Published in final edited form as:

Biometrics. 2013 March ; 69(1): 41–51. doi:10.1111/j.1541-0420.2012.01808.x.

Corrected Confidence Bands for Functional Data Using Principal Components

J. Goldsmith^{1,*}, S. Greven², and C. Crainiceanu³

¹Department of Biostatistics, Columbia University, New York, New York 10032, U.S.A.

²Department of Statistics, Ludwig-Maximilians-Universität München, 80539 München, Germany

³Department of Biostatistics, Johns Hopkins University, Baltimore, Maryland 21205, U.S.A.

Abstract

Functional principal components (FPC) analysis is widely used to decompose and express functional observations. Curve estimates implicitly condition on basis functions and other quantities derived from FPC decompositions; however these objects are unknown in practice. In this article, we propose a method for obtaining correct curve estimates by accounting for uncertainty in FPC decompositions. Additionally, pointwise and simultaneous confidence intervals that account for both model- and decomposition-based variability are constructed. Standard mixed model representations of functional expansions are used to construct curve estimates and variances conditional on a specific decomposition. Iterated expectation and variance formulas combine model-based conditional estimates across the distribution of decompositions. A bootstrap procedure is implemented to understand the uncertainty in principal component decomposition quantities. Our method compares favorably to competing approaches in simulation studies that include both densely and sparsely observed functions. We apply our method to sparse observations of CD4 cell counts and to dense white-matter tract profiles. Code for the analyses and simulations is publicly available, and our method is implemented in the R package *refund* on CRAN.

Keywords

Bootstrap; Functional principal components analysis; Iterated expectation and variance; Simultaneous bands

1. Introduction

Functional principal components (FPC) analysis is a key tool in functional data analysis that has applications in dimension reduction, estimation of curves in the presence of measurement error, and the expression of traditional longitudinal data as sparsely observed functional data. Despite the ubiquity of FPC analysis (hereafter FPCA) in the functional data literature, principal components expansions are poorly understood in terms of their sources of variability. In this article, we develop a method for the estimation of curves that accounts for uncertainty in the model-based estimation of curve-specific scores for a particular FPC

decomposition as well as uncertainty in the FPC decomposition itself. Further, we construct confidence intervals for the curves that include both model- and decomposition-based sources of variability.

Using FPCA, individual curves are expressed as a linear combination of population-level basis functions and curve-specific scores. These scores can be predicted as the best linear unbiased predictors (BLUPs) in a mixed model framework for a particular FPC decomposition, which includes the basis functions and score variances. Thus FPCA allows a parsimonious, nonparametric representation of functional data. Additionally, mixed model results allow the construction of confidence intervals for estimated curves, again for a particular decomposition. However, in practice the decomposition is based on the estimated covariance structure of the population of curves and is subject to variability. This added variability can be substantial in many settings, including when a small number of curves is observed and when the curves are observed sparsely at the subject level, yet is often overlooked. Accurate estimation and inference for functional expansions is important in many scientific settings, such as understanding CD4 cell count trajectories in HIV positive individuals based on infrequent measurements or quantifying intracranial white-matter structure using diffusion tensor imaging (DTI).

Our proposed method explicitly conditions on a particular FPC decomposition when predicting curve-specific scores and estimating the associated variability. Next, we use iterated expectation and variance formulas to combine uncertainty in the conditional estimates across the distribution of FPC decomposition objects (the basis functions, mean function, measurement error and score variances, and the truncation lag). We use the bootstrap to derive an empirical distribution of decomposition objects by resampling curves with replacement and, for each sample, constructing an FPC decomposition. Thus, our method accounts for two major sources of variability: (1) the model-based uncertainty conditional on a particular decomposition; and (2) the uncertainty in the estimated FPC decomposition objects. Moreover, it is applicable when curves are sparsely or densely sampled, for small and large sample sizes and when the curves are measured with error.

FPCA has a long history in the statistical literature, dating at least to the treatment of growth curves in Rao (1958). For a modern introduction to the topic, see Ramsay and Silverman (2005, Chapter 8) or Jolliffe (2002, Section 12.3). Recently, Yao, Müller, and Wang (2005) proposed the use of BLUPs to predict curve-specific scores for FPC expansions in a mixed model framework; this approach is related to that of Shi, Weiss, and Taylor (1996) and Rice and Wu (2001), in which curves are expressed using penalized spline bases in a mixed model, but can be used when curves are sparsely observed. BLUPs have also been used in multilevel and longitudinal FPC analysis by Di et al. (2009) and Greven et al. (2010), respectively, to predict scores in more complex decompositions. The mixed model framework has allowed the use of Bayesian techniques for FPCA (Crainiceanu and Goldsmith 2010) and the joint modeling of curves and scalar outcomes in functional regression models (Goldsmith et al. 2012). Asymptotic consistency and distributions of scores predicted in the mixed model framework are developed by Yao et al. (2005), and asymptotic properties of FPC decompositions are considered in Hall and Hosseini-Nasab (2006). James, Hastie, and Sugar (2000) use a penalized spline approach to estimate FPC basis functions directly, rather than through decomposing an estimated covariance matrix. This article also uses the bootstrap to generate pointwise confidence intervals by taking the empirical quantiles of estimated functions; importantly, this approach accounts for uncertainty in the FPC decompositions, but not for model-based uncertainty. The bootstrap is similarly used with functional data by Hall and Hosseini-Nasab (2006) to construct confidence intervals for FPC basis functions, and by Crainiceanu et al. (2012) for inference on the mean difference curve between two groups.

As curves can be viewed as realizations from stochastic processes and thus as random objects, it may be more appropriate to refer to “prediction intervals” than to “confidence intervals” when providing bands for predicted curves. However, in keeping with existing nomenclature (e.g., Yao et al. 2005), we will use the term confidence interval when referring to bands for the curves without additional white noise error, in contrast to prediction intervals including this additional variability source. This is also in line with the use of these terms in the penalized spline literature (e.g., Ruppert, Wand, and Carroll 2003, compare page 139), where random effects are used as a device in the modeling of curves, and the term confidence rather than prediction interval is used when predicting curves without additional noise.

The remainder of the article is organized as follows. In Section 2 we review FPCA and the estimation of curves in a mixed model framework. Section 3 introduces our method for including uncertainty in principal component expansions in the estimation of curves and understanding the associated variability. Simulation results comparing several methods are presented and discussed in Section 4. Two statistically and scientifically distinct applications are considered in Section 5. We close with a discussion in Section 6.

2. Mixed Model Framework for FPCA

Given curves $Y_i(d)$, $d \in [0, 1]$, for $1 \leq i \leq I$, define the covariance operator $\Sigma Y(d, d') = \text{Cov}[Y_i(d), Y_i(d')]$ so that $\Sigma Y(d, d')$ is a bivariate function giving the covariance between two locations on the same curve. Let $\sum_{k=1}^{\infty} \lambda_k \phi_k(d) \phi_k(d')$ be the spectral decomposition of $\Sigma Y(d, d')$, where $\phi_k(d) = \{\phi_k(d) : k \in \mathbb{Z}^+\}$ are orthonormal eigenfunctions and $\lambda_1, \lambda_2, \dots$ are the corresponding nonincreasing eigenvalues. Based on this decomposition, a Karhunen–Loève expansion for $Y_i(d)$ is $Y_i(d) = \mu(d) + \sum_{k=1}^{\infty} \xi_{ik} \phi_k(d)$, where the $\xi_{ik} = \int_0^1 \{Y_i(d) - \mu(d)\} \phi_k(d) dd$ are uncorrelated random variables with mean 0 and variance λ_k , and $\mu(d) = E[Y(d)]$.

In practice, we observe the functions $Y_i(d)$ with error—that is, we observe $\tilde{Y}_i(d) = Y_i(d) + \varepsilon_i(d)$ where we assume that $\varepsilon_i(d) \sim N[0, \sigma^2]$. Moreover, functions $Y_i(d)$ are observed on finite grids $\{d_{ij}\}_{j=1}^{J_i}$ that are often sparse at the subject level. These practical concerns impact the estimation of the mean function, the FPC decomposition, and the Karhunen–Loève expansion. First, the mean function is estimated using penalized splines fit to the pooled observations under working independence; the smoothing parameter is selected via restricted maximum likelihood (REML). Next, we use a method of moments approach to construct a raw estimate of the covariance matrix and smooth the off-diagonal elements to obtain $\tilde{\Sigma}^Y(d, d')$ (Staniswalis and Lee 1998; Yao et al. 2003). Smoothing of the raw covariance matrix is performed using bivariate penalized splines, again with smoothing parameters chosen via REML. For both the mean function and the covariance matrix, other smoothing approaches could be used without modification of the underlying technique. The retained estimated principal component basis functions $\hat{\phi}(d) = \{\hat{\phi}_k(d) : k \in \{1, \dots, \hat{K}\}\}$ and score variances $\{\hat{\lambda}_k : k \in \{1, \dots, \hat{K}\}\}$ are based on the decomposition of $\tilde{\Sigma}^Y(d, d')$; the estimate $\hat{K}$ of the truncation K lag is the minimum number of components needed to explain 99% of the variability in the observed curves, although other methods for estimating K are possible. The estimate of the covariance matrix $\Sigma Y(d, d')$ is given by

$$\widehat{\Sigma}^Y(d, d') = \sum_{k=1}^{\widehat{K}} \widehat{\lambda}_k \widehat{\phi}_k(d) \widehat{\phi}_k(d') = \widehat{\phi}(d) \widehat{\Lambda} \widehat{\phi}^T(d'), \quad (1)$$

where $\widehat{\Lambda}$ is a diagonal matrix with elements $\widehat{\lambda}_1, \dots, \widehat{\lambda}_{\widehat{K}}$ based on the retained basis functions and score variances. The measurement error variance σ^2 is estimated as the average difference between the middle 60% of diagonal elements of the raw covariance matrix and $\widehat{\Sigma}^Y(d, d')$; omitting the extreme diagonal elements provides more stable estimates of the measurement error variance.

Numeric integration for the estimation of the scores may be inaccurate when the curves are sparsely observed. With this in mind, Yao et al. (2005) predict scores $\xi_i = \{\xi_{ik} : k \in K\}$ using BLUPs, effectively basing inference on the mixed model

$$\begin{aligned} \bar{Y}_i(d_{ij}) &= \mu(d_{ij}) + \sum_{k=1}^K \xi_{ik} \phi_k(d_{ij}) + \epsilon_i(d_{ij}), \\ \xi_i &\stackrel{iid}{\sim} N[0, \Lambda], \epsilon_i(d_{ij}) \stackrel{iid}{\sim} N[0, \sigma^2], \end{aligned} \quad (2)$$

where ξ_i and $\epsilon_i = \{\epsilon_i(d_{i1}), \dots, \epsilon_i(d_{iJ_i})\}$ are independent. Notationally, let $\theta = \{\phi, \mu, \Lambda, \sigma^2, K\}$ be the collection of unobserved FPC decomposition objects: the basis functions ϕ evaluated over the domain of d ; the mean function μ evaluated over the domain of d ; the matrix of score variances $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_K\}$; the measurement error variance σ^2 ; and the truncation lag K . Similarly, let $\widehat{\theta} = \{\widehat{\phi}, \widehat{\mu}, \widehat{\Lambda}, \widehat{\sigma}^2, \widehat{K}\}$ be the estimated FPC decomposition objects. Also, let $\mathbf{d}_i = \{d_{ij}\}_{j=1}^{J_i}$ and $\mathbf{d}_g$ be a dense grid on the domain of d , often taken as the union of the $\mathbf{d}_i$. In the following, $\tilde{Y}_i(\mathbf{d}_i)$ will represent the vector of observations for curve i , the vector $\phi_k(\mathbf{d}_i)$ will be the k th basis function evaluated at locations $\mathbf{d}_i$, and the $J_i \times K$ matrix $\boldsymbol{\phi}(\mathbf{d}_i)$ will be the collection of basis functions evaluated at locations $\mathbf{d}_i$. Analogous notation will represent functions evaluated over the dense grid $\mathbf{d}_g$. We emphasize the conditioning on the FPC decomposition objects notationally by conditioning on $\widehat{\theta}$.

The estimated BLUPs for the loadings are

$$\widehat{\xi}_{\widehat{\theta},i} = E[\xi_i | \tilde{Y}_i(\mathbf{d}_i), \widehat{\theta}] = \left(\widehat{\phi}^T(\mathbf{d}_i) \widehat{\phi}(\mathbf{d}_i) + \widehat{\sigma}^2 \widehat{\Lambda}^{-1} \right)^{-1} \times \widehat{\phi}^T(\mathbf{d}_i) \left(\tilde{Y}_i(\mathbf{d}_i) - \widehat{\mu}(\mathbf{d}_i) \right). \quad (3)$$

A truncated Karhunen–Loève expansion for the i th curve over the dense grid $\mathbf{d}_g$, based on the objects estimated in the FPC decomposition and the predicted scores, is

$$\widehat{Y}_{\widehat{\theta},i}(\mathbf{d}_g) = E[\tilde{Y}_i(\mathbf{d}_g) | \widehat{\xi}_{\widehat{\theta},i}, \widehat{\theta}] = \widehat{\mu}(\mathbf{d}_g) + \sum_{k=1}^{\widehat{K}} \widehat{\xi}_{\widehat{\theta},ik} \widehat{\phi}_k(\mathbf{d}_g) = \widehat{\mu}(\mathbf{d}_g) + \widehat{\phi}(\mathbf{d}_g) \widehat{\xi}_{\widehat{\theta},i}. \quad (4)$$

Model-based covariances for the predicted scores are used to obtain the covariance operator for the estimated curve

$$\text{Var}[\widehat{Y}_{\widehat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) | \widehat{\theta}] \approx \widehat{\phi}(\mathbf{d}_g) \left(\frac{1}{\widehat{\sigma}^2} \widehat{\phi}^T(\mathbf{d}_i) \widehat{\phi}(\mathbf{d}_i) + \widehat{\Lambda}^{-1} \right)^{-1} \times \widehat{\phi}^T(\mathbf{d}_g); \quad (5)$$

the curve-specific covariance in (5) is similar to the estimate of ΣY in (1), but is decreased according to the term $\frac{1}{\widehat{\sigma}^2} \widehat{\phi}(\mathbf{d}_i)^T \widehat{\phi}(\mathbf{d}_i)$. This covariance focuses on the variability in the

curve-specific expansion, and does not include uncertainty in the estimated overall mean function $\hat{\mu}(\mathbf{d}_g)$. Next, approximate $(1 - \alpha)\%$ level pointwise confidence intervals for the predicted curves are given by

$$\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) \pm \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) \times \sqrt{\text{diag}\left\{\text{Var}\left[\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \mid \hat{\theta}\right]\right\}}, \quad (6)$$

where $\Phi(\cdot)$ is the standard Gaussian cumulative distribution function. Similarly, approximate $(1 - \alpha)\%$ level simultaneous intervals are given by

$$\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) \pm \sqrt{\chi_{K,(1-\alpha)}^2 \text{diag}\left\{\text{Var}\left[\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \mid \hat{\theta}\right]\right\}}, \quad (7)$$

where $\chi_{K,(1-\alpha)}^2$ is the $(1 - \alpha)$ quantile of a chi-squared distribution with K degrees of freedom. The intervals in (6) and (7) are justified by asymptotic arguments in Yao et al. (2005).

3. Correct Expansions and Inference

The mixed model framework developed in Section 2 explicitly conditions on the FPC decomposition objects θ and, given the basis functions, mean function, score and measurement error variances, and truncation lag, derives best linear predictors for curve-specific loadings and variance estimates for functional expansions. When these elements are fixed at their true values, standard mixed model results support the construction described (Ruppert et al. 2003). In practice, the elements of the FPC decomposition are unknown and must be estimated from the data. As noted, asymptotic work in Yao et al. (2005) can be used to justify the above-mentioned construction for large samples, but this approach can lead to a substantial understatement of the total variability in the estimated curves when the uncertainty in the decomposition is high.

Our proposed method accounts for both model- and decomposition-based uncertainty in estimating curves and in assessing the variability of these estimates. Briefly, the framework described in Section 2 is used to construct estimates of curves and their associated variability given a particular decomposition. Iterated expectation and variance formulas combine model-based curve estimates and variances across the distribution of decompositions to account for uncertainty in FPC decompositions. We implement a bootstrap procedure to examine uncertainty in the distribution of decompositions; that is, we resample curves with replacement and construct FPC decompositions based on the bootstrap sample. Because we estimate all objects of the FPC decomposition $\theta = \{\phi, \mu, \Lambda, \sigma^2, K\}$ within each bootstrap sample, uncertainty in all of these objects is accounted for by our method. Note that other approaches for assessing uncertainty in θ could be used without altering the fundamental concept of combining two major sources of uncertainty in FPC expansions.

More specifically, we draw a sample $\tilde{\mathcal{Y}}_b$ of I functions from the full data $\tilde{\mathcal{Y}}$ with replacement. Based on this sample, we derive an estimate of the covariance operator $\hat{\Sigma}_b^{\mathcal{Y}}$ by smoothing the off-diagonal elements of the raw covariance matrix and obtain an estimate of the FPC decomposition objects $\hat{\theta}_b = \{\hat{\phi}_b, \hat{\mu}_b, \hat{\Lambda}_b, \hat{\sigma}_b^2, \hat{K}_b\}$. Conditioning on $\hat{\theta}_b$, we proceed as in Section 2 and estimate the scores $\hat{\xi}_{\hat{\theta}_b,i}$, curves $\hat{Y}_{\hat{\theta}_b,i}(\mathbf{d}_g)$, and variances $\text{Var}\left[\hat{Y}_{\hat{\theta}_b,i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \mid \hat{\theta}_b\right]$ for each element of the full data set $\tilde{\mathcal{Y}}$. Thus, resampling is

used only to assess uncertainty in the decomposition—given each bootstrap decomposition, estimates and variances are constructed for each curve Y_i . Model-based expectations and variances condition on a particular decomposition, but can be estimated for curves Y_i that are not contained in the bootstrap sample $\tilde{\mathbf{Y}}_b$.

Next, we combine information across bootstrap samples to estimate curves and construct variability estimates that account for uncertainty in the FPC decompositions. First, we use iterated expectations to estimate curves over the dense grid $\mathbf{d}_g$:

$$\hat{Y}_i(\mathbf{d}_g) = E_{\hat{\theta}} \left\{ E_{\tilde{\mathbf{Y}}_i | \hat{\theta}} \left[\tilde{Y}_i(\mathbf{d}_g) | \hat{\xi}_{\hat{\theta},i}, \hat{\theta} \right] \right\}. \quad (8)$$

Thus, estimated curves are the mean of model-based estimates taken over the distribution of decomposition objects in θ . This is estimated taking the average of the model-based estimates across our bootstrap samples. The total covariance operator of the estimated curves is based on the iterated variance formula:

$$\text{Var} \left[\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \right] = E_{\hat{\theta}} \left[\text{Var}_{\tilde{\mathbf{Y}} | \hat{\theta}} \left(\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) | \hat{\theta} \right) \right] + \text{Var}_{\hat{\theta}} \left[E_{\tilde{\mathbf{Y}} | \hat{\theta}} \left(\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) | \hat{\theta} \right) \right]. \quad (9)$$

We make several remarks regarding (9). We first emphasize that $\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g)$ and $Y_i(\mathbf{d}_g)$ are expanded using different mean and basis functions: the former is expanded in terms of estimated objects in $\hat{\theta}$ while the latter is expanded using true objects in θ . In practice we approximate the conditional variance $\text{Var}_{\tilde{\mathbf{Y}} | \hat{\theta}} \left(\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) | \hat{\theta} \right)$ using the model-based variance in (5) for each bootstrap sample, understanding that this form incorrectly assumes the same expansion for both $\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g)$ and $Y_i(\mathbf{d}_g)$. We then average the model-based variances across bootstrap samples to reflect the outer expectation with respect to $\hat{\theta}$. Similarly the conditional expectation $E_{\tilde{\mathbf{Y}} | \hat{\theta}} \left(\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g) | \hat{\theta} \right)$ includes the bias that is induced by expanding $\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g)$ in terms of the estimated objects $\hat{\theta}$. We estimate this bias by taking the difference of $\hat{Y}_{\hat{\theta}_b,i}(\mathbf{d}_g)$ and $\hat{Y}_{\hat{\theta}_{\text{full}},i}(\mathbf{d}_g)$, where $\hat{\theta}_b$ and $\hat{\theta}_{\text{full}}$ are the decompositions based on the bootstrap and full data, respectively, and consider the variance of this bias across bootstrap samples. Taken together, (9) combines the expected variation in the estimation of curves given a particular decomposition and the variability of the bias that stems from the use of estimated decomposition objects.

Pointwise confidence intervals for the predicted curves are given by

$$\hat{Y}_i(\mathbf{d}_g) \pm \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\text{diag} \left\{ \text{Var} \left[\hat{Y}_i(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \right] \right\}}. \quad (10)$$

Simultaneous confidence intervals for FPC expansions are constructed similarly to intervals for penalized splines (Ruppert et al. 2003). Specifically, a $100(1 - \alpha)\%$ simultaneous confidence interval for $Y_i(\mathbf{d}_g)$ is given by

$$\hat{Y}_i(\mathbf{d}_g) \pm m_{(1-\alpha)} \sqrt{\text{diag} \left\{ \text{Var} \left[\hat{Y}_i(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \right] \right\}}, \quad (11)$$

where $m_{(1-\alpha)}$ is the $(1 - \alpha)$ quantile of the random variable

$$\max_{d \in \mathbf{d}_g} \left| \frac{\hat{Y}_i(d) - Y_i(d)}{\sqrt{\text{diag} \left\{ \text{Var} \left[\hat{Y}_i(d) - Y_i(d) \right] \right\}}} \right|. \quad (12)$$

To estimate $m_{(1-\alpha)}$, we note that $\hat{Y}_i(\mathbf{d}_g) - Y_i(\mathbf{d}_g)$ is approximately multivariate normal centered at $\mathbf{0}$ and with the total covariance:

$$\hat{Y}_i(\mathbf{d}_g) - Y_i(\mathbf{d}_g) \stackrel{\text{approx.}}{\sim} N \left[\mathbf{0}, \text{Var} \left(\hat{\mathbf{Y}}_i(\mathbf{d}_g) - \mathbf{Y}_i(\mathbf{d}_g) \right) \right]. \quad (13)$$

We repeatedly draw functions with the estimated covariance from this distribution, and for each draw calculate the resulting value of (12); the $(1 - \alpha)$ quantile of these values is used as an estimate of $m_{(1-\alpha)}$. Because this sampling is done only once after all bootstrap samples have been used, estimating simultaneous intervals does not add considerably to the total computation time of this approach.

The construction of pointwise intervals in (10) and simultaneous intervals in (11) is based on the variance $\text{Var}[\hat{Y}_i(\mathbf{d}_g) - Y_i(\mathbf{d}_g)]$ of the curve estimate constructed using iterated expectations in (8). In practice, we approximate this by the total variability of the model-

based estimate $\text{Var}[\hat{Y}_{\hat{\theta},i}(\mathbf{d}_g) - Y_i(\mathbf{d}_g)]$. Construction of confidence intervals also assumes that the estimated curve averaged over the distribution of FPC decompositions is approximately normally distributed around the true curve. However the distribution of curves may exhibit skewness, overdispersion, or multimodality; in this case, the expected value and variance defined in (8) and (9) may not completely specify the distribution. Alternatively, one might pose more flexible parametric distributions or use empirical distributions to describe the behavior of curves accounting for variability in FPC decompositions. However, asymptotic results in Yao et al. (2005) justify the normal distribution for large sample sizes, and simulation studies in Section 4 indicate that the normality assumption is reasonable and provides confidence intervals that achieve the nominal coverage in a variety of situations, including small sample sizes and sparse observations.

4. Simulations

In this section, we conduct a simulation study to examine the accuracy of the estimated curves and the coverage probabilities of confidence intervals based on our proposed method. We construct data sets using the following model

$$\tilde{Y}_i(d) = \mu(d) + \sum_{k=1}^4 \xi_{ik} \phi_k(d) + \epsilon_i(d)$$

for d observed on the dense equally spaced grid $\mathbf{d}_g = \{(t - 0.5)/50, t = 1, \dots, 50\}$, $\mu(d) = \frac{1}{4}d$ and $\epsilon_i(d) \sim N[0, \sigma^2]$; two scenarios for the distribution of scores ξ_i are given below. The eigenfunctions are taken to be

$$\phi = \left\{ \phi_1(d) = 1; \phi_2(d) = \sqrt{3}(2d - 1); \phi_3(d) = \sqrt{5}(6d^2 - 6d + 1); \phi_4(d) = \sqrt{7}(20d^3 - 30d^2 + 12d - 1) \right\}$$

and the variances are $\lambda_k = \left\{ .75^{k-1} \right\}_{k=1}^4$. Two hundred such data sets are constructed for all possible combinations of the scenarios below:

1. Sample sizes (a) $I = 20$; (b) $I = 50$; (c) $I = 100$;

2. Measurement error variances (a) $\sigma^2 = 0.0005$; (b) $\sigma^2 = 0.0025$; (c) $\sigma^2 = 0.01$;
3. Number of points observed per curve (a) $J_i = 50$; (b) $J_i = 20$; (c) an unbalanced design. For the unbalanced design, the number of points observed per curve follow a $\text{Pois}[\lambda = 15]$ distribution; for sparse designs (b) and (c), the locations of observed points are uniformly distributed within d_g .
4. 4. a) Normally distributed scores: $\xi_{ik} \sim N[0, \lambda_k]$; b) Non-normally distributed scores (mixture of normals): ξ_{ik} drawn with equal probability from either

$$N\left[-\sqrt{\lambda_k/2}, \lambda_k/2\right] \text{ or } N\left[\sqrt{\lambda_k/2}, \lambda_k/2\right].$$

This gives a total of 54 possible designs. Additional simulations presented in the supplementary material use $K = 2$ and $K = 6$ to examine the effect of changing the number of components on the relative performance of several inferential approaches; the results of these analyses are similar to those given below.

For each data set, we estimate curves and construct 95% pointwise confidence intervals using the method described in Section 3 with 100 bootstrap samples per data set; additional unreported simulations indicate that this number of bootstrap samples is sufficient to understand the uncertainty in the estimation of FPC decomposition objects in the scenarios we consider. We estimate curves and confidence intervals using three methods: (i) the iterated expectation and variance method proposed in Section 3, labeled “IV” in Figures, Tables, and the text below; (ii) the mixed model approach estimating the FPC decomposition objects θ from the data, labeled “ $\text{MM} - \hat{\theta}$ ”; and (iii) the mixed model approach fixing the FPC decomposition objects θ at their true values, labeled “ $\text{MM} - \theta$ ”. In addition, when examining pointwise intervals we consider using the 2.5% and 97.5% quantiles of the curves estimated within each bootstrap sample, labeled “Naive Bootstrap.”

Table 1 provides the average mean squared error

$\text{AMSE} = \frac{1}{I} \sum_{i=1}^I \left[\frac{1}{50} \sum_{d \in d_g} \left(\hat{Y}_i(d) - Y_i(d) \right)^2 \right]$ of the curves estimated from IV, $\text{MM} - \hat{\theta}$, and $\text{MM} - \theta$ approaches using non-normally distributed scores (results for normally distributed scores are similar and reported in the online Appendix). For reference on the values in this table, Figure 1 displays two examples of simulated curves, the observed data points, and curve estimates with the MSE for each estimate.

The example in the left panel comes from a simulation with $I = 20$, $J_i = 20$, and $\sigma^2 = 0.01$, while the example in the right comes from a simulation with $I = 50$, $J_i = 50$, and $\sigma^2 = 0.0025$.

For densely observed curves, Table 1 indicates that the IV and $\text{MM} - \hat{\theta}$ approaches have similar performances, but that the IV method modestly outperforms the model-based approach when curves are sparsely observed. For the unbalanced design, some MSE's are quite large; this stems from estimating the measurement error variance to be zero and dramatically overfitting the observed data. In practice, these situations would be clear aberrations and other solutions (such as fixing σ^2 to be positive) could be implemented. As expected, performance improves as sample size increases and when the measurement error variance is smaller.

Next we consider inference on estimated curves. Table 2 displays the average coverage, taken across all grid points and all subjects, of pointwise confidence intervals for simulation scenarios with non-normally distributed scores (again, results for normally distributed scores are similar and reported in the online Appendix).

The results in Table 2 demonstrate the reliability of the IV method, with coverages between 92.8 and 98.5 for all simulation scenarios. Lower coverages are observed for both the naive bootstrap method and the $\text{MM} - \hat{\theta}$ approach, particularly for smaller sample sizes and sparser observations.

The naive bootstrap and $\text{MM} - \hat{\theta}$ approaches each consider only one important source of variability: the naive bootstrap approach captures uncertainty in the FPC decomposition objects θ but does not consider the model-based uncertainty in the estimation of the scores, while the $\text{MM} - \hat{\theta}$ approach includes the model-based variability in the subject-specific scores but does not consider uncertainty in the decomposition objects. As the sample size increases, θ is better estimated, so the coverage of the $\text{MM} - \hat{\theta}$ approach that conditions on this estimate improves. For the naive bootstrap interval, coverage decreases as sample size increases because θ is better estimated within each bootstrap sample, which provides more stable estimates of individual curves across samples. Thus, the bootstrap estimates themselves show lower variability, but do not reflect the uncertainty in the predicted curve-specific scores. Interestingly, coverage for the $\text{MM} - \hat{\theta}$ method decreases as the measurement error variance decreases. For scenarios in which the true measurement error variance σ^2 is small, this parameter is often estimated to be zero. When this happens, confidence intervals estimated using the $\text{MM} - \hat{\theta}$ approach are degenerate and coverage probabilities are zero. Finally, the $\text{MM} - \theta$ method attains the nominal coverage in all situations, which illustrates that the low coverages in the asymptotic approach are not caused by the mixed model framework, but by uncertainty in the FPC decomposition objects.

Table 3 provides the average coverage of simultaneous confidence intervals taken over all grid points and subjects. For the IV approach, these intervals are constructed using the quantile estimate procedure described in equations (11)–(13) in Section 3. For the $\text{MM} - \hat{\theta}$ and $\text{MM} - \theta$ methods, simultaneous intervals are given by equation (7) in Section 2.

The coverage of the IV method is somewhat lower for simultaneous intervals than for pointwise intervals, with values ranging between 84.9% and 97.3%; in general, however, the intervals approach nominal coverage. On the other hand, the $\text{MM} - \hat{\theta}$ approach has substantially worse performance in many scenarios with coverages as low as 13.2%. The effect of low measurement error variance on decreasing coverage probabilities using the mixed model approach is more evident for simultaneous intervals than for pointwise intervals, particularly for sparse data scenarios. Results for the $\text{MM} - \theta$ method indicate that intervals constructed using equation (7) are in fact conservative, with coverages of approximately 98.2% for all simulation scenarios. In the Web Appendix, we discuss the adaptation of the quantile estimation procedure used to construct simultaneous intervals for the IV method to the $\text{MM} - \hat{\theta}$ and $\text{MM} - \theta$ methods. Additional simulations indicate that using the quantile estimation procedure, rather than the χ^2 quantile in equation (7), produces confidence intervals with nominal coverage for the $\text{MM} - \theta$ approach, although coverages for the $\text{MM} - \hat{\theta}$ method drop to as low as 8.8%.

Figure 2 illustrates the results in Tables 2 and 3 by examining a single curve in a data set with $I = 20$, $J_i = 20$, and $\sigma^2 = 0.01$. Shown in the left panel of this figure is the true curve $Y_i(d)$, the observed points and the estimated curves within each bootstrap sample. This panel illustrates the amount of variability in the estimated curve based on uncertainty in the FPC decomposition objects θ . Displayed in the middle panel for the same curve are the 95% pointwise confidence intervals based on the IV approach, the $\text{MM} - \hat{\theta}$ method, and the quantiles of the bootstrapped curve estimates. The naive bootstrap approach does not consider model-based variance in the estimated scores, while the $\text{MM} - \hat{\theta}$ approach ignores

variability in $\hat{\theta}$. In the right panel are the simultaneous confidence intervals given by the IV and $\text{MM} - \hat{\theta}$ approaches. Pointwise and simultaneous confidence intervals based on the model-only approach are narrower than those based on the iterated variance approach, and do not achieve nominal coverage level.

The improvements in average mean-squared errors and coverage probabilities for the IV approach over the $\text{MM} - \hat{\theta}$ method incur computational costs. Estimation of decomposition objects requires bivariate smoothing of a covariance matrix, which can be computationally expensive and must be repeated for each bootstrap sample. The time needed to complete this step will increase with the number of points in $\mathbf{d}_g$ and with the number of knots used in the bivariate smoothing. However, in the scenarios we consider estimating the decomposition objects takes between 4 and 6 seconds for each iteration, making the proposed IV approach computationally feasible.

5. Applications

In this section we consider two statistically and scientifically distinct applications. First, we treat longitudinal observations of CD4 cell counts as sparsely observed functional data; second, we consider densely observed white-matter tract profiles. Both data sets and the code implementing our analyses are available online.

5.1 CD4 Data

Human immune deficiency virus (HIV) attacks an immune cell called the CD4 cell and thereby decreases an affected individual's response to infectious agents. Because of this, CD4 cell count per milliliter of blood is a useful surrogate measure for the progression of HIV. In this data set from the Multicenter AIDS Cohort Study (MACS), we observe CD4 cell counts for 366 infected individuals in a traditional longitudinal study (Kaslow et al. 1987). That is, subjects were observed at roughly semi-annual visits with varying numbers of visits per subjects: there are a total of 1888 data points, with between 1 and 11 observations per subject and a median of 5. CD4 counts are plotted against months since seroconversion (the time at which HIV becomes detectable) in the top-left panel of Figure 3, with a randomly selected subset of patients highlighted in red. Longitudinal data analysis approaches to this study are presented in Diggle et al. (1994), and the treatment of a related data set as sparsely observed functional data is given in Yao et al. (2005).

In this application, we wish to estimate the continuous trajectory of CD4 cell counts over time based on a limited number of observations for each subject. FPCA offers a powerful tool for estimating population-level trends by using the dense collection of observations across subjects and predicting subject-specific functions using a mixed model framework. The predicted functions are subject to model-based variability as well as uncertainty in the FPC decomposition. This second source of variability can be particularly important when functions are sparsely observed at the subject level. Here we analyze the CD4 cell count data using both the IV method developed in Section 3 and the $\text{MM} - \hat{\theta}$ method. Both approaches are implemented on the full data set as well as the randomly selected subset of 25 individuals to demonstrate the effect of sample size in the analysis of sparse functional data.

Results of the CD4 cell count application are shown in Figure 3. The top-left panel plots the CD4 cell counts against months since seroconversion, highlighting the sample used in the subset analysis. Curve estimates and intervals for a single subject in the subset analysis are provided in the top-right panel of Figure 3. Simultaneous intervals are constructed using resampling procedure for the IV approach and using the asymptotic interval given in (7) for the $\text{MM} - \hat{\theta}$ method. For this subject, the estimated CD4 cell count trajectory using the IV

method is lower for months -18 to -10 and higher for months 0 to 12 than the trajectory estimated using the mixed model approach. In these times, the pointwise and simultaneous confidence intervals constructed using the IV approach are much wider than those constructed using the $\text{MM} - \hat{\theta}$ framework. For example, using the method proposed in Section 3, at month -18 this subject's CD4 cell count is estimated to be approximately 2600 CD4 cells per milliliter of blood with upper and lower bounds of 3300 and 1900; using the mixed model approach, the CD4 cell count is estimated to be 2800 with bounds 3100 and 2500. Further, using the upper bound of the simultaneous interval to estimate the month in which CD4 cell count drops below 1500 per milliliter would yield month 10 under the IV method and month 4 under the $\text{MM} - \hat{\theta}$ approach. Thus, inference for a subject's CD4 trajectory can change dramatically depending on the procedure used to construct mean and variance estimates.

In the bottom-left panel, we compare the widths of point-wise confidence intervals yielded by the IV and $\text{MM} - \hat{\theta}$ approaches. For the full data analysis, intervals resulting from the proposed method have a similar width as those from the mixed model approach, but are up to 1.18 times as wide for months (36, 42) where there are fewer observations across subjects. In the analysis of the randomly selected 25 subjects, pointwise confidence intervals resulting from the IV approach are up to 1.52 times wider than those from the $\text{MM} - \hat{\theta}$ method, and are on average 1.21 times wider. Similarly, a comparison of the average widths of simultaneous confidence intervals given by the two methods is shown in the bottom-right panel of Figure 3. The relationship between the widths of simultaneous intervals is similar to that for pointwise intervals for the sampled data. For the full data analysis, the intervals given by the IV approach are in fact slightly narrower than those for the $\text{MM} - \hat{\theta}$ method; recall that simulations indicated that the simultaneous intervals in (7) may in fact be conservative asymptotically, while the IV simultaneous intervals could attain coverages slightly under the nominal level. Generally, these results indicate that ignoring the variability in the FPC decomposition can substantially understate the total variability in estimated curves for small samples and sparse data.

5.2 White-Matter Tract Data

Next, we consider a neuroimaging study of white matter tracts in multiple sclerosis (MS) patients and healthy controls. Intracranial white-matter consists of tracts or bundles of myelinated axons that carry electrical signals between areas of the brain. Major examples of white matter tracts are the corticospinal tracts, which connect the motor cortex to the spinal cord, and the corpus callosum, which bridges the left and right hemispheres. MS is an autoimmune disorder that results in white matter lesions, axonal damage and demyelination, and can lead to severe patient disability. DTI is an MRI modality able to resolve individual tracts by tracing the diffusion of water, which tends to be anisotropic within white matter tracts and isotropic elsewhere in the brain (Basser, Mattiello, and LeBihan 1994; Mori and Barker 1999; Basser et al. 2000; LeBihan et al. 2001). Fractional anisotropy (FA) is one of several quantities derived from DTI scans that measures overall anisotropy, and may be decreased in MS patients.

Our data set consists of 42 healthy controls and 100 MS patients, for whom we observe FA values sampled densely at 55 locations along the right corticospinal tract. The top-left panel of Figure 4 displays these FA profiles, separating MS patients and healthy controls. There is significant variability in tract profiles comparing subjects, as well as apparent measurement error in the observed FA values. In addition, 35% of subjects are missing observations for at least one tract location, and 17% of subjects are missing observations for five or more locations. For all subjects, any missing values are consecutive and begin at tract location 0, so that missingness is clustered at the beginning of the tract.

Because tract profiles are affected by disease progression, they are useful biomarkers and have been related to patient disability (Goldsmith et al. 2011). The goal of this analysis is to understand the total variability in tract profile estimates using FPCA. We compare the proposed IV method to the standard $\text{MM} - \hat{\theta}$ approach to assess the contribution of decomposition-based uncertainty on overall variability in estimated curves. Additionally, to examine the effect of sample size when curves are densely sampled, we conduct our analyses on both the full data set and on the subset of healthy controls.

Figure 4 shows the results from the analysis of the DTI tract data.

A comparison of the relative widths of pointwise confidence intervals constructed using the IV and $\text{MM} - \hat{\theta}$ methods is given in the top-right panel of Figure 4. In the full data analysis, the IV method results in confidence intervals that are on average 1.12 times wider than those given by the $\text{MM} - \hat{\theta}$ approach. When considering the 42 healthy controls only, the IV approach produces intervals that are on average 1.40 times wider than those for the $\text{MM} - \hat{\theta}$ approach. The comparison of the intervals yielded by these methods illustrates that, even for densely sampled observations, the mixed model approach that conditions on a single decomposition can significantly understate the total variability by neglecting the uncertainty in the FPC decomposition objects.

A demonstration of the difference between the IV and $\text{MM} - \hat{\theta}$ approaches is given in the bottom panels of Figure 4 using the tract profile of a single healthy control. On the left are the estimated curves and pointwise confidence intervals from the analysis based on the 42 healthy controls only; the interval based on the IV method is substantially wider at several locations on the tract, particularly near 0.0, 0.6, and 0.8. In addition, the curve estimated using iterated expectations falls outside the pointwise confidence interval given by the $\text{MM} - \hat{\theta}$ approach at several locations between 0.5 and 0.8. On the right are the estimates and pointwise intervals from the full data analysis, showing that the two approaches provide more similar intervals when a larger sample size is used. Thus for small sample sizes, accounting for the uncertainty in the FPC decomposition by using the IV method results in noticeably wider confidence intervals; as the sample size increases, the decomposition uncertainty decreases and the IV and $\text{MM} - \hat{\theta}$ methods yield more similar results.

6. Discussion

FPC analysis is ubiquitous in functional data analysis. FPC decompositions are estimated from the data, and curves are expressed using a truncated Karhunen–Loève expansion with curve-specific scores predicted from a mixed model that conditions (often implicitly) on the estimated decomposition. Previous work on estimation and inference for FPC expansions has focused on the model-based prediction of scores and the associated variability without including the uncertainty in the FPC decompositions. In many cases, including small samples or sparsely observed curves, this second source of uncertainty is non-negligible.

In this article, we propose a method for estimating curves and understanding the variability of those estimates that accounts for both model- and decomposition-based uncertainty. Specifically, we use iterated expectations to average model-based estimates that condition on specific decompositions across the distribution of those decompositions; similarly, the iterated variance formula provides correct inference by combining both sources of uncertainty. A bootstrap procedure provides an empirical distribution of the FPC decompositions. Simulations demonstrate the accuracy of our method, both in terms of the MSEs of estimated curves and in terms of the coverage of associated confidence intervals. We apply our method to two statistically and scientifically distinct data sets, one using

longitudinally observed CD4 cell counts as sparse functional data and the other using partially unobserved tract profiles from a neuroimaging study. An R implementation of our method is given in the refund package on CRAN (Crainiceanu et al. 2011).

Future work may proceed in several directions. The proposed method assesses uncertainty in the FPC decomposition through the bootstrap; alternatively, one could use a probabilistic model for the unknown objects in the decomposition. Such an approach could be integrated seamlessly with the methods developed here, but the specification of an appropriate model for the unknown objects is a difficult challenge. Inference for more complex FPC-based methods may be improved through the extension of iterated expectation and variance formulas. For instance, longitudinal functional principal components analysis (LFPCA) decomposes repeated observations of curves into baseline subject-specific components, time-varying subject-specific components, and subject-visit-specific components (Greven et al. 2010); including model-and decomposition-based variability will improve the understanding of uncertainty in this scientifically useful setting. Additionally, the proposed method has focused on uncertainty in the Karhunen–Loève expansion arising from both model-and decomposition-based variability, but it does not address the model-based uncertainty in the estimation of the mean function. Correctly including this source of variability may increase the accuracy of the estimated curves and confidence intervals, particularly in cases where the estimate of the mean function may be poor.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The work of Goldsmith and Crainiceanu was supported by Award R01NS060910 from the National Institute Of Neurological Disorders And Stroke. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of Neurological Disorders and Stroke or the National Institutes of Health. The work of Greven was funded by Emmy Noether grant GR 3793/1-1 from the German Research Foundation.

References

- Basser P, Mattiello J, LeBihan D. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*. 1994; 66:259–267. [PubMed: 8130344]
- Basser P, Pajevic S, Pierpaoli C, Duda J. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*. 2000; 44:625–632. [PubMed: 11025519]
- Crainiceanu CM, Goldsmith J. Bayesian functional data analysis using WinBUGS. *Journal of Statistical Software*. 2010; 32:1–33.
- Crainiceanu C, Reiss R, Goldsmith J, Greven S, Scheipl F, Huang L. Refund: Regression with Functional Data. R package version 0. 2011:1–5.
- Crainiceanu CM, Staicu A-M, Ray S, Punjabi N. Bootstrap-based inference on the difference in the means of two correlated functional processes. *Statistics in Medicine*. 2012 doi: 10.1002/sim.5439.
- Di C-Z, Crainiceanu CM, Caffo BS, Punjabi NM. Multilevel functional principal component analysis. *Annals of Applied Statistics*. 2009; 4:458–488. [PubMed: 20221415]
- Diggle, PJ.; Heagerty, P.; Liang, K-Y.; Zeger, SL. *Analysis of Longitudinal Data*. Oxford University Press; Oxford, U.K.: 1994.
- Goldsmith J, Crainiceanu CM, Caffo B, Reich D. Penalized functional regression analysis of white-matter tract profiles in multiple sclerosis. *NeuroImage*. 2011; 57:431–439. [PubMed: 21554962]
- Goldsmith J, Crainiceanu CM, Caffo B, Reich D. Longitudinal penalized functional regression for cognitive outcomes on neuronal tract measurements. *Journal of the Royal Statistical Society: Series C*. 2012; 61:1–17.

- Greven S, Crainiceanu CM, Caffo B, Reich D. Longitudinal functional principal component analysis. *Electronic Journal of Statistics*. 2010; 4:1022–1054. [PubMed: 21743825]
- Hall P, Hosseini-Nasab B. On properties of functional principal components analysis. *Journal of the Royal Statistical Society: Series B*. 2006; 68:109–126.
- James GM, Hastie TJ, Sugar CA. Principal component models for sparse functional data. *Biometrika*. 2000; 87:587–602.
- Jolliffe, IT. *Principal Component Analysis*. Springer; New York: 2002.
- Kaslow RA, Ostrow DG, Detels R, Phair JP, Polk BF, Rinaldo CR. The Multicenter AIDS Cohort Study: Rationale, organization and selected characteristics of the participants. *American Journal of Epidemiology*. 1987; 126:310–318. [PubMed: 3300281]
- LeBihan D, Mangin J, Poupon C, Clark C. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*. 2001; 13:534–546. [PubMed: 11276097]
- Mori S, Barker P. Diffusion magnetic resonance imaging: Its principle and applications. *The Anatomical Record*. 1999; 257:102–109. [PubMed: 10397783]
- Ramsay, JO.; Silverman, BW. *Functional Data Analysis*. Springer; New York: 2005.
- Rao CR. Some statistical methods for comparison of growth curves. *Biometrics*. 1958; 14:1–17.
- Rice JA, Wu CO. Nonparametric mixed effects models for unequally sampled noisy curves. *Biometrics*. 2001; 57:253–259. [PubMed: 11252607]
- Ruppert, D.; Wand, MP.; Carroll, RJ. *Semiparametric Regression*. Cambridge University Press; Cambridge, U.K.: 2003.
- Shi M, Weiss R, Taylor J. An analysis of paediatric CD4 counts for acquired immune deficiency syndrome using flexible random curves. *Journal of the Royal Statistical Society: Series C*. 1996; 45:151–163.
- Staniswalis J, Lee J. Nonparametric regression analysis of longitudinal data. *Journal of the American Statistical Association*. 1998; 444:1403–1418.
- Yao F, Müller HG, Clifford A, Dueker S, Follett J, Lin Y, Buchholz B, Vogel J. Shrinkage estimation for functional principal component scores with application to the population. *Biometrics*. 2003; 59:676–685. [PubMed: 14601769]
- Yao F, Müller H, Wang J. Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*. 2005; 100:577–590.

Figure 1.

Comparison of three methods for estimating curves: the proposed iterated expectation method ("IV"), the model-only approach estimating θ from the data ("MM - $\hat{\theta}$ "), and the mixed model approach fixing θ at the truth ("MM - θ "). The left panel displays a curve from a simulated data set with $I = 20$, $J_i = 20$, and $\sigma^2 = 0.01$; the right panel displays a curve from a simulation with $I = 50$, $J_i = 50$, and $\sigma^2 = 0.0025$. MSEs are provided for each estimate. This figure appears in color in the electronic version of this article.

Figure 2.

Illustration of the proposed inference method using a single function. The left panel shows the true curve, the observed points as dots, and the within-bootstrap sample estimated curves overlaid. The middle panel shows pointwise confidence intervals based on the iterated variance (“IV”) approach, the mixed model method (“ $\text{MM} - \hat{\theta}$ ”) and the quantiles of the bootstrapped estimates (“Bootstrap”). In the right panel are simultaneous confidence intervals from the IV and $\text{MM} - \hat{\theta}$ approaches. This figure appears in color in the electronic version of this article.

Figure 3.

Results from analysis of CD4 cell count data. The top-left panel shows the observed CD4 counts; a random sample of 25 subjects is highlighted. The top-right panel shows the estimated curves ("Est") for a single CD4 trajectory based on the sampled subjects and full data set using the iterated approach ("IV") and the mixed model approach (" $\text{MM} - \hat{\theta}$ "); also shown are pointwise ("PW Int") and simultaneous ("Simul Int") confidence intervals. The bottom-left and bottom-right panels compare the width of pointwise and simultaneous confidence intervals, respectively, for the IV and $\text{MM} - \hat{\theta}$ approaches averaged over all subjects; comparisons for the full and sampled data are given. Simultaneous intervals are constructed using the resampling procedure in Section 3 for the IV method and using the interval in (7) for the $\text{MM} - \hat{\theta}$ method. This figure appears in color in the electronic version of this article.

Figure 4.

Results from analysis of DTI tract data. The top-left panel shows the observed profiles for all subjects with healthy controls highlighted. The top-right panel compares the width of confidence intervals for the $\text{MM} - \hat{\theta}$ and IV approaches averaged over all subjects; analyses of the full data and healthy controls only are provided. The bottom-left and bottom-right panels show the estimated curves and pointwise confidence intervals for a single function based on the healthy controls and full data set, respectively; estimates are constructed using both the iterated approach (“IV”) and mixed model (“ $\text{MM} - \hat{\theta}$ ”) approaches. This figure appears in color in the electronic version of this article.

Table 1

100 $\times$ AMSE, taken over all grid points and subjects, of estimated curves constructed using the proposed iterated expectation method (“IV”), the mixed model approach estimating θ from the data $MM - \hat{\theta}$, and the mixed model approach with θ fixed at the true value (“MM – θ ”). Curves are sampled with $J_i = 50$, $J_i = 20$, and using an unbalanced design (“Unbal”)

	IV			$MM - \hat{\theta}$			MM – θ		
	$J_i = 50$	$= 20$	Unbal	$J_i = 50$	$= 20$	Unbal	$J_i = 50$	$= 20$	Unbal
$I = 20; \sigma^2 = 0.0005$	0.006	0.200	0.535	0.006	0.222	0.805	0.004	0.013	0.035
$\sigma^2 = 0.0025$	0.026	0.260	0.971	0.026	0.284	0.886	0.020	0.065	0.136
$\sigma^2 = 0.01$	0.100	0.478	0.801	0.099	0.485	0.867	0.079	0.242	0.407
$I = 50; \sigma^2 = 0.0005$	0.005	0.089	1.246	0.005	0.105	3.729	0.004	0.013	0.039
$\sigma^2 = 0.0025$	0.023	0.144	0.396	0.023	0.162	1.576	0.020	0.063	0.142
$\sigma^2 = 0.01$	0.087	0.335	1.009	0.088	0.349	0.643	0.079	0.237	0.419
$I = 100; \sigma^2 = 0.0005$	0.005	0.055	0.280	0.005	0.066	1.161	0.004	0.013	0.040
$\sigma^2 = 0.0025$	0.021	0.109	0.775	0.021	0.120	1.752	0.020	0.064	0.138
$\sigma^2 = 0.01$	0.083	0.292	0.515	0.083	0.301	0.530	0.078	0.237	0.415

Table 2

Average pointwise coverage, taken over all grid points and subjects, of 95% confidence intervals constructed using the proposed iterated variance method (“IV”), the approach that uses the quantiles of bootstrapped curve estimates (“Naive Bootstrap”), the mixed model approach estimating θ from the data $MM - \hat{\theta}$, and the mixed model approach with θ fixed at its true value (“ $MM - \theta$ ”). Coverages are expressed as percents

	IV			Naive Bootstrap			$MM - \hat{\theta}$			$MM - \theta$		
	$J_i = 50$	$J_i = 20$	Unbal	$J_i = 50$	$J_i = 20$	Unbal	$J_i = 50$	$J_i = 20$	Unbal	$J_i = 50$	$J_i = 20$	Unbal
$I = 20; \sigma^2 = 0.0005$	95.4	98.5	98.4	68.0	94.8	94.2	89.2	29.0	33.2	95.1	94.8	94.9
$\sigma^2 = 0.0025$	95.3	96.4	96.3	63.5	88.6	88.0	90.5	50.7	46.6	95.1	94.8	94.9
$\sigma^2 = 0.01$	95.3	94.5	93.4	61.1	76.0	74.9	91.0	80.7	75.6	95.1	94.7	94.8
$I = 50; \sigma^2 = 0.0005$	94.9	96.4	97.5	51.5	89.9	89.7	92.6	36.2	35.3	95.0	94.9	95.0
$\sigma^2 = 0.0025$	94.8	93.5	94.2	46.7	78.7	79.1	93.0	66.4	57.2	95.0	95.0	95.1
$\sigma^2 = 0.01$	94.9	94.4	93.7	44.3	61.8	62.3	93.2	88.7	85.4	95.0	95.0	95.2
$I = 100; \sigma^2 = 0.0005$	95.0	94.4	96.0	39.2	86.0	85.4	94.0	35.2	42.1	95.1	95.1	95.0
$\sigma^2 = 0.0025$	95.0	92.8	92.9	35.2	71.5	71.2	94.1	76.0	69.2	95.1	95.1	95.0
$\sigma^2 = 0.01$	94.9	94.9	94.3	33.5	52.4	51.9	94.1	91.7	90.7	95.1	95.1	95.1

Table 3

Average coverage of 95% simultaneous confidence intervals constructed using the proposed iterated variance method (“IV”), the mixed model approach estimating θ from the data $MM - \hat{\theta}$ and the mixed model approach with θ fixed at its true value (“MM – θ ”). Coverages are expressed as percents

	IV			$MM - \hat{\theta}$			MM – θ		
	$J_i = 50$	$= 20$	Unbal	$J_i = 50$	$= 20$	Unbal	$J_i = 50$	$= 20$	Unbal
$I = 20; \sigma^2 = 0.0005$	93.4	98.2	97.3	81.7	13.2	17.4	98.0	98.1	98.3
$\sigma^2 = 0.0025$	94.2	94.4	93.6	87.2	26.1	25.7	98.0	98.1	98.3
$\sigma^2 = 0.01$	94.7	92.0	88.3	89.1	60.0	53.8	98.0	98.2	98.1
$I = 50; \sigma^2 = 0.0005$	92.4	93.1	95.6	91.7	19.1	25.6	98.2	98.2	98.3
$\sigma^2 = 0.0025$	94.0	86.6	87.9	94.3	42.7	40.9	98.3	98.2	98.2
$\sigma^2 = 0.01$	94.4	92.7	90.3	94.9	83.1	77.3	98.3	98.2	98.2
$I = 100; \sigma^2 = 0.0005$	92.4	85.7	90.0	94.8	20.2	31.5	98.3	98.3	98.3
$\sigma^2 = 0.0025$	94.1	85.7	84.9	96.4	57.0	52.5	98.3	98.3	98.3
$\sigma^2 = 0.01$	94.4	94.0	92.4	96.6	91.8	88.9	98.4	98.4	98.3