Published in final edited form as:

J Stat Theory Pract. 2013 April 1; 7(2): . doi:10.1080/15598608.2013.772811.

Effect on Prediction when Modeling Covariates in Bayesian Nonparametric Models

Alejandro Cruz-Marcelo,

Capital One, Dallas TX 77030

Gary L. Rosner,

Division of Oncology Biostatistics & Bioinformatics, Sidney Kimmel Comprehensive Cancer Center at Johns Hopkins, Baltimore MD 21205

Peter Müller, and

Department of Mathematics, The University of Texas, Austin TX 78712

Clinton F. Stewart

Pharmaceutical Sciences, St. Jude Children's Research Hospital, Memphis TN 38105

Alejandro Cruz-Marcelo: alejandro.cruzmarcelo@gmail.com; Gary L. Rosner: grosner1@jhmi.edu; Peter Müller: pmueller@math.utexas.edu; Clinton F. Stewart: Clinton.Stewart@stjude.org

Abstract

In biomedical research, it is often of interest to characterize biologic processes giving rise to observations and to make predictions of future observations. Bayesian nonparametric methods provide a means for carrying out Bayesian inference making as few assumptions about restrictive parametric models as possible. There are several proposals in the literature for extending Bayesian nonparametric models to include dependence on covariates. Limited attention, however, has been directed to the following two aspects. In this article, we examine the effect on fitting and predictive performance of incorporating covariates in a class of Bayesian nonparametric models by one of two primary ways: either in the weights or in the locations of a discrete random probability measure. We show that different strategies for incorporating continuous covariates in Bayesian nonparametric models can result in big differences when used for prediction, even though they lead to otherwise similar posterior inferences. When one needs the predictive density, as in optimal design, and this density is a mixture, it is better to make the weights depend on the covariates. We demonstrate these points via a simulated data example and in an application in which one wants to determine the optimal dose of an anticancer drug used in pediatric oncology.

Keywords

Covariates modeling; Dependent Dirichlet process; Dirichlet process mixture; Hierarchical model; Nonparametric Bayes; Predictive distribution

1 Introduction

A modeling framework that has become increasingly popular is the use of Bayesian nonparametric methods (Müller et al. 2004). The Dirichlet process (DP) (Ferguson 1973) is a particularly popular prior model for an unknown random measure. The popularity of the DP is due to its elegance, simplicity, and the existence of computationally efficient Markov chain Monte Carlo (MCMC) posterior simulation algorithms (MacEachern and Müller 1998). In addition, hierarchical models based on random effects distributions with DP priors are able to accommodate outliers, multimodality, and dependence in multivariate, longitudinal, and functional data (Dunson 2009). Examples of nonparametric methods based

on the DP include applications in pharmacokinetics (Rosner and Müller 1997), econometrics (Griffin and Steel 2004), spatial modelling (Gelfand et al. 2005), meta-analysis (Burr and Doss 2005), variable selection (Kim et al. 2006), genetics (Xing et al. 2007), density estimation (Dunson et al. 2007; Rodriguez and ter Horst 2008), and survival analysis (De Iorio et al. 2009).

The DP, as described by Ferguson (1973), does not incorporate covariates. It specifies a prior model $p(G)$ for a random probability measure G , without any notion of a regression on covariates. Hence, an active area of research is extending nonparametric models to allow the unknown distribution to depend on covariates. A popular approach starts with the Sethuraman (1994) representation of a DP. This representation states that a random measure G that follows a DP with total mass parameter M and base measure G_0 , denoted $DP(M, G_0)$, can be represented as

$$G = \sum_{h=1}^{\infty} w_h \delta(\theta_h), \quad w_h = v_h \prod_{i=1}^{h-1} (1 - v_i), \quad v_h \sim \text{Beta}(1, M), \quad \theta_h \sim G_0, \quad (1)$$

where $\delta(\theta)$ is a point mass at θ . Generalizations of (1) achieve the desired dependence of G on covariates by making the weights, w_h , and/or the locations, θ_h , vary with the covariates according to a stochastic process. Such extensions of the Sethuraman representation are probability models on a collection of dependent random probability measures $\{G_x, x \in X\}$, where X is the corresponding covariate space. MacEachern (1999) defines the Dependent DP (DDP) as

$$G_x = \sum_{h=1}^{\infty} w_h \delta(\theta_{xh}), \quad \text{for any } x. \quad (2)$$

Other generalizations of the Sethuraman representation that introduce covariates via the locations have been applied to the analysis of variance (De Iorio et al. 2004), spatial modeling (Gelfand et al. 2005), time series (Caron et al. 2006), and regression (De Iorio et al. 2009).

On the other hand, there are also generalizations based on making the weights covariate-dependent. (Griffin and Steel 2006) proposed an order-based dependent Dirichlet process (DP) that makes the order of v_h in the stick-breaking construction be a function of the covariates. Dunson and Park (2008) express v_h as a covariate-dependent kernel multiplied by beta weights, and Fuentes-García et al. (2009) model w_h with a nonparametric mixture model that depends on covariates.

The literature also includes approaches for modeling covariates that are not based on the Sethuraman representation. Cifarelli and Regazzini (1978) introduce dependence on covariates by defining a covariate-dependent regression model as the base measure of independent DPs at each level of the covariates (Carota and Parmigiani 2002; Griffin and Steel 2004). Another approach for modeling dependence is forming convex combinations of independent DP (Dunson et al. 2007; Müller et al. 2004). Finally, Müller et al. (1996) and Müller and Rosner (1998) include covariates x_i in an augmented response vector (y_i, x_i) and obtain the desired dependence by focusing on the conditional distribution given x_i . We refer to this last method as Conditional DP.

Although modeling dependence on covariates has been a very active area of research, limited research has examined the relative performance of such methods or improved

understanding of which features are suitable in order to produce better results. There is a lack of guidelines for partitioners that would allow them to choose among the many alternative models. This article considers such a comparison, focusing on predictive inference and continuous covariates. We show that different approaches for modeling dependence on continuous covariates can lead to very similar posterior fits yet produce very different results when used for prediction. In addition, when the predictive density is a mixture, we show that making the weights depend on the covariates improves the quality of the predictions. We illustrate our recommendation by comparing the Linear DDP (De Iorio et al. 2009) to the Conditional-DP (Müller et al. 1996). As we show, each of these two Bayesian nonparametric models incorporates continuous covariates in different ways, illustrating two main themes that also characterize many other approaches. We apply these two nonparametric models to a simulated, albeit extreme, data set and to data from a pharmacokinetic meta-analysis in pediatric oncology. Section 2 describes and compares both methods. Implementation and empirical results are reported in Section 3. Finally, conclusion and discussion appear in Section 4.

2 Modeling Approaches

This section describes the Linear DDP and the Conditional DP, points out their differences regarding the form of the predictive density, and explains how such differences affect the performance of each method.

Both methods introduce continuous covariates into the typical DP mixture (DPM) model given by

$$\mathbf{y}_i \stackrel{iid}{\sim} \int f(\mathbf{y}|\mu) dG(\mu), \quad G \sim DP(M, G_0), \quad (3)$$

where f is a probability density. The DPM model (3) is a mixture model with a DP prior on the mixing measure G . In many practical applications, the kernel $f(\mathbf{y}|\mu)$ is set to be a normal multivariate density with mean μ and common covariance matrix $\mathbf{S}$. The discrete nature of G implies that (3) reduces to a sum over the point masses θ_h in (1).

2.1 Linear DDP

The Linear DDP introduced in (De Iorio et al. 2009) models the relationship between continuous covariates and the unknown distribution by replacing the random probability measure G in the DPM model (3) with a collection of random probability measures indexed by $\mathbf{x}$, $\{G_{\mathbf{x}}, \mathbf{x} \in X\}$, where $\mathbf{x} = (x_1, \dots, x_d)$ denotes a d -dimensional vector of continuous covariates and X is the corresponding covariate space.

De Iorio et al. (2009) define a prior on $\{G_{\mathbf{x}}, \mathbf{x} \in X\}$. The construction is a special case of the DDP model in (2). A prior for the point masses $\theta_{\mathbf{x}h}$ is defined by assuming that $\theta_h = \{\theta_{\mathbf{x}h}, \mathbf{x} \in X\}$ are i.i.d. realizations of a stochastic process in $\mathbf{x}$. Thus the point masses $\theta_{\mathbf{x}h}$ are independent across h (for fixed $\mathbf{x}$). The weights, w_h , follow a stick-breaking prior as in

Sethuraman (1994), that is, $w_h = v_h \prod_{i=1}^{h-1} (1 - v_i)$, $v_h \sim \text{Beta}(1, M)$. The DDP model (2) implies that $G_{\mathbf{x}}$ follows a DP for each $\mathbf{x}$. Specifically, $G_{\mathbf{x}} \sim DP(M, G_{0\mathbf{x}})$, where the base measure $G_{0\mathbf{x}}$ is the marginal distribution at $\mathbf{x}$ of the stochastic process on the point masses $\theta_{\mathbf{x}h}$.

In summary, the DDP model induces dependence of the random measures $G_{\mathbf{x}}$ by assuming that the sample paths θ_h are dependent across $\mathbf{x}$. De Iorio et al. (2009), in particular, impose a linear model on $\theta_{\mathbf{x}h}$ given by

$$\theta_{xh} = \mathbf{m}_h + \sum_{i=1}^d \beta_{ih} x_i, \quad (4)$$

where d is the number of continuous covariates, $\mathbf{m}_h \stackrel{iid}{\sim} p_{\mathbf{m}}^0$ and $\beta_{ih} \stackrel{iid}{\sim} p_{\beta_i}^0$, $i = 1 \dots d$. It follows that the base measure, G_{0x} , is given by the convolution of $p_{\mathbf{m}}^0$ and $p_{\beta_i}^0$, $i = 1 \dots d$. The distributional assumptions on θ_{xh} proposed by De Iorio et al. (2009) imply that the random measures G_{0x} share the common main effect given by the random $\mathbf{m}_h$. In addition, for each i , β_i represents a slope coefficient as in a standard linear model.

The prior (2) with the linear model on the locations as in (4) is the Linear DDP of De Iorio et al. (2009). When such a prior is used to introduce continuous covariates into the DPM model (3) with a normal kernel, it leads to models given by

$$(\mathbf{y}_i | \mathbf{x}_i = \mathbf{x}) \stackrel{iid}{\sim} \int N(\mathbf{y}; \mu, \mathbf{S}) dG_{\mathbf{x}}(\mu), \quad \{G_{\mathbf{x}}, \mathbf{x} \in X\} \sim \text{Linear DDP}(M, G_{0x}), \quad (5)$$

where G_{0x} denotes the set of distributions for the main effect and slope coefficients in (4). Appropriate hyperpriors for S , M , and G_{0x} complete the model.

It is possible to rewrite model (5) in terms of a mixture of linear models. Let $\mathbf{x}_i$ denote the continuous covariate vector of subject i . Now, define the random matrix $\mathbf{\Gamma}_h = [\mathbf{m}_h, \beta_{1h}, \dots, \beta_{ph}]$ and design vectors $\mathbf{d}_i = (1, \mathbf{x}_i)$, $i = 1, \dots, n$, such that $\theta_{xih} = \mathbf{\Gamma}_h \mathbf{d}_i$. Model (5) can then be rewritten as

$$(\mathbf{y}_i | \mathbf{x}_i) \stackrel{iid}{\sim} \int N(\mathbf{y}; \tilde{\mathbf{\Gamma}} \mathbf{d}_i, \mathbf{S}) dG(\tilde{\mathbf{\Gamma}}), \quad G \sim DP(M, G_0), \quad (6)$$

with base measure $G_0 = (p_{\mathbf{m}}^0, p_{\beta_1}^0, \dots, p_{\beta_p}^0)$. The reformulation (6) is convenient because it has the form of a DPM model (3). Hence, it is possible to carry out posterior inference in the Linear DDP model by using well-known MCMC algorithms designed for DPM models, such as those described by MacEachern and Müller (1998).

For later reference we replace the mixture in (6) by an additional level in the hierarchical model

$$(\mathbf{y}_i | \mathbf{x}_i, \mathbf{\Gamma}_i, \mathbf{S}) \stackrel{iid}{\sim} N(\mathbf{y}; \mathbf{\Gamma}_i \mathbf{d}_i, \mathbf{S}), \quad \mathbf{\Gamma}_i \sim G, \quad \text{and} \quad G \sim DP(M, G_0). \quad (6a)$$

Since a DP is almost surely discrete, there is a positive probability for ties among the $\mathbf{\Gamma}_i$. Let $\{\mathbf{\Gamma}_1^* \dots \mathbf{\Gamma}_k^*\}$, $k \leq n$, be the set of unique values in $\{\mathbf{\Gamma}_1 \dots \mathbf{\Gamma}_n\}$, and let n_j denote the number of $\mathbf{\Gamma}_i$ equal to $\mathbf{\Gamma}_j^*$. Also for later reference, we state the conditional prior $p(\mathbf{\Gamma}_{n+1} | \mathbf{\Gamma}_1, \dots, \mathbf{\Gamma}_n)$ using this notation:

$$p(\mathbf{\Gamma}_{n+1} | \mathbf{\Gamma}_1, \dots, \mathbf{\Gamma}_n) \propto \sum_{j=1}^k n_j I(\mathbf{\Gamma}_{n+1} = \mathbf{\Gamma}_j^*) + M G_0(\mathbf{\Gamma}_{n+1}). \quad (6b)$$

In particular, note that the conditional distribution is marginalized with respect to the unknown probability measure G . The prior predictive model (6b) is also known as the Polya urn. It is one of the many alternative defining properties of the DP prior (Ferguson 1973).

We conclude the description of the Linear DDP by deriving the predictive density of model (5). Let Θ denote the set of all model parameters in (6a), and let $\Theta^{(t)}$, $t = 1, \dots, n$ denote a posterior Monte Carlo sample. With Y representing the data in hand and $\mathbf{x}_{n+1}$ the covariate vector for a new subject, the predictive density can be approximated as follows:

$$\begin{aligned} p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, Y) &= E[p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, Y, \Theta)|Y] \\ &\approx \frac{1}{N} \sum_{t=1}^N p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, Y, \Theta^{(t)}) \\ &= \frac{1}{N} \sum_{t=1}^N p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}). \end{aligned}$$

The specific expressions for the probabilities in the last average are easily obtained from (6b):

$$p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}) = \sum_{j=1}^{k^{(t)}+1} \alpha_j^{(t)} p_j(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}). \quad (7)$$

In (7),

$$\alpha_j^{(t)} \propto \begin{cases} n_j^{(t)} & j=1 \dots k^{(t)}, \\ M^{(t)} & j=k^{(t)}+1, \end{cases}$$

with $\sum_{j=1}^{k^{(t)}+1} \alpha_j^{(t)} = 1$, and

$$p_j(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}) = \begin{cases} N(\mathbf{y}_{n+1}; \mathbf{\Gamma}_j^{*(t)} \mathbf{d}_{n+1}, \mathbf{S}^{(t)}) & j=1 \dots k^{(t)}, \\ \int N(\mathbf{y}_{n+1}; \mathbf{\Gamma} \mathbf{d}_{n+1}, \mathbf{S}^{(t)}) dG_0(\mathbf{\Gamma}) & j=k^{(t)}+1. \end{cases}$$

2.2 Conditional DP

The Conditional DP approach introduces regression into the DPM model (3) by including the continuous covariates in an augmented response vector $\mathcal{Y} = (\mathbf{y}, \mathbf{x})$ in the nonparametric model. Specifically, Müller and Rosner (1998) make the unknown distribution $p(\mathbf{y})$ depend on covariates $\mathbf{x}$ by defining a DP mixture model for the joint model $p(\mathbf{y}, \mathbf{x})$. Considering a normal kernel, the Conditional DP modifies the DPM model as follows:

$$(\mathbf{y}_i, \mathbf{x}_i) \stackrel{iid}{\sim} \int N((\mathbf{y}, \mathbf{x}); \mu, \mathbf{S}) dG(\mu), \quad G \sim DP(M, G_0), \quad (8)$$

which is equivalent to the hierarchical model

$$\begin{aligned} (\mathbf{y}_i, \mathbf{x}_i) &\stackrel{iid}{\sim} N(\mathbf{y}; \boldsymbol{\mu}_i, \mathbf{S}), \\ \boldsymbol{\mu}_i &\sim G, \quad \text{and} \quad G \sim DP(M, G_0). \end{aligned} \quad (8a)$$

We refer to this modeling approach as Conditional DP, because the implied conditional distribution $p(\mathbf{y}|\mathbf{x})$ formalizes the desired regression on $\mathbf{x}$. Particularly, as explained in Müller and Rosner (1998), the mixture of normals $\int N((\mathbf{y}, \mathbf{x}); \boldsymbol{\mu}, \mathbf{S}) dG(\boldsymbol{\mu})$ implies a locally weighted mixture of normal linear regressions for $E[\mathbf{y}_i|\mathbf{x}_i]$. Implementation of posterior inference for the Conditional DP model (8) is straightforward, because it has the form of a DPM model. Hence, the MCMC algorithms described by MacEachern and Müller (1998) can be used.

The predictive density for a new $(n+1)$ -th subject, conditional on covariates $\mathbf{x}_{n+1}$, can be estimated using

$$p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, Y) \approx \frac{1}{N} \sum_{t=1}^N p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}). \quad (9)$$

The conditional density in the last average can be derived from the density $p(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}|\Theta^{(t)})$ corresponding to the DPM model (8a) as follows. Let $\{\boldsymbol{\mu}_1^* \dots \boldsymbol{\mu}_k^*\}$, $k = n$, be the set of distinct vectors from the set $\{\boldsymbol{\mu}_1 \dots \boldsymbol{\mu}_n\}$, with n_j the number of $\boldsymbol{\mu}_i$ equal to $\boldsymbol{\mu}_j^*$. Similar to (7), the predictive density for $(\mathbf{y}_{n+1}, \mathbf{x}_{n+1})$ is

$$p(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}|\Theta^{(t)}) = \sum_{j=1}^{k^{(t)}+1} \alpha_j^{(t)} p_j(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}|\Theta^{(t)}) \quad (10)$$

with

$$\alpha_j^{(t)} \propto \begin{cases} n_j^{(t)} & j=1 \dots k^{(t)}, \\ M^{(t)} & j=k^{(t)}+1, \end{cases}$$

where $\sum_{j=1}^{k^{(t)}+1} \alpha_j^{(t)} = 1$, and

$$p_j(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}|\Theta^{(t)}) = \begin{cases} N(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}; \boldsymbol{\mu}_j^*, \mathbf{S}^{(t)}) & j=1 \dots k^{(t)}, \\ \int N(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}; \boldsymbol{\mu}, \mathbf{S}^{(t)}) dG_0(\boldsymbol{\mu}) & j=k^{(t)}+1. \end{cases}$$

It follows that the conditional density needed in (9) is given by

$$p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}) = \sum_{j=1}^{k^{(t)}+1} \left(\frac{\alpha_j^{(t)} p_j(\mathbf{x}_{n+1}|\Theta^{(t)})}{\sum_{\ell=1}^{k^{(t)}+1} \alpha_\ell^{(t)} p_\ell(\mathbf{x}_{n+1}|\Theta^{(t)})} \right) p_j(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)}), \quad (11)$$

where $p_j(\mathbf{x}_{n+1}|\Theta^{(t)})$ is obtained by integrating out $\mathbf{y}_{n+1}$ from the joint distribution $p_j(\mathbf{y}_{n+1}, \mathbf{x}_{n+1}|\Theta^{(t)})$ in the mixture (10).

2.3 Comparing Approaches

Both extensions of the DPM model, the Linear DDP and the Conditional DP, estimate the predictive density by averaging the mixture distribution $p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)})$ with respect to a Monte Carlo sample from the posterior distribution. We claim, however, that the specific factors determining the weights in such mixtures can affect the resulting predictive inference.

As shown in (7), the weights in the mixture density $p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)})$ corresponding to the Linear DDP are solely determined by $n_j^{(t)}$ and $M^{(t)}$. That is, the relative importance of each component in the mixture is mainly determined by the size of the “clusters” induced by the unique values of $\Gamma_i^{(t)}$. This feature is an important limitation of the Linear DDP. It implies that, for each t , the weights of the mixture remain the same for any new subject rather than change as a function of the subject’s particular covariates $\mathbf{x}_{n+1}$. In other words, there is no built-in mechanism in the predictive density to favor those components in the mixture that are more likely to provide a better fit to the specific characteristics of the new subject.

Unlike the formulas in the Linear DDP, the weights in the mixture $p(\mathbf{y}_{n+1}|\mathbf{x}_{n+1}, \Theta^{(t)})$ corresponding to the Conditional DP are a function of the covariates via the marginal density $p_j(\mathbf{x}_{n+1}|\Theta^{(t)})$, as shown in (11). It follows that components in the mixture with a high marginal density on $\mathbf{x}_{n+1}$ will tend to have higher weights. Hence, the mixture is adjusted to reflect the specific characteristics of a given new subject. The inclusion of covariates in the weights increases the chance of using a regression structure suitable for the given subject.

In summary, the differences mentioned above suggest that, in terms of predictive inference, the Conditional DP should outperform the Linear DDP. We present empirical evidence corroborating such a claim in the next section.

Finally, we comment on a weakness of the Conditional DP that somewhat offsets the discussed features. The sampling model in (7) can be factored as $p(y_i|x_i) * p(x_i)$, highlighting the fact that the likelihood includes an additional factor for the covariates x_i . This is technically inappropriate when, for example, the covariates are chosen and fixed by design.

3 Empirical Implementation

This section illustrates the superior performance of the Conditional DP over the Linear DDP for predictive inference. Two data sets are considered; the first one is a simulated data set, while the second derives from a population pharmacokinetic study.

3.1 Simulation Example

The simulated data set corresponds to a multiple regression model with two dependent variables and a single predictor variable. Specifically, the data are generated from a model similar to (6) as follows. Let $\Gamma = [\mathbf{m}, \boldsymbol{\beta}]$ be a 2×2 matrix where $\mathbf{m}$ is a vector of constants and $\boldsymbol{\beta}$ is a vector of slope coefficients. Our simulated data set includes a sample of bivariate observations $\mathbf{y}_i = (y_{i1}, y_{i2})^T$, $i = 1, \dots, 100$, generated from the bivariate distribution $N(\Gamma_i \mathbf{d}_i, s\mathbf{I})$, where $\mathbf{I}$ is the 2×2 identity matrix, $s = 0.10$, and $\mathbf{d}_i = (1, x_i)^T$ is a design vector including the subject-specific covariate x_i . We introduce two underlying regression structures into the simulated data set by randomly setting Γ_i equal to one of the following:

$$\left\{ \begin{array}{l} \Lambda_1 = \begin{bmatrix} 0 & 3 \\ 1 & 0 \end{bmatrix}, \quad w.p. 1/2 \\ \Lambda_2 = \begin{bmatrix} 0 & 0 \\ 1 & 3 \end{bmatrix}, \quad w.p. 1/2 \end{array} \right. \quad (12)$$

Finally, each x_i is generated from a uniform distribution. If $\Gamma_i = \Lambda_1$ then $x_i \sim U(0, 1)$, otherwise $x_i \sim U(-1, 0)$. Hence, the value of the covariate x_i provides information about the specific underlying regression structure. As shown later, the ability to incorporate such information is a crucial difference between the Linear DDP and the Conditional DP.

The resulting simulated data set reflects two regression structures given by the product $\Gamma_i d_i$; one describes the horizontal line segment $(0, 1)^T + x(3, 0)^T$, $x \in (0, 1)$, while the other follows the vertical line segment $(0, 1)^T + x(0, 3)^T$, $x \in (-1, 0)$ (see Figure 1). The simulated data are generated from a model that resembles (6), matching the modeling assumptions of the Linear DDP. Such a feature was chosen in order to rule out the characteristics of the simulated data as an explanation for the poor performance of the Linear DDP.

We used posterior MCMC simulation to generate posterior Monte Carlo samples under Linear DDP (6) and the Conditional DP (8), respectively. In both cases the uncertainty on the common matrix S is modeled by adopting a conjugate inverse Wishart prior $S^{-1} \sim \text{Wishart}(r, (r\mathbf{R})^{-1})$ with r degrees of freedom and mean $r(r\mathbf{R})^{-1} = \mathbf{R}^{-1}$. In addition, the base measure is assumed to be multivariate normal, $G_0 \sim N(\mathbf{b}, \mathbf{B})$, and M is given a gamma distribution, $M \sim \text{Ga}(a_m, b_m)$. We ran 10,000 iterations of the MCMC algorithm for both models; convergence was reached after 2000 iterations. In general, the posterior fit under the Linear DDP is slightly better than the one obtained with the Conditional DP. For example, we estimated the first, second, and third quartiles, (Q1, Q2, Q3), of the distribution of the Euclidean distance between the posterior mean of y_i and its sample value. Those statistics were equal to (0.07, 0.12, 0.18) for the Linear DDP, while the Conditional DP showed somewhat greater values, (0.09, 0.16, 0.32).

Substituting the posterior Monte Carlo samples in (7) and (11), we evaluated the posterior predictive density corresponding to a new subject with covariate $x \in (-1, 1) \setminus \{0\}$ (i.e., $x \neq 0$). The results for $x = 0.5$ and $x = -0.5$ highlight the shortcoming of the Linear DDP. It wrongly assigns positive probability to regions in the plane without any sample points (see Figure 1). Such results can be explained as follows. Since the Linear DDP correctly recognizes the two regression structures, the mixture in (7) is dominated by two components, each one corresponding to one of the regression structures. It follows that for any given covariate x , the mixture in (7) gives positive probability to two regions, one for each regression structure. In contrast, the Conditional DP incorporates the values of the covariates into the weights of the mixture (11). Therefore, the Conditional DP is able to identify those regression structures that are more likely, given the value of x .

3.2 Pharmacokinetics Example

We consider a data set that includes clinical covariates and longitudinal data consisting of drug concentration measurements. The measurements came from patients receiving the anti-cancer drug topotecan in several different studies. We fit a Bayesian population pharmacokinetic (PK) model to estimate and predict plasma concentration-time curves. This section compares the performance of the Linear DDP and Conditional DP when introducing dependence on patient covariates into the PK model. The inclusion of covariates enables us to predict individual concentration-time curves appropriate for an individual patient instead

of the group average. Such estimated curves provide valuable information that can be used, for example, to design dose individualization schemes for future patients who are starting topotecan treatment.

Before describing the Bayesian PK model, here are some details about the data set. The population consisted of 138 children enrolled in seven clinical studies (see Table 1). These data have previously been used by Schaiquevich et al. (2007) to characterize the population pharmacokinetics of topotecan lactone in children with cancer and to identify covariates related to topotecan disposition. The data include concentration-time measurements corresponding to several treatment occasions for each patient. We use only those corresponding to the first treatment in our sample. This is because we are interested in making inference (and prediction) for patients with no previous topotecan treatment. In all cases, topotecan was infused intravenously over 30 minutes. See Schaiquevich et al. (2007) and references therein for more details regarding the eligibility criteria for the clinical studies, drug administration, blood collection, and patient demographics.

The structure of a Bayesian population PK model is as follows. Let y_{ij} denote the j -th measurement for the i th patient, and θ_i the vector of random effects of patient i . The probability model for the PK data is given by

$$p(y_{ij}|\theta_i), \quad p(\theta_i|\varphi), \quad p(\varphi). \quad (13)$$

Here, $p(y_{ij}|\theta_i)$ is a parametric, and typically non-linear, regression model for the concentration-time curve, $p(\theta_i|\varphi)$ is the prior for θ_i , and, finally, $p(\varphi)$ denotes the probability model of the hyperparameters. Bayesian models similar to (13) have been considered in Zeger and Karim (1991) for generalized linear mixed models and in Wakefield (1994) using a multivariate normal population distribution as population model.

The Linear DDP and the Conditional DP can be used to introduce the covariates, $\mathbf{x}_i$, into the PK model (13) via the prior distribution on θ_i . The resulting prior has the form (5) or (8), respectively, where $\mathbf{y}_i$ is replaced with θ_i . That is, using the Linear DDP model, the prior on $(\theta_i|\mathbf{x}_i)$ is a mixture of normals with a mixing measure given by a family of random measures indexed by $\mathbf{x}$. If the Conditional DP is used, the vector of parameters is augmented to include the covariates. Hence, the prior $p(\theta_i, \mathbf{x}_i)$ is a mixture of normals

$$p(\theta_i, \mathbf{x}_i) = \int N((\theta_i, \mathbf{x}_i); \mu, \mathbf{S}) dG(\mu),$$

with a DP prior on the mixing measure G . The model achieves the desired nonlinear, semi-parametric regression of the parameters on the covariates via the conditional distribution

$$p(\theta_i|\mathbf{x}_i) = \frac{p(\theta_i, \mathbf{x}_i)}{\int p(\theta_i, \mathbf{x}_i) d\theta_i} \propto p(\theta_i, \mathbf{x}_i).$$

Both modeling approaches assign a flexible non-parametric prior distribution to the population parameters. Hence, they both are able to accommodate heterogeneity in the patient population, such as outliers, over-dispersion, and multimodality.

Since we want to emphasize the differences between the Linear DDP and the Conditional DP, similar criteria are used to specify $p(y_{ij}|\theta_i)$ and $p(\phi)$, regardless of the prior on θ_i being considered.

The model for the concentration-time curve, $p(y_{ij}|\theta_i)$, is determined by the nonlinear regression

$$\log(y_{ij}) = \log(f(\theta_i, \tau_{ij})) + \varepsilon_{ij},$$

where y_{ij} is the j th concentration measurement for the i th patient at time τ_{ij} , $\varepsilon_{ij} \sim N(0, \lambda^{-1})$ is the noise term with precision λ , and the function f is a two-compartment model with constant rate intravenous infusion (Gibaldi and Perrier 1982) that describes the concentration at time τ as

$$f(\theta, \tau) = \frac{D/\gamma}{V_1 K_2} \left\{ 1 - \left[\left(\frac{K_2 - \alpha}{\beta - \alpha} \right) e^{-\beta\tau + \beta(\tau - \gamma)\mathbf{1}_{(\tau \geq \gamma)}} - \left(\frac{K_2 - \beta}{\beta - \alpha} \right) e^{-\alpha\tau + \alpha(\tau - \gamma)\mathbf{1}_{(\tau \geq \gamma)}} \right] \right\} \\ \times \left[\left(\frac{K_2 - \alpha}{\beta - \alpha} \right) e^{-\beta(\tau - \gamma)} - \left(\frac{K_2 - \beta}{\beta - \alpha} \right) e^{-\alpha(\tau - \gamma)} \right] \mathbf{1}_{(\tau \geq \gamma)} \quad (14)$$

with

$$\alpha = \frac{1}{2} \left[(K_1 + K_2 + K_{-1}) + \sqrt{(K_1 + K_2 + K_{-1})^2 - 4K_{-1}K_2} \right], \\ \beta = \frac{1}{2} \left[(K_1 + K_2 + K_{-1}) - \sqrt{(K_1 + K_2 + K_{-1})^2 - 4K_{-1}K_2} \right].$$

Here, the four parameters $\{V_1, K_2, K_1, K_{-1}\}$ are all positive, γ is the duration of infusion, D is the dose, and $\mathbf{1}_{(\cdot)}$ denotes an indicator function. Finally, we parameterize (14) with

$$\theta = (\log(V_1), \log(K_2), \log(K_1), \log(K_{-1}))^T.$$

Such a parameterization guarantees that the subject-specific parameters can take any value, and thus the use of a mixture of normals as a prior is appropriate.

Finally, we introduce distributional assumptions for the hyperparameters, ϕ , that allow the implementation of a MCMC algorithm that is computationally efficient. In the discussion that follows, it is assumed that the Linear DDP is written in the equivalent form (6). Regarding the parameters of the DP, the total mass parameter M is given a gamma prior $Ga(a_m, b_m)$, while the base measure G_0 follows a multivariate normal distribution $N(\mathbf{b}, \mathbf{B})$. The moments of G_0 are assumed random with hyperpriors $\mathbf{b} \sim N(\mathbf{b}_0, \mathbf{B}_0)$ and $\mathbf{B}^{-1} \sim \text{Wishart}(w, (w\mathbf{W})^{-1})$, where w is the degrees of freedom and $\mathbf{W}^{-1}$ is the mean of $\mathbf{B}^{-1}$. Finally, a Wishart prior is also used for $\mathbf{S}$, $\mathbf{S}^{-1} \sim \text{Wishart}(r, (r\mathbf{R})^{-1})$, and λ is given a gamma prior $G(a_\lambda, b_\lambda)$.

Unlike $p(y_{ij}|\theta_i)$, the characteristics of $p(\phi)$ change in accordance with the prior being used for θ . Specifically, the dimension of G_0 changes as follows. Let p denote the number of model parameters; that is, p equals the dimension of θ . Let d be the number of covariates

being modeled. For the Conditional DP, G_0 is a distribution on a vector with dimension $p+d$, while under the Linear DDP, G_0 is a distribution on the columns of the matrix Γ . In the latter case, it is useful to think of G_0 as the distribution of the vector with dimension $p(d+1)$ that results from the matrix Γ by placing (i.e., stacking) the columns of the matrix end to end.

Putting together the assumptions described above for each component in (13), it follows that two population PK models have been completely specified, each one implementing a different prior on θ . For brevity, we will refer to those PK models by using only the name of the modeling approach used for the prior on θ , that is, Conditional DP or Linear DDP. The analysis that follows includes the covariates age, body surface area (BSA), and glomerular filtration rate (GFR). Such covariates were chosen because, as shown in Schaiquevich et al. (2007), they are significant for explaining topotecan disposition across patients. Covariate values were centered at zero.

Implementation of both PK models requires a MCMC scheme to sample from the corresponding posterior distribution. Such sampling schemes can be efficiently implemented, since both the kernel of the mixture and the base measure G_0 are normally distributed (MacEachern and Müller 1998). Since the full conditional of all the parameters, with the exception of θ , have a closed form, they can be updated via Gibbs sampling. For θ , the non-linearity in (14) implies that its full conditional is not a known distribution. Therefore, θ is updated using the adaptive Metropolis Hasting algorithm introduced by Haario et al. (2001). For each model, the corresponding MCMC sampling scheme is used to draw a Monte Carlo sample of size 30, 000 from the posterior distribution, after a burn-in period of 20, 000 iterations.

Despite the differences between the Linear DDP and the Conditional DP, they lead to very similar posterior fits on the concentration-time curves per child. Specifically, posterior mean estimates of each curve, along with pointwise 95% probability intervals, are practically indistinguishable (see Figure 2). Although the curves correspond to a single patient in the data set, similar results are found for all of the 138 children in the seven studies.

It is in terms of prediction, however, that the Linear DDP and the Conditional DP show different performance. In particular, we compare the predictive distribution for the concentration-time curves of the 15 patients belonging to study 1. We obtained a Monte Carlo sample from those distributions by first fitting each PK model to the other patients in the data set (123 children in studies 2–7), then generating a Monte Carlo sample from the predictive distribution of θ for each patient belonging to study 1 (using formulas (7) and (11), along with the patient-specific covariates), and, finally, using each of those samples to evaluate (14).

Comparison of the results by method shows that only the Conditional DP was able to produce sensible estimators. As seen in Figure 3, the Linear DDP had greater predictive uncertainty than the Conditional DP and led to unrealistic predicted concentration-time profiles. When using the Linear DDP, the estimated predictive distribution provided no information about the real concentration-time curves. Although the curves shown correspond to a single patient, similar results were found for all the patients in study 1.

4 Discussion

In this paper we have focused on extensions of nonparametric Bayesian models that introduce dependence on covariates. We have shown that good posterior fits with those extensions do not necessarily translate to good prediction. In addition, we have shown that, when the predictive density of such extensions is estimated by averaging mixture distributions, better predictions are produced when the weights in that mixture depend on the

covariates. The arguments above have been illustrated by comparing the Linear DDP (De Iorio et al. 2009) and the Conditional DP (Müller and Rosner 1998).

When using the Linear DDP, the weights in the predictive density are solely determined by the size of the clusters induced by the DP, with each cluster representing a different regression structure. Such a feature is a drawback, because it leads to the Linear DDP being unable to identify those regression structures that are more likely for a new subject, given that subject's covariates relative to the other observed covariates. Combining the covariates with inadequate regression structures leads to unrealistic estimators. Evidence of such behavior has been shown in this paper. In the simulation example, the Linear DDP wrongly introduced an additional mode to the estimated predictive density, while in the pharmacokinetic example, it predicted unrealistic concentration curves.

In contrast, the predictive density of the Conditional DP is a mixture with weights that depend on the covariates of the new subject. This results in an agreement between the covariates and the regression structures. Such an agreement leads to considerable improvement in the predictive inference when compared to the Linear DDP. In spite of the clear differences in terms of prediction, both methods produce identical results when used for posterior inference.

Although we have focused on continuous covariates, it is also desirable when modeling categorical and discrete covariates to be able to identify the components in the predictive mixture that produce the best fit. Hence, in those cases it is also advantageous to have predictive densities with covariate-dependent weights. Still, under some modeling approaches for introducing dependence on finite covariates (numeric or categorical), such as the ANOVA DDP (De Iorio et al. 2004) or the Spatial DDP (Gelfand et al. 2005), the limitations from not having covariate-dependent weights may not be as noticeable as those shown by the Linear DDP. This happens because, in those models, the mechanism through which the covariates determine the form of the components in the mixture can vary, if needed, depending on the specific values of the covariates. Such adaptability is not achieved by the Linear DDP, because in each component the corresponding regression structure remains the same regardless of the specific values of the continuous covariates. In particular, the ANOVA DDP models dependence on categorical covariates in an ANOVA fashion. Assuming, for simplicity, that there is only one categorical covariate, the ANOVA DDP implies that the mean of each component in the predictive mixture distribution is equal to a mean effect plus an offset vector. Since such an offset vector can vary according to the value of the covariate, the means of the mixture components automatically account for component-specific associations with the discrete covariate level. Hence, it is possible to generate reasonable inferences, even though the weights in the predictive density are only based on the size of the clusters induced by the DP.

Finally, the results in this paper provide insight regarding which strategies for introducing dependence on covariates can lead to better posterior inferences. Specifically, it provides evidence in favor of extensions of the Sethuraman representation (1) that make the weights vary with the covariates, because such extensions result in nonparametric Bayesian models with a prediction rule based on a mixture with covariate-dependent weights. Examples of those extensions include the order-based DDP of Griffin and Steel (2006) and the kernel stick-breaking process in Dunson and Park (2008).

Acknowledgments

This research was supported by the National Cancer Institute grant R01CA075981; The Brown Foundation Fellowship, Center for Computational Finance and Economic Systems; and the NSF VIGRE grant DSM-0739420.

References

- Burr D, Doss H. A Bayesian semiparametric model for random-effects meta-analysis. *Journal of the American Statistical Association*. 2005; 100(469):242–251.
- Caron F, Davy M, Doucet A, Duflos E, Vanheeghe P. Bayesian inference for dynamic models with Dirichlet process mixtures. *IEEE Transactions on Signal Processing*. 2006; 56(1):71–84.
- Carota C, Parmigiani G. Semiparametric regression for count data. *Biometrika*. 2002; 89(2):265–285.
- Cifarelli D, Regazzini E. Nonparametric statistical problems under partial exchangeability. the use of associative means. *Annali del Istituto di Matematica Finanziaria dell Università di Torino*. 1978; 12:1–36.
- De Iorio M, Johnson W, Müller P, Rosner G. Bayesian nonparametric nonproportional hazards survival modeling. *Biometrics*. 2009; 65(3):762–771. [PubMed: 19210742]
- De Iorio M, Müller P, Rosner G, MacEachern S. An ANOVA model for dependent random measures. *Journal of the American Statistical Association*. 2004; 99(465):205–215.
- Dunson D. Bayesian nonparametric hierarchical modeling. *Biometrical Journal*. 2009; 51(2):273–284. [PubMed: 19358217]
- Dunson D, Park J. Kernel stick-breaking processes. *Biometrika*. 2008; 95(2):307–323. [PubMed: 18800173]
- Dunson D, Pillai N, Park J. Bayesian density regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2007; 69(2):163–183.
- Ferguson T. A Bayesian analysis of some nonparametric problems. *Annals of Statistics*. 1973; 1(2):209–230.
- Fuentes-García R, Mena R, Walker S. A nonparametric dependent process for Bayesian regression. *Statistics and Probability Letters*. 2009; 79(8):1112–1119.
- Gelfand A, Kottas A, MacEachern S. Bayesian nonparametric spatial modeling with Dirichlet process mixing. *Journal of the American Statistical Association*. 2005; 100(471):1021–1036.
- Gibaldi, M.; Perrier, D. *Pharmacokinetics*. 2. Marcel Dekker, Inc; New York: 1982.
- Griffin J, Steel M. Semiparametric Bayesian inference for stochastic frontier models. *Journal of Econometrics*. 2004; 123(1):121–152.
- Griffin J, Steel M. Order-based dependent Dirichlet processes. *Journal of the American Statistical Association*. 2006; 101(473):179–194.
- Haario H, Saksman E, Tamminen J. An adaptive Metropolis algorithm. *Bernoulli*. 2001; 7(2):223–242.
- Kim S, Tadesse M, Vannucci M. Variable selection in clustering via Dirichlet process mixture models. *Biometrika*. 2006; 93(4):877–893.
- MacEachern, S. *ASA Proceedings of the Section on Bayesian Statistical Science*. American Statistical Association; Alexandria, VA: 1999. Dependent nonparametric processes; p. 50–55.
- MacEachern S, Müller P. Estimating mixture of Dirichlet process models. *Journal of Computational and Graphical Statistics*. 1998; 7(2):223–238.
- Müller P, Erkanli A, West M. Bayesian curve fitting using multivariate normal mixtures. *Biometrika*. 1996; 83(1):67–79.
- Müller P, Quintana F, Rosner G. A method for combining inference across related nonparametric Bayesian models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*. 2004; 66(3):735–749.
- Müller, P.; Rosner, G. Semi-parametric PK/PD models. In: Dey, D.; Müller, P.; Sinha, D., editors. *Practical Nonparametric and Semiparametric Bayesian Statistics*. Springer-Verlag; New York: 1998. p. 323–337.
- Rodríguez A, ter Horst E. Bayesian dynamic density estimation. *Bayesian Analysis*. 2008; 3(2):339–366.
- Rosner G, Müller P. Bayesian population pharmacokinetic and pharmacodynamic analyses using mixture models. *Journal of Pharmacokinetics and Pharmacodynamics*. 1997; 25(2):209–233.
- Schaiquevich P, Panetta JC, Iacono LC, Freeman BB, Santana VM, Gajjar A, Stewart CF. Population pharmacokinetic analysis of topotecan in pediatric cancer patients. *Clinical Cancer Research*. 2007; 13(22):6703–6711. [PubMed: 18006771]

- Sethuraman J. A constructive definition of Dirichlet priors. *Statistica Sinica*. 1994; 4(2):639–650.
- Wakefield J. An expected loss approach to the design of dosage regimens via sampling-based methods. *The Statistician*. 1994; 43(1):13–29.
- Xing E, Jordan M, Sharan R. Bayesian haplotype inference via the Dirichlet process. *Journal of Computational Biology*. 2007; 14(3):267–284. [PubMed: 17563311]
- Zeger S, Karim M. Generalized linear models with random effects; a Gibbs sampling approach. *Journal of the American Statistical Association*. 1991; 86(413):79–86.

Figure 1.

Contours of the predictive density. The elements of the simulated data set are marked with dots. The plots are organized by modeling approach (columns) and by covariate level used to perform predictive inference (rows). The estimators produced by the Linear DDP are not acceptable because they wrongly introduce an extra mode to the predictive density. Such results reflect the fact that the Linear DDP selects a regression structure based solely on the cluster size, which by design is approximately half for each one of the two underlying regression structures, rather than using the covariate level of the new subject.

(a) Conditional DP

(b) Linear DDP

Figure 2.

Posterior concentration-time curves by method. The dashed lines denote the posterior mean concentration curves, the solid lines refer to the 95% pointwise probability intervals, and the “x”s denote the observed concentration-time combinations. The estimated curves show that the posterior fits produced by both methods are practically the same.

(a) Conditional DP

(b) Linear DDP

Figure 3. Predictive concentration-time curves by method. The curves denote the pointwise quartiles Q1 Q2 Q3, and the “x”s mark the observed concentration-time combinations.

Table 1

Characteristics of the clinical trials from which PK data was obtained

Trial no.	No. of patients	Trial type	Topotecan dosage
1	15	Phase I recurrent solid tumors	Target AUC = 120–180 ng · h/mL
2	21	Phase I recurrent acute leukemia	Fixed dosage = 2.4 mg/m ²
3	28	Phase I recurrent solid tumors	0.8 and 1.1 mg/m ²
4	10	Phase II newly diagnosed medulloblastoma	Target AUC = 120–160 ng · h/mL
5	22	Phase I recurrent solid tumors	1.4, 1.7, 2.0, and 2.4 mg/m ²
6	30	Phase II newly diagnosed high-risk medulloblastoma	Target AUC = 80–120 ng · h/mL
7	12	Phase II recurrent Wilms tumor	Target AUC = 70–90 ng · h/mL