

Multiresolution Broad Area Search: Monitoring Spatial Characteristics of Gapless Remote Sensing Data

LAURA J. WENDELBERGER^{1,2,*}, JOSH M. GRAY^{3,4}, ALYSON G. WILSON¹,
RASMUS HOUBORG⁵, AND BRIAN J. REICH¹

¹*Department of Statistics, North Carolina State University, Raleigh, NC, USA*

²*Lawrence Livermore National Laboratory, Livermore, CA, USA (current affiliation)*

³*Department of Forestry, North Carolina State University, Raleigh, NC, USA*

⁴*Center for Geospatial Analytics, North Carolina State University, Raleigh, NC, USA*

⁵*Planet Labs PBC, San Francisco, CA, USA*

Abstract

Global earth monitoring aims to identify and characterize land cover change like construction as it occurs. Remote sensing makes it possible to collect large amounts of data in near real-time over vast geographic areas and is becoming available in increasingly fine temporal and spatial resolution. Many methods have been developed for data from a single pixel, but monitoring pixel-wise spectral measurements over time neglects spatial relationships, which become more important as change manifests in a greater number of pixels in higher resolution imagery compared to moderate resolution. Building on our previous robust online Bayesian monitoring (roboBayes) algorithm, we propose monitoring multiresolution signals based on a wavelet decomposition to capture spatial change coherence on several scales to detect change sites. Monitoring only a subset of relevant signals reduces the computational burden. The decomposition relies on gapless data; we use 3 m Planet Fusion Monitoring data. Simulations demonstrate the superiority of the spatial signals in multiresolution roboBayes (MR roboBayes) for detecting subtle changes compared to pixel-wise roboBayes. We use MR roboBayes to detect construction changes in two regions with distinct land cover and seasonal characteristics: Jacksonville, FL (USA) and Dubai (UAE). It achieves site detection with less than two thirds of the monitoring processes required for pixel-wise roboBayes at the same resolution.

Keywords *Bayesian; change detection; monitoring; online; remote sensing; wavelet*

1 Introduction

Satellite signals from a variety of sensors can observe any location in the world, making remote sensing data suitable for global monitoring. As part of a multistage detection and characterization procedure, a broad area search should provide prompt change detection and localization as the basis for further investigation. We chose to analyze Planet Fusion Monitoring data, which imputes missing observations to create gapless, daily, 3 m images. We aim to present a methodology that leverages multiresolution spatiotemporal information enabled by a gapless dataset to monitor a set of spectral signals and propose change sites that can be queried for further refinement and characterization.

*Corresponding author. Email: wendelberger1@llnl.gov.

Pixel-wise algorithms monitor the signal set for a particular location over time, accounting for seasonal and gradual temporal trends present in the data to differentiate sustained change from transient deviations. Pixel-based analysis lends itself to straightforward parallelization, which is important for large-scale analysis. Control charts based on Cumulative Sums (CUSUM; Page, 1954) and Exponentially Weighted Moving Averages (EWMA; Roberts, 1959) flag change based on sustained deviation from the expectation of the monitoring process. Continuous Change Detection and Classification (CCDC; Zhu and Woodcock, 2014) flags land cover change based on successive deviations from a harmonic model, but struggles to monitor short time series and is insensitive to low magnitude changes. Continuous monitoring of Land Disturbance (COLD; Zhu et al., 2020) extends CCDC on several fronts, including input signal analysis and variable selection. The roboBayes algorithm presented in Wendelberger et al. (2021), based on Bayesian Online Changepoint Detection (BOCPD; Adams and MacKay, 2007), is built on the same harmonic model, but its hierarchical Bayesian approach lends strength to its estimation and it can model multiple signals along with their covariance. It views change based on how far back the current model fits, diminishing the need for a long historical stable period and facilitating detection of subtle changes based on cumulative evidence.

As high-resolution imagery proliferates, pixel-based methods are susceptible to over-detection because real, small changes that are normally blurred out become possible to detect (Niemeyer et al., 2008). Further, even when real change is detected, the shape can be irregular, segmented, or full of holes, presenting a change region as multiple smaller ones or failing to capture the entire expanse of the change region (Bontemps et al., 2008). This makes it difficult to translate the results to unified change regions, even with post-processing methods like smoothing and clumping.

The spatial information is underutilized in the above pixel-based temporal approaches to global remote sensing. Pixels tend to change in spatially coherent groups, relative to their surroundings, i.e., a change area is portrayed by the spectral signal change of multiple adjacent pixels in unison. Mei (2010) acknowledges space by combining and monitoring locally aggregated CUSUM statistics. Methods based on linear models broaden to the spatial domain by regressing on a spatial (or spatiotemporal) neighborhood. Both spatial Breaks for Additive Season and Trend (BFAST; Lu et al., 2016) and spatiotemporal Bayesian Online Changepoint Detection with Model Selection (BOCPDMS; Knoblauch and Damoulas, 2018) extend monitoring in this way. The results are sensitive to the neighborhood size, which trades off with computational cost. BOCPDMS introduces a model selection module to adjust the neighborhood size, but does not have a natural mechanism to localize change within a region aside from local applications. A Gaussian Process (GP) model incorporates multi-dimensional information through a correlation function that depends on distance between observation locations. Saatçi et al. (2010) propose the extension of BOCPD to a GP framework. Developments related to GPs for change monitoring include simultaneous inference and change detection (Imbiriba et al., 2019), sparse GPs for faster computation (Gu et al., 2020), multi-scale GP regression (Susiluoto et al., 2020), and covariance monitoring (Kuhn et al., 2014; Horvath et al., 2022).

A spatial signal representation over time occupies the middle ground between a temporal pixel model and a more abstract algorithm. A wavelet decomposition represents a signal as a combination of spatially localized orthogonal basis functions (Daubechies, 1988, 1990). These signals incorporate information from a neighborhood of pixels, amplifying spatially coherent change. A multi-resolution analysis uses Discrete Wavelet Transform (DWT) to represent a signal in terms of a series of spatially resolved basis functions. Efficient decomposition algorithms (Mallat, 1989; Shensa, 1992) make multiresolution analysis (MRA) computationally feasible.

Wavelet decomposition over time (Killick et al., 2013), while useful for detecting second-order structural change, does not support a sequential updating scheme. Instead, a two-dimensional decomposition over space is considered. Rinoshika and Rinoshika (2020) illustrates MRA for multiple dimensions. Wavelets transformations are used in image processing to denoise data (Mallat, 1989; Donoho, 1995; Wang et al., 2010) and segment images (Xu et al., 2017; Huang et al., 2018). Differences in successive segmented images, or similarly, when a difference image itself is segmented (Celik and Ma, 2011; Li et al., 2017), is one way to detect change. These methods focus on pairs of images, but with a broad monitoring algorithm, it is necessary to update results at many sequential time steps and to consider seasonal structure to avoid inadvertently flagging regular seasonal variation as change. Image differencing is particularly susceptible to ephemeral real change and corrupted data that may present as change.

Wavelet representations require a complete grid of data. With traditional imagery use, cloud masks flag unusable data, leaving gaps in images and precluding the use of wavelet decomposition on a broad monitoring scale without consideration of imputation. Planet Fusion Monitoring data supplements PlanetScope images with publicly available satellite data from Sentinel-2 and Landsat-8/9 using the Cubesat enabled Spatio-Temporal Enhancement Method (CESTEM; Houborg and McCabe, 2018a; Houborg and McCabe, 2018b). The end result is a radiometrically and geometrically harmonized, cloud-masked, and gap-filled data cube with daily temporal and 3 m spatial resolution. In addition, Quality Assurance (QA) data characterizes each pixel at each time, making it possible to differentiate pixels based on the amount of observed data that goes into the response value. With this substantial increase in temporal and spatial resolution, Planet Fusion contains information about disturbances at a very fine resolution. Attempting to monitor large regions over long periods using existing, pixel-based methods would scale up the computation time with the increase in resolution; i.e., monitoring a scene using 3 m data would take 100 times longer than the 30 m equivalent. The Planet Fusion Monitoring product necessitates a change in analysis methods to take advantage of the information that it contains while remaining computationally feasible.

In this paper, we use Planet Fusion Monitoring data to scan regions for construction change from 2018 through early 2022. Balafas et al. (2018) scans for earthquake damage change by modeling wavelet coefficients at each time as a GP realization, then identifying model parameter change. Similarly, but without the need to fit a GP, we propose using a wavelet representation by decomposing incoming Planet Fusion images with DWT and then monitoring the coefficients of the multiresolution wavelet basis functions to detect construction changes in a process called Multi-Resolution roboBayes monitoring (MR roboBayes). We judiciously select which coefficients to monitor, which reduces the number of monitoring processes. We evaluate MR roboBayes on simulated datasets to demonstrate its ability to detect subtly changed cohesive sites in a timely manner better than pixel-based roboBayes. We apply MR roboBayes to two different regions: Jacksonville, FL (USA) and Dubai (UAE), which have very different landscapes and evaluate performance to examine how MR roboBayes generalizes across land cover types. MR roboBayes is able to achieve comparable detection with substantially fewer monitoring processes.

The remainder of the paper proceeds as follows. The motivating data are described in Section 2. Section 3 reviews wavelet representation and introduces MR roboBayes. Computation is discussed in Section 4. MR roboBayes is evaluated using a simulation study in Section 5 and applied to the test regions in Section 6. Section 7 provides discussion and conclusions.




Figure 1: A $6.144 \times 6.144 \text{ km}^2$ region in Jacksonville, FL at (left) the beginning and (right) end of the study period. Construction sites beginning before 2018, but ending after are annotated in green. Construction sites beginning after 2018 are annotated in yellow.

2 Construction in Jacksonville, FL (USA) and Dubai (UAE)

Automated monitoring facilitates surveillance of large geographical areas for new construction. In this case, we are interested in identifying major construction projects, i.e., airports, industrial areas, etc., rather than single homes. Planet Fusion Monitoring data has 3 m spatial resolution and daily temporal resolution (Planet Fusion Team, 2022). The available bands are the three visible bands (red, green, and blue) as well as a near infrared band (NIR). Visible bands are highly correlated, so a single band is chosen to represent brightness. For Jacksonville, we monitor brightness via the blue band along with a complementary signal, normalized difference vegetation index (NDVI; Tucker, 1979). NDVI is a unitless red and near-infrared based index commonly used to measure vegetation properties that discriminates well between vegetation and cleared land. In Dubai, only the red band is monitored since it is less susceptible to atmospheric effects than blue, and NDVI is omitted since vegetation is not prevalent. QA values quantify the “percent synthetic” content for each pixel, i.e., the proportion of the pixel response created by the gap-filling algorithm. Entirely missing pixels have QA of 100% while completely observed pixels have QA of 0%.

We examine data for two distinct regions to explore generalizability of the method. The first region is a $6.144 \times 6.144 \text{ km}^2$ area in Jacksonville, FL (USA). The second is a $12.288 \times 12.288 \text{ km}^2$ region south of Dubai (UAE). Annotations with spatial location as well as start/end dates for areas undergoing heavy construction are provided by the Space-Based Machine Automated Recognition Technique (SMART) program for both regions. Some of the annotated construction sites begin before 2018-01-01, when Planet Fusion becomes available. They are included in the annotations because substantial change that can be detected occurs after the initial site clearing and may be present in the data during the study period.

Data in Jacksonville is typically seasonal because of its natural vegetation as shown in Figure 1; both NDVI and brightness are sensitive to construction changes. In this region, some small sites can be hard to differentiate from other false local flags.



Figure 2: A $12.288 \times 12.288 \text{ km}^2$ region in Dubai, UAE at (left) the beginning and (right) end of the study period. Construction sites beginning before 2018, but ending after are annotated in green. Construction sites beginning after 2018 are annotated in yellow.

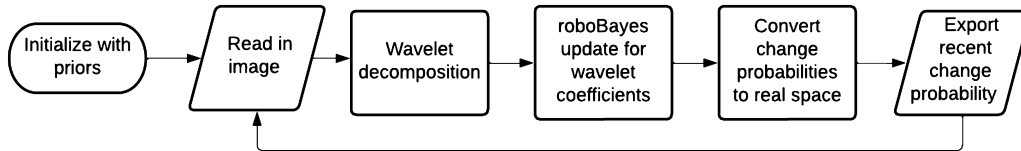


Figure 3: Workflow for MR roboBayes.

In the Dubai region shown in Figure 2, it can be difficult to differentiate land which has undergone construction from the bright surrounding sand. In addition to the subtle nature of the changes, this region contains many relatively large sites because it is an industrial area. Shifting dunes present a unique challenge in this region; they are a sustained change over a large region (much like a building), but not one that we are interested in sending on for further scrutiny.

3 Statistical Methods

A summary of the steps involved in the analysis are visualized in Figure 3. As each multidimensional image signal is observed, a wavelet transformation is applied to create independent spatial signals. We then monitor the spatial signals for change using the roboBayes algorithm. Finally, the spatial signals are recombined to quantify change in real space. The rest of the section details the steps in the algorithm.

Denote the data at spatial location $\mathbf{s} = [s_1, s_2]$, where s_1 and s_2 are Cartesian coordinates, of spectral signal $z = 1, \dots, d$ at time t as $Y_{zt}(\mathbf{s})$. Here $d = 2$ for the two spectral signals: brightness and NDVI. It is assumed that data are observed on a $N \times N$ grid of $n = N^2$ unique locations with $\mathbf{s} \in \{1, \dots, N\}^2$, where $N = 2^J$. The entire image is denoted $Y_{zt} = \{Y_{zt}(\mathbf{s}_1), \dots, Y_{zt}(\mathbf{s}_n)\}^T$. We

use a basis representation,

$$Y_{zt}(\mathbf{s}) = \sum_{i=1}^n y_{zti} B_i(\mathbf{s}), \quad (1)$$

where $B_i(\mathbf{s})$ are the 2D wavelet basis images evaluated at \mathbf{s} with weights y_{zti} . We propose to use wavelets to transform spatial images Y_{zt} to independent signals $y_{zt} = \{y_{zt1}, \dots, y_{ztn}\}$ and perform change detection on the wavelet coefficients, monitoring multiple resolutions simultaneously. Wavelet coefficients both encode spatial information and can easily be monitored in parallel.

3.1 Wavelet Basis Representation

Instead of analyzing individual pixels, it is desirable to interpret the image in a way that recognizes spatial patterns. Wavelets are localized in space, so they can summarize the behavior of groups of pixels. First, we introduce wavelet representation in one dimension, then extend to two dimensions and define the representation of Y_{zt} using a wavelet basis. Then we apply roboBayes to monitor wavelet coefficients and finally define change sites in real space based on posterior change probabilities.

1-D Multi-Resolution Wavelet Basis A wavelet basis is determined by the mother wavelet $\psi(s)$ and father wavelet $\phi(s)$ (also called a scaling function) that can be evaluated for a 1-D location s . We choose the Haar wavelet basis, whose 1D mother function $\psi(s)$ and scaling function $\phi(s)$ are:

$$\psi(s) = \begin{cases} 1 & \text{if } 1 \leq s < \frac{N}{2} + 1, \\ -1 & \text{if } \frac{N}{2} + 1 \leq s < N + 1, \\ 0 & \text{otherwise,} \end{cases} \quad \phi(s) = \begin{cases} 1 & \text{if } 0 \leq s < N + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Haar wavelets can represent abrupt changes, like the edge of a building, and lend themselves to convenient interpretation as each wavelet basis function is active on only a subset of the signal, and is 0 elsewhere.

One can generate the entire family of basis functions by shifting and scaling these generating functions. For any location $s \in \mathbb{R}$, scale $j \in \{1, \dots, J\}$, and shift $k \in \{1, \dots, K_j\}$, where $K_j = N2^{-j}$:

$$\begin{aligned} \psi_{jk}(s) &= 2^{-j/2} \psi(2^{-j}s - k), \\ \phi_{jk}(s) &= 2^{-j/2} \phi(2^{-j}s - k). \end{aligned} \quad (3)$$

A multiresolution discrete wavelet transform represents a signal as a combination of wavelet basis functions of varying resolutions j . For a complete orthonormal basis $\{\phi_{J1}, \psi_{jk} : j \in \{1, \dots, J\}, k \in \{1, \dots, K_j\}\}$, an arbitrary signal $Y(s)$ can be represented as

$$Y(s) = c\phi_{J1}(s) + \sum_{j=1}^J \sum_{k=1}^{K_j} d_{jk} \psi_{jk}(s). \quad (4)$$

The coefficients c and d_{jk} can be computed for $Y(s)$ using the discrete wavelet transform (DWT; Mallat, 1989). Mallat's pyramid algorithm makes the DWT fast, on the order $N \log N$ (Lange and Lange, 2017), and the original signal can be recovered exactly from the unique collection of wavelet basis coefficients.

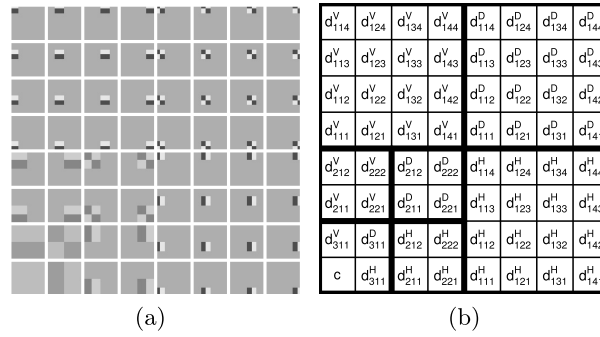


Figure 4: 2D separable Haar basis functions (a) and their coefficients (b).

2D Multiresolution Analysis To extend the wavelet representation to two dimensions, the scaling and wavelet basis functions are applied across both dimensions to produce a 2D scaling function Φ and three mother wavelets in the horizontal (Ψ^H), vertical (Ψ^V), and diagonal (Ψ^D) directions:

$$\begin{aligned}
 \Phi(\mathbf{s}) &= \Phi(s_1, s_2) = \phi(s_1)\phi(s_2), \\
 \Psi^H(\mathbf{s}) &= \Psi^H(s_1, s_2) = \phi(s_1)\psi(s_2), \\
 \Psi^V(\mathbf{s}) &= \Psi^V(s_1, s_2) = \psi(s_1)\phi(s_2), \\
 \Psi^D(\mathbf{s}) &= \Psi^D(s_1, s_2) = \psi(s_1)\psi(s_2).
 \end{aligned} \tag{5}$$

The shifting and scaling of these functions follows from the 1D equations. For a scale of j and shift of $\mathbf{k} = [k_1, k_2]$,

$$\Phi_{jk}(\mathbf{s}) = 2^{-j}\phi(2^{-j}s_1 - k_1)\phi(2^{-j}s_2 - k_2), \tag{6}$$

and similarly for the mother wavelets. An arbitrary 2D spatial process $Y(\mathbf{s})$ can be represented as

$$Y(\mathbf{s}) = c\Phi(\mathbf{s}) + \sum_{j=1}^J \sum_{k_1}^{K_j} \sum_{k_2}^{K_j} \sum_{m \in \{H, V, D\}} d_{jk}^m \Psi_{jk}^m(\mathbf{s}), \tag{7}$$

where Φ is the scaling function, Ψ_{jk}^m is the transformed mother wavelet, c is the scaling coefficient, and d_{jk}^m are the wavelet coefficients such that j indexes the decomposition level, \mathbf{k} indexes the shift in location, and m indexes the direction type (horizontal, vertical, and diagonal). The 2D separable Haar mother and scaling functions are visualized in Figure 4 and their expressions are available in the supplementary material.

To connect with (1) and streamline notation, let $\mathbf{y} = (y_1, \dots, y_n)^T$ be the vector of wavelet coefficients c and d_{jk}^m and let $B_1(\mathbf{s}), \dots, B_n(\mathbf{s})$ be the collection of corresponding scaling (Φ) and mother (Ψ_{jk}^m) functions. The n terms inherit support and resolution inferred as

$$\begin{aligned}
 A_s &= \left\{ i; \left\lceil \frac{s}{K_j} \right\rceil = \mathbf{k} \right\}, \\
 Q_i &= \left\{ \mathbf{s}; \left\lceil \frac{s}{K_j} \right\rceil = \mathbf{k} \right\},
 \end{aligned} \tag{8}$$

where each i indexes a wavelet component that is uniquely defined by its scale, shift, and direction $\{j, \mathbf{k}, m\}$. A_s contains the indices of all basis functions that contribute to pixel \mathbf{s} and Q_i identifies all real space pixels that contribute to the wavelet basis function with coefficient y_i . The equivalence $\left\lceil \frac{s}{K_j} \right\rceil = \mathbf{k}$ holds when the location \mathbf{s} has a nonzero value in the wavelet basis component.

Image Representation The signal $Y_{zt}(\mathbf{s})$ can be decomposed based on (7) using the DWT for each image to obtain the wavelet basis function coefficients y_{zti} . Then $\mathbf{y}_{ti} = [y_{1ti}, \dots, y_{dti}]^T$ is the vector time series for basis function i . We will monitor the multivariate time series $\{y_{1i}, y_{2i}, \dots\}$, which is assumed independent (over i) using the Bayesian change point detection algorithm defined below. Change monitoring in wavelet space takes advantage of the spatial coherence of changed areas, where pixel-by-pixel signal monitoring does not.

3.2 Online Changepoint Detection

Stable Period Model To assess different components for change, we monitor \mathbf{y}_{ti} . Suppose that, given data up to time t_{max} , groups of data belong to consecutive states $\tilde{g}_{ti} = h \in \{1, 2, \dots\}$, as in a Product Partition Model (Barry and Hartigan, 1993). The multi-dimensional data can be modeled as

$$\mathbf{y}_{ti} | \boldsymbol{\beta}_{hi}, \boldsymbol{\Sigma}_{hi}, \tilde{g}_{ti} = h \sim \text{Normal}(\mathbf{x}_t^T \boldsymbol{\beta}_{hi}, \boldsymbol{\Sigma}_{hi}), \quad (9)$$

where state-specific parameters $\boldsymbol{\beta}_{hi}$ and $\boldsymbol{\Sigma}_{hi}$ are $k \times d$ and $d \times d$, respectively. To capture the seasonal (third-order harmonic) and temporal (linear slope) trends, the covariates are taken to be

$$\mathbf{x}_t^T = \left[1, \sin\left(\frac{2\pi t}{1 * 365}\right), \cos\left(\frac{2\pi t}{1 * 365}\right), \dots, \sin\left(\frac{2\pi t}{3 * 365}\right), \cos\left(\frac{2\pi t}{3 * 365}\right), t \right].$$

The unknown state parameters $\boldsymbol{\theta}_{hi} = [\boldsymbol{\beta}_{hi}, \boldsymbol{\Sigma}_{hi}]$ have prior distributions

$$\begin{aligned} \boldsymbol{\beta}_{hi} | \boldsymbol{\Sigma}_{hi} &\sim \text{Matrix Normal}_{d,k}(\mathbf{B}_0, \boldsymbol{\Lambda}_0^{-1}, \boldsymbol{\Sigma}_h), \\ \boldsymbol{\Sigma}_{hi} &\sim \text{Inverse Wishart}(\mathbf{V}_0, \nu_0), \end{aligned} \quad (10)$$

with fixed hyperparameters $\boldsymbol{\eta} = [\mathbf{B}_0, \boldsymbol{\Lambda}_0, \mathbf{V}_0, \nu_0]$. The dimensions of the hyperparameters are: \mathbf{B}_0 is $k \times d$, $\boldsymbol{\Lambda}_0$ is $k \times k$, \mathbf{V}_0 is $d \times d$, and ν_0 is a scalar.

Bayesian Online Changepoint Detection (BOCPD) The number of models for every possible segmentation of data into estimated states with an unknown number of changepoints grows quickly as new data becomes available. Instead, Adams and MacKay (2007) propose BOCPD to retain only data relevant for the most recent changepoint. Run length at time t is defined as the number of time points that belong to the state \tilde{g}_{ti} up to time t , i.e., the number of time steps since the most recent change. We denote the run length for wavelet component i as \tilde{r}_{ti} :

$$\tilde{r}_{ti} = r \text{ if } \tilde{g}_{ti} = \tilde{g}_{ti} - r + 1 = \tilde{g}_{(t-r),i} + 1. \quad (11)$$

Then for each possible run length $\tilde{r}_{ti} = r$, the state parameters for the final state are $\boldsymbol{\beta}_{ti}^{(r)}$ and $\boldsymbol{\Sigma}_{ti}^{(r)}$. We discuss converting run length from wavelet space to real space in Section 3.3.

The run length prior is

$$P(\tilde{r}_{ti} | \tilde{r}_{(t-1),i}) = \begin{cases} 1 - \lambda & \text{if } \tilde{r}_{ti} = \tilde{r}_{(t-1),i} + 1, \\ \lambda & \text{if } \tilde{r}_{ti} = 0. \end{cases} \quad (12)$$

Then the posterior run length distribution is

$$f(\tilde{r}_{ti} | \mathbf{y}_{(1:t)i}) = \frac{f(\tilde{r}_{ti}, \mathbf{y}_{(1:t)i})}{f(\mathbf{y}_{(1:t)i})}, \quad (13)$$

where the joint distribution undergoes recursive updates based on the predictive probability for the current observation, the prior run length distribution, and the preceding joint distribution:

$$f(\mathbf{y}_{(1:t),i}, \tilde{r}_{ti}) = \sum_{\tilde{r}_{(t-1),i}} f(\mathbf{y}_{ti} | \mathbf{y}_{(t-\tilde{r}_{ti}): (t-1),i}) f(\tilde{r}_{ti} | \tilde{r}_{t-1}) f(\mathbf{y}_{1:(t-1),i}, \mathbf{r}_{(t-1),i}). \quad (14)$$

roboBayes Occasionally in monitoring data, a single anomalous observation occurs, but the data before and after the anomaly are really in the same state. For example, a cloud, shadow, or even vehicle may manifest as a signal different from the true land cover. To avoid flagging an erroneous change triggered by these transient events, roboBayes introduces an outlier state $\tilde{g}_{ti} = 0$ where the prior outlier probability is $f(\tilde{g}_{ti} = 0) = p_a$ and calculates the probability distribution of the most recent outlier $o_{ti} \in \mathcal{O}_t = \{\emptyset, 1, \dots, r-1\}$ where $f(o_{ti} = s) = p_a$ with $o_{ti} = a$ if $\tilde{g}_{(a-1),i} = \tilde{g}_{(a+1),i}$ and $\tilde{g}_{ai} = g_0$. Along with the run length and outlier prior distributions, the full hierarchical model specification is

$$\mathbf{y}_{ti} | \boldsymbol{\beta}_{ti}^{(r)}, \boldsymbol{\Sigma}_{ti}^{(r)}, o_{ti} \neq \emptyset \sim \text{Normal}(\mathbf{x}_t^T \boldsymbol{\beta}_{ti}^{(r)}, \boldsymbol{\Sigma}_{ti}^{(r)}), \quad (15)$$

$$\mathbf{y}_{ti} | o_{ti} = \emptyset \sim \text{Normal}(\boldsymbol{\mu}_0, \boldsymbol{\Omega}_0), \quad (16)$$

$$\boldsymbol{\beta}_{ti}^{(r)} | \tilde{r}_{ti}, o_{ti} \sim \text{Matrix Normal}(\mathbf{B}_0, \boldsymbol{\Lambda}_0^{-1}, \boldsymbol{\Sigma}_{ti}^{(r)}), \quad (17)$$

$$\boldsymbol{\Sigma}_{ti}^{(r)} \sim \text{Inverse Wishart}(\mathbf{V}_0, \nu_0). \quad (18)$$

Then, with fixed parameters $\boldsymbol{\eta}$, $\boldsymbol{\mu}_0$, and $\boldsymbol{\Omega}_0$, fixed probabilities λ and p_a , data $\mathbf{y}_{1i}, \dots, \mathbf{y}_{ti}$ and $\mathbf{x}_1, \dots, \mathbf{x}_t$, it is possible to estimate the posterior distributions for \tilde{r}_{ti} and o_{ti} . An approximation removes outliers based on a threshold as the algorithm runs.

It is recommended to generate level (j) and type (m) specific prior hyperparameters $\boldsymbol{\eta}$ based on linear models of the wavelet coefficients from past observations of the region of interest during a stable period. Similarly, the values governing the outlier distribution can be chosen based on expert opinion or known outlier behavior. The run length increment probability λ can be chosen based on the expected number of changes in space and time and outlier probability p_a can be chosen based on expected number of outliers.

3.3 Change Detection Based on Multiresolution Run Length Distributions

Next, we monitor the response using the roboBayes algorithm, which requires the covariates $\mathbf{X}_{1:t} = [\mathbf{x}_1^T, \dots, \mathbf{x}_t^T]$ and the $t_{\max} \times d$ dimensional response \mathbf{y}_i for each wavelet coefficient. Then, using level/type specific prior information and a recursively updating monitoring algorithm, it produces a posterior distribution for the run length.

Denote the real space run length at time t and pixel location s as r_{ts} . The run length distribution of pixel s will be defined below as a function of the run length distributions of all coefficients \tilde{r}_{ti} in A_s from (8). Since Haar wavelets are sparse, only a small fraction of them affect a particular pixel location. For continuous bases like Daubechies wavelets, thresholding can be applied to approximate the spatial range of each basis function. Suppose we define the probability of recent change for a pixel as the probability that the contribution of any of its wavelet components have changed in the last L observations:

$$P(r_{ts} \leq L) = 1 - P(r_{ts} > L) = 1 - \prod_{i \in A_s} \{1 - P(\tilde{r}_{ti} \leq L)\}. \quad (19)$$

Without data, $P(\tilde{r}_{ti} \leq L) = (1 - \lambda)^L$ based on the prior run length distribution. Conditioning (19) on the data, the posterior run length distribution (13) determines the right-hand-side of the

equation. This definition is vulnerable to assigning high change probability to real space pixels that have not changed, but share the same wavelet domain as pixels that have. Redundant detection, or high probability that multiple components have changed, can confirm or even localize real change. For example, change probability in a coarse resolution alone proposes too large of an area and in fine resolutions is susceptible to creating many false positive sites. Probable change in both is good evidence of real change in the local area and is robust to scattered false positives.

If c is the number of components necessary to declare change in real space, change conditioned on c can be defined as the probability that the contribution of at least c components have changed in the last L observations. Then (19) is the case where $c = 1$. Extending the concept in (19) to $c = 2$ components:

$$P(r_{ts} \leq L | c = 2) = 1 - \prod_{i \in A_s} P(\tilde{r}_{ti} \leq L) - \sum_{i \in A_s} \left\{ P(\tilde{r}_{ti} \leq L) \prod_{i' \in A_s \setminus i} [1 - P(\tilde{r}_{ti'} \leq L)] \right\}, \quad (20)$$

where $A \setminus i$ is the set A_s , excluding i . The number c should be chosen such that real change is probable in at least c components and less than c probable components comprise superfluous site area/FPs. To create change pixels, we threshold the pixel-wise probabilities of recent change:

$$C_t(\mathbf{s}) = I\{P(r_{st} \leq L | c) > \tau\}. \quad (21)$$

When combining wavelet space run lengths to form real space run lengths, the number of components c and change threshold τ must be chosen *a priori*. The change threshold applies to the real-space probabilities.

While defining and calculating change probability for each real space pixel is convenient for interpretation, instead, we could threshold probabilities in wavelet space. From this, a count of the total number of changed components affecting a pixel is recorded and then a second threshold is applied on the count of changed components. Let $\tilde{C}_{ti} = I(P(\tilde{r}_{ti} \leq L) \geq \tau_p)$ be an indicator of a change for coefficient i at time t . Then a pixel \mathbf{s} is declared to have changed if at least τ_c coefficients in its support have changed, i.e.,

$$C_t(\mathbf{s}) = I\left(\sum_{i \in A_s} \tilde{C}_{ti} \geq \tau_c\right), \quad (22)$$

where $\tau_p \in [0, 1]$ is the probability threshold and $\tau_c \in \mathbb{Z}$ is the count threshold (both chosen *a priori*). To transform changed pixels in real space to change sites, we group pixel-wise detections in C_t based on connected pixels to produce a set of detected change polygons $\mathcal{P}_t = \{P_{t1}, \dots, P_{tn}\}$ at each time point t .

3.4 Joint Run Length

Intuitively, given that change is spatially coherent, it is desirable that the prior change probability for two pixels in close proximity is similar. We derive the induced joint run length distribution of two pixels \mathbf{s} and \mathbf{s}' and show under what circumstances the recombination method in (19) supports this intuition. Under the model, the prior probability of remaining in the same state is $P(r_{ts} > L) = p_0$, where $p_0 = (1 - \lambda)^L$ as the prior probability of no recent change. The definition

of change in wavelet space induces a joint distribution among the pixels in real space. The prior joint probability of no change for pixels \mathbf{s} and \mathbf{s}' is

$$P(r_{st} > L, r_{s't} > L | c = 1) = p_0^{|A_s|+|A_{s'}|-|A_s \cap A_{s'}|}. \quad (23)$$

The intersection term can be written

$$|A_s \cap A_{s'}| = 1 + 3 \sum_{j=1}^J I\left(\left\lceil \frac{s_1}{2^j} \right\rceil = \left\lceil \frac{s_1'}{2^j} \right\rceil\right) I\left(\left\lceil \frac{s_2}{2^j} \right\rceil = \left\lceil \frac{s_2'}{2^j} \right\rceil\right),$$

and

$$|A_s| = 3J + 1.$$

See the supplementary material for details. In the trivial case $\mathbf{s} = \mathbf{s}'$, we recover the prior probability $p_0^{|A_s|}$. In the case where $[s_1, s_2]$ and $[s_1', s_2']$ share contributions from all wavelets, the induced prior probability is the same. However, if two pixels share only the coarsest resolution wavelet, the probability reduces to $p_0^{3(2J-1)}$. The induced joint prior probability of two pixels remaining in the same state is based on how many wavelet resolutions both pixels are present in. In most cases, two spatially close pixels will contribute to multiple wavelet resolutions, so they are more likely to change (or not change) as a pair. On the other hand, pixels that are further apart belong to few of the same resolutions, so their change behavior is less related to each other.

4 Computational Details

Initial prior hyperparameters are estimated for each decomposition level based on initial regression fits of wavelet coefficient data for the region of interest. The roboBayes algorithm was implemented using the R package `roboBayes`, which can be found on CRAN. Only wavelet observations whose average synthetic percentage $\tilde{q}_{ti} = \frac{1}{|Q_i|} \sum_{Q_i} q_{ts}$ is less than 20% are monitored (based on contributing real space pixels with synthetic percentage q_{ts}). Similarly, in a pixel-wise analysis, only observations with $q_{ts} \leq 20\%$ are monitored. The bulk of computation occurs at the highest resolution wavelets, where quality is still very spatially localized. Still, there may be more data removed during the wavelet analysis than in the pixel-wise analysis, which would further increase the computational advantage at the cost of using all the information in the data. The probability of recent change in the most recent $L = 15$ time points (after a 3 time point delay to allow time for outlier recognition) is recorded. The choice of L is mainly to quantify what qualifies as “recent”. Too short of an L may fail to flag real change; there may not be enough evidence to clear the threshold. As L increases, it summarizes temporal information more coarsely; flagging change in the past year, but with no more precise indication of when, might not be very useful information in practice. The best L is one that is just long enough to detect the change, but not longer. A pilot study of latency of change detection (time points between the true change and when roboBayes might detect it) guided the choice of L . To maintain computational efficiency, only those run lengths less than the truncation limit of 35 or with probability greater than $1e - 4$ are retained.

5 Simulation

Here we apply the methods described in Section 3. The objectives of the simulation are to compare detection metrics for classic roboBayes applied to single pixels versus MR roboBayes.




Figure 5: Simulated data with five regions (cyan boxes) that change at time points indicated by the number next to each region.

5.1 Data Generation

A $256 \times 256 \times 2$ array response $Y_t = \{Y_{t1}, Y_{t2}\}$ (a 2-dimensional signal measured at $\mathbf{s} \in \{1, \dots, N = 256\}^2$ pixel locations) with a mean process $\mu = \{\mu_1, \mu_2\}$ was simulated at 80 time steps. Each mean process is a realization of a GP with mean zero and with Matérn covariance with range 1 and differentiability parameter 0.1. Autocorrelation $\rho = 0.4$ is present among the time steps. Five separate spatially coherent changes, indexed by k , were introduced at various time steps as shown in Figure 5. The data generating model is

$$\begin{aligned} Y_{tz} &= \mu_z + \sum_{k=1}^5 \beta_k I(t \geq t_k^*) \Delta_k + \varepsilon_t, \\ \varepsilon_t &\sim \text{Normal}(0.4\varepsilon_{t-1}, \sigma^2 I_N), \end{aligned} \quad (24)$$

where μ_z is an underlying mean process, $\beta_k = 1 \forall k \in \{1, \dots, 5\}$, $\mathbf{t}^* = [20, 30, 40, 50, 60]^T$, the areas covered by Δ_k are shown in Figure 5, $\sigma = 0.5$, and $\varepsilon_0 = 0$.

5.2 Competing Methods and Metrics

We compare the spatial MR roboBayes algorithm to the pixel-wise implementation of roboBayes. Pixel-wise roboBayes is applied to the 2-dimensional signal over time for each pixel in the region and site grouping is done using thresholding, then associating connected regions to produce site polygons. MR roboBayes is applied as described in Section 3. Change declaration strategies under consideration vary based on subset (either levels 1–4 or 3–5 of a possible 8), type (recombining run lengths versus count threshold), and the number of components that need to be flagged. The first case considers components with levels 1–4 and combines them into real space run lengths using the probability of recent change of one component. The second case considers the probability of recent change of two components from levels 1–4. The third and fourth case are the same, but are instead applied to levels 3–5. The remaining cases threshold each wavelet component based on the probability of change exceeding τ_p to achieve a real space count of changed components, then threshold that count on $\tau_c = 1, \dots, 6$. The count strategy is applied to the same sets as described above.

The prior values for both algorithms are determined based on MLE estimates of the parameters for fits of a random sample of the simulated data during a stable period, with the exception of setting $\Lambda_0 = 0.1 * \hat{\Lambda}$. A search window of $L = 30$ was used. The metrics to compare the two methods are the precision, recall, F_1 score, and latency applied with respect to sites. The first step to evaluate the metrics is to associate detected change with annotated sites.

Site Association Let $\mathcal{A}_t = \{A_{t1}, \dots, A_{tn}\}$ be the set of true annotation polygons at time t . Recall $\mathcal{P}_t = \{P_{t1}, \dots, P_{tn}\}$ is the set of detected polygons at time t . The first attempt to associate detected to annotated polygons uses the intersection over union (IoU), or Jaccard index (Arbelaez et al., 2010; van den Burg and Williams, 2020), for each pair of polygons A_t and P_t :

$$IoU(A_t, P_t) = \frac{A_t \cap P_t}{A_t \cup P_t}. \quad (25)$$

Two site polygons are associated if their IoU is greater than the association threshold $\tau_{IoU} = 0.2$: $IoU(A_t, P_t) \geq \tau_{IoU}$. Using this metric, a single detected polygon that covers two sites can be inadvertently omitted because the intersection compared to the union of *either* site may fail to meet the threshold, i.e., under-segmentation. However, in the context of a broad area search, it is desirable to count such a detection as a success, leaving further refinement to more precise characterization methods down the line.

Therefore, we also match sites based on the proportion of the annotated polygon covered by the detected polygon using the Intersection over Truth (IoT):

$$IoT(A_t, P_t) = \frac{A_t \cap P_t}{Area(A_t)}. \quad (26)$$

A TP is recorded for each A_t such that $IoT(A_t, P_t) \geq \tau_{IoT} = 0.5$.

Metrics If a true change occurred in that recent time period, the detection is a true positive (TP). If no real change occurred in the preceding period, the detection is a false positive (FP). If the recent change probability does not exceed τ within $L = 15$ time points of a true change, a false negative (FN) is recorded. Precision, recall, and F1 score (Rijsbergen, 1979) are calculated from the TP, FN, and FP values over all the pixels.

$$\begin{aligned} \text{Precision} &= \frac{TP}{TP + FP}, \\ \text{Recall} &= \frac{TP}{TP + FN}, \\ \text{F1} &= 2 \frac{PR}{P + R}. \end{aligned}$$

Latency is calculated only for TP detections and is the difference between the first time the probability of recent change exceeds τ and the time of true change.

5.3 Simulation Results

Table 1 records the metrics for roboBayes and MR roboBayes. Pixel-wise roboBayes has poor precision, recall, and F1. The detections for an illustrative example in Figure 6 show that the poor recall for pixel-wise roboBayes stems mainly from oversegmentation, or the breaking up of the detection polygon into many smaller polygons. In the context of a broad area search, this makes it challenging to identify and query coherent sites for further examination. Precision suffers substantially from many of these small sites within true annotations being counted as FPs. F1 score is better for MR roboBayes than pixel-wise roboBayes. Latency for MR roboBayes is fairly short. High latency likely contributes to the difference in recall between pixelwise and MR roboBayes; by the time pixel-wise roboBayes would have been able to detect the change, the search period is over.

Table 1: Metrics for 100 simulations given as mean (standard error) comparing roboBayes with MR roboBayes. MR roboBayes, $c = 1$ and $c = 2$ refer to the methods that calculate run length based on the probability of recent change in 1 and 2 components, respectively. MR roboBayes subset applies methods to run lengths based on only levels 3, 4, and 5. MR roboBayes count uses a threshold on the number of wavelet components to declare change sites.

	Precision	Recall	F1	Latency
roboBayes	0.00 (0.00e+00)	0.00 (0.00e+00)	0.00 (0.00e+00)	–
MR roboBayes, $c=1$	0.03 (1.80e-03)	0.89 (1.00e-02)	0.06 (3.22e-03)	3.79 (4.50e-02)
MR roboBayes, $c=2$	0.02 (6.39e-04)	0.77 (9.59e-03)	0.04 (1.21e-03)	4.35 (1.34e-01)
MR roboBayes subset, $c=1$	0.88 (1.97e-02)	1.00 (0.00e+00)	0.92 (1.40e-02)	4.06 (4.22e-02)
MR roboBayes subset, $c=2$	0.89 (2.12e-02)	1.00 (2.00e-03)	0.92 (1.68e-02)	4.77 (1.21e-01)
MR roboBayes count1	0.03 (1.78e-03)	0.89 (1.00e-02)	0.05 (3.19e-03)	3.81 (5.01e-02)
MR roboBayes count2	0.02 (6.52e-04)	0.77 (9.73e-03)	0.04 (1.24e-03)	4.40 (1.43e-01)
MR roboBayes count3	0.01 (5.42e-04)	0.46 (1.37e-02)	0.03 (1.04e-03)	6.25 (3.36e-01)
MR roboBayes count4	0.00 (3.23e-04)	0.02 (5.45e-03)	0.00 (6.06e-04)	21.60 (7.16e-01)
MR roboBayes count5	0.00 (0.00e+00)	0.00 (0.00e+00)	0.00 (0.00e+00)	–
MR roboBayes count6	0.00 (0.00e+00)	0.00 (0.00e+00)	0.00 (0.00e+00)	–
MR roboBayes count1 subset	0.88 (1.97e-02)	1.00 (0.00e+00)	0.92 (1.40e-02)	4.08 (4.86e-02)
MR roboBayes count2 subset	0.89 (2.12e-02)	1.00 (2.00e-03)	0.92 (1.70e-02)	4.81 (1.26e-01)
MR roboBayes count3 subset	0.06 (4.01e-03)	0.67 (1.74e-02)	0.10 (6.16e-03)	6.02 (2.21e-01)
MR roboBayes count4 subset	0.04 (2.06e-03)	0.40 (1.24e-02)	0.07 (3.36e-03)	5.34 (3.03e-01)
MR roboBayes count5 subset	0.79 (3.22e-02)	0.26 (9.38e-03)	0.36 (1.22e-02)	4.17 (3.16e-01)
MR roboBayes count6 subset	0.00 (0.00e+00)	0.00 (0.00e+00)	0.00 (0.00e+00)	–

Within the MR roboBayes variants, metrics are sensitive to the subset of components that are examined. The subset that excludes levels 1 and 2 shows superior performance to the more inclusive set of components. Applying the run length conversion requiring recent change in at least two components to only the coarse subset (levels 3–5) of components produces the best precision, recall, and F1 score. The latency is higher than some of the alternatives, but is still competitive.

6 Construction Broad Area Search

In this section, we analyze the Jacksonville and Dubai construction data described in Section 2 using pixel-wise roboBayes and MR roboBayes. The response \mathbf{Y}_t containing the signals is observed at each time t . We model \mathbf{Y}_t using the state distribution in (9) about a third-order

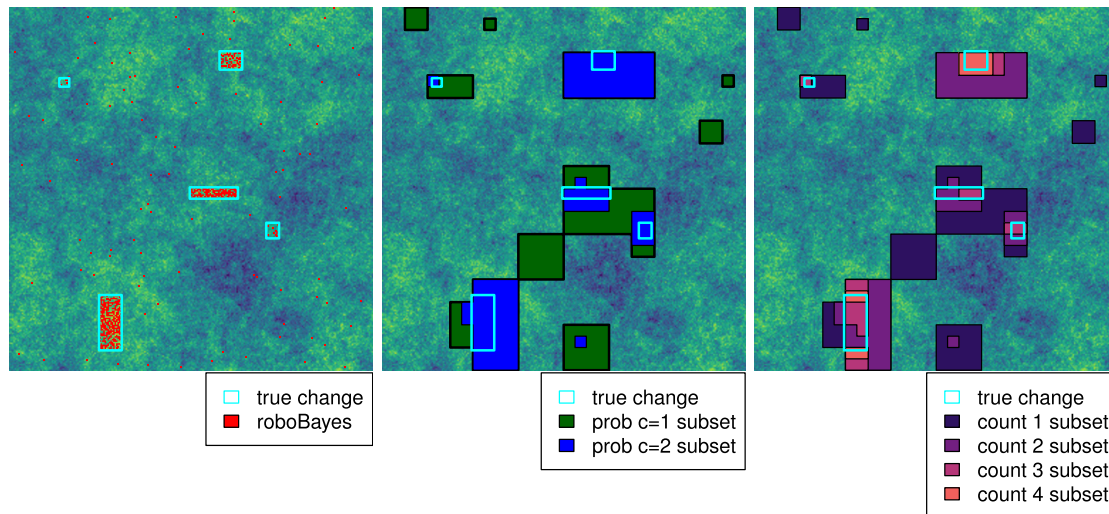


Figure 6: Site polygons for roboBayes and MR roboBayes on one signal from the simulated data.

seasonal model and calculate posterior run lengths using each algorithm. In a true monitoring scenario, historical data would be used to estimate prior hyperparameters. In this case, there are only 4 years of data and few change sites occurring near the end of the period, making temporal splitting of the data problematic. Instead, priors are calculated for each region based on regression fits to data in that region over the entire time period. The hyperparameters B_0 , V_0 , v_0 , and Λ_0 are estimated based on the distribution of the regression parameter estimates. The component of $\Lambda_0 = \text{diag}(\hat{\Lambda})$ governing the inverse variance of the linear trend is set to 100 to limit its possible magnitude. When the estimated parameters deliver poor performance, further adjusted settings are considered (any adjustments are specified in the following subsections). The prior state change probability is set to $\lambda = 1e - 8$, the search window is set to $L = 30$, and the other parameters are set to the defaults in **roboBayes**.

The sites under investigation are expected to be at least 9000 m², so the data is aggregated to 24 m for pixel-wise analysis of both regions. Similarly, the lowest resolution wavelets are not monitored in the multiresolution approach; levels 3 to 7 (24 × 24 m² to 384 × 384 m²) are monitored. Small sites less than 9000 m² in area are filtered out for both methods. For MR roboBayes, only the horizontal and vertical direction wavelets are monitored because they (or their combination, depending on the image orientation) are interpretable as the edges of buildings. The diagonal components do not add much new information, so they are not monitored, reducing monitoring computation by one third. Based on the simulation result performance and the relative ease of tuning with counts (after the monitoring computation) compared to pre-monitoring tuning parameters, the count threshold approach is chosen for MR roboBayes.

Since this method is intended as an initial step that scans an area for change, identifying sites for further inspection, measured detections may include change unrelated to construction. The available annotations in Figures 1 and 2 are construction-specific; real land cover change exists in other parts of the image. Recall that only the yellow annotations indicate that the entire construction activity occurred during the period of data availability. Failing to detect a green annotation is not necessarily a false negative.

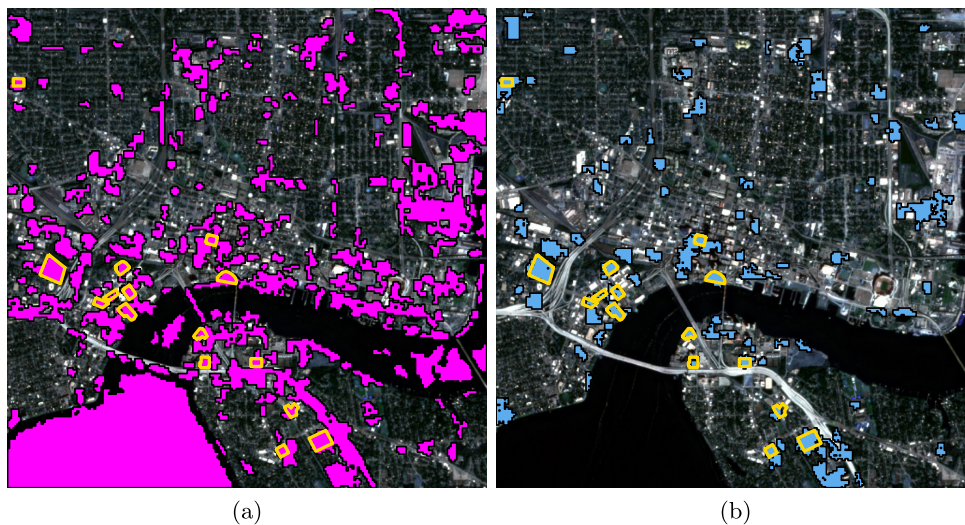


Figure 7: Detected change sites in Jacksonville for (a) pixel-wise roboBayes and (b) MR roboBayes with count threshold 9.

6.1 Jacksonville

Pixel-wise roboBayes was run for several sets of prior hyperparameters and the one with the best performance (captured all the known TP sites) inflates the estimated \hat{V}_0 by a factor of 2 and sets $\nu_0 = 30$ (the estimated parameters resulted in over detection). MR roboBayes was run using the estimated hyperparameters from the region. For a single update, MR roboBayes takes less than three quarters of the time that pixel-wise roboBayes does (a reduction from 27 to 19 seconds) due to the smaller number of monitoring processes.

A summary of all change sites detected over the entire Jacksonville region from 2018-01-01 to 2022-03-15 for both pixel-wise roboBayes and MR roboBayes are shown in Figure 7. Change sites for MR roboBayes were determined based on a count threshold of 9, which was the highest count threshold that captured all of the available annotated construction events. Many additional, non-annotated events are detected by both algorithms. These sites are either too small or contain changes of a different type than those in the “truth” set. Pixel-wise roboBayes is particularly susceptible to FPs over water. MR roboBayes captures compact change sites better than long, skinny ones like the road near the bottom of the image.

The evolution of change detection metrics for a construction site over time is shown in Figure 8. The site is undisturbed at 2019-07-04. In 2019-08-13, after site clearing, the pixel-wise change proposal appears oversegmented while the site proposed by MR roboBayes is more cohesive. On 2019-10-12, as construction progresses, both methods produce a strong change signal for the entire site. As activity in the rest of the site has settled down in 2020-01-30, both methods easily identify the completion of the building and distinguish it from its surroundings.

6.2 Dubai

In Dubai, only the red band was monitored. Pixel-wise roboBayes was run with settings $\nu_0 = 60$, an inflated variance $V_0 = 10 * \hat{V}$, and a search period of 15 time points. MR roboBayes was run with estimated prior parameters except for setting $\nu_0 = 50$. Settings were chosen to achieve similar overall TP detection performance.

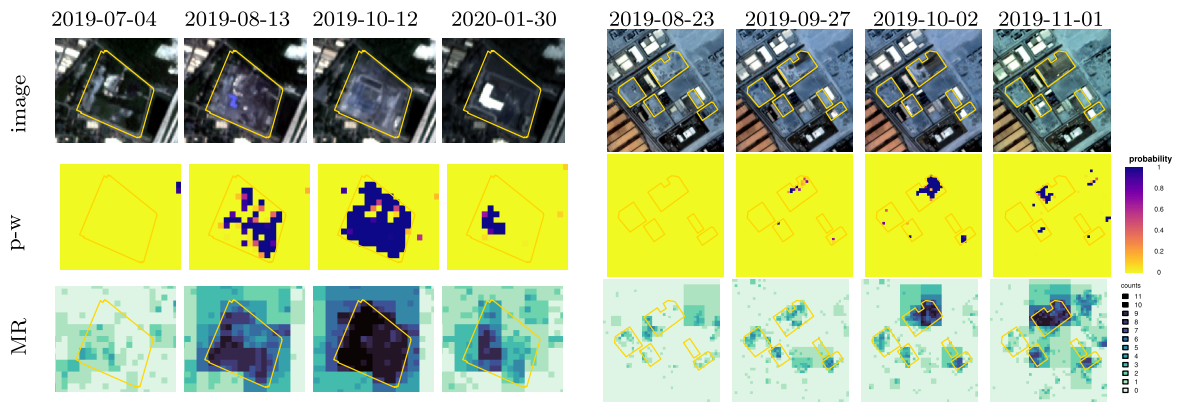


Figure 8: Change detection evolution for construction in (left) Jacksonville and (right) Dubai for pixel-wise (p-w) and MR roboBayes.

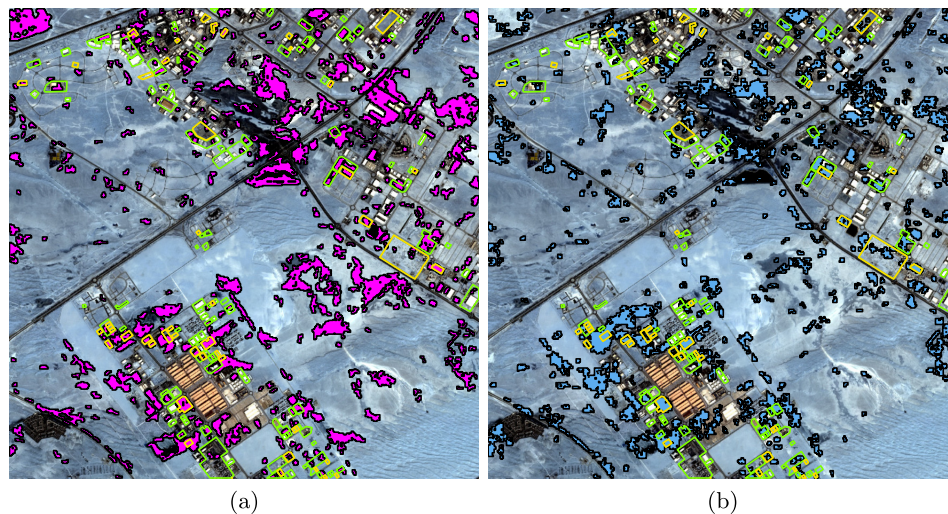


Figure 9: Detected change sites in Dubai for (a) pixel-wise roboBayes and (b) MR roboBayes with count threshold 7.

All change sites detected in Dubai from 2018-01-01 to 2022-04-15 for both pixel-wise roboBayes and MR roboBayes are shown in Figure 7. At these settings, MR roboBayes is more robust to dune shifting changes than the pixel-wise algorithm, but struggles to detect all of the annotated construction sites. MR roboBayes achieves similar detection performance using only two thirds of the monitoring runs by omitting diagonal terms.

Change detection probabilities and counts for a subset of the region are shown in Figure 8. The sites are undisturbed at 2019-08-23. In 2019-09-27, noticeable change is reflected in spotty pixel-wise change proposals and relatively low magnitude, but cohesive, MR roboBayes change proposals. On 2019-10-02, both methods produce a strong change signal for the changed areas. On 2019-11-01, both methods capture recent change in several sites, but pixel-wise roboBayes marks several detections in the desert as well. With an appropriate threshold, MR roboBayes may not detect these.

7 Discussion

Monitoring large geographical regions for change must quickly identify cohesive candidate sites while remaining computationally feasible. The pixel-wise roboBayes algorithm, like other pixel-based methods, fails to take into account the spatial patterns that exist in change sites. This leads to estimated change regions with fragmented sites that underestimate the spatial range of real change and introduce many false positive detections. We propose MR roboBayes to introduce a multiresolution spatial component to the application of roboBayes while preserving the opportunity for distributed computing. Compared to a pixel-wise analysis based on the same change detection algorithm, MR roboBayes generates more cohesive sites that translate into better detection metrics. In a simulation study, we show that the sites detected by MR roboBayes better capture real change sites in the context of broad area search. Further, the simulations show that considering detections in multiple resolutions, whether that is through run length recombination or a count of flagged wavelets components, can greatly improve precision without hurting recall and leaving latency similar. We apply MR roboBayes to regions in Jacksonville and Dubai to detect heavy construction sites.

The results from MR roboBayes are sensitive to the set of wavelet basis functions that are monitored. Including coarse resolutions can lead to very large proposal areas, so care must be taken in choosing the coarsest component. Further, in this application, the wavelet coefficient signals can be somewhat noisier than in pixel-wise signals, hiding some types of change. However, MR roboBayes behaves fairly well with empirical hyperparameters and the count threshold allows for post-hoc tuning, making it more wieldy than pixel-wise roboBayes. It has a niche application as a search mechanism that prioritizes presenting cohesive sites where pixel-based methods give delineation as a by-product. Where pixel-based methods might require post-processing to form larger connected sites, MR roboBayes could benefit from site shape refinement. Refinement based on changed fine resolutions or nonzero fine resolutions (edges) are opportunities for presenting more precise change sites.

In this application, Haar wavelets are a natural choice to represent edges of artificial structures. Construction is a special case where we expect rectilinear change, as are various anthropogenic land cover changes like cultivation. Land cover changes like deforestation would not necessarily be rectilinear, so other wavelet basis functions, or even other decompositions entirely, may prove to be appropriate inputs for the roboBayes algorithm. Further investigation into translating wavelet basis change into real space change is warranted, using rules like the distribution of change in multiple components to define change.

Supplementary Material

The construction of 2D Haar wavelets, a calculation for induced prior change probability, and sensitivity results are available in the Supplementary Material. The Jacksonville and Dubai dataset is proprietary, but code to generate and analyze a simulated dataset is included to demonstrate the algorithm.

Acknowledgement

We extend our gratitude to Dr. Shannon Holloway for her continued maintenance of the roboBayes R package.

Funding

This research is based upon work supported in part by the Office of the Director of National Intelligence (Intelligence Advanced Research Projects Activity) via 2021-20111000006. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This material is based upon work supported in part by the National Science Foundation under Grant No. DMS-2152887.

References

- Adams RP, MacKay DJC (2007). Bayesian online changepoint detection. arXiv preprint: <https://arxiv.org/abs/0710.3742>.
- Arbelaez P, Maire M, Fowlkes C, Malik J (2010). Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33: 898–916.
- Balafas K, Kiremidjian AS, Rajagopal R (2018). The wavelet transform as a gaussian process for damage detection. *Structural Control and Health Monitoring*, 25: e2087.
- Barry D, Hartigan JA (1993). A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, 88: 309–319.
- Bontemps S, Bogaert P, Titeux N, Defourny P (2008). An object-based change detection method accounting for temporal dependences in time series with medium to coarse spatial resolution. *Remote Sensing of Environment*, 112: 3181–3191.
- Celik T, Ma KK (2011). Multitemporal image change detection using undecimated discrete wavelet transform and active contours. *IEEE Transactions on Geoscience and Remote Sensing*, 49: 706–716.
- Daubechies I (1988). Orthonormal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 41: 909–996.
- Daubechies I (1990). The wavelet transform, time-frequency localization and signal analysis. *IEEE Transactions on Information Theory*, 36: 961–1005.
- Donoho DL (1995). De-noising by soft-thresholding. *IEEE Transactions on Information Theory*, 41: 613–627.
- Gu M, Fei J, Sun S (2020). Online anomaly detection with sparse gaussian processes. *Neurocomputing*, 403: 383–399.
- Horvath L, Rice G, Zhao Y (2022). Change point analysis of covariance functions: A weighted cumulative sum approach. *Journal of Multivariate Analysis*, 189: 104877.
- Houborg R, McCabe MF (2018a). A cubesat enabled spatio-temporal enhancement method (ces-tem) utilizing planet, landsat and modis data. *Remote Sensing of Environment*, 209: 211–226.
- Houborg R, McCabe MF (2018b). Daily retrieval of ndvi and lai at 3 m resolution via the fusion of cubesat, landsat, and modis data. *Remote Sensing*, 10: 890.
- Huang Y, Bortoli VD, Zhou F, Gilles J (2018). Review of wavelet-based unsupervised texture segmentation, advantage of adaptive wavelets. *IET Image Processing*, 12: 1626–1638.
- Imbiriba T, Lamountain G, Wu P, Erdogmus D, Closas P (2019). Change detection and gaussian process inference in piecewise stationary environments under noisy inputs. In: *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*,

- CAMSAP 2019 – Proceedings*, 530–534.
- Killick R, Eckley IA, Jonathan P (2013). A wavelet-based approach for detecting changes in second order structure within nonstationary time series. *Electronic Journal of Statistics*, 7: 1167–1183.
- Knoblauch J, Damoulas T (2018). Spatio-temporal Bayesian on-line changepoint detection with model selection. In: *Proceedings of the 35th International Conference on Machine Learning*, volume 80, 2718–2727.
- Kuhn J, Ellens W, Mandjes M (2014). Detecting changes in the scale of dependent gaussian processes: A large deviations approach. *Lecture Notes in Computer Science (including sub-series Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, LNCS, 8499: 170–184.
- Lange MM, Lange AM (2017). A hierarchical algorithm for approximate nearest neighbor searching in a dataset of pyramid-based image representations. *Procedia Engineering*, 201: 302–311.
- Li Z, Shi W, Zhang H, Hao M (2017). Change detection based on gabor wavelet features for very high resolution remote sensing images. *IEEE Geoscience and Remote Sensing Letters*, 14: 783–787.
- Lu M, Pebesma E, Sanchez A, Verbesselt J (2016). Spatio-temporal change detection from multidimensional arrays: Detecting deforestation from MODIS time series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 117: 227–236.
- Mallat SG (1989). A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11: 674–693.
- Mei Y (2010). Efficient scalable schemes for monitoring a large number of data streams. *Biometrika*, 97: 419–433.
- Niemeyer I, Marpu P, Nussbaum S (2008). Change detection using object features. In: *Object-Based Image Analysis: Spatial Concepts for Knowledge-Driven Remote Sensing Applications* (T Blaschke, S Lang, GJ Hay, eds.), 185–201. Springer, Berlin Heidelberg, Berlin, Heidelberg.
- Page ES (1954). Continuous inspection schemes. *Biometrika*, 41: 100–115.
- Planet Fusion Team (2022). Planet Fusion Monitoring technical specification, version 1.0.0, San Francisco, CA.
- Rijsbergen CJV (1979). *Information Retrieval*. Butterworths.
- Rinoshika A, Rinoshika H (2020). Application of multi-dimensional wavelet transform to fluid mechanics. *Theoretical and Applied Mechanics Letters*, 10: 98–115.
- Roberts SW (1959). Control chart tests based on geometric moving averages. *Technometrics*, 1: 239–250.
- Saatçi Y, Turner R, Rasmussen CE (2010). Gaussian process change point models. In: *Proceedings of the 27th International Conference on Machine Learning*, 927–934.
- Shensa MJ (1992). The discrete wavelet transform: Wedding the Å Trouns and Mallat algorithms. *IEEE Transactions on Signal Processing*, 40: 2464–2482.
- Susiluoto J, Spantini A, Haario H, Härkönen T, Marzouk Y (2020). Efficient multi-scale gaussian process regression for massive remote sensing data with satgp v0.1.2. *Geoscientific Model Development*, 13: 3439–3463.
- Tucker CJ (1979). Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sensing of Environment*, 8: 127–150.
- van den Burg GJJ, Williams CKI (2020). An evaluation of change point detection algorithms. arXiv preprint: <https://arxiv.org/abs/2003.06222>.
- Wang XY, Yang HY, Fu ZK (2010). A new wavelet-based image denoising using undecimated

- discrete wavelet transform and least squares support vector machine. *Expert Systems with Applications*, 37: 7040–7049.
- Wendelberger LJ, Gray JM, Reich BJ, Wilson AG (2021). Monitoring deforestation using multivariate Bayesian online changepoint detection with outliers. arXiv preprint: <https://arxiv.org/abs/2112.12899>.
- Xu F, Liu J, Dong C, Wang X (2017). Ship detection in optical remote sensing images based on wavelet transform and multi-level false alarm identification. *Remote Sensing*, 9: 985.
- Zhu Z, Woodcock CE (2014). Continuous change detection and classification of land cover using all available landsat data. *Remote Sensing of Environment*, 144: 152–171.
- Zhu Z, Zhang J, Yang Z, Aljaddani AH, Cohen WB, Qiu S, et al. (2020). Continuous monitoring of land disturbance based on landsat time series. *Remote Sensing of Environment*, 238: 111116.