

FUNCTIONAL CLUSTERING IN NESTED DESIGNS: MODELING VARIABILITY IN REPRODUCTIVE EPIDEMIOLOGY STUDIES

BY ABEL RODRIGUEZ¹ AND DAVID B. DUNSON²

University of California, Santa Cruz and Duke University

We discuss functional clustering procedures for nested designs, where multiple curves are collected for each subject in the study. We start by considering the application of standard functional clustering tools to this problem, which leads to groupings based on the average profile for each subject. After discussing some of the shortcomings of this approach, we present a mixture model based on a generalization of the nested Dirichlet process that clusters subjects based on the distribution of their curves. By using mixtures of generalized Dirichlet processes, the model induces a much more flexible prior on the partition structure than other popular model-based clustering methods, allowing for different rates of introduction of new clusters as the number of observations increases. The methods are illustrated using hormone profiles from multiple menstrual cycles collected for women in the Early Pregnancy Study.

1. Introduction. The literature on functional data analysis has seen a spectacular growth in the last twenty years, showing promise in applications ranging from genetics [Luan and Li (2003), Ramoni, Sebastiani and Kohane (2002), Wakefield, Zhou and Self (2003)] to proteomics [Ray and Mallick (2006)], epidemiology [Bigelow and Dunson (2009)] and oceanography [Rodríguez, Dunson and Gelfand (2009)]. Because functional data are inherently complex, functional clustering is useful as an exploratory tool in characterizing variability among subjects; the resulting clusters can be used as a predictive tool or simply as a hypothesis-generating mechanism that can help guide further research. Some examples of functional clustering methods include Abraham et al. (2003), who use B-spline fitting coupled with k -means clustering; Tarpey and Kinateder (2003), who apply k -means clustering via the principal points of random functions; James and Sugar (2003), who develop methods for sparsely sampled functional data that employ spline representations; García-Escudero and Gordaliza (2005), where the robust k -means method for functional clustering is developed; Serban and

Received July 2010; revised February 2014.

¹Supported in part by Grant R01 GM090201-01 from the National Institute of General Medical Sciences of the National Institutes of Health.

²Supported in part by Grant R01 ES017240-01 from the National Institute of Environmental Health Sciences of the National Institutes of Health.

Key words and phrases. Nonparametric Bayes, nested Dirichlet process, functional clustering, hierarchical functional data, hormone profile.

Wasserman (2005), who use a Fourier representation for the functions along with k -means clustering; Heard, Holmes and Stephens (2006), where a Bayesian hierarchical clustering approach that relies on spline representations is proposed; Ray and Mallick (2006), who build a hierarchical Bayesian model that employs a Bayesian nonparametric mixture model on the coefficients of the wavelet representations; and Chiou and Li (2007), where a k -centers functional clustering approach is developed that relies on the Karhunen–Loève representation of the underlying stochastic process generating the curves and accounts for both the means and the modes of variation differentials between clusters.

All of the functional clustering methods described above have been designed for situations where a single curve is observed for each subject or experimental condition. Extensions to nested designs where multiple curves are collected per subject typically assume that coefficients describing subject-specific curves arise from a common parametric distribution, and clustering procedures are then applied to the parameters of this underlying distribution. The result is a procedure that generates clusters of subjects based on their average response curve, which is not appropriate in applications in which subjects vary not only in the average but also in the variability of the replicate curves. For example, in studies of trajectories in reproductive hormones that collect data from repeated menstrual cycles, the average trajectory may provide an inadequate summary of a woman's reproductive functioning. Some women have regular cycles with little variability across cycles in the hormone trajectories, while other women vary substantially across cycles, with a subset of the cycles having very different trajectory shapes. In fact, one indication of impending menopause and a decrease in fecundity is an increase in variability across the cycles. Hence, in forming clusters and characterizing variability among women and cycles in hormone trajectories, it is important to be flexible in characterizing both the mean curve and the distribution about the mean. This situation is not unique to hormone data, and similar issues arise in analyzing repeated medical images as well as other applications.

This paper discusses hierarchical Bayes models for clustering nested functional data in the context of the Early Pregnancy Study (EPS) [Wilcox et al. (1998)], where progesterone profiles were collected for both conceptive and nonconceptive women from multiple menstrual cycles. Our models use spline bases along with mixture priors to create sparse but flexible representations of the hormone profiles, and can be applied directly to other basis systems such as wavelets. We start by introducing a hierarchical random effects model on the spline coefficients which, along with a generalization of the Dirichlet process mixture (DPM) prior [Escobar and West (1995), Ferguson (1973), Sethuraman (1994)], allows for mean-response-curve clustering of women, in the spirit of Ray and Mallick (2006). Then, we extend the model to generate distribution-based clusters using a nested Dirichlet process (NDP) [Rodríguez, Dunson and Gelfand (2008)]. The resulting model simultaneously clusters both curves and subjects, allowing us to identify outlier

curves within each group of women, as well as outlying women whose distribution of profiles differs from the rest. To the best of our knowledge, there is no classical alternative for this type of distribution-based multilevel clustering.

In order to provide some insight into the challenges associated with functional clustering in the context of the EPS, consider the hormonal profiles depicted in Figure 1. Frames (a) to (c) depict the hormone profiles for 3 women, while frame (d) shows the mean profile corresponding to each one of them, obtained by simply averaging all available observations at a given day within the cycle. When looking

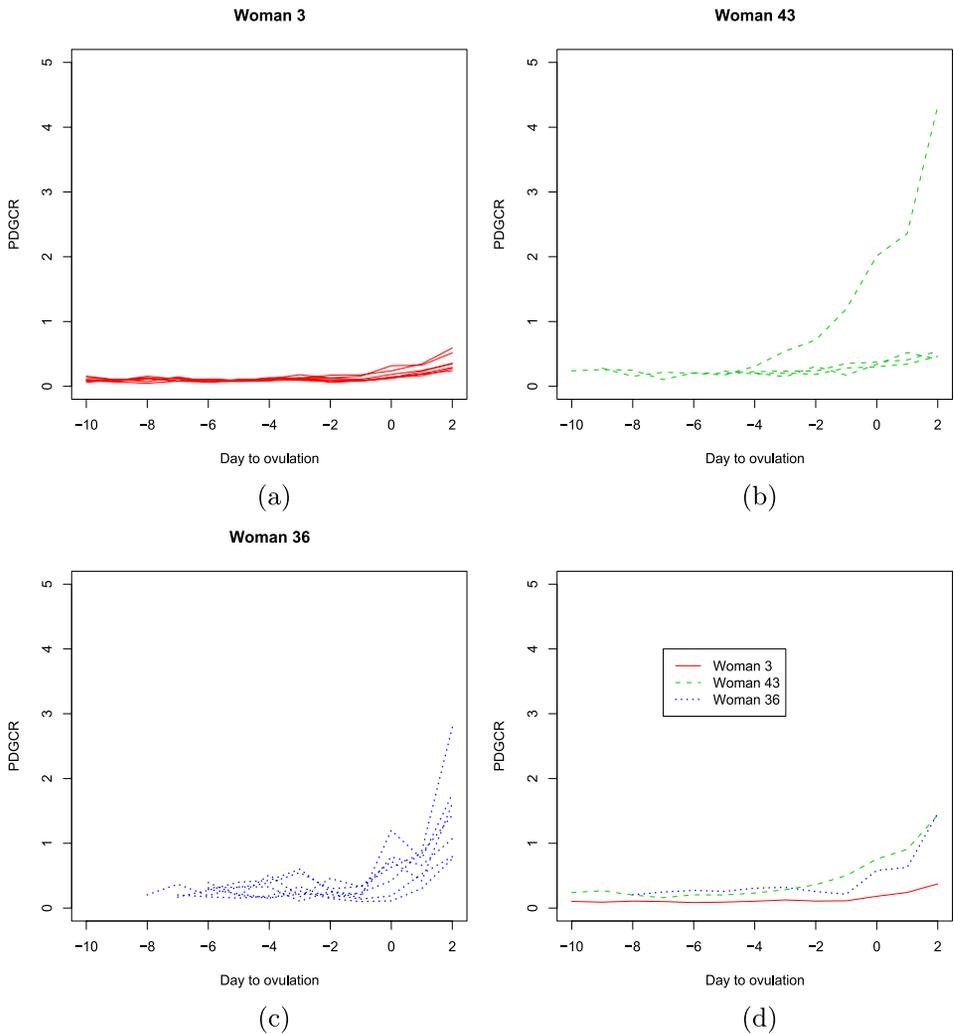


FIG. 1. Comparison of hormone profiles for three women in the Early Pregnancy Study. Frames (a) to (c) show multiple profiles for each woman, while frame (d) shows the average profile for each woman.

at the mean profiles in (d), women 43 and 36 seem to have very similar hormonal responses, which are different from those of woman 3. However, when the individual profiles are considered, it is clear that most of the cycles of woman 43 look like those of woman 3 and that the big difference in the means is driven by the single abnormal cycle. This result, although not surprising if we consider that means are notoriously nonrobust, suggests that simple approaches that average profiles over individuals (or, equivalently, use Gaussian distributions to describe subject-specific variability) might not be the most appropriate for this type of data.

The use of Bayesian nonparametric mixture models for clustering has a long history [Lau and Green (2007), Medvedovic and Sivaganesan (2002), Quintana and Iglesias (2003)] and presents a number of practical advantages over other model-based clustering techniques. Nonparametric mixtures induce a probability distribution on the space of partitions of the data, therefore, we do not need to specify in advance the number of clusters in the sample. Once updated using the data, this distribution on partitions allows us to assess variability, and hence characterize uncertainty in the clustering structure (including that associated with the estimation of the curves), providing a more complete picture than classical methods. In this paper, we work with a generalized Dirichlet process (GDP) first introduced by Hjort (2000) and study some of its properties as a clustering tool. In particular, we show that the GDP generates a richer prior on data partitions than those induced by popular models such as the Dirichlet process [Ferguson (1973)] or the two-parameter Poisson–Dirichlet process [Pitman (1996)], as it allows for an asymptotically bounded number of clusters in addition to logarithmic and power-law rates of growth.

The paper is organized as follows: Section 2 reviews the basics of nonparametric regression and functional clustering, while Section 3 explores the design of nonparametric mixture models for functional clustering. Building on these brief reviews, Section 4 describes two Bayesian approaches to functional clustering in nested designs, while Section 5 describes Markov chain Monte Carlo algorithms for this problem. An illustration focused on the EPS is presented in Section 6. Finally, Section 7 presents a brief discussion and future research directions.

2. Model-based functional clustering. To introduce our notation, consider first a simple functional clustering problem where multiple noisy observations are collected from functions f_1, \dots, f_I . More specifically, for subjects $i = 1, \dots, I$ and within-subject design points $t = 1, \dots, T_i$, observations consist of ordered pairs (x_{it}, y_{it}) , where

$$y_{it} = f_i(x_{it}) + \varepsilon_{it}, \quad \varepsilon_{it} \sim \mathbf{N}(0, \sigma_i^2).$$

For example, in the EPS, y_{it} corresponds to the level of progesterone in the blood of woman i collected at day x_{it} of the menstrual cycle, and f_i denotes a smooth trajectory in progesterone for woman i (initially supposing a single menstrual cycle of data from each woman), and clusters in $\{f_i\}_{i=1}^I$ could provide insight into the

variability in progesterone curves across women, while potentially allowing us to identify abnormal or outlying curves.

If all curves are observed at the same covariate levels (i.e., $T_i = T$ and $x_{it} = x_t$ for every i), a natural approach to functional clustering is to apply standard clustering methods to the data vectors, $\mathbf{y}_i = (y_{i1}, \dots, y_{iT})'$. For example, in the spirit of Ramsay and Silverman (2005), one could apply hierarchical or k -means clustering to the first few principal components [Yeung and Ruzzo (2001)]. Alternatively, with uneven spacings or missing observations, nonparametric regression techniques such as kernel regression [e.g., see Li and Racine (2004), Racine and Li (2004)] could be used to interpolate the value of the curves to a common grid, and then traditional clustering techniques could be applied. From a model-based perspective, one could instead suppose that \mathbf{y}_i is drawn from a mixture of k multivariate Gaussian distributions, with each Gaussian corresponding to a different cluster [Fraley and Raftery (2002), Yeung et al. (2001)]. The number of clusters could then be selected using the BIC criteria [Fraley and Raftery (2002), Li (2005)] or a nonparametric Bayes approach could be used to bypass the need for this selection, while allowing the number of clusters represented in a sample of I individuals to increase stochastically with sample size [Medvedovic and Sivaganesan (2002)].

A more general approach that allows us to deal with unevenly spaced and missing observations is to fit a nonparametric model to each curve and then project all the curves onto a common space. For example, we can represent the unknown function f_i as a linear combination of prespecified basis functions $\{b_k\}_{k=1}^p$, that is, we can write

$$f_i(x_{it}) = \theta_{i0} + \sum_{k=1}^p \theta_{ik} b_k(x_{it}),$$

where $\boldsymbol{\theta}_i = (\theta_{i0}, \theta_{i1}, \dots, \theta_{ip})$ are basis coefficients specific to subject i , with variability in these coefficients controlling variability in the curves $\{f_i\}_{i=1}^I$. A common approach to functional clustering is to induce clustering of the curves through clustering of the basis coefficients [Abraham et al. (2003), Heard, Holmes and Stephens (2006)]. Then the methods discussed above for clustering of the data vectors $\{\mathbf{y}_i\}_{i=1}^I$ in the balanced design case can essentially be applied directly to the basis coefficients $\{\boldsymbol{\theta}_i\}_{i=1}^I$.

Although the methods apply directly to other choices of basis functions, our focus will be on splines, which have been previously used in the context of hormone profiles [Bigelow and Dunson (2009), Brumback and Rice (1998)]; given a set of knots τ_1, \dots, τ_p , the k th member of the basis system is defined as

$$b_k(x) = (x - \tau_k)_+^q,$$

where $(\cdot)_+ = \max\{\cdot, 0\}$. Given the knot locations, inferences on $\boldsymbol{\theta}_i$ and σ_i^2 can be carried out using standard linear regression tools, however, selecting the number and location of the nodes τ_1, \dots, τ_p can be a challenging task. A simple solution is

to use a large number of equally spaced knots, together with a penalty term on the coefficients to prevent overfitting. From a Bayesian perspective, this penalty term can be interpreted as a prior on the spline coefficients; for example, the maximum likelihood estimator (MLE) obtained under an L^2 penalty on the spline coefficients is equivalent to the maximum a posteriori estimates for a Bayesian model under a normal prior, while the MLE under an L^1 penalty is equivalent to the maximum a posterior estimate under independent double-exponential priors on the spline coefficients.

Instead of the more traditional Gaussian and double exponential priors, in this paper we focus on zero-inflated priors, in the spirit of [Smith and Kohn \(1996\)](#). Priors of this type enforce sparsity by zeroing out some of the spline coefficients and, by allowing us to select a subset of the knots, provides adaptive smoothing. In their simpler form, zero-inflated priors assume that the coefficients are independent from each other and that

$$(1) \quad \theta_{ik} | \gamma, \sigma_i^2 \sim \gamma \mathcal{N}(0, \omega_k \sigma_i^2) + (1 - \gamma) \delta_0, \quad \sigma_i^2 \sim \text{IGam}(\nu_1, \nu_2),$$

where δ_x denotes the degenerate distribution putting all its mass at x , ω_k controls the overdispersion of the coefficients with respect to the observations and γ is the prior probability that the coefficient θ_{ik} is different from zero. In order to incorporate a priori dependence across coefficients, we can reformulate the hierarchical model by introducing Bernoulli random variables $\lambda_{i1}, \dots, \lambda_{ip}$ such that

$$\mathbf{y}_i | \boldsymbol{\theta}_i, \sigma_i^2, \boldsymbol{\Lambda}_i \sim \mathcal{N}(\mathbf{B}(\mathbf{x}_i) \boldsymbol{\Lambda}_i \boldsymbol{\theta}_i, \sigma_i^2 \mathbf{I}), \quad \boldsymbol{\theta}_i | \sigma_i^2 \sim \mathcal{N}(\mathbf{0}, \sigma_i^2 \boldsymbol{\Omega}),$$

$$\sigma_i^2 \sim \text{IGam}(\nu_1, \nu_2),$$

where $\mathbf{y}_i = (y_{i1}, \dots, y_{iT_i})$ and $\mathbf{x}_i = (x_{i1}, \dots, x_{iT_i})$ are, respectively, the vectors of responses and covariates associated with subject i , $\mathbf{B}(\mathbf{x}_i)$ is the matrix of basis functions also associated with subject i with entries $[\mathbf{B}(\mathbf{x}_i)]_{tk} = b_k(x_{it})$, $\boldsymbol{\Lambda}_i = \text{diag}\{\lambda_{i1}, \dots, \lambda_{ip}\}$ and λ_i equals 1 independently with probability γ , and $\boldsymbol{\Omega}$ is a $p \times p$ covariance matrix. Note that if $\boldsymbol{\Omega}$ is a diagonal matrix and $[\boldsymbol{\Omega}]_{kk} = \omega_k$, we recover the independent priors in (1). For the single curve case, choices for $\boldsymbol{\Omega}$ based on the regression matrix $\mathbf{B}(\mathbf{x}_i)$ are discussed in [DiMatteo, Genovese and Kass \(2001\)](#), [Liang et al. \(2008\)](#) and [Paciorek \(2006\)](#).

Although the preceding two-stage approach is simple to implement using off-the-shelf software, it ignores the uncertainty associated with the estimation of the basis coefficients while clustering the curves. In the spirit of [Fraley and Raftery \(2002\)](#), an alternative that deals with this issue is to employ a mixture model of the form

$$(2) \quad \mathbf{y}_i | \{\boldsymbol{\theta}_k^*\}, \{\sigma_k^{*2}\}, \{\boldsymbol{\Lambda}_k^*\} \sim \sum_{k=1}^K w_k \mathcal{N}(\mathbf{B}(\mathbf{x}_j) \boldsymbol{\Lambda}_k^* \boldsymbol{\theta}_k^*, \sigma_k^{*2} \mathbf{I}), \quad \sum_{k=1}^K w_k = 1,$$

where $\boldsymbol{\theta}_k^*$ is the vector of coefficients associated with the k th cluster, $\boldsymbol{\Lambda}_k^*$ is the diagonal selection matrix for the k th cluster, σ_k^{*2} is the observational variance associated with observations collected in the k th cluster, w_k can be interpreted as

the proportion of curves associated with cluster k , and K is the maximum number of clusters in the sample. From a frequentist perspective, estimation of this model can be performed using expectation–maximization (EM) algorithms, while selection of the number of mixture components can be carried out using the BIC. Alternatively, Bayesian inference can be performed for this model using Markov chain Monte Carlo (MCMC) algorithms once appropriate priors for the vector $\mathbf{w} = (w_1, \dots, w_K)$ and the cluster-specific parameters $(\boldsymbol{\theta}_k^*, \boldsymbol{\Lambda}_k^*, \sigma_k^{*2})$ have been chosen, opening the door to simple procedures for the estimation of the number of clusters in the sample.

3. Bayesian nonparametric mixture models for functional data. Note that the model in (2) can be rewritten as a hierarchical model by introducing latent variables $\{(\boldsymbol{\theta}_i, \sigma_i^2, \boldsymbol{\Lambda}_i)\}_{i=1}^I$ so that

$$(3) \quad \mathbf{y}_i | \boldsymbol{\theta}_i, \sigma_i^2, \boldsymbol{\Lambda}_i \sim \mathbf{N}(\mathbf{B}(\mathbf{x}_i) \boldsymbol{\Lambda}_i \boldsymbol{\theta}_i, \sigma_i^2 \mathbf{I}), \quad \boldsymbol{\theta}_i, \sigma_i^2, \boldsymbol{\Lambda}_i | G \sim G,$$

$$G(\cdot) = \sum_{k=1}^K w_k \delta_{(\boldsymbol{\theta}_k^*, \sigma_k^{*2}, \boldsymbol{\Lambda}_k^*)}(\cdot).$$

Therefore, specifying a joint prior on \mathbf{w} and $\{(\boldsymbol{\theta}_k^*, \sigma_k^{*2}, \boldsymbol{\Lambda}_k^*)\}_{k=1}^K$ is equivalent to specifying a prior on the discrete distribution G generating the latent variables $\{(\boldsymbol{\theta}_i, \sigma_i^2, \boldsymbol{\Lambda}_i)\}_{i=1}^I$. In this section we discuss strategies to specify a flexible prior distribution on this mixing distribution in the context of functional clustering. In particular, we concentrate on nonparametric specifications for G through the class of stick-breaking distributions.

A stick-breaking prior [Ishwaran and James (2001), Ongaro and Cattaneo (2004)] with baseline measure G_0 and precision parameters $\{a_l\}_{l=1}^L$ and $\{b_l\}_{l=1}^L$ is defined as

$$(4) \quad G(\cdot) = \sum_{k=1}^K w_k \delta_{\boldsymbol{\vartheta}_k}(\cdot),$$

where the atoms $\{\boldsymbol{\vartheta}_k\}_{k=1}^K$ are independent and identically distributed samples from G_0 and the weights $\{w_k\}_{k=1}^K$ are constructed as $w_k = u_k \prod_{s < k} (1 - u_s)$, with $\{u_k\}_{k=1}^K$ another independent and identically distributed sequence of random variables such that $u_k \sim \text{Beta}(a_k, b_k)$ for $k < K$ and $u_K = 1$. For example, taking $K = \infty$, $a_k = 1 - a$ and $b_k = b + ka$ for $0 \leq a < 1$ and $b > -a$ yields the two-parameter Poisson–Dirichlet process [Ishwaran and James (2001), Pitman (1995)], denoted $\text{PY}(a, b, G_0)$, with the choice $a = 0$ resulting in the Dirichlet process [Ferguson (1973), Sethuraman (1994)], denoted $\text{DP}(b, G_0)$. In mixture models such as (3), G_0 acts as the common prior for the cluster-specific parameters $\{\boldsymbol{\vartheta}_k\}_{k=1}^K$, while the sequences $\{a_k\}_{k=1}^K$ and $\{b_k\}_{k=1}^K$ control the a priori expected number and size of the clusters.

The main advantage of nonparametric mixture models such as the Poisson–Dirichlet process as a clustering tool is that they allow for automatic inferences on the number of components in the mixture. Indeed, these models induce a prior probability on all possible partitions of the set of observations, which is updated based on the information contained in the data. However, Poisson–Dirichlet processes have two properties that might be unappealing in our EPS application; first, they place a relatively large probability on partitions that include many small clusters and, second, they imply that the number of clusters will tend to grow logarithmically (if $a = 0$) or as a power law (if $a > 0$) as more observations are included in the data set. However, priors that favor [Introduction](#) of increasing numbers of clusters without bound as the number of subjects increase have some disadvantages in terms of interpretability and sparsity in characterizing high-dimensional data. For example, in applying DP mixture models for clustering of the progesterone curves in EPS, [Bigelow and Dunson \(2009\)](#) obtained approximately 32 different clusters, with half of these clusters singletons. Many of the clusters appeared similar, and it may be that this large number of clusters was partly an artifact of the DP prior. [Dunson \(2009\)](#) proposed a local partition process prior to reduce dimensionality in characterizing the curves, but this method does not produce easily interpretable functional clusters. Hence, it is appealing to use a more flexible global clustering prior that allows the number of clusters to instead converge to a finite constant.

With this motivation, we focus on the generalized Dirichlet process (GDP) introduced by [Hjort \(2000\)](#), denoted $\text{GDP}(a, b, G_0)$. The GDP corresponds to a stick-breaking prior with $K = \infty$, $a_k = a$ and $b_k = b$ for all k . When compared against the Poisson–Dirichlet process, the GDP has quite distinct properties.

THEOREM 1. *Let Z_n be the number of distinct observations in a sample of size n from a distribution G , where $G \sim \text{GDP}(a, b, G_0)$. The expected number of clusters $\mathbf{E}(Z_n)$ is given by*

$$(5) \quad \mathbf{E}(Z_n) = \sum_{i=1}^n \frac{ia\Gamma(a+b)\Gamma(b+i-1)}{\Gamma(b)\Gamma(a+b+i) - \Gamma(a+b)\Gamma(b+i)}.$$

The proof can be seen in [Appendix A](#). Note that for $a = 1$, this expression simplifies to $\mathbf{E}(Z_n) = \sum_{i=1}^n \frac{b}{b+i-1} \sim o(\log n)$, a well-known result for the Dirichlet process [[Antoniak \(1974\)](#)]. Letting $W_n = Z_n - Z_{n-1}$ denote the change in the number of clusters in adding the n th individual to a sample with $n - 1$ subjects, Stirling’s approximation can be used to show that

$$\mathbf{E}(W_n) = \frac{na\Gamma(a+b)\Gamma(b+n-1)}{\Gamma(b)\Gamma(a+b+n) - \Gamma(a+b)\Gamma(b+n)} \approx C(a, b)n^{-a},$$

where $C(a, b) = \{a\Gamma(a+b)/\Gamma(b)\} \exp\{-2(a+1)\}$. Hence, $\mathbf{E}(W_n) \rightarrow 0$ as $n \rightarrow \infty$ and new clusters become increasingly rare as the sample size increases. Note that for $a \leq 1$, the number of clusters will grow slowly but without bound

as n increases, with $E(Z_n) \rightarrow \infty$. The rate of growth in this case is proportional to n^{1-a} , which is similar to what is obtained by using the Poisson Dirichlet prior [Sudderth and Jordan (2009)]. However, when $a > 1$ the expected number of clusters instead converges to a finite constant, which is a remarkable difference compared with the Dirichlet and Poisson–Dirichlet process. As mentioned above, there may be a number of practical advantages to bounding the number of clusters. In addition, a finite bound on the number clusters seems to be more realistic in many applications, including the original species sampling applications that motivated much of the early development in this area [McCloskey (1965), Pitman (1995)].

In order to gain further insight into the clustering structure induced by the $GDP(a, b, G_0)$, we present in Figure 2 the relationship between the size of the largest cluster and the mean number of clusters in the partition (left panel), and the mean cluster size and the number of clusters (right panel) for a sample of size $n = 1000$. The numbers presented in this figure are based on the results from 20,000 simulations from the stick-breaking prior. Each continuous line corresponds to a combination of shape parameters such that $a/(a + b)$ is constant, while the dashed line in the plots corresponds to the combinations available under a Dirichlet process. The plots demonstrate that the additional parameter in the GDP allows us to simultaneously control the number of clusters and the relative size of the clusters, increasing the flexibility of the model as a clustering procedure.

The previous discussion focused on the impact of the prior distribution for the mixture weights on the clustering structure. Another important issue in the specification of the model is the selection of the baseline measure G_0 . Note that in the

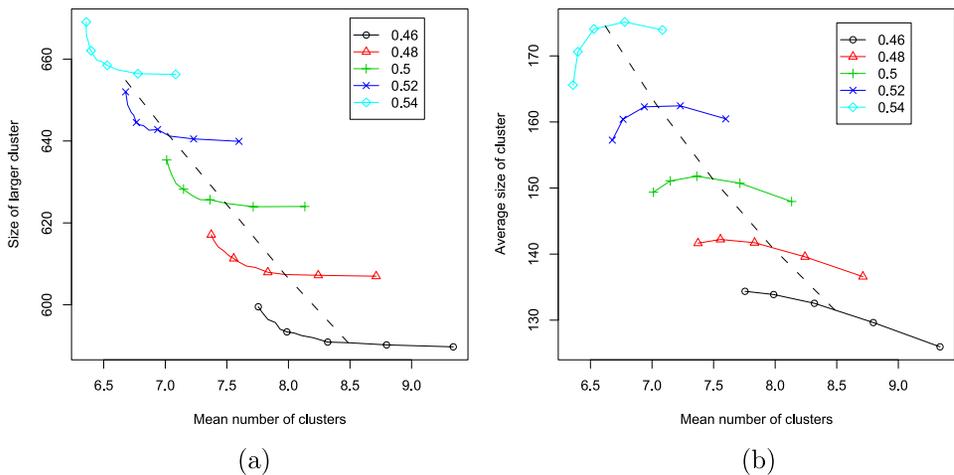


FIG. 2. Clustering structure induced by a $GDP(a, b, G_0)$ for a sample of size $n = 1000$. Panel (a) shows the relationship between the size of the largest cluster and the mean number of clusters for different GDPs, where each curve shares a common $E(u_k) = a/(a + b)$. Panel (b) shows the relationship between the average cluster size and the mean number of clusters. The dashed lines correspond to the combinations available under a standard Dirichlet process.

functional clustering setting $\boldsymbol{\vartheta}_k = (\boldsymbol{\theta}_k^*, \sigma_k^{*2}, \boldsymbol{\Lambda}_k^*)$ and, therefore, a computationally convenient choice that is in line with our previous discussion on basis selection and zero-inflated priors is to write

$$(6) \quad G_0(\boldsymbol{\theta}, \sigma^2, \boldsymbol{\Lambda}) = \mathbf{N}(\boldsymbol{\theta}|\mathbf{0}, \sigma^2\boldsymbol{\Omega}) \times \text{IGam}(\sigma^2|v_1, v_2) \times \prod_{s=1}^p \text{Ber}(\lambda_s|\gamma).$$

A prior of this form allows differential adaptive smoothing for each cluster in the data; the level of smoothness is controlled by γ (the prior probability of inclusion for each of the spline coefficients) and, therefore, it is convenient to assign to it a hyperprior such as $\gamma \sim \text{Beta}(\eta_1, \eta_2)$.

4. Functional clustering in nested designs. Consider now the case where multiple curves are collected for each subject in the study. In this case, the observations consist of ordered pairs (y_{ijt}, x_{ijt}) where

$$y_{ijt} = f_{ij}(x_{ijt}) + \varepsilon_{ijt},$$

where f_{ij} is the j th functional replicate for subject i , with $i = 1, \dots, I$, $j = 1, \dots, n_i$ and $t = 1, \dots, T_{ij}$. For example, in the EPS, f_{ij} is the measurement error-corrected smooth trajectory in the progesterone metabolite PdG over the j th menstrual cycle from woman i , with t indexing the sample number and x_{ijt} denoting the day within the i, j menstrual cycle relative to a marker of ovulation day.

A natural extension of (3) to nested designs arises by modeling the expected evolution of progesterone in time for cycle j of woman i as $f_{ij} = \mathbf{B}(\mathbf{x}_{ij})\boldsymbol{\theta}_{ij}$ and using a hierarchical model for the set of curve-specific parameters $\{\boldsymbol{\theta}_{ij}\}$ in order to borrow information across subjects and/or replicates. In the following subsections, we introduce two alternative nonparametric hierarchical priors that avoid parametric assumptions on the distribution of the basis coefficients, while inducing hierarchical functional clustering.

4.1. *Mean-curve clustering.* As a first approach, we consider a Gaussian mixture model, which characterizes the basis coefficients for functional replicate j from subject i as conditionally independent draws from a Gaussian distribution with subject-specific mean and variance, in the spirit of Booth, Casella and Hobert (2008):

$$(7) \quad \begin{aligned} \mathbf{y}_{ij}|\boldsymbol{\theta}_{ij}, \sigma_i &\sim \mathbf{N}(\mathbf{B}(\mathbf{x}_{ij})\boldsymbol{\theta}_{ij}, \sigma_i^2\mathbf{I}), & \boldsymbol{\theta}_{ij}|\boldsymbol{\theta}_i, \boldsymbol{\Lambda}_i, \sigma_i^2 &\sim G_i, \\ G_i &= \mathbf{N}(\boldsymbol{\Lambda}_i\boldsymbol{\theta}_i, \sigma_i^2\boldsymbol{\Sigma}), \end{aligned}$$

where $\boldsymbol{\Lambda}_i, \boldsymbol{\theta}_i, \sigma_i^2$ are as described in expression (3) and $\boldsymbol{\Sigma}$ is a $p \times p$ covariance matrix. In this model, the average curve for subject i is obtained as $\mathbf{E}\{f_{ij}(x)|\boldsymbol{\Lambda}_i, \boldsymbol{\theta}_i, \sigma_i^2\} = \mathbf{B}(x)\boldsymbol{\Lambda}_i\boldsymbol{\theta}_i$, with $\boldsymbol{\Lambda}_i$ providing a mechanism for subject-specific basis selection, so that the curves from subject i only depend on the basis

functions corresponding to nonzero diagonal elements of Λ_i . The variability in the replicate curves for the same subject is controlled by $\sigma_i^2 \Sigma$, with the subject-specific multiplier allowing subjects to vary in the degree of variability across the replicates. The need to allow such variability is well justified in the hormone curve application.

In order to borrow information across women, we need a hyperprior for the woman specific parameters $\{(\Lambda_i, \sigma_i^2, \theta_i)\}_{i=1}^I$. Since we are interested in clustering subjects, a natural approach is to specify this hyperprior nonparametrically through a generalized Dirichlet process centered around the baseline measure in (6), just as we did for the single curve case. This yields

$$(\theta_i, \sigma_i^2, \Lambda_i) | G \sim G, \quad G \sim \text{GDP}(a, b, G_0)$$

with G_0 given in (6). Since the distribution G is almost surely discrete, the model identifies clusters of women with similar average curves. This is clearer if we integrate the curve-specific coefficients $\{\theta_{ij}\}$ and the unknown distribution G out of the model to obtain

$$(8) \quad \mathbf{y}_{i1}, \dots, \mathbf{y}_{in_i} | \{w_k\}, \{\theta_k^*\}, \{\sigma_k^{*2}\}, \{\Lambda_k^*\} \\ \sim \sum_{k=1}^K w_k \left\{ \prod_{j=1}^{n_i} \text{N}(\mathbf{B}(\mathbf{x}_{ij}) \Lambda_k^* \theta_k^*, \sigma_k^{*2} (\mathbf{I} + \Sigma)) \right\}.$$

By incorporating the distribution of the selection matrices $\Lambda_1, \dots, \Lambda_I$ in the random distribution G , this model allows for a different smoothing pattern for each cluster of curves. This is an important difference with a straight generalization of the model in Ray and Mallick (2006), who instead treat the selection matrix as a hyperparameter in the baseline measure G_0 and therefore induce a common smoothing pattern across all clusters.

The model is completed by assigning priors for the hyperparameters. For the random effect variances we take inverse-Wishart priors:

$$\Omega \sim \text{IWis}(\nu_\Omega, \Omega_0), \quad \Sigma \sim \text{IWis}(\nu_\Sigma, \Sigma_0).$$

In the spirit of the unit information priors [Paciorek (2006)], the hyperparameters for these priors can be chosen so that Ω_0 and Σ_0 are proportional to

$$\sum_{i=1}^I \sum_{j=1}^{n_i} \mathbf{B}(\mathbf{x}_{ij})' \mathbf{B}(\mathbf{x}_{ij}).$$

Finally, the concentration parameters a and b are given gamma priors $a \sim \text{Gam}(\kappa_a, \tau_a)$ and $b \sim \text{Gam}(\kappa_b, \tau_b)$ and the probability of inclusion γ is assigned a beta prior, $\gamma \sim \text{Beta}(\eta_1, \eta_2)$.

4.2. *Distribution-based clustering.* Because the subject-specific distributions $\{G_i\}_{i=1}^I$ were assumed to be Gaussian and the nonparametric prior was placed on their means, the model in the previous section clusters subjects based on their average profile. However, as we discussed in Section 1, clustering based on the mean profiles might be misleading in studies such as the EPS in which there are important differences among subjects in not only the mean curve but also the distribution about the mean. In hormone curve applications, it is useful to identify clusters of trajectories over the menstrual cycle to study variability in the curves and identify outlying cycles that may have reproductive dysfunction. It is also useful to cluster women based not simply on the average curve but on the distribution of curves. With this motivation in mind, we generalize our hierarchical nonparametric specification to construct a model with additional parameters that enables clustering based on more than proximity between the mean curves. This generalization further enables clustering within subjects as well as subjects.

To motivate our nonparametric construction, consider first the simpler case in which there are only two types of curves in each cluster of women (say, normal and abnormal), so that it is natural to model the subject-specific distribution as a two-component mixture where

$$(9) \quad \begin{aligned} & \mathbf{y}_{ij} | \varpi_i, \mathbf{\Lambda}_{1i}, \boldsymbol{\theta}_{1i}, \sigma_{1i}^2, \mathbf{\Lambda}_{2i}, \boldsymbol{\theta}_{2i}, \sigma_{2i}^2 \\ & \sim \varpi_i \mathbf{N}(\mathbf{B}(\mathbf{x}_{ij}) \mathbf{\Lambda}_{1i} \boldsymbol{\theta}_{1i}, \sigma_{1i}^2 \mathbf{I}) + (1 - \varpi_i) \mathbf{N}(\mathbf{B}(\mathbf{x}_{ij}) \mathbf{\Lambda}_{2i} \boldsymbol{\theta}_{2i}, \sigma_{2i}^2 \mathbf{I}), \end{aligned}$$

ϖ_i can be interpreted as the proportion of curves from subject i that are in group 1 (say, normal), and $(\boldsymbol{\theta}_{1i}, \sigma_{1i}^2)$ are the parameters that describe curves from a normal cycle, $\mathbf{\Lambda}_{1i}$ is a diagonal variable selection matrix for the normal cycles that zeroes out unnecessary coefficients to avoid overfitting, $(\boldsymbol{\theta}_{2i}, \sigma_{2i}^2)$ are the parameters describing the curves from an abnormal cycle, and $\mathbf{\Lambda}_{2i}$ is the variable selection matrix for the abnormal cycles. Note that in this case we have not one but two variance parameters for each individual, which provide additional flexibility by allowing each cluster of curves to present a different level of observational noise. This feature is desirable in the EPS because, for a given woman, observational noise in abnormal cycles tends to be larger than in normal cycles.

Under this formulation, the subject-specific distribution is described by the vector of parameters $(\varpi_i, \mathbf{\Lambda}_{1i}, \boldsymbol{\theta}_{1i}, \sigma_{1i}^2, \mathbf{\Lambda}_{2i}, \boldsymbol{\theta}_{2i}, \sigma_{2i}^2)$, and clustering subjects could be accomplished by clustering these vectors. We can accomplish this by using another mixture model that mimics (2) and (8), so that

$$(10) \quad \begin{aligned} & \mathbf{y}_{i1}, \dots, \mathbf{y}_{in_i} | \{\pi_k\}, \{\varpi_k\}, \{\boldsymbol{\theta}_{1k}^*\}, \{\sigma_{1k}^{*2}\}, \{\mathbf{\Lambda}_{1k}^*\}, \{\boldsymbol{\theta}_{2k}^*\}, \{\sigma_{2k}^{*2}\}, \{\mathbf{\Lambda}_{2k}^*\} \\ & \sim \sum_{k=1}^K \pi_k \prod_{j=1}^{n_i} \{ \varpi_k \mathbf{N}(\mathbf{B}(\mathbf{x}_{ij}) \mathbf{\Lambda}_{1k}^* \boldsymbol{\theta}_{1k}^*, \sigma_{1k}^{*2} \mathbf{I}) \\ & \quad + (1 - \varpi_k) \mathbf{N}(\mathbf{B}(\mathbf{x}_{ij}) \mathbf{\Lambda}_{2k}^* \boldsymbol{\theta}_{2k}^*, \sigma_{2k}^{*2} \mathbf{I}) \}. \end{aligned}$$

To formulate our Bayesian nonparametric model for clustering, we start by rewriting (10) as a general mixture model where

$$(11) \quad \mathbf{y}_{ij} | \boldsymbol{\theta}_{ij}, \sigma_{ij}^2, \boldsymbol{\Lambda}_{ij} \sim \mathbf{N}(\mathbf{B}(\mathbf{x}_{ij}) \boldsymbol{\Lambda}_{ij} \boldsymbol{\theta}_{ij}, \sigma_{ij}^2 \mathbf{I}), \quad \boldsymbol{\theta}_{ij}, \sigma_{ij}^2, \boldsymbol{\Lambda}_{ij} | G_i \sim G_i$$

and G_i is a discrete distribution which is assigned a nonparametric prior. Note that this is analogous to the formulation in (7), but by replacing the Gaussian distribution with a random distribution with a nonparametric prior we are modeling the within-subject variability by clustering curves into groups with homogeneous shape.

Now, we need to define a prior over the collection $\{G_i\}_{i=1}^I$ that induces clustering among the distributions. For example, we could use a discrete distribution whose atoms are in turn random distributions, for example,

$$G_i \sim \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*},$$

where $\pi_k = v_k \prod_{s < k} (1 - v_s)$, $v_k \sim \text{Beta}(a_1, b_1)$ and $G_k^* \sim \text{GDP}(a_2, b_2, G_0)$ independently. This implies that

$$G_k^* = \sum_{l=1}^{\infty} \varpi_{lk} \delta_{(\boldsymbol{\theta}_{lk}^*, \sigma_{lk}^{2*}, \boldsymbol{\Lambda}_{lk}^*)}, \quad (\boldsymbol{\theta}_{lk}^*, \sigma_{lk}^{2*}, \boldsymbol{\Lambda}_{lk}^*) \sim G_0,$$

with $\varpi_{lk} = u_{lk} \prod_{s < l} (1 - u_{sk})$ and $u_{lk} \sim \text{Beta}(a_2, b_2)$ and G_0 as in (6). Therefore, if we were to replace the collection $\{G_k^*\}_{k=1}^{\infty}$ with random discrete distributions with only two atoms and we were to integrate over the random distributions $\{G_i\}_{i=1}^I$, this model would be equivalent to (10) with $K = \infty$.

This model on the collection $\{G_i\}_{i=1}^I$ is a generalization of the nested Dirichlet process introduced in Rodríguez, Dunson and Gelfand (2008) and, as with other models based on nested nonparametric processes, interesting special cases can be obtained by considering the limit of the precision parameters. For example, letting $b_2 \rightarrow 0$ while keeping a_2 fixed induces a model where all menstrual cycles within a woman are assumed to have the same profile, and subjects are clustered according to their mean cycle. Such a model is equivalent to the one obtained by taking $\boldsymbol{\Sigma} \rightarrow 0$ in (7). On the other hand, by letting $b_1 \rightarrow \infty$ while keeping a_1 constant, we obtain a model where all subjects are treated as different and menstrual cycles are clustered within each women. In this case, information is borrowed across the menstrual cycles of each women, but not across women.

Intuitively, we can think of the model we just described as first clustering curves within a subject via a model-based version of k -means applied to the subject-specific basis coefficients. Then, two subjects having similar distributions of curve clusters will be clustered together. Because we use a joint hierarchical model, these stages are done simultaneously.

Again, the model is completed by specifying prior distributions on the remaining parameters. As before, we let $\boldsymbol{\Omega} \sim \text{IWis}(v_{\boldsymbol{\Omega}}, \boldsymbol{\Omega}_0)$, $v_2 \sim \text{Gam}(\rho, \psi)$ and

$\gamma \sim \text{Beta}(\eta_1, \eta_2)$, providing a conditionally conjugate specification amenable for simple computational implementation. Finally, for the precision priors of the GDPs we set

$$\begin{aligned} a_1 &\sim \text{Gam}(\kappa_{a_1}, \tau_{a_1}), & b_1 &\sim \text{Gam}(\kappa_{b_1}, \tau_{b_1}), \\ a_2 &\sim \text{Gam}(\kappa_{a_2}, \tau_{a_2}), & b_2 &\sim \text{Gam}(\kappa_{b_2}, \tau_{b_2}). \end{aligned}$$

5. Computation. As is commonplace in Bayesian inference, we resort to Markov chain Monte Carlo (MCMC) algorithms [Robert and Casella (1999)] for computation in our functional clustering models. Given an initial guess for all unknown parameters in the model, the algorithms proceed by sequentially sampling blocks of parameters from their full conditional distributions. In particular, we design our algorithms using truncated versions of the GDP and the nested GDP, where a large but finite number of atoms are used to approximate the nonparametric mixture distributions. Well-known results on the convergence of truncations as the number of atoms grows that were originally presented in Ishwaran and James (2001) and Rodríguez, Dunson and Gelfand (2008) can be directly extended to this problem (see Appendix B). Furthermore, because most components of the model are conditionally conjugate, most of the full conditional distributions can be directly sampled using Gibbs steps. Full details of the algorithm can be seen in the online supplementary materials [Rodríguez and Dunson (2014)].

Convergence of the MCMC algorithms was assessed using the multi-chain method described in Gelman and Rubin (1992), which was applied to monitor the (unnormalized) posterior distribution, as well as the number of occupied clusters on each level of the model and the Frobenius norm of the matrices Ω and Σ .

6. The Early Pregnancy Study. Progesterone plays a crucial role in controlling different aspects of reproductive function in women, from fertilization to early development and implantation. Therefore, understanding the variability of hormonal profiles across the menstrual cycle and across subjects is important in understanding mechanisms of infertility and early pregnancy loss, as well as for developing approaches for identifying abnormal menstrual cycles and women for diagnostic purposes. Our data, extracted from the Early Pregnancy Study [Wilcox et al. (1998)], consists of log daily creatinine-corrected concentrations of pregnanediol-3-glucuronide (PdG) for 60 women along multiple menstrual cycles, measured in micrograms per milligram of creatinine ($\mu\text{g}/\text{ml Cr}$). We focus on 13-day intervals extending from 10 days before ovulation to 2 days after ovulation. According to the results in Dunson et al. (1999), this interval should include the fertile window of the menstrual cycle during which noncontracepting intercourse has a nonnegligible probability of resulting in a conception. Also, for this illustration we considered only nonconceptive cycles and women with at least four cycles in each record. Therefore, the number of curves per woman varies between 4 and 9.

We analyzed the EPS data using both the mean-based clustering model described in Section 4.1 and the distribution-based clustering model of Section 4.2. These models were fitted using the algorithms from Section 5. In the mean-based clustering algorithm, the GDP was truncated so that $K = 40$, while in the distribution-based algorithm the nested GDP was truncated so that $K = 40$ and $L = 30$. Although these numbers might seem large given the sample sizes involved, a large number of empty components are helpful in improving the mixing of the algorithms. In both cases, we used piecewise linear splines ($q = 1$) and $p = 13$ knots, corresponding to each of the days considered in the study. All inferences presented in this section are based on 100,000 samples obtained after a burn-in period of 10,000 iterations.

Prior distributions in the mean-based clustering algorithm were set as follows. For the concentration parameters, we used proper priors $a \sim \text{Gam}(3, 3)$ and $b \sim \text{Gam}(3, 3)$, and for the observational variance, we set $\sigma^2 \sim \text{IGam}(2, 0.04)$, so that $E(\sigma^2) = 0.04$. Note that setting $E(a) = 1$ a priori is natural because (as we discussed in Section 3) the asymptotic behavior of the expected number of clusters is different if $a > 1$ or $a \leq 1$. On the other hand, the prior mean for σ^2 is based on information from previous studies. To allow uncertainty in the probability of basis selection within the base measure, we let $\gamma \sim \text{Beta}(2, 4)$, implying that we expect about one-third of the spline basis functions to be used in any given cluster. Priors for the distribution-based clustering algorithm were chosen in a similar way, with $a_1 \sim \text{Gam}(3, 3)$, $b_1 \sim \text{Gam}(3, 3)$, $a_2 \sim \text{Gam}(3, 3)$ and $b_2 \sim \text{Gam}(3, 3)$, while for the baseline measure we picked a prior with the inclusion probabilities $\gamma \sim \text{Beta}(2, 4)$ and the prior on the group specific variances as given by $\text{IGam}(2, 0.04)$.

To better understand the effect of our prior on a , b , a_1 , b_1 , a_2 and b_2 , we show in Figure 3 the prior expected number of clusters implied by our choices. In particular, Figure 3(a) shows that, a priori, we expect between 1 and 8 clusters of subjects with high probability, with the most likely prior value being 3 clusters (note that this applies to both the mean-based and the distribution-based clustering methods). On the other hand, Figure 3(b) shows that, for a subject for which 7 cycles have been observed, we expect between 1 and 3 clusters with high probability. On the other hand, to assess the effect of the prior choices on the results, we conducted a small sensitivity analysis. In particular, the priors for a , b , a_1 , a_2 , b_1 and b_2 were replaced with exponential distributions with mean 2. The induced priors on the number of clusters have similar modes to our original specifications but have higher variability, placing substantially more mass on larger number of clusters. Although the posterior distribution for these parameters was somewhat affected, we did not see any substantial change in our posterior inferences on the clustering structure or the curve shapes (which are the main focus of our analysis). Similarly, we explored the effect of a $\text{Beta}(1, 1)$ prior on γ and an $\text{IGam}(2, 0.1)$ prior on σ^2 without any significant change in the estimates of the hormonal profiles.

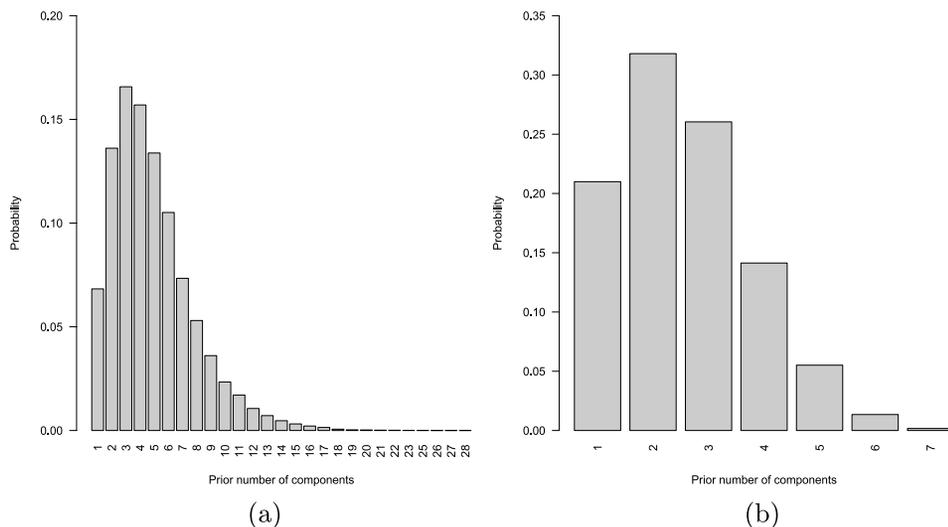


FIG. 3. Prior number of clusters implied by our specification of a , b , a_1 , b_1 , a_2 and b_2 . (a) $n = 60$ (subject-level clustering), (b) $n = 7$ (curve-level clustering).

We start our analysis by comparing the clustering structure generated by the mean-based and distribution-based models considered in Section 4. For this purpose, we show in Figures 4 and 5 heatmaps of the average pairwise clustering probability matrix under these two models. Entry (i, j) of the matrix contains the posterior probability that observations i and j are assigned to the same cluster. The black squares in the plots correspond to point estimates of the clustering structure obtained through the method described in Lau and Green (2007). In our case, the point estimate is obtained by minimizing a loss function that assigns equal weights to all pairwise misclassification errors. Therefore, the resulting plots provide information about the optimal clustering structure for the data as well as the uncertainty associated with it.

Figures 4 and 5 show that the structure of the clusters generated by both models are similar. For example, cluster 1 from mean-based clustering (counting from the bottom left corner of Figure 4) contains similar subjects as cluster 1 from distribution-based clustering (counting from the bottom left corner of Figure 5). The same is true for cluster 2 from both algorithms, and for cluster 5 from the mean-based clustering model and cluster 4 from the distribution-based model. There are also similarities in the structure of the outlier clusters (corresponding to the small clusters at the top and right of both heat maps). Both models assign subject 25 to a singleton cluster, while grouping subjects 21 and 53 into a pair, with subjects 40 and 51 in a separate pair. There are also some important differences between the two methods. Mean-based clustering groups subjects 36 and 43 together and assigns subject 3 to a different cluster (Figure 1 shows differences in the mean curve of subject 3), while distribution-based cluster groups subjects

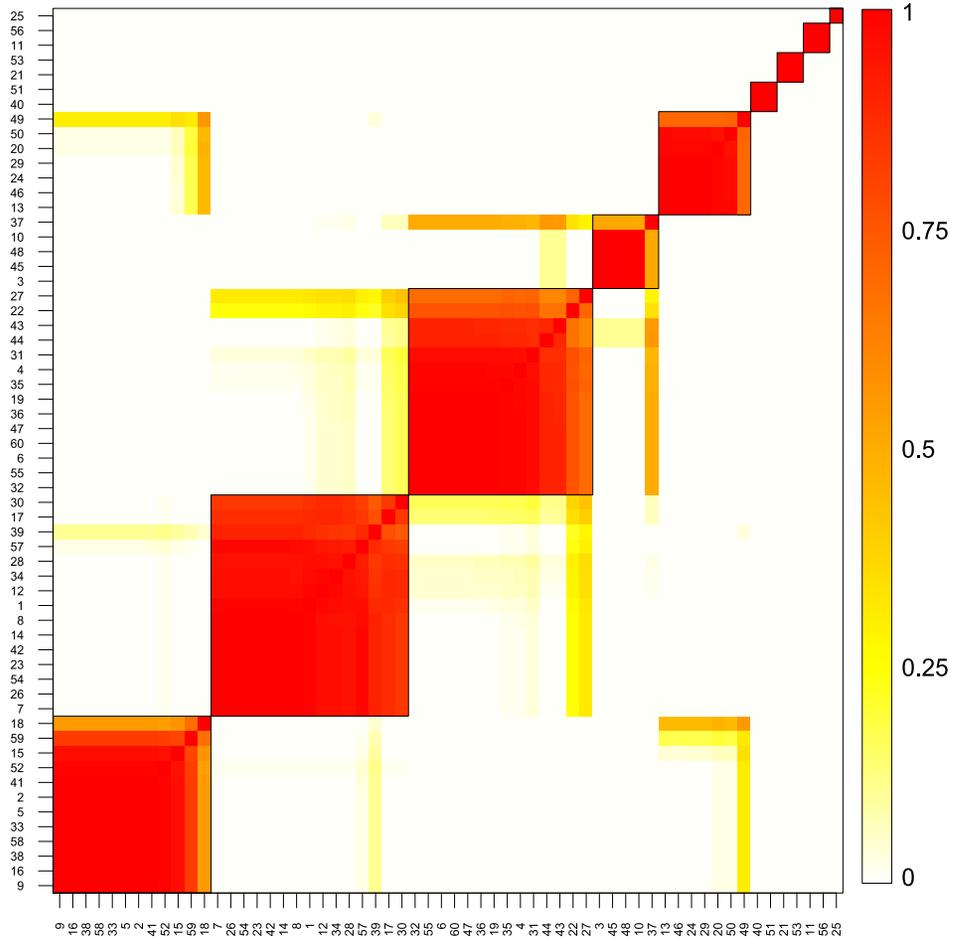


FIG. 4. Average incidence matrix, illustrating probabilities of joint pairwise classification for the 60 women in the EPS under the mean-curve clustering procedure described in Section 4.1. White corresponds to zero probability, while red corresponds to 1. The squares correspond to a point estimate of the cluster structure in the data.

3 and 43 together and assigns 36 to a different group. The behavior illustrates the robustness of distribution-based clustering, because except for a single outlier cycle for subject 43, the curves for this subject are much more similar to those of subject 3 than those of 36. Mean-based clustering also treats subjects 11 and 56 as singletons, while distribution-based clustering assigns them to the large cluster 4 (counting from the bottom left). Visually the raw profiles for subjects 11 and 56 are not very different from other profiles in cluster 4 [see panels (d), (e) and (f) of Figure 6], so this behavior seems reasonable.

Figure 7 shows reconstructed profiles under the distribution-based clustering model for some representative women in each of the main four groups. Most pro-

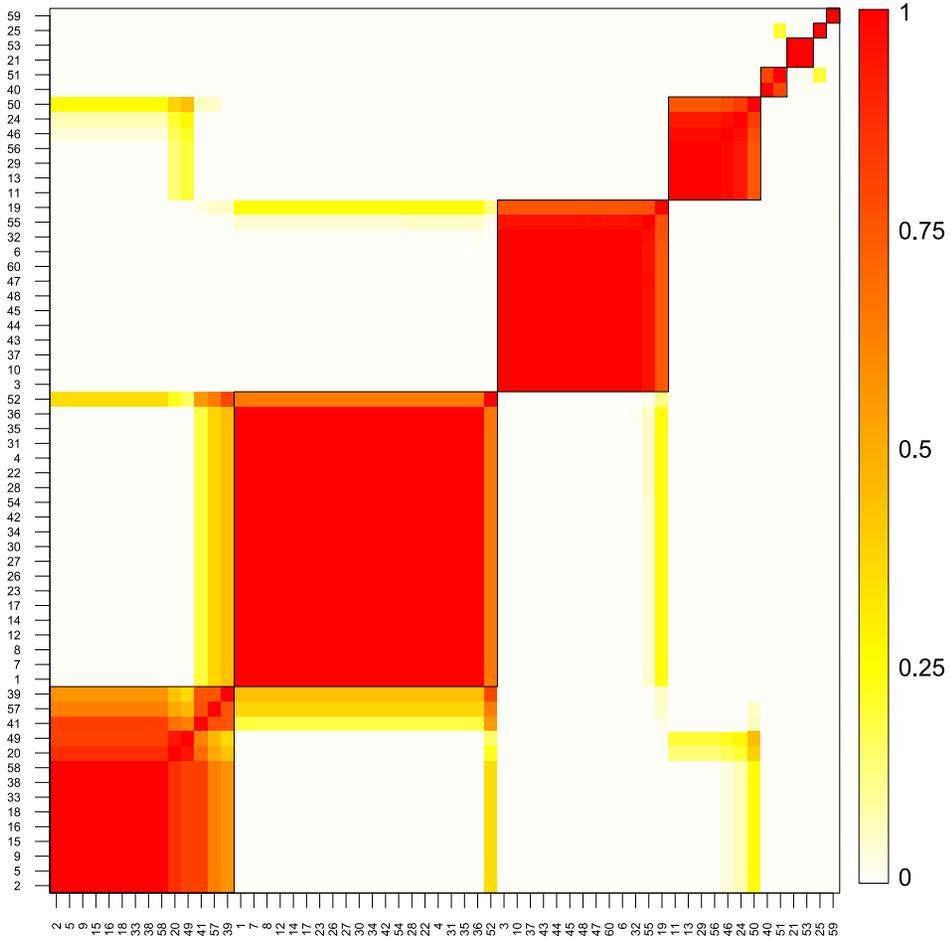


FIG. 5. Average incidence matrix, illustrating probabilities of joint pairwise classification for the 60 women in the EPS under the distribution-based clustering procedure described in Section 4.2. White corresponds to zero probability, while red corresponds to 1. The squares correspond to a point estimate of the cluster structure in the data.

files are flat before ovulation, when hormone levels start to increase. Also, in most clusters the profiles tend to be relatively consistent for any single woman. However, we can see some outliers, typically corresponding to elevated post-ovulation levels and/or early increases in the hormone levels. Cluster 3 corresponds to women with very low hormonal levels, even after ovulation. This group has few outliers, and those present are characterized by a slightly larger increase in progesterone after ovulation, which is still under $1 \mu\text{g/ml}$ Cr. Group 2 shows much more diversity in the hormonal profiles, as well as a slightly higher baseline level in progesterone level and an earlier rise in progesterone than group 3. Group 1 tends to show few outliers, and otherwise differs from the previous ones in a higher baseline level

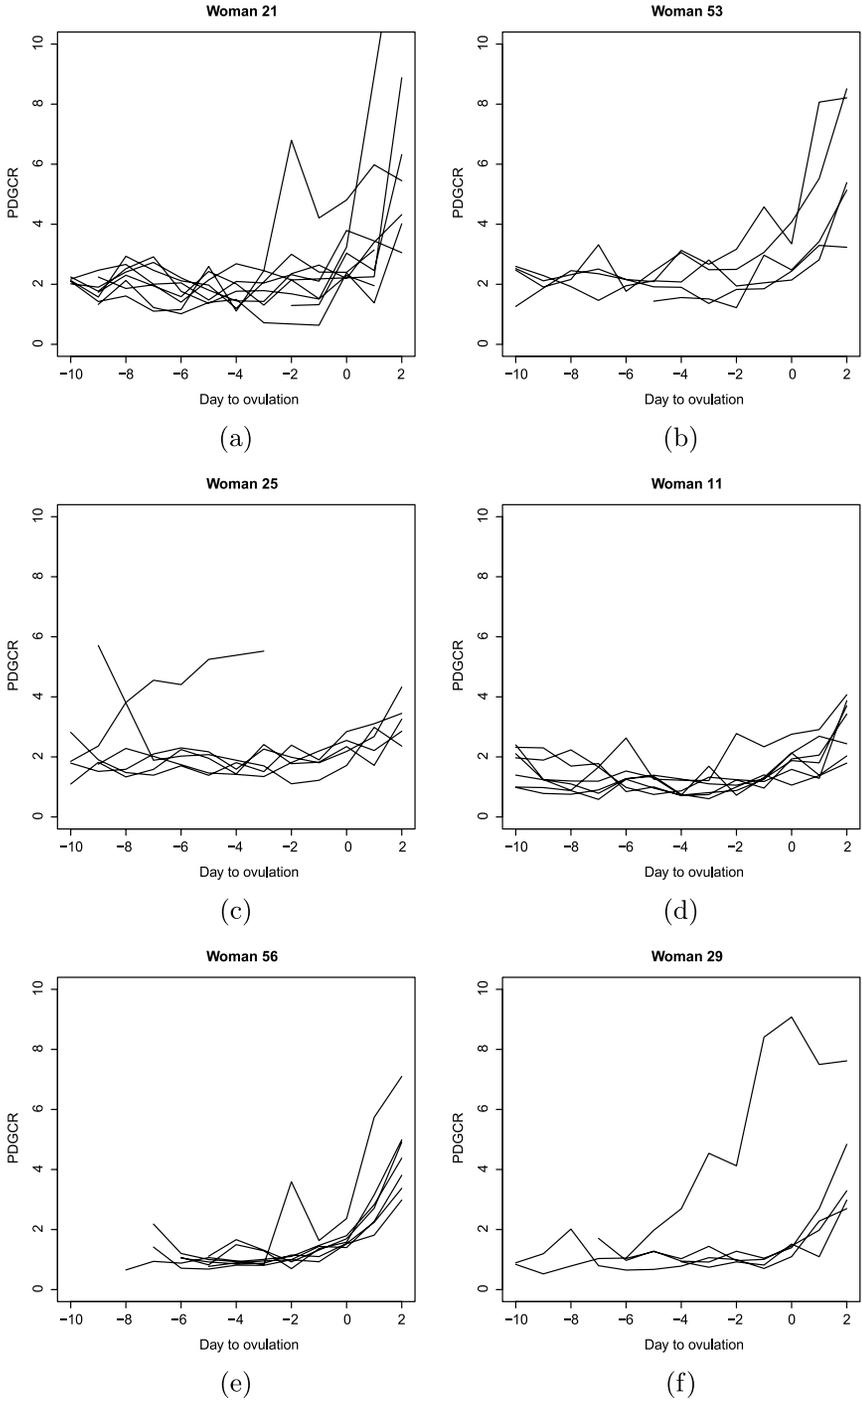


FIG. 6. Raw data associated with the hormonal profiles for three outlier subjects in the Early Pregnancy Study.

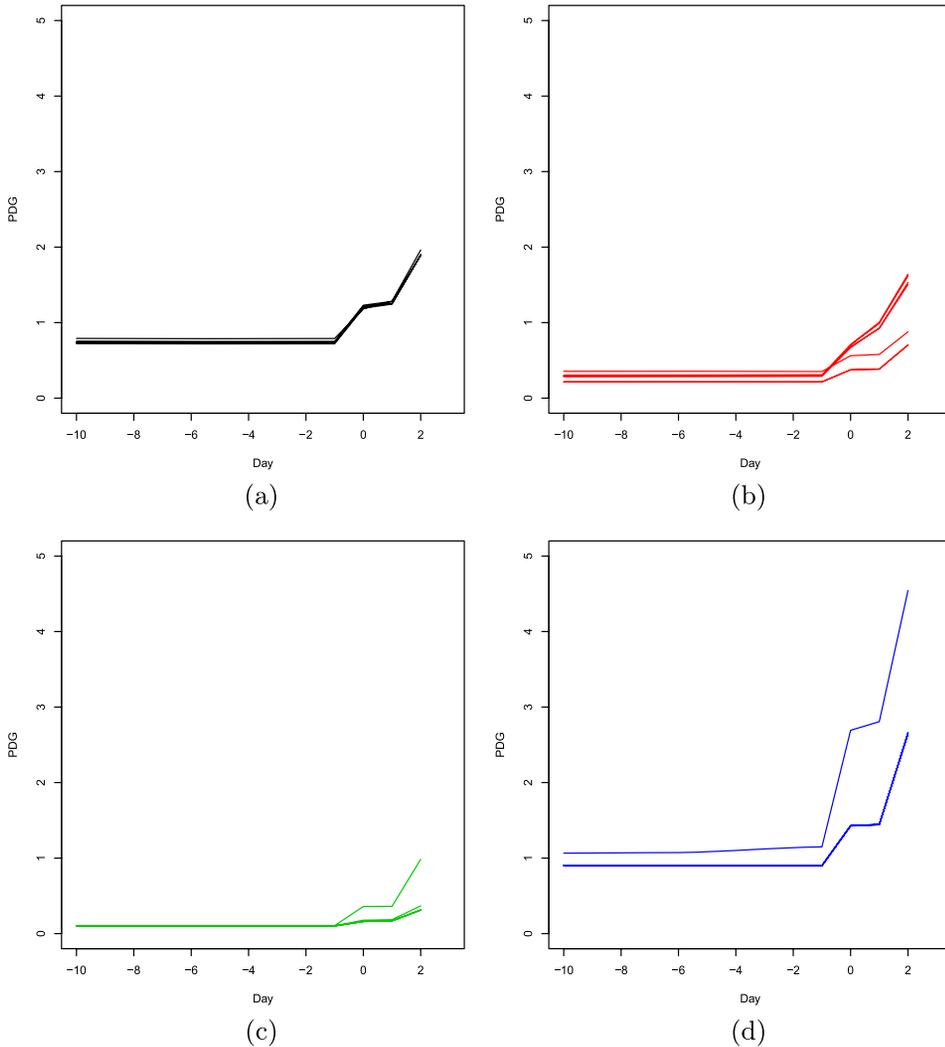


FIG. 7. Reconstructed profiles for some representative subjects in the study. Panel (a) corresponds to patient 9 (who was chosen from cluster 1, counting from the bottom left), panel (b) to patient 36 (who was chosen from cluster 2), panel (c) to patient 45 (who was chosen from cluster 3), and panel (d) corresponds to patient 13 (who was chosen from cluster 4).

and an early and very fast increase in progesterone. Finally, group 4 presents “normal” cycles with the highest baseline level of progesterone ($1 \mu\text{g/ml Cr}$) and the fastest increase in progesterone after ovulation, along with “abnormal” cycles with even higher baseline levels and very extreme levels of progesterone after ovulation (close to $5 \mu\text{g/ml Cr}$).

For comparison we also applied a simple functional clustering approach based on kernel smoothing and hierarchical clustering. We first fitted a Gaussian ker-

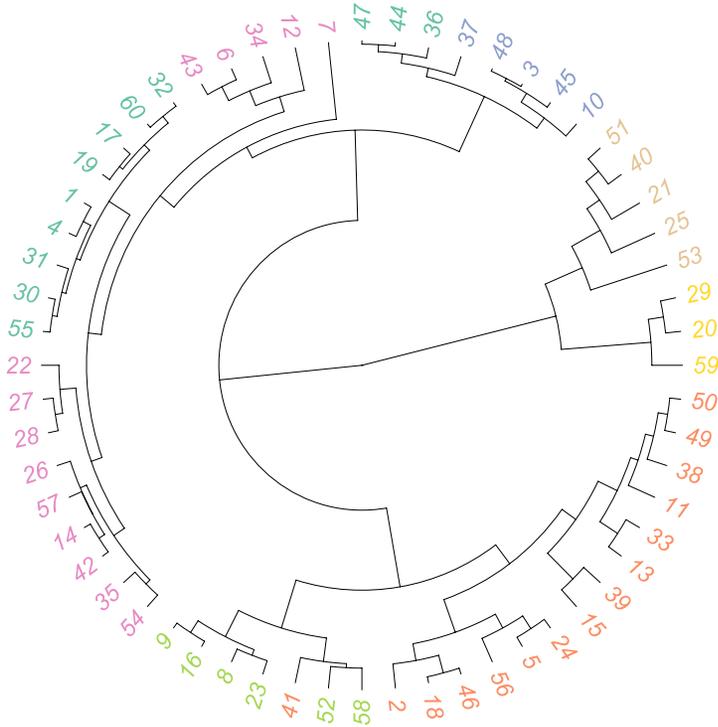


FIG. 8. Dendrogram for an alternative functional clustering algorithm that uses kernel smoothing to project all curves to a common set of grid knots and then applies complete-linkage hierarchical clustering and BIC to create a partition of the data.

nel smoother [Li and Racine (2004), Racine and Li (2004)] to each of the curves, generated fitted values over a common grid (in this example we used 6 equally-spaced knots), then computed the average predicted values for each subject at each point of the grid, and finally applied complete-linkage hierarchical clustering (as implemented in the R package `mclust`), with the number of clusters selected using BIC. This approach identified 7 clusters; Figure 8 shows the associated dendrogram, where colors are used to represent the clusters. Some of these clusters are similar to the ones identified by the mean-based clustering model. For example, the small cluster of 5 subjects (3, 10, 37, 45 and 48), corresponding to the fourth cluster from the bottom left in Figure 3, is perfectly identified by hierarchical clustering. Similarly, subjects 21, 25, 40, 51 and 53, which are all identified as outliers by the mean-based clustering model, are allocated to a single cluster by the hierarchical clustering approach. However, the majority of the clustering pattern generated by this method is quite different from the one obtained with mean-based or the distribution-based clustering. Furthermore, the number and structure of clusters depends heavily on the number and location of grid points used to interpolate the functions. For example, when the interpolation grid contains 5 knots we obtain

8 clusters, while 18 clusters are obtained when 11 knots are used for interpolation. Part of this difference could be due to the fact that, because most of the missing values tend to concentrate at the beginning of the cycle, many of the fitted functional values correspond to extrapolations rather than interpolations when a large number of knots are used.

We also note that posterior estimates of the precision parameters on the GDP suggest that a logarithmic rate of growth for the expected number of clusters might be reasonable for this data. For the mean-based clustering, the posterior mean for a was 1.06 and the 95% posterior symmetric credible interval was (0.65, 1.49), while the posterior mean for b was 0.77 with 95% credible interval (0.17, 1.72). For the distribution-based clustering, the corresponding estimates are 1.03 (0.62, 1.56) and 0.72 (0.20, 1.73) for a_1 and b_1 , and 1.12 (0.71, 1.43) and 0.27 (0.15, 1.51) for a_2 and b_2 .

Finally, we investigated the relationship between the clusters we identified using our distribution-based clustering algorithm and a series of covariates available for each subject, including age at the beginning of the study (Age), age of first menses (Age_Mense), average length of menses (Mense_Length), race (Race), body mass index (BMI), whether the subject had been taking contraceptives before the start of the study (EV_Pill) and whether the subject self reports as a consumer of marijuana during the period of the study (Marijuana_Use). To investigate this relationship, we used a classification tree, as implemented in the R package `rpart`. To simplify interpretation, we merged all the small outlier clusters into a single group (which we label cluster 5). The resulting classification tree (obtained using a Cp value of 0.03 for pruning) can be seen in Figure 9. Note that race and marijuana consumption do not appear to play any role in predicting the shape of the hormonal profiles. On the other hand, BMI, along with age, menses length and menses age, seems to have some predictive capability. For example, a BMI higher than 22 and an age of menses under 12 years old seem to predict progesterone profiles like the ones observed in cluster 3.

7. Discussion. We have presented two approaches to functional clustering in nested designs. These approaches look into different features of the nested samples and are therefore applicable in different circumstances. Our mean-based clustering approach is easier to interpret and provides an excellent alternative when within-subject samples are homogenous. However, when within-subject curves are heterogeneous, mean-based clustering can lead to biased results. Therefore, in studies such as the EPS, distribution-based models such as the one described here provide a viable alternative that acknowledges the heterogeneity in the function replicates from a subject.

One interesting insight that can be gathered from the results of the EPS data is that, for small numbers of functional replicates per subject and rare outliers, the effect of the distribution-based clustering is to perform clustering based on the modal rather than the mean profile. That is, the distribution-based clustering model

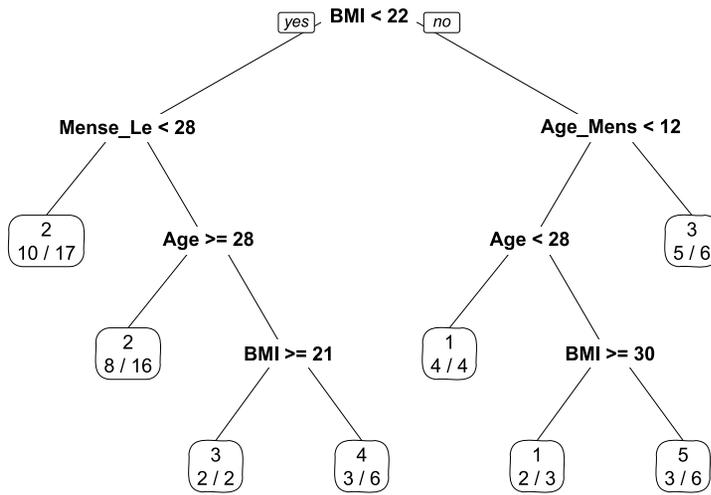


FIG. 9. Classification tree explaining cluster membership on the basis of six covariates: age at the beginning of the study (Age), age of first menses (Age_Mense), average length of menses (Mense_Length), race (Race), body mass index (BMI), whether the subject had been taking contraceptives before the start of the study (EV_Pill) and whether the subject self reports as a consumer of marijuana during the period of the study (Marijuana_Use).

is able to automatically discount the abnormal curves, leading to more appropriate clustering patterns if the effect of outliers needs to be removed. Naturally, this perceived advantage of the distribution-based clustering method implicitly assumes that abnormal curves should be discounted. Although that assumption is justified in our application, users should be aware of it when applying our method to other data sets.

APPENDIX A: PROOF OF THEOREM 1

Let $\theta_1^*, \theta_2^*, \dots$ be a sequence of independent and identically distributed samples from a random distribution G , which follows a $GDP(a, b, G_0)$ distribution. Also, let W_i be 1 if θ_i^* is different from every $\theta_1^*, \dots, \theta_{i-1}^*$, and zero otherwise. Clearly, $Z_n = \sum_{i=1}^n W_i$ is the number of distinct values among the first n samples from a $GDP(a, b, G_0)$. Hjort (2000) shows that

$$\begin{aligned}
 E(W_i) &= i \frac{E\{u(1-u)^{i-1}\}}{1 - E\{(1-u)^i\}} \\
 &= i \frac{((\Gamma(a+b))/(\Gamma(a)\Gamma(b)))((\Gamma(a+1)\Gamma(b+i-1))/(\Gamma(a+b+i)))}{1 - ((\Gamma(a+b))/(\Gamma(a)\Gamma(b)))((\Gamma(a)\Gamma(b+i))/(\Gamma(a+b+i)))} \\
 &= \frac{ia\Gamma(a+b)\Gamma(b+i-1)}{\Gamma(b)\Gamma(a+b+i) - \Gamma(a+b)\Gamma(b+i)},
 \end{aligned}$$

which completes the proof.

APPENDIX B: TRUNCATIONS OF GENERALIZED DIRICHLET PROCESSES

THEOREM 2. Assume that samples of n observations have been collected for each of J distributions and are contained in vector $\mathbf{y} = (\mathbf{y}'_1, \dots, \mathbf{y}'_J)$. Also, let

$$P^{\infty\infty}(\boldsymbol{\theta}) = \int \int P(\boldsymbol{\theta}|G_j)P^\infty(dG_j|Q)P^\infty(dQ),$$

$$P^{LK}(\boldsymbol{\theta}) = \int \int P(\boldsymbol{\theta}|G_j)P^L(dG_j|Q)P^K(dQ)$$

be, respectively, the prior distribution of the model parameters under the nested GDP model and its corresponding truncation after integrating out the random distributions, and $P^{\infty\infty}(\mathbf{y})$ and $P^{LK}(\mathbf{y})$ be the prior predictive distribution of the observations derived from these priors. Then

$$\int |P^{LK}(\mathbf{y}) - P^{\infty\infty}(\mathbf{y})| d\mathbf{y} \leq \int |P^{LK}(\boldsymbol{\theta}) - P^{\infty\infty}(\boldsymbol{\theta})| \leq \varepsilon^{LK}(\alpha, \beta),$$

where

$$\varepsilon^{LK}(\alpha, \beta) = \begin{cases} 4\left(1 - \left[1 - \left(\frac{b_1}{a_1 + b_1}\right)^{K-1}\right]^J\right), & \text{if } L = \infty, K < \infty, \\ 4\left(1 - \left[1 - \left(\frac{b_2}{a_2 + b_2}\right)^{L-1}\right]^{nJ}\right), & \text{if } L < \infty, K = \infty, \\ 4\left(1 - \left[1 - \left(\frac{b_1}{a_1 + b_1}\right)^{K-1}\right]^J \left[1 - \left(\frac{b_2}{a_2 + b_2}\right)^{L-1}\right]^{nJ}\right), & \text{if } L < \infty, K < \infty. \end{cases}$$

The proof is a direct extension of results in Ishwaran and James (2001, 2002) and Rodríguez, Dunson and Gelfand (2008) and it is omitted for reasons of space. This result is particularly important since it justifies the use of computational algorithms based on finite mixtures and allows us to choose adequate truncation levels.

Acknowledgments. We would like to thank Nils Hjort for his helpful comments and Allen Wilcox for providing access to the EPS data.

SUPPLEMENTARY MATERIAL

Supplement to “Functional clustering in nested designs: Modeling variability in reproductive epidemiology studies” (DOI: [10.1214/14-AOAS751SUPP.pdf](https://doi.org/10.1214/14-AOAS751SUPP.pdf)). The supplementary materials contain the details of the Markov chain Monte Carlo algorithm used to fit the models introduced in the paper.

REFERENCES

- ABRAHAM, C., CORNILLON, P. A., MATZNER-LØBER, E. and MOLINARI, N. (2003). Unsupervised curve clustering using B-splines. *Scand. J. Stat.* **30** 581–595. [MR2002229](#)
- ANTONIAK, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Ann. Statist.* **2** 1152–1174. [MR0365969](#)
- BIGELOW, J. L. and DUNSON, D. B. (2009). Bayesian semiparametric joint models for functional predictors. *J. Amer. Statist. Assoc.* **104** 26–36. [MR2663031](#)
- BOOTH, J. G., CASELLA, G. and HOBERT, J. P. (2008). Clustering using objective functions and stochastic search. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 119–139. [MR2412634](#)
- BRUMBACK, B. A. and RICE, J. A. (1998). Smoothing spline models for the analysis of nested and crossed samples of curves. *J. Amer. Statist. Assoc.* **93** 961–994. [MR1649194](#)
- CHIOU, J.-M. and LI, P.-L. (2007). Functional clustering and identifying substructures of longitudinal data. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69** 679–699. [MR2370075](#)
- DIMATTEO, I., GENOVESE, C. R. and KASS, R. E. (2001). Bayesian curve-fitting with free-knot splines. *Biometrika* **88** 1055–1071. [MR1872219](#)
- DUNSON, D. B. (2009). Nonparametric Bayes local partition models for random effects. *Biometrika* **96** 249–262. [MR2507141](#)
- DUNSON, D. B., BAIRD, D. D., WILCOX, A. J. and WEINBERG, C. R. (1999). Day-specific probabilities of clinical pregnancy based on two studies with imperfect measures of ovulation. *Hum. Reprod.* **14** 1835–1839.
- ESCOBAR, M. D. and WEST, M. (1995). Bayesian density estimation and inference using mixtures. *J. Amer. Statist. Assoc.* **90** 577–588. [MR1340510](#)
- FERGUSON, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1** 209–230. [MR0350949](#)
- FRALEY, C. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc.* **97** 611–631. [MR1951635](#)
- GARCÍA-ESCUADERO, L. A. and GORDALIZA, A. (2005). A proposal for robust curve clustering. *J. Classification* **22** 185–201. [MR2231171](#)
- GELMAN, A. and RUBIN, D. (1992). Inferences from iterative simulation using multiple sequences. *Statist. Sci.* **7** 457–472.
- HEARD, N. A., HOLMES, C. C. and STEPHENS, D. A. (2006). A quantitative study of gene regulation involved in the immune response of anopheline mosquitoes: An application of Bayesian hierarchical clustering of curves. *J. Amer. Statist. Assoc.* **101** 18–29. [MR2252430](#)
- HJORT, N. (2000). Bayesian analysis for a generalized Dirichlet process prior. Technical report, Univ. Oslo.
- ISHWARAN, H. and JAMES, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *J. Amer. Statist. Assoc.* **96** 161–173. [MR1952729](#)
- ISHWARAN, H. and JAMES, L. F. (2002). Approximate Dirichlet process computing in finite normal mixtures: Smoothing and prior information. *J. Comput. Graph. Statist.* **11** 508–532. [MR1938445](#)
- JAMES, G. M. and SUGAR, C. A. (2003). Clustering for sparsely sampled functional data. *J. Amer. Statist. Assoc.* **98** 397–408. [MR1995716](#)
- LAU, J. W. and GREEN, P. J. (2007). Bayesian model-based clustering procedures. *J. Comput. Graph. Statist.* **16** 526–558. [MR2351079](#)
- LI, J. (2005). Clustering based on a multilayer mixture model. *J. Comput. Graph. Statist.* **14** 547–568. [MR2170201](#)
- LI, Q. and RACINE, J. (2004). Cross-validated local linear nonparametric regression. *Statist. Sinica* **14** 485–512. [MR2059293](#)
- LIANG, F., PAULO, R., MOLINA, G., CLYDE, M. A. and BERGER, J. O. (2008). Mixtures of g priors for Bayesian variable selection. *J. Amer. Statist. Assoc.* **103** 410–423. [MR2420243](#)

- LUAN, Y. and LI, H. (2003). Clustering of time-course gene expression data using a mixed effects model with b-splines. *Bioinformatics* **19** 474–482.
- MCCLOSKEY, J. W. T. (1965). A model for the distribution of individuals by species in an environment. Ph.D. thesis, Michigan State Univ. [MR2615013](#)
- MEDVEDOVIC, M. and SIVAGANESAN, S. (2002). Bayesian infinite mixture model based clustering of gene expression profiles. *Bioinformatics* **18** 1194–1206.
- ONGARO, A. and CATTANEO, C. (2004). Discrete random probability measures: A general framework for nonparametric Bayesian inference. *Statist. Probab. Lett.* **67** 33–45. [MR2039931](#)
- PACIOREK, C. J. (2006). Misinformation in the conjugate prior for the linear model with implications for free-knot spline modelling. *Bayesian Anal.* **1** 375–383 (electronic). [MR2221270](#)
- PITMAN, J. (1995). Exchangeable and partially exchangeable random partitions. *Probab. Theory Related Fields* **102** 145–158. [MR1337249](#)
- PITMAN, J. (1996). Some developments of the Blackwell–MacQueen urn scheme. In *Statistics, Probability and Game Theory. Papers in Honor of David Blackwell* (T. S. Ferguson, L. S. Shapley and J. B. MacQueen, eds.) 245–267. IMS, Hayward, CA. [MR1481784](#)
- QUINTANA, F. A. and IGLESIAS, P. L. (2003). Bayesian clustering and product partition models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **65** 557–574. [MR1983764](#)
- RACINE, J. and LI, Q. (2004). Nonparametric estimation of regression functions with both categorical and continuous data. *J. Econometrics* **119** 99–130. [MR2041894](#)
- RAMONI, M. F., SEBASTIANI, P. and KOHANE, I. S. (2002). Cluster analysis of gene expression dynamics. *Proc. Natl. Acad. Sci. USA* **99** 9121–9126 (electronic). [MR1909705](#)
- RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer, New York. [MR2168993](#)
- RAY, S. and MALLICK, B. (2006). Functional clustering by Bayesian wavelet methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68** 305–332. [MR2188987](#)
- ROBERT, C. P. and CASELLA, G. (1999). *Monte Carlo Statistical Methods*. Springer, New York. [MR1707311](#)
- RODRIGUEZ, A. and DUNSON, D. B. (2014). Supplement to “Functional clustering in nested designs: Modeling variability in reproductive epidemiology studies.” DOI:10.1214/14-AOAS751SUPP.
- RODRÍGUEZ, A., DUNSON, D. B. and GELFAND, A. E. (2008). The nested Dirichlet process. *J. Amer. Statist. Assoc.* **103** 1131–1144. [MR2528831](#)
- RODRÍGUEZ, A., DUNSON, D. B. and GELFAND, A. E. (2009). Bayesian nonparametric functional data analysis through density estimation. *Biometrika* **96** 149–162. [MR2482141](#)
- SERBAN, N. and WASSERMAN, L. (2005). CATS: Clustering after transformation and smoothing. *J. Amer. Statist. Assoc.* **100** 990–999. [MR2201025](#)
- SETHURAMAN, J. (1994). A constructive definition of Dirichlet priors. *Statist. Sinica* **4** 639–650. [MR1309433](#)
- SMITH, M. and KOHN, R. (1996). Nonparametric regression using Bayesian variable selection. *J. Econometrics* **75** 317–343.
- SUDDERTH, E. B. and JORDAN, M. I. (2009). Shared segmentation of natural scenes using dependent Pitman–Yor processes. In *Advances in Neural Information Processing Systems* 21 (D. Koller, D. Schuurmans, Y. Bengio and L. Bottou, eds.).
- TARPEY, T. and KINATEDER, K. K. J. (2003). Clustering functional data. *J. Classification* **20** 93–114. [MR1983123](#)
- WAKEFIELD, J. C., ZHOU, C. and SELF, S. G. (2003). Modelling gene expression data over time: Curve clustering with informative prior distributions. In *Bayesian Statistics, 7 (Tenerife, 2002)* (J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith and M. West, eds.) 721–732. Oxford Univ. Press, New York. [MR2003536](#)

WILCOX, A. J., WEINBERG, C. R., O'CONNOR, J. F., BAID, D. D., SCHLATTERER, J. P., CANFIELD, R. E., ARMSTRONG, E. G. and NISULA, B. C. (1998). Incidence of early loss of pregnancy. *N. Engl. J. Med.* **319** 189–194.

YEUNG, K. Y. and RUZZO, W. L. (2001). Principal component analysis for clustering gene expression data. *Bioinformatics* **17** 763–774.

YEUNG, K. Y., FRALEY, C., MURUAN, A., RAFTERY, A. E. and RUZZO, W. L. (2001). Model-based clustering and data transformation for gene expression data. *Bioinformatics* **17** 977–987.

DEPARTMENT OF APPLIED MATHEMATICS
AND STATISTICS
UNIVERSITY OF CALIFORNIA, SANTA CRUZ
1156 HIGH STREET
MAILSTOP SOE2
SANTA CRUZ, CALIFORNIA 95064
USA
E-MAIL: abel@soe.ucsc.edu

DEPARTMENT OF STATISTICS
AND DECISION SCIENCES
DUKE UNIVERSITY
BOX 90251
DURHAM, NORTH CAROLINA 27708
USA