



Published in final edited form as:

J Comput Graph Stat. 2019 ; 28(2): 350–361. doi:10.1080/10618600.2018.1529601.

MM Algorithms For Variance Components Models

Hua Zhou^a, Liuyi Hu^b, Jin Zhou^c, Kenneth Lange^d

^aDepartment of Biostatistics, University of California, Los Angeles, CA

^bDepartment of Statistics, North Carolina State University, Raleigh, NC

^cDivision of Epidemiology and Biostatistics, University of Arizona, Tucson, AZ

^dDepartment of Human Genetics, University of California, Los Angeles, CA

Abstract

Variance components estimation and mixed model analysis are central themes in statistics with applications in numerous scientific disciplines. Despite the best efforts of generations of statisticians and numerical analysts, maximum likelihood estimation and restricted maximum likelihood estimation of variance component models remain numerically challenging. Building on the minorization-maximization (MM) principle, this paper presents a novel iterative algorithm for variance components estimation. Our MM algorithm is trivial to implement and competitive on large data problems. The algorithm readily extends to more complicated problems such as linear mixed models, multivariate response models possibly with missing data, maximum a posteriori estimation, and penalized estimation. We establish the global convergence of the MM algorithm to a Karush-Kuhn-Tucker (KKT) point and demonstrate, both numerically and theoretically, that it converges faster than the classical EM algorithm when the number of variance components is greater than two and all covariance matrices are positive definite.

Keywords

global convergence; matrix convexity; linear mixed model (LMM); maximum a posteriori (MAP) estimation; minorization-maximization (MM); multivariate response; penalized estimation; variance components model

1 Introduction

Variance components and linear mixed models are among the most potent tools in a statistician's toolbox, finding numerous applications in agriculture, biology, economics, genetics, epidemiology, and medicine. Given an observed $n \times 1$ response vector \mathbf{y} and $n \times p$ predictor matrix \mathbf{X} , the simplest variance components model postulates that $\mathbf{Y} \sim N(\mathbf{X}\boldsymbol{\beta}, \boldsymbol{\Omega})$, where $\boldsymbol{\Omega} = \sum_{i=1}^m \sigma_i^2 \mathbf{V}_i$, and the $\mathbf{V}_1, \dots, \mathbf{V}_m$ are m fixed positive semidefinite matrices. The parameters of the model can be divided into mean effects $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ and variance

components $\sigma^2 = (\sigma_1^2, \dots, \sigma_m^2)$. Throughout we assume $\mathbf{\Omega}$ is positive definite. The extension to singular $\mathbf{\Omega}$ will not be pursued here. Estimation revolves around the log-likelihood function

$$L(\beta, \sigma^2) = -\frac{1}{2} \ln \det \mathbf{\Omega} - \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta)^T \mathbf{\Omega}^{-1} (\mathbf{y} - \mathbf{X}\beta). \quad (1)$$

Among the commonly used methods for estimating variance components, maximum likelihood estimation (MLE) (Hartley and Rao, 1967) and restricted (or residual) MLE (REML) (Harville, 1977) are the most popular. REML first projects \mathbf{y} to the null space of \mathbf{X} and then estimates variance components based on the projected responses. If the columns of a matrix \mathbf{B} span the null space of \mathbf{X}^T , then REML estimates the σ_i^2 by maximizing the log-likelihood of the redefined response vector $\mathbf{B}^T \mathbf{Y}$, which is normally distributed with mean $\mathbf{0}$ and covariance $\mathbf{B}^T \mathbf{\Omega} \mathbf{B} = \sum_{i=1}^m \sigma_i^2 \mathbf{B}^T \mathbf{V}_i \mathbf{B}$.

There exists a large literature on iterative algorithms for finding MLE and REML (Laird and Ware, 1982; Lindstrom and Bates, 1988, 1990; Harville and Callanan, 1990; Callanan and Harville, 1991; Bates and Pinheiro, 1998; Schafer and Yucel, 2002). Fitting variance components models remains a challenge in models with a large sample size n or a large number of variance components m . Newton's method (Lindstrom and Bates, 1988) converges quickly but is numerically unstable owing to the non-concavity of the log-likelihood. Fisher's scoring algorithm replaces the observed information matrix in Newton's method by the expected information matrix and yields an ascent algorithm when safeguarded by step halving. However the calculation and inversion of expected information matrices cost $\mathcal{O}(mn^3) + \mathcal{O}(m^3)$ flops and quickly become impractical for large n or m , unless \mathbf{V}_j are low rank, block diagonal, or have other special structures. The expectation-maximization (EM) algorithm initiated by Dempster et al. (1977) is a third alternative (Laird and Ware, 1982; Laird et al., 1987; Lindstrom and Bates, 1988; Bates and Pinheiro, 1998). Compared to Newton's method, the EM algorithm is easy to implement and numerically stable, but painfully slow to converge. In practice, a strategy of priming Newton's method by a few EM steps leverages the stability of EM and the faster convergence of second-order methods.

In this paper we derive a novel minorization-maximization (MM) algorithm for finding the MLE and REML estimates of variance components. We prove global convergence of the MM algorithm to a Karush-Kuhn-Tucker (KKT) point and explain why MM generally converges faster than EM for models with more than two variance components. We also sketch extensions of the MM algorithm to the multivariate response model with possibly missing responses, the linear mixed model (LMM), maximum a posteriori (MAP) estimation, and penalized estimation. The numerical efficiency of the MM algorithm is illustrated through simulated data sets and a genomic example with 200 variance components.

2 Preliminaries

Background on MM algorithms

Throughout we reserve Greek letters for parameters and indicate the current iteration number by a superscript t . The MM principle for maximizing an objective function $f(\boldsymbol{\theta})$ involves minorizing the objective function $f(\boldsymbol{\theta})$ by a surrogate function $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ around the current iterate $\boldsymbol{\theta}^{(t)}$ of a search (Lange et al., 2000). Minorization is defined by the two conditions

$$\begin{aligned} f(\boldsymbol{\theta}^{(t)}) &= g(\boldsymbol{\theta}^{(t)} | \boldsymbol{\theta}^{(t)}) \\ f(\boldsymbol{\theta}) &\geq g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)}), \quad \boldsymbol{\theta} \neq \boldsymbol{\theta}^{(t)}. \end{aligned} \quad (2)$$

In other words, the surface $\boldsymbol{\theta} \mapsto g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ lies below the surface $\boldsymbol{\theta} \mapsto f(\boldsymbol{\theta})$ and is tangent to it at the point $\boldsymbol{\theta} = \boldsymbol{\theta}^{(t)}$. Construction of the minorizing function $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ constitutes the first M of the MM algorithm. The second M of the algorithm maximizes the surrogate $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ rather than $f(\boldsymbol{\theta})$. The point $\boldsymbol{\theta}^{(t+1)}$ maximizing $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ satisfies the ascent property $f(\boldsymbol{\theta}^{(t+1)}) \geq f(\boldsymbol{\theta}^{(t)})$. This fact follows from the inequalities

$$f(\boldsymbol{\theta}^{(t+1)}) \geq g(\boldsymbol{\theta}^{(t+1)} | \boldsymbol{\theta}^{(t)}) \geq g(\boldsymbol{\theta}^{(t)} | \boldsymbol{\theta}^{(t)}) = f(\boldsymbol{\theta}^{(t)}), \quad (3)$$

reflecting the definition of $\boldsymbol{\theta}^{(t+1)}$ and the tangency and domination conditions (2). The ascent property makes the MM algorithm remarkably stable. The validity of the descent property depends only on increasing $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$, not on maximizing $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$. With obvious changes, the MM algorithm also applies to minimization rather than to maximization. To minimize a function $f(\boldsymbol{\theta})$, we majorize it by a surrogate function $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ and minimize $g(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)})$ to produce the next iterate $\boldsymbol{\theta}^{(t+1)}$. The acronym should not be confused with the maximization-maximization algorithm in the variational Bayes context (Jeon, 2012).

The MM principle (De Leeuw, 1994; Heiser, 1995; Kiers, 2002; Lange et al., 2000; Hunter and Lange, 2004) finds applications in multidimensional scaling (Borg and Groenen, 2005), ranking of sports teams (Hunter, 2004), variable selection (Hunter and Li, 2005; Yen, 2011), optimal experiment design (Yu, 2010), multivariate statistics (Zhou and Lange, 2010), geometric programming (Lange and Zhou, 2014), survival models (Hunter and Lange, 2002; Ding et al., 2015), sparse covariance estimation (Bien and Tibshirani, 2011), and many other areas (Lange, 2016). The celebrated EM principle (Dempster et al., 1977) is a special case of the MM principle. The Q function produced in the E step of an EM algorithm minorizes the log-likelihood up to an irrelevant constant. Thus, both EM and MM share the same advantages: simplicity, stability, graceful adaptation to constraints, and the tendency to avoid large matrix inversion. The more general MM perspective frees algorithm derivation from the missing data straitjacket and invites wider applications (Wu and Lange, 2010). Figure 1 shows the minorization functions of EM and MM for a variance components model with $m = 2$ variance components.

Convex matrix functions

For symmetric matrices we write $A \leq B$ when $B - A$ is positive semidefinite and $A < B$ if $B - A$ is positive definite. A matrix-valued function f is said to be (matrix) convex if

$$f[\lambda A + (1 - \lambda)B] \leq \lambda f(A) + (1 - \lambda)f(B)$$

for all A, B , and $\lambda \in [0, 1]$. Our derivation of the MM variance components algorithm hinges on the convexity of the two functions mentioned in the next lemma. See standard text Boyd and Vandenberghe (2004) for the verification of both facts.

Lemma 1. (a) The matrix fractional function $f(A, B) = A^T B^{-1} A$ is jointly convex in the $m \times n$ matrix A and the $m \times m$ positive definite matrix B . (b) The log determinant function $f(B) = \ln \det B$ is concave on the set of positive definite matrices.

3 Univariate response model

Our strategy for maximizing the log-likelihood (1) is to alternate updating the mean parameters β and the variance components σ^2 . Updating β given σ^2 is a standard general least squares problem with solution

$$\beta^{(t+1)} = (X^T \Omega^{-(t)} X)^{-1} X^T \Omega^{-(t)} y, \quad (4)$$

where $\Omega^{-(t)}$ represents the inverse of $\Omega^{(t)} = \sum_{i=1}^m \sigma_i^{2(t)} V_i$. Updating σ^2 given $\beta^{(t)}$ depends on two minorizations. If we assume that all of the V_j are positive definite, then the joint convexity of the map $(X, Y) \mapsto X^T Y^{-1} X$ for positive definite Y implies that

$$\begin{aligned} \Omega^{(t)} \Omega^{-1} \Omega^{(t)} &= \left(\sum_{i=1}^m \sigma_i^{2(t)} V_i \right) \left(\sum_{i=1}^m \sigma_i^2 V_i \right)^{-1} \left(\sum_{i=1}^m \sigma_i^{2(t)} V_i \right) \\ &\leq \sum_{i=1}^m \frac{\sigma_i^{2(t)}}{\sum_j \sigma_j^{2(t)}} \left(\frac{\sum_j \sigma_j^{2(t)}}{\sigma_i^{2(t)}} \sigma_i^{2(t)} V_i \right) \left(\frac{\sum_j \sigma_j^{2(t)}}{\sigma_i^{2(t)}} \sigma_i^2 V_i \right)^{-1} \left(\frac{\sum_j \sigma_j^{2(t)}}{\sigma_i^{2(t)}} \sigma_i^{2(t)} V_i \right) \\ &= \sum_{i=1}^m \frac{\sigma_i^{4(t)}}{\sigma_i^2} V_i. \end{aligned}$$

When one or more of the V_j are rank deficient, we replace each V_j by $V_{j,\epsilon} = V_j + \epsilon I$ for $\epsilon > 0$ small and let $\Omega_\epsilon^{(t)} = \sum_i \sigma_i^{2(t)} V_{i,\epsilon}$. Sending ϵ to 0 in $\Omega_\epsilon^{(t)} \Omega_\epsilon^{-1} \Omega_\epsilon^{(t)} \leq \sum_{i=1}^m (\sigma_i^{4(t)} / \sigma_i^2) V_{i,\epsilon}$ now gives the desired majorization $\Omega^{(t)} \Omega^{-1} \Omega^{(t)} \leq \sum_{i=1}^m (\sigma_i^{4(t)} / \sigma_i^2) V_i$ in the general case. Negating both sides leads to the minorization

$$-(\mathbf{y} - \mathbf{X}\beta)^T \boldsymbol{\Omega}^{-1} (\mathbf{y} - \mathbf{X}\beta) \geq -(\mathbf{y} - \mathbf{X}\beta)^T \boldsymbol{\Omega}^{-\langle t \rangle} \left(\sum_{i=1}^m \frac{\sigma_i^{4\langle t \rangle}}{\sigma_i^2} \mathbf{V}_i \right) \boldsymbol{\Omega}^{-\langle t \rangle} (\mathbf{y} - \mathbf{X}\beta) \quad (5)$$

that effectively separates the variance components $\sigma_1^2, \dots, \sigma_m^2$ in the quadratic term of the log-likelihood (1).

The convexity of the function $A \mapsto -\log \det A$ is equivalent to the supporting hyperplane minorization

$$-\text{Indet} \boldsymbol{\Omega} \geq -\text{Indet} \boldsymbol{\Omega}^{\langle t \rangle} - \text{tr} \left[\boldsymbol{\Omega}^{-\langle t \rangle} (\boldsymbol{\Omega} - \boldsymbol{\Omega}^{\langle t \rangle}) \right] \quad (6)$$

that separates $\sigma_1^2, \dots, \sigma_m^2$ in the log determinant term of the log-likelihood (1). Combination of the minorizations (5) and (6) gives the overall minorization

$$\begin{aligned} g(\boldsymbol{\sigma}^2 | \boldsymbol{\sigma}^{2\langle t \rangle}) & \quad (7) \\ &= -\frac{1}{2} \text{tr}(\boldsymbol{\Omega}^{-\langle t \rangle} \boldsymbol{\Omega}) - \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta^{\langle t \rangle})^T \boldsymbol{\Omega}^{-\langle t \rangle} \left(\sum_{i=1}^m \frac{\sigma_i^{4\langle t \rangle}}{\sigma_i^2} \mathbf{V}_i \right) \boldsymbol{\Omega}^{-\langle t \rangle} (\mathbf{y} - \mathbf{X}\beta^{\langle t \rangle}) + c^{\langle t \rangle} \\ &= \sum_{i=1}^m \left[-\frac{\sigma_i^2}{2} \text{tr}(\boldsymbol{\Omega}^{-\langle t \rangle} \mathbf{V}_i) - \frac{1}{2} \frac{\sigma_i^{4\langle t \rangle}}{\sigma_i^2} (\mathbf{y} - \mathbf{X}\beta^{\langle t \rangle})^T \boldsymbol{\Omega}^{-\langle t \rangle} \mathbf{V}_i \boldsymbol{\Omega}^{-\langle t \rangle} (\mathbf{y} - \mathbf{X}\beta^{\langle t \rangle}) \right] + c^{\langle t \rangle}, \end{aligned}$$

where $c^{\langle t \rangle}$ is an irrelevant constant. Maximization of $g(\boldsymbol{\sigma}^2 | \boldsymbol{\sigma}^{2\langle t \rangle})$ with respect to σ_i^2 yields the simple multiplicative update

$$\sigma_i^{2\langle t+1 \rangle} = \sigma_i^{2\langle t \rangle} \sqrt{\frac{(\mathbf{y} - \mathbf{X}\beta^{\langle t \rangle})^T \boldsymbol{\Omega}^{-\langle t \rangle} \mathbf{V}_i \boldsymbol{\Omega}^{-\langle t \rangle} (\mathbf{y} - \mathbf{X}\beta^{\langle t \rangle})}{\text{tr}(\boldsymbol{\Omega}^{-\langle t \rangle} \mathbf{V}_i)}}, \quad i = 1, \dots, m. \quad (8)$$

As a sanity check on our derivation, consider the partial derivative

$$\frac{\partial}{\partial \sigma_i^2} L(\beta, \boldsymbol{\sigma}^2) = -\frac{1}{2} \text{tr}(\boldsymbol{\Omega}^{-1} \mathbf{V}_i) + \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta)^T \boldsymbol{\Omega}^{-1} \mathbf{V}_i \boldsymbol{\Omega}^{-1} (\mathbf{y} - \mathbf{X}\beta). \quad (9)$$

Given $\sigma_i^{2\langle t \rangle} > 0$, it is clear from the update formula (8) that $\sigma_i^{2\langle t+1 \rangle} < \sigma_i^{2\langle t \rangle}$ when $\frac{\partial}{\partial \sigma_i^2} L < 0$.

Conversely $\sigma_i^{2\langle t+1 \rangle} > \sigma_i^{2\langle t \rangle}$ when $\frac{\partial}{\partial \sigma_i^2} L > 0$.

```

Input :  $\mathbf{y}, \mathbf{X}, \mathbf{V}_1, \dots, \mathbf{V}_m$ 
Output: MLE  $\hat{\beta}, \hat{\sigma}_1^2, \dots, \hat{\sigma}_m^2$ 
1 Initialize  $\sigma_i^{(0)} > 0, i = 1, \dots, m$ 
2 repeat
3    $\mathbf{\Omega}^{(t)} \leftarrow \sum_{i=1}^m \sigma_i^{2(t)} \mathbf{V}_i$ 
4    $\hat{\beta}^{(t)} \leftarrow \arg \min_{\beta} (\mathbf{y} - \mathbf{X}\beta)^T \mathbf{\Omega}^{-t} (\mathbf{y} - \mathbf{X}\beta)$ 
5    $\mathbf{r}^{(t)} \leftarrow \mathbf{y} - \mathbf{X}\hat{\beta}^{(t)}$ 
6    $\sigma_i^{2(t+1)} \leftarrow \sigma_i^{2(t)} \sqrt{(\mathbf{r}^{(t)T} \mathbf{\Omega}^{-t} \mathbf{V}_i \mathbf{\Omega}^{-t} \mathbf{r}^{(t)}) / \text{tr}(\mathbf{\Omega}^{-t} \mathbf{V}_i)}, i = 1, \dots, m$ 
7 until objective value converges
    
```

Algorithm 1: MM algorithm for MLE of the variance components of model (1).

Algorithm 1 summarizes the MM algorithm for MLE of the univariate response model (1). The update formula (8) assumes that the numerator under the square root sign is nonnegative and the denominator is positive. The numerator requirement is a consequence of the positive semidefiniteness of \mathbf{V}_i . The denominator requirement is not obvious but can be verified through the Hadamard (elementwise) product representation $\text{tr}(\mathbf{\Omega}^{-t} \mathbf{V}_i) = \mathbf{1}^T (\mathbf{\Omega}^{-t} \odot \mathbf{V}_i) \mathbf{1}$. The following lemma of Schur (1911) is crucial. We give a self-contained probabilistic proof in Supplementary Materials S.1.

Lemma 2 (Schur). *The Hadamard product of a positive definite matrix with a positive semidefinite matrix with positive diagonal entries is positive definite.*

We can now obtain the following characterization of the MM iterates.

Proposition 1. *Assume \mathbf{V}_i has strictly positive diagonal entries. Then $\text{tr}(\mathbf{\Omega}^{-t} \mathbf{V}_i) > 0$ for all t . Furthermore if $\sigma_i^{2(0)} > 0$ and $\mathbf{\Omega}^{-t}(\mathbf{y} - \mathbf{X}\hat{\beta}^{(t)}) \notin \text{null}(\mathbf{V}_i)$ for all t , then $\sigma_i^{2(t)} > 0$ for all t . When \mathbf{V}_i is positive definite, $\sigma_i^{2(t)} > 0$ holds if and only if $\mathbf{y} \neq \mathbf{X}\hat{\beta}^{(t)}$.*

Proof. The first claim follows easily from Schur’s lemma. The second claim follows by induction. The third claim follows from the observation that $\text{null}(\mathbf{V}_i) = \{\mathbf{0}\}$.

In most applications, $\mathbf{V}_m = \mathbf{I}$. Proposition 1 guarantees that if $\sigma_m^{2(0)} > 0$ and the residual vector $\mathbf{y} - \mathbf{X}\hat{\beta}^{(t)}$ is nonzero, then $\sigma_m^{2(t)}$ remains positive and thus $\mathbf{\Omega}^{(t)}$ remains positive definite throughout all iterations. This fact does not prevent any of the sequences $\sigma_i^{2(t)}$ from converging to 0. In this sense, the MM algorithm acts like an interior point method, approaching the optimum from inside the feasible region.

Univariate response: two variance components

The major computational cost of Algorithm 1 is inversion of the covariance matrix $\mathbf{\Omega}^{(t)}$ at each iteration. The special case of $m = 2$ variance components deserves attention as repeated matrix inversion can be avoided by invoking the simultaneous congruence decomposition for two symmetric matrices, one of which is positive definite (Rao, 1973; Horn and Johnson, 1985). This decomposition is also called the generalized eigenvalue decomposition (Golub and Van Loan, 1996; Boyd and Vandenberghe, 2004). If one assumes $\mathbf{\Omega} = \sigma_1^2 \mathbf{V}_1 + \sigma_2^2 \mathbf{V}_2$ and

```

Input :  $y, X, V_1, V_2$ 
Output: MLE  $\hat{\beta}, \hat{\sigma}_1^2, \hat{\sigma}_2^2$ 
1 Simultaneous congruence decomposition:  $(D, U) \leftarrow (V_1, V_2)$ 
2 Transform data:  $\tilde{y} \leftarrow U^T y, \tilde{X} \leftarrow U^T X$ 
3 Initialize  $\sigma_1^{(0)}, \sigma_2^{(0)} > 0$ 
4 repeat
5    $w_i^{(t)} \leftarrow (\sigma_1^{2(t)} d_i + \sigma_2^{2(t)})^{-1}, \quad i = 1, \dots, n$ 
6    $\beta^{(t)} \leftarrow \arg \min_{\beta} \sum_{i=1}^n w_i^{(t)} (\tilde{y}_i - \tilde{x}_i^T \beta)^2$ 
7    $r^{(t)} \leftarrow \tilde{y} - \tilde{X} \beta^{(t)}$ 
8    $\sigma_1^{2(t+1)} \leftarrow \sigma_1^{2(t)} \sqrt{\frac{r^{(t)T} (\sigma_1^{2(t)} D + \sigma_2^{2(t)} I)^{-1} D (\sigma_1^{2(t)} D + \sigma_2^{2(t)} I)^{-1} r^{(t)}}{\text{tr}[(\sigma_1^{2(t)} D + \sigma_2^{2(t)} I)^{-1} D]}}$ 
9    $\sigma_2^{2(t+1)} \leftarrow \sigma_2^{2(t)} \sqrt{\frac{r^{(t)T} (\sigma_1^{2(t)} D + \sigma_2^{2(t)} I)^{-2} r^{(t)}}{\text{tr}[(\sigma_1^{2(t)} D + \sigma_2^{2(t)} I)^{-1}]}}$ 
10 until objective value converges
    
```

Algorithm 2: Simplified MM algorithm for MLE of model (1) with $m = 2$ variance components and $\Omega = \sigma_1^2 V_1 + \sigma_2^2 V_2$.

lets $(V_1, V_2) \mapsto (D, U)$ be the decomposition with U nonsingular, $U^T V_1 U = D$ diagonal, and $U^T V_2 U = I$, then

$$\begin{aligned}
 \Omega^{(t)} &= U^{-T} (\sigma_1^{2(t)} D + \sigma_2^{2(t)} I_n) U^{-1} & (10) \\
 \Omega^{- (t)} &= U (\sigma_1^{2(t)} D + \sigma_2^{2(t)} I_n)^{-1} U^T \\
 \det(\Omega^{(t)}) &= \det(\sigma_1^{2(t)} D + \sigma_2^{2(t)} I_n) \det(V_2).
 \end{aligned}$$

With the revised responses $\tilde{y} = U^T y$ and the revised predictor matrix $\tilde{X} = U^T X$, the update (8) requires only vector operations and costs $\mathcal{O}(n)$ flops. Updating the fixed effects is a weighted least squares problem with the transformed data (\tilde{y}, \tilde{X}) and observation weights

$w_i^{(t)} = (\sigma_1^{2(t)} d_i + \sigma_2^{2(t)})^{-1}$. Algorithm 2 summarizes the simplified MM algorithm for two variance components.

Numerical experiments

This section compares the numerical performance of MM, EM, Fisher scoring, and the lme4 package in R (Bates et al., 2015) on simulated data from a two-way ANOVA random effects model and a genetic model. For ease of comparison, all algorithm runs start from $\sigma^{2(0)} = \mathbf{1}$ and terminate when the relative change $(L^{(t+1)} - L^{(t)}) / (L^{(t)} + 1)$ in the log-likelihood is less than 10^{-8} .

Two-way ANOVA: We simulated data from a two-way ANOVA random effects model

$$y_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \epsilon_{ijk}, \quad 1 \leq i \leq a, 1 \leq j \leq b, 1 \leq k \leq c,$$

where $\alpha_i \sim N(0, \sigma_1^2)$, $\beta_j \sim N(0, \sigma_2^2)$, $(\alpha\beta)_{ij} \sim N(0, \sigma_3^2)$, and $\epsilon_{ijk} \sim N(0, \sigma_e^2)$ are jointly independent. Here i indexes levels in factor 1, j indexes levels in factor 2, and k indexes observations in the (i, j) -combination. This corresponds to $m = 4$ variance components. In the simulation, we set $\sigma_1^2 = \sigma_2^2 = \sigma_3^2$ and varied the ratio σ_1^2 / σ_e^2 ; the numbers of levels a and b in factor 1 and factor 2, respectively; and the number of observations c in each combination of factor levels. For each

simulation scenario, we simulated 50 replicates. The sample size was $n = abc$ for each replicate.

Tables 1 and 2 show the average number of iterations and the average runtimes when there are $a = b = 5$ levels of each factor. Based on these results and further results not shown for other combinations of a and b , we draw the following conclusions: Fisher scoring takes the fewest iterations; the MM algorithm always takes fewer iterations than the EM algorithm; the faster rate of convergence of Fisher scoring is outweighed by the extra cost of evaluating and inverting the information matrix. Table 1 in Supplementary Materials S.2 shows that all algorithms converged to same objective values.

Genetic model: We simulated a quantitative trait \mathbf{y} from a genetic model with two variance components and covariance matrix $\boldsymbol{\Omega} = \sigma_a^2 \widehat{\boldsymbol{\Phi}} + \sigma_e^2 \mathbf{I}$, where $\widehat{\boldsymbol{\Phi}}$ is a full-rank empirical kinship matrix estimated from the genome-wide measurements of 212 individuals using Option 29 of the Mendel software (Lange et al., 2013). In this example, MM had run times similar to Fisher scoring, and both were much faster than EM and lme4.

In summary, the MM algorithm appears competitive even in small-scale examples. Many applications involve a large number of variance components. In this setting, the EM algorithm suffers from slow convergence and Fisher scoring from an extremely high cost per iteration. Our genomic example in Section 7 reinforces this point.

4 Global convergence of the MM algorithm

The Karush-Kuhn-Tucker (KKT) necessary conditions for a local maximum $\boldsymbol{\sigma}^2 = (\sigma_1^2, \dots, \sigma_m^2)$ of the log-likelihood (1) require each component of the score vector to satisfy

$$\frac{\partial}{\partial \sigma_i^2} L(\boldsymbol{\sigma}^2) \in \begin{cases} \{0\} & \sigma_i^2 > 0 \\ (-\infty, 0] & \sigma_i^2 = 0. \end{cases}$$

In this section we establish the global convergence of Algorithm 1 to a KKT point. To reduce the notational burden, we assume that \mathbf{X} is null and omit estimation of fixed effects $\boldsymbol{\beta}$. The analysis easily extends to the nontrivial \mathbf{X} case. Our convergence analysis relies on characterizing the properties of the objective function $L(\boldsymbol{\sigma}^2)$ and the MM algorithmic mapping $\boldsymbol{\sigma}^2 \mapsto M(\boldsymbol{\sigma}^2)$ defined by equation (8). Special attention must be paid to the boundary values $\sigma_i^2 = 0$. We prove convergences for two cases, which cover most applications. For example, the genetic model in Section 3 satisfies Assumption 1, while the two-way ANOVA model satisfies Assumption 2.

Assumption 1. All \mathbf{V}_j are positive definite.

Assumption 2. \mathbf{V}_1 is positive definite, each \mathbf{V}_j is nontrivial, $\mathcal{H} = \text{span}\{\mathbf{V}_2, \dots, \mathbf{V}_m\}$ has dimension $q < n$, and $\mathbf{y} \notin \mathcal{H}$.

The key condition $\mathbf{y} \notin \text{span}\{\mathbf{V}_2, \dots, \mathbf{V}_m\}$ in the second case is also necessary for the existence of an MLE or REML (Demidenko and Massam, 1999; Grzadziel and Michalski, 2014). In Supplementary Materials S.4, we derive a sequence of lemmas en route to the global convergence result declared in Theorem 1.

Theorem 1. *Under either Assumption 1 or 2, the MM sequence $\{\sigma^{2(t)}\}_{t \geq 0}$ has at least one limit point. Every limit point is a fixed point of $M(\sigma^2)$. If the set of fixed points is discrete, then the MM sequence converges to one of them. Finally, when the iterates converge, their limit is a KKT point.*

5 MM versus EM

Examination of Tables 2 and 3 suggests that the MM algorithm usually converges faster than the EM algorithm. We now provide an explanation for this observation. Again for notational convenience, we consider the REML case where \mathbf{X} is null. Since the EM principle is just a special instance of the MM principle, we can compare their convergence properties in a unified framework. Consider an MM map $M(\theta)$ for maximizing the objective function $f(\theta)$ via the surrogate function $g(\theta | \theta^{(t)})$. Close to the optimal point θ^∞ ,

$$\theta^{(t+1)} - \theta^\infty \approx dM(\theta^\infty)(\theta^{(t)} - \theta^\infty),$$

where $dM(\theta^\infty)$ is the differential of the mapping M at the optimal point θ^∞ of $f(\theta)$. Hence, the local convergence rate of the sequence $\theta^{(t+1)} = M(\theta^{(t)})$ coincides with the spectral radius of $dM(\theta^\infty)$. Familiar calculations (Lange, 2010) demonstrate that

$$dM(\theta^\infty) = \mathbf{I} - [d^2g(\theta^\infty | \theta^\infty)]^{-1} d^2f(\theta^\infty).$$

In other words, the local convergence rate is determined by how well the surrogate surface $g(\theta | \theta^\infty)$ approximates the objective surface $f(\theta)$ near the optimal point θ^∞ . In the EM literature, $dM(\theta^\infty)$ is called the *rate matrix* (Meng and Rubin, 1991). Fast convergence occurs when the surrogate $g(\theta | \theta^\infty)$ hugs the objective $f(\theta)$ tightly around θ^∞ . Figure 1 shows a case where the MM surrogate locally dominates the EM surrogate. We demonstrate that this is no accident.

The Q-function in the EM algorithm

$$s_{EM}(\sigma^2 | \sigma^{2(t)}) = -\frac{1}{2} \sum_{i=1}^m \left[\text{rank}(\mathbf{V}_i) \cdot \ln \sigma_i^2 + \text{rank}(\mathbf{V}_i) \frac{\sigma_i^{2(t)}}{\sigma_i^2} - \frac{\sigma_i^{4(t)}}{\sigma_i^2} \text{tr}(\mathbf{\Omega}^{-(t)} \mathbf{V}_i) \right] - \frac{1}{2} \sum_{i=1}^m \frac{\sigma_i^{4(t)}}{\sigma_i^2} \mathbf{y}^T \mathbf{\Omega}^{-(t)} \mathbf{V}_i \mathbf{\Omega}^{-(t)} \mathbf{y}$$

minimizes the log-likelihood up to an irrelevant constant. Supplementary Materials S.6 gives a detailed derivation for the more general multivariate response case. Both surrogates $g_{EM}(\boldsymbol{\sigma}^2 | \boldsymbol{\sigma}^{2(\infty)})$ and $g_{MM}(\boldsymbol{\sigma}^2 | \boldsymbol{\sigma}^{2(\infty)})$ are parameter separated. This implies that both second differentials $d^2 g_{EM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})$ and $d^2 g_{MM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})$ are diagonal. A small diagonal entry of either matrix indicates fast convergence of the corresponding variance component. Our next result shows that, under Assumption 1, *on average* the diagonal entries of $d^2 g_{EM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})$ dominate those of $d^2 g_{MM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})$ when $m > 2$. Thus, the EM algorithm tends to converge more slowly than the MM algorithm, and the difference is more pronounced as the number of variance components m grows. See Supplementary Materials S.4 for the proof.

Theorem 2. *Let $\boldsymbol{\sigma}^{2(\infty)} > 0_m$ be a common limit point of the EM and MM algorithms. Then both second differentials $d^2 g_{MM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})$ and $d^2 g_{EM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})$ are diagonal with*

$$d^2 g_{EM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})_{ii} = -\frac{\text{rank}(\mathbf{V}_i)}{2\sigma_i^{4(\infty)}}$$

$$d^2 g_{MM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})_{ii} = -\frac{\mathbf{y}^T \boldsymbol{\Omega}^{-(\infty)} \mathbf{V}_i \boldsymbol{\Omega}^{-(\infty)} \mathbf{y}}{\sigma_i^{2(\infty)}} = -\frac{\text{tr}(\boldsymbol{\Omega}^{-(\infty)} \mathbf{V}_i)}{\sigma_i^{2(\infty)}}.$$

Furthermore, the average ratio

$$\frac{1}{m} \sum_{i=1}^m \frac{d^2 g_{MM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})_{ii}}{d^2 g_{EM}(\boldsymbol{\sigma}^{2(\infty)} | \boldsymbol{\sigma}^{2(\infty)})_{ii}} = \frac{2}{mn} \sum_{i=1}^m \text{tr}(\boldsymbol{\Omega}^{-(\infty)} \sigma_i^{2(\infty)} \mathbf{V}_i) = \frac{2}{m} < 1$$

for $m > 2$ when all \mathbf{V}_j have full rank n .

It is not clear whether a similar result holds under Assumption 2. Empirically we observed faster convergence of MM than EM, for example, in the two-way ANOVA example (Table 1). Also note that both the EM and MM algorithms must evaluate the traces $\text{tr}(\boldsymbol{\Omega}^{-(\ell)} \mathbf{V}_j)$ and quadratic forms $(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(\ell)})^T \boldsymbol{\Omega}^{-(\ell)} \mathbf{V}_j \boldsymbol{\Omega}^{-(\ell)} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(\ell)})$ at each iteration. Since these quantities are also the building blocks of the approximate rate matrices $d^2 g(\boldsymbol{\sigma}^{2(\ell)} | \boldsymbol{\sigma}^{2(\ell)})$, one can rationally choose either the EM or MM updates based on which has smaller diagonal entries measured by the ℓ_1 , ℓ_2 or ℓ_∞ norms. At negligible extra cost, this produces a hybrid algorithm that retains the ascent property and enjoys the better of the two convergence rates under either Assumption 1 or 2.

6 Extensions

Besides its competitive numerical performance, Algorithm 1 is attractive for its simplicity and ease of generalization. In this section, we outline MM algorithms for multivariate response models possibly with missing data, linear mixed models, MAP estimation, and penalized estimation.

6.1 Multivariate response model

Consider the multivariate response model with $n \times d$ response matrix \mathbf{Y} , which has no missing entries, mean $E \mathbf{Y} = \mathbf{X}\mathbf{B}$, and covariance

$$\mathbf{\Omega} = \text{Cov}(\text{vec}\mathbf{Y}) = \sum_{i=1}^m \mathbf{\Gamma}_i \otimes \mathbf{V}_i.$$

The $p \times d$ coefficient matrix \mathbf{B} collects the fixed effects, the $\mathbf{\Gamma}_i$ are unknown $d \times d$ variance components, and the \mathbf{V}_i are known $n \times n$ covariance matrices. If the vector $\text{vec}\mathbf{Y}$ is normally distributed, then \mathbf{Y} equals a sum of independent matrix normal distributions (Gupta and Nagar, 1999). We now make this assumption and pursue estimation of \mathbf{B} and the $\mathbf{\Gamma}_i$, which we collectively denote as $\mathbf{\Gamma}$. Under the normality assumption, Roth's Kronecker product identity $\text{vec}(\mathbf{CDE}) = (\mathbf{E}^T \otimes \mathbf{C})\text{vec}(\mathbf{D})$ yields the log-likelihood

$$\begin{aligned} L(\mathbf{B}, \mathbf{\Gamma}) &= -\frac{1}{2} \ln \det \mathbf{\Omega} - \frac{1}{2} \text{vec}(\mathbf{Y} - \mathbf{X}\mathbf{B})^T \mathbf{\Omega}^{-1} \text{vec}(\mathbf{Y} - \mathbf{X}\mathbf{B}) \\ &= -\frac{1}{2} \ln \det \mathbf{\Omega} - \frac{1}{2} [\text{vec}\mathbf{Y} - (\mathbf{I}_d \otimes \mathbf{X})\text{vec}\mathbf{B}]^T \mathbf{\Omega}^{-1} [\text{vec}\mathbf{Y} - (\mathbf{I}_d \otimes \mathbf{X})\text{vec}\mathbf{B}]. \end{aligned} \tag{11}$$

Updating \mathbf{B} given $\mathbf{\Gamma}^{(t)}$ is accomplished by solving the general least squares problem met earlier in the univariate case. Update of $\mathbf{\Gamma}_i$ given $\mathbf{B}^{(t)}$ is difficult due to the positive semidefiniteness constraint. Typical solutions involve reparameterization of the covariance matrix (Pinheiro and Bates, 1996). The MM algorithm derived in this section gracefully accommodates the covariance constraints.

Updating $\mathbf{\Gamma}$ given $\mathbf{B}^{(t)}$ requires generalizing the minorization (5). In view of Lemma 1 and the identities $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD})$ and $(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1}$, we have

$$\begin{aligned} \mathbf{\Omega}^{(t)} \mathbf{\Omega}^{-1} \mathbf{\Omega}^{(t)} &= m \left[\frac{1}{m} \sum_{i=1}^m \mathbf{\Gamma}_i^{(t)} \otimes \mathbf{V}_i \right] \left[\frac{1}{m} \sum_{i=1}^m \mathbf{\Gamma}_i \otimes \mathbf{V}_i \right]^{-1} \left[\frac{1}{m} \sum_{i=1}^m \mathbf{\Gamma}_i^{(t)} \otimes \mathbf{V}_i \right] \\ &\leq m \frac{1}{m} \sum_{i=1}^m (\mathbf{\Gamma}_i^{(t)} \otimes \mathbf{V}_i) (\mathbf{\Gamma}_i \otimes \mathbf{V}_i)^{-1} (\mathbf{\Gamma}_i^{(t)} \otimes \mathbf{V}_i) \\ &= \sum_{i=1}^m (\mathbf{\Gamma}_i^{(t)} \mathbf{\Gamma}_i^{-1} \mathbf{\Gamma}_i^{(t)}) \otimes \mathbf{V}_i \end{aligned}$$

or equivalently

$$\mathbf{\Omega}^{-1} \leq \mathbf{\Omega}^{-(t)} \left[\sum_{i=1}^m (\mathbf{\Gamma}_i^{(t)} \mathbf{\Gamma}_i^{-1} \mathbf{\Gamma}_i^{(t)}) \otimes \mathbf{V}_i \right] \mathbf{\Omega}^{-(t)}. \tag{12}$$

This derivation relies on the invertibility of the matrices \mathbf{V}_i . One can relax this assumption by substituting $\mathbf{V}_{\epsilon,i} = \mathbf{V}_i + \epsilon \mathbf{I}_n$ for \mathbf{V}_i and sending s to 0.

The majorization (12) and the minorization (6) jointly yield the surrogate

$$g(\Gamma|\Gamma^{(t)}) = -\frac{1}{2} \sum_{i=1}^m \left\{ \text{tr}[\Omega^{-(t)}(\Gamma_i \otimes V_i)] + (\text{vec} \mathbf{R}^{(t)})^T \left[(\Gamma_i^{(t)} \Gamma_i^{-1} \Gamma_i^{(t)}) \otimes V_i \right] (\text{vec} \mathbf{R}^{(t)}) \right\} + c^{(t)},$$

where $\mathbf{R}^{(t)}$ is the $n \times d$ matrix satisfying $\text{vec} \mathbf{R}^{(t)} = \Omega^{-(t)} \text{vec}(\mathbf{Y} - \mathbf{X}\mathbf{B}^{(t)})$ and $c^{(t)}$ is an irrelevant constant. Based on the Kronecker identities $(\text{vec} \mathbf{A})^T \text{vec} \mathbf{B} = \text{tr}(\mathbf{A}^T \mathbf{B})$ and $\text{vec}(\mathbf{CDE}) = (\mathbf{E}^T \otimes \mathbf{C}) \text{vec}(\mathbf{D})$, the surrogate can be rewritten as

$$\begin{aligned} g(\Gamma|\Gamma^{(t)}) &= -\frac{1}{2} \sum_{i=1}^m \left\{ \text{tr}[\Omega^{-(t)}(\Gamma_i \otimes V_i)] + \text{tr}(\mathbf{R}^{(t)T} V_i \mathbf{R}^{(t)} \Gamma_i^{(t)} \Gamma_i^{-1} \Gamma_i^{(t)}) \right\} + c^{(t)} \\ &= -\frac{1}{2} \sum_{i=1}^m \left\{ \text{tr}[\Omega^{-(t)}(\Gamma_i \otimes V_i)] + \text{tr}(\Gamma_i^{(t)} \mathbf{R}^{(t)T} V_i \mathbf{R}^{(t)} \Gamma_i^{-1} \Gamma_i^{(t)}) \right\} + c^{(t)}. \end{aligned}$$

The first trace is linear in Γ_i with the coefficient of entry $(\Gamma_i)_{jk}$ equal to

$$\text{tr}(\Omega_{jk}^{-(t)} V_i) = \mathbf{1}_n^T (V_i \odot \Omega_{jk}^{-(t)}) \mathbf{1}_n,$$

where $\Omega_{jk}^{-(t)}$ is the (j, k) -th $n \times n$ block of $\Omega^{-(t)}$ and \odot indicates elementwise product. The matrix M_i of these coefficients can be written as

$$M_i = (I_d \otimes \mathbf{1}_n)^T \left[(\mathbf{1}_n \mathbf{1}_d^T \otimes V_i) \odot \Omega^{-(t)} \right] (I_d \otimes \mathbf{1}_n).$$

The directional derivative of $g(\Gamma | \Gamma^{(t)})$ with respect to Γ_i in the direction Δ_i is

$$\begin{aligned} &-\frac{1}{2} \text{tr}(M_i \Delta_i) + \frac{1}{2} \text{tr}(\Gamma_i^{(t)} \mathbf{R}^{(t)T} V_i \mathbf{R}^{(t)} \Gamma_i^{-1} \Delta_i \Gamma_i^{-1}) \\ &= -\frac{1}{2} \text{tr}(M_i \Delta_i) + \frac{1}{2} \text{tr}(\Gamma_i^{-1} \Gamma_i^{(t)} \mathbf{R}^{(t)T} V_i \mathbf{R}^{(t)} \Gamma_i \Gamma_i^{-1} \Delta_i). \end{aligned}$$

Because all directional derivatives of $g(\Gamma | \Gamma^{(t)})$ vanish at a stationarity point, the matrix equation

$$M_i = \Gamma_i^{-1} \Gamma_i^{(t)} \mathbf{R}^{(t)T} V_i \mathbf{R}^{(t)} \Gamma_i \Gamma_i^{-1} \quad (13)$$

holds. Fortunately, this equation admits an explicit solution. For positive scalars a and b , the solution to the equation $b = X^{-1} a X^{-1}$ is $X = \pm \sqrt{a/b}$. The matrix analogue of this equation is the Riccati equation $\mathbf{B} = \mathbf{X}^{-1} \mathbf{A} \mathbf{X}^{-1}$, whose solution is summarized in the next lemma.

Lemma 3. Assume \mathbf{A} and \mathbf{B} are positive definite and \mathbf{L} is the Cholesky factor of \mathbf{B} . Then $\mathbf{Y} = \mathbf{L}^{-T} (\mathbf{L}^T \mathbf{A} \mathbf{L})^{1/2} \mathbf{L}^{-1}$ is the unique positive definite solution to the matrix equation $\mathbf{B} = \mathbf{X}^{-1} \mathbf{A} \mathbf{X}^{-1}$.

The Cholesky factor \mathbf{L} in Lemma 3 can be replaced by the symmetric square root of \mathbf{B} . The solution, which is unique, remains the same. The Cholesky decomposition is preferred for its cheaper computational cost and better numerical stability.

```

input :  $\mathbf{Y}, \mathbf{X}, \mathbf{V}_1, \dots, \mathbf{V}_m$ 
output: MLE  $\hat{\mathbf{B}}, \hat{\Gamma}_1, \dots, \hat{\Gamma}_m$ 
1 Initialize  $\Gamma_i^{(0)}$  positive definite,  $i = 1, \dots, m$ 
2 repeat
3    $\Omega^{(t)} \leftarrow \sum_{i=1}^m \Gamma_i^{(t)} \otimes \mathbf{V}_i$ 
4    $\mathbf{B}^{(t)} \leftarrow \arg \min_{\mathbf{B}} [\text{vec} \mathbf{Y} - (\mathbf{I}_d \otimes \mathbf{X}) \text{vec} \mathbf{B}]^T \Omega^{-(t)} [\text{vec} \mathbf{Y} - (\mathbf{I}_d \otimes \mathbf{X}) \text{vec} \mathbf{B}]$ 
5    $\mathbf{R}^{(t)} \leftarrow \text{reshape}(\Omega^{-(t)} \text{vec}(\mathbf{Y} - \mathbf{X} \mathbf{B}^{(t)}), n, d)$ 
6   for  $i = 1, \dots, m$  do
7     Cholesky  $\mathbf{L}_i^{(t)} \mathbf{L}_i^{(t)T} \leftarrow (\mathbf{I}_d \otimes \mathbf{1}_n)^T [(\mathbf{1}_d \mathbf{1}_d^T \otimes \mathbf{V}_i) \otimes \Omega^{-(t)}] (\mathbf{I}_d \otimes \mathbf{1}_n)$ 
8      $\Gamma_i^{(t+1)} \leftarrow \mathbf{L}_i^{-(t)T} [\mathbf{L}_i^{(t)T} (\Gamma_i^{(t)} \mathbf{R}^{(t)T} \mathbf{V}_i \mathbf{R}^{(t)} \Gamma_i^{(t)}) \mathbf{L}_i^{(t)}]^{1/2} \mathbf{L}_i^{-(t)}$ 
9   end
10 until objective value converges
    
```

Algorithm 3: The MM algorithm for MLE of the multivariate response model (11).

Algorithm 3 summarizes the MM algorithm for fitting the multi-response model (3). Each iteration invokes m Cholesky decompositions and symmetric square roots of $d \times d$ positive definite matrices. Fortunately in most applications, d is a small number. The following result guarantees the non-singularity of the Cholesky factor throughout the iterations. See Supplementary Materials S.8 for the proof.

Proposition 2. Assume \mathbf{V}_i has strictly positive diagonal entries. Then the symmetric matrix $\mathbf{M}_i = (\mathbf{I}_d \otimes \mathbf{1}_n)^T [(\mathbf{1}_d \mathbf{1}_d^T \otimes \mathbf{V}_i) \otimes \Omega^{-(t)}] (\mathbf{I}_d \otimes \mathbf{1}_n)$ is positive definite for all t . Furthermore if $\Gamma_i^{(0)} > 0$ and no column of $\mathbf{R}^{(t)}$ lies in the null space of \mathbf{V}_i for all t , then $\Gamma_i^{(t)} > 0$ for all t .

Multivariate response, two variance components—When there are $m = 2$ variance components $\mathbf{\Omega} = \Gamma_1 \otimes \mathbf{V}_1 + \Gamma_2 \otimes \mathbf{V}_2$, repeated inversion of the $nd \times nd$ covariance matrix $\mathbf{\Omega}$ reduces to a single $n \times n$ simultaneous congruence decomposition and, per iteration, two $d \times d$ Cholesky decompositions and one $d \times d$ simultaneous congruence decomposition. The simultaneous congruence decomposition of the matrix pair $(\mathbf{V}_1, \mathbf{V}_2)$ involves generalized eigenvalues $\mathbf{d} = (d_1, \dots, d_n)$ and a nonsingular matrix \mathbf{U} such that $\mathbf{U}^T \mathbf{V}_1 \mathbf{U} = \mathbf{D} = \text{diag}(\mathbf{d})$ and $\mathbf{U}^T \mathbf{V}_2 \mathbf{U} = \mathbf{I}$. If the simultaneous congruence decomposition of $(\Gamma_1^{(t)}, \Gamma_2^{(t)})$ is $(\mathbf{\Lambda}^{(t)}, \Phi^{(t)})$ with $\Phi^{(t)T} \Gamma_1^{(t)} \Phi^{(t)} = \mathbf{\Lambda}^{(t)} = \text{diag}(\lambda^{(t)})$ and $\Phi^{(t)T} \Gamma_2^{(t)} \Phi^{(t)} = \mathbf{I}_d$, then

$$\begin{aligned}
 \mathbf{\Omega}^{(t)} &= (\Phi^{-(t)} \otimes \mathbf{U}^{-1})^T (\mathbf{\Lambda}^{(t)} \otimes \mathbf{D} + \mathbf{I}_D \otimes \mathbf{I}_n) (\Phi^{-(t)} \otimes \mathbf{U}^{-1}) \\
 \mathbf{\Omega}^{-t} &= (\Phi^{(t)} \otimes \mathbf{U}) (\mathbf{\Lambda}^{(t)} \otimes \mathbf{D} + \mathbf{I}_D \otimes \mathbf{I}_n)^{-1} (\Phi^{(t)} \otimes \mathbf{U})^T \\
 \det \mathbf{\Omega}^{(t)} &= \det(\mathbf{\Lambda}^{(t)} \otimes \mathbf{D} + \mathbf{I}_D \otimes \mathbf{I}_n) \det [(\Phi^{-(t)} \otimes \mathbf{U}^{-1})^T (\Phi^{-(t)} \otimes \mathbf{U}^{-1})] \\
 &= \det(\mathbf{\Lambda}^{(t)} \otimes \mathbf{D} + \mathbf{I}_D \otimes \mathbf{I}_n) \det(\Gamma_2^{(t)} \otimes \mathbf{V}_2) \\
 &= \det(\mathbf{\Lambda}^{(t)} \otimes \mathbf{D} + \mathbf{I}_D \otimes \mathbf{I}_n) \det(\Gamma_2^{(t)})^n \det(\mathbf{V}_2)^D.
 \end{aligned}$$

Updating the fixed effects reduces to a weighted least squares problem for the transformed responses $\tilde{\mathbf{Y}} = \mathbf{U}^T \mathbf{Y}$, transformed predictor matrix $\tilde{\mathbf{X}} = \mathbf{U}^T \mathbf{X}$, and observation weights

$(\lambda_k^{(t)} d_i + 1)^{-1}$. Algorithm 4 summarizes the simplified MM algorithm. The lengthy derivations are relegated to Supplementary Materials S.5.

```

Input :  $Y, X, V_1, V_2$ 
Output: MLE  $\hat{B}, \hat{\Gamma}_1, \hat{\Gamma}_2$ 
1 Simultaneous congruence decomposition:  $(D, U) \leftarrow (V_1, V_2)$ 
2 Transform data:  $\tilde{Y} \leftarrow U^T Y, \tilde{X} \leftarrow U^T X$ 
3 Initialize  $\Gamma_1^{(0)}, \Gamma_2^{(0)}$  positive definite
4 repeat
5   Simultaneous congruence decomposition  $(\Lambda^{(t)}, \Phi^{(t)}) \leftarrow (\Gamma_1^{(t)}, \Gamma_2^{(t)})$ 
6    $B^{(t)} \leftarrow \arg \min_B [\text{vec}(\tilde{Y}\Phi^{(t)}) - (\Phi^{(t)T} \otimes \tilde{X})\text{vec}B]^T (\Lambda^{(t)} \otimes D + I_d \otimes I_n)^{-1} [\text{vec}(\tilde{Y}\Phi^{(t)}) - (\Phi^{(t)T} \otimes \tilde{X})\text{vec}B]$ 
7   Cholesky  $L_1^{(t)} L_1^{(t)T} \leftarrow \Phi^{(t)} \text{diag} \left( \text{tr} \left( D(\lambda_k^{(t)} D + I_n)^{-1} \right), k = 1, \dots, d \right) \Phi^{(t)T}$ 
8   Cholesky  $L_2^{(t)} L_2^{(t)T} \leftarrow \Phi^{(t)} \text{diag} \left( \text{tr} \left( (\lambda_k^{(t)} D + I_n)^{-1} \right), k = 1, \dots, d \right) \Phi^{(t)T}$ 
9    $N_1^{(t)} \leftarrow D^{1/2} [(\tilde{Y} - \tilde{X}B^{(t)})\Phi^{(t)}] \oslash (d\lambda^{(t)T} + \mathbf{1}_n \mathbf{1}_d^T) \Lambda^{(t)} \Phi^{- (t)}$ 
10   $N_2^{(t)} \leftarrow [(\tilde{Y} - \tilde{X}B^{(t)})\Phi^{(t)}] \oslash (d\lambda^{(t)T} + \mathbf{1}_n \mathbf{1}_d^T) \Phi^{- (t)}$ 
11   $\Gamma_i^{(t+1)} \leftarrow L_i^{- (t)T} (L_i^{(t)T} M_i^{(t)T} M_i^{(t)} L_i^{(t)})^{1/2} L_i^{- (t)}, i = 1, 2$ 
12 until objective value converges
    
```

Algorithm 4: MM algorithm for multivariate response model $\Omega = \Gamma_1 \otimes V_1 + \Gamma_2 \otimes V_2$ with two variance components matrices. Note that \oslash denotes a Hadamard quotient.

6.2 Multivariate response model with missing responses

In many applications the multivariate response model (11) involves missing responses. For instance, in testing multiple longitudinal traits in genetics, some trait values y_{ij} may be missing due to dropped patient visits, while their genetic covariates are complete. Missing data destroys the symmetry of the log-likelihood (11) and complicates finding the MLE. Fortunately, MM algorithm 3 easily adapts to this challenge.

The familiar EM argument (McLachlan and Krishnan, 2008, Section 2.2) shows that

$$-\frac{n}{2} \text{Indet} \Omega^{(t)} - \frac{1}{2} \text{tr} \{ \Omega^{- (t)} [\text{vec}(Z^{(t)} - XB^{(t)})\text{vec}(Z^{(t)} - XB^{(t)})^T + C^{(t)}] \} \quad (14)$$

minorizes the observed log-likelihood at the current iterate $(B^{(t)}, \Gamma_1^{(t)}, \dots, \Gamma_m^{(t)})$. Here $Z^{(t)}$ is the completed response matrix given the observed responses $Y_{obs}^{(t)}$ and the current parameter values. The complete data Y is assumed to be normally distributed $N(\text{vec}(XB^{(t)}), \Omega^{(t)})$. The block matrix $C^{(t)}$ is 0 except for a lower-right block consisting of a Schur complement.

To maximize the surrogate (14), we invoke the familiar minorization (6) and majorization (12) to separate the variance components Γ_i . At each iteration we impute missing entries by their conditional means, compute their conditional variances and covariances to supply the Schur complement, and then update the fixed effects and variance components by the explicit updates of Algorithm 3. The required conditional means and conditional variances can be conveniently obtained in the process of inverting $\Omega^{(t)}$ by the sweep operator of computational statistics (Lange, 2010, Section 7.3).

6.3 Linear mixed model (LMM)

The linear mixed model plays a central role in longitudinal data analysis. Consider the single-level LMM (Laird and Ware, 1982; Bates and Pinheiro, 1998) for n independent data clusters (y_i, X_i, Z_i) with

$$Y_i = X_i\beta + Z_i\gamma_i + \epsilon_i, \quad i = 1, \dots, n,$$

where β is a vector of fixed effects, the $\gamma_i \sim N(\mathbf{0}, \mathbf{R}_i(\theta))$ are independent random effects, and $\epsilon_i \sim N(0, \sigma^2 I_{n_i})$ captures random noise independent of γ_i . We assume the matrices Z_i have full column rank. The within-cluster covariance matrices $\mathbf{R}_i(\theta)$ depend on a parameter vector θ , typical choices for $\mathbf{R}_i(\theta)$ impose autocorrelation, compound symmetry, or unstructured correlation. It is clear that Y_i is normal with mean $X_i\beta$, covariance $\Omega_i = Z_i\mathbf{R}_i(\theta)Z_i^T + \sigma^2 I_{n_i}$, and log-likelihood

$$L_i(\beta, \theta, \sigma^2) = -\frac{1}{2} \ln \det \Omega_i - \frac{1}{2} (y_i - X_i\beta)^T \Omega_i^{-1} (y_i - X_i\beta).$$

The next three technical facts about pseudo-inverses are used in deriving the MM algorithm for LMM and their proofs are in Supplementary Materials S.9-S.11.

Lemma 4. *If A has full column rank and B has full row rank, then $(AB)^+ = B^+A^+$.*

Lemma 5. *If A and B are positive semidefinite matrices with the same range, then*

$$\lim_{\epsilon \downarrow 0} (B + \epsilon I)(A + \epsilon I)^{-1}(B + \epsilon I) = BA^+B.$$

Lemma 6. *If R and S are positive definite matrices, and the conformable matrix Z has full column rank, then the matrices ZRZ^T and ZSZ^T share a common range.*

The convexity of the map $(X, Y) \mapsto X^T Y^{-1} X$ and Lemmas 4–6 now yield via the obvious limiting argument the majorization

$$\begin{aligned} \Omega^{(t)} \Omega^{-1} \Omega^{(t)} &= \left(Z_i \mathbf{R}_i(\theta^{(t)}) Z_i^T + \sigma^{2(t)} I_{n_i} \right) \left(Z_i \mathbf{R}_i(\theta) Z_i^T + \sigma^2 I_{n_i} \right)^{-1} \left(Z_i \mathbf{R}_i(\theta^{(t)}) Z_i^T + \sigma^{2(t)} I_{n_i} \right) \\ &\leq \left(Z_i \mathbf{R}_i(\theta^{(t)}) Z_i^T \right) \left(Z_i \mathbf{R}_i(\theta) Z_i^T \right)^+ \left(Z_i \mathbf{R}_i(\theta^{(t)}) Z_i^T \right) + \frac{\sigma^{4(t)}}{\sigma^2} I_{n_i} \\ &= \left[Z_i \mathbf{R}_i(\theta^{(t)}) Z_i^T Z_i^T + \right] \mathbf{R}_i^{-1}(\theta) \left[Z_i^+ Z_i \mathbf{R}_i(\theta^{(t)}) Z_i^T \right] + \frac{\sigma^{4(t)}}{\sigma^2} I_{n_i}. \end{aligned}$$

In combination with the minorization (6), this gives the surrogate

$$g_i(\theta, \sigma^2 | \theta^{(t)}, \sigma^{2(t)}) = -\frac{1}{2} \text{tr}(\mathbf{Z}_i^T \boldsymbol{\Omega}_i^{-2(t)} \mathbf{Z}_i \mathbf{R}_i(\theta)) - \frac{1}{2} r_i^{(t)T} \mathbf{R}_i^{-1}(\theta) r_i^{(t)} - \frac{1}{2} \text{tr}(\boldsymbol{\Omega}_i^{-2(t)}) - \frac{\sigma^{4(t)}}{2\sigma^2} (y_i - X_i \beta^{(t)})^T \boldsymbol{\Omega}_i^{-2(t)} (y_i - X_i \beta^{(t)}) + c^{(t)},$$

for the log-likelihood $L_i(\theta, \sigma^2)$, where

$$r_i^{(t)} = (\mathbf{Z}_i^+ \mathbf{Z}_i \mathbf{R}_i(\theta^{(t)}) \mathbf{Z}_i^T) \boldsymbol{\Omega}_i^{-2(t)} (y_i - X_i \beta^{(t)}) = \mathbf{R}_i(\theta^{(t)}) \mathbf{Z}_i^T \boldsymbol{\Omega}_i^{-2(t)} (y_i - X_i \beta^{(t)}).$$

The parameters θ and σ^2 are nicely separated. To maximize the overall minorization function $\sum_i g_i(\theta, \sigma^2 | \theta^{(t)}, \sigma^{2(t)})$, we update σ^2 via

$$\sigma^{2(t+1)} = \sigma^{2(t)} \sqrt{\frac{\sum_i (y_i - X_i \beta^{(t)})^T \boldsymbol{\Omega}_i^{-2(t)} (y_i - X_i \beta^{(t)})}{\sum_i \text{tr}(\boldsymbol{\Omega}_i^{-2(t)})}}.$$

For structured models such as autocorrelation and compound symmetry, updating θ is a low-dimensional optimization problem that can be approached through the stationarity condition

$$\sum_i \text{vec}(\mathbf{Z}_i^T \boldsymbol{\Omega}_i^{-2(t)} \mathbf{Z}_i - \mathbf{R}_i^{-1}(\theta) r_i^{(t)} r_i^{(t)T} \mathbf{R}_i^{-1}(\theta)) \frac{\partial}{\partial \theta_j} \text{vec} \mathbf{R}_i(\theta) = 0$$

for each component θ_j . For the unstructured model with $\mathbf{R}_i(\theta) = \mathbf{R}$ for all i , the stationarity condition reads

$$\sum_i \mathbf{Z}_i^T \boldsymbol{\Omega}_i^{-2(t)} \mathbf{Z}_i = \mathbf{R}^{-1} \left(\sum_i r_i^{(t)} r_i^{(t)T} \right) \mathbf{R}^{-1}$$

and admits an explicit solution based on Lemma 3.

The same tactics apply to a multilevel LMM (Bates and Pinheiro, 1998) with responses

$$Y_i = X_i \beta + Z_{i1} \gamma_{i1} + \dots + Z_{im} \gamma_{im} + \epsilon_i.$$

Minorization separates parameters for each level (variance component). Depending on the complexity of the covariance matrices, maximization of the surrogate can be accomplished analytically. For the sake of brevity, details are omitted.

6.4 MAP estimation

Suppose β follows an improper flat prior, the variance components σ_i^2 follow inverse gamma priors with shapes $\alpha_i > 0$ and scales $\gamma_i > 0$, and these priors are independent. The log-posterior density then reduces to

$$-\frac{1}{2} \ln \det \Omega - \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta)^T \Omega^{-1} (\mathbf{y} - \mathbf{X}\beta) - \sum_{i=1}^m (\alpha_i + 1) \ln \sigma_i^2 - \sum_{i=1}^m \frac{\gamma_i}{\sigma_i^2} + c, \quad (15)$$

where c is an irrelevant constant. The MAP estimator of (β, σ^2) is the mode of the posterior distribution. The update (4) of β given σ^2 remains the same. To update σ^2 given β , apply the same minorizations (5) and (6) to the first two terms of equation (15). This separates parameters and yields a convex surrogate for each σ_i^2 . The minimum of the σ_i^2 surrogate is defined by the stationarity condition

$$0 = -\frac{1}{2} \text{tr}(\Omega^{-(t)} \mathbf{v}_i) + \frac{\sigma_i^{4(t)}}{2\sigma_i^4} (\mathbf{y} - \mathbf{X}\beta^{(t)})^T \Omega^{-(t)} \mathbf{v}_i \Omega^{-(t)} (\mathbf{y} - \mathbf{X}\beta^{(t)}) - \frac{\alpha_i + 1}{\sigma_i^2} + \frac{\gamma_i}{\sigma_i^4}.$$

Multiplying this by σ_i^4 gives a quadratic equation in σ_i^2 . The positive root should be taken to meet the nonnegativity constraint on σ_i^2 .

For the multivariate response model (11), we assume the variance components Γ_i follow independent inverse Wishart distributions with degrees of freedom $\nu_i > d - 1$ and scale matrix $\Psi_i > 0$. The log density of the posterior distribution is

$$-\frac{1}{2} \ln \det \Omega - \frac{1}{2} \text{vec}(\mathbf{Y} - \mathbf{X}\mathbf{B})^T \Omega^{-1} \text{vec}(\mathbf{Y} - \mathbf{X}\mathbf{B}) \quad (16)$$

$$- \frac{1}{2} \sum_{i=1}^m (\nu_i + d + 1) \ln \det \Gamma_i - \frac{1}{2} \sum_{i=1}^m \text{tr}(\Psi_i \Gamma_i^{-1}) + c,$$

where c is an irrelevant constant. Invoking the minorizations (6) and (12) for the first two terms and the supporting hyperplane minorization for $-\ln \det \Gamma_i$ gives the surrogate function

$$g(\Gamma | \Gamma^{(t)}) = -\frac{1}{2} \sum_{i=1}^m \text{tr}(\Omega^{-(t)} (\Gamma_i \otimes \mathbf{v}_i)) - \frac{1}{2} \sum_{i=1}^m \text{tr}(\Gamma_i^{(t)} \mathbf{R}^{(t)T} \mathbf{v}_i \mathbf{R}^{(t)} \Gamma_i^{(t)} \Gamma_i^{-1})$$

$$- \frac{1}{2} \sum_{i=1}^m (\nu_i + d + 1) \text{tr}(\Gamma_i^{-(t)} \Gamma_i) - \frac{1}{2} \sum_{i=1}^m \text{tr}(\Psi_i \Gamma_i^{-1}) + c^{(t)}.$$

The optimal Γ_i satisfies the stationarity condition

$$\begin{aligned} & (I_d \otimes 1_n)^T \left[(1_d 1_d^T \otimes V_i) \odot \Omega^{-(t)} \right] (I_d \otimes 1_n) + (V_i + d + 1) \Gamma_i^{-(t)} \\ &= \Gamma_i^{-1} \left(\Gamma_i^{(t)} R^{(t)T} V_i R^{(t)} \Gamma_i^{(t)} + \Psi_i \right) \Gamma_i^{-1}, \end{aligned}$$

which can be solved by Lemma 3.

6.5 Variable selection

In the statistical analysis of high-dimensional data, the imposition of sparsity leads to better interpretation and more stable parameter estimation. MM algorithms mesh well with penalized estimation. The simple variance components model (1) illustrates this fact. For the selection of fixed effects, minimizing the lasso-penalized log-likelihood $-L(\beta, \sigma^2) + \lambda \sum_j |\beta_j|$ is often recommended (Schelldorfer et al., 2011). The only change to the MM Algorithm 1 is that in estimating β , one solves a lasso penalized general least squares problem rather than an ordinary general least squares problem. The updates of the variance components σ_i^2 remain the same. For estimation of a large number of variance components, one can minimize the ridge-penalized log-likelihood

$$-L(\beta, \sigma^2) + \lambda \sum_{i=1}^m \sigma_i^2$$

subject to the nonnegativity constraints $\sigma_i^2 \geq 0$. The variance update (8) becomes

$$\sigma_i^{2(t+1)} = \sigma_i^{2(t)} \sqrt{\frac{(y - X\beta^{(t)})^T \Omega^{-(t)} V_i \Omega^{-(t)} (y - X\beta^{(t)})}{\text{tr}(\Omega^{-(t)} V_i) + 2\lambda}}, \quad i = 1, \dots, m,$$

which clearly exhibits shrinkage but no thresholding. The lasso penalized log-likelihood

$$-L(\beta, \sigma^2) + \lambda \sum_{i=1}^m \sigma_i \quad (17)$$

subject to nonnegativity constraint $\sigma_i \geq 0$ achieves both ends. The update of σ_i is chosen among the positive roots of a quartic equation and the boundary 0, whichever yields a lower objective value. Next section illustrates variance component selection using lasso penalty on a real genetic data set.

7 A numerical example

Quantitative trait loci (QTL) mapping aims to identify genes associated with a quantitative trait. Current sequencing technology measures millions of genetic markers in study subjects. Traditional single-marker tests suffer from low power due to the low frequency of many

markers and the corrections needed for multiple hypothesis testing. Region-based association tests are a powerful alternative for analyzing next generation sequencing data with abundant rare variants.

Suppose \mathbf{y} is a $n \times 1$ vector of quantitative trait measurements on n people, \mathbf{X} is an $n \times p$ predictor matrix (incorporating predictors such as sex, smoking history, and principal components for ethnic admixture), and \mathbf{G} is an $n \times m$ genotype matrix of m genetic variants in a pre-defined region. The linear mixed model assumes

$$Y = X\beta + G\gamma + \epsilon, \quad \gamma \sim N(0, \sigma_g^2 \mathbf{I}), \quad \epsilon \sim N(0, \sigma_e^2 \mathbf{I}_n),$$

where β are fixed effects, γ are random genetic effects, and σ_g^2 and σ_e^2 are variance components for the genetic and environmental effects, respectively. Thus, the phenotype vector \mathbf{Y} has covariance $\sigma_g^2 \mathbf{G}\mathbf{G}^T + \sigma_e^2 \mathbf{I}_n$, where $\mathbf{G}\mathbf{G}^T$ is the kernel matrix capturing the overall effect of the m variants. Current approaches test the null hypothesis $\sigma_g^2 = 0$ for each region separately and then adjust for multiple testing (Lee et al., 2014; Zhou et al., 2016). Instead of this marginal testing strategy, we consider the joint model

$$y = X\beta + s_1^{-1/2} \mathbf{G}_1 \gamma_1 + \dots + s_m^{-1/2} \mathbf{G}_m \gamma_m + \epsilon, \\ \gamma_i \sim N(0, \sigma_i^2 \mathbf{I}), \quad \epsilon \sim N(0, \sigma_e^2 \mathbf{I}_n)$$

and select the variance components σ_i^2 via the penalization (17). Here s_i is the number of variants in region i , and the weights $s_i^{-1/2}$ put all variance components on the same scale.

We illustrate this approach using the COPDGene exome sequencing study (Regan et al., 2010). After quality control, 399 individuals and 646,125 genetic variants remain for analysis. Genetic variants are grouped into 16,619 genes to expose those genes associated with the complex trait height. We include age, sex, and the top 3 principal components in the mean effects. Because the number of genes vastly exceeds the sample size $n = 399$, we first pare the 16,619 genes down to 200 genes according to their marginal likelihood ratio test p-values and then carry out penalized estimation of the 200 variance components in the joint model (17). This is similar to the sure independence screening strategy for selecting mean effects (Fan and Lv, 2008). Genes are ranked according to the order they appear in the lasso solution path. Table 4 lists the top 10 genes together with their marginal LRT p-values. Figure 1 in Supplementary Materials displays the corresponding segment of the lasso solution path. It is noteworthy that the ranking of genes by penalized estimation differs from the ranking according to marginal p-values. The same phenomenon occurs in selection of highly correlated mean predictors. This penalization approach for selecting variance components warrants further theoretical study.

8 Discussion

The current paper leverages the MM principle to design powerful and versatile algorithms for variance components estimation. The MM algorithms derived are notable for their simplicity, generality, numerical efficiency, and theoretical guarantees. Both ordinary MLE and REML are apt to benefit. Other extensions are possible. In nonlinear models (Bates and Watts, 1988; Lindstrom and Bates, 1990), the mean response is a nonlinear function in the fixed effects $\boldsymbol{\beta}$. One can easily modify the MM algorithms to update $\boldsymbol{\beta}$ by a few rounds of Gauss-Newton iteration. The variance components updates remain unchanged.

One can also extend our MM algorithms to elliptically symmetric densities

$$f(\mathbf{y}) = \frac{e^{-\frac{1}{2}\kappa(\delta^2)}}{(2\pi)^{\frac{n}{2}}(\det\boldsymbol{\Omega})^{\frac{1}{2}}}$$

defined for $\mathbf{y} \in \mathbb{R}^n$, where $\delta^2 = (\mathbf{y} - \boldsymbol{\mu})^T \boldsymbol{\Omega}^{-1}(\mathbf{y} - \boldsymbol{\mu})$ denotes the Mahalanobis distance between \mathbf{y} and $\boldsymbol{\mu}$. Here we assume that the function $\kappa(s)$ is strictly increasing and strictly concave. Examples of elliptically symmetric densities include the multivariate t , slash, contaminated normal, power exponential, and stable families. Previous work (Huber and Ronchetti, 2009; Lange and Sinsheimer, 1993) has focused on using the MM principle to convert parameter estimation for these robust families into parameter estimation under the multivariate normal. One can chain the relevant majorization $\kappa(s) \leq \kappa(s^{(t)}) + \kappa'(s^{(t)})(s - s^{(t)})$ with our previous minorizations and simultaneously split variance components and pass to the more benign setting of the multivariate normal. These extensions are currently under investigation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The research is partially supported by NIH grants R01HG006139, R01GM53275, R01GM105785 and K01DK106116. The authors thank Michael Cho, Dandi Qiao, and Edwin Silverman for their assistance in processing and assessing COPDGene exome sequencing data. COPDGene is supported by NIH grants R01HL089897 and R01HL089856.

References

- Bates D, Mächler M, Bolker B, and Walker S (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.
- Bates D and Pinheiro J (1998). Computational methods for multilevel models. Technical Report Technical Memorandum BL0112140–980226-01TM, Bell Labs, Lucent Technologies, Murray Hill, NJ.
- Bates DM and Watts DG (1988). *Nonlinear Regression Analysis and Its Applications*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics John Wiley & Sons, Inc, New York.

- Bien J and Tibshirani RJ (2011). Sparse estimation of a covariance matrix. *Biometrika*, 98(4):807–820. [PubMed: 23049130]
- Borg I and Groenen PJ (2005). *Modern Multidimensional Scaling: Theory and Applications* Springer Science & Business Media.
- Boyd S and Vandenberghe L (2004). *Convex Optimization* Cambridge University Press, Cambridge.
- Callanan TP and Harville DA (1991). Some new algorithms for computing restricted maximum likelihood estimates of variance components. *J. Statist. Comput. Simulation*, 38(1–4):239–259.
- De Leeuw J (1994). Block-relaxation algorithms in statistics. In *Information Systems and Data Analysis*, pages 308–324. Springer.
- Demidenko E and Massam H (1999). On the existence of the maximum likelihood estimate in variance components models. *Sankhyā Ser. A*, 61(3):431–443.
- Dempster A, Laird N, and Rubin D (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B*, 39(1–38).
- Ding J, Tian G-L, and Yuen KC (2015). A new MM algorithm for constrained estimation in the proportional hazards model. *Comput. Statist. Data Anal*, 84:135–151.
- Fan J and Lv J (2008). Sure independence screening for ultrahigh dimensional feature space (with discussion). *J. R. Statist. Soc. B*, 70:849–911.
- Golub GH and Van Loan CF (1996). *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences Johns Hopkins University Press, Baltimore, MD, third edition.
- Grzadziel M and Michalski A (2014). A note on the existence of the maximum likelihood estimate in variance components models. *Discuss. Math. Probab. Stat*, 34(1–2):159–167.
- Gupta A and Nagar D (1999). *Matrix Variate Distributions*. Monographs and Surveys in Pure and Applied Mathematics Taylor & Francis.
- Hartley HO and Rao JNK (1967). Maximum-likelihood estimation for the mixed analysis of variance model. *Biometrika*, 54:93–108. [PubMed: 6049561]
- Harville D and Callanan T (1990). Computational aspects of likelihood-based inference for variance components. In Gianola D and Hammond K, editors, *Advances in Statistical Methods for Genetic Improvement of Livestock*, volume 18 of *Advanced Series in Agricultural Sciences*, pages 136–176. Springer Berlin Heidelberg.
- Harville DA (1977). Maximum likelihood approaches to variance component estimation and to related problems. *J. Amer. Statist. Assoc*, 72(358):320–340. With a comment by J. N. K. Rao and a reply by the author.
- Heiser WJ (1995). Convergent computation by iterative majorization: theory and applications in multidimensional data analysis. *Recent Advances in Descriptive Multivariate Analysis*, pages 157–189.
- Horn RA and Johnson CR (1985). *Matrix Analysis* Cambridge University Press, Cambridge.
- Huber PJ and Ronchetti EM (2009). *Robust Statistics*. Wiley Series in Probability and Statistics John Wiley & Sons, Inc, Hoboken, NJ, second edition.
- Hunter DR (2004). MM algorithms for generalized Bradley-Terry models. *Ann. Statist*, 32(1):384–406.
- Hunter DR and Lange K (2002). Computing estimates in the proportional odds model. *Ann. Inst. Statist. Math*, 54(1):155–168.
- Hunter DR and Lange K (2004). A tutorial on MM algorithms. *Amer. Statist*, 58(1):30–37.
- Hunter DR and Li R (2005). Variable selection using MM algorithms. *Ann. Statist*, 33(4):1617–1642.
- Jeon M (2012). *Estimation of Complex Generalized Linear Mixed Models for Measurement and Growth* PhD thesis, University of California, Berkeley.
- Kiers HA (2002). Setting up alternating least squares and iterative majorization algorithms for solving various matrix optimization problems. *Computational Statistics & Data Analysis*, 41(1):157–170.
- Laird N, Lange N, and Stram D (1987). Maximum likelihood computations with repeated measures: application of the EM algorithm. *J. Amer. Statist. Assoc*, 82(397):97–105.
- Laird NM and Ware JH (1982). Random-effects models for longitudinal data. *Biometrics*, 38(4):963–974. [PubMed: 7168798]

- Lange K (2010). *Numerical Analysis for Statisticians*. Statistics and Computing Springer, New York, second edition.
- Lange K (2016). *MM Optimization Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Lange K, Hunter DR, and Yang I (2000). Optimization transfer using surrogate objective functions. *J. Comput. Graph. Statist*, 9(1):1–59. With discussion, and a rejoinder by Hunter and Lange.
- Lange K, Papp J, Sinsheimer J, Sripracha R, Zhou H, and Sobel E (2013). Mendel: the Swiss army knife of genetic analysis programs. *Bioinformatics*, 29:1568–1570. [PubMed: 23610370]
- Lange K and Sinsheimer JS (1993). Normal/independent distributions and their applications in robust regression. *Journal of Computational and Graphical Statistics*, 2:175–198.
- Lange K and Zhou H (2014). MM algorithms for geometric and signomial programming. *Mathematical Programming Series A*, 143:339–356.
- Lee S, Abecasis GR, Boehnke M, and Lin X (2014). Rare-variant association analysis: Study designs and statistical tests. *The American Journal of Human Genetics*, 95(1):5–23. [PubMed: 24995866]
- Lindstrom MJ and Bates DM (1988). Newton-Raphson and EM algorithms for linear mixed-effects models for repeated-measures data. *J. Amer. Statist. Assoc*, 83(404):1014–1022.
- Lindstrom MJ and Bates DM (1990). Nonlinear mixed effects models for repeated measures data. *Biometrics*, 46(3):673–687. [PubMed: 2242409]
- McLachlan GJ and Krishnan T (2008). *The EM Algorithm and Extensions*. Wiley Series in Probability and Statistics Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition.
- Meng X-L and Rubin DB (1991). Using EM to obtain asymptotic variance-covariance matrices: the SEM algorithm. *Journal of the American Statistical Association*, 86(416):899–909.
- Pinheiro J and Bates D (1996). Unconstrained parametrizations for variance-covariance matrices. *Statistics and Computing*, 6(3):289–296.
- Rao CR (1973). *Linear Statistical Inference and its Applications*, 2nd ed John Wiley & Sons.
- Regan EA, Hokanson JE, Murphy JR, Make B, Lynch DA, Beaty TH, Curran-Everett D, Silverman EK, and Crapo JD (2010). Genetic epidemiology of COPD (COPDGene) study designs. *COPD*, 7:32–43. [PubMed: 20214461]
- Schafer JL and Yucel RM (2002). Computational strategies for multivariate linear mixed-effects models with missing values. *J. Comput. Graph. Statist*, 11(2):437–457.
- Schelldorfer J, Bühlmann P, and van de Geer S (2011). Estimation for high-dimensional linear mixed-effects models using A1-penalization. *Scand. J. Stat*, 38(2):197–214.
- Schur J (1911). Bemerkungen zur Theorie der beschränkten Bilinearformen mit unendlich vielen Veränderlichen. *J. Reine Angew. Math*, 140:1–28.
- Wu TT and Lange K (2010). The MM alternative to EM. *Statistical Science*, 25:492–505.
- Yen T-J (2011). A majorization-minimization approach to variable selection using spike and slab priors. *Ann. Statist*, 39(3):1748–1775.
- Yu Y (2010). Monotonic convergence of a general algorithm for computing optimal designs. *Ann. Statist*, 38(3):1593–1606.
- Zhou H and Lange K (2010). MM algorithms for some discrete multivariate distributions. *Journal of Computational and Graphical Statistics*, 19:645–665. [PubMed: 20877446]
- Zhou JJ, Hu T, Qiao D, Cho MH, and Zhou H (2016). Boosting gene mapping power and efficiency with efficient exact variance component tests of single nucleotide polymorphism sets. *Genetics*, 204(3):921–931. [PubMed: 27646141]




Figure 1: Surrogate functions of EM and MM minorize the log-likelihood surface of a 2-variance component model at point $(\sigma_1^{2(t)}, \sigma_2^{2(t)}) = (18.5, 0.7)$. MM surrogate function hugs the log-likelihood surface tighter than EM.

Table 1:

Fisher scoring converges fastest and MM takes fewer iterations than EM. Shown below are average number of iterations until convergence for MM, EM, and FS for fitting a two-way ANOVA model with $a = b = 5$ levels of both factors. Standard errors are given in parentheses.

σ_1^2/σ_e^2	Method	$c = \# \text{ observations per combination}$			
		5	10	20	50
0.00	MM	143.12(99.76)	118.26(62.91)	96.26(50.61)	81.10(33.42)
	EM	2297.72(797.95)	1711.70(485.92)	1170.06(365.48)	788.10(216.60)
	FS	25.64(11.20)	21.10(7.00)	16.46(4.37)	13.88(2.88)
0.05	MM	121.86(98.52)	69.38(50.23)	55.88(37.34)	29.50(18.80)
	EM	1464.26(954.27)	538.04(504.42)	254.90(253.86)	104.98(157.97)
	FS	16.78(9.13)	12.62(6.22)	9.68(3.22)	8.10(1.34)
0.10	MM	84.74(59.33)	62.98(50.48)	40.46(31.43)	25.86(18.79)
	EM	985.46(830.49)	360.32(462.62)	157.70(231.91)	68.26(107.85)
	FS	15.20(10.10)	10.58(5.92)	8.58(3.56)	7.50(1.72)
1.00	MM	31.04(33.27)	29.60(27.66)	25.32(25.39)	24.90(20.76)
	EM	130.18(299.03)	161.14(290.23)	64.20(135.38)	84.88(137.88)
	FS	6.62(4.72)	6.32(3.64)	5.12(1.87)	5.36(1.50)
10.00	MM	29.80(35.42)	34.16(38.25)	28.82(28.44)	20.90(14.28)
	EM	115.94(274.33)	177.30(301.71)	80.12(155.67)	75.02(127.38)
	FS	12.72(5.14)	12.86(4.94)	11.66(3.95)	11.76(3.66)
20.00	MM	30.10(32.92)	32.72(39.02)	23.70(21.20)	19.62(15.67)
	EM	148.04(318.40)	85.86(180.28)	61.74(140.84)	37.36(83.89)
	FS	18.76(7.51)	17.40(5.21)	17.22(5.67)	16.28(5.03)

Table 2:

MM shows shortest run times than EM, Fisher scoring (FS), and lme4. Shown below are average run times (milliseconds) for fitting a two-way ANOVA model with $a = b = 5$ levels of both factors. Standard errors are given in parentheses.

σ_1^2/σ_e^2	Method	$c = \# \text{ observations per combination}$			
		5	10	20	50
0.00	MM	11.46(7.77)	10.06(5.29)	11.93(6.35)	10.44(3.99)
	EM	189.32(71.32)	148.20(48.13)	147.87(49.97)	96.28(24.97)
	FS	34.27(33.47)	24.89(8.55)	23.70(14.15)	20.46(4.54)
	lme4	25.84(12.10)	22.32(1.25)	27.34(4.06)	36.14(5.59)
0.05	MM	9.79(7.72)	6.19(4.22)	6.87(4.37)	4.45(2.20)
	EM	116.03(75.57)	47.72(45.35)	30.60(29.88)	14.23(19.68)
	FS	19.18(10.23)	15.37(7.48)	12.78(4.06)	12.39(2.35)
	lme4	22.76(1.96)	24.88(2.60)	28.72(3.10)	47.34(16.29)
0.10	MM	7.07(4.78)	6.29(4.94)	5.14(3.72)	3.95(2.23)
	EM	78.96(66.19)	35.48(45.81)	19.53(27.71)	9.67(13.56)
	FS	17.36(11.26)	14.44(9.00)	12.08(6.31)	11.47(2.40)
	lme4	22.66(1.83)	28.90(8.70)	30.16(4.43)	44.58(4.89)
1.00	MM	2.66(2.61)	3.22(2.91)	3.57(3.15)	3.85(2.50)
	EM	10.71(23.93)	15.88(27.52)	8.35(16.26)	11.34(16.65)
	FS	7.88(5.44)	9.10(4.95)	7.12(2.42)	8.46(2.27)
	lme4	23.12(1.75)	30.22(9.37)	29.96(4.47)	42.82(8.32)
10.00	MM	2.48(2.72)	3.24(3.19)	3.84(3.35)	3.35(1.71)
	EM	9.66(22.02)	15.98(26.57)	10.24(18.78)	10.27(15.40)
	FS	15.19(6.05)	16.39(6.11)	15.81(5.15)	18.14(5.46)
	lme4	35.02(3.83)	47.12(8.10)	63.24(15.33)	102.78(34.49)
20.00	MM	2.57(2.49)	3.13(3.53)	3.13(2.44)	3.07(1.81)
	EM	12.28(25.71)	8.44(16.89)	8.01(17.12)	5.47(9.76)
	FS	22.09(8.53)	22.03(6.14)	23.08(7.21)	23.99(7.38)
	lme4	37.34(12.91)	50.24(8.59)	63.62(17.39)	91.14(28.39)

Table 3:

MM and Fisher scoring (FS) show superior performance than EM and lme4. Shown below are average performance for fitting a genetic model. Standard errors are given in parentheses.

σ_d^2/σ_e^2	Method	Iteration	Runtime (ms)	Objective
0.00	MM	198.02(102.23)	133.61(822.67)	-375.59(9.63)
	EM	1196.10(958.51)	29.71(12.34)	-375.60(9.64)
	FS	7.60(3.07)	19.34(33.77)	-375.59(9.63)
	lme4	-	401.02(142.04)	-375.59(9.64)
0.05	MM	185.86(99.41)	17.26(1.76)	-377.39(10.52)
	EM	1227.62(1030.07)	29.82(12.74)	-377.40(10.52)
	FS	7.84(2.74)	14.97(1.55)	-377.39(10.52)
	lme4	-	425.04(144.00)	-377.39(10.52)
0.10	MM	169.24(99.75)	16.97(1.59)	-378.40(11.44)
	EM	924.80(912.23)	26.06(11.26)	-378.41(11.45)
	FS	7.32(2.75)	15.06(1.38)	-378.40(11.44)
	lme4	-	435.14(128.87)	-378.40(11.44)
1.00	MM	58.96(23.69)	15.53(0.75)	-409.54(10.90)
	EM	105.10(79.65)	15.49(0.96)	-409.54(10.90)
	FS	5.80(1.05)	14.66(0.89)	-409.54(10.90)
	lme4	-	493.14(52.80)	-409.54(10.90)
10.00	MM	110.00(63.13)	16.22(1.12)	-532.48(8.77)
	EM	642.48(1470.38)	22.32(18.37)	-532.57(8.75)
	FS	14.98(5.21)	14.78(0.97)	-531.72(8.92)
	lme4	-	2897.12(15006.38)	-532.48(8.77)
20.00	MM	110.52(34.81)	16.07(0.91)	-590.87(7.15)
	EM	1014.22(1775.40)	27.03(22.33)	-590.89(7.15)
	FS	17.72(3.13)	14.79(0.93)	-588.46(7.27)
	lme4	-	5059.24(20692.67)	-590.79(7.15)

Table 4:

Top 10 genes selected by the lasso penalized variance component model (17) are tallied with their marginal p-values in an association study of 200 genes and the complex trait height.

Lasso Rank	Gene	Marginal P-value	# Variants
1	DOLPP1	2.35×10^{-6}	2
2	C9orf21	3.70×10^{-5}	4
3	PLS1	2.29×10^{-3}	5
4	ATP5D	6.80×10^{-7}	3
5	ADCY4	1.01×10^{-3}	11
6	SLC22A25	3.95×10^{-3}	14
7	RCSD1	9.04×10^{-4}	4
8	PCDH7	1.20×10^{-4}	7
9	AVIL	8.34×10^{-4}	11
10	AHR	1.14×10^{-3}	7

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript