



A new REML (parameter expanded) EM algorithm for linear mixed models

S. M. Diffey¹ A. B. Smith^{1,*} A. H. Welsh² and B. R. Cullis¹

University of Wollongong, Australian National University

Summary

Linear mixed models are regularly applied to animal and plant breeding data to evaluate genetic potential. Residual maximum likelihood (REML) is the preferred method for estimating variance parameters associated with this type of model. Typically an iterative algorithm is required for the estimation of variance parameters. Two algorithms which can be used for this purpose are the expectation-maximisation (EM) algorithm and the parameter expanded EM (PX-EM) algorithm. Both, particularly the EM algorithm, can be slow to converge when compared to a Newton-Raphson type scheme such as the average information (AI) algorithm. The EM and PX-EM algorithms require specification of the complete data, including the incomplete and missing data. We consider a new incomplete data specification based on a conditional derivation of REML. We illustrate the use of the resulting new algorithm through two examples: a sire model for lamb weight data and a balanced incomplete block soybean variety trial. In the cases where the AI algorithm failed, a REML PX-EM based on the new incomplete data specification converged in 28% to 30% fewer iterations than the alternative REML PX-EM specification. For the soybean example a REML EM algorithm using the new specification converged in fewer iterations than the current standard specification of a REML PX-EM algorithm. The new specification integrates linear mixed models, Henderson's mixed model equations, REML and the REML EM algorithm into a cohesive framework.

Key words: variance components; mixed model equations

1. Introduction

Linear mixed models are routinely applied to biological data. The preferred method for estimating the parameters associated with these models is residual maximum likelihood (REML) (Patterson & Thompson 1971). Many statistical software packages available for the REML estimation of parameters associated with linear mixed models implement a Newton-Raphson type algorithm such as the Fisher scoring algorithm or the AI algorithm (Gilmour, Thompson & Cullis 1995). There are two problems with these types of algorithms. Firstly, successive iterations of these algorithms are not guaranteed to increase the residual log-likelihood function. Secondly, parameter updates may not remain in their parameter space. Either problem may result in the algorithm failing to converge to a solution unless some modification of a Newton-Raphson type algorithm is implemented (see Smyth 1998 for details on one such modification).

*Author to whom correspondence should be addressed.

¹Centre for Bioinformatics and Biometrics, University of Wollongong, Wollongong, NSW 2522, Australia

²Mathematical Sciences Institute, Australian National University, Canberra, ACT 0200, Australia

e-mail: simon@apexbiometry.com

Acknowledgements. Simon Diffey gratefully acknowledges the financial support provided by the Grains Research Development Corporation and the New South Wales Department of Primary Industries while undertaking a PhD programme.

The EM algorithm (Dempster, Laird & Rubin 1977) and the parameter expanded EM (PX-EM) algorithm (Liu, Rubin & Wu 1998) have the desirable properties of monotonic convergence and the parameter updates remaining in the parameter space. Dempster, Laird & Rubin (1977) and Liu, Rubin & Wu (1998) only consider maximum likelihood estimation although REML estimation is briefly mentioned by Dempster, Laird & Rubin (1977) in their variance components example. A REML EM algorithm to estimate the parameters associated with a linear mixed model applied to longitudinal data was considered by Laird & Ware (1982) and Foulley, Jaffrezic & Robert-Cranie (2000). A REML PX-EM algorithm for linear mixed models was considered by Foulley & Van Dyk (2000). In all the aforementioned papers REML estimates of variance parameters were obtained by considering the vector of fixed effects to have a variance tending to infinity. This is the approach briefly described by Dempster, Laird & Rubin (1977).

The purpose of this paper is to present an alternative REML (PX-)EM algorithm for linear mixed models. The derivation of this algorithm can be used as a foundation for key concepts associated with the estimation of linear mixed model variance parameters. These concepts include Henderson’s mixed model equations, REML, and iterative schemes. Two examples are considered and in both cases the alternative algorithm converges in fewer iterations than the standard algorithm.

This paper proceeds in four parts. First, we provide statistical background on the linear mixed model, Henderson’s mixed model equations (Henderson et al. 1959), and a conditional derivation of REML (Verbyla 1990). Second, we consider two alternative complete data specifications for a REML EM algorithm. One is the complete data specification used in current specifications of a REML EM algorithm for linear mixed models and the other is a REML EM algorithm based on a new incomplete data specification. We outline a proof that this new algorithm will converge to a (local) maximum of the residual log-likelihood function. Third, we consider a REML PX-EM algorithm for linear mixed models using the new specification. Fourth, we illustrate the use of the current and new REML EM and REML PX-EM algorithms for linear mixed models by considering two examples.

2. Statistical background

2.1. Linear mixed model formulation

The linear mixed model can be written as

$$y = X\beta + Zu + e, \tag{1}$$

where y is an $n \times 1$ vector of observed data, β is a $t \times 1$ vector of fixed effects and X is its associated design matrix which is assumed to be of full column rank. The $n \times b$ design matrix Z is associated with the vector of random effects $u \sim N\{\theta, G(\gamma)\}$. The matrix $G(\gamma)$ is positive definite and depends on the parameter vector γ . It is assumed that $e \sim N(\theta, R)$, where $R = \sigma_e^2 \Sigma(\phi)$, and the matrix $\Sigma(\phi)$ is assumed to be positive definite and depends on the parameter vector ϕ . We will assume that n is larger than b but recognise that this is not always the case. The joint distribution of y , u , and e is

$$\begin{pmatrix} y \\ u \\ e \end{pmatrix} \sim N \left\{ \begin{pmatrix} X\beta \\ \theta \\ \theta \end{pmatrix}, \begin{pmatrix} H & ZG & R \\ GZ^T & G & \theta \\ R & \theta & R \end{pmatrix} \right\}, \tag{2}$$

where $H = ZGZ^T + R$. The superscript \top denotes matrix transpose. The aim is to estimate the parameter vector $\theta = (\beta^T, \kappa^T)^T$, where $\kappa = (\gamma^T, \sigma_\epsilon^2, \phi^T)^T$ is a vector of variance parameters.

2.2. Henderson's mixed model equations

Components of Henderson's mixed model equations are widely used in schemes for computing REML estimates of variance parameters. A discussion on their convenience for this purpose can be found in section 7.6 of Searle, Casella & McCulloch (2006) and references therein.

Henderson's mixed model equations are derived by considering the log joint density function of y and u , i.e.

$$\log\{f(y, u; \theta)\} = \log\{f(y|u; \theta)\} + \log\{f(u; \gamma)\}.$$

Obtaining the conditional distribution $f(y|u; \theta)$ requires a well known result in multivariate normal statistics and the joint distribution given in (2). The conditional distribution $f(y|u; \theta)$ is

$$y|u \sim N(X\beta + Zu, R).$$

The log joint density function of y and u is therefore (excluding constants)

$$\begin{aligned} \log\{f(y, u; \theta)\} = & -\frac{1}{2} \left[\log\{\det(R)\} + \log\{\det(G)\} + u^T G^{-1} u \right. \\ & \left. + (y - X\beta - Zu)^T R^{-1} (y - X\beta - Zu) \right]. \end{aligned} \quad (3)$$

Differentiating (3) with respect to β and u and equating to zero gives Henderson's mixed model equations

$$\begin{aligned} X^T R^{-1} X \hat{\beta} + X^T R^{-1} Z \tilde{u} &= X^T R^{-1} y, \\ Z^T R^{-1} X \hat{\beta} + (Z^T R^{-1} Z + G^{-1}) \tilde{u} &= Z^T R^{-1} y, \end{aligned}$$

where $\hat{\beta}$ is the best linear unbiased estimator (BLUE) of β and \tilde{u} is the best linear unbiased predictor (BLUP) of u . When expressed in matrix notation the coefficient matrix of Henderson's mixed model equations can be written as

$$C = \begin{pmatrix} X^T R^{-1} X & X^T R^{-1} Z \\ Z^T R^{-1} X & Z^T R^{-1} Z + G^{-1} \end{pmatrix} = \begin{pmatrix} C_{XX} & C_{XZ} \\ C_{ZX} & C_{ZZ} \end{pmatrix}$$

with the inverse of the coefficient matrix as

$$\begin{aligned} C^{-1} &= \begin{pmatrix} (X^T H^{-1} X)^{-1} & -(X^T H^{-1} X)^{-1} X^T H^{-1} Z G \\ -G^T Z^T H^{-1} X (X^T H^{-1} X)^{-1} & (Z^T S Z + G^{-1})^{-1} \end{pmatrix} \\ &= \begin{pmatrix} C^{XX} & C^{XZ} \\ C^{ZX} & C^{ZZ} \end{pmatrix}, \end{aligned}$$

where $S = R^{-1} - R^{-1} X (X^T R^{-1} X)^{-1} X^T R^{-1}$. The matrix C^{ZZ} is often referred to as the prediction error variance. We will define the conditional variances required for computation of the E-step in the REML EM and REML PX-EM algorithm for linear mixed models in terms related to the coefficient matrix of Henderson's mixed model equations.

2.3. A conditional derivation of REML

The preferred method for estimating the variance parameters associated with a linear mixed model is REML. The original reference for REML is the paper by Patterson & Thompson (1971); however the new implementation of a REML EM and REML PX-EM algorithm for linear mixed models is based on the conditional derivation of REML presented by Verbyla (1990). The conditional derivation of REML begins by considering the transformation

$$L^T y = \begin{pmatrix} L_1^T y \\ L_2^T y \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix},$$

where $L = (L_1, L_2)$ is a non-singular matrix, L_1 and L_2 are $n \times t$ and $n \times (n - t)$ matrices respectively, both of full-column rank and chosen to satisfy $L_1^T X = I_t$ and $L_2^T X = \theta$. The distribution of the transformed data is

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \sim N \left\{ \begin{pmatrix} \beta \\ \theta \end{pmatrix}, \begin{pmatrix} L_1^T H L_1 & L_1^T H L_2 \\ L_2^T H L_1 & L_2^T H L_2 \end{pmatrix} \right\}.$$

The joint distribution of y_1 and y_2 can be expressed as the product of the marginal distribution of y_2 and the conditional distribution $y_1|y_2$. Therefore, the log-likelihood function of $L^T y$ can be expressed as

$$\ell(\theta; L^T y) = \ell(\kappa; y_2) + \ell(\theta; y_1|y_2). \tag{4}$$

The vector y_1 is of length t and the conditional log-likelihood function $\ell(\theta; y_1|y_2)$ is used to estimate the fixed effects, after which there is no information left for the estimation of variance parameters. The log-likelihood function associated with the marginal density of κ , i.e., $\ell(\kappa; y_2)$ is used for the REML estimation of variance parameters and can be expressed as

$$\ell(\kappa; y_2) = -\frac{1}{2} \left[\log\{\det(L_2^T H L_2)\} + y_2^T (L_2^T H L_2)^{-1} y_2 \right] \tag{5}$$

Using the identities (see Verbyla 1990) we can show that

$$P = L_2 (L_2^T H L_2)^{-1} L_2^T = H^{-1} - H^{-1} X (X^T H^{-1} X)^{-1} X^T H^{-1},$$

$$\log\{\det(L_2^T H L_2)\} = \log\{\det(L^T L)\} + \log\{\det(H)\} + \log\{\det(X^T H^{-1} X)\},$$

and then if we ignore the constant $\log\{\det(L^T L)\}$, (5) can be written as

$$\ell(\kappa; y_2) = -\frac{1}{2} \left[\log\{\det(H)\} + \log\{\det(X^T H^{-1} X)\} + y^T P y \right], \tag{6}$$

which is the the form of the residual log-likelihood function most commonly presented.

3. Two alternative incomplete data specifications

3.1. Current derivations of a REML EM algorithm for the linear mixed model

Current derivations of REML EM and REML PX-EM algorithms for linear mixed models are implemented by considering the vector of fixed effects β to be a random effect with variance tending to infinity. We will refer to this implementation as the random effects approach. Using this approach we will assume that $\beta \sim N(\theta, B)$ and consider the complete data to be defined as $y_c^{(P)} = (y^T, \beta^T, u^T)^T$ where the superscript P in parentheses on y_c is used to distinguish this complete data specification from the new complete data specification

to be presented shortly. The random effects approach is the approach briefly described by Dempster, Laird & Rubin (1977) and followed by others including Laird & Ware (1982); Foulley, Jaffrezic & Robert-Cranie (2000) and Foulley & Van Dyk (2000). We define the joint distribution of y , β , u , and e as

$$\begin{pmatrix} y \\ \beta \\ u \\ e \end{pmatrix} \sim N \left\{ \begin{pmatrix} \theta \\ \theta \\ \theta \\ \theta \end{pmatrix}, \begin{pmatrix} F & XB & ZG & R \\ BX^\top & B & \theta & \theta \\ GZ^\top & \theta & G & \theta \\ R & \theta & \theta & R \end{pmatrix} \right\}, \quad (7)$$

where

$$F = XBX^\top + ZGZ^\top + R = XBX^\top + H$$

and $F^{-1} = H^{-1} - H^{-1}X(B^{-1} + X^\top H^{-1}X)^{-1}X^\top H^{-1}$. The random effect approach assumes that the variance of β tends to infinity, i.e., $B^{-1} \rightarrow \theta$. Therefore

$$F^{-1} \rightarrow H^{-1} - H^{-1}X(X^\top H^{-1}X)^{-1}X^\top H^{-1} = P.$$

We will assume that if $B^{-1} \rightarrow \theta$ then F^{-1} can be approximated by P . The complete data log density function can be written as

$$\log\{f(y_c^{(P)}; \theta)\} = -\frac{1}{2} \left[\log\{\det(R)\} + \log\{\det(G)\} + u^\top G^{-1}u + e^\top R^{-1}e \right], \quad (8)$$

which is equivalent to the log joint density function of y and u used to derive Henderson's mixed model equations and presented in (3) except we have made the substitution $e = y - X\beta - Zu$.

The E-step of a REML EM algorithm using $y_c^{(P)}$ involves taking the conditional expectation of the complete data log density function in (8) with respect to y and at the current iterate of θ , denoted $\theta^{(w)}$, i.e.,

$$\begin{aligned} Q^{(P)}(\theta; \theta^{(w)}) &= E \left[\log\{f(y_c^{(P)}; \theta)\} | y; \theta^{(w)} \right] \\ &= E \left[\log\{f(e; \theta)\} | y; \theta^{(w)} \right] + E \left[\log\{f(u; \gamma)\} | y; \theta^{(w)} \right] \\ &= Q^{(P)}(\sigma_e^2, \phi; \theta^{(w)}) + Q^{(P)}(\gamma; \theta^{(w)}). \end{aligned}$$

Computation of $Q^{(P)}(\sigma_e^2, \phi; \theta^{(w)})$ and $Q^{(P)}(\gamma; \theta^{(w)})$ requires the conditional distributions $e|y$ and $u|y$ respectively. From (7) it can be shown that

$$\begin{aligned} e|y &\sim N(RPy, WC^{-1}W^\top), \\ u|y &\sim N(GZ^\top Py, C^{ZZ}), \end{aligned}$$

where $W = (X, Z)$ is an $n \times (t + b)$ matrix. Using these two results we can write

$$\begin{aligned} Q^{(P)}(\sigma_e^2, \phi; \theta^{(w)}) &= -\frac{1}{2} \left[\log\{\det(R)\} + \tilde{e}^{(w)\top} R^{-1} \tilde{e}^{(w)} + \text{tr}(R^{-1}WC^{-1(w)}W^\top) \right], \\ Q^{(P)}(\gamma; \theta^{(w)}) &= -\frac{1}{2} \left[\log\{\det(G)\} + \tilde{u}^{(w)\top} G^{-1} \tilde{u}^{(w)} + \text{tr}(G^{-1}C^{ZZ(w)}) \right], \end{aligned}$$

where $\tilde{e}^{(w)} = R^{(w)}P^{(w)}y$ and $\tilde{u}^{(w)} = G^{(w)}Z^\top P^{(w)}y$.

The M-step of a REML EM algorithm for linear mixed models using $y_c^{(P)}$ involves maximising $Q^{(P)}(\sigma_e^2, \phi; \theta^{(w)})$ and $Q^{(P)}(\gamma; \theta^{(w)})$ with respect to σ_e^2 , ϕ and γ respectively. Noting that $R = \sigma_e^2 \Sigma$ the updating equation for σ_e^2 is

$$\sigma_e^{2(w+1)} = \frac{1}{n} [\tilde{\mathbf{e}}^{(w)\top} \boldsymbol{\Sigma}^{-1(w)} \tilde{\mathbf{e}}^{(w)} + \text{tr}(\boldsymbol{\Sigma}^{-1(w)} \mathbf{W} \mathbf{C}^{-1(w)} \mathbf{W}^\top)].$$

The updating equations for $\phi_{ij}, i = 1, \dots, n, j = 1, \dots, n, i \geq j$ and $\gamma_{gh}, g = 1, \dots, b, h = 1, \dots, b, g \geq h$, depend on the specific form of $\boldsymbol{\Sigma}$ and \mathbf{G} but can be found by equating the following partial derivatives to zero,

$$\begin{aligned} \frac{\partial Q^{(P)}(\sigma_e^2, \boldsymbol{\phi}; \boldsymbol{\theta}^{(w)})}{\partial \phi_{ij}} &= -\frac{1}{2} \left[\text{tr} \left(\boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \phi_{ij}} \right) - \frac{1}{\sigma_e^2} \tilde{\mathbf{e}}^{(w)\top} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \phi_{ij}} \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{e}}^{(w)} \right. \\ &\quad \left. - \frac{1}{\sigma_e^2} \text{tr} \left(\boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \phi_{ij}} \boldsymbol{\Sigma}^{-1} \mathbf{W} \mathbf{C}^{-1(w)} \mathbf{W}^\top \right) \right], \\ \frac{\partial Q^{(P)}(\boldsymbol{\gamma}; \boldsymbol{\theta}^{(w)})}{\partial \gamma_{gh}} &= -\frac{1}{2} \left[\text{tr} \left(\mathbf{G}^{-1} \frac{\partial \mathbf{G}}{\partial \gamma_{gh}} \right) - \tilde{\mathbf{u}}^{(w)\top} \mathbf{G}^{-1} \frac{\partial \mathbf{G}}{\partial \gamma_{gh}} \mathbf{G}^{-1} \tilde{\mathbf{u}}^{(w)} \right. \\ &\quad \left. - \text{tr} \left(\mathbf{G}^{-1} \frac{\partial \mathbf{G}}{\partial \gamma_{gh}} \mathbf{G}^{-1} \mathbf{C}^{ZZ(w)} \right) \right]. \end{aligned}$$

An estimate for $\boldsymbol{\beta}$ can be obtained at the end of the REML EM algorithm by using the generalised least squares estimate of $\boldsymbol{\beta}$. Although we have presented the random effects approach to implementing a REML EM algorithm it is worth noting that Cullis, Smith & Thompson (2004) and Knight (2008) derived exactly the same parameter updates by considering $\boldsymbol{\beta}$ as fixed. The difference between their fixed effects approach and the random effects approach was that in the E-step the conditional expectation is with respect to y_2 rather than \mathbf{y} .

3.2. A new incomplete data specification for a REML EM algorithm for the linear mixed model

We consider a linear mixed model specified in the same way as in (1). We define a new complete data specification $\mathbf{y}_c^{(U)} = (\mathbf{y}_2^\top, \mathbf{u}^\top)^\top$. This differs from the previously published specification $\mathbf{y}_c^{(P)} = (\mathbf{y}^\top, \boldsymbol{\beta}^\top, \mathbf{u}^\top)^\top$ by considering the incomplete data to be the transformed observed data vector associated with the marginal log density function in (6) which is used for REML estimation of the variance parameters. To form the complete data log density function requires the joint distribution of y_2 and \mathbf{u} , i.e.,

$$\begin{pmatrix} y_2 \\ \mathbf{u} \end{pmatrix} \sim \text{N} \left\{ \begin{pmatrix} \boldsymbol{\theta} \\ \boldsymbol{\theta} \end{pmatrix}, \begin{pmatrix} \mathbf{L}_2^\top \mathbf{H} \mathbf{L}_2 & \mathbf{L}_2^\top \mathbf{Z} \mathbf{G} \\ \mathbf{G} \mathbf{Z}^\top \mathbf{L}_2 & \mathbf{G} \end{pmatrix} \right\}.$$

The complete data log density function using $\mathbf{y}_c^{(U)}$ can be written (excluding constants) as

$$\begin{aligned} \log\{f(\mathbf{y}_c^{(U)}; \boldsymbol{\kappa})\} &= -\frac{1}{2} \left[\log\{\det(\mathbf{L}_2^\top \mathbf{R} \mathbf{L}_2)\} + \log\{\det(\mathbf{G})\} + \mathbf{u}^\top \mathbf{G}^{-1} \mathbf{u} \right. \\ &\quad \left. + (\mathbf{y} - \mathbf{Z} \mathbf{u})^\top \mathbf{S} (\mathbf{y} - \mathbf{Z} \mathbf{u}) \right]. \end{aligned} \tag{9}$$

The relationship between the complete data log density functions based on $\mathbf{y}_c^{(U)}$ and $\mathbf{y}_c^{(P)}$ can be shown to be

$$\log\{f(\mathbf{y}_c^{(P)}; \boldsymbol{\theta})\} = \log\{f(\mathbf{y}_c^{(U)}; \boldsymbol{\kappa})\} + \log\{f(\mathbf{y}_1 | y_2; \boldsymbol{\theta})\},$$

Table 1 Summary of REML EM algorithm complete data specifications for the linear mixed model presented in (1).

Complete data	β fixed or random	References
$\mathbf{y}_c^{(P)} = (\mathbf{y}^\top, \beta^\top, \mathbf{u}^\top)^\top$	random	Dempster, Laird & Rubin (1977); Laird & Ware (1982); Foulley, Jaffrezic & Robert-Cranie (2000)
$\mathbf{y}_c^{(P)} = (\mathbf{y}^\top, \mathbf{u}^\top)^\top$	fixed	Knight (2008); Cullis, Smith & Thompson (2004)
$\mathbf{y}_c^{(U)} = (\mathbf{y}_2^\top, \mathbf{u}^\top)^\top$	fixed	Not published previously as far as we are aware

which is similar to the relation presented in (4) in the Verbyla (1990) conditional derivation of REML. This suggests that using the complete data specification $\mathbf{y}_c^{(P)}$ should be similar to using $\ell(\theta; \mathbf{L}^\top \mathbf{y})$ for the REML estimation of variance parameters.

For the complete data specification $\mathbf{y}_c^{(U)}$ the E-step involves taking the conditional expectation of the complete data log density function in (9) with respect to \mathbf{y}_2 and at the current iterate of κ , denoted $\kappa^{(w)}$, i.e.,

$$\begin{aligned} Q^{(U)}(\kappa; \kappa^{(w)}) &= E[\log\{f(\mathbf{y}_c^{(U)}; \kappa)\} | \mathbf{y}_2; \kappa^{(w)}] \\ &= E[\log\{f(\mathbf{L}_2^\top \mathbf{e}; \sigma_e^2, \phi)\} | \mathbf{y}_2; \kappa^{(w)}] + E[\log\{f(\mathbf{u}; \gamma)\} | \mathbf{y}_2; \kappa^{(w)}] \\ &= Q^{(U)}(\sigma_e^2, \phi; \kappa^{(w)}) + Q^{(U)}(\gamma; \kappa^{(w)}). \end{aligned}$$

Computation of $Q^{(U)}(\sigma_e^2, \phi; \kappa^{(w)})$ and $Q^{(U)}(\gamma; \kappa^{(w)})$ only requires the conditional distribution $\mathbf{u} | \mathbf{y}_2$ which is equivalent to $\mathbf{u} | \mathbf{y}$ when applying the random effects approach to $\mathbf{y}_c^{(P)} = (\mathbf{y}^\top, \beta^\top, \mathbf{u}^\top)^\top$. We can write

$$\begin{aligned} Q^{(U)}(\sigma_e^2, \phi; \kappa^{(w)}) &= -\frac{1}{2} \left[\log\{\det(\mathbf{L}_2^\top \mathbf{R} \mathbf{L}_2)\} + (\mathbf{y} - \mathbf{Z}\tilde{\mathbf{u}}^{(w)})^\top \mathbf{S}(\mathbf{y} - \mathbf{Z}\tilde{\mathbf{u}}^{(w)}) \right. \\ &\quad \left. + \text{tr}(\mathbf{Z}^\top \mathbf{S} \mathbf{Z} \mathbf{C}^{\text{ZZ}(w)}) \right], \\ Q^{(U)}(\gamma; \kappa^{(w)}) &= Q^{(P)}(\gamma; \theta^{(w)}). \end{aligned}$$

The M-step of a REML EM algorithm for linear mixed models using $\mathbf{y}_c^{(U)}$ involves maximising $Q^{(U)}(\sigma_e^2, \phi; \kappa^{(w)})$ and $Q^{(U)}(\gamma; \kappa^{(w)})$ with respect to σ_e^2, ϕ and γ respectively.

The updating equation for σ_e^2 is

$$\sigma_e^{2(w+1)} = \frac{1}{n-t} \left[(\mathbf{y} - \mathbf{Z}\tilde{\mathbf{u}}^{(w)})^\top \mathbf{U}^{(w)} (\mathbf{y} - \mathbf{Z}\tilde{\mathbf{u}}^{(w)}) + \text{tr}(\mathbf{Z}^\top \mathbf{U}^{(w)} \mathbf{Z} \mathbf{C}^{\text{ZZ}(w)}) \right].$$

where $\mathbf{U}^{(w)} = \Sigma^{-1(w)} - \Sigma^{-1(w)} \mathbf{X}(\mathbf{X}^\top \Sigma^{-1(w)} \mathbf{X})^{-1} \mathbf{X}^\top \Sigma^{-1(w)}$. The updating equation for ϕ_{ij} , $i = 1, \dots, n$, $j = 1, \dots, n$, $i \geq j$ can be found by equating the following partial derivative to zero,

$$\begin{aligned} \frac{\partial Q^{(U)}(\sigma_e^2, \phi; \kappa^{(w)})}{\partial \phi_{ij}} &= -\frac{1}{2} \left[\text{tr} \left(\mathbf{U} \frac{\partial \Sigma}{\partial \phi_{ij}} \right) - \frac{1}{\sigma_e^2} (\mathbf{y} - \mathbf{Z}\tilde{\mathbf{u}}^{(w)})^\top \mathbf{U} \frac{\partial \Sigma}{\partial \phi_{ij}} \mathbf{U} (\mathbf{y} - \mathbf{Z}\tilde{\mathbf{u}}^{(w)}) \right. \\ &\quad \left. - \frac{1}{\sigma_e^2} \text{tr} \left(\mathbf{Z}^\top \mathbf{U} \frac{\partial \Sigma}{\partial \phi_{ij}} \mathbf{U} \mathbf{Z} \mathbf{C}^{\text{ZZ}(w)} \right) \right]. \end{aligned}$$

A summary of complete data specifications for a REML EM algorithm is provided in Table 1.

It is worth noting that firstly, although y_2 constitutes the incomplete data of $y_c^{(U)}$ it does not need to be explicitly computed. Secondly, the observed information matrix can be obtained using the method of Louis (1982) or Oakes (1999).

3.3. Outline of a proof that using $y_c^{(U)} = (y_2^\top; u^\top)^\top$ will converge to a (local) maximum of the residual log-likelihood function

The proof can be obtained by following the steps in the example of Section 3 in Dempster, Laird & Rubin (1977). We begin by defining $g(y_2; \kappa)$ as the marginal density function of y_2 and $k(u|y_2; \kappa)$ as the conditional density function of u given y_2 so that we can write

$$\log\{f(y_c^{(U)}; \kappa)\} = \log\{g(y_2; \kappa)\} + \log\{k(u|y_2; \kappa)\}. \tag{10}$$

Rearranging and taking the conditional expectation of both sides of (10) with respect to y_2 and at the current iterate of κ , denoted $\kappa^{(w)}$, i.e., the E-step of the EM algorithm, we have

$$\ell(\kappa; y_2) = Q^{(U)}(\kappa; \kappa^{(w)}) - H^{(U)}(\kappa; \kappa^{(w)}),$$

where $Q^{(U)}(\kappa; \kappa^{(w)}) = E[\log\{f(y_c^{(U)}; \kappa)\}|y_2; \kappa^{(w)}]$ and $H^{(U)}(\kappa; \kappa^{(w)}) = E[\log\{k(u|y_2; \kappa)\}|y_2; \kappa^{(w)}]$. Choosing an update of $\kappa^{(w)}$, denoted $\kappa^{(w+1)}$, such that $Q^{(U)}(\kappa; \kappa^{(w+1)}) \geq Q^{(U)}(\kappa; \kappa^{(w)})$, i.e., the M-step of the EM algorithm, will result in convergence to a (local) maximum of the residual log-likelihood function since the difference $H^{(U)}(\kappa; \kappa^{(w+1)}) - H^{(U)}(\kappa; \kappa^{(w)})$ can be bounded by Jensen’s inequality.

4. A new REML PX-EM algorithm for linear mixed models

Practical implementations imply that a REML PX-EM algorithm is to be preferred to a REML EM algorithm. The REML PX-EM algorithm, like the REML EM algorithm, has monotonic convergence and parameter updates remain in the parameter space. Furthermore, the E-step of a REML PX-EM algorithm is exactly the same as the E-step of a REML EM algorithm. The gains in using the PX-EM algorithm compared to the EM algorithm are in the rates of convergence. Liu, Rubin & Wu (1998) show that the PX-EM algorithm has a rate of convergence that is at least as fast as the EM algorithm, and often much faster.

Foulley & Van Dyk (2000) considered a REML PX-EM algorithm for linear mixed models using the random effects approach, i.e. the approach where y comprises the incomplete data and where the vector of fixed effects β is considered a random effect with variance tending to infinity. However a REML PX-EM algorithm where y_2 constitutes the incomplete data has not been considered previously.

4.1. Linear mixed model formulation when using a REML PX-EM algorithm

For the REML PX-EM algorithm the linear mixed model is reformulated by introducing an auxiliary parameter λ (say) so that the linear mixed model in (1) is expanded to

$$y = X\beta + Z\Lambda f + e, \tag{11}$$

where X and Z are $n \times t$ and $n \times b$ design matrices respectively, and $\Lambda = \Lambda(\lambda)$ is a $b \times b$ real invertible matrix which is a function of the $v \times 1$ auxiliary parameter vector λ . It is assumed that the joint distribution of f and ϵ is

$$\begin{pmatrix} f \\ e \end{pmatrix} \sim N \left\{ \begin{pmatrix} \theta \\ \theta \end{pmatrix}, \begin{pmatrix} D & \theta \\ \theta & R_* \end{pmatrix} \right\},$$

where $D = D(d)$ and $R_* = \sigma_{e_*}^2 \Sigma_*(\phi_*)$ are symmetric positive definite matrices. The marginal distribution of y under the parameter expanded model is

$$y \sim N(X\beta, H_*),$$

where $H_* = Z\Lambda D\Lambda^T Z^T + R_*$. The expanded parameter set for the model in (11) is $\mathcal{K} = (\kappa_*^T, \lambda^T)^T$ where $\kappa_* = (\sigma_{e_*}^2, d^T, \phi_*^T)^T$. The subscript $*$ is used to distinguish variance parameters associated with the linear mixed model specified in (11) and the linear mixed model specified in (1). The expanded parameter vector \mathcal{K} must satisfy two conditions:

- \mathcal{K} can be reduced to the original parameter vector $\kappa = (\sigma_e^2, \gamma^T, \phi^T)^T$ by a many-to-one reduction function. For the model specified in (11) this reduction function, denoted by \mathcal{R} is $\kappa = (\sigma_e^2, \gamma^T, \phi^T)^T = \mathcal{R}(\mathcal{K}) = (\sigma_{e_*}^2, \text{vech}(\Lambda D \Lambda^T), \phi_*^T)^T$.
- When the auxiliary parameter λ is set to its null value λ_0 (say) the expanded complete data model is reduced to the original complete data model. If we define $u = \Lambda f$ then $G = \Lambda D \Lambda^T$ and when $\Lambda(\lambda_0) = I_b$ the expanded model in (11) is reduced to the original model in (1) with $u = f$ and $G = D$.

Liu, Rubin & Wu (1998) note that the idea of parameter expansion is to perform a covariance adjustment between the imputed missing data and the known incomplete data vector. In the context of a REML PX-EM algorithm for a linear mixed model the auxiliary parameter λ is used to “correct” the estimate of G to produce an adjusted estimate. At convergence there is no longer any need to adjust the estimate of G and $\Lambda(\lambda) = I_b$. Another way to consider the auxiliary parameter in the REML PX-EM algorithm is that it performs a rotation on the vector of random effects in an attempt to orthogonalise the space in which variance parameters are estimated.

We have used a notation similar to that of Liu, Rubin & Wu (1998), i.e., the subscript $*$ identifies variance parameters in the parameter expanded model. However it is worth noting that $\phi_* = \phi$ and $\sigma_{e_*}^2 = \sigma_e^2$.

4.2. REML PX-EM algorithm using y_2 as the incomplete data

For a REML PX-EM algorithm for the linear mixed model we consider the complete data specification $y_{c\mathcal{X}}^{(U)} = (y_2^T, f^T)^T$ where the subscript \mathcal{X} is used to distinguish between a REML PX-EM algorithm and a REML EM algorithm. For the parameter expanded model based on $y_{c\mathcal{X}}^{(U)}$ the joint distribution of the incomplete and missing data is

$$\begin{pmatrix} y_2 \\ f \end{pmatrix} \sim N \left\{ \begin{pmatrix} \theta \\ \theta \end{pmatrix}, \begin{pmatrix} L_2^T H_* L_2 & L_2^T Z \Lambda D \\ D \Lambda^T Z^T L_2 & D \end{pmatrix} \right\}, \tag{12}$$

where $H_* = Z\Lambda D\Lambda^T Z^T + R_*$. The complete data log density function for $y_c^{(U)}$ is

$$\begin{aligned} \log\{f_{\mathcal{X}}(y_c^{(U)}; \mathcal{K})\} = & -\frac{1}{2} \left[\log\{\det(L_2^T R_* L_2)\} + \log\{\det(D)\} + f^T D^{-1} f \right. \\ & \left. + (y - Z\Lambda f)^T S_*(y - Z\Lambda f) \right], \tag{13} \end{aligned}$$

where $S_* = R_*^{-1} - R_*^{-1}X(X^\top R_*^{-1}X)^{-1}X^\top R_*^{-1}$. The E-step of the REML PX-EM algorithm involves taking the conditional expectation of the complete data log density function in (13) with respect to y_2 and at the current iterate of $\mathcal{K}^{(w)} = \{\kappa_*^{(w)\top}, (\lambda = \lambda_0)^\top\}^\top$. Setting the auxiliary parameter equal to its null value at each iteration reduces the computations in the REML PX-EM algorithm E-step to those of the REML EM algorithm E-step. We can write

$$\begin{aligned} Q_{\mathcal{X}}^{(U)}(\mathcal{K}; \mathcal{K}^{(w)}) &= E[\log\{f_{\mathcal{X}}(y_c^{(U)}; \mathcal{K})\} | y_2; \mathcal{K}^{(w)}] \\ &= E[\log\{f_{\mathcal{X}}(L_2^\top \epsilon; \mathcal{K})\} | y_2; \mathcal{K}^{(w)}] + E[\log\{f_{\mathcal{X}}(f; d)\} | y_2; \mathcal{K}^{(w)}] \\ &= Q_{\mathcal{X}}^{(U)}(\sigma_{e_*}^2, \phi_*, \lambda; \mathcal{K}^{(w)}) + Q_{\mathcal{X}}^{(U)}(d; \mathcal{K}^{(w)}), \end{aligned}$$

where

$$\begin{aligned} Q_{\mathcal{X}}^{(U)}(\sigma_{e_*}^2, \phi_*, \lambda; \mathcal{K}^{(w)}) &= -\frac{1}{2} \left[\log\{\det(L_2^\top R_* L_2)\} \right. \\ &\quad \left. + E\left\{ (y - Z\Lambda f)^\top S_* (y - Z\Lambda f) | y_2; \mathcal{K}^{(w)} \right\} \right], \\ Q_{\mathcal{X}}^{(U)}(d; \mathcal{K}^{(w)}) &= -\frac{1}{2} \left[\log\{\det(D)\} + E\{f^\top D^{-1} f | y_2; \mathcal{K}^{(w)}\} \right]. \end{aligned}$$

The M-step of the REML PX-EM algorithm involves maximising $Q_{\mathcal{X}}^{(U)}(\mathcal{K}; \mathcal{K}^{(w)})$ with respect to $\mathcal{K} = (\sigma_{e_*}^2, d^\top, \phi_*^\top, \lambda^\top)^\top$ and then applying the reduction function $\mathcal{R}(\mathcal{K}) = (\sigma_{e_*}^2, \text{vech}(\Lambda D \Lambda^\top), \phi_*^\top)^\top$, where vech refers to vector-half (see Section 12.9 in Searle 1982), to obtain estimates of $\kappa = (\sigma_e^2, \gamma^\top, \phi^\top)^\top$. Regardless of the form of R_* and D an update for the auxiliary parameter λ can be written

$$\lambda^{(w+1)} = A^{-1(w)} b^{(w)},$$

where the elements of the $b^2 \times b^2$ matrix $A^{(w)}$ are

$$a_{gh,qr}^{(w)} = \tilde{u}^{(w)\top} J_{qr}^\top Z^\top S^{(w)} Z \frac{\partial \Lambda^{(w)}}{\partial \lambda_{gh}^{(w)}} \tilde{u}^{(w)} + \text{tr} \left(J_{qr}^\top Z^\top S^{(w)} Z \frac{\partial \Lambda^{(w)}}{\partial \lambda_{gh}^{(w)}} C^{ZZ(w)} \right),$$

where g, h, q , and r are $1, \dots, b$ and J_{qr} is a $b \times b$ indicator matrix with a unit entry in row q , column r and zeros elsewhere. The elements of the $b^2 \times 1$ vector $b^{(w)}$ are

$$b_{gh}^{(w)} = y^\top S^{(w)} Z \frac{\partial \Lambda^{(w)}}{\partial \lambda_{gh}^{(w)}} \tilde{u}^{(w)}.$$

5. Results

We consider two examples: a sire model for lamb weight data and a balanced incomplete block soybean variety trial. We have chosen these two examples because the data are easily obtained and because it is relatively straight-forward to use a statistical computing environment such as R (R Core Team 2017) to implement the REML EM and REML PX-EM algorithms presented below.

For both examples we applied a REML EM algorithm using the complete data specifications $y_c^{(P)}$ and $y_c^{(U)}$, and a REML PX-EM algorithm using the complete data specifications $y_{c\mathcal{X}}^{(P)}$ and $y_{c\mathcal{X}}^{(U)}$. We do not present the parameter updating equations for the REML EM algorithms

Table 2 Number of iterations to achieve convergence for the REML EM algorithms based on the complete data specifications $\mathbf{y}_c^{(P)}$ and $\mathbf{y}_c^{(U)}$ and REML PX-EM algorithms based on $\mathbf{y}_{c\mathcal{X}}^{(P)}$ and $\mathbf{y}_{c\mathcal{X}}^{(U)}$ applied to the lamb weight and soybean variety trial yield data.

Algorithm	Spec.	Iter. (SV1)	Iter. (SV2)
Sire model for lamb weight data			
REML EM	$\mathbf{y}_c^{(P)}$	339	341
REML EM	$\mathbf{y}_c^{(U)}$	339	340
REML PX-EM	$\mathbf{y}_{c\mathcal{X}}^{(P)}$	76	77
REML PX-EM	$\mathbf{y}_{c\mathcal{X}}^{(U)}$	54	54
AI		11	fail
Soybean variety trial yield data			
REML EM	$\mathbf{y}_c^{(P)}$	18	18
REML EM	$\mathbf{y}_c^{(U)}$	14	16
REML PX-EM	$\mathbf{y}_{c\mathcal{X}}^{(P)}$	17	18
REML PX-EM	$\mathbf{y}_{c\mathcal{X}}^{(U)}$	12	13
AI		8	fail

Note: The first set of starting values (SV1) for the lamb weight data were $\sigma_e^{2(0)} = 2$ and $\sigma_u^{2(0)} = 2$ and for the soybean variety trial yield data were $\sigma_e^{2(0)} = 1$ and $\sigma_u^{2(0)} = 1$. The second set of starting values (SV2) for the lamb weight data were $\sigma_e^{2(0)} = 2$ and $\sigma_u^{2(0)} = 3$ and for the soybean variety trial yield data were $\sigma_e^{2(0)} = 8$ and $\sigma_u^{2(0)} = 4$. The abbreviations ‘‘Spec.’’ and ‘‘Iter.’’ mean specification and iterations respectively.

using the complete data specifications $\mathbf{y}_c^{(P)}$ and $\mathbf{y}_c^{(U)}$ as they can be obtained by making the substitution $\mathbf{A} = \mathbf{I}_b$ in the updating equations for the REML PX-EM algorithms based on $\mathbf{y}_{c\mathcal{X}}^{(P)}$ and $\mathbf{y}_{c\mathcal{X}}^{(U)}$ respectively. We also considered the AI algorithm described by Gilmour, Thompson & Cullis (1995) which is a popular Newton-Raphson type algorithm for fitting these types of models. We have used the ASRem1-R (Version 3) software package (Butler 2009) within the R computing environment to implement the AI algorithm. The standard implementation of the AI algorithm within the ASRem1-R software package involves a rescaling of the AI matrix when computing parameter updates. To implement the algorithm as described by Gilmour, Thompson & Cullis (1995), i.e., without a rescaling of the AI matrix, the ASRem1-R package was used to obtain the AI matrix and the score for a given set of parameter estimates and parameter updates were computed manually within the R computing environment. The R code and data files for reproducing the convergence information in Table 2 are available at www.github.com/SimonDiffey/ANZJS_REMLPXEM.

In all cases we used the convergence criterion proposed by Foulley, Jaffrezic & Robert-Cranie (2000), i.e., the criterion that algorithm is stopped when

$$\sqrt{(\boldsymbol{\kappa}^{(w+1)} - \boldsymbol{\kappa}^{(w)})^\top (\boldsymbol{\kappa}^{(w+1)} - \boldsymbol{\kappa}^{(w)}) / \boldsymbol{\kappa}^{(w)\top} \boldsymbol{\kappa}^{(w)}} < 10^{-8},$$

where $\boldsymbol{\kappa}$ is the vector of variance parameters.

5.1. Example: lamb weight data

We analysed the lamb weight data presented in Callanan & Harville (1991) which consists of birth weights (in pounds) of single birth male lambs that are the progeny of 62

ewes. The ages of the ewes are categorised into three groups. The lambs were sired by one of 23 rams that belonged to five different population lines. The number of observations per sire is highly unbalanced. At one extreme there are seven sires that have a single observation each and at the other extreme there is a single sire that has nine observations.

A possible way of modelling these data would be to fit sire as a random effect and the overall mean, ewe age and line as fixed effects. We define y as a 62×1 vector of observed birth weights and $X = (I, X_a, X_l)$ as a 62×7 design matrix of full column rank parameterised with the corner point constraint. The term I is a 62×1 vector of 1s. The design matrices X_a and X_l are indicator matrices for the second and third age groups and the second to the fifth population line respectively. The vector of fixed effects is similarly partitioned as $\beta = (\mu, \beta_a^\top, \beta_l^\top)^\top$. We assumed that $u \sim N(0, \sigma_u^2 I_{23})$, the vector of rescaled random effects is $f \sim N(0, dI_{23})$, and the 62×23 matrix Z is their associated design matrix. It is assumed that the vector of residuals is distributed $\epsilon \sim N(0, \sigma_{e^*}^2 I_{62})$. The invertible matrix associated with the working parameter is $\Lambda = \lambda I_{23}$ where λ is scalar. We define the variance parameter vector of interest to be $\kappa = (\sigma_e^2, \sigma_u^2)^\top$ and the expanded parameter vector to be $\mathcal{K} = (\sigma_{e^*}^2, d, \lambda)^\top$.

The updating equations for a REML PX-EM algorithm based on the complete data specification $y_{c\mathcal{X}}^{(P)} = (y^\top, \beta^\top, f^\top)^\top$ are

$$\begin{aligned} d^{(w+1)} &= \frac{1}{b} \left\{ \tilde{u}^{(w)\top} \tilde{u}^{(w)} + tr(C^{ZZ^{(w)}}) \right\}, \\ \sigma_{e^*}^{2(w+1)} &= \frac{1}{n} \left\{ \tilde{e}^{(w)\top} \tilde{e}^{(w)} + tr(WC^{-1^{(w)}}W^\top) \right\}, \\ \lambda_P^{(w+1)} &= \frac{\tilde{u}^{(w)\top} Z^\top (y - X\beta^{(w)}) - tr(Z^\top X C^{XZ^{(w)}})}{\tilde{u}^{(w)\top} Z^\top Z \tilde{u}^{(w)} + tr(Z^\top Z C^{ZZ^{(w)}})}, \end{aligned}$$

where we use the subscript P to distinguish between the update for λ based on the complete data specification $y_{c\mathcal{X}}^{(P)}$ and $y_{c\mathcal{X}}^{(U)}$. Note that $b = 23$, $n = 62$, and $\beta^{(w)}$ is the generalised least squares estimate of β at the w -th iterate. The updates for $\sigma_{e^*}^2$ can be re-expressed as

$$\sigma_{e^*}^{2(w+1)} = \frac{1}{n} \left\{ (y - Z\tilde{u}^{(w)})^\top K (y - Z\tilde{u}^{(w)}) + tr(Z^\top K Z C^{ZZ^{(w)}}) + t \sigma_{e^*}^{2(w)} \right\} \tag{14}$$

where $K = I_n - X(X^\top X)^{-1}X^\top$ and $t = 7$. When using the complete data specification $y_{c\mathcal{X}}^{(U)} = (y_2^\top, f^\top)^\top$ the updating equation for d is the same as that employed when using $y_{c\mathcal{X}}^{(P)}$. The updating equations for the other parameters are

$$\begin{aligned} \sigma_{e^*}^{2(w+1)} &= \frac{1}{n-t} \left\{ (y - Z\tilde{u}^{(w)})^\top K (y - Z\tilde{u}^{(w)}) + tr(Z^\top K Z C^{ZZ^{(w)}}) \right\}, \\ \lambda_U^{(w+1)} &= \frac{y^\top K Z \tilde{u}^{(w)}}{\tilde{u}^{(w)\top} Z^\top K Z \tilde{u}^{(w)} + tr(Z^\top K Z C^{ZZ^{(w)}})}. \end{aligned} \tag{15}$$

The difference between (14) and (15) is subtle but illuminating. In the latter the degrees of freedom associated with the estimation of the fixed effects are incorporated directly into the divisor, whereas in the former they are incorporated in the term $t \sigma_{e^*}^{2(w)}$. In the early iterations of a REML (PX)EM algorithm (14) and (15) will result in different updates although the two will be equivalent at convergence.

In regards to the relationship between λ_U and λ_P it can be shown, using identities arising from the mixed model equations and the fact that $\text{tr}(\mathbf{C}\mathbf{C}^{-1}) = \mathbf{I}_{t+b}$, that

$$\lambda_U^{(w+1)} - 1 = \frac{a}{b}(\lambda_P^{(w+1)} - 1), \quad (16)$$

i.e., the left hand side of (16) is a rescaling of the right hand side where $a = \tilde{\mathbf{u}}^{(w)\top} \mathbf{Z}^\top \mathbf{Z} \tilde{\mathbf{u}}^{(w)} + \text{tr}(\mathbf{Z}^\top \mathbf{Z} \mathbf{C} \mathbf{Z} \mathbf{Z}^{(w)})$ and $b = \tilde{\mathbf{u}}^{(w)\top} \mathbf{Z}^\top \mathbf{K} \mathbf{Z} \tilde{\mathbf{u}}^{(w)} + \text{tr}(\mathbf{Z}^\top \mathbf{K} \mathbf{Z} \mathbf{C} \mathbf{Z} \mathbf{Z}^{(w)})$.

We used two sets of starting values for σ_e^2 and σ_u^2 . The first consisted of $\sigma_e^{2(0)} = 2$ and $\sigma_u^{2(0)} = 2$, and the second of $\sigma_e^{2(0)} = 2$ and $\sigma_u^{2(0)} = 3$. Using the updating equations above, applying the reduction function $\kappa = \mathcal{R}(\mathcal{K}) = (\sigma_{e*}^2, d\lambda^2)^\top$, and applying the convergence criteria to the parameter vector $\kappa = (\sigma_e^2, \sigma_u^2)^\top$, we obtained REML estimates of the variance parameters σ_e^2 and σ_u^2 given by $\hat{\sigma}_e^2 = 2.961$ and $\hat{\sigma}_u^2 = 0.517$. The number of iterations to achieve convergence for all algorithms applied to the lamb weight data are presented in Table 2.

5.2. Example: a balanced incomplete block soybean variety trial

The soybean variety trial data originate from table 1 of a technical report by Weiss & Cox (1939). The experimental design is a balanced incomplete block design consisting of 6 replicate plots of 31 soybean varieties and 31 blocks. The plots were 6 feet by 16 feet and each contained two rows of soybeans.

We fitted a model where the overall mean and the effect of soybean variety were taken to be fixed effects. Block was fitted as a random effect. We define \mathbf{y} to be a 186×1 vector of observed yields (bushels per acre) and $\mathbf{X} = (\mathbf{I}, \mathbf{X}_v)$ as a 186×31 matrix of full column rank parameterised with the corner point constraint. The design matrix \mathbf{X}_v is an indicator matrix for variety number two through to variety number thirty one. The vector of fixed effects is $\boldsymbol{\beta} = (\mu, \boldsymbol{\beta}_v^\top)^\top$ and we assume that $\mathbf{u} \sim \mathbf{N}(\mathbf{0}, \sigma_u^2 \mathbf{I}_{31})$ where σ_u^2 is a variance component associated with between block variation. We also assume that $\mathbf{e} \sim \mathbf{N}(\mathbf{0}, \sigma_e^2 \mathbf{I}_{186})$ and that \mathbf{u} and \mathbf{e} are independent of each other.

For the REML PX-EM algorithms we set the vector of rescaled random effects associated with blocks equal to $\mathbf{f} \sim \mathbf{N}(\mathbf{0}, d\mathbf{I}_{31})$ and we assume that the vector of residuals is $\boldsymbol{\epsilon} \sim \mathbf{N}(\mathbf{0}, \sigma_{e*}^2 \mathbf{I}_{186})$. The invertible matrix associated with the working parameter is $\boldsymbol{\Lambda} = \lambda \mathbf{I}_{31}$. We define the variance parameter vector of interest to be $\kappa = (\sigma_e^2, \sigma_u^2)^\top$ and the expanded parameter vector to be $\mathcal{K} = (\sigma_{e*}^2, d, \lambda)^\top$.

The updating equations for the REML PX-EM algorithm based on the complete data specification $\mathbf{y}_{c\mathcal{X}}^{(P)} = (\mathbf{y}^\top, \boldsymbol{\beta}^\top, \mathbf{f}_a^\top)^\top$ and $\mathbf{y}_{c\mathcal{X}}^{(U)} = (\mathbf{y}_2^\top, \boldsymbol{\beta}^\top, \mathbf{f}_a^\top)^\top$ have the same form as those presented for the lamb weight data and we note that $n = 186$, $t = 31$ and $b = 31$.

We used two sets of starting values for σ_e^2 and σ_u^2 . The first consisted of $\sigma_e^{2(0)} = 1$ and $\sigma_u^{2(0)} = 1$. The second consisted of $\sigma_e^{2(0)} = 8$ and $\sigma_u^{2(0)} = 4$. Using the updating equations, applying the reduction function $\kappa = \mathcal{R}(\mathcal{K}) = (\sigma_{e*}^2, d\lambda^2)^\top$, and applying the convergence criteria to the parameter vector $\kappa = (\sigma_e^2, \sigma_u^2)^\top$, we found the REML estimates of the variance parameters σ_e^2 and σ_u^2 to be $\hat{\sigma}_e^2 = 3.585$ and $\hat{\sigma}_u^2 = 5.267$. The number of iterations to achieve convergence for all algorithms applied to the soybean variety trial data are presented in Table 2.

6. Discussion

The AI algorithm is a Newton-Raphson type algorithm and in cases where the algorithm converges it will outperform REML EM algorithms and their variants. However, a problem




Figure 1. Contour plot of the residual log-likelihood surface for the lamb weight data and the first 5 iterations of the REML PX-EM algorithms using the complete data specification $\mathbf{y}_{c\mathcal{X}}^{(U)}$ and $\mathbf{y}_{c\mathcal{X}}^{(P)}$. The starting values of $\sigma_e^{2(0)} = 2$ and $\sigma_u^{2(0)} = 3$ are represented by the black circle. The REML estimates of $\hat{\sigma}_e^2 = 2.962$ and $\hat{\sigma}_u^2 = 0.517$ are represented by the “target”.

with Newton-Raphson type algorithms is that the variance parameter updates are not guaranteed to remain in the parameter space and the algorithm can fail to converge. For the sire model for the lamb weight data and the incomplete block model for the soybean variety yield data, we considered a set of starting values where this is the case (Table 2). REML EM algorithms and their variants have the desirable properties of monotonic convergence and parameter updates remaining in the parameter space. On the other hand they can be slow to converge. In Table 2 we considered the performance of four variants of the REML EM algorithm.

We introduce a new incomplete data specification, referred to as y_2 which is based on the transformed observed data vector associated with the marginal distribution in the conditional derivation of REML (Verbyla 1990). This is a useful alternative to the current practice of considering the incomplete data to be y . For both the lamb weight and the soybean variety trial data the number of iterations to convergence was less when using the incomplete data specification y_2 compared to y . There was only one instance where the number of iterations to convergence was the same for the two alternative REML EM algorithms (Table 2). In the soybean variety trial example a REML EM algorithm using the new incomplete data specification converged in fewer iterations than a REML PX-EM algorithm using the current practice of considering the incomplete data to be y . A REML PX-EM algorithm where the incomplete data are defined as y_2 resulted in fewer iterations to convergence for both the lamb weight and soybean variety trial data. Depending on starting values, convergence was achieved in 28% to 30% fewer iterations in both examples. We have conducted simulation studies based on a one-way random effects model and in more than 99% of cases a REML PX-EM algorithm based on specifying the incomplete data as y_2 converged in fewer iterations than a REML PX-EM algorithm using y as the incomplete data.

A contour plot of the REML log-likelihood surface for the lamb weight data and the first five iterations of the REML PX-EM algorithms using the complete data specification $y_{c\mathcal{X}}^{(U)}$ and $y_{c\mathcal{X}}^{(P)}$ with the starting values $\sigma_e^{2(0)} = 2$ and $\sigma_u^{2(0)} = 3$ is provided in Figure 1. With this set of starting values the AI algorithm fails. There are two outstanding features of Figure 1; firstly, the REML PX-EM algorithm based on $y_{c\mathcal{X}}^{(U)}$ comes closer to the solution after every iterate than does a REML PX-EM algorithm based on $y_{c\mathcal{X}}^{(P)}$ and secondly, a REML PX-EM algorithm based on $y_{c\mathcal{X}}^{(U)}$ is close to the solution after a small number of iterations. This second feature suggests that a hybrid AI and REML PX-EM algorithm might be of some benefit in certain cases. Where the AI algorithm fails an alternative might be to consider a hybrid algorithm that starts with a small number of REML PX-EM algorithm iterations before switching to the AI algorithm. This is an approach considered by Meyer (2006), Knight (2008) and Zhou & Stephens (2014).

References

- BUTLER, D. (2009). asreml: asreml() fits the linear mixed model. Available from URL: www.vsnl.co.uk. R package version 3.0.
- CALLANAN, T.P. & HARVILLE, D.A. (1991). Some new algorithms for computing restricted maximum likelihood estimates of variance components. *Journal of Statistical Computation and Simulation* **38**, 239–259.
- CULLIS, B., SMITH, A. & THOMPSON, R. (2004). Perspectives of ANOVA, REML, and a general linear mixed model. In *Methods and Models in Statistics*, eds. N. Adams, M. Crowder, D.J. Hand & D. Stephens, pp. 53–94. London: Imperial College Press. In honour of Professor John Nelder, FRS.
- DEMPSTER, A., LAIRD, N. & RUBIN, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* **39**, 1–38.
- FOULLEY, J.L., JAFFREZIC, F. & ROBERT-CRANIE, C. (2000). EM-REML estimation of covariance parameters in Gaussian mixed models for longitudinal data analysis. *Genetics Selection Evolution* **32**, 129–141.
- FOULLEY, J.L. & VAN DYK, D. (2000). The PX-EM algorithm for fast stable fitting of Henderson's mixed model. *Genetics Selection Evolution* **32**, 143–163.
- GILMOUR, A., THOMPSON, R. & CULLIS, B. (1995). Average Information REML: an efficient algorithm for variance parameter estimation in linear mixed models. *Biometrics* **51**, 1440–1450.
- HENDERSON, C., KEMPTHORNE, O., SEARLE, S. & VON KROSIGK, C. (1959). The estimation of environmental and genetic trends from records subject to culling. *Biometrics* **15**, 192–218.
- KNIGHT, E. (2008). Improved iterative schemes for REML estimation of variance parameters in linear mixed models. Ph.D. thesis, The University of Adelaide, School of Agriculture, Food and Wine.

- LAIRD, N. & WARE, J. (1982). Random-effects models for longitudinal data. *Biometrics* **38**, 963–974.
- LIU, C., RUBIN, D. & WU, Y. (1998). Parameter expansion to accelerate the EM algorithm: the PX-EM algorithm. *Biometrika* **85**, 755–770.
- LOUIS, T. (1982). Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society, B* **44**, 226–233.
- MEYER, K. (2006). PX \times AI: Algorithmics for better convergence in restricted maximum likelihood estimation. Proceedings of the 8th World Congress on Genetics Applied to Livestock Production, Belo Horizonte, Brasil. Available from URL: <http://didgeridoo.une.edu.au/km/StuffToDownload/8WCGALP/pxai.pdf>. Last downloaded 28/08/2017.
- OAKES, D. (1999). Direct calculation of the information matrix via the EM algorithm. *Journal of the Royal Statistical Society, B*, **61**, 479–482.
- PATTERSON, H. & THOMPSON, R. (1971). Recovery of inter-block information when block sizes are unequal. *Biometrika* **58**, 545–554.
- R Core Team. (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. Available from URL: <https://www.R-project.org/>.
- SEARLE, S.R. (1982). *Matrix Algebra Useful for Statistics*. Wiley Series in Probability and Mathematical Statistics: Applied probability and Statistics. New York, NY: Wiley.
- SEARLE, S., CASELLA, G. & MCCULLOCH, C. (2006). *Variance Components*. Hoboken, New Jersey: John Wiley & Sons.
- SMYTH, G. (1998). Optimization and Nonlinear Equations. In *Encyclopedia of Biostatistics*, eds. P. Armitage and T. Colton, pp. 3174–3180. Chichester: John Wiley & Sons.
- VERBYLA, A. (1990). A conditional derivation of residual maximum likelihood. *Australian Journal of Statistics* **32**, 227–230.
- WEISS, M.G. & COX, G.M. (1939). Balanced incomplete block and lattice square designs for testing yield differences among large numbers of soybean varieties. *Research Bulletin (Iowa Agriculture and Home Economics Experiment Station)* **23**, 1.
- ZHOU, X. & STEPHENS, M. (2014). Efficient algorithms for multivariate linear mixed models in genomewide association studies. *Nature Methods* **11**, 407.