

Bayesian Clustering with Regression

Peter Müller

University of Texas M.D. Anderson Cancer Center, Houston, TX 77030, U.S.A.

Fernando Quintana

Departamento de Estadística, Pontificia Universidad Católica de Chile, Santiago, Chile.

Gary Rosner

University of Texas M.D. Anderson Cancer Center, Houston, TX 77030, U.S.A.

Summary. We consider clustering with regression, i.e., we develop a probability model for random partitions that is indexed by covariates. The motivating application is predicting time to progression for patients in a breast cancer trial. We proceed by reporting a weighted average of the responses of clusters of earlier patients. The weights should be determined by the similarity of the new patient's covariate with the covariates of patients in each cluster. We achieve the desired inference by defining a random partition model that includes a regression on covariates. Patients with similar covariates are a priori more likely to be clustered together. Posterior predictive inference in this model formalizes the desired prediction.

We build on product partition models (PPM). We define an extension of the PPM to include a regression on covariates by including in the cohesion function a new factor that increases the probability of experimental units with similar covariates to be included in the same cluster. We discuss implementations suitable for any combination of continuous, categorical, count and ordinal covariates.

1. Introduction

We develop a probability model for clustering with covariates, that is a probability model for partitioning a set of experimental units, where the probability of any particular partition is allowed to depend on covariates. The motivating application is inference in a clinical trial. The outcome is time to progression for breast cancer patients. The covariates include treatment dose, initial tumor burden, an indicator for menopause and more. We wish to define a probability model for clustering patients with the specific feature that patients with equal or similar covariates should be more likely to co-cluster than others.

Let $i = 1, \dots, n$, index experimental units, and let $\rho_n = \{S_1, \dots, S_k\}$ denote a partition of the n experimental units into k subsets S_j . When n is obvious from the context we omit the subindex and write ρ . Let x_i and y_i denote the covariates and response reported for the i -th unit. Let $x^n = (x_1, \dots, x_n)$ and $y^n = (y_1, \dots, y_n)$ denote the entire set of recorded covariates and data, and let $x_j^* = (x_i, i \in S_j)$ and $y_j^* = (y_i, i \in S_j)$ denote covariates and data corresponding to units in the j -th cluster. Sometimes it is convenient to introduce indicators $e_i \in \{1, \dots, k\}$ with $e_i = j$ if $i \in S_j$, and use (k, e_1, \dots, e_n) to describe the partition. We call a probability model $p(\rho)$ a clustering model, excluding in particular purely constructive definitions of clustering as a specific algorithm (without reference to probability models). Many clustering models include, implicitly or explicitly, a sampling model $p(y \mid \rho)$. Probability models for $p(\rho)$ and inference for clustering have been extensively discussed over the past few years. See Quintana (2006) for a recent review. In this paper we are interested in adding a regression to replace $p(\rho)$ with $p(\rho \mid x)$.

We focus on the product partition models (PPM). The PPM (Hartigan, 1990; Barry and Hartigan, 1993) constructs $p(\rho)$ by introducing cohesion functions $c(A) \geq 0$ for $A \subset \{1, \dots, n\}$ that measure how tightly grouped the elements in

A are thought to be, and defines a probability model for a partition ρ and data y as

$$p(\rho) \propto \prod c(S_i) \quad \text{and} \quad p(y | \rho) = \prod_{j=1}^k p_j(y_{S_j}). \quad (1)$$

Model (1) is conjugate. The posterior $p(\rho | y)$ is again in the same product form.

Alternatively, the species sampling model (SSM) (Pitman, 1996; Ishwaran and James, 2003) defines an exchangeable probability model $p(\rho)$ that depends on ρ only indirectly through the cardinality of the partitioning subsets, $p(\rho) = p(|S_1|, \dots, |S_k|)$. The SSM can be alternatively characterized by a sequence of predictive probability functions (PPF) that describe how individuals are sequentially assigned to either already formed clusters or to start new ones. The choice of the PPF is not arbitrary. One has to make sure that a sequence of random variables that are sampled by iteratively applying the PPF is exchangeable. The popular Dirichlet process (DP) model (Ferguson, 1973; Antoniak, 1974) is a special case of a SSM. Moreover, the marginal distribution that a DP induces on partitions is also a PPM with cohesions $c(A) = M \times (|A| - 1)!$ (Quintana and Iglesias, 2003; Dahl, 2003). Here M denotes the total mass parameter of the DP prior.

Model based clustering (Banfield and Raftery, 1993; Dasgupta and Raftery, 1998; Fraley and Raftery, 2002) implicitly defines a probability model on clustering by assuming a mixture model

$$p(y_i | \eta, k) = \sum_{j=1}^k \tau_j p_j(y_i | \theta_j),$$

where $\eta = (\theta_1, \dots, \theta_k, \tau_1, \dots, \tau_k)$ are the parameters of a size k mixture model. Together with a prior $p(k)$ on k and $p(\theta, \tau | k)$, the mixture defines a probability model on clustering. Consider the equivalent hierarchical model

$$p(y_i | e_i = j, k, \eta) = p_j(y_i | \theta_j) \quad \text{and} \quad Pr(e_i = j | k, \eta) = \tau_j. \quad (2)$$

The implied posterior distribution on (e_1, \dots, e_n) and k defines a probability model on ρ_n . Richardson and Green (1997) develop posterior simulation strategies for mixture of normal models. Green and Richardson (1999) discuss the relationship to DP mixture models.

Especially in the context of spatial data, a popular clustering model is based on Voronoi tessellations (Green and Sibson, 1978; Okabe et al., 2000; Kim et al., 2005). For a review see, for example, Denison et al. (2002). Although not usually considered as a probability model on clustering, it can be defined as such by means of the following construction. We include it in this brief review of clustering models because of its similarity to the model proposed in this paper. Assume x_i are spatial coordinates for observation i . Let $d(x_1, x_2)$ denote a distance measure, for example Euclidean distance in \mathbb{R}^2 . Clusters are defined with latent cluster centers T_j , $j = 1, \dots, k$, by setting $e_i = \arg \min_j d(x_i, T_j)$. That is, a partition is defined by allocating each observation to the cluster whose center T_j is closest to x_i . Conditional on T , cluster allocation is deterministic. Marginalizing with respect to a prior on T we could define a probability model $p(\rho | k)$.

In this paper we build on the PPM (1) to define a covariate-dependent random partition model by augmenting the PPM with an additional factor that induces the desired dependence on the covariates. We refer to the additional factor as similarity function. Focusing on continuous covariates a similar approach is being proposed in independent current work by Park and Dunson (2007). Shahbaba and Neal (2007) introduce a related model for categorical outcomes, i.e., classification. They use a logistic regression for the categorical outcome on covariates and a random partition of the samples, with cluster-specific regression parameters in each partitioning subset. The random partition is defined by a DP prior and includes the covariates as part of the response vector.

In section 2 we state the proposed model and considerations in choosing the similarity function. In section 3 we show that the computational effort of posterior simulation remains essentially unchanged from PPM models without covariates. In 4 we propose specific choices of the similarity function for common data formats. In section 5 we show a simulation study and a data analysis example. In the context of the simulation study we contrast the proposed approach to an often used ad-hoc implementation of including the covariates in an augmented response vector.

2. Clustering with Covariates

We build on the PPM (1), modifying the cohesion function $c(S_j)$ with an additional factor that achieves the desired regression. Recall that $x_j^* = (x_i, i \in S_j)$ denotes all covariates for units in the j -th cluster. Let $g(x_j^*)$ denote a non-negative function of x_j^* that formalizes similarity of the x_i with larger values $g(x_j^*)$ for sets of covariates that are judged to be similar. We define the model

$$p(\rho | x) \propto \prod_{j=1}^k g(x_j^*) \cdot c(S_j) \quad (3)$$

with the normalization constant $g_n(x^n) = \sum_{\rho_n} \prod_{j=1}^k g(x_j^*) c(S_j)$. By a slight abuse of notation we include x behind the conditioning bar even when x is not a random variable. We will later discuss specific choices for the similarity function g . As a default choice we propose to define $g(\cdot)$ as the marginal probability in an auxiliary probability model q , even if x_i are not considered random,

$$g(x_j^*) \equiv \int \prod_{i \in S_j} q(x_i | \xi_j) q(\xi_j) d\xi_j. \quad (4)$$

The function (4) satisfies the following two properties that are desirable for a similarity function in (3). First, we require symmetry with respect to permutations of the sample indices i . The probability model must not depend on the order

of introducing the experimental units. This implies that the similarity function $g(\cdot)$ must be symmetric in its arguments. Second, we require that the similarity function scales across sample size, in the sense that $\int g(x^*) = \int g(x^*, x) dx$. In words, the similarity of any cluster is the average of any augmented cluster.

Under these two requirements (4) is not only technically convenient. It is also the only possible similarity function that satisfies these two constraints.

Proposition 1: Assume that (i) a similarity function $g(x^*)$ satisfies the two constraints; (ii) $q(x_j^*)$ integrates over the covariate space, $\int g(x_j^*) dx_j^* < \infty$. Then $g(x_j^*)$ is proportional to the marginal distribution on x_j^* under a hierarchical auxiliary model as in (4).

The result follows from De Finetti's representation theorem for exchangeable probability measures, applied to a normalized version of $q(\cdot)$. See, for example, Bernardo and Smith (1994, chapter 4.3). The representation theorem applies for an infinite sequence of variables that are subject to the symmetry constraint. The result establishes that all similarity functions that satisfy the symmetry constraint are of the form (4).

The definition of the similarity function with the auxiliary model $q(\cdot)$ also implies another important property. The random partition model (3) is coherent across sample sizes. The model for the first n experimental units follows from the model for $(n + 1)$ observations by appropriate marginalization. Without covariates we would simply require $p(\rho_n) = \sum_{e_{n+1}} p(\rho_{n+1})$, with the summation defined over all possible values of e_{n+1} . With covariates we consider the following condition.

Corollary: Assume that the similarity function is defined by means of an auxiliary model, as in (4). The covariate dependent PPM (3) is coherent across sample sizes as formalized by the following relationship of $p(\rho_n | x^n)$ and $p(\rho_{n+1} |$

x^n, x_{n+1}):

$$p(\rho_n | x^n) = \sum_{e_{n+1}} \int p(\rho_{n+1} | x^n, x_{n+1}) q(x_{n+1} | x^n) dx_{n+1}, \quad (5)$$

for the probability model $q(x_{n+1} | x^n) \propto g_{n+1}(x^{n+1})/g_n(x^n)$.

We complete the random partition model (3) with a sampling model that defines independence across clusters and exchangeability within each cluster. We include cluster-specific parameters θ_j and common hyperparameters η :

$$p(y^n | \rho^n, \theta, \eta, x^n) = \prod_{j=1}^k \prod_{i \in S_j} p(y_i | x_i, \theta_j, \eta) \quad \text{and} \quad p(\theta | \eta) = \prod_j p(\theta_j | \eta). \quad (6)$$

The resulting model extends PPMs of the type (1) while keeping the product structure.

3. Posterior Inference

A practical advantage of the proposed default choice for $g(x_j^*)$ is that it greatly simplifies posterior simulation. In words, posterior inference in (3) and (6) is identical to the posterior inference that we would obtain if x_i were part of the random response vector y_i . Formally, define an auxiliary model $q(\cdot)$ by replacing (6) with

$$q(y^n, x^n | \rho^n, \theta, \eta) = \prod_j \prod_{i \in S_j} p(y_i | x_i, \theta_j, \eta) q(x_i | \xi_j) \quad \text{and} \quad q(\theta, \xi | \eta) = \prod_j p(\theta_j | \eta) q(\xi_j) \quad (7)$$

and replace the covariate-dependent prior $p(\rho_n | x^n)$ by the PPM $q(\rho_n) \propto \prod c(S_j)$. We continue to use $q(\cdot)$ for the auxiliary probability model that is introduced as a computational device only, and $p(\cdot)$ for the proposed inference model. The posterior distribution $q(\rho | y^n, x^n)$ under the auxiliary model is identical to $p(\rho | y^n, x^n)$ under the proposed model. An important caveat is that

ξ_j and θ_j must be a priori independent. In particular the prior on θ_j and $q(\xi_j)$ must not include common hyperparameters. If we had $p(\theta_j | \eta)$ and $q(\xi_j | \eta)$ depend on a common hyperparameter η , then the posterior distribution under the auxiliary model (7) would differ from the posterior under the original model. Let $q(x^n | \eta) = \sum_{\rho_n} \prod_j q(x_j^* | \xi_j) q(\xi_j | \eta)$. The two posterior distributions would differ by a factor $1/q(x^n | \eta)$. The implication on the desired inference can be substantial. See the results in section 5 for an example and more discussion. However, the restriction that θ_j and ξ_j be independent is natural. The function $g(x_j^*)$ defines similarity of the experimental units in one cluster and is chosen by the user. There is no notion of learning about the similarity.

A minor complication arises with posterior predictive inference, i.e., reporting $p(y_{n+1} | x^n, x_{n+1})$. Using $\tilde{x} = x_{n+1}$, $\tilde{y} = y_{n+1}$ and $\tilde{e} = e_{n+1}$ to simplify notation, we find $p(\tilde{y} | \tilde{x}, x^n, y^n) = \int p(\tilde{y} | \tilde{x}, \rho_{n+1}, y^n) dp(\rho_{n+1} | \tilde{x}, y^n, x^n)$. The integral is simply a sum over all configurations ρ_{n+1} . But it is not immediately recognizable as a posterior integral with respect to $p(\rho_n | y^n)$. This can easily be overcome by an importance sampling reweighting step. The prior on $\rho_{n+1} = (\rho_n, \tilde{e})$ can be written as

$$p(\tilde{e} = \ell, \rho_n | \tilde{x}, x^n) \propto \prod_{j \neq \ell} c(S_j) g(x_j^*) c(S_\ell \cup \{n+1\}) g(x_\ell^*, \tilde{x}) = p(\rho_n | x^n) \frac{g(x_\ell^*, \tilde{x})}{g(x_\ell^*)} \cdot \frac{c(S_\ell \cup \{n+1\})}{c(S_\ell)}$$

Let $q_\ell(\tilde{x} | x_\ell^*) \equiv g(x_\ell^*, \tilde{x})/g(x_\ell^*)$. The posterior predictive distribution becomes

$$p(\tilde{y} | \tilde{x}, y^n, x^n) \propto \int \sum_{\ell=1}^{k_n+1} p(\tilde{y} | y_\ell^*, \tilde{e} = \ell) q_\ell(\tilde{x} | x_\ell^*) \frac{c(S_\ell \cup \{n+1\})}{c(S_\ell)} p(\rho_n | y^n, x^n) d\rho_n. \quad (8)$$

Therefore, sampling from (8) is implemented on top of posterior simulation for $\rho_n \sim p(\rho_n | y^n)$. For each imputed ρ_n , generate $\tilde{e} = \ell$ with probabilities propor-

tional to

$$w_\ell = q_\ell(\tilde{x} | x_\ell^*) \frac{c(S_\ell \cup \{n+1\})}{c(S_\ell)}$$

and generate \tilde{y} from $p(\tilde{y} | y_\ell^*, \tilde{e} = \ell)$, weighted with w_ℓ , noting that in the special case $\ell = k_n + 1$ we get $w_{k_n+1} = g(x^*)c(\{n+1\})$.

Reporting posterior inference for random partition models is complicated by problems related to the label switching problem. See Jasra et al. (2005) for a recent summary and review of the literature. Usually the focus of inference in semiparametric mixture models like (6) is on density estimation and prediction, rather than inference for specific clusters and parameter estimation. Predictive inference is not subject to the label switching problem, as the posterior predictive distribution marginalizes over all possible partitions. However, in some examples we want to highlight the choice of the clusters. We briefly comment on the implementation of such inference. We use an approach that we found particularly meaningful to report inference about the regression $p(\rho | x)$. We report inference stratified by k , the number of clusters. For given k we first find a set of k indices $I = (i_1, \dots, i_k)$ with high posterior probability of the corresponding cluster indicators being distinct, i.e., $Pr(D_I | y)$ is high for $D_I = \{e_i \neq e_j \text{ for } i \neq j \text{ and } i, j \in I\}$. To avoid computational complications, we do not insist on finding the k -tuple with the highest posterior probability. We then use i_1, \dots, i_k as anchors to define cluster labels by restricting posterior simulations to k clusters with the units i_j in distinct clusters. Post-processing MCMC output, this is easily done by discarding all imputed parameter vectors that do not satisfy the constraint. We re-label clusters, indexing the cluster that contains unit i_j as cluster j . Now it is meaningful to report posterior inference on specific clusters. The proposed post-processing is similar to the pivotal reordering suggested in Marin and Robert (2007, chapter 6.4). An alternative loss function based method to formalize inference on partitions is discussed in Lau and Green (2007).

4. Similarity Functions

For continuous covariates we suggest as a default choice for $g(x_j^*)$ the marginal distribution of x_j^* under a normal sampling model. Let $N(x; m, V)$ denote a normal model for the random variable x , with moments m and V , and let $\text{Ga}(x; a, b)$ denote a Gamma distributed random variable with mean a/b . We use $q(x_j^* | \xi_j) = \prod_{i \in S_j} N(x_i; m_j, s_j)$, with a conjugate prior for $\xi_j = (m_j, s_j)$ as $p(\xi_j) = N(m_j; m, B) \cdot \text{Ga}(s_j^{-1}; \nu, S_0)$, with fixed m, B, ν, S_0 . The main reason for this choice is operational simplicity. A simplified version uses fixed $s_j \equiv s$. The resulting function $g(x_j^*)$ is the joint density of a collection of correlated multivariate normal vectors, with mean m and $\text{Cor}(x_i, x_{i'}) = B/(V + B)$. The fixed variance s specifies how strong we weigh similarity of the x . In implementations we used $s = c_1 \widehat{S}$, where \widehat{S} is the empirical variance of the covariate, and $c_1 = 0.5$ is a scaling factor that specifies over what range we consider values of this covariate as similar. Analogously we set $B = c_2 \widehat{S}$ with $c_2 = 10$. The choice between fixed s versus variable s_j should reflect prior judgement on the variability of clusters. Variable s_j allows for some clusters to include a wider range of x values than others.

When constructing a cohesion function for categorical covariates, a default choice is based on a Dirichlet prior. Assume x_i is a categorical covariate, $x_i \in \{1, \dots, C\}$. To define $g(x_j^*)$, let $q(x_i = c) = \xi_{jc}$ denote the probability mass function. Together with a conjugate Dirichlet prior, $q(\xi_j) = \text{Dir}(\alpha_1, \dots, \alpha_C)$ we define the similarity function as a Dirichlet-categorical probability

$$g(x_j^*) = \int \prod_{i \in S_j} \xi_{j, x_i} dq(\xi_j) = \int \prod_{c=1}^C \xi_{jc}^{n_{jc}} dq(\xi_j), \quad (9)$$

with $n_{jc} = \sum_{i \in S_j} I(x_i = c)$. This is a Dirichlet-multinomial model without the multinomial coefficient. For binary covariates the similarity function becomes a Beta-Binomial probability without the Binomial coefficient. The choice of the

hyperparameters α needs some care. To facilitate the formation of clusters that are characterized by the categorical covariates we recommend Dirichlet hyperparameters $\alpha_c < 1$. For example, for $C = 2$, the bimodal nature of a Beta distribution with such parameters assigns high probability to binomial success probabilities ξ_{j1} close to 0 or 1. Similarly, the Dirichlet distribution with parameters $\alpha_c < 1$ favors clusters corresponding to specific levels of the covariate.

A convenient specification of $g(\cdot)$ for ordinal covariates is an ordinal probit model. The model can be defined by means of latent variables and cutoffs. Assume an ordinal covariate x with C categories. Following Johnson and Albert (1999), consider a latent trait Z and cutoffs $-\infty = \gamma_0 < \gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_{C-1} < \gamma_C = \infty$, so that $x = c$ if and only if $\gamma_{c-1} < Z \leq \gamma_c$. Consider fixed cutoff values $\gamma_c = c - 1$, $c = 1, \dots, C - 1$, and a normally distributed latent trait, $Z_i \sim N(m_j, s_j)$. Let $\Phi_c = \Pr(\gamma_{c-1} < Z \leq \gamma_c \mid m_j, s_j)$ denote the normal quantiles. We define $q(x_i \mid s_i = j, m_j, s_j) = \Phi_c$. The definition of the similarity function is completed with $q(m_j, s_j)$ as a normal-inverse gamma distribution as for the continuous covariates. We have

$$q(x_j^*) = \int \prod_{c=1}^C \Phi_c^{n_{cj}} N(m_j; m, B) \text{Ga}(s_j^{-1}; \nu, S_0) dm_j ds_j,$$

where $n_{cj} = \sum_{i \in S_j} I\{x_i = c\}$.

Finally, to define $g(\cdot)$ for count-type covariates, we use a mixture of Poisson distributions

$$g(x_j^*) = \frac{1}{\prod_{i \in S_j} x_i!} \int \xi_j^{\sum_{i \in S_j} x_i} \exp(-\xi_j |S_j|) dq(\xi_j). \quad (10)$$

With a conjugate gamma prior, $q(\xi_j) = \text{Ga}(\xi_j; a, b)$, the similarity function $g(x_j^*)$ allows easy analytic evaluation. As a default, we suggest choosing $a = 1$ and $a/b = c\hat{S}$, where \hat{S} is the empirical variance of the covariate, and a/b is the expectation of the gamma prior.

The main advantage of the proposed similarity functions is computational

simplicity of posterior inference. A minor limitation is the fact that the proposed default similarity functions, in addition to the desired dependence of the random partition model on covariates also include a dependence on the cluster size $n_j = |S_j|$. From a modeling perspective this is undesirable. The mechanism to define size dependence should be the underlying PPM and the cohesion functions $c(S_j)$ in (3). However, we argue that the additional cluster size penalty that is introduced through $g(\cdot)$ can be ignored in comparison to, for example, the cohesion function $c(S_j) = M(n_j - 1)!$ that is implied by the popular DP prior.

Proposition 2: The similarity function introduces an additional cluster size penalty in model (3). Consider the case of constant covariates, $x_i \equiv x$, and let n_j denote the size of cluster j . The default choices of $g(x_j^*)$ for continuous, categorical and count covariates introduce a penalty for large n_j , with $\lim_{n_j \rightarrow \infty} g(x_j^*) = 0$. But the rate of decrease is ignorable compared to the cohesion $c(S_j)$. Let $f(n_j)$ be such that $\lim_{n_j \rightarrow \infty} g(x_j^*)/f(n_j) \geq M$ with $0 < M < \infty$. For continuous covariates the rate is $f(n_j) = (2\pi)^{-\frac{n_j}{2}} V^{-\frac{n_j-1}{2}} (\rho + n_j)^{\frac{1}{2}}$ with $\rho = V/B$. For categorical covariates it is $f(n_j) = (A + n_j)^{A-\alpha_x}$, with $A = \sum_c \alpha_c$. For count covariates it is $f(n_j) = C^{-\frac{n_j}{2}} (\alpha + n_j x)^{\frac{1}{2}}$ with $C = 2\pi x \exp(1/6x)$.

Proof: see appendix. \square

Another important concern besides the dependence of $g(\cdot)$ on n_j is the dependence on $q(\xi_j)$ in the auxiliary model in (4). In particular, we focus on the dependence on possibly hyperparameters ϕ that index $q(\xi_j)$. We write $q(\xi_j | \phi)$ when we want to highlight the use of such (fixed) hyperparameters. The auxiliary models $q(\cdot)$ for all proposed default similarity functions include such hyperpa-

rameters. Model (3) and (4) implies conditional cluster membership probabilities

$$p(e_{n+1} = j \mid x_{n+1}, x^n, \rho_n) \propto \frac{c(S_j \cup \{n+1\})}{c(S_j)} \frac{g(x_j^*, x_{n+1})}{\underbrace{g(x_j^*)}_{\equiv q_j(x_{n+1} \mid x_j^*, \phi)}} \quad (11)$$

$j = 1, \dots, k_n$, with the convention $g(\emptyset) = c(\emptyset) = 1$. The cluster membership probabilities (11) are asymptotically independent of ϕ . Noting that $q_j(x_{n+1} \mid x_j^*, \phi)$ can be written as a posterior predictive distribution in the auxiliary model, the statement follows from the asymptotic agreement of posterior predictive distributions.

Proposition 3: Consider any two similarity functions $g^h(x_j^*)$, $h = 1, 2$, based on an auxiliary model (4) with different probability models $q^h(\xi_j)$, $s = 1, 2$, but the same model $q(x_i \mid \xi_j)$. For example, $q^h(\xi_j) = q(\xi_j \mid \phi_h)$. Assume that both auxiliary models satisfy the regularity conditions of Schervish (1995, Section 7.4.2) and assume $q(x_i \mid \xi) \leq K$ is bounded. Let $q_j^h(x_{n+1} \mid x_j^*)$ denote the un-normalized cluster membership probability in equation (11) under $q^h(\xi_j)$. Then

$$\lim_{n_j \rightarrow \infty} q_j^1(x_{n+1} \mid x_j^*) - q_j^2(x_{n+1} \mid x_j^*) = 0 \quad (12)$$

The limit is a.s. under i.i.d. sampling from $q(x_i \mid \xi_0)$.

See the appendix for a straightforward proof making use of asymptotic posterior normality only. The statement is about asymptotic cluster membership probabilities for a future unit only and should not be overinterpreted. The marginal probability for a set of elements forming a cluster does depend on hyperparameters. An example are the expressions for $g(x_j^*)$ in the proof of proposition 2.

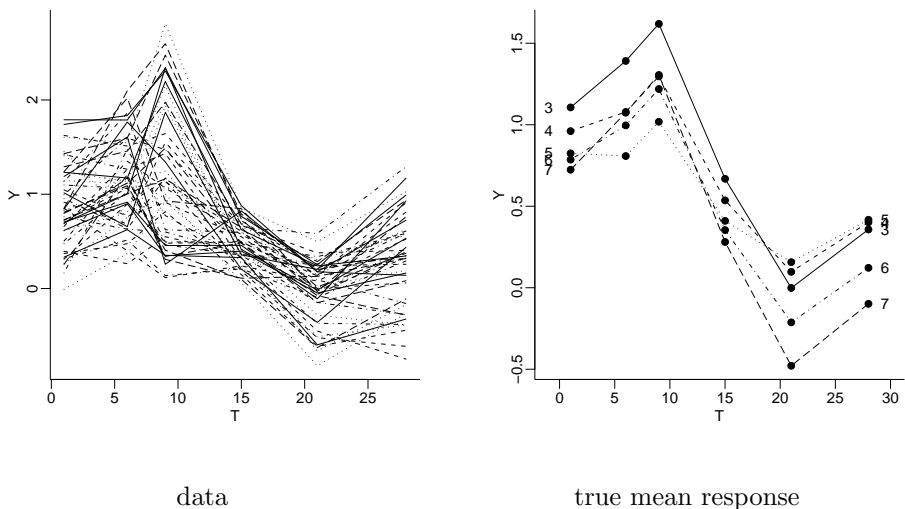


Fig. 1. Simulation example. Data (left panel) and simulation truth for the mean response $E(y_i | x_i)$ as a function of the covariate x_i (right panel).

5. Examples

5.1. Simulation Study

We set up a simulation study with a 6-dimensional response y_i , and a continuous covariate x_i , and for $n = 47$ data points (y_i, x_i) , $i = 1, \dots, n$. The simulation is set up to mimic responses of blood counts over time for patients undergoing chemo-immunotherapy. Each data point corresponds to one patient. The six responses for each patient could be for six key time points over the course of one cycle of the therapy. The covariate is a treatment dose. In the simulation study we sampled x_i from 5 possible values, $x_i \in \{3, 4, 5, 6, 7\}$. For each value of x we defined a different mean vector. Figure 1b shows the 5 distinct mean vectors for the 6-dimensional response (plotted against days T). The bullets indicate the six responses. Adding a 6-dimensional normal residual we generated the data shown in Figure 1a.

We then implemented the proposed model for covariate-dependent clustering

to define a random partition model with a regression on the covariate x_i , as proposed earlier. Let $N(x; m, s)$ indicate a normal distribution for the random variable x , with moments (m, s) , let $W(V; \nu, S)$ denote Wishart prior for a random matrix V with scalar parameter ν and matrix variate parameter S , and let $Ga(s; a, b)$ denote a gamma distribution parameterized to have mean a/b . Let $\theta_j = (\mu_j, V_j)$. We used a multivariate normal model $p(y_i | \theta_j) = N(\mu_j, V_j)$, with a conditionally conjugate prior for θ_j , i.e., $p(\theta_j) = N(\mu_j; m_y, B_y) W(V_j^{-1}; s, S_y^{-1})$. The model is completed with conjugate hyperpriors for m_y, B_y and S_y . The similarity function is a normal kernel. We use

$$g(x_j^*) = \int \prod_{i \in S_j} N(x_i; m_j, s_j) N(m_j; m, B) Ga(s_j^{-1}; \nu, S_0).$$

The hyperparameters m, B, S_0 are fixed.

Figure 2 summarizes posterior inference by the posterior predictive distribution for y_{n+1} arranged by x_{n+1} . The marginal posterior distribution for the number of clusters assigns posterior probabilities $p(k = 3 | data) = 76\%$, $p(k = 4 | data) = 20\%$ and $p(k = 5 | data) = 3\%$. We used the procedure described at the end of section 3 to report cluster-specific summaries. Conditional on $k = 3$, the three clusters are characterized by $x_i = 3.4(0.2)$ for the first cluster, $x_i = 5(0)$ for the 2nd cluster and $x_i = 7(0)$ for the third cluster (posterior means and standard deviations of x_i assigned to each cluster). The average cluster sizes are 24, 11 and 11.

For comparison, the right panel of the same figure shows posterior predictive inference in a model using a PPM prior on clustering, without the use of covariates. In this case, the inference is by construction the same for all covariate values.

A technically convenient ad-hoc solution to include covariates in a sampling model is to proceed as if the covariates were part of the response vector. This approach is used, for example in Mallet et al. (1988) or Müller et al. (2004).




Fig. 2. Simulation example. Estimated mean response under the proposed covariate regression (left panel) and without (right panel).

Doing so introduces an additional factor in the likelihood. Let y generically denote the response, x the covariate, and θ the model parameters. Treating the covariate as part of the response vector is equivalent to replacing the sampling model $p(y | x, \theta)$ by $p(y | x, \theta) \cdot p(x | \theta)$. Equivalently, one can interpret the additional factor $p(x | \theta)$ as part of the prior. The reported posterior inference is as if we had changed the prior model from the original $p(\theta)$ to $\tilde{p}(\theta) \propto p(\theta)p(x | \theta)$. Wong et al. (2003) use this interpretation for a similar construction in the context of prior probability models for a positive definite matrix.

The implied modification of the prior probability model could be less innocuous than what it seems. We implemented inference under the PPM prior model (1), without covariates, and a sampling model as before, but now for an extended response vector augmented by the covariate x_i . We assume a multivariate normal sampling model $p(y_i, x_i | \theta_j) = N(\mu_j, V_j)$, with a conditionally conjugate prior for θ_j , i.e., $p(\theta_j) = N(\mu_j; m_y, B_y) W(V_j^{-1}; s, S_y^{-1})$. As before the model is completed with conjugate hyperpriors for m_y, B_y and S_y . The hyperparameters




Fig. 3. Prediction for \tilde{y} arranged by \tilde{x} for the covariate response model with the additional factors $p(x_j^n | \theta_j, \eta)$ in the prior (left panel). Posterior distribution on the size of the largest cluster under the covariate response model (right panel). The example has a total sample size of $n = 47$.

are chosen exactly as before. The additional 7-th row and column of the prior means for B_y and S_y are all zero except for the (7,7) diagonal element which we fix at $E(m_j)$ and $E(s_j)$ to match the moments in the auxiliary model $q(\cdot)$ from before. We refer to the model as *covariate response* model, in contrast to the proposed *covariate regression* model. Figure 3 shows inference under this modified model which differs from the proposed model by the additional factors $p(x_j^n | \theta_j)$ in the prior. In this example, the covariate response model leads to a much reduced size of the partition. With posterior probability 97% the partition includes only one cluster. This in turn leads to essentially a simple linear regression implied by the dominating multivariate normal $p(x_i, y_i | \theta_1)$, as clearly seen in Figure 3. In contrast, under the proposed covariate regression model the size of the largest cluster is estimated between 16 and 26 (not shown). The covariate response model could be modified to better match the simulation

truth. This could be achieved, for example, by fixing the hyperparameters in a way that avoids correlation of the first 6 and the last dimension of μ_j . The resulting model would have exactly the format of (3).

We thus do not consider the principled way of introducing the covariates in the PPM to be the main feature of the proposed model. Rather, we argue that the proposed model greatly simplifies the inclusion of covariates including a variety of data formats. It would be unnecessarily difficult to attempt the construction of a joint model for continuous responses and categorical, binary and continuous covariates. We argue that it is far easier to focus on a modification of the cohesion function and that this does not imply restricting the scope of the proposed models. This is illustrated in the following data analysis example.

5.2. A Survival Model with Patient Baseline Covariates

We consider data from a high-dose chemotherapy treatment of women with breast cancer. The data for this particular study have been discussed in Rosner (2005) and come from Cancer and Leukemia Group B (CALGB) Study 9082. It consists of measurements taken from 763 randomized patients, available as of October 1998 (enrollment had occurred between January 1991 to May 1998). The response of interest is the survival time, defined as the time until death from any cause, relapse, or diagnosis with a second malignancy. There are two treatments, one involving a low dose of the anticancer drugs, and the other consisting of aggressively high dose chemotherapy. The high-dose patients were given considerable regenerative blood-cell supportive care (including marrow transplantation) to help decreasing the impact of opportunistic infections rising from the severely-affected immune system. The number of observed failures was 361, with 176 under high dose and 185 under low dose chemotherapy.


The dataset also includes information on the following covariates for each

patient: a treatment indicator defined as 1 if a high-dose was administered and 0 otherwise (HI); age in years at baseline (AGE); the number of positive lymph nodes found at diagnosis (POS) (the more the worse the prognosis, i.e. the more likely the cancer has spread); tumor size in millimeters (TS), a one-dimensional measurement; an indicator of whether the tumor is positive for the estrogen or progesterone receptor (ER+) (patients who were positive also received the drug tamoxifen and are expected to have better risk) and an indicator of the woman’s menopausal status, defined as 1 if she is either perimenopausal or postmenopausal or 0 otherwise (MENO). Two of these six covariates are continuous (AGE,TS), three are binary (ER+,MENO,HI) and one is a count (POS).

First we carried out inference in a model using the indicator for high-dose as the only covariate, i.e. $x_i = \text{HI}$. We implemented model (6) with a similarity function for the binary covariate based on the beta-binomial model (9). We used $\alpha = (0.1, 0.1)$ to favor clusters with homogenous dose assignment. Conditional on an assumed partition ρ_n we use a normal sampling model $p(y_i | \theta_j) = N(\mu_j, V_j)$, with a conjugate normal-inverse gamma prior $p(\theta_j | \eta) = N(\mu_j | m_y, B_y) \text{Ga}(V_j^{-1} | s/2, sS_y/2)$ and hyperprior $m_y \sim N(a_m, A_m)$, $S_y \sim \text{Ga}(q, q/R)$. Here $\eta = (a_m, A_m, B_y, q, R)$ are fixed hyperparameters. We use $a_m = \hat{m}$, the sample average of y_i , $A_m = 100$, $B_y = 100^2$, $s = q = 4$, and $R = 100$. Figure 4a shows inference summaries. The posterior distribution $p(k | \text{data})$ for the number of clusters is shown in Figure 5. The three largest clusters contain 28%, 23%, and 14% of the experimental units.

Next we extended the covariate vector to include all six covariates. Denote by $x_i = (x_{i1}, \dots, x_{i6})$ the 6-dimensional covariate vector. We implement random clustering with regression on covariates as in model (6). The similarity function is defined as

$$g(x_j^*) = \prod_{\ell=1}^6 g^\ell(x_{j\ell}^*). \quad (13)$$

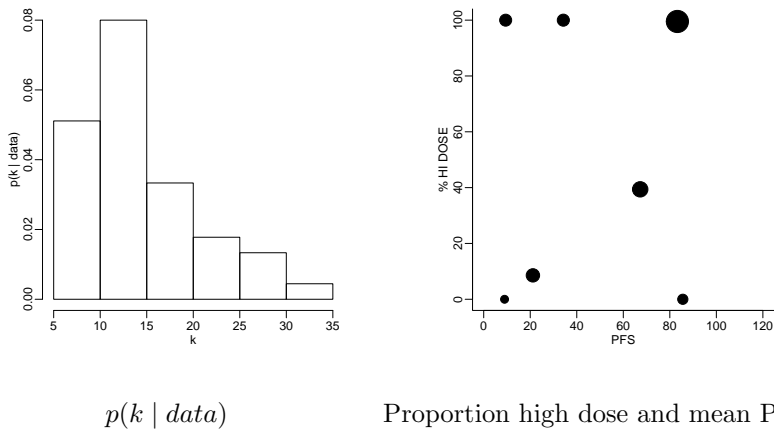


$$S(t) \equiv p(y_{n+1} \geq t \mid \text{data})$$

hazard $h(t)$

data (KM)

Fig. 4. Survival example: Estimated survival function (left panel) and hazard (center panel), arranged by $x \in \{HI, LO\}$. The grey shades show pointwise one posterior predictive standard deviation uncertainty. The right panel shows the data for comparison (Kaplan-Meier curve by dose).



$p(k \mid \text{data})$

Proportion high dose and mean PFS


Fig. 5. Survival example: Posterior for the number of clusters k (left panel), and proportion of patients with high dose (%HI) and average progression free survival (PFS) by cluster. The size (area) of the bullets is proportional to the average cluster size.

For each covariate we follow the suggestion in Section 4 to define a factor $g^\ell(\cdot)$ of the similarity function, using hyperparameters specified as follows. The similarity function for the three binary covariates is defined as in (9) with $\alpha^o = (0.1, 0.1)$ for HI, and $\alpha^o = (0.5, 0.5)$ for ER+, and MENO. The two continuous covariates AGE and TS were standardized to sample mean 0 and unit standard deviation. The similarity functions were specified as described in section 4, with fixed $s = 0.25$, $m = 0$ and $B = 1$. Finally, for the count covariate POS we used the similarity function (10) with $(a, b) = (1.5, 0.1)$. The sampling model is unchanged from before.

We assume that censoring times are independent of the event times and all parameters in the model. Posterior predictive survival curves for various covariate combinations are shown in Figure 6. In the figure, “baseline” refers to HI = 0, tumor size 38mm (the empirical median), ER = 0, Meno = 0, average age (44 years), and POS = 15 (empirical mean). Other survival curves are labeled to indicate how the covariates change from baseline, with TS− indicating tumor size 26mm (the empirical first quartile), TS+ indicating tumor size 50mm (third quartile), HI referring to high-dose chemotherapy, and ER+ indicating positive estrogen or progesterone receptor status. The inference suggests that treated patients with tumor size below the empirical median and that were positive for estrogen or progesterone receptor have almost uniformly highest predicted survival curves than any other combination of covariates.


Figure 7 summarizes features of the posterior clustering. Interestingly, clusters are typically highly correlated to the postmenopausal status, as seen in the right panel. The high-dose indicator is also seen to be positively correlated to the progression-free survival (PFS).

The model used so far consisted of a mixture of normal sampling model for the event times y_i . A more parsimonious sampling model is proposed in Rosner




$$S(t | x)$$


Fig. 6. Survival example: Posterior predictive survival function $S(t | x) \equiv p(y_{n+1} \geq t | x_{n+1} = x, \text{data})$, arranged by x . The “baseline” case refers to all continuous and count covariates equal to the empirical mean, and all binary covariates equal 0. The legend indicates $TS-$ and $TS+$ for tumor size equal 26mm and 50mm (first and third empirical quartile), HI for $HI = 1$ and $ER+$ for $ER = 1$. The legend is sorted by the survival probability at 5 years, indicated by a thin vertical line.



$p(k | data)$



PFS and % high dose



PFS and % postmenopausal

Fig. 7. Survival example: Posterior distribution for the number of clusters (left panel), mean PFS and % high dose patients per cluster (center panel), mean PFS and % postmenopausal patients per cluster (right panel).




Fig. 8. Survival example with the piecewise exponential sampling model: estimated survival $S(t | x)$ arranged by covariates x .

(2005) who uses a piece-wise exponential sampling model with survival function $Pr(y \geq t) =$

$$S(t \mid \lambda_1, \lambda_2, \lambda_3, \tau_1, \tau_2) = \begin{cases} e^{-\lambda_1 t} & \text{if } t \leq \tau_1, \\ e^{-\lambda_1 \tau_1} e^{-\lambda_2 (t - \tau_1)} & \text{if } \tau_1 < t \leq \tau_2, \\ e^{-\lambda_1 \tau_1} e^{-\lambda_2 (\tau_2 - \tau_1)} e^{-\lambda_3 (t - \tau_2)} & \text{if } \tau_2 < t, \end{cases} \quad (14)$$

where $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ defines the exponential rates for each piece, and the change-points τ_1 and τ_2 are fixed as 1 and 2 years, respectively.

Inference under sampling model (14) is summarized in Figure 8. The posterior distribution on the number of clusters (not shown) now shows a posterior mode for $k = 7$, and positive posterior probability for $4 \leq k \leq 14$. The average cluster sizes of the 3 largest clusters are 43%, 30% and 18% of all observations. Compare with Figure 7a. The simple normal sampling model underlying the inference in Figure 7 requires larger size mixtures than the more parsimonious model (14). On the other hand the fixed change-points τ_1 and τ_2 lead to discontinuities in the densities that are visible as the corners in the estimated survival functions in Figure 8. The terms in the mixture are models of the form (14). Each mixture term is a distribution with support over the entire real line. This is in contrast to the mixture of normal kernels with more localized support. 6.

6. Conclusion

We have proposed a novel model for random partitions with a regression on covariates. The model builds on the popular PPM random partition models by introducing an additional factor to modify the cohesion function. We refer to the additional factor as similarity function. It increases the prior probability that experimental units with similar covariates are co-clustered. We provide default choices of the similarity function for popular data formats.

The main features of the model are the possibility to include additional prior

information related to the covariates, the principled nature of the model construction, and a computationally efficient implementation. We compared the model with a popular ad-hoc implementation that includes the covariates as part of an augmented response vector. This approach is feasible in particular when both, the response variable y_i and the covariate x_i are continuous. A joint multivariate normal model for the augmented response vector implicitly introduces the desired regression on covariates. Besides issues of appropriate likelihood specification, the main limitation of such an approach is the practical restriction to continuous responses and covariates. Building a joint model for other data formats is possible, but would be more complicated than a straightforward principled approach that treats x_i as covariates.

Among the limitations of the proposed method is an implicit penalty for the cluster size that is implied by the similarity function. Consider all equal covariates $x_i \equiv x$. The value of the similarity functions proposed in section 4 decreases across cluster size. This limitation could be mitigated by allowing an additional factor $c^*(|S_j|)$ in (6) to compensate the size penalty implicit in the similarity function.

The programs are available as a function in the R package `clusterX` at <http://odin.mdacc.tmc.edu/~pm/prog.html> The function `clusterx(.)` implements the proposed covariate dependent random partition model for an arbitrary combination of continuous, categorical, binary and count covariates, using a mixture of normal sampling model for y_i .

Acknowledgment

Research was partially supported by NIH under grant 1R01CA75981, by FONDECYT under grant 1060729 and the Laboratorio de Análisis Estocástico PBCT-ACT13. Most of the research was done while the first author was visiting at

Appendix

Proof of Proposition 2

For simplicity we drop the j index in n_j , x_j^* etc., relying on the context to prevent ambiguity. Let $\mathbf{1}$ denote a $(n \times 1)$ vector of all ones. For continuous covariates, evaluation of the $g(x^*)$ gives

$$g(x^*) = (2\pi)^{-\frac{n}{2}} [(V + nB)V^{n-1}]^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (x - m)^2 \frac{1}{V} \mathbf{1}' \left(I - J \frac{B}{V + nB} \right) \mathbf{1} \right\}$$

which after some further simplification becomes $\left(2\pi V^{\frac{n-1}{n}} \right)^{-\frac{n}{2}} (\rho + n)^{\frac{1}{2}} M(n)$ with $\lim_{n \rightarrow \infty} M(n) = M$, $0 < M < \infty$. For categorical covariates, Stirling's approximation and some further simplifications gives

$$g(x^*) = \frac{\Gamma(A)}{\prod \Gamma(\alpha_c)} \frac{\prod \Gamma(\alpha_c) \frac{\Gamma(\alpha_x + n)}{\Gamma(\alpha_x)}}{\Gamma(A + n)} \geq \frac{1}{A + n^{A - \alpha_x}} M(n).$$

Similarly, for count covariates we find

$$g(x^*) = \left(\frac{1}{x!} \right)^n \frac{b^a}{\Gamma(a)} \frac{\Gamma(a) + nx}{(b + n)^{a + nx}} \geq (2\pi x)^{-\frac{n}{2}} e^{-\frac{n}{12x}} (\alpha + nx)^{\frac{1}{2}} \cdot M(n).$$

Proof of Proposition 3

The ratio $g^h(x_j^*, x_{n+1})/g^h(x_j^*)$ defines the conditional probability $q_j^h(x_{n+1} | x_j^*)$ under i.i.d. sampling in the auxiliary model, and thus $q_j^h(x_{n+1} | x_j^*) = \int q(x_{n+1} | \xi_j) q^h(\xi_j | x_j^*) d\xi_j$.

The result follows from asymptotic normality of $q^h(\xi_j | x_j^*)$, as $n_j \rightarrow \infty$. Let $\widehat{\xi}_j$ denote the m.l.e. for ξ_j based on n_j observations in cluster j . Let $LL' = [-I(\widehat{\xi}_j)]^{-1}$ denote a Choleski decomposition of the negative inverse of the observed Fisher information matrix (Schervish, 1995, equation 7.88), and let

$\psi_n = L(\xi_j - \widehat{\xi}_j)$. Theorem 7.89 of Schervish (1995) implies for any $\epsilon > 0$ and any compact subset B of the parameter space:

$$\lim_{n_j \rightarrow \infty} P_{\xi_{j0}} \underbrace{\left(\sup_{\psi_n \in B} |q^1(\psi_n | x_j^*) - q^2(\psi_n | x_j^*)| > \epsilon \right)}_{\pi_n} = 0.$$

The limit is in the cluster size n_j . The probability is under an assumed true sampling model $q(x_i | \xi_j^o)$, and the supremum is over B .

Consider a sequence of reparametrizations of $q(x_i | \xi_j)$ to $q(x_i | \psi_n)$. Let $B_M = \{\xi_j : |\xi_j^o - \xi_j| < M\}$ be an increasing sequence of compact sets, and recall the definition of $\pi_n = P_{\xi_j^o}(\sup_B \dots > \epsilon)$. Then

$$\lim_n \int q(x_{n+1} | \xi_j) (q^2(\xi_j | x_j^*) - q^1(\xi_j | x_j^*)) d\xi_j \leq \lim_M \lim_n \epsilon K(1 - \pi_n) + \pi_n K = \epsilon K,$$

for any $\epsilon > 0$.

References

- Antoniak, C. E. (1974) Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The Annals of Statistics*, **2**, 1152–1174.
- Banfield, J. D. and Raftery, A. E. (1993) Model-based Gaussian and non-Gaussian clustering. *Biometrics*, **49**, 803–821.
- Barry, D. and Hartigan, J. A. (1993) A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, **88**, 309–319.
- Bernardo, J.-M. and Smith, A. F. M. (1994) *Bayesian theory*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. Chichester: John Wiley & Sons Ltd.
- Dahl, D. B. (2003) Modal clustering in a univariate class of product partition models. *Tech. Rep. 1085*, Department of Statistics, University of Wisconsin.

- Dasgupta, A. and Raftery, A. E. (1998) Detecting Features in Spatial Point Processes With Clutter via Model-Based Clustering. *Journal of the American Statistical Association*, **93**, 294–302.
- Denison, D. G. T., Holmes, C. C., Mallick, B. K. and Smith, A. F. M. (2002) *Bayesian methods for nonlinear classification and regression*. Wiley Series in Probability and Statistics. Chichester: John Wiley & Sons Ltd.
- Ferguson, T. S. (1973) A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, **1**, 209–230.
- Fraley, C. and Raftery, A. E. (2002) Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc.*, **97**, 611–631.
- Green, P. J. and Richardson, S. (1999) Modelling Heterogeneity with and without the Dirichlet Process. *Tech. rep.*, University of Bristol, Department of Mathematics.
- Green, P. J. and Sibson, R. (1978) Computing Dirichlet tessellations in the plane. *Comput. J.*, **21**, 168–173.
- Hartigan, J. A. (1990) Partition models. *Communications in Statistics, Part A – Theory and Methods*, **19**, 2745–2756.
- Ishwaran, H. and James, L. F. (2003) Generalized weighted Chinese restaurant processes for species sampling mixture models. *Statist. Sinica*, **13**, 1211–1235.
- Jasra, A., Holmes, C. C. and Stephens, D. A. (2005) Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling. *Statist. Sci.*, **20**, 50–67.
- Johnson, V. E. and Albert, J. H. (1999) *Ordinal data modeling*. Statistics for Social Science and Public Policy. New York: Springer-Verlag.

- Kim, H.-M., Mallick, B. K. and Holmes, C. C. (2005) Analyzing nonstationary spatial data using piecewise Gaussian processes. *Journal of the American Statistical Association*, **100**, 653–668.
- Lau, J. W. and Green, P. J. (2007) Bayesian Model-Based Clustering Procedures. *Journal of Computational and Graphical Statistics*, **16**, 526–558.
- Mallet, A., Mentré, F., Gilles, J., Kelman, A., Thomson, A., S.M., Bryson and Whiting, B. (1988) Handling covariates in population pharmacokinetics with an application to gentamicin. *Biomedical Measurement Informatics and Control*, **2**, 138–146.
- Marin, J.-M. and Robert, C. P. (2007) *Bayesian Core. A Practical Approach to Computational Bayesian Statistics*. New York: Springer-Verlag.
- Müller, P., Quintana, F. and Rosner, G. (2004) A method for combining inference across related nonparametric Bayesian models. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **66**, 735–749.
- Okabe, A., Boots, B., Sugihara, K. and Chiu, S. N. (2000) *Spatial tessellations: concepts and applications of Voronoi diagrams*. Wiley Series in Probability and Statistics. Chichester: John Wiley & Sons Ltd., second edn. With a foreword by D. G. Kendall.
- Park, J.-H. and Dunson, D. (2007) Bayesian generalized product partition models. *Tech. rep.*, Duke University.
- Pitman, J. (1996) Some Developments of the Blackwell-MacQueen Urn Scheme. In *Statistics, Probability and Game Theory. Papers in Honor of David Blackwell* (eds. T. S. Ferguson, L. S. Shapeley and J. B. MacQueen), 245–268. Haywar, California: IMS Lecture Notes - Monograph Series.

- Quintana, F. A. (2006) A predictive view of Bayesian clustering. *Journal of Statistical Planning and Inference*, **136**, 2407–2429.
- Quintana, F. A. and Iglesias, P. L. (2003) Bayesian Clustering and Product Partition Models. *Journal of The Royal Statistical Society Series B*, **65**, 557–574.
- Richardson, S. and Green, P. J. (1997) On Bayesian analysis of mixtures with an unknown number of components (with discussion). *Journal of the Royal Statistical Society, Series B*, **59**, 731–792.
- Rosner, G. L. (2005) Bayesian monitoring of clinical trials with failure-time endpoints. *Biometrics*, **61**, 239–245.
- Shahbaba, B. and Neal, R. M. (2007) Nonlinear models using dirichlet process mixtures. *Tech. rep.*, University of Toronto.
- Wong, F., Carter, C. K. and Kohn, R. (2003) Efficient estimation of covariance selection models. *Biometrika*, **90**, 809–830.