A Geometric Approach for Computing Tolerance Bounds for Elastic Functional Data

J. Derek Tucker^a, John R Lewis^a, Caleb King^a, and Sebastian Kurtek^b

^aStatistical Sciences, Sandia National Laboratories, Albuquerque, NM ^bDepartment of Statistics, The Ohio State University, Columbus, OH

ARTICLE HISTORY

Compiled May 22, 2019

ABSTRACT

We develop a method for constructing tolerance bounds for functional data with random warping variability. In particular, we define a generative, probabilistic model for the amplitude and phase components of such observations, which parsimoniously characterizes variability in the baseline data. Based on the proposed model, we define two different types of tolerance bounds that are able to measure both types of variability, and as a result, identify when the data has gone beyond the bounds of amplitude and/or phase. The first functional tolerance bounds are computed via a bootstrap procedure on the geometric space of amplitude and phase functions. The second functional tolerance bounds utilize functional Principal Component Analysis to construct a tolerance factor. This work is motivated by two main applications: process control and disease monitoring. The problem of statistical analysis and modeling of functional data in process control is important in determining when a production has moved beyond a baseline. Similarly, in biomedical applications, doctors use long, approximately periodic signals (such as the electrocardiogram) to diagnose and monitor diseases. In this context, it is desirable to identify abnormalities in these signals. We additionally consider a simulated example to assess our approach and compare it to two existing methods.

KEYWORDS

Compositional noise; functional data analysis; functional tolerance bounds; functional Principal Component Analysis

1. Introduction

A significant amount of data collected in biomedical applications, process monitoring and reliability engineering is in the form of functions where each data object is a collection of data points over some index (e.g., time or frequency). In these applications, the goal is often to provide estimates on the range in which a certain proportion of the population of functions is expected to fall while accounting for sampling uncertainty. With this goal in mind, this work focuses on developing theoretically sound methods for constructing tolerance bounds for functional data. Tolerance bounds are confidence bounds on quantiles and can be used to construct ranges within which

CONTACT J. Derek Tucker. Email: jdtuck@sandia.gov

'there is $(1 - \alpha)100\%$ confidence that (1 - p)100% of the population of functional data falls' [9]. For an equal-tailed tolerance bound, the upper bound is an upper confidence bound on the 1 - p/2-quantile, while the lower bound is a lower confidence bound on the p/2-quantile. Such a notion of tolerance bounds is important in many different applications.

In applied settings that involve collection and analysis of functional data, it is common to ignore the dataset's functional nature and only extract key scalar features on which to base inferences. For example, in monitoring of an electrocardiogram of a patient, it is often desirable to identify periods when a heart beat is outside of normal rhythm. Key scalar features in this case could be the periodicity of the electrocardiogram or the time between heartbeats. The statistical analysis would then proceed by making inferential statements, such as constructing tolerance bounds, on the key features only. Common examples of other mathematical features extracted for such analyses include local maxima and minima, the number of peaks, or a rate of change at a particular point on the function. As an alternative to extracting and analyzing a finite set of key features, the functional data can be first discretized, and then treated as a finite vector. In this approach, many tools from standard multivariate analysis, including Principal Component Analysis (PCA), become available. However, standard multivariate data analysis approaches ignore the intrinsic infinite-dimensional nature of the data, as well as the strong dependence between neighboring points on the functions. On the other hand, we aim to develop theory/methods that respect these important properties of functional data, and discretize only at the end of the overall process as a necessary step for computer implementation. Another important aspect of the data ignored by standard techniques that discretize at the outset of the analvsis is potential horizontal (warping) variation, explained in the following paragraph. Ignoring such variation can make the analysis of functional data less meaningful.

Alternatively, there has been considerable effort in statistics to develop methods that can analyze functional data objects without significant loss of information due to summarization. Such methodology is known as functional data analysis and has a rich history in Statistics. An excellent introduction to this field is given in several books including [24], [6] and [31]. An interesting aspect of most functional data is that the underlying variability can be ascribed to two sources. These two sources are termed the amplitude (or y or vertical) variability and the phase (or x or horizontal or warping) variability. Capturing these two sources of variability is crucial when modeling and monitoring functional data in a process control architecture, and can greatly affect the construction of tolerance bounds. In this work, we refer to functional data that contains both amplitude and phase variability as *elastic*. This important concept is illustrated in Figure 1 through a simulated example. The observed functions are generated according to the equation $y_i(t) = z_i e^{-(t-a_i)^2/2}, t \in [0,1], i = 1, 2, \dots, 29,$ where z_i are *i.i.d.* $\mathcal{N}(1, (0.05)^2)$ and a_i are *i.i.d.* $\mathcal{N}(0, (1.25)^2)$. The top left panel shows the simulated functions; each sample function is unimodal with slight variability in the height of the peak and large variability in its placement. The relative heights of the peak can be attributed to the amplitude variability, while the different locations of the peak constitute the phase variability. The cross-sectional (pointwise) mean of this data is shown in the top middle panel. This mean ignores the phase variability which results in averaging out of the main unimodal amplitude feature. If tolerance bounds were to be constructed using this cross-sectional approach, the intrinsic shape of the data would be lost, and the bounds would not capture the true underlying variability in the data. Alternatively, the phase variability can be accounted for by first aligning the functions. As an example, the top right panel shows time-aligned functions. The



Figure 1. Demonstration of amplitude and phase variability in functional data. Top left: Original functions. Top middle: Cross-sectional mean without alignment. Top right: Aligned functions (amplitude). Bottom left: Warping functions (phase). Bottom middle: Cross-sectional mean after alignment.

alignment involves a transformation of the horizontal axis via warping functions shown in the bottom left panel. The aligned functions capture the amplitude variability while the warping functions capture the phase variability. The cross-sectional mean of the aligned functions (amplitude) is shown in the last panel, where the sharp unimodal structure of the original data is retained. The tolerance bounds developed in this work provide bounds that maintain the shape of the data by accounting for both directions of variability.

1.1. Past Work and Contributions

Recently, Storlie et al. [33] developed a method to test the shape of a population of curves using a B-Spline basis, and a hierarchical Gaussian process approach to form confidence intervals. Rathnayake and Choudhary [25] developed tolerance bounds for functional data using functional Principal Component Analysis (fPCA). Further, Sun and Genton [34] developed a boxplot display for functional data, which provides a nice visualization technique for a sample of functions. This approach can also detect functional outliers. A drawback of all of the aforementioned methods is that they do not account for potential phase variability in functional data, i.e., they assume that the data either (1) does not need to be aligned, or (2) has already been aligned in a pre-processing step (generally using some criterion that is unrelated to subsequent data analysis). The first assumption is unrealistic in process control and biomedical applications, while the second approach results in suboptimal solutions due to the disjointedness of the alignment and data analysis procedures. A more systematic approach is to develop methods that build the alignment step into the statistical procedure of interest.

A couple of papers in recent literature have taken into account the phase variability in the application of monitoring functional data. Lewis et al. [19] expanded upon the generative model presented in [37], and used a bootstrap approach to generate tolerance bounds. However, the phase variability in this approach was discarded, and as a result, the generated tolerance bounds did not maintain the shape of the original data. Grasso et al. [8] developed an idea similar to Storlie et al. [33] for process monitoring. In this work, they built the alignment of the functional data into their procedure, where they used a parametric model for warping functions in conjunction with fPCA. However, the \mathbb{L}^2 metric they used for alignment and construction of the subsequent bounds has serious theoretical limitations (e.g., the pinching effect) as described in [22]. Additionally, the use of a parametric model for warping functions may not be flexible enough to achieve good alignment in general applications.

In a recent paper, Xie et al. [39] developed an alternative visualization approach to the method of [34], that directly takes warping variability into account in the construction of the boxplot displays. Their method is based on the Riemannian geometry of the amplitude and phase representation spaces and builds upon the general elastic functional data analysis framework presented in [32], [15], and [37]. Our approach also builds upon that work by generalizing to arbitrary quantile estimates while accounting for sampling uncertainty.

In this paper, we present two methods for computing tolerance bounds for functional data observed under random warping variability. First, using a modification of the joint fPCA approach developed in [18], we create a generative model for amplitude-phase functions. One can then sample from this model and use the bootstrap approach to construct tolerance bounds on the amplitude and phase components separately. The benefit of this approach is that the tolerance bounds maintain the shape of the original data while capturing both amplitude and phase variability. This approach is similar to [25], but their method fails to account for the phase variability, which greatly affects the structure of the tolerance bounds.

The second approach also uses the previously mentioned joint fPCA method, and constructs the tolerance bounds in the joint amplitude-phase coefficient space (after projecting onto the lower dimensional space spanned by the eigenvectors of the covariance). To do this, a multivariate Gaussian model is assumed for joint amplitude-phase fPCA coefficients and a tolerance factor is computed. A new function, which one would like to test, can be projected onto the same coefficient space and its tolerance score can be compared to the tolerance factor. This approach is similar in spirit to the elastic functional statistical process control (fSPC) presented in [36].

The rest of this paper is organized as follows. In Section 2, we review the relevant material from elastic functional data analysis and develop a joint amplitude-phase fPCA model. Section 3 describes the two methods for constructing tolerance bounds for elastic functional data. In Sections 4 and 5, we report the results of applying the proposed approach to a simulated dataset and two real datasets from different application domains. Finally, we close with a brief summary and some ideas for future work in Section 6.

2. Combined Phase-Amplitude fPCA

We begin by giving a short review of the combined phase-amplitude fPCA method of [18], with a slight modification which will be described clearly in later sections. Their method is based on the functional data analysis approach outlined in [32], [15] and [37]; see those references for more details.

Let f be a real-valued function with the domain [0, 1]; this domain can be easily generalized to any other compact subinterval of \mathbb{R} . For concreteness, only functions that are absolutely continuous on [0, 1] will be considered and we let \mathcal{F} denote the set of all such functions. In practice, since the observed data are discrete anyway, this assumption is not a restriction. Also, let Γ be the set of orientation-preserving diffeomorphisms of the unit interval [0, 1]: $\Gamma = \{\gamma : [0, 1] \to [0, 1] | \gamma(0) = 0, \gamma(1) =$ $1, \gamma$ is a diffeomorphism}. Elements of Γ play the role of warping functions. For any $f \in \mathcal{F}$ and $\gamma \in \Gamma$, the composition $f \circ \gamma$ denotes the time warping of f by γ . With the composition operation, the set Γ is a Lie group with the identity element $\gamma_{id}(t) = t$. This is an important observation since the group structure of Γ is seldom utilized in past papers on functional data analysis.

As described in [37], there are two metrics to measure the amplitude and phase variability of functions. These metrics are proper distances, one on the quotient space \mathcal{F}/Γ (i.e., amplitude) and the other on the group Γ (i.e., phase). The amplitude or y-distance for any two functions $f_1, f_2 \in \mathcal{F}$ is defined to be

$$d_a(f_1, f_2) = \inf_{\gamma \in \Gamma} \| q_1 - (q_2 \circ \gamma) \sqrt{\dot{\gamma}} \|,$$
(2.1)

where $q(t) = \operatorname{sign}(\dot{f}(t))\sqrt{|\dot{f}(t)|}$ is known as the square-root slope function (SRSF) (\dot{f} is the time derivative of f). The optimization problem in Equation 2.1 is most commonly solved using a Dynamic Programming algorithm; see [26] for a detailed description. If f is absolutely continuous, then $q \in \mathbb{L}^2([0, 1], \mathbb{R})$ [26], henceforth denoted by \mathbb{L}^2 . For properties of the SRSF and the reason for its use in this setting, we refer the reader to [30], [22] and [17]. Moreover, it can be shown that for any $\gamma_1, \gamma_2 \in \Gamma$, we have $d_a(f_1 \circ \gamma_1, f_2 \circ \gamma_2) = d_a(f_1, f_2)$, i.e., the amplitude distance is invariant to function warping.

2.1. 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 to model the warping functions. The following framework was previously used in various settings including (1) modeling re-parameterizations of curves [29], (2) putting prior distributions on warping functions [13] and [20], (3) studying execution rates of human activities in videos [38], and many others. It is also very closely related to the squareroot representation of probability density functions introduced by [2], and later used for various statistical tasks (see e.g., [14], [27] and [28]).

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 SRSF defined earlier, and 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



Figure 2. Depiction of the SRSF space of warping functions as a sphere and a tangent space at ψ_1 .

 $\gamma(t) = \int_0^t \psi(s)^2 ds$. 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 ψ s 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, i.e., the phase distance, is exactly the arc-length between their corresponding SRSFs on Ψ :

$$d_p(\gamma_1, \gamma_2) = d_{\psi}(\psi_1, \psi_2) \equiv \cos^{-1}\left(\int_0^1 \psi_1(t)\psi_2(t)dt\right) .$$
 (2.2)

Figure 2 depicts the SRSF space of warping functions as a unit sphere [35].

2.2. Mapping to the Tangent Space at Identity Element

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

$$\exp_{\psi}(v) = \cos(\|v\|)\psi + \sin(\|v\|)\frac{v}{\|v\|},$$
(2.3)

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_{ψ}^{-1} : $\Psi \mapsto T_{\psi}(\Psi)$, is defined as

$$\exp_{\psi}^{-1}(\psi_1) = \frac{\theta}{\sin(\theta)}(\psi_1 - \cos(\theta)\psi), \qquad (2.4)$$

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 2. 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. For details on the definition of the sample Karcher mean and how to compute it, please refer to [37].

2.3. Model for Combined Functional Principal Components

To model the association between the amplitude of a function and its phase, Lee and Jung [18] use a concatenated function g^C on the extended domain [0, 2] (for some C > 0). The domain is extended to accommodate a combination of the aligned (amplitude) function and the warping (phase) function. The motivation for the combined function is to be able to properly handle properly the correlation between the phase and amplitude components in functional data. Since the domain is [0, 1] for both, Lee and Jung [18] define the function g^C on the extended domain as follows:

$$g^{C}(t) = \begin{cases} f^{*}(t), & t \in [0, 1) \\ Cv(t-1), & t \in [1, 2], \end{cases}$$
(2.5)

where f^* only contains the function's amplitude (i.e., after groupwise alignment to the mean via SRSFs). Furthermore, they assume that $g^C \in \mathbb{L}^2([0, 2], \mathbb{R})$. The parameter C is introduced to adjust for the scaling imbalance between f^* and v. In the current work, we make a slight modification to their method. In particular, it seems more appropriate to construct the function g^C using the SRSF q^* of the aligned function f^* , since q^* is guaranteed to be an element of \mathbb{L}^2 . Thus, with a slight abuse in notation, we proceed with the following joint representation of amplitude and phase:

$$g^{C}(t) = \begin{cases} q^{*}(t), & t \in [0, 1) \\ Cv(t-1), & t \in [1, 2], \end{cases}$$
(2.6)

where C is again used to adjust for the scaling imbalance between q^* and v.

Henceforth, we assume that q^* and v are both sampled using T points, making the dimensionality of $g^C \in \mathbb{R}^{2T}$. Then, given a sample of amplitude-phase functions $\{g_1^C, \ldots, g_n^C\}$, and their sample mean $\hat{\mu}_g^C = [\hat{\mu}_{q^*} \quad \hat{\mu}_v^C]$, we can compute the sample covariance matrix as

$$K_g^C = \frac{1}{n-1} \sum_{i=1}^n (g_i^C - \hat{\mu}_g^C) (g_i^C - \hat{\mu}_g^C)^\mathsf{T} \in \mathbb{R}^{(2T) \times (2T)} .$$
 (2.7)

The Singular Value Decomposition $K_g^C = U_g^C \Sigma_g^C (U_g^C)^{\mathsf{T}}$ provides the joint principal directions of variability in the given amplitude-phase functions as the first $p \leq n$

columns of U_g^C . These can be converted back to the original representation spaces (\mathcal{F} and γ) using the mappings defined earlier. Moreover, one can calculate the observed principal coefficients as $\langle g_i^C, U_{g,j}^C \rangle$, for the i^{th} function and the j^{th} principal direction of variability. The superscript of C is used to denote the dependence of the principal coefficients on the scaling factor.

This framework can be used to visualize the joint principal geodesic paths. First, the matrix U_g^C is partitioned into the pair $(U_{q^*}^C, U_v^C)$. Then, the amplitude and phase paths within one standard deviation of the mean are computed as

$$q_{\tau,j}^{*C} = \hat{\mu}_{q^*} + \tau \sqrt{\Sigma_{g,jj}^C} U_{q^*,j}^C, \qquad (2.8)$$

$$v_{\tau,j}^C = \tau \frac{\sqrt{\sum_{g,jj}^C} U_{v,j}^C}{C} U_{v,j}^C , \qquad (2.9)$$

where $\tau \in [-1, 1]$, $\Sigma_{g,jj}$ and U_j^C are the j^{th} principal component variance and direction of variability, respectively (note that the mean $\hat{\mu}_v^C$ is always zero). Then, one can obtain a joint amplitude-phase principal path by composing $f_{\tau,j}^{*C}$ (this is the function corresponding to SRSF $q_{\tau,j}^{*C}$) with $\gamma_{\tau,j}^C$ (this is the warping function corresponding to $v_{\tau,j}^C$).

 $v_{\tau,j}^C$). The results of the above procedure clearly differ for variations of C. For example, using small values of C, the first few principal directions of variability will capture more amplitude variation, while for large values of C, the leading directions reflect more phase variation. Lee and Jung [18] present a data-driven method for estimating C for a given sample of functions. We use this approach in the current work to determine an appropriate value of C. Other approaches to choosing the value of C include (1) cross-validation metrics such as prediction performance, or (2) manual tuning based on which variability the user wants to emphasize in the statistical analysis.

2.4. Statistical Model of Functions via fPCA

There are several possibilities to develop statistical models for capturing the phase and amplitude variability in functional data. Once we have obtained the fPCA coefficients for the combined phase and amplitude variability we can impose probability models directly on the coefficients. This in turn induces a distribution on the function space \mathcal{F} . Let $c = (c_1, \ldots, c_k)$ be the k dominant principal coefficients of the combined model as described in the previous two sections. Recall that the coefficients are constructed using $c_j = \left\langle g^C, U_{g,j}^C \right\rangle$. The number k is determined by the user and can be selected in different ways: (1) by minimizing cross-validated fPCA reconstruction error, (2) by retaining the smallest number of fPCA coefficients that explain at least X% of the variability in the given data (X is usually chosen as a large number, e.g., 90 or 95), or (3) by iteratively comparing tolerance bounds using an increasing number of fPCA components and stopping when the bounds change negligibly. In our approach, for simplicity, we use method (2) with X = 90.

The vector c is modeled using a multivariate Gaussian probability distribution with zero mean and covariance Σ , i.e., $c \sim \mathcal{N}_k(0, \Sigma)$. By construction of the principal coefficients, the mean vector is zero and the covariance is a $k \times k$ diagonal matrix. The diagonal elements of the covariance are estimated directly using the eigenvalues of the sample covariance matrix, $\hat{\sigma}_1^C, \ldots, \hat{\sigma}_k^C$. The model on the fPCA scores induces a probability model on \mathcal{F} and provides a means of efficiently sampling functions that exhibit the amplitude and phase variability of the original data.

3. Functional Tolerance Bounds

In this section, we provide two methods for calculating tolerance bounds for functional data in the presence of warping variability. In general, we seek bounds which guarantee that with $(1 - \alpha)100\%$ confidence, (1 - p)100% of the data falls within these bounds [9]. Our first approach uses the parametric bootstrap to construct tolerance bounds sampling from the model on the fPCA coefficients. The second method provides the tolerance bound in the fPCA coefficient space using a tolerance factor based on the multivariate Gaussian model. We provide a detailed description of both of these procedures next.

3.1. Method 1: Bootstrapped Geometric Tolerance Bounds

First, we construct statistical bounds using bootstrapping from the fPCA-based model described in Section 2.4, and provide a means of characterizing the uncertainty in the original functional data. Bootstrapping refers to repeated sampling from the model, and this process can be used to construct confidence bounds for essentially any quantity of interest. For a detailed overview of statistical bootstrap techniques see [4] and [5].

For two-sided tolerance bounds, there is both an upper and a lower bound. The upper tolerance bound is an upper confidence bound on an upper population quantile. The lower tolerance bound is a lower confidence bound on a lower population quantile. In this sense, tolerance bounds are simply confidence bounds on population quantiles. Construction of equal-tailed tolerance bounds using the bootstrap approach is described as the following procedure:

- 1. Sample n functions from the constructed generative model described in Section 2.4, resulting in a random sample g_i^C , i = 1, ..., n. 2. From g_i^C extract the amplitude functions (SRSFs) q_i^* and vectors v_i . The random
- warping functions can be constructed using $\gamma_i(t) = \int_0^t (\exp_{\hat{\mu}_{\psi}}(v_i(s)))^2 ds.$
- 3. Estimate the p/2 and (1-p/2) quantiles of the set of random SRSF-based amplitudes and random warping functions, denoted by $(q_{p/2}^*, q_{1-p/2}^*)$ and $(\gamma_{p/2}, \gamma_{1-p/2})$, respectively. There are many ways to estimate quantiles for functional data, with the cross-sectional (pointwise) approach being most popular [19]. We propose to use the geometric method of Xie et al. [39], which relies on the Riemannian geometry of the amplitude and phase spaces. In that paper, the authors only compute quartiles, but their method can be easily extended to calculate general quantiles.
- 4. Repeat steps 1-3 S times for a large S (S should be large enough relative to α to provide stable bounds). This results in a collection (of size S) of $(q_{p/2}^*, q_{1-p/2}^*)$
- and $(\gamma_{p/2}, \gamma_{1-p/2})$. 5. Calculate the $\alpha/2$ and $(1 \alpha/2)$ quantiles of the S samples of $(q_{p/2}^*, q_{1-p/2}^*)$ and $(\gamma_{p/2}, \gamma_{1-p/2})$. These quantiles form $(1-\alpha)100\%$ bootstrap tolerance bounds with (1-p)100% coverage.
- 6. The amplitude tolerance bounds are pulled back to the quotient space \mathcal{F}/Γ via

integration. For display of the quantiles we use the surface method of Xie et al. [39] for both the amplitude and phase components.

3.2. Method 2: Combined fPCA-based Tolerance Region

The second approach to construct tolerance bounds is to directly use the fPCA-based multivariate Gaussian model that is imposed on the principal coefficients. Then, one can construct multivariate tolerance regions on the coefficient space using the model given in [12]. We provide more details of this approach next.

Let $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{R}^d$ be a random sample from a k-variate Gaussian distribution with mean vector μ and covariance matrix Σ . The sample mean vector $\bar{\mathbf{x}}$ and the sums of squares and cross-product matrix A are defined as:

$$\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{x}_i, \quad A = \sum_{i=1}^{n} (\mathbf{x}_i - \bar{\mathbf{x}}) (\mathbf{x}_i - \bar{\mathbf{x}})^{\mathsf{T}}.$$

A tolerance region that contains at least p proportion of the data from a $\mathcal{N}_k(\mu, \Sigma)$ distribution with β confidence is given by

$$\{\mathbf{x}: (n-1)(\mathbf{x}-\bar{\mathbf{x}})^{\mathsf{T}}A^{-1}(\mathbf{x}-\bar{\mathbf{x}}) \le b\}$$
(3.1)

The parameter b is known as the tolerance factor, and is determined by the probability condition $P_{\bar{\mathbf{x}},A}(P_{\mathbf{x}}((n-1)(\mathbf{x}-\bar{\mathbf{x}})^{\mathsf{T}}A^{-1}(\mathbf{x}-\bar{\mathbf{x}}) \leq b|\bar{\mathbf{x}},A) \geq p) = \beta$. The exact method of computing b is known to be extremely difficult and there are multiple approximations that have been proposed in the literature (see [11] for multiple methods). In this work, we use the approach of Krishnamoorthy and Mondal [12] due to its known accuracy and precision. Note that in the fPCA coefficient space $\bar{x} = 0$ by construction, which simplifies the computation. Once the tolerance factor is computed, a new sample function can be tested against this factor using Equation 3.1.

4. Simulation Results

In this section, we present results on a simulated dataset and compare the proposed approaches to the recent methods of Lewis et al. [19] and Rathnayake and Choudhary [25]. The method of Lewis et al. is most closely related to our work and is thus important to compare to. In all of the plots in the following sections, we re-scale the domain of the functional observations and warping functions to [0, 1] for simplicity.

4.1. Method 1: Bootstrapped Geometric Tolerance Bounds

First, we provide a numerical simulation for the bootstrapped geometric tolerance bounds method. For this purpose, we generate data previously used in [10]. The individual functions are given by:

$$y_i(t) = z_{i,1}e^{-(t-1.5)^2/2} + z_{i,2}e^{-(t+1.5)^2/2}, \ t \in [-3,3], \ i = 1, 2, \dots, 21, \dots, 21$$



Figure 3. Alignment of the simulated data set. (a) Original functions. (b) Aligned functions (amplitude). (c) Warping functions (phase).

where $z_{i,1}$ and $z_{i,2}$ are *i.i.d.* $\mathcal{N}(1, (0.25)^2)$. Each of the simulated functions are then warped according to:

$$\gamma_i(t) = 6\left(\frac{e^{a_i(t+3)/6} - 1}{e^{a_i} - 1}\right) - 3, \text{ if } a_i \neq 0, \text{ otherwise } \gamma_i = \gamma_{id}$$

where $\gamma_{id}(t) = t$ is the identity warping. Here, a_i are equally spaced between -1 and 1, and the observed functions are computed by composition using $f_i = y_i \circ \gamma_i$. A set of 30 such functions forms the original data and is shown in Figure 3(a). The aligned functions (amplitude) and corresponding warping functions (phase) are shown in Figure 3(b) and (c), respectively.

Given that we have the phase-amplitude separation of the simulated data, we can calculate the principal directions of variability using the combined fPCA method defined in Section 2. Figure 4 shows the results of this procedure on the simulated dataset presented in Figure 3. We plot the joint amplitude-phase principal geodesic paths using Equations (2.8) and (2.9), for $\tau = -1, 0, 1$ and j = 1, 2, 3. The first three singular values for this data are 4.97, 4.24 and 0.77, with the rest being fairly small. The first four principal directions of variability accounted for 90% of the overall variability in the data. The first principal direction of variability displayed in panel (a) mostly corresponds to phase variation, while the second direction in panel (b) captures the height (amplitude) variability in panel (c).

Next, we calculate tolerance bounds for the simulated data using the bootstrap approach presented in Section 3.1. Using the combined amplitude and phase fPCA computed in the previous section, we impose a multivariate Gaussian model using the first four principal directions. We perform 500 bootstrap re-samples; within each iteration, we sample 30 functions to calculate the tolerance bounds. We use this procedure to compute the tolerance bound with 99% coverage with a confidence level of 95%.

Figure 5 presents the tolerance bounds for the (a) amplitude and (b) phase. The upper bound is shown in red and the lower bound is shown in blue (the median is also plotted in green). Given the structure of the original data, the bounds are intuitive for the phase component. However, the amplitude bounds overlap, i.e., for the first peak the lower bound is above the upper bound while for the second peak, the lower bound is below the upper bound. While this result may seem counterintuitive at first, it stems from the geometric approach to generate the quantiles. In particular, the amplitude



Figure 4. Joint principal directions (PD) of variability for the simulated dataset for $\tau = -1$ (blue), 0 (red), 1 (green).



Figure 5. Bootstrapped geometric tolerance bounds for the simulated dataset.

component captures the "shape" of the functions, and the quantiles are constructed directly on the amplitude space. Thus, it is more appropriate to think of the bounds as points on that space rather than pointwise (cross-sectional) bounds as currently done in the literature. As a result, there is no guarantee that the bounds will not overlap as they are constructed to capture the global "shape" tolerance region on the amplitude function space. This will also be seen in the real data applications.

Xie et al. [39] provide a detailed commentary of this in their paper and propose a surface plot using the proper metrics to display the quantiles. We thus use the same approach here, and present such surface plots for the amplitude and phase components in Figure 6(a) and (b), respectively, with the bounds shown in red. In panel (a), the amplitude component of the median and bounds is shown along the z-axis. Furthermore, the functions are separated according to the pairwise amplitude distances between them. Based on this surface plot we now have a natural view of the amplitude variability in the data, with the bounds separated by appropriate distances for a clear display. Panel (b) provides a similar display for the phase component. In this case, it is more effective to display the difference between the median/bounds and the identity element γ_{id} . Thus, as expected, the median is very close to the constant function 0. The red warping bounds reflect natural variability in the original data and are approximately equidistant from the median.



Figure 6. Surface plots of the bootstrapped geometric tolerance bounds for the simulated dataset.



Figure 7. Tolerance bounds for the simulated dataset constructed using the methods in [25] and [19].

Figure 7(a) presents the tolerance bounds calculated using the approach of Rathnayake and Choudhary [25]. Comparing these bounds with those presented in Figure 5, we see some striking structural differences. In particular, in Figure 7(a), it is difficult to determine the relative contributions of amplitude and phase to the tolerance bounds. Moreover, the upper and lower bounds are not representative of the actual shape of the original functions, and both exhibit more than just the two peaks found in the data. This is due to the fact that the bounds were computed in a cross-sectional manner without accounting for warping variability. Figure 7(b) presents the tolerance bounds calculated using the approach of Lewis et al. [19]. Again, comparing these bounds with those presented in Figure 5, we see some structural differences. Both the upper and lower bounds do not accurately capture the true underlying shape of the given data.

Next, we generate an additional set of random functions from our combined fPCA model to check the coverage of the calculated tolerance bounds. We study the per-

	99%	95%	90%
Amplitude	100.00	97.60	93.80
Phase	99.80	98.40	90.00
Rathnayake and Choudhary [25]	75.23	70.12	67.43
Lewis et al. $[19]$	71.01	67.26	64.58

Table 1. Simulated confidence values of 90% coverage tolerance using the bootstrap-based approach (top). Corresponding estimated confidence values using the methods in [25] and [19].



Figure 8. Surface plots of the bootstrapped geometric tolerance bounds for the simulated dataset for different confidence levels (90%=green, 95%=blue, 99%=red). An outlier amplitude and phase function is shown in purple.

formance of our method based on three sets of tolerance bounds at the 90% coverage rate with confidence levels of 90%, 95% and 99%. We first generate 100 random functions; for each function, we compute its SRSF (q) and warp it to $\hat{\mu}_q$ to extract the corresponding warping function γ and the aligned SRSF (q^{*}). We then compute the 90% quantile, and compare q^* and γ to the corresponding tolerance bounds (the entire function must fall within the tolerance bounds). This process is repeated 500 times and the corresponding estimated confidence values are listed in Table 1 for each of the corresponding true confidence values. For both amplitude and phase tolerance bounds, the calculated confidence values are slightly higher than expected; this suggests that the bounds we compute are conservative. Nonetheless, the simulated values are very close to the correct confidence levels. For comparison, we also list the estimated confidence values computed using the methods defined in [25] and [19]. For both methods, the estimated confidence values are much lower than expected. This is most likely due to the fact that these approaches do not appropriately account for the amplitude and phase variabilities in the given data.

Next, we show the effect of the confidence level on the tolerance bounds. Figure 8 presents surface plots of the tolerance bounds for (a) amplitude and (b) phase for three different settings. The red curves are the 99% confidence tolerance bounds, the blue curves are the 95% confidence tolerance bounds, and the green curves are the 90% confidence tolerance bounds. All of these bounds were generated for 90% coverage. For



Figure 9. Effects of the number of principal directions of variability retained by the combined amplitude and phase fPCA model on the resulting tolerance bounds for 99% coverage with 95% confidence.

both amplitude and phase, as the confidence level increases the bounds move outward in distance from the median as one would expect. For the phase, the 95% and 90% confidence tolerance bounds are extremely close. The purple amplitude function in panel (a) falls outside the amplitude tolerance bounds for all three confidence levels; in panel (b), its phase component also falls outside the 90% confidence tolerance bounds for phase. The surface plots are able to accurately display the bounds by nicely separating them according to amplitude or phase distance.

To show the effects of the number of principal directions of variability retained by the combined amplitude and phase fPCA model on the resulting tolerance bounds, we applied the bootstrap method for 1-4 principal directions. The results are shown in Figure 9. As the dimensionality of the model increases to 3 or 4, we see the bounds remain relatively unchanged showing good stability. With enough principal components (approximately four in this example), the constructed tolerance bounds are able to more accurately represent the variability found in the original data, for both phase and amplitude.

Lastly, we compare the proposed tolerance bounds to those of Rathnayake and Choudhary [25] and Lewis et al. [19] for different confidence levels at 99% coverage. These results are presented in Figure 10 for 90% (top), 95% (middle) and 99% (bottom) confidence levels. The proposed approach (labeled Elastic) is able to accurately capture the shape of the given data in the tolerance bounds for all confidence levels; we also note that, overall, the bounds spread outward as the confidence level increases. As seen previously, both of the competing methods (labeled Rathnayake for [25] and Lewis for [19]) produce multimodal tolerance bounds, where the various modes increase or decrease as the confidence level changes. This comparison was performed only on the amplitude, as the methods in [25] and [19] do not consider phase and amplitude separately. Figure 11 provides the phase tolerance bounds for the same varying levels of confidence at 99% coverage for the proposed method.

4.2. Method 2: Combined fPCA-based Tolerance Region

In this section, we calculate tolerance bounds for the simulated dataset using the fPCA basis approach presented in Section 3.2. Again, we impose a multivariate Gaussian model using the first four principal directions of variability. We calculate a tolerance



Figure 10. Comparison of the proposed method (Elastic) to [25] (Rathnayake) and [19] (Lewis). We use the simulated data to construct tolerance bounds for confidence levels 90% (top), 95% (middle), and 99% (bottom) at 99% coverage.



Figure 11. Proposed phase tolerance bounds, constructed using the simulated data, for confidence levels 90%, 95% and 99% at 99% coverage.

region with 99% coverage with a confidence level of 95%. In contrast to the bootstrap approach described in the previous section, this method constructs the tolerance region directly on the fPCA coefficient space.

First, we must calculate the tolerance factor, which is done using Algorithm 2 in [12] with 100,000 iterations. The resulting tolerance factor b for this dataset, and the retained dimension of four, is 32.0027. For different dimensions of the multivariate Gaussian model, the tolerance factor can be easily calculated using the above-mentioned algorithm or via the tables provided by [12]. Each new function that needs to be tested can be projected onto the fPCA basis. Based on the resulting fPCA coefficients, we



Figure 12. Histogram of tolerance scores for each function in the simulated dataset after projection onto the first four fPCA basis functions.

	99%	95%	90%
fPCA	99.20	98.40	97.00

Table 2. Simulated confidence values of 90% coverage tolerance using the combined fPCA-based approach.

compute the function's tolerance score using Equation 3.1 and compare it to the tolerance factor. Figure 12 presents a histogram of the tolerance scores for all 21 functions in the simulated dataset. All of the scores fall well-below the calculated tolerance factor; this is expected as the tolerance bound has 99% coverage, and was calculated using the same data.

Again, we generate an additional set of random functions from our model to check the coverage of the calculated tolerance region. As previously, we study the performance of our method based on three sets of tolerance regions at the 90% coverage rate with confidence levels of 90%, 95%, and 99%. We first generate 100 random functions, and for each function, we compute its SRSF (q). The SRSF is then warped to $\hat{\mu}_q$ to extract the corresponding warping function γ and aligned SRSF (q^*). The warping function is mapped to the tangent space at the warping mean, which can be computed using the algorithm described in [39]. The function g^C is then calculated and projected into the fPCA coefficient space. We calculate the 90% quantile in the coefficient space as well as the tolerance scores for each of the functions; the tolerance scores are then compared to the tolerance factor. We repeat this process 500 times, and report the estimated confidence values in Table 2 for each tolerance bound. Similarly to the bootstrap method, the simulated confidence values are conservative.

5. Applications to Real Data

Here, we present results on two real datasets: (1) axial weld data and (2) PQRST complexes extracted from electrocardiogram (ECG) signals [16]. For each of the examples, we study the effectiveness of the proposed approach in the calculation of the tolerance bounds, specifically in capturing the amplitude and phase variabilities.



Figure 13. Alignment of the WRS axial contour weld data. (a) Original functions. (b) Aligned functions (amplitude). (c) Warping functions (phase).



Figure 14. Bootstrapped geometric tolerance bounds for the WRS axial contour weld data.

5.1. Tolerance Bounds for Weld Data

Weld residual stress (WRS) is a main component crack formation in safety critical pipe welds within nuclear power plants [1]. The stress cannot be measured directly and involves releasing strain from a weld specimen and converting measured displacements to stress using a mathematical model. Two measurement methods, both destructive to the specimen, are the deep hole drilling method and the contour method [23], [21]. The measurements can help assess the validity of WRS predictions for new welds [19]. Figure 13(a) shows original contour method measurements along the axial direction of a weld for five different locations along the circumference of the weld. The values represent stress as a function of normalized weld depth. Though the phase variability is small relative to the amplitude variability, it is important to account for both since WRS predictions affect crack-growth calculations used in plant safety assessments [1]. With a small dataset, it is difficult to assess the validity of tolerance bound coverage rates. However, our motivation for tolerance bounds is to provide a preliminary estimate of the uncertainty in the WRS measurements which would be assessed by subject matter experts before applying them to crack-growth calculations. Figure 13(b) and (c) show the aligned functions (amplitude) and warping functions (phase), respectively. The main source of phase variability in this data comes from the difference in location of the first peak across measurements.



Figure 15. Surface plots of the bootstrapped geometric tolerance bounds for the WRS axial contour weld data.

	Tolerance Factor	No of PCs	Tolerance score mean (sd)	Sample Size
Weld	154.96	2	1.60(0.80)	5
ECG	17.29	4	3.95~(2.68)	80
Table 2	Summany of regults obta	ined using the ac	mbined fPCA based televence region	approach

 Table 3.
 Summary of results obtained using the combined fPCA-based tolerance region approach.

Figure 14 presents the tolerance bounds for (a) amplitude and (b) phase. The tolerance bounds were again calculated using 500 bootstrap resamples with a sample size of 5 functions in each iteration. We chose the two leading principal directions of variability as they captured over 90% of the overall variability. Both tolerance bounds have 99% coverage with a 95% confidence level. Figure 15 shows the surface plot of the tolerance bounds for the phase and amplitude components. This is the same dataset used in [19], and it is clear that our tolerance bounds are more representative of the shape in the original data. Furthermore, the separate bounds on phase and amplitude show the overall contribution of each source of variability. This impact is lost in the tolerance bounds generated by Lewis et al. [19] as shown in Figure 16(b).

Figure 16(a) presents the tolerance bounds calculated using the approach of Rathnayake and Choudhary [25]. The first peak and valley in these bounds are not as sharp as those computed using the proposed method and show in Figure 14(a). Figure 16(b) presents the tolerance bounds calculated using the approach of Lewis et al. [19]. While the features of these bounds are sharper than those in Figure 16(a), they fail to capture the phase variability in the given data. It also appears that the main source of variability in these bounds is simply vertical translation, with very little variability in the shape.

The first row in Table 3 presents the calculated tolerance factor for this dataset for a tolerance region with 99% coverage and 95% confidence. As before, we again used two principal directions of variability in the combined fPCA model to compute the tolerance factor. In this case, since the dataset contains only five samples, the tolerance factor is quite large due to sampling uncertainty. The mean of the tolerance scores computed for each of the functions in this data is smaller than the tolerance



Figure 16. Tolerance bounds for the WRS axial contour weld dataset constructed using the methods in [25] and [19].

factor. However, with a larger sample size, we would expect the tolerance factor to be much smaller.

5.2. PQRST Complexes from ECG Biosignals

The electrocardiogram (ECG) data used in this work was obtained from the PTB Diagnostic ECG Database [3] on PhysioNet [7]. The ECG is a medical diagnostic tool that is routinely used to monitor the function of the heart, and is standard for diagnosing and monitoring various heart diseases and conditions, including myocardial infarction. The dataset considered in this work consists of 80 PQRST complexes segmented from a long ECG signal using the method presented in Kurtek et al. [16]. Each complex corresponds to a single heartbeat where PQRST refer to the five peak and valley features (P=slight first peak, Q=sharp first valley, R=sharp second peak, S=sharp second valley and T=slight third peak). While each PQRST complex corresponding to a healthy heartbeat contains these five features, the magnitude of the peaks and valleys can be significantly different across subjects (i.e., amplitude variability); these magnitudes can also vary for different heartbeats within a single subject, albeit not to the same degree. Furthermore, the timing of these features can also be quite different (i.e., phase variability). Phase variability in this application corresponds to different timings and durations of heartbeats across individuals. This motivates the use of the proposed method to construct separate tolerance bounds for the amplitude and phase components based on this data. Figure 17(a) displays the 80 segmented PQRST complex functions. It is clear that there is significant phase variability in this data as the five features are not well-aligned. The aligned functions (amplitude) and corresponding warping functions (phase) are shown in panels (b) and (c), respectively.

Figure 18 presents the bootstrapped tolerance bounds for (a) amplitude and (b) phase. We again use 500 bootstrap resamples with a sample size of 50. The number of principal directions of variability chosen from the combined fPCA model was four; these directions captured over 90% of the overall variability in the data. Both tolerance bounds have 99% coverage with a 95% confidence level. Figure 19 shows the corresponding surface plots. The amplitude tolerance bounds capture the relative sizes of the three peaks and two valleys. This is well-demonstrated in the surface plot in Figure 19(a). The phase tolerance bounds exhibit the variability in the location of the



Figure 17. Alignment of the PQRST ECG data. (a) Original functions. (b) Aligned functions (amplitude). (c) Warping functions (phase).



Figure 18. Bootstrapped geometric tolerance bounds for the PQRST ECG dataset.

three peaks.

Figure 20(a) presents the tolerance bounds calculated using the approach of Rathnayake and Choudhary [25]. The bounds computed using this approach completely lose the structure present in the original data. This is precisely due to the fact that there is considerable phase variability in the given PQRST ECG signals. As before, it is difficult to determine the relative contributions of amplitude and phase to the computed tolerance bounds, and what the lower bound actually means in terms of the semantic features of the PQRST complex. Figure 20(b) presents the tolerance bounds calculated using the approach of Lewis et al. [19]. Qualitatively, these bounds appear better than those in panel (a). However, comparing these bounds with those presented in Figure 18, we still see major distortions of the PQRST features.

The last row in Table 3 reports the calculated tolerance factor for the PQRST ECG dataset for 99% coverage with 95% confidence. We used four principal directions of variability in the combined amplitude-phase fPCA model to compute the tolerance factor. Again, the mean tolerance score of the functions in this dataset is smaller than the computed tolerance factor.



Figure 19. Surface plots of the bootstrapped geometric tolerance bounds for the PQRST ECG data.



Figure 20. Tolerance bounds for the PQRST ECG dataset constructed using the methods in [25] and [19].

6. Discussion and Future Work

We presented two methods for computing tolerance bounds for elastic functional data, i.e., functional data with random warping variability. For both methods, we used a combined amplitude and phase functional Principal Component Analysis model. The fPCA was used to define a convenient generative model, which is easy to sample from. This enabled the implementation of an efficient bootstrapping procedure to generate geometrically-motivated tolerance bounds. The second approach used the multivariate Gaussian fPCA model directly to define tolerance regions, and to compute a corresponding cutoff value called the tolerance factor. Therefore, one can easily test whether a function falls inside or outside this tolerance region by computing a simple tolerance score on the fPCA coefficient space. We demonstrated the applicability of these two approaches on a simple simulated example as well as two real data examples wherein the observed functional data has clear amplitude and phase variation. In this work, we have focused on accounting for phase and amplitude variability when constructing tolerance bounds for elastic functional data. However, pointwise noise is always troublesome for the separation of phase and amplitude in functional data. The preprocessing required to create functions from raw data can also have an impact on the constructed tolerance bounds. The proposed method relies on the assumption that the input functions are at least absolutely continuous. However, the numerical procedures used in this work have better behavior for smoother functional data. The theory itself is stated for infinite-dimensional data objects and therefore the eventual results dependent on the density of sampling of the functions under study. In the sparse functional data setting, one first has to perform a smoothing or fitting step, which can affect the final results. In future work, we plan to quantify the robustness of the proposed procedures to these various preprocessing steps.

Since the proposed methods rely on fPCA to construct tolerance bounds, certain irregularities that fall outside of the space spanned by the leading fPCA basis functions are ignored. In practice, one should select a sufficient number of fPCA basis functions to capture all relevant directions of variability; if one believes that small scale variability is important, then even basis functions with small eigenvalues should be included in the analysis. On the other hand, if the small scale irregularities are precisely the ones that should be flagged, then leaving them out during the construction of the tolerance bounds is beneficial. While we have outlined some approaches to selecting an appropriate number of fPCA basis functions to define the tolerance bounds, our future work will focus on the effects of this choice in real applied settings.

We have additionally identified several other directions for future work. First, we will explore the influence of the weight C in the combined amplitude and phase fPCA model on the resulting tolerance bounds and tolerance factor. In particular, we want to assess the effects on simulated confidence values. Second, our method relies on tangent space approximations for the phase component, and recently Yu et al. [40] showed that in some cases the method of Principal Nested Spheres for dimension reduction provides more intuitive results. Third, in many applications, the functional data of interest may be more complex than the simple univariate functions considered in this work; some examples include shapes of curves, surfaces, and images. These more complicated data objects often exhibit additional sources of variability beyond amplitude and phase, which must be taken into account when computing tolerance bounds.

Acknowledgments

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 research was in part supported by the National Technical Nuclear Forensics Center (NTNFC) of the U.S. Department of Homeland Security (DHS). Sebastian Kurtek's work was partially supported by NSF grants DMS-1613054, CCF-1740761 and CCF-1839252, and by NIH grant R37 CA214955. The authors would like to thank Dr. Marc Welliver at Sandia National Laboratories for his technical support during this work. They would also like to acknowledge the Associate Editor and Reviewer for providing constructive comments that have significantly improved the content of this manuscript.

References

- M.L. Benson, M.N. Tran, and M.R. Hill, Phase 2b weld residual stress round robin: Mockup design and comparisons of measurement and simulation results, in Proceedings of the ASME Pressure Vessels and Piping Conference. 2015.
- [2] A. Bhattacharya, On a measure of divergence between two statistical populations defined by their probability distributions, Bulletin of Calcutta Mathematical Society 35 (1943), pp. 99–109.
- [3] R. Bousseljot, D. Kreiseler, and A. Schnabel, Nutzung der EKG-signaldatenbank CAR-DIODAT der PTB uber das internet, Biomedizinische Technik 40 (1995), pp. S317–S318.
- [4] A.C. Davison and D.V. Hinkley, Bootsrap Methods and their Application, Cambridge University Press, 1997.
- [5] A.C. Davison, D.V. Hinkley, and G.A. Young, *Recent developments in bootstrap methodology*, Statistical Science 18 (2003), pp. 141–157.
- [6] F. Ferraty and P. Vieu, Nonparametric Functional Data Analysis: Theory and Practice, Springer-Verlag New York, Inc., 2006.
- [7] A.L. Goldberger, L.A.N. Amaral, L. Glass, J.M. Hausdorff, P.C. Ivanov, R.G. Mark, J.E. Mietus, G.B. Moody, C. Peng, and H.E. Stanley, *Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals.*, Circulation 101 (2000), pp. e215–e220. Available at http://www.physionet.org.
- [8] M. Grasso, A. Menafoglio, B.M. Colosimo, and P. Secchi, Using curve-registration information for profile monitoring, Journal of Quality Technology 48 (2016), pp. 99–127.
- [9] G.J. Hahn and W.Q. Meeker, Statistical Intervals: A Guide for Practitioners, John Wiley & Sons, Inc., 2011.
- [10] A. Kneip and J.O. Ramsay, *Combining registration and fitting for functional models*, Journal of the American Statistical Association 103 (2008).
- [11] K. Krishnamoorthy and T. Matthew, Statistical Tolerance Regions: Theory, Applications, and Computation, Wiley: New York, 2009.
- [12] K. Krishnamoorthy and S. Mondal, Improved tolerance factors for multivariate normal distributions, Communications in Statistics - Simulation and Computation 35 (2006), pp. 461–478.
- [13] S. Kurtek, A geometric approach to pairwise Bayesian alignment of functional data using importance sampling, Electronic Journal of Statistics 11 (2017), pp. 502–531.
- [14] S. Kurtek and K. Bharath, Bayesian sensitivity analysis with Fisher-Rao metric, Biometrika 102 (2015), pp. 601–616.
- [15] S. Kurtek, A. Srivastava, and W. Wu, Signal Estimation Under Random Time-Warpings and Nonlinear Signal Alignment, in Proceedings of Neural Information Processing Systems (NIPS). 2011.
- [16] S. Kurtek, W. Wu, G.E. Christensen, and A. Srivastava, Segmentation, alignment and statistical analysis of biosignals with application to disease classification, Journal of Applied Statistics 40 (2013), pp. 1270–1288.
- [17] S. Lahiri, D. Robinson, and E. Klassen, Precise matching of PL curves in \mathbb{R}^n in the Square Root Velocity framework, Geometry, Imaging and Computing 2 (2015), pp. 133–186.
- [18] S. Lee and S. Jung, Combined analysis of amplitude and phase variations in functional data, arXiv:1603.01775 [stat.ME] (2017), pp. 1–21. Available at https:

//arxiv.org/abs/1603.01775.

- [19] J.R. Lewis, D. Brooks, and M.L. Benson, Methods for Uncertainity Quantification and Comparison of Weld Residual Stress Measurements and Predicitions, in Proceedings of the ASME Pressure Vessels and Piping Conference. 2017.
- [20] Y. Lu, R. Herbei, and S. Kurtek, Bayesian registration of functions with a Gaussian process prior, Journal of Computational and Graphical Statistics 26 (2017), pp. 894–904.
- [21] A.H. Mahmoudi, S. Hossain, C.E. Truman, D.J. Smith, and M.J. Pavier, A new procedure to measure near yield residual stresses using the deep hole drilling technique, Experimental Mechanics 49 (2008), pp. 595–604.
- [22] J. Marron, J. Ramsay, L. Sangalli, and A. Srivastava, Functional data analysis of amplitude and phase variation, Statistical Science 30 (2015), pp. 468–484.
- [23] M.B. Prime, R.J. Sebring, J.M. Edwards, D.J. Hughes, and P.J. Webster, Laser surface-contouring and spline data-smoothing for residual-stress measuremen, Experimental Mechanics 44 (2004), pp. 176–184.
- [24] J.O. Ramsay and B.W. Silverman, Functional Data Analysis, Springer, 2005.
- [25] L.N. Rathnayake and P.K. Choudhary, *Tolerance bands for functional data*, Biometrics 72 (2016), pp. 503–512.
- [26] D. Robinson, Functional analysis and partial matching in the square root velocity framework, Ph.D. diss., Florida State University, 2012.
- [27] A. Saha, K. Bharath, and S. Kurtek, *Geometric variational approach to Bayesian inference*, Journal of the American Association, To Appear (2019).
- [28] A. Saha and S. Kurtek, Sensitivity measures for Bayesian nonparametric density estimation models, Sankhya A, To Appear (2019).
- [29] A. Srivastava and I.H. Jermyn, Looking for shapes in two-dimensional, cluttered point clouds, IEEE Trans. Pattern Analysis and Machine Intelligence 31 (2009), pp. 1616–1629.
- [30] A. Srivastava, E. Klassen, S. Joshi, and I. Jermyn, Shape analysis of elastic curves in Euclidean spaces, IEEE Trans. Pattern Analysis and Machine Intelligence 33 (2011), pp. 1415–1428.
- [31] A. Srivastava and E.P. Klassen, Functional and Shape Data Analysis, Springer-Verlag, 2016.
- [32] A. Srivastava, W. Wu, S. Kurtek, E. Klassen, and J.S. Marron, Registration of functional data using Fisher-Rao metric, arXiv:1103.3817v2 [math.ST] (2011). Available at http://arxiv.org/abs/1103.3817v2.
- [33] C.B. Storlie, M.L. Fugate, D.M. Higdon, A.V. Huzurbazar, E.G. Francois, and D.C. McHugh, *Methods for characterizing and comparing populations of shock* wave curves, Technometrics 55 (2013), pp. 436–449.
- [34] Y. Sun and M.G. Genton, *Functional boxplots*, Journal of Computational and Graphical Statistics 20 (2011), pp. 316–334.
- [35] J.D. Tucker, Functional component analysis and regression using elastic methods, Ph.D. diss., Florida State University, 2014.
- [36] J.D. Tucker, Functional Statistical Process Control using Elastic Methods, in Proceedings of Joint Statistical Meetings. 2016.
- [37] J.D. Tucker, W. Wu, and A. Srivastava, Generative models for functional data using phase and amplitude separation, Computational Statistics and Data Analysis 61 (2013), pp. 50–66.
- [38] A. Veeraraghavan, A. Srivastava, A.K. Roy-Chowdhury, and R. Chellappa, *Rate-invariant recognition of humans and their activities*, IEEE Trans. on Image Processing 8 (2009), pp. 1326–1339.

- [39] W. Xie, S. Kurtek, K. Bharath, and Y. Sun, A geometric approach to visualization of variability in functional data, Journal of American Statistical Association 112 (2017), pp. 979–993.
- [40] Q. Yu, X. Lu, and J.S. Marron, Principal nested spheres for time-warped functional data analysis, Journal of Computational and Graphical Statistics 26 (2017), pp. 144–151.