JOHNS HOPKINS
BLOOMBERG
SCHOOL *of* PUBLIC HEALTH

Johns Hopkins University, Dept. of Biostatistics Working Papers

11-7-2012

LIKELIHOOD RATIO TESTS FOR THE MEAN STRUCTURE OF CORRELATED FUNCTIONAL PROCESSES

Ana-Maria Staicu
Department of Statistics, North Carolina State University

Yingxing Li
The Wang Yanan Institute for Studies in Economics, Xiamen University, Xiament, China

Ciprian Crainiceanu
Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics, ccrainic@jsph.edu

David M. Ruppert
Cornell University

Likelihood Ratio Tests for the Mean Structure of Correlated Functional Processes

Ana-Maria Staicu[*]   Yingxing Li[†]   Ciprian M. Crainiceanu[‡]

David Ruppert[§]

August 7, 2012

Abstract

The paper introduces a general framework for testing hypotheses about the structure of the mean function of complex functional processes. Important particular cases of the proposed framework are: 1) testing the null hypotheses that the mean of a functional process is parametric against a nonparametric alternative; and 2) testing the null hypothesis that the means of two possibly correlated functional processes are equal or differ by only a simple parametric function. A global pseudo likelihood ratio test is proposed and its asymptotic distribution is derived. The size and power properties of the test are confirmed in realistic simulation scenarios. Finite sample power results indicate that the proposed test is much more powerful than competing alternatives. Methods are applied to testing the equality between the means of normalized δ-power of sleep electroencephalograms of subjects with sleep-disordered breathing and matched controls.

1

Some Key Words: functional data, longitudinal data, pseudo likelihood, Sleep Health Heart Study, two sample problem.

1   Introduction

We introduce pseudo likelihood ratio testing (pseudo LRT) for hypotheses about the structure of the mean of complex functional or longitudinal data. The main theoretical results are: 1) the asymptotic distribution of the pseudo LRT under general assumptions; and 2) simple sufficient conditions for these general assumptions to hold in the cases of longitudinal and functional data. The methods are applied to testing whether there is a difference between the average normalized δ-power of 51 subjects with sleep-disordered breathing (SDB) and 51 matched controls.

Tests of a parametric null hypothesis against a nonparametric alternative when the errors are independent and identically distributed has been under intense methodological development. For example, Fan, Zhang, and Zhang (2001) introduced a generalized likelihood ratio test, while Crainiceanu and Ruppert (2004) and Crainiceanu *et al.* (2005) introduced a likelihood ratio test. In contrast, development for non-independent errors has received less attention, although there are some results. For example, Zhang and Chen (2007) proposed hypothesis testing about the mean of functional data based on discrepancy measures between the estimated means under the null and alternative models. Their approach requires a dense sampling design. We propose a pseudo LRT for various parametric null hypotheses about the mean function against nonparametric alternatives when errors are correlated. The pseudo LRT can be applied to dense or sparse functional data, with or without missing observations. Our simulation results show that in cases where the approach of Zhang and Chen applies, the pseudo LRT is considerably more powerful.

Here we consider a wider spectrum of null hypotheses, which includes the hypothesis that the means of two functional processes are the same. There are many recent methodological developments that address this problem. For example, Fan and Lin (1998) developed an adjusted Neyman testing procedure for independent stationary linear Gaussian processes; Cuevas, Febrero, and Fraiman (2004) proposed an F-test for independent processes; Staicu, Lahiri, and Carroll (2012) considered an L^2-norm-based global testing procedure for dependent processes; Crainiceanu *et al.* (2012) introduced bootstrap-based procedures using joint confidence intervals. Our global pseudo LRT procedure has the advantage that it is appli-

2

cable to independent or dependent samples of curves with both dense and sparse sampling design.

Our approach is based on modeling the mean function as a penalized spline with a mixed effect representation (Ruppert *et al.* 2003). Various hypotheses of interest can then be formulated as a combination of assumptions that variance components and fixed effects parameters are zero. Testing for a zero variance component in this context is non-standard as the parameter is on the boundary of the parameter space (Self and Liang, 1987) and the vector of observations cannot be partitioned into independent subvectors. Our approach is inspired by Crainiceanu and Ruppert (2004) who derived the finite sample and asymptotic null distributions of the LRT in linear mixed models (LMM) when the errors are independent. Here we allow the errors to have a general covariance structure, which is treated as a nuisance parameter. We discuss a pseudo LRT obtained from the LRT by replacing the error covariance by a consistent estimator. Pseudo LRTs with parameters of interest or nuisance parameters on the boundary are discussed by Liang and Self (1996) and Chen and Liang (2010), respectively. Their derivations of the asymptotic null distributions require that the estimated nuisance parameters are $\sqrt{n}$-consistent—this assumption does not usually hold when the nuisance parameters have infinite dimension, e.g., for functional data.

We demonstrate that, if an appropriate consistent estimator of the error covariance is used, then the asymptotic null distribution of the pseudo LRT statistic is the same as the distribution of the LRT using the true covariance. For longitudinal data, we discuss some commonly used models and show that under standard assumptions one obtains a suitable consistent estimator of the covariance. For both densely and sparsely sampled functional data, we use smoothness assumptions to derive appropriate consistent estimators of the covariance function. The methodology is extended to testing for differences between group means of two dependent or independent samples of curves, irrespective of their sampling design. The main innovations of this paper are the extension of pseudo likelihood to infinite dimensional nuisance parameters, suitable for modeling correlation functions in longitudinal data analysis (LDA) and functional data analysis (FDA), and the development of a rigorous asymptotic theory for testing null hypotheses about the structure of the population mean for clustered data.

The remainder of the paper is organized as follows. Section 2 presents the general methodology and the null asymptotic distribution of the pseudo LRT for dependent data. Section 3 discusses applications of the pseudo LRT for LDA and FDA. The pseudo LRT properties

are evaluated by a simulation study in Section 4. Testing equality of two mean curves is presented in Section 5 and illustrated using the Sleep Heart Health Study data in Section 6. A brief discussion is found in Section 7.

2 Pseudo LRT for dependent data

In this section we describe the models and hypotheses considered, introduce the pseudo LRT, and derive this test's finite-sample and asymptotic null distributions. Let Y_{ij} be the jth measurement of the response on the ith subject at time point t_{ij}, $1 \leq j \leq m_i$ and $1 \leq i \leq n$, and let $\mu_{ij} = \mu(t_{ij})$ where μ is the population mean curve. We are interested in hypothesis testing about the mean function μ when the within-subject correlation is complex. We model the population mean function $\mu(\cdot)$ nonparametrically as a degree pth truncated spline $\mu(t) = \sum_{k=0}^{p} \beta_k t^k + \sum_{k=1}^{K} b_k (t - \kappa_k)_+^p$, where $x_+^p = \max(0, x)^p$. Here $\kappa_1, \ldots, \kappa_K$ are knots and K is assumed to be large enough to ensure the desired flexibility (see Ruppert, 2002; Ruppert *et al.*, 2003). The coefficients β_k are fixed effects and smoothing is induced by assuming that b_k are independent $N(0, \sigma_b^2)$ random coefficients. We focus here on truncated polynomial splines because they are easy to explain, but methods can be applied to any type of penalized spline.

Denote by $\boldsymbol{X}_i$ the $m_i \times (p+1)$ dimensional matrix with the jth row equal to $(1, t_{ij}, \ldots, t_{ij}^p)$, by $\boldsymbol{\beta}$ the $(p+1) \times 1$ dimensional vector of β_k, by $\boldsymbol{Z}_i$ the $m_i \times K$ dimensional matrix with jth row equal to $\{(t_{ij} - \kappa_1)_+^p, \ldots, (t_{ij} - \kappa_K)_+^p\}$, and by $\mathbf{b}$ the K-dimensional vector of b_k. If $\mathbf{Y}_i$ is the $m_i \times 1$ dimensional vector of Y_{ij}, then the model for $\mathbf{Y}_i$ can be written as a LMM

$$\mathbf{Y}_i = \boldsymbol{X}_i \boldsymbol{\beta} + \boldsymbol{Z}_i \mathbf{b} + \mathbf{e}_i, \text{ for } i = 1, \ldots, n, \tag{1}$$

where $\mathbf{b}$ is assumed $N(0, \sigma_b^2 \boldsymbol{I}_K)$, $\mathbf{e}_i$ has mean zero and covariance matrix $\boldsymbol{\Sigma}_i$, and $\mathbf{b}, \mathbf{e}_1, \ldots, \mathbf{e}_n$ are mutually independent. Within each vector $\mathbf{e}_i$ the errors are allowed to be correlated to capture the within-cluster variability. Note that (1) is not the Laird and Ware model for longitudinal data (Laird and Ware, 1982), which requires $\mathbf{b}$ to depend on the cluster i and for $\mathbf{b}_1, \ldots, \mathbf{b}_n$ to be mutually independent. Thus, unlike standard LMMs the data in model (1) cannot be partitioned into independent subvectors. Therefore, standard asymptotic theory of mixed effects models does not directly apply to model (1), and different asymptotic distributions are obtained than in the Laird and Ware model (Crainiceanu and Ruppert, 2004). Many hypotheses of interest about the structure of the mean function μ are equivalent

4

to hypotheses about the fixed effects $\beta_0, \ldots, \beta_p$ and the variance component σ_b^2. Let Q denote a subset of $\{0, 1, \ldots, p\}$; the null hypothesis that μ has a specified polynomial form can be formulated as

$$H_0 : \beta_q = 0 \text{ for } q \in Q \text{ and } \sigma_b^2 = 0 \quad \text{versus} \quad H_A : \exists\, q_0 \in Q \text{ such that } \beta_{q_0} \neq 0 \text{ or } \sigma_b^2 > 0 \quad (2)$$

When $\boldsymbol{\Sigma}_i = \sigma_e^2 \mathbf{I}_{m_i}$ such hypotheses have been tested by Crainiceanu and Ruppert (2004) and Crainiceanu *et al.* (2005) using LRTs. Here we extend results to the case when $\boldsymbol{\Sigma}_i$ is not necessary diagonal to capture the complex correlation structures of longitudinal and functional data; see Sections 3 and 5 for examples of commonly used $\boldsymbol{\Sigma}_i$. In Section 5 we also extend testing to include null hypotheses of no difference between the means of two groups. For now we focus on the simpler case, which comes with its own set of subtleties.

Our methodology is based on the assumption that the joint distribution of $\mathbf{b}$ and $\mathbf{e}_i$'s is multivariate Gaussian, but the simulation results in Section 4 indicates that our pseudo LRT is robust to this assumption. Let $\mathbf{e}$ be the stacked vector of $\mathbf{e}_i$'s, $\mathbf{Y}$ the stacked vector of $\mathbf{Y}_i$'s, and $\boldsymbol{X}$ and $\boldsymbol{Z}$ be the stacked matrices of $\boldsymbol{X}_i$'s and $\boldsymbol{Z}_i$'s, respectively. Also let $N = \sum_{i=1}^n m_i$ be the total number of observations and $\boldsymbol{\Sigma}$ be an $N \times N$ block diagonal matrix, where the ith block is equal to $\boldsymbol{\Sigma}_i$, for $i = 1, \ldots, n$. When $\boldsymbol{\Sigma}$ is known, twice the log-likelihood of $\mathbf{Y}$ is, up to an additive constant, $2 \log L_{\mathbf{Y}}(\boldsymbol{\beta}, \sigma_b^2) = -\log(|\boldsymbol{\Sigma} + \sigma_b^2 \boldsymbol{Z}\boldsymbol{Z}^T|) - (\mathbf{Y} - \boldsymbol{X}\boldsymbol{\beta})^T (\boldsymbol{\Sigma} + \sigma_b^2 \boldsymbol{Z}\boldsymbol{Z}^T)^{-1} (\mathbf{Y} - \boldsymbol{X}\boldsymbol{\beta})$, and the LRT statistic is $LRT_N = \sup_{H_0 \cup H_A} 2 \log L_{\mathbf{Y}}(\boldsymbol{\beta}, \sigma_b^2) - \sup_{H_0} 2 \log L_{\mathbf{Y}}(\boldsymbol{\beta}, \sigma_b^2)$. Here $|\cdot|$ is the determinant of a square matrix.

In practice, $\boldsymbol{\Sigma}$ is typically unknown, so we consider testing the hypothesis using pseudo LRT obtained by replacing $\boldsymbol{\Sigma}$ in the LRT by an estimate $\widehat{\boldsymbol{\Sigma}}$. Denote by $\boldsymbol{A}^{-1/2}$ a matrix square root of $\boldsymbol{A}^{-1}$, where $\boldsymbol{A}$ is a positive definite matrix, and let $\widehat{\mathbf{Y}} = \widehat{\boldsymbol{\Sigma}}^{-1/2}\mathbf{Y}$, $\widehat{\boldsymbol{X}} = \widehat{\boldsymbol{\Sigma}}^{-1/2}\boldsymbol{X}$, $\widehat{\boldsymbol{Z}} = \widehat{\boldsymbol{\Sigma}}^{-1/2}\boldsymbol{Z}$. Thus, twice the pseudo log likelihood is, up to a constant,

$$2 \log \widehat{L}_{\widehat{\mathbf{Y}}}(\boldsymbol{\beta}, \sigma_b^2) = -\log |\widehat{\mathbf{H}}_{\sigma_b^2}| - (\widehat{\mathbf{Y}} - \widehat{\mathbf{X}}\boldsymbol{\beta})^T \widehat{H}_{\sigma_b^2}^{-1} (\widehat{\mathbf{Y}} - \widehat{\mathbf{X}}\boldsymbol{\beta}), \quad (3)$$

where $\widehat{\boldsymbol{H}}_{\sigma_b^2} = \boldsymbol{I}_N + \sigma_b^2 \widehat{\boldsymbol{Z}}\widehat{\boldsymbol{Z}}^T$, and the pseudo LRT statistic for testing (2) is $pLRT_N = \sup_{H_0 \cup H_A} 2 \log \widehat{L}_{\widehat{\mathbf{Y}}}(\boldsymbol{\beta}, \sigma_b^2) - \sup_{H_0} 2 \log \widehat{L}_{\widehat{\mathbf{Y}}}(\boldsymbol{\beta}, \sigma_b^2)$. The asymptotic null distribution of the pseudo LRT is discussed next.

PROPOSITION 2.1. *Suppose that* $\mathbf{Y}$ *is obtained from model (1), and assume a Gaussian joint distribution for* $\mathbf{b}$ *and* $\mathbf{e}$, *where* $\mathbf{e} = (\mathbf{e}_1^T, \ldots, \mathbf{e}_n^T)^T$. *In addition, assume the following:*

(C1) The null hypothesis H_0 defined in (2) holds.

<div align="center">5</div>

(C2) The minimum eigenvalue of $\boldsymbol{\Sigma}$ is bounded away from 0 as $n \to \infty$. Let $\widehat{\boldsymbol{\Sigma}}$ be an estimator of $\boldsymbol{\Sigma}$ satisfying $\mathbf{a}^T\widehat{\boldsymbol{\Sigma}}^{-1}\mathbf{a} - \mathbf{a}^T\boldsymbol{\Sigma}^{-1}\mathbf{a} = o_p(1)$, $\mathbf{a}^T\widehat{\boldsymbol{\Sigma}}^{-1}\mathbf{e} - \mathbf{a}^T\boldsymbol{\Sigma}^{-1}\mathbf{e} = o_p(1)$, where $\mathbf{a}$ is any $N \times 1$ non random normalized vector.

(C3) There exists positive constants ϱ and ϱ' such that $N^{-\varrho}\mathbf{Z}^T\mathbf{Z}$ and $N^{-\varrho'}\mathbf{X}^T\mathbf{X}$ converge to nonzero matrices. For every eigenvalue $\tilde{\xi}_{k,N}$ and $\tilde{\zeta}_{k,N}$ of the matrices $N^{-\varrho}\mathbf{Z}^T\boldsymbol{\Sigma}^{-1}\mathbf{Z}$ and $N^{-\varrho}\{\mathbf{Z}^T\boldsymbol{\Sigma}^{-1}\mathbf{Z} - \mathbf{Z}^T\boldsymbol{\Sigma}^{-1}\mathbf{X}(\mathbf{X}^T\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}^T\boldsymbol{\Sigma}^{-1}\mathbf{Z}\}$ respectively, we have $\tilde{\xi}_{k,N} \xrightarrow{P} \xi_k$ and $\tilde{\zeta}_{k,N} \xrightarrow{P} \zeta_k$ for some $\xi_1, \ldots, \xi_K$, $\zeta_1, \ldots, \zeta_K$ that are not all 0.

Let $\#Q$ be the cardinality of the set Q in the null hypothesis (2). Then

$$pLRT_N \xrightarrow{D} \sup_{\lambda \geq 0} LRT_\infty(\lambda) + \sum_{j=1}^{\#Q} \nu_j^2, \tag{4}$$

where $LRT_\infty(\lambda) = \sum_{k=1}^{K} \frac{\lambda}{1+\lambda\zeta_k} w_k^2 - \sum_{k=1}^{K} \log(1 + \lambda\xi_k)$, $w_k \sim N(0, \zeta_k)$ for $k = 1, \ldots, K$, $\nu_j \sim N(0,1)$ for $j = 1, \ldots, \#Q$, and the w_k's and ν_j's are mutually independent.

Here we used $\xrightarrow{P}$ to denote convergence in probability and $\xrightarrow{D}$ to denote convergence in distribution. The proof of Proposition 2.1, like all proofs, is given in the Web Supplement. The proposition shows that when using an appropriate covariance estimator, the asymptotic null distribution of the pseudo LRT is the same as the null distribution of the corresponding LRT when the covariance is known. Assumption (C2) provides a necessary condition for how close the estimated $\widehat{\boldsymbol{\Sigma}}^{-1}$ and the true $\boldsymbol{\Sigma}^{-1}$ precision matrices have to be. This condition (see also Cai, Liu and Luo, 2011) is related to the rate of convergence between the estimator and the true precision matrix in the spectral norm. For example, if $\|\widehat{\boldsymbol{\Sigma}}^{-1} - \boldsymbol{\Sigma}^{-1}\|_2 = o_p(1)$ then the first part of (C2) holds, where $\|\boldsymbol{A}\|_2$ denotes the spectral norm of a matrix $\boldsymbol{A}$ defined by $\|\boldsymbol{A}\|_2 = \sup_{|\mathbf{x}|_2 \leq 1} |\boldsymbol{A}\mathbf{x}|_2$ and $|\mathbf{a}|_2 = \sqrt{\sum_{i=1}^{r} a_i^2}$ for $\mathbf{a} \in R^r$. Such an assumption may seem difficult to verify, but in Sections 3 and 5 we show that it is satisfied by many estimators of covariance structures commonly employed in LDA and FDA. Assumption (C3) is standard in LRT; for example, when $\mathbf{Z}$ is the design matrix for truncated power polynomials with equally spaced knots (see Section 3.2), taking $\varrho = 1$ is a suitable choice (Crainiceanu, 2003).

Consider the particular case when there are m observations per subject and identical design points across subjects, i.e., $t_{ij} = t_j$, so that $\boldsymbol{X}_i$ and $\boldsymbol{Z}_i$ do not depend on i and $\boldsymbol{\Sigma} = \boldsymbol{I}_n \otimes \boldsymbol{\Sigma}_0$ where $\otimes$ is the Kronecker product. Then (C2) is equivalent to:

(C2$'$) *The minimum eigenvalue of* $\boldsymbol{\Sigma}_0$ *is bounded away from* 0. *Let* $\widehat{\boldsymbol{\Sigma}}_0$ *be its consistent estimator satisfying* $\mathbf{a}^T \widehat{\boldsymbol{\Sigma}}_0^{-1} \mathbf{a} - \mathbf{a}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{a} = o_p(1)$, *and* $\mathbf{a}^T \widehat{\boldsymbol{\Sigma}}_0^{-1} \mathbf{e}_0 - \mathbf{a}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{e}_0 = o_p(1)$, *where* $\mathbf{a}$ *is any* $m \times 1$ *non random normalized vector and* $\mathbf{e}_0 = n^{-1/2} \sum_{i=1}^{n} \mathbf{e}_i$.

The asymptotic null distribution of $pLRT_N$ is not standard. However, as Crainiceanu and Ruppert (2004) point out, the null distribution can easily be simulated, once the eigenvalues ξ_k's and ζ_k's are determined. For completeness, we review their proposed algorithm.

Step 1 For a sufficiently large L, define a grid $0 = \lambda_1 < \lambda_2 < \cdots < \lambda_L$ of possible values for λ.

Step 2 Simulate independent $N(0, \zeta_k)$ random variables w_k, $k = 1, \ldots, K$.

Step 3 Compute $LRT_\infty(\lambda)$ in (4) and determine its maximizer $\lambda_{\max}$ on the grid.

Step 4 Compute $pLRT = LRT_\infty(\lambda_{\max}) + \sum_{j=1}^{\#Q} \nu_j^2$, where the ν_j's are i.i.d. $N(0,1)$.

Step 5 Repeat Steps 2–4.

The R package `RLRsim` (Scheipl, Greven, and Küchenhoff, 2008) or a MATLAB function `http://www.biostat.jhsph.edu/~ccrainic/software.html` can be used for implementation of Algorithm 1. It takes roughly 1.8 seconds to simulate 100,000 simulations from the null distribution using `RLRsim` on a standard computer (64-bit Windows with 2.8 GHz Processors and 24 GB random access memory).

3   Applications to longitudinal and functional data

We now turn our attention to global tests of parametric assumptions about the mean function in LDA and FDA and describe simple sufficient conditions under which assumption (C2) or (C2$'$) holds. This will indicate when results in Section 2 can be applied for testing.

3.1   Longitudinal data

Statistical inference for the mean function has been one of the main foci of LDA research (Diggle *et al.* 2002). Longitudinal data are characterized by repeated measurements over time on a set of individuals. Observations on the same subject are likely to remain correlated even after covariates are included to explain observed variability. Accounting for this

correlation in LDA is typically done using several families of covariances. Here we focus first on the case of commonly used parametric covariance structures. Consider the general model

$$Y_{ij} = \mu(t_{ij}) + e_i(t_{ij}), \qquad \text{cov}\{e_i(t_{ij}), e_i(t_{ij'})\} = \sigma_e^2 \varphi(t_{ij}, t_{ij'}; \boldsymbol{\theta}), \tag{5}$$

where t_{ij} is the time point at which Y_{ij} is observed and $\mu(t)$ is a smooth mean function. The random errors $e_{ij} = e_i(t_{ij})$ are assumed to have a covariance structure that depends on the variance parameter, σ_e^2, and the function $\varphi(\cdot, \cdot; \boldsymbol{\theta})$, which is assumed to be a positive definite function known up to the parameter $\boldsymbol{\theta} \in \Theta \subset R^d$.

Using the penalized spline representation of the mean function, $\mu(t_{ij}) = \boldsymbol{X}_{ij}\boldsymbol{\beta} + \boldsymbol{Z}_{ij}\mathbf{b}$, the model considered here can be written in a LMM framework $\mathbf{Y}_i = \boldsymbol{X}_i\boldsymbol{\beta} + \boldsymbol{Z}_i\mathbf{b} + \mathbf{e}_i$, where $\mathbf{e}_i = \{e_i(t_{i1}), \ldots, e_i(t_{im})\}^T$ has covariance matrix $\boldsymbol{\Sigma}_i = \sigma_e^2 \boldsymbol{C}_i(\boldsymbol{\theta})$, and $\boldsymbol{C}_i(\boldsymbol{\theta})$ is an $m_i \times m_i$ dimensional matrix with the (j, j')th entry equal to $\varphi(t_{ij}, t_{ij'}; \boldsymbol{\theta})$. Hypothesis testing can then be carried out as in Section 2. Proposition 3.1 below provides simpler sufficient conditions for the assumption (C2) to hold.

PROPOSITION 3.1. *Suppose that for model (5) the number of observation per subject m_i is bounded, for all $i = 1, \ldots, n$, the regularity conditions (A1)-(A3) in the Appendix hold, $\sigma_\epsilon^2 > 0$, $\sqrt{n}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}) = O_p(1)$, and $\widehat{\sigma}_e^2 - \sigma_e^2 = o_p(1)$. Then condition (C2) holds for $\widehat{\Sigma} = \widehat{\sigma}_e^2 \text{diag}\{C_1(\widehat{\boldsymbol{\theta}}), \ldots, C_n(\widehat{\boldsymbol{\theta}})\}$.*

One approach that satisfies these assumption is quasi-maximum likelihood estimation, as considered in Fan and Wu (2008). The authors proved that, under regularity assumptions that include (A1)-(A3), the quasi-maximum likelihood estimator $\widehat{\boldsymbol{\theta}}$, and the nonparametric estimator $\widehat{\sigma}^2$ are asymptotically normal, with $\sqrt{n}$ convergence rates.

3.2 Functional data

In contrast to longitudinal data, where the number of time points is small, and simple correlation structures are warranted, functional data have a much larger number of sampling points and require flexible correlations structures; see Rice (2004) for a thorough discussion of longitudinal and functional data and analytic methods. It is theoretically and practically useful to think of functional data as realizations of an underlying stochastic process.

Let $V_1, \ldots, V_n$ denote independent and identically distributed random functions on a bounded and closed time interval $\mathcal{T}$, satisfying $\int_{\mathcal{T}} V_i(t)^2 \, dt < \infty$. For simplicity take

8

$\mathcal{T} = [0, 1]$. The mean function $\mu(t) = EV_i(t)$ is considered unknown and assumed to have continuous second order derivatives. The covariance function $\text{cov}\{V_i(t), V_i(t')\} = \Gamma(t, t')$, also unknown, is assumed to be continuous over $[0, 1]$. Mercer's lemma (see for example Section 1.2 of Bosq, 2000) now implies a spectral decomposition of the function Γ, in terms of a sequence of continuous eigenfunctions, also known as functional principal components, θ_k, and decreasing sequence σ_k^2 of non-negative eigenvalues, $\Gamma(t, t') = \sum_k \sigma_k^2 \theta_k(t) \theta_k(t')$, where $\sum_k \sigma_k^2 < \infty$. Following the usual convention, we assume that $\sigma_1^2 > \sigma_2^2 > \ldots \geq 0$. The eigenfunctions form an orthonormal sequence in the space of squared integrable functions and we may represent each curve using the Karhunen-Loève (KL) expansion (Karhunen, 1947; Loève, 1945) as $V_i(t) = \mu(t) + \sum_{k \geq 1} \xi_{ik} \theta_k(t), t \in [0, 1]$, where ξ_{ik} are uncorrelated random variables with mean zero and variance $E[\xi_{ik}^2] = \sigma_k^2$.

We consider the general case, where the random curves are observed with noise that is independent of the curves. Let $t_{i1}, \ldots, t_{im_i}$ be the times at which the ith random curve is observed. Furthermore, let $Y_i(t_{ij})$ be the jth observation of the random function $V_i(\cdot)$ observed at time t_{ij} and let $\epsilon_{ij} = \epsilon_i(t_{ij})$ be the additional measurement errors that are assumed independent and identically distributed. Then, the model we consider is

$$Y_i(t_{ij}) = \mu(t_{ij}) + \sum_{k \geq 1} \xi_{ik} \theta_k(t_{ij}) + \epsilon_{ij}, \tag{6}$$

where ϵ_{ij} are assumed to have mean zero and finite variance $E[\epsilon_{ij}^2] = \sigma_\epsilon^2$. Our objective is to carry out different hypothesis tests about the population mean function $\mu(\cdot)$ by employing the pseudo LRT. As argued in Proposition 2.1 this testing procedure relies on an accurate estimator of the model covariance, and, thus, of the covariance function Γ and the noise variance σ_ϵ^2.

The FDA literature contains several methods for obtaining consistent estimators of both the eigenfunctions/eigenvalues and the error variance; see for example Ramsay and Silverman (2005), Yao, Müller and Wang (2005). Furthermore, properties of the functional principal component estimators, including their convergence rates, have been investigated by a number of researchers (Hall and Hosseini-Nasab, 2006; Hall, Müller and Wang, 2006; Li and Hsing, 2010, etc.) for a variety of sampling design scenarios. In particular for a dense sampling design, where $m_i = m$, Hall *et al.* (2006) argue that one can first construct de-noised trajectories $\widehat{Y}_i(t)$ by running a local linear smoother over $\{t_{ij}, Y_i(t_{ij})\}_j$, and then estimate all eigenvalues and eigenfunctions by conventional PCA as if $\widehat{Y}_i(t)$ were generated from the true model and without any error. They point out that when $m = n^{1/4+\nu}$ for

9

$\nu > 0$ and the smoothing parameter is appropriately chosen, one can obtain estimators of eigenfunctions/eigenvalues with $\sqrt{n}$ consistency. Of course, for a sparse sampling design, the estimators enjoy different convergence rates.

For our theoretical developments we assume that in (6), ξ_{ik} and ϵ_{ij} are jointly Gaussian. This assumption has been commonly employed in functional data analysis; see for example Yao, et al. (2005). Simulation results, reported in Section 4.1, indicate that the proposed method is robust in regard to violations of the Gaussian assumption. Moreover, we assume that the covariance function Γ has M non-zero eigenvalues, where $1 \le M < \infty$. The number of eigenvalues M is considered unknown and it can be estimated using the percentage of variance explained, AIC, BIC or testing for zero variance components, as discussed Staicu, Crainiceanu and Carroll (2010). We use the percentage variance explained in the simulation experiment and the data analysis. Next, we discuss the pseudo LRT procedure separately for the dense sampling design and for the sparse sampling design. More specifically we discuss conditions such that the requirement (C2) of Proposition 2.1 holds.

Dense sampling design. This design refers to the situation where the times, at which the trajectories are observed, are regularly spaced in $[0, 1]$ and increase to ∞ with n. We assume that each curve i is observed at common time points, i.e., $t_{ij} = t_j$ for all $j = 1, \ldots, m$. Thus $\boldsymbol{\Sigma}_i$ is the same for all subjects, say $\boldsymbol{\Sigma}_i = \boldsymbol{\Sigma}_0$ for all i.

PROPOSITION 3.2. *Consider that the above assumptions for model (6) hold. Assume the following conditions hold:*

(F1) *If $\widehat{\theta}_k(t)$, $\widehat{\sigma}_k^2$, and $\widehat{\sigma}_\epsilon^2$ denote the estimators of the eigenfunctions, eigenvalues, and noise variance correspondingly, then*

$\||\widehat{\theta}_k - \theta_k\|| = O_p(n^{-\alpha})$, $\widehat{\sigma}_k^2 - \sigma_k^2 = O_p(n^{-\alpha})$, *and* $\widehat{\sigma}_\epsilon^2 - \sigma_\epsilon^2 = O_p(n^{-\alpha})$.

(F2) *We have $m \sim n^\delta$ where $0 < \delta < 2\alpha$.*

Then (C2) of Proposition 2.1 holds for the estimator $\widehat{\boldsymbol{\Sigma}} = \boldsymbol{I}_n \otimes \widehat{\boldsymbol{\Sigma}}_0$ of $\boldsymbol{\Sigma}$, where $\widehat{\boldsymbol{\Sigma}}_0$ is defined by

$$[\widehat{\boldsymbol{\Sigma}}_0]_{jj'} = \sum_{k=1}^{M} \widehat{\sigma}_k^2 \widehat{\theta}_k(t_j) \widehat{\theta}_k(t_{j'}) + \widehat{\sigma}_\epsilon^2 1(t_j = t_{j'}), \qquad 1 \le j, j' \le m. \qquad (7)$$

Assumption (F1) concerns the L^2 convergence rate of the estimators; for local linear smoothing, Hall, et al. (2006) showed that the optimal L^2 convergence rate is $n^{-\alpha}$ where

10

$\alpha = 1/2$. Condition (F2) imposes an upper bound on the number of repeated measurements per curve: this requirement is needed in the derivation of the asymptotic null distribution of the pseudo LRT. In particular, when linear smoothing is used and $\alpha = 1/2$ (see Hall et al., 2006), condition (F2) reduce to $m = n^\delta$, for $1/4 < \delta < 1$. Nevertheless, empirical results showed that the pseudo LRT performs well, even when applied to settings where the number of repeated measurements is much larger than the number of curves. In particular, Section 4.1 reports reliable results for the pseudo LRT applied to data settings where m is up to eight times larger than n.

Remark 1. An alternative approach for situations where m is much larger than n is to use the following two-step procedure. First estimate the eigenfunctions / eigenvalues and the noise variance using the whole data, and then apply the pseudo LRT procedure only to a subset of the data that corresponds to suitably chosen subset of time points $\{\tilde{t}_1, \ldots, \tilde{t}_{\tilde{m}}\}$ where $\tilde{m}$ is such that it satisfies assumption, (F2). Our empirical investigation of this approach shows that the power does not change with m and that there is some loss of power for smaller sample sizes n. However, the power loss decreases as n increases. The alternative approach is designed for use with large m and can be used for, say, $m > 1000$ with only a negligible loss of power. Even for smaller value of m, we find that our test is more powerful than its competitor, the test due to Zhang and Chen (2007).

Remark 2. The result in Proposition 3.2 can accommodate situations when data are missing at random. More precisely, let $t_1, \ldots, t_m$ be the grid of points in the entire data and denote by n_j the number of observed responses Y_{ij} corresponding to time t_j. Under the assumption that $n_j/n \to 1$ for all j, the conclusion of Proposition 2.1 still holds.

Sparse sampling design. Sparse sampling refers to the case when observation times vary between subjects and the number of observations per subject, m_i, is bounded and small. Examples of sparse sampling are auction bid prices (Jank and Shmueli, 2006), growth data (James, Hastie and Sugar, 2001), and many observational studies. The following proposition presents simplified conditions under which the requirement (C2) of Proposition 2.1 is satisfied. The main idea is to view the sparsely observed functional data as incomplete observations from dense functional data.

PROPOSITION 3.3. *Consider that the above assumptions about the model (6) are met. In addition assume the following conditions:*

(F1') The number of measurements per subject is finite, i.e., $\sup_i m_i < \infty$. Furthermore it is

assumed that, for each subject i, the corresponding design points $\{t_{ij} : j = 1, \ldots, m_i\}$ are generated uniformly and without replacement from a set $\{t_1, \ldots, t_m\}$, where $t_k = (k - 1/2)/m$, for $k = 1, \ldots, m$ and m diverges with n.

(F2') $\sup_{t \in \mathcal{T}} |\widehat{\theta}_k(t) - \theta_k(t)| = O_p(n^{-\alpha})$, $\widehat{\sigma}_k^2 - \sigma_k^2 = O_p(n^{-\alpha})$, and $\widehat{\sigma}_\epsilon^2 - \sigma_\epsilon^2 = O_p(n^{-\alpha})$.

(F3') We have $m \sim n^\delta$ where $0 < \delta < 2\alpha$.

Then condition (C2) holds for the estimator $\widehat{\boldsymbol{\Sigma}} = \text{diag}\{\widehat{\boldsymbol{\Sigma}}_1, \ldots, \widehat{\boldsymbol{\Sigma}}_n\}$ of $\boldsymbol{\Sigma}$, where the $m_i \times m_i$ matrix $\widehat{\boldsymbol{\Sigma}}_i$ is defined similarly to (7) with $(t_j, t_{j'})$ replaced by $(t_{ij}, t_{ij'})$ and m replaced by m_i.

Condition (F1') can be weaken for design points that are generated from a uniform distribution. In such cases, the design points are rounded to the nearest $t_k = (k - 1/2)/m$, and can be viewed as being sampled uniformly without replacement from $\{t_1, \ldots, t_m\}$ for some $m \to \infty$. Because of the smoothness intrinsic to functional data (observed without noise), the effect of this rounding is asymptotically negligible when $m \to \infty$ at a rate faster than $n^{-\alpha}$. Thus condition (F1') can be relaxed to assuming that t_{ij}'s are uniformly distributed between 0 and 1. Because $0 < \delta < 2\alpha$, if m to grows at rate n or faster then the alternative approach is needed.

Condition (F2') regards uniform convergence rates of the covariance estimator; see also Li and Hsing (2010). For local linear estimators Yao *et al.* (2005) showed that, under various regularity conditions, the uniform convergence rate is of order $n^{-1/2}h_\Gamma^{-2}$, where h_Γ is the bandwidth for the two-dimensional smoother and is selected such that $nh_\Gamma^{2\ell+4} < \infty$, $\ell > 0$. When the smoothing parameter is chosen appropriately, and $\ell = 4$, the convergence rate is of order $O_p(n^{-1/3})$; thus conditions (F1') and (F3') reduce to $m = n^\delta$, for $\delta < 2/3$.

In summary, tests of the mean function in both densely and sparsely observed functional data can be carried out in the proposed pseudo LRT framework. Under the assumptions required by Propositions 3.2 and 3.3 respectively, and under the additional assumptions (C1) and (C3) of Proposition 2.1, the asymptotic null distribution of the pseudo LRT with the covariance estimator $\widehat{\boldsymbol{\Sigma}}$ is the same as if the true covariance were used and is given by expression (4).

4 Simulation study

In this section we investigate the finite sample Type I error rates and power of the pseudo LRT. Each simulated data set has n subjects. The data, $Y_i(t)$, for subject i, $i = 1, \ldots, n$, and timepoint t, $t \in \mathcal{T} = [0, 1]$, are generated from model (6) with scores ξ_{ik} that have mean zero and variance $E[\xi_{ik}^2] = \sigma_k^2$, where $\sigma_1^2 = 1$, $\sigma_2^2 = 0.5$, $\sigma_3^2 = 0.25$, and $\sigma_k^2 = 0$ for all $k \geq 4$. Also $\theta_{2k-1}(t) = \sqrt{2}\cos(2k\pi t)$ and $\theta_{2k}(t) = \sqrt{2}\sin(2k\pi t)$ for all $k \geq 1$. The interest is in testing the hypothesis H_0: $\mu(t) = 0, \forall t \in [0, 1]$, versus H_A: $\mu(t) \neq 0$ for some t. We varied μ in a family of functions parameterized by a scalar parameter $\rho \geq 0$ that controls the departure from H_0, with $\rho = 0$ corresponding to H_0. This family consists of increasing and symmetric functions $\mu_\rho(t) = \rho/\{1 + e^{10(0.5-t)}\} - \rho/2$. We used two noise variances: $\sigma_\epsilon^2 = 0.125$ (small) and $\sigma_\epsilon^2 = 2$ (large). All results are based on 1000 simulations.

4.1 Dense functional data

In this scenario, each subject is observed at m equally spaced time points $t_j = (j - 1/2)/m$, for $j = 1, \ldots, m$. We consider two types of generating distributions for the scores, ξ_{ik}: in one setting they are generated from a Normal distribution, $N(0, \sigma_k^2)$, while in another setting they are generated from a mixture distribution of two Normals $N(-\sqrt{\sigma_k^2/2}, \sigma_k^2/2)$ and $N(\sqrt{\sigma_k^2/2}, \sigma_k^2/2)$ with equal probability. We model the mean function using linear splines with K knots. The choice of K is not important, as long as it is large enough to ensure the desired flexibility (Ruppert, 2002). We selected the number of knots, based on the simple default rule of thumb $K = \max\{20, \min(0.25 \times \text{number of unique } t_j, 35)\}$ inspired from Ruppert et al. (2003). The pseudo LRT requires estimation of the covariance function, Σ, or, equivalently, Σ_0; see Section 3.2. This step is crucial as the accuracy of the covariance estimator has a sizeable impact on the performance of the pseudo LRT.

Let $\widetilde{G}(t_j, t_{j'})$ be the sample covariance estimator of $\text{cov}\{Y_i(t_j), Y_i(t_{j'})\}$, and let $\widehat{G}(\cdot, \cdot)$ be obtained by smoothing $\{\widetilde{G}(t_j, t_{j'}) : t_j \neq t_{j'}\}$ using a bivariate thin-plate spline smoother. We used the R package mgcv (Wood, 2006), with the smoothing parameter selected by restricted maximum likelihood (REML). The noise variance is estimated by $\widehat{\sigma}_\epsilon^2 = \int_0^1 \{\widetilde{G}(t, t) - \widehat{G}(t, t)\}_+ dt$; if this estimate is not positive then it is replaced by a small positive number. Denote by $\widehat{\sigma}_k^2$ and $\widehat{\theta}_k$ the kth eigenvalue and eigenfunction of the covariance $\widehat{G}$, for $k \geq 1$. The smoothing-based covariance estimator, $\widehat{\Sigma}_0$, is determined using expression (7), where M, the number of eigenvalues/eigenfunction is selected using the cumulative percentage criterion

13

(see for example Di et al., 2009). In our simulation study, we used M corresponding to 99% explained variance. Once $\widehat{\boldsymbol{\Sigma}}_0$ is obtained, the data are "pre-whitened" by multiplication with $\widehat{\Sigma}_0^{-1/2}$. Then, the pseudo LRT is applied to the transformed data. The p-value of the test is automatically obtained from the function `exactLRT` (based on 10^5 replications) of the `R` package `RLRsim` (Scheipl, et al., 2008), which implements Algorithm A1, given in Section 2.

Table 1: *Type I error rates, based on* 1000 *simulations, of the pseudo LRT for testing H_0 : $\mu \equiv 0$ in the context of dense functional data generated by model (6) with $\sigma_\epsilon^2 = 0.125$, for various n and m, and when the scores ξ_{ik} are generated from a Normal distribution (normal) or mixture distribution of two Normals (non-normal). In the pseudo LRT, the mean function is modeled using linear splines.*

(n, m)	scores distribution	$\alpha = 0.20$	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.01$
(50, 100)	normal	0.216	0.111	0.057	0.021
(50, 100)	non-normal	0.209	0.126	0.060	0.012
(50, 400)	normal	0.236	0.124	0.068	0.016
(50, 400)	non-normal	0.223	0.129	0.076	0.010
(100, 100)	normal	0.209	0.115	0.054	0.009
(100, 100)	non-normal	0.222	0.112	0.053	0.010
(100, 400)	normal	0.220	0.112	0.059	0.013
(100, 400)	non-normal	0.215	0.127	0.062	0.016
(200, 80)	normal	0.217	0.099	0.054	0.012
(200, 80)	non-normal	0.199	0.103	0.052	0.009

Table 1 shows the Type I error rates of the pseudo LRT corresponding to nominal levels $\alpha = 0.20, 0.10, 0.05$ and 0.01, and for various sample sizes ranging between $n = 50$ and $n = 200$ and m ranging between 80 and 400. Table 1 shows that the pseudo LRT using a smooth estimator of the covariance has Type I error rates that are close to the nominal level, for all significance levels. The results indicate that the performance of the pseudo LRT is robust in regard to violations of the Gaussian assumption on the scores; see the lines corresponding to 'non-normal' for the distribution of the scores.

Figure 1 shows the power functions for testing the null hypothesis $H_0 : \mu \equiv 0$. The results are only little affected by the magnitude of noise, and for brevity we only present the case of

14

low noise level. The solid lines correspond to pseudo LRT with smooth covariance estimator, the dashed lines correspond to the LRT test with known covariance matrix, and the dotted lines correspond to the global L^2-norm-based test of Zhang and Chen (2007), henceforth denoted ZC test. The performance of the pseudo LRT with the smooth covariance estimator is very close to its counterpart based on the true covariance; hence the pronounced overlap between the solid and dashed lines of the Figure 1. Overall, the results indicate that the pseudo LRT has excellent power properties, and furthermore that the power slightly improves as the number of measurements per subject m increases. Intuitively, this should be expected as a larger number of sampling curves per curve, m, corresponds more available information about the process, and thus about the mean function. By comparison, the power of the L^2 norm-based test is very low and it barely changes with m. In further simulations not reported here in the interest of space, the only situation we found where the ZC test becomes competitive for the pseudo LRT is when the deviation of the mean function from the function specified by the null hypothesis is confined to the space spanned by the eigenfunctions of the covariance function of the curves. In fact, the asymptotic theory in Zhang and Chen's Theorem 7 suggests that this would be the case where their test is most powerful.

4.2 Sparse functional data

We now consider the case when each subject is observed at m_i time points $t_{ij} \in [0,1]$, $j = 1, \ldots, m_i$, generated uniformly from the set $\{t_j = (j-1/2)/m : j = 1, \ldots, m\}$, where $m = 75$. There are $n = 250$ subjects and an equal number $m_i = 10$ time points per subject. The main difference from the dense sampling case is the calculation of the covariance estimator. For sparse data we start with a raw undersmooth covariance estimator based on the pooled data. Specifically, we first center the data $\{Y_i(t_{ij}) - \widetilde{\mu}(t_{ij})\}$, using a pooled undersmooth estimator of the mean function, $\widetilde{\mu}(t_j)$, and then construct the sample covariance of the centered data, using complete pairs of observations. At the second step, the raw estimator is smoothed using the R package mgcv (Wood, 2006). The smoothing parameter is selected via a modified generalized cross validation (GCV) or the un-biased risk estimator (UBRE) using $\gamma > 1$ to increases the amount of smoothing (Wood, 2006). The data $\{Y_i(t_{ij}) - \widetilde{\mu}(t_{ij})\}$ are correlated which causes undersmoothing, but using $\gamma > 1$ counteracts this effect. Reported results are based on $\gamma = 1.5$, a choice which was observed to yield good covariance estimators in simulations for various sample sizes. Further investigation of the choice of γ would be useful but is beyond the scope of this paper.

Table 2 illustrates the size performance of the pseudo LRT for sparse data, indicating results similar to the ones obtained for the dense sampling scenario. Figure 2 (bottom panels) shows the power functions for testing $H_0 : \mu \equiv 0$ when the true mean function is from the family described earlier. For the large noise scenario, $\sigma_\epsilon^2 = 2$, the results of the pseudo LRT with the smooth covariance estimator are very close to the counterparts based on the true covariance. This is expected, as the noise is relatively easier to estimate, and thus when the noise is a large part of the total random variation, then a better estimate of the covariance function is obtained. On the other hand, having large noise affects the power negatively. When the noise has a small magnitude, the power when the covariance estimator is used is still very good and relatively close to the power when the true covariance function is used. Results for ZC are not shown because their approach requires densely sampled data.
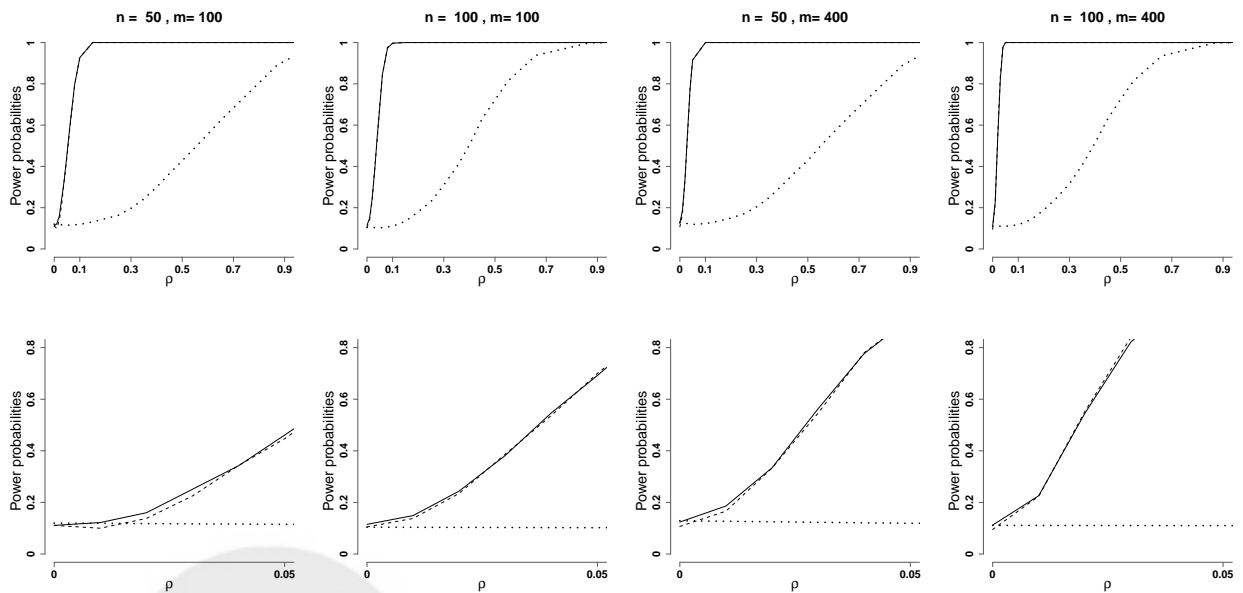

Figure 1: The power functions for testing $H_0 : \mu \equiv 0$ for dense functional data generated from model (6) with true mean function parameterized by ρ, for low noise variance $\sigma_\epsilon^2 = 0.125$. Top panels: power probabilities for different sample sizes n and number of measurements per curve m. Bottom panels: power probabilities for the same scenarios as the top panels, for $\rho \in [0, 0.05]$ to show detail in the low power region. Results are for the pseudo LRT based on the true covariance (dashed line), the smooth covariance estimator (solid line), ZC's L^2 norm-based test (dotted line) and for a nominal level $\alpha = 0.10$.

16

Table 2: *Type I error rates based on* 1000 *simulations when testing* $H_0 : \mu \equiv 0$ *with sparse functional data,* $n = 250$ *subjects and* $m_i = 10$ *observations per subject. Pseudo LRT with the true covariance (true) and a smoothing-based estimator of the covariance (smooth) are compared. The mean function is modeled using linear splines with* $K = 20$ *knots.*

σ_ϵ^2	method	cov. choice	$\alpha = 0.2$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.01$
0.125	LRT	true	0.213	0.109	0.055	0.009
	pLRT	smooth	0.210	0.113	0.062	0.017
2	LRT	true	0.207	0.092	0.051	0.011
	pLRT	smooth	0.196	0.089	0.043	0.012

5 Two samples of functional data

As with scalar or multivariate data, functional data are often collected from two or more populations, and we are interested in hypotheses about the differences between the population means. Here we consider only the case of two samples both for simplicity and because the example in Section 6 has two samples.

Again as with scalar or multivariate data, the samples can be independent or paired. The experimental cardiology study discussed in Cuevas *et al.* (2004), where calcium overload was measured at a frequency of 10s for one hour in two independent groups (control and
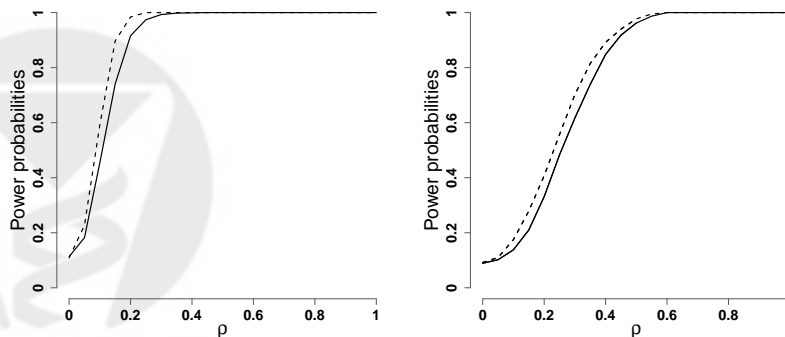
Figure 2: The power functions for testing $H_0 : \mu \equiv 0$ for sparse functional data generated from model (6) with true mean function parameterized by ρ, for two noise magnitudes $\sigma_\epsilon^2 = 0.125$ (left panel) and $\sigma_\epsilon^2 = 2$ (right panel). The results are for the pseudo LRT based on the true covariance (dashed line) and the smooth covariance estimator (solid line) and for a nominal level $\alpha = 0.10$.

treatment), is an example of independent samples of functional data. In the matched case-control study considered in Section 6, Electroencephalogram (EEG) data collected at a frequency of 125Hz for over 4 hours for an apneic group and a matched healthy control group; the matching procedure induces dependence between cases and controls. For other examples of dependent samples of functional data see, for example, Morris and Carroll (2006), Di et $al.$ (2009), and Staicu et $al.$ (2010).

We discuss global testing of the null hypothesis of equality of the mean functions in two samples of curves. Results are presented separately for independent and dependent functional data. Testing for the structure of the mean difference in two independent samples of curves can be done by straightforwardly extending the ideas presented in Section 3.2. In the interest of space, the details are described in the Web Supplement. Here we focus on the case when the two sets of curves are dependent, and furthermore when in each set, the curves are sparsely sampled.

5.1 Dependent samples of functional data

We use the functional ANOVA framework introduced by Di et $al.$ (2009) and discuss inference for the population-level curves. Let $Y_{idj} = Y_{id}(t_{idj})$ be response for cluster i and group d at time point t_{idj}. For example, in the application in Section 6, the clusters are the matched pairs and the groups are subjects with SDB and controls. Let Y_{idj} be modeled as

$$Y_{id}(t_{idj}) = \mu(t) + \mu_d(t_{idj}) + \eta_i(t_{idj}) + \upsilon_{id}(t_{idj}) + \epsilon_{idj}, \tag{8}$$

where $\mu(t)$ is the overall mean function, $\mu_d(\cdot)$ is the group-specific mean function, $\eta_i(t)$ is the cluster-specific deviation at time point t, $\upsilon_{id}(t)$ is the cluster-group deviation at t, ϵ_{idj} is the measurement error and $t_{idj} \in \mathcal{T}$ for $i = 1, \ldots, n$, $d = 1, 2$, and $j = 1, \ldots, m_{id}$. For identifiability we assume that $\mu_1 + \mu_2 \equiv 0$. It is assumed that level 1 (subject) random functions, η_i, and level 2 (subject-group) random functions, υ_{id}, are uncorrelated mean zero stochastic processes with covariance functions Γ_1 and Γ_2 respectively (Di et $al.$, 2009). Furthermore, it is assumed that the ϵ_{idj}'s are independent and identically distributed with mean zero and variance $E[\epsilon_{idj}] = \sigma_\epsilon^2$ and independent of all η_i's and υ_{id}'s. As in Section 3.2, we assume that Γ_1 and Γ_2 admit orthogonal expansions: $\Gamma_1(t, t') = \sum_{k \geq 1} \sigma_{1,k}^2 \theta_{1,k}(t) \theta_{1,k}(t')$, and $\Gamma_2(t, t') = \sum_{l \geq 1} \sigma_{2,l}^2 \theta_{2,l}(t) \theta_{2,l}(t')$. Here $\sigma_{1,1}^2 > \sigma_{1,2}^2 > \ldots$ are the level 1 ordered eigenvalues and $\sigma_{2,1}^2 > \sigma_{2,2}^2 > \ldots$ are the level 2 ordered eigenvalues. Then, the functions η_i and υ_{id} can be approximated by the KL expansion: $\eta_i(t) = \sum_{k \geq 1} \xi_{ik} \theta_{1,k}(t)$, $\upsilon_{id}(t) = \sum_{l \geq 1} \zeta_{idl} \theta_{2,l}(t)$,

18

where ξ_{ik} and ζ_{idl} are principal component scores withe mean zero and variance equal to $\sigma_{1,k}^2$ and $\sigma_{2,l}^2$. As before it is assumed that the covariance functions have finite non-zero eigenvalues and in addition that ξ_{ik}, ζ_{idl} and ϵ_{idj} are mutually independent and they are jointly Gaussian distributed.

The main objective is to test that the group mean functions are equal, or equivalently that $\mu_1 \equiv 0$. Irrespective of the sampling design (dense or sparse), we assume that the set of *pooled* time points, $\{t_{idj} : i, j\}$ is dense in $\mathcal{T}$ for each d. Our methodology requires that the same sampling scheme is maintained for the two samples of curves, e.g., the curves are not densely observed in one sample and sparsely observed in the other sample. (One could extend the theory to the case of one sample being densely observed and the other sparse, but data of this type would be rare so we did not attempt such an extension.) We use quasi-residuals, $\widetilde{Y}_{idj} = Y_{id}(t_{idj}) - \bar{\mu}(t_{idj})$, where $\bar{\mu} = (\widetilde{\mu}_1 + \widetilde{\mu}_2)/2$ is the average of the estimated mean functions, $\widetilde{\mu}_d$ for $d = 1, 2$, which are obtained using the pooled data in each group. Because of the identifiability constraint, the estimated $\bar{\mu}$ can be viewed as a smooth estimate of the overall mean function μ. We assume that the overall mean function is estimated well enough (Kulasekera, 1995), so that $\widetilde{Y}_{idj}$ can be modeled similarly to (A.1), but without μ. Thus, we assume that $\mu \equiv 0$ and that the null hypothesis is $\mu_1 \equiv 0$. The pseudo LRT methodology differs according to the sampling design. Here we focus on the setting of sparse sampled curves; the Web Supplement details the methods for the dense sampled curves.

Assume that the functions are observed at irregularly spaced time points, $\{t_{idj} : i, j\}$, and that the set of pooled time points is dense in $\mathcal{T}$ for each d. As pointed out in Crainiceanu *et al.* (2012), taking pairwise differences is no longer realistic. Nevertheless, we assume that $\mu_1(t)$ can be approximated by pth degree truncated polynomials: $\mu_1(t) = x_t \boldsymbol{\beta} + z_t \mathbf{b}$. Let $\mathbf{X}_{id}$ denote the $m_{id} \times (p+1)$ dimensional matrix with the jth row equal to $x_{t_{idj}}$, and let $\widetilde{\mathbf{X}}_i = [\mathbf{X}_{i1}^T \mid -\mathbf{X}_{i2}^T]^T$, and analogously define the $m_{id} \times K$ matrices $\mathbf{Z}_{id}$'s for $d = 1, 2$ and construct $\widetilde{\mathbf{Z}}_i = [\mathbf{Z}_{i1}^T \mid -\mathbf{Z}_{i2}^T]^T$ respectively.

Denote by $\widetilde{\mathbf{Y}}_i$ the m_i-dimensional vector obtained by stacking first $\widetilde{Y}_{i1j}$'s over $j = 1, \ldots, m_{i1}$, and then $\widetilde{Y}_{i2j}$'s over $j = 1, \ldots, m_{i2}$, where $m_i = m_{i1} + m_{i2}$. It follows that, the $m_i \times m_i$-dimensional covariance matrix of $\widetilde{\mathbf{Y}}_i$, denoted by $\boldsymbol{\Sigma}_i$ can be partitioned as

$$\boldsymbol{\Sigma}_i = \begin{pmatrix} \boldsymbol{\Sigma}_{i,11} & \boldsymbol{\Sigma}_{i,12} \\ \boldsymbol{\Sigma}_{i,21} & \boldsymbol{\Sigma}_{i,22} \end{pmatrix}, \tag{9}$$

where $\boldsymbol{\Sigma}_{i,dd}$ is $m_{id} \times m_{id}$-dimensional matrix with the (j, j') element equal to $\Gamma_1(t_{idj}, t_{idj'}) + \Gamma_2(t_{idj}, t_{idj'}) + \sigma_\epsilon^2 1(j = j')$, and $\boldsymbol{\Sigma}_{i,dd'}$ is $m_{id} \times m_{id'}$-dimensional matrix with the (j, j') element

19

equal to $\Gamma_1(t_{i1j}, t_{i2j'})$ for $d, d' = 1, 2$, $d \neq d'$. We can rewrite the model $\widetilde{\mathbf{Y}}_i$ using a LMM framework as $\widetilde{\mathbf{Y}}_i = \widetilde{\mathbf{X}}_i \boldsymbol{\beta} + \widetilde{\mathbf{Z}}_i \mathbf{b} + \mathbf{e}_i$, where $\mathbf{e}_i$ is m_i-dimensional vector, independent, with mean zero, and covariance matrix given by $\boldsymbol{\Sigma}_i$ described above. The hypothesis $\mu_1 \equiv 0$ can be tested as in Section 3.2. The required covariance estimators $\widehat{\boldsymbol{\Sigma}}_i$ are obtained by replacing $\boldsymbol{\Sigma}_{i,dd'}$ with $\widehat{\boldsymbol{\Sigma}}_{i,dd'}$ respectively for $1 \leq d, d' \leq 2$, which in turn are based on estimators of eigenfunctions, eigenvalues at each of the two levels, and the noise variance. For example, Di, Crainiceanu and Jank (2011) developed estimation methods for Γ_1, Γ_2 and σ_ϵ^2, $\{\sigma_{1,k}^2, \theta_{1,k}(t)\}_k$, and $\{\sigma_{2,l}^2, \theta_{2,l}(t)\}_l$. The next proposition presents conditions for these estimators, under which the assumption (C3), of Proposition 2.1 holds. It follows that, under the additional assumptions (C1) and (C3) of Proposition 2.1, the asymptotic null distribution of the pseudo LRT statistic is given by (4).

PROPOSITION 5.1. *Assume the following conditions for model (8) hold:*

(M1') *The number of measurements per subject per visit is finite, i.e., $\sup_i m_{id} < \infty$ for $d = 1, 2$. Furthermore it is assumed that, for each subject i, the corresponding observation points $\{t_{idj} : j = 1, \ldots, m_{id}\}$ are generated uniformly and without replacement from a set $\{t_1, \ldots, t_m\}$, where $t_k = (k - 1/2)/m$, for $k = 1, \ldots, m$ and m diverges with n*

(M2') *If $\widehat{\theta}_{d,k}(t)$, $\widehat{\sigma}_{d,k}^2$, and $\widehat{\sigma}_\epsilon^2$ denote the estimators of the group-specific eigenfunctions, eigenvalues, and of the noise variance correspondingly, then $\sup_{t \in \mathcal{T}} |\widehat{\theta}_{1,k}(t) - \theta_{1,k}(t)| = O_p(n^{-\alpha})$, $\widehat{\sigma}_{1,k}^2 - \sigma_{1,k}^2 = O_p(n^{-\alpha})$, $\sup_{t \in \mathcal{T}} |\widehat{\theta}_{2,l}(t) - \theta_{2,l}(t)| = O_p(n^{-\alpha})$, $\widehat{\sigma}_{2,l}^2 - \sigma_{2,l}^2 = O_p(n^{-\alpha})$, for all k, l, and $\widehat{\sigma}_\epsilon^2 - \sigma_\epsilon^2 = O_p(n^{-\alpha})$.*

(M3') *We have $m \sim n^\delta$ where $0 < \delta < 2\alpha$.*

Then condition (C2) holds for the estimator $\widehat{\boldsymbol{\Sigma}} = \text{diag}\{\widehat{\boldsymbol{\Sigma}}_1, \ldots, \widehat{\boldsymbol{\Sigma}}_n\}$ of $\boldsymbol{\Sigma}$, whose ith block $\boldsymbol{\Sigma}_i$ is defined in (9).

Conditions (M1')–(M3') are analogous to (F1')–(F3') and are concerned with the sampling design, the regularity of the true covariance functions, and the accuracy of the different covariance components estimation. We conclude that the sampling design assumptions can be relaxed at the cost of accurate estimation of the level 1 covariance function, Γ_1.

6 The Sleep Heart Health Study

The Sleep Heart Health Study (SHHS) is a large-scale comprehensive multi-site study of sleep and its impacts on health outcomes. Detailed descriptions of this study can be found in Quan et al. (1997), Crainiceanu et al. (2009), and Di et al. (2009). The principal goal of the study is to evaluate the association between sleep measures and cardiovascular and non-cardiovascular health outcomes. In this paper, we focus on comparing the brain activity as measured by sleep electroencephalograms (EEG) between subjects with and without sleep-disordered breathing (SDB). The SHHS collected in-home polysomnogram (PSG) data on thousands of subjects at two visits. For each subject and visit, two-channel Electroencephalograph (EEG) data were recorded at a frequency of 125Hz (125 observations/second). Here we focus on a particular characteristic of the spectrum of the EEG data, the proportion of δ-power, which is a summary measure of the spectral representation of the EEG signal. For more details on its definition and interpretations, see Borbely and Achermann (1999), Crainiceanu *et al.* (2009) and Di *et al.* (2009). In our study we use percent δ-power calculated in 30-second intervals. Figure 3 shows the sleep EEG percent δ-power in adjacent 30-second intervals for the first 4 hours after sleep onset, corresponding to 3 matched pairs of subjects; missing observations indicate wake periods. In each panel the percent δ-power is depicted in black lines for the SDB subjects and in gray lines for the corresponding matched controls.
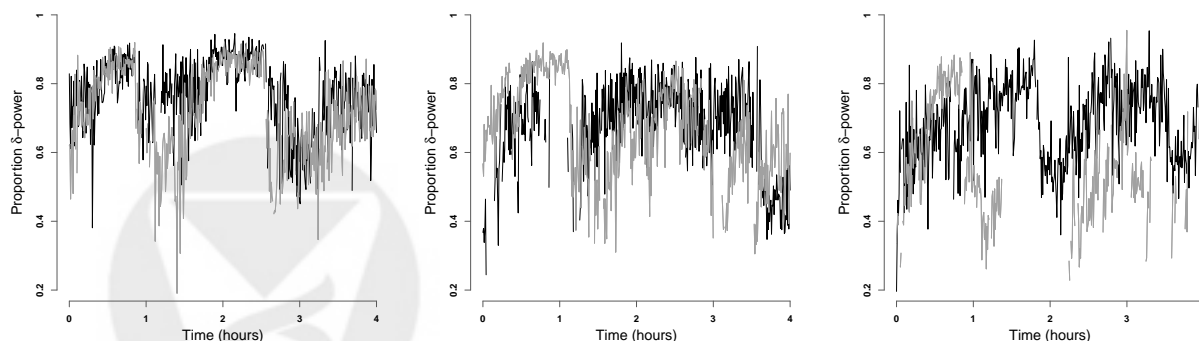

Figure 3: Sleep EEG percent δ power for the first 4 hours after sleep onset, corresponding to 3 matched pairs of controls (gray) and SDB (black).

Our interest is to compare the proportion of δ-power between the severe SDB subjects and healthy individuals, i.e., subjects without SDB, while controlling for various demographic factors. Subjects with severe SDB are identified as those with respiratory disturbance index

(RDI) greater than 30 events/hour, while subjects without SDB are identified as those with an RDI smaller than 5 events/hour. Propensity score matching (Swihart, *et al.* 2012) was used to balance the groups and minimize confounding. SDB subjects were matched with no-SDB subjects on age, BMI, race, and sex to obtain a total of 51 matched pairs. In this study missing data patterns are subject-specific, with the proportion of missingness varying dramatically across subjects. Thus, simply taking the within-group differences would be inefficient. We use pseudo LRT for dependent samples of sparse functional data, as described in Section 5.1, to test for the equality of the proportion of δ-power in the two groups.

To be specific, let $\{Y_{iA}(t), Y_{iC}(t)\}$ be the proportion of δ-power measured at the tth 30 seconds interval from sleep onset, where $t = t_1, \ldots, t_T = t_{480}$, for the ith pair of matched subjects, where A refers to the SDB and C refers to the control. For each subject some of the observations might be missing. Following Crainiceanu, et al. (2012), we model each set of curves $Y_{id}(t)$ by (8) for $d = A, C$. We are interested in testing the hypothesis H_0: $\mu_{AC} \equiv 0$, where $\mu_{AC}(t) = \mu_A(t) - \mu_C(t)$ is the difference mean function. As a preliminary step we obtain initial estimators of the group mean functions, for each of the SDB and control groups, say $\widetilde{\mu}_A(t)$ and $\widetilde{\mu}_C(t)$. We use penalized spline smoothing of all pairs $\{t, Y_{id}(t)\}$. Pseudo-residuals are calculated as $\widetilde{Y}_{id}(t) = Y_{id}(t) - \{\widetilde{\mu}_A(t) + \widetilde{\mu}_C(t)\}/2$. It is assumed that $\widetilde{Y}_{id}(t)$ can be modeled as (8), where the mean functions are $\mu_{AC}(t)$, for $d = A$ and is $-\mu_{AC}(t)$ for $d = C$ respectively. Linear splines with $K = 35$ knots are used to model the difference mean function, μ_{AC}. Pseudo LRT is applied to the pseudo-residuals, with an estimated covariance $\widehat{\Sigma}$ based on the methods in Di, et al (2011).

The pseudo LRT statistic for the null hypothesis that $\mu_{AC} \equiv 0$ is 27.74, which corresponds to a p-value $< 10^{-5}$. This indicates strong evidence against the null hypothesis of no differences between the proportion δ-power in the SDB and control group. We also tested the null hypothesis on a constant difference, that is, $\mu_{AC} \equiv a$ for some constant a; the pseudo LRT statistic is 25.63 with a p-value nearly 0. Thus, there is strong evidence that the two mean functions differ by more than a constant shift. Using a pointwise confidence intervals approach, Crainiceanu, *et al.* (2012) found that differences between the apneic and control group were not significant, indicating that their local test is less powerful than pseudo LRT when testing for global differences. The global pseudo LRT does find strong evidence against the null of no difference, but cannot pinpoint where these differences are located. We suggest using the pseudo LRT introduced in this paper to test for difference and, if differences are significant by the pseudo LRT, then locating them with the methods described

in Crainiceanu, *et al.* (2012). This combination of methods allows a more nuanced analysis and either method alone could provide.

7 Discussion

This paper develops a pseudo LRT procedure for testing the structure of the mean function and derives its asymptotic distribution when data exhibit complex correlation structure. In simulations pseudo LRT maintained its nominal level very well when a smooth estimator of the covariance was used and exhibited excellent power performance. Pseudo LRT was applied to test for the equality of mean curves in the context of two dependent or independent samples of curves. The close relation between the LRT and restricted LRT (RLRT) seems to imply that one should expect similar theoretical properties of the pseudo RLTR, obtained by substituting the true covariance by a consistent estimator, when data exhibit complex correlation structure.

References

BILLINGSLEY, P. (1968). *Convergence of Probability Measures.* Wiley, New York.

BORBELY, A. A. and ACHERMANN, P. (1999). Sleep homeostasis and models of sleep regulation. *J. Biological Rhythms* **14**, 557–568.

BOSQ, D. (2000) *Linear Processes in Function Spaces: Theory and Applications.* Lectures notes in statistics, Springer Verlag

CAI, T., LIU, W. and LUO, X. (2011). A constrained L1 minimization approach to sparse precision matrix estimation. *J. Amer. Statist. Assoc* **494**, 594–607.

CHEN, Y. and LIANG, K. Y. (2010). On the asymptotic behavior of the pseudolikelihood ratio test statistic with boundary problems. *Biometrika* **97**, 603–602.

CRAINICEANU, C. M. (2003). Nonparametric Likelihood Ratio Testing. *PhD Dissertation.* Cornell University.

CRAINICEANU, C. M., CAFFO, B. , DI, C., and PUNJABI, N. M. (2009). Nonparametric signal extraction and measurement error in the analysis of electroencephalographic activity during sleep. *J. Amer. Statist. Assoc.* **104**, 541–555.

CRAINICEANU, C. M. and RUPPERT, D. (2004). Likelihood Ratio Tests in Linear Mixed Models with One Variance Component. *Jour. Royal. Statist. Series B* **66**, 165–185.

CRAINICEANU, C. M., RUPPERT, D., CLAESKENS, G. and WAND, M. P. (2005). Exact Likelihood Ratio Tests for Penalized Splines. *Biometrika* **92**, 91–103.

CRAINICEANU, C. M., STAICU, A.-M., RAY, S., and PUNJABI, N. M. (2012). Bootstrap-based inference on the difference in the means of two correlated functional processes, *Statistics in Medicine*, to appear.

CUEVAS, A. , FEBRERO, M. and FRAIMAN, R. (2004). An anova test for functional data. *Comput. Statist. and Data Analysis* **47**, 111–122.

DI, C., CRAINICEANU, C. M., CAFFO, B. and PUNJABI, N. M. (2009). Multilevel functional principal component analysis. *Ann. Appl. Statist.* **3**, 458–488.

DI, C., CRAINICEANU, C. M. and JANK, W. (2011). Multilevel sparse functional principal component analysis. *Under review.*

DIGGLE, P., HEAGERTY, P.J., LIANG, K. Y. and ZEGER, S. (2002). *Analysis of Longitudinal Data*, 2nd Ed. Oxford University Press.

FAN, J. and LIN, S. (1998). Test of significance when data are curves. *J. Amer. Statist. Assoc.* **93**, 1007–1021.

FAN, J. and WU, Y. (2008). Semiparametric estimation of covariance matrices for longitudinal data. *J. Amer. Statist. Assoc.* **103**, 1520–1533.

FAN, J., ZHANG, C. and ZHANG, J. (2001). Generalized Likelihood Ratio Statistics and Wilks Phenomenon. *Ann. Statist.* **29**, 153–193.

GOLUB, G. H. and VAN LOAN, C. F. (1983). *Matrix Computations.* The Johns Hopkins University Press.

HARVILLE, D. A. (1997). *Matrix Algebra From A Statistician's Perspective.* Springer, New York.

HALL, P., MÜLLER, H. G. and WANG, J. L. (2006). Properties of principal component methods for functional and longitudinal data analysis. *Ann. Statist.* **34**, 1493–1517.

24

HALL, P., HOSSEINI-NASAB, M. (2006). On properties of functional principal components analysis. *J. R. Stat. Soc. Ser. B Statist. Methodol.* **68** 109–126.

JAMES, G., HASTIE, T. G., and SUGAR, C. A. (2001). Principal Component Models for Sparse Functional Data. *Biometrika* **87**, 587–602.

JANK, W. and SHMUELI, G. (2006). Functional data analysis in electronic commerce research. *Statis. Sci.* **21**, 155–166.

KARHUNEN, K. (1947). Über lineare Methoden in der Wahrscheinlichkeitsrechnung. *Ann. Acad. Sci. Fennicae. Ser. A. I. Math.-Phys.* **37** 3–79.

KULASEKERA, K. B. (1995). Comparison of regression curves using quasi-residuals. *J. Amer. Statist. Assoc.* **90**, 1085–1093.

LAIRD, N. M. and WARE, J. H. (1982). Random-Effects Models for Longitudinal Data. *Biometrics* **38**, 963–974.

LI, Y. and HSING, T. (2010). Uniform Convergence Rates for Nonparametric Regression and Principal Component Analysis in Functional/Longitudinal Data. *Ann. Statist.* **38**, 3321–3351.

LIANG, K. Y. and SELF, S. G. (1996). On the asymptotic behaviour of the pseudolikelihood ratio test statistic. *J. R. Statist. Soc. B.* **58**, 785–796.

LOÈVE, M. (1945). Fonctions aléatoires de second ordre. *C. R. Acad. Sci.* 220–469.

MORRIS, J. S. and CARROLL , R. J. (2006). Wavelet-based functional mixed models. *J. Roy. Statist. Soc. Ser. B,* **68**, 179–199.

PATTERSON, H. D. and THOMPSON, R. (1971). Recovery of Inter-Block Information when Block Sizes are Unequal. *Biometrika* **58**, 545–554.

QUAN, S. F., HOWARD, B. V., IBER, C., KILEY, J. P., NIETO, F. J., O'CONNOR, G. T., RAPOPORT, D. M., REDLINE, S., ROBBINS, J., SAMET, J. M. and WAHL, P. W. (1997). The sleep heart health study: Design, rationale, and methods. *Sleep* **20**, 1077–1085.

RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional data analysis (2nd Ed.).* New York: Springer-Verlag.

RICE, J. A. (2004). Functional and Longitudinal Data Analysis: Perspectives on smoothing. *Statist. Sinica* **14**, 631–647.

RUPPERT, D. (2002). Selecting the number of knots for penalized splines. *J. Comput. Graph. Statist.* **11**, 735–757.

RUPPERT, D., WAND, W. P., and CARROLL, R. J. (2003), *Semiparametric Regression*, New York: Cambridge University Press.

SCHEIPL, F., GREVEN, S. and KÜCHENHOFF, H. (2008). Size and power of tests for a zero random effect variance or polynomial regression in additive and linear mixed models. *Comput. Statist. and Data Analysis*, **52**, 3283–3299.

SELF, S. G. and LIANG, K. Y. (1987). Asymptotic properties of maximum likelihood and likelihood rati otests under nonstandard conditions. *J. Amer. Statist. Assoc.* **82**, 605–610.

STAICU, A.-M. , CRAINICEANU, C. M. and CARROLL, R. J. (2010). Fast methods for spatially correlated multilevel functional data. *Biostatistics* **11**, 177–194.

STAICU, A.-M., LAHIRI, S. N. and CARROLL, R. J. (2012). Tests of significance for spatially correlated multilevel functional data. *Under review.*

SWIHART, B.J., CAFFO, B., CRAINICEANU, C. M., and PUNJABI, N.M. (2012). Modeling multilevel sleep transitional data via poisson log-linear multilevel models. *Working paper*, Johns Hopkins University.

WOOD, S. N. (2006). *Generalized Additive Models: An Introduction with R.* Chapman and Hall/CRC.

WOODBURY, M. A. (1950). Inverting modified matrices, *Memorandum Report*, Princeton University.

YAO, F., MÜLLER, H. G. and WANG, J. L. (2005). Functional data analysis for sparse longitudinal data. *J. Amer. Statist. Assoc.* **100**, 577–590.

ZHANG, J. T. and CHEN, J. W. (2007). Statistical Inference for Functional Data. *Ann. Statist.* **35**, 1052–1079.