

Streamlined Variance Calculations for Semiparametric Mixed Models

BY ANDREW D.A.C. SMITH AND M.P. WAND

*School of Mathematics and Statistics,
University of New South Wales, Sydney 2052, AUSTRALIA*

29th March, 2007

ABSTRACT

Semiparametric mixed model analysis benefits from variability estimates such as standard errors of effect estimates and variability bars to accompany curve estimates. We show how the underlying variance calculations can be done extremely efficiently compared with the direct naïve approach. These streamlined calculations are linear in the number of subjects, representing a two orders of magnitude improvement.

Some keywords: Additive mixed models; Longitudinal data analysis; Penalised splines; Semiparametric regression; Subject-specific curves.

1 Introduction

A current vibrant area of research is the use of nonparametric regression, or smoothing, techniques in the analysis of longitudinal data. Prominent examples include [1,2,3,4,5]. Summaries may be found in books such as [6,7,8].

FIGURE 1 NEAR HERE

Figure 1: *Spinal bone mineral density data broken down according to ethnicity of the subjects.*

Figure 1 shows an example of data that benefits from such methodology. It consists of longitudinal measurements on the spinal bone mineral density (SBMD) of a cohort of young female subjects (source: Reference [9]). One question of interest concerns differences in mean SBMD among the 4 ethnic groups after accounting for age. An appropriate model is the *additive mixed model*

$$\text{SBMD}_{ij} = f(\text{age}_{ij}) + \beta_2 \text{black}_i + \beta_3 \text{hispanic}_i + \beta_4 \text{white}_i + U_i + \varepsilon_{ij}, \quad (1)$$

Here SBMD_{ij} denotes the j th ($1 \leq j \leq n_i$) SBMD measurement on subject i ($1 \leq i \leq m$), f is a smooth, but otherwise unspecified, function for the mean effect of age and black_i , hispanic_i and white_i are indicator variables for ethnicity. In addition, the U_i i.i.d. $N(0, \sigma_U^2)$ are random subject intercepts, and the ε_{ij} i.i.d. $N(0, \sigma_\varepsilon^2)$, independent of the U_i 's, account for within-subject variability. If f is modelled using penalised splines, such as

$$f(\text{age}) = \beta_0 + \beta_1 \text{age} + \sum_{k=1}^K u_k (\text{age} - \kappa_k)_+, \quad u_k \text{ i.i.d. } N(0, \sigma_u^2),$$

where $\kappa_1, \dots, \kappa_K$ is a dense set of knots, then (1) reduces to a linear mixed model. The details are given in Section 2. Standard fitting leads to the summary table given in Table 1 showing, for example, a highly significant difference between black and Asian females in terms of mean SBMD.

TABLE 1 NEAR HERE

Table 1: *Summary of ethnicity effects for linear mixed model fit of (1).*

Figure 2 provides a summary of the estimate of f and its variability. The curves correspond to the function $\hat{f}(\text{age})$ with vertical shifting according to $\hat{\beta}_i$, $i = 2, 3, 4$. The variability bars, shown as dashed lines, are $\pm 2 \times$ the estimated standard errors and correspond to approximate pointwise 95% confidence intervals of mean SBMD.

FIGURE 2 NEAR HERE

Figure 2: *Data from Figure 1 with estimates of $f(\text{age})$ added. The dashed lines correspond to $\pm 2 \times$ the estimated standard error.*

Underlying Table 1 and Figure 2 is the estimated covariance matrix of the coefficient estimates:

$$[\hat{\beta}_0, \dots, \hat{\beta}_4, \hat{u}_1, \dots, \hat{u}_K, \hat{U}_1, \dots, \hat{U}_m]^T.$$

This covariance matrix involves the inversion of a $(5 + K + m) \times (5 + K + m)$ matrix \mathbf{M} . Typically K is in the range 15–40 regardless of the sample size variables. However, the number of subjects m can be arbitrarily large. For the SBMD example $m = 230$ which does not pose serious problems. But other studies involve much larger m . For example the *Six Cities Study of Air Pollution and Health*, described in [10] (p.210), has $m = 13,379$. Thus, for many situations, the number of subjects m is the dominant term in the dimension of \mathbf{M} . Henceforth, let K be fixed and small but allow m to be arbitrarily large. Then it is well-known from numerical linear algebra (e.g. Reference [11]) that naïve computation of \mathbf{M}^{-1} is $O(m^3)$. Hence direct variance calculations can be very costly, or even prohibitive, in large longitudinal studies.

The purpose of this article is to show that variance calculations of interest, such as those required for the standard errors of Table 1 and the variability bars of Figure 2, can be done in $O(m)$ operations. The key is recognition that the contribution to \mathbf{M} from the random intercept component is an $m \times m$ diagonal matrix. Such streamlining essentially removes computational obstacles involving variances for models such as (1) for most practical values of m and, thus, greatly benefits semiparametric mixed model analysis.

Section 2 gives the details of our streamlined approach to variance calculations for models like (1). In Section 3 we describe the extension of the approach to subject-specific curve models. Closing remarks are made in Section 4.

2 Additive Mixed Models

In this section we consider a more general version of (1):

$$y_{ij} = f(s_{ij}) + \mathbf{x}_{ij}^T \boldsymbol{\beta}_x + U_i + \varepsilon_{ij}. \quad (2)$$

Here y_{ij} is the j th ($1 \leq j \leq n_i$) measurement of the response of the i th subject ($1 \leq i \leq m$), s_{ij} is a predictor with a possibly non-linear effect, \mathbf{x}_{ij} is a $p \times 1$ vector of predictors with

a linear effect, with corresponding coefficient vector $\beta_{\mathbf{x}}$. Again U_i i.i.d $N(0, \sigma_U^2)$ is a random intercept and ε_{ij} i.i.d $N(0, \sigma_\varepsilon^2)$, independently of the U_i 's. The smooth function f is modelled using penalised splines of the form:

$$f(s) = \beta_0 + \beta_1 s + \sum_{k=1}^K u_k z_k(s),$$

where z_k , $1 \leq k \leq K$, is an appropriate spline basis. The simplest choice is $z_k(s) = (s - \kappa_k)_+$ for a dense set of knots $\kappa_1, \dots, \kappa_K$ but many other options exist (e.g. Reference [6], Section 3.7). If we let u_k i.i.d $N(0, \sigma_u^2)$ then (2) becomes a linear mixed model $\mathbf{y} = \mathbf{X}\beta + \mathbf{Z}\mathbf{u} + \varepsilon$, where

$$\begin{aligned} \beta &= [\beta_0, \beta_1, \beta_{\mathbf{x}}^T]^T, & \mathbf{Z} &= [\mathbf{Z}_G \quad \mathbf{Z}_R], & \mathbf{u} &= [\mathbf{u}_G^T \quad \mathbf{u}_R^T]^T, \\ \mathbf{X} &= \begin{bmatrix} 1 & s_{11} & \mathbf{x}_{11}^T \\ \vdots & \vdots & \vdots \\ 1 & s_{1n_1} & \mathbf{x}_{1n_1}^T \\ \vdots & \vdots & \vdots \\ 1 & s_{m1} & \mathbf{x}_{m1}^T \\ \vdots & \vdots & \vdots \\ 1 & s_{mn_m} & \mathbf{x}_{mn_m}^T \end{bmatrix}, & \mathbf{Z}_G &= \begin{bmatrix} z_1(s_{11}) & \cdots & z_K(s_{11}) \\ \vdots & & \vdots \\ z_1(s_{1n_1}) & \cdots & z_K(s_{1n_1}) \\ \vdots & & \vdots \\ z_1(s_{m1}) & \cdots & z_K(s_{m1}) \\ \vdots & & \vdots \\ z_1(s_{mn_m}) & \cdots & z_K(s_{mn_m}) \end{bmatrix}, & (3) \\ \mathbf{Z}_R &= \text{blockdiag}(\mathbf{1}_{n_i}), & \mathbf{u}_G &= [u_1, \dots, u_K]^T, & \mathbf{u}_R &= [U_1, \dots, U_m]^T, \\ & & & & & \mathbf{G}_G = \text{Cov}(\mathbf{u}_G) = \sigma_u^2 \mathbf{I}_K, & \mathbf{G}_R &= \text{Cov}(\mathbf{u}_R) = \sigma_U^2 \mathbf{I}_m \end{aligned}$$

and the vectors \mathbf{y} and ε are defined analogously. Note that the random effects have been partitioned into spline coefficients (subscript G) and subject effects (subscript R).

Let $\hat{\mathbf{G}}_G$ and $\hat{\mathbf{G}}_R$ be the restricted maximum likelihood estimates of \mathbf{G}_G and \mathbf{G}_R and $\hat{\beta}$ and $\hat{\mathbf{u}}$ be the empirical best linear unbiased predictors of β and \mathbf{u} (e.g. Reference [6], Section 4.5). Then all variance calculations can be done using the estimated covariance matrix

$$\widehat{\text{Cov}} \begin{bmatrix} \hat{\beta} \\ \hat{\mathbf{u}} - \mathbf{u} \end{bmatrix} = \mathbf{M}^{-1}$$

where

$$\mathbf{M} = \hat{\sigma}_\varepsilon^{-2} \mathbf{C}^T \mathbf{C} + \hat{\mathbf{B}}, \quad \mathbf{C} = [\mathbf{X} \quad \mathbf{Z}] \quad \text{and} \quad \hat{\mathbf{B}} = \text{blockdiag}(\mathbf{0}, \hat{\mathbf{G}}_G^{-1}, \hat{\mathbf{G}}_R^{-1})$$

(e.g. Reference [6], Section 4.7).

The direct approach to obtaining standard errors of the entries of $\hat{\beta}$ (as in Table 1) and variability bars for the smooth function estimate (as in Figure 2) involves inversion of \mathbf{M} . As mentioned in Section 1, the matrix \mathbf{M} will increase in dimension as the number of subjects, m , grows and for very large m it will become too computationally intensive to invert practically, since inversion is an $O(m^3)$ process. However we do not need to find all the entries of the matrix. For example, some entries relate to the correlation between different subjects' responses, which would rarely be useful. We can also exploit the fact that, for large m , most of \mathbf{M} is diagonal. Hence, we propose a streamlined approach based around a block decomposition of \mathbf{M} :

$$\mathbf{M} = \hat{\sigma}_\varepsilon^{-2} \begin{bmatrix} \mathbf{X}^T \mathbf{X} & \mathbf{X}^T \mathbf{Z}_G & \mathbf{X}^T \mathbf{Z}_R \\ \mathbf{Z}_G^T \mathbf{X} & \mathbf{Z}_G^T \mathbf{Z}_G + \hat{\sigma}_\varepsilon^2 \hat{\mathbf{G}}_G^{-1} & \mathbf{Z}_G^T \mathbf{Z}_R \\ \mathbf{Z}_R^T \mathbf{X} & \mathbf{Z}_R^T \mathbf{Z}_G & \mathbf{Z}_R^T \mathbf{Z}_R + \hat{\sigma}_\varepsilon^2 \hat{\mathbf{G}}_R^{-1} \end{bmatrix} \equiv \hat{\sigma}_\varepsilon^{-2} \begin{bmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} \\ \mathbf{M}_{21} & \mathbf{M}_{22} \end{bmatrix}.$$

Using standard results on the inverse of a block-partitioned matrix (e.g. Reference [12])

$$\begin{aligned} \mathbf{M}^{-1} &= \hat{\sigma}_\varepsilon^2 \begin{bmatrix} \mathbf{M}^{11} & \mathbf{M}^{12} \\ \mathbf{M}^{21} & \mathbf{M}^{22} \end{bmatrix} \\ \text{where } \mathbf{M}^{11} &= (\mathbf{M}_{11} - \mathbf{M}_{12}\mathbf{M}_{22}^{-1}\mathbf{M}_{21})^{-1}, \\ \mathbf{M}^{12} &= -\mathbf{M}^{11}\mathbf{M}_{12}\mathbf{M}_{22}^{-1}, \quad \mathbf{M}^{21} = (\mathbf{M}^{12})^T \\ \text{and } \mathbf{M}^{22} &= \mathbf{M}_{22}^{-1} - \mathbf{M}_{22}^{-1}\mathbf{M}_{21}\mathbf{M}^{11}\mathbf{M}_{12}\mathbf{M}_{22}^{-1}. \end{aligned} \quad (4)$$

Straightforward matrix multiplication finds $\mathbf{X}^T\mathbf{X}$, $\mathbf{X}^T\mathbf{Z}_G$ and $\mathbf{Z}_G^T\mathbf{Z}_G$ in $O(m)$ steps, but $\mathbf{Z}_R^T\mathbf{X}$ and $\mathbf{Z}_R^T\mathbf{Z}_G$ are more complicated since \mathbf{Z}_R is at least $m \times m$. However the special structure of \mathbf{Z}_R means that the matrix products amount to within-subject row sums, which are $O(m)$, since for $\mathbf{A} = \mathbf{X}$ or $\mathbf{A} = \mathbf{Z}_G$

$$\mathbf{Z}_R^T\mathbf{A} = \begin{bmatrix} 1 & \cdots & 1 & & & & \mathbf{0} \\ & & & \ddots & & & \\ & & & & & & \\ \mathbf{0} & & & & 1 & \cdots & 1 \end{bmatrix} \begin{bmatrix} \mathbf{a}_{11}^T \\ \vdots \\ \mathbf{a}_{1n_1}^T \\ \vdots \\ \mathbf{a}_{m1}^T \\ \vdots \\ \mathbf{a}_{mn_m}^T \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^{n_1} \mathbf{a}_{1j}^T \\ \vdots \\ \sum_{j=1}^{n_m} \mathbf{a}_{mj}^T \end{bmatrix}.$$

In addition, \mathbf{M}_{22} is diagonal:

$$\mathbf{Z}_R^T\mathbf{Z}_R + \hat{\sigma}_\varepsilon^2\hat{\sigma}_U^{-2}\mathbf{I}_m = \text{diag}_{1 \leq i \leq m} (n_i + \hat{\sigma}_\varepsilon^2/\hat{\sigma}_U^2).$$

This feature is crucial. It enables us to find the biggest inverse required in (4) in $O(m)$ steps since $\mathbf{M}_{22}^{-1} = \text{diag}_{1 \leq i \leq m} \left(\frac{\hat{\sigma}_U^2}{\hat{\sigma}_U^2 n_i + \hat{\sigma}_\varepsilon^2} \right)$. Thus we can write \mathbf{M} as

$$\hat{\sigma}_\varepsilon^{-2} \begin{bmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} \\ \mathbf{M}_{21} & \mathbf{M}_{22} \end{bmatrix} = \hat{\sigma}_\varepsilon^{-2} \left[\begin{array}{c|ccc} \mathbf{M}_{11} & \mathbf{h}_1 & \cdots & \mathbf{h}_m \\ \mathbf{h}_1^T & n_1 + \frac{\hat{\sigma}_\varepsilon^2}{\hat{\sigma}_U^2} & & \mathbf{0} \\ \vdots & & \ddots & \\ \mathbf{h}_m^T & \mathbf{0} & & n_m + \frac{\hat{\sigma}_\varepsilon^2}{\hat{\sigma}_U^2} \end{array} \right],$$

where $\mathbf{h}_i^T = \sum_{j=1}^{n_i} [1 \quad s_{ij} \quad \mathbf{x}_{ij}^T \quad z_1(s_{ij}) \quad \cdots \quad z_K(s_{ij})]$. Then using (4) we see that:

$$\mathbf{M}^{11} = \left(\mathbf{M}_{11} - \sum_{i=1}^m \frac{\hat{\sigma}_U^2 \mathbf{h}_i \mathbf{h}_i^T}{\hat{\sigma}_U^2 n_i + \hat{\sigma}_\varepsilon^2} \right)^{-1}.$$

This inverse is $(2+p+K) \times (2+p+K)$ and is relatively easy to compute. The summation term renders this whole process $O(m)$. This matrix is all we require to plot global or group error bars, or perform hypothesis tests as in Table 1. However to find the error covariance of the individual subjects' fitted responses we need \mathbf{M}^{12} and \mathbf{M}^{22} , where

$$\mathbf{M}^{12} = \left[-\frac{\hat{\sigma}_U^2}{\hat{\sigma}_U^2 n_1 + \hat{\sigma}_\varepsilon^2} \mathbf{M}^{11} \mathbf{h}_1 \quad \cdots \quad -\frac{\hat{\sigma}_U^2}{\hat{\sigma}_U^2 n_m + \hat{\sigma}_\varepsilon^2} \mathbf{M}^{11} \mathbf{h}_m \right].$$

Each submatrix is a constant times $(2+p+K) \times (2+p+K)$ times $(2+p+K) \times 1$. However we must calculate m of them, so \mathbf{M}^{12} (and consequently \mathbf{M}^{21}) requires $O(m)$ calculations.

Finally, we only need to find the diagonal entries of \mathbf{M}^{22} since we would rarely be interested in the correlation between two subjects' fitted responses. These are given by

$$\mathbf{M}_{ii}^{22} = \frac{\hat{\sigma}_U^2}{\hat{\sigma}_U^2 n_i + \hat{\sigma}_\varepsilon^2} \left(1 + \frac{\hat{\sigma}_U^2 \mathbf{h}_i^T \mathbf{M}^{11} \mathbf{h}_i}{\hat{\sigma}_U^2 n_i + \hat{\sigma}_\varepsilon^2} \right), \quad 1 \leq i \leq m.$$

The relevant diagonal entries can be calculated in $O(m)$ steps. Therefore the total asymptotic complexity of this process is $O(m)$, representing an improvement of order m^2 over the naïve approach to matrix inversion. As the number of subjects, m , increases the improvements due to streamlining become enormous.

We could alter the model in (2) to include more random subject effects, such as random slopes. This would only affect the final stage of the calculation as it would alter the structure of \mathbf{M}_{22} . However the result of this is that \mathbf{M}_{22} becomes block-diagonal, so it can still be inverted in $O(m)$ calculations.

The practical benefits of streamlined variance calculations were explored in a simulation study. Data were generated according to

$$y_{ij} = -\sin(2\pi s_{ij}) + 0.3 x_{ij} + U_i + \varepsilon_{ij}, \quad 1 \leq i \leq n_i, \quad 1 \leq i \leq m \quad (5)$$

where the s_{ij} were generated from the uniform distribution on $(0, 1)$, the x_{ij} were generated from the Bernoulli distribution with $P(x_{ij} = 0) = P(x_{ij} = 1) = \frac{1}{2}$, U_i i.i.d. $N(0, 0.5^2)$ and ε_{ij} i.i.d. $N(0, 0.2^2)$. The $0.3 x_{ij}$ term represents a binary offset from the smooth function in s_{ij} . The within-subject sample sizes n_i were generated uniformly from $\{1, 2, 3, 4\}$ and values of $m \in \{100, 500, 2500, 12500\}$ were considered. For each value of m , twenty-five replicate data sets were generated. All computations were done in the R language [13] on a Dell Optiplex SX280 PC with a 2.8GHz Intel Pentium 4 Processor. The Appendix contains code used for the streamlined variance calculations. The time taken to compute variability bars and the standard error of the offset estimate was recorded. Table 2 summarises the results. Note that the naïve method approach failed for $m = 12500$ due to required storage for \mathbf{M} exceeding memory restrictions.

TABLE 2 NEAR HERE

Table 2: Average times in seconds (standard errors) for computing variability bars and standard error of the binary offset.

There is little practical difference between the two methods for $m = 100$ and $m = 500$. However for $m = 2500$ streamlined variance calculation is much faster — taking about one tenth of a second on average compared with almost 4 minutes for the naïve approach. For $m = 12500$ the streamlined approach is still well under one second, while the naïve approach is not viable for typical 2006 computing environments.

3 Extension to Subject Specific Curves

Models (1) and (2) featured *random intercepts*: the difference between the fitted subject response and estimated population mean curve is constant. This may not in general be realistic; the subject-specific difference may be as complicated as the underlying function f . Durban, Harezlak, Wand & Carroll [14] develop a *subject-specific curves* model, based on penalised splines, in which the subject-specific difference is modelled by a random semiparametric function:

$$y_{ij} = f(s_{ij}) + \mathbf{x}_{ij}^T \boldsymbol{\beta}_x + g_i(s_{ij}) + \varepsilon_{ij}. \quad (6)$$

Earlier work on models of this type includes [15,2,3,5,16]. In this new model U_i is replaced by

$$g_i(s) = U_{i0} + U_{i1}s + \sum_{i=1}^{K_g} v_{ik} z_{gk}(s), \quad v_{gk} \text{ i.i.d. } N(0, \sigma_v^2),$$

where z_{gk} , $1 \leq k \leq K_g$, is an appropriate spline basis. It is an advantage of our setup that g_i need not share the same spline basis as f . Therefore the splines of f are labelled z_{fk} , $1 \leq k \leq K_f$. The model (6) can be written as a linear mixed model $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\varepsilon}$ where

$$\begin{aligned} \mathbf{X} &= \begin{bmatrix} \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_m \end{bmatrix}, & \mathbf{Z} &= \begin{bmatrix} \mathbf{Z}_1 & \mathbf{T}_1 \mathbf{W}_1 & & \mathbf{0} \\ \vdots & & \ddots & \\ \mathbf{Z}_m & \mathbf{0} & & \mathbf{T}_m \mathbf{W}_m \end{bmatrix}, \\ \mathbf{X}_i &= \begin{bmatrix} 1 & s_{i1} & \mathbf{x}_{i1}^T \\ \vdots & \vdots & \vdots \\ 1 & s_{in_i} & \mathbf{x}_{in_i}^T \end{bmatrix}, & \mathbf{Z}_i &= \begin{bmatrix} z_{f1}(s_{i1}) & \cdots & z_{fK_f}(s_{i1}) \\ \vdots & & \vdots \\ z_{f1}(s_{in_i}) & \cdots & z_{fK_f}(s_{in_i}) \end{bmatrix}, \\ \mathbf{T}_i &= \begin{bmatrix} 1 & s_{i1} \\ \vdots & \vdots \\ 1 & s_{in_i} \end{bmatrix}, & \mathbf{W}_i &= \begin{bmatrix} z_{g1}(s_{i1}) & \cdots & z_{gK_g}(s_{i1}) \\ \vdots & & \vdots \\ z_{g1}(s_{in_i}) & \cdots & z_{gK_g}(s_{in_i}) \end{bmatrix}, \\ \mathbf{u} &= [\mathbf{u}_G^T, \mathbf{U}_1^T, \mathbf{v}_1^T, \dots, \mathbf{U}_m^T, \mathbf{v}_m^T]^T, & \mathbf{U}_i &= [U_{i0}, U_{i1}]^T, & \mathbf{v}_i &= [v_{i1}, \dots, v_{iK_g}]^T, \\ \mathbf{G}_G &= \text{Cov} [\mathbf{U}_1^T, \mathbf{v}_1^T, \dots, \mathbf{U}_m^T, \mathbf{v}_m^T]^T = \text{blockdiag}(\mathbf{D}, \sigma_v^2 \mathbf{I}_{K_g}, \dots, \mathbf{D}, \sigma_v^2 \mathbf{I}_{K_g}), \end{aligned}$$

and \mathbf{y} , $\boldsymbol{\beta}$, $\boldsymbol{\varepsilon}$, \mathbf{u}_G and \mathbf{G}_G are the same as in (3). \mathbf{D} is a general, symmetric, 2×2 matrix. Note the \mathbf{Z} matrix takes a different form from that given in [14] and Section 9.3 of [6]. It has been changed to make the calculations more manageable.

As before the variance calculations require the estimated covariance matrix \mathbf{M}^{-1} . This matrix is now even larger than that of Section 2. Therefore we propose a streamlined approach, again based around a block decomposition of \mathbf{M} :

$$\begin{aligned} \widehat{\sigma}_\varepsilon^{-2} \begin{bmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} \\ \mathbf{M}_{21} & \mathbf{M}_{22} \end{bmatrix} &\equiv \widehat{\sigma}_\varepsilon^{-2} \left[\begin{array}{c|ccc} \mathbf{M}_{11} & \mathbf{H}_1 & \cdots & \mathbf{H}_m \\ \mathbf{H}_1^T & \mathbf{L}_1 & & \mathbf{0} \\ \vdots & & \ddots & \\ \mathbf{H}_m^T & \mathbf{0} & & \mathbf{L}_m \end{array} \right], \\ \text{where } \mathbf{M}_{11} &= \begin{bmatrix} \sum_{i=1}^m \mathbf{X}_i^T \mathbf{X}_i & \sum_{i=1}^m \mathbf{X}_i^T \mathbf{Z}_i \\ \sum_{i=1}^m \mathbf{Z}_i^T \mathbf{X}_i & \sum_{i=1}^m \mathbf{Z}_i^T \mathbf{Z}_i + \lambda \mathbf{I}_{K_f} \end{bmatrix}, & \lambda &= \frac{\widehat{\sigma}_\varepsilon^2}{\widehat{\sigma}_u^2}, \\ \mathbf{H}_i &= \begin{bmatrix} \mathbf{X}_i^T \mathbf{T}_i & \mathbf{X}_i^T \mathbf{W}_i \\ \mathbf{Z}_i^T \mathbf{T}_i & \mathbf{Z}_i^T \mathbf{W}_i \end{bmatrix} \\ \text{and } \mathbf{L}_i &= \begin{bmatrix} \mathbf{T}_i^T \mathbf{T}_i + \widehat{\sigma}_\varepsilon^2 \hat{\mathbf{D}}^{-1} & \mathbf{T}_i^T \mathbf{W}_i \\ \mathbf{W}_i^T \mathbf{T}_i & \mathbf{W}_i^T \mathbf{W}_i + \tau \mathbf{I}_{K_g} \end{bmatrix}, & \tau &= \frac{\widehat{\sigma}_\varepsilon^2}{\widehat{\sigma}_v^2}. \end{aligned}$$

In this model the matrix \mathbf{M}_{22} , the contribution from the random subject component, is block-diagonal. It can therefore be inverted in $O(m)$ steps, enabling streamlined variance calculations, again using (4):

$$\begin{aligned} \mathbf{M}^{-1} &= \widehat{\sigma}_\varepsilon^2 \begin{bmatrix} \mathbf{M}^{11} & \mathbf{M}^{12} \\ \mathbf{M}^{21} & \mathbf{M}^{22} \end{bmatrix} \\ \text{where } \mathbf{M}^{11} &= \left(\mathbf{M}_{11} - \sum_{i=1}^m \mathbf{H}_i \mathbf{L}_i^{-1} \mathbf{H}_i^T \right)^{-1}, \end{aligned}$$

$$\begin{aligned} \mathbf{M}^{12} &= [-\mathbf{M}^{11}\mathbf{H}_1\mathbf{L}_1^{-1} \ \cdots \ -\mathbf{M}^{11}\mathbf{H}_m\mathbf{L}_m^{-1}], \\ \text{and } \mathbf{M}_{ii}^{22} &= \mathbf{L}_i^{-1} + \mathbf{L}_i^{-1}\mathbf{H}_i^T\mathbf{M}^{11}\mathbf{H}_i\mathbf{L}_i^{-1}, \quad 1 \leq i \leq m. \end{aligned}$$

Since \mathbf{M}_{11} , \mathbf{H}_i and \mathbf{L}_i have dimensions much smaller than m , the complexity of the matrix calculations required in these submatrices does not increase as m increases. Therefore the calculations with the highest order of complexity are the summations in \mathbf{M}_{11} and \mathbf{M}^{11} , and the calculation of the m relevant submatrices of \mathbf{M}^{12} , \mathbf{M}^{21} and \mathbf{M}^{22} . This renders the whole process as $O(m)$, preserving the improvement over naïve inversion that we already witnessed for random intercepts.

FIGURE 3 NEAR HERE

Figure 3: *Heights of 190 girls with Acute Lymphoblastic Leukaemia. We have highlighted two subjects (one column for each) with different shaped responses. The top row shows the subjects' individual response profiles, and the bottom row the estimated subject-specific fitted responses (solid lines). The bottom panels also show $\pm 2 \times$ estimated standard error bars from both model (7) (dashed lines) and a random intercepts model (dotted lines).*

Durban *et al.* [5] fit a semiparametric model with subject-specific curves to longitudinal data on the heights of 190 girls with Acute Lymphoblastic Leukaemia. We fit a similar model in which height is modelled as a smooth function of age:

$$\text{height}_{ij} = f(\text{age}_{ij}) + g_i(\text{age}_{ij}) + \varepsilon_{ij}, \quad (7)$$

where f and g_i are modelled using radial cubic splines with 15 and 10 knots respectively.

Figure 3 shows some fitted functions and error bars from this model, superimposed on the data. We have drawn out two subjects' responses that differ from each other in order to show how the subject-specific curves model can produce different estimated smooth curves for each subject. We have plotted error bars according to the subject-specific curves model (7) and also a random intercepts model such as (2), both obtained via the streamlined method. The random-intercepts error bars are much wider than those of the subject-specific curves model.

4 Closing Remarks

The use of semiparametric regression in the analysis of longitudinal data has become commonplace in the last decade. The sample sizes of datasets arising from longitudinal studies are also on the increase, and semiparametric regression methods are susceptible to breakdown if implemented naïvely. The methods developed in this article overcome such problems, even for massive datasets, and therefore are recommended for general practice.

Appendix: R implementation

In this Appendix we provide R code that demonstrates the streamlined variance calculations described in Section 2, for data simulated according to (5). In December 2006 this code successfully ran on version 2.4.0 of R under the Linux operating system on the second author's computer and produced the output

Estimated effect of x=1 and standard error is:

```
est.    s.e.
0.3080 0.0661
```

and the plot shown in Figure 4. A text file containing this code is available on request from the second author. At the time of writing his e-mail address is `mwand@uow.edu.au`.

FIGURE 4 NEAR HERE

Figure 4: *Plot obtained from running the code in this Appendix.*

Load required packages:

```
library(nlme) ; library(lattice)
```

Generate synthetic data according to (5):

```
set.seed(39402) ; m <- 250 ; nVals <- sample(1:4,m,replace=TRUE)
betaVal <- 0.3 ; sigU <- 0.5 ; sigEps <- 0.2
f <- function(x) return(-sin(2*pi*x))
U <- rnorm(m,0,sigU)
sVals <- NULL ; xVals <- NULL
Uvals <- NULL ; idNum <- NULL
for (i in 1:m)
{
  idNum <- c(idNum,rep(i,nVals[i]))
  stt <- runif(1,0,1-0.05*(nVals[i]-1))
  sVals <- c(sVals,seq(stt,by=0.05,length=nVals[i]))
  xCurr <- sample(c(0,1),1)
  xVals <- c(xVals,rep(xCurr,nVals[i]))
  Uvals <- c(Uvals,rep(U[i],nVals[i]))
}
epsVals <- rnorm(sum(nVals),0,sigEps)
yVals <- f(sVals) + betaVal*xVals + Uvals + epsVals
```

Fit an additive mixed model:

```
numKnots <- 15 ; p <- 1
knots <- quantile(unique(sVals),seq(0,1,length=
  (numKnots+2))[-c(1,(numKnots+2))])
X <- cbind(rep(1,length(yVals)),sVals,xVals)
svd.Omega <- svd(abs(outer(knots,knots,"-"))^3)
matrix.sqrt.Omega <- t(svd.Omega$V%*%(t(svd.Omega$U)*sqrt(svd.Omega$d)))
ZSpline <- t(solve(matrix.sqrt.Omega,t(abs(outer(sVals,knots,"-"))^3)))
groupVec <- factor(rep(1,length(yVals)))
ZBlock <- list(list(groupVec=pdIdent(~ZSpline-1)),list(idNum=pdIdent(~1)))
ZBlock <- unlist(ZBlock,recursive=FALSE)
dataFr <- groupedData(yVals~xVals|groupVec,
  data=data.frame(yVals,X,ZSpline,idNum))
fit <- lme(yVals~-1+X,data=dataFr,random=ZBlock)
betaHat <- fit$coef$fixed
uHat <- unlist(fit$coef$random)
uSplineHat <- uHat[1:ncol(ZSpline)]
```

Perform streamlined variance estimation for variability bars and standard error of binary predictor:

```
sig.eps.hat <- fit$sigma
sig.u.hat <- intervals(fit)$reStruct$groupVec$est
sig.U.hat <- intervals(fit)$reStruct$idNum$est
Mllinv <- diag(c(rep(0,p+2),rep(sig.eps.hat^2/sig.u.hat^2,numKnots)))
Mllinv <- Mllinv + crossprod(cbind(X,ZSpline))
M.21 <- rowsum(cbind(X,ZSpline),idNum)
for (i in 1:m)
```



```

{
  h.i <- M.21[i,] ; n.i <- M.21[i,1]
  M11inv <- (M11inv - sig.U.hat^2*tcrossprod(h.i)/
            (sig.U.hat^2*n.i+sig.eps.hat^2))
}
covMat <- (sig.eps.hat^2)*solve(M11inv)

```

Report results:

```

cat("\n Estimated effect of x=1 and standard error is:\n\n")
outp <- c(betaHat[3],sqrt(covMat[3,3]))
names(outp) <- c("est.", "s.e.")
print(signif(outp,3)) ; cat("\n\n")
ng <- 101 ; sValsg <- seq(0,1,length=ng)
ZgSpline <- t(solve(matrix.sqrt.Omega,t(abs(outer(sValsg,knots,"-")^3))))
plotMatrix0 <- cbind(rep(1,ng),sValsg,rep(0,ng),ZgSpline)
fhatgREML <- plotMatrix0 %*% c(betaHat,uSplineHat)
Var0 <- diag(plotMatrix0%*%covMat%*%t(plotMatrix0))
Var1 <- 2*plotMatrix0%*%covMat[,3] + covMat[3,3]
xLabs <- paste("x =",as.character(xVals))
pobj <- xyplot(yVals~sVals|xLabs,groups=idNum,xlab="value of s",
             ylab="value of y",subscripts=TRUE,
             panel=function(x,y,subscripts,groups)
             {
               panel.grid() ; panel.superpose(x,y,subscripts,groups,
                                               type="b",col="grey60",pch=16)
               panelInd <- any(xVals[subscripts]==1)
               panel.xyplot(sValsg,fhatgREML+panelInd*betaHat[3],
                           lwd=3,type="l",col="black")
               for (multFac in c(-2,2))
                 panel.xyplot(sValsg,fhatgREML+panelInd*betaHat[3]+multFac
                              *sqrt(Var0+panelInd*Var1),lty=2,lwd=2,type="l",col="black")
             })
print(pobj)

```

References

1. Wang, Y. (1998). Smoothing spline models with correlated random errors. *Journal of the American Statistical Association*, **93**, 341–348.
2. Brumback, B.A. and Rice, J.A. (1999). Smoothing spline models for the analysis of nested and crossed samples of curves (with discussion). *Journal of the American Statistical Association*, **93**, 961–994.
3. Verbyla, A.P., Cullis, B.R., Kenward, M.G. and Welham, S.J. (1999). The analysis of designed experiments and longitudinal data by using smoothing splines (with discussion). *Applied Statistics*, **48**, 269–312.
4. Lin, X. and Carroll, R.J. (2001). Semiparametric regression for clustered data. *Biometrika*, **88**, 1179–1865.
5. Rice, J.A. and Wu, C.O. (2001). Nonparametric mixed effect models for unequally sampled noisy curves. *Biometrics*, **57**, 253–259.
6. Ruppert, D., Wand, M. P. and Carroll, R.J. (2003). *Semiparametric Regression*. New York: Cambridge University Press.
7. Wood, S.N. (2006). *Generalized Additive Models: An Introduction with R*. Boca Raton: Chapman & Hall/CRC.

8. Wu, H. and Zhang, J.-T. (2006). *Nonparametric Regression Methods for Longitudinal Data Analysis*, New York: Wiley.
9. Bachrach, L.K., Hastie, T., Wang, M-C, Narasimhan, B., Marcus R. (1999). Bone mineral acquisition in healthy Asian, Hispanic, Black and Caucasian youth. A longitudinal study. *Journal of Clinical Endocrinology and Metabolism*, **84**, 4702–12.
10. Fitzmaurice, G.M., Laird, N.M. and Ware, J.H. (2004). *Applied Longitudinal Analysis*, Hoboken, New Jersey: Wiley.
11. Cormen, T.H., Leiserson, C.E., Rivest, R.L. and Stein, C. (2001). *Introduction to Algorithms*, Cambridge, Massachusetts: MIT Press.
12. Harville, D.A. (2000). *Matrix Algebra From a Statistician's Perspective*. New York: Springer.
13. R Development Core Team (2006). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org>.
14. Durban, M., Harezlak, J., Wand, M.P. and Carroll, R.J. (2005). Simple fitting of subject-specific curves for longitudinal data. *Statistics in Medicine*, **24**, 1153–1167.
15. Donnelly, C.A., Laird, N.M. and Ware, J.H. (1995). Prediction and creation of smooth curves for temporally correlated longitudinal data. *Journal of the American Statistical Association*, **90**, 984–989.
16. Guo, W. (2002). Functional mixed effects models. *Biometrics*, **58**, 121–128.