Nonlinear random effects mixture models: Maximum likelihood estimation via the EM algorithm

Xiaoning Wang[a], Alan Schumitzky[b], David Z. D'Argenio[a],[*]

[a]*Department of Biomedical Engineering, University of Southern California, 1042 Downey Way, DRB 140, Los Angeles, CA 90089, USA*
[b]*Department of Mathematics, University of Southern California, Los Angeles, CA 90089, USA*

Abstract

Nonlinear random effects models with finite mixture structures are used to identify polymorphism in pharmacokinetic/pharmacodynamic (PK/PD) phenotypes. An EM algorithm for maximum likelihood estimation approach is developed and uses sampling-based methods to implement the expectation step, that results in an analytically tractable maximization step. A benefit of the approach is that no model linearization is performed and the estimation precision can be arbitrarily controlled by the sampling process. A detailed simulation study illustrates the feasibility of the estimation approach and evaluates its performance. Applications of the proposed nonlinear random effects mixture model approach to other population PK/PD problems will be of interest for future investigation.
© 2007 Elsevier B.V. All rights reserved.

Keywords: Finite mixture models; Mixed effects models; Pharmacokinetics/pharmacodynamics

1. Introduction

There is substantial variability in the way individuals respond to medications, in both treatment efficacy and toxicity. The sources of a drug's underlying pharmacokinetic (PK) and pharmacodynamic (PD) variability can include demographic factors (such as age, sex, weight), physiological status (such as renal, liver, cardiovascular function), disease states, genetic differences, interactions with other drugs and environmental factors. In their seminal work, Sheiner et al. (1972) proposed a parametric nonlinear mixed-effects modeling framework for quantifying both within and between subject variability in a drug's PKs, and developed an approximate maximum likelihood solution to the problem. Since the introduction by Beal and Sheiner (1979) of the general purpose software package NONMEM implementing this approach, other approximate maximum likelihood algorithms have been introduced to solve the nonlinear random and mixed effects modeling problem (see Davidian and Giltinan, 1995 for an extensive review). An exact maximum likelihood (i.e., no linearization) solution to the parametric population modeling problem based on the EM algorithm has also been proposed by Schumitzky (1995) and fully developed and implemented by Walker (1996). The population modeling framework has had a significant impact on how PK (and PD) variability is quantified and studied during drug

---

development, and on the identification of important covariates associated with a drug's inter-individual kinetic/dynamic variability.

While population models incorporating measured covariates have proven to be useful in drug development, it is recognized that genetic polymorphisms in drug metabolism and in the molecular targets of drug therapy, for example, can also have a significance influence on the efficacy and toxicity of medications (Evans and Relling, 1999). There is, therefore, a need for population modeling approaches that can extract and model important subpopulations using PK/PD data collected in the course of drug development trials and other clinical studies, in order to help identify otherwise unknown genetic determinants of observed PK/PD phenotypes. The nonparametric maximum likelihood approach for nonlinear random effects modeling developed by Mallet (1986), as well as the nonparametric Bayesian approaches of Wakefield and Walker (1997) and Rosner and Muller (1997), and the smoothed nonparametric maximum likelihood method of Davidian and Gallant (1993) all address this important problem. In this paper we propose a parametric approach using finite mixture models to identify subpopulations with distinct PK/PD properties.

An EM algorithm for exact maximum likelihood estimation of nonlinear random effects finite mixture models is introduced, extending the previous work of Schumitzky (1995) and Walker (1996). The EM algorithm has been used extensively for linear mixture model applications (see McLachlan and Peel, 2000 for a review). The algorithm for nonlinear mixture models presented below has an analytically tractable M step, and uses sampling-based methods to implement the E step. Section 2 of this paper describes the finite mixture model within a nonlinear random effects modeling framework. Section 3 gives the EM algorithm for the maximum likelihood estimation of the model. Section 4 addresses individual subject classification, while an error analysis is presented in Section 5. A detailed simulation study of a PK model is presented in Section 6. Section 7 contains a discussion.

2. Nonlinear random effects finite mixture models

A two-stage nonlinear random effects model that incorporates a finite mixture model is given by

$$\mathbf{Y}_i|\boldsymbol{\theta}_i, \boldsymbol{\beta} \sim N(\mathbf{h}_i(\boldsymbol{\theta}_i), \mathbf{G}_i(\boldsymbol{\theta}_i, \boldsymbol{\beta})), \quad i = 1, \ldots, n \tag{1}$$

and

$$\boldsymbol{\theta}_i, \ldots, \boldsymbol{\theta}_n \sim_{i.i.d} \sum_{k=1}^{K} w_k N(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \tag{2}$$

where $i = 1, \ldots, n$ indexes the individuals and $k = 1, \ldots, K$ indexes the mixing components.

At the first stage represented by (1), $\mathbf{Y}_i = (y_{1i}, \ldots, y_{m_i i})^{\mathrm{T}}$ is the observation vector for the ith individual ($\mathbf{Y}_i \in R^{m_i}$); $\mathbf{h}_i(\boldsymbol{\theta}_i)$ is the function defining the PK/PD (PK/PD) model, including subject specific variables (e.g., drug doses), and $\boldsymbol{\theta}_i$ is the vector of model parameters (random effects) ($\boldsymbol{\theta}_i \in R^p$). In (1) $\mathbf{G}_i(\boldsymbol{\theta}_i, \boldsymbol{\beta})$ is a positive definite covariance matrix ($\mathbf{G}_i \in R^{m_i \times m_i}$) that may depend upon $\boldsymbol{\theta}_i$ as well as on other parameters $\boldsymbol{\beta}$ (fixed effects) ($\boldsymbol{\beta} \in R^q$).

At the second stage given by (2), a finite mixture model with K multivariate normal components is used to describe the population distribution. The weights $\{w_k\}$ are nonnegative numbers summing to one, denoting the relative size of each mixing component (subpopulation), for which $\boldsymbol{\mu}_k(\boldsymbol{\mu}_k \in R^p)$ is the mean vector and $\boldsymbol{\Sigma}_k(\boldsymbol{\Sigma}_k \in R^{p \times p})$ is the positive definite covariance matrix.

Letting ϕ represent the collection of parameters, $\{\boldsymbol{\beta}, (w_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), k = 1, \ldots, K\}$, the population problem involves estimating ϕ given the observation data $\{\mathbf{Y}_1, \ldots, \mathbf{Y}_n\}$. The maximum likelihood estimate (MLE) can be obtained by maximizing the overall data likelihood L with respect to ϕ. Under the $i.i.d.$ assumption of the individual parameters $\{\boldsymbol{\theta}_i\}$, L is given by the expression

$$L(\phi) = \prod_{i=1}^{n} \int p(\mathbf{Y}_i|\boldsymbol{\theta}_i, \boldsymbol{\beta}) \sum_{k=1}^{K} w_k p(\boldsymbol{\theta}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, d\boldsymbol{\theta}_i. \tag{3}$$

The MLE of ϕ is defined as ϕ_{ML} with $L(\phi_{ML}) \geqslant L(\phi)$ for all ϕ in the parameter space.

3. Solution via the EM algorithm

The EM algorithm, originally introduced by Dempster et al. (1977), is a widely applicable approach to the iterative computation of MLEs. It was used by Schumitzky (1995) and Walker (1996) to solve the nonlinear random effects maximum likelihood problem for a second stage model consisting of a single normal distribution. The EM algorithm is typically formulated in terms of "complete" versus "missing" data structure. Consider the model given by (1) and (2) for the important case

$$\mathbf{G}_i(\boldsymbol{\theta}_i, \boldsymbol{\beta}) = \sigma^2 \mathbf{H}_i(\boldsymbol{\theta}_i), \tag{4}$$

where $\mathbf{H}_i(\boldsymbol{\theta}_i)$ is a known function and $\boldsymbol{\beta} = \sigma^2$. The component label vector $\mathbf{z}_i$ is introduced as a K-dimensional indicator such that $z_i(k)$ is one or zero depending on whether or not the parameter $\boldsymbol{\theta}_i$ arises from the kth mixing component. The individual subject parameters $(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n)$ are regarded as unobserved random variables. The "complete" data are then represented by $\mathbf{Y}_c = \{(\mathbf{Y}_i, \boldsymbol{\theta}_i, \mathbf{z}_i), i = 1, \ldots, n\}$ with $\{\boldsymbol{\theta}_i, \mathbf{z}_i\}$ representing the "missing" data.

The algorithm starts with $\phi^{(0)}$ and moves from $\phi^{(r)}$ to $\phi^{(r+1)}$ at the rth iteration. At the E-step, define

$$Q(\phi, \phi^{(r)}) = E\{\log L_c(\phi)|\mathbf{Y}, \phi^{(r)}\},$$

where $\log L_c(\phi)$ is the complete data likelihood given by

$$\log L_c(\phi) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_i(k) \, \log \, p(\mathbf{Y}_i, \boldsymbol{\theta}_i|\sigma^2, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \tag{5}$$

Now

$$E\{z_i(k)|\mathbf{Y}, \phi\} = \mathrm{pr}\{z_i(k) = 1|\mathbf{Y}, \phi\} = \tau_i(k),$$

and by Bayes' Theorem,

$$\tau_i(k) = \frac{w_k \, p(\mathbf{Y}_i|\sigma^2, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{k=1}^{K} w_k \, p(\mathbf{Y}_i|\sigma^2, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)} = \frac{w_k \int p(\mathbf{Y}_i|\sigma^2, \boldsymbol{\theta}_i) p(\boldsymbol{\theta}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, \mathrm{d}\boldsymbol{\theta}_i}{\sum_{k=1}^{K} w_k \int p(\mathbf{Y}_i|\sigma^2, \boldsymbol{\theta}_i) p(\boldsymbol{\theta}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, \mathrm{d}\boldsymbol{\theta}_i}.$$

Introducing the notation

$$g_{ik}(\boldsymbol{\theta}_i, \phi) = \frac{w_k \, p(\mathbf{Y}_i|\sigma^2, \boldsymbol{\theta}_i) p(\boldsymbol{\theta}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{k=1}^{K} w_k \int p(\mathbf{Y}_i|\sigma^2, \boldsymbol{\theta}_i) p(\boldsymbol{\theta}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, \mathrm{d}\boldsymbol{\theta}_i},$$

then

$$\tau_i(k) = \int g_{ik}(\boldsymbol{\theta}_i, \phi) \, \mathrm{d}\boldsymbol{\theta}_i.$$

The expected value of (5) is given by

$$Q(\phi, \phi^{(r)}) = \sum_{i=1}^{n} \int g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)}) \, \log \, p(\mathbf{Y}_i, \boldsymbol{\theta}_i|\sigma^2, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, \mathrm{d}\boldsymbol{\theta}_i,$$

where

$$\log \, p(\mathbf{Y}_i, \boldsymbol{\theta}_i|\sigma^2, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = C - \frac{m_i}{2} \log(\sigma^2) - \frac{1}{2\sigma^2} (\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))^{\mathrm{T}} \mathbf{H}_i(\boldsymbol{\theta}_i)^{-1} (\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))$$

$$- \frac{1}{2}(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)^{\mathrm{T}} \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{\theta}_i - \boldsymbol{\mu}_k) - \frac{1}{2} \log |\boldsymbol{\Sigma}_k|$$

for some constant C.

The M-step takes $\phi^{(r)} \to \phi^{(r+1)}$ where $\phi^{(r+1)}$ is the unique optimizer of $Q(\phi, \phi^{(r)})$ such that $\phi^{(r+1)} = \arg \max_{\phi} Q$ $(\phi, \phi^{(r)})$. Let $\phi' = \{\boldsymbol{\beta}, (\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), k = 1, \ldots, K\}$, then the optimizer of $Q(\phi, \phi^{(r)})$ relative to ϕ' occurs at interior points, and the corresponding components of $\phi^{(r+1)}$ are the unique solution to

$$\frac{\partial}{\partial \phi'} Q(\phi, \phi^{(r)}) \bigg|_{\phi^{(r+1)}} = 0. \tag{6}$$

From the expression of $\log p(\mathbf{Y}_i, \boldsymbol{\theta}_i | \sigma^2, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$,

$$\frac{\partial}{\partial \boldsymbol{\mu}_k} \log L_c(\phi) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_i(k)\{(\boldsymbol{\Sigma}_k)^{-1}(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)\},$$

so

$$E\left\{\frac{\partial}{\partial \boldsymbol{\mu}_k} \log L_c(\phi) | \mathbf{Y}, \phi^{(r)}\right\} = \sum_{i=1}^{n} \int g_{ik}(\boldsymbol{\theta}_i \phi^{(r)})\{\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)\}\, \mathrm{d}\boldsymbol{\theta}_i.$$

Also,

$$\frac{\partial}{\partial \boldsymbol{\Sigma}_k} \log L_c(\phi) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_i(k)\left\{\frac{1}{2}\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)^{\mathrm{T}}\boldsymbol{\Sigma}_k^{-1} - \frac{1}{2}\boldsymbol{\Sigma}_k^{-1}\right\}$$

and

$$\frac{\partial}{\partial (\sigma^2)} \log L_c(\phi) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_i(k)\left\{-\frac{m_i}{2}\frac{1}{\sigma^2} + \frac{1}{2}\frac{(\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))^{\mathrm{T}}\mathbf{H}_i(\boldsymbol{\theta}_i)^{-1}(\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))}{\sigma^4}\right\}.$$

The unique solution of (6) is thus given by

$$\boldsymbol{\mu}_k^{(r+1)} = \frac{\sum_{i=1}^{n} \int \boldsymbol{\theta}_i g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i}{\sum_{i=1}^{n} \int g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i}, \tag{7}$$

$$\boldsymbol{\Sigma}_k^{(r+1)} = \frac{\sum_{i=1}^{n} \int (\boldsymbol{\theta}_i - \boldsymbol{\mu}_k^{(r+1)})(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k^{(r+1)})^{\mathrm{T}} g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i}{\sum_{i=1}^{n} \int g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i}, \tag{8}$$

and

$$(\sigma^2)^{(r+1)} = \frac{\sum_{i=1}^{n} \sum_{k=1}^{K} \int (\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))^{\mathrm{T}}\mathbf{H}_i(\boldsymbol{\theta}_i)^{-1}(\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i}{\sum_{i=1}^{n} m_i}. \tag{9}$$

The updated estimates $\{w_k^{(r+1)}\}$ are calculated independently. If $\mathbf{z}_i$ were observable, then the MLE of w_k would be $\hat{w}_k = (1/n)\sum_{i=1}^{n} z_i(k)$. By replacing each $\mathbf{z}_i$ by its conditional expectation from the E step, the updating for w_k is given by (see McLachlan and Peel, 2000)

$$w_k^{(r+1)} = \frac{1}{n} \sum_{i=1}^{n} \tau_i(k) = \frac{1}{n} \int g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i. \tag{10}$$

Dempster et al. (1977) showed that the resulting sequence $\{\phi^{(r+1)}\}$ has the likelihood improving property $L(\phi^{(r+1)}) \geqslant L(\phi^{(r)})$. It can be shown that the above updates are well defined, that is for all $1 \leqslant k \leqslant K$, if $w_k^{(0)} > 0$ then $w_k^{(r+1)} > 0$ so that $\sum_{i=1}^{n} \int g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)})\, \mathrm{d}\boldsymbol{\theta}_i > 0$ and $\boldsymbol{\Sigma}_k^{(r+1)}$ are positive definite. Wu (1983) and Tseng (2005) gave the sufficient conditions for the convergence of $\phi^{(r)}$ to a stationary point of the likelihood function $L(\phi)$. A number of starting positions are suggested, however, in an effort to ensure convergence to a global maximum.

In order to implement the algorithm all the integrals in (7)–(10) must be evaluated at each iterative step. For the non-mixture problem involving a relatively simple PK model, Walker (1996) proposed Monte Carlo integration to evaluate the required integrals. We and others (Ng et al., 2005) have found that importance sampling is preferable to the Monte Carlo integration for approximating the integrals in the EM algorithm for a number of representative models of interest in PK/PD. We have also applied importance sampling to the current mixture model.

We note that all the integrals above have the following form:

$$\int f(\boldsymbol{\theta}_i) g_{ik}(\boldsymbol{\theta}_i, \phi) \, \mathrm{d}\boldsymbol{\theta}_i = \frac{\int f(\boldsymbol{\theta}_i) w_k \, p(\mathbf{Y}_i | \sigma^2, \boldsymbol{\theta}_i) \, p(\boldsymbol{\theta}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, \mathrm{d}\boldsymbol{\theta}_i}{\sum_{k=1}^K \int w_k \, p(\mathbf{Y}_i | \sigma^2, \boldsymbol{\theta}_i) \, p(\boldsymbol{\theta}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \, \mathrm{d}\boldsymbol{\theta}_i}.$$

For each mixing component, a number of samples are taken from an envelope distribution, $\boldsymbol{\theta}_{i(k)}^{(1)}, \ldots, \boldsymbol{\theta}_{i(k)}^{(T)} \sim_{i.i.d.} p_{e(k)}(\boldsymbol{\theta}_i)$, and used to approximate the integrals as follows:

$$\int f(\boldsymbol{\theta}_i) g_{ik}(\boldsymbol{\theta}_i, \phi) \, \mathrm{d}\boldsymbol{\theta}_i \cong \frac{\sum_{l=1}^{\mathrm{T}} f(\boldsymbol{\theta}_{i(k)}^{(l)}) w_k \, p(\mathbf{Y}_i | \sigma^2, \boldsymbol{\theta}_{i(k)}^{(l)}) \, p(\boldsymbol{\theta}_{i(k)}^{(l)} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) / p_{e(k)}(\boldsymbol{\theta}_{i(k)}^{(l)})}{\sum_{k=1}^K \sum_{l=1}^{\mathrm{T}} w_k \, p(\mathbf{Y}_i | \sigma^2, \boldsymbol{\theta}_{i(k)}^{(l)}) \, p(\boldsymbol{\theta}_{i(k)}^{(l)} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) / p_{e(k)}(\boldsymbol{\theta}_{i(k)}^{(l)})}. \tag{11}$$

For each mixing component for each subject, the envelope distribution is taken to be a multivariate normal density using the subject's previously estimated conditional mean and conditional covariance as its mean and covariance. Therefore all the random samples are independent and individual specific. For details of the importance sampling approach in general, see Geweke (1989). The number of independent samples, T, will depend on the complexity of the model and the required accuracy in the integral approximations.

4. Classification of subjects

It is of interest to assign each individual subject to a subpopulation. Such a classification will allow further investigation into the genetic basis of any identified PK/PD polymorphism. The quantity $\tau_i(k) = E\{z_i(k) | \mathbf{Y}, \phi\}$ in the E step is the posterior probability that the ith individual belongs to the kth mixing component. The classification involves assigning an individual to the subpopulation associated with the highest posterior probability of membership. For example, for each i, $(i = 1, \ldots, n)$, set

$$\hat{z}_i(k) = 1 \quad \text{if } k = \arg \max_c \tau_i(c),$$

or to zero otherwise. No additional computation is required since all the $\tau_i(k)$ are evaluated during each EM step.

5. Standard errors

Assuming the regularity conditions from Philppou and Roussas (1975), it can be shown that asymptotically as $n \to \infty$,

$$Cov(\phi_{ML}) \approx \left(\sum_{i=1}^n \mathbf{V}_i(\phi_{ML}) \right)^{-1},$$

where $\mathbf{V}_i(\phi) = ((\partial/\partial\phi) \log p(\mathbf{Y}_i | \phi))((\partial/\partial\phi) \log p(\mathbf{Y}_i | \phi))^{\mathrm{T}}$. Now

$$\frac{\partial}{\partial\phi} \log p(\mathbf{Y}_i | \phi) = \sum_{k=1}^K \int \left\{ \frac{\partial}{\partial\phi} \log p(\mathbf{Y}_i, \boldsymbol{\theta}_i | \phi) \right\} g_{ik}(\boldsymbol{\theta}_i, \phi) \, \mathrm{d}\boldsymbol{\theta}_i,$$

and the gradient components are calculated for $k = 1, \ldots, K$ as

$$\mathbf{s}_{\boldsymbol{\mu}_k} = \frac{\partial}{\partial\boldsymbol{\mu}_k} \log p(\mathbf{Y}_i | \phi) = \int g_{ik}(\boldsymbol{\theta}_i, \phi) \{\boldsymbol{\Sigma}_k^{-1} (\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)\} \, \mathrm{d}\boldsymbol{\theta}_i,$$

$$\mathbf{s}'_{\boldsymbol{\Sigma}_k} = \frac{\partial}{\partial\boldsymbol{\Sigma}_k} \log p(\mathbf{Y}_i | \phi) = \int g_{ik}(\boldsymbol{\theta}_i, \phi) \left\{ \left(-\frac{1}{2}\right) \boldsymbol{\Sigma}_k^{-1} + \frac{1}{2}\boldsymbol{\Sigma}_k^{-1} (\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)^{\mathrm{T}} \boldsymbol{\Sigma}_k^{-1} \right\} \, \mathrm{d}\boldsymbol{\theta}_i,$$

$$\mathbf{s}_{\boldsymbol{\Sigma}_k} = \text{lower triangular part of } \left\{ 2\mathbf{s}'_{\boldsymbol{\Sigma}_k} - \text{diag}(\mathbf{s}'_{\boldsymbol{\Sigma}_k}) \right\},$$

$$s_{\sigma^2} = \frac{\partial}{\partial \sigma^2} \log p(\mathbf{Y}_i|\phi) = \sum_{k=1}^{K} \int g_{ik}(\boldsymbol{\theta}_i, \phi) \left\{ -\frac{1}{2}\frac{m_i}{\sigma^2} + \frac{1}{2}\frac{(\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))^{\mathrm{T}}\mathbf{H}_i^{-1}(\boldsymbol{\theta}_i)(\mathbf{Y}_i - \mathbf{h}_i(\boldsymbol{\theta}_i))}{\sigma^4} \right\} \mathrm{d}\boldsymbol{\theta}_i,$$

and for $k = 1, \ldots, K - 1$,

$$s_{w_k} = \frac{\partial}{\partial w_k} \log p(\mathbf{Y}_i|\phi) = \frac{1}{w_k} \int g_{ik}(\boldsymbol{\theta}_i, \phi) \, \mathrm{d}\boldsymbol{\theta}_i - \frac{1}{w_K} \int g_{iK}(\boldsymbol{\theta}_i, \phi) \, \mathrm{d}\boldsymbol{\theta}_i.$$

Introduce the notation $\mathbf{s}_{\boldsymbol{\varpi}_k} = ((\mathbf{s}_{\boldsymbol{\Sigma}_k})_{1,1}, (\mathbf{s}_{\boldsymbol{\Sigma}_k})_{2,1}, \ldots, (\mathbf{s}_{\boldsymbol{\Sigma}_k})_{p,p})$, where $(\mathbf{s}_{\boldsymbol{\Sigma}_k})_{i,j}$ is the component of $\mathbf{s}_{\boldsymbol{\Sigma}_k}$ in the (i, j) position. Put these results together to produce the vector

$$\mathbf{s}_i = (\mathbf{s}_{\boldsymbol{\mu}_1}, \ldots, \mathbf{s}_{\boldsymbol{\mu}_K}, \mathbf{s}_{\boldsymbol{\varpi}_1}, \ldots, \mathbf{s}_{\boldsymbol{\varpi}_K}, s_{w_1}, \ldots, s_{w_{K-1}}, s_{\sigma^2}),$$

so

$$Cov(\phi_{ML}) = \left(\sum_{i=1}^{n} \mathbf{s}_i \mathbf{s}_i^{\mathrm{T}} \right)^{-1}.$$

All the computations can be performed during the importance sampler calculation at the final iteration of the EM algorithm.

6. Example

In this section a simulation study is conducted to evaluate the proposed algorithm for calculating the exact MLEs for a population finite mixture model. A one compartment PK model is used, with the observations of plasma concentration given by

$$y_{ji} = \frac{D}{V_i} \exp(-k_i * t_j)(1 + \varepsilon_{ji}),$$

where D is a bolus drug administration with 100 units of dose, V represents the volume of distribution and k is the elimination rate constant. For all the individuals, $m_i = 5$ with $t_1 = 1.5$, $t_2 = 2$, $t_3 = 3$, $t_4 = 4$ and $t_5 = 5.5$. The within-individual error is assumed to be $i.i.d.$ with variance 0.01. Data sets were simulated from this model for each of 100 subjects sampled from the following population model:

$$V_i \sim_{i.i.d.} N(20, 2^2),$$

$$k_i \sim_{i.i.d.} 0.8N(0.3, 0.06^2) + 0.2N(0.6, 0.06^2),$$

where V_i and k_i are assumed to be independent. A total of 200 such population data sets were generated. This model represents the PK of a drug with an elimination that can be characterized by two distinct subpopulations.

The formulation of Section 2 has been modified to accommodate the important case where a subset of parameters are modeled by a multivariate normal distribution and the remaining parameters follow a mixture of normals (see modified updating formulas in the Appendix). The MLEs were obtained for each of the 200 population sets using the EM algorithm with importance sampling described above. For each of the estimated parameters ϕ, its percent prediction error was calculated for each population data set as

$$\text{pe}_j = 100 \times (\phi_j^{\text{ML}} - \phi_j)/\phi_j, \quad j = 1, \ldots, 200.$$

These percent prediction errors were used to calculate the mean prediction error and root mean square prediction error for each parameter. In addition, for each population data set, the calculated standard errors were used to construct 95% confidence intervals for all estimated parameters. The percent coverage of these confidence intervals was then tabulated over the 200 population data sets. Finally, the individual subject classification accuracy was evaluated for each population data set.
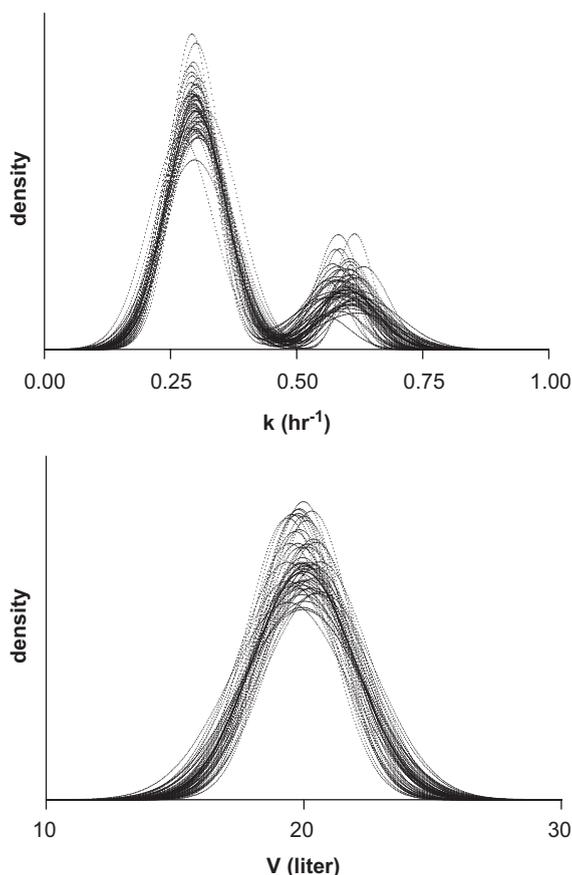
Fig. 1. True (solid line) and estimated (dotted line) population densities of k (upper panel) and V (lower panel) from the population simulation analysis.

Table 1
Mean of parameter estimates (over 200 simulations); mean percent prediction error (PE) and root mean square percent prediction error (RMSE); percent coverage of 95% confidence interval

Parameter	Population values	Mean of estimates	Mean PE (%)	RMSE (%)	Coverage of 95% Cl (%)
μ_V	20	19.978	0.043586	1.0399	94.5
μ_{k1}	0.3	0.29982	−0.09045	1.6491	96.5
μ_{k2}	0.6	0.60029	0.10042	2.6455	90.5
w_1	0.8	0.80448	0.55991	5.4248	94.5[a]
σ_V^2	4	3.7605	−5.0867	23.822	94.5
σ_{k1}^2	0.0036	0.003573	−1.0539	14.88	91.0
σ_{k2}^2	0.0036	0.003195	−10.857	40.236	83.5
σ	0.1	0.099931	−0.06857	4.0618	95.5

[a]*The coverage of the transformed variable* $\log w_1/(1 - w_1)$ *is shown.*

Fig. 1 provides a graphical illustration of the results showing the true population distributions of V and k along with the estimated distribution obtained from the 200 simulated population data sets. Quantitative results are presented in Table 1, which gives mean and root mean square prediction errors (RMSE) as well as the percent coverage of the calculated confidence intervals for each of the estimated parameters. The parameter estimates, overall, match the population values and the percent coverage of confidence intervals is reasonable. The estimates of population variance

Table 2
Parameter estimates by importance sampling with 1000, 2000 and 3000 samples

Parameter	$T = 1000$	$T = 2000$	$T = 3000$
μ_V	19.85039	19.84843	19.85201
μ_{k1}	0.30840	0.30837	0.30853
μ_{k2}	0.60038	0.60029	0.60079
w_1	0.75164	0.75157	0.75269
σ_V^2	5.73544	5.72227	5.70365
σ_{k1}^2	0.00327	0.00323	0.00328
σ_{k2}^2	0.00202	0.00203	0.00200
σ	0.09706	0.09687	0.09692
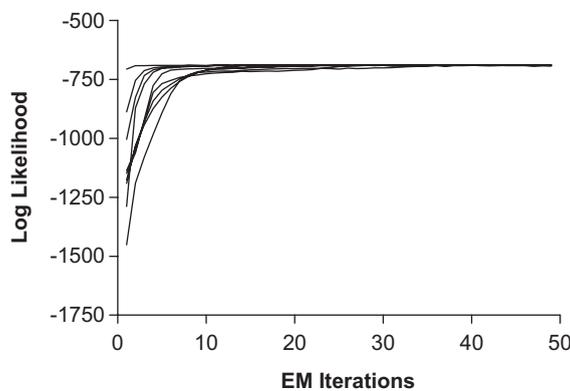
Fig. 2. Convergence of log-likelihood by starting the EM algorithm from nine positions.

have relatively greater biases and RMSE. Over the 200 population data sets, on average 1.54 out of 100 subjects were classified in the wrong subpopulation. The largest number of subjects misclassified was 4, while all the subjects were correctly classified in 83 of the 200 population data sets.

Central to the calculation of the MLEs is the computation of the integrals in (7)–(9), as approximated by importance sampling in our implementation. Using one of the 200 population data sets we examined the influence of the number of samples (T) used in the importance sampler, as well as the number of EM iterations required to achieve two digits of accuracy for each of the estimated parameters. Table 2 presents the parameter estimates from 50 EM iterations by using 1000, 2000 and 3000 samples in the important sampling. Accuracy to two digits was obtained with 1000 samples. Based on this experience, T was taken to be 1000 in this simulations study and 50 EM iterations were run on each data set.

Fig. 2 demonstrates the convergence of log-likelihood values for a particular data set by starting the 50 EM iterations from nine different positions. The log-likelihood values were approximated via Monte Carlo integration:

$$\log\{L(\phi)\} \cong \sum_{i=1}^{n} \log \left\{ \frac{1}{T} \sum_{k=1}^{K} \sum_{l=1}^{T} w_k \, p(\mathbf{Y}_i | \sigma^2, \boldsymbol{\theta}_{i(k)}^{(l)}) p(\boldsymbol{\theta}_{i(k)}^{(l)} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) / p_{e(k)}(\boldsymbol{\theta}_{i(k)}^{(l)}) \right\},$$

where $\boldsymbol{\theta}_{i(k)}^{(1)}, \ldots, \boldsymbol{\theta}_{i(k)}^{(T)} \sim_{i.i.d.} p_{e(k)}(\boldsymbol{\theta}_i)$. In any particular example, of course, the number of EM iterations, the number of samples T used to approximate integrals, as well as the use of different starting guess will depend on the experiment design and complexity of the model.

7. Discussion

In this paper, an EM algorithm for maximum likelihood estimation of nonlinear random effects mixture models is presented that has application in pharmacokinetic/pharmacodynamic population modeling studies. It extends the previous work on the use of the EM algorithm for MLE of nonlinear random effects models, to the case of finite mixture models, and reinforces the practicability of using exact (no linearizing approximation) MLE estimation in PK/PD modeling studies (see also, Kuhn and Lavielle, 2005 for a stochastic EM variation). The parametric mixture model MLE approach presented also complements previous work on nonparametric Bayesian and smoothed nonparametric MLE, in addressing the increasingly important problem of identifying subpopulations with distinct PK/PD properties in drug development trials. We note that approximate maximum likelihood methods using mixture models are also available in NONMEM.

The EM algorithm has been used extensively for fitting linear mixture models in numerous applications in diverse fields of study. Even for linear problems involving mixture of normal components, a number of challenges attend the use of the EM algorithm for maximum likelihood estimation (McLachlan and Peel, 2000) that are also relevant to the nonlinear random effects problem. These include: potential unboundness of the likelihood for heteroscedastic covariance components; local maxima of the likelihood function; and choice of the number of mixing components. Application of the algorithm for nonlinear random effects mixture models presented here can be guided by the extensive work related to these issues for linear mixture modeling.

We have investigated the possible unboundness of the likelihood for the example considered in this paper. If, for example, in the first component of the mixture, μ_1 satisfies $\mathbf{h}_i(\mu_1) = \mathbf{Y}_i$ for any i and $\Sigma_1 \to \mathbf{0}$, then the likelihood will tend to infinity, and the global maximizer will not exist. By restricting the covariance matrices $\Sigma_k, k = 1, \ldots, K$ to be equal (homoscedastic components), as is often done in mixture modeling, the unboundness of the likelihood will be eliminated. In our example with heteroscedastic variance components, each individual has five error-associated observations, while the parameter space is of dimension two. The condition for likelihood singularity is therefore very unlikely to be satisfied.

Future work is also needed to extend the algorithm to include important practical cases involving more general error variance models and random effects covariates.

Appendix

For the example in Section 6, as in other PK/PD problem, it is often reasonable to assume that the mechanism of genetic polymorphism applies to only part of the system, for example, drug metabolism or drug target. It is therefore desirable to partition the parameter θ_i into two components, one (α_i) that follows a mixture of multivariate normals and the second (β_i) defined by a single multivariate normal distribution: $\theta_i = \{\alpha_i, \beta_i\}$, where α_i and β_i are independent. The EM updates for this special case are given by

$$(\mu\_\alpha)_k^{(r+1)} = \frac{\sum_{i=1}^n \int \alpha_i g_{ik}(\theta_i, \phi^{(r)}) \, \mathrm{d}\theta_i}{\sum_{i=1}^n \int g_{ik}(\theta_i, \phi^{(r)}) \, \mathrm{d}\theta_i},$$

$$(\Sigma\_\alpha)_k^{(r+1)} = \frac{\sum_{i=1}^n \int (\alpha_i - (\mu\_\alpha)_k^{(r+1)})(\alpha_i - (\mu\_\alpha)_k^{(r+1)})^{\mathrm{T}} g_{ik}(\theta_i, \phi^{(r)}) \, \mathrm{d}\theta_i}{\sum_{i=1}^n \int g_{ik}(\theta_i, \phi^{(r)}) \, \mathrm{d}\theta_i} \quad \text{for } 1 \leqslant k \leqslant K,$$

$$(\mu\_\beta)^{(r+1)} = \frac{\sum_{i=1}^n \{\sum_{k=1}^K \int \beta_i g_{ik}(\theta_i, \phi^{(r)}) \, \mathrm{d}\theta_i\}}{n}$$

and

$$(\mathbf{\Sigma}\_\boldsymbol{\beta})^{(r+1)} = \frac{\sum_{i=1}^{n} \left\{ \sum_{k=1}^{K} \int (\boldsymbol{\beta}_i - (\boldsymbol{\mu}\_\boldsymbol{\beta})^{(r+1)})(\boldsymbol{\beta}_i - (\boldsymbol{\mu}\_\boldsymbol{\beta})^{(r+1)})^{\mathrm{T}} g_{ik}(\boldsymbol{\theta}_i, \phi^{(r)}) \, \mathrm{d}\boldsymbol{\theta}_i \right\}}{n}.$$

The updates for $\{w_k, 1 \leqslant k \leqslant K\}$ and σ^2 are the same as in Section 3.

References

Beal, S.L., Sheiner, L.B., 1979. NONMEM User's Guide, Part I. San Francisco: Division of Clinical Pharmaology, University of California.

Davidian, M., Gallant, A.R., 1993. The non-linear mixed effects model with a smooth random effects density. Biometrika 80, 475–488.

Davidian, M., Giltinan, M., 1995. Nonlinear Models for Repeated Measurement Data. Chapman and Hall, New York.

Dempster, A.P., Laird, N., Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm. J. Roy. Statist. Soc. B 39, 1–38.

Evans, W.E., Relling, M.V., 1999. Pharmacogenomics: translating functional genomics into rational therapeutics. Science 286, 487–491.

Geweke, J., 1989. Bayesian inference in econometric models using Monte Carlo integration. Econometrica 57, 1317–1340.

Kuhn, E., Lavielle, M., 2005. Maximum likelihood estimation in nonlinear mixed effects models. Comput. Statist. Data Anal. 49, 1020–1038.

Mallet, A., 1986. A maximum likelihood estimation method for random coefficient regression models. Biometrika 73, 645–656.

McLachlan, G.J., Peel, D., 2000. Finite Mixture Models. Wiley, New York.

Ng, C.M., Joshi, A., Dedrick, R.L., Garovoy, M.R., Bauer, R.J., 2005. Pharmacokinetic-pharmacodynamic-efficacy analysis of Efalizumab in patients with moderate to severe psoriasis. Pharm. Res. 77, 1088–1100.

Philppou, A., Roussas, G., 1975. Asymptotic normality of the maximum likelihood estimate in the independent but not identically distributed case. Ann. Inst. Statist. Math. 27, 45–55.

Rosner, G.L., Muller, P., 1997. Bayesian population pharmacokinetic and pharmacodynamic analyses using mixture models. J. Pharmacokinet. Biopharm. 25, 209–234.

Schumitzky, A., 1995. EM Algorithms and two stage methods in pharmacokinetic population analysis. In: D.Z. D'Argenio (Ed.), Advanced methods of Pharmacokinetic and Pharmacodynamic Systems Analysis, vol. II, Plenum Press, New York, pp. 145-160.

Sheiner, L.B., Rosenberg, B., Melmon, K.L., 1972. Modeling of individual pharmacokinetics for computer-aided drug dosing. Comput. Biomed. Res. 5, 441–459.

Tseng, P., 2005. An analysis of the EM algorithm and entropy-like proximal point methods. Math. Oper. Res. 29, 27–44.

Wakefiled, J.C., Walker, S.G., 1997. Bayesian nonparametric population models: formulation and comparison with likelihood approaches. J. Pharmacokinet. Biopharm. 25, 235–253.

Walker, S., 1996. An EM algorithm for nonlinear random effects models. Biometrics. 52, 934–944.

Wu, C.F., 1983. On the convergence properties of the EM algorithm. Ann. Statist. 11, 95–103.