

Nonparametric Bayesian Data Analysis

Peter Müller* Fernando A. Quintana†

Abstract. We review the current state of nonparametric Bayesian inference. The discussion follows a list of important statistical inference problems, including density estimation, regression, survival analysis, hierarchical models and model validation. For each inference problem we review relevant nonparametric Bayesian models and approaches including Dirichlet process (DP) models and variations, Polya trees, wavelet based models, neural network models, spline regression, CART, dependent DP models, and model validation with DP and Polya tree extensions of parametric models.

1 INTRODUCTION

Nonparametric Bayesian inference is an oxymoron and misnomer. Bayesian inference by definition always requires a well defined probability model for observable data y and any other unknown quantities θ , i.e., parameters. Nonparametric Bayesian inference traditionally refers to Bayesian methods that result in inference comparable to classical nonparametric inference, like kernel density estimation, scatterplot smoothers, etc. Such flexible inference is typically achieved by models with massively many parameters. In fact, a commonly used technical definition of nonparametric Bayesian models are probability models with infinitely many parameters (Bernardo and Smith 1994). Equivalently, nonparametric Bayesian models are probability models on function spaces. Nonparametric Bayesian models are used to avoid critical dependence on parametric assumptions, to robustify parametric models, and to define model diagnostics and sensitivity analysis for parametric models by embedding them in a larger encompassing nonparametric model. The latter two applications are technically simplified by the fact that many nonparametric models allow to center the probability distribution at a given parametric model.

*Department of Biostatistics, Box 447, University of Texas M. D. Anderson Cancer Center, Houston, TX 77030-4009, USA. e-mail: pm@odin.mdacc.tmc.edu

†Departamento de Estadística, Pontificia Universidad Católica de Chile, Casilla 306, Santiago 22, CHILE. e-mail: quintana@mat.puc.cl. Partially supported by grant FONDECYT 1020712. First author supported by NIH/NCI under grant NIH R01CA75981.

In this article we review the current state of Bayesian nonparametric inference. The discussion follows a list of important statistical inference problems, including density estimation, regression, survival analysis, hierarchical models and model validation. The list is not exhaustive. In particular, we will not discuss nonparametric Bayesian approaches in time series analysis, and in spatial and spatio-temporal inference.

Other recent surveys of nonparametric Bayesian models appear in Walker et al. (1999) and Dey et al. (1998). Nonparametric models based on Dirichlet process mixtures are reviewed in MacEachern and Müller (2000). A recent review of nonparametric Bayesian inference in survival analysis can be found in Sinha and Dey (1997).

2 DENSITY ESTIMATION

The density estimation problem starts with a random sample $x_i \stackrel{\text{iid}}{\sim} F(x_i)$, $i = 1, \dots, n$, generated from some unknown distribution F . A Bayesian approach to this problem requires a probability model for the unknown F . Traditional parametric inference considers models that can be indexed by a finite dimensional parameter, for example, the mean and covariance matrix of a multivariate normal distribution of the appropriate dimension. In many cases, however, constraining inference to a specific parametric form may limit the scope and type of inferences that can be drawn from such models. In contrast, under a nonparametric perspective we consider a prior probability model $p(F)$ for the unknown density F , for F in some infinite dimensional function space. This requires the definition of probability measures on a collection of distribution functions. Such probability measures are generically referred to as *random probability measures* (RPM). Ferguson (1973) states two important desirable properties for this class of measures (see also Antoniak 1974): (I) their support should be large and (II) posterior inference should be “analytically manageable.” In the parametric case, the development of MCMC methods (see, e.g. Gelfand and Smith 1990) allows to largely overcome the restrictions posed by (II). In the nonparametric context, however, computational aspects are still the subject of much research.

We next describe some of the most common random probability measures adopted in the literature.

2.1 The Dirichlet Process

Motivated by properties (I) and (II), Ferguson (1973) introduced the Dirichlet process (DP) as an RPM. A random probability distribution F is generated by a DP if for any partition A_1, \dots, A_k of the sample space the vector of random probabilities $F(A_i)$ follows a Dirichlet distribution: $(F(A_1), \dots, F(A_k)) \sim D(M \cdot F_0(A_1), \dots, M \cdot F_0(A_k))$. We denote this by $F \sim \mathcal{D}(M, F_0)$. Two parameters need to be specified: the weight parameter M , and the

base measure F_0 . The base measure F_0 defines the expectation, $E(B) = F_0(B)$, and M is a precision parameter that defines variance. For more discussion of the role of these parameter see Walker et al. (1999). A fundamental motivation for the DP construction is the simplicity of posterior updating. Assume

$$x_1, \dots, x_n | F \stackrel{\text{iid}}{\sim} F, \quad \text{and} \quad F \sim \mathcal{D}(M, F_0). \quad (1)$$

Let $\delta_x(\cdot)$ denote a point mass at x . The posterior distribution is $F|x_1, \dots, x_n \sim \mathcal{D}(M+n, F_1)$ with $F_1 \propto F_0 + \sum_{i=1}^n \delta_{x_i}$.

More properties of the DP are discussed, among others, in Ferguson (1973), Korwar and Hollander (1973), Antoniak (1974), Diaconis and Freedman (1986), Rolin (1992), Diaconis and Kemperman (1996) and in Cifarelli and Melilli (2000). Of special relevance for computational purposes is the Polya urn representation by Blackwell and MacQueen (1973). Another very useful result is the construction by Sethurman (1994): Any $F \sim \mathcal{D}(M, F_0)$ can be represented as

$$F(\cdot) = \sum_{h=1}^{\infty} w_h \delta_{\mu_h}(\cdot),$$

$$\mu_h \stackrel{\text{iid}}{\sim} F_0 \text{ and } w_h = U_h \prod_{j < h} (1 - U_j) \text{ with } U_h \stackrel{\text{iid}}{\sim} \text{Beta}(1, M) \quad (2)$$

In words, realizations of the DP can be represented as infinite mixtures of point masses. The locations μ_h of the point masses are a sample from F_0 , and the random weights w_h are generated by a “stick-breaking” procedure. In particular, the DP is an almost surely (a.s.) discrete RPM.

The DP is by far the most popular nonparametric model in the literature (for a recent review, see MacEachern and Müller 2000). However, the a.s. discreteness is in many applications inappropriate. A simple extension to remove the constraint to discrete measures is to introduce an additional convolution, representing the RPM F as

$$F(x) = \int f(x|\theta) dG(\theta) \quad \text{with} \quad G \sim \mathcal{D}(M, G_0). \quad (3)$$

Such models are known as DP mixtures (MDP) (Escobar 1988, MacEachern 1994, Escobar and West 1995). Using a Gaussian kernel, $f(x|\mu, S) = \phi_{\mu, S}(x) \propto \exp[-(x-\mu)^T S^{-1}(x-\mu)/2]$, and mixing with respect to $\theta = (\mu, S)$ we obtain density estimates resembling traditional kernel density estimation. Related models have been studied in Lo (1984), Escobar and West (1995) and in Gasparini (1996). Posterior consistency is discussed in Ghosal, Ghosh and Ramamoorthi (1999).

Posterior inference in MDP models is based on MCMC posterior simulation. Most approaches proceed by breaking the mixture in (3) with the introduction of latent variables θ_i

as $x_i|\theta_i \sim f(x|\theta)$ and $\theta_i \sim G$. Efficient MCMC simulation for general MDP models is discussed, among others, in Bush and MacEachern (1996), MacEachern and Müller (1998), Neal (2000) and West, Müller and Escobar (1994). For related algorithms in a more general setting, see Ishwaran and James (2001). Alternatively to MCMC simulation, sequential importance sampling-based methods have been proposed for MDP models. Examples can be found in Liu (1996), Quintana (1998), MacEachern, Clyde and Liu (1999), Ishwaran and Takahara (2002) and references therein. A third class of methods for MDP models, called the *predictive recursion*, was proposed by Newton and Zhang (1999). Consider the posterior predictive distribution in model (3). Let $F_n(B) \stackrel{\text{def}}{=} E(F(B)|x_1, \dots, x_n)$ denote the posterior mean of the RPM. The posterior mean is identical to the predictive distribution, $F_n(B) = P(\theta_{n+1} \in B|x_1, \dots, x_n)$ for any Borel set B in the appropriate space. The Polya urn representation implies

$$F_1(B) = \frac{M}{M+1}F_0(B) + \frac{1}{M+1}P(\theta_1 \in B|x_1).$$

Newton and Zhang (1999) extrapolate this representation to a recursion in the general case:

$$F_i(B) = (1 - w_i)F_{i-1}(B) + w_iP_{i-1}(\theta_i \in B|x_i), \quad (4)$$

where the probability in the second term in the right-hand side of (4) is computed under the current approximation F_{i-1} , and the nominal values for the weights are $w_i = 1/(M+i)$, $i \geq 1$. The approximation is exact for $i = 1$. In general $F_n(B)$ depends on the order in which x_1, \dots, x_n are processed, but this dependence is rather weak, and in practice, it is recommended to average over a number of permutations of the data. The method is very fast to execute and produces very good approximations, although it tends to over-smooth the results. For a comparison of the computational strategies mentioned here, see Quintana and Newton (2000).

Model (1) has the advantage of the conjugate form. However, getting exact draws from a DP is impossible because this requires the generation of an infinite mixture of point masses. Typical MCMC schemes are based on integrating out the DP via Blackwell and MacQueen's (1973) representation. This makes it difficult to produce inference on functionals of the posterior DP. A similar problem is found in the more general MDP models. Some authors propose MCMC strategies where, instead of integrating out the DP, an approximation to the DP is considered. This is usually done by drawing from $\sum_{h=1}^N w_h \delta_{\mu_h}(\cdot)$ for large enough N . Examples of this strategy can be found in Muliere and Tardella (1998), Ishwaran and James (2002), Kottas and Gelfand (2001), and Gelfand and Kottas (2002).

2.2 Other Discrete Random Probability Measures

An interesting extension of the DP that has been used in the context of density estimation is the invariant DP introduced by Dalal (1979). The idea is to define a prior process on the space of distribution functions that have a structure that can be characterized via invariance, for example, symmetry or exchangeability. Dalal's (1979) construction is based on invariance under a finite group, essentially by restricting Ferguson's (1973) definition to invariant centering measures and partitions. This guarantees that the posterior process is also invariant. Dalal (1979) uses this setup to estimate distribution functions that are symmetric with respect to a known value μ , using F_0 such that $F_0(t) = 1 - F_0(2\mu - t)$ for all $t \leq \mu$ and the group $\mathcal{G} = \{g_1, g_2\}$ where $g_1(x) = x$ and $g_2(x) = 2\mu - x$.

An alternative model to (1) or (3) is obtained by replacing the prior DP with a convenient approximation. Natural candidates follow from truncating Sethurman's (1994) construction (2). In this setup, the prior $\sum_{h=1}^{\infty} w_h \delta_{\mu_h}(\cdot)$ is replaced by $\sum_{h=1}^N w_h \delta_{\mu_h}(\cdot)$ for some appropriately chosen value of N . An example of this procedure is the ϵ -DP proposed by Muliere and Tardella (1998), where N is chosen such that the total variation distance between the DP and the truncation is bounded by a given ϵ . Another variation is the Dirichlet-multinomial process introduced by Muliere and Secchi (1995). Here the RPM is, for some finite N ,

$$F(\cdot) = \sum_{h=1}^N w_h \delta_{\mu_h}(\cdot),$$

$$(w_1, \dots, w_N) \sim D(M \cdot N^{-1}, \dots, M \cdot N^{-1}) \text{ and } \mu_h \stackrel{\text{iid}}{\sim} F_0.$$

More generally, Pitman (1996) described a class of models

$$F(\cdot) = \sum_{h=1}^{\infty} w_h \delta_{\mu_h}(\cdot) + \left(1 - \sum_{h=1}^{\infty} w_h\right) F_0(\cdot), \quad (5)$$

where, for a continuous distribution F_0 , we have $\mu_h \stackrel{\text{iid}}{\sim} F_0$, assumed independent of the non-negative random variables w_h . The weights w_h are constrained by $\sum_{h=1}^{\infty} w_h \leq 1$. The model is known as *Species Sampling Model* (SSM), with the interpretation of w_h as the relative frequency of the h -th species in a list of species present in a certain population, and μ_h as the tag assigned to that species. If $\sum_{h=1}^{\infty} w_h = 1$ the SSM is called *proper* and the corresponding prior RPM is discrete. The stick-breaking priors studied by Ishwaran and James (2001) are a special case of (5), adopting the form $\sum_{h=1}^N w_h \delta_{\mu_h}(\cdot)$, where $1 \leq N \leq \infty$. The weights are defined as $w_h = \prod_{j=1}^{h-1} (1 - U_j) U_h$ with $U_h \sim \text{Beta}(a_h, b_h)$, independently, for a given sequences (a_1, a_2, \dots) and (b_1, b_2, \dots) . Stick-breaking priors are quite general, including not only the Dirichlet-multinomial process and the DP as special cases, but also a two-parameter DP extension, known as the Pitman-Yor process (Pitman and Yor 1987),

and the beta two-parameter process (Ishwaran and Zarepour 2000). Additional examples and MCMC implementation details for stick-breaking RPMs can be found in Ishwaran and James (2001). Further discussion on SSMs appears in Pitman (1996) and Ishwaran and James (2003).

An interesting property of MDP models is that any exchangeable sequence of random variables can be well approximated in the sense of the Prokhorov metric by a certain sequence of mixtures of DPs (Regazzini 1999). In practice, however, this result has limited use. We review next some methods for defining RPMs supported on the set of continuous distributions that have been used in density estimation problems.

2.3 Polya Trees

Polya trees (PT) are proposed in Lavine (1992, 1994) as a generalization of the DP. Like the DP, the PT model satisfies conditions (I) and (II). The PT includes DP models as a special case. But in contrast to the DP, an appropriate choice of the PT parameters allows to generate continuous distributions with probability 1. The definition requires a nested sequence $\Pi = \{\pi_m, m = 1, 2, \dots\}$ of partitions of the sample space Ω . Without loss of generality, we assume the partitions are binary. We start with a partition $\pi_1 = \{B_0, B_1\}$ of the sample space, $\Omega = B_0 \cup B_1$, and continue with nested partitions defined by $B_0 = B_{00} \cup B_{01}$, $B_1 = B_{10} \cup B_{11}$, etc. Thus the partition at level m is $\pi_m = \{B_\epsilon, \epsilon = \epsilon_1 \dots \epsilon_m\}$, where ϵ are all binary sequences of length m . We say that F has a PT (prior) distribution, denoted by $F \sim \mathcal{PT}(\Pi, \mathcal{A})$ if there is a sequence of nonnegative constants $\mathcal{A} = \{\alpha_\epsilon\}$ and independent random variables $\mathcal{Y} = \{Y_\epsilon\}$ such that $Y_\epsilon \sim \text{Beta}(\alpha_{\epsilon_0}, \alpha_{\epsilon_1})$ and for every $\epsilon = (\epsilon_1, \dots, \epsilon_m)$ and $m \geq 1$

$$F(B_{\epsilon_1, \dots, \epsilon_m}) = \left(\prod_{j=1; \epsilon_j=0}^m Y_{\epsilon_1 \dots \epsilon_{j-1}} \right) \left(\prod_{j=1; \epsilon_j=1}^m (1 - Y_{\epsilon_1 \dots \epsilon_{j-1}}) \right).$$

The type of models used for density estimation now replace the DP in (1) and (3) by the $\mathcal{PT}(\Pi, \mathcal{A})$ prior. For a description of samples from a PT prior, see Walker et al. (1999). Posterior consistency issues for density estimation using PT priors have been discussed in Barron, Schervish and Wasserman (1999).

Polya trees have some practical limitations. First, the resulting RPM is dependent on the specific partition adopted. Second, the fixed partitioning scheme results in discontinuities in the predictive distributions. Third, implementations for higher dimensional distributions require extensive housekeeping and are impractical. To mitigate problems related to the discontinuities Paddock et al. (2003) and Hanson and Johnson (2002) introduced randomized Polya trees. The idea is based on dyadic rational partitions, but instead of taking the nominal half-point Paddock et al. (2003) randomly choose a “close” cutoff. This construction is

shown to reduce the effect of the binary tree partition on the first two points noted above. On the other hand, Hanson and Johnson (2002) consider instead a mixture with respect to a hyperparameter that defines the partitioning tree. The problem concerning high dimension persists though.

2.4 Bernstein Polynomials

For a distribution function F on the unit interval, the corresponding Bernstein polynomial is defined as

$$B(x, k, F) = \sum_{j=0}^k F(j/k) \cdot \binom{k}{j} x^j (1-x)^{k-j}.$$

A remarkable property of $B(x, k, F)$ is that it converges uniformly to F as $k \rightarrow \infty$. The definition for $B(x, k, F)$ takes the form of a mixture of Beta densities. Petrone(1999a, 1999b) exploits this property to propose a class of prior distributions on the set of densities defined on $(0, 1]$. Petrone and Wasserman (2002) consider the following model. Assume x_1, \dots, x_n are conditionally i.i.d. given k and w_k with common density

$$f(x|k, w_k) = \sum_{j=1}^k w_{jk} \left\{ \frac{k!}{(j-1)!(k-j)!} \right\} x^{j-1} (1-x)^{k-j},$$

where k is the number of components in the mixture of Beta densities and the weights $w_k = (w_{1k}, \dots, w_{kk})$ satisfy $w_{jk} \geq 0$ and $\sum_{j=1}^k w_{jk} = 1$. We call f a Bernstein polynomial density (BPD). The model is completed by assuming a prior distribution $p(k)$ for k and a distribution $H_k(\cdot)$ given k on the $(k-1)$ -dimensional simplex. Petrone (1999a) showed that if $p(k) > 0$ for all $k \geq 1$ then every distribution on $(0, 1]$ is the (weak) limit of some sequence of BPD, and every continuous density on $(0, 1]$ can be well approximated in the Kolmogorov-Smirnov distance by BPD. Petrone and Wasserman (2002) discuss MCMC strategies for fitting the above model and prove consistency of posterior density estimation under mild conditions. Rates of such convergence are given in Ghosal (2001).

2.5 Other Random Distributions

Lenk (1988) introduces the logistic normal process. The construction of a logistic normal process starts with a Gaussian process $Z(x)$ with mean function $\mu(x)$ and covariance function $\sigma(x, y)$. The transformed process $W = \exp(Z)$ is a lognormal process. Stopping the construction here, and defining a random density $f(x) \propto W$ would be impractical. The lognormal process is not closed under prior to posterior updating, i.e., the posterior on f conditional on observing $y_i \sim f$, $i = 1, \dots, n$ is not proportional to a lognormal process. Instead Lenk (1988) proceeds by defining the generalized lognormal process $LN_X(\mu, \sigma, \zeta)$,

defined essentially by weighting realizations under the lognormal process with the random integral $(\int W d\lambda)^\zeta$. Let $f(x) \propto V(x)$ for $V \sim LN_X(\mu, \sigma, \zeta)$. The density f is said to be logistic normal process $LNS_X(\mu, \sigma, \zeta)$. The posterior on f , conditional on a random sample $y \sim f$, is again a logistic normal process $LNS_X(\mu^*, \sigma, \zeta^*)$. The updated parameters are $\mu^*(s) = \mu(s) + \sigma(s, y)$ and $\zeta^* = \zeta - 1$.

3 REGRESSION

The generic regression problem seeks to estimate an unknown mean function $g(x)$ based on data with i.i.d. measurement errors: $y_i = g(x_i) + \epsilon_i$, $i = 1, \dots, n$. Bayesian inference on g starts with a prior probability model for the unknown function g . If restrictive parametric assumptions for g are inappropriate we are led to consider nonparametric Bayesian models. Many approaches proceed by considering some basis $\mathcal{B} = \{f_1, f_2, f_3, \dots\}$ for an appropriate function space, like the space of square integrable functions. Typical examples are the Fourier basis, wavelet bases, and spline bases. Given a chosen basis \mathcal{B} , any function g can be represented as $g(\cdot) = \sum_h b_h f_h(\cdot)$. A random function g is parametrized by the sequence $\theta = (b_1, b_2, \dots)$ of basis coefficients. Assuming a prior probability model for θ we implicitly put a prior probability model on the random function.

3.1 Spline Models

A commonly used class of basis functions are splines, for example cubic regression splines $\mathcal{B} = \{1, x, x^2, x^3, (x - \xi_1)_+^3, \dots, (x - \xi_T)_+^3\}$, where $(x)_+ = \max(x, 0)$ and $\xi = (\xi_1, \dots, \xi_T)$ is a set of knots. Together with a normal measurement error $\epsilon_i \sim N(0, \sigma)$ this defines a nonparametric regression model

$$y_i = \sum b_h f_h(x_i) + \epsilon_i. \quad (6)$$

The model is completed with a prior $p(\xi, c, \sigma)$ on the set of knots and corresponding coefficients. Smith and Kohn (1996), Denison, Mallick and Smith (1998b), and DiMatteo, Genovese and Kass (2001) are typical examples of such models. Approaches differ mainly in the choice of priors and the implementation. Typically the prior is assumed to factor $p(\xi, b, \sigma) = p(\xi)p(\sigma)p(b|\sigma)$. Smith and Kohn (1996) use the Zellner g-prior (Zellner 1986) for $p(b)$. The prior covariance matrix $Var(b|\sigma)$ is assumed to be proportional to $(B'B)^{-1}$, where B is the design matrix for the given data set. Assuming a conjugate normal prior $b \sim N(0, c\sigma(B'B)^{-1})$ the conditional posterior mean $E(b|\xi, \sigma)$ is a simple linear shrinkage of the least squares estimate \hat{b} . DiMatteo, Genovese and Kass (2001) use a unit-information prior which is defined as a Zellner g-prior with the scalar c chosen such that the prior variance

is equivalent to one observation. Denison et al. (1998b) prefer a ridge prior $p(b) = N(0, V)$ with $V = \text{diag}(\infty, v, \dots, v)$.

Posterior simulation in (6) is straightforward except for the computational challenge of updating ξ , the number and location of knots. This typically involves reversible jump MCMC (Green 1995). Denison et al. (1998a) propose “birth,” “death” and “move” proposals to add, delete and change knots from the currently imputed set ξ of knots. In the implementation of these moves it is important to marginalize with respect to the coefficients b_h . In the conditionally conjugate setup with a normal prior $p(b|\sigma)$ the marginal posterior $p(\xi|\sigma, y)$ can be evaluated analytically. DiMatteo et al. (2001) propose an approximate evaluation of the relevant Bayes factors based on BIC (Bayesian information criterion). An interesting alternative, called focused sampling, is discussed in Smith and Kohn (1998).

3.2 Multivariate Regression

Extensions of spline regression to multiple covariates are complicated by the curse of dimensionality. Smith and Kohn (1997) define a spline based bivariate regression model. General, higher dimensional regression models require some simplifying assumptions about the nature of interactions to allow a practical implementation. One approach is to assume additive effects

$$y_i = \sum_j g_j(x_{ij}) + \epsilon_i,$$

and proceed with each g_j as before. Shively, Kohn and Wood (1999) and Denison, Mallick and Smith (1998b) propose such implementations. Denison, Mallick and Smith (1998c) explore an alternative extension of univariate splines, following the idea of MARS (multivariate adaptive regression splines, Friedman 1991). MARS uses basis functions that are constructed as products of univariate functions. Let $x_i = (x_{i1}, \dots, x_{ip})$ denote the multivariate covariate vector. MARS assumes

$$g(x_i) = b_0 + \sum_{h=1}^k b_h f_h(x_i) \text{ with } f_h(x) = \prod_{j=1}^{J_h} s_{hj}(x_{w_{hj}} - t_{hj})_+.$$

Here we used linear spline terms $(x - t_{hj})_+$ to construct the basis functions f_h . Each basis function defines an interaction of J_h covariates. The indices w_{hj} specify the covariates and t_{hj} gives the corresponding knots.

Another intuitively appealing multivariate extension are CART (classification and regression tree) models. Chipman, George and McCulloch (1998) and Denison, Mallick and Smith (1998a) discuss Bayesian inference in CART models. A regression tree is parametrized by a pair (T, θ) describing a binary tree T with b terminal nodes, and a parameter vector $\theta = (\theta_1, \dots, \theta_b)$ with θ_i defining the sampling distribution for observations that are assigned

to terminal node i . Let y_{ik} , $k = 1, \dots, n_i$ denote the observations assigned to the i -th node. In the simplest case the sampling distribution for the i -th node might be i.i.d. sampling, $y_{ik} \sim N(\theta_i, \sigma)$, $k = 1, \dots, n_i$, with a node-specific mean. The tree T describes a set of rules that decide how observations are assigned to terminal nodes. Each internal node of the tree has an associated splitting rule that decides whether an observation is assigned to the right or to the left branch. Let x_j , $j = 1, \dots, p$ denote the covariates of the regression. The splitting rule is of the form $(x_j > s)$ for some threshold s . Thus each splitting node is defined by a covariate index and threshold. The leaves of the tree are the terminal nodes. Chipman, George and McCulloch (1998) and Denison, Mallick and Smith (1998a) propose Bayesian inference in regression trees by defining a prior probability model for (θ, T) and implementing posterior MCMC. The MCMC scheme includes the following types of moves: (a) splitting a current terminal node (“grow”); (b) removing a pair of terminal nodes and making the parent into a terminal node (“prune”); (c) changing a splitting variable or threshold (“change”). Chipman, George and McCulloch (1998) use an additional swap move to propose a swap of splitting rules among internal nodes. The complex nature of the parameter space makes it difficult to achieve a well mixing Markov chain simulation. Chipman, George and McCulloch (1998) caution against using one long run, and instead advise to use frequent restarts. MCMC posterior simulation in CART models should be seen as stochastic search for high posterior probability trees. Achieving practical convergence in the MCMC simulation is not typically possible.

An interesting special case of multivariate regression arises in spatial inference problems. The spatial coordinates (x_{i1}, x_{i2}) are the covariates for a response surface $g(x_i)$. Wolpert and Ickstadt (1998a) propose a nonparametric model for a spatial point process. At the top level of a hierarchical model they assume a Poisson process as sampling model for the observed data. Let x_i denote the coordinates of an observed event. For example, x_i could be the recorded occurrence of a species in a species sampling problem. The model assumes a Poisson process $x_i \sim Po(\Lambda(x))$ with intensity function $\Lambda(x)$. The intensity function in turn is modeled as a convolution of a normal kernel $k(x, s)$ and a Gamma process, $\Lambda(x) = \int k(x, s)\Gamma(ds)$ and $\Gamma(ds) \sim Gamma(\alpha(ds), \beta(ds))$. With constant $\beta(s) = \beta$ and rescaling the Gamma process to total mass one, the model for $\Lambda(x)$ reduces to a Dirichlet process mixture of normals.

Arjas and Heikkinen (1997) propose an alternative approach to inference for a spatial Poisson process. The prior probability model is based on Voronoi tessellations with a random number and location of knots.

3.3 Wavelet based modeling

Wavelets provide an orthonormal basis in L^2 representing $g \in L^2$ as $g(x) = \sum_j \sum_k d_{jk} \psi_{jk}(x)$, with basis functions $\psi_{jk}(x) = 2^{j/2} \psi(2^j x - k)$ that can be expressed as shifted and scaled versions of one underlying function ψ . The practical attraction of wavelet bases is the availability of super-fast algorithms to compute the coefficients d_{jk} given a function, and vice versa. Assuming a prior probability model for the coefficients d_{jk} implicitly puts a prior probability model on the random function g . Typical prior probability models for wavelet coefficients include positive probability mass at zero. Usually this prior probability mass depends on the “level of detail” j , $Pr(d_{jk} = 0) = \pi_j$. Given a non-zero coefficient, an independent prior with level dependent variances is assumed, for example, $p(d_{jk} | d_{jk} \neq 0) = N(0, \tau_j^2)$. Appropriate choice of π_j and τ_j achieves posterior rules for the wavelet coefficients d_{jk} , which closely mimic the usual wavelet thresholding and shrinkage rules (Chipman et al. 1997, Vidakovic 1998). Clyde and George (2000) discuss the use of empirical Bayes estimates for the hyperparameters in such models.

Posterior inference is greatly simplified by the orthonormality of the wavelet basis. Consider a regression model $y_i = g(x_i) + \epsilon_i$, $i = 1, \dots, n$, with equally spaced data x_i , for example, $x_i = i/n$. Substitute a wavelet basis representation $g(\cdot) = \sum_j \sum_k d_{jk} \psi_{jk}(x)$, let y, d and ϵ denote the data vector, the vector of all wavelet coefficients and the residual vector, respectively. Also, let $B = [\psi_{jk}(x_i)]$ denote the design matrix of the wavelet basis functions evaluated at the x_i . Then we can write the regression in matrix notation as $y = Bd + \epsilon$. The discrete wavelet transform of the data finds, in a computationally highly efficient algorithm, $\hat{d} = B^{-1}y$. Assuming independent normal errors, $\epsilon_i \sim N(0, \sigma^2)$, orthogonality of the design matrix B implies $\hat{d}_{jk} \sim N(d_{jk}, \sigma^2)$, independently across (j, k) . Assuming *a priori* independent d_{jk} leads to *a posteriori* independence of the wavelet coefficients d_{jk} . In other words, we can consider one univariate inference problem $p(d_{jk} | y)$ at a time. Even if the prior probability model $p(d)$ is not marginally independent across d_{jk} , it typically assumes independence conditional on hyperparameters, still leaving a considerable simplification of posterior simulation.

The above detailed explanation serves to highlight two critical assumptions. Posterior independence, conditional on hyperparameters or marginally, only holds for equally spaced data and under *a priori* independence over d_{jk} . In most applications prior independence is a technically convenient assumption, but does not reflect genuine prior knowledge. However, incorporating assumptions about prior dependence is not excessively difficult either. Starting with an assumption about dependence of $g(x_i)$, $i = 1, \dots, n$, Vannucci and Corradi (1999) show that a straightforward two dimensional wavelet transform can be used to derive the corresponding covariance matrix for the wavelet coefficients d_{jk} .

In the absence of equally spaced data the convenient mapping of the raw data y_i to the empirical wavelet coefficients \hat{d}_{jk} is lost. The same is true for inference problems other than regression where wavelet decomposition is used to model random functions. Typical examples are the unknown density in a density estimation (Müller and Vidakovic 1998), or the spectrum in a spectral density estimation (Müller and Vidakovic 1999). In either case evaluation of the likelihood $p(y|d)$ requires reconstruction of the random function $g(\cdot)$. Although a technical inconvenience, this does not hinder the practical use of a wavelet basis. The super-fast wavelet decomposition and reconstruction algorithms still allow computationally efficient likelihood evaluation even with the original raw data.

3.4 Neural Networks

Neural networks are another popular approach following the general theme of defining random functions by probability models for coefficients with respect to an appropriate basis. Now the basis are rescaled versions of logistic functions. Let $\Psi(\eta) = \exp(\eta)/(1 + \exp(\eta))$, then $g(x) = \sum_{j=1}^M \beta_j \Psi(x' \gamma_j)$ can be used to represent a random function g . The random function is parametrized by $\theta = (\beta_1, \gamma_1, \dots, \beta_M, \gamma_M)$. Bayesian inference proceeds by assuming an appropriate prior probability model and considering posterior updating conditional on the observed data. Recent reviews of statistical inference for neural networks in regression models appear in Cheng and Titterton (1994) and Stern (1996). Neal (1996) and Müller and Ríos-Insua (1998) discuss specifically Bayesian inference in such models. Ríos-Insua and Müller (1998) argue to include the number of components M in the parameter vector and consider inference over “variable architecture” neural network models. Lee (2001) compares alternative Bayesian model selection criteria for neural networks.

3.5 Other Nonparametric Regression Methods

Alternatively to modeling the random function g , the nonparametric regression problem can be reduced to a density estimation problem by proceeding as if the pairs (x_i, y_i) were an i.i.d. sample, $(x_i, y_i) \sim F(x, y)$, from some unknown distribution F . Inference on F implies inference on the conditional means process $g_F(x) \equiv E_F(y|x)$. Müller, Erkanli and West (1996) propose this approach using a DP mixture model for inference on the unknown joint distribution F . Regression curves g estimated under this approach take the form of locally weighted linear regression lines, similar to traditional kernel regression in classical nonparametric inference. Considering (x_i, y_i) as an i.i.d. sample – wrongly – introduces an additional factor $\prod F(x_i)$ in the likelihood $\prod F(x_i, y_i) = \prod F(x_i) F(y_i|x_i)$ and thus provides only approximate inference.

An interesting approach to isotonic regression is pursued in Lavine and Mockus (1995)

who use a rescaled cumulative density function F to model a regression mean curve $g(x) = a + b F(x)$. Assuming a DP prior for F they implement nonparametric Bayesian inference.

Newton, Czado and Chappell (1996) propose a modified DP, constraining the random probability measure to median 0 and fixed length central interval (like the interquartile range, for example). The modified DP is used to define a link F in a nonparametric binary regression model with $P(y_i = 1) = F(x'_i\beta)$.

4 SURVIVAL ANALYSIS

Survival analysis involves modeling the time until a certain event occurs (survival times), often including a regression on covariates. In most applications, the data is subject to right-censoring. Let x_1, \dots, x_n denote the survival times, $x_i \sim F(\cdot)$. Let C_1, \dots, C_n denote the (possibly random) censoring times. The actually observed data is a collection of pairs $(T_1, I_1), \dots, (T_n, I_n)$ with censored observations $T_i = \min\{x_i, C_i\}$ and censoring indicators $I_i = I\{x_i \leq C_i\}$. Interval and other types of censoring could be also considered in a similar fashion. Two quantities are of primary interest in survival analysis: the survival function $S(t) = 1 - F(t)$ and the hazard rate function $\lambda(t) = F'(t)/S(t)$. It turns out that the integrated or cumulative hazard function $\Lambda(t) = \int_0^t \lambda(s) ds$ is simpler to estimate, and there is a one-to-one correspondence between $S(t)$ and $\Lambda(t)$, given by $S(t) = \exp(-\Lambda(t))$.

Assuming C_1, \dots, C_n to be constant, Susarla and Van Ryzin (1976) discuss inference with a DP prior on F . The posterior mean converges to Kaplan and Meier's (1958) product limit estimate as the total mass parameter $M \rightarrow 0^+$. More recently, Florens and Rolin (2001) provided a closed form description of the posterior process under a DP prior and random censoring times. The characterization is quite useful for posterior simulation of functionals of the posterior distribution of F . For a review of related approaches applying the DP to similar problems see Ferguson, Phadia and Tiwari (1992). Doss (1994) studied an MDP model for survival data subject to more general censoring schemes. Evaluation of the posterior mean of F is done through an interesting MCMC scheme that involves DP draws using a composition method. Convergence of the algorithm is also discussed.

4.1 Neutral to the Right Processes

Many stochastic process priors that have been proposed as nonparametric prior distributions for survival data analysis belong to the class of neutral to the right (NTTR) processes. An RPM $F(t)$ is an NTTR process on the real line, if it can be expressed as $F(t) = 1 - \exp(-Y(t))$, where $Y(t)$ is a stochastic process with independent increments, almost surely right-continuous and non-decreasing with $P\{Y(0) = 0\} = 1$ and $P\{\lim_{t \rightarrow \infty} Y(t) = \infty\} = 1$ (Doksum 1974). Walker et al. (1999) call $Y(t)$ an NTTR Lévy process. Doksum

(1974) showed that the posterior for a NTTR prior and i.i.d. sampling is again a NTTR process. Ferguson and Phadia (1979) showed that for right censored data the class of NTTR process priors remains closed, i.e., the posterior is still a NTTR process.

NTTR processes are used in many approaches that construct probability models for $\lambda(t)$ or $\Lambda(t)$, rather than directly for F . Dykstra and Laud (1981) define the extended Gamma process, generalizing the Gamma process studied in Ferguson (1973). The idea is to consider first an NTTR process $\{Y(t)\}$ such that for $Y(t) - Y(s) \sim \Gamma(\alpha(t) - \alpha(s), 1)$ for all $t > s \geq 0$, where $\alpha(t)$ is a nondecreasing left-continuous function on $[0, \infty)$. The new process is defined as $\int_0^t \beta(s) dY(s)$ for a positive right-continuous function $\beta(t)$. Dykstra and Laud (1981) consider such processes on the hazard function $\lambda(t)$, studying their properties and obtaining estimates of the posterior hazard function without censoring and with right-censoring. In particular, the resulting function $\lambda(t)$ is monotone.

An alternative model was proposed by Hjort (1990), by placing a Beta process prior on $\Lambda(t)$. To understand this construction, let us look at a discrete version of the process first. Following Nieto-Barajas and Walker (2002b), consider a partition of the time axis $0 = \tau_0 < \tau_1 < \tau_2 \dots$, and failures occurring at times chosen from the set $\{\tau_1, \tau_2, \dots\}$. Let λ_j denote the hazard at time τ_j , $\lambda_j = P(x = \tau_j | x \geq \tau_j)$. Hjort (1990) assumes independent, beta-distributed priors for $\{\lambda_j\}$. This generates a discrete process with independent increments for the cumulative hazard function $\Lambda(\tau_j) = \sum_{i=0}^j \lambda_i$. The class is closed under prior to posterior updating as the posterior process is again of the same type. The continuous version of this discrete Beta process is derived by a limit argument as the interval lengths $\tau_j - \tau_{j-1}$ approach zero (for details, see Hjort 1990). Full Bayesian inference for a model with a Beta process prior for the cumulative hazard function using Gibbs sampling can be found in Damien, Laud and Smith (1996). A variation of this idea was used by Walker and Mallick (1997). Specifically, they assumed $\lambda(t)$ to be constant at $\lambda_1, \lambda_2, \dots$ over the intervals $[0, \tau_1], (\tau_1, \tau_2], \dots$ with independently distributed gamma priors on $\{\lambda_j\}$. As pointed out in Nieto-Barajas and Walker (2002b), there is no limit version of this process.

Since a NTTR process $Y(t)$ has at most a countable number of discontinuity points, it turns out that every NTTR process can be decomposed as the sum of a continuous component and a purely jump component. This observation is very useful for simulation purposes (Walker and Damien 1998, Walker et al. 1999). To simulate from the jump component, Walker and Damien (1998) suggest using methods discussed in Walker (1995) or the latent variables method of Damien, Wakefield and Walker (1999), depending on the specific form adopted by the density to sample from. To simulate from the continuous part Walker and Damien (1998) note that a random variable arising from this component is infinitely divisible and build on a method originally proposed by Bondesson (1982), but discarded by the same author due to the practical implementation difficulties arising at that time. Wolpert

and Ickstadt (1998a) proposed an alternative method for approximately sampling from the continuous part, called the Inverse Lévy Measure (ILM) algorithm. It is based on the result that any nonnegative infinitely divisible distribution can be represented as the distribution at time $t = 1$ of an increasing stochastic process X_t (called *subordinator*) with stationary and independent increments. The Lévy-Khintchine Theorem (e.g. page 163 of Durrett 1996) states that the characteristic function of such distribution satisfies

$$\log(\varphi(t)) = ict - \frac{\sigma^2 t^2}{2} + \int_{\mathbb{R}} \left(e^{itx} - 1 - \frac{itx}{1+x^2} \right) \nu(dx),$$

where ν is called the Lévy measure and is such that $\nu(\{0\}) = 0$ and $\int_{\mathbb{R}} \frac{x^2}{1+x^2} \nu(dx) < \infty$. Therefore, to simulate the process X_t over an interval $[0, T]$ we can proceed as follows: generate independent jump times σ_m from the uniform distribution on $[0, T]$, jumps τ_m from a unit-rate Poisson process, define $\nu_m = \inf\{u \geq 0 : \nu([u, \infty)) \leq \tau_m/T\}$, and set $X_t = \sum\{\nu_m : \sigma_m \leq t\}$. This summation defining X_t will have a finite number of terms if and only if $\nu([0, \infty)) < \infty$. Thus, in general the method leads to an approximate simulation. The name ILM comes from the fact that $\nu_m = L^{-1}(\tau_m/T)$ where $L(u) = \nu([u, \infty))$. See additional details in Wolpert and Ickstadt (1998b).

4.2 Dependent Increments Models

We have already discussed independent increments models for the cumulative hazard function $\Lambda(t)$. In the discrete version this implies independence for the hazards $\{\lambda_j\}$. A different modeling perspective is obtained by assuming dependence. A convenient way to introduce dependence is a Markovian process prior on $\{\lambda_k\}$. Gamerman (1991) proposes the following model: $\log(\lambda_j) = \log(\lambda_{j-1}) + \epsilon_j$ for $j \geq 2$, where $\{\epsilon_j\}$ are independent with $E(\epsilon_j) = 0$ and $Var(\epsilon_j) = \sigma^2 < \infty$. In the linear Bayesian method of Gamerman (1991) only a partial specification of the $\{\epsilon_j\}$ is required. The resulting model extends Leonard's (1978) smoothness prior for density estimation, stated also in terms of a discrete survival formulation, but under the assumption that $\epsilon_j \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$.

Later, Gray (1994) used a similar prior process but directly on the hazards $\{\lambda_j\}$, without the log transformation. A further generalization involving a martingale process was proposed in Arjas and Gasbarra (1994). More recently, Nieto-Barajas and Walker (2002b) proposed a model based on a latent process $\{u_k\}$ such that $\{\lambda_j\}$ is included as

$$\lambda_1 \rightarrow u_1 \rightarrow \lambda_2 \rightarrow u_2 \rightarrow \dots$$

and the pairs (u, λ) are generated from conditional densities $f(u|\lambda)$ and $f(\lambda|u)$ implied by a specified joint density $f(u, \lambda)$. The main idea is to ensure linearity in the conditional expectation: $E(\lambda_{k+1}|\lambda_k) = a_k + b_k \lambda_k$. Nieto-Barajas and Walker (2002b) show that both, the

gamma process of Walker and Mallick (1997) and the discrete Beta process of Hjort (1990) are obtained as special cases of their construction, under appropriate choices of $f(u, \lambda)$.

In the continuous case, Nieto-Barajas and Walker (2002b) proposed a Markovian model where the hazard rate function is modelled as

$$\lambda(t) = \int_0^t \exp\{-a(t-u)\} dL(u), \quad (7)$$

for $a > 0$, and where $L(t)$ is a pure jump process, i.e., an independent increments process on $[0, \infty)$ without Gaussian components (Ferguson and Klass 1972, Walker and Damien 2000). This model, called Lévy driven Markov process, extends Dykstra and Laud's (1981) proposal by allowing non-monotone sample paths for $\lambda(t)$. In addition, the sample paths are piecewise continuous functions. Nieto-Barajas and Walker (2002b) obtain posterior distributions under (7) for different types of censoring and discuss applications in several special cases, including the Markov-Gamma process.

4.3 Competing Risks Model

An interesting extension of survival models considers a system with r components arranged in series. Here x_1, \dots, x_r are the failure times of the components and we observe (T, I) , where $T = \min\{x_1, \dots, x_r\}$ and $I = j$ if $T = x_j$. This setup is known as the competing risks model with r sources of failure. The survival function for the j th component is $S_j(t) = P(x_j > t)$ and the subsurvival function is $S_j^*(t) = P(T > t, I = j)$. The system survival function is $S(t) = P(T > t) = \sum_{j=1}^r S_j^*(t)$. Let $\mathbf{x}_i = (x_{i1}, \dots, x_{ir})$, $i = 1, \dots, n$ be a sample from the latent x_1, \dots, x_r failure times. The actual observed data are $(T_1, I_1), \dots, (T_n, I_n)$. Salinas-Torres, Pereira and Tiwari (1997) introduced the multivariate DP as a nonparametric model for the joint distribution of the failure times $\mathbf{x}_1, \dots, \mathbf{x}_n$. Let F_{01}, \dots, F_{0r} be distribution functions on the appropriate space and M_1, \dots, M_r be positive mass parameters, and let $v = (v_1, \dots, v_r) \sim D(M_1, \dots, M_r)$. Then $\mathbf{P} = (v_1 P_1, \dots, v_r P_r)$ is called a multivariate DP or dimension r if $P_j \sim \mathcal{D}(M_j, F_{0j})$.

Consider now a given risk subset $\Delta \subset \{1, \dots, r\}$ and let Δ^c be its complement. The corresponding subsurvival and survival functions are given by $S_\Delta^*(t) = P(T > t, I \in \Delta)$ and $S_\Delta(t) = P(\min_{j \in \Delta} x_j > t)$. The data structure obtained for the case $r = 2$, $\Delta = \{1\}$ and $\Delta^c = \{2\}$ reduces to the usual right-censored problem with random censoring times. Peterson (1977) gives an expression for the survival function $S_\Delta(t)$ in terms of the subsurvival functions $S_\Delta^*(t)$ and $S_{\Delta^c}^*$:

$$S_\Delta(t) = \varphi(S_\Delta^*(t), S_{\Delta^c}^*(t)), \quad \text{for } t \leq t^* = \min\{t_{S_\Delta}, t_{S_{\Delta^c}}\}, \quad (8)$$

where

$$\varphi(H, G; t) = \exp \left(\int_0^t \frac{dH(s)}{H(s) + G(s)} \right) \prod_t \left\{ \frac{H(s_+) + G(s_+)}{H(s_-) + G(s_-)} \right\},$$

and $t_{S_\Delta} = \sup\{t : S_\Delta(t) > 0\}$. Here, \int_0^t represents integration over the union of intervals of continuity points of H that are less than or equal to t , and \prod_t represents a product over the discontinuity points of H that are less than or equal to t (we assume that $S_\Delta^*(t)$ and $S_{\Delta^c}^*(t)$ have no common discontinuities). In this setting, Salinas-Torres, Pereira and Tiwari (2002) derived Bayes estimates of $S_\Delta(t)$ under quadratic loss function. The estimate has the property that it can be obtained by substituting the Bayes estimates for S_Δ^* and $S_{\Delta^c}^*$ into (8).

4.4 Models Based on Proportional Hazards

So far we have discussed survival analysis models without covariates. To incorporate covariates, the most popular choice is the proportional hazards model, introduced in Cox (1972). Assuming T_1, \dots, T_n are the failure times of n individuals, the hazard rate functions are modelled as

$$\lambda_i(t) = \lambda_0(t) \exp \{ \mathbf{Z}_i(t)^T \boldsymbol{\beta} \}, \quad i = 1, \dots, n \quad (9)$$

where $\mathbf{Z}_i(t)$ is the p -dimensional vector of covariates for the i th individual at time $t > 0$, $\boldsymbol{\beta}$ is the vector of regression coefficients, and $\lambda_0(t)$ is the baseline hazard rate function.

Semi-parametric approaches to inference in (9) consider a nonparametric specification of $\lambda_0(t)$. A model based on an independent increments gamma process was proposed by Kalbfleisch (1978) who studied its properties and estimation. Extensions of this model to neutral to the right processes was discussed in Wild and Kalbfleisch (1981). In the context of multiple event time data, Sinha (1993) considered an extension of Kalbfleisch's (1978) model for $\lambda_0(t)$. The proposal assumes the events are generated by a counting process with intensity given by a multiplicative expression similar to (9), but including an indicator of the censoring process, and individual frailties to accommodate the multiple events occurring per subject. Sinha (1993) discusses posterior inference for this model using Gibbs sampling, under the assumption of gamma distributed frailties. Extensions of this model to the case of positive stable frailty distributions and a correlated prior process with piecewise exponential hazards can be found in Qiou, Ravishanker and Dey (1999). See additional comments, details on computational strategies and extensions to multivariate survival data in Sinha and Dey (1998).

Other modeling approaches based on (9) have been studied in the literature. Laud, Damien and Smith (1998) consider (9) using a Beta process prior for $\Lambda(t)$, and proposing an MCMC implementation for full posterior inference. Nieto-Barajas and Walker (2001)

propose their flexible Lévy drive Markov process (Nieto-Barajas and Walker 2002a) to model $\lambda_0(t)$, and allowing for time dependent covariates. Full posterior inference is achieved via substitution sampling.

Accelerated failure time models are an alternative framework to introduce regression in survival analysis. Instead of introducing the regression in the log hazard, as in (9), the generic accelerated failure time model assumes that failure times T_i arise as $\log T_i = -\mathbf{Z}'_i\boldsymbol{\beta} + \log(x_i)$. Nonparametric approaches assume a probability model for the unknown distribution of $\log(x_i)$. Models based on DP priors appear in Johnson and Christensen (1989) and Kuo and Mallick (1997). Walker and Mallick (1999) propose an alternative PT prior model.

5 HIERARCHICAL MODELS

An important application of nonparametric approaches arises in modeling random effects distributions in hierarchical models. Often little is known about the specific form of the random effects distribution. Assuming a specific parametric form is typically motivated by technical convenience rather than by genuine prior beliefs. Although inference about the random effects distribution itself is rarely of interest, it can have implications on the inference of interest. Thus it is important to allow for population heterogeneity, outliers, skewness etc.

In the context of a traditional randomized block ANOVA model with subject specific random effects z_i a Bayesian nonparametric model can be used to allow for more general random effects distributions. Bush and MacEachern (1996) propose a DP prior for $z_i \sim G$, $G \sim \mathcal{D}(G_0, M)$. Kleinman and Ibrahim (1998) propose the same approach in a more general framework for a linear model with random effects. They discuss an application to longitudinal random effects models. Müller and Rosner (1997) use DP mixture of normals to avoid the awkward discreteness of the implied random effects distribution. Also, the additional convolution with a normal kernel significantly simplifies posterior simulation for sampling distributions beyond the normal linear model. Mukhopadhyay and Gelfand (1997) implement the same approach in generalized linear models with linear predictor $z_i + x'_i\boldsymbol{\beta}$ and a DP mixture model for the random effect z_i . In Wang and Taylor (2001) random effects W_i are entire longitudinal paths for each subject in the study. They use integrated Ornstein-Uhlenbeck stochastic process priors for $W_i S(t)$.

A further complication arises when the model hierarchy in a hierarchical model continues beyond the nonparametric model, i.e., if the nonparametric model appears in a submodel of the larger hierarchical model. For example, in a hierarchical analysis of related clinical studies there might be a different random effects distribution in each of the related clinical

trials. Let G_i denote the random distribution or random function in submodel i . Assuming a nonparametric model $p(G_i)$ for the i -th submodel, model completion requires an additional assumption about the joint distribution of $\{G_i, i \in I\}$. Using DP priors, $G_i \sim \mathcal{D}(G_i^o, M)$, marginally for each G_i , a conceptually straightforward approach is to link the base measures G_i^o . For example, the base measure G_i^o could include a regression on covariates specific to the i -th submodel. This construction is introduced in Cifarelli and Regazzini (1978) as mixture of products of Dirichlet process. The model is used, for example, in Muliere and Petrone (1993) who define dependent nonparametric models $F_x \sim \mathcal{D}(M, F_x^o)$ by assuming a regression in the base measure $F_x^o = N(\beta x, \sigma^2)$. Similar models are discussed in Mira and Petrone (1996) and Giudici et al. (2003). Carota and Parmigiani (2002) and Dominici and Parmigiani (2001) use the same approach to model random distributions $G_i \sim \mathcal{D}(G_i^o, M_i)$ centered around, among other choices, a Binomial base measure $G_i^o = \text{Bin}(\theta_i^o, N_i)$, including the total mass parameter M_i in the hierarchy. Both, the binomial success probability θ_i^o and the total mass parameter M_i are modelled as a regression on covariates d_i , specific to submodel i .

Linking the related nonparametric models through a regression on the parameters of the nonparametric models limits the nature of the dependence to the structure of this regression. MacEachern (1999) proposes the dependent DP (DDP) as an alternative approach to define a dependent prior model for a set of random measures $\{G_x\}$, with $G_x \sim \mathcal{D}$ marginally. Recall the stick breaking representation (2) for the DP random measure, $G_x = \sum_h w_{xh} \delta(\mu_{xh})$. The key idea behind the DDP is to introduce dependence across the measures G_x by assuming the distribution of the point masses μ_{xh} dependent across different levels of x , but still independent across h . In the basic version of the DDP the weights are assumed the same across x , i.e., $w_{xh} = w_h$. To introduce dependence of μ_{xh} across x MacEachern (1999) uses a Gaussian process. DeIorio et al. (2002) construct the ANOVA DDP as a joint probability model for dependent random measures. They consider a family of unknown probability measures F_x indexed by categorical factors x . For example, in a clinical trial, F_x might be the random effects distribution for patients with categorical covariates x . Covariates might include treatment levels, etc. Dependence across $\{F_x\}$ is induced by assuming ANOVA models on μ_{xh} across x .

6 MODEL VALIDATION

An interesting use of nonparametric Bayesian inference arises in model validation. One way to validate a proposed parametric model is to consider a nonparametric extension and report appropriate summaries of a comparison of the parametric and the nonparametric fit.

Carota and Parmigiani (1996) and Carota et al. (1996) discuss such approaches using DP extensions and point out the limitations of formalizing the comparison with a Bayes factor. Due to the discrete nature of the Dirichlet process RPM inference is driven by the number of duplicates in the data set. They suggest, among other approaches, to consider KL-divergence of prior to posterior on the random probability model. Conigliani, Castro and O’Hagan (2000) discuss a similar approach, using fractional Bayes factors to summarize the comparison.

Berger and Guglielmi (2001) take up the same theme, but replace the DP prior with a PT model. To center the PT model at a parametric model $f(x|\theta)$ they construct PTs with mean measure $f(x|\theta)$. They fix the nested partition sequence and set the parameters α_ϵ for the random probabilities such that the desired mean is achieved. Computation of Bayes factors for the model validation is greatly simplified by the availability of a closed form expression for the marginal distribution under such PT models:

$$m(x_1, \dots, x_n|\theta) = \prod f(x_i|\theta) \prod_{j=2}^n \prod_{m=1}^{m^*(x_j)} \frac{\alpha'_{\epsilon_m(x_j)}(\alpha_{\epsilon_{m-1}0(x_j)} + \alpha_{\epsilon_{m-1}1(x_j)})}{\alpha_{\epsilon_m(x_j)}(\alpha'_{\epsilon_{m-1}0(x_j)} + \alpha'_{\epsilon_{m-1}1(x_j)})}$$

The α_ϵ are the Beta distribution parameters in the definition of the PT, as defined in Section 2.3. The indices $\epsilon_m(x_j) = \epsilon_1 \dots \epsilon_m$ identifies the partitioning subset $B_{\epsilon_1 \dots \epsilon_m}$ of level m that contains x_j , i.e., $x_j \in B_\epsilon$, and α'_ϵ are the parameters of the posterior PT, given the observations (x_1, \dots, x_{j-1}) . The upper bound $m^*(x_j)$ in the product is the smallest level m such that no $x_i, i < j$ belongs to the same partitioning subset $B_{\epsilon_m(x_j)}$ as x_j at level m . The α sequences depend on the parameter θ . Evaluation of Bayes factors of the parametric model versus nonparametric extension requires one more step of marginalization to marginalize w.r.t. θ . Berger and Guglielmi (2001) describe suitable numerical methods.

A related approach is pursued in Mazzuchi, Soofi and Soyer (2000). They consider parametric models defined as maximum entropy models in a moment class. This includes the exponential, gamma, Weibull, normal, etc. By considering the posterior expected KL divergence between the parametric model and a nonparametric extension centered at that parametric model they define a diagnostic of fit. For the nonparametric extension they use a DP model centered at the maximum entropy parametric model.

7 CONCLUSION

We have reviewed some important aspects of nonparametric Bayesian inference. Rather than attempting a complete catalog of existing methods we focused on typical modeling strategies in important inference problems. Also, we emphasized recent developments over a historical perspective. The chosen classification of Bayesian nonparametric approaches into the listed

application areas is an arbitrary subjective choice, leading us to miss some interesting nonparametric Bayesian methods that did not fit cleanly into one of these arbitrary categories. Typical examples are Quintana (1998) and Lee and Berger (1999), discussing nonparametric approaches to modeling contingency tables and selection sampling, respectively.

An important aspect of nonparametric Bayesian inference that we excluded from the discussion are computational issues. Many approaches are driven by what are essentially computational concerns. Another important line of research that we excluded from the discussion are the many methods that are nonparametric in flavor even if they are not technically inference in infinite dimensional parameter spaces. Typical examples are finite mixture model. Such models often provide flexible inference very much like corresponding nonparametric extensions.

Finally, we did not discuss methods that are nonparametric Bayes in the literal sense, rather than in the sense of the technical definition we gave in the introduction. A typical example is Lavine (1995) who discusses inference based on a partial likelihood argument.

References

- Antoniak, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems, *The Annals of Statistics* **2**: 1152–1174.
- Arjas, E. and Gasbarra, D. (1994). Nonparametric Bayesian inference from right censored survival data, using the Gibbs sampler, *Statistica Sinica* **4**: 505–524.
- Arjas, E. and Heikkinen, J. (1997). An algorithm for nonparametric bayesian estimation of a poisson intensity, *Computational Statistics* **12**: 385–402.
- Barron, A., Schervish, M. J. and Wasserman, L. (1999). The consistency of posterior distributions in nonparametric problems, *The Annals of Statistics* **27**: 536–561.
- Berger, J. and Guglielmi, A. (2001). Bayesian testing of a parametric model versus nonparametric alternatives, *Journal of the American Statistical Association* **96**: 174–184.
- Bernardo, J. M. and Smith, A. F. M. (1994). *Bayesian theory*, New York: Wiley.
- Blackwell, D. and MacQueen, J. B. (1973). Ferguson distributions via Pólya urn schemes, *The Annals of Statistics* **1**: 353–355.
- Bondesson, L. (1982). On simulation from infinitely divisible distributions, *Advances in Applied Probability* **14**: 855–869.

- Bush, C. A. and MacEachern, S. N. (1996). A semiparametric Bayesian model for randomised block designs, *Biometrika* **83**: 275–285.
- Carota, C. and Parmigiani, G. (1996). On Bayes factors for nonparametric alternatives, in J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 5 – Proceedings of the Fifth Valencia International Meeting*, pp. 507–511.
- Carota, C. and Parmigiani, G. (2002). Semiparametric regression for count data, *Biometrika* **89**: 265–281.
- Carota, C., Parmigiani, G. and Polson, N. G. (1996). Diagnostic measures for model criticism, *Journal of the American Statistical Association* **91**: 753–762.
- Cheng, B. and Titterton, D. M. (1994). Neural networks: A review from a statistical perspective (Disc: p31-54), *Statistical Science* **9**: 2–30.
- Chipman, H. A., George, E. I. and McCulloch, R. E. (1998). Bayesian CART model search (C/R: p948-960), *Journal of the American Statistical Association* **93**: 935–948.
- Chipman, H. A., Kolaczyk, E. D. and McCulloch, R. E. (1997). Adaptive Bayesian wavelet shrinkage, *Journal of the American Statistical Association* **92**: 1413–1421.
- Cifarelli, D. M. and Melilli, E. (2000). Some New Results for Dirichlet Priors, *The Annals of Statistics* **28**: 1390–1413.
- Cifarelli, D. and Regazzini, E. (1978). Problemi statistici non parametrici in condizioni di scambiabilità parziale e impiego di medie associative, *Technical report*, Quaderni dell’Istituto di Matematica Finanziaria, Torino.
- Clyde, M. and George, E. (2000). Flexible empirical bayes estimation for wavelets, *Journal of the Royal Statistical Society, Series B* **62**: 681–698.
- Conigliani, C., Castro, J. I. and O’Hagan, A. (2000). Bayesian assessment of goodness of fit against nonparametric alternatives, *Canadian Journal of Statistics* **28**: 327–342.
- Cox, D. R. (1972). Regression models and life-tables (with discussion), *Journal of the Royal Statistical Society, Series B* **34**: 187–220.
- Dalal, S. R. (1979). Dirichlet invariant processes and applications to nonparametric estimation of symmetric distribution functions, *Stochastic Processes and their Applications* **9**: 99–108.

- Damien, P., Laud, P. and Smith, A. (1996). Implementation of Bayesian Nonparametric Inference Using Beta Processes, *Scandinavian Journal of Statistics* **23**: 27–36.
- Damien, P., Wakefield, J. and Walker, S. (1999). Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables, *Journal of the Royal Statistical Society, Series B* **61**: 331–344.
- DeIorio, M., Müller, P., Rosner, G. L. and MacEachern, S. N. (2002). ANOVA DDP models: A review, in D. D. Denison, M. H. Hansen, C. C. Holmes, B. Mallick and B. Yu (eds), *Nonlinear Estimation and Classification*, Springer-Verlag, p. 467.
- Denison, D. G. T., Mallick, B. K. and Smith, A. F. M. (1998a). A Bayesian CART algorithm, *Biometrika* **85**: 363–377.
- Denison, D. G. T., Mallick, B. K. and Smith, A. F. M. (1998b). Automatic Bayesian Curve Fitting, *Journal of the Royal Statistical Society Series B* **60**: 333–350.
- Denison, D. G. T., Mallick, B. K. and Smith, A. F. M. (1998c). Bayesian MARS, *Statistics and Computing* **8**: 337–346.
- Dey, D., Müller, P. and Sinha, D. (eds) (1998). *Practical Nonparametric and Semiparametric Bayesian Statistics*, Springer-Verlag.
- Diaconis, P. and Freedman, D. (1986). On the consistency of Bayes estimates (with discussion), *The Annals of Statistics* **14**: 1–67.
- Diaconis, P. and Kemperman, J. (1996). Some new tools for Dirichlet priors, in J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 5. Proceedings of the Fourth Valencia International Meeting*, Oxford University Press, pp. 97–106.
- DiMatteo, I., Genovese, C. R. and Kass, R. (2001). Bayesian Curve Fitting with Free-Knot Splines, *Biometrika* **88**: 1055–1071.
- Doksum, K. (1974). Tailfree and Neutral Random Probabilities and Their Posterior Distributions, *Annals of Probability* **2**: 183–201.
- Dominici, F. and Parmigiani, G. (2001). Bayesian semi-parametric analysis of developmental toxicology data, *Biometrics* **57**: 150–157.
- Doss, H. (1994). Bayesian Nonparametric Estimation for Incomplete Data Via Successive Substitution Sampling, *Annals of Statistics* **22**: 1763–1786.

- Durrett, R. (1996). *Probability: Theory and Examples, Second Edition*, Belmont, California: Duxbury Press.
- Dykstra, R. L. and Laud, P. (1981). A Bayesian nonparametric approach to reliability, *The Annals of Statistics* **9**: 356–367.
- Escobar, M. (1988). *Estimating the means of several normal populations by estimating the distribution of the means*, PhD thesis, Yale University.
- Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures, *Journal of the American Statistical Association* **90**: 577–588.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems, *The Annals of Statistics* **1**: 209–230.
- Ferguson, T. S. and Klass, M. J. (1972). A representation of independent increment processes without Gaussian components, *The Annals of Mathematical Statistics* **43**: 1634–1643.
- Ferguson, T. S. and Phadia, E. G. (1979). Bayesian nonparametric estimation based on censored data, *The Annals of Statistics* **7**: 163–186.
- Ferguson, T. S., Phadia, E. G. and Tiwari, R. C. (1992). Bayesian nonparametric inference, in M. Ghosh and P. K. Pathak (eds), *Current Issues in Statistical Inference: Essays in Honor of D. Basu*, Hayward, California: IMS Lecture Notes - Monograph Series, Vol. 17.
- Florens, J.-P. and Rolin, J.-M. (2001). Simulation of posterior distributions in nonparametric censored analysis, *Technical report*, Université des Sciences Sociales de Toulouse.
- Friedman, J. H. (1991). Multivariate adaptive regression splines (Disc: p67-141), *The Annals of Statistics* **19**: 1–67.
- Gamerman, D. (1991). Dynamic bayesian models for survival data, *Applied Statistics* **40**: 63–79.
- Gasparini, M. (1996). Bayesian density estimation via mixture of Dirichlet processes, *Journal of Nonparametric Statistics* **6**: 355–366.
- Gelfand, A. E. and Kottas, A. (2002). A Computational Approach for Full Nonparametric Bayesian Inference under Dirichlet Process Mixture Models, *Journal of Computational and Graphical Statistics* **11**: 289–305.
- Gelfand, A. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities, *Journal of the American Statistical Association* **85**: 398–409.

- Ghosal, S. (2001). Convergence rates for density estimation with Bernstein polynomials, *Annals of Statistics* **29**: 1264–1280.
- Ghosal, S., Ghosh, J. K. and Ramamoorthi, R. V. (1999). Posterior consistency of Dirichlet mixtures in density estimation, *Annals of Statistics* **27**: 143–158.
- Giudici, P., Mezzetti, M. and Muliere, P. (2003). Mixtures of dirichlet process priors for variable selection in survival analysis, *Journal of statistical planning and inference* **111**: 101–115.
- Gray, R. J. (1994). A Bayesian analysis of institutional effects in a multicenter cancer clinical trial, *Biometrics* **50**: 244–253.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination, *Biometrika* **82**: 711–732.
- Hanson, T. and Johnson, W. (2002). Modeling regression error with a mixture of polya trees, *Journal of the American Statistical Association* **97**: 1020–1033.
- Hjort, N. L. (1990). Nonparametric Bayes estimators based on beta processes in models for life history data, *The Annals of Statistics* **18**: 1259–1294.
- Ishwaran, H. and James, L. F. (2001). Gibbs Sampling Methods for Stick-Breaking Priors, *Journal of the American Statistical Association* **96**: 161–173.
- Ishwaran, H. and James, L. F. (2002). Approximate Dirichlet process computing in finite normal mixtures: smoothing and prior information, *Journal of Computational and Graphical Statistics* **11**: 508–532.
- Ishwaran, H. and James, L. F. (2003). Generalized weighted Chinese restaurant processes for species sampling mixture models, *Statistica Sinica* . In press.
- Ishwaran, H. and Takahara, G. (2002). Independent and identically distributed Monte Carlo algorithms for semiparametric linear mixed models, *Journal of the American Statistical Association* **97**: 1154–1166.
- Ishwaran, H. and Zarepour, M. (2000). Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models, *Biometrika* **87**: 371–390.
- Johnson, W. and Christensen, R. (1989). Bayesian nonparametric survival analysis for the accelerated failure time model, *Statistics and Probability Letters* **8**: 179–184.

- Kalbfleisch, J. D. (1978). Nonparametric Bayesian analysis of survival time data, *Journal of the Royal Statistical Society, Series B* **40**: 214–221.
- Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations, *Journal of the American Statistical Association* **53**: 457–481.
- Kleinman, K. and Ibrahim, J. (1998). A semi-parametric bayesian approach to the random effects model, *Biometrics* **54**: 921–938.
- Korwar, R. M. and Hollander, M. (1973). Contributions to the theory of Dirichlet processes, *The Annals of Probability* **1**: 705–711.
- Kottas, A. and Gelfand, A. E. (2001). Bayesian Semiparametric Median Regression Modeling, *Journal of the American Statistical Association* **96**: 1458–1468.
- Kuo, L. and Mallick, B. (1997). Bayesian semiparametric inference for the accelerated failure time model, *Canadian Journal of Statistics* **25**: 457–472.
- Laud, P., Damien, P. and Smith, A. F. M. (1998). Bayesian Nonparametric and Covariate Analysis of Failure Time Data, *in* Dey et al. (1998), pp. 213–225.
- Lavine, M. (1992). Some aspects of Polya tree distributions for statistical modelling, *The Annals of Statistics* **20**: 1222–1235.
- Lavine, M. (1994). More aspects of Polya tree distributions for statistical modelling, *The Annals of Statistics* **22**: 1161–1176.
- Lavine, M. (1995). On an approximate likelihood for quantiles, *Biometrika* **82**(1): 220–222.
- Lavine, M. and Mockus, A. (1995). A nonparametric bayes method for isotonic regression, *Journal for Statistical Planning and Inference* **46**: 235–248.
- Lee, H. (2001). Model selection for neural network classification, *Journal of Classification* **18**(2): 227–243.
- Lee, J. and Berger, J. (1999). Semiparametric bayesian analysis of selection models, *J. American Statistical Assoc.* **96**: 1397–1409.
- Lenk, P. (1988). The logistic normal distribution for bayesian,nonparametric predictive densities, *Journal of the American Statistical Association* **83**: 509–516.
- Leonard, T. (1978). Density estimation, stochastic processes and prior information, *Journal of the Royal Statistical Society, Series B* **40**: 113–116.

- Liu, J. S. (1996). Nonparametric hierarchical Bayes via sequential imputations, *The Annals of Statistics* **24**: 911–930.
- Lo, A. Y. (1984). On a class of Bayesian nonparametric estimates: I. Density estimates, *The Annals of Statistics* **12**: 351–357.
- MacEachern, S. (1994). Estimating normal means with a conjugate style Dirichlet process prior, *Communications in Statistics: Simulation and Computation* **23**: 727–741.
- MacEachern, S. (1999). Dependent nonparametric processes, *ASA Proceedings of the Section on Bayesian Statistical Science*, American Statistical Association, Alexandria, VA.
- MacEachern, S. N., Clyde, M. and Liu, J. S. (1999). Sequential Importance Sampling for Nonparametric Bayes Models: The Next Generation, *Canadian Journal of Statistics* **27**: 251–267.
- MacEachern, S. N. and Müller, P. (1998). Estimating Mixture of Dirichlet Process Models, *Journal of Computational and Graphical Statistics* **7**(2): 223–338.
- MacEachern, S. N. and Müller, P. (2000). Efficient MCMC Schemes for Robust Model Extensions using Encompassing Dirichlet Process Mixture Models, in F. Ruggeri and D. Ríos-Insua (eds), *Robust Bayesian Analysis*, New York:Springer-Verlag, pp. 295–316.
- Mazzuchi, T. A., Soofi, E. S. and Soyer, R. (2000). Computations of Maximum Entropy Dirichlet for Modeling Lifetime Data, *Computational Statistics and Data Analysis* **32**: 361–378.
- Mira, A. and Petrone, S. (1996). Bayesian hierarchical nonparametric inference for change-point problems, in J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 5*, Oxford University Press, Oxford.
- Mukhopadhyay, S. and Gelfand, A. (1997). Dirichlet process mixed generalized linear models, *Journal of the American Statistical Association* **92**: 633–639.
- Muliere, P. and Petrone, S. (1993). A Bayesian predictive approach to sequential search for an optimal dose: Parametric and nonparametric models, *Journal of the Italian Statistical Society* **2**: 349–364.
- Muliere, P. and Secchi, P. (1995). A note on a proper Bayesian Bootstrap, *Technical Report 18*, Università degli Studi di Pavia, Dipartimento di Economia Politica e Metodi Quantitativi.

- Muliere, P. and Tardella, L. (1998). Approximating distributions of random functionals of Ferguson-Dirichlet priors, *Canadian Journal of Statistics* **26**: 283–297.
- Müller, P., Erkanli, A. and West, M. (1996). Bayesian curve fitting using multivariate normal mixtures, *Biometrika* **83**: 67–79.
- Müller, P. and Ríos-Insua, D. (1998). Issues in bayesian analysis of neural network models, *Neural Computation* **10**: 571–592.
- Müller, P. and Rosner, G. (1997). A bayesian population model with hierarchical mixture priors applied to blood count data, *Journal of the American Statistical Association* **92**: 1279–1292.
- Müller, P. and Vidakovic, B. (1998). Bayesian inference with wavelets: Density estimation, *Journal of Computational and Graphical Statistics* **7**: 456–468.
- Müller, P. and Vidakovic, B. (1999). Mcmc methods in wavelet shrinkage: Non-equally spaced regression, density and spectral density estimation, *in* P. Müller and B. Vidakovic (eds), *Practical Nonparametric and Semiparametric Bayesian Statistics*, Springer-Verlag, pp. 187–202.
- Neal, R. M. (1996). *Bayesian learning for neural networks*, New York: Springer-Verlag.
- Neal, R. M. (2000). Markov Chain Sampling Methods for Dirichlet Process Mixture Models, *Journal of Computational and Graphical Statistics* **9**: 249–265.
- Newton, M. A., Czado, C. and Chappell, R. (1996). Semiparametric bayesian inference for binary regression, *Journal of the American Statistical Association* **91**: 142–153.
- Newton, M. A. and Zhang, Y. (1999). A recursive algorithm for nonparametric analysis with missing data, *Biometrika* **86**: 15–26.
- Nieto-Barajas, L. and Walker, S. G. (2001). A Semiparametric Bayesian Analysis of Survival Data Based on Markov Gamma Processes, *Technical report*, Department of Mathematical Sciences, University of Bath.
- Nieto-Barajas, L. and Walker, S. G. (2002a). Bayesian Nonparametric Survival Analysis via Lévy Driven Markov Processes, *Technical report*, Department of Mathematical Sciences, University of Bath.
- Nieto-Barajas, L. and Walker, S. G. (2002b). Markov beta and gamma processes for modelling hazard rates, *Scandinavian Journal of Statistics* **29**: 413–424.

- Paddock, S., Ruggeri, F., Lavine, M. and West, M. (2003). Randomised Polya Tree Models for Nonparametric Bayesian Inference, *Statistica Sinica* **13**: 443–460.
- Peterson, A. V. (1977). Expressing the Kaplan-Meier estimator as a function of empirical subsurvival functions, *Journal of the American Statistical Association* **72**: 854–858.
- Petrone, S. (1999a). Bayesian density estimation using Bernstein polynomials, *Canadian Journal of Statistics* **27**: 105–126.
- Petrone, S. (1999b). Random Bernstein polynomials, *Scandinavian Journal of Statistics* **26**: 373–393.
- Petrone, S. and Wasserman, L. (2002). Consistency of Bernstein Polynomial Posteriors, *Journal of the Royal Statistical Society, Series B* **64**: 79–100.
- Pitman, J. (1996). Some Developments of the Blackwell-MacQueen Urn Scheme, in T. S. Ferguson, L. S. Shapley and J. B. MacQueen (eds), *Statistics, Probability and Game Theory. Papers in Honor of David Blackwell*, Haywar, California: IMS Lecture Notes - Monograph Series, pp. 245–268.
- Pitman, J. and Yor, M. (1987). The Two-Parameter Poisson-Dirichlet Distribution Derived from a Stable Subordinator, *The Annals of Probability* **25**: 855–900.
- Qiou, Z., Ravishanker, N. and Dey, D. K. (1999). Multivariate survival analysis with positive stable frailties, *Biometrics* **55**: 637–644.
- Quintana, F. A. (1998). Nonparametric Bayesian Analysis for Assessing Homogeneity in $k \times l$ Contingency Tables with Fixed Right Margin Totals, *Journal of the American Statistical Association* **93**: 1140–1149.
- Quintana, F. A. and Newton, M. A. (2000). Computational aspects of Nonparametric Bayesian analysis with applications to the modeling of multiple binary sequences, *Journal of Computational and Graphical Statistics* **9**(4): 711–737.
- Regazzini, E. (1999). Old and recent results on the relationship between predictive inference and statistical modeling either in nonparametric or parametric form, in J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 6*, Oxford University Press, pp. 571–588.
- Ríos-Insua, D. and Müller, P. (1998). Feedforward Neural Networks for Nonparametric Regression, in D. Dey, P. Müller and D. Sinha (eds), *Practical Nonparametric and Semiparametric Bayesian Statistics*, New York: Springer-Verlag, pp. 181–193.

- Rolin, J.-M. (1992). Some useful properties of the Dirichlet process, *Technical Report 9207*, Center for Operations Research & Econometrics, Université Catholique de Louvain.
- Salinas-Torres, V. H., Pereira, C. A. B. and Tiwari, R. (1997). Convergence of Dirichlet measures arising in context of Bayesian analysis of competing risks models, *Journal of Multivariate Analysis* **62**: 24–35.
- Salinas-Torres, V. H., Pereira, C. A. B. and Tiwari, R. (2002). Bayesian Nonparametric Estimation in a Series System or a Competing-Risks Model, *Journal of Nonparametric Statistics* **14**: 449–458.
- Sethurman, J. (1994). A constructive definition of dirichlet priors, *Statistica Sinica* **4**: 639–650.
- Shively, T. S., Kohn, R. and Wood, S. (1999). Variable selection and function estimation in additive nonparametric regression using a data-based prior (C/R: p794-806), *Journal of the American Statistical Association* **94**: 777–794.
- Sinha, D. (1993). Semiparametric Bayesian analysis of multiple event time data, *Journal of the American Statistical Association* **88**: 979–983.
- Sinha, D. and Dey, D. K. (1997). Semiparametric Bayesian analysis of survival data, *Journal of the American Statistical Association* **92**: 1195–1212.
- Sinha, D. and Dey, D. K. (1998). Survival analysis using semiparametric Bayesian methods, in D. Dey, P. Müller and D. Sinha (eds), *Practical Nonparametric and Semiparametric Bayesian Statistics*, New York: Springer-Verlag, pp. 195–211.
- Smith, M. and Kohn, R. (1996). Nonparametric regression using Bayesian variable selection, *Journal of Econometrics* **75**: 317–343.
- Smith, M. and Kohn, R. (1997). A Bayesian approach to nonparametric bivariate regression, *Journal of the American Statistical Association* **92**: 1522–1535.
- Smith, M. and Kohn, R. (1998). Nonparametric Estimation of Irregular Functions with Independent or Autocorrelated Errors, in D. Dey, P. Müller and D. Sinha (eds), *Practical Nonparametric and Semiparametric Bayesian Statistics*, New York: Springer-Verlag, pp. 157–179.
- Stern, H. S. (1996). Neural networks in applied statistics (Disc: p215-220), *Technometrics* **38**: 205–214.

- Susarla, V. and Van Ryzin, J. (1976). Nonparametric Bayesian estimation of survival curves from incomplete observations, *Journal of the American Statistical Association* **71**: 897–902.
- Vannucci, M. and Corradi, F. (1999). Covariance structure of wavelet coefficients: Theory and models in a Bayesian perspective, *Journal of the Royal Statistical Society, Series B, Methodological* **61**: 971–986.
- Vidakovic, B. (1998). Nonlinear wavelet shrinkage with Bayes rules and Bayes factors, *Journal of the American Statistical Association* **93**: 173–179.
- Walker, S. (1995). Generating random variates from D-distributions via substitution sampling, *Statistical Computing* **5**: 311–315.
- Walker, S. and Damien, P. (1998). Sampling methods for Bayesian nonparametric inference involving stochastic processes, in D. Dey, P. Müller and D. Sinha (eds), *Practical Nonparametric and Semiparametric Bayesian Statistics*, New York: Springer-Verlag, pp. 243–254.
- Walker, S., Damien, P., Laud, P. and Smith, A. (1999). Bayesian nonparametric inference for distributions and related functions (with discussion), *Journal of the Royal Statistical Society, Series B* **61**: 485–527.
- Walker, S. G. and Damien, P. (2000). Representation of Lévy processes without Gaussian components, *Biometrika* **87**: 447–483.
- Walker, S. G. and Mallick, B. K. (1997). Hierarchical generalized linear models and frailty models with Bayesian nonparametric mixing, *Journal of the Royal Statistical Society, Series B* **59**: 845–860.
- Walker, S. and Mallick, B. (1999). A bayesian semiparametric accelerated failure time model, *Biometrics* **55**: 477–483.
- Wang, Y. and Taylor, J. M. (2001). Jointly modeling longitudinal and event time data with application to acquired immunodeficiency syndrome, *Journal of the American Statistical Association* **96**: 895–903.
- West, M., Müller, P. and Escobar, M. (1994). Hierarchical priors and mixture models with applications in regression and density estimation, in A. Smith and P. Freeman (eds), *Aspects of Uncertainty: A Tribute to D. V. Lindley*, New York: Wiley, pp. 363–386.

- Wild, C. J. and Kalbfleisch, J. D. (1981). A note on a paper by Ferguson and Phadia, *The Annals of Statistics* **9**: 1061–1065.
- Wolpert, R. L. and Ickstadt, K. (1998a). Poisson/gamma random field models for spatial statistics, *Biometrika* **85**: 251–267.
- Wolpert, R. L. and Ickstadt, K. (1998b). Simulation of Lévy random fields, *in* D. Dey, P. Müller and D. Sinha (eds), *Practical Nonparametric and Semiparametric Bayesian Statistics*, New York: Springer-Verlag, pp. 227–242.
- Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g -prior distributions, *in* P. K. Goel and A. Zellner (eds), *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, Amsterdam: North-Holland Publishing Company, pp. 233–243.