

Sparse Matrix Representations of Linear Mixed Models

Douglas Bates
R Development Core Team
Douglas.Bates@R-project.org

June 7, 2004

Abstract

We describe a representation of linear mixed-effects models using positive semidefinite, symmetric, compressed, column-oriented, sparse matrices. This representation provides for efficient evaluation of the profiled log-likelihood or profiled restricted log-likelihood of the model, given the relative precision parameters of the random effects. The evaluation is based upon the LDL^T form of the Cholesky decomposition of the augmented sparse representation. Additionally, we can use information from this representation to evaluate ECME updates and the gradient of the criterion being optimized.

The sparse matrix methods that we employ have both a symbolic phase, in which the number and positions of the off-diagonal elements are determined, and a numeric phase, in which the actual numeric values are determined. The symbolic phase need only be done once and it can be accomplished knowing only the grouping factors with which the random effects are associated. An important part of the symbolic phase is determination of a fill-minimizing permutation of the rows and columns of a sparse semi-definite matrix. This matrix have a special structure in the linear mixed-effects problem and we provide a new fill-minimizing algorithm tuned to this structure.

1 Introduction

Mixed-effects models, also called multilevel models, panel data models, and frailty models, are widely used in many areas of statistical applications (Pinheiro and Bates, 2000). The basic form of the model, the linear mixed model, also serves as an approximation in iterative estimation of the parameters in more general forms such as the generalized linear mixed model (GLMM) and the nonlinear mixed model (NMM).

In §2 we define a general form of a linear mixed model using grouping factors and model matrices that are associated with the grouping factors. This form, which can be used for multiple levels of random effects in either nested or crossed configurations, can be represented and manipulated using a sparse, symmetric, semidefinite matrix and several dense matrices. We show that a profiled log-likelihood can be evaluated from that solution of a penalized least squares problem and that this solution can be obtained from the Cholesky decomposition of an augmented form of the sparse, symmetric matrix.

Many implementations of the Cholesky decomposition of sparse, symmetric, semidefinite matrices have both a symbolic phase, in which the number and positions of the off-diagonal elements are determined, and a numeric phase, in which the actual numeric values are determined. In §3 we show that the symbolic analysis for the matrices we consider need only be done once and can be accomplished knowing only the grouping factors. An important part of the symbolic phase is determination of a fill-reducing permutation of the rows and columns of the symmetric matrix. We show that by suitably ordering the grouping factors and by restricting ourselves to permutations that correspond to reorderings of the levels within the grouping factors we can determine effective fill-reducing orderings.

Finally, in §?? we show how these methods can be used to implement general penalized least squares approaches to models such as the GLMM and the NMM and then to implement more accurate approximations to the marginal likelihood using Laplacian integration or adaptive Gauss-Hermite integration.

2 Linear mixed models

We describe the form of the linear mixed-effects model that we consider and restate some of the formulas from [Bates and DebRoy \(2004\)](#) using the LDL form of the Cholesky decomposition of a sparse, semi-definite matrix.

2.1 Form of the model

We consider linear mixed-effects models that can be written as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{b} + \boldsymbol{\epsilon} \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \mathbf{b} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{\Omega}^{-1}), \boldsymbol{\epsilon} \perp \mathbf{b} \quad (1)$$

where \mathbf{y} is the n -dimensional response vector, \mathbf{X} is an $n \times p$ model matrix for the p -dimensional fixed-effects vector $\boldsymbol{\beta}$, \mathbf{Z} is the $n \times q$ model matrix for the q -dimensional random-effects vector \mathbf{b} that has a Gaussian distribution with mean $\mathbf{0}$ and relative precision matrix $\boldsymbol{\Omega}$ (i.e., $\boldsymbol{\Omega}$ is the precision of \mathbf{b} relative to the precision of $\boldsymbol{\epsilon}$), and $\boldsymbol{\epsilon}$ is the random noise assumed to have a spherical Gaussian distribution. The symbol \perp indicates independence of random variables. We assume that \mathbf{X} has full column rank and that $\boldsymbol{\Omega}$, which is a function of an (unconstrained) parameter vector $\boldsymbol{\theta}$, is positive definite.

2.1.1 Grouping factors for the random effects

Although q , the dimension of the vector \mathbf{b} (and, correspondingly, the number of columns in \mathbf{Z} and the number of rows and columns in $\mathbf{Z}^T \mathbf{Z}$ and $\boldsymbol{\Omega}$) can be very large, these vectors and matrices can be divided into components associated with grouping factors $\mathbf{f}_i, i = 1, \dots, k$ in the data. Each grouping factor is of length n , the same as the length of \mathbf{y} . The number of distinct values in \mathbf{f}_i , also called the number of *levels* of \mathbf{f}_i , is $m_i, i = 1, \dots, k$. In the general form of the model, a model matrix \mathbf{Z}_i of size $n \times q_i$ is associated with grouping factor $\mathbf{f}_i, i = 1, \dots, k$. Typically the q_i are very small. In fact, in one common form of the model, called a *variance components* model, $q_1 = q_2 = \dots = q_k = 1$ and each of the $\mathbf{Z}_i, i = 1, \dots, k$ consist of a single column of 1's.

In the general form, the random effects vector \mathbf{b} , of length $q = \sum_{i=1}^k m_i q_i$, is partitioned into k “outer blocks” where the i 'th outer block is of size $m_i q_i, i = 1, \dots, k$. The columns of \mathbf{Z} and the rows and columns of $\mathbf{Z}^T \mathbf{Z}$ and $\boldsymbol{\Omega}$ are similarly partitioned. Each of the outer blocks is further subdivided into m_i inner blocks of size q_i . The grouping factors determine the outer blocks and the levels of each grouping factor determine the inner blocks.

In the models that we will consider, the random effects associated with different grouping factors are independent. That is, $\mathbf{\Omega}$ is block-diagonal in k blocks of sizes $m_i q_i \times m_i q_i, i = 1, \dots, k$. Furthermore, the random effects associated with the levels of a given blocking factor are independent and identically distributed. Thus the i 'th diagonal block in $\mathbf{\Omega}$ is itself block diagonal and these diagonal blocks are m_i repetitions of a $q_i \times q_i$ matrix $\mathbf{\Omega}_i, i = 1, \dots, k$, providing

$$|\mathbf{\Omega}| = \sum_{i=1}^k m_i |\mathbf{\Omega}_i| \quad (2)$$

For a variance components model the matrices $\mathbf{\Omega}_i, i = 1, \dots, k$ are 1×1 positive definite matrices which we can consider to be positive scalars $\omega_i, i = 1, \dots, k$. The matrix $\mathbf{\Omega}$ is block-diagonal of size $\sum_{i=1}^k m_i$ and the diagonal blocks are $\omega_i \mathbf{I}_{m_i}$ where \mathbf{I}_{m_i} is the $m_i \times m_i$ identity matrix. Thus $|\mathbf{\Omega}| = \sum_{i=1}^k m_i \omega_i$. The k -dimensional vector $\boldsymbol{\theta}$ where $\theta_i = \log \omega_i, i = 1, \dots, k$ can be used as the unconstrained parameter vector.

The columns of the matrix \mathbf{Z} are similarly divided into blocks. For the variance components model the i th block is the set of indicator columns for the m_i levels of $\mathbf{f}_i, i = 1, \dots, k$. Because each block is a set of indicators, the diagonal blocks of $\mathbf{Z}^T \mathbf{Z}$ are themselves diagonal. However, unlike the corresponding blocks in $\mathbf{\Omega}$, these blocks are not necessarily a multiple of the identity. The diagonal elements of the i th diagonal block are the m_i frequencies of occurrence of each the levels of the i th grouping factor in the data. (Because all the elements of \mathbf{Z} are zero or one, the diagonals of $\mathbf{Z}^T \mathbf{Z}$ are simply the counts of the number of ones in the corresponding column of \mathbf{Z} .)

The off-diagonal blocks of $\mathbf{Z}^T \mathbf{Z}$ in a variance components model are the pairwise cross-tabulations of the corresponding grouping factors.

2.1.2 The Scottish secondary school example

An example may help to clarify these descriptions.

Data on achievement scores of Scottish secondary school students are described in [Paterson \(1991\)](#) and are analyzed in [Rasbash et al. \(2002, ch. 18\)](#) and other references. In the `Matrix` package for R these data are available as the data set `ScotsSec` containing the achievement scores (`attain`), some demographic data (`sex` and `social class`), a `verbal` reasoning score based on

tests taken at entry to secondary school, and the *primary* and secondary (*second*) schools attended by 3435 students.

The grouping factors for the random effects are *primary* (148 distinct schools) and *second* (19 distinct schools). The locations of the non-zeros in the 167×167 matrix $\mathbf{Z}^\top \mathbf{Z}$ are shown in Figure 1. For a variance components model with these grouping factors. In the figure darker greys indicate larger magnitudes.

2.1.3 General structure of the sparse matrix

For the variance components model $b\mathbf{Z}^\top \mathbf{Z}$ is based on the pairwise cross-tabulation of the grouping factors. In the more general model where some of the \mathbf{Z}_i can have multiple columns, the structure of $\mathbf{Z}^\top \mathbf{Z}$ can be derived from the structure of the pairwise cross-tabulation matrix. Both $\mathbf{Z}^\top \mathbf{Z}$ and the pairwise cross-tabulation can be divided into a $k \times k$ grid of blocks. The pattern of non-zeros in the (i, j) block of $\mathbf{Z}^\top \mathbf{Z}$ is obtained by replacing each non-zero in the (i, j) block of the cross-tabulation by a $q_i \times q_j$ matrix. Notice that we can determine the patterns of non-zeros in $\mathbf{Z}^\top \mathbf{Z}$ knowing only the $q_i, i = 1, \dots, k$ and the cross-tabulation of the grouping factors.

2.1.4 Crossed and nested grouping factors

In the Scottish secondary school example if all the students from a given primary school attended the same secondary school we would say that *primary* is *nested within second*. That is not the case. We can see in Figure 1 that there is a moderate amount of *crossing* of these two grouping factors. If there was at least one student in the study from each combination of primary school and secondary school we would describe the grouping factors *primary* and *second* as being *fully crossed*. Again, that is not the case for the Scottish secondary data. Grouping factors like these, which are neither nested nor fully crossed, are said to be *partially crossed*.

2.2 Estimation criteria and related quantities

For ease of reference we restate some of the results from [Bates and DebRoy \(2004\)](#) in the form in which they will be calculated.

Given the observed responses \mathbf{y} and the model matrices \mathbf{X} and \mathbf{Z} , we wish to determine either the maximum likelihood (ML) or the restricted maximum

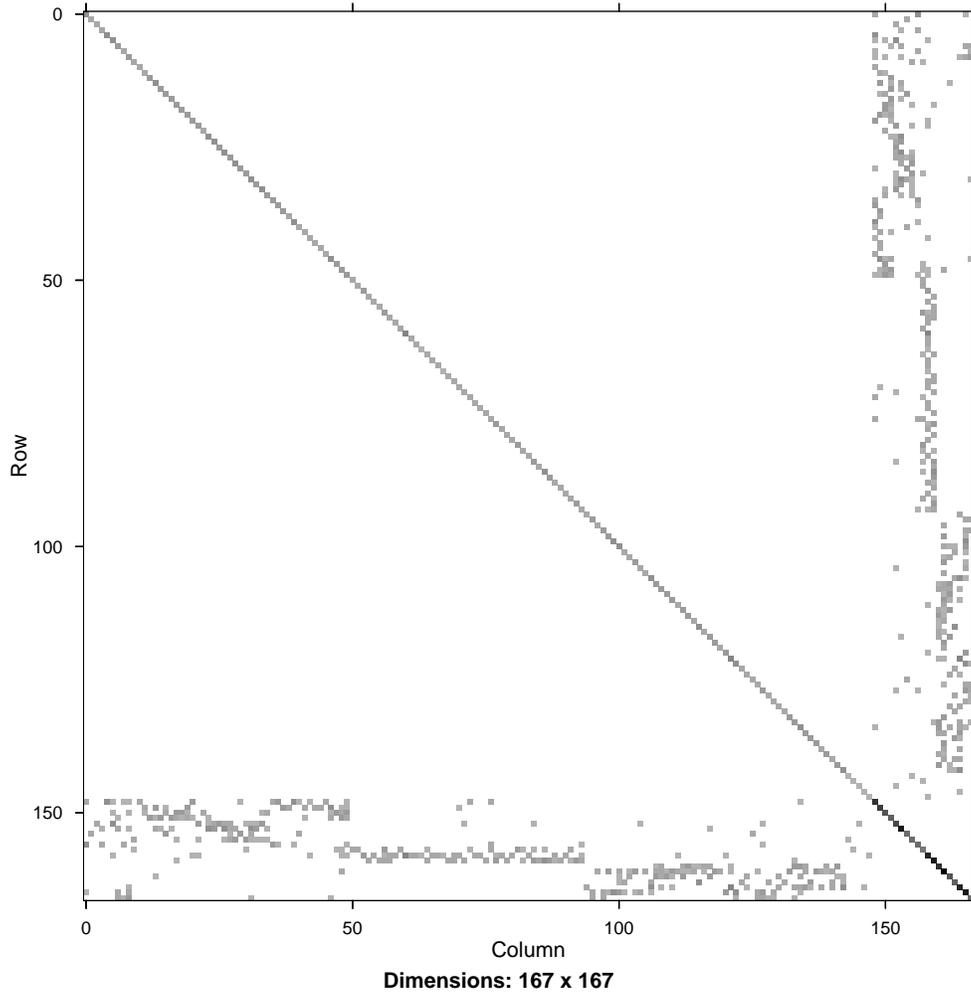


Figure 1: Location of non-zero elements in $\mathbf{Z}^T \mathbf{Z}$ for a variance components model for the *ScotsSec* data. Darker squares indicate larger magnitudes. Rows and columns are numbered from zero. The first 148 rows and columns correspond to the levels of the *primary* grouping factor and the last 19 rows and columns correspond to levels of the *second* grouping factor.

likelihood (REML) estimates of the parameters $\boldsymbol{\theta}$, $\boldsymbol{\beta}$, and σ^2 . Because the conditional estimates of $\boldsymbol{\beta}$ and σ^2 , given a value of $\boldsymbol{\theta}$, for either criterion can be determined from the solution to penalized least squares problem, we can reduce the optimization problem to one involving $\boldsymbol{\theta}$ only. This reduction of the dimension of the optimization problem is called *profiling* the objective.

The conditional, penalized least squares problem can be solved using the Cholesky decomposition

$$\begin{bmatrix} \mathbf{Z}^\top \mathbf{Z} + \boldsymbol{\Omega} & \mathbf{Z}^\top \mathbf{X} & \mathbf{Z}^\top \mathbf{y} \\ \mathbf{X}^\top \mathbf{Z} & \mathbf{X}^\top \mathbf{X} & \mathbf{X}^\top \mathbf{y} \\ \mathbf{y}^\top \mathbf{Z} & \mathbf{y}^\top \mathbf{X} & \mathbf{y}^\top \mathbf{y} \end{bmatrix} = \mathbf{R}^\top \mathbf{R} \quad \text{where} \quad \mathbf{R} = \begin{bmatrix} \mathbf{R}_{ZZ} & \mathbf{R}_{ZX} & \mathbf{r}_{Zy} \\ \mathbf{0} & \mathbf{R}_{XX} & \mathbf{r}_{Xy} \\ \mathbf{0} & \mathbf{0} & r_{yy} \end{bmatrix}. \quad (3)$$

where the matrices \mathbf{R}_{ZZ} and \mathbf{R}_{XX} are upper triangular of dimension $q \times q$ and $p \times p$ respectively. The corresponding vectors, \mathbf{r}_{Zy} and \mathbf{r}_{Xy} , are of dimension q and p , and r_{yy} is a scalar. The conditions that $\boldsymbol{\Omega}$ be positive definite and \mathbf{X} have full column rank ensure that \mathbf{R}_{ZZ} and \mathbf{R}_{XX} are nonsingular.

In our implementation we do not form the upper triangular Cholesky factor \mathbf{R}_{ZZ} . Instead we use Tim Davis's LDL package (Davis, 2004) to factor

$$\mathbf{Z}^\top \mathbf{Z} + \boldsymbol{\Omega} = \mathbf{L} \mathbf{D} \mathbf{L}^\top \quad (4)$$

where \mathbf{L} is a sparse, unit, lower triangular matrix and \mathbf{D} is diagonal with positive diagonal elements. Because the diagonal elements of the unit triangular matrix \mathbf{L} are, by definition, unity, they are not explicitly stored.

In general the matrices $\mathbf{Z}^\top \mathbf{X}$ and $\mathbf{X}^\top \mathbf{X}$ are dense. We use functions from the LDL package to solve for \mathbf{R}_{ZX} in

$$\mathbf{D}^{1/2} \mathbf{L}^\top \mathbf{R}_{ZX} = \mathbf{Z}^\top \mathbf{X} \quad (5)$$

Having solved for \mathbf{R}_{ZX} we can downdate $\mathbf{X}^\top \mathbf{X}$ and determine the dense Cholesky factor \mathbf{R}_{XX} in

$$\mathbf{X}^\top \mathbf{X} - \mathbf{R}_{ZX}^\top \mathbf{R}_{ZX} = \mathbf{R}_{XX}^\top \mathbf{R}_{XX} \quad (6)$$

Similar relationships are used to determine \mathbf{r}_{Zy} , \mathbf{r}_{Xy} , and r_{yy} . In fact, in our implementation we append \mathbf{y} to \mathbf{X} when forming $\mathbf{Z}^\top \mathbf{X}$ and $\mathbf{X}^\top \mathbf{X}$ so that (??) provides both \mathbf{R}_{ZX} and \mathbf{r}_{Zy} and (6) provides \mathbf{R}_{XX} , \mathbf{r}_{Xy} , and r_{yy} .

The conditional estimates of $\boldsymbol{\beta}$ satisfy

$$\mathbf{R}_{XX} \widehat{\boldsymbol{\beta}}(\boldsymbol{\theta}) = \mathbf{r}_{Xy} \quad (7)$$

and the conditional modes of the random effects satisfy

$$\mathbf{D}^{1/2} \mathbf{L}^\top \widehat{\mathbf{b}}(\boldsymbol{\theta}) = \mathbf{r}_{zy} - \mathbf{R}_{ZX} \widehat{\boldsymbol{\beta}}. \quad (8)$$

The conditional ML estimate of σ^2 is $\widehat{\sigma}^2(\boldsymbol{\theta}) = r_{yy}^2/n$ and the conditional REML estimate is $\widehat{\sigma}_R^2(\boldsymbol{\theta}) = r_{yy}^2/(n-p)$.

The profiled optimization problem, expressed in terms of the deviance, is

$$\begin{aligned} \widehat{\boldsymbol{\theta}} &= \arg \min_{\boldsymbol{\theta}} -2\tilde{\ell}(\boldsymbol{\theta}) \\ &= \arg \min_{\boldsymbol{\theta}} \left\{ \log \left(\frac{|\mathbf{D}|}{|\boldsymbol{\Omega}|} \right) + n \left[1 + \log \left(\frac{2\pi r_{yy}^2}{n} \right) \right] \right\} \end{aligned} \quad (9)$$

$$\begin{aligned} \widehat{\boldsymbol{\theta}}_R &= \arg \min_{\boldsymbol{\theta}} -2\tilde{\ell}_R(\boldsymbol{\theta}) \\ &= \arg \min_{\boldsymbol{\theta}} \left\{ \log \left(\frac{|\mathbf{D}| |\mathbf{R}_{XX}|^2}{|\boldsymbol{\Omega}|} \right) + (n-p) \left[1 + \log \left(\frac{2\pi r_{yy}^2}{n-p} \right) \right] \right\} \end{aligned} \quad (10)$$

for ML and REML estimation, respectively. The gradients of these criteria are

$$\nabla(-2\tilde{\ell}) = \text{tr} \left[\mathbf{D} \boldsymbol{\Omega} \left((\mathbf{Z}^\top \mathbf{Z} + \boldsymbol{\Omega})^{-1} - \boldsymbol{\Omega}^{-1} + \frac{\widehat{\mathbf{b}} \widehat{\mathbf{b}}^\top}{\widehat{\sigma} \widehat{\sigma}} \right) \right] \quad (11)$$

$$\nabla(-2\tilde{\ell}_R) = \text{tr} \left[\mathbf{D} \boldsymbol{\Omega} \left(\mathbf{V}_b - \boldsymbol{\Omega}^{-1} + \frac{\widehat{\mathbf{b}} \widehat{\mathbf{b}}^\top}{\widehat{\sigma}_R \widehat{\sigma}_R} \right) \right] \quad (12)$$

where

$$\mathbf{V}_b = \mathbf{L}^{-\top} \mathbf{D}^{-1/2} (\mathbf{I} + \mathbf{R}_{ZX} \mathbf{R}_{XX}^{-1} \mathbf{R}_{XX}^{-\top} \mathbf{R}_{ZX}^\top) \mathbf{D}^{-1/2} \mathbf{L}^{-1} \quad (13)$$

and \mathbf{D} denotes the Frechet derivative.

If good starting estimates of $\boldsymbol{\theta}$ are not available, the initial Newton iterations for (9) or (10) can be unstable. We can refine our initial estimates with a moderate number of ECME steps for which $\boldsymbol{\theta}_{i+1}$ satisfies

$$\text{tr} \left[\mathbf{D} \boldsymbol{\Omega} \left(\frac{\widehat{\mathbf{b}}(\boldsymbol{\theta}_i) \widehat{\mathbf{b}}(\boldsymbol{\theta}_i)^\top}{\widehat{\sigma}(\boldsymbol{\theta}_i) \widehat{\sigma}(\boldsymbol{\theta}_i)} + (\mathbf{Z}^\top \mathbf{Z} + \boldsymbol{\Omega}(\boldsymbol{\theta}_i))^{-1} - \boldsymbol{\Omega}(\boldsymbol{\theta}_{i+1})^{-1} \right) \right] = \mathbf{0} \quad (14)$$

for ML estimates or

$$\text{tr} \left[\mathbf{D} \boldsymbol{\Omega} \left(\frac{\widehat{\mathbf{b}}(\boldsymbol{\theta}_i) \widehat{\mathbf{b}}(\boldsymbol{\theta}_i)^\top}{\widehat{\sigma}_R(\boldsymbol{\theta}_i) \widehat{\sigma}_R(\boldsymbol{\theta}_i)} + \mathbf{V}_b(\boldsymbol{\theta}_i) - \boldsymbol{\Omega}(\boldsymbol{\theta}_{i+1})^{-1} \right) \right] = \mathbf{0} \quad (15)$$

for REML.

At this point it is easy to formulate a general method of obtaining ML or REML estimates for a linear mixed model:

1. Given the data \mathbf{y} and the model matrices \mathbf{X} and \mathbf{Z} , formulate initial estimates $\boldsymbol{\theta}_0$. Some heuristics for doing so are given in [Pinheiro and Bates \(2000, ch. 3\)](#).
2. Use a moderate number of ECME steps, (14) or (15), to refine these starting estimates. Each ECME step requires evaluating $\boldsymbol{\Omega}(\boldsymbol{\theta})$ followed by the decomposition (3) and the solutions to (5), (6), (7) and (8).
3. Use a Newton method to optimize the criterion (9) or (10) with gradient (11) or (12). Each evaluation of the criterion requires evaluating $\boldsymbol{\Omega}(\boldsymbol{\theta})$ followed by the decomposition (3) and the solutions to (5), and (6). Gradient evaluations require the solutions to (7) and (8).

In [Bates and DebRoy \(2004\)](#) we show that similar calculations can be used to evaluate the Hessian of the profiled criteria and that the deviance forms of the criteria are bounded below throughout the parameter space. Reasonable starting values determined by the ECME iterations and analytic expressions for the gradients and Hessians help to make (9) and (10) very well controlled optimization problems. The most difficult computational step in the ECME or Newton iterations is the sparse Cholesky decomposition (3).

3 Symbolic analysis

Although the decomposition (3) will be performed many times for different trial values of $\boldsymbol{\theta}$, the structure of $\mathbf{Z}^T \mathbf{Z} + \boldsymbol{\Omega}$ – in particular, the number and positions of the non-zeros in $\mathbf{Z}^T \mathbf{Z} + \boldsymbol{\Omega}$ and in \mathbf{L} – will be the same for each evaluation. Because the LDL package provides one C function to perform the symbolic analysis that determines the number and position of the non-zeros in \mathbf{L} and another C function that determines the numerical values in the decomposition, we can do the symbolic analysis separately.

The number and positions of the nonzeros in \mathbf{L} depends on the positions of the nonzeros in $\mathbf{Z}^T \mathbf{Z}$. Any nonzero position in the lower triangle of $\mathbf{Z}^T \mathbf{Z}$ is potentially nonzero in \mathbf{L} . However other positions in \mathbf{L} can become nonzero during the course of the decomposition. This is called “fill-in”. The

extent of the fill-in can be changed by reordering the components of \mathbf{b} and, correspondingly, the columns of \mathbf{Z} .

Although there are general approaches, such as approximate minimal degree (Davis, 1996) or graph-partitioning algorithms (Karapis, 2003), for determining a fill-minimizing permutation, it is more effective for us to exploit the special structure of $\mathbf{Z}^\top \mathbf{Z}$ in searching for such a permutation.

As mentioned above, when considering the structure of $\mathbf{Z}^\top \mathbf{Z} + \mathbf{\Omega}$ we need only consider the structure for the variance components model because the structure for the general model is obtained from the structure for the variance components model by replacing each nonzero in the (i, j) block of the variance components model by a $q_i \times q_j$ nonzero matrix. Similarly we can derive the structure of the \mathbf{L} matrix for the general model from that of the variance components model provided we restrict our attention to permutations that do not mix levels from different grouping factors. That is, we consider only those fill-reducing permutations that consist of, at most, a permutation of the grouping factors and permutations of the levels within each grouping factor.

In what follows we restrict our attention to the variance components model.

Fill-in is determined by the elimination tree (Liu, 1990) for the symmetric matrix. We can determine the Cholesky decomposition, and hence the elimination tree and the extent of the fill-in, column-wise starting with the first column. We know that there will be “original” nonzeros in \mathbf{L} wherever there are nonzeros in the lower triangle of $\mathbf{Z}^\top \mathbf{Z} + \mathbf{\Omega}$ and, possibly, some additional, “induced” nonzeros. If there are nonzeros, either original or induced, in column j below the diagonal, say, in rows i and k , then a nonzero is induced in the (i, k) position of \mathbf{L} . Consider again the division of $\mathbf{Z}^\top \mathbf{Z} + \mathbf{\Omega}$ and \mathbf{L} into a $k \times k$ array of blocks determined by the grouping factors. For a variance components model, the diagonal blocks are themselves diagonal. Because the $(1, 1)$ block is diagonal the row numbers of any nonzeros below the diagonal must be greater than m_1 . That is, there will not be any induced nonzeros in the first m_1 columns. Because this first block of m_1 rows and columns will not experience any fill-in, we choose the first grouping factor to make m_1 as large as possible.

There will be no need to permute the levels of the first grouping factor. This can be a considerable savings in the effort required to determine a fill-reducing permutation. For the Scottish secondary school example we can leave the first 148 columns in their original order and consider only permutations of the last 19 columns.

We obtain the matrix to determine the permutation of the second and subsequent groups by “projecting” the first m_1 columns onto the last $q - m_1$ columns as described above. If there are only two grouping factors this corresponds to recording the positions (i, k) with $i, k > m_1$ and both (i, j) and (i, k) are non-zero for some $j \leq m_1$. If there are more than two grouping factors then we record all of these plus any of the original nonzeros below the m_1 st row and to the right of the m_1 st column.

In the case of two nested grouping factors there will only be one nonzero element below the diagonal in the first m_1 columns, hence there is no fill-in. For more than two nested grouping factors, any pair of nonzeros occurring in the first m_1 columns must be in rows that correspond to different grouping factors and that off-diagonal in the projected matrix must already be nonzero. That is, nested grouping factors do not generate any fill-in. Not only can the matrix \mathbf{L} be created “in place” (that is, with exactly the same positions of nonzero off-diagonal elements as in $\mathbf{Z}^\top \mathbf{Z}$) but also \mathbf{L}^{-1} has the same pattern of nonzeros. This is unusual. In most cases \mathbf{L}^{-1} has many more nonzeros than does \mathbf{L} .

Notice that a single grouping factor is, trivially, a nested sequence.

3.1 Examples

In Figure 2 we show the projection of the pairwise crosstabulation of *primary* and *second* in the *ScotsSec* data, as shown in Figure 1, onto the 19×19 block for the *second* factor.

4 Acknowledgements

This work was supported by U.S. Army Medical Research and Materiel Command under Contract No. DAMD17-02-C-0119. The views, opinions and/or findings contained in this report are those of the authors and should not be construed as an official Department of the Army position, policy or decision unless so designated by other documentation.

I thank Deepayan Sarkar and Tim Davis for helpful discussions and suggestions and Harold Doran for his suggestion of using sparse matrix techniques for linear mixed-effects models.

The `Matrix` package for R is based on code from the LDL (Davis, 2004), TAUCS (Toledo, 2003), and Metis (Karaplis, 2003) packages.

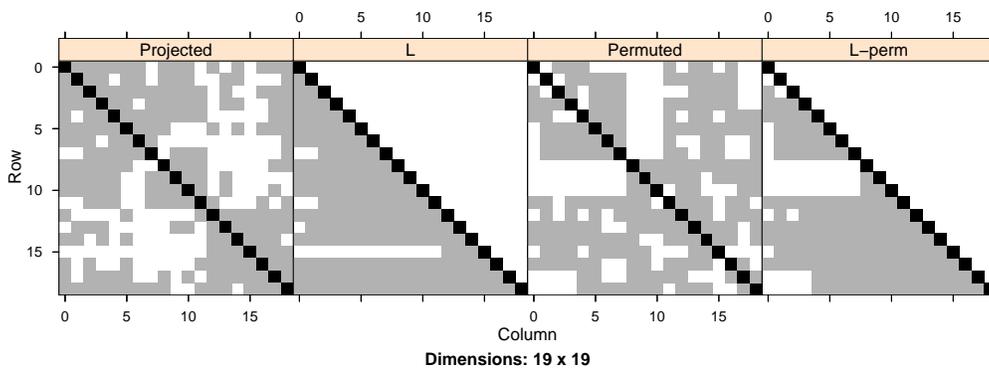


Figure 2: Projection of the pairwise crosstabulation for the Scottish secondary school data onto the 19×19 lower right block for the *second* factor. The “Projected” panel shows the original (black squares) and projected (gray squares) non-zero positions. The “L” panel shows the implicit diagonal (black squares) and the non-zero off-diagonal (gray squares) for this ordering of the levels of the factor. The “Permuted” panel shows the symmetric matrix after permuting the rows and columns according to the fill-reducing permutation determined by the Metis package and the “L-perm” panel shows the \mathbf{L} matrix from this ordering.

References

- Douglas M. Bates and Saikat DebRoy. Linear mixed models and penalized least squares. *J. of Multivariate Analysis*, 2004. to appear. 3, 5, 9
- Tim Davis. An approximate minimal degree ordering algorithm. *SIAM J. Matrix Analysis and Applications*, 17(4):886–905, 1996. 10
- Timothy A. Davis. Algorithm 8xx: A concise sparse Cholesky factorization package. Technical report, Department of Computer and Information Science and Engineering, University of Florida, 2004. 7, 11
- George Karapis. Metis: Family of multilevel partitioning algorithms. <http://www-users.cs.umn.edu/karypis/metis/>, 2003. 10, 11
- J. W. H. Liu. The role of elimination trees in sparse factorization. *SIAM J. Matrix Analysis and Applications*, 11(1):134–172, 1990. 10
- L. Paterson. Socio economic status and educational attainment: a multidimensional and multilevel study. *Evaluation and Research in Education*, 5: 97–121, 1991. 4
- José C. Pinheiro and Douglas M. Bates. *Mixed-Effects Models in S and S-PLUS*. Springer, 2000. ISBN 0-387-98957-9. 2, 9
- Jon Rasbash, William Browne, Harvey Goldstein, Min Yand, Ian Plewis, Michael Healy, Geoff Woodhouse, David Draper, Ian Longford, and Toby Lewis. *A User's Guide to MLwiN*. Institute of Education, University of London, version 2.1d edition, 2002. ISBN 0-85473-612-3. 4
- Sivan Toledo. Taucs: A library of sparse linear solvers. <http://www.tau.ac.il/stoledo/taucs/>, 2003. 11