

RESEARCH ARTICLE

Scalable multiple changepoint detection for functional data sequences

Trevor Harris¹ | Bo Li¹ | J. Derek Tucker^{1,2}

¹Department of Statistics, University of Illinois at Urbana-Champaign, Champaign, Illinois, USA

²Sandia National Laboratories, Albuquerque, New Mexico, USA

Correspondence

Trevor Harris, Department of Statistics, University of Illinois at Urbana-Champaign, Champaign, IL, USA.
Email: trevorh2@illinois.edu

Funding information

Laboratory Directed Research and Development Program at Sandia National Laboratories, Grant/Award Number: DE-NA0003525; National Science Foundation, Grant/Award Numbers: NSF-DMS-1830312, NSF-DGE-1922758, NSF-DMS-2124576

Abstract

We propose the multiple changepoint isolation (MCI) method for detecting multiple changes in the mean and covariance of a functional process. We first introduce a pair of projections to represent the variability “between” and “within” the functional observations. We then present an augmented fused lasso procedure to split the projections into multiple regions robustly. These regions act to isolate each changepoint away from the others so that the powerful univariate CUSUM statistic can be applied region-wise to identify the changepoints. Simulations show that our method accurately detects the number and locations of changepoints under many different scenarios. These include light and heavy tailed data, data with symmetric and skewed distributions, sparsely and densely sampled changepoints, and mean and covariance changes. We show that our method outperforms a recent multiple functional changepoint detector and several univariate changepoint detectors applied to our proposed projections. We also show that MCI is more robust than existing approaches and scales linearly with sample size. Finally, we demonstrate our method on a large time series of water vapor mixing ratio profiles from atmospheric emitted radiance interferometer measurements.

KEYWORDS

atmospheric radiance, CUSUM, functional change points, fused lasso, robust procedures, time domain, time series

1 | INTRODUCTION

The statistical analysis of functional time series has become increasingly important to many scientific fields including climatology (Shang & Hyndman, 2011), finance (Kokoszka & Zhang, 2012), geophysics (Hörmann & Kokoszka, 2012), demography (Hyndman & Booth, 2008), manufacturing (Woodall, 2007), and environmental modeling (Finazzi et al., 2019; Fortuna et al., 2020; Qu et al., 2021). A functional time series is a sequence of infinite dimensional objects, such as curves and surfaces, observed over time. Functional time series are analogous to univariate or multivariate time series, except that we observe a continuous function at each point in time (Bosq, 2012). Just as in univariate and multivariate time series, a functional time series can experience abrupt changes in its generating process. These abrupt changes, or changepoints, can complicate statistical analysis by invalidating stationarity assumptions. However, they can also be interesting in their own right by revealing unexpected heterogeneous patterns.

Recently, changepoint analysis has become particularly important in climatological and environmental process modeling (Beaulieu et al., 2020; Jarušková & Antoch, 2020; Lee et al., 2020). Although these works focus on important

univariate changepoint detection applications, environmental profiles, such as water vapor columns (Sakai et al., 2019), are increasingly the object of study. Thus changepoint detection in functional time series are a natural and timely extension to the finite dimensional methods.

Within the functional data analysis (FDA) literature, changepoint detection has largely focused on the at most one change (AMOC) problem. In Berkes et al. (2009) a cumulative sum (CUSUM) test was proposed for independent functional data, which was further studied in Aue et al. (2009), where its asymptotic properties were developed. This test was then extended to weakly dependent functional data by Hörmann and Kokoszka (2010) and epidemic changes by Aston and Kirch (2012). Zhang et al. (2011) introduced a test for changes in the mean of weakly dependent functional data using self-normalization to alleviate the use of asymptotic control. Later, Sharipov et al. (2016) similarly developed a sequential block bootstrap procedure for these methods. Recently, Gromenko et al. (2017) considered changes in spatially correlated functional data, and Aue et al. (2018) proposed a fully functional method for finding a mean change without losing information due to dimension reduction.

Detecting multiple changepoints in a functional time series has received relatively scant attention compared to the AMOC problem. Recently, Li and Ghosal (2018) proposed a Bayesian method for identifying multiple changepoints in the mean by transforming the functional data into wavelets and identifying changes in the wavelet coefficient processes. In addition, Chiou et al. (2019) introduced a dynamic segmentation method for multiple changepoints in the mean based on dynamic programming and backward elimination to find an optimal set of changepoints. Multiple changepoints can also be identified by augmenting AMOC methods with a recursive binary segmentation algorithm to partition the functional time series (Aue et al., 2018; Berkes et al., 2009). The consistency of binary segmentation approaches was established by Rice and Zhang (2019). Despite these advances, there are several outstanding issues with these approaches that we hope to address. Namely, suboptimal computational scalability, insufficient power to detect covariance and shape based alternatives, and a lack of robustness.

The computational complexity of functional multiple changepoints detection methods has been an obstacle for their wide application to large functional time series. Bayesian methods that rely on Markov Chain Monte Carlo sampling are intrinsically burdensome because they typically require an enormous number of samples to reach convergence. Ordinary dynamic programming and binary segmentation algorithms scale quadratically and log-linearly, respectively, with the data's sample size. As larger and larger functional time series data sets are curated, methods that scale linearly with sample size are called for to meet the computational demand. In the univariate and multivariate changepoint detection literature, computationally efficient methods for multiple changepoint detection have already emerged. These include the cumulative segmentation method (Muggeo & Adelfio, 2011), the pruned exact linear time (PELT) algorithm (Killick et al., 2012), the functional pruning optimal partitioning (FPOP) algorithm (Maidstone et al., 2017), and the robust functional pruning optimal partitioning (r-FPOP) (Fearnhead & Rigail, 2019).

Another limitation of the existing functional changepoint detection methods is that most methods can only detect changes in the functional process's mean. While mean changes are the most conspicuous, covariance changes are equally important and can also occur. The need for detecting covariance changes has already been noticed and tackled in the univariate time series literature, where methods targeting variance changes have been developed (Adelfio, 2012; Chapman et al., 2020). Recently, covariance changepoint detection has been considered in functional data (Aue et al., 2020), however, methods that target both mean and covariance changes in a functional process are still not available. Lastly, many previous methods were developed under the assumption that the data follows a Gaussian process (GP). Their performance on non-Gaussian, skewed, or heavy tailed data, which may be encountered in practice, is not well studied and could potentially be suboptimal.

We propose the multiple changepoint isolation (MCI) method to robustly detect multiple changepoints in a functional time series' distribution. Our method uses a pair of projections to minimally represent the variability "between" and "within" individual functions. Although there is no clear direct connection between the projections and specific parameters, we show that the projections largely preserve the mean and covariance information (Section 3) We then detect changepoints in each projection using an augmented fused lasso procedure based on the CUSUM statistic (Page, 1954). After that, changepoints are aggregated across all projections into a final list. This approach combines robust segmentation through the fused lasso (Tibshirani et al., 2005) with optimal detection through the CUSUM to achieve very low error rates compared with existing methods (see Section 3 and Supplement Section B).

Our method also remedies three outstanding issues in the functional changepoint detection literature. First, our method detects a broader range of changepoint types, beyond mean changes, than existing methods. Second, our method is computationally efficient, having only linear time computational complexity. Finally, our method is robust to asymmetry and heavier than Gaussian tails.

2 | MULTIPLE CHANGEPOINT ISOLATION

Let $\{f_t, t \in 1, \dots, n\}$ be a sequence of continuous functions in $L^2([0, 1])$, hereafter L^2 , the Hilbert space of square integrable functions with domain $[0, 1]$ satisfying $\int_0^1 |f_t(s)|^2 ds < \infty$. The domain $[0, 1]$ is assumed without loss of generality, and we further assume that each function $f_t \in \{f_t, t \in 1, \dots, n\}$ is observed on the same finite grid of points $0 < s_1 < \dots < s_m < 1$. Suppose that the observations $\{f_t, t \in 1, \dots, n\}$ are distributed according to

$$f_t \sim F(\mu_t, \Sigma_t), \quad (1)$$

where $F(\mu_t, \Sigma_t)$ refers to a functional process with mean function μ_t and covariance function Σ_t . We assume that $F(\cdot)$ is piecewise weakly stationary, that is $\mu_{t+1} = \mu_t$ and $\Sigma_{t+1} = \Sigma_t$ for all $t \in 1, \dots, n$ except when t is a changepoint. We assume there are M_1 time points where $\mu_{t+1} \neq \mu_t$ and M_2 time points where $\Sigma_{t+1} \neq \Sigma_t$. Both M_1 and M_2 are unknown integers between 1 and n . Let τ_j^μ with $j = 1, \dots, M_1$ and τ_k^Σ with $k = 1, \dots, M_2$ denote the times where μ_t and Σ_t change, that is, where $\mu_{\tau_j^\mu+1} \neq \mu_{\tau_j^\mu}$ and $\Sigma_{\tau_k^\Sigma+1} \neq \Sigma_{\tau_k^\Sigma}$. For simplicity and ease of notation, we union the two sets of changepoints into a combined sequence τ_1, \dots, τ_M , where $\max(M_1, M_2) \leq M < n$. The goal of our method is to estimate *all* changepoint locations τ_1, \dots, τ_M from the sequence $\{f_t, t \in 1, \dots, n\}$.

2.1 | Main algorithm

We divide our detection algorithm into four distinct steps. The first two are proposed to give initial estimates of changepoints so that the third step can finely tune the initial estimates using the standard CUSUM statistic with a false discovery rate (FDR) control. The final step aggregates all changepoints detected across all projections together. We specify the four steps of our detection algorithm below:

1. *Univariate projection.* Project the functional data, $\{f_t, t \in 1, \dots, n\}$ onto \mathbb{R} with two projections, one representing variability “between” different curves, denoted by P_T , and the other representing variability “within” individual curves, denoted by P_{E1} . The two projections yield two projected time series, $Y^{(1)} = \{Y_t^{(1)}, t \in 1, \dots, n\}$ and $Y^{(2)} = \{Y_t^{(2)}, t \in 1, \dots, n\}$.
2. *Changeset regionalization.* Represent each projected process as $Y^{(i)} = \theta^{(i)} + \epsilon^{(i)}$ for $i = 1, 2$, where $\theta^{(i)}$ is a piecewise constant vector with entries $\theta_t^{(i)}$ and $\epsilon^{(i)}$ is i.i.d. noise. Estimate $\hat{\theta}^{(i)} = \arg \min_{\theta} ||Y^{(i)} - \theta^{(i)}||_2^2 + \lambda \sum_{t=1}^{n-1} |\theta_{t+1}^{(i)} - \theta_t^{(i)}|$, that is, with the fused lasso estimator. Group nearby jumps in $\hat{\theta}^{(i)}$ within c time steps into individual sets, called “changesets.” The changesets define subregions of $Y^{(i)}$, with at least $2c + 1$ data points each, that are assumed to contain AMOC.
3. *CUSUM.* Apply CUSUM regionwise to detect the potential changepoint in each subregion of the two projections. Apply an FDR correction to control the number of false detections at a specified level α . Retain the changepoints whose p -values after correction are below α . Let $\hat{\tau}_1^{(1)}, \dots, \hat{\tau}_{m_1}^{(1)}$ and $\hat{\tau}_1^{(2)}, \dots, \hat{\tau}_{m_2}^{(2)}$ denote the estimated changepoints from the two projections, respectively.
4. *Finalization.* Concatenate, sort, and deduplicate the retained changepoints from both projections into a single list, denoted by $\hat{\tau}_1, \dots, \hat{\tau}_m$, which are our final estimates of the true changepoints.

Three hyper-parameters are involved in the algorithm: the fused lasso penalty term λ , the linkage parameter c , and the false positive rate α . We use the data to optimize the λ and c parameters through grid search. Details are given in Section A.2. The false positive level α is considered fixed, as is common in hypothesis testing, although this parameter could be optimized as well. A small simulation study in Section B.2 in the supplement shows that MCI is relatively insensitive to the prespecified α level, as long as α is below 0.01. We elaborate the first two steps of our algorithm in the following two subsections.

2.2 | Univariate projections

Step one projects the functional data sequence $\{f_t, t \in 1, \dots, n\}$ onto \mathbb{R} through two projections P_T and P_{E1} , which are defined as

$$P_T(f_t) = \int_0^1 |\nabla f_t(s)| ds, \quad (2)$$




FIGURE 1 Estimated changepoints, using CUSUM applied to each FPC, under two types of changes (smoothness and range—see Equation 6) for a functional time series with $n = 1000$ and one changepoint at $t = 500$. CUSUM estimates (black dots) are generally centered around the true changepoint location (red line) but show high variability. No black dot indicates that no changepoint was detected in that FPC. This illustrates a major challenge in detecting changepoints, particularly covariance-based changepoints, purely through FPC analysis since the estimate can vary wildly depending on which FPC is used

$$P_{E1}(f_t) = \int_0^1 f_t(s) \phi_1(s) ds, \quad (3)$$

where ∇ is the differential operator, $|\cdot|$ is the L^2 norm, and ϕ_1 is the eigenfunction of the covariance operator of f_t corresponding to the largest singular value.

The first projection P_T is a total variation norm (TVN) type projection of f_t . The TVN based projection measures the variability “within” the function f_t by measuring the total squared length of f_t , and it is invariant to rigid and elastic transformations. The projection P_T is sensitive to the high-frequency features in $\{f_t, t \in 1, \dots, n\}$, because features at higher frequencies contribute more to the length than lower frequencies. Therefore, P_T can be more helpful in detecting changes in the covariance operators Σ_t (Equation 1), such as changes in variance or smoothness of the functions (Section 3.3).

The second projection P_{E1} is the first functional principal component, which measures the dominant mode of variation in the data and can be estimated through standard functional principal component analysis (Ramsay, 2004) or elastic functional principal component analysis if the data contain phase variability (Srivastava & Klassen, 2016; Tucker et al., 2013). Whereas P_T is sensitive to high frequency features, P_{E1} is sensitive to low frequency features such as the mean, slope, orientation, and sign of functional curves. Combining P_T and P_{E1} will retain the information in functional sequences at both high and low frequencies.

We could include further FPC coefficients, say $P_{E2}(f_t) = \int_0^1 f_t(s) \phi_2(s) ds$ and $P_{E3}(f_t) = \int_0^1 f_t(s) \phi_3(s) ds$, but we find this is unnecessary since P_T is more sensitive to high frequency changes than any individual FPC (Section B.2 of the supplement). Using only the top FPC also avoids having to select the right number of FPCs through a potentially expensive cross-validation procedure (Hyun et al., 2016).

Univariate projections are a classic and powerful technique in functional data for extracting features and compressing a functional process to a finite dimensional space. Our system of only two projections is based on the principle that the number of projections should be minimized and that each projection should provide a high amount of independent information. Overly projecting the data can lead to two adverse effects in change point detection: increased estimation error and reduced test power. Estimation error increases because changes will often partially manifest in multiple projections simultaneously, meaning changepoints may be repeatedly estimated on noisy representations. This can lead to many spurious changepoint estimates near the true ones (Figure 1). Test power decreases because more projections mean more tests are being performed, so the necessary FDR or family wise error rate correction will have to control for a higher number of tests. This directly leads to a power loss for each test, which can result in missing some changepoints entirely.

2.3 | Changeset regionalization

Let $Y^{(i)} = \{Y_t^{(i)}, t \in 1, \dots, n\}$, $i = 1, 2$, denote the two projections of the functional sequence $\{f_t, t \in 1, \dots, n\}$ by P_T and P_{E1} . Each projection conveniently converts changes in the spectrum of f_t into simple mean changes in $Y_t^{(i)}$. We model the mean

changes, or jumps, with a signal-plus-noise model, $Y^{(i)} = \theta^{(i)} + \epsilon^{(i)}$, where $\theta^{(i)}$ is a piecewise constant $n \times 1$ vector with entries $\theta_t^{(i)}$, and the i.i.d. noise term $\epsilon^{(i)}$ has a finite second moment. Because $\theta^{(i)}$ is piecewise constant, the changepoints in $Y^{(i)}$ are reflected in the jumps, or nonzero differences, of $\theta^{(i)}$, that is, where $\theta_{t+1}^{(i)} \neq \theta_t^{(i)}$. Without any risk of confusion, we suppress the superscript (i) to ease the notation in what follows.

Ideally, we could estimate $\hat{\theta}$ with an L^0 norm penalized empirical risk estimator, since L^0 minimization asymptotically recovers the jump locations (Boysen et al., 2009). However, L^0 minimization is often computationally infeasible. We, therefore, instead start with a convex approximation, L^1 minimization, and try to modify the approximation into an L^0 -like solution. The L^1 minimization, in this case known as the fused lasso (Tibshirani et al., 2005), has received a flurry of recent interest (Chan et al., 2014; Hyun et al., 2016; Lin et al., 2016; Rojas & Wahlberg, 2014) for changepoint estimation. The fused lasso is an optimally adaptive solution to piecewise signal recovery (Donoho & Johnstone, 1998; Guntuboyina et al., 2017; Mammen & van de Geer Mammen, 1997) that can be solved for with non-iterative convex optimization (Condat, 2013). Thus it can be used to accurately and quickly approximate θ . The fused lasso estimates θ via

$$\hat{\theta} = \arg \min_{\theta} \|Y - \theta\|_2^2 + \lambda \sum_{t=1}^{n-1} |\theta_{t+1} - \theta_t|, \quad (4)$$

where λ is a tuning parameter that controls the degree of regularization. We optimize λ from the data as part of our overall fitting procedure based on BIC minimization (Section A.2). The estimator $\hat{\theta}$ will necessarily be a piecewise constant vector, due to the penalty term $\lambda \sum_{t=1}^{n-1} |\theta_{t+1} - \theta_t|$ that induces sparsity on the differences of θ (Tibshirani et al., 2005).

It seems we could simply identify changepoints with the fused lasso by the breaks in $\hat{\theta}$, that is, elements of the set $A(\hat{\theta}) = \{t : \hat{\theta}_t \neq \hat{\theta}_{t+1}\}$ for a given λ . However, the set $A(\hat{\theta})$ is suboptimal for changepoint detection because the fused lasso tends to overestimate the number of changepoints and split single large breaks into multiple, tightly grouped, and smaller breaks (Fryzlewicz, 2014; Rojas & Wahlberg, 2014). Moreover, while the fused lasso is optimally adaptive to piecewise constant signals (Guntuboyina et al., 2017), it lacks consistent support recovery (Rojas & Wahlberg, 2014), meaning changepoint detection with $A(\hat{\theta})$ is inconsistent without further refinement (Chan et al., 2014; Hyun et al., 2016).

We propose a changeset regionalization procedure to further sparsify and improve the fused lasso estimates as a solution for changepoint detection. Recognizing that the fused lasso tends to estimate a single changepoint with a sequence of closely grouped estimates around the true value (Rojas & Wahlberg, 2014), the regionalization procedure aims to collapse nearby estimates together and create a sequence of “changesets” that likely correspond to only one changepoint. This allows the estimated changepoints to be reevaluated using a more powerful CUSUM statistics for each changeset and thus corrects the overestimation and inaccuracy issues of the fused lasso.

We first construct the changesets from the fused lasso estimated changepoints. We define the linkage parameter c as the maximal distance any two neighboring estimated changepoints, which were supposed to correspond to a common actual change point, can be apart. That is, we assume the true changepoints are at least $c + 1$ time steps apart. Based on this assumption, we agglomerate the elements of $A(\hat{\theta})$ into sets, that are at least $c + 1$ time steps apart, by linking together estimated changepoints that are within c time steps of each other. Like λ , we optimize c from the data as part of our overall fitting procedure based on BIC minimization (Section A.2).

To illustrate the agglomeration process, suppose we observe Y_1, \dots, Y_{200} and the fused lasso estimates changepoints at the locations

$$A(\hat{\theta}) = \{10, 51, 54, 58, 100, 103, 106, 108, 112, 123, 140, 141\}.$$

We use our fitting procedure to estimate the optimal linkage parameter as $c = 5$. Agglomerating $A(\hat{\theta})$ with linkage parameter $c = 5$ yields the unique set of sets

$$B(\hat{\theta}, c = 5) = \{\{10\}, \{51, 54, 58\}, \{100, 103, 106, 108, 112\}, \{123\}, \{140, 141\}\}.$$

We call the agglomerated changepoints the changesets and denote their collection as $B(\hat{\theta}, c)$. Let $M = |B(\hat{\theta}, c)|$ denote the number of estimated changesets. Each changeset, denoted $b_j \in B(\hat{\theta}, c)$ for $j \in 1, \dots, M$, is intended to represent at most one true changepoint, since clusters of fused lasso based changepoint estimates tend to correspond to a single changepoint and we assumed that each true changepoint is at least $c + 1$ time steps apart. For instance, the set $\{51, 54, 58\}$ represents the true changepoint at $t = 55$, while $\{100, 103, 106, 108, 112\}$ represents the true changepoint at $t = 107$.

The changesets in $B(\hat{\theta}, c)$ sparsify the fused lasso estimate, but they do not precisely estimate changepoints. To convert the changesets into changepoint estimates, we further introduce regionalization. Using the boundaries of the changesets, we split the data into regions assumed to contain AMOC each. We define the j th region, R_j , as the interval:

$$R_j = [\sup b_{j-1} + 1, \inf b_{j+1} - 1], \quad (5)$$

where $b_j \in B(\hat{\theta}, c)$ is the j th changeset, $\sup b_0 = 0$, and $\inf b_{M+1} = n + 1$. Following the above example, $B(\hat{\theta}, c = 5)$ would generate the following regions

$$R_1 = [1, 50], R_2 = [11, 99], R_3 = [59, 122], R_4 = [113, 139], R_5 = [124, 200].$$

Each interval R_j is the largest possible interval that contains b_j and no other changesets. We, therefore, assume that each R_j contains AMOC since the changeset b_j is intended to represent a single changepoint. Note that, since each changeset is at least $c + 1$ timesteps apart, each region has at least $2c + 1$ data points.

We then apply the standard CUSUM test region-wise to identify changepoint within each region. That is, we apply the CUSUM test to $Y_{R_j} = \{Y_{\sup b_{j-1}+1}, \dots, Y_{\inf b_{j+1}-1}\}$ for all $j \in 1, \dots, M$. Unlike the fused lasso, the CUSUM is a uniformly most powerful test for single changepoint detection and it is a consistent estimator of the changepoint location. Because each region is tested independently, resulting in M tests, we adjust the M p -values using the Benjamini–Hochberg procedure (Benjamini & Hochberg, 1995) to control the CUSUM's FDR at a prespecified level α . We retain all estimated changepoints with an adjusted p -value below the prespecified α level.

3 | SIMULATIONS

We investigate MCI's empirical performance in detecting changes in the mean and covariance of a functional process. We consider the case when there are no changepoints, when there are only a few changepoints ("sparse" setting), and when there are many changepoints ("dense" setting). We also consider functional data with light and heavy tails and with symmetric and skewed distributions.

3.1 | Simulation setup

We start by simulating symmetric functional data. We use a GP model and a t -process (TP) model to simulate symmetric light-tailed and heavy-tailed functional data, respectively. Let Z_t denote the symmetric process. We have

$$Z_t = \mu_t + \epsilon,$$

where $\mu_t \in L^2([0, 1])$ is the mean and ϵ follows a zero-mean GP or TP with Matérn covariance function (Stein, 2012),

$$C(x, x') = \frac{\sigma^2 \sqrt{\pi} r^{2\nu}}{2^{\nu-1} \Gamma(\nu + 1/2)} \left(\frac{\|x - x'\|}{r} \right)^\nu K_\nu \left(\frac{\|x - x'\|}{r} \right), \quad (6)$$

with variance parameter σ^2 , range parameter r , smoothness parameter ν , and $K_\nu(\cdot)$ as a modified Bessel function of the 2nd kind of order ν . If ϵ follows a TP, then an additional degrees of freedom parameter df is required; we set $df = 3$ in all simulations. We set the smoothness parameter $\nu = 1$ to ensure mean square differentiability of the sample paths so that the TVN projections are meaningful.

To generate a skewed functional process $\{Y_t, t \in 1, \dots, n\}$ on the domain $[0, 1]$, we let

$$Y_t = \log(1 + e^{Z_t(s)}).$$

We call the transformed GP and transformed t -process a log-sum Gaussian process (LS-GP) and a log-sum t -process (LS-TP), respectively. All results presented in this section are based on LS-GP and LS-TP simulated data because many real datasets, particularly the atmospheric emitted radiance interferometer data we analyze here, are skewed. Also, the asymmetric data represent a more challenging situation for changepoint detection. Results under a GP or TP show similar

TABLE 1 Possible values for the LS-GP (LS-TP) parameters

Parameter	Possible values
μ	$\psi_1, \psi_2, \psi_3, \psi_4, \psi_5$
σ^2	0.50, 0.66, 0.83, 1.00, 1.16, 1.33, 1.50, 1.66, 1.83, 2.00
r	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
ν	1.0

Note: μ indicates the mean function, while σ , r , and ν are scalar valued parameters in the Matérn covariance function. The ψ functions are taken from Chiou et al. (2019) and are defined in Equation (7).

patterns as the LS-GP and LS-TP simulations, but unsurprisingly all changepoints detectors improved under symmetric distributions with light tails. We defer the GP and TP results to Section B in the supplement.

Table 1 provides a list of candidate values for the mean μ , the variance σ^2 , and the range r to take in our simulation. Specific parameter setting are described in Section 3.3. The parameter ν is not varied because it acts on the process in nearly the same way as r . We specify possible mean functions μ , for the LS-GP or LS-TP, using the ψ functions from the simulations in Chiou et al. (2019):

$$\begin{aligned}
 \psi_1(t) &= 5t^2 - \exp(1 - 20t), \\
 \psi_2(t) &= 0.5 - 100(t - 0.1)(t - 0.3)(t - 0.5)(t - 0.9), \\
 \psi_3(t) &= \psi_2(t) + 0.8 \sin(1 + 10\pi t), \\
 \psi_4(t) &= 1 + 3t^2 - 5t^3 + 0.6 \sin(1 + 10\pi t), \\
 \psi_5(t) &= 1 + 3t^2 - 5t^3.
 \end{aligned} \tag{7}$$

Differences between ψ functions, as measured by the L^2 norm, yield different scales of change between the mean functions. The change from ψ_1 to ψ_2 is the largest, ψ_3 to ψ_4 is moderate, and ψ_2 to ψ_3 and ψ_4 to ψ_5 are both small. Additionally, these ψ functions represent changes in both magnitude and the shape of the functional process.

To generate M randomly spaced changepoints in either μ , σ , or r of an i.i.d. LS-GP (or LS-TP) sequence, we first randomly sample $M + 1$ parameter values and $M + 1$ segment lengths (i.e., segment sample sizes). We generically denote the sequence of parameter values as $\theta_1, \dots, \theta_{M+1}$ and the sequence of segment lengths as n_1, \dots, n_{M+1} . Only one parameter is varied at a time, while the others are fixed at prespecified values. To generate the functional time series, we sample n_1 LS-GP's (or LS-TPs) with parameter $\theta = \theta_1$, then n_2 LS-GP's (or LS-TPs) with parameter $\theta = \theta_2$, and so on for the $M + 1$ segments. Each segment is concatenated together so that the boundaries between segments represent changepoints in the functional time series. Parameter values are sampled from Table 1, and sampling is done so that consecutive values are not the same. Segment lengths are samples from either a Unif[500, 1000] or Unif[5000, 10,000] for the dense or sparse setting, respectively.

3.2 | Assessment criterion

We ran 500 simulations for each parameter setting and computed the annotation error (Truong et al., 2019) and the energy distance (Székely, 2003), also called energy error, between the true changepoints and the estimated changepoints. The annotation error, widely used for assessing changepoints detection, measures the difference in the number of detected changepoints and the number of true changepoints. The energy error measures the distance between the set of estimated changepoints and the set of true changepoints. Let $X = \{X_1, \dots, X_n\}$ and $Y = \{Y_1, \dots, Y_m\}$ be two sets. The annotation distance between X and Y is calculated as

$$d_A(X, Y) = |n - m|,$$

and the energy distance between two sets is calculated as

$$d_E(X, Y) = \frac{2}{nm} \sum_{i=1}^n \sum_{j=1}^m |X_i - Y_j| - \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n |X_i - X_j| - \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m |Y_i - Y_j|.$$




FIGURE 2 Annotation error when no changepoints are present in the data. MCI and DSBE nearly always detect 0 changepoints in both the light tail (LS-GP) and the heavy tail (LS-TP) setting. PELT and WBS detect an average of 55 and 25 changepoints, respectively, when no changepoints exist and the functional process is heavy tailed. FPOP performs better than PELT and WBS, but still detects around six false changepoints on average

A low annotation distance (error) means that the algorithm consistently estimates the number of changepoints correctly, while a low energy distance (error) means that the estimated and true changepoints are very similar. In all simulation settings, we compare the different detectors on their ability to achieve low annotation and low energy errors.

Because the DSBE method proposed by Chiou et al. (2019) is the only multiple changepoint detection method that is computationally comparable to MCI, we compare our method to DSBE. Additionally, we include comparisons against three univariate, linear time changepoint detectors applied to the two projections of the TVN and the first functional principal component: PELT algorithm (Killick et al., 2012), the r-FPOP algorithm (Fearnhead & Rigai, 2019), and the WBS procedure (Fryzlewicz, 2014). This will allow us to learn the advantages of our detection method on top of the projections.

For PELT, we used the `cpt.meanvar()` function in the `changepoint` package with default settings and method = "PELT." For r-FPOP, we used the `Rob_seg.std()` function from the `robseg` package with $L1$ loss and tuning parameter $\lambda = 5 \log(n)$. The authors recommended λ to scale with the log of the sample size and the multiplier 5 was found to have the overall best results. For WBS, we used the `wbs()` function, from the `wbs` package, with default settings and changepoints found via `ssc.penalty` (strengthened Schwarz information criterion) minimization. For DSBE, we used the author's provided code with the number of changepoint candidates $K = 50$, minimum segment length $b = 5$, and significance threshold $\alpha = 0.05/K$. DSBE is only applied to the mean change simulations because DSBE was designed to only detect mean changes.

Although we have tried to select the parameters to favor the above-mentioned methods, some default settings of those functions may not lead to optimal results. Nevertheless, We think it is still a fair comparison with our MCI method because we did not tune the parameters to favor our results. Instead, we estimated the λ and c parameters each time using a quick grid search (Section A.2). We set $\alpha = 0.001$ to keep the false positive rate low and avoid many false changepoints, similar to what was done for DSBE, which controls $\alpha = 0.05/K$. It is a common practice for changepoint detection to control the false positive rate at a low level. We set $\alpha = 0.001$, but users can choose other small α values though extremely small α may lead to missing changepoints. We conducted a small simulation study (Section B.2 in the supplement) to test the sensitivity of MCI to α . We found that MCI's power is generally insensitive to α as long as α is small, say smaller than 0.001.

Note that we are not evaluating the three methods, PELT, FPOP, and rFPOP in the general settings for univariate time series. The results below only show their performance for projections of skewed and heavy tailed functional data, on which we focus. Conclusions from the following simulations should not be extended to the general univariate setting.

3.3 | Assessment results

3.3.1 | No changepoints

We first consider the situation where there are no changepoints, that is, $M = 0$. Functional observations are generated from either a LS-GP or LS-TP with constant mean function $\mu = 0$ and Matérn covariance with $\sigma = 1$, $r = 0.2$, and $\nu = 1$. Figure 2 shows distribution of annotation errors under the LS-GP and LS-TP for each method.

All methods have an almost uniformly zero annotation error under the LS-GP, meaning that all methods have an essentially zero false positive rate. However, under the LS-TP, only MCI and DSBE maintain their almost uniformly

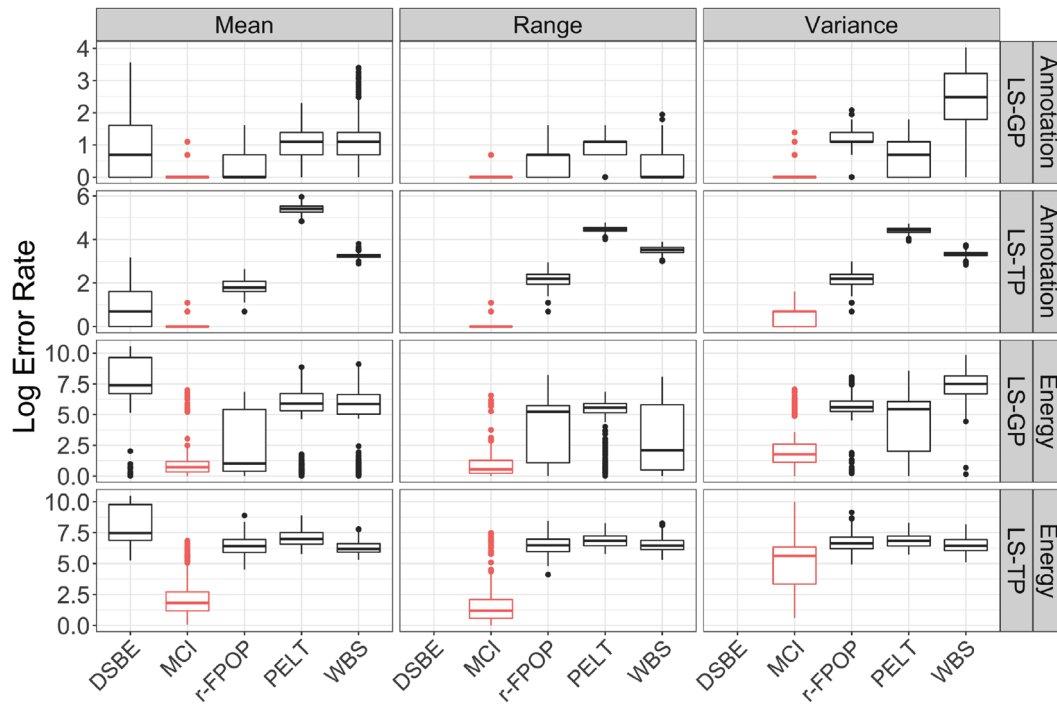


FIGURE 3 Annotation and energy error rates for sparse changes in the mean function, range, and variance. Both error types are plotted on the $\log(1 + \text{error})$ scale. Under both the log-sum Gaussian process (LS-GP) and the log-sum t -process (LS-TP), the MCI method has far lower error rates than the other methods for detecting all changepoint types

near-zero error rates, while the annotation error for FPOP, PELT, and WBS all increases dramatically. PELT and WBS have particularly high annotation error rates, possibly due to their sensitivity to larger random fluctuations caused by the heavy tailed generating process. FPOP seems more robust than PELT and WBS, likely because FPOP uses the robust $L1$ loss function, but it is still less robust than either DSBE or MCI.

3.3.2 | Sparse changepoints

We next consider the situation when the changepoints are relatively far apart, that is, the changepoints are sparse. We simulate data as described in Section 3.1 with $M = 5$ changepoints and segment lengths sampled from $\text{Unif}[5000, 10,000]$. To study changepoints in μ , we fix the covariance parameters to $\sigma = 1$, $\nu = 1$, and $r = 0.2$. For changepoints in variance, we fix $\mu = 0$, $\nu = 1$, and $r = 0.2$, and for changepoints in range we fix $\mu = 0$, $\nu = 1$, and $\sigma = 1$.

Figure 3 summarizes each method's annotation error and energy error in detecting changes in the mean, range, and variance. Our MCI method maintains the overall lowest annotation error and the overall lowest energy error. Together, this shows that MCI improves changepoint estimation in terms of both the number and location, whether the error process is light or heavy tailed. The univariate methods applied to our projections have reasonably low annotation error rates on the LS-GP. However, they show high annotation error rates on the LS-TP due to overestimation, while their energy error rates are high on both. DSBE sees the most negligible deterioration from LS-GP to LS-TP compared to the alternative methods, although its energy error rates were already high under LS-GP.

3.3.3 | Dense changepoints

Finally, we consider the situation when the changepoints are relatively close to each other, that is, when changepoints are dense. To simulate data with dense changepoints, we set $M = 50$ and sample segment lengths from $\text{Unif}[500, 1000]$. This design results in changepoints that are, on average, ten times as dense as the sparse changepoints in Section 3.3.2. The parameter setting for studying changes in the mean, variance, and range are the same as in Section 3.3.2.




FIGURE 4 Annotation and energy error rates for dense changepoints in the mean, range, and variance parameters. Error values are plotted on the $\log(1 + \text{error})$ scale. MCI performs the best for detecting mean and variance changes in the LS-GP and for detecting range changes in the LS-TP. No method dominates for detecting variance changes in the LS-TP

We summarize each method's ability to detect dense mean, range, and variance changes in Figure 4. MCI again attains the lowest overall annotation and energy error rates. However, we find that the gap between MCI and other methods is smaller than when the changepoints were sparse. On the one hand, this is because MCI's error rates increase due to the CUSUM having fewer data per segment to estimate the location of changepoints. On the other hand, the competing methods see a decrease in their error rates for two reasons. One is that PELT and FPOP tend to estimate many changepoints, so when the actual number of changepoints is also high, their annotation and energy errors will naturally drop. The other reason is that WBS uses randomly sized intervals and a binary segmentation algorithm to find changepoints, both of which make the method less sensitive to changepoints density.

In summary, Figures 3 and 4 show that MCI is generally more accurate and more skillful for a broader type of changepoints than the other approaches. This is because MCI effectively corrects the overestimation of fused lasso solution through changeset regionalization and application of CUSUM. Under the light tailed LS-GP data, MCI has the lowest annotation and energy error rates across all changepoint types and densities. However, the three other algorithms' performance varies greatly depending on the changepoints' type and density. With the heavy tailed LS-TP data, the error rate of MCI increases across all changepoint detection compared to LS-GP data. Nevertheless, the other methods show more deteriorated skill for heavy-tailed data.

3.4 | Linear computation

Our method's computational complexity grows linearly with sample size because each sub-step of the method grows linearly in time, and the number of sub-steps does not grow with sample size. For the FPC decomposition, we use the FACE algorithm (Xiao et al., 2016), which is linear in sample size and function length. Computing the TVN of each function requires only a single pass over the data, so the TVN projection is linear as well. For fused lasso estimation, we use the Condat algorithm (Condat, 2013), which is linear in time. Refining the segmentation over the data with changeset regionalization can be done with a single iteration. Finally, the CUSUM is computable in linear time, and the overlapping regions only cause CUSUM to be run (nearly) twice over the data set. The optimization procedure for λ and c uses a grid search over a fixed sequence of λ and c values, ensuring that the optimization step is also linear in sample size.




FIGURE 5 Median run time of MCI (black points) under seven sample size settings (10,000–70,000 observations). Trend line (red) shows the claimed linear relationship between sample size and computational time. Simulations were performed on single core of a 2017 MacBook Pro with a 2.33 GHz i5 processor

To demonstrate MCI's linear computation time, we conduct a simulation study to empirically assess the run time growth of MCI over an increasing sequence of sample sizes. We take the simulation with mean zero GP and covariance parameters $\sigma = 1$, $r = 0.2$, and $\nu = 1$ as example. We consider sample sizes from $n = 1000$ to $n = 7000$ in 1000 unit increments, and run simulation 1000 times for each sample size. We show the run time of MCI under each of the seven sample sizes in Figure 5. The run time of MCI exhibits a linear trend with the sample size.

4 | APPLICATION TO PROFILES OF WATER VAPOR

The water vapor mixing ratio is the water vapor density over the dry air density in a given atmospheric unit. It is an important variable in Meteorology for distinguishing individual air masses, monitoring the effects of soil evapotranspiration and large water body evaporation (North et al., 2014), and for the early detection of heavy precipitation events (Sakai et al., 2019). Changepoint detection helps identify sudden changes in an air parcel's water vapor content due to precipitation events and air parcels mixing. Retrospectively identifying sudden changes in the water vapor profiles' structure is often necessary before constructing statistical models for identifying precipitation events.

We apply our functional changepoint detector to the water vapor mixing ratio profiles collected from the Atmospheric Emitted Radiance Interferometer (AERI) instrument at the Lamont, Oklahoma Facility. The AERI instruments are maintained by the U.S. Department of Energy's (DOE) Atmospheric Radiation Measurement (ARM) Program to collect high-resolution atmospheric profile data (Stokes & Schwartz, 1994). The raw data are openly available in "aeri01prof3feltz" at <http://dx.doi.org/10.5439/1027271>. In this dataset, each profile consists of 58 measurements of the water vapor mixing ratio along a single atmospheric column from 0 to 44,000 m above ground level in Lamont, Oklahoma. Complete profiles were collected every 8 min, thus providing near-continuous monitoring of atmospheric conditions. We removed the top 18 altitude points representing the 11,000–44,000 m range due to the extremely high rate of measurement errors in this range. Therefore, we only consider the 40 measurements from 0 to 10,000 m.

For our analysis, we consider the entire time series of water vapor profiles from January 4, 2007 to March 10, 2014. This period corresponds to 234,062 profiles, each sampled at the same 40 altitudes. To illustrate the data, we plotted profiles of water vapor mixing ratios from August 9 through September 6 in 2008 in Figure 6. Each vertical line represents an individual profile, with colors indicating the value of the profile at each altitude. Abrupt increases and decreases in water vapor along time are visible, indicating rapid changes from high density to low density and vice versa. The changes could be caused by sudden precipitation events or air mass mixing. We examined the marginal distribution of this data and found non-Gaussian behavior, including heavy tails and kurtosis.

4.1 | Identification of changepoints

We applied our MCI method to the water vapor mixing ratio profiles and found 210 changepoints. Figure 7 shows four examples of the changepoints identified in the first functional principal component projection (P_{E1}), and Figure 8 shows four examples of changepoints identified in the TVN projection (P_T).




FIGURE 6 AERI water vapor mixing ratio profiles from August 9 to September 6, 2008. The x -axis is time, the y -axis is altitude, and the color values indicate the density of water vapor at each altitude and time. Each vertical line is an entire function

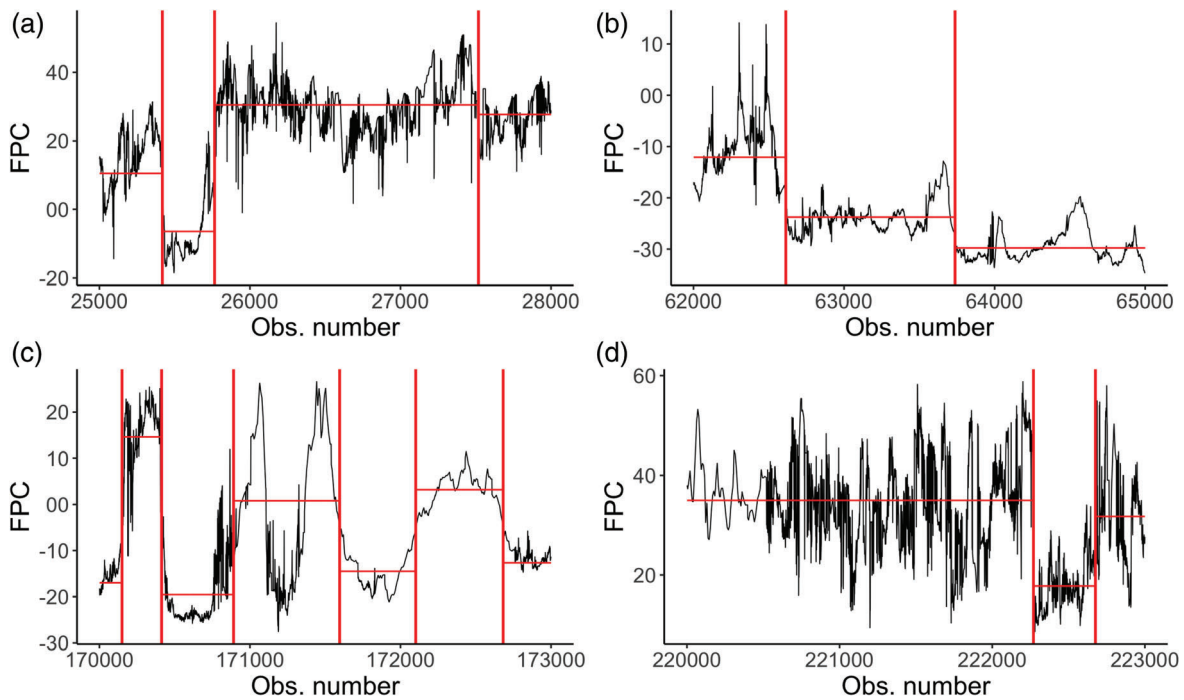


FIGURE 7 Four examples of the changepoints detected by MCI in the P_{E1} projection of the AERI profiles. The black curve is the projected values, the red vertical lines mark the estimated changepoints, and the red horizontal lines indicate the segment means between changepoints

Several commonalities between the two plots are apparent. The first is that panels A and D in both figures are highly similar in appearance, and the changepoints identified in these regions are similar between the two projections. This happens because the P_{E1} and P_T are not necessarily orthogonal to each other, and changepoints may manifest in both low and high frequency spectrums. Another common feature is that MCI is robust to independence violations and heavy tails in P_{E1} and P_T . This can be seen in the right half of panel C in both Figures 7 and 8, where both time series exhibit an autoregressive structure, yet MCI does not seem to yield overly dense changepoints that might be caused by the correlation. Heavy tailed behavior can be observed in panels A, B, and D, where large spikes in the time series are observed. MCI does not detect these anomalies as changepoints in either projection.

Figures 7 and 8 also show that there are many differences between the two projections, meaning that P_{E1} and P_T measure very different aspects of the data. For example, panels B and C are almost entirely different in the two plots, including the locations of the detected changepoints. MCI was, therefore, able to pick up on a broader range of changepoints than P_{E1} would allow because of the P_T projection.

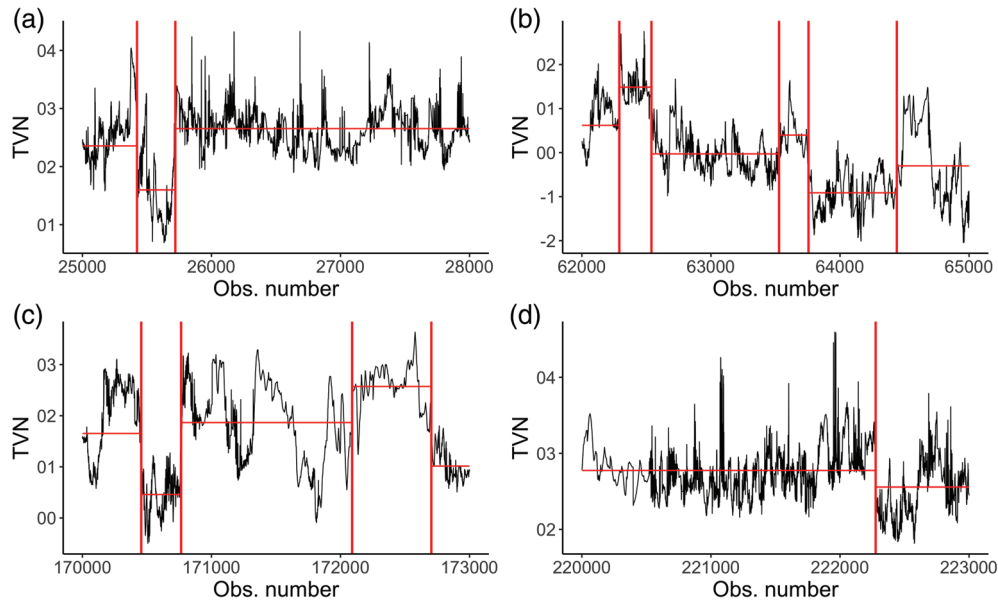


FIGURE 8 Four examples of the changepoints detected by MCI in the P_T projection of the AERI profiles. The black curve is the projected values, the red vertical lines mark the estimated changepoints, and the red horizontal lines indicate the segment means between changepoints

5 | DISCUSSION

We propose the MCI method for detecting multiple changepoints of a functional time series. The changes can be either in the mean or in the dependence structure of the functional data. We first introduce a minimal system of projections to represent the variability “between” and “within” each function. Motivated by the non-normal behavior of the TVN projection, we then introduce an augmented fused lasso based strategy to robustly segment the time series into regions likely containing at most one changepoint. Finally, CUSUM is applied region-wise to detect and identify each region’s potential changepoint. Our extensive simulations show that the MCI method is accurate, computationally efficient, and robust to the underlying data distribution. Finally, we demonstrate MCI on water vapor mixing ratios over time.

Our two projections, the total variation norm and the first FPC, efficiently represent the major variability “within” and “between” each function, respectively. Our method is, therefore, able to detect a broad range of changepoints stemming from changes in the mean and covariance structure of the data. In contrast, entirely functional metric based approaches are, in general, only powerful against changes in the mean (Aue et al., 2018). Our minimal data reduction, compared with using many FPCs, also has computational and theoretical benefits. The computational burden is lessened by only needing to compute a single FPC, and we conduct far fewer tests than if we had used numerous FPCs. Furthermore, conducting fewer tests helps our detector maintain higher power since more testing means harsher multiple testing corrections.

Our changepoint detection differs from the existing fused lasso based approaches in two significant ways. First, we used the fused lasso only as a segmentation procedure to identify regions likely to contain only at most one changepoint. Second, changepoint estimation is conducted via CUSUM testing and is not based on the specific jumps in the fused lasso estimate. This strategy is more efficient and powerful than directly using the fused lasso for changepoint detection since the fused lasso is shown only to be ϵ -consistent (Rojas & Wahlberg, 2014) and CUSUM is the UMP test.

In future work, we would like to consider the theoretical underpinnings of the MCI method more rigorously. For instance, it remains to be shown whether the MCI is a consistent estimator of the changepoints or whether the MCI is asymptotically powerful, although our simulations seem to imply both properties. We may further study the MCI’s robustness and compare its current form with alternatives using robust changepoint statistics rather than CUSUM. The procedure to optimize the tuning parameters λ and c (Section A.2) is also heuristic and not guaranteed to find globally optimal parameters. A better algorithm with optimality guarantees may be possible. However, we found through testing that the basin of optimal solutions is generally quite large, so more precise estimation algorithms may not be necessary.

Finally, we may also consider extensions to higher dimensional functional processes, such as multivariate functional time, continuous surfaces, and spatial fields.

ACKNOWLEDGEMENTS

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This work was also supported in part by the National Science Foundation through the awards NSF-DMS-1830312, NSF-DGE-1922758 and NSF-DMS-2124576.

REFERENCES

- Adelfio, G. (2012). Change-point detection for variance piecewise constant models. *Communications in Statistics-Simulation and Computation*, 41(4), 437–448.
- Aston, J. A., & Kirch, C. (2012). Detecting and estimating changes in dependent functional data. *Journal of Multivariate Analysis*, 109, 204–220.
- Aue, A., Gabrys, R., Horváth, L., & Kokoszka, P. (2009). Estimation of a change-point in the mean function of functional data. *Journal of Multivariate Analysis*, 100(10), 2254–2269.
- Aue, A., Rice, G., & Sönmez, O. (2018). Detecting and dating structural breaks in functional data without dimension reduction. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3), 509–529.
- Aue, A., Rice, G., & Sönmez, O. (2020). Structural break analysis for spectrum and trace of covariance operators. *Environmetrics*, 31(1), e2617.
- Beaulieu, C., Killick, R., Ireland, D., & Norwood, B. (2020). Considering long-memory when testing for changepoints in surface temperature: A classification approach based on the time-varying spectrum. *Environmetrics*, 31(1), e2568.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300.
- Berkes, I., Gabrys, R., Horváth, L., & Kokoszka, P. (2009). Detecting changes in the mean of functional observations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(5), 927–946.
- Bosq, D. (2012). *Linear processes in function spaces: Theory and applications* (Vol. 149). Springer Science & Business Media.
- Boysen, L., Kempe, A., Liebscher, V., Munk, A., & Wittich, O. (2009). Consistencies and rates of convergence of jump-penalized least squares estimators. *The Annals of Statistics*, 37(1), 157–183.
- Chan, N. H., Yau, C. Y., & Zhang, R.-M. (2014). Group lasso for structural break time series. *Journal of the American Statistical Association*, 109(506), 590–599.
- Chapman, J.-L., Eckley, I., & Killick, R. (2020). A nonparametric approach to detecting changes in variance in locally stationary time series. *Environmetrics*, 31(1), e2576.
- Chiou, J.-M., Chen, Y.-T., & Hsing, T. (2019). Identifying multiple changes for a functional data sequence with application to freeway traffic segmentation. *The Annals of Applied Statistics*, 13(3), 1430–1463.
- Condat, L. (2013). A direct algorithm for 1-d total variation denoising. *IEEE Signal Processing Letters*, 20(11), 1054–1057.
- Donoho, D. L., & Johnstone, I. M. (1998). Minimax estimation via wavelet shrinkage. *The Annals of Statistics*, 26(3), 879–921.
- Fearnhead, P., & Rigall, G. (2019). Changepoint detection in the presence of outliers. *Journal of the American Statistical Association*, 114(525), 169–183.
- Finazzi, F., Fassò, A., Madonna, F., Negri, I., Sun, B., & Rosoldi, M. (2019). Statistical harmonization and uncertainty assessment in the comparison of satellite and radiosonde climate variables. *Environmetrics*, 30(2), e2528.
- Fortuna, F., Gattone, S. A., & Di Battista, T. (2020). Functional estimation of diversity profiles. *Environmetrics*, 31(8), e2645.
- Fryzlewicz, P. (2014). Wild binary segmentation for multiple change-point detection. *The Annals of Statistics*, 42(6), 2243–2281.
- Gromenko, O., Kokoszka, P., & Reimherr, M. (2017). Detection of change in the spatiotemporal mean function. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(1), 29–50.
- Guntuboyina, A., Lieu, D., Chatterjee, S., & Sen, B. (2017). Adaptive risk bounds in univariate total variation denoising and trend filtering. *arXiv preprint arXiv:1702.05113*.
- Hörmann, S., & Kokoszka, P. (2010). Weakly dependent functional data. *The Annals of Statistics*, 38(3), 1845–1884.
- Hörmann, S., & Kokoszka, P. (2012). *Functional time series*. In T Subba Rao (Ed.), *Handbook of statistics* (Vol. 30, pp. 157–186). Elsevier.
- Hyndman, R. J., & Booth, H. (2008). Stochastic population forecasts using functional data models for mortality, fertility and migration. *International Journal of Forecasting*, 24(3), 323–342.
- Hyun, S., G'Sell, M., & Tibshirani, R. J. (2016). Exact post-selection inference for changepoint detection and other generalized lasso problems. *arXiv preprint arXiv:1606.03552*.
- Jarušková, D., & Antoch, J. (2020). Changepoint analysis of klementinum temperature series. *Environmetrics*, 31(1), e2570.
- Killick, R., Fearnhead, P., & Eckley, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500), 1590–1598.

- Kokoszka, P., & Zhang, X. (2012). Functional prediction of intraday cumulative returns. *Statistical Modelling*, 12(4), 377–398.
- Lee, J., Lund, R., Woody, J., & Xu, Y. (2020). Trend assessment for daily snow depths with changepoint considerations. *Environmetrics*, 31(1), e2580.
- Li, X. & Ghosal, S. (2018). Bayesian change point detection for functional data. *arXiv preprint arXiv:1808.01236*.
- Lin, K., Sharpnack, J., Rinaldo, A., & Tibshirani, R. J. (2016). Approximate recovery in changepoint problems, from ℓ_2 estimation error rates. *arXiv preprint arXiv:1606.06746*.
- Louizos, C., Welling, M., & Kingma, D. P. (2018). *Learning sparse neural networks through l_0 regularization*. Proceedings of the International Conference on Learning Representations.
- MacEachern, S. N., Rao, Y., & Wu, C. (2007). A robust-likelihood cumulative sum chart. *Journal of the American Statistical Association*, 102(480), 1440–1447.
- Maidstone, R., Hocking, T., Rigai, G., & Fearnhead, P. (2017). On optimal multiple changepoint algorithms for large data. *Statistics and Computing*, 27(2), 519–533.
- Mammen, E., & van de Geer, S. (1997). Locally adaptive regression splines. *The Annals of Statistics*, 25(1), 387–413.
- Muggeo, V. M., & Adelfio, G. (2011). Efficient change point detection for genomic sequences of continuous measurements. *Bioinformatics*, 27(2), 161–166.
- North, G. R., Pyle, J. A., & Zhang, F. (2014). *Encyclopedia of atmospheric sciences* (Vol. 1). Elsevier.
- Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41(1/2), 100–115.
- Qu, Z., Dai, W., & Genton, M. G. (2021). Robust functional multivariate analysis of variance with environmental applications. *Environmetrics*, 32(1), e2641.
- Ramsay, J. O. (2004). Functional data analysis. *Encyclopedia of Statistical Sciences*, 4, 85–109.
- Rice, G. & Zhang, C. (2019). Consistency of binary segmentation for multiple change-points estimation with functional data. *arXiv preprint arXiv:2001.00093*.
- Rojas, C. R. & Wahlberg, B. (2014). On change point detection using the fused lasso method. *arXiv preprint arXiv:1401.5408*.
- Sakai, T., Nagai, T., Izumi, T., Yoshida, S., & Shoji, Y. (2019). Automated compact mobile Raman lidar for water vapor measurement: Instrument description and validation by comparison with radiosonde, GNSS, and high-resolution objective analysis. *Atmospheric Measurement Techniques*, 12(1), 313–326.
- Shang, H. L., & Hyndman, R. J. (2011). Nonparametric time series forecasting with dynamic updating. *Mathematics and Computers in Simulation*, 81(7), 1310–1324.
- Shao, X., & Zhang, X. (2010). Testing for change points in time series. *Journal of the American Statistical Association*, 105(491), 1228–1240.
- Sharipov, O., Tewes, J., & Wendler, M. (2016). Sequential block bootstrap in a hilbert space with application to change point analysis. *Canadian Journal of Statistics*, 44(3), 300–322.
- Srivastava, A., & Klassen, E. P. (2016). *Functional and shape data analysis*. Springer.
- Stein, M. L. (2012). *Interpolation of spatial data: Some theory for kriging*. Springer Science & Business Media.
- Stokes, G. M., & Schwartz, S. E. (1994). The atmospheric radiation measurement (ARM) program: Programmatic background and design of the cloud and radiation test bed. *Bulletin of the American Meteorological Society*, 75(7), 1201–1222.
- Székel, G. J. (2003). E-statistics: The energy of statistical samples. *Bowling Green State University, Department of Mathematics and Statistics Technical Report*, 3(05), 1–18.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., & Knight, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1), 91–108.
- Tibshirani, R. J., & Taylor, J. (2011). The solution path of the generalized lasso. *The Annals of Statistics*, 39(3), 1335–1371.
- Truong, C., Oudre, L., & Vayatis, N. (2019). Selective review of offline change point detection methods. *Signal Processing*, 167, 107299.
- Tucker, J. D., Wu, W., & Srivastava, A. (2013). Generative models for functional data using phase and amplitude separation. *Computational Statistics & Data Analysis*, 61, 50–66.
- Woodall, W. H. (2007). Current research on profile monitoring. *Production*, 17(3), 420–425.
- Xiao, L., Zipunnikov, V., Ruppert, D., & Crainiceanu, C. (2016). Fast covariance estimation for high-dimensional functional data. *Statistics and Computing*, 26(1-2), 409–421.
- Zhang, X., Shao, X., Hayhoe, K., & Wuebbles, D. J. (2011). Testing the structural stability of temporally dependent functional observations and application to climate projections. *Electronic Journal of Statistics*, 5, 1765–1796.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Harris, T., Li, B., & Tucker, J. D. (2021). Scalable multiple changepoint detection for functional data sequences. *Environmetrics*, e2710. <https://doi.org/10.1002/env.2710>

APPENDIX A

A.1 CUSUM statistic

The CUSUM statistic is a powerful and classical statistic for detecting a *single* change in the mean of a univariate time series. (MacEachern et al., 2007; Page, 1954). Let $Y = \{Y_t, t \in 1, \dots, n\}$ generically denote a univariate sequence. We can use the CUSUM statistic to test for a single change in the mean of the univariate sequence. That is, we can test the hypothesis that

$$\begin{aligned} H_0 : E(Y_1) &= \dots = E(Y_n), \\ H_A : E(Y_1) &= \dots = E(Y_k) \neq E(Y_{k+1}) = \dots = E(Y_n), \end{aligned}$$

for an unknown $1 < k < n$. A CUSUM process is defined as

$$T_n(\lfloor nr \rfloor) = n^{-1/2} \sum_{t=1}^{\lfloor nr \rfloor} (Y_t - \bar{y}_n), \quad r \in [0, 1],$$

where $\bar{y}_n = n^{-1} \sum_{t=1}^n Y_t$. The CUSUM statistic is defined as the supremum of the scaled CUSUM process

$$G_n = \sup_{r \in [0, 1]} |T_n(\lfloor nr \rfloor) / \hat{\sigma}_n|, \quad (\text{A1})$$

where $\hat{\sigma}_n^2$ is a consistent estimator of the long run variance $\sigma^2 = \lim_{n \rightarrow \infty} \text{nvar}(\bar{X}_n)$. We approximate $\hat{\sigma}_n^2$ with $\text{Var}(Y - \hat{\theta})$, where $\hat{\theta}$ a nonparametric estimate of the mean of Y . Under H_0 ,

$$G_n \xrightarrow{D} \sup_{r \in [0, 1]} |B(r) - rB(1)|,$$

where $B(\cdot)$ is a Brownian motion so $B(r) - rB(1)$ is a Brownian bridge on $[0, 1]$. Critical values for $\sup_{r \in [0, 1]} |B(r) - rB(1)|$ can be computed using the well-known Kolmogorov distribution (Shao & Zhang, 2010):

$$P(G_n < t) = 1 - 2 \sum_{j=1}^{\infty} (-1)^j e^{-2j^2 t^2}. \quad (\text{A2})$$

If the test rejects H_0 , then the estimated changepoint location is

$$\hat{k} = n \arg \sup_{r \in [0, 1]} |T_n(\lfloor nr \rfloor)|. \quad (\text{A3})$$

The CUSUM estimator (A3) is highly powerful, in fact uniformly most powerful, for detecting a *single* change in the mean of a univariate time series. However, in our formulation, we need to detect *multiple* changes in the mean and covariance of functional data.

As mentioned in Section 1, existing strategies for multiple changepoint detection include dynamic segmentation and binary segmentation to augment CUSUM. However, these strategies may not be robust to heavy tailed or asymmetric data, such as our data (Figure 6). Another avenue for multiple changepoint detection, not relying on CUSUM, are the fused lasso (Tibshirani & Taylor, 2011) based strategies (Chan et al., 2014; Hyun et al., 2016; Lin et al., 2016; Rojas & Wahlberg, 2014). These methods use the breaks or jumps in a fused lasso estimate to detect multiple changepoints. The fused lasso is more robust to heavy tails and anomalies, but its power is suboptimal compared with CUSUM, and it typically overestimates the number of changepoints (Chan et al., 2014; Fryzlewicz, 2014).

In Section 2, we will introduce a new strategy for multiple changepoint that combines the strengths of the CUSUM and fused lasso into a single procedure for detecting changes mean and covariance of functional time series.

A.2 BIC optimization

As mentioned in Section 2, we need optimize out the fused lasso penalty term λ and the linkage parameter c . Following the methodology in Chan et al., (2014), we consider Bayesian information criteria (BIC) minimization. Given a set of

estimated changepoints τ_1, \dots, τ_m we define θ^* as

$$\theta_t^* = \sum_{i=0}^M \mathbb{1}(\tau_i < t \leq \tau_{i+1}) \left[\frac{1}{\tau_{i+1} - \tau_i} \sum_{t=\tau_i}^{\tau_{i+1}} Y_t \right],$$

where $\tau_0 = 1$, $\tau_{M+1} = n$, and n is the sample size. We define $S_n(\tau_{i-1}, \tau_i) = \sum_{t=\tau_i}^{\tau_{i+1}} (Y_t - \theta_t^*)^2$ and $S_n(\tau_1, \dots, \tau_m) = \sum_{i=1}^m S_n(\tau_{i-1}, \tau_i)$. Finally, we can define the BIC of the estimated changepoints as

$$BIC(\lambda, c) = S_n(\tau_1, \dots, \tau_m) + m \log(n),$$

and estimate λ and c by solving

$$\hat{\lambda}, \hat{c} = \arg \min_{\lambda, c} BIC(\lambda, c). \quad (\text{A4})$$

Minimizing Equation (A4) is a proxy for finding θ^* with a minimal L^0 norm (Louizos et al., 2018), or, equivalently, a minimal set of changepoints, τ_1, \dots, τ_m , sufficient to explain the variation in the projected data.

The objective in Equation (A4) is non-differentiable, so we use grid search to brute force estimating $\hat{\lambda}, \hat{c}$. We restrict our search to multiples of \sqrt{n} for each parameter. That is, we allow each parameter to grow as a multiple of the square root of the functional time series' length, $\lambda = r\sqrt{n}$ and $c = k\sqrt{n}$. In our implementation, we allow r and c to range from 0.2 to 5 in increments of 0.1.

The \sqrt{n} scaling for λ was used in Lin et al. (2016) for their fused lasso estimator to ensure the solution is well estimated while still filtering out a majority of the extraneous changepoints. Rojas and Wahlberg (2014) showed that c depends on sample size n as $c = kn$. However, directly using n made the numerical estimation of c challenging due to the small magnitudes of k . Instead, we made c scale with \sqrt{n} and found this to work better in our simulations.