RESEARCH ARTICLE

WILEY

Elastic functional changepoint detection of climate impacts from localized sources

J. Derek Tucker⦿ | **Drew Yarger**⦿

Statistical Sciences, Sandia National
Laboratories, Albuquerque, New Mexico,
USA

Correspondence
J. Derek Tucker, Statistical Sciences,
Sandia National Laboratories,
Albuquerque, NM, USA.
Email: jdtuck@sandia.gov

Funding information
U.S. Department of Energy, Grant/Award
Number: DE-NA0003525

Abstract

Detecting changepoints in functional data has become an important problem as interest in monitoring of climate phenomenon has increased, where the data is functional in nature. The observed data often contains both amplitude (y-axis) and phase (x-axis) variability. If not accounted for properly, true changepoints may be undetected, and the estimated underlying mean change functions will be incorrect. In this article, an elastic functional changepoint method is developed which properly accounts for these types of variability. The method can detect amplitude and phase changepoints which current methods in the literature do not, as they focus solely on the amplitude changepoint. This method can easily be implemented using the functions directly or can be computed via functional principal component analysis to ease the computational burden. We apply the method and its nonelastic competitors to both simulated data and observed data to show its efficiency in handling data with phase variation with both amplitude and phase changepoints. We use the method to evaluate potential changes in stratospheric temperature due to the eruption of Mt. Pinatubo in the Philippines in June 1991. Using an epidemic changepoint model, we find evidence of a increase in stratospheric temperature during a period that contains the immediate aftermath of Mt. Pinatubo, with most detected changepoints occurring in the tropics as expected.

KEYWORDS

changepoint analysis, functional data analysis, functional principal component analysis

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 (Fortuna et al., 2020). A functional time series is a sequence of functions (i.e., infinite dimensional objects), 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). In Figure 1, we plot daily stratospheric temperatures at one location as a functional time series, one of the motivating data sets for this work. In each year, we observe a function that describes the temperature over the course of that year. Although not obvious in Figure 1, stratospheric temperature trends experienced a significant change from normal following the 1991 Mt. Pinatubo eruption.
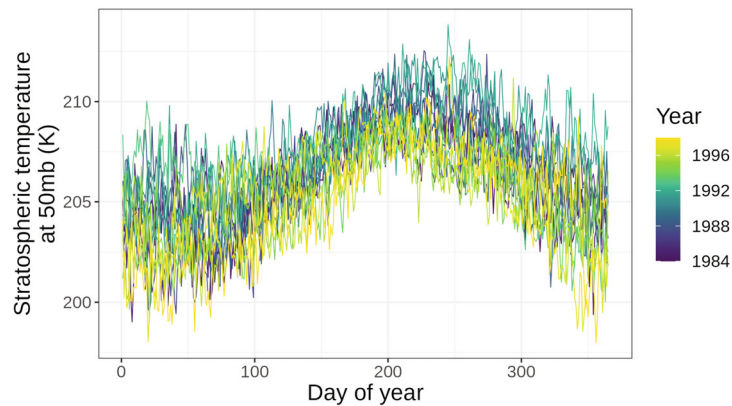
FIGURE 1 Daily stratospheric temperature data presented as a functional time series at 115°E and 10°N, years 1984–1998 obtained from the Modern-Era Retrospective analysis for Research and Applications, version 2 (MERRA-2) reanalysis database (Global Modeling and Assimilation Office (GMAO), 2015).

Similar to univariate and multivariate time series, a functional time series can experience abrupt changes in its data-generating mechanism. These abrupt changes are known as changepoints and can complicate further analysis by invalidating stationarity assumptions. They also can lead to interesting analysis by determining when data has patterns which are not homogeneous over time, or when a forcing has caused the data-generation mechanism to change. We consider the problem of detecting changepoints after observing an entire sequence of functions rather than considering a "streaming" or "real-time" changepoint detection problem. That is, we are interested in evaluating (1) whether a change occurred in the functional time series or not and (2) if a change was detected, when the change occurred and how the data-generating mechanism changed. This is particularly useful for changepoints that can traverse multiple years, which is the case for our motivating problem.

Current methods in the Functional Data Analysis (FDA) literature on changepoint detection have been solely focused on cross-sectional analysis, or amplitude variability only, and thus completely ignore phase variability that often exist in the data, especially climate data. 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). 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. Sharipov et al. (2016) similarly developed a sequential block bootstrap procedure for these methods. More 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 change in the mean without losing information due to functional principal component analysis (fPCA) truncation.

Phase variability could be just a nuisance parameter or actually contain information. In the latter case, if phase variability is ignored, the distance metric between functions will be degenerate and the estimated underlying mean function will not be representative of the data-generating mechanism (Srivastava & Klassen, 2016; Tucker et al., 2013). An example of this is shown in Figure 2. In Panel (b), it is clear that the mean is not representative of the underlying data-generating mechanism demonstrated in Panel (a). After alignment shown in panel (c), the mean in Panel (d) accounts for phase variability and is representative of the underlying data-generating mechanism.

In this article, we present a functional changepoint method that accounts for the amplitude and phase variability. We demonstrate two different approaches for changepoint detection in this setting: (1) by utilizing the appropriate fPCA (Lee & Jung, 2017; Tucker et al., 2013) or (2) by utilizing a fully functional approach similar to Aue et al. (2018) where the mean is computed using the Karcher mean under the Fisher-Rao metric (Srivastava & Klassen, 2016, and described later). This Fisher-Rao metric is a proper distance and is *elastic* (that is, it is invariant to random warping). Through this metric we separate out the amplitude and phase variability for subsequent changepoint analysis. This allows for the identification of changepoints where the underlying mean changes in amplitude, phase, or a combination of both. We evaluate performance in two simulation studies, showing that the proposed test statistics properly identify changes and perform appropriately under the null hypothesis.

A motivation for developing this methodology is to introduce methodology to evaluate changes the temporal pattern of stratospheric temperatures after the eruption of Mt. Pinatubo in June 1991, as presented in Figure 1. Using the
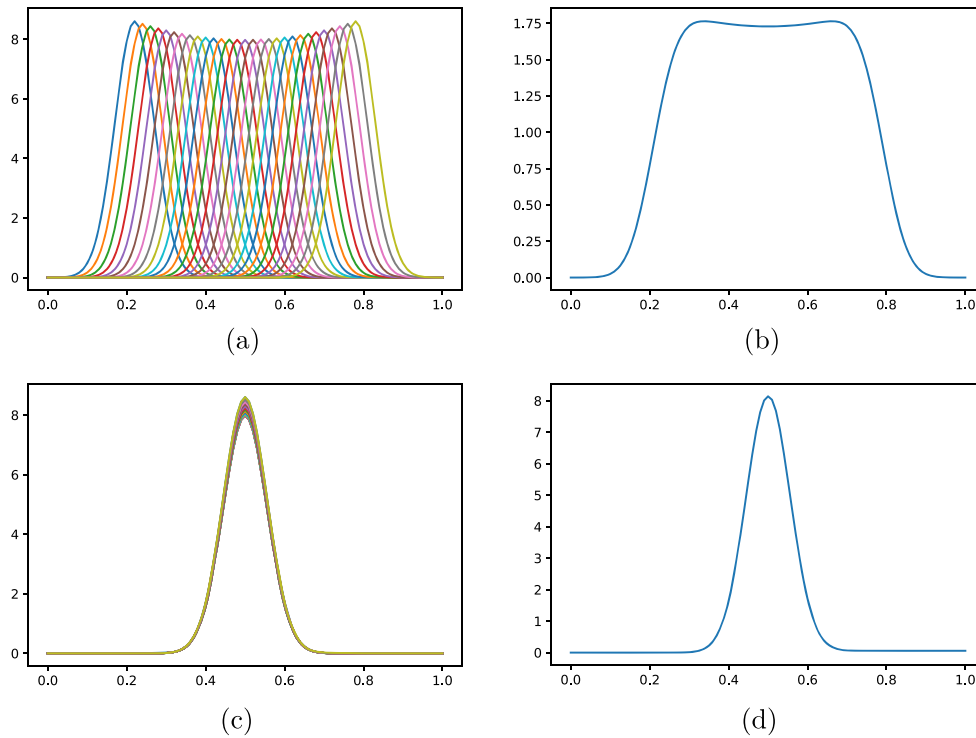
FIGURE 2 Example of using proper mean calculation on functional data. Panel (a) presents a set of functional data that contains a large amount of phase variability. Panel (b) shows the computed mean if the phase variability is ignored. Panel (c) shows the alignment of the data using the method in Tucker et al. (2013) and the resulting Karcher mean in Panel (d). (a) Original data. (b) Original mean. (c) Aligned data. (d) Karcher mean.

years 1984–1998, we aim to estimate a change in this sequence of functional data during these $n = 15$ years. There is evidence that the eruption of Mt. Pinatubo and its injection of aerosols into the atmosphere increased temperatures in the stratosphere, leading to many downstream shocks to the climate system (Robock, 2003). Climate data is known to exhibit strong phase variability in the form of transient or cyclic variability from year to year. In our analysis of stratospheric temperature, we find that controlling for phase variability isolates changes in the amplitude while filtering out transient weather variability from different years. When identifying changes in climate data, controlling for seasonality is an important aspect of many analyses; see Reeves et al. (2007) and more recently Lund et al. (2022) for an application, overview, and common challenges for detecting changes in general (non-functional) climate time series data. The data in Figure 1 demonstrates substantial seasonality in stratospheric temperature. By describing seasonality using smooth functions, the proposed model handles general seasonality characteristics of the data and retains the high-resolution nature of daily-averaged data.

The remainder of the article is organized as follows: Section 2 presents the elastic amplitude and phase changepoint methods, and outlines the implementation using fully functional or functional principal component analysis. Section 3 provides analysis of the methodologies on simulated and multiple climate variables from the MERRA-2 data base. Finally, concluding remarks and future work are presented in Section 4.

2 | ELASTIC FUNCTIONAL CHANGEPOINT ANALYSIS

2.1 | The functional changepoint problem

Assume we have real-valued functions $f_1, \ldots, f_n$ that are absolutely continuous on the interval $[0, 1]$. Using this interval is not a restriction and can be easily extended to any interval. The standard changepoint problem (Aue et al., 2018; Berkes et al., 2009) assumes the data is from the following model:

$$f_i = \mu + \delta \mathbb{1}(i > k^*) + \epsilon_i,$$

where $\|\mu\|^2 < \infty$, $\|\delta\|^2 < \infty$, and $\mathbb{E}\|\epsilon_i\|^2 < \infty$ and $\|\cdot\|$ is the standard $\mathbb{L}^2$ functional norm. The point k^* labels the time of the unknown mean change and ϵ_i is a zero mean random error term with may be i.i.d. or potentially weakly dependent as in Aue et al. (2018). The changepoint detection problem then becomes a hypothesis testing problem of

$$H_0 : \delta = 0 \quad \text{vs.} \quad H_A : \delta \neq 0.$$

Aue et al. (2018) considers the test statistic of $T_n = \max_{1 \leq k \leq n} \|S_{n,k}\|^2$, where

$$S_{n,k} = \frac{1}{\sqrt{n}}(k\mu_k - k\mu_n),$$

$\mu_k = k^{-1}\sum_{i=1}^{k} f_i$. Aue et al. (2018) demonstrate asymptotic properties of this test statistic. In addition to determining whether there is evidence of a change or not (the detection problem), Aue et al. (2018) also give an estimate of k^* using $\hat{k}^* = \mathrm{argmax}_{1 \leq k \leq n} \|S_{n,k}\|^2$ (the estimation problem).

In addition to the fully-functional statistic, Berkes et al. (2009) and subsequent work propose a functional principal component-based test statistic that includes dimension reduction. This test statistic evaluates evidence of a change in the mean of the functional data, and Berkes et al. (2009) also propose an estimate of k^* which was later adopted in Aue et al. (2009). These two approaches will motivate our test statistics for functional data with phase and amplitude variability.

Since the test statistics of Aue et al. (2018) and Berkes et al. (2009) are based on functional statistics that do not use functional alignment, this cross-sectional approach assumes there is no phase variability and that there is only a random amplitude term. We know from Tucker et al. (2013) and Srivastava and Klassen (2016) that functional data contains two types of variation, amplitude and phase, and both are important to consider when evaluating functional data for statistical analysis and modeling.

2.2 | The elastic functional changepoint problem

Before introducing the changepoint problem for functional data with amplitude and phase variation, we review key concepts regarding the separation of amplitude and phase variation in functional data. First, we will introduce the group of warping functions that represent actions on the phase of functions. Let γ denote a warping function that is from the following group Γ, the phase space of the unit interval $[0, 1]$ defined as

$$\Gamma = \{\gamma : [0,1] \mapsto [0,1] \mid \gamma(0) = 0, \ \gamma(1) = 1, \ \gamma \ \text{is a diffeomorphism}\}.$$

This diffeomorphic constraint gives rise to the notion of elasticity because the elements of Γ, that is, phase functions, can only smoothly stretch and contract portions of the unit interval so that it maps back to itself. It should also be noted, that defining on the interval $[0, 1]$ is not a restriction and all the following construction can be extended to any arbitrary interval. The set Γ is a group under the composition operation $\circ$ and has identity element $\gamma(t) = t$. Additionally, we will use the Fisher-Rao metric and the square-root velocity function (SRVF) for its attractive properties (see Srivastava & Klassen, 2016 for properties and appropriate proofs). The square-root velocity function (SRVF) is defined as:

$$q = \mathrm{sign}(\dot{f})\sqrt{|\dot{f}|},$$

where $\dot{f}$ is the time derivative of the function f. The Fisher-Rao metric and for functions f_1 and f_2 with corresponding SRVFs q_1 and q_2, is defined as

$$d_{\mathrm{FR}}(f_1, f_2) = \|q_1 - q_2\|^2,$$

which is the $\mathbb{L}^2$ distance between the SRVFs. Define $(q, \gamma) = q \circ \gamma \sqrt{\dot{\gamma}}$ to represent the effect of a warping function to a square-root velocity function. That is, if the SRVF of f is q, then the SRVF of $f \circ \gamma$ is (q, γ). An attractive property of the Fisher-Rao distance is its invariance to random warping, so that for two SRVFs q_1 and q_2 and a warping function γ, $\|(q_1, \gamma) - (q_2, \gamma)\|^2 = \|q_1 - q_2\|^2$.

To consider a changepoint problem for the amplitude of functional data that has phase variability, we propose the model

$$q_i = \left((\mu_q + \delta_q \mathbb{1}(i > k^*) + \epsilon_i), \gamma_i^{-1} \right).$$

Here μ_q is the mean of the square-root velocity functions before a change k^*, δ is its change, $\{\epsilon_i\}_{i=1}^n$ are random, zero-mean functions representing amplitude variability, and $\{\gamma_i\}_{i=1}^n$ are random warping functions that represent phase variability. Here, we now aim to detect a changepoint in the sequence of SRVF functions while controlling for phase variability. We can also view the data as

$$(q_i, \gamma_i) = \mu_q + \delta_q \mathbb{1}(i > k^*) + \epsilon_i.$$

For the changepoint detection problem, we wish to test the hypothesis

$$H_0 : \delta_q = 0 \quad \text{vs.} \quad H_A : \delta_q \neq 0.$$

We are also interested in the estimation problem of k^* and the mean elements μ_q and δ_q.

We next turn to a different model that hypothesizes a change in the phase of the functional data. Now let

$$(q_i, \gamma_i) = \mu_q + \epsilon_i.$$

In this setting, we consider a model for γ_i as

$$\gamma_i = \begin{cases} \eta_i(\mu_\gamma) & \text{if } i \leq k^* \\ \eta_i(\nu_\gamma) & \text{if } i > k^*, \end{cases}$$

where $\eta_i(\mu)$ denotes a random element of Γ with mean $\mu \in \Gamma$. Therefore, $\mu_\gamma \in \Gamma$ and $\nu_\gamma \in \Gamma$ are the mean warping functions before and after the change, respectively. We then aim to test the hypothesis

$$H_0 : \mu_\gamma = \nu_\gamma \quad \text{vs.} \quad H_A : \mu_\gamma \neq \nu_\gamma.$$

Again, we also consider the estimation problem of k^* as well as estimating μ_q, μ_γ, and ν_γ.

In the next two subsections, we propose a test procedure for changepoints in the amplitude mean and phase mean of functional data, respectively. While we have separated the amplitude and phase changepoint detection problems, there may be changes in both the amplitude and phase of the functional data. In this setting, by properly disentangling these two different sources of variability, one can evaluate changes without the sources of variability potentially impeding each other. If the practitioner has no prior knowledge on what type of changepoint may exist in the data, one may test for both an amplitude and a phase change and use a Bonferroni correction to the resulting p-values. However, in many cases, the type of change may be known a priori. For example, we are primarily interested in an amplitude change in our data application.

2.3 | Amplitude changepoint

From Srivastava and Klassen (2016) the amplitude distance (which is a proper distance) between two functions is defined as

$$d_a(f_1, f_2) = \inf_{\gamma \in \Gamma} \|q_1 - (q_2, \gamma)\|.$$

Notice that $d_a(f_1, f_2)$ is the infimum of the Fisher-Rao distance between the square-root velocity function q_1 and a warped version of the SRVF q_2. From a proper distance we can define the Karcher mean of functions as

$$\mu_q^k = \arg \min_q \sum_{i=1}^k d_a(q, q_i)^2.$$

It should be noted that since the solution to this optimization problem is invariant to random warping and technically produces an orbit (i.e., a set of solutions). We pick the element of the orbit similar to Tucker et al. (2013) where the optimizer is the one corresponding to where the mean of γ_i is γ_{id}. Notice that the alignment process and the computation of a Karcher mean do not take into account potential dependence of elements in the functional time series.

To determine if an amplitude changepoint has occurred, we can define a test statistic using the (scaled) functional cumulative sum. This is similar to the approach used in Aue et al. (2018). The test statistic is defined as

$$S_{n,k} = \frac{1}{\sqrt{n}}\left(k\mu_q^k - k\mu_q^n\right), \tag{1}$$

where μ_q^n is the Karcher mean of the entire data set and for a specific value of k we can compute the Karcher mean μ_q^k. We can also define a test statistic using the $\mathbb{L}^2$ norm to also test for differences in means which is similar to the test statistic proposed in Gromenko et al. (2017). The test statistic is defined as

$$\Lambda_2 = \frac{1}{n}\sum_{k=1}^{n} d_a\left(k\mu_q^k, k\mu_q^n\right)^2. \tag{2}$$

One can use the max-point structural break detector

$$T_n = \max_{1 \le k \le n} \|S_{n,k}\|^2$$

to test the hypothesis H_0 versus H_A. One may also estimate the changepoint $\hat{k}^*$ where

$$\hat{k}^* = \min\{k : \|S_{n,k}\| = \max_{1 \le k \le n}\|S_{n,k}\|\}.$$

For the asymptotic properties of this statistic the reader is referred to Aue et al. (2018). Under this model T_n is distributed according to

$$T_n \xrightarrow{d} \sup_{0 \le x \le 1} \sum_{l=1}^{\infty} \lambda_l B_l^2(x),$$

where B_l are i.i.d. standard Brownian bridges defined on [0, 1] and $\{\lambda_l\}$ are the eigenvalues on the long-run covariance kernel of the error sequence ϵ_i. For this theorem, we assume that the errors $\epsilon_i(u)$ are $L^p - m$ approximable, and that again $\|\mu\|^2 < \infty$, $\|\delta\|^2 < \infty$, and $\mathbb{E}\|\epsilon_i\|^2 < \infty$; see Aue et al. (2018) for the definition of $L^p - m$ approximable, which allows for some dependence between elements in the functional time series. The estimation of $\{\lambda_\ell\}$ can then be done using the estimator in Section 3.1 of Aue et al. (2018). Computation of critical values of T_n can be done by the simulation of independent Brownian bridges and truncating the infinite sum; see Section 3.2 of Aue et al. (2018) for more details on this computation.

The test statistic T_n evaluates potential mean shape changes for functional data under phase variability. We can also easily implement this using functional principal component analysis (fPCA) as in Berkes et al. (2009). In order to properly account for the variability, we can use the vertical fPCA and horizontal fPCA presented in Tucker et al. (2013). These fPCA methods account for the variability by first separating the phase and amplitude and then performing the fPCA on the spaces separately. Using these methods one can construct the changepoint test in the amplitude space using the SRVF, q, or specifically the aligned SRVF, $\tilde{q}$, and in the phase space using the warping functions, γ.

We can use the *vertical* fPCA from Tucker et al. (2013) to test for an amplitude change. We will first give a short review of this fPCA method below, and more details can be found in Tucker et al. (2013).

2.3.1 | Vertical functional principal components

Let $f_1, \ldots, f_n$ be a given set of functions, and $q_1, \ldots, q_n$ be the corresponding SRVFs, μ_q be their Karcher mean, and let $\tilde{q}_i$s be the corresponding aligned SRVFs using Algorithm 1 from Tucker et al. (2013). In performing vertical fPCA, one also needs to include the variability associated with the initial values, that is, $\{f_i(0)\}$, of the given functions. Since

representing functions by their SRVFs ignores vertical translation, this variable is treated separately. That is, a functional variable f is analyzed using the pair $(q, f(0))$ rather than just q. This way, the mapping from the function space $\mathcal{F}$ to $\mathbb{L}^2 \times \mathbb{R}$ is a bijection. In practice, where q is represented using a finite partition of $[0, 1]$, say with cardinality T, the combined vector $h_i = [q_i \ f_i(0)]$ simply has dimension $(T + 1)$ for fPCA. We can define a sample covariance operator for the aligned combined vector $\tilde{h}_i = [\tilde{q}_i \ f_i(0)]$ as

$$K_h = \frac{1}{n-1} \sum_{i=1}^{n} (\tilde{h}_i - \mu_h)(\tilde{h}_i - \mu_h)^\top \in \mathbb{R}^{(T+1) \times (T+1)},$$

where $\mu_h = [\mu_q \ \overline{f}(0)]$. Taking the SVD, $K_h = U_h \Lambda_h V_h^\top$, we can calculate the directions of principal variability in the given SRVFs using the first $p \leq n$ columns of U_h and can then convert back to the function space $\mathcal{F}$, via integration, to find the principal components of the original functional data. Moreover, we can calculate the observed principal scores as $\eta_{i,l} = \langle \tilde{h}_i, U_{h,l} \rangle$.

The test statistic would then follow directly from Berkes et al. (2009) with test statistic

$$T_N(x) = \frac{1}{N} \sum_{l=1}^{d} \lambda_l^{-1} \left(\sum_{1 \leq i \leq Nx} \eta_{i,l} - x \sum_{1 \leq i \leq N} \eta_{i,l} \right)^2, \tag{3}$$

where λ are the eigenvalues from Λ_h and η are the scores from the fPCA. The test statistic is distributed as

$$T_N(x) \xrightarrow{d} \sum_{1 \leq l \leq d} B_l^2(x), \quad 0 \leq x \leq 1,$$

where $B_l(\cdot)$ are independent Brownian Bridges on $[0, 1]$.

2.4 | Phase changepoint

From the calculation of the Karcher mean, μ_q^k, we have the set of warping functions $\{\gamma_i\}^k$. To perform a changepoint test on the space of warping functions directly is difficult, as Γ is a infinite-dimensional nonlinear manifold and computation of the mean is not straightforward due to the unknown geometry. To facilitate computing a mean, we can simplify the geometry of Γ.

2.4.1 | Simplifying geometry of Γ

The space of warping functions, Γ, is an infinite-dimensional nonlinear manifold and therefore cannot be treated as a standard Hilbert space. To overcome this problem, we will use tools from differential geometry to perform statistical analyses and model the warping functions. The following framework was previously used in various settings including (1) modeling re-parameterizations of curves (Srivastava & Jermyn, 2009) and (2) putting prior distributions on warping functions (Kurtek, 2017; Tucker et al., 2021), and many others. It is also very closely related to the square-root representation of probability density functions introduced by Bhattacharya (1943).

We represent an element $\gamma \in \Gamma$ by the square-root of its derivative $\psi = \sqrt{\dot{\gamma}}$. Note that this is the same as the SRVF defined earlier, and it takes this form since $\dot{\gamma} > 0$. The identity γ_{id} maps to a constant function with value $\psi_{id}(t) = 1$. Since $\gamma(0) = 0$, the mapping from γ to ψ is a bijection and one can reconstruct γ from ψ using $\gamma(t) = \int_0^t \psi(s)^2 ds$. An important advantage of this transformation is that since $\|\psi\|^2 = \int_0^1 \psi(t)^2 dt = \int_0^1 \dot{\gamma}(t) dt = \gamma(1) - \gamma(0) = 1$, the set of all such ψs is the positive orthant of the unit Hilbert sphere in $\mathbb{L}^2$: $\Psi = \mathbb{S}_\infty^+$. In other words, the square-root representation simplifies the complicated geometry of Γ to a (subset of a) unit sphere. The distance between any two warping functions, that is, the phase distance, is exactly the arc-length between their corresponding SRVFs on Ψ:

$$d_p(\gamma_1, \gamma_2) = d_\psi(\psi_1, \psi_2) \equiv \cos^{-1} \left(\int_0^1 \psi_1(t) \psi_2(t) dt \right). \tag{4}$$

Figure 3 depicts the SRVF space of warping functions as a unit sphere (Tucker, 2014).
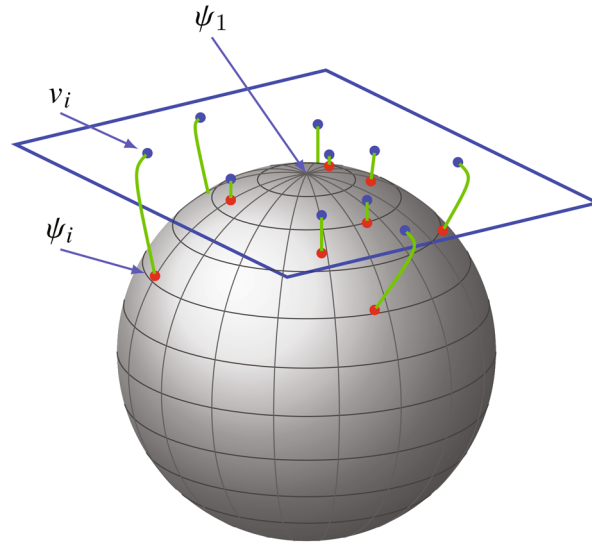
FIGURE 3 Depiction of the SRVF space of warping functions as a sphere and a tangent space at ψ_1.

While the geometry of $\Psi \subset \mathbb{S}_\infty$ is more tractable, it is still a nonlinear manifold and computing standard statistics remains difficult. Instead, we use a tangent (vector) space at a certain fixed point for further analysis. The tangent space at any point $\psi \in \Psi$ is given by: $T_\psi(\Psi) = \{v \in \mathbb{L}^2 | \int_0^1 v(t)\psi(t)dt = 0\}$. To map between the representation space Ψ and tangent spaces, one requires the exponential and inverse-exponential mappings. The exponential map at a point $\psi \in \Psi$ denoted by $\exp_\psi : T_\psi(\Psi) \mapsto \Psi$, is defined as

$$\exp_\psi(v) = \cos(\|v\|)\psi + \sin(\|v\|)\frac{v}{\|v\|}, \tag{5}$$

where $v \in T_\psi(\Psi)$. Thus, $\exp_\psi(v)$ maps points from the tangent space at ψ to the representation space Ψ. Similarly, the inverse-exponential map, denoted by $\exp_\psi^{-1} : \Psi \mapsto T_\psi(\Psi)$, is defined as

$$\exp_\psi^{-1}(\psi_1) = \frac{\theta}{\sin(\theta)}(\psi_1 - \cos(\theta)\psi), \tag{6}$$

where $\theta = d_p(\gamma_1, \gamma)$. This mapping takes points from the representation space to the tangent space at ψ.

The tangent space representation v is sometimes referred to as a *shooting vector*, as depicted in Figure 3. The remaining question is which tangent space should be used to represent the warping functions. A sensible point on Ψ to define the tangent space is at the sample Karcher mean $\hat{\mu}_\psi$ (corresponding to $\hat{\mu}_\gamma$) of the given warping functions. We can compute the mean of the warping functions using Algorithm 2 from Tucker et al. (2013).

We then can perform a test if we have a changepoint in the phase space (on the tangent space of $\mathbb{S}_\infty$) using the test statistic

$$S_{n,k} = \frac{1}{\sqrt{n}}\left(k\bar{v}^k - k\bar{v}^n\right), \tag{7}$$

which is similar as defined previously, where $\bar{v}^k$ is the shooting vector of the Karcher mean of the warping functions, $\hat{\mu}_\psi$, computed using the first k warping functions. We can also define the test using the $\mathbb{L}^2$ norm similar to Gromenko et al. (2017) as

$$\Lambda_2 = \frac{1}{n}\sum_{k=1}^n \|k\bar{v}^k - k\bar{v}^n\|^2.$$

Under this test we are testing if there is a phase shift or changepoint. The definition of T_n, its distribution, and finding $\hat{k}^*$ are the same as the amplitude changepoint in Section 2.3. As mentioned above we can also implement this using *horizontal* fPCA from Tucker et al. (2013).
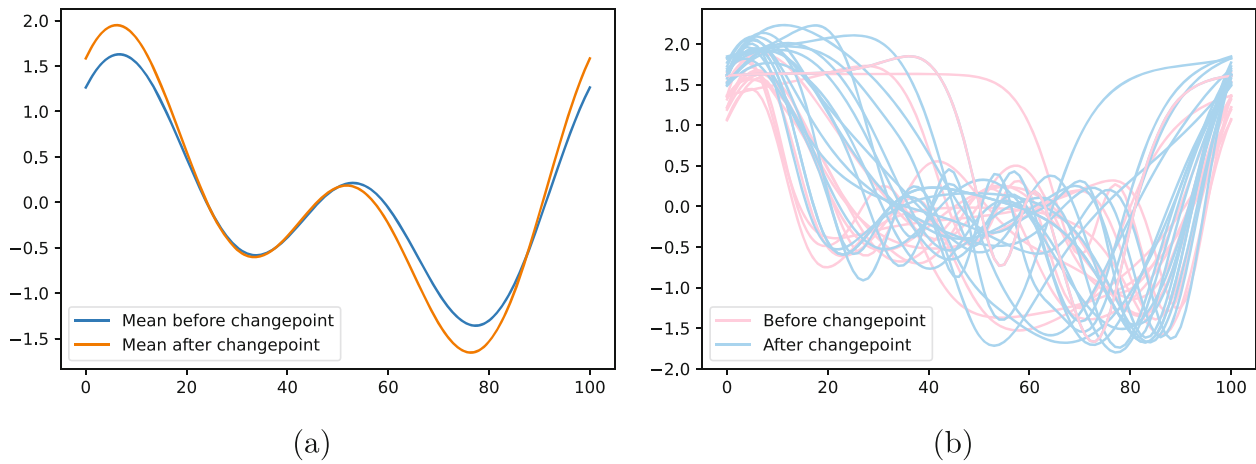
FIGURE 4 Example of simulated data (a) mean before changepoint μ and mean after changepoint $\mu + \delta$. (b) Simulated data, with color corresponding to whether the function was before or after the change. For this example, both elastic amplitude fully-functional and dimension-reduction-based approaches detected a changepoint with $\hat{k}^* = 12 = k^*$ and p-values of 0.000. The cross-sectional approaches estimated $\hat{k}^*$ as 11 (fully-functional) or 13 (dimension-reduction) with p-values of 0.142 and 0.319, respectively.

2.4.2 | Horizontal functional principal components

To perform horizontal fPCA we will use the tangent space at μ_ψ (Srivastava & Klassen, 2016) to perform analysis, where μ_ψ is the mean of the transformed warping functions. In this tangent space, we can define a sample covariance function:

$$K_\psi = \frac{1}{n-1} \sum_{i=1}^{n} v_i v_i^\mathsf{T} \in \mathbb{R}^{T \times T}.$$

The singular value decomposition (SVD) of $K_\psi = U_\psi \Lambda_\psi V_\psi^\mathsf{T}$ provides the estimated principal components of $\{\psi_i\}$: the principal directions $U_{\psi,l}$ and the observed scores $\eta_{i,l} = \langle v_i, U_{\psi,l} \rangle$. This analysis on $\mathbb{S}_\infty$ is similar to the ideas presented in Srivastava et al. (2005) although one can also use the idea of principal nested sphere for this analysis (Jung et al., 2012).

We can then use the test defined in Equation (3) to perform the changepoint test. The distribution to find critical point and confidence interval apply directly.

3 | APPLICATION

To demonstrate our proposed method, we will compare our elastic method to the cross-section changepoint methods of Aue et al. (2009) and Aue et al. (2018). First, we will assess our method's performance on simulated data where an amplitude changepoint exists as well as simulated data where a phase changepoint exists. We then apply our method to a global temperature observations dataset based on the MERRA-2 reanalysis product (Global Modeling and Assimilation Office (GMAO), 2015) during the time period (1984–1998) that contains the eruption of Mt. Pinatubo in the Philippines in June 1991.

3.1 | Simulation data

3.1.1 | Amplitude changepoint

To discuss the basics of the simulation setup, we consider sample sizes of $n \in \{15, 30, 50, 75\}$, with an amplitude changepoint $k^* = 0.40 \cdot n$ so that $k^* \in \{6, 12, 20, 30\}$, respectively. For each simulation, the mean functions are generated as

$$\mu(t) = a_0 \cos(2\pi t) + b_0 \sin(2\pi t) + a_1 \cos(4\pi t) + b_1 \sin(4\pi t)$$

$$\delta(t) = \Delta \cos(2\pi t) + \Delta \sin(2\pi t) + \Delta \cos(4\pi t) + \Delta \sin(4\pi t),$$

where $a_i \sim U[0, 1]$ and $b_i \sim U[0, 1]$ for $i \in \{0, 1\}$, and Δ is a signal strength parameter. The error term is taken as

$$\epsilon_i(t) = Z_{1,i} + Z_{2,i} \frac{1}{\sqrt{2}} \cos(8\pi t) + Z_{3,i} \frac{1}{\sqrt{2}} \sin(8\pi t),$$

where $Z_{l,i} \overset{\text{ind}}{\sim} \mathcal{N}(0, \sigma_l^2)$, with $(\sigma_1, \sigma_2, \sigma_3) = (0.14, 0.10, 0.08)$. Finally, each function was warped with an independent, random warping function with variance 0.2, generated similarly to Section 5.1 of Harris et al. (2021). We consider results over 100 simulations for each value of n and each signal strength $\Delta \in \{0, 0.04, 0.08, 0.12, 0.16\}$. In Figure 4, we plot an example of the simulated data with $n = 30$ and $\Delta = 0.16$, as well as its results using these test statistics. For each simulation, we compute score-based and fully-functional test statistics for both the cross-sectional (Aue et al., 2009; 2018) and elastic amplitude changepoint methods.

Overall changepoint detection results are presented for the dimension-reduction based methodology in Table 1. The $\Delta = 0$ column (no change) represents the estimated Type I error of the detection procedure. We see that both approaches control Type I error well. As Δ and n increase, it becomes easier to detect changepoints under the alternative hypothesis. The elastic approach uniformly improves power over the cross-sectional approach, detecting close to all changepoints when $\delta \geq 0.12$ and $n \geq 30$. In contrast, the cross-sectional approach only detects some changepoints and may not be consistent in this setting, as increasing n does not always substantially improve the number of changepoints detected.

We present respective results for the fully-functional approaches in Table 2. While Type I error is slightly inflated for both approaches, similar gaps in power between the elastic and cross-sectional approaches hold when the alternative is true, again suggesting that the cross-sectional approach will miss potential changes in the data. In general,

TABLE 1 Proportion of simulations with estimated changepoint at the $\alpha = 0.05$ level for the amplitude change.

Elastic PCA	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.01	0.07	0.58	0.77	0.86
$n = 30$	0.02	0.22	0.92	0.99	1.00
$n = 50$	0.03	0.25	0.95	1.00	1.00
$n = 75$	0.01	0.39	0.98	1.00	1.00
Cross-sectional PCA	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.01	0.02	0.03	0.05	0.07
$n = 30$	0.02	0.09	0.17	0.18	0.30
$n = 50$	0.02	0.08	0.13	0.34	0.36
$n = 75$	0.05	0.04	0.24	0.24	0.40

Note: (Top) Using the elastic PCA test statistic (bottom) using the cross-sectional PCA test statistic of Aue et al. (2009).

TABLE 2 Proportion of simulations with estimated changepoint at the $\alpha = 0.05$ level for the amplitude change.

Elastic fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.19	0.18	0.35	0.54	0.85
$n = 30$	0.07	0.16	0.27	0.85	1.00
$n = 50$	0.06	0.11	0.52	1.00	1.00
$n = 75$	0.09	0.20	0.86	1.00	1.00
Cross-sectional fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.13	0.10	0.13	0.09	0.15
$n = 30$	0.04	0.12	0.07	0.16	0.15
$n = 50$	0.06	0.06	0.12	0.18	0.20
$n = 75$	0.04	0.08	0.13	0.21	0.28

Note: (Top) Using the elastic fully-functional test statistic (bottom) using the cross-sectional fully-functional test statistic of Aue et al. (2018).

the fully-functional test statistics tend to have more less power than the dimension-reduction-based test statistics. We also compare the estimation of k^* using a sample-size-normalized RMSE based on the estimated changepoints $\hat{k}^*_{mc}$ over each simulation $mc = 1, \ldots, MC = 100$:

$$\sqrt{\frac{1}{MC} \sum_{mc=1}^{MC} \left(\frac{\hat{k}^*_{mc} - k^*}{n} \right)^2},$$

with results in Table 3 for the fully-functional test statistics. The results mirror the testing results, with the elastic approach consistently estimating k^* better compared to the cross-sectional approach. The cross-sectional approach only partially improves estimates as Δ increases or n increases, again suggesting that cross-sectional approaches will be inadequate when there is phase variability in the data even if the sample size and signal strength are large.

3.1.2 | Phase changepoint

To evaluate changepoint detection in the mean phase of functional data, We consider the same simulation setup with the exception of the following. We take $\delta(t) = 0$, and the warping functions have mean $\gamma(t) = t$ before k^*. After k^*, the mean of the warping functions is a randomly-generated warping function with variance Δ_γ, where $\Delta_\gamma \in \{0, 0.075, 0.15, 0.225, 0.30\}$. The variance of the noise in each warping function was set to 0.05.

We give results in Tables 4 and 5. The elastic test for a phase change and the cross-sectional test perform similarly for both testing and changepoint estimation. The cross-sectional approach has a bit higher Type I error and slightly more power when $n = 15$. The results are essentially same between the two approaches for the estimation problem.

TABLE 3 Sample-size-normalized RMSE of estimated change time for the amplitude change.

Elastic fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.21	0.21	0.17	0.10	0.06
$n = 30$	0.23	0.18	0.10	0.05	0.03
$n = 50$	0.24	0.16	0.06	0.01	0.01
$n = 75$	0.23	0.15	0.05	0.01	0.00
Cross-sectional fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.19	0.21	0.21	0.19	0.18
$n = 30$	0.21	0.21	0.21	0.18	0.19
$n = 50$	0.22	0.21	0.18	0.17	0.14
$n = 75$	0.20	0.21	0.19	0.14	0.14

Note: (Top) Using the elastic fully-functional test statistic (bottom) using the cross-sectional fully-functional test statistic of Aue et al. (2018).

TABLE 4 Proportion of simulations with estimated changepoint at the $\alpha = 0.05$ level for the phase change.

Elastic fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.06	0.41	0.81	0.81	0.93
$n = 30$	0.03	0.61	0.82	0.95	0.97
$n = 50$	0.04	0.75	0.95	0.95	0.98
$n = 75$	0.00	0.88	0.92	0.97	1.00
Cross-sectional fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.14	0.65	0.88	0.93	0.95
$n = 30$	0.07	0.65	0.85	0.94	0.99
$n = 50$	0.08	0.76	0.96	0.92	0.98
$n = 75$	0.08	0.88	0.93	0.97	1.00

Note: (Top) Using the elastic fully-functional test statistic (bottom) using the cross-sectional fully-functional test statistic of Aue et al. (2018).

TABLE 5 Sample-size-normalized RMSE of estimated change time for the phase change.

Elastic fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.21	0.15	0.04	0.05	0.05
$n = 30$	0.21	0.12	0.05	0.04	0.04
$n = 50$	0.19	0.11	0.04	0.04	0.01
$n = 75$	0.21	0.04	0.05	0.04	0.00
Cross-sectional fully-functional	$\Delta = 0$	$\Delta = 0.04$	$\Delta = 0.08$	$\Delta = 0.12$	$\Delta = 0.16$
$n = 15$	0.22	0.15	0.07	0.07	0.03
$n = 30$	0.22	0.12	0.06	0.04	0.02
$n = 50$	0.22	0.11	0.05	0.04	0.01
$n = 75$	0.22	0.05	0.05	0.04	0.01

Note: (Top) Using the elastic fully-functional test statistic (bottom) using the cross-sectional fully-functional test statistic of Aue et al. (2018).
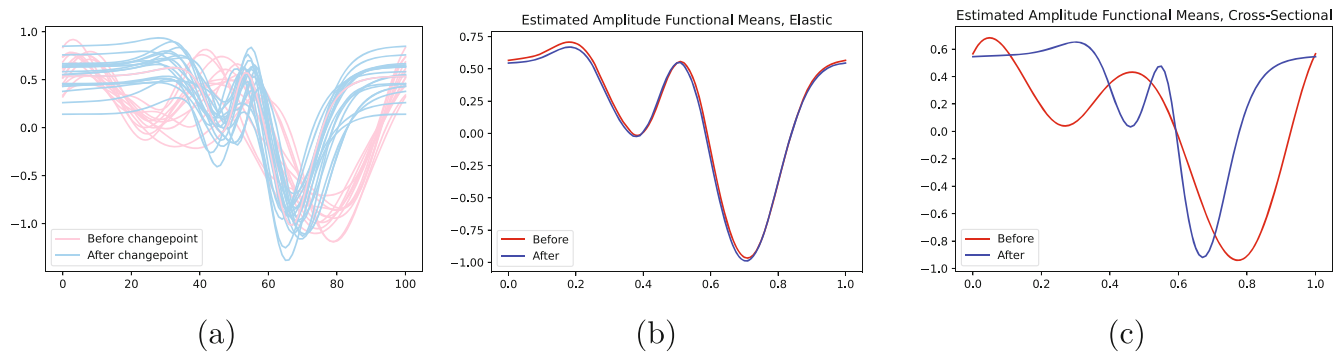
FIGURE 5 A phase simulation example with $\Delta_\gamma = 0.3$ and $n = 30$. (a) Simulated functional data (b) estimated amplitude means using the elastic fully-functional test statistic. (c) Estimated means using the cross-sectional fully-functional test statistic. Both approaches detected the correct changepoint $k^* = 12$ with a *p*-value of 0.000.

While the cross sectional approach is able to detect changes, it will not be able to identify the data-generating processes, as it interprets all variability in the data as amplitude variability. To demonstrate this, we plot an example in Figure 5. While the mean amplitude functions are the same before and after the changepoint, the data warping in (a) shows considerably different observed functions. The estimated means based on the elastic and cross-sectional fully-functional statistics are shown in (b) and (c) respectively. While the elastic approach estimates nearly the same amplitude function before and after the changepoint, phase variability is present and the cross-sectional detection procedure does not recognize the similarity between the mean amplitude functions of the data.

3.2 | MERRA-2 stratospheric temperature

We now present our data analysis on stratospheric temperature of the climate reanalysis data MERRA-2 (Global Modeling and Assimilation Office (GMAO), 2015). Using data from the years 1984–1998, we aim to evaluate changes related to the eruption of Mt. Pinatubo in June 1991. See, for example, Thompson et al. (2009) for information on variability and variables related to Mt. Pinatubo's eruption. For this study, we focus on daily stratospheric temperature near the 50 millibar pressure surface. We run each detection procedure at 3312 different locations on a grid. Since we are primarily interested in an amplitude changepoint, we do not present results based on a change in the phase function.

We begin by presenting fully-functional results for a single location in Figure 6. In (a) and (b), we plot the original functional data along with the estimated mean functions for the elastic approach (a) and the cross-sectional
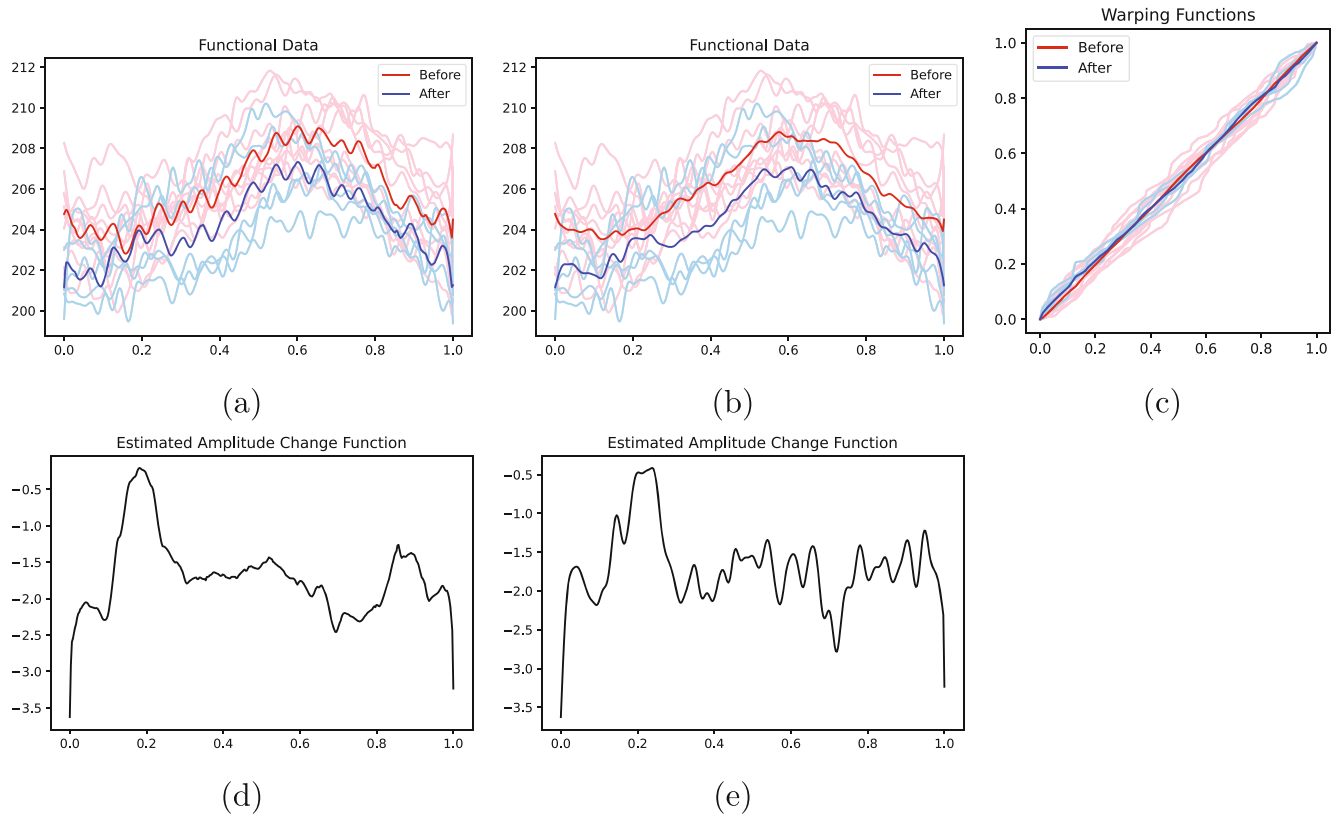
FIGURE 6 An example of MERRA-2 stratospheric temperature (in Kelvin) fully-functional results at 0°W and 2°N. The elastic approach had $\hat{k}^* = 1993$ with a p-value of 0.028, and the cross-sectional had $\hat{k}^* = 1993$ with a p-value of 0.067. (a) Functional data and estimated means after alignment. (b) Functional data and estimated means before alignment. (c) Estimated warping functions. (d) Estimated change for elastic approach. (e) Estimated change for cross-sectional approach. (e) Warping functions.

approach (b) before and after the estimated change. The elastic approach appears to align weather patterns, maintaining cyclical behavior in the temperature throughout the year. In contrast, the cross-sectional approach averages these over years while ignoring phase variability. The distinction between these two estimates is filtered into estimates of the change function in (d) and (e). The estimated change function from the elastic approach (d) does not have major oscillations since the cyclical behavior is shared between the two mean functions. For the cross-sectional approach of Aue et al. (2018), however, this variability is propagated into the estimate of the change function, resulting in an estimate that is more noisy. By accounting for phase variability in the original data, the estimate of the change is more easily isolated and estimated. In panel (c), we plot the warping functions, colored based on if they were before or after the changepoint. While some phase variability is present and substantially affects the estimates, functions do not have to be warped drastically, and there does not appear to be a trend or change in the warping functions.

We present the results for all locations in Figure 7. Locations were only plotted if there was an associated p-value less than 0.05. For the most part, the elastic and cross-sectional approaches have similar results. Both detect changes in the tropics in 1993 and 1994, where the impact of Mt. Pinatubo is particularly strong. The elastic approach detects fewer changes in the poles, which are less likely to be associated with the eruption of Mt. Pinatubo. To interpret the changes detected, we see that the changes in the tropics correspond to a decrease in temperature. Combined with the time of the changepoint, both methodologies detect when the stratospheric temperature "returns to normal" after a period of increase in 1991 and 1992 immediately following the eruption. The estimated changes are comparable between the two approaches. We summarize the detected changepoints in Table 6. The vast majority of the detected changepoints had $\hat{k}^* = 1992$ or $\hat{k}^* = 1993$, with an associated decrease in temperature on average. In Figure 8, we plot results based on the dimension reduction test statistics. Although fewer changepoints were detected, overall the pattern and nature of the changes are similar to those of the fully-functional test statistics.

In our analysis, we appear to primarily detect a decrease in temperature after the year 1992 or so. Since Mt. Pinatubo is expected to temporarily increase stratospheric temperatures immediately after the eruption, we appear to be detecting

(a)

(b)
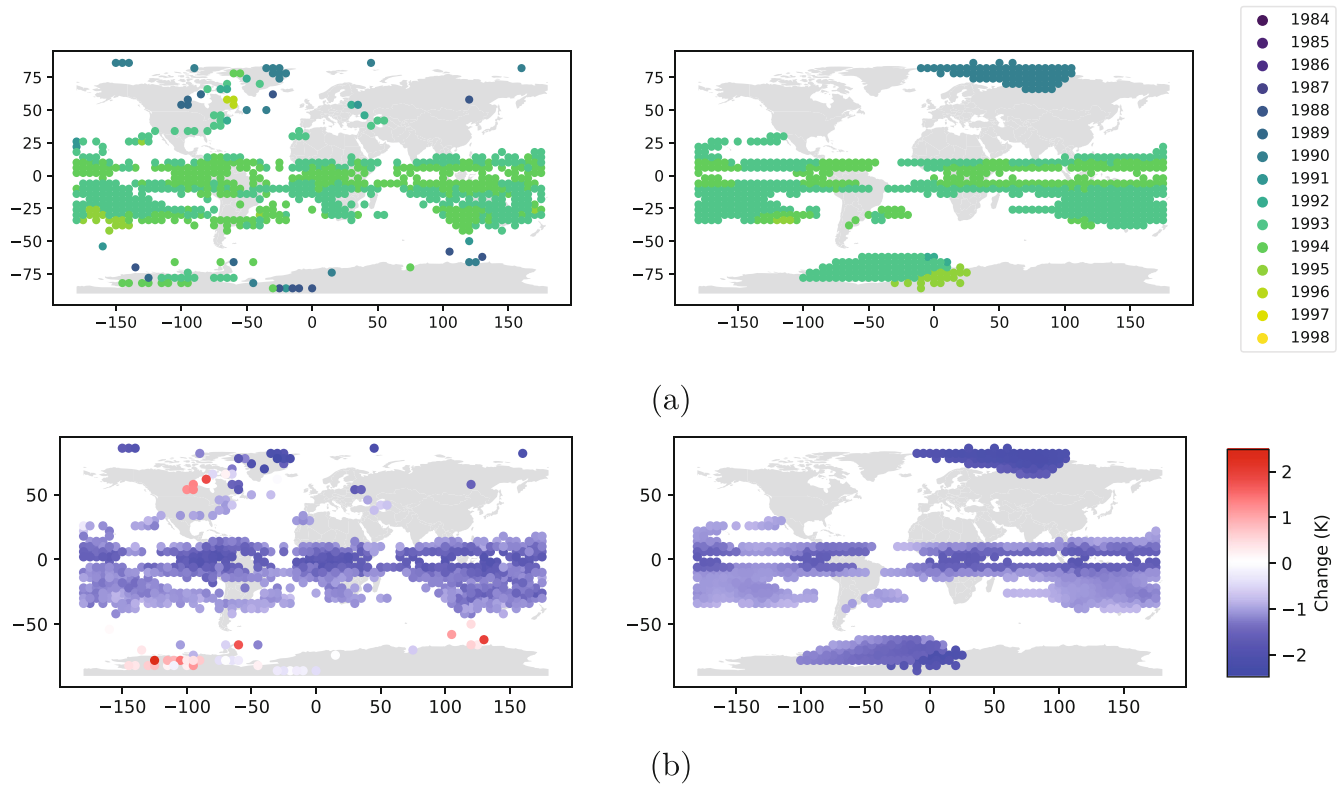
FIGURE 7 MERRA-2 stratospheric temperature fully-functional results. (a) Estimated change years $\hat{k}^*$ with p-values less than 0.05 using elastic approach (left) and cross-sectional approach (right). (b) Averaged estimated change $\int_0^1 \hat{\delta}(t)dt$ with p-values less than 0.05 for elastic approach (left) and cross-sectional approach (right).

**TABLE 6** Summary of detected test statistics by the elastic fully-functional method.

Year	# of detections	Avg. change (K)	Year	# of detections	Avg. change (K)
1984	0	N/A	1992	350	−1.12
1985	0	N/A	1993	291	−1.41
1986	0	N/A	1994	17	−1.18
1987	9	0.03	1995	3	−1.64
1988	6	1.60	1996	0	N/A
1989	19	−1.36	1997	0	N/A
1990	6	−0.35	1998	0	N/A
1991	6	−0.97			

when the temperature "returns to normal" as suggested before. To more directly target the eruption, we consider the epidemic-type changepoint model studied in Aston and Kirch (2012):

$$f_i = \mu + \delta \mathbb{1}\left(k_1^* < i \le k_2^*\right) + \epsilon_i,$$

where the only difference with the "at-most-one-change" model $f_i = \mu + \delta \mathbb{1}(i > k^*) + \epsilon_i$ is the indicator $\mathbb{1}\left(k_1^* < i \le k_2^*\right)$. Fully-functional and dimension-reduction elastic approaches for detecting this type of changepoint and for estimating k_1^* (when the change period begins) and k_2^* (when the change period ends) follow naturally from the test statistics and results in Aston and Kirch (2012) and Section 2. We present results for the epidemic model in Figure 9. While fewer
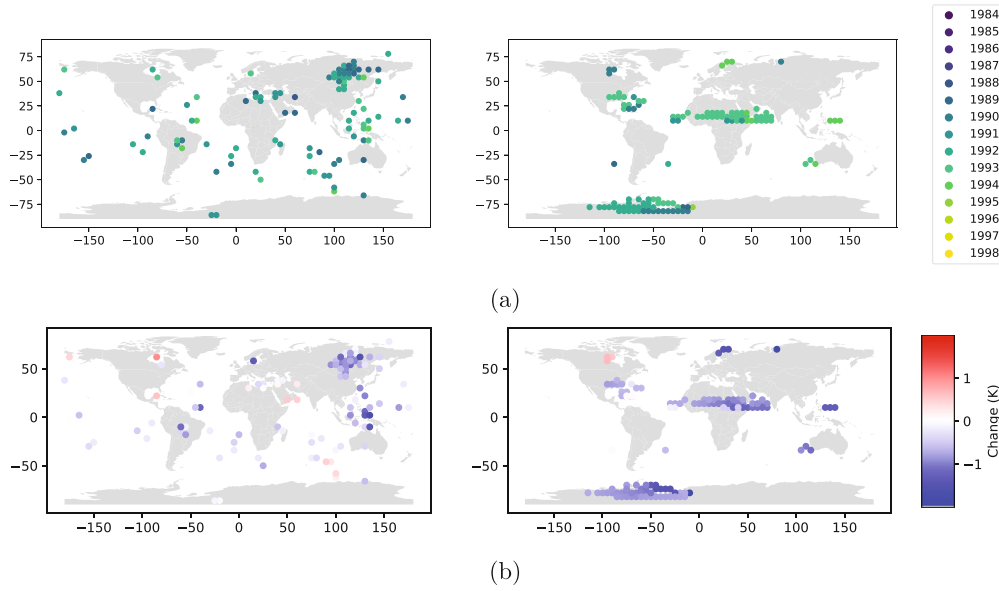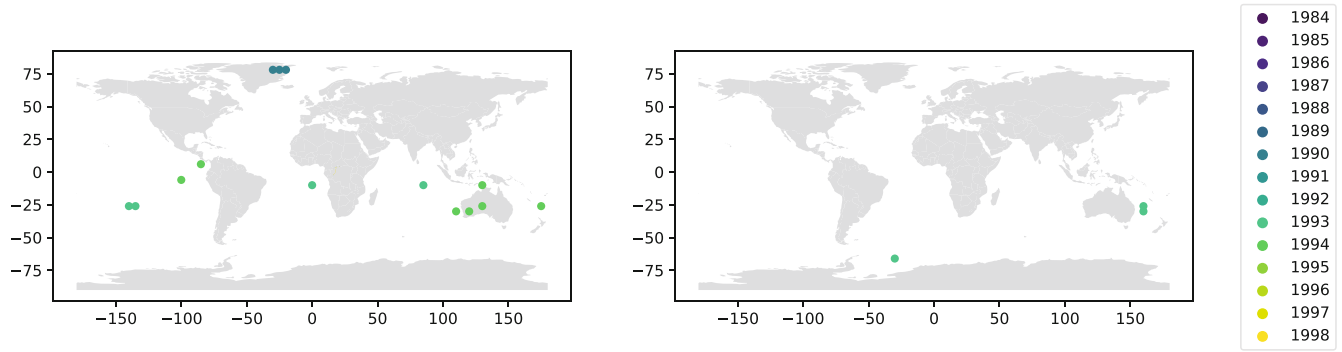
FIGURE 8 MERRA-2 stratospheric score-based results. (a) Estimated change years $\hat{k}^*$ with *p*-values less than 0.05 using elastic approach (left) and cross-sectional approach (right). (b) Averaged estimated change $\int_0^1 \hat{\delta}(t)dt$ with *p*-values less than 0.05 for elastic approach (left) and cross-sectional approach (right).
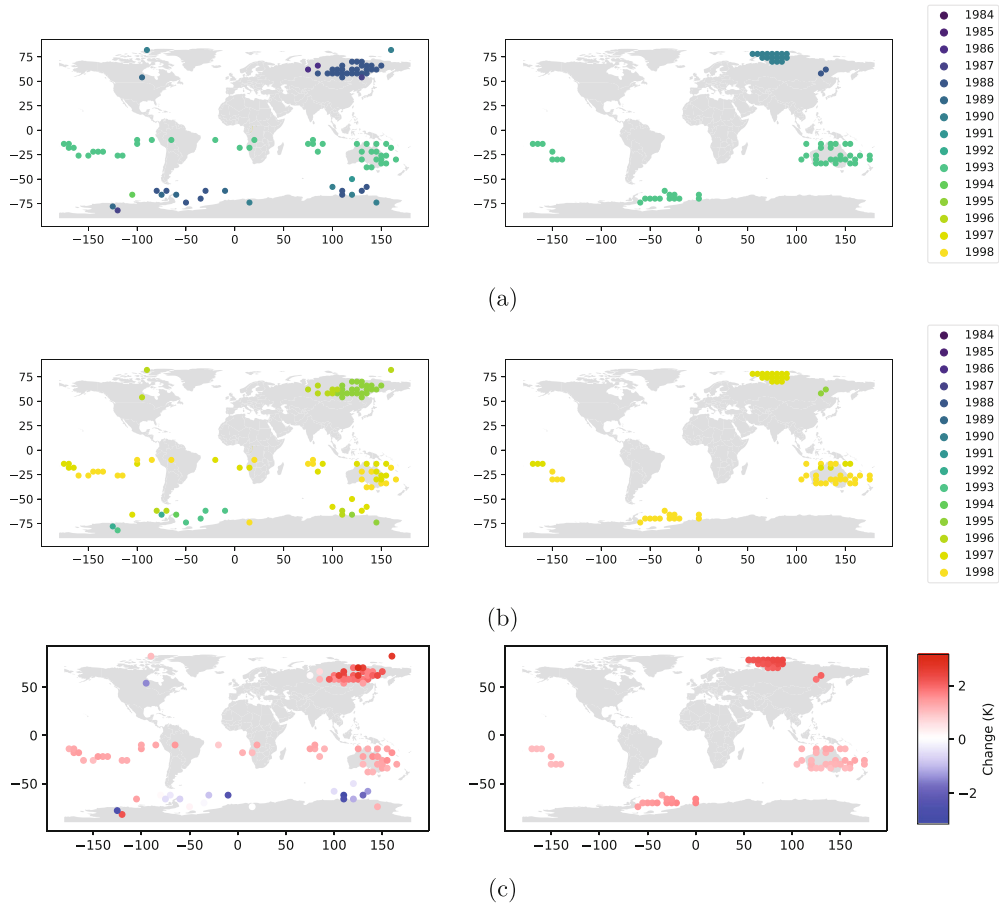
FIGURE 9 Fully-functional epidemic results. (a) Estimated change years $\hat{k}_1^*$ with *p*-values less than 0.05 using elastic approach (left) and cross-sectional approach (right) (b) the same as (a) with $\hat{k}_2^*$ (c) averaged estimated change $\int_0^1 \hat{\delta}(t)dt$ for locations with *p*-values less than 0.05.

FIGURE 10 Fully-functional results from the at-most-one-change model with a Benjamini–Hochberg p-value correction. Estimated change years $\hat{k}^*$ with Benjamini–Hochberg p-values less than 0.05 using elastic approach (left) and cross-sectional approach (right).

changepoints are detected compared to the at-most-one-change model, the beginning of the changepoints k_1^* are estimated earlier in the tropics, and these changes are associated with an increase in temperature rather than a decrease.

Finally, we evaluate the sensitivity of the detected changepoints to multiple testing correction. In Figure 10, we plot results with Benjamini–Hochberg corrected p-values for the fully-functional test statistics analogous to Figure 7. We see that very few changepoints are detected, but the elastic approach detects substantially more changepoints after p-value adjustment.

4 | CONCLUSION

We have proposed a new flexible elastic approach to functional changepoint detection that handles phase and amplitude variability in the data accordingly. We have demonstrated its advantages over the state-of-the-art methods in the functional changepoint literature using both simulated and real datasets. Unlike the current cross-sectional methods, the elastic method accurately estimates the mean of the underlying data-generating mechanism and also handles phase as a nuisance or as information suggesting a potential changepoint in the functional data. Correctly estimating the underlying mean before and after the change can have large impacts on interpretability of the data and implementation of downstream analysis.

In a simulation study, the procedures correctly detect a changepoint and estimate its time when there is either an amplitude or phase changepoint in the functional time series. In contrast, cross-sectional approaches miss detecting changepoints when there is a change in the mean amplitude, and the time of changepoint is not estimated well. In an analysis of stratospheric temperature, both the cross-sectional approaches and our introduced test statistics detect changes is the tropics after the eruption of Mt. Pinatubo. However, the test statistics that take into account phase variability result in a less noisy estimate of the change function, increasing its interpretability and smoothness.

For future work, we look to extend the method to the multiple change-point problem. This extension will not be trivial as scalability is a problem as shown in Harris et al. (2022). Multiple changepoints can be detected in $S_{n,k}$ if using a multiple peak detection scheme, but determining the p-value is not straightforward, and the test statistics will have to be defined carefully. Another interesting question is the performance of the test statistics in the presence of outliers. For example, an analysis of aerosol optical depth, which immediately increases after a volcanic eruption and returns to lower levels relatively quickly, would be illustrative. Other work has proposed robustified fully-functional statistics (Wegner & Wendler, 2022), and application of their test statistics to our setting would be relatively straightforward. In addition, while we have treated each location individually, comprehensively studying the spatial or multivariate problem of how functional time series at different locations relate to each other would be a fruitful direction for future research that the authors are currently studying. Paynabar et al. (2016), Wang et al. (2022), and Moradi et al. (2023) approach this problem in three different ways. Such an approach will also lend itself more to the multiple testing problem. Additionally, we can extend this from curves to trajectories that lie on Riemannian manifolds $\mathcal{M}$. In this case, one has to account for the non-zero curvature of the space and in particular, the calculation of the gradient in the Fisher Rao metric and computation of the required statistics.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

J. Derek Tucker https://orcid.org/0000-0001-8844-2169
Drew Yarger https://orcid.org/0000-0002-0420-7346

REFERENCES

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., & Sonmez, O. (2018). Detecting and dating structural breaks in functional data without dimension reduction. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, *80*, 509–529.

Berkes, I., Gabrys, R., Horvath, L., & Kokoszka, P. (2009). Detecting changes in the mean of functional observations. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, *71*, 927–946.

Bhattacharya, A. (1943). On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, *35*, 99–109.

Bosq, D. (2012). *Linear processes in function spaces: Theory and applications* (Vol. *149*). Springer Science & Business Media.

Fortuna, F., Gattone, S. A., & Di Battista, T. (2020). Functional estimation of diversity profiles. *Environmetrics*, *31*(8), e2645.

Global Modeling and Assimilation Office (GMAO). (2015). tavg1_2d_aer_nx_daily: Merra-2 3d iau state, meteorology instantaneous 3-hourly (p-coord, 0.625x0.5l42). https://doi.org/10.5067/KLICLTZ8EM9D

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.

Harris, T., Li, B., & Tucker, J. D. (2022). Scalable multiple changepoint detection for functional data sequences. *Environmetrics*, *33*(2), e2710.

Harris, T., Tucker, J. D., Li, B., & Shand, L. (2021). Elastic depths for detecting shape anomalies in functional data. *Technometrics*, *63*(4), 466–476.

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. S. Rao, S. S. Rao, & C. R. Rao (Eds.), *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.

Jung, S., Dryden, I. L., & Marron, J. S. (2012). Analysis of principal nested spheres. *Biometrika*, *99*(3), 551–568.

Kokoszka, P., & Zhang, X. (2012). Functional prediction of intraday cumulative returns. *Statistical Modelling*, *12*(4), 377–398.

Kurtek, S. (2017). A geometric approach to pairwise Bayesian alignment of functional data using importance sampling. *Electronic Journal of Statistics*, *11*(1), 502–531.

Lee, S., & Jung, S. (2017). Combined analysis of amplitude and phase variations in functional data. arXiv:1603.01775 [stat.ME].

Lund, R. B., Beaulieu, C., Killick, R., Lu, Q., & Shi, X. (2022). Good practices and common pitfalls in climate time series changepoint techniques: A review. *Journal of Climate*, 1–38.

Moradi, M., Cronie, O., Pérez-Goya, U., & Mateu, J. (2023). Hierarchical spatio-temporal change-point detection. *The American Statistician*, 1–11.

Paynabar, K., Zou, C., & Qiu, P. (2016). A change-point approach for phase-i analysis in multivariate profile monitoring and diagnosis. *Technometrics*, *58*(2), 191–204.

Reeves, J., Chen, J., Wang, X. L., Lund, R., & Lu, Q. Q. (2007). A review and comparison of changepoint detection techniques for climate data. *Journal of Applied Meteorology and Climatology*, *46*(6), 900–915.

Robock, A. (2003). Introduction: Mount Pinatubo as a test of climate feedback mechanisms. *Geophysical Monograph-American Geophysical Union*, *139*, 1–8.

Shang, H. L., & Hyndman, R. J. (2011). Nonparametric time series forecasting with dynamic updating. *Mathematics and Computers in Simulation*, *81*(7), 1310–1324.

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., & Jermyn, I. H. (2009). Looking for shapes in two-dimensional, cluttered point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *31*(9), 1616–1629.

Srivastava, A., Joshi, S. H., Mio, W., & Liu, X. (2005). Statistical shape analysis: Clustering, learning and testing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *27*(4), 590–602.

Srivastava, A., & Klassen, E. P. (2016). *Functional and shape data analysis*. Springer.

Thompson, D., Wallace, J., Jones, P., & Kennedy, J. (2009). Identifying signatures of natural climate variability in time series of global-mean surface temperature: Methodology and insights. *Journal of Climate*, *22*, 6120–6141.

Tucker, J. D. (2014). Functional component analysis and regression using elastic methods [Ph.D. thesis]. Florida State University.

Tucker, J. D., Shand, L., & Chowdhary, K. (2021). Multimodal Bayesian registration of noisy functions using Hamilitonian Monte Carlo. *Computational Statistics and Data Analysis*, *163*, 1–24.

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.

Wang, M., Harris, T., & Li, B. (2022). Asynchronous changepoint estimation for spatially correlated functional time series. *Journal of Agricultural, Biological and Environmental Statistics*, *28*, 157–176.

Wegner, L., & Wendler, M. (2022). Robust change-point detection for functional time series based on *u*-statistics and dependent wild bootstrap. arXiv preprint arXiv:2206.01458.

Woodall, W. H. (2007). Current research on profile monitoring. *Production*, *17*(3), 420–425.

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.