

Mixed models in R using the lme4 package

Part 6: Theory of linear mixed models

Douglas Bates

University of Wisconsin - Madison
and R Development Core Team
<Douglas.Bates@R-project.org>

Merck, Sharp & Dohme; Rahway, NJ
Sept 24, 2010

1 Definition of linear mixed models

Outline

- 1 Definition of linear mixed models
- 2 The penalized least squares problem

Outline

- 1 Definition of linear mixed models
- 2 The penalized least squares problem
- 3 The sparse Cholesky factor

Outline

- 1 Definition of linear mixed models
- 2 The penalized least squares problem
- 3 The sparse Cholesky factor
- 4 Evaluating the likelihood

Definition of linear mixed models

- As previously stated, we define a linear mixed model in terms of two random variables: the n -dimensional \mathbf{Y} and the q -dimensional \mathbf{B}
- The probability model specifies the conditional distribution

$$(\mathbf{Y}|\mathbf{B} = \mathbf{b}) \sim \mathcal{N}(\mathbf{Z}\mathbf{b} + \mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I}_n)$$

and the unconditional distribution

$$\mathbf{B} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\theta).$$

These distributions depend on the parameters $\boldsymbol{\beta}$, $\boldsymbol{\theta}$ and σ .

- The probability model defines the *likelihood* of the parameters, given the observed data, \mathbf{y} . In theory all we need to know is how to define the likelihood from the data so that we can maximize the likelihood with respect to the parameters. In practice we want to be able to evaluate it quickly and accurately.

Properties of Σ_θ ; generating it

- Because it is a variance-covariance matrix, the $q \times q$ Σ_θ must be symmetric and *positive semi-definite*, which means, in effect, that it has a “square root” — there must be another matrix that, when multiplied by its transpose, gives Σ_θ .
- We never really form Σ ; we always work with the *relative covariance factor*, Λ_θ , defined so that

$$\Sigma_\theta = \sigma^2 \Lambda_\theta \Lambda_\theta^\top$$

where σ^2 is the same variance parameter as in $(\mathbf{y}|\mathbf{B} = \mathbf{b})$.

- We also work with a q -dimensional “spherical” or “unit” random-effects vector, \mathbf{u} , such that

$$\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_q), \quad \mathbf{B} = \Lambda_\theta \mathbf{u} \Rightarrow \text{Var}(\mathbf{B}) = \sigma^2 \Lambda \Lambda^\top = \Sigma.$$

- The linear predictor expression becomes

$$\mathbf{Z}\mathbf{b} + \mathbf{X}\boldsymbol{\beta} = \mathbf{Z}\Lambda_\theta \mathbf{u} + \mathbf{X}\boldsymbol{\beta}$$

The conditional mean $\mu_{\mathbf{u}|\mathbf{y}}$

- Although the probability model is defined from $(\mathbf{y}|\mathbf{u} = \mathbf{u})$, we observe \mathbf{y} , not \mathbf{u} (or \mathbf{b}) so we want to work with the other conditional distribution, $(\mathbf{u}|\mathbf{y} = \mathbf{y})$.
- The joint distribution of \mathbf{y} and \mathbf{u} is Gaussian with density

$$\begin{aligned} f_{\mathbf{y},\mathbf{u}}(\mathbf{y}, \mathbf{u}) &= f_{\mathbf{y}|\mathbf{u}}(\mathbf{y}|\mathbf{u}) f_{\mathbf{u}}(\mathbf{u}) \\ &= \frac{\exp(-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\boldsymbol{\Lambda}_\theta\mathbf{u}\|^2)}{(2\pi\sigma^2)^{n/2}} \frac{\exp(-\frac{1}{2\sigma^2} \|\mathbf{u}\|^2)}{(2\pi\sigma^2)^{q/2}} \\ &= \frac{\exp(-[\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\boldsymbol{\Lambda}_\theta\mathbf{u}\|^2 + \|\mathbf{u}\|^2] / (2\sigma^2))}{(2\pi\sigma^2)^{(n+q)/2}} \end{aligned}$$

- $(\mathbf{u}|\mathbf{y} = \mathbf{y})$ is also Gaussian so its mean is its mode. I.e.

$$\mu_{\mathbf{u}|\mathbf{y}} = \arg \min_{\mathbf{u}} \left[\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\boldsymbol{\Lambda}_\theta\mathbf{u}\|^2 + \|\mathbf{u}\|^2 \right]$$

Minimizing a penalized sum of squared residuals

- An expression like $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\boldsymbol{\Lambda}_\theta\mathbf{u}\|^2 + \|\mathbf{u}\|^2$ is called a *penalized sum of squared residuals* because $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\boldsymbol{\Lambda}_\theta\mathbf{u}\|^2$ is a sum of squared residuals and $\|\mathbf{u}\|^2$ is a penalty on the size of the vector \mathbf{u} .
- Determining $\mu_{\mathbf{u}|\mathbf{y}}$ as the minimizer of this expression is a *penalized least squares* (PLS) problem. In this case it is a *penalized linear least squares problem* that we can solve directly (i.e. without iterating).
- One way to determine the solution is to rephrase it as a linear least squares problem for an extended residual vector

$$\mu_{\mathbf{u}|\mathbf{y}} = \arg \min_{\mathbf{u}} \left\| \begin{bmatrix} \mathbf{y} - \mathbf{X}\boldsymbol{\beta} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{Z}\boldsymbol{\Lambda}_\theta \\ \mathbf{I}_q \end{bmatrix} \mathbf{u} \right\|^2$$

This is sometimes called a *pseudo-data* approach because we create the effect of the penalty term, $\|\mathbf{u}\|^2$, by adding “pseudo-observations” to \mathbf{y} and to the predictor.

Solving the linear PLS problem

- The conditional mean satisfies the equations

$$(\mathbf{Z}\mathbf{\Lambda}_\theta\mathbf{\Lambda}_\theta^\top\mathbf{Z}^\top + \mathbf{I}_q) \boldsymbol{\mu}_{\mathcal{U}|\mathcal{Y}} = \mathbf{\Lambda}_\theta^\top\mathbf{Z}^\top(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}).$$

- This would be interesting but not very important were it not for the fact that we actually can solve that system for $\boldsymbol{\mu}_{\mathcal{U}|\mathcal{Y}}$ even when its dimension, q , is very, very large.
- Because \mathbf{Z} is generated from indicator columns for the grouping factors, it is sparse. $\mathbf{Z}\mathbf{\Lambda}_\theta$ is also very sparse.
- There are sophisticated and efficient ways of calculating a sparse Cholesky factor, which is a sparse, lower-triangular matrix \mathbf{L}_θ that satisfies

$$\mathbf{L}_\theta\mathbf{L}_\theta^\top = \mathbf{\Lambda}_\theta^\top\mathbf{Z}^\top\mathbf{Z}\mathbf{\Lambda}_\theta + \mathbf{I}_q$$

and, from that, solving for $\boldsymbol{\mu}_{\mathcal{U}|\mathcal{Y}}$.

The sparse Choleksy factor, \mathbf{L}_θ

- Because the ability to evaluate the sparse Cholesky factor, \mathbf{L}_θ , is the key to the computational methods in the `lme4` package, we consider this in detail.
- In practice we will evaluate \mathbf{L}_θ for many different values of θ when determining the ML or REML estimates of the parameters.
- As described in Davis (2006), §4.6, the calculation is performed in two steps: in the *symbolic decomposition* we determine the position of the nonzeros in \mathbf{L} from those in $\mathbf{Z}\mathbf{\Lambda}_\theta$ then, in the *numeric decomposition*, we determine the numerical values in those positions. Although the numeric decomposition may be done dozens, perhaps hundreds of times as we iterate on θ , the symbolic decomposition is only done once.

A fill-reducing permutation, P

- In practice it can be important while performing the symbolic decomposition to determine a *fill-reducing permutation*, which is written as a $q \times q$ permutation matrix, P . This matrix is just a re-ordering of the columns of I_q and has an orthogonality property, $PP^T = P^T P = I_q$.
- When P is used, the factor L_θ is defined to be the sparse, lower-triangular matrix that satisfies

$$L_\theta L_\theta^T = P [\Lambda_\theta^T Z_\theta^T Z_\theta \Lambda_\theta + I_q] P^T$$

- In the Matrix package for R, the Cholesky method for a sparse, symmetric matrix (class `dsCMatrix`) performs both the symbolic and numeric decomposition. By default, it determines a fill-reducing permutation, P . The update method for a Cholesky factor (class `CHMfactor`) performs the numeric decomposition only.

Applications to models with simple, scalar random effects

- Recall that, for a model with simple, scalar random-effects terms only, the matrix Σ_{θ} is block-diagonal in k blocks and the i th block is $\sigma_i^2 \mathbf{I}_{n_i}$ where n_i is the number of levels in the i th grouping factor.
- The matrix Λ_{θ} is also block-diagonal with the i th block being $\theta_i \mathbf{I}_{n_i}$, where $\theta_i = \sigma_i / \sigma$.
- Given the grouping factors for the model and a value of θ we produce $\mathbf{Z}\Lambda_{\theta}$ then \mathbf{L} , using Cholesky the first time then update.
- To avoid recalculating we assign
 - `flist` a list of the grouping factors
 - `nlev` number of levels in each factor
 - `Zt` the transpose of the model matrix, \mathbf{Z}
 - `theta` current value of θ
 - `Lambda` current Λ_{θ}
 - `Ut` transpose of $\mathbf{Z}\Lambda_{\theta}$

Cholesky factor for the Penicillin model



```
> flist <- subset(Penicillin, select = c(plate, sample))
> Zt <- do.call(rBind, lapply(flist, as, "sparseMatrix"))
> (nlev <- sapply(flist, function(f) length(levels(factor(f)))))
```

```
plate sample
    24     6
```

```
> theta <- c(1.2, 2.1)
> Lambda <- Diagonal(x = rep.int(theta, nlev))
> Ut <- crossprod(Lambda, Zt)
> str(L <- Cholesky(tcrossprod(Ut), LDL = FALSE, Imult = 1))
```

```
Formal class 'dCHMsimpl' [package "Matrix"] with 10 slots
..@ x      : num [1:189] 3.105 0.812 0.812 0.812 0.812 ...
..@ p      : int [1:31] 0 7 14 21 28 35 42 49 56 63 ...
..@ i      : int [1:189] 0 24 25 26 27 28 29 1 24 25 ...
..@ nz     : int [1:30] 7 7 7 7 7 7 7 7 7 7 ...
..@ nxt    : int [1:32] 1 2 3 4 5 6 7 8 9 10 ...
..@ prv    : int [1:32] 31 0 1 2 3 4 5 6 7 8 ...
..@ colcount: int [1:30] 7 7 7 7 7 7 7 7 7 7 ...
..@ perm   : int [1:30] 23 22 21 20 19 18 17 16 15 14 ...
..@ type   : int [1:4] 2 1 0 1
```

Images of $\Lambda_\theta^\top \mathbf{Z}^\top \mathbf{Z} \Lambda_\theta + \mathbf{I}$ and \mathbf{L}



- Note that there are nonzeros in the lower right of \mathbf{L} in positions that are zero in the lower triangle of $\Lambda_\theta^\top \mathbf{Z}^\top \mathbf{Z} \Lambda_\theta + \mathbf{I}_q$. This is described as “fill-in”.

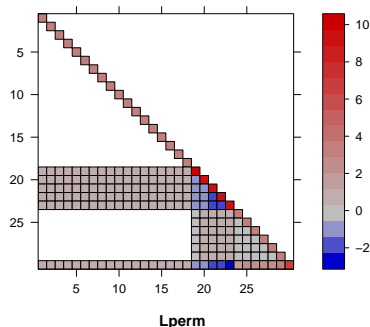

Reversing the order of the factors

- To show the effect of a fill-reducing permutation, we reverse the order of the factors and calculate the Cholesky factor with and without a fill-reducing permutation.
- We evaluate `nnzero` (number of nonzeros) for `L`, from the original factor order, and for `Lnoperm` and `Lperm`, the reversed factor order without and with permutation

```
> Zt <- do.call(rBind, lapply(flist[2:1], as, "sparseMatrix"))
> Lambda <- Diagonal(x = rep.int(theta[2:1], nlev[2:1]))
> Ut <- crossprod(Lambda, Zt)
> Lnoperm <- Cholesky(tcrossprod(Ut), perm = FALSE, LDL = FALSE)
> Lperm <- Cholesky(tcrossprod(Ut), LDL = FALSE, Imult = 1)
> sapply(lapply(list(L, Lnoperm, Lperm), as, "sparseMatrix"),
```

```
[1] 189 450 204
```


Images of the reversed factor decompositions



- Without permutation, we get the worst possible fill-in. With a fill-reducing permutation we get much less but still not as good as the original factor order.
- This is why the permutation is called “fill-reducing”, not “fill-minimizing”. Getting the fill-minimizing permutation in the general case is a very hard problem.



Cholesky factor for the Pastes data

- For the special case of nested grouping factors, such as in the Pastes and classroom data, there is no fill-in, regardless of the permutation.
- A permutation is nevertheless evaluated but it is a “post-ordering” that puts the nonzeros near the diagonal.

```
> Zt <- do.call(rBind, lapply(flist <- subset(Pastes, , c(sample
+                               as, "sparseMatrix")))
> nlev <- sapply(flist, function(f) length(levels(factor(f))))
> theta <- c(0.4, 0.5)
> Lambda <- Diagonal(x = rep.int(theta, nlev))
> Ut <- crossprod(Lambda, Zt)
> L <- Cholesky(tcrossprod(Ut), LDL = FALSE, Imult = 1)
> str(L@perm)
```

```
int [1:40] 2 1 0 30 5 4 3 31 8 7 ...
```

Image of the factor for the Pastes data



- The image for the Cholesky factor from the classroom data model is similar but, with more than 400 rows and columns, the squares for the nonzeros are difficult to see.

The conditional density, $f_{\mathbf{u}|\mathbf{y}}$

- We know the joint density, $f_{\mathbf{y},\mathbf{u}}(\mathbf{y}, \mathbf{u})$, and

$$f_{\mathbf{u}|\mathbf{y}}(\mathbf{u}|\mathbf{y}) = \frac{f_{\mathbf{y},\mathbf{u}}(\mathbf{y}, \mathbf{u})}{\int f_{\mathbf{y},\mathbf{u}}(\mathbf{y}, \mathbf{u}) d\mathbf{u}}$$

so we almost have $f_{\mathbf{u}|\mathbf{y}}$. The trick is evaluating the integral in the denominator, which, it turns out, is exactly the likelihood, $L(\boldsymbol{\theta}, \boldsymbol{\beta}, \sigma^2|\mathbf{y})$, that we want to maximize.

- The Cholesky factor, \mathbf{L}_θ is the key to doing this because

$$\mathbf{P}^\top \mathbf{L}_\theta \mathbf{L}_\theta^\top \mathbf{P} \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}} = \boldsymbol{\Lambda}_\theta^\top \mathbf{Z}^\top (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}).$$

Although the `Matrix` package provides a one-step solve method for this, we write it in stages:

Solve $\mathbf{L} \mathbf{c}_u = \mathbf{P} \boldsymbol{\Lambda}_\theta^\top \mathbf{Z}^\top (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})$ for \mathbf{c}_u .

Solve $\mathbf{L}^\top \mathbf{P} \boldsymbol{\mu} = \mathbf{c}_u$ for $\mathbf{P} \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}}$ and $\boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}}$ as $\mathbf{P}^\top \mathbf{P} \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}}$.

Evaluating the likelihood

- The exponent of $f_{\mathbf{y}, \mathbf{u}}(\mathbf{y}, \mathbf{u})$ can now be written

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\boldsymbol{\Lambda}_\theta \mathbf{u}\|^2 + \|\mathbf{u}\|^2 = r^2(\boldsymbol{\theta}, \boldsymbol{\beta}) + \|\mathbf{L}^\top \mathbf{P}(\mathbf{u} - \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}})\|^2.$$

where $r^2(\boldsymbol{\theta}, \boldsymbol{\beta}) = \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{U}\boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}}\|^2 + \|\boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}}\|^2$. The first term doesn't depend on \mathbf{u} and the second is relatively easy to integrate.

- Use the change of variable $\mathbf{v} = \mathbf{L}^\top \mathbf{P}(\mathbf{u} - \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}})$, with $d\mathbf{v} = \text{abs}(|\mathbf{L}||\mathbf{P}|) d\mathbf{u}$, in

$$\begin{aligned} \int \frac{\exp\left(\frac{-\|\mathbf{L}^\top \mathbf{P}(\mathbf{u} - \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}})\|^2}{2\sigma^2}\right)}{(2\pi\sigma^2)^{q/2}} d\mathbf{u} \\ = \int \frac{\exp\left(\frac{-\|\mathbf{v}\|^2}{2\sigma^2}\right)}{(2\pi\sigma^2)^{q/2}} \frac{d\mathbf{v}}{\text{abs}(|\mathbf{L}||\mathbf{P}|)} = \frac{1}{\text{abs}(|\mathbf{L}||\mathbf{P}|)} = \frac{1}{|\mathbf{L}|} \end{aligned}$$

because $\text{abs}|\mathbf{P}| = 1$ and $\text{abs}|\mathbf{L}|$, which is the product of its diagonal elements, all of which are positive, is positive.

Evaluating the likelihood (cont'd)

- As is often the case, it is easiest to write the log-likelihood. On the deviance scale (negative twice the log-likelihood)

$\ell(\boldsymbol{\theta}, \boldsymbol{\beta}, \sigma | \mathbf{y}) = \log L(\boldsymbol{\theta}, \boldsymbol{\beta}, \sigma | \mathbf{y})$ becomes

$$-2\ell(\boldsymbol{\theta}, \boldsymbol{\beta}, \sigma | \mathbf{y}) = n \log(2\pi\sigma^2) + \frac{r^2(\boldsymbol{\theta}, \boldsymbol{\beta})}{\sigma^2} + \log(|\mathbf{L}_\theta|^2)$$

- We wish to minimize the deviance. Its dependence on σ is straightforward. Given values of the other parameters, we can evaluate the conditional estimate

$$\widehat{\sigma}^2(\boldsymbol{\theta}, \boldsymbol{\beta}) = \frac{r^2(\boldsymbol{\theta}, \boldsymbol{\beta})}{n}$$

producing the *profiled deviance*

$$-2\tilde{\ell}(\boldsymbol{\theta}, \boldsymbol{\beta} | \mathbf{y}) = \log(|\mathbf{L}_\theta|^2) + n \left[1 + \log \left(\frac{2\pi r^2(\boldsymbol{\theta}, \boldsymbol{\beta})}{n} \right) \right]$$

- However, an even greater simplification is possible because the deviance depends on $\boldsymbol{\beta}$ only through $r^2(\boldsymbol{\theta}, \boldsymbol{\beta})$.

Profiling the deviance with respect to β

- Because the deviance depends on β only through $r^2(\theta, \beta)$ we can obtain the conditional estimate, $\hat{\beta}_\theta$, by extending the PLS problem to

$$r_\theta^2 = \min_{\mathbf{u}, \beta} \left[\|\mathbf{y} - \mathbf{X}\beta - \mathbf{Z}\Lambda_\theta\mathbf{u}\|^2 + \|\mathbf{u}\|^2 \right]$$

with the solution satisfying the equations

$$\begin{bmatrix} \Lambda_\theta^\top \mathbf{Z}^\top \mathbf{Z} \Lambda_\theta + \mathbf{I}_q & \mathbf{U}_\theta^\top \mathbf{X} \\ \mathbf{X}^\top \mathbf{Z} \Lambda_\theta & \mathbf{X}^\top \mathbf{X} \end{bmatrix} \begin{bmatrix} \mu_{\mathbf{u}|\mathbf{y}} \\ \hat{\beta}_\theta \end{bmatrix} = \begin{bmatrix} \Lambda_\theta^\top \mathbf{Z}^\top \mathbf{y} \\ \mathbf{X}^\top \mathbf{y} \end{bmatrix}$$

- The profiled deviance, which is a function of θ only, is

$$-2\tilde{\ell}(\theta) = \log(|\mathbf{L}_\theta|^2) + n \left[1 + \log \left(\frac{2\pi r_\theta^2}{n} \right) \right]$$

Solving the extended PLS problem

- For brevity we will no longer show the dependence of matrices and vectors on the parameter θ .
- As before we use the sparse Cholesky decomposition, with L and P satisfying $LL^T = P(\Lambda_\theta^T Z^T Z \Lambda_\theta + I)$ and c_u , the solution to $Lc_u = P\Lambda_\theta^T Z^T y$.
- We extend the decomposition with the $q \times p$ matrix R_{ZX} , the upper triangular $p \times p$ matrix R_X , and the p -vector c_β satisfying

$$\begin{aligned}LR_{ZX} &= P\Lambda_\theta^T Z^T X \\ R_X^T R_X &= X^T X - R_{ZX}^T R_{ZX} \\ R_X^T c_\beta &= X^T y - R_{ZX}^T c_u\end{aligned}$$

so that

$$\begin{bmatrix} P^T L & 0 \\ R_{ZX}^T & R_X^T \end{bmatrix} \begin{bmatrix} L^T P & R_{ZX} \\ 0 & R_X \end{bmatrix} = \begin{bmatrix} \Lambda_\theta^T Z^T Z \Lambda_\theta + I & \Lambda_\theta^T Z^T X \\ X^T Z \Lambda_\theta & X^T X \end{bmatrix}.$$

Solving the extended PLS problem (cont'd)

- Finally we solve

$$\begin{aligned}\mathbf{R}_X \hat{\boldsymbol{\beta}}_{\theta} &= \mathbf{c}_{\beta} \\ \mathbf{L}^T \mathbf{P} \boldsymbol{\mu}_{\mathbf{u}|\mathbf{y}} &= \mathbf{c}_{\mathbf{u}} - \mathbf{R}_{ZX} \hat{\boldsymbol{\beta}}_{\theta}\end{aligned}$$

- The profiled REML criterion also can be expressed simply. The criterion is

$$L_R(\boldsymbol{\theta}, \sigma^2 | \mathbf{y}) = \int L(\boldsymbol{\theta}, \boldsymbol{\beta}, \sigma^2 | \mathbf{y}) d\boldsymbol{\beta}$$

The same change-of-variable technique for evaluating the integral w.r.t. \mathbf{u} as $1/\text{abs}(|\mathbf{L}|)$ produces $1/\text{abs}(|\mathbf{R}_X|)$ here and removes $(2\pi\sigma^2)^{p/2}$ from the denominator. On the deviance scale, the profiled REML criterion is

$$-2\tilde{\ell}_R(\boldsymbol{\theta}) = \log(|\mathbf{L}|^2) + \log(|\mathbf{R}_x|^2) + (n - p) \left[1 + \log\left(\frac{2\pi r_{\theta}^2}{n - p}\right) \right]$$

- These calculations can be expressed in a few lines of **R** code. Assume `rho` contains `y`, `X`, `Zt`, `REML`, `L`, `nlev` and `XtX` ($\mathbf{X}^T \mathbf{X}$).

Code for evaluating the profiled deviance

```
profDev <- function(rho, theta) {
  stopifnot(is.numeric(theta), length(theta)==length(rho$nlev))
  Ut <- crossprod(Diagonal(x=rep.int(theta,rho$nlev)),rho$Zt)
  L <- update(rho$L, Ut, mult = 1)
  cu <- solve(L, solve(L, Ut %*% rho$y, sys = "P"), sys = "L")
  RZX <- solve(L, solve(L, Ut %*% rho$X, sys = "P"), sys = "L")
  RX <- chol(rho$XtX - crossprod(RZX))
  cb <- solve(t(RX),crossprod(rho$X,rho$y)- crossprod(RZX, cu))
  beta <- solve(RX, cb)
  u <- solve(L,solve(L,cu - RZX %*% beta, sys="Lt"), sys="Pt")
  fitted <- as.vector(crossprod(Ut, u) + rho$X %*% beta)
  prss <- sum(c(rho$y - fitted, as.vector(u))^2)
  n <- length(fitted);  p <- ncol(RX)
  if (rho$REML) return(determinant(L)$mod +
                       2 * determinant(RX)$mod +
                       (n-p) * (1+log(2*pi*prss/(n-p))))
  determinant(L)$mod + n * (1 + log(2*pi*prss/n))
}
```

Checking profDev, lmer version of fit

```
> invisible(lmer(mathgain ~ mathkind + minority + ses + (1|class
```

```
npt = 5 , n = 2
```

```
rhobeg = 0.2 , rhoend = 2e-07
```

```
0.020: 11: 11396.1;0.465012 0.298249
```

```
0.0020: 17: 11391.5;0.340940 0.314970
```

```
0.00020: 19: 11391.5;0.336337 0.313884
```

```
2.0e-05: 23: 11391.5;0.336012 0.314470
```

```
2.0e-06: 26: 11391.5;0.335979 0.314536
```

```
2.0e-07: 29: 11391.5;0.335979 0.314536
```

```
At return
```

```
34: 11391.532: 0.335980 0.314536
```

Checking profDev, bobyqa applied to profDev

```
> ls.str(rho <- with(classroom, simplemer(list(classid, school)))
L : Formal class 'dCHMsimpl' [package "Matrix"] with 10 slots
nlev : int [1:2] 312 107
REML : logi FALSE
X : num [1:1190, 1:4] 1 1 1 1 1 1 1 1 1 1 ...
XtX : num [1:4, 1:4] 1190 555324 806 -15.4 555324 ...
y : int [1:1190] 32 109 56 83 53 65 51 66 88 -7 ...
Zt : Formal class 'dgCMatrix' [package "Matrix"] with 6 slots

> invisible(bobyqa(c(0.836158, 0.489669), function(x) profDev(x)))
npt = 5 , n = 2
rhobeg = 0.1672316 , rhoend = 1.672316e-07
  0.017:    5:    11249.5;0.836158 0.489669
  0.0017:   6:    11249.5;0.836158 0.489669
  0.00017: 12:    11249.4;0.839657 0.481151
  1.7e-05: 17:    11249.4;0.840209 0.480882
  1.7e-06: 23:    11249.4;0.840261 0.480806
  1.7e-07: 27:    11249.4;0.840261 0.480803
At return
32:    11249.448: 0.840261 0.480802
```

How lmer works

- Essentially `lmer` takes its arguments and creates a structure like the `rho` environment shown above. The optimization of the profiled deviance or the profiled REML criterion happens within this environment.
- The creation of Λ_θ is somewhat more complex for models with vector-valued random effects but not excessively so.
- Some care is taken to avoid allocating storage for large objects during each function evaluation. Many of the objects created in `profDev` are updated in place within `lmer`.
- Once the optimizer, `bobyqa`, has converged some additional information for the summary is calculated.

Summary

- For a linear mixed model, even one with a huge number of observations and random effects like the model for the grade point scores, evaluation of the ML or REML profiled deviance, given a value of θ , is straightforward. It involves updating Λ_θ , L_θ , R_{ZX} , R_X , calculating the penalized residual sum of squares, r_θ^2 and two determinants of triangular matrices.
- The profiled deviance can be optimized as a function of θ only. The dimension of θ is usually very small. For the grade point scores there are only three components to θ .