Bayesian bandwidth estimation and semi-metric selection for a functional partial linear model with unknown error density

Han Lin Shang*
Australian National University

October 26, 2019

Abstract

This study examines the optimal selections of bandwidth and semi-metric for a functional partial linear model. Our proposed method begins by estimating the unknown error density using a kernel density estimator of residuals, where the regression function, consisting of parametric and nonparametric components, can be estimated by functional principal component and functional Nadaraya-Watson estimators. The estimation accuracy of the regression function and error density crucially depends on the optimal estimations of bandwidth and semi-metric. A Bayesian method is utilized to simultaneously estimate the bandwidths in the regression function and kernel error density by minimizing the Kullback-Leibler divergence. For estimating the regression function and error density, a series of simulation studies demonstrate that the functional partial linear model gives improved estimation and forecast accuracies compared with the functional principal component regression and functional nonparametric regression. Using a spectroscopy dataset, the functional partial linear model yields better forecast accuracy than some commonly used functional regression models. As a by-product of the Bayesian method, a pointwise prediction interval can be obtained, and marginal likelihood can be used to select the optimal semi-metric.

Keywords: functional Nadaraya-Watson estimator; scalar-on-function regression; Gaussian kernel mixture; Markov chain Monte Carlo; error-density estimation; spectroscopy

*Postal address: Research School of Finance, Actuarial Studies and Statistics, Level 4, Building 26C, Kingsley Street, Australian National University, Acton ACT 2601, Canberra, Australia; Telephone number: +61(2) 612 50535; Fax number: +61(2) 612 50087; Email: hanlin.shang@anu.edu.au

1 Introduction

In scalar-on-function regression, parametric regression models are useful for interpreting the linear relationship between scalar response and functional predictor, while nonparametric regression models may capture a possible nonlinear relationship, and thus may improve the estimation and prediction accuracies of the regression function. By combining the advantages of parametric and nonparametric regression models, the semiparametric regression models, such as the functional partial linear model, have received increasing attention in the literature (e.g., [Lian, 2011](#)).

Despite rapid developments in the estimations of functional partial linear models, the optimal selections of semi-metric and bandwidth remain largely unexplored. To address this, we consider the optimal parameter selections from the perspective of error-density estimation in functional regression. The estimation of error density is important for understanding the residual behavior of regression models and assessing the adequacy of the error distribution assumption (e.g., [Akritas and Van Keilegom, 2001](#); [Cheng and Sun, 2008](#)). The estimation of error density is also useful for testing the symmetry of the residual distribution (e.g., [Ahmad and Li, 1997](#); [Neumeyer and Dette, 2007](#)); for the estimation of the density of response variable (e.g., [Escanciano and Jacho-Chávez, 2012](#)); and is important for statistical inference, prediction and model validation (e.g., [Efromovich, 2005](#); [Muhsal and Neumeyer, 2010](#)).

In a nonparametric scalar-on-function regression, [Shang \(2013\)](#) proposes a Bayesian bandwidth estimation method for determining the bandwidth of the Nadaraya-Watson (NW) estimator for estimating regression mean function and the bandwidth of the kernel-form error density. Also, in nonparametric scalar-on-function regression, [Shang \(2016\)](#) proposes a Bayesian bandwidth estimation method for determining the bandwidth of the NW estimator for estimating regression quantile function and the bandwidth of the kernel-form error density. Further, [Shang \(2016\)](#) uses marginal likelihood as a means of selecting the optimal semi-metric. Building on the early work by [Zhang et al. \(2011\)](#), [Zhang and King \(2011\)](#) and [Shang \(2013, 2014a,b, 2016, 2019\)](#), we consider a kernel error-density estimator which explores data-driven features, such as asymmetry, skewness, and multi-modality, and relies on residuals obtained from the estimated regression function and bandwidth of residuals. Differing from those early work, we derive an approximate likelihood and a posterior for the functional partial linear model (a semiparametric model). We present a Markov chain Monte Carlo (MCMC) sampling algorithm for simultaneously sampling bandwidth parameters, linear regression coefficient function $\beta(t)$, and nonlinear NW estimator of $m(z)$. Through a series of simulation studies, the functional partial linear model gives improved estimation and forecast accuracies of the regression mean function compared with the functional principal component regression and nonparametric scalar-on-function regression.

While the selection of bandwidth determines the amount of smoothing, the selection of semi-metric plays an important role in measuring distances among functions. In nonparametric func-

tional data analysis, the optimal bandwidth can be determined either by functional cross-validation (Rachdi and Vieu, 2007) or by a Bayesian method incorporating the information on error density (Shang, 2013). However, the optimal selection of semi-metric is rather arbitrary. As noted by Ferraty and Vieu (2006, Chapters 3 and 13), a semi-metric based on a derivative should be used for a set of smooth functional data; a semi-metric based on a dimension-reduction technique, such as functional principal component analysis (FPCA), should be used for a set of rough functional data. To the best of our knowledge, little attention has been paid for selecting which semi-metric is adequate based on rigorous statistical criteria.

This lack of rigorous statistical criteria motivates our investigation of a Bayesian method for estimating bandwidths in the regression function and error density simultaneously, and for selecting the optimal semi-metric, based on the notion of marginal likelihood (also referred to Kass and Raftery (1995) as the evidence). The marginal likelihood reflects a summary of evidence provided by the data supporting a choice of semi-metric as opposed to its competing semi-metric. The optimal semi-metric is that which has the largest marginal likelihood. With the marginal likelihood, it is straightforward to calculate the Bayes factor for measuring the strength of evidence between any two semi-metrics (e.g., Kass and Raftery, 1995). Selecting the optimal semi-metric and bandwidths often improves the estimation and prediction accuracies of the regression function, including the functional partial linear model considered here.

The remainder of this paper is organized as follows. In Section 2, we introduce the partial linear model with the functional covariate, previously studied by Lian (2011), and the estimation procedure in Aneiros-Pérez and Vieu (2015), Aneiros-Pérez et al. (2015) and Ling et al. (2019). In Section 3, we review the Bayesian bandwidth-estimation method and introduce the notion of marginal likelihood computed as a by-product of the MCMC. The optimal semi-metric is determined by having the largest marginal likelihood among several possible semi-metrics. Using a series of simulation studies in Section 4, we evaluate and compare the estimation accuracies of the regression function and error density. Also, we compare the forecast accuracy of regression function among the functional partial linear model, functional principal component regression, and functional nonparametric regression. In Section 5, we apply the proposed method to a spectroscopy dataset in the food quality control and compare its forecast accuracy with 13 commonly used functional regression models. Conclusions are presented in Section 6, along with some reflections on how the method presented here could be extended.

2 Model and estimator

There is an increasing amount of literature on the development of nonparametric functional estimators, such as the functional NW estimator (Ferraty and Vieu, 2006), the functional local-

linear estimator (Berlinet et al., 2011), the functional k -nearest neighbor estimator (Burba et al., 2009), and the distance-based local-linear estimator (Boj et al., 2010). In the functional partial linear model, we choose to estimate the *conditional mean* by the functional NW estimator because of its simplicity and robustness against a large gap in design points.

2.1 A functional partial linear model

When modeling scalar responses and functional predictors, some functional variables are related to responses linearly while other variables have nonlinear relationships with the response. We consider a random data triplet $(\mathcal{X}, \mathcal{Z}, y)$, where y is a real-valued response and the functional random variable $\mathcal{X}$ is valued in Hilbert space containing square-integrable functions and $\mathcal{Z}$ is valued in some infinite-dimensional semi-metric vector space $(\mathcal{F}, d(\cdot, \cdot))$. Let $(\mathcal{X}_i, \mathcal{Z}_i, y_i)_{i=1,2,\dots,n}$ be a sample of data triplets that are independent and identically distributed (iid) as $(\mathcal{X}, \mathcal{Z}, y)$. We consider a functional partial linear model with homoscedastic errors. Given a set of observations $(\mathcal{X}_i, \mathcal{Z}_i, y_i)$, the regression model can be expressed as

$$y_i = \int_{t \in \mathcal{I}} \mathcal{X}_i(t) \beta(t) dt + m(\mathcal{Z}_i) + \varepsilon_i,$$

where $\mathcal{I}$ represents function support range, $\beta(\cdot)$ is a regression coefficient function, $m(\cdot)$ is an unknown smooth real function that captures the possible nonlinear relationship between $\mathcal{Z}$ and y and $(\varepsilon_1, \dots, \varepsilon_n)$ are iid random error satisfying

$$E(\varepsilon_i | \mathcal{X}_i, \mathcal{Z}_i) = 0.$$

Following the estimation procedure of Aneiros-Pérez and Vieu (2006), it is natural to estimate the regression coefficient function $\beta(\cdot)$ and the smooth function $m(\cdot)$ by the ordinary least squares and functional NW estimators, given by

$$\hat{\beta}_h(t) = [\widetilde{\mathcal{X}}_h^\top(t) \widetilde{\mathcal{X}}_h(t)]^+ \widetilde{\mathcal{X}}_h^\top(t) \widetilde{\mathbf{y}}_h, \quad (1)$$

$$\hat{m}_h(z) = \sum_{i=1}^n w_h(z, \mathcal{Z}_i) [y_i - \langle \mathcal{X}_i(t), \hat{\beta}_h(t) \rangle], \quad (2)$$

where $^+$ represents generalized inverse and $^\top$ represents the column vector; $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_n)^\top$ and $\mathbf{y} = (y_1, \dots, y_n)^\top$, $\widetilde{\mathcal{X}}_h = (\mathbf{I} - \mathbf{W}_h) \mathcal{X}$ and $\widetilde{\mathbf{y}}_h = (\mathbf{I} - \mathbf{W}_h) \mathbf{y}$, $\mathbf{W}_h = [w_h(\mathcal{Z}_i, \mathcal{Z}_j)]_{i,j}$ is a weight matrix with $w_h(z, \mathcal{Z}_i)$ being the NW-type weights,

$$w_h(z, \mathcal{Z}_i) = \frac{K[d(z, \mathcal{Z}_i)/h]}{\sum_{j=1}^n K[d(z, \mathcal{Z}_j)/h]},$$

where $K(\cdot)$ is a kernel function, such as the Gaussian kernel function considered in this paper, h is a positive real-valued estimated bandwidth parameter, controlling the trade-off between squared

bias and variance in the mean squared error (MSE) and $d(\cdot, \cdot)$ is a semi-metric used to quantify differences among curves.

The regression coefficient function $\hat{\beta}_h(t)$ in (1) can be seen as the ordinary least squares estimator obtained by regressing the partial residual vector $\tilde{y}_h$ on the partial residual vector of functions $\tilde{\mathcal{X}}_h$. To avoid possible non-singularity, it is common to use a dimension-reduction technique, such as functional principal component regression (e.g., [Reiss and Ogden, 2007](#); [Reiss et al., 2017](#); [Lian, 2011](#)). However, $\hat{m}_h(z)$ in (2) can be seen as the functional NW estimator, with the partial residual vector as the response variable. Therefore, the estimation accuracy of $\hat{\beta}_h(t)$ and $\hat{m}_h(z)$ crucially depends on the optimal estimation of h . **With an estimated bandwidth of the NW estimator, we obtain an estimate of $\beta(t)$ at each iteration of the MCMC sampling. Then, we take average over all iterations to obtain the final estimates, namely $\hat{h}$, $\hat{\beta}_h(t)$ and $\hat{m}_h(z)$.**

In functional nonparametric regression, the bandwidth parameter is commonly determined by functional cross-validation (e.g., [Ferraty and Vieu, 2009](#); [Burba et al., 2009](#); [Barrientos-Marin et al., 2010](#); [Ferraty et al., 2010a](#)). Functional cross-validation aims to select a bandwidth that minimizes the squared loss function and has the appealing feature that no estimation of the error variance is required. However, as stated, an accurate estimation of error density is important. Thus, we consider a Bayesian method to estimate bandwidth parameters in the regression function and error density simultaneously. The Bayesian method aims to select a bandwidth that minimizes the Kullback-Leibler (KL) divergence.

2.2 Choice of semi-metric

The choice of semi-metric has effects on the size and form of neighborhoods and can thus control the concentration properties of functional data. Although [Ferraty and Vieu \(2006, p.193\)](#) note that the selection of the semi-metric remains an open question, there is no method to quantify which semi-metric is adequate, or how to select the optimal semi-metric, using statistical criteria.

From a practical perspective, the semi-metric based on derivative should be utilized for a set of smooth functional data. This semi-metric can be expressed as

$$d_q^{\text{deriv}}(\mathcal{Z}_i, \mathcal{Z}) = \sqrt{\int_t [\mathcal{Z}_i^{(q)}(t) - \mathcal{Z}^{(q)}(t)]^2 dt},$$

where $\mathcal{Z}^{(q)}$ is the q^{th} -order derivative of $\mathcal{Z}$, in which first and second derivatives are commonly used in practice (e.g., [Goutis, 1998](#); [Ferraty and Vieu, 2002, 2009](#)). Computationally, the semi-metric based on the q^{th} -order derivative uses a B -spline approximation for each curve, and a derivative of a B -spline function can be directly computed by differencing the B -spline coefficients and B -spline basis functions of one order lower (for more detail, see [de Boor, 2001](#); [Ferraty and Vieu, 2006](#), Section 3.4).

For a set of rough functional data, a semi-metric based on FPCA is commonly used (for detail

on the choice of semi-metric from the practical and theoretical perspectives, see [Ferraty and Vieu, 2006](#), Chapters 3 and 13). The semi-metric based on FPCA computes proximities among a set of rough curves using the principal component scores. FPCA reduces the functional data to a finite-dimensional space (i.e., K number of components). The finite-dimensional representation can be expressed as

$$d_K(\mathcal{Z}_i, \mathcal{Z}) \approx \sqrt{\sum_{\zeta=1}^K (\beta_{\zeta,i} - \beta_{\zeta})^2 \|\phi_{\zeta}(t)\|^2},$$

where K represents the number of retained functional principal components, and $\|v\| = \sqrt{\langle v, v \rangle}$ represents the induced norm.

3 Bayesian method

In Section 3.1, we briefly review [Shang's \(2013\)](#) Bayesian method for selecting optimal bandwidth. As a by-product of the MCMC, the optimal semi-metric can be determined with the largest marginal likelihood among a set of semi-metrics (as presented in Section 3.3).

3.1 Bayesian bandwidth estimation

[Zhang et al. \(2011\)](#) consider a kernel-form error density in a nonparametric regression model, while [Zhang and King \(2011\)](#) consider the kernel-form error density in a nonlinear parametric model where Bayesian sampling approaches are used. Following the early work by [Zhang et al. \(2011\)](#) and [Zhang and King \(2011\)](#), the unknown error density $f(\varepsilon)$ can be approximated by a location-mixture Gaussian density, given by

$$f(\varepsilon; b) = \frac{1}{n} \sum_{j=1}^n \frac{1}{b} \phi\left(\frac{\varepsilon - \varepsilon_j}{b}\right),$$

where $\phi(\cdot)$ is the probability density function of the standard Gaussian distribution and Gaussian densities have means at ε_j , for $j = 1, 2, \dots, n$ and a common standard deviation b .

Given that errors are unknown in practice, we approximate them by residuals obtained from the functional principal component and functional NW estimators of the conditional mean. Given bandwidths h and b , the kernel likelihood of $\mathbf{y} = (y_1, \dots, y_n)^\top$ is given by

$$\hat{L}(\mathbf{y}|h, b) = \prod_{i=1}^n \left[\frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \frac{1}{b} \phi\left(\frac{\hat{\varepsilon}_i - \hat{\varepsilon}_j}{b}\right) \right], \quad (3)$$

where $\hat{\varepsilon}_i = y_i - \int_{t \in \mathcal{I}} \mathcal{X}_i(t) \hat{\beta}(t) dt - \hat{m}(\mathcal{Z}_i)$ represents i^{th} residual obtained from the estimated regression function. This likelihood is not proper since some terms are left out. Instead, the likelihood is a pseudo-likelihood. The pseudo-likelihood is a likelihood function associated

with a family of probability distributions, which does not necessarily contain the true function (Gourieroux et al., 1984). As a consequence, the resulting Bayesian estimators, while consistent, may have an inaccurate posterior variance, and subsequent credible sets constructed by this posterior may not be asymptotically valid (see also Chernozhukov and Hong, 2003; Müller, 2013).

By assuming that the independent priors of squared bandwidths h^2 and b^2 follow inverse Gamma distribution $IG(10^{-3}, 10^{-3})$, the posterior density can be obtained by multiplying the kernel likelihood in (3) with the prior density. The posterior density can be expressed as (up to a normalizing constant)

$$\pi(h^2, b^2 | \mathbf{y}) \propto \widehat{L}(\mathbf{y} | h, b) \times \pi(h^2) \times \pi(b^2).$$

3.2 Sampling algorithm

Computationally, an adaptive random-walk Metropolis algorithm can be used for sampling both parameters (h, b) iteratively (see Garthwaite et al., 2016). The sampling procedure is briefly described as follows:

- 1) Specify a Gaussian proposal distribution and begin the sampling iteration process by choosing an arbitrary value of (h^2, b^2) and denoting it as $(h_{(0)}^2, b_{(0)}^2)$. The starting points can be drawn from a uniform distribution $U(0, 1)$.
- 2) From the initial bandwidths $h_{(0)}^2$ and $b_{(0)}^2$, we obtain an initial estimate of linear regression coefficient function, $\widehat{\beta}_{h_{(0)}}(t)$ from (1) and an initial estimate of nonlinear regression function, $\widehat{m}_{h_{(0)}}(z)$ from (2).
- 3) At the k^{th} iteration, the current state $h_{(k)}^2$ is updated as $h_{(k)}^2 = h_{(k-1)}^2 + \tau_{(k-1)}u$, where u is drawn from the proposal density which can be the standard Gaussian density and $\tau_{(k-1)}$ is an adaptive tuning parameter with an arbitrary initial value $\tau_{(0)}$. The updated $h_{(k)}^2$ is accepted with a probability given by

$$\min \left\{ \frac{\pi(h_{(k)}^2, b_{(k-1)}^2 | \mathbf{y})}{\pi(h_{(k-1)}^2, b_{(k-1)}^2 | \mathbf{y})}, 1 \right\}.$$

- 4) The tuning parameter for the next iteration is set to

$$\tau_{(k)} = \begin{cases} \tau_{(k-1)} + c(1 - \xi)/k & \text{if } h_{(k)}^2 \text{ is accepted} \\ \tau_{(k-1)} - c\xi/k & \text{if } h_{(k)}^2 \text{ is rejected} \end{cases}$$

where $c = \tau_{(k-1)} / (\xi - \xi^2)$ is a constant and ξ is the optimal target acceptance probability, which is 0.44 for univariate updating (e.g., Roberts and Rosenthal, 2009).

- 5) Conditional on $h_{(k)}^2$ and $\mathbf{y}$, we obtain an estimator of linear regression coefficient function, $\hat{\beta}_{h_{(k)}}(t)$, and an NW estimator of nonlinear regression function, $\hat{m}_{h_{(k)}}(z)$ (see, e.g., [Chen et al., 2017](#)).
- 6) Repeat Steps 2-5 for $b_{(k)}^2$, conditional on $h_{(k)}^2$ and $\mathbf{y}$.
- 7) Repeat Steps 2-6 for $M + N$ times, discard $(h_{(0)}^2, \beta_{h_{(0)}}(t), m_{h_{(0)}}(z), b_{(0)}^2), (h_{(1)}^2, \beta_{h_{(1)}}(t), m_{h_{(1)}}(z), b_{(1)}^2), \dots, (h_{(M)}^2, \beta_{h_{(M)}}(t), m_{h_{(M)}}(z), b_{(M)}^2)$ for burn-in to let the effects of the transients wear off, estimate $\hat{h}^2 = \frac{\sum_{k=M+1}^{M+N} h_{(k)}^2}{N}$, $\hat{\beta}_h(t) = \frac{\sum_{k=M+1}^{M+N} \beta_{h_{(k)}}(t)}{N}$, $\hat{m}_h(z) = \frac{\sum_{k=M+1}^{M+N} m_{h_{(k)}}(z)}{N}$, and $\hat{b}^2 = \frac{\sum_{k=M+1}^{M+N} b_{(k)}^2}{N}$. The burn-in period is taken to be $M = 1,000$ iterations and the number of iterations after the burn-in is $N = 10,000$ iterations.

3.3 Bayesian model selection

In Bayesian inference, model selection is often conducted through the marginal likelihood of the model of interest against a competing model. The marginal likelihood is the expectation of likelihood with respect to the prior density of parameters and reflects a summary of evidence provided by the data supporting the model as opposed to its competing model. It is seldom calculated as the integral of the product of the likelihood and prior density of parameters, but instead, is often computed numerically (e.g., [Gelfand and Dey, 1994](#); [Chib, 1995](#)). We use the method proposed by [Chib \(1995\)](#) to compute the marginal likelihood.

Let $\boldsymbol{\theta} = (h, b)$ be the parameter vector and $\mathbf{y} = (y_1, \dots, y_n)$ be the data. [Chib \(1995\)](#) demonstrated that the marginal likelihood under model A can be expressed as

$$L_A(\mathbf{y}) = \frac{L_A(\mathbf{y}|\boldsymbol{\theta})\pi_A(\boldsymbol{\theta})}{\pi_A(\boldsymbol{\theta}|\mathbf{y})},$$

where $L_A(\mathbf{y}|\boldsymbol{\theta})$, $\pi_A(\boldsymbol{\theta})$ and $\pi_A(\boldsymbol{\theta}|\mathbf{y})$ denote the likelihood, prior and posterior, respectively. $L_A(\mathbf{y})$ is often computed at the posterior estimate of $\boldsymbol{\theta}$. The numerator has a closed form and can be computed analytically. The denominator can be estimated by its kernel-density estimator based on the simulated samples of $\boldsymbol{\theta}$ through a posterior sampler. Based on marginal likelihoods, the Bayes factor of model A against model B is defined as $L_A(\mathbf{y})/L_B(\mathbf{y})$, which can be used to decide whether model A or B is preferred, along with its degree of preference ([Kass and Raftery, 1995](#)).

3.4 Adaptive estimation of error density

In the kernel-density estimation of directly observed data, it has been noted that the leave-one-out estimator is profoundly affected by extreme observations in the sample (e.g., [Bowman, 1984](#)). When the actual error density has sufficient long tails, the leave-one-out kernel-density estimator with a global bandwidth estimated under the KL divergence is likely to overestimate the tails

of the density (Shang, 2013). To address this issue, Zhang and King (2011) and Shang (2016) suggested that localized bandwidths should improve the estimation accuracy of error density for the symmetrical unimodal distribution. The kernel density estimator with localized bandwidths assigns small bandwidths to the observations in the high-density region and large bandwidths to the observations in the low-density region. The localized error-density estimator can be written as

$$\hat{f}(\varepsilon_i; \tau, \tau_\varepsilon) = \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \frac{1}{\tau (1 + \tau_\varepsilon |\hat{\varepsilon}_j|)} \phi \left(\frac{\hat{\varepsilon}_i - \hat{\varepsilon}_j}{\tau (1 + \tau_\varepsilon |\hat{\varepsilon}_j|)} \right),$$

where $\tau(1 + \tau_\varepsilon |\hat{\varepsilon}_j|)$ is the bandwidth assigned to $\hat{\varepsilon}_j$, for $j = 1, 2, \dots, n$ and the vector of parameters is now $(h, \tau, \tau_\varepsilon)$. The adaptive random-walk Metropolis algorithm, described in Section 3.2, can be used for sampling parameters, where the prior density of $\tau_\varepsilon \sim U(0, 1)$.

4 Simulation

The principal aim of this section is to illustrate the proposed method through simulated data. We consider two cases in which simulated curves are smooth in Section 4.3 and in which simulated curves are rough in Section 4.5. Using the error criteria in Sections 4.1 and 4.2, we assess the estimation accuracies of the regression function and error density, respectively. In Section 4.4, we conduct a prior sensitivity analysis.

4.1 Criteria for assessing estimation and prediction accuracies of the regression function

To measure the estimation accuracy of the regression function, we first calculate the averaged mean squared error (AMSE) between the true regression function $g(\cdot)$ and the estimated regression function $\hat{g}(\cdot)$. This is expressed as

$$\begin{aligned} \text{AMSE} &= \text{E} \left\{ [g(\mathcal{X}; \mathcal{Z}) - \hat{g}(\mathcal{X}; \mathcal{Z})]^2 \right\} = \text{E} \left(\text{E} \left\{ [g(\mathcal{X}; \mathcal{Z}) - \hat{g}(\mathcal{X}; \mathcal{Z})]^2 \mid \mathcal{X}, \mathcal{Z} \right\} \right) \\ &\approx \frac{1}{n} \sum_{i=1}^n \text{E} \left\{ [g(\mathcal{X}; \mathcal{Z}) - \hat{g}(\mathcal{X}; \mathcal{Z})]^2 \mid \mathcal{X} = \mathcal{X}_i, \mathcal{Z} = \mathcal{Z}_i \right\} \\ &\approx \frac{1}{n} \frac{1}{B} \sum_{i=1}^n \sum_{s=1}^B [g(\mathcal{X}_i; \mathcal{Z}_i) - \hat{g}^s(\mathcal{X}_i; \mathcal{Z}_i)]^2, \end{aligned}$$

where $B = 100$ represents the number of replications.

To measure the prediction accuracy of the regression function, we calculate the averaged mean squared prediction error (AMSPE) between the holdout regression function $g(\mathcal{X}_{\text{new}}; \mathcal{Z}_{\text{new}})$ and

the predicted regression function $\hat{g}(\mathcal{X}_{\text{new}}; \mathcal{Z}_{\text{new}})$. This is expressed as

$$\begin{aligned} \text{AMSPE} &= \text{E} \left\{ [g(\mathcal{X}_{\text{new}}; \mathcal{Z}_{\text{new}}) - \hat{g}(\mathcal{X}_{\text{new}}; \mathcal{Z}_{\text{new}})]^2 \right\} \\ &\approx \frac{1}{\eta} \frac{1}{B} \sum_{j=1}^{\eta} \sum_{s=1}^B [g(\mathcal{X}_j; \mathcal{Z}_j) - \hat{g}^s(\mathcal{X}_j; \mathcal{Z}_j)]^2, \end{aligned}$$

where η represents the length of prediction samples.

4.2 Criteria for assessing error-density estimation

To measure the difference between the true error density $f(\varepsilon)$ and the estimated error density $\hat{f}(\varepsilon)$, we consider two risk functions. First, we calculate the mean integrated squared error (MISE), which is given by

$$\text{MISE} [\hat{f}(\varepsilon)] = \int_{\varepsilon=a}^b [f(\varepsilon) - \hat{f}(\varepsilon)]^2 d\varepsilon,$$

for $\varepsilon \in [a, b]$. For each replication, the MISE can be approximated at 1,001 grid points bounded between a closed interval, such as $[-10, 10]$. These can be expressed as

$$\text{MISE} [\hat{f}(\varepsilon)] \approx \frac{1}{50} \sum_{i=1}^{1001} \left\{ f \left[-10 + \frac{(i-1)}{50} \right] - \hat{f} \left[-10 + \frac{(i-1)}{50} \right] \right\}^2.$$

From an average of 100 replications, the approximated mean integrated squared error (AMISE) is used to assess the estimation accuracy of error density. The AMISE is defined as

$$\text{AMISE} = \frac{1}{B} \sum_{b=1}^B \text{MISE}_b.$$

The second risk function is the KL divergence, given by

$$\text{KL}(\hat{f}, f) = -\text{E} \int_{\varepsilon=a}^b f(\varepsilon) \ln \hat{f}(\varepsilon) d\varepsilon.$$

With f fixed in simulation, $\text{KL}(\hat{f}, f)$ is minimized for $f = \hat{f}$ (Kullback, 1959).

4.3 Simulation of smooth curves

We describe the construction of the simulated data. First, we build simulated discretized curves as

$$\mathcal{X}_i(t_j) = a_i \cos(2t_j) + b_i \sin(4t_j) + c_i \left(t_j^2 - \pi t_j + \frac{2}{9} \pi^2 \right), \quad i = 1, 2, \dots, n, \quad (4)$$

where t represents the function support range and $0 \leq t_1 \leq t_2 \leq \dots \leq t_{100} \leq \pi$ are equispaced points within the function support range, a_i , b_i , c_i are independently drawn from a uniform distribution on $[0, 1]$, and n represents the sample size. The functional form of (4) is taken from Ferraty et al. (2010b) and Shang (2013, 2014b). Figure 1 presents a replication of 100 simulated smooth curves, along with the first-order derivative of the curves approximated by B-splines.

Figure 1: 100 simulated smooth curves.

Once the curves are formed, we define $\mathcal{Z}$ as the first-order derivative of the curves (see also [Mas and Pumo, 2009](#)). We compute the response variable through the following steps:

- 1) Construct a regression-function operator g , which performs the mapping from function-valued space to real-valued space. The regression-function operator g is expressed as

$$g(\mathcal{X}_i; \mathcal{Z}_i) = 10 \times \left(a_i^2 - b_i^2 \right).$$

- 2) Generate $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n$, which are independently drawn from mixture normal distributions. To highlight the possible non-normality of the error density, we consider three different error densities previously introduced by [Marron and Wand \(1992\)](#):

- (a) unimodal symmetric distribution, such as t_5 ;
- (b) skewunimodal distribution: $0.2 \times N(0, 1) + 0.2 \times N\left(0.5, \left(\frac{2}{3}\right)^2\right) + 0.6 \times N\left(13/12, \left(\frac{5}{9}\right)^2\right)$;
- (c) skewbimodal distribution: $0.75 \times N(0, 1) + 0.25 \times N\left(1.5, \left(\frac{1}{3}\right)^2\right)$.

- 3) Compute the corresponding responses: $y_i = g(\mathcal{X}_i; \mathcal{Z}_i) + \varepsilon_i$, for $i = 1, 2, \dots, n$.

4.3.1 Estimating the regression function

For a given data triplet $(\mathcal{X}, \mathcal{Z}, y)$ and a bandwidth h , we compute the discrepancy between $g(\cdot)$ and $\hat{g}(\cdot)$. We use the following Monte-Carlo scheme:

- Build 100 replications $[(\mathcal{X}_i^s, \mathcal{Z}_i^s, y_i^s)_{i=1, \dots, n}]_{s=1, \dots, 100}$;
- Compute 100 estimates $[g(\cdot) - \hat{g}^s(\cdot)]_{s=1, \dots, 100}$, where $\hat{g}^s(\cdot)$ is the functional NW estimator of the regression function computed over the s^{th} replication;

- Obtain the AMSE and AMSPE by averaging over 100 replications of MSE and MSPE.

Table 1 presents the AMSEs and AMSPEs for the estimated conditional mean, where the error density is estimated by a kernel error density with a global bandwidth and localized bandwidths. The functional partial linear model considers two function-valued random variables, $\mathcal{X}$ and $\mathcal{Z}$. In contrast, functional principal component regression and functional nonparametric regression models consider only the original functional curves, $\mathcal{X}$. Thus, the functional partial linear model produces improved estimation and forecast accuracies than functional principal component regression and functional nonparametric regression. It is also advantageous to estimate error density with localized bandwidths to achieve the best estimation and forecast accuracies of the regression function. An exception is in the skewbimodal error density, where the difference between the two bandwidth procedures is quite small. Among the three types of semi-metrics, the AMSEs are smaller for the two semi-metrics based on derivatives, compared with the semi-metric based on FPCA. Following the early work by Shang (2014b), we use three retained principal components that seem to be sufficient to capture the primary mode of variation in this example. The semi-metric based on the second derivative has the smallest AMSE and AMSPE among the three semi-metrics.

Table 1: Comparisons of the estimation and forecast accuracies of the regression function for the functional partial linear model, functional principal component regression, and functional nonparametric regression for the three choices of semi-metric. The underscore G and L represent a global bandwidth or localized bandwidths for estimating error density. The smallest errors are highlighted in bold.

	FPCR	Type of semi-metric								
		1 st derivative			2 nd derivative			FPCA _{Z=3}		
		FPLM _G	FPLM _L	FNP	FPLM _G	FPLM _L	FNP	FPLM _G	FPLM _L	FNP
<u>t_5</u>										
AMSE	1.148	1.039	0.993	5.382	1.032	0.969	5.039	1.164	1.097	5.713
AMSPE	1.332	1.361	1.285	5.247	1.329	1.235	5.055	1.541	1.452	5.484
<u>Skewunimodal</u>										
AMSE	1.600	1.237	1.216	5.847	1.148	1.136	5.353	1.337	1.320	6.182
AMSPE	1.783	1.366	1.371	5.746	1.313	1.297	5.383	1.551	1.519	5.980
<u>Skewbimodal</u>										
AMSE	1.235	1.010	1.006	5.450	0.971	0.973	5.017	1.130	1.116	5.771
AMSPE	1.416	1.213	1.199	5.334	1.144	1.145	5.041	1.363	1.353	5.569

Figure 2 presents the log marginal likelihoods (LMLs) for the semi-metrics in all three error densities considered. Again, the semi-metric based on the second derivative has the largest LMLs and is thus considered the optimal semi-metric in this example.

Figure 2: Comparisons of the LMLs in the functional partial linear model with a global bandwidth and localized bandwidths for the three choices of semi-metric. The semi-metric based on the first derivative is shown in white, the semi-metric based on the second derivative is shown in dark grey, and the semi-metric based on the FPCA is shown in light grey.

4.3.2 Estimating the error density

With a set of residuals and a fixed residual bandwidth, it is possible to apply a univariate kernel-density estimator and compute the discrepancy between $f(\varepsilon)$ and $\hat{f}(\varepsilon)$ by using the following Monte-Carlo scheme:

- Compute 100 replications of residuals $[y_i^s - \hat{g}_i^s(\cdot)]_{s=1, \dots, 100}$;
- Apply a univariate kernel density to estimate error density, where the residual bandwidth is estimated by the Bayesian method for 100 replications;
- Compute the MISE and KL divergence between the true error density $f^s(\varepsilon)$ and estimated error density $\hat{f}^s(\varepsilon)$ for $s = 1, 2, \dots, 100$;
- Obtain the AMISE and averaged KL divergence by averaging over 100 replications of MISE and KL divergence.

Table 2 presents AMISEs and averaged KL divergences for the kernel error density with bandwidths estimated by Bayesian methods with a global bandwidth and localized bandwidths. For the three error densities, the kernel error density with localized bandwidths produces smaller AMISEs and smaller averaged KL divergences than the kernel error density with a global bandwidth.

Table 2: AMISEs and averaged KL divergences for the error-density estimation in the functional partial linear model with a global bandwidth and localized bandwidths for the three choices of semi-metric. The underscore G and L represent a global bandwidth and localized bandwidths for estimating error density. The smallest errors are highlighted in bold.

$f(\varepsilon)$	Type of semi-metrics (G)			Type of semi-metrics (L)		
	1 st derivative	2 nd derivative	FPCA _{Z=3}	1 st derivative	2 nd derivative	FPCA _{Z=3}
<u>AMISE</u>						
t_5	0.0015	0.0015	0.0019	0.0008	0.0008	0.0008
skewunimodal	0.0088	0.0084	0.0088	0.0074	0.0073	0.0075
skewbimodal	0.0026	0.0024	0.0026	0.0024	0.0022	0.0026
<u>Averaged KL divergence</u>						
t_5	0.1427	0.1421	0.1957	0.1299	0.1272	0.1391
skewunimodal	0.6205	0.5486	0.6737	0.4996	0.4636	0.5695
skewbimodal	0.4247	0.3905	0.4736	0.2990	0.2756	0.4112

4.4 Diagnostic check

As a demonstration with one replication, we plot the MCMC sample paths of the parameters on the left panel of Figure 3 and the autocorrelation functions (ACFs) of these sample paths on the right panel of Figure 3. With the t_5 error density, these plots demonstrate that the sample paths are mixed reasonably well. Table 3 summarizes the ergodic averages, 95% Bayesian credible intervals (CIs), standard error (SE), batch mean SE, and simulation inefficiency factor (SIF) values. Considered previously by Kim et al. (1998) and Meyer and Yu (2000), the SIF can be interpreted as the number of draws needed to have iid observations. Based on the SIF values, we find that the simulation chain converges very well.

Table 3: MCMC results of the bandwidth estimation under the prior density of $IG(\alpha = 1, \beta = 0.05)$ with t_5 error density.

Prior density: $IG(\alpha = 1, \beta = 0.05)$			SE	Batch-mean SE	SIF
Parameter	Mean	95% Bayesian CIs			
h	0.4656	(0.2136, 0.7026)	0.1217	0.3175	6.80
b	0.5000	(0.3377, 0.7078)	0.0967	0.2300	5.65

Figure 3: MCMC sample paths and ACFs of the sample paths with t_5 error density.

4.5 Simulation of rough curves

In this simulation study, we consider the same functional form as given in (4), but add one extra variable $d_j \sim U(-0.1, 0.1)$ in the construction of curves. This functional form is taken from [Shang \(2014b, Section 4.2\)](#). Figure 4 presents the simulated curves for one replication, along with the first-order derivative of the curves.

Figure 4: 100 simulated rough curves.

From the smooth curves presented in Figure 2, we found the difference in LML between a global bandwidth and localized bandwidths to be quite small for the functional partial linear model. Therefore, in Figure 5, we report only the LMLs for the functional partial linear model with a global bandwidth. The LMLs are computed for the three semi-metrics in all three error densities considered, and the overall optimal semi-metric is determined based on 100 replications. With the median LML, the semi-metric based on the first derivative has the largest LML for the t_5 error density, followed closely by the semi-metric based on the FPCA with three retained components. In the skewunimodal and skewbimodal error densities, the semi-metric based on the FPCA gives the largest LML and is thus considered the optimal semi-metric in these two error densities.

Figure 5: Comparisons of the LMLs in the functional partial linear model with a global bandwidth for the three choices of semi-metric. The semi-metric based on the first derivative is shown in white, the semi-metric based on the second derivative is shown in dark grey, and the semi-metric based on the FPCA is shown in light grey.

5 Spectrometric data analysis

This dataset focuses on the prediction of the fat content in meat samples based on near-infrared (NIR) absorbance spectra. The dataset was obtained from <http://lib.stat.cmu.edu/datasets/tecator> and has been studied by many researchers (e.g., Ferraty and Vieu, 2006; Aneiros-Pérez and Vieu, 2006). Each food sample contains finely chopped pure meat with a percentage of fat content. For each unit i (among 215 pieces of finely chopped meat), we observe one spectrometric curve,

denoted by $\mathcal{X}_i$. The spectrometric curve measures the absorbance at a grid of 100 wavelengths, denoted by $\mathcal{X}_i = [\mathcal{X}_i(t_1), \dots, \mathcal{X}_i(t_{100})]$. For each unit i , we consider the first-order derivative of $\mathcal{X}_i$ and observe the fat content $y \in R$ by analytical chemical processing. Given a new spectrometric curve $(\mathcal{X}_{\text{new}}, \mathcal{X}'_{\text{new}})$, we aim to predict the corresponding fat content y_{new} . A graphical display of the spectrometric curves is presented in Figure 6.

Figure 6: Graphical display of spectrometric curves.

To assess the point forecast accuracy of the functional partial linear model, we split the original samples into two subsamples (see also [Ferraty and Vieu, 2006](#), p.105). The first subsample is named the ‘learning sample’, which contains the first 160 units $\{\mathcal{X}_i, \mathcal{Z}_i, y_i\}_{i=1,2,\dots,160}$. The second subsample is named the ‘testing sample’, which contains the last 55 units $\{\mathcal{X}_l, \mathcal{Z}_l, y_l\}_{l=161,\dots,215}$. The learning sample allows us to build the regression model with the optimal smoothing parameter. The testing sample allows us to evaluate the prediction accuracy.

To measure estimation and prediction accuracies, we consider the root mean squared error (RMSE) and root mean squared prediction error (RMSPE). The errors are expressed as

$$\text{RMSE} = \sqrt{\frac{1}{160} \sum_{\omega=1}^{160} (y_{\omega} - \hat{y}_{\omega})^2},$$

$$\text{RMSPE} = \sqrt{\frac{1}{55} \sum_{\delta=161}^{215} (y_{\delta} - \hat{y}_{\delta})^2}.$$

For comparison, we also consider the functional principal component regression examined by [Reiss and Ogden \(2007\)](#) and the functional nonparametric regression examined by [Ferraty and Vieu \(2006\)](#).

As presented in Table 4, the functional partial linear model produces the smallest RMSE and RMSPE, in particular when the semi-metric is chosen based on the second derivative. The semi-metric based on the second derivative not only has the smallest RMSE and RMSPE, but it also has the largest LML.

Table 4: Estimation and forecast accuracies for the functional partial linear model, functional principal component regression, and functional nonparametric regression. The smallest error measures and largest LML are shown in bold.

Criterion	FPCR	Type of semi-metrics					
		1 st derivative		2 nd derivative		FPCA _{Z=3}	
		FPLM	FNP	FPLM	FNP	FPLM	FNP
RMSE	8.0832	2.4281	4.6014	1.5993	2.2899	7.3236	10.8779
RMSPE	9.1156	1.9191	4.7755	1.4075	1.9429	8.3378	11.2603
LML		-360.21		-291.37		-524.48	

To verify the model adequacy, we present several scatter plots of the holdout responses against the predicted responses in Figure 7. Among the three regression models, the functional partial linear model provides the best prediction accuracy.

Figure 7: Diagnostic check of the three functional regression models.

The original functional data being iid allows sampling with replacement to obtain a set of bootstrap samples. For each bootstrap sample, we use the first 160 pairs of data for estimation and the remaining 55 pairs of data for assessing the prediction accuracy. By repeating this procedure 100 times, we obtain the averaged RMSEs, and RMSPEs presented in Table 5. The functional partial linear model produces more accurate forecasts than functional principal component regression and functional nonparametric regression. In particular, the functional partial linear model produces the most accurate forecasts when the semi-metric is based on the second derivative. This optimal selection of semi-metric is further confirmed by the boxplot of LMLs presented in Figure 8.

Table 5: Estimation and forecast accuracy averaged over 100 replications for the functional partial linear model, functional principal component regression, and functional nonparametric regression. The smallest error measures are shown in bold.

Criterion	FPCR	Type of semi-metrics					
		1 st derivative		2 nd derivative		FPCA _{Z=3}	
		FPLM	FNP	FPLM	FNP	FPLM	FNP
RMSE	8.1011	2.2858	4.6449	1.5269	2.2343	7.3675	10.7080
RMSPE	8.2234	2.3100	4.5655	1.6259	2.1947	7.8725	10.8047

Figure 8: Comparisons of the LMLs in the functional partial linear model with a global bandwidth for the three choices of semi-metric.

There has been rapid development in functional regression models. Here, we consider a large set of functional regression models collected in Goldsmith and Scheipl (2014) and Ferraty and Vieu (2011). These models include: (1) restricted maximum likelihood-based functional linear model with a locally adaptive penalty (Cardot et al., 2003); (2) penalized functional regression (Goldsmith et al., 2011); (3) functional principal component regression on the first few functional principal components (Reiss and Ogden, 2007); (4) linear model on the first K functional principal components, where optimal K is estimated by 20-fold bootstrap (Ramsay and Silverman, 2005); (5) REML-based single-index signal regression with locally adaptive penalty (Wood, 2011); (6) cross-validation based single-index signal regression (Eilers et al., 2009); (7) penalized partial least squares (Krämer et al., 2008); (8) least absolute shrinkage and selection operator penalized linear model on the first few functional principal components (Friedman et al., 2010); (9) functional

nonparametric regression with NW estimator for estimating conditional mean (Ferraty and Vieu, 2006); (10) functional nonparametric regression with NW estimator for estimating conditional median (Laksaci et al., 2011); (11) functional nonparametric regression with NW estimator for estimating conditional mode (Ferraty et al., 2005); (12) functional nonparametric regression with k nearest neighbor estimator (Burba et al., 2009); (13) functional nonparametric regression with most-predictive design points (Ferraty et al., 2010a); (14) functional partial linear model.

The boxplot of RMSPEs is presented in Figure 9. The functional partial linear model with the semi-metric and bandwidth selected by the Bayesian method has the smallest RMSPE among all methods considered.

Figure 9: Comparisons of forecast accuracy for the 14 functional regression models listed above.

A by-product of the Bayesian method is that it is possible to compute the pointwise prediction interval nonparametrically. To make this computation, we first compute the cumulative density function of the error density, over a set of grid points within a range, such as between -10 and 10. Then, we take the inverse of the cumulative density function and find two grid points that are closest to the 10% and 90% quantiles. The 80% pointwise prediction interval of a holdout sample is obtained by adding the two grid points to a point forecast. In Figure 10, the holdout fat content in percentage is presented as diamond-shaped dots, the point forecasts of the fat content are presented as round circles, and the 80% pointwise prediction intervals are presented as vertical bars. At the nominal coverage probability of 80%, the empirical coverage probability is 87%; at the nominal coverage probability of 50%, the empirical coverage probability is 51%.

Figure 10: A plot of predicted fat content in percentage and their 80% pointwise prediction intervals. The actual holdout data are presented as diamond-shaped dots; the point forecasts of the fat content are presented as circles; the 80% pointwise prediction intervals are presented as vertical bars.

6 Conclusion

We propose a Bayesian method to estimate optimal bandwidths in a functional partial linear model with homoscedastic errors and unknown error density. As a by-product of the MCMC sampling algorithm, it allows us to compute log marginal likelihood for a semi-metric. The optimal semi-metric is the one that provides the smallest log marginal likelihood, as evident from the results in Table 4. Given that no closed-form expression for our bandwidth estimator exists, establishing the mathematical properties, such as the asymptotic optimality of Shibata (1981), of the bandwidth estimator has been difficult. However, we have developed an approximate solution to the bandwidth estimator through MCMC. As a by-product of the Bayesian method, a prediction interval can be obtained, and marginal likelihood can also be used to determine the optimal choice of semi-metric among a set of semi-metrics.

Through a series of simulation studies, we find the functional partial linear model has better estimation and forecast accuracies than functional principal component regression and functional nonparametric regression. For a set of smooth curves, we further confirm the optimality of the semi-metric based on the second derivative; for a set of rough curves, we affirm the overall optimal semi-metric is the semi-metric based on the FPCA, followed closely by the semi-metric based on

the first derivative. The kernel density estimator with localized bandwidths performs similarly to the kernel density estimator with a global bandwidth for estimating regression function and outperforms the one with a global bandwidth for estimating error density.

Using the spectroscopy dataset, the functional partial linear model produces the smallest forecast error of some commonly used functional regression models. The Bayesian method not only allows determination of the optimal semi-metric via marginal likelihood, but it also allows the nonparametric construction of pointwise prediction intervals for measuring prediction uncertainty.

There are many ways in which the Bayesian method and functional partial linear model can be extended, and we briefly outline five:

- 1) Combine different semi-metrics with weights based on their marginal likelihoods; this leads to the idea of ensemble forecasting.
- 2) Consider other functional regression estimators, such as the functional local-linear estimator of [Benhenni et al. \(2007\)](#) or the k -nearest neighbor estimator of [Burba et al. \(2009\)](#) – the functional local-linear estimator can improve the estimation accuracy of the regression function by using a higher-order kernel function; the k -nearest neighbor estimator considers the local structure of the data and gives better forecasts when the functional data are heterogeneously concentrated.
- 3) Extend to a functional partial linear model with heterogeneous errors; the covariate-dependent variance can be modeled by another kernel-density estimator (e.g., [Chen et al., 2009](#)).
- 4) Extend to a functional partial linear model with autoregressive errors (e.g., [You and Zhou, 2006](#)).
- 5) Consider a J test of [Davidson and MacKinnon \(1981\)](#) for selecting non-nested functional regression models.

Acknowledgments

The author would like to thank a reviewer for insightful comments and suggestions, which substantially improve the early version of the paper. The author acknowledges a faculty research grant from the College of Business and Economics at the Australian National University.

References

- Ahmad, I. and Li, Q. (1997). Testing symmetry of an unknown density function by kernel method. *Journal of Nonparametric Statistics*, 7(3):279–293.
- Akritis, M. G. and Van Keilegom, I. (2001). Non-parametric estimation of the residual distribution. *Scandinavian Journal of Statistics*, 28(3):549–567.
- Aneiros-Pérez, G., Ferraty, F., and Vieu, P. (2015). Variable selection in partial linear regression with functional covariate. *Statistics: A Journal of Theoretical and Applied Statistics*, 49(6):1322–1347.
- Aneiros-Pérez, G. and Vieu, P. (2006). Semi-functional partial linear regression. *Statistics and Probability Letters*, 76(11):1102–1110.
- Aneiros-Pérez, G. and Vieu, P. (2015). Partial linear modelling with multi-functional covariates. *Computational Statistics*, 30(3):647–671.
- Barrientos-Marin, J., Ferraty, F., and Vieu, P. (2010). Locally modelled regression and functional data. *Journal of Nonparametric Statistics*, 22(5):617–632.
- Benhenni, K., Ferraty, F., Rachdi, M., and Vieu, P. (2007). Local smoothing regression with functional data. *Computational Statistics*, 22(3):353–369.
- Berlinet, A., Elamine, A., and Mas, A. (2011). Local linear regression for functional data. *Annals of the Institute of Statistical Mathematics*, 63(5):1047–1075.
- Boj, E., Delicado, P., and Fortiana, J. (2010). Distance-based local linear regression for functional predictors. *Computational Statistics & Data Analysis*, 54(2):429–437.
- Bowman, A. W. (1984). An alternative method of cross-validation for the smoothing of density estimates. *Biometrika*, 71(2):353–360.
- Burba, F., Ferraty, F., and Vieu, P. (2009). k -nearest neighbour method in functional nonparametric regression. *Journal of Nonparametric Statistics*, 21(4):453–469.
- Cardot, H., Ferraty, F., and Sarda, P. (2003). Spline estimators for the functional linear model. *Statistica Sinica*, 13:571–591.
- Chen, H., Smyth, R., and Zhang, X. (2017). A Bayesian sampling approach to measuring the price responsiveness of gasoline demand using a constrained partially linear model. *Energy Economics*, 67:346–354.
- Chen, L.-H., Cheng, M.-Y., and Peng, L. (2009). Conditional variance estimation in heteroscedastic regression models. *Journal of Statistical Planning and Inference*, 139(2):236–245.

- Cheng, F. and Sun, S. (2008). A goodness-of-fit test of the errors in nonlinear autoregressive time series models. *Statistics & Probability Letters*, 78(1):50–59.
- Chernozhukov, V. and Hong, H. (2003). An MCMC approach to classical estimation. *Journal of Econometrics*, 115(2):293–346.
- Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association: Theory and Method*, 90(432):1313–1321.
- Davidson, R. and MacKinnon, J. G. (1981). Several tests for model specification in the presence of alternative hypotheses. *Econometrica*, 49(3):781–793.
- de Boor, C. (2001). *A Practical Guide to Splines*, volume 27 of *Applied Mathematical Sciences*. Springer, revised edition.
- Efromovich, S. (2005). Estimation of the density of regression errors. *The Annals of Statistics*, 33(5):2194–2227.
- Eilers, P. H. C., Li, B., and Marx, B. D. (2009). Multivariate calibration with single-index signal regression. *Chemometrics and Intelligent Laboratory Systems*, 96(2):196–202.
- Escanciano, J. C. and Jacho-Chávez, D. T. (2012). $\sqrt{n}$ uniformly consistent density estimation in nonparametric regression models. *Journal of Econometrics*, 167(2):305–316.
- Ferraty, F., Hall, P., and Vieu, P. (2010a). Most-predictive design points for functional data predictors. *Biometrika*, 97(4):807–824.
- Ferraty, F., Laksaci, A., and Vieu, P. (2005). Functional time series prediction via conditional mode estimation. *C. R. Acad. Sci. Paris, Ser. I*, 340:389–392.
- Ferraty, F., Van Keilegom, I., and Vieu, P. (2010b). On the validity of the bootstrap in non-parametric functional regression. *Scandinavian Journal of Statistics*, 37(2):286–306.
- Ferraty, F. and Vieu, P. (2002). The functional nonparametric model and application to spectrometric data. *Computational Statistics*, 17(4):545–564.
- Ferraty, F. and Vieu, P. (2006). *Nonparametric Functional Data Analysis: Theory and Practice*. New York: Springer.
- Ferraty, F. and Vieu, P. (2009). Additive prediction and boosting for functional data. *Computational Statistics & Data Analysis*, 53(4):1400–1413.
- Ferraty, F. and Vieu, P. (2011). Richesse et complexité des données fonctionnelles. *Revue MODULAD*, 43:25–43.

- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1).
- Garthwaite, P. H., Fan, Y., and Sisson, S. A. (2016). Adaptive optimal scaling of Metropolis-Hastings algorithms using the Robbins-Monro process. *Communications in Statistics – Theory and Methods*, 45(17):5098–5111.
- Gelfand, A. E. and Dey, D. K. (1994). Bayesian model choice: Asymptotics and exact calculations. *Journal of the Royal Statistical Society, Series B*, 56(3):501–514.
- Goldsmith, J., Bobb, J., Crainiceanu, C. M., Caffo, B., and Reich, D. (2011). Penalized functional regression. *Journal of Computational and Graphical Statistics*, 20(4):830–851.
- Goldsmith, J. and Scheipl, F. (2014). Estimator selection and combination in scalar-on-function regression. *Computational Statistics & Data Analysis*, 70:362–372.
- Gourieroux, C., Monfort, A., and Trognon, A. (1984). Pseudo maximum likelihood methods: Theory. *Econometrica*, 52(3):681–700.
- Goutis, C. (1998). Second-derivative functional regression with applications to near infra-red spectroscopy. *Journal of the Royal Statistical Society, Series B*, 60(1):103–114.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association: Review Paper*, 90(430):773–795.
- Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393.
- Krämer, N., Boulesteix, A.-L., and Tutz, G. (2008). Penalized partial least squares with application to B-spline transformations and functional data. *Chemometrics and Intelligent Laboratory Systems*, 94(1):60–69.
- Kullback, S. (1959). *Information Theory and Statistics*. Wiley, New York.
- Laksaci, A., Lemdani, M., and Saïd, E. O. (2011). Asymptotic results for an L^1 -norm kernel estimator of the conditional quantile for functional dependent data with application to climatology. *Sankhyā: The Indian Journal of Statistics*, 73-A(1):125–141.
- Lian, H. (2011). Functional partial linear model. *Journal of Nonparametric Statistics*, 23(1):115–128.
- Ling, N., Aneiros, G., and Vieu, P. (2019). k NN estimation in functional partial linear modeling. *Statistical Papers*, in press.

- Marron, J. S. and Wand, M. P. (1992). Exact mean integrated squared error. *Annals of Statistics*, 20(2):712–736.
- Mas, A. and Pumo, B. (2009). Functional linear regression with derivatives. *Journal of Nonparametric Statistics*, 21(1):19–40.
- Meyer, R. and Yu, J. (2000). BUGS for a Bayesian analysis of stochastic volatility models. *Econometrics Journal*, 3(2):198–215.
- Muhsal, B. and Neumeyer, N. (2010). A note on residual-based empirical likelihood kernel density estimation. *Electronic Journal of Statistics*, 4:1386–1401.
- Müller, U. (2013). Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix. *Econometrica*, 81(5):1805–1849.
- Neumeyer, N. and Dette, H. (2007). Testing for symmetric error distribution in nonparametric regression models. *Statistica Sinica*, 17(2):775–795.
- Rachdi, M. and Vieu, P. (2007). Nonparametric regression for functional data: Automatic smoothing parameter selection. *Journal of Statistical Planning and Inference*, 137(9):2784–2801.
- Ramsay, J. and Silverman, B. (2005). *Functional Data Analysis*. Springer, New York, 2nd edition.
- Reiss, P. T., Goldsmith, J., Shang, H. L., and Ogden, R. T. (2017). Methods for scalar-on-function regression. *International Statistical Review*, 85(2):228–249.
- Reiss, P. T. and Ogden, R. T. (2007). Functional principal component regression and functional partial least squares. *Journal of the American Statistical Association*, 102(479):984–996.
- Roberts, G. O. and Rosenthal, J. S. (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367.
- Shang, H. L. (2013). Bayesian bandwidth estimation for a nonparametric functional regression model with unknown error density. *Computational Statistics & Data Analysis*, 67:185–198.
- Shang, H. L. (2014a). Bayesian bandwidth estimation for a nonparametric functional regression model with mixed types of regressors and unknown error density. *Journal of Nonparametric Statistics*, 26(3):599–615.
- Shang, H. L. (2014b). Bayesian bandwidth estimation for a semi-functional partial linear regression model with unknown error density. *Computational Statistics*, 29(3-4):829–848.
- Shang, H. L. (2016). A Bayesian approach for determining the optimal semi-metric and bandwidth in scalar-on-function quantile regression with unknown error density and dependent functional data. *Journal of Multivariate Analysis*, 146:95–104.

- Shang, H. L. (2019). Estimation of a functional single index model with dependent errors and unknown error density. *Communications in Statistics – Simulation and Computation*, in press.
- Shibata, R. (1981). An optimal selection of regression variables. *Biometrika*, 68(1):45–54.
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B*, 73(1):3–36.
- You, J. and Zhou, X. (2006). Statistical inference in a panel data semiparametric regression model with serially correlated errors. *Journal of Multivariate Analysis*, 97(4):844–873.
- Zhang, X. and King, M. L. (2011). Bayesian semiparametric GARCH models. Working paper 24, Monash University. URL: <https://www.monash.edu/business/econometrics-and-business-statistics/research/publications/ebs/wp24-11.pdf>.
- Zhang, X., King, M. L., and Shang, H. L. (2011). Bayesian estimation of bandwidths for a non-parametric regression model with a flexible error density. Working paper 10, Monash University. URL: <https://www.monash.edu/business/econometrics-and-business-statistics/research/publications/ebs/wp10-11.pdf>.