Published in final edited form as:

J R Stat Soc Ser C Appl Stat. 2007 March ; 56(2): . doi:10.1111/j.1467-9876.2007.00569.x.

Semiparametric Bayesian classification with longitudinal markers

Rolando De la Cruz-Mesía,

Pontificia Universidad Católica de Chile, Santiago, Chile

Fernando A. Quintana, and

Pontificia Universidad Católica de Chile, Santiago, Chile

Peter Müller

University of Texas, Houston, USA

Summary

We analyse data from a study involving 173 pregnant women. The data are observed values of the β human chorionic gonadotropin hormone measured during the first 80 days of gestational age, including from one up to six longitudinal responses for each woman. The main objective in this study is to predict normal *versus* abnormal pregnancy outcomes from data that are available at the early stages of pregnancy. We achieve the desired classification with a semiparametric hierarchical model. Specifically, we consider a Dirichlet process mixture prior for the distribution of the random effects in each group. The unknown random-effects distributions are allowed to vary across groups but are made dependent by using a design vector to select different features of a single underlying random probability measure. The resulting model is an extension of the dependent Dirichlet process model, with an additional probability model for group classification. The model is shown to perform better than an alternative model which is based on independent Dirichlet processes for the groups. Relevant posterior distributions are summarized by using Markov chain Monte Carlo methods.

Keywords

Dependent non-parametric model; Discriminant analysis; Longitudinal data; Markov chain Monte Carlo sampling; Non-parametric modelling; Random-effects models; Species sampling models

1. Introduction

We develop a semiparametric Bayesian approach for classification based on longitudinal markers. We define a suitable extension of hierarchical models to allow such classification. We introduce a new class of models building on the dependent Dirichlet process (DDP) models that were proposed in MacEachern (1999). In a motivating example we compare the performance of the proposed model with parametric Bayesian inference and with traditional maximum-likelihood-based classification.

In many disease areas longitudinal markers allow early detection of a specific disease. Atypical example is the use of prostate specific antigen profiles over time as a marker for prostate cancer (Morrell *et al.*, 1995; Inoue *et al.*, 2004). A common feature of inference

related to such data is the need for classification rules that allow coherent and easy sequential updating as the data for a new patient accrue over time. In this paper we propose a model-based semiparametric Bayesian approach to classification that facilitates such sequential updating. The motivating application concerns the classification of pregnancies into normal and abnormal. To detect a number of complications during pregnancy, a variety of quantities are measured at the antenatal examinations. One of these clinical variables is the beta subunit of human chorionic gonadotropin (β -HCG) which shows dramatic changes in women during pregnancy. It has been established that values of β -HCG are different in women who have normal pregnancies with terminal deliveries from those women who have spontaneous abortions or other types of adverse pregnancy outcomes (France *et al.*, 1996). This association has made it possible to classify, with some degree of uncertainty, the outcome of pregnancy. The inference problem is formally described as a discriminant analysis based on the longitudinal β -HCG outcome.

Classical linear discriminant analysis classifies subjects into one of g groups or populations by using multivariate observations. Usually, these vector-valued observations are obtained from cross-sectional studies and represent different subject characteristics such as age, gender or other relevant factors. In general, a common and unrestricted covariance matrix is assumed in the g different groups. Modifications of this method have also been used to classify subjects when the vector of multivariate observations represents repeated measures collected in a longitudinal study. Azen and Afifi (1972) introduced a two-stage model in which a discriminant function is obtained at each time point. In the second stage, the coefficients enter a linear regression *versus* time to obtain a slope and intercept. These slopes and intercepts are then used as input for a final linear discriminant function. This method is limited by the fact that multiple observations per subject are required to allocate a subject to one of g groups at any point in time.

Albert (1983) extended the classical concepts of discriminant analysis to multivariate response curves observed over fixed time intervals. Using interpolation or curve fitting procedures, a time-varying distance measure between the individual curve and group-specific curves is used to allocate a subject to a group. This methodology requires that the response curves in the training sample are fully observed over the time interval considered.

Albert and Kshirsagar (1993) proposed an exploratory method based on a growth curve structure embedded in a canonical variate analysis to achieve dimension reduction in a discriminant analysis framework. They suggested this approach for classification but did not apply it in that setting. No longitudinal data structures other than growth curves were considered.

An important limitation in the use of linear discriminant analysis for longitudinal data is that the method is applicable only for essentially balanced data, an increasingly exceptional situation in longitudinal studies. Therefore, an approach is needed that does not rely on complete observations over time. In recent years some work has been done regarding discriminant analysis for longitudinal data by using both linear and non-linear random-effects models. Tomasko *et al.* (1999) modified linear discriminant analysis by using mixed model multivariate analysis of variance for the estimation of fixed effects and for a determination of various structures of covariance matrices, including unstructured, compound symmetry and autoregressive of order 1. Brown *et al.* (2001) discussed Bayesian methods in discriminant analysis using linear random-effects models. Marshall and Barón (2000) considered non-linear random-effects models to describe profiles in different groups and stated the optimal allocation rule. Fieuws *et al.* (2003) used linear as well as non-linear random-effects models for the description of group-specific profiles. Recently, De la Cruz-

Mesía and Quintana (2006) have given a Bayesian version to the classification problem for longitudinal data.

All these approaches consider parametric models for the random effects. Unrelated to the classification problem, several recent references generalize restrictive parametric models for longitudinal data by placing a non-parametric prior on the random-effects distribution. The literature includes, among many others, Bush and MacEachern (1996), Davidian and Gallant (1993), Ishwaran and Takahara (2002), Kleinman and Ibrahim (1998), Mentré and Mallet (1994), Müller *et al.* (2004), Müller and Rosner (1997), Walker and Wakefield (1998) and Zhang and Davidian (2001). In this paper we develop a variation of these semiparametric Bayesian longitudinal data models that is suitable for sequential classification as desired for inference with longitudinal markers. Specifically, we use an analysis-of-variance-DDP model (De Iorio *et al.*, 2004) to introduce semiparametric random-effects models that include dependence across the subpopulations of women with normal and abnormal pregnancies. We complete the model by adding a probability for group indicators. The augmented model for the repeated measurements and group indicators allows us to formalize the classification desired.

The paper is organized as follows. We first give a brief description of the data set in Section 2. In Section 3, we extend the framework of traditional classification methods to the longitudinal hierarchical setting. Section 4 provides a discussion of non-parametric models based on the Dirichlet process, including methods for introducing dependence across related random probability measures. In Section 5 we illustrate the proposed longitudinal classification method by using data on β -HCG measured in women with normal and abnormal pregnancy outcomes. An appropriate posterior simulation scheme based on the Gibbs sampling algorithm is described. Lastly, Section 6 concludes with a final discussion.

2. Data

We consider a data set reporting repeated measurements on β -HCG for $n = 173$ young women, representing 173 different pregnancies over a period of 2 years in a private obstetrics clinic in Santiago, Chile. The values of β -HCG were measured during the first 80 days of gestational age. The women were classified as having normal pregnancies if they had a normal delivery, or as having abnormal pregnancies if they had any complication resulting in a non-terminal delivery and loss of the foetus. The 173 women altogether contribute a total of 375 observations. Each woman is measured from one up to six times. These data were originally presented in Marshall and Barón (2000). Approximately 30% of the women had one β -HCG measurement, 31% had two, 33% had three and 6% had four or more measurements.

Fig. 1 presents the subject-specific $\log(\beta\text{-HCG})$ profiles for women with normal and abnormal pregnancies. The two populations appear clearly distinct when considering the ensemble of profiles. However, for any one of the profiles the classification into one or the other subpopulation is far less certain, in particular when considering series of partial responses. The main inference goal in analysing these data is to provide a rule to classify a new patient. The rule should allow sequential updating as data accrue for the new patient. The classification will critically hinge on the implied inference on the distribution of profiles for each of the two subpopulations. The semiparametric model proposed defines a richer class of random-effects distributions than other models.

3. Classification using hierarchical models

We use an augmentation scheme of semiparametric longitudinal data models to develop the desired model-based classification for longitudinal marker data.

Let $y_i = (y_{i1}, \dots, y_{in_i})'$ represent the observed response vector for the i th patient, recorded at known times $t'_i = (t_{i1}, t_{i2}, \dots, t_{in_i})$. Here n_i is the number of repeated measurements recorded for patient i . Let $x_i \in \{0, 1, \dots, g-1\}$ denote the known class label for the i th patient. In our application $g = 2$, with $x_i = 0$ and $x_i = 1$ indicating normal and abnormal pregnancy respectively. The label x_i is known for women with already reported delivery, but unknown for women with partial data before delivery. Without loss of generality we assume that x_i , $i = 1, \dots, m$, is known, and x_{m+1} is unknown. Also without loss of generality we assume that $x_i \in \{0, 1\}$ takes only two possible values. Let $y^m = (y_1, \dots, y_m, x_1, \dots, x_m)$ denote all data, including the recorded class memberships x_i , up to the m th patient. The classification problem is formalized as reporting $p(x_{m+1}|y^m, y_{m+1})$. Here y_{m+1} is the currently available partial response vector for the new patient $m+1$. We now construct a probability model to allow evaluation of the classification probabilities desired.

Consider a generic semiparametric hierarchical model of the form

$$(y_i|\theta_i) \sim p(y_i|\theta_i), \quad (\theta_i|x_i, \varphi, G_0, G_1) \sim G_{x_i}(\theta_i|\varphi), \quad (G_0, G_1|\psi) \sim F_\psi. \quad (1)$$

In words, data y_i for the i th experimental unit are sampled from a probability model parameterized by a random-effects vector θ_i . The θ_i are generated from a random-effects distribution G_x , with $x = x_i$. The random-effects distribution depends on a covariate x_i that is specific to the i th sampling unit and possibly additional hyperparameters φ . In general, the parameter vector θ_i might be partitioned into common fixed effects θ^F and unit-specific random-effects θ_i^R . Fixed effects are common to all patients and have no patient index i . In our example we use this partition. The model is completed by assuming a prior model for the unknown G_x . If G_x were indexed by a finite dimensional vector of hyperparameters, e.g. normal moments, then the model would reduce to a traditional parametric hierarchical model. In contrast, in a non-parametric Bayesian approach G_x is assumed to be a random probability measure with an appropriate prior probability model F_ψ for the unknown distribution. In other words, F_ψ is a distribution on distributions. Here ψ indicates hyperparameters in the definition of F_ψ . A popular approach is to assume that each G_x arises from a Dirichlet process prior, independently across x , conditional on ψ . The random measures could be linked at the level of the hyper-parameters. We discuss more details of this construction and the proposed alternative model in the next section.

For the top level sampling model $p(y_i|\theta_i)$ in model (1) we assume a non-linear regression

$$y_{ij} = f(\theta_i; t_{ij}) + \varepsilon_{ij}, \quad (2)$$

with a mean function $f(\theta; \cdot)$ parameterized by θ and evaluated at known times t_{ij} , $j = 1, \dots, n_i$. The residual term ε_{ij} is assumed to be normally distributed with mean 0 and variance σ^2 .

Model (1) specifies a joint probability model

$$p(y_1, \dots, y_m | \varphi, x_1, \dots, x_m, \psi),$$

after marginalizing with respect to G_0 , G_1 and θ_i , $i = 1, \dots, m$. To facilitate classification we augment the model with a marginal probability for x_i :

$$\Pr(x_i = x) = \pi_x. \quad (3)$$

The augmented model implies the desired classification as a conditional probability $p(x_{m+1}|y^m, y_{m+1})$, marginalizing with respect to the unknown θ_i , G_x and other possibly unknown hyper-parameters.

In maximum likelihood classification theory, the probability that a future unit y_{m+1} belongs to group or population x is estimated as

$$p(x_{m+1}=x|y_{m+1}, y^m, \hat{\Theta}) \propto \pi_x p(y_{m+1}|x_{m+1}=x, \hat{\Theta})$$

where $\hat{\Theta}$ indicates the maximum likelihood estimate of the fixed effect parameters that remain after integrating out all the random effects. The unit is then classified in that group for which the highest probability is attained.

From a Bayesian viewpoint the classification probabilities are obtained by weighting with the posterior distributions of the parameters. Let $\Theta = (\phi, \psi, \theta_1, \dots, \theta_m, \theta_{m+1})$ denote the vector of all parameters in the model, including those for the new $(m+1)$ th patient. Using Bayes's rule we find the probability that a new unit y_{m+1} belongs to group x as

$$P(x_{m+1}=x|y_{m+1}, y^m) = \int p(x_{m+1}=x|y_{m+1}, y^m, \Theta) p(\Theta|y_{m+1}, y^m) d\Theta \\ \propto \int \pi_x p(y_{m+1}|\Theta, x_{m+1}=x) p(\Theta|y_m) d\Theta. \quad (4)$$

To verify expression (4) use

$$p(x_{m+1}=x|y_{m+1}, y^m, \Theta) = \pi_x p(y_{m+1}|\Theta, x_{m+1}=x) / p(y_{m+1}|\Theta)$$

and

$$p(\Theta|y_{m+1}, y^m) \propto p(y_{m+1}|\Theta) p(\Theta|y_m).$$

The integration is usually analytically intractable. Therefore, we shall construct a set of Markov chain Monte Carlo (MCMC) samples $\{\Theta^{(b)}, b = 1, \dots, B\}$ from the posterior distribution and use the Rao–Blackwellization

$$\hat{p}_x \equiv \frac{1}{B} \sum_{b=1}^B \pi_x p(y_{m+1}|\Theta^{(b)}, x_{m+1}=x) \quad (5)$$

to approximate expression (4). If the prevalences π_x are unknown hyperparameters as well, then expression (5) would use the imputed values $\pi_x^{(b)}$.

Using a *percentage correctly classified* loss function (McLachlan, 2004), the Bayes classification of a future y_{m+1} is given by

$$\hat{x}_{m+1} = \arg \max_x \{p(x_{m+1}=x|y_{m+1}, y^m)\}.$$

The unit is classified in that group for which the highest posterior probability is attained.

4. Semiparametric models for longitudinal classification

We now discuss specific choices for the random probability measure F_ψ in model (1). We start with a review of the DP and some extensions.

The DP is a probability measure on the space of distribution functions defined on some space $\mathcal{X}$ (equipped with a σ -field $\mathcal{B}$). We use $\text{DP}(M, G^*)$ to denote the DP, where $M > 0$ is a scalar (precision parameter) and G^* is a specified base-line distribution that is defined on $(\mathcal{X}, \mathcal{B})$. A random distribution function G on $(\mathcal{X}, \mathcal{B})$ generated from $\text{DP}(M, G^*)$ is almost surely discrete and admits the following representation. Letting δ_a denote a Dirac measure at a we have

$$G = \sum_{l=1}^{\infty} \omega_l \delta_{\mu_l}. \quad (6)$$

The weights ω_l and locations μ_l are generated by the following stick breaking scheme: $\omega_1 = z_1$,

$$\omega_l = z_l \prod_{r=1}^{l-1} (1 - z_r), \quad l=2, 3, \dots,$$

with $z_l \sim \text{IID Be}(1, M)$ and $\mu_l \sim \text{IID } G^*$, independently of the ω_l (Sethuraman, 1994).

The use of DPs to model random distributions entails some limitations. In particular, the random probability measure G is almost surely discrete. A commonly used extension to mitigate this limitation is the DP mixture (DPM) model (Antoniak, 1974). DPM models avoid the discreteness by introducing an additional convolution with a continuous kernel. This model has become popular in applied Bayesian non-parametric work. The typical DPM model assumes

$$\theta_1, \dots, \theta_m \stackrel{\text{IID}}{\sim} G^M(\theta), \quad \text{with } G^M(\theta) = \int f(\theta|\mu) dG(\mu), \text{ and } G \sim \text{DP}(M, G^*), \quad (7)$$

i.e. a mixture with a DP prior on the random mixing measure G . We use G^M to denote the random mixture model with mixing measure G . The mixture model (7) can be equivalently written as a hierarchical model by introducing latent variables μ_i and breaking the mixture as $\theta_i|\mu_i \sim f(\theta|\mu)$ and $\mu_i \sim \text{IID } G$, $i = 1, \dots, m$, and finally $G \sim \text{DP}(M, G^*)$. One of the attractive features of DPM models is the straightforward nature of posterior MCMC simulation. The computational effort is, in principle, independent of the dimensionality of μ_i . Efficient MCMC simulation for general DPM models is discussed, among others, in Bush and MacEachern (1996), Escobar and West (1995), MacEachern and Müller (1998), Neal (2000) and Jain and Neal (2004).

Several references have considered extensions of DP and DPM models to hierarchical models over related random distributions, as needed to model the joint prior on (G_0, G_1) in model (1). Some of the earliest developments of dependent DP models appeared in Cifarelli and Regazzini (1978), who defined dependence across related random measures $\{G_x\}$ by introducing a regression for the base-line measure G_x^* of marginally DP-distributed random

measures, $G_x \sim \text{DP}(M, G_x^*)$. The model was used, for example, in Muliere and Petrone (1993), who defined dependent non-parametric models $G_x \sim \text{DP}(M, G_x^*)$ by assuming a regression in the base-line measure $G_x^* = N(\beta x, \tau^2)$. Comparing with the notation in model (1), the hyperparameters here are $\psi = (M, \beta, \tau)$. Similar models are discussed in Mira and Petrone (1996) and Giudici *et al.* (2003).

Linking the related non-parametric models through a regression on the parameters of the non-parametric models limits the nature of the dependence to the structure of this regression. MacEachern (1999) proposed the dependent DP (DDP) as an alternative approach to define a dependent prior model for a set of random measures $\{G_x\}$, with $G_x \sim \text{DP}$ marginally. Recall Sethuraman's stick breaking representation (6) for the DP random measure,

$$G_x = \sum_h \omega_{xh} \delta_{\mu_{xh}}.$$

The key idea behind the DDP is to introduce dependence across the measures G_x by assuming the distribution of the point masses μ_{xh} to be dependent across different levels of x , but still independent across h . In the basic version of the DDP the weights are assumed to be the same across x , i.e. $\omega_{xh} = \omega_h$. To introduce dependence of μ_{xh} across x MacEachern (1999) used a Gaussian process. An application to spatial modelling is further developed in Gelfand *et al.* (2005) by allowing the locations θ to be drawn from a random field (a Gaussian process). The same method to induce dependence is used in De Iorio *et al.* (2004) to achieve an analysis-of-variance type of structure on μ_{xh} across x . Griffin and Steel (2006) introduced dependence in non-parametric distributions by making the weights in the Sethuraman representation dependent on the covariates. We chose to fix the weights w_{xh} across covariates and to introduce the dependence through the point mass locations μ_{xh} , mainly because of computational simplicity.

The construction that was introduced in De Iorio *et al.* (2004) is a natural approach to introduce DDP measures to implement model (1). Specifically, let $d'_i = (1, 0)$ if $x_i = 0$ and $d'_i = (1, 1)$ if $x_i = 1$. We assume that

$$\theta_i \sim G_x^M(\theta_i), \quad \text{with } G_x^M(\theta) = \int \mathcal{N}(\theta | \alpha d_i, \tau^2) dG(\alpha), \quad G \sim \text{DP}(M, G_\psi^*). \quad (8)$$

Here ψ indicates hyperparameters in the definition of the base measure. In words, the trick to construct dependent random measures G_x^M is to start with a random measure on the coefficients α . Depending on x_i , a design vector d_i selects a linear function of the α . Finally, using an additional convolution with a normal kernel we define continuous and dependent random measures G_x^M . Introducing latent variables α_i , model (8) can be equivalently rewritten as a hierarchical model:

$$\theta_i = \alpha_i d_i + \eta_i, \quad \alpha_i \sim G, \quad G \sim \text{DP}(M, G_\psi^*), \quad (9)$$

with $\eta_i \sim \mathcal{N}(0, \tau^2)$. Let p denote the dimension of θ_i . The latent variable α_i is a $p \times 2$ random matrix. The first column α_{i0} is the random-effects vector for a patient from group $x = 0$. The second column α_{i1} is the offset to generate a random effect for a patient from group $x = 1$. The modelling strategy proposed implies that the α_{i0} -parameters are estimated from data coming from both groups. At the same time, we can learn about possible dependences between α_{i0} and α_{i1} that may be group specific. Learning about such features is not possible

with alternative models involving *a priori* independent non-parametric models, e.g. two independent DPs. Under a model with two independent DPs we would learn only about α_{i0} for patients from the group $x_i = 0$, and about $\alpha_{i0} + \alpha_{i1}$ for patients from the group $x_i = 1$. Inference about the dependence of α_{i0} and α_{i1} for future patients would not be possible. Later, in Section 5.2 and Table 2, we shall show how in the example the increased borrowing of strength in the dependent model leads to a small improvement in the misclassification rate, from 21.9% to 19.6%. The model is completed with a sampling model for y_i , $y_i \sim p(y_i|\theta_i)$, a marginal prior, $\Pr(x_i = 1) = \pi_1$, for x_i , and hyperpriors on unknown hyperparameters, including τ^2 , M , ψ and π_1 . See Section 5 for an example of specific choices in an application.

The equivalent hierarchical model (9) highlights the nature of the model as a DPM model, allowing the use of any of the posterior MCMC simulation methods that are proposed for such models. Compared with MCMC sampling for DPM models, as summarized, for example, in MacEachern and Müller (2000), the only additional step is the imputation of the latent group indicators x_i . We briefly summarize key features of the MCMC implementation. The discrete nature of the DP random measure G implies a positive probability for ties among the latent quantities α_i in model (9). The configuration of ties determines many details of the MCMC method. Let k denote the number of distinct values among $\{\alpha_i, i = 1, \dots, m\}$ and let $\{\alpha_j^*, j = 1, \dots, k, j = 1, \dots, k\}$ denote such values. Recall from the discussion after equation (6) that $\alpha_j^* \sim G^*$, independent and identically distributed. We define configuration indicators s_i with $s_i = j$ if and only if $\alpha_i = \alpha_j^*$. The unique values $(\alpha_1^*, \dots, \alpha_k^*)$ and configuration indicators s together provide an alternative representation of $(\alpha_1, \dots, \alpha_m)$. The marginal prior for $(s_1, \dots, s_m|G^*, M)$, marginalizing in particular with respect to the random probability measure G , can easily be described. It is known as the Polya urn scheme (Blackwell and MacQueen, 1973). This fact greatly simplifies posterior MCMC simulation for the DPM models, such as model (9) together with equation (3) and the sampling model (2). We outline the transition probabilities that are used in the MCMC implementation. We use the notation $[x|y, z]$ to indicate that the parameter x is updated conditionally on currently imputed values for y and z . We use Y to denote all data generically, θ to indicate the set of all θ_i and $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_m)$, etc. Also, $\varphi = (\tau^2, \beta_{1x}, \beta_{2x}, \sigma_x^2, x=0, 1)$ and $\psi = (\xi, R, M)$. Each iteration of the MCMC algorithm consists of the transition probabilities $[\alpha_j^*|s, \theta, \psi, \varphi]$, $[\alpha_i|\theta_i, \alpha_{-i}, \psi, \varphi]$, $[\theta_i|\alpha_i, \varphi, Y]$, $[\beta|\theta, \varphi, Y]$, $[x_{m+1}|\dots]$, and transition probabilities to change the remaining parameters σ_x^2 , τ^2 , ξ , R and M .

Note that changes in the sampling model (2) would not affect the transition probabilities for α_i and the parameters that are specific to the DP model. Updating θ_i proceeds like inference in a fully parametric model with sampling model (2) and normal prior $\theta_i \sim N(\alpha_i d_i, \tau^2)$. In other words, the computational effort that is related to the longitudinal model is the same as in a fully parametric model.

5. Application

5.1. Model specification

We apply the model proposed to the analysis of the longitudinal β -HCG data. Mean values of the $\log(\beta\text{-HCG})$ for the 173 women show a non-linear relationship with days of pregnancy. Fig. 1 shows time profiles for normal and abnormal pregnancies. The analysis in Marshall and Barón (2000) suggests that woman-to-woman variation is adequately accounted for by the introduction of random effects to model the asymptotic behaviour of the $\log(\beta\text{-HCG})$ level (θ_i below). They proposed the following non-linear random-effects model. Recall that $y_i = (y_{i1}, \dots, y_{in_i})'$ are the observed $\log(\beta\text{-HCG})$ measurements at

occasions $t_i = (t_{i1}, \dots, t_{in_i})'$ for woman $i = 1, \dots, m = 173$, and $x = 0$ and $x = 1$ indicate respectively normal and abnormal pregnancy groups.

$$y_i | (x_i = x) \sim \mathcal{N}(\mu_{ix}, \sigma_x^2 \mathbf{I}_{n_i}), \quad \text{with } \mu_{ix} = \theta_i [1 + \exp\{-(t_i - \beta_{1x})/\beta_{2x}\}]^{-1}. \quad (10)$$

Here θ_i is a scalar subject-specific random effect, and $\beta_x = (\beta_{1x}, \beta_{2x})$, $x = 0, 1$, are bivariate fixed effects for the abnormal and the normal group respectively. In model (10), the vector $(\theta_i, \beta_{1x}, \beta_{2x})$ characterizes the profile for the i th woman in group x . Marshall and Barón (2000) and De la Cruz-Mesía and Quintana (2006) assumed that $\theta_i \sim N(\theta_x, \tau_x^2)$.

A simple parametric model with a normal random-effects distribution is adequate to describe subject-specific profiles and to fit smooth profiles to observed data. However, detailed features of the random-effects model can critically change the predictive classification probabilities for patients with random effects that are imputed away from the centre of the estimated random-effects distributions. This leads us to consider the semiparametric model (8), or its equivalent version (9). We shall later compare the proposed non-parametric inference with a comparable parametric model and show how the non-parametric extension changes critical predictions.

In model (9) we assume for the base-line distribution G^* a two-dimensional normal distribution. Specifically, we take $G^* = \mathcal{N}_2(\xi, R)$. To complete the model specification, we assume independent hyperpriors

$$\begin{aligned} \beta_x &\sim \mathcal{N}_2(\beta_{0x}, B_{0x}), \quad \sigma_x^2 \sim \mathcal{IG}(a_{0x}, b_{0x}), \quad \tau^2 \sim \mathcal{IG}(c_0, d_0), \\ \xi &\sim \mathcal{N}_2(\xi_0, \Sigma_\xi), \quad R^{-1} \sim \text{Wishart}_2\{q, (qR_0)^{-1}\}. \end{aligned} \quad (11)$$

Here, $\mathcal{IG}(a, b)$ denotes the inverse gamma distribution, parameterized to have mean $1/\{b(a-1)\}$. The Wishart prior on R^{-1} is parameterized such that $E(R^{-1}) = R_0^{-1}$. The first parameter of the Wishart distribution is the scalar parameter; the second is the matrix parameter.

The implementation of model (9) requires adopting specific values for $M, a_{0x}, b_{0x}, c_0, d_0, q, \beta_{0x}, \xi_0, B_{0x}, \Sigma_\xi$ and R_0 . The parameter M of the DP prior $\text{DP}(M, G^*)$ controls how close a realization of the process is to the base-line distribution G^* . Additionally, in the DPM model, M controls the distribution of the number of distinct elements of the vector $(a_1, \dots, a_m)$ and hence the number of distinct components of the mixture (see Antoniak (1974) and De Iorio *et al.* (2004) for more details). Treating M as an unknown hyperparameter and assuming a gamma prior, $M \sim \mathcal{G}(a_M, b_M)$, Escobar and West (1995) derived an efficient posterior sampling scheme for M . We follow this approach, using $a_M = b_M = 1$.

The values of the other hyperparameters in distributions (11) were taken as $\beta_{00} = \beta_{01} = \xi_0 = (0, 0)'$, $B_{00} = B_{01} = \Sigma_\xi = 10000\mathbf{I}_2$, $q = 3$, $R_0 = \mathbf{I}_2$, $a_{00} = a_{01} = c_0 = 3$ and $b_{00} = b_{01} = d_0 = 0.01$. These choices imply a prior mean variance of σ_k^2 and τ^2 equal to 2500. Here $\mathbf{I}_2$ is the 2×2 identity matrix. Prior probabilities of group membership were assumed to be proportional to the size of the groups in the training sample. We also performed the analysis with different hyperparameter values, obtaining very similar results. This suggests robustness to the choices of hyperparameter.

Updating the latent mixture parameters a_i and the hyperparameters $\beta_x, \sigma_x^2, \tau^2$ and M proceeds with standard posterior simulation methods for DPMs. See, for example,

MacEachern and Müller (1998) and De Iorio *et al.* (2004) for a full description of the Gibbs sampling scheme.

The full conditionals for implementing the Gibbs sampler are not available in closed form for β_{1x} and β_{2x} . To update β_{1x} and β_{2x} we thus use a Metropolis–Hastings step with a normal approximation to the full conditional as the candidate distribution. Resampling M is done by introducing a latent beta-distributed variable, as described by Escobar and West (1995), based on West (1992).

To perform the Gibbs sampling, we chose starting-points in a neighbourhood of the maximum likelihood estimates of model parameters. In theory the Markov chain convergence and ergodic properties are independent of the initial values. In practice, however, a good choice of starting-points shortens the number of iterations that are required until practical convergence. We generated 100000 iterations. After 10000 iterations, samples were collected, at a spacing of 90 iterations, to obtain approximately independent samples, leaving us with a total of $B = 1000$ posterior Monte Carlo samples for calculating posterior quantities of interest.

To diagnose convergence, we used methods that are available in the BOA package (Smith, 2004). Because of the high dimensional parameter vector, we prefer to use diagnostics, such as those proposed by Geweke (1992), which do not require multiple parallel chains.

5.2. Results

Figs 2(a) and 2(b) show histograms of the subject-specific parameters θ_i estimated by using the empirical Bayes methods as implemented in the SAS system. Specifically, they show the posterior means of θ_i conditional on all the other hyperparameters being evaluated at their maximum likelihood estimates for model (10) with normally distributed random effects (see, for example, Vonesh and Chinchilli (1997)). Figs 2(c) and 2(d) show posterior predictive draws under the Bayesian semiparametric approach BSP for both abnormal and normal groups. In Figs 2(c) and 2(d) a smooth curve shows the posterior estimated random-effects distribution $\overline{G^M}(\theta) = E\{G^M(\theta)|y^m\}$. For comparison, Figs 2(a) and 2(b) show a kernel density estimate based on the histogram of the corresponding estimates. To evaluate the posterior mean $\overline{G^M}$ we exploit the identity $\overline{G^M}(\theta) = p(\theta_{m+1}|y^m)$, which follows from

$$\begin{aligned} p(\theta_{m+1}|y^m, x_{m+1}=x) &= \int p(\theta_{m+1}|G_x^M) dp(G_x^M|y^m) \\ &= \int G_x^M(\theta_{m+1}) dp(G_x^M) = \overline{G^M}. \end{aligned}$$

We can, therefore, approximate $\overline{G^M}$ by a kernel density estimate of posterior predictive draws, $\theta_{m+1} \sim p(\theta_{m+1}|y^m)$. The maximum likelihood estimates show asymmetry in the normal group and bimodality in the abnormal group. A non-parametric specification of the distribution of the random effects allows for the flexibility to estimate such features. See Figs 2(c) and 2(d).

The parameter M induces a distribution on the number of clusters into which the observations fall. Recall the definition of configuration indicators s_i in the discussion following model (9). We refer to sets of observations with equal configuration indicators, i.e. a common value α_i , as clusters. The DDP model that we use to implement inference in this paper relies on a single mass parameter M . For this model, clusters of observations occur both within and across groups. The number of clusters is stochastically increasing with the number of observations (see De Iorio *et al.* (2004)). Recall that k was defined as the

number of clusters. Let k_x denote the number of clusters of observations in group x . We find the posterior mean $E(k_x|y^m)$ (with standard deviations $SD(k_x|y^m)$ in parentheses) to be 5.9 (1.8) and 5.2 (1.5) for $x = 0$ and $x = 1$ respectively. The posterior mean $E(k|y^m)$ (with standard deviation $SD(k|y^m)$ in parentheses) is 6.3 (1.9).

As part of the analysis we estimated individual β -HCG profiles and standard errors. These profiles can be used to assess the goodness of fit. Fitted profiles with ± 2 posterior standard deviations curves are displayed for six selected patients in Fig. 3: three of them in the normal group (patients 2, 66 and 75) and the remaining three in the abnormal group (patients 15, 29 and 45). On the basis of these plots we informally assess the goodness of fit of the model to the data. The posterior inference captures the varying observation error between subjects.

Next, we consider the problem of evaluating the classification rule. This is naturally carried out through an estimate of the associated misclassification rates. At this point we could apply the rule to the observed data and count the (relative) frequencies of misclassified observations. In doing so, we conclude that the BSP yields the best results (the data are not shown). However, it can be argued that this yields overly optimistic misclassification error rates as the same observations are used to determine and to evaluate the classification rule (McLachlan, 2004). Another traditional approach is cross-validation (Lachenbruch and Mickey, 1968). It computes the classification rule by leaving out one subject at a time and records whether this observation is correctly classified or not. Here, we classify an observation as abnormal if the posterior probability for $x_i = 0$ is greater than $\frac{1}{2}$. Table 1 presents the results by using cross-validation based on the method that was described in Section 4, the Bayesian parametric approach developed by De la Cruz-Mesía and Quintana (2006) and a frequentist method that was developed by Marshall and Barón (2000). We found interesting differences between the three approaches. The misclassification rate under the BSP model is 14.5% ($25 = 173$), which is less than under the Bayesian parametric (BP) model and the method based on maximum likelihood estimation, 17.3% ($30/173$) and 18.5% ($32/173$) respectively. A traditional way to summarize these results is a receiver operating characteristic curve, which plots the true positive rate against the false positive rate for the various possible cut points of the classification rule (0.5 was used when calculating the results that are displayed in Table 1). Fig. 4 shows this curve for both Bayesian models. We see how the BSP model improves on the BP method (higher area under the curve). To understand further the corresponding classification, Fig. 5 shows estimated classification probabilities for all 173 women, arranged by true x_i , and within each group sorted in decreasing order. We see how the BSP model dominates the BP model for most of the range, in the sense of implying higher and lower probabilities for normal and abnormal pregnancies respectively. The most noticeable exception is the rightmost part for the abnormal cases (lowest classification probabilities), where this trend is reversed. But this is of little concern, as at that range of values for the probabilities almost any rule would classify these women as abnormal.

The reported receiver operating characteristic curve provides a conservative comparison in the following sense. It is based on classification of patients with complete data recorded over the first 80 days of gestational age. More important for an informed clinical treatment decision are differences in early prediction, based on early responses only. To illustrate this use, we generate from the posterior predictive distribution data for one future patient for each group and evaluate expression (4) for up to five possible observations. Fig. 6 shows how the classification probabilities change as we accrue more data. Fig. 6 compares inference under the proposed semiparametric model and a corresponding parametric model fixing the random probability measure G at the base measure G^* . For the normal pregnancy patient, we observe a steady growth of the probabilities. In contrast, for the abnormal

pregnancy patient this probability first increases and starts to decrease to values that leave no question about the classification. A possible explanation for this is the rather heterogeneous patterns that are found for abnormal pregnancy patients. Indeed, many of these show an initial increase in the $\log_{10}(\beta\text{-HCG})$ responses (just as all the normal pregnancy patients do) followed by a decrease in some of the patients. Thus the classification probabilities for abnormal pregnancy patients require a few more observations than the normal ones to reflect the correct outcome. For the abnormal pregnancy patient the predictive classification probabilities by using the BSP model decrease more rapidly than under the BP model. After two observations we find a difference greater than 10% in predictive probabilities. From a clinical perspective, a 10% difference in predicted probabilities can be key to making the right treatment decision at this critical early time. The receiver operating characteristic curve that is shown in Fig. 4 evaluates classification based on profiles over the entire observation period. As shown in Fig. 6, the improvement for classification based on the first two or three observations is even larger.

To assess the model fit and to compare different models, we calculate the conditional predictive ordinate (CPO) (Gelfand *et al.*, 1995) for each observation. Chen *et al.* (2000) showed in detail how to obtain Monte Carlo estimates of the CPO statistics. We can compare different models by using sums of $\log(\text{CPOs})$ of the individual observations.

Define $\widehat{\text{CPO}}_i$ to be the Monte Carlo estimates of the i th subject's CPO statistic. Greater values of $S = \sum \log(\widehat{\text{CPO}}_i)$ indicate a better fit. We found $S = -117.2$ for the BSP model. For the BP model we found $S = -124.1$. The difference suggests that the BSP model provides a marginally better fit to the $\log(\beta\text{-HCG})$ data than its parametric counterpart.

We next investigate the effect of the dependence that is introduced in the DDP compared with a model with two independent DPMs. Fig. 7 displays the results of 500 posterior predictive draws from the bivariate distribution $p(a_{m+1}|Y)$. We can identify two large clusters, each suggesting negative correlation between main effect and abnormal pregnancy offset parameters. The resulting covariance structure clearly differs between these components. Note that such findings would not be possible under a model with two independent DPs. To compare our model with that defined by two independent DPMs we changed model (8) by using $d'_i = (1, 0)$ and $d'_i = (0, 1)$ or $x_i = 0$ and $x_i = 1$ respectively. We refer to the new model as model iBSP. For a fair comparison we use the same choices of hyperparameter as before, implying in particular that the marginal probability models for the random-effects distributions G_x^M , $x = 0, 1$, remain unchanged under the BSP model and model iBSP. We carried out the same inference as described in Fig. 6, focusing on the classification for a future woman, $m+1$, after the first $n_{m+1} = 2$ observations, assuming that the unknown truth is $x_{m+1} = 0$, i.e. an abnormal pregnancy. Fig. 6 reports the classification probabilities $\Pr(x_{m+1} = 1|y_{m+1,1}, y_{m+1,2}, y^m) = 50\%$ for the proposed BSP model, and 63% for the BP model. For model iBSP we find a probability of 55%, justifying the minor additional effort to implement the DDP model. However, this depends on a single patient, as just described. We investigated this issue further, considering the classification of *every* patient, on the basis of only the first two observations, and assuming the same proportion of normal pregnancies as was empirically observed. This is essentially equivalent to a cross-validation of the inference for all patients. Table 2 summarizes the classification as normal or abnormal under each of the competing models. The reported misclassification rates show an improvement under the dependent model compared with the independent model.

Finally, we investigate the effect of choices of hyperparameter on the reported inference. Again, consider the predicted classification for a future woman (assuming the unknown truth to be $x_{m+1} = 0$) based on one or two observations. Table 3 shows these probabilities for

various combinations of M , $E(\tau^2)$ and $E(\sigma_x^2)$. The corresponding probabilities do exhibit some variation, but the implied classifications remain unchanged. In fact, the estimated error rates are the same as reported in Table 1 in all cases (the data are not shown).

6. Discussion

We have proposed a model-based approach to classification of longitudinal profiles. The underlying models in each group or population are given by non-linear semiparametric models. Flexibility for relaxing the distributional assumptions is introduced by using a non-parametric specification on the random-effects models. Dependence in the growth curves is introduced through a design vector indicating group membership and selecting appropriate features of a common underlying random probability measure. The approach is appropriate for classifying longitudinal profiles of data sets with unbalanced data structure. It uses all available information for classifying subjects over time, regardless of the number or timing of the observations. Moreover, the influence on discrimination of both the between-group and within-group components variability can be readily quantified, and the posterior simulation scheme is straightforwardly implemented. The approach is particularly appropriate for decision-making in clinical practice where the number and times of observations are often arbitrary and depend on the progression of the patient.

A key feature of our approach is the flexibility that is provided by the non-parametric model for random effects. A straightforward generalization of our approach could accommodate more information that is available. This can be done by inclusion of more covariates or by considering other markers, thus extending the framework to a multivariate one.

Limitations of the proposed model are the reliance on posterior simulation and the nature of the non-parametric generalization. Although posterior simulation is straightforward, it requires the development of some problem-specific software. The non-parametric modelling is on the random-effects distribution only, but it still requires the user to choose a parametric model for $p(y_i|\theta_i)$. Alternative models could use the available data from patients $i = 1, \dots, m$, to learn about the nature of the longitudinal dependence, using, for example, methods that were reviewed in Denison *et al.* (2002). Another possible extension is the use of problem-specific decision rules. In the inference reported we classified patients by maximum posterior predictive probability of group membership. Alternatively, one could imagine an approach that takes into account the sequential nature of the decision problem. It is conceivable that even with high probability of abnormal pregnancy a clinician might decide to wait for one more measurement, trading off the additional information with a possible loss in treatment options.

Finally model (8) and (9) allows an easy generalization to more general non-parametric priors on G . In particular, we can easily replace the DP model by a species sampling model (Pitman, 1996), which allows more general prior distributions on configurations of the α -parameters. See further discussion of such models in Ishwaran and James (2003) and in Quintana (2006).

Acknowledgments

We thank Professor Guillermo Marshall for giving us access to his β -HCG data set. The first author thanks the Comisión Nacional de Investigación Científica y Tecnológica for partially supporting his doctoral studies at the Pontificia Universidad Católica de Chile. Part of the research was done while the first author was visiting the Department of Biostatistics, the University of Texas, M. D. Anderson Cancer Center, under grant MECESUP PUC 0103. The research was supported by National Institutes of Health–National Cancer Institute grant 2 R01 CA75981-04A1 and by Fondo Nacional de Desarrollo Científico y Tecnológico grants 1020712, 1060729 and 3060071. We also thank the Joint Editor, Associate Editor and the reviewers for their constructive comments, which helped to improve this manuscript substantially.

References

- Albert A. Discriminant analysis based on multivariate response curve: a descriptive approach to dynamic allocation. *Statist Med.* 1983; 2:95–106.
- Albert JM, Kshirsagar AM. The reduced-rank growth curve model for discriminant analysis of longitudinal data. *Aust J Statist.* 1993; 35:345–357.
- Antoniak CE. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Ann Statist.* 1974; 2:1152–1174.
- Azen SP, Afifi AA. Two models for assessing prognosis on the basis of successive observations. *Math Biosci.* 1972; 14:169–176.
- Blackwell D, MacQueen JB. Ferguson distributions via Pólya urn schemes. *Ann Statist.* 1973; 1:353–355.
- Brown PJ, Kenward MG, Bassett EE. Bayesian discrimination with longitudinal data. *Biostatistics.* 2001; 2:417–432. [PubMed: 12933633]
- Bush CA, MacEachern SN. A semiparametric Bayesian model for randomized block designs. *Biometrika.* 1996; 83:275–285.
- Chen, MH.; Shao, QM.; Ibrahim, JG. Monte Carlo Methods in Bayesian Computation. Vol. ch. 10. New York: Springer; 2000.
- Cifarelli DM, Regazzini E. Nonparametric statistical problems under partial exchangeability: the use of associative means (in Italian). *Ann Inst Mat Finian Univ Tor III.* 1978; 12:1–36.
- Davidian M, Gallant AR. The nonlinear mixed effects model with a smooth random effects density. *Biometrika.* 1993; 80:475–488.
- De la Cruz-Mesía R, Quintana F. A model-based approach to Bayesian classification with applications to predicting pregnancy outcomes from longitudinal β -hCG profiles. *Biostatistics.* 2006 to be published.
- De Iorio M, Müller P, Rosner GL, MacEachern SN. An ANOVA model for dependent random measures. *J Am Statist Ass.* 2004; 99:205–215.
- Denison, D.; Holmes, C.; Mallick, B.; Smith, A. Bayesian Methods for Nonlinear Classification and Regression. New York: Wiley; 2002.
- Escobar MD, West M. Bayesian density estimation and inference using mixtures. *J Am Statist Ass.* 1995; 90:577–588.
- Fieuws, S.; Verbeke, G.; Brant, LJ. Technical Report. Biostatistical Centre, Katholieke Universiteit Leuven; Leuven: 2003. Classification of longitudinal profiles using nonlinear mixed-effects models.
- France JT, Keelan J, Song L, Liddell H, Zanderigo A, Knox B. Serum concentrations of human chorionic gonadotrophin and immunoreactive inhibin in early pregnancy and recurrent miscarriage: a longitudinal study. *Aust New Zeal J Obstet Gyn.* 1996; 36:325–330.
- Gelfand, AE.; Dey, DK.; Chang, H. Model determination using predictive distributions with implementation via sampling-based methods (with discussion). In: Bernardo, JM.; Berger, JO.; Dawid, AP.; Smith, AFM., editors. *Bayesian Statistics. Vol. 4.* Oxford: Oxford University Press; 1995.
- Gelfand AE, Kottas A, MacEachern SN. Bayesian nonparametric spatial modelling with Dirichlet processes mixing. *J Am Statist Ass.* 2005; 100:1021–1035.
- Geweke, J. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In: Bernardo, JM.; Berger, JO.; Dawid, AP.; Smith, AFM., editors. *Bayesian Statistics. Vol. 4.* Oxford: Oxford University Press; 1992. p. 169-194.
- Giudici P, Mezzetti M, Muliere P. Mixtures of Dirichlet process priors for variable selection in survival analysis. *J Statist Planng Inf.* 2003; 111:101–115.
- Griffin JE, Steel MFJ. Order-based dependent Dirichlet processes. *J Am Statist Ass.* 2006; 101:179–194.
- Inoue LYT, Etzioni R, Slate EH, Morrell C, Penson DF. Combining longitudinal studies of PSA. *Biostatistics.* 2004; 5:483–500. [PubMed: 15208207]

- Ishwaran H, James L. Gibbs sampling methods for stick-breaking priors. *J Am Statist Ass.* 2001; 96:161–173.
- Ishwaran H, James LJ. Generalized weighted Chinese restaurant processes for species sampling mixture models. *Statist Sin.* 2003; 13:1211–1235.
- Ishwaran H, Takahara G. Independent and identically distributed Monte Carlo algorithms for semiparametric linear mixed models. *J Am Statist Ass.* 2002; 97:1154–1166.
- Jain S, Neal RM. A split-merge Markov Chain Monte Carlo procedure for the Dirichlet process mixture model. *J Computnl Graph Statist.* 2004; 13:158–182.
- Kleinman KP, Ibrahim JG. A semi-parametric Bayesian approach to the random effects model. *Biometrics.* 1998; 54:265–278. [PubMed: 9544521]
- Lachenbruch PA, Mickey MR. Estimation of error rates in discriminant analysis. *Technometrics.* 1968; 10:1–11.
- MacEachern SN. Dependent nonparametric process. *Proc Bayes Statist Sci Sect Am Statist Ass.* 1999
- MacEachern SN, Müller P. Estimating mixture of Dirichlet process models. *J Computnl Graph Statist.* 1998; 7:223–238.
- MacEachern, SN.; Müller, P. Efficient MCMC schemes for robust model extensions using encompassing Dirichlet process mixture models. In: Ríos-Insua, D.; Ruggeri, F., editors. *Robust Bayesian Analysis*. New York: Springer; 2000. p. 295–315.
- Marshall G, Barón AE. Linear discriminant models for unbalanced longitudinal data. *Statist Med.* 2000; 19:1969–1981.
- McLachlan, GJ. *Discriminant Analysis and Statistical Pattern Recognition*. New York: Wiley; 2004.
- Mentré F, Mallet A. Handling covariates in population pharmacokinetics. *Int J Biomed Comput.* 1994; 36:25–33. [PubMed: 7927857]
- Mira, A.; Petrone, S. Bayesian hierarchical nonparametric inference for change-point problems. In: Bernardo, JM.; Berger, JO.; Dawid, AP.; Smith, AFM., editors. *Bayesian Statistics*. Vol. 5. Oxford: Oxford University Press; 1996.
- Morrell CH, Pearson JD, Carter HB, Brant LJ. Estimating unknown transition times using a piecewise non-linear mixed effects model in men with prostate cancer. *J Am Statist Ass.* 1995; 90:45–53.
- Muliere P, Petrone S. A Bayesian predictive approach to sequential search for an optimal dose: parametric and nonparametric models. *J Ital Statist Soc.* 1993; 2:349–364.
- Müller P, Quintana F, Rosner G. A method for combining inference across related nonparametric Bayesian models. *J R Statist Soc B.* 2004; 66:735–749.
- Müller P, Rosner GL. A Bayesian population model with hierarchical mixture priors applied to blood count data. *J Am Statist Ass.* 1997; 92:1279–1292.
- Neal RM. Markov chain sampling methods for Dirichlet process mixture models. *J Computnl Graph Statist.* 2000; 9:249–265.
- Pitman, J. Some developments of the Blackwell-MacQueen urn scheme. In: Ferguson, TS.; Shapeley, LS.; MacQueen, JB., editors. *Statistics, Probability and Game Theory: Papers in Honor of David Blackwell*. Hayward: Institute of Mathematical Statistics; 1996. p. 245–268.
- Quintana FA. A predictive view of Bayesian clustering. *J Statist Planng Inf.* 2006; 136:2407–2429.
- Sethuraman J. A constructive definition of Dirichlet priors. *Statist Sin.* 1994; 4:639–650.
- Smith, BJ. Bayesian Output Analysis Program (BOA), version 1.1.2 for S-PLUS and R. College of Public Health, University of Iowa; Iowa City: 2004. (Available from <http://www.public-health.uiowa.edu/boa>.)
- Tomasko L, Helms RW, Snapinn SM. A discriminant analysis extension to mixed models. *Statist Med.* 1999; 18:1249–1260.
- Vonesh, EF.; Chinchilli, VM. *Linear and Nonlinear Models for the Analysis of Repeated Measurements*. New York: Dekker; 1997.
- Walker SG, Wakefield JC. Population models with a nonparametric random coefficient distribution. *Sankhya B.* 1998; 60:196–212.
- West, M. Technical Report 92–A03. Institute of Statistics and Decision Sciences, Duke University; Durham: 1992. Hyperparameter estimation in Dirichlet process mixture models.

Zhang D, Davidian M. Linear mixed models with flexible distributions of random effects for longitudinal data. *Biometrics*. 2001; 57:795–802. [PubMed: 11550930]

Fig. 1.
Observed profiles of β -HCG for all 173 women: (a) normal pregnancies; (b) abnormal pregnancies

Fig. 2.

Estimated subject-specific (θ_i) parameters by using (a), (b) maximum likelihood estimation and (c), (d) posterior predictive draws for the BSP model with a smooth curve overlaid on each plot: (a) normal group; (b) abnormal group; (c) normal group; (d) abnormal group

Fig. 3.

Fitted curves for (a)–(c) three patients in the normal group and (d)–(f) three in the abnormal group (O, actual observations; —, fitted curves; ---, ± 2 posterior standard deviations): (a) patient 2; (b) patient 15; (c) patient 29; (d) patient 45; (e) patient 66; (f) patient 75

Fig. 4. Receiver operating characteristic curves for classification under the BSP model (—) and the BP model (-----)

Fig. 5. Estimated probabilities of classification under the BSP model (—) and the BP model (-----) in the normal group for all individuals in decreasing order within the normal and abnormal groups

Fig. 6.

Evolution of classification probabilities for one normal (○) and one abnormal (×) future patient as a function of the number of observations. (——, BSP model; -----, BP model); the difference in inference is critical; for example, consider inference after the second observation—if the (unknown) truth is an abnormal pregnancy, the probability of a correct classification under the model proposed is 10% higher under a parametric model

Fig. 7.
Scatterplot of 500 posterior predictive draws of analysis-of-variance coefficients α_{m+1} for a future patient

Table 1

Classification results by using Bayesian parametric (BP), semiparametric (BSP) and classical methods

Group	Classifications by the following methods and groups:						Total
	Classical		BP		BSP		
	Normal	Abnormal	Normal	Abnormal	Normal	Abnormal	
Normal	113	11	115	9	117	7	124
Abnormal	21	28	21	28	18	31	49
	134	39	136	37	134	39	173

Table 2
Cross-validation using the first two observations for the DDP and DP models

Group	Classifications by the following methods and groups:						Total
	DDP		DP				
	Normal	Abnormal	Normal	Abnormal			
Normal	118	6	116	8		124	
Abnormal	28	21	30	19		49	
	146	27	146	27		173	

Table 3

Effect of choices of hyperparameter on the classification probabilities[†]

$E(\sigma_x^2)$		$E(\tau^2)$ Probabilities for the following values of M :			
		1*	5	10	
5	5	0.819	0.816	0.820	p_1
		0.954	0.950	0.950	p_2
	50	0.822	0.818	0.818	p_1
		0.955	0.964	0.955	p_2
	500	0.818	0.816	0.814	p_1
		0.955	0.961	0.956	p_2
50	5	0.822	0.821	0.817	p_1
		0.955	0.954	0.951	p_2
	50	0.824	0.818	0.814	p_1
		0.957	0.961	0.959	p_2
	500	0.819	0.819	0.814	p_1
		0.954	0.955	0.958	p_2
500	5	0.819	0.813	0.815	p_1
		0.954	0.942	0.947	p_2
	50	0.824	0.818	0.819	p_1
		0.959	0.958	0.957	p_2
	500	0.818	0.807	0.804	p_1
		0.955	0.949	0.947	p_2

[†]We report, for various combinations of M , $E(\sigma_x^2)$ and $E(\tau^2)$, the classification probability for a normal pregnancy patient (i.e. the unknown truth is $x_{M+1} = 0$) given one (p_1) and two (p_2) observations. Here, 1* denotes that for this case $M \sim \mathcal{G}(1, 1)$.