

Mixed models in R using the lme4 package

Part 8: Nonlinear 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 Nonlinear mixed models

Outline

- 1 Nonlinear mixed models
- 2 Statistical theory, applications and approximations

Outline

- 1 Nonlinear mixed models
- 2 Statistical theory, applications and approximations
- 3 Model definition

Outline

- 1 Nonlinear mixed models
- 2 Statistical theory, applications and approximations
- 3 Model definition
- 4 Comparing estimation methods


Outline

- 1 Nonlinear mixed models
- 2 Statistical theory, applications and approximations
- 3 Model definition
- 4 Comparing estimation methods
- 5 Fitting NLMMs in lme4 and nlme

Nonlinear mixed models

- Population pharmacokinetic data are often modeled using *nonlinear mixed-effects models* (NLMMs).
- These are *nonlinear* because pharmacokinetic parameters - rate constants, clearance rates, etc. - occur nonlinearly in the model function.
- In statistical terms these are *mixed-effects models* because they involve both *fixed-effects parameters*, applying to the entire population or well-defined subsets of the population, and *random effects* associated with particular experimental or observational units under study.
- Many algorithms for obtaining parameter estimates, usually “something like” the *maximum likelihood estimates* (MLEs), for such models have been proposed and implemented.
- Comparing different algorithms is not easy. Even understanding the definition of the model and the proposed algorithm is not easy.

An example: Theophylline pharmacokinetics



- These are serum concentration profiles for 12 volunteers after injection of an oral dose of Theophylline, as described in Pinheiro and Bates (2000).

Modeling pharmacokinetic data with a nonlinear model

- These are longitudinal repeated measures data.
- For such data the time pattern of an individual's response is determined by pharmacokinetic parameters (e.g. rate constants) that occur nonlinearly in the expression for the expected response.
- The form of the nonlinear model is determined by the pharmacokinetic theory, not derived from the data.

$$d \cdot k_e \cdot k_a \cdot C \frac{e^{-k_e t} - e^{-k_a t}}{k_a - k_e}$$

- These pharmacokinetic parameters vary over the population. We wish to characterize typical values in the population and the extent of the variation.
- Thus, we associate random effects with the parameters, k_a , k_e and C in the nonlinear model.

Statistical theory and applications - why we need both

- For 30 years, I have had the pleasure of being part of the U. of Wisconsin-Madison Statistics Dept. This year we celebrate the 50th anniversary of the founding of our department by George Box (who turned 90 earlier this year).
- George's approach, emphasizing **both** the theory and the applications of statistics, has now become second-nature to me.
- We are familiar with the dangers of practicing theory without knowledge of applications. As George famously said, "All models are wrong; some models are useful." How can you expect to decide if a model is useful unless you use it?
- We should equally be wary of the application of statistical techniques for which we know the "how" but not the "why". Despite the impression we sometimes give in courses, applied statistics is not just a "black box" collection of formulas into which you pour your data, hoping to get back a p-value that is less than 5%. (In the past many people felt that "applied statistics is the use of SAS" but now we know better.)

The evolving role of approximation

- When Don Watts and I wrote a book on nonlinear regression we included a quote from Bertrand Russell, “Paradoxically, all exact science is dominated by the idea of approximation”. In translating statistical theory to applied techniques (computing algorithms) we almost always use some approximations.
- Sometimes the theory is deceptively simple (maximum likelihood estimates are the values of the parameters that maximize the likelihood, given the data) but the devil is in the details (so exactly how do I maximize this likelihood?).
- Decades of work by many talented people have provided us with a rich assortment of computational approximations and other tricks to help us get to the desired answer - or at least close to the desired answer.
- It is important to realize that approximations, like all aspects of computing, have a very short shelf life. Books on theory can be useful for decades; books on computing may be outmoded in a few years.

Failure to revisit assumptions leads to absurdities

- Forty years ago, when I took an intro engineering stats class, we used slide rules or pencil and paper for calculations. Our text took this into account, providing short-cut computational formulas and “rules of thumb” for the use of approximations, plus dozens of pages of tables of probabilities and quantiles.
- Today’s computing resources are unimaginably more sophisticated yet the table of contents of most introductory text hasn’t changed.
- The curriculum still includes using tables to evaluate probabilities, calculating coefficient estimates of a simple linear regression by hand, creating histograms (by hand, probably) to assess a density, approximating a binomial by a Poisson or by a Gaussian for cases not available in the tables, etc.
- Then we make up PDF slides of this content and put the file on a web site for the students to download and follow on their laptops during the lecture. Apparently using the computer to evaluate the probabilities or to fit a model would be cheating - you are supposed to do this by hand.

And what about nonlinear mixed-effects models?

- Defining the statistical model is subtle and all methods proposed for determining parameter estimates use approximations.
- Often the many forms of approximations are presented as different “types” of estimates from which one can pick and choose.
- In 2007-2008 a consortium of pharma companies, the NLMEc, discussed “next generation” simulation and estimation software for population PK/PD modeling. They issued a set of user requirements for such software including, in section 4.4 on estimation

The system will support but not be limited to the following estimation methods: FO, FOI, FOCE, FOCEI, Laplacian, Lindstrom and Bates, MCMC, MCPM, SAEM, Gaussian quadrature, and nonparametric methods.

- Note the emphasis on estimation methods (i.e. algorithms). All of these techniques are supposed to approximate the mle's but that is never mentioned.

Linear and nonlinear mixed-effects models

- Both linear and nonlinear mixed-effects models, are based on the n -dimensional response random variable, \mathcal{Y} , whose value, \mathbf{y} , is observed, and the q -dimensional, unobserved random effects variable, \mathcal{B} .
- In the models we will consider $\mathcal{B} \sim \mathcal{N}(\mathbf{0}, \Sigma_\theta)$. The variance-covariance matrix Σ_θ can be huge but it is completely determined by a small number of *variance-component parameters*, θ .
- The conditional distribution of the response, \mathcal{Y} , is

$$(\mathcal{Y} | \mathcal{B} = \mathbf{b}) \sim \mathcal{N}(\mu_{\mathcal{Y}|\mathcal{B}}, \sigma^2 \mathbf{I}_n)$$

- The conditional mean, $\mu_{\mathcal{Y}|\mathcal{B}}$, depends on \mathbf{b} and on the fixed-effects parameters, β , through a *linear predictor* expression, $\mathbf{Z}\mathbf{b} + \mathbf{X}\beta$.
- For a linear mixed model (LMM), $\mu_{\mathcal{Y}|\mathcal{B}}$ is exactly the linear predictor. For an NLMM the linear predictor determines the parameter values in the nonlinear model function which then determines the mean.

Conditional mode and profiled Laplace approximation for NLMMs

- As previously stated, determining the conditional mode

$$\tilde{\mathbf{u}}_{\theta, \beta} = \arg \min_{\mathbf{u}} \left[\|\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}|\mathbf{u}}\|^2 + \|\mathbf{u}\|^2 \right]$$

in an NLMM is a penalized nonlinear least squares (PNLS) problem.

- It is a nonlinear optimization problem but a comparatively simple one. The penalty term *regularizes* the optimization.
- The Laplace approximation to the profiled deviance (profiled over σ^2) is, as before,

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

where $\mathbf{L}_{\boldsymbol{\theta}}$ is the sparse Cholesky factor evaluated at the conditional mode.

- The motivation for this approximation is that it replaces the conditional distribution, $(\mathbf{U}|\mathbf{Y} = \mathbf{y})$, for parameters $\boldsymbol{\beta}$, $\boldsymbol{\theta}$ and σ , by a multivariate Gaussian approximation. *evaluated at the mode.*

Laplace approximation and adaptive Gauss-Hermite quadrature

- The Laplace approximation

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

is a type of *smoothing objective* consisting of two terms: $n \left[1 + \log \left(\frac{2\pi r^2(\boldsymbol{\theta}, \boldsymbol{\beta})}{n} \right) \right]$, which measures fidelity to the data, and $\log(|\mathbf{L}_{\boldsymbol{\theta}}|^2)$, which measures the complexity of the model.

- For models with a simple structure for the random effects (the matrices $\boldsymbol{\Sigma}_{\boldsymbol{\theta}}$ and $\boldsymbol{\Lambda}_{\boldsymbol{\theta}}$ are block diagonal consisting of a large number of small blocks) a further enhancement is to use *adaptive Gauss-Hermite quadrature*, requiring values of the RSS at several points near $\tilde{\boldsymbol{u}}_{\boldsymbol{\theta}, \boldsymbol{\beta}}$
- Note that the modifier *adaptive*, meaning evaluating at the conditional mode, is important. Gauss-Hermite quadrature without first determining the conditional mode is not a good idea.

Consequences for comparisons of methods

- We should distinguish between an algorithm, which is a sort of a black box, and a criterion, such as maximizing the likelihood (or, equivalently, minimizing the deviance).
- The criterion is based on the statistical model and exists outside of any particular implementation or computing hardware. It is part of the theory, which has a long shelf life.
- A particular approximation, algorithm and implementation has a short shelf life.
- I claim it does not make sense to regard the FO, FOI, ... methods as producing well-defined types of “estimates” in the same sense that maximum likelihood estimates, or maximum *a posteriori* estimates are defined.
- If you use a criterion to define an estimation method then implementations should be compared on the basis of that criterion, not on something like mean squared error.

Specifying the nonlinear model function

- We must specify the nonlinear model function and the linear predictor expression in the model formula for an NLMM. We do this with a *3-part formula* expression.
- At present nonlinear model function must return both the conditional mean and the gradient expression (derivative of the conditional mean with respect to the nonlinear model parameters). It is helpful to use the `deriv()` function to symbolically differentiate the model function.
- Some common models have been encapsulated as “self-starting” nonlinear regression models. For example, the first-order pharmacokinetic model used for the `Theoph` data is available as `SSfo1`. Run `example(SSfo1)` to see what is meant.

Specifying the mixed-effects formula

- The mixed-effects formula for an `nlmex` model has a similar form to that for `lmer` or `glmer` but with new constraints.
- In an NLMM all of the fixed-effects parameters and all of the random effects are with respect to the nonlinear model parameters, which are `lKe`, `lKa` and `lCl` in this case.
- For the purpose of specifying the model, these names are defined as indicator variables.
- In the `lme4` package, the default fixed-effects expression is `0 + lKe + lKa + lCl`, indicating that the intercept term is suppressed and that there is a single fixed effect for each nonlinear model parameter. In `lme4a` this must be specified explicitly (although that may change).
- Random effects must also be specified with respect to the nonlinear model parameters. In the `lme4a` version terms look like `(0 + lKe|Subject)`.

Model building

- Even more so that for GLMMs and NLMMs, it is important to start with simple specification of the random effects. Expecting to estimate many variance-covariance parameters for random effects filtered through a nonlinear model function is unrealistic.
- Use `verbose=TRUE` in all but the simplest cases.
- In `lme4a` the `verbose` option shows two sets of iterations, one over θ only and one over θ and β .
- You can suppress the second optimization, which often does little to lower the deviance, by setting `nAGQ=0`.
- We will begin with independent scalar random effects for each nonlinear model parameter.

Initial fit

```
> Th.start <- c(lKe = -2.5, lKa = 0.5, lCl = -3)
> nm1 <- nlmer(conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~
+             0+lKe+lKa+lCl+(0+lKe|Subject)+(0+lKa|Subject)
+             +(0+lCl|Subject), nAGQ=0, Theoph,
+             start = Th.start, verbose=TRUE)
```

```
npt = 7 , n = 3
```

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

```
0.020: 12:      369.069;0.330364  1.03590 0.628374
0.0020: 26:      355.652;0.0987868  1.12332 0.239366
0.00020: 52:      354.040; 0.00000 0.925125 0.236817
2.0e-05: 60:      354.040; 0.00000 0.926861 0.236828
2.0e-06: 68:      354.040;3.63175e-06 0.927020 0.236817
2.0e-07: 73:      354.040; 0.00000 0.927017 0.236816
```

```
At return
```

```
78:      354.04035: 4.74002e-09 0.927016 0.236816
```

Results of initial fit

Nonlinear mixed model fit by maximum likelihood ['merMod']

Formula: conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~ 0 + lKe + lKa + lCl +

Data: Theoph

AIC	BIC	logLik	deviance
368.0404	388.2200	-177.0202	354.0404

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	lKe	1.125e-17	3.354e-09
Subject	lKa	4.303e-01	6.560e-01
Subject	lCl	2.808e-02	1.676e-01
Residual		5.007e-01	7.076e-01

Number of obs: 132, groups: Subject, 12

Fixed effects:

	Estimate	Std. Error	t value
lKe	-2.45467	0.01406	-174.53
lKa	0.46644	0.19459	2.40
lCl	-3.22717	0.04968	-64.96

Reducing the model

The variance of the random effect for 1K3 is essentially zero. Eliminate it.

```
> nm2 <- nlmer(conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~  
+             0+lKe+lKa+lCl+(0+lKa|Subject)  
+             +(0+lCl|Subject), Theoph, nAGQ=0,  
+             start = Th.start, verbose=TRUE)
```

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

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

```
 0.020: 9: 354.776; 1.00379 0.200012  
 0.0020: 14: 354.150;0.991448 0.248465  
 0.00020: 19: 354.119;0.992540 0.237716  
 2.0e-05: 43: 354.040;0.927023 0.236819  
 2.0e-06: 47: 354.040;0.927023 0.236819  
 2.0e-07: 50: 354.040;0.927016 0.236816
```

```
At return
```

```
55: 354.04035: 0.927016 0.236816
```

Allowing within-subject correlation of random effects

Now allow for possible correlation of these random effects

```
> nm3 <- nlmer(conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~  
+               0+lKe+lKa+lCl+(0+lKa+lCl|Subject),  
+               Theoph, start = Th.start, verbose=TRUE)
```

```
npt = 7 , n = 3
```

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

```
0.020: 11:      354.779; 1.00390 0.00155136 0.200015  
0.0020: 21:      354.154;0.990900 -0.00881029 0.228757  
0.00020: 51:      354.044;0.940578 -0.00287833 0.237085  
2.0e-05: 67:      354.040;0.926993 -0.00181432 0.236803  
2.0e-06: 74:      354.040;0.926990 -0.00176709 0.236768  
2.0e-07: 79:      354.040;0.926989 -0.00176645 0.236768
```

```
At return
```

```
85:      354.03976: 0.926989 -0.00176668 0.236769
```

```
npt = 12 , n = 6
```

```
rhobeg = 0.6454427 , rhoend = 6.454427e-07
```

```
0.065: 13:      366.775;0.926989 -0.00176668 0.645443 -2.45474 0.466472  
0.0065: 29:      358.859;0.959051 -0.119105 0.314781 -2.41046 0.303550 -  
0.00065: 71:      354.665; 1.10706 0.0111648 0.242826 -2.46804 0.401075 -  
6.5e-05: 117:      353.972;0.941246 -4.31282e-05 0.237685 -2.46866 0.484115
```


Model nm3

Nonlinear mixed model fit by maximum likelihood [`'merMod'`]

Formula: `conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~ 0 + lKe + lKa + lCl +`

Data: `Theoph`

AIC	BIC	logLik	deviance
367.9718	388.1514	-176.9859	353.9718

Random effects:

Groups	Name	Variance	Std.Dev.	Corr
Subject	lKa	0.44258	0.6653	
	lCl	0.02821	0.1680	0.001
Residual		0.49953	0.7068	

Number of obs: 132, groups: Subject, 12

Fixed effects:

	Estimate	Std. Error	t value
lKe	-2.46845	0.01403	-175.94
lKa	0.48436	0.19730	2.45
lCl	-3.23034	0.04979	-64.89

Anova comparison

```
> anova(nm2, nm3)
```

Data: Theoph

Models:

```
nm2: conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~ 0 + lKe + lKa + lCl +
```

```
nm2:      (0 + lKa | Subject) + (0 + lCl | Subject)
```

```
nm3: conc ~ SSfol(Dose, Time, lKe, lKa, lCl) ~ 0 + lKe + lKa + lCl +
```

```
nm3:      (0 + lKa + lCl | Subject)
```

	Df	AIC	BIC	logLik	deviance	Chisq	Chi	Df	Pr(>Chisq)
nm2	6	366.04	383.34	-177.02	354.04				
nm3	7	367.97	388.15	-176.99	353.97	0.0686		1	0.7934

- This comparison is not quite fair because `nm2` was fit with `nAGQ=0` and `nm3` allowed both phases of the optimization.
- But we already conclude that the more complex model `nm3` is not a significantly better fit.

Summary

- The theory of NLMMs follows fairly closely the theory of LMMs
- Model specification is more complex because of an additional level of parameters to specify.
- The progress of the iterations should be carefully monitored. Often variance component estimates that are very close to zero but not exactly zero are provided.
- There is a tendency to incorporate too much complexity in such models. As Einstein said, “A scientific theory should be as simple as possible, but no simpler.”