

ADDITIVE MODELS WITH PREDICTORS SUBJECT TO MEASUREMENT ERROR

BHASWATI GANGULI¹, JOHN STAUDENMAYER^{2*} AND M.P. WAND³

*Indian Institute of Management, University of Massachusetts
and the University of New South Wales*

Summary

This paper develops a likelihood-based method for fitting additive models in the presence of measurement error. It formulates the additive model using the linear mixed model representation of penalized splines. In the presence of a structural measurement error model, the resulting likelihood involves intractable integrals, and a Monte Carlo expectation maximization strategy is developed for obtaining estimates. The method's performance is illustrated with a simulation study.

Key words: Metropolis–Hastings; mixed models; Monte Carlo expectation maximization; nested EM; penalized splines; restricted maximum likelihood.

1. Introduction

Additive models are multi-predictor regression models that assume little more than additivity. Their advantage over ordinary regression models is that the functional form of continuous predictors is not restricted to be in a parametric family and therefore is more flexible.

For regression models in general it is well established that measurement error in a predictor distorts its estimated relationship to the response (e.g. Fuller, 1987 Chapter 1). A number of studies have devised schemes for redressing this situation for parametric regression models (e.g. Carroll, Ruppert & Stefanski, 1995; Spiegelman, Rosner & Logan, 2000) and single-predictor non-parametric models (Fan & Truong, 1993; Carroll, Maca & Ruppert, 1999; Berry, Carroll & Ruppert, 2002; Staudenmayer & Ruppert, 2004). However, to our knowledge, there is no available methodology for handling measurement error in additive models.

Carroll *et al.* (1999) formulated and compared several estimators that have small, persistent, asymptotic biases but perform quite well in simulations. These estimators address the covariate measurement error with either simulation extrapolation or extended regression calibration, and they use penalized splines or local linear regression for the non-parametric estimation. Berry *et al.* (2002) proposed Bayesian estimation methods for this problem, and their simulations indicate improvements over the methods of Carroll *et al.* (1999).

Received August 2003; revised January 2004; accepted January 2004.

* Author to whom correspondence should be addressed.

¹ Indian Institute of Management, Calcutta, India.

² Dept of Mathematics and Statistics, Lederle Graduate Research Tower, University of Massachusetts, Amherst MA, 01003-9305, USA. e-mail: jstauden@math.umass.edu

³ Dept of Statistics, The University of New South Wales, Sydney NSW 2052, Australia.

Acknowledgments. This work was partially supported by NIH grant T32 ES07142-18 and NSF-DMS 0306227. The authors thank the editor, an associate editor, and an anonymous referee for helpful suggestions and comments.

In this paper we exploit connections between penalized spline smoothing and likelihood-based mixed models (e.g. Ruppert, Wand & Carroll, 2003). With this approach the additive model extension is straightforward, so we work at this level of generality. The measurement error structure involves adjustment of the likelihood, although the resulting likelihood function is intractable. We overcome this problem using a nested Monte Carlo expectation maximization (MCEM) algorithm (Wei & Tanner, 1990; Smyth, 1996; van Dyk, 2000).

Although our paper focuses on smoothing, our methods generally apply to linear mixed models of the form $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{b} + \boldsymbol{\varepsilon}$ where \mathbf{b} and $\boldsymbol{\varepsilon}$ are multivariate normally distributed random vectors and the matrices \mathbf{X} and \mathbf{Z} both have mismeasured components. We believe that these are the first methods to generally address the case when the random effect's covariates (\mathbf{Z}) have measurement error. The case when only \mathbf{X} is allowed to be mismeasured is addressed in Buonaccorsi, Demidenko & Tosteson (2000) for instance.

In Section 2 we give the model formulation and likelihood function, and review briefly the connections between penalized spline smoothing and likelihood based mixed models. Section 3 describes the estimation scheme. Section 4 illustrates our approach, through a simulation study and an application of the method to air pollution or mortality data. We conclude with some discussion and mention of extensions in Section 5.

2. Formulation

Let \mathbf{y} be an $n \times 1$ vector of responses, and $\mathbf{x}_1, \dots, \mathbf{x}_d$ be $n \times 1$ vectors of predictors for which the Gaussian additive model,

$$\mathbf{y} = \beta_0 \mathbf{1} + \sum_{j=1}^d m_j(\mathbf{x}_j) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \stackrel{d}{=} N_n(\mathbf{0}, \sigma_\varepsilon^2 \mathbf{I}),$$

is reasonably assumed. The m_j are general smooth functions, modelled through the mixed model formulation of penalized regression splines: $m_j(\mathbf{x}_j) = \beta_j \mathbf{x}_j + \sum_{k=1}^{K_j} b_k^j (\mathbf{x}_j - \kappa_k^j \mathbf{1})_+$, $\mathbf{b}_j \stackrel{d}{=} N_{K_j}(\mathbf{0}, \sigma_{b_j}^2 \mathbf{I})$, independent ($j = 1, \dots, d$), where $\mathbf{b}_j = (b_1^j, \dots, b_{K_j}^j)$ (e.g. Eilers & Marx, 1996; Wand, 2003). Note that $\sigma_{b_j}^2 = 0$ corresponds to linearity of m_j . A fixed number of knots are placed on a grid in each dimension $j = 1, \dots, d$. See Ruppert (2002) for a discussion of knot (κ_k^j) selection.

As an aside, a link between penalized splines and mixed models is reviewed in Ruppert *et al.* (2003 Section 4.9). In the context of a scatterplot smoother ($d = 1$), for ease of explanation, the penalized spline fitting criterion finds $\boldsymbol{\beta}$ and \mathbf{b} to minimize

$$\frac{1}{\sigma_\varepsilon^2} \sum_{i=1}^n \left(y_i - \beta_0 - \beta_1 x_i - \sum_{k=1}^K b_k (x_i - \kappa_k)_+ \right)^2 + \frac{\lambda^2}{\sigma_\varepsilon^2} \sum_{k=1}^K b_k^2.$$

When \mathbf{b} is treated as a random variable with mean $\mathbf{0}$ and covariance $\sigma_u^2 \mathbf{I}$ where $\sigma_u^2 = \sigma_\varepsilon^2 / \lambda^2$, this criterion is equivalent to the best linear unbiased predictor (BLUP) criterion in a linear mixed model. Additional connections between mixed models and splines have been explored and exploited by a number of other authors including Verbyla *et al.* (1999) who drew upon Green & Silverman (1994 Section 3.8) and Speed (1991).

In our situation, some of the \mathbf{x}_j are not directly observable, and instead we observe

$$\mathbf{w}_j = \mathbf{x}_j + \mathbf{u}_j, \quad \text{where } \mathbf{u}_j \stackrel{d}{=} N_n(\mathbf{0}, \sigma_{u_j}^2 \mathbf{I}) \text{ independent } (j = 1, \dots, d),$$

where the $\sigma_{u_j}^2$ are known ($\sigma_{u_j}^2 = 0$ corresponds to \mathbf{x}_j being observed). Let $f(\mathbf{x}_1, \dots, \mathbf{x}_d; \boldsymbol{\theta})$ denote the joint density of the predictors, where f is determined by the parameter vector $\boldsymbol{\theta}$; this is sometimes called a structural assumption. We also assume that the dependence of \mathbf{y} on the predictors is non-differential:

$$(\mathbf{y} \mid \mathbf{x}_1, \dots, \mathbf{x}_d, \mathbf{w}_1, \dots, \mathbf{w}_d) \stackrel{d}{=} (\mathbf{y} \mid \mathbf{x}_1, \dots, \mathbf{x}_d).$$

With these specifications,

$$(\mathbf{y} \mid \mathbf{x}_1, \dots, \mathbf{x}_d) \stackrel{d}{=} N_n(\mathbf{X}\boldsymbol{\beta}, \mathbf{V}_x),$$

where
$$\mathbf{V}_x = \sum_{j=1}^d \sigma_{b_j}^2 \mathbf{Z}_j \mathbf{Z}_j^\top + \sigma_\varepsilon^2 \mathbf{I}, \quad \boldsymbol{\beta} = (\beta_0, \dots, \beta_d),$$

$$\mathbf{X} = [\mathbf{1} \ \mathbf{x}_1 \ \dots \ \mathbf{x}_d] \quad \text{and} \quad \mathbf{Z}_j = [(\mathbf{x}_j - \kappa_1^j \mathbf{1})_+ \ \dots \ (\mathbf{x}_j - \kappa_{K_j}^j \mathbf{1})_+].$$

We also use the notation

$$\mathbf{x} = [\mathbf{x}_1 \ \dots \ \mathbf{x}_d], \quad \mathbf{b} = [\mathbf{b}_1 \ \dots \ \mathbf{b}_d], \quad \mathbf{w} = [\mathbf{w}_1 \ \dots \ \mathbf{w}_d] \quad \text{and} \quad \mathbf{Z} = [\mathbf{Z}_1 \ \dots \ \mathbf{Z}_d].$$

Further, let \mathbf{x}_{-j} represent \mathbf{x} without its j th component. Finally, let $\sigma_b^2 = (\sigma_{b_1}^2, \dots, \sigma_{b_d}^2)$, $\sigma_u^2 = (\sigma_{u_1}^2, \dots, \sigma_{u_d}^2)$, and $\boldsymbol{\psi} = (\boldsymbol{\theta}, \boldsymbol{\beta}, \sigma_b^2, \sigma_\varepsilon^2)$. The likelihood is

$$\begin{aligned} L(\boldsymbol{\psi}) &= f(\mathbf{y}, \mathbf{w}; \boldsymbol{\psi}) \\ &= \int_{\mathbb{R}^n} \dots \int_{\mathbb{R}^n} f(\mathbf{y} \mid \mathbf{x}; \boldsymbol{\beta}, \sigma_b^2, \sigma_\varepsilon^2) f(\mathbf{w} \mid \mathbf{x}; \sigma_u^2) f(\mathbf{x}_1, \dots, \mathbf{x}_d; \boldsymbol{\theta}) \, d\mathbf{x} \\ &= \int_{\mathbb{R}^n} \dots \int_{\mathbb{R}^n} \varphi_{\mathbf{V}_x}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) f(\mathbf{w} \mid \mathbf{x}; \sigma_u^2) f(\mathbf{x}; \boldsymbol{\theta}) \, d\mathbf{x}, \end{aligned}$$

where
$$\varphi_{\boldsymbol{\Sigma}}(\mathbf{x}) = (2\pi)^{-n/2} |\boldsymbol{\Sigma}|^{-1/2} \exp(-\frac{1}{2} \mathbf{x}^\top \boldsymbol{\Sigma}^{-1} \mathbf{x})$$

denotes the $N_n(\mathbf{0}, \boldsymbol{\Sigma})$ density.

To better appreciate the complexity that measurement error brings to this likelihood, suppose that

$$\mathbf{x}_j \stackrel{d}{=} N(\mu_j \mathbf{1}, \sigma_j^2 \mathbf{I}), \quad \text{independent } (j = 1, \dots, d).$$

Then the likelihood is

$$L(\boldsymbol{\psi}) = \int_{\mathbb{R}^n} \dots \int_{\mathbb{R}^n} \varphi_{\mathbf{V}_x}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) \prod_{j=1}^d \varphi_{\sigma_{u_j}^2 \mathbf{I}}(\mathbf{w}_j - \mathbf{x}_j) \prod_{j=1}^d \varphi_{\sigma_j^2 \mathbf{I}}(\mathbf{x}_j - \mu_j \mathbf{1}) \, d\mathbf{x},$$

and, despite the fact that everything is normal, the integrals cannot be reduced due to the presence of $\mathbf{x}_1, \dots, \mathbf{x}_d$ in \mathbf{V}_x . While a more flexible structural model is often desirable (e.g. Carroll *et al.*, 1984), we present the case of a normal structural model to illustrate the method.

This problem appears to be similar to the one encountered in maximum likelihood estimation in the generalized linear mixed model (e.g. McCulloch & Searle, 2001 Chapter 8). We emphasize that the integrals here are a result of the measurement error and not the mixed model. This observation highlights an interesting connection between the somewhat separate statistical subfields of mixed models and measurement error.

Apart from maximization of $L(\boldsymbol{\psi})$ for estimation of the parameters in $\boldsymbol{\psi}$, prediction of \mathbf{b} is required for estimation of the function $m_j(\cdot)$, $1 \leq j \leq K$. The best predictor (BP) of \mathbf{b} is

$$\hat{\mathbf{b}} = E(\mathbf{b} \mid \mathbf{y}, \mathbf{x}).$$

A commonly used estimator is the empirical best predictor (eBP), $\hat{\mathbf{b}}^j$, the BP calculated with σ_b^2 and σ_ε^2 fixed at their current estimates, $\hat{\sigma}_b^2$ and $\hat{\sigma}_\varepsilon^2$, which we estimated by REML. Algebraically, the eBP is

$$\hat{\mathbf{b}} = (\mathbf{Z}^\top \mathbf{Z} + \mathbf{D})^{-1} \mathbf{Z}^\top (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}), \quad \text{where } \mathbf{D} = \hat{\sigma}_\varepsilon^2 \text{blockdiag}_{1 \leq j \leq d} \left(\frac{1}{\hat{\sigma}_{b_j}^2} \mathbf{I}_{K_j} \right).$$

In the case of normal $\boldsymbol{\varepsilon}$ and \mathbf{b} , the eBP is also unbiased. As a result of this and its form, the eBP is also the empirical best linear unbiased estimator (eBLUP) in this case (e.g. McCulloch & Searle, 2001 Chapter 9).

3. Estimation

As discussed in Section 2, since $L(\boldsymbol{\psi})$ involves intractable integrals, direct determination of $\hat{\boldsymbol{\psi}}$ to maximize $L(\boldsymbol{\psi})$ is very difficult. In the somewhat similar case of the generalized linear mixed model (e.g. McCulloch & Searle, 2001 Chapters 8, 10), a tractable approach uses the nested Monte Carlo Expectation Maximization algorithm (nested MC–EM algorithm). Section 3.2 introduces the nested MC–EM algorithm and describes how to generate the required Monte Carlo sample. Section 3.3 then enumerates our algorithm in detail.

3.1. EM algorithm

Denote the observed data as $\mathbf{Y}_{\text{obs}} = (\mathbf{y}, \mathbf{w})$ with log-likelihood $\ell_{\text{obs}}(\boldsymbol{\psi})$. Let $\mathbf{Y}_{\text{comp}} = (\mathbf{y}, \mathbf{w}, \mathbf{x}, \mathbf{b})$ be an augmentation of \mathbf{Y}_{obs} , which we refer to as the complete data; let the associated log-likelihood be denoted by $\ell_{\text{comp}}(\boldsymbol{\psi})$. The EM algorithm maximizes $\ell_{\text{obs}}(\boldsymbol{\psi})$ in $\boldsymbol{\psi}$ using an iterative process comprising two steps. Let the current estimates of the parameters be $\boldsymbol{\psi}^{(t)}$ where (t) indexes iterations. The expectation or E-step consists of computing the ‘best guess’ (conditional expectation with the expectation indexed by the most recent parameter estimates) of the complete data log-likelihood given the observed data,

$$Q(\boldsymbol{\psi}; \boldsymbol{\psi}^{(t)}) = E(\ell_{\text{comp}}(\boldsymbol{\psi}) \mid \mathbf{Y}_{\text{obs}}; \boldsymbol{\psi}^{(t)}).$$

Next, the maximization or M-step consists of calculating

$$\boldsymbol{\psi}^{(t+1)} = \underset{\boldsymbol{\psi}}{\operatorname{argmax}} Q(\boldsymbol{\psi}; \boldsymbol{\psi}^{(t)}).$$

When possible, the procedure can be streamlined by computing only the conditional expectations of the sufficient statistics for the complete data problem. Under regularity conditions, $\boldsymbol{\psi}^{(t)} \rightarrow \hat{\boldsymbol{\psi}}_{\text{MLE}}$ as $t \rightarrow \infty$.

3.2. MCEM algorithm and nesting

One extension to the EM algorithm is Monte Carlo (MC) estimation of the E-step (Wei & Tanner, 1990). This extension replaces analytic (or numerical) determination of the required expectation with a Monte Carlo estimate. We use this method since our E-step does not have a closed form.

The E-step could be implemented by sampling from the distribution of $(\mathbf{x}, \mathbf{b} \mid \mathbf{y}, \mathbf{w}; \boldsymbol{\psi}^{(t)})$, using a Gibbs sampler that alternates between samples from $(\mathbf{b} \mid \mathbf{x}, \mathbf{y}, \mathbf{w}; \boldsymbol{\psi}^{(t)})$ (multivariate normal) and $(\mathbf{x}_j \mid \mathbf{b}, \mathbf{x}_{-j}, \mathbf{y}, \mathbf{w}; \boldsymbol{\psi}^{(t)})$, $j = 1, \dots, d$ (see Fact 1 below), but this approach would probably be quite inefficient. Smyth (1996) and van Dyk (2000) suggest a more efficient approach called the nested EM algorithm. The idea behind that approach is to alternate one of the Gibbs chain elements with a ‘nested’ EM algorithm that holds Monte Carlo estimates from the Gibbs output fixed. van Dyk (2000) likens the approach to an EM version of blocked Gibbs sampling.

In our case, this approach consists of two steps for each EM iteration (indexed by t) and each additive component (indexed by j). The first step is the Monte Carlo step. It consists of generating a sample from $(\mathbf{x}_j \mid \mathbf{b}, \mathbf{x}_{-j}, \mathbf{y}, \mathbf{w}; \boldsymbol{\psi}^{(t)})$ and using this sample to compute Monte Carlo estimates of the expectations of the complete data sufficient statistics for $\boldsymbol{\psi}^{(t)}$ that involve \mathbf{x}_j . The second step is followed by the nested EM step, where we fix the estimates of the complete data sufficient statistics at their current values and compute $\boldsymbol{\psi}^{(t+1)}$ from several iterations of an EM algorithm that treats only \mathbf{b} as latent. van Dyk (2000) shows that convergence of the nested EM in each nested step is not necessary for convergence of the complete algorithm. The benefit of this procedure is that it makes more efficient use of computationally expensive Metropolis–Hastings samples. Since the nested step is aimed at fitting a linear mixed model, another method such as Newton–Raphson could be used in the nested step instead of EM.

The following fact shows one way of drawing from the conditional densities of

$$(\mathbf{x}_j \mid \mathbf{b}, \mathbf{x}_{-j}, \mathbf{y}, \mathbf{w}, \boldsymbol{\psi}^{(t)}) \quad (j = 1, \dots, d).$$

This method was also used by Berry *et al.* (2002). It is a direct consequence of the likelihoods that specify our model.

Fact 1. *The density of $(\mathbf{x}_j \mid \mathbf{b}, \mathbf{x}_{-j}, \mathbf{y}, \mathbf{w}, \boldsymbol{\psi}^{(t)})$ is proportional to*

$$\exp \left(-\frac{1}{2\sigma_\varepsilon^2} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\mathbf{b})^\top (\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\mathbf{b}) - \sum_{j=1}^d \frac{1}{2\sigma_{\mathbf{x}_j}^2} (\mathbf{x}_j - \boldsymbol{\mu}_j)^\top (\mathbf{x}_j - \boldsymbol{\mu}_j) - \sum_{j=1}^d \frac{1}{2\sigma_{\mathbf{u}_j}^2} (\mathbf{w}_j - \mathbf{x}_j)^\top (\mathbf{w}_j - \mathbf{x}_j) \right).$$

As a result, generating from the conditional distribution of \mathbf{x}_j can be done using the Metropolis–Hastings algorithm; see e.g. Casella & Robert (1999 Chapter 6).

3.3. Complete algorithm

Now we specify the complete algorithm.

- (i) Start with an initial $\boldsymbol{\psi}^{(0)}$. To establish the starting values, we fit a fixed degree-of-freedom additive model using $E(\mathbf{x}_j | \mathbf{w}_j)$ in place of the observed covariates. These conditional expectations require estimates of μ_{x_j} and $\sigma_{x_j}^2$, and we estimate them using method of moments estimators: $\hat{\mu}_{x_j} = \mathbf{1}_n^\top \mathbf{w}_j / n$ and $\hat{\sigma}_{x_j}^2 = (s_{w_j}^2 - \sigma_{u_j}^2) / s_{w_j}^2$. With $j = 1, \dots, d$ indexing the additive components, the $(t + 1)$ th parameter update involves step (ii).
- (ii) Sample from $(\mathbf{x}_j | \mathbf{b}, \mathbf{x}_{-j}, \mathbf{y}, \mathbf{w}; \boldsymbol{\psi}^{(t)})$ using the Metropolis–Hastings algorithm and Fact 1.
- (iii) Let

$$\mathbf{P} = \begin{bmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{X}^\top \mathbf{Z} \\ \mathbf{Z}^\top \mathbf{X} & \mathbf{Z}^\top \mathbf{Z} + \mathbf{D} \end{bmatrix}.$$

With the results from step (ii), compute Monte Carlo estimates of

$$\mathbf{P}, \mathbf{W}^\top \mathbf{y}, \quad \text{and} \quad \mathbf{W}^\top \mathbf{W}, \quad \text{where} \quad \mathbf{W} = [\mathbf{X} \ \mathbf{Z}].$$

We denote estimates of these quantities by

$$\hat{\mathbf{P}}, \hat{\mathbf{W}}^\top \mathbf{y}, \quad \text{and} \quad \hat{\mathbf{W}}^\top \hat{\mathbf{W}}.$$

- (iv) Holding the estimates from step (iii) fixed, run several iterations of an update scheme. For instance, using the standard EM algorithm to compute REML estimates (e.g. Dempster, Rubin & Tsutakawa, 1981), the $(k + 1)$ th nested updates are:
- (a) $\mathbf{b}^{(k+1)} = (\boldsymbol{\beta}^{(k+1)}, \mathbf{b})$ and $\mathbf{b}^{(k+1)} = \hat{\mathbf{P}}^{-1} \hat{\mathbf{W}}^\top \mathbf{y}$.
- (b) Letting $\mathbf{P}_{[b_j]}^{-1}$ correspond to the block submatrix of \mathbf{P}^{-1} associated with b_j ,

$$\sigma_{b_j}^{2(k+1)} = \frac{1}{K_j} = (\mathbf{b}_j^\top \mathbf{b}_j + \text{tr}(\sigma_\varepsilon^{2(k)} \hat{\mathbf{P}}_{[b_j]}^{-1})) \quad (j = 1, \dots, d).$$

$$(c) \quad \sigma_\varepsilon^{2(k+1)} = \frac{1}{n} (\mathbf{y}^\top \mathbf{y} - 2\mathbf{y}^\top \hat{\mathbf{W}} \mathbf{b}^{(k+1)} + \mathbf{b}^{(k+1)\top} \hat{\mathbf{W}}^\top \hat{\mathbf{W}} \mathbf{b}^{(k+1)} + \sigma_\varepsilon^{2(k)} \text{tr}(\hat{\mathbf{W}}^\top \hat{\mathbf{W}} \hat{\mathbf{P}}^{-1})).$$

- (v) Calculate $\mu_{x_j}^{(t+1)}$ and $\sigma_{x_j}^{2(t+1)}$ using standard point estimates based on the Monte Carlo data in step (iii).

We terminate the algorithm after plots of the current estimates of the additive components ($\hat{m}_j^{(t+1)}(\cdot)$) appear to stabilize (Wei & Tanner, 1990).

4. Illustrations

4.1. Simulation

Figure 1 contains the results of a simulation study to illustrate the performance of our method on scatterplot data and explore its sensitivity to different levels of regression and measurement error variance. The simulation design was a 2×2 factorial on measurement error variance and regression error variance. The measurement error variance levels were determined by setting a scale-free measure of measurement error (the reliability ratio $\lambda =$

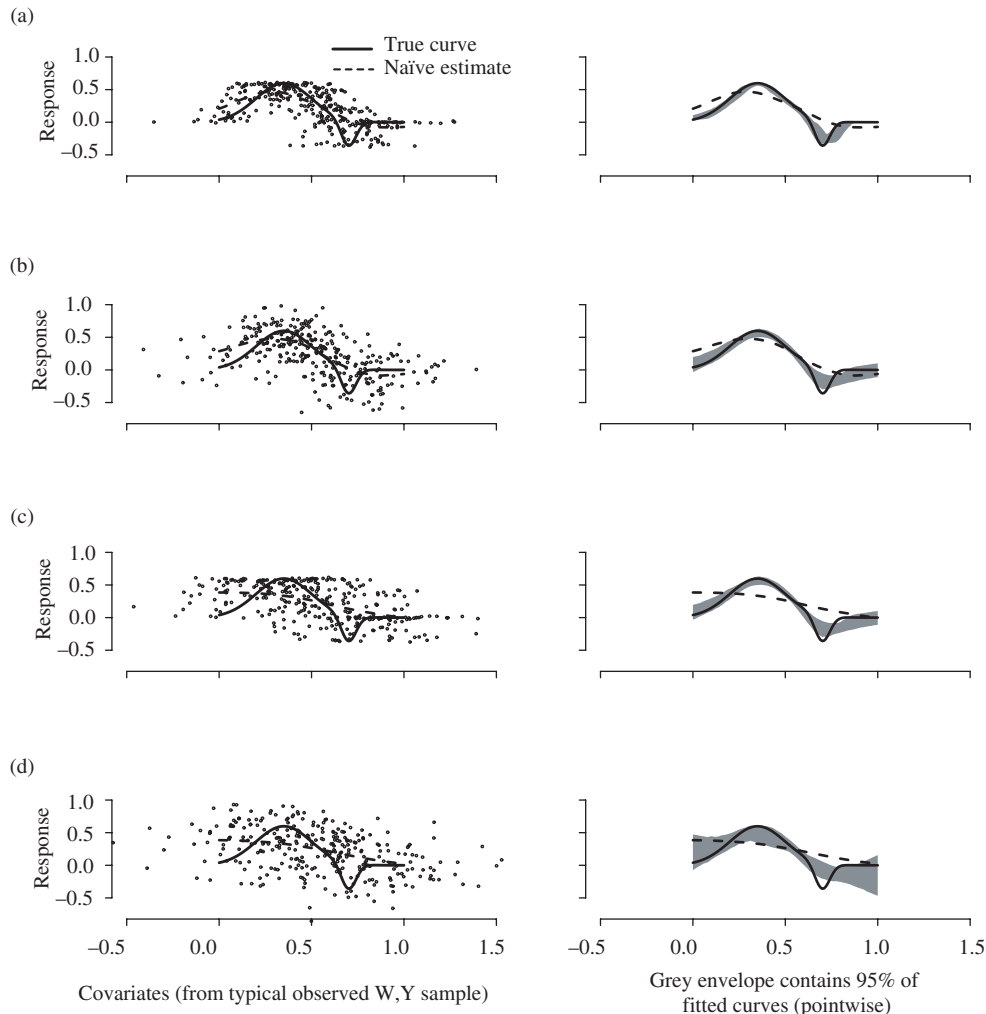


Figure 1. Output from a Monte Carlo simulation study to illustrate the performance of the method. The test cases consist of a 2×2 factorial combination of two levels of measurement error and regression error; 200 replicates; sample size = 300. The left column of panels shows the true curve, the mean estimate that ignores measurement error (naïve estimate), and the observed data. The right column shows the naïve estimate, true curve, and a shaded envelope that covers 95% (pointwise) of the EM-spline estimates. (a) Moderate measurement error and low regression error, (b) moderate measurement error and high regression error, (c) high measurement error and low regression error, and (d) high measurement error and high regression error.

$\sigma_x^2 / (\sigma_u^2 + \sigma_x^2)$ to moderate = 0.8 and high = 0.5. The regression error variance was set to a very small amount, $\sigma_\varepsilon^2 = 0.01^2$, and a higher amount, $\sigma_\varepsilon^2 = 0.2^2$. In each case, true $m(x) = 1.5\varphi((x - 0.35)/0.15) - \varphi((x - 0.7)/0.04)$ where φ denotes the normal probability density function, and the sample size was 300. We ran 200 Monte Carlo iterations, and we computed average pointwise squared error for each fitted curve to assess the estimates. Within each simulation case, the 5th, 50th, or 95th percent best fit in Figure 1 is the curve associated with the empirical 5th, 50th, or 95th percent smallest average pointwise squared error.




Figure 2. The sensitivity of the fits to four levels of additive covariate measurement error in $\log(\text{TSP})$, for an additive model that describes the relationship between mortality and the factors: (a) total suspended particles ($\log(\text{TSP})$), (b) day, (c) temperature, and (d) humidity in Milan, Italy 1980–1989 (Zanobetti *et al.*, 2000)

The figures illustrate that our proposed method identifies features of the true curve that were missed by the estimator that ignores covariate measurement error (the naïve estimator). Even in the presence of ‘moderate’ measurement error, the naked eye does not immediately identify the true curve from the observed scatterplot data.

4.2. Air pollution example

Next, we fit a simple additive model to an air pollution/mortality dataset recorded in Milan, Italy, from 1980 to 1989 (Zanobetti *et al.*, 2000) and analyse its sensitivity to measurement error in total suspended particles. Given the lack of validation data, our purpose is to show that models such as this one can be sensitive to covariate measurement error. See Dominici & Zeger (2000) for a discussion of measurement error in air pollution/mortality studies.

Let i index day and TSP_i , day_i , temp_i , humid_i , and mort_i denote the measured total suspended particles, sequential day number, average temperature, average relative humidity, and mortality (from natural causes) count on day i . Our additive model regression equation is:

$$E(\ln(\text{mort}_i)) = \beta_0 + m_1(\ln(\text{TSP}_i)) + m_2(\text{day}_i) + m_3(\text{temp}_i) + m_4(\text{humid}_i).$$

To assess the sensitivity of estimates to measurement error in TSP_i , we fit the model assuming that the reliability ratio (see Section 4.1) for $\ln(\text{TSP}_i)$ is 1.0 (no measurement error), 0.9, 0.8, and 0.7.

Figure 2 shows that the estimate of the function that relates $\ln(\text{TSP})$ to mortality appears to be sensitive to measurement error. This suggests that an analysis of these data would benefit from measurement error modelling and the collection of validation data.

5. Discussion and extensions

We have described an effective mechanism for additive modelling when predictors are subject to measurement error. Unlike Berry *et al.* (2002), our approach is relative frequentist, and our simulations point to evidence that good performance can be achieved without the specification of priors. We have also considered multipredictor additive models rather than single predictor scatterplot smoothing.

Three straightforward augmentations to our method include calculating approximate standard errors for the curve estimates, incorporating validation data, and allowing a more flexible model for the structural distribution of (x_1, \dots, x_d) . Since the curve estimates are linear combinations of estimates of β and b , it is sufficient to estimate the asymptotic covariance matrix of $\hat{\beta}$ and \hat{b} . This can be done using the method described in Section 5 of Chan & Kuk (1997). Validation data can be incorporated, using the method described in Spiegelman *et al.* (2000) for instance. The distribution of the unobserved covariates can be modelled with a mixture of normals, as described in Carroll *et al.* (1999) for instance.

References

- BERRY, S., CARROLL, R.J. & RUPPERT, D. (2002). Bayesian smoothing and regression splines for measurement error problems. *J. Amer. Statist. Assoc.* **97**, 160–169.
- BUONACCORSI, J., DEMIDENKO, E. & TOSTESON, T. (2000). Estimation in longitudinal random effects models with measurement error. *Statist. Sinica* **10**, 885–903.
- CARROLL, R.J., MACA, J.D. & RUPPERT, D. (1999). Nonparametric regression in the presence of measurement error. *Biometrika* **86**, 541–554.
- CARROLL, R.J., RUPPERT, D. & STEFANSKI, L.A. (1995). *Measurement Error in Nonlinear Models*. London: Chapman and Hall.
- CARROLL, R.J., SPIEGELMAN, C.H., LAN, K.K.G., GAILEY, K.T. & ABBOTT, R.D. (1984). On errors in variables for binary regression models. *Biometrika* **71**, 19–25.
- CASELLA, G. & ROBERT, C.P. (1999). *Monte Carlo Statistical Methods*. New York: Springer.
- CHAN, J.S.K. & KUK, A.Y.C. (1997). Estimation for probit–linear mixed models with correlated random effects. *Biometrics* **53**, 86–97.
- DEMPSTER, A.P., RUBIN, D.B. & TSUTAKAWA, R.K. (1981). Estimation in covariance components models. *J. Amer. Statist. Assoc.* **76**, 341–353.
- DOMINICI, F. & ZEGER S.L. (2000). A measurement error model for time–series studies of air pollution and mortality. *Biostatistics* **1**, 157–175.
- EILERS, P.H.C. & MARX, B.D. (1996). Flexible smoothing with B-splines and penalties (with discussion). *Statist. Sci.* **89**, 89–121.
- FAN, J. & TRUONG, Y.K. (1993). Nonparametric regression with errors in variables. *Ann. Statist.* **21**, 1900–1925.
- FULLER, W.A. (1987). *Measurement Error Models*. New York: John Wiley & Sons.
- GREEN, P.J. & SILVERMAN, B.W. (1994). *Nonparametric Regression and Generalized Linear Models*. London: Chapman and Hall.
- MCCULLOCH, C.E. & SEARLE, S.R. (2001). *Generalized, Linear, and Mixed Models*. New York: Wiley.

- RUPPERT, D. (2002). Selecting the number of knots for penalized splines. *J. Comput. Graph. Statist.* **11**, 735–757.
- RUPPERT, D., WAND, M.P. & CARROLL, R.J. (2003). *Semiparametric Regression*. Cambridge, UK: Cambridge University Press.
- SMYTH, G.K. (1996). Partitioned algorithms for maximum likelihood and other non-linear estimation. *Statist. Comput.* **6**, 201–216.
- SPEED, T. (1991). Discussion of “That BLUP is a good thing: the estimation of random effects” by G. Robinson. *Statist. Sci.* **6**, 42–44.
- SPIEGELMAN, D., ROSNER, B. & LOGAN, R. (2000). Estimation and inference for logistic regression with covariate misclassification and measurement error, in main study/validation study designs. *J. Amer. Statist. Assoc.* **95**, 51–61.
- STAUDENMAYER, J. & RUPPERT, D. (2004). Local polynomial regression and simulation extrapolation. *J. Roy. Statist. Soc. Ser. B* **66**, 17–30.
- VAN DYK, D.A. (2000). Nesting EM algorithms for computational efficiency. *Statist. Sinica* **10**, 203–226.
- VERBYLA, A.P., CULLIS, B.R., KENWARD, M.G. & WELHAM, S.J. (1999). The analysis of designed experiments and longitudinal data by using smoothing splines (with discussion). *J. Roy. Statist. Soc. Ser. C* **48**, 269–300.
- WAND, M.P. (2003). Smoothing and mixed models. *Comput. Statist.* **18**, 223–249.
- WEI, G.C.G. & TANNER, M.A. (1990). A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *J. Amer. Statist. Assoc.* **85**, 699–704.
- ZANOBETTI, A., WAND, M.P., SCHWARTZ, J. & RYAN, L.M. (2000). Generalized additive distributed lag models: quantifying mortality displacement. *Biostatistics* **1**, 279–292.