



Bounding uncertainty in functional data: A case study

Caleb King , Nevin Martin & James Derek Tucker

To cite this article: Caleb King , Nevin Martin & James Derek Tucker (2021) Bounding uncertainty in functional data: A case study, Quality Engineering, 33:1, 178-188, DOI: [10.1080/08982112.2020.1851710](https://doi.org/10.1080/08982112.2020.1851710)

To link to this article: <https://doi.org/10.1080/08982112.2020.1851710>



Published online: 29 Jan 2021.



Submit your article to this journal [↗](#)



Article views: 17



View related articles [↗](#)



View Crossmark data [↗](#)

CASE REPORT



Bounding uncertainty in functional data: A case study

Caleb King^a , Nevin Martin^b, and James Derek Tucker^b 

^aJMP Division, SAS Institute, Cary, North Carolina; ^bStatistical Sciences, Sandia National Laboratories, Albuquerque, New Mexico

ABSTRACT

Functional data are fast becoming a preeminent source of information across a wide range of industries. A particularly challenging aspect of functional data is bounding uncertainty. In this unique case study, we present our attempts at creating bounding functions for selected applications at Sandia National Laboratories (SNL). The first attempt involved a simple extension of functional principal component analysis (fPCA) to incorporate covariates. Though this method was straightforward, the extension was plagued by poor coverage accuracy for the bounding curve. This led to a second attempt utilizing elastic methodology which yielded more accurate coverage at the cost of more complexity.

KEYWORDS

Functional data analysis; functional principal components; reliability specifications; tolerance bounds; uncertainty quantification

Introduction

Functional data have become more prevalent as a primary source of information in a variety of industries, thanks in part to the rise of improved sensor technology and sampling methods. There has also been a large amount of research to develop better methods for analyzing such data, such as functional principal component analysis (fPCA), functional alignment, and functional regression models (see Ramsay and Silverman (2005); Srivastava and Klassen (2016) for an introduction to the most common methods). However, there is still space in this area for additional research. One such area is the construction of bounding curves, such as confidence or tolerance bounds. The purpose for creating such bounding curves can be varied, but one specific purpose we discuss here is for use in defining and/or assessing product specifications.

Constructing bounding curves can be a difficult task since functional data do not exhibit random variation in the same way as univariate data. Initial attempts have involved converting the functional data to univariate data through summary metrics, such as a mean or maximum value. However, these data reduction techniques obfuscate the true nature of the functional variation and may lead to specifications that are either insufficient or grossly conservative. While there are times where a univariate metric may be appropriate (i.e., a maximum threshold for a voltage curve), in general a bounding curve that respects

the functional variation of the data, however complex it may be, will always prove superior to univariate metrics in capturing the different types of variation that may occur.

Recently, Storlie et al. (2013) 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 (2016) developed tolerance bounds for functional data using fPCA. Sun and Genton (2011) developed a boxplot display for functional data, which provides a nice visualization technique for a sample of functions and is also useful for detecting functional outliers. However, these methods do not take into account the amplitude and phase variability present in functional data, that is, they assume that the data (1) does not need to be aligned or (2) has already been aligned using some unrelated criterion. The first assumption is unrealistic in process control applications while the second approach results in sub-optimal solutions due to the disjointedness of the two sources of variation. A more systematic approach is to develop methods that build the alignment step into the statistical procedure.

To illustrate the difficulty of assessing functional uncertainty as well as highlight some of the novel approaches to overcome this difficulty, we present a selection of applications from Sandia National Laboratories (SNL) in which functional data analysis (FDA) techniques were used to help assess and bound

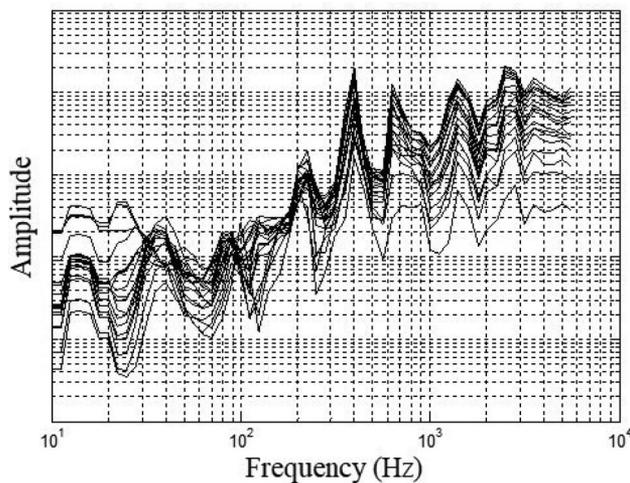


Figure 1. Example of a sampling of ASD curves.

uncertainty in functional data. Our case study will proceed as follows. First, we discuss the motivating functional data application that initiated the study. Next, we present our first attempt at providing a solution based on a simple extension of fPCA to include external factors combined with bootstrap techniques. We then discuss how a simulation study revealed some unforeseen drawbacks to an initial solution and what steps we took to try and understand the issue. Ultimately, a second attempt was made using a different approach involving elastic FDA techniques. Although this alternative method was not implemented in the motivating application, it has and continues to be implemented in similar settings across SNL, which we illustrate through two additional applications in the final sections.

By this point, it should be clear that this case study digresses from the traditional approach. Our goal in presenting the entire process, including the pivot from one method to another, is to provide insight into the entire problem-solving process rather than focus solely on a single successful solution. Success in research is almost never gained in a single step, but often through several smaller steps, each building upon the strengths of the previous attempt while also trying to address any setbacks along the way.

In addition, showing a method that did not yield optimal results can be informative to other researchers, and not just to avoid “reinventing the wheel”. Publication bias is a pre-iminent issue as of late and has given rise to a replication crisis in scientific research (Ioannidis 2005). Among the many proposed solutions, one approach is to encourage the publication of null or even negative results (Jena 2017). While the case study presented here is not necessarily a “null” result, the idea of being open to publishing

case studies in which the outcome was unexpected may be one worth considering. After all, as this very case study illustrates, there is no such thing as failure, only opportunity.

Motivating application

The initial motivating application involved the need for a more data-driven methodology in specifying environmental requirements for a munitions system. Over the course of their lifetimes, military munition systems are exposed to a variety of environments. Of particular interest are environments encountered during transportation, most notably shocks and vibrations. In this application, the focus was on mechanical vibrations experienced by munitions systems carried by aircraft during flight. Engineers wished to characterize the nature of these mechanical vibrations as a function of air speed and pressure, among other factors. From this, they could then compute a bounding curve to serve as the basis for creating product specifications for the munitions system.

The data used in the computation were collected using a carefully designed series of test flights at selected speeds and altitudes. A mock munitions system is carried on the wings of the aircraft, with accelerometers placed at key points across the test object. Once a flight test is completed, the data are downloaded and converted to Acceleration Spectral Densities (ASDs), which express the energy per unit time of the vibration as a function of frequency (Lalanne 2014). An example of the data is shown in Figure 1. Note that this data is different from more common forms of functional data in that the domain of the data is not time or location.

Common practice for the engineers was to create specifications by marking a series of lines that bound the maximum values observed at each frequency. These lines would be drawn either right at the maximum values or some specified value above them, with this value being conservatively chosen, yet not necessarily based on any observed variation in the data. One key weakness of this approach, in addition to its “ad-hoc” nature, was that it could not easily be applied to environments not seen during test flights. It was these extrapolated environments that were of primary interest to the engineers and stakeholders, so computing these bounds were vital. One alternative approach to achieve bounds at extrapolated environments was to compute the square root of the area under the curves, resulting in a metric called the root mean square acceleration (denoted by G_{rms} , see

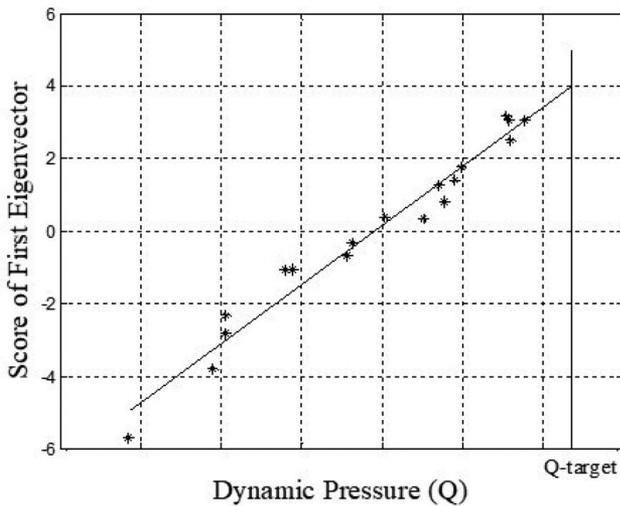


Figure 2. Example of the initial modeling approach. Here, the scores from the first eigenvector hold a linear relationship with dynamic pressure, Q . Q -target represents the specific environment for which tolerance bounds and ultimately product specifications need to be computed.

Lalanne (2014) for further details). This metric could then be used as a univariate response in a standard regression model, which could then be used to extrapolate to the environments of interest and any corresponding uncertainty accounted for through prediction intervals and the like. While this did resolve the extrapolation issue, the resulting specifications could only be in terms of G_{rms} , resulting in a loss of functional information.

First attempt: Functional PCA with bootstrapping

Realizing these limitations, the engineers sought out a methodology that could compute a general data-driven bounding curve on the ASDs that was a function of environmental factors. They had already been provided a method by a colleague who relied on the work of Rathnayake and Choudhary (2016), which relies on fPCA to break down the mean-centered functional data into its principal component vectors. The *loadings* or scores associated with each principal component vector were then used in a bootstrap procedure to estimate sampling variation. From the bootstrapped scores, a tolerance bound curve is created, designed to capture a given quantile p of the population of curves with confidence level $1 - \alpha$.

The method of Rathnayake and Choudhary (2016) only applies to functional data collected independent of experimental factors. That is, there is no built-in mechanism to directly handle functional variation due to one or more factors. Thus, a first attempt at

generating bounding curves was to extend the method of Rathnayake and Choudhary (2016) to allow for experimental factors (Thomas et al. 2016b). The procedure is outlined as follows.

The functional PCA method

Consider a collection of functional curves $\mathbf{Y}_i(t)$, $i = 1, 2, \dots, n$, where t represents a general input, and mean $\boldsymbol{\mu}(t) = 1/n \sum_{i=1}^n \mathbf{Y}_i(t)$. Then, using fPCA, we have

$$\mathbf{Y}_i(t) - \boldsymbol{\mu}(t) \approx \sum_{k=1}^K s_{ki} \phi_k(t), \quad [3.1]$$

where $\phi_k(t)$ is a principal component vector representing a principal mode of functional variation and s_{ki} is the i th score representing an individual scaling of the k th principal component vector for the i th functional response curve (Ramsay and Silverman 2005). The number of principal components K is specified by the user and can be determined by how many components are needed to capture a specified percentage of the variation in the data.

In examining the principal component scores, it typically happened that the first and possibly second principal components showed clear trends with one or more factors, as illustrated in Figure 2 with the dynamic pressure (Q), a quantity related to both aircraft speed and altitude. To account for the presence of experimental factors, we decided to fit a regression model directly to these scores.

The choice to use regression models naturally led us to consider using a parametric bootstrap, wherein the scores were assumed to have standard normal distributions (once the regressors were accounted for). This distributional assumption was based on the property that the scores have mean 0 and variance equal to the k th eigenvalue of the principal component decomposition of the functional covariance matrix $E\{(\mathbf{Y}(t) - \boldsymbol{\mu}(t))(\mathbf{Y}(t) - \boldsymbol{\mu}(t))'\}$. Furthermore, the columns of the matrix $\mathbf{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_K\}$ are orthogonal, so a multivariate normal distribution with diagonal covariance matrix was used. The goal, as with the original method, was to create a tolerance bound curve at some specified level of the environmental factor(s) on a given quantile p with confidence level $1 - \alpha$. The detailed steps of the parametric bootstrap algorithm are summarized below.

1. If the k th principal component has a linear model associated with it, compute $\hat{\mu}_k = \mathbf{x}_0^T \hat{\boldsymbol{\beta}}$. Otherwise, compute $\hat{\mu}_k = \frac{1}{n} \sum_{i=1}^n s_{ki}$.

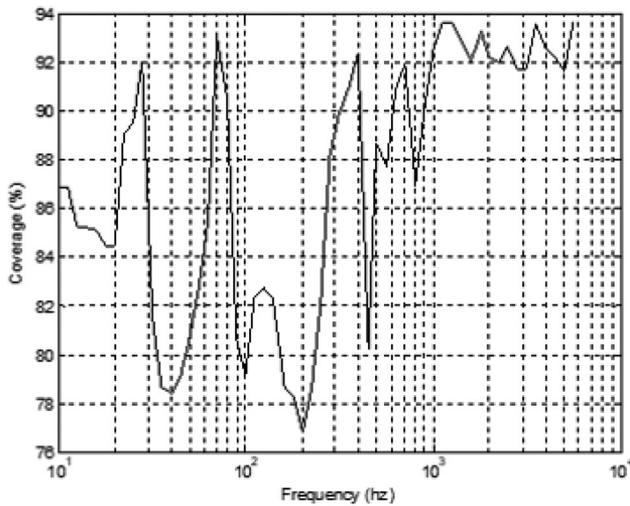


Figure 3. Example of simulation results. The intended coverage on the 95th percentile is 90 percent. However, as can be seen here, the actual coverage varies significantly.

2. For $b = 1, 2, \dots, B$,

(a) Simulate $\hat{\sigma}_{kb} \sim \hat{\sigma}_k \sqrt{\frac{df_k}{\chi_{df_k}^2}}$, $k = 1, 2, \dots, p$, where

$\hat{\sigma}_k$ is the estimated standard deviation of the scores for the k th principal component, df_k are the appropriate degrees of freedom, and $\chi_{df_k}^2$ is a Chi-squared random variable with df_k degrees of freedom.

(b) If the k th principal component has a linear model associated with it, compute $\tilde{\sigma}_\mu^2 = \tilde{\sigma}_{kb}^2 \mathbf{x}_0^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_0$, where \mathbf{X} is the model matrix and \mathbf{x}_0 is the target level. Otherwise, compute $\tilde{\sigma}_\mu^2 = \tilde{\sigma}_{kb}^2 / n$.

(c) Simulate $\tilde{\mu}_{kb} \sim \text{Normal}(\hat{\mu}_{kb}, \tilde{\sigma}_\mu^2)$.

(d) Compute $\bar{\mathbf{Y}}_b(t) = \bar{\mathbf{Y}}(t) + \sum_{k=1}^K \tilde{\mu}_{kb} \phi_k(t)$ and $\mathbf{V}_b(t) = \sqrt{\sum_{k=1}^K \phi_k(t)^2 \tilde{\sigma}_{kb}^2}$, where $\bar{\mathbf{Y}}$ is the pointwise mean of the original data and $\phi_k(t)$ is the k th eigenfunction.

(e) Compute $\mathbf{P}_b(t) = \bar{\mathbf{Y}}_b(t) + z_p \mathbf{V}_b(t)$, where z_p is the p th percentile of the standard normal distribution.

3. The pointwise $(1 - \alpha)$ -level upper tolerance bound is given by the $(1 - \alpha)$ -percentile of the P_{bj} , $j = 1, 2, \dots, J$.

Simulation study and results

To assess the accuracy of the tolerance bound coverage, a simulation study was conducted. Functional data were simulated by first simulating scores from a multivariate normal distribution, using the estimated model means from the ASD data as the true means,

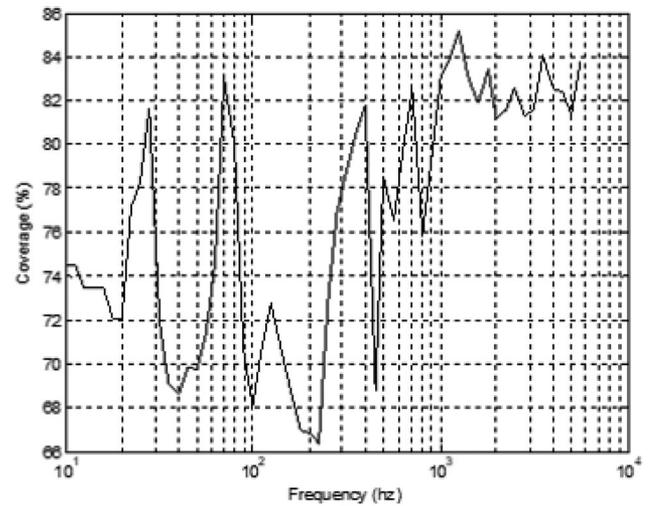


Figure 4. Example of simulation results for lower coverage tolerance bound. The intended coverage on the 95th percentile is 75 percent. Note that the pointwise bounds are more stable here.

and then combining them with the eigenfunctions from the original ASD data. The simulation study was divided into two parts: one which simulated the same number of curves as the ASD data and one which used 10 replicates of each of the original curves. The purpose of the latter simulation was to imitate an increased sample size consisting of replicate data points.

The range of tolerance bounds considered were 99-90, 99-50, 99-75, 95-90, and 95-50, where the first number represents the quantile of interest and the second number represents the confidence level. We then computed the actual coverage by determining the percentage of simulated boundary curves that covered the known quantile at each point in the domain. That is, our computation was of pointwise coverage. It was at this point that we discovered a surprising result. The simulation results seemed to indicate that the actual coverage tended to fall below the nominal coverage and could actually vary widely over the domain of the functional data. For example, in one simulation illustrated in Figure 3, the mean true pointwise coverage for a 99-90 tolerance bound (bound on the 99th percentile with 90 percent confidence) was actually closer to 85 percent, fluctuating between 77 percent and 96 percent across the function domain. In a pathological twist, the true coverage only seemed to stabilize with lower confidence levels, as illustrated in Figure 4. This behavior was also seen with the replicate data.

Upon seeing these results, we found ourselves in a bit of a quandary. Our method was certainly an improvement over those currently being used by the

engineers. Yet, at the same time, one would not be able to accurately claim that the resulting bounding curve covered a given quantile with $1 - \alpha$ confidence, which could be problematic since the purpose of the curve was to bound the population with a known level of uncertainty. Furthermore, there seemed to be no clear explanation for this strange behavior as the curves did not seem to exhibit significant phase variation, ruling that out in our minds as a possible cause. An attempt was made to try a non-parametric bootstrap approach, but this only resulted in nonphysical bounds and did not resolve the coverage inaccuracy.

We presented this work at the 2016 Joint Statistical Meetings (Thomas et al. 2016a). When looking for solutions, it's often helpful to seek outside assistance and so we decided to reach out to the statistics community for any insight into ways we could improve the coverage accuracy and consistency of our method. At the presentation, a colleague suggested that, by applying the regression model directly to the scores, we had in fact eliminated those particular components as a source of variation. That is, while the bootstrap did account for the model uncertainty, it did not adequately account for the variation in the eigenvector loadings. This colleague then suggested the elastic method as an alternative, which we now discuss in the following section.

Second attempt: The elastic method

The basis for the this tolerance bound methodology stems from the fact that most functional data have two types of variability — amplitude (vertical) and phase (horizontal). Most standard FDA approaches do not account for both types of variability in a principled fashion and so an *elastic* method is needed (Srivastava and Klassen 2016). Note that, upon closer inspection, there is in fact some small phase variability in the ASD data presented in Figure 1. Specifically, the “peaks” and “valleys” of the functions are not perfectly aligned across frequency (i.e., phase variability), nearly drowned out by the variation across amplitude (i.e., amplitude variability). The previous attempt at applying fPCA on the original data only considered the amplitude variability. The elastic method uses a more natural approach to characterizing functional data by using fPCA on a joint representation of the phase and amplitude spaces.

Elastic functional data approach

The elastic methodology begins by characterizing the phase and amplitude variabilities of the data. This is

an important first step as this characterization is used to fit an fPCA model that generates new functions that represent the phase and amplitude variability in the original data. These generated functions are then used in a bootstrap procedure to estimate tolerance bounds on the phase and amplitude spaces separately, resulting in bounds that appropriately handle the variabilities seen in functional data. The details of the procedure are provided below with more detail available to the reader in Tucker et al. (2020). Methods used in the results are available in R package *fdasrvf* available on CRAN.¹

To characterize the phase and amplitude variabilities in the original data, the functions are first aligned along the x-axis using *warping functions* (Srivastava and Klassen 2016). The resulting aligned functions are used to characterize amplitude variability, while the warping functions are used to characterize the phase variability. These warping functions, Γ , are defined as a set of functions (γ) on the interval $[0, 1]$ that have the following properties:

- γ has one-to-one correspondence from $[0, 1]$ into $[0, 1]$,
- $\gamma(0) = 0$ and $\gamma(1) = 1$, and
- both γ and γ^{-1} are differentiable.

For any function $f \in F$, where F is the set of all real-valued functions on $[0, 1]$, f is warped by γ through composition: $f \circ \gamma = f(\gamma(t))$. The goal is to align any two functions f_1 and f_2 using the following amplitude distance metric:

$$d_a(f_1, f_2) = \inf_{\gamma \in \Gamma} \left\| q_1 - (q_2 \circ \gamma) \sqrt{\dot{\gamma}} \right\|, \quad [4.1]$$

where

$$q_i(t) = \text{sign}(\dot{f}_i(t)) \sqrt{|\dot{f}_i(t)|} \quad [4.2]$$

is known as the square-root slope function (SRSF) and \dot{f} is the derivative of f with respect to t . By using the SRSF for alignment (instead of the original functions), the amplitude distance is a proper distance; that is, this distance metric is non-negative, symmetric, and satisfies the property $d_a(f_1, f_3) \leq d_a(f_1, f_2) + d_a(f_2, f_3)$.² Because of this, it is now possible to define statistics on the functional space (e.g., mean and variance) that naturally handle the known variabilities. Once the SRSFs are aligned, they can then be mapped back to F to obtain the

¹MATLAB code is also available at https://github.com/jdtuck/fdaSRSF_MATLAB.

²Mathematical proof provided in Srivastava and Klassen (2016).

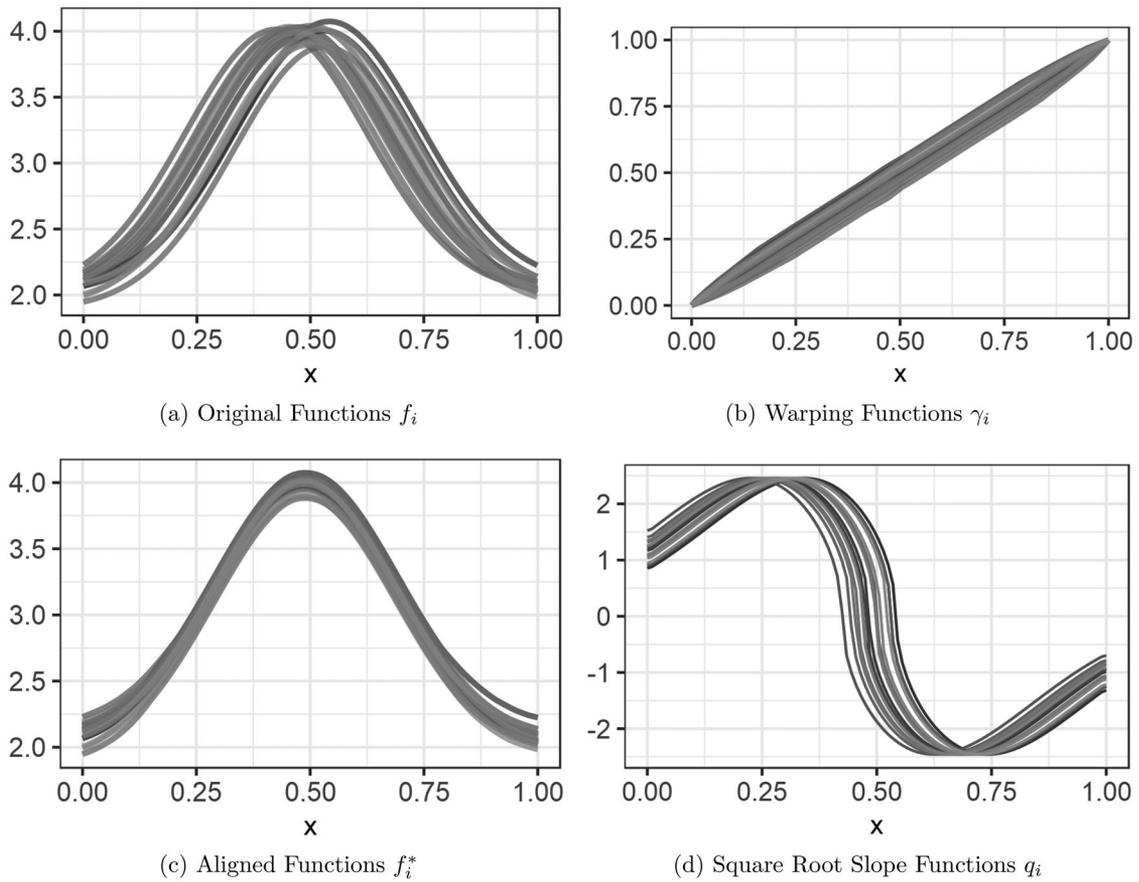


Figure 5. Notional example of warping functional data.

aligned original functions. Additional details about SRSFs are given in Marron et al. (2015); Srivastava et al. (2011); Tucker, Wu, and Srivastava (2013) and Srivastava and Klassen (2016).

After aligning the functions, the phase variability can be characterized using the same functions γ that were used to align the data. The space of Γ is a non-linear manifold which creates some difficulty in defining a distance between two warping functions γ_1 and $\gamma_2 \in \Gamma$. In order to define a distance, we use the square-root of the derivative of γ :

$$\psi = \sqrt{\dot{\gamma}}. \quad [4.3]$$

This representation allows the set of all $\psi \in \Psi$ to be the positive orthant of a Hilbert sphere (\mathbb{S}_∞^+). Although this is still an infinite dimensional manifold, it now has a known geometry which we can utilize. Therefore, the phase distance between γ_1 and γ_2 can then be calculated as the arc-length between ψ_1 and ψ_2 on the Hilbert sphere:

$$d_p(\gamma_1, \gamma_2) = d_\psi(\psi_1, \psi_2) \equiv \cos^{-1} \left(\int_0^1 \psi_1(t) \psi_2(t) dt \right). \quad [4.4]$$

To be able to compute standard statistics, the geometry of Ψ is then further simplified by analyzing the warping functions on a tangent space defined as

$$T_\psi(\Psi) = \left\{ v \in \mathbb{L}^2 \mid \int_0^1 v(t) \psi(t) dt = 0 \right\}. \quad [4.5]$$

Here, v is a tangent space vector and allows us to perform fPCA in a linear space and then map back to Γ . For details see Tucker et al. (2020).

Figure 5 shows an example of how phase and amplitude variabilities are decomposed using simulated data. Figure 5(a) gives the original data y_i , where amplitude and phase variation are both present, noted by the fact that the “peaks” of the curves do not align along either the x or y axes. Figure 5(b,c) give the warping γ_i and aligned f_i^* functions, respectively, while Figure 5(d) shows the SRSFs q_i . The f_i^* functions have been aligned such that the only variability remaining is variability along the y-axis (amplitude). The corresponding γ_i provide a measure of the phase variability in the original data.

Once the variabilities have been characterized, an fPCA model can then be fit on a joint representation of the aligned and warping functions. This idea

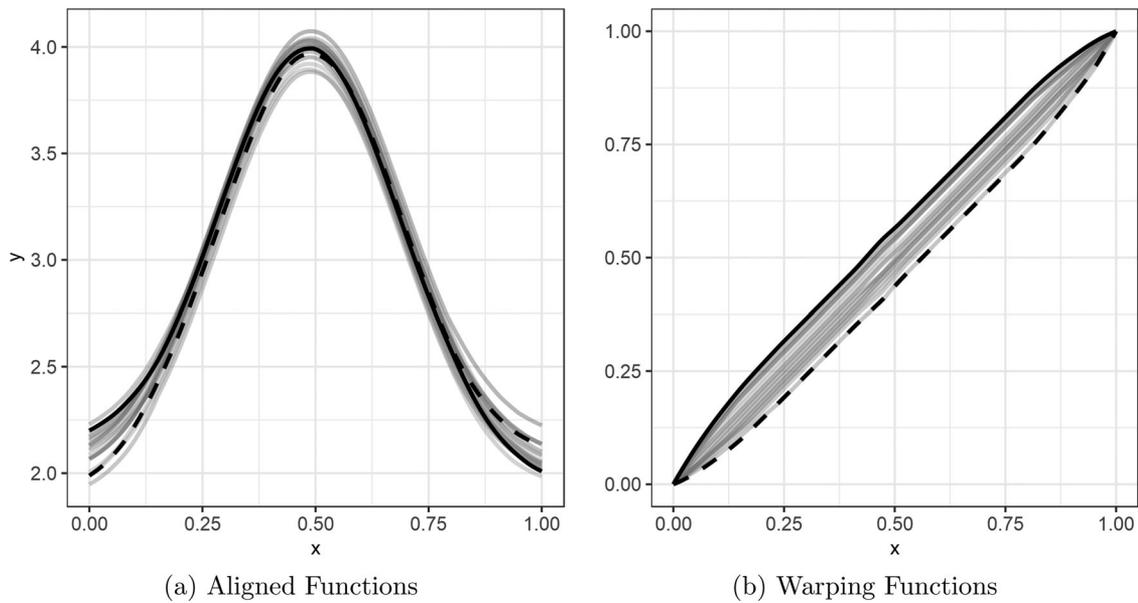


Figure 6. Tolerance bounds on aligned (left) and warping (right) functions.

expands on the work of Lee and Jung (2017) by using a concatenated function (g^C) on the extended domain $[0, 2]$ for some $C > 0$, defined as

$$g^C(t) = \begin{cases} q^*(t), & t \in [0, 1) \\ Cv(t - 1), & t \in [1, 2], \end{cases} \tag{4.6}$$

where C is a scaling factor and q^* is the SRSF of the aligned function f^* and v is the corresponding tangent vector computed from γ . Then, for a sample of amplitude-phase functions $\{g_1^C, \dots, g_n^C\}$, their sample mean can be defined as $\hat{\mu}_g^C = [\hat{\mu}_{q^*} \quad \hat{\mu}_v^C]$ with a sample covariance matrix of

$$K_g^C = \frac{1}{n-1} \sum_{i=1}^n (g_i^C - \hat{\mu}_g^C)(g_i^C - \hat{\mu}_g^C)^T \in \mathbb{R}^{(2T) \times (2T)}. \tag{4.7}$$

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

After the fPCA coefficients have been estimated, a probability model can be applied to the coefficients, allowing for a distribution on F from which to sample new functions. Let $c = (c_1, \dots, c_K)$ be the K dominant principal coefficients. These c can then be modeled using a multivariate Gaussian probability distribution with zero mean and covariance Σ , that is, $c \sim \mathcal{N}_K(0, \Sigma)$, where Σ is a $K \times K$ diagonal matrix with diagonal elements estimated using the eigenvalues of the sample covariance matrix, $\hat{\sigma}_1^C, \dots, \hat{\sigma}_K^C$. The superscript C shows the dependence on the scaling factor.

The fPCA model is then used in a bootstrapping procedure to estimate tolerance bounds. For the purposes of illustration, we will describe the construction of a symmetric tolerance interval intended to capture $100p\%$ of the population. An extension to one-sided tolerance bounds is straightforward. The fPCA model is repeatedly sampled using the above described coefficient model and the $(1-p)/2$ and $1-(1-p)/2$ quantiles of the set of random SRSF-based amplitudes and warping functions (denoted as $(q_{(1-p)/2}^*, q_{1-(1-p)/2}^*)$ and $(\gamma_{(1-p)/2}, \gamma_{1-(1-p)/2})$, respectively) are estimated. The quantiles are estimated using an extension of the boxplot procedure of Xie et al. (2017) This process is repeated S times such that the $\alpha/2$ and $(1-\alpha/2)$ quantiles of $(q_{(1-p)/2}^*, q_{1-(1-p)/2}^*)$ and $(\gamma_{(1-p)/2}, \gamma_{1-(1-p)/2})$ can be estimated. These quantiles form $(1-\alpha)100\%$ bootstrap tolerance intervals with $100p\%$ coverage.

Figure 6 shows the resulting tolerance bounds on aligned functions (left) and the warping functions (right) of the simulated data. While the bounds on the warping functions are intuitive, the bounds on the aligned functions may be unexpected as they overlap near $x=0.60$. This is due to the geometric approach in which the bounds are estimated. For the functions in Figure 6, the tolerance region indicates that the variation is not solely in amplitude (as is typically expected) but rather as more of a “twisting” variation. The aligned functions represent the overall “shape” of the original data and therefore their tolerance bounds capture a global “shape” tolerance region of the aligned function space, allowing for the possibility of the bounds crossing. Further discussion

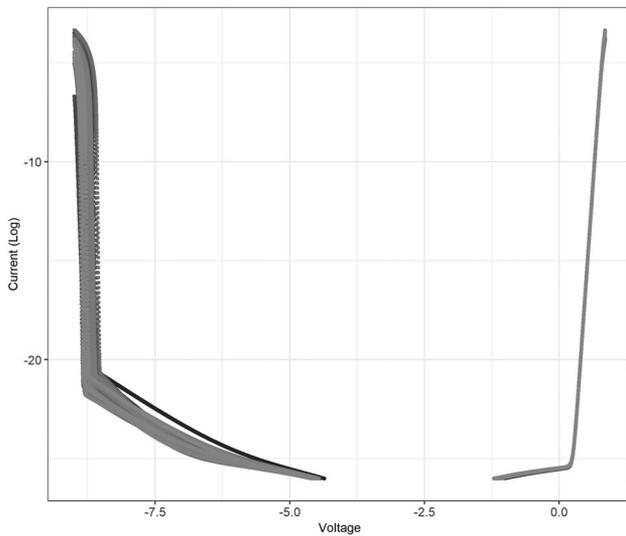


Figure 7. Zener diode electrical data from 44 devices.

of this idea is provided in Tucker et al. (2020) and Xie et al. (2017).

Tucker et al. (2020) describe a coverage study that was performed using the new method that showed the tolerance bounds were close to their true coverage levels, albeit somewhat conservative. In particular, for a tolerance bound on the upper 90 percent quantile with 95 percent confidence, the approach achieved actual confidence levels of 97.6 percent and 98.4 percent for the aligned and warping functions, respectively. The confidence coverage was greatly improved compared to the original method proposed in Section 3 as it accounted for variability in both the amplitude and phase directions. Note that while the elastic functional data approach does not currently account for external factors, the method could be extended using the approach similar to Cardot (2007), by adding factors into the elastic fPCA representation.

Further extensions

By the time a tolerance bound procedure had been developed using the elastic methodology, the engineers on the ASD project had decided to revert back to their original procedures. It seemed another drawback of our first attempt was that it had led in some of their applications to bounds that were nonphysical. Though we were still able to publish the new method for elastic functional tolerance bounds (Tucker et al. 2020), a different application had to be used. More details on that application can be found there, but for the purposes of this case study we present another project which led us (“us” now referring to Derek and Nevin; the other researchers had left Sandia by this

point) to an important extension of the elastic tolerance bound methodology.

Engineers at the labs were working on compact model calibration for electrical device design. Their data were collected by measuring Zener diode devices, capturing log current vs. voltage for approximately 44 nominally identical diodes, as shown in Figure 7. Differences in the log current-voltage behavior across diodes were due to manufacturing process variation and it was of interest to bound a percentage of the population of log current-voltage functions with a given level of confidence. These bounds would then be used in the calibration process of their compact models; that is, models for individual devices within an electrical circuit design.

As these bounds were to be used in compact model calibration, they needed to be defined in the original data space (as opposed to the warping and aligned spaces separately) while still accounting for both types of variability. It was determined that the elastic FDA method could be extended to produce bounds in the original data space by composing four pairwise combinations of the bounds, as follows:

1. Lower amplitude bound with the inverse of the lower phase bound: $f_{\alpha/2}^* \circ \gamma_{\alpha/2}^{-1}$.
2. Lower amplitude bound with the inverse of the upper phase bound: $f_{\alpha/2}^* \circ \gamma_{1-\alpha/2}^{-1}$.
3. Upper amplitude bound with the inverse of the lower phase bound: $f_{1-\alpha/2}^* \circ \gamma_{\alpha/2}^{-1}$.
4. Upper amplitude bound with the inverse of the upper phase bound: $f_{1-\alpha/2}^* \circ \gamma_{1-\alpha/2}^{-1}$.

While estimating four tolerance bounds may not be intuitive, each bound represents a combination of phase and amplitude variability that is important to capture. Two of the bounds bound lower phase variability, while the other two bound upper phase variability. Similarly, two bound lower amplitude variability, while the other two bound upper amplitude variability. Each combination of upper and lower bound is important to consider as it corresponds to some bounding behavior that may be relevant to capture in compact modeling.

An fPCA model was fit to the Zener diode data and 500 bootstrap samples were generated to calculate tolerance bounds on the aligned and warping functions, as shown in Figure 8(a,b). These bounds were constructed such that they represent where we would expect 90 percent ($p = 0.90$) of the devices to fall with 95 percent confidence ($\alpha = 0.05$). The bounds were

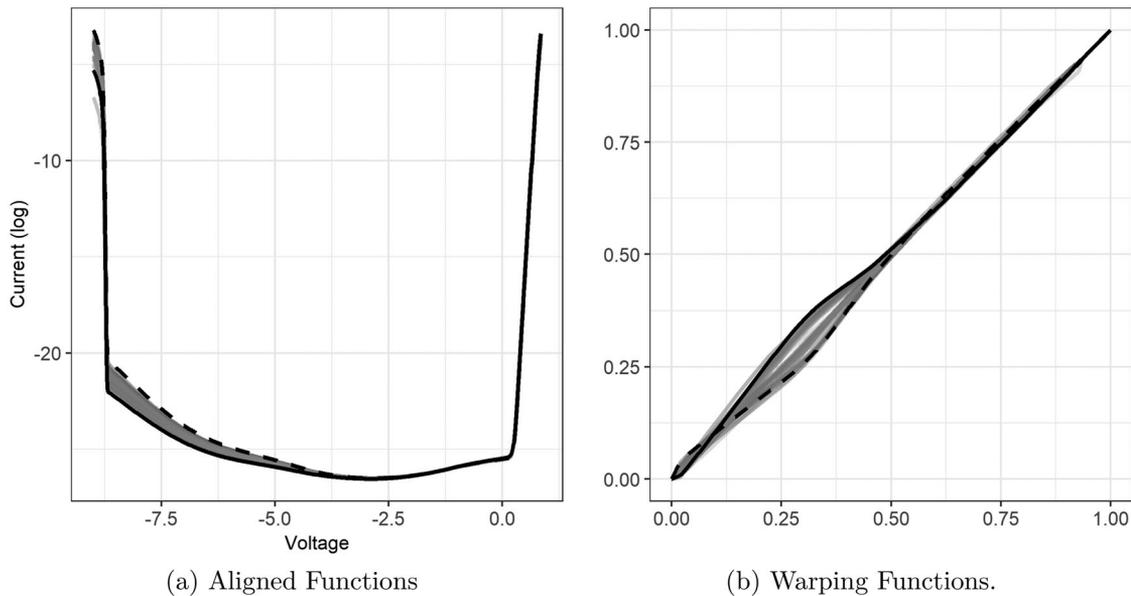


Figure 8. Tolerance bounds on aligned (left) and warping (right) Zener diode functions.

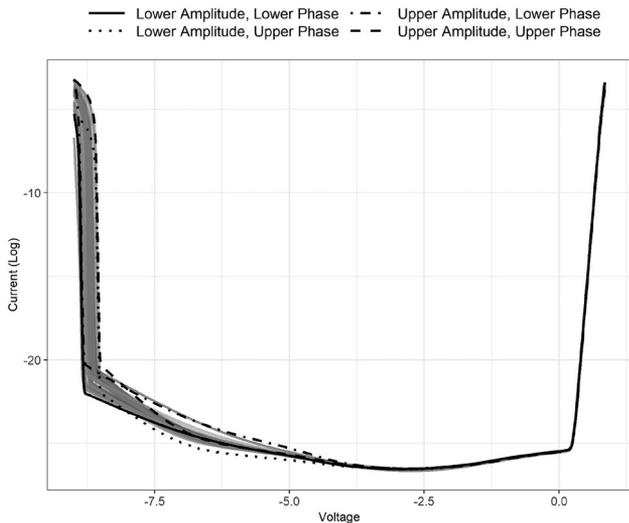


Figure 9. Tolerance bounds on original Zener diode data.

then transformed back to the original data space, as shown in Figure 9. Figure 10 provides a closer view of the current-voltage curve that is critical to estimate accurately for compact model calibration. The engineers were pleased with the resulting four bounds, as they represent combinations of high and low current and voltage that are important to characterize. They also recognized that if phase variation had been ignored, the bounds that would have been produced would result in a nonphysical representation of Zener diode behavior (i.e., they would lead to physically-impossible parameters in the model calibration).

Additionally, a coverage study was conducted to assess whether these bounds are performing as intended. This study involved sampling a new set of

44 diode devices from the original data with replacement (i.e., the same device can be sampled more than once). The warping functions and aligned SRSFs were then calculated for each device. For amplitude, the upper 90 percent quantile of the SRSFs was calculated and compared to the tolerance bound to see if the entire SRSF fell within that bound. This was repeated 500 times to estimate the coverage level of this bound. This same process was performed for the phase tolerance bounds using the warping functions. The actual amplitude and phase confidence levels were estimated to be 97.7 percent and 99.3 percent, respectively. While somewhat conservative, these values are relatively close to the expected value of 95 percent confidence, and are consistent with the results seen in Tucker et al. (2020).

Further details on the application of this method to Zener diode data are described in Reza et al. (2020). This work illustrated a successful application and extension of the new tolerance bound approach. Future work for the engineers on this project involves incorporating this approach into their semi-automated model calibration workflow.

Lasting impacts

In addition to Zener diode data, the elastic method has also been applied to weld residual stress data for nuclear power plant systems and many other application areas at SNL. As of the writing of this case study, there has been renewed interest in applying the elastic method to new ASD applications. Overall, there has been a sharp increase in awareness of the importance

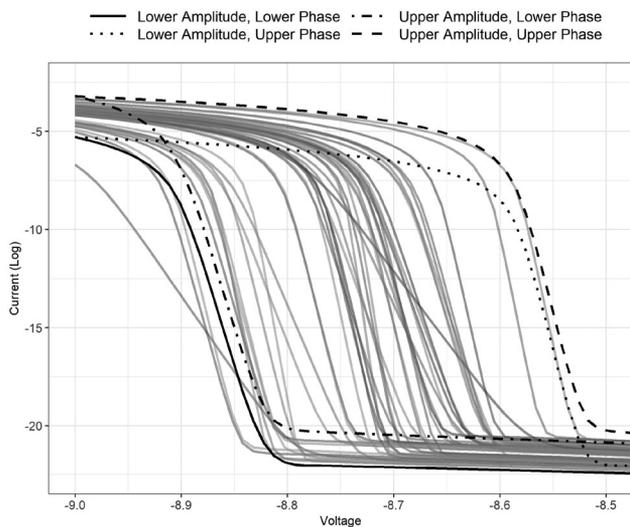


Figure 10. Zoomed-in tolerance bounds on the reverse breakdown region of the log current-voltage curve.

of FDA in the Statistical Sciences Department at SNL which has led to continued research to extend the methodology. Additionally, as more engineers are becoming familiar with the concept of FDA, they are recognizing the added benefit that can come from analyzing functions in their entirety.

While the initial fPCA/bootstrap approach was ultimately replaced by the elastic method, this should not imply it should be abandoned entirely. With the exception of the coverage anomaly, it remains a sound, practical approach to modeling previously aligned functional data. In this case study, the choice was made to switch to the elastic method, but an alternative choice could easily have been made to revise the initial approach to more accurately estimate the true variation in the functional data once the regression models have been estimated. As it stands, this remains an open research question.

The purpose of this case study has been to openly portray the twists and turns of the collaboration and research process. Rarely is the path to a successful solution a straight one, and there is no guarantee that it will even end. The best we can do is try and improve on what has come before, however windy that path may be.

Acknowledgments

The authors would like to thank Thomas Buchheit, Shahed Reza, Biliana Paskaleva and Andrew Sandoval for their contribution to the Zener diode work. The authors would also like to thank two anonymous reviewers for their helpful comments in improving the article.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology &

Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DENA0003525.

This article describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

About the authors

Caleb King is a Research Statistician Developer for the Design of Experiments platform in JMP Software. Prior to working at JMP, he worked for 3 years as a Statistical Scientist at Sandia National Laboratories. He received his BS in Mathematics and Statistics from Calvin University in 2009 and his MS and PhD in Statistics in 2011 and 2015 from Virginia Tech. His research interests include design of experiments, reliability, accelerated testing, and small-sample theory.

Nevin Martin is a former Member of the Technical Staff in the Statistical Sciences department. She received a BS degree in Finance from the University of Arizona and a MS degree in Statistics from the University of New Mexico in 2010 and 2016, respectively. Her research interests include design of experiment and uncertainty quantification (UQ) methods for computer simulations, functional data analysis, statistical computing and machine learning.

James Derek Tucker is a Principal Member of the Technical Staff at Sandia National Laboratories. He received his BS in Electrical Engineering Cum Laude and MS in Electrical Engineering from Colorado State University in 2007 and 2009, respectively. In 2014 he received the PhD. degree in Statistics from Florida State University in Tallahassee, FL under the co-advisement of Dr. Anuj Srivastava and Dr. Wei Wu. He currently is leading research projects in the area of satellite image registration and point processes modeling for monitory applications. His research is focused on pattern theoretic approaches to problems in image analysis, computer vision, signal processing, and functional data analysis. In 2017, he received the Director of National Intelligence Team Award for his contributions to the Signal Location in Complex Environments (SLICE) team.

ORCID

Caleb King  <http://orcid.org/0000-0003-3060-6604>

James Derek Tucker  <http://orcid.org/0000-0001-8844-2169>

References

- Cardot, H. 2007. Conditional functional principal components analysis. *Scandinavian Journal of Statistics* 34 (2): 317–35. doi:10.1111/j.1467-9469.2006.00521.x.
- Ioannidis, J. 2005. Why most published research findings are false. *PLoS Medicine* 2 (8):e124. doi:10.1371/journal.pmed.0020124.
- Jena, A. B. 2017. 'Null' research findings aren't empty of meaning. Let's publish them. <https://www.statnews.com/>

- [2017/11/10/null-research-findings/](#) (accessed December 10, 2019).
- Lalanne, C. 2014. *Random vibration*. Great Britain and United States: John Wiley & Sons, Ltd.
- Lee, S., and S. Jung. 2017. Combined analysis of amplitude and phase variations in functional data. 1–21. arXiv: 1603.01775 [stat.ME].
- Marron, J., J. Ramsay, L. Sangalli, and A. Srivastava. 2015. Functional data analysis of amplitude and phase variation. *Statistical Science* 30 (4):468–84. doi:10.1214/15-STS524.
- Ramsay, J. O., and B. W. Silverman. 2005. *Functional data analysis*, 2nd ed. New York: Springer.
- Rathnayake, L. N., and P. K. Choudhary. 2016. Tolerance bands for functional data. *Biometrics* 72 (2):503–12. doi:10.1111/biom.12434.
- Reza, S., N. Martin, T. Buchheit, and J. D. Tucker. 2020. Tolerance bound calculation for compact model calibration using functional data analysis. Proceedings of IEEE Electron Devices Technology and Manufacturing Conference.
- Srivastava, A., E. Klassen, S. Joshi, and I. Jermyn. 2011. Shape analysis of elastic curves in Euclidean spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33 (7):1415–28. doi:10.1109/TPAMI.2010.184.
- Srivastava, A., and E. P. Klassen. 2016. *Functional and shape data analysis*. New York: Springer-Verlag.
- Storlie, C. B., M. L. Fugate, D. M. Higdon, A. V. Huzurbazar, E. G. Francois, and D. C. McHugh. 2013. Methods for characterizing and comparing populations of shock wave curves. *Technometrics* 55 (4):436–49. doi:10.1080/00401706.2013.805662.
- Sun, Y., and M. G. Genton. 2011. Functional boxplots. *Journal of Computational and Graphical Statistics* 20 (2): 316–34. doi:10.1198/jcgs.2011.09224.
- Thomas, E., C. King, J. Cap, and A. Montoya. 2016a. Construction of tolerance bounds for a multivariate response associated with a covariate: A case study. Paper presented at the Joint Statistical Meetings, Chicago.
- Thomas, E. V., C. King, J. Cap, and A. Montoya. 2016b. Construction of tolerance bounds for a multivariate response associated with a covariate: A case study. Technical Report SAND2016-0859, Albuquerque, NM: Sandia National Laboratories.
- Tucker, J. D., J. R. Lewis, C. King, and S. Kurtek. 2020. A geometric approach for computing tolerance bounds for elastic functional data. *Journal of Applied Statistics* 47 (3): 481–505. doi:10.1080/02664763.2019.1645818.
- Tucker, J. D., W. Wu, and A. Srivastava. 2013. Generative models for functional data using phase and amplitude separation. *Computational Statistics & Data Analysis* 61: 50–66. doi:10.1016/j.csda.2012.12.001.
- Xie, W., S. Kurtek, K. Bharath, and Y. Sun. 2017. A geometric approach to visualization of variability in functional data. *Journal of the American Statistical Association* 112 (519):979–93. doi:10.1080/01621459.2016.1256813.