

JOSÉ C. PINHEIRO

Model building using covariates in nonlinear mixed-effects models

Journal de la société française de statistique, tome 143, n° 1-2 (2002), p. 79-101

http://www.numdam.org/item?id=JSFS_2002__143_1-2_79_0

© Société française de statistique, 2002, tous droits réservés.

L'accès aux archives de la revue « Journal de la société française de statistique » (<http://publications-sfds.math.cnrs.fr/index.php/J-SFdS>) implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/conditions>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques

<http://www.numdam.org/>

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

José C. PINHEIRO *

ABSTRACT

Nonlinear mixed-effects (NLME) models are useful in describing nonlinear relationships between a response variable and parameters and covariates in data that are grouped according to a classification factor. Examples of such *grouped data* include longitudinal and repeated measures data, which frequently arise in many areas of application, such as pharmacokinetics, biostatistics, and economics. By associating common random effects to observations sharing the same level of a classification factor, NLME models flexibly represent the covariance structure induced by the grouping of the data. The random effects in an NLME model account for differences in the parameter values among groups. In many applications, additional covariates are collected along with the response and model covariates and can be used to explain, at least partially, the variation among groups. Including covariates in an NLME model to explain inter-group variation generally leads to simplifications in the random effects model and to a better understanding of the model producing the response. This paper describes a model-building strategy for identifying and including covariates in an NLME model, using the capabilities available in the NLME library for S-PLUS and R. The use of the proposed methodology and the capabilities in the NLME software are illustrated with real-life examples from ecophysiology and pharmacokinetics.

RÉSUMÉ

Les modèles non-linéaires à effets mixtes (NLME) sont très utiles pour décrire des relations non-linéaires entre une variable réponse et des paramètres ou covariables lorsque les données observées sont groupées selon un niveau de facteur. De telles données englobent entre autres les données longitudinales, les mesures répétées et concernent beaucoup de domaines applicatifs telles la phamaco-cinétique, la biostatistique, l'économie. En associant aux observations des effets aléatoires communs à chaque niveau au facteur de groupement, la structure de covariance associée est aisément représentée par les modèles NLME : les effets aléatoires caractérisent les différences des valeurs des paramètres entre groupes. Dans beaucoup d'applications, des covariables additionnelles sont mesurées parallèlement à la variable réponse et elles peuvent être utilisées pour expliquer des variations même partielles entre les groupes. Inclure ces covariables dans un modèle NLME pour expliquer des variations inter-groupes amène généralement à une simplification de la modélisation des effets aléatoires et à une meilleure compréhension du modèle explicatif. Ce papier décrit une stratégie de modélisation pour l'identification et la prise en compte de covariables utilisant les capacités de la librairie NLME disponible sous S-PLUS et sous R. La mise en œuvre de cette méthodologie est illustrée sur des exemples concrets en écophysologie et en pharmo-cinétique.

* Department of Biostatistics, Novartis Pharmaceuticals, One Health Plaza, East Hanover, New Jersey, U.S.A.; E-mail : jose.pinhoiro@pharma.novartis.com

1. Introduction

Nonlinear mixed-effects (NLME) models are useful in describing nonlinear relationships between a response variable and parameters and covariates in data that are grouped according to a classification factor. Examples of such *grouped data* include longitudinal and repeated measures data, which frequently arise in many areas of application, such as pharmacokinetics, biostatistics, and economics. NLME models assume that the form of the intra-group model relating the response variable to covariates is common to all groups, but some of the parameters that define the model are allowed to vary with group, through the use of random effects. By associating common random effects to observations in the same group, NLME models flexibly represent the covariance structure induced by the grouping of the data.

The random effects in an NLME model account for individual deviations in the parameters among groups. In many applications, these deviations can be at least partially explained by differences in covariate values among groups. Including covariates in an NLME model to explain inter-group variation often leads to simplifications in the random effects model and generally allows a better understanding of the mechanism producing the response.

During the process of adding covariates to an NLME model, several model building questions need to be addressed, such as :

- Among the candidate covariates, which are potentially useful in explaining the random effects variation ?
- Which random effects have their variation best explained by covariates ?
- How should the potentially useful covariates be tested for inclusion in the model ?
- Should random effects be included in, or eliminated from, the model, after covariates have been included ?

This paper describes a model-building strategy for addressing these questions in the context of NLME models, using the capabilities available in the **NLME** library for **S-PLUS** and **R** (Pinheiro and Bates, 2000). The use of the proposed methodology and the capabilities in the **NLME** software are illustrated with real-life examples from ecophysiology and pharmacokinetics.

The rest of the paper is organized as follows. The nonlinear mixed-effects model is described in Section 2. Section 3 introduces examples from ecophysiology (CO_2 uptake) and pharmacokinetics (clinical study of Quinidine) which are later used to illustrate the model building methodology, described in Section 4. Conclusions and suggestions for further research are presented in Section 5.

2. A nonlinear mixed-effects model

Nonlinear mixed-effects (NLME) models are mixed-effects models in which the response function is nonlinear in at least some of the underlying parameters. Several different nonlinear mixed-effects models have been proposed in the

literature (Sheiner and Beal, 1980; Mallet *et al.*, 1988; Lindstrom and Bates, 1990; Vonesh and Carter, 1992; Davidian and Gallant, 1992; Wakefield *et al.*, 1994). We adopt here the NLME model proposed by Lindstrom and Bates (1990), which can be viewed as a hierarchical model that generalizes both the linear mixed-effects model of Laird and Ware (1982) and the usual nonlinear regression model for independent data (Bates and Watts, 1988). In the first stage, the j th observation on the i th group is described as

$$y_{ij} = f(\phi_{ij}, \mathbf{x}_{ij}) + \epsilon_{ij}, \quad i = 1, \dots, M, \quad j = 1, \dots, n_i \quad (1)$$

where f is a nonlinear function of a group-specific vector of parameters ϕ_{ij} and the vector of covariates \mathbf{x}_{ij} , the ϵ_{ij} are normally distributed, independent within-group error terms, M is the total number of groups, and n_i is the number of observations in the i th group. In the second stage the group-specific parameters are modeled as

$$\phi_{ij} = \mathbf{A}_{ij}\boldsymbol{\beta} + \mathbf{B}_{ij}\mathbf{b}_i \quad (2)$$

where $\boldsymbol{\beta}$ represents the *fixed effects*; \mathbf{b}_i the *random effects* (varying with i but not with j), which are assumed to be independently distributed as $\mathcal{N}(\mathbf{0}, \boldsymbol{\Psi})$. \mathbf{A}_{ij} and \mathbf{B}_{ij} are design matrices for the fixed and random effects respectively, which may depend on the values of some covariates at the j th observation. It is further assumed that the \mathbf{b}_i are independent of the ϵ_{ij} .

Inclusion of covariates in the NLME model is done primarily through the \mathbf{A}_{ij} matrices. As described in Section 4, candidate covariates for inclusion in the model are screened using plots of the estimated random effects against available covariates. The most promising covariates are incorporated in the model by adding corresponding columns to the \mathbf{A}_{ij} , with resulting estimated fixed effects being tested for significance. This model building strategy is described and illustrated in detail in Section 4.

Different methods have been proposed to estimate the parameters in the NLME model (Ramos and Pantula, 1995; Davidian and Giltinan, 1995; Vonesh and Chinchilli, 1997); we concentrate here on methods based on the likelihood function. Because the random effects \mathbf{b}_i are unobserved quantities, maximum likelihood estimation in NLME models is based on the marginal density of the responses \mathbf{y} , which is calculated as

$$\ell(\boldsymbol{\beta}, \boldsymbol{\Psi}, \sigma | \mathbf{y}) = \int p(\mathbf{y} | \mathbf{b}, \boldsymbol{\beta}, \sigma) p(\mathbf{b} | \boldsymbol{\Psi}) d\mathbf{b} \quad (3)$$

Because the model function f in (1) can be nonlinear in the random effects, the integral in (3) generally does not have a closed-form expression. To make the numerical optimization of the likelihood function a tractable problem, different approximations to (3) have been proposed. Some of these methods consist of taking a first-order Taylor expansion of the model function f around the expected value of the random effects (*i.e.*, $\mathbf{0}$) (Sheiner and Beal, 1980; Vonesh and Carter, 1992), or around the conditional modes of the random effects ($\hat{\mathbf{b}}_i$) (Lindstrom and Bates, 1990). We adopt here the approximation suggested

by Lindstrom and Bates (1990), which is implemented via an alternating algorithm comprising a linear mixed-effects (LME) step and a penalized nonlinear least squares (PNLS) step. This is the algorithm used in the `nlme` function of the **NLME** library. Inferences on the model parameters, including hypothesis testing, are based on asymptotic results for the linear mixed-effects log-likelihood used in the LME step of the alternating algorithm (Pinheiro and Bates, 2000).

3. Examples

In this section we introduce the two examples that will be used to illustrate the model building methodology for incorporating covariates in an NLME model described in Section 4. We also describe how to fit the associated NLME models using the tools in the **NLME** library.

3.1. Carbon dioxide uptake

Data from a study of the cold tolerance of a C_4 grass species, *Echinochloa crus-galli* is reported in Potvin *et al.* (1990). A total of 12 four-week-old plants, 6 from Québec and 6 from Mississippi, were divided into two groups – control plants that were kept at 26°C and chilled plants that were subject to 14 h of chilling at 7°C. After 10 h of recovery at 20°C, carbon dioxide (CO_2) uptake rates (in $\mu\text{mol}/\text{m}^2\text{s}$) were measured for each plant at seven increasing concentrations of ambient CO_2 ($\mu\text{L}/\text{L}$). The objective of the experiment was to evaluate the effect of plant type and chilling treatment on the CO_2 uptake. The CO_2 data, displayed in Figure 1, are available in the **NLME** library as the `groupedData` object `CO2` (Pinheiro and Bates, 2000).

```
> CO2
```

```
Grouped Data : uptake ~ conc | Plant
```

	Plant	Type	Treatment	conc	uptake
1	Qn1	Quebec	nonchilled	95	16.0
2	Qn1	Quebec	nonchilled	175	30.4
. . .					
83	Mc3	Mississippi	chilled	675	18.9
84	Mc3	Mississippi	chilled	1000	19.9

It is clear from Figure 1 that the CO_2 uptake rates of Québec plants are greater than those of Mississippi plants and that chilling the plants reduces their CO_2 uptake rates, with this decrease being more pronounced in Mississippi plants than in Québec plants.

An asymptotic regression model with an offset is used in Potvin *et al.* (1990) to represent the expected CO_2 uptake rate $U(c)$ as a function of the ambient CO_2 concentration c

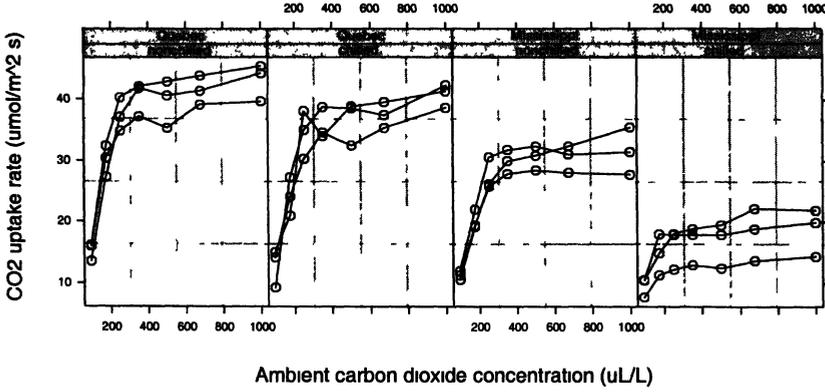


FIG 1 – CO₂ uptake versus ambient CO₂ by chilling treatment and plant type for *Echinochloa crus-galli* plants

$$U(c) = \phi_1 \{1 - \exp[-\exp(\phi_2)(c - \phi_3)]\}, \quad (4)$$

where ϕ_1 is the asymptotic CO₂ uptake rate, ϕ_2 is the logarithm of the rate constant, and ϕ_3 is the maximum ambient concentration of CO₂ at which there is no uptake. The logarithm of the rate constant is used to enforce the positivity of the estimated rate constant, while keeping the optimization problem unconstrained.

We initially consider an NLME version of the CO₂ uptake model (4) with all parameters as mixed effects and no treatment covariates. The corresponding model for the CO₂ uptake u_{ij} of plant i at ambient CO₂ concentration c_{ij} is

$$u_{ij} = \phi_{1i} \{1 - \exp[-\exp(\phi_{2i})(c_{ij} - \phi_{3i})]\} + \epsilon_{ij},$$

$$\phi_i = \begin{bmatrix} \phi_{1i} \\ \phi_{2i} \\ \phi_{3i} \end{bmatrix} = \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix} + \begin{bmatrix} b_{1i} \\ b_{2i} \\ b_{3i} \end{bmatrix} = \boldsymbol{\beta} + \mathbf{b}_i, \quad (5)$$

$$\mathbf{b}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Psi}), \quad \epsilon_{ij} \sim \mathcal{N}(0, \sigma^2),$$

where ϕ_{1i} , ϕ_{2i} , and ϕ_{3i} have the same interpretation as in model (4), but are now allowed to vary with plant. The fixed effects, $\boldsymbol{\beta}$, represent the population average of the individual parameters, ϕ_i , and the random effects, \mathbf{b}_i , represent the deviations of the ϕ_i from their population average. In the notation of the general NLME model (2), $\mathbf{A}_{ij} = \mathbf{B}_{ij} = \mathbf{I}$.

The NLME library includes an `SSasymptOff` function with a self-starting implementation of model (4), which is used to automatically generate starting estimates for the parameters in the model. We use it to fit the NLME model (5).

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

```

> fm1CO2 <- nlme(uptake ~ SSasymOff(conc, Asym, lrc, c0),
+   data = CO2, fixed = Asym + lrc + c0 ~ 1)
> fm1CO2
. . .
  Fixed : list(Asym ~ 1, lrc ~ 1, c0 ~ 1)

  Asym      lrc      c0
32.474 -4.6362 43.543

Random effects :
  Formula : list(Asym ~ 1, lrc ~ 1, c0 ~ 1)
  Level : Plant
  Structure : General positive-definite

           StdDev   Corr
  Asym  9.50999  Asym      lrc
  lrc   0.12828 -0.160
  c0  10.40519  0.999 -0.139
  Residual 1.76641
. . .

```

with $\phi_1 = \text{Asym}$, $\phi_2 = \text{lrc}$, and $\phi_3 = \text{c0}$.

The very high correlation between Asym and c0 gives indication that the fit corresponds to a numerically unstable solution. In general, from practical experience, this occurs when the random effects model is over-parameterized and the optimization algorithm attempts to converge to a lower dimension $\hat{\Psi}$ but, because of the parameterization used in the NLME software, ends up converging to a boundary solution. In this particular case, the estimated $\hat{\Psi}$ suggests that a single random effect could be used to represent both the Asym and the c0 random effects. The scatter plot matrix of the estimated random effects (not shown) confirms that Asym and c0 are in almost perfect linear alignment. A possible model for the random effects in this case would express the c0 random effect as a multiple δ of the Asym random effect. This model, however, leads to numerical difficulties in the estimation of δ using the alternating algorithm implemented in the NLME software, resulting in singular gradient matrices.

An alternative approach is to consider simpler models with one of the highly correlated random effects removed. Dropping the Asym random effect results in a highly significant decrease in the log-likelihood of the fitted model (p-value < 0.0001), while dropping the c0 random effect gives an equivalent fit of the data (p-value of 0.41 for the likelihood ratio test comparing the two nested models). Therefore, we chose the model with Asym and lrc random effects only, which can be fit in NLME with

```

> fm2CO2 <- update(fm1CO2.nlme, random = Asym + lrc ~ 1)
> fm2CO2
. . .
Random effects :
  Formula : list(Asym ~ 1, lrc ~ 1)
  Level : Plant
  Structure : General positive-definite

          StdDev   Corr
  Asym  9.65939   Asym
  lrc   0.19951  -0.777
Residual 1.80792
. . .

```

This is the model that will be used as a starting point for covariate model building in Section 4.

Formally, the hypotheses being tested when deciding whether a random effect could be dropped from model (5) are

$$H_0 : \Psi = \begin{pmatrix} \Psi_0 & \mathbf{0}' \\ \mathbf{0} & 0 \end{pmatrix}, \Psi_0 \in M_2^+ \quad \text{vs.} \quad H_1 : \Psi \in M_3^+, \quad (6)$$

where M_q^+ represents the cone of symmetric positive-definite matrices of order q . This corresponds to a *non-classical* likelihood ratio test scenario, as the null hypothesis H_0 lies in the boundary of the parameter space. As discussed in Stram and Lee (1994) (see also the correction to the original paper in Stram and Lee (1995)), using the results of Self and Liang (1987), classical likelihood ratio tests under this scenario tend to be conservative. That is, the p -value calculated from the χ_1^2 distribution is greater than it should be, under the correct asymptotic distribution. As described in Stram and Lee (1994), the asymptotic distribution of the likelihood ratio test under the boundary conditions (6) is given by a mixture of χ^2 distributions with degrees-of-freedom ≤ 1 . The conclusions regarding the choice of which random effects to keep in the NLME model for the CO₂ data are unchanged when the corrected likelihood ratio test distribution is used (the corresponding p -value for eliminating the c_0 random effect is 0.20).

3.2. Clinical study of quinidine

Routine clinical data on patients receiving the drug quinidine as a treatment for cardiac arrhythmia (atrial fibrillation of ventricular arrhythmias) were reported in Verme *et al.* (1992). All patients were receiving oral quinidine doses. At irregular intervals blood samples were drawn and serum concentrations of quinidine were determined. These data, shown in Figure 2, are available as the `groupedData` object `Quinidine` in the `NLME` library.

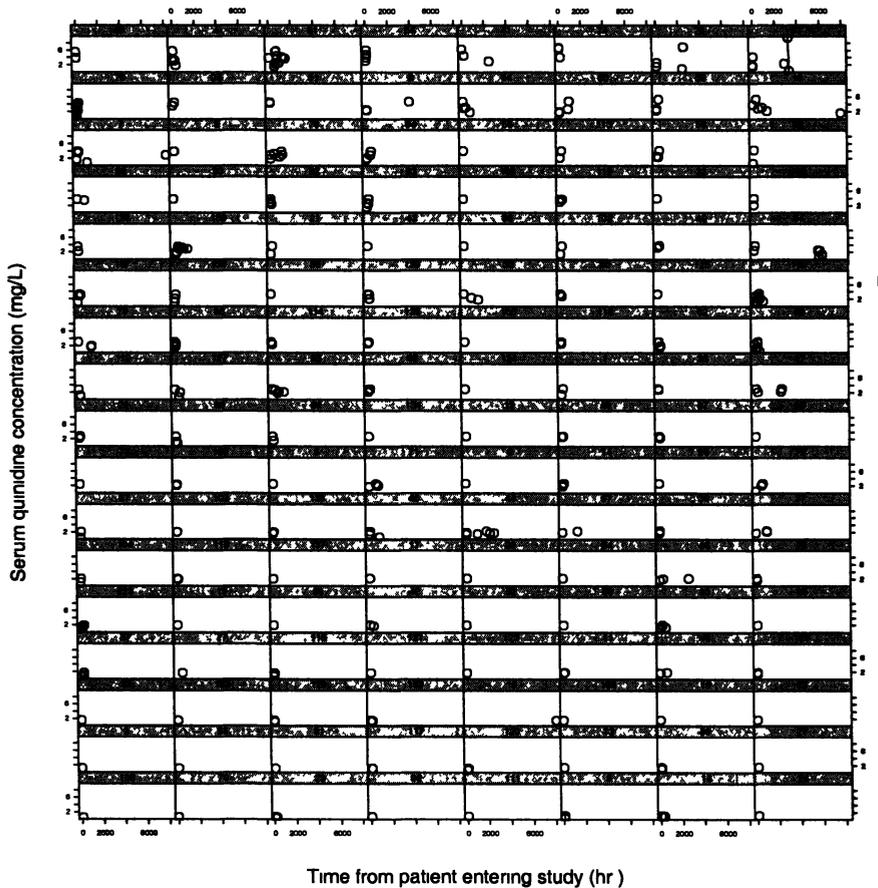


FIG 2 – Serum concentrations of quinidine in 136 hospitalized patients under varying dosage regimens versus time since entering the study.

A total of 361 quinidine concentration measurements were made on 136 hospitalized patients under varying dosage regimens. The times since hospitalization at which the quinidine concentrations were measured varied between 0.13 and 8095.5 hours. Most patients have only a few concentration measurements – 34% have only one and 80% have three or fewer. Only 5% of the patients have seven or more observations. Additional demographic and physiological data were collected for each subject. The additional available covariates are described in Table 1. Some of these covariates, such as age, body weight, and creatinine clearance, were “time-varying.” That is, their value for a particular patient could change during the course of the study. Others, such as race, remained constant. One of the main objectives of the study was to investigate relationships between the individual pharmacokinetic parameters and the covariates. Statistical analyses of these data using different modeling approaches

are given in Davidian and Gallant (1992), Davidian and Giltinan (1995) and Wakefield (1996).

TABLE 1. – Demographic and physiological covariates in the quinidine data.

Age (yr)	42–92
Glycoprotein concentration (mg/100 dL)	0.39–3.16
Body weight (kg)	41–119
Congestive heart failure	no/mild, moderate, severe
Creatinine clearance (ml/min)	< 50, ≥ 50
Ethanol abuse	none, current, former
Height (in.)	60–79
Race	Caucasian, Latin, Black
Smoking status	no, yes

The model that has been suggested for the quinidine data is the one-compartment open model with first-order absorption. This model can be defined recursively as follows. Suppose that, at time t , the patient receives a dose d_t and prior to that time the last dose was given at time t' . The expected concentration in the serum compartment, C_t , and in the absorption compartment, Ca_t , are given by

$$C_t = C_{t'} \exp[-K(t-t')] + \frac{Ca_{t'} k_a}{k_a - K} \{ \exp[-K(t-t')] - \exp[-k_a(t-t')] \}$$

$$Ca_t = Ca_{t'} \exp[-k_a(t-t')] + \frac{d_t}{V} \tag{7}$$

where V is the apparent volume of distribution, k_a is the absorption rate constant, and K is the elimination rate constant.

When a patient receives the same dose d at regular time intervals Δ , model (7) converges to the steady state model

$$C_t = \frac{d k_a}{V(k_a - K)} \left[\frac{1}{1 - \exp(-K\Delta)} - \frac{1}{1 - \exp(-k_a\Delta)} \right]$$

$$Ca_t = \frac{d}{V[1 - \exp(-k_a\Delta)]} \tag{8}$$

Patients considered to be in steady state conditions have concentrations modeled as above.

Finally, for a between-dosages time t , the model for the expected concentration C_t , given that the last dose was received at time t' , is identical to (7).

Using the fact that the elimination rate constant K is equal to the ratio between the clearance (Cl) and the volume in distribution (V), we can reparametrize models (7) and (8) in terms of V , k_a , and Cl . To ensure that the estimates of V , k_a , and Cl are positive, we can rewrite models (7) and (8) in terms of $lV = \log(V)$, $lk_a = \log(k_a)$ and $lCl = \log(Cl)$.

The initial conditions for the recursive model (7) and (8) are $C_0 = 0$ and $Ca_0 = d_0/V$, with d_0 denoting the first dose received by the patient. It has been assumed throughout the model's definition that the bioavailability of the drug, *i.e.*, the percentage of the administered dose that reaches the measurement compartment, is equal to one.

The function `quinModel` in the **NLME** library implements the recursive models (7) and (8), parameterized in terms of lV , lKa and lCl . This is not a self-starting model, so initial values for the fixed effects need to be provided when calling `nlme`. We used values reported in the literature as starting estimates for the fixed effects.

Preliminary analyses of the data, without using any covariates to explain inter-subject variation, indicate that only lCl and lV need random effects to account for their variability in the patient population, and that the corresponding random effects can be assumed to be independent. The corresponding models for the fixed and random effects are

$$lCl_i = \beta_1 + b_{1i}, \quad lV_i = \beta_2 + b_{2i}, \quad lKa_i = \beta_3, \\ \mathbf{b}_i = \begin{bmatrix} b_{1i} \\ b_{2i} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \psi_1 & 0 \\ 0 & \psi_2 \end{bmatrix} \right), \quad (9)$$

which is fitted in **S-PLUS** or **R** with

```
> fm1Quin <-
+ nlme(conc ~ quinModel(Subject, time, conc, dose, interval,
+                       lV, lKa, lCl),
+       data = Quinidine, fixed = lV + lKa + lCl ~ 1,
+       random = pdDiag(lV + lCl ~ 1), groups = ~ Subject,
+       start = list(fixed = c(5, -0.3, 2)),
+       na.action = na.include, naPattern = ~ !is.na(conc) )
> fm1Quin
. . .
Fixed : lV + lKa + lCl ~ 1
      lV      lKa      lCl
5.3796 -0.20535 2.4687

Random effects :
Formula : list(lV ~ 1, lCl ~ 1)
Level : Subject
Structure : Diagonal

          lV      lCl Residual
StdDev : 0.31173 0.32276 0.73871
. . .
```

This will be the starting model used for covariate model building of the quinidine data in the next section. Note that the 1V and 1C1 are assumed to be independent, so only estimates for the corresponding standard deviations are presented in the **S** output.

4. Incorporating covariates in the NLME model

The general model building approach to be used for incorporating covariates in the model consists in starting with an NLME model with no covariates to explain inter-group variation, and using plots of the estimated random effects \hat{b}_i versus the candidate covariates to identify *interesting* patterns. Because the random effects accommodate individual departures from the population mean, plotting the estimated random effects against the candidate covariates provides useful information for the model-building process. A systematic pattern in a given random effect with respect to a covariate would indicate that the covariate should be included in the model.

If no interesting patterns are observed, the current model is kept, else, the covariate-coefficient pair with the most promising pattern is tested for inclusion in the model. The procedure is then applied sequentially, until no further interesting patterns are found.

The number of additional parameters to be estimated tends to grow considerably with the inclusion of covariates and, possibly, their associated random effects in the model. If the number of covariate-coefficient combinations is large, we suggest using a *forward stepwise* approach in which covariate-coefficient pairs are included in the model one at a time and the potential importance of the remaining covariates is graphically assessed at each step. The significance of the fixed-effects associated with a covariate included in the model is assessed using the Wald-type tests, based on asymptotic results for LME models (Pinheiro and Bates, 2000).

The inclusion of new random effects in the model when a covariate is added is rare, but should be investigated. The more common situation is that random effects can be eliminated from the model after covariates are included to account for inter-group variation. In both cases we proceed by comparing nested models using either likelihood ratio tests, or information criterion statistics (AIC and BIC). We illustrate the use of the proposed model-building strategy with the CO₂ uptake and Quinidine examples described in Section 3.

4.1. CO₂ uptake

The primary question of interest for the CO₂ data is the effect of plant type and chilling treatment on the individual model parameters ϕ_i . To plot the estimated random effects against the covariates, we first need to extract the \hat{b}_i from the fitted model and combine them with the covariates. The `ranef` function in **NLME** accomplishes that.

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

```
> fm2CO2.RE <- ranef(fm2CO2, augFrame = T)
> fm2CO2.RE
      Asym      lrc      Type  Treatment  conc  uptake
Qn1  6.17160  0.0483563  Quebec nonchilled  435  33.229
Qn2 10.53264 -0.1728531  Quebec nonchilled  435  35.157
Qn3 12.21810 -0.0579930  Quebec nonchilled  435  37.614
. . .
```

The `augFrame` argument, when `TRUE`, indicates that summary values for all the variables in the data frame should be returned along with the estimated random effects. When a covariate is constant within a group, such as `Treatment` and `Type` in the `CO2` data, its unique values per group are returned. Otherwise, if the covariate varies within the group and is numeric, such as `conc` and `uptake` in `CO2`, the group means are returned; if it is a categorical variable, the most frequent values (modes) within each group are used.

The `plot` method for objects produced by the `ranef` function is the most useful tool for identifying relationships between individual parameters and covariates. The simple call below produces the plot in Figure 3.

```
> plot(fm2CO2.RE, form = ~ Type * Treatment)
```

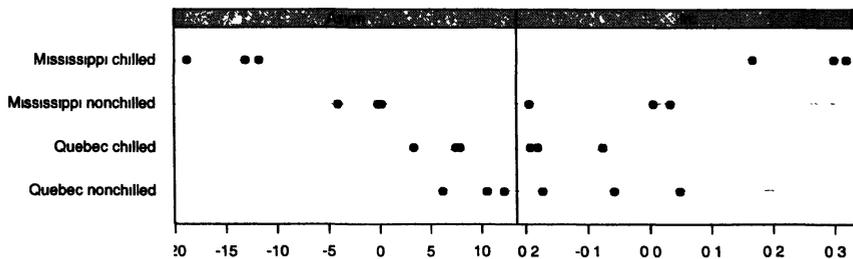


FIG 3. – Dotplots of estimated random effects corresponding to `fm2CO2` versus all combinations of plant type and chilling treatment.

Figure 3 shows a strong relationship between the estimated random effects and the covariates – `Asym` decreases when the plants are chilled and is higher among Québec plants than Mississippi plants, with the increase in `Asym` from chilled to nonchilled plants being larger among Mississippi plants than among Québec plants, suggesting an interaction between `Type` and `Treatment`. There is also evidence of a `Type` by `Treatment` interaction on the log-rate – `lrc` increases with chilling for Mississippi plants and decreases with chilling for Québec plants. We include both covariates in the model to explain the `Asym` and `lrc` plot-to-plot variation. This leads to the following changes in ϕ_{1i} and ϕ_{2i} in model (5).

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

$$\begin{aligned} \phi_{1i} &= \beta_1 + \gamma_{11}x_{1i} + \gamma_{21}x_{2i} + \gamma_{31}x_{1i}x_{2i} + b_{1i}, \\ \phi_{2i} &= \beta_2 + \gamma_{12}x_{1i} + \gamma_{22}x_{2i} + \gamma_{32}x_{1i}x_{2i} + b_{2i}, \end{aligned} \tag{10}$$

$$x_{1i} = \begin{cases} -1, & \text{Type}_i = \text{Québec}, \\ 1, & \text{Type}_i = \text{Mississippi}, \end{cases} \quad x_{2i} = \begin{cases} -1, & \text{Treatment}_i = \text{nonchilled}, \\ 1, & \text{Treatment}_i = \text{chilled}, \end{cases}$$

where β_1 represents the average asymptotic uptake rate, γ_{11} and γ_{12} represent the plant type main effects, γ_{21} and γ_{22} represent the chilling treatment main effects, and γ_{31} and γ_{32} represent the plant type–chilling treatment interactions.

We update the fitted model, incorporating the covariates as in (10) through the fixed argument.

```
> fm3CO2 <- update(fm2CO2,
+ fixed = list(Asym + lrc ~ Type * Treatment, c0 ~ 1),
+ start = c(32.4, 0, 0, 0, -4.6, 0, 0, 0, 49.3))
```

Because the fixed-effects model has changed, new starting values must be provided. We use the previous estimates for β_1 , β_2 and β_3 and set the initial values for the new fixed effects to zero. Wald-type tests are then used to assess the significance of the individual fixed effects.

```
> summary( fm3CO2.nlme )
. . .
Random effects :
. . .
                StdDev   Corr
Asym.(Intercept) 2.3496 Asym.(
  lrc.(Intercept) 0.0796 -0.92
                Residual 1.7920

Fixed effects : list(Asym + lrc ~ Type * Treatment, c0 ~ 1)
                Value Std.Error DF t-value p-value
Asym.(Intercept) 32.342    0.7849 64  41.208 <.0001
  Asym.Type      -7.990    0.7785 64 -10.264 <.0001
  Asym.Treatment -4.210    0.7781 64  -5.410 <.0001
Asym.Type :Treatment -2.725    0.7781 64  -3.502 0.0008
  lrc.(Intercept) -4.509    0.0809 64 -55.743 <.0001
    lrc.Type       0.133    0.0552 64   2.417 0.0185
    lrc.Treatment  0.100    0.0551 64   1.812 0.0747
  lrc.Type :Treatment 0.185    0.0554 64   3.345 0.0014
                c0    50.511    4.3646 64  11.573 <.0001
. . .
```

All fixed effects introduced in the model to explain the variability in *Asym* are highly significant and the *Type* by *Treatment* interaction for *lrc* is also highly significant, confirming the previous conclusions from Figure 3.

The inclusion of the experimental factors in the model resulted in a reduction in the estimated standard deviation for the *Asym* random effects from 9.66 to 2.35 and in the estimated standard deviation for the *lrc* random effects from 0.20 to 0.08, indicating that a substantial part of the plot-to-plot variation in these two coefficients is explained by differences in plant type and chilling treatment.

Figure 4 displays the plots of the estimated random effects corresponding to *fm3CO2* versus the experimental factors. As expected, no systematic patterns can be found.

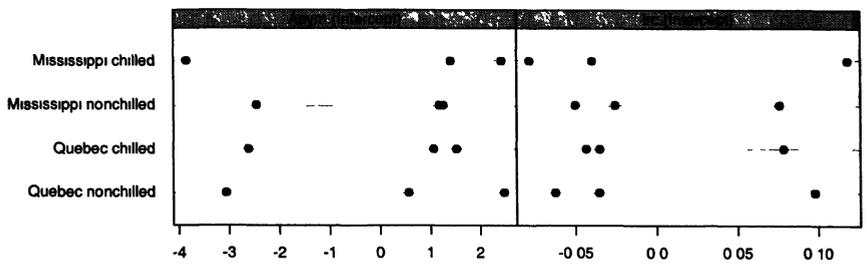


FIG. 4. – Dotplots of estimated random effects corresponding to *fm3CO2* versus all combinations of plant type and chilling treatment.

After covariates have been introduced in the model to account for inter-group variation, a natural question to ask is which random effects, if any, are still needed. The ratio between a random-effect's standard deviation and the absolute value of the corresponding fixed effect gives an idea of the relative inter-group variability for the coefficient and is useful in deciding which random effects should be tested for deletion from the model. For the *fm3CO2* fit these ratios are 7.3% for *Asym*. (Intercept) and 1.8% for *lrc*. (Intercept), suggesting that the latter should be tested for exclusion first. The likelihood ratio test comparing *fm3CO2* to a model with a single *Asym* random effect gives a non-significant p-value of 0.27, indicating that the *lrc* random effect could be dropped after the inclusion of the covariates. A subsequent test comparing the simpler NLME model to a model without any random effects resulted in a highly significant p-value, indicating that the *Asym* random effect needs to be kept in the model to account for unexplained plot-to-plot variation. The same remarks about the conservativeness of the classical likelihood ratio test for evaluating the removal of random effects from the model, included at the end of Section 3.1, also apply here. As before, the conclusions are unchanged when the corrected distribution of the test is used.

4.2. Quinidine

To investigate which covariates may account for patient-to-patient variation in the pharmacokinetic parameters, we first extract the estimated random effects, augmented with summary values for the available covariates (the modal value is used for time-varying categorical variables and the mean for time-varying numeric variables).

```
> fm1Quin.RE <- ranef( fm1Quin, aug = T )
> fm1Quin.nlmeRE[1 :3,]
```

	lV	lCl	time	conc	dose	interval	Age	Height
109	0.0005212	-0.0028369	61.58	0.50000	NA	NA	70	67
70	0.0362214	0.3227614	1.50	0.60000	NA	NA	68	69
23	-0.0254211	0.4402551	91.14	0.56667	NA	NA	75	72

	Weight	Race	Smoke	Ethanol	Heart	Creatinine	glyco
109	58	Caucasian	no	none	No/Mild	>= 50	0.46000
70	75	Caucasian	no	former	No/Mild	>= 50	1.15000
23	108	Caucasian	yes	none	No/Mild	>= 50	0.83667

The `dotplot` displays used to visualize the relationships between the estimated random effects and the covariates in the CO₂ example do not scale up well when there are a large number of groups, or a large number of covariates in the data, as in the quinidine study. Also, they cannot be used with numeric covariates, like `Weight` and `Age`. The `plot` method for objects returned by the `ranef` function actually allows a more flexible type of display for these situations. Relationships between estimated random effects and categorical variables are displayed using boxplots, while scatter plots are used for displaying the relationships between the estimated random effects and numeric covariates. Specifying a two-sided formula in the `form` argument, with the random effect on left-hand side and the desired covariates, separated by the `+` operator, on the right-hand side, indicates to the `plot` method that the more general display should be used. For example, to plot the estimated `lCl` random effects against the available covariates we use

```
> plot( fm1Quin.RE, form = lCl ~ Age + Smoke + Ethanol +
+      Weight + Race + Height + glyco + Creatinine + Heart)
```

The resulting plot, shown in Figure 5, indicates that clearance decreases with glycoprotein concentration and age, and increases with creatinine clearance and weight. There is also evidence that clearance decreases with severity of congestive heart failure and is smaller in Blacks than in both Caucasians and Latins. The glycoprotein concentration is clearly the most important covariate for explaining the `lCl` inter-individual variation. A straight line seems adequate to model the observed relationship.

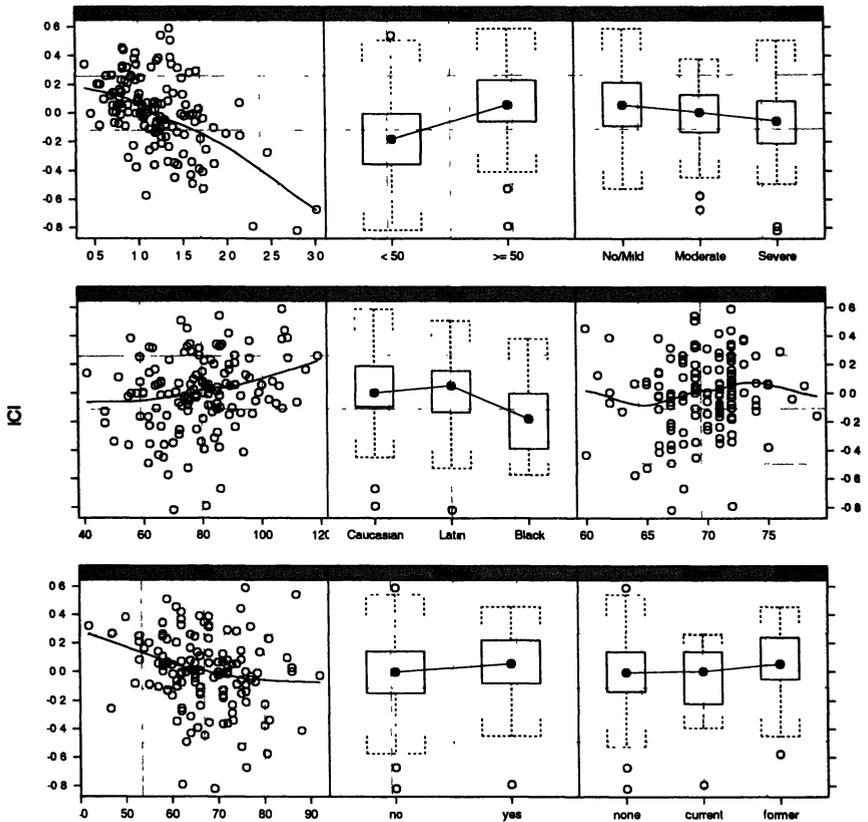


FIG. 5. – Estimated log-clearance random effects from model `fmlQuin` versus demographic and physiological covariates in the quinidine data. A loess smoother is included in the scatter plots of the continuous covariates to aid in visualizing possible trends.

Because of the number of observations per individual varies considerably in the quinidine study, the random effects are, as a result, estimated with different precisions. An alternative plot, taking the different precisions into account, would use *standardized* estimates $\hat{b}_i^* = \hat{b}_i / sd(\hat{b}_i)$, where $sd(\hat{b}_i)$ represents the estimated standard deviation of \hat{b}_i . This type of plot is not currently implemented in NLME, so we will stick here to plots based on the unnormalized \hat{b}_i .

Figure 6 presents the plots of the estimated 1V random effects versus the available covariates. None of the covariates seems helpful in explaining the variability of this random effect and we do not pursue the modeling of its variability any further.

Following the forward stepwise approach mentioned earlier, initially only the glycoprotein concentration is included in the model to explain the 1C1

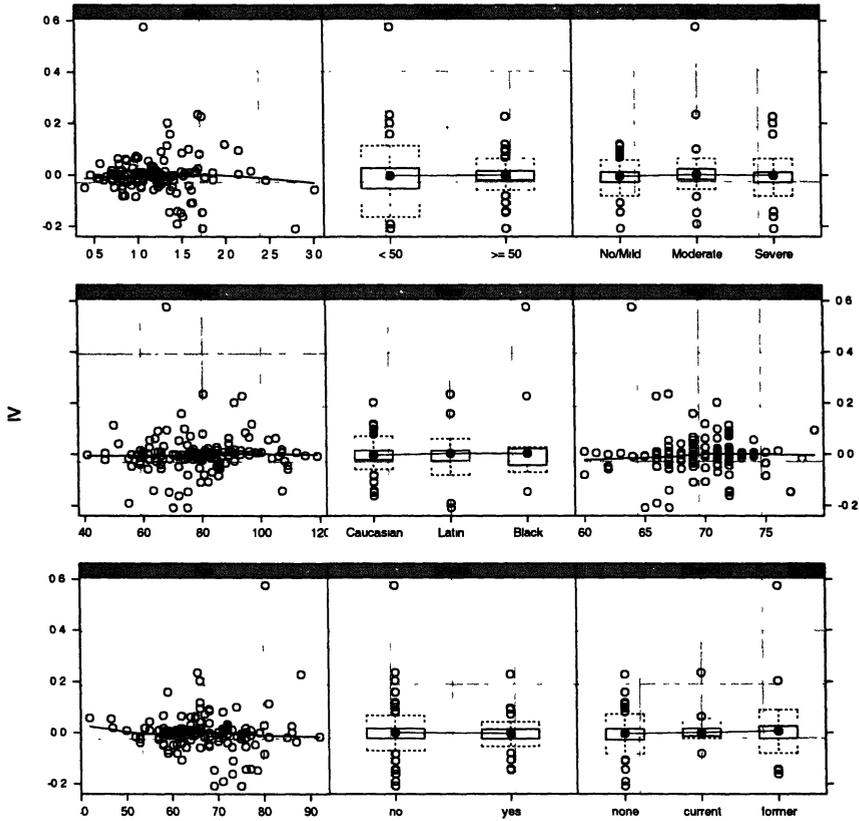


FIG. 6. – Estimated log-volume random effects from model fm1Quin versus demographic and physiological covariates in the quinidine data.

inter-subject variation according to a linear model. The modified version of model (9) is

$$lCl_{i,j} = (\beta_1 + b_{1i}) + \beta_4 glyco_{i,j}.$$

Because the glycoprotein concentration may change with time on the same patient, the random effects for 1Cl need to be indexed by both patient i and time j . The corresponding model is fitted with

```
> fm2Quin <- update(fm1Quin, fixed = list(lV + lKa ~ 1, lCl ~ glyco),
+ start = c(5.38, -0.21, 2.47, 0))
> summary(fm2Quin)
. . .
Random effects :
. . .
          lV  lCl.(Intercept)  Residual
StdDev : 0.26766             0.27037  0.63745
```

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

Fixed effects : list(lV + lKa ~ 1, lCl ~ glyco)

	Value	Std.Error	DF	t-value	p-value
lV	5.3085	0.10244	222	51.818	<.0001
lKa	-0.6662	0.30251	222	-2.202	0.0287
lCl.(Intercept)	3.1067	0.06473	222	47.997	<.0001
lCl.glyco	-0.4914	0.04263	222	-11.527	<.0001

As expected, the estimated lCl.glyco fixed effect is very significant, indicating that the glycoprotein concentration should be kept in the model.

To search for further covariates to include in the model, we investigate the plots of the estimated lCl.(Intercept) random effects from the fm2Quin fit versus the covariates, presented in Figure 7.

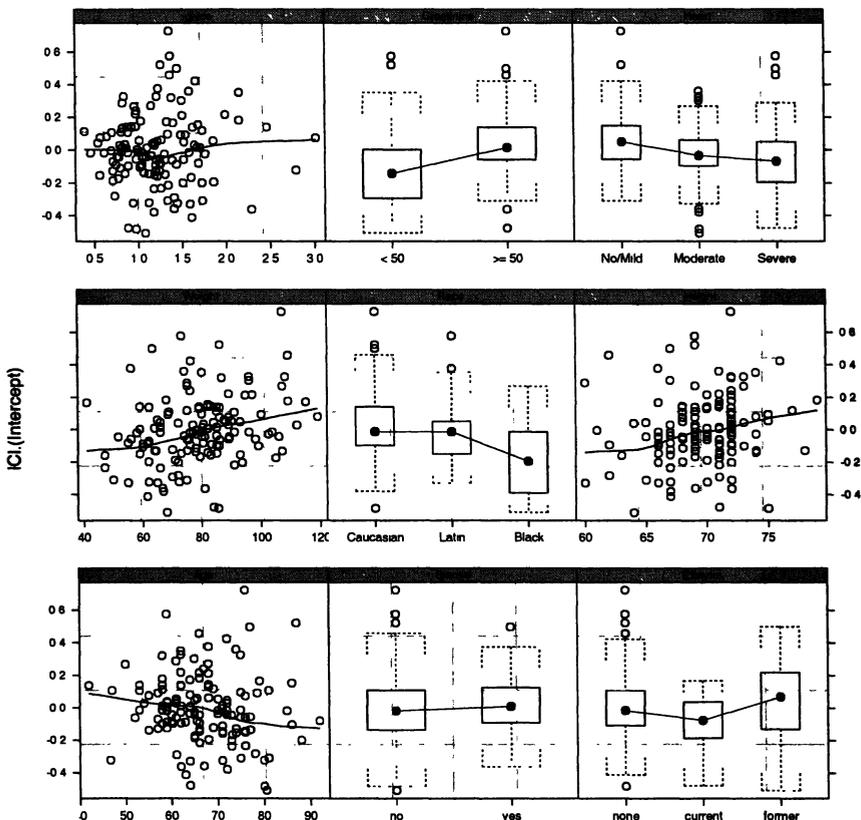


FIG. 7. - Estimated log-clearance random effects from model fm2Quin versus demographic and physiological covariates in the quinidine data.

As expected, there is no relation between the estimated random effects and glyco. The plots indicate that the estimated lCl.(Intercept) random

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

effects increase with creatinine clearance, weight, and height, decrease with age and severity of congestive heart failure, and are smaller in Blacks than in Caucasians and Latins. The most relevant variable appears to be the creatinine clearance, which is included in the model as a binary variable taking value 0 when creatinine is < 50 and 1 when creatinine is ≥ 50.

```
> options(contrasts = c("contr.treatment", "contr.poly"))
> fm3Quin <- update( fm2Quin,
+   fixed = list(lV + lKa ~ 1, lCl ~ glyco + Creatinine),
+   start = c(5.31, -0.67, 3.11, -0.49, 0))
> summary( fm3Quin.nlme )
. . .
Fixed effects : list(lV + lKa ~ 1, lCl ~ glyco + Creatinine)
              Value Std.Error  DF t-value p-value
      lV      5.2900   0.10631 221  49.761  <.0001
      lKa     -0.7462   0.29635 221  -2.518  0.0125
lCl.(Intercept)  2.9229   0.07221 221  40.477  <.0001
      lCl.glyco -0.4632   0.04117 221 -11.251  <.0001
      lCl.Creatinine  0.2125   0.04491 221   4.733  <.0001
. . .
```

The fixed effect corresponding to lCl.Creatinine is very significant, as expected.

The final model produced by this stepwise model-building approach includes extra terms for indicator functions of the events *race = Black*, *heart = No/Mild*, and *ethanol = former* to explain the clearance variation. The corresponding model for the log-clearance is expressed as

$$lCl_{i,j} = (\beta_1 + b_{1i}) + \beta_4 glyco_{i,j} + \beta_5 Creatinine_{i,j} + \beta_6 I(\text{Race} = \text{Black})_i + \beta_7 I(\text{Heart} = \text{No/Mild})_i + \beta_8 I(\text{Ethanol} = \text{former})_i$$

and the resulting fit gives

```
> summary(fm4Quin.nlme)
. . .
Random effects :
. . .
              lV lCl.(Intercept) Residual
StdDev : 0.29855                0.222  0.62565
```

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

Fixed effects : list(lV+lKa ~ 1, lCl ~ glyco+Creatinine+I(Race == "Black")+
I(Heart == "No/Mild")+I(Ethanol == "former"))

	Value	Std.Error	DF	t-value	p-value
lV	5.3036	0.10494	218	50.540	<.0001
lKa	-0.7288	0.29299	218	-2.487	0.0136
lCl.(Intercept)	2.8585	0.07459	218	38.322	<.0001
lCl.glyco	-0.4617	0.03987	218	-11.582	<.0001
lCl.Creatinine	0.2086	0.04366	218	4.779	<.0001
lCl.I(Race == "Black")	-0.2697	0.09409	218	-2.867	0.0046
lCl.I(Heart == "No/Mild")	0.1473	0.05397	218	2.729	0.0069
lCl.I(Ethanol == "former")	0.1528	0.06531	218	2.340	0.0202
. . .					

All lCl coefficients are significant at the usual levels.

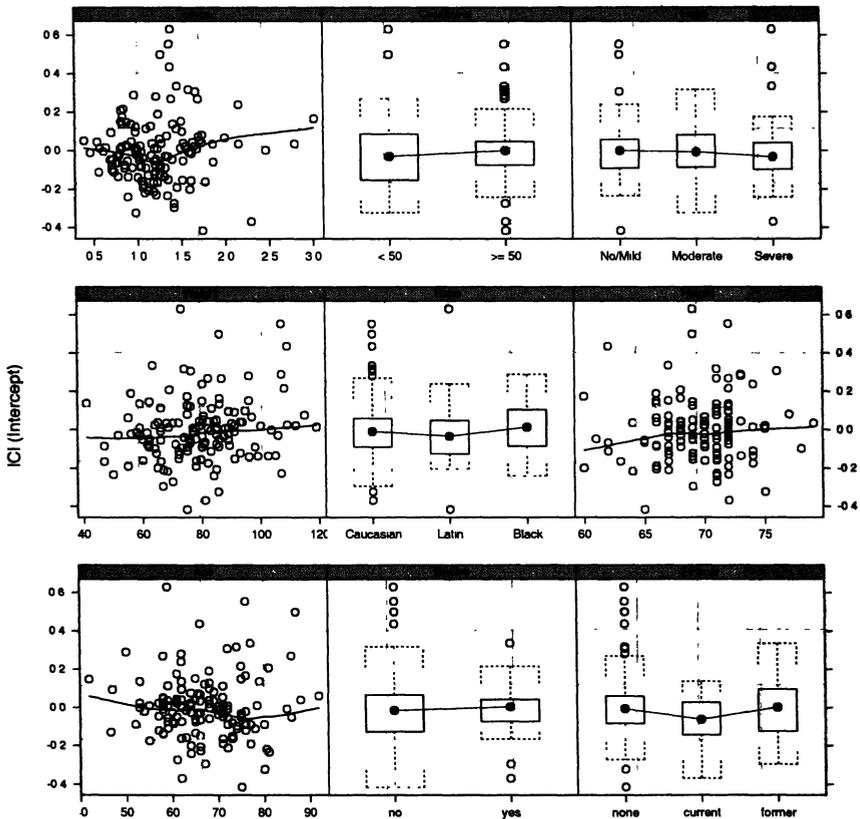


FIG. 8. - Estimated log-clearance random effects from model fm4Quin versus demographic and physiological covariates in the quinidine data.

The plots of the estimated random effects for model fm4Quin versus covariates, displayed in Figure 8, do not suggest any further covariates to be included in the model. Even though there was a reduction in the estimated standard deviations for the 1C1 and 1V random effects, neither could be removed from the model.

As reported in previous analyses of the quinidine data (Davidian and Giltinan, 1995), there is evidence that the variability in the concentration measurements increases with the quinidine concentration. This can be investigated using the tools available in the NLME library, but we will not pursue such analysis here.

5. Conclusions

Incorporating covariates in an NLME model allows inter-group variation to be explained, at least partially, via fixed effects. Besides enhancing the understanding of the model and reducing the dependency on random effects, this allows better predictions to be made from the model, especially for groups not previously observed. Often times, the inclusion of covariates in the model leads to the removal of random effects, or a simplification of the random effects model.

The NLME library in S-PLUS and R provides a flexible and rich environment for model building with covariates in NLME models, including powerful graphical tools for identifying interesting covariates and useful functions for summarizing covariates and combining them with estimated random effects.

The general model building strategy proposed in this paper is to, when feasible, start with a model with all coefficients having both fixed and random effects. This may lead to over-parameterization in the random effects model and the resulting $\hat{\Psi}$ should be investigated for possible indications of numerical instability (the most common problem being very high correlations, in absolute value, between some of the random effects). When this happens, a more parsimonious random effects model should be identified via testing of nested models (information criteria such as the AIC and the BIC could also be used at this stage), in order to achieve greater numerical stability. This model should then be used as a starting point to identify potentially useful covariates to explain inter-group variation. Plots of the estimated random effects versus available covariates are particularly useful at this step. A forward stepwise approach is proposed, with the most promising covariate identified in the plots being included in the model at each step, then tested for significance, and plots of the new estimated random effects versus covariates investigated. The procedure continues until no further covariates are available, or promising covariates can be identified in the plots. Finally, after all useful covariates have been identified and included in the model, the remaining random effects should be re-evaluated for their need in the model. Comparisons of coefficients of variation and confidence intervals for standard deviations corresponding to the random effects are useful at this step.

With genetic data being increasingly collected in a variety of areas where NLME models are used, such as pharmacogenetic data in clinic trials, thousands of potentially useful covariates are becoming available for model building. This creates challenging statistical and computational problems, such as the need for automatic methods of selection that account for multiple testing, algorithms that scale-up with the number of fixed effects, etc. Further research is needed before the methods and software described here can be used in this type of applications.

Acknowledgments

I would like to thank Robert Faivre, Didier Concordet, and an anonymous referee for many helpful comments and suggestions that improved the quality and readability of this article.

REFERENCES

- BATES D.M., WATTS D.G. (1988), *Nonlinear Regression Analysis and Its Applications*, Wiley, New York.
- DAVIDIAN M., GALLANT A.R. (1992), Smooth Nonparametric Maximum Likelihood Estimation for Population Pharmacokinetics, with Application to Quinidine, *Journal of Pharmacokinetics and Biopharmaceutics*, 20, 529-556.
- DAVIDIAN M., GILTINAN D.M. (1995), *Nonlinear Models for Repeated Measurement Data*, Chapman & Hall, London.
- LAIRD N.M., WARE J.H. (1982), Random-effects Models for Longitudinal Data, *Biometrics*, 38, 963-974.
- LINDSTROM M.J., BATES D.M. (1990), Nonlinear Mixed-Effects Models for Repeated Measures Data, *Biometrics*, 46, 673-687.
- MALLET A., MENTRE F., STEIMER J.-L., LOKIEK F. (1988), Nonparametric maximum likelihood estimation for population pharmacokinetics, with applications to Cyclosporine, *Journal of Pharmacokinetics and Biopharmaceutics*, 16, 311-327.
- PINHEIRO J.C., BATES D.M. (2000), *Mixed-Effects Models in S and S-PLUS*, Springer-Verlag, New York.
- POTVIN C., LECHOWICZ M.J., TARDIF S. (1990), The statistical analysis of ecophysiological response curves obtained from experiments involving repeated measures, *Ecology*, 71, 1389-1400.
- RAMOS R.Q., PANTULA S.G. (1995), Estimation of Nonlinear Random Coefficient Models, *Statistics & Probability Letters*, 24, 49-56.
- SELF S.G., LIANG K.Y. (1987), Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions, *Journal of the American Statistical Association*, 82, 605-610.
- SHEINER L.B., BEAL S.L. (1980), Evaluation of Methods for Estimating Population Pharmacokinetic Parameters. I. Michaelis-Menten Model: Routine Clinical Pharmacokinetic Data, *Journal of Pharmacokinetics and Biopharmaceutics*, 8, (6), 553-571.
- STRAM D.O., LEE J.W. (1994), Variance components testing in the longitudinal mixed-effects models, *Biometrics*, 50, 1171-1177.

MODEL BUILDING USING COVARIATES IN NONLINEAR MIXED-EFFECTS MODELS

- STRAM D.O., LEE J.W. (1995), Correction to "Variance components testing in the longitudinal mixed-effects models", *Biometrics*, 51, 1196.
- VERME C.N., LUDDEN T.M., CLEMENTI W.A., HARRIS S.C. (1992), Pharmacokinetics of quinidine in male patients : A population analysis, *Clinical Pharmacokinetics*, 22, 468-480.
- VONESH E.F., CARTER R.L. (1992), Mixed-Effects Nonlinear Regression for Unbalanced Repeated Measures, *Biometrics*, 48, 1-18.
- VONESH E.F., CHINCHILLI V.M. (1997), Linear and Nonlinear Models for the Analysis of Repeated Measures, *Marcel Dekker*, New York.
- WAKEFIELD J. (1996), The Bayesian Analysis of Population Pharmacokinetic Models, *Journal of the American Statistical Association*, 91, 62-75.
- WAKEFIELD J., SMITH A.F.M., RACINE-POON A., GELFAND A.E. (1994), Bayesian Analysis of Linear and Nonlinear Population Models Using the Gibbs Sampler, *Applied Statistics*, 43, 201-221.