resources and may therefore leave gaps in the coherent collection aperture devoted to SAR imaging. Such gapped collections with missing pulses—often referred to as interrupted SAR[1]—typically suffer significant image distortion that can substantially degrade subsequent exploitation tasks. While degraded imagery is problematic, the problem is further complicated in applications like change detection relying on multiple images potentially collected under different interrupt patterns and containing different levels of distortion. In this paper, we consider a Bayesian approach to change detection that exploits prior information and joint processing to mitigate the missing data and different collection geometries in interrupted SAR.

For limited aperture sizes, given by a sufficiently large carrier frequency and a sufficiently small azimuth expanse, the k-space data transduced by the radar is approximately rectangular, and separability of the 2D Fourier Transform dictates that the impact of losing data at certain slow-time pulses does not impact range resolution but does degrade the cross-range point-spread-function (PSF) in a well-defined manner given by the Fourier Transform of the interrupt pattern itself[2]

$$PSF(x) = \mathcal{F}(g(k_x)), \tag{1}$$

where $g(k_x) \in \{0, 1\}$ is the "interrupt pattern" equal to 1 if a measurement was made at the azimuth position corresponding to k_x and equal to 0 if no measurement was obtained. In change detection, one of the primary problems is that different interrupt patterns between time-1 and time-2 produce variations in the PSF which may appear to be changes and increase the false alarm rate.

Corresponding author: J. Ash, E-mail: ashj@ece.osu.edu

Early work considering the reconstruction of SAR imagery from missing data focused on "gap-filling," where a preprocessing step focused on filling k-space gaps, and a subsequent step applied standard Fourier inversion to generate an image. Work in this category includes use of the Burg algorithm for autoregressive linear prediction of missing data[1] , and a least-squares method that iteratively performs gap-filling and image estimation.[3] Alternative approaches utilized sparse reconstruction techniques from the field of compressive sensing, effectively exploiting the sparsity of dominant scattering centers in a scene for reconstruction from undersampled data. Examples include algorithms based on ℓ_p regularization[4–6] and greedy approaches.[7,8] The survey paper by Potter et al.[9] contains a comprehensive account of sparse methods applied to radar imaging.

Change detection under different reference and mission interrupt patterns was considered by Stojanovic et al.[2] Here the authors considered two missing-data mitigation strategies for both coherent and incoherent change detection. The first approach utilized an ℓ_1 regularized estimator, basis-pursuit denoising (BPDN), to independently form SAR images for time-1 and time-2, and then applied traditional change-detection algorithms to the resulting images. The second approach formed time-1 and time-2 images by matched-filtering (MF) where only k-spaces samples common to both time-1 and time-2 were utilized. After image formation, traditional change-detection algorithms were applied. We refer to this method as common support (CS) because the data are truncated to common support in k-space before matched filtering. For change detection, this has the advantage of producing equivalent imaging artifacts in both images, thereby lessening the number of false alarms that arise due to different sidelobe structures. The disadvantage is that additional phase history data is discarded in an operating environment that is already data-limited.

In this paper, we adopt an alternative approach to interrupted SAR change detection that does not form images as a preliminary step and does not discard data. Our formulation utilizes a large-scale Bayesian model to incorporate prior information about the scenes and properties of changes, and then *jointly* produces a posterior change map and scene estimates for time-1 and time-2. The proposed framework inherently accounts for different interrupt patterns between the reference and mission passes.

The remainder of this paper is organized as follows. In Section 2 we present the model for joint imaging and change detection, while Section 3 describes the method of large-scale inference over the model. In Section 4 we present results for measured and synthetic data. Conclusions are provided in Section 5.

2. MODEL FOR JOINT IMAGING AND COHERENT CHANGE DETECTION

We consider a spotlight-mode SAR observation model where, after demodulation, filtering, and sampling we obtain the tomographic system model[10,11]

$$y_\Omega = F_\Omega x + n \tag{2}$$

where $x \in \mathbb{C}^N$ denotes a (lexicographically ordered) complex reflectivity image of the scene, and $F_\Omega \in \mathbb{C}^{M\times N}$ is the 2D Fourier transform operator mapping the image into k-space observations. The subscript Ω indicates that the transform is consistent with a set of operating parameters Ω—such as radar bandwidth, slow-time pulse availability, and flight geometry—that determine the support of the k-space samples. The associated noisy phase history data are denoted y_Ω, and $n \in \mathbb{C}^M$ represents observation noise which is assumed to be independent and identically distributed (IID) as zero-mean circular complex Gaussian with variance v_y, $p(n_i) = \mathcal{CN}(n_i; 0, v_y)$, where $\mathcal{CN}(q; \mu, v)$ denotes the circular Gaussian pdf of variable q with mean μ and variance v.

The change detection problem consists of a set of measurements from a reference pass (time-1) and a mission pass (time-2). The measurements may be collected under different conditions at time-1 and time-2, Ω_1 and Ω_2, respectively. Although the subsequent development is general, in this paper we only consider differences resulting from different cross-range pulse support. Respectively, we denote the time-1 and time-2 data as

$$y_1 = F_1 x_1 + n_1 \in \mathbb{C}^{M_1} \quad \text{(reference pass)} \tag{3}$$

$$y_2 = F_2 x_2 + n_2 \in \mathbb{C}^{M_2} \quad \text{(mission pass)} \tag{4}$$

where subscripts indicate time-indexed quantities and the symbol Ω has been omitted for notational brevity. Importantly, the scene image at time-1 (x_1) and time-2 (x_2) may differ, and the measurements y_1 and y_2 may

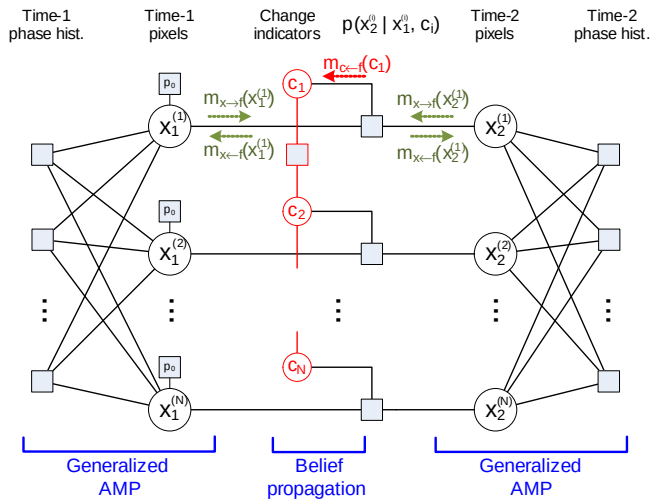

Figure 1. Factor graph model for joint imaging and change detection.

correspond to different k-space sampling locations, which may also be of different sizes; $M_1 \neq M_2$, in general. Let $x_1^{(i)}$ and $x_2^{(i)}$ denote corresponding time-1 and time-2 pixel values at location i, then the mixed-support Bayesian change detection problem seeks to determine the posterior probability that the pixel values differ:

$$p(x_1^{(i)} \neq x_2^{(i)} | y_1, y_2). \tag{5}$$

In the absence of measurements, we assume an IID circular complex Gaussian *a priori* distribution for all pixel values

$$p_0(x_t^{(i)}) = \mathcal{CN}(x_t^{(i)}; 0, v_0), \ t \in \{1, 2\}, \ i \in \{1, \ldots, N\}. \tag{6}$$

We further propose a conditional distribution on the time-2 pixel values

$$p(x_2^{(i)} | x_1^{(i)}, c_i) = c_i p_0(x_2^{(i)}) + (1 - c_i)\mathcal{CN}(x_2^{(i)}; x_1^{(i)}, v_d), \tag{7}$$

where $c_i \in \{0, 1\}$ is an unknown "change indicator" signaling whether the time-2 pixel value $x_2^{(i)}$ differs from the time-1 value $x_1^{(i)}$. The setting $c_i = 1$ indicates a change at pixel i, and $x_2^{(i)}$ is distributed according to the prior $p_0(\,)$. Under the no-change case, $c_i = 0$, $x_2^{(i)}$ is related to $x_1^{(i)}$. We model the no-change distribution of $x_2^{(i)}$ as $\mathcal{CN}(x_2^{(i)}; x_1^{(i)}, v_d)$, indicating that under this condition $x_2^{(i)} = x_1^{(i)} + d_i$, with $d_i \sim \mathcal{CN}(0, v_d)$ being a small *distortion* term allowing for the possibility that $x_2^{(i)}$ and $x_1^{(i)}$ do not have to be precisely identical in order for us to declare no change. This distortion tolerance could be used, for example, to accommodate slight errors in the assumed platform positions with non-isotropic scatterers in the scene. Because (7) considers the phase of $x_1^{(i)}$, the change detection algorithm is a form of coherent change detection. Other conditional distributions may be used to realize alternative forms of change detection, such as incoherent change detection.

We choose a model for the collection $c = \{c_1, \ldots, c_N\}$ of change bits than encodes the notions that 1) changes in the scene are generally sparse, and 2) changes generally exhibit spatial clustering (e.g., many pixels on an inserted target). We denote the marginal distribution of the change bits as $p(c_i = 1) = \rho_1$ and $p(c_i = 0) = \rho_0 = 1 - \rho_1$. Setting $\rho_1 < 0.5$ enforces the notion that change pixels are less likely to occur in the scene than no-change pixels. We use a pairwise Markov random field with 4-connected neighborhood structure to encode spatial correlation among changes. In particular, let

$$p(c_i | c_j) = \begin{cases} \psi_\Delta, & c_i \neq c_j \\ 1 - \psi_\Delta, & c_i = c_j \end{cases} \tag{8}$$

denote the pairwise marginal transition probability between c_i and any of its four neighbors c_j, where $0 \leq \psi_\Delta \leq 1$ denotes the change probability. Values of ψ_Δ less than 0.5 discourage transitions and favor spatial continuity among the change states c. The total distribution of the change bits takes the form

$$p(c) = \frac{1}{Z} \prod_{i=1}^{N} p(c_i) \prod_{j \in \mathcal{N}(i)} p(c_i|c_j), \tag{9}$$

where Z is a normalization constant, and $\mathcal{N}(i)$ denotes the set of neighbors of pixel i.

3. CHANGE DETECTION

Combining the elements of the model described in Section 2 provides a complete probabilistic description of the SAR change detection system. The posterior distribution takes the form

$$p(x_1, x_2, c|y_1, y_2) \propto p(y_1|x_1)\ p(y_2|x_2)\, p(c) \prod_i p(x_2^{(i)}|x_1^{(i)} c_c) p(x_1^{(i)}), \tag{10}$$

with the conditional likelihood functions provided by the Fourier measurement equations (3) and (4), $p(y_t|x_t) = \mathcal{CN}(y_t; F_t x_t, v_y I),\ t \in \{1, 2\}$. The factor graph in Fig. 1 graphically depicts the posterior distribution (10) using circular nodes to represent unknown random variables and square nodes to represent factors of the distribution. Ideally, we would seek the posterior distribution of the change indicators marginalized over the potential image states $p(c|y_1, y_2)$, however this is computationally intractable. Instead, we evaluate the posterior marginal distribution of the indicators on a per-pixel basis $p(c_i|y_1, y_2)$ which is efficiently computed using belief propagation.[12] A byproduct of this approach is that we also obtain distributions over each time-1 and time-2 pixel as well: $p(x_1^{(i)}|y_1, y_2)$ and $p(x_2^{(i)}|y_1, y_2)$.

Belief propagation (BP) is a form of message passing based inference wherein messages are passed along the edges of a factor graph in an attempt to marginalize a product of functions (e.g., (10)). Messages in the graph can intuitively be interpreted as local beliefs about the state of associated variables. Compared to naive marginalization, the BP algorithm exploits the factorization properties of the graph and prescribes a set of rules describing how messages are formulated and utilized. When the graph contains no cycles, it is known that two rounds of message passing yield exact marginal distributions. When BP is applied to graphs with cycles, an application known as loopy BP (LBP), only approximate marginal distributions are generated; however, in many application domains, including for example, error-correcting codes[13] and computer vision,[14] LBP still provides state-of-the art performance. Here, we demonstrate that LBP provides excellent change detection performance for interrupted SAR.

The message passing rules between variables and factors in the graph are summarized below:[15]

Variable x to factor f:

$$m_{x \to f}(\tilde{x}) = \prod_{f' \in N(x) \backslash f} m_{x \leftarrow f'}(\tilde{x}), \tag{11}$$

where $N(x) \backslash f$ denotes the set of all neighboring factors of variable x, except for f.

Factor f to variable x:

$$m_{x \leftarrow f}(\tilde{x}) = \sum_{X \backslash x} f(X) \prod_{x' \in X \backslash x} m_{x' \to f}(x'), \tag{12}$$

where $X = N(f)$ denotes the set of neighboring variables on which the factor f depends, and the summation is over all elements of X other than the destination variable x, which is fixed as $x = \tilde{x}$. In (12), the summation is replaced by integration for continuous variables, such as the complex pixel values. After convergence of message passing, the posterior marginal probability that variable x is equal to $\tilde{x}$ is given by

$$p(x = \tilde{x}|y_1, y_2) \propto \prod_{f \in N(x)} m_{x \leftarrow f}(\tilde{x}). \tag{13}$$

While the BP rules (11), (12) could be applied directly to change-detection factor graph of Fig. 1, we utilize a recently developed form of approximate message passing called generalized approximate message passing (GAMP)[16] to simplify the message passing process for portions of the graph governed by linear transformations. For large linear problems, the GAMP algorithm uses approximations based on the central limit theorem and Taylor series truncation to reduce the complexity of LBP. In particular, GAMP efficiently generates the messages associated with the problem of estimating a vector x from a system where elements $\{z_i\}$ of $z = Ax$ are observed via an "output channel" producing y_i according to a known distribution $p(y_i|z_i)$, and where each component of x is described by an input density $p(x_i)$. Noting the connections to the measurement equations (3), (4), with A given by the Fourier transform and the output channel simply adding Gaussian noise, we observe that the GAMP algorithm can be used to efficiently generate messages over two portions (subgraphs) of the total model corresponding to the measurements and pixel values at time-1 and time-2. Specifically, for time-t, the GAMP algorithm takes as inputs the measurements y_t, the appropriate Fourier operator F_t, and incoming messages $\{m_{x\leftarrow f}(x_t^{(i)})\}$ and produces the outgoing messages $\{m_{x\rightarrow f}(x_t^{(i)})\}$. The GAMP associated messages are shown in green in Fig. 1.

The the time-1 and time-2 output messages from GAMP combine at the $p(x_2^{(i)}|x_1^{(i)}, c_i)$ nodes (see Eq. (7)) to collectively produce a data-dependent local belief (Fig. 1: red $m_{c\leftarrow f}(c_i)$) about the state of c_i. This belief is combined with the change-bit priors $p(c_i)$ and spatial transition probabilities (8) according to the BP rules (11), (12). In LBP, message updates may be scheduled in a number of ways, ranging from simple synchronous updates to fully adaptive asynchronous updates prioritizing nodes bearing more information.[17] In this work, we adopt a simple block-wise round-robin update schedule where we cyclically apply GAMP to update the time-1 image and outgoing messages, apply GAMP to update the time-2 image and outgoing messages, and then apply standard LBP to update the change bits. Analogous to block coordinate descent in traditional function optimization, we let the messages of each block converge before cyclically moving on to the next block. After convergence of all messages, the posterior marginals are computed via (13).

4. NUMERICAL RESULTS

4.1 Gotcha data

Here we evaluate our proposed change detection algorithm using measured X-band complex imagery provided by the GOTCHA change detection dataset.[18] From the dataset, we choose as our time-1 reference scene x_1 a portion of image `FP0120` corresponding to a grassy area where foot traffic subsequently occurs. As the time-2 mission pass scene x_2, we choose the same spatial region of image `FP0124`. Magnitude images of these two scenes are shown in Fig.2(a) and (b). The images appear relatively clutter-like, without any dominant scattering, and the phase and low-amplitude changes induced by the foot traffic are generally undetectable by eye in the magnitude images.

We compare our method to the traditional technique of first forming image estimates, $\hat{x}_1$ for time-1 and $\hat{x}_2$ for time-2, and then performing coherent change detection (CCD) from the resultant complex imagery. For pixel i, we take the sample coherence estimate as the measure of change[10, 19]

$$\gamma_i = \frac{\left|\sum_{j\in N(i)} \hat{x}_1^{(j)} \hat{x}_2^{(j)*}\right|}{\sqrt{\sum_{j\in N(i)} \left|\hat{x}_1^{(j)}\right|^2 \sum_{j\in N(i)} \left|\hat{x}_2^{(j)}\right|^2}}, \tag{14}$$

where $N(i)$ denotes a neighborhood around pixel i. Throughout this paper, we use a 5×5 window centered on pixel i. The CCD metric γ_i detects both magnitude and phase changes between the two images. In the absence of change around pixel i, $\gamma_i = 1$, while changes result in coherence loss giving $0 \leq \gamma_i \leq 1$.

Figures 2(c) through (g) illustrate traditional CCD using (14) applied to five different methods of image reconstruction. Figure 2(c) evaluates γ for the original (full-data) images, x_1 and x_2, and serves as a benchmark to compare the interrupted cases to. The oval-shaped path of the foot traffic is very apparent in this figure. Figures 2(d) through (h) all illustrate change detection performance on interrupted data which was synthetically generated by Fourier transforming the complex imagery and deleting portions of k-space corresponding

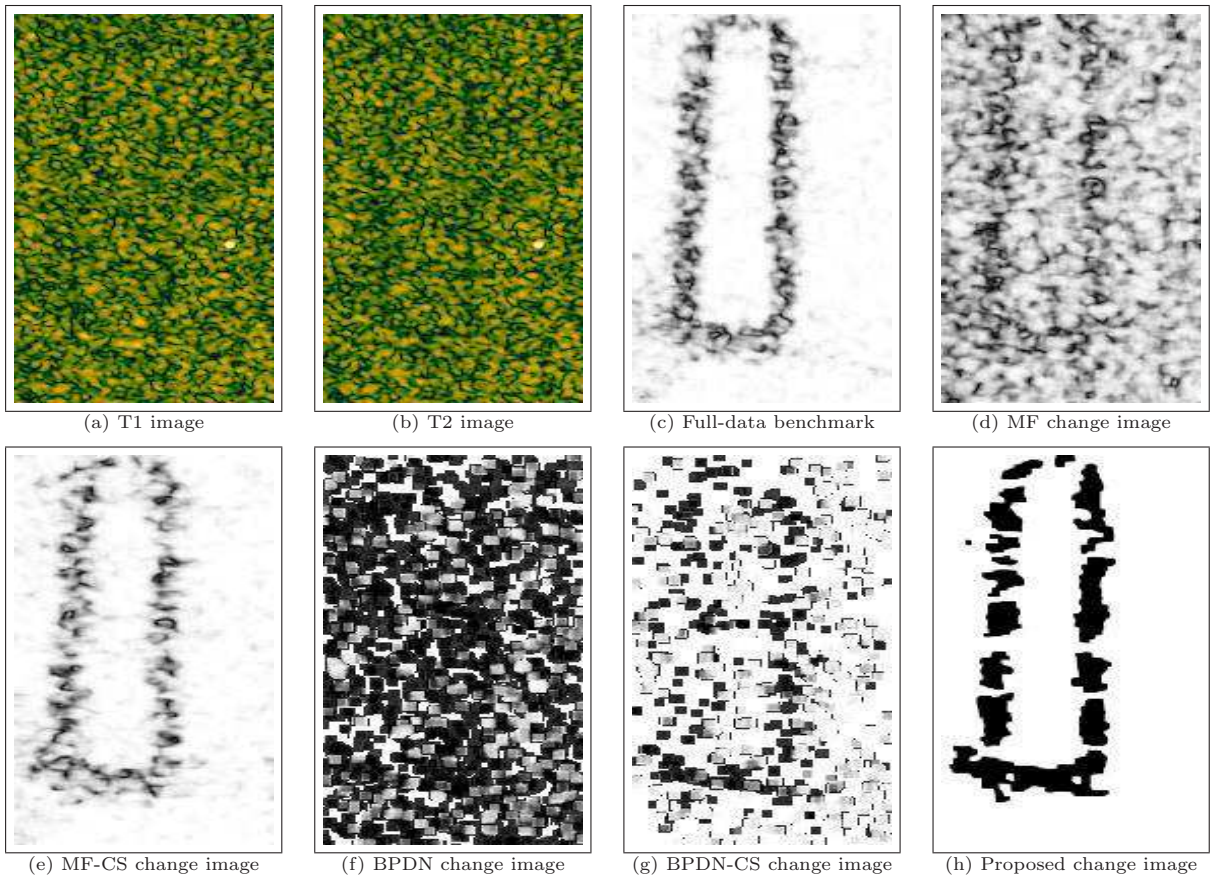

(a) T1 image (b) T2 image (c) Full-data benchmark (d) MF change image

(e) MF-CS change image (f) BPDN change image (g) BPDN-CS change image (h) Proposed change image

Figure 2. *Change-detection results: measured GOTCHA imagery.* (a),(b) time-1 and time-2 magnitude images of scene; (c) full-data (non-interrupted) CCD benchmark image (white=no change, black=change); (d)-(g) traditional change estimates based on sequential imaging + CCD; (h) proposed joint Bayesian method.

to particular cross-range pulses. Starting with Nyquist-sampled Fourier data, the interrupt patterns used were $g_1 = \{+48\%, -5\%, +47\%\}$ for time-1, and $g_2 = \{-14\%, +14\%, -14\%, +14\%, -14\%, +14\%, -16\%\}$ for time-2. Here, $-x\%$ indicates that a contiguous block of pulses spanning $x\%$ of the aperture was deleted, and $+y\%$ indicates that $y\%$ was retained. Therefore, in total, 5% of the time-1 pulses where discarded (from the middle of the aperture), and 58% of the time-2 pulses were deleted in four roughly equal-sized blocks.

Figures 2(d) and (e) consider matched filter (MF) reconstructions. In the first case, the matched filter is applied to all available post-interrupt measurements, and in the second case the time-1 and time-2 measurements are truncated such that they have common k-space support, labeled as MF-CS. The MF case, which has different sidelobe structures for time-1 and time-2, exhibits relatively poor performance with increased false alarms and generally reduced contrast in change identification (compared to the full-data benchmark, Fig. 2(c)), while the MF-CS case compares favorably to the benchmark.

Following Stojanovic et al.,[2] we consider sparsified image reconstructions for use in CCD. In particular, we consider the Bayesian interpretation of basis pursuit denoising (BPDN) to estimate each image as

$$\hat{x}_t = \arg\min_{x_t} \frac{1}{v_y} ||y_t - F_t x_t||_2^2 + \lambda ||x_t||_1, \tag{15}$$

which produces the maximum *a posteriori* (MAP) estimate of x_t given the linear Gaussian measurements (3),(4) and a complex Laplacian prior on x_t, $p_0(x_t) = (\frac{\lambda^2}{2\pi})^N e^{-\lambda||x_t||_1}$. To optimally tune λ to the problem at hand, we use the maximum likelihood estimate of λ obtained from the true scenes x_1, x_2. We note that even though we

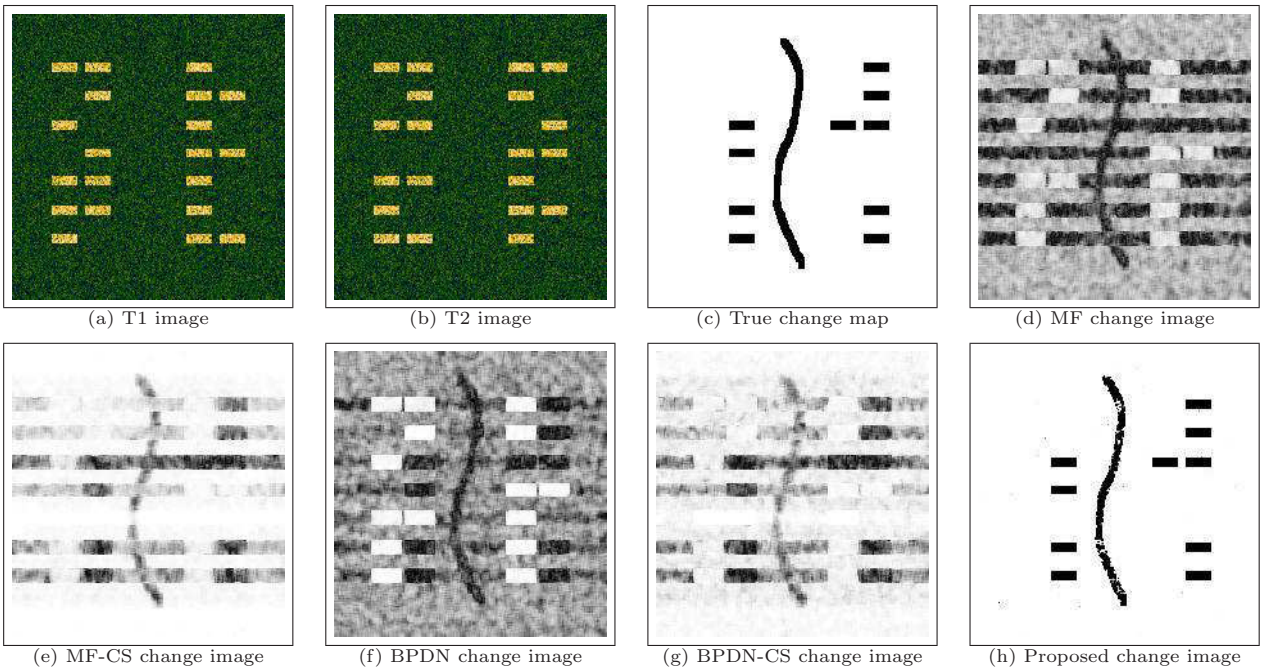

(a) T1 image (b) T2 image (c) True change map (d) MF change image

(e) MF-CS change image (f) BPDN change image (g) BPDN-CS change image (h) Proposed change image

Figure 3. *Change-detection results: synthetic scene.* (a),(b) time-1 and time-2 magnitude images of scene; (c) true change map (white=no change, black=change); (d)-(g) traditional change estimates based on sequential imaging + CCD; (h) proposed joint Bayesian method.

have unrealistically optimized λ to maximize BPDN performance, the results remain relatively poor because this particular scene is primarily noise-like and is a poor match to the Laplacian prior. This not true for the synthetic example considered in the following section. Figures 2(f) and (g) illustrate the CCD performance using BPDN and BPDN with common support (BPDN-CS), respectively. The blockiness of the change estimates results from sparsity in the image reconstructions and the 5×5 windows used in the CCD estimate. BPDN-CS appears slightly better than BPDN, although both estimates compare poorly to the benchmark.

Finally, Fig. 2(h) illustrates the change estimates for the proposed joint Bayesian method, where for each pixel i we have plotted the posterior probability of change $p(c_i = 1|y_1, y_2)$. The estimated change map compares very favorably to the full-data benchmark and even appears as a denoised version of the latter. Similar to the BPDN case, we have optimized the prior to the scene by fitting v_0 in (6) to x_1, x_2. We estimated v_y from the original images and chose $v_d = v_y/100$ for the distortion parameter, and $\rho_1 = 0.05$, $\psi_\Delta = 0.05$ for the change prior and spatial transition probabilities, respectively.

Although the estimated change map appears similar to the full-data benchmark, the dataset lacks the precise ground truth needed to numerically quantify performance. As such, in the next section we consider synthetically generated data for this purpose.

4.2 Synthetic data

In this section we present results for a synthetically generated scene that contains 1) significant magnitude changes corresponding to the movement of cars in a parking lot, and 2) phase-only changes representative of small perturbations, such as the walking path above. The cars and background are clearly discernible in the time-1 and time-2 images shown in Figures 3(a) and (b). The simulated data mimics radar parameters found in the dataset used above:[18] X-band circular SAR, BW=640MHz, 0.2 m range and cross-range resolution. Further, we have used a signal-to-noise ratio of 34 dB and a signal-to-clutter ratio of 18 dB.

Figure 3(c) depicts the locations of induced change and serves as ground truth in subsequent analysis. The rectangles in this figure highlight locations of magnitude changes (insertion or deletion) corresponding to car

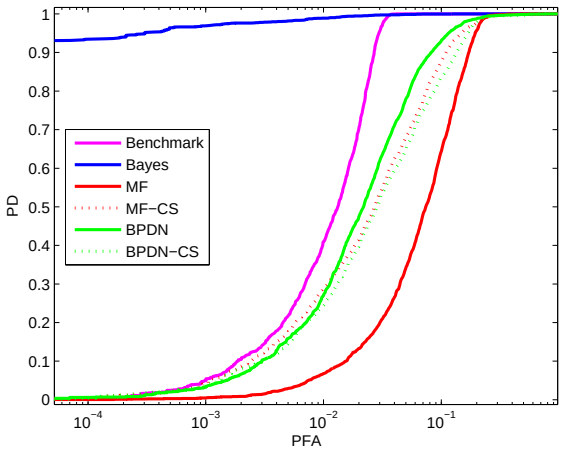

Figure 4. ROC curve for synthetic data with 30% random pulse discard rate. The proposed joint Bayesian method exhibits significantly lower false alarm rates than traditional methods.

movements, and the curved path illustrates locations where the time-1 and time-2 images differ only in phase. The interrupted measurements were generated by starting with Nyquist sampling and randomly discarding 30% of the pulses from the time-1 and time-2 collections. Pulses were discarded independently for time-1 and time-2, and the common support (CS) between the collections corresponded to a 50% discard rate relative to Nyquist.

Figures 3(d) through (h) are analogous to the change estimates considered in the previous section. With some large-magnitude scatterers in the scene, the MF reconstructions now contain substantial blurring in the cross-range dimension resulting in a significant number of false alarms. As with the Gotcha data, the MF-CS estimate is considerably better due to equivalent side-lobe patterns between time-1 and time-2. Compared to the noise-like scene above, the magnitude-sparse Laplacian prior is better suited to the current parking lot scene and BPDN and BPDN-CS perform reasonably well, each detecting most changes but failing to suppress many false alarms. Lastly, we see in Figure 3(h) that the proposed joint Bayesian method does a good job localizing the magnitude and phase-only changes without incurring significant false alarms.

Finally, in Figure 4 we quantify performance of the algorithms by plotting ROC curves corresponding to the detection statistic images of Fig. 3(d) through (h). The advantage of common support when using the matched filter is readily apparent, increasing the probability of detection (PD) from 0.07 to 0.30 when the probability of false alarm (PFA) is equal to 0.01, for example. For PFA=0.01, both BPDN algorithms achieve approximately PD=0.25. The similarity in performance of the two BPDN variants is consistent with phase transition phenomenon from compressive sensing relating the probability of correct recovery to the number of available measurements. Finally, for PFA=0.01, we observe that the proposed Bayesian method achieves PD=0.99. Consistent with Fig. 3, we see that the proposed joint Bayesian method significantly outperforms the other methods, particularly with respect to reduced false alarm rates for a given PD.

5. CONCLUSIONS

We have presented an approach for jointly estimating imagery and coherent changes in an interrupted SAR environment. The approach is Bayesian in nature and allows the algorithm to mitigate missing data through the use of prior information about the time-1 and time-2 scenes as well as anticipated structure in the changes, such as sparsity and spatial clustering. We demonstrated that loopy belief propagation, assisted by generalized approximate message passing, provides an effective and computationally tractable method of Bayesian inference for this problem. Unlike some traditional approaches to the problem which discard data, the Bayesian method fully exploited all available data and demonstrated orders of magnitude reduction in false alarm rates for the examples considered. In future work, we will consider an extension of the method to support incoherent joint Bayesian processing and automated methods of selecting the parameters of the algorithm.

6. ACKNOWLEDGEMENT

This work was supported in part by the Air Force Research Laboratory Sensors Directorate award FA8650-07-D-1220-0006 and by DARPA/ONR award N66001-11-1-4090. We thank Steve Scarborough (AFRL) and Phil Schniter (OSU) for discussions pertaining to this work.

REFERENCES

[1] Salzman, J., Akamine, D., Lefevre, R., and Kirk, J., "Interrupted synthetic aperture radar," *IEEE Aero. Elect. Sys. Magazine* **17**(5), 33–39 (2002).

[2] Stojanovic, I., Karl, W., and Novak, L., "Reconstruction of interrupted SAR imagery for persistent surveillance change detection," in [*Proc. SPIE, Algorithms for Synthetic Aperture Radar Imagery XIX*], **8394** (2012).

[3] Larsson, E., Stoica, P., and Li, J., "Amplitude spectrum estimation for two-dimensional gapped data," *IEEE Trans. Sig. Proc.* **50**(6), 1343–1354 (2002).

[4] Cetin, M. and Karl, W., "Feature-enhanced synthetic aperture radar image formation based on nonquadratic regularization," *IEEE Trans. Image Proc.* **10**, 623–631 (April 2001).

[5] Vu, D., Xu, L., Xue, M., and Li, J., "Nonparametric missing sample spectral analysis and its applications to interrupted SAR," *IEEE J. of Selected Topics in Sig. Proc.* **6**, 1–14 (Feb. 2012).

[6] Austin, C., Ertin, E., and Moses, R., "Sparse signal methods for 3-D radar imaging," *IEEE Journal of Selected Topics in Signal Processing* **5**, 408–423 (June 2011).

[7] Burns, J., Subotic, N., and Pandelis, D., "Adaptive decomposition in electromagnetics," in [*Proc. Int. Symp. Antennas Propag. Soc.*], 1984–1987 (Jul. 1997).

[8] Bhattacharya, S., Blumensath, T., Mulgrew, B., and Davies, M., "Synthetic aperture radar raw data encoding using compressed sensing," in [*IEEE Radar Conference (RADAR)*], 1–5 (2008).

[9] Potter, L., Ertin, E., Parker, J., and Cetin, M., "Sparsity and compressed sensing in radar imaging," *Proceedings of the IEEE* **98**, 1006–1020 (June 2010).

[10] 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 (1996).

[11] Munson, D., O'Brien, J., and Jenkins, W., "A tomographic formulation of spotlight-mode synthetic aperture radar," *Proc. of the IEEE* **71**(8), 917–925 (1983).

[12] Pearl, J., [*Probabilistic Reasoning in Intelligent Systems.*], Morgan Kaufman, San Mateo, CA (1988).

[13] Frey, B. and MacKay, D., "A revolution: Belief propagation in graphs with cycles," in [*Proc. Neural Inform. Process. Syst. Conf.*], 479–485 (1997).

[14] Freeman, W., Pasztor, E., and Carmichael, O., "learning low-level vision," *International Journal of Computer Vision* **40**(1), 25–47 (2000).

[15] Kschischang, F., Frey, B., and Loeliger, H., "Factor graphs and the sum-product algorithm," *IEEE Trans. In* **47**, 498–519 (Feb. 2001).

[16] Rangan, S., "Generalized approximate message passing for estimation with random linear mixing,," in [*Proc. IEEE Int. Symp. Inform. Theory (ISIT)*], 2168–2172 (2011).

[17] Elidan, G., McGraw, I., and Koller, D., "Residual belief propagation: Informed scheduling for asynchronous message passing," in [*Proc. Twenty-Second Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*], 165–173, AUAI Press, Arlington, Virginia (2006).

[18] Scarborough, S., Gorham, L., Minardi, M., Majumder, U., Judge, M., Moore, L., Novak, L., Jaroszewksi, S., Spoldi, L., and Pieramico, A., "A challenge problem for SAR change detection and data compression," in [*Proc. SPIE, Algorithms for Synthetic Aperture Radar Imagery XVII*], **7699** (2010).

[19] Novak, L., "Cohernt change detection for multi-polarization SAR," in [*Proc. ASILOMAR Conf. on Circuits, Systems, and Computers*], (2005).