

Modelling and Computation Using NCoRM Mixtures for Density Regression

Jim Griffin* and Fabrizio Leisen†

Abstract. Normalized compound random measures are flexible nonparametric priors for related distributions. We consider building general nonparametric regression models using normalized compound random measure mixture models. Posterior inference is made using a novel pseudo-marginal Metropolis-Hastings sampler for normalized compound random measure mixture models. The algorithm makes use of a new general approach to the unbiased estimation of Laplace functionals of compound random measures (which includes completely random measures as a special case). The approach is illustrated on problems of density regression.

Keywords: dependent random measures, mixture models, multivariate Lévy measures, pseudo-marginal samplers, Poisson estimator.

1 Introduction

The problem of Bayesian nonparametric inference for distributions at different regressor values has been an extremely active area of research. Many approaches use dependent nonparametric mixture models and build on the idea of dependent Dirichlet process mixture models (MacEachern, 1999), which generalized the commonly-used Dirichlet process mixture model. A generic dependent nonparametric mixture model assumes that a sample y_1, \dots, y_n observed at regressor values x_1, \dots, x_n (where $x_i \in \mathbb{X}$ for some measurable space \mathbb{X}) is modelled as

$$y_i|x_i \sim q(y_i|\theta_{c_i}(x_i)), \quad p(c_i = k) = w_k(x_i), \quad k = 1, \dots, \infty, \quad (1)$$

where $q(y|\theta)$ is a distribution for y (where $y \in \mathbb{Y}$ for some measurable space \mathbb{Y}) with parameter θ , $w_k(x) \geq 0$ for all k and $x \in \mathbb{X}$, $\sum_{k=1}^{\infty} w_k(x) = 1$ almost surely for all $x \in \mathbb{X}$ and $\theta_1(x), \theta_2(x), \theta_3(x), \dots$ are independent realisations of a stochastic process. We refer to $\theta_1(x), \theta_2(x), \theta_3(x), \dots$ as the locations of the mixture components. The model simplifies to a nonparametric mixture model if the sample is observed at a single regressor value.

Many approaches to constructing specific models in the form of (1) generalize the stick-breaking construction of the Dirichlet process (Sethuraman, 1994) and these were reviewed by Dunson (2010). Alternatively, models can be constructed by normalising dependent random measures. This generalizes the approach introduced by Regazzini et al. (2003) to an arbitrary dimension. These constructions have several advantages.

*SMSAS, University of Kent, Canterbury, CT2 7FS, U. K., jeg28@kent.ac.uk

†SMSAS, University of Kent, Canterbury, CT2 7FS, U. K., F.Leisen@kent.ac.uk

Firstly, the weights $w_1(x), w_2(x), \dots$ are not ordered, as is the case with many stick-breaking constructions. Secondly, dependence is defined at the level of the weights $w_k(x)$ rather than, as is typical in stick-breaking constructions, through a non-linear transformation of the weights. Foti and Williamson (2012) defined a wide-class of such process using normalized kernel-weighted random measures, which generalize the approach to time-dependent random measures in Griffin (2011). Griffin et al. (2013) developed an approach to modelling a finite set of dependent random measures using superpositions of completely random measure (see also Lijoi and Nipoti, 2014; Lijoi et al., 2014a,b; Chen et al., 2013). Alternatively, dependence can be modelled through a Lévy copula (Leisen and Lijoi, 2011; Leisen et al., 2013; Zhu and Leisen, 2015). Compound random measures (CoRM) (Griffin and Leisen, 2017) are a unifying framework for many dependent random measures including many of the superposition and Lévy copula approaches. They have been applied to modelling graphs for overlapping communities by Todeschini and Caron (2016). Griffin and Leisen (2017) described posterior sampling methods for a particular class of normalized compound random measure mixtures which exploits a representation of the Laplace transform of a CoRM through a univariate integral of a moment generating function. Ranganath and Blei (2017) independently developed a normalized CoRM model where the weights depend on a Gaussian process and described a variational Bayesian algorithm for inference.

In this paper, we will consider extending the class of compound random measures (CoRM) from finite collections of distributions to infinite collections of distributions. This allows us to define CoRM models where the weights follow a time series model, the weights follow a regression model or the weights are defined through a hierarchical model. The computational algorithms in Griffin and Leisen (2017) cannot be used in this wider class of models since moment generating functions are not available in closed form. Therefore, we develop a new Markov chain Monte Carlo (MCMC) algorithm for nonparametric mixture models with Completely Random Measure (CRM) priors which uses a novel pseudo-marginal MCMC method (Andrieu and Roberts, 2009).

The paper is organized in the following way. Section 2 discusses defining NCoRM mixture models for distributions indexed by continuous covariates. Section 3 introduces a novel computational algorithm for NCoRM mixtures which can be widely applied. Section 4 illustrates how NCoRM can be used in density regression problems and how the computational algorithm performs. Section 5 concludes. Matlab code for the examples in this paper is available from <http://www.kent.ac.uk/smsas/personal/jeg28/index.htm>.

2 Modelling with normalized compound random measure mixtures

For simplicity, we will consider mixture models of the form in (1) with $\theta_k(x) = \theta_k$ for all $x \in \mathbb{X}$, leading to a mixture model with weights which vary over \mathbb{X} (many of the ideas in this paper could be extended to the model where $\theta_k(x)$ follows a stochastic process, such as a Gaussian process, over \mathbb{X}). The model is

$$y_i|x_i \sim q(y_i|\theta_{c_i}), \quad p(c_i = k) = w_k(x_i), \quad k = 1, \dots, \infty. \quad (2)$$

We consider the weights

$$w_k(x) = \frac{m_k(x)J_k}{\sum_{l=1}^{\infty} m_l(x)J_l}, \tag{3}$$

where $m_k(x)$ is a random function on \mathbb{X} for which $m_k(x) \geq 0$ for all $x \in \mathbb{X}$ and the function m_k is independent of m_l , $J_1, J_2, J_3 \dots$, are the jumps of the process with directing Lévy process ν^* and $\theta_1, \theta_2, \theta_3, \dots$ are i.i.d. We will refer to $m_k(x)$ as a *score* or *score function* and to ν^* as the *directing Lévy process*. The model reduces to the NCoRM models considered by Griffin and Leisen (2017) if \mathbb{X} is a finite set. In particular, they introduced a class of dependent random probability measures $\tilde{p}_1, \tilde{p}_2, \dots, \tilde{p}_d$ which can be represented as

$$\tilde{p}_i = \sum_{k=1}^{\infty} w_{ik} \delta_{\theta_k}$$

with

$$w_{ik} = \frac{m_{ik}J_k}{\sum_{k=1}^{\infty} m_{ik}J_k},$$

where (m_{1k}, \dots, m_{dk}) are i.i.d. draws from a d -variate *score* distribution h (for $k = 1, 2, \dots$), $J_1, J_2, J_3 \dots$, are the jumps of the process with *directing Lévy process* ν^* and $\theta_k \stackrel{i.i.d.}{\sim} \tilde{\alpha}$, with $\tilde{\alpha} = \alpha/\alpha(\mathbb{Y})$ where α is a positive finite measure. Under suitable conditions, the vector $(\tilde{p}_1, \tilde{p}_2, \dots, \tilde{p}_d)$ can be seen as a vector of normalized completely random measures, i.e.

$$\tilde{p}_j = \frac{\tilde{\mu}_j}{\tilde{\mu}_j(\mathbb{Y})} \quad j = 1, \dots, d,$$

where $(\tilde{\mu}_1, \tilde{\mu}_2, \dots, \tilde{\mu}_d)$ is a *Compound Random Measure* (CoRM). The model introduced in (2) assumes that \mathbb{X} could potentially be a countable set. In this case, we assume that, for every finite subset $S = \{s_1, \dots, s_l\}$ of covariates, the $w_k(s_1), \dots, w_k(s_l)$ displayed in (3) are the weights of a l -dimensional NCoRM process.

The specification of the weights displayed in (3) has several attractive features which motivate our choice. Firstly, the nonparametric approach allows the definition of a flexible model for density regression. Secondly, the dependence between $w_k(x)$ and $w_k(x')$ for $x, x' \in \mathbb{X}$ can be controlled by the choice of the distribution for the random function m_k . Many methods have been developed to model such random functions and can be used to define a suitable dependent nonparametric mixture model. Thirdly, the weights are not *a priori* stochastically ordered (as with many stick-breaking processes). Lastly, the structure of the model allows simpler computational methods to be developed than many other dependent extensions of normalized random measures.

We will concentrate on models where $m_l(x) = \exp\{r_l(x)\}$ and $r_l(x)$ is a random function on \mathbb{X} taking value on \mathbb{R} . Griffin and Leisen (2017) considered using the variance of the ratio of the same jump at values $x, x' \in \mathbb{X}$ as a simple measure of the strength of dependence between the (unnormalized) random measure at values x and x' . In this case, the ratio is $\zeta(x, x') = m_l(x)/m_l(x') = \exp\{r_l(x) - r_l(x')\}$ and the distribution of $\zeta(x, x')$ will often be easy to work with. For example, $\zeta(x, x')$ will be log normally distributed if $r_l(x)$ and $r_l(x')$ have a bivariate normal marginal distribution.

In this paper, we will consider models in which $r_l(x)$ is a stochastic process for which $\mathbb{E}[r_l(x)] = 0$ for all $x \in \mathbb{X}$. This gives CoRM models a high degree of flexibility. To illustrate the use of NCoRM mixtures in a regression context, we will consider a choice of $r_l(x)$ which is suitable for continuous regressors and a choice of $r_l(x)$ which is suitable for categorical regressors:

- **Continuous regressors:** In this case, we define $r_1(x), r_2(x), r_3(x), \dots$ to be independent Gaussian processes with covariance function $\sigma_0^2 \kappa(\cdot, \cdot)$ where $\kappa(\cdot, \cdot)$ is a correlation function. This implies that $\log \zeta(\cdot, \cdot)$ follows a normal distribution with mean zero and variance $2\sigma_0^2$.
- **Categorical regressors:** Suppose that we have two categorical regressors, $x_i = (x_{i,1}, x_{i,2})$ then we could assume a different parameter for each combination of levels so that $r_l(x_i) = \gamma_{x_{i,1}, x_{i,2}}^{(l)}$. Alternatively, we could use the specification $r_l(x_i) = \alpha_{x_{i,1}}^{(l)} + \beta_{x_{i,2}}^{(l)} + \gamma_{x_{i,1}, x_{i,2}}^{(l)}$ where

$$\alpha_j^{(k)} \sim N(0, \sigma_1^2), \quad \beta_j^{(k)} \sim N(0, \sigma_2^2), \quad \gamma_{i,j}^{(k)} \sim N(0, \sigma_{1,2}^2).$$

Then, the $\alpha^{(l)}$ and $\beta^{(l)}$ parameters act as main effects and $\gamma^{(l)}$ as interactions which can be interpreted in a similar way to a logistic regression model. For example, $\log \zeta(x, x')$ is normally distributed with mean 0 and variance $2(\sigma_1^2 + \sigma_2^2 + \sigma_{1,2}^2)$ if both levels of x are different to the levels of x' . Whereas, $\log \zeta(x, x')$ is normally distributed with mean 0 and variance $2(\sigma_2^2 + \sigma_{1,2}^2)$ if only the second level of x and x' are different. This shows how the dependence of jump sizes depends on the levels of the regressors.

Posterior inference is impossible using existing methods and the following section describes a general purpose algorithm for NCoRM mixture models.

3 Computational methods

Posterior inference for nonparametric mixture models is challenging due to the infinite-dimensional random probability measure in the model. To address this problem, two main MCMC approaches to defining a finite-dimensional target have been developed. Firstly, marginal methods integrate the random probability measure from the posterior. Secondly, conditional methods truncate the random probability measure. These methods can be further divided into exact methods which use a random truncation to sample exactly from the posterior and methods which fix the level of truncation leading to some truncation error. Griffin and Leisen (2017) suggest a marginal method and an exact conditional method (a slice sampler). The availability of an analytical expression for the moment generating function for the score distribution is key to their sampling methods but this is impossible to evaluate in closed form for the more general NCoRM models described in this paper. We propose a hybrid conditional-marginal sampler using a pseudo-marginal Metropolis-Hastings algorithm (Andrieu and Roberts, 2009).

We assume that we observe data $(x_1, y_1), \dots, (x_n, y_n)$ and wish to fit the model in (2). Without loss of generality, we also assume that the values x_1, x_2, \dots, x_n are distinct and write $m_{k,i} = m_k(x_i)$ and $m_k = (m_{k,1}, \dots, m_{k,n})$. Following Griffin and Leisen (2017), it is convenient to use an augmented form of the likelihood which introduces an allocation variable for each observation. Let n_k be the number of observations allocated to the k -th jump, we order the jumps so that J_1, \dots, J_K have points allocated to them (*i.e.* $n_k > 0$ for $1 \leq k \leq K$) and J_{K+1}, J_{K+2}, \dots have no points allocated to them (*i.e.* $n_k = 0$ for $k > K$). Marginalizing over jumps which have no points allocated and the location of all atoms and writing $M = \alpha(\mathbb{Y})$ and $\tilde{\alpha} = \alpha/\alpha(\mathbb{Y})$ gives

$$M^K \prod_{k=1}^K \left\{ J_k^{n_k} \left[\prod_{i=1; s_i=k}^n m_{k,i} \right] \exp \left\{ - \sum_{i=1}^n v_i J_k m_{k,i} \right\} h(m_k) \nu^*(J_k) \right\} L \prod_{k=1}^K q(y^{(k)}), \tag{4}$$

where

$$\begin{aligned} L &= \mathbb{E} \left[\exp \left\{ - \sum_{i=1}^n v_i \sum_{k=1}^{\infty} J_k m_{k,i} \right\} \right] \\ &= \exp \left\{ - \int_{(\mathbb{R}^+)^n} \int_0^{\infty} \left(1 - \exp \left\{ -z \sum_{i=1}^n v_i m_i \right\} \right) h(m^*) \nu^*(z) dz dm^* \right\}, \end{aligned} \tag{5}$$

$$q(y) = \int \prod q(y_i|\theta) \tilde{\alpha}(\theta) d\theta,$$

and $y^{(k)} = \{y|s_i = k, 1 \leq i \leq n\}$. The expression in (4) is the (marginalized) likelihood of the data. Griffin and Leisen (2017) use the analytical expression for L and integrals over J_1, \dots, J_K to define a marginal sampler. In general, these integrals are not analytically available to us. We replace L by an unbiased estimate \hat{L} (a possible unbiased estimator is discussed in the next Section) to define the following target

$$M^K \prod_{k=1}^K \left\{ J_k^{n_k} \left[\prod_{i=1; s_i=k}^n m_{k,i} \right] \exp \left\{ - \sum_{i=1}^n v_i J_k m_{k,i} \right\} h(m_k) \nu^*(J_k) \right\} \hat{L} \prod_{k=1}^K q(y^{(k)}).$$

Finally, we assume that h has parameters τ and ν^* has parameters ξ on which we want to make inference and define the target

$$\begin{aligned} &p(\tau)p(\xi)p(M)M^K \prod_{k=1}^K \left\{ J_k^{n_k} \left[\prod_{i=1; s_i=k}^n m_{k,i} \right] \exp \left\{ - \sum_{i=1}^n v_i J_k m_{k,i} \right\} h(m_k|\tau) \nu_{\xi}^*(J_k) \right\} \\ &\times \hat{L} \prod_{k=1}^K q(y^{(k)}). \end{aligned}$$

We propose a novel sampling strategy for the variable s in a nonparametric mixture model and a novel computational algorithm to deal with the Laplace transform component of the target above. This algorithm can be applied to posterior inference for a wide variety of Bayesian nonparametric processes beyond NCoRM processes.

Updating c

To update c_i , we write the full conditional distribution as proportional to

$$\prod_{k=1}^{K_i^-} (J_k m_{k,i} k(\{y_j | c_j = k\} \cup y_i))^{I(c_i=k)} \left(M m_{K_i^-+1,i} \gamma \left(m_{K_i^-+1} \right) k(y_i) \right)^{I(c_i=K_i^-+1)} \\ \times p \left(J_{K_i^-+1} \mid m_{K_i^-+1} \right) h \left(m_{K_i^-+1} \mid \tau \right),$$

where K_i^- is the number of distinct values in $c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n$, $(J_1, m_1), \dots, (J_K, m_K)$ are ordered so that $c_i = K$,

$$\gamma(m_k) = \int z \exp \left\{ -z \sum_{i=1}^n v_i m_{k,i} \right\} \nu_\xi^*(z) dz$$

and

$$p \left(J_{K_i^-+1} \mid m_{K_i^-+1} \right) = \frac{J_{K_i^-+1} \exp \left\{ -J_{K_i^-+1} \sum_{i=1}^n v_i m_{K_i^-+1,i} \right\} \nu_\xi^*(J_{K_i^-+1})}{\gamma(m_{K_i^-+1})}.$$

A new value of $J_{K_i^-+1}$ is sampled from this full conditional distribution leading to an algorithm which is similar to Algorithm 8 of Neal (2000). See James et al. (2009), Lijoi and Prünster (2010) and Favaro and Teh (2013) for extension to non-conjugate normalized random measure mixtures.

If the i -th observation was allocated to a singleton cluster in the previous iteration, the full conditional distribution of c_i is

$$p(c_i = j) \propto \begin{cases} J_j m_{j,i} \frac{q(\{y_k | c_k = j\} \cup y_i)}{q(\{y_k | c_k = j\})} & j = 1, \dots, K_i^-, \\ M m_{j,i} \gamma(m_j) q(\{y_i\}) & j = K_i^- + 1. \end{cases}$$

If the i -th observation was not allocated to a singleton cluster in the previous iteration, we propose $m_{K_i^-+1} \sim h(m_{K_i^-+1} | \tau)$ and $J_{K_i^-+1} \sim p(J_{K_i^-+1} | m_{K_i^-+1})$, then

$$p(c_i = j) \propto \begin{cases} J_j m_{j,i} \frac{q(\{y_k | c_k = j\} \cup y_i)}{q(\{y_k | c_k = j\})} & j = 1, \dots, K_i^-, \\ M m_{K_i^-+1,i} \gamma \left(m_{K_i^-+1} \right) q(\{y_i\}) & j = K_i^- + 1. \end{cases}$$

In Appendix B (Griffin and Leisen, 2017) we provide the details of the full conditional distributions for the variables J_k, m_k, v_i, ξ, M and τ . The next Section will introduce the novel pseudo-marginal Metropolis-Hastings algorithm used to address the intractability of the Laplace transform part of the target distribution.

3.1 Unbiased estimation of the Laplace functional

Andrieu and Roberts (2009) introduced a sampling scheme, called *pseudo-marginal* Metropolis-Hastings, which allows sampling from distributions which cannot be evalu-

ated pointwise. The main idea of the method is to replace the target distribution with a nonnegative unbiased estimator.

In our framework, we are often interested in evaluating objects such as the expectation in (5),

$$L = \exp \left\{ - \int_{(\mathbb{R}^+)^p} \int_0^\infty \left(1 - \exp \left\{ -z \sum_{i=1}^n v_i m(x_i) \right\} \right) p(m) \nu^*(z) dz dm \right\}.$$

We will use the Poisson estimator (Papaspiliopoulos, 2009) which has been successfully used in MCMC approaches for diffusions (see *e.g.* Fearnhead et al., 2010). Consider, the equation

$$L_\phi = \exp \left\{ - \int_D \phi(x) dx \right\} \leq 1 \tag{6}$$

for $D \subset \mathbb{R}^p$ where $\phi(x) \geq 0$ for all $x \in D$ and $\int_D \phi(x) dx < \infty$. The Poisson estimator of (6) is introduced in the following Theorem where some properties are described. The proof of the Theorem can be found in the Appendix. We denote the Poisson distribution with parameter λ by $Pn(\lambda)$.

Theorem 1. *Consider the following estimator,*

$$\hat{L}_\phi = \prod_{i=1}^K \left(1 - \frac{\phi(x_i)}{a C \kappa(x_i)} \right), \tag{7}$$

where κ is a p.d.f. on D , $C > \frac{\phi(x)}{\kappa(x)}$ for $x \in D$, $a > 1$, $K \sim Pn(a C)$ and $x_i \stackrel{i.i.d.}{\sim} \kappa$. Then,

$$\mathbb{E}[\hat{L}_\phi] = \exp \left\{ - \int_D \phi(x) dx \right\}$$

and

$$\begin{aligned} V[\hat{L}_\phi] &= L_\phi^2 \left(\exp \left\{ \frac{1}{a C} \int_D \frac{\phi(x)^2}{\kappa(x)} dx \right\} - 1 \right) \\ &\leq L_\phi^2 \left(\exp \left\{ \frac{1}{a} \int_D \phi(x) dx \right\} - 1 \right) < \infty. \end{aligned}$$

The estimator has the useful property that it is always positive. This contrasts with other approaches which define unbiased estimators of infinite sums using random truncation where it is difficult to ensure that estimates are always positive (see *e.g.* Rhee and Glynn, 2015; Lyne et al., 2015).

Returning to the expression in (5) and, again, assuming that x_1, x_2, \dots, x_n are distinct, this can be re-expressed as

$$L = \exp \left\{ - \int_{(\mathbb{R}^+)^n} \int_0^\infty \left(1 - \exp \left\{ -z \sum_{i=1}^n v_i m_i^* \right\} \right) h(m^*) \nu^*(z) dz dm^* \right\}$$

$$\begin{aligned}
 &= \exp \left\{ - \int_{(\mathbb{R}^+)^n} \int_0^\infty \sum_{k=1}^n v_k m_k^* \int_0^z \exp \left\{ -t \sum_{i=1}^n v_i m_i^* \right\} dt h(m^*) \nu^*(z) dz dm^* \right\} \\
 &= \exp \left\{ - \int_{(\mathbb{R}^+)^n} \int_0^\infty \sum_{k=1}^n v_k m_k^* h(m^*) \exp \left\{ -t \sum_{i=1}^n v_i m_i^* \right\} \int_t^\infty \nu^*(z) dz dt dm^* \right\} \\
 &= \exp \left\{ - \int_{(\mathbb{R}^+)^n} \int_0^\infty \sum_{k=1}^n v_k m_k^* h(m^*) \exp \left\{ -t \sum_{i=1}^n v_i m_i^* \right\} T_{\nu^*}(t) dt dm^* \right\} \\
 &= \prod_{k=1}^n L_k, \tag{8}
 \end{aligned}$$

where $T_{\nu^*}(t) = \int_t^\infty \nu^*(z) dz$ is the tail mass function for the Lévy process with Lévy intensity ν^* and

$$L_k = \exp \left\{ - \int_{(\mathbb{R}^+)^n} \int_0^\infty v_k m_k^* h(m^*) \exp \left\{ -t \sum_{i=1}^n v_i m_i^* \right\} T_{\nu^*}(t) dt dm^* \right\}.$$

The expression for L_k has the form of (6) with $x = (z, m_k^*)$, $D = (0, \infty) \times \mathbb{R}^{+n}$ and

$$\phi(z, m_k^*) = v_k m_k^* h(m_k^*) \exp \left\{ -z \sum_{i=1}^n v_i m_i^* \right\} T_{\nu^*}(z).$$

Clearly $\int_D \phi(x) dx < \infty$. To use the Poisson estimator, a suitable density is

$$\kappa(z, m_k^*) = \kappa_{\nu^*}(z) \frac{m_k^* h(m_k^*)}{\mathbb{E}[m_k^*]}, \tag{9}$$

where $\kappa_{\nu^*}(z) > T_{\nu^*}(z)$ for all $z \in \mathbb{R}^+$. We use $C = v_k \mathbb{E}[m_k^*] B$ where $\frac{T_{\nu^*}(z)}{\kappa_{\nu^*}(z)} < B$ for all z . Suitable forms of κ_{ν^*} for some popular nonparametric processes are given in Section 3.1.

In computation for more usual normalized random measures (Griffin and Walker, 2011; Favaro and Teh, 2013), we are interested in

$$\mathbb{E}[\exp\{-vJ\}] = \exp \left\{ - \int_0^\infty (1 - \exp\{-vz\}) \nu(z) dz \right\} \leq 1, \tag{10}$$

where $\nu(z)$ is a Lévy process and the expectation is taken over all jumps on \mathbb{R}^+ . This expectation can, similarly, be re-expressed as

$$= \exp \left\{ - \int_0^\infty (1 - \exp\{-vz\}) \nu(z) dz \right\} = \exp \left\{ -v \int_0^\infty T_\nu(t) \exp\{-vt\} dt \right\},$$

which is (6) with $D = \mathbb{R}^+$ and $\phi(x) = vT_\nu(x) \exp\{-vx\}$. The estimator in (7) provides an unbiased estimator of (8) and (10) which can be used in the sampler described in this section.

Controlling the variability of \hat{L}_ϕ

Pseudo-marginal Metropolis-Hastings algorithms converge to the correct distribution but the asymptotic variance of an average calculated using the algorithm depends on the variance of the unbiased estimator. For example, suppose that the unbiased estimator is an importance sampler. Andrieu and Vihola (2016) show that the asymptotic variance of the pseudo-marginal sampler decreases as the number of samples in the importance sampler increases (leading to an importance sampler with a lower asymptotic variance). Although we do not use an importance sampler, the Poisson estimator is closely related and it is intuitively reasonable that the asymptotic variance of averages calculated using the pseudo-marginal Metropolis-Hastings sampler will decrease as the variance of the Poisson estimator in (7) decreases.

The variability of the Poisson estimator is controlled by a with larger values of a leading to a smaller variance. However, larger values of a will also lead to longer computational times since the mean number of terms in \hat{L}_ϕ is aC . In this section, we will assume that the expected number of evaluations of the ratio $\phi(x)/\kappa(x)$ is d . Therefore, $d = aC$ for the estimator in (7). An alternative method for controlling the variability involves defining the estimator $\hat{L}_\phi^{AVE} = \frac{1}{N} \sum_{i=1}^N \hat{L}_\phi^{(i)}$ where $\hat{L}_\phi^{(1)}, \dots, \hat{L}_\phi^{(N)}$ are independent realisations of \hat{L}_ϕ . The estimator has variance

$$V[\hat{L}_\phi^{AVE}] = \frac{L_\phi^2}{N} \left(\exp \left\{ \frac{1}{aC} \int_D \frac{\phi(x)^2}{\kappa(x)} dx \right\} - 1 \right) = \frac{L_\phi^2}{N} \left(\exp \left\{ \frac{N}{d} \int_D \frac{\phi(x)^2}{\kappa(x)} dx \right\} - 1 \right),$$

since $d = aCN$. It is straightforward to show that

$$\frac{L_\phi^2}{N+1} \left(\exp \left\{ \frac{N+1}{d} \int_D \frac{\phi(x)^2}{\kappa(x)} dx \right\} - 1 \right) > \frac{L_\phi^2}{N} \left(\exp \left\{ \frac{N}{d} \int_D \frac{\phi(x)^2}{\kappa(x)} dx \right\} - 1 \right)$$

and so the variance of the estimator grows with N for fixed d . This suggests that we should use \hat{L}_ϕ^{AVE} with $N = 1$ which is the Poisson estimator in (7). In this case, the choice of $\kappa(x)$ which minimizes the variance for fixed d is $\kappa(x) = \frac{\phi(x)}{-\log L_\phi}$ which provides a criterion for choosing $\kappa(x)$.

As we have already mentioned the asymptotic variance of averages calculated using the pseudo-marginal algorithm will typically decrease as a increases but the computational time will increase. Therefore, there is an optimal value of a which is able to provide the lowest asymptotic variance for a fixed computational budget (number of evaluations of $\phi(x)/\kappa(x)$). Doucet et al. (2015) established an upper bound for the asymptotic variance under certain assumptions which allows this optimal value to be derived. They demonstrated that this value can be close to optimal when the assumptions are violated. In our context, their main assumption is

$$\hat{L}_\phi = \epsilon L_\phi,$$

where $\log \epsilon \sim N(-\sigma^2/2, \sigma^2)$. They refer to σ^2 as the noise variance and it is straightforward to show that

$$\sigma^2 = \frac{1}{aC} \int_D \frac{\phi(x)^2}{\kappa(x)} dx.$$

Doucet et al. (2015) showed that the optimal value of the noise variance (in terms of asymptotic variance), σ_{opt}^2 , depends on the properties of the chain but provide guidelines on how this can be approximated. Following the derivation of Doucet et al. (2015), the optimal value of a , for fixed κ , is

$$a_{opt} = \frac{1}{\sigma_{opt}^2 C} \int_D \frac{\phi(x)^2}{\kappa(x)} dx.$$

In practice, we have found that the value $a = 8$ works well for the processes considered in this paper.

Examples

Brix (1999) provided a bound for the tail-mass integral of the generalized gamma process which is extended to the stable-Beta process by Arbel and Prünster (2016). However, both bounds are not tight and we suggest tighter bounds for both processes. Indeed, the estimator introduced in Theorem 1 with the proposal in (9) requires draws from a Poisson distribution whose mean is proportional to $v_k \mathbb{E}[m_k^*] D$ where $\frac{T_{\nu^*}(z)}{\kappa_{\nu^*}(z)} < B$ for all z . Therefore, better choices of $\kappa_{\nu^*}(x)$ can improve the computational efficiency of the method by requiring a smaller value of B .

Generalized gamma process

The generalized gamma process has Lévy density $\nu^*(y) = \frac{1}{\Gamma(1-\sigma)} y^{-1-\sigma} \exp\{-\lambda y\}$ and the tail-mass functions $T_{\nu^*}(t)$ is an incomplete gamma function. It is straightforward to show that $T_{\nu^*}(t) < \tilde{\kappa}_{\nu^*}(t)$ where

$$\tilde{\kappa}_{\nu^*}(t) = \begin{cases} \frac{1}{\sigma\Gamma(1-\sigma)}(t^{-\sigma} - 1) & t < b, \\ \frac{1}{\sigma\Gamma(1-\sigma)}(b^{-\sigma} - 1) \exp\{-\lambda(t - b)\} & t \geq b, \end{cases}$$

and $b = 0.65$. A suitable choice of D is $D = \frac{1}{\sigma\Gamma(1-\sigma)}[\frac{b^{1-\sigma}}{1-\sigma} - b + b^{-\sigma} - 1]$ and the bounding p.d.f. is $\kappa_{\nu^*}(t) = \frac{1}{D} \tilde{\kappa}_{\nu^*}(t)$. Consider $\kappa_{\nu^*}(t)$ truncated to $t < b$, taking the transformation $y = \frac{1}{\sigma}(t^{-\sigma} - 1)$ leads to the density $\kappa(y) \propto y(\sigma y + 1)^{-1/\sigma-1}$ truncated to $y > \frac{1}{\sigma}(b^{-\sigma} - 1)$. This can be expressed as a mixture of gamma distributions where $y|\Xi \sim \text{Ga}(2, \Xi)$, $\Xi \sim \text{Ga}(1/\sigma - 1, 1/\sigma)$.

As $\sigma \rightarrow 0$ for $\lambda = 1$, the generalized gamma process converges to the gamma process which has Lévy density $\nu(z) = z^{-1} \exp\{-z\}$ and $T_\nu(t) = \int_t^\infty \nu(y) dy = E_1(t)$ where $E_1(t)$ is the exponential-integral function. Both the bounding p.d.f. and simulation scheme for $t < b$ also converge. It is straightforward to show that the limit is

$$\tilde{\kappa}_{\nu^*}(t) = \begin{cases} -\log t & t < b, \\ -(\log b) \exp\{-(t - b)\} & t \geq b, \end{cases}$$

with $D = b - b \log b - \log b$. The appropriate transformation for $t < b$ is $y = -\log(t)$ which has p.d.f. $z \exp\{-z\}$, *i.e.* $z \sim \text{Ga}(2, 1)$ truncated to $y > -\log b$.

If $\lambda = 0$, the generalized gamma process is stable process. However, the tail mass function is infinite for a stable process and this simulation scheme is not possible.

Stable-Beta process

The stable-Beta process has Lévy density $\nu^*(z) = \frac{\Gamma(\phi)}{\Gamma(\sigma+\phi)\Gamma(1-\sigma)}z^{-\sigma-1}(1-\lambda z)^{\sigma+\phi-1}$ for $0 < z < 1/\lambda$. It is straightforward to show that $T_{\nu^*}(t) < \tilde{\kappa}_{\nu^*}(t)$ where

$$\tilde{\kappa}_{\nu^*}(t) = \begin{cases} \frac{\Gamma(\phi)}{\sigma\Gamma(\sigma+\phi)\Gamma(1-\sigma)}(t^{-\sigma} - 1) & t < b, \\ \frac{\Gamma(\phi)}{\sigma\Gamma(\sigma+\phi)\Gamma(1-\sigma)}(b^{-\sigma} - 1) \frac{(1-\lambda t)^{\gamma+\sigma}}{(1-\lambda b)^{\gamma+\sigma}} & t \geq b, \end{cases}$$

and $b = 0.65$. A suitable choice of D is $D = \frac{\Gamma(\phi)}{\Gamma(\sigma+\phi)\Gamma(1-\sigma)}[\frac{b^{1-\gamma}}{1-\gamma} - b + \frac{b^{-\sigma}-1}{\lambda(\gamma+\sigma+1)}(1-\lambda b)]$ and the bounding p.d.f. is $\kappa_{\nu^*}(t) = \frac{1}{D}\tilde{\kappa}_{\nu^*}(t)$. As $\sigma \rightarrow 0$ for $\lambda = 1$, the stable-Beta process converges to the Beta process which has Lévy density $\nu(y) = \frac{1}{\Gamma(\gamma)}y^{-1}(1-y)^{\gamma-1}$ for $0 < y < 1$. In this case, the limit of $\tilde{\kappa}_{\nu^*}(t)$ is

$$\tilde{\kappa}_{\nu^*}(t) = \begin{cases} -\log t & t < b, \\ \frac{-\log b}{(1-b)^\gamma}(1-t)^\gamma & t \geq b, \end{cases}$$

and the limit of D is $D = b - b \log b - \frac{\log b}{\gamma+1}(1-b)$.

4 Illustrations

4.1 Example 1: Discrete regressors

The algorithms developed in this paper are illustrated using an analysis of hematological data arising from a dose-escalation study which has previously been analysed by Müller and Rosner (1997). The data are white blood cell counts over time for a sample of 52 patients receiving different levels of two treatments: cyclophosphamide (CTX) and a second drug (GM-CSF). The data for each patient is summarized as the maximum likelihood estimates from a non-linear regression model with seven parameters fitted to that patient’s time profile. The model assumes that the mean response at time t with parameters $\theta = (z_1, z_2, z_3, \tau_1, \tau_2, \beta_1)$ is given by

$$f(\theta, t) = \begin{cases} z_1 & t < \tau_1, \\ rz_1 + (1-r)g(\theta, \tau_2) & \tau_1 \leq t < \tau_2, \\ g(\theta, t) & t \geq \tau_2, \end{cases}$$

where $r = (\tau_2 - t)/(\tau_2 - \tau_1)$ and $g(\theta, t) = z_2 + z_3/[1 + \exp\{2.0 - \beta_1(t - \tau_2)\}]$. The model implies that the white blood cell count is constant (at level z_1) before τ_1 followed by a linear progression between τ_1 and τ_2 and a logistic recovery after τ_2 . The parameters z_2 and z_3 control the white blood cell count at the start and end of recovery. De Iorio et al. (2004) applied an ANOVA-DDP model to these data which assumes a mixture model with constant weights and an ANOVA model for the locations for each treatment. In contrast, we fitted a mixture model with weights that vary with the treatment combination but with locations that do not depend on the treatment level. Specifically, we

assume that y_i are the estimated parameters for the i -th patient and that $x_{i,1}$ is the level of CTX and $x_{i,2}$ is the level of GM-CSF. The model is

$$y_i | s_i \sim N(\mu_{s_i}, \Sigma), \quad \mu_k \sim N(\mu_0, \lambda^{-1}\Sigma),$$

$$\Sigma^{-1} \sim W(\nu, \Psi),$$

$$p(c_i = k) = \frac{J_k \exp \left\{ \alpha_{x_{i,1}}^{(k)} + \beta_{x_{i,2}}^{(k)} + \gamma_{x_{i,1}, x_{i,2}}^{(k)} \right\}}{\sum_{l=1}^{\infty} J_l \exp \left\{ \alpha_{x_{i,1}}^{(l)} + \beta_{x_{i,2}}^{(l)} + \gamma_{x_{i,1}, x_{i,2}}^{(l)} \right\}},$$

$$\alpha_j^{(k)} \sim N(0, \sigma_1^2), \quad \beta_j^{(k)} \sim N(0, \sigma_2^2), \quad \gamma_{i,j}^{(k)} \sim N(0, \sigma_{1,2}^2)$$

where $W(\nu, \Psi)$ is a Wishart distribution with ν degrees of freedom and mean $\nu\Psi$. The directing Lévy process is taken to be a gamma process. The model assumes a two-way ANOVA model with interaction for the logarithm of the weights. This does not place restriction on the combination of weights but does encourage similar weights for similar combinations of levels. The priors were $\sigma_1^2 \sim \text{Ga}(1, 2)$, $\sigma_2^2 \sim \text{Ga}(1, 2)$, $\sigma_{1,2}^2 \sim \text{Ga}(1, 2)$ and $M \sim \text{Ga}(1, 1)$. For the purposes of illustration, we set μ_0 equal to the sample mean of the data, $\Psi = \frac{1}{9(\nu-8)}\hat{\Sigma}$ where $\hat{\Sigma}$ is the covariance of the data which implies that the prior mean of Σ is $\frac{1}{9}\hat{\Sigma}$ and we choose $\lambda = 0.01$. The MCMC algorithms was run for a total of 35 000 iterations. The first 5 000 were used as a burn-in with the subsequent values thinned every fifth sample. This gave a sample of 6 000 values.




Figure 1: Posterior mean probability density of z_1 for three levels of GM-CSF and four levels of CTX shown as: solid line (CTX=1.5), dashed line (CTX=3.0), dot-dashed line (CTX=4.5) and dashed line (CTX=6.0).

The inference about the marginal probability of two parameters z_1 and z_2 are shown in Figures 1 and 2. The parameter z_1 is the initial white blood cell count. The distribution is bi-modal with the size of the smaller mode increasing with GM-CSF. This indicates that there are differences in the proportion of patients with lower white blood




Figure 2: Posterior mean probability density of z_2 for three levels of GM-CSF and four levels of CTX shown as: solid line (CTX=1.5), dashed line (CTX=3.0), dot-dashed line (CTX=4.5) and dashed line (CTX=6.0).




Figure 3: Trace plots of the parameters M , σ_1^2 , σ_2^2 and $\sigma_{1,2}^2$.

cell count across the different treatments. The parameter z_2 controls the level of white-blood cells when recovery begins and this is again bi-modal with the smaller mode decreasing with GM-CSF for z_2 .

Figure 3 shows trace plots for the total mass parameter M and the three parameters controlling the differences between jumps at each treatment level. These clearly show good performance of the sampler for this problem.

4.2 Example 2: Continuous Regressors

A regression model is used to define an infinite mixture model with regressor dependent weights. We observe pairs $(x_1, y_1), \dots, (x_n, y_n)$ where $x_i \in \mathbb{R}^p$ and $y_i \in \mathbb{R}$ and use the model

$$y_i \sim N(\theta_{s_i}, a\sigma^2), \quad \theta_k \sim N(\mu, (1-a)\sigma^2),$$

$$p(c_i = k) = \frac{J_k \exp\{r_k(x_i)\}}{\sum_{l=1}^{\infty} J_l \exp\{r_l(x_i)\}},$$

where $r_1(x), r_2(x), r_3(x), \dots$ are independent Gaussian processes. A generalized gamma directing Lévy process is used with $\lambda = 1$ and three values of σ : $\sigma = 0$ (a gamma

process), $\sigma = 0.1$ and $\sigma = 0.5$.

We apply the model to data from a simulated motorcycle accident used to test crash helmets (Silverman, 1985), which are available as the `mcycle` data frame in the R package MASS. The data are head accelerations (in g) measured at different times in milliseconds after impact. We assume that the Gaussian processes have covariance function $C(x, y) = \phi \exp\{-\frac{\|x-y\|}{L}\}$ where $\|\cdot\|$ is Euclidean distance and L is the lengthscale. The priors are $\alpha \sim U(0, 1)$, $p(\mu, \sigma^2) \propto \sigma^{-2}$, $L \sim \text{Ga}(1, 1)$, $M \sim \text{Ga}(1, 1)$, and $\phi^{-1} \sim \text{Ga}(1, 4)$. The prior for ϕ is chosen so that $\log m_k(x_i)$ typically takes values in $(-4, 4)$. The MCMC algorithms was a total of 33 000 iterations. The first 3 000 were used as a burn-in with the subsequent values thinned every third sample. This gave a sample of 10 000 values.




Figure 4: Motorcycle data: Data (dots) and posterior mean density of $y|x$ (darker colours show larger density) with $\sigma = 0$, $\sigma = 0.1$ and $\sigma = 0.5$.

Figure 4 shows the posterior mean of the conditional density of head acceleration given time from impact for the three values of σ with the data superimposed. In each case, the model was able to follow the data and capture the changing the heterogeneity in the variance. The inference seems robust to the choice of σ .




Figure 5: Motorcycle data: Trace plots of the parameters M , ϕ and L for $\sigma = 0$ (the gamma process).

Trace plots for the three parameters M , ϕ and L for the case $\sigma = 0$ (the gamma process) are shown in Figure 5. These clearly show that the parameters are mixing well across the MCMC chain.

4.3 Comparison of predictive performance

We ran a simulation exercise to understand how the NCoRM regression model developed in this paper compared to two commonly used dependent nonparametric priors: the single- p dependent Dirichlet process (De Iorio et al., 2004) and a probit stick-breaking

process mixture (Rodriguez and Dunson, 2011). The methods were compared by 10-fold cross-validation using simulated data sets of size 100 (leading to training data sets with 90 observations) according to the out-of-sample log-predictive scores (LPS), *i.e.*

$$\text{LPS} = -\frac{1}{100} \sum_{i=1}^{10} \sum_{j=1}^{10} \log p(y_{i,j}^{test} | y_{i,1}^{train}, \dots, y_{i,90}^{train}),$$

where $y_{i,1}^{test}, \dots, y_{i,10}^{test}$ is the i -th testing sample and $y_{i,1}^{train}, \dots, y_{i,90}^{train}$ is the i -th training sample. Three sets of data were simulated to cover different modelling situations. The first two data sets used regressor-dependent mixture models covering the simple case of a two component mixture and a more complicated scenario with four components. The third dataset used a non-linear regression model which depends on four parameters a , b , c and d . Different values of the parameters lead to different features of the data such as homoscedascity or heteroscedascity, jumps or different levels of smoothness. In all cases, there was a single regressor which was generated uniformly on $(0, 1)$. The detailed descriptions of the data sets are given below.

- Simulated Data Sets I

The responses were simulated as

$$y \sim \begin{cases} N(-1, \sigma^2) & \text{if } s = 1, \\ N(1, \sigma^2) & \text{if } s = 2, \end{cases}$$

where

$$p(s = 1) = \frac{\exp\{r \sin(2\pi x)\}}{1 + \exp\{r \sin(2\pi x)\}}, \quad p(s = 2) = \frac{1}{1 + \exp\{r \sin(2\pi x)\}}.$$

We consider two values of σ (0.1 and 0.5) which allow for different levels of separation between the two mixture components and two values of r (1 and 2) which control the rate at $p(s = 1)$ changes over the range of x .

- Simulated Data Sets II

The responses were simulated as

$$y \sim \begin{cases} N(-1, \sigma^2) & \text{if } s = 1, \\ N(1, \sigma^2) & \text{if } s = 2, \\ N(-2, \sigma^2) & \text{if } s = 3, \\ N(2, \sigma^2) & \text{if } s = 4, \end{cases}$$

where $p(s = 1) \propto \exp\{2 \sin(2\pi x)\}$, $p(s = 2) \propto 1/2 + 2/5(x - 1/2)$, $p(s = 3) \propto 1/2 - 2(x - 1/2)^2$ and $p(s = 4) \propto 1$. Again, we consider two values of σ (0.1 and 0.5) to give different levels of separation between the clusters.

- Simulated Data Sets III

The responses were simulated as

$$y = g_{a,b}(x) + I(c = 1)h(x) + 0.1 \epsilon k(x)^{I(d=1)},$$

where

$$g_{a,b}(x) = \begin{cases} 0 & x < a, \\ \sin(2\pi(x-a)/(b-a)) & a \leq x \leq b, \\ 0 & x > b, \end{cases}$$

$$h(x) = \begin{cases} -2 & x < 1/3, \\ -3/4 & 1/3 < x < 3/4, \\ 0 & x > 3/4, \end{cases}$$

$k(x) = 0.15(1 + 19|\sin(2\pi)|)$ and $\epsilon \sim N(0, 1)$. Different choices of a , b , c and d lead to responses which have a nonlinear mean and potentially heteroscedasticity. If $c = 0$ and $d = 0$, the means of responses have a sine wave scaled to the interval (a, b) (with zero mean outside the interval) with homoscedastic noise. Additionally, the mean of the responses jumps at $1/3$ and $3/4$ if $c = 1$ and the errors are heteroscedastic if $d = 1$.

σ	r	Probit SB	DDP	NCoRM
0.1	2	-0.33	-0.18	-0.36
0.1	1	0.30	0.51	0.26
0.5	2	1.35	1.41	1.26
0.5	1	1.53	1.54	1.50

Table 1: LPS of the probit stick-breaking, dependent Dirichlet process and NCoRM mixtures for Simulated Data Sets I.

The LPS for the three different nonparametric priors with Simulated Data Sets I are given in Table 1. The NCoRM mixture outperformed both the DDP and probit stick-breaking processes for the four combination of σ and r . The LPS's for Simulated

σ	Probit SB	DDP	NCoRM
0.1	-0.22	-0.15	-0.22
0.5	1.49	1.54	1.48

Table 2: LPS of the probit stick-breaking, dependent Dirichlet process and NCoRM mixtures for Simulated Data Sets II.

Data Sets II (Table 2) show that the DDP outperformed the other methods for $\sigma = 0.1$ and the NCoRM outperformed the other methods for $\sigma = 0.5$.

The LPS for Simulated Data Sets III are shown in Table 3. For these data sets, the methods were ranked in the same order with the DDP giving the best performance and the NCoRM mixture outperforming the probit stick-breaking mixtures. The difference between the DDP and NCoRM was largest for the model without jumps and with homoscedasticity ($c = 0$ and $d = 0$). This is not surprising since a Gaussian process with normal errors would provide good approximation of the sine curve and models which only allow dependence through the weights can only approximate the curve using piecewise constant fits. If there are jumps, the advantage of the DDP over the NCoRM was reduced. In all case, the NCoRM mixture substantially outperformed the probit stick-breaking mixture. This reflects the construction of the probit

a	b	c	d	Probit SB	DDP	NCoRM
0	1	0	0	0.80	-0.81	-0.31
0	1	0	1	0.66	-0.42	-0.14
0	1	1	0	0.99	-0.01	0.25
0	1	1	1	0.80	0.12	0.25
1/4	3/4	0	0	0.58	-0.81	-0.27
1/4	3/4	0	1	0.87	-0.55	-0.43
1/4	3/4	1	0	0.48	-0.24	-0.16
1/4	3/4	1	1	0.89	-0.06	-0.08

Table 3: LPS of the probit stick-breaking, dependent Dirichlet process and NCoRM mixtures for Simulated Data Sets III.

stick-breaking processes. In all stick-breaking processes, the weights are stochastically ordered and the probit stick-breaking process assumes that the atom with largest *a priori* expected weight does not depend on x . Although the data can change the order *a posteriori*, this ordering persists in data sets of the size considered in these simulated examples.

The results of these simulations suggest some guidelines which can be used in more general situations. The model displayed in (1) is very general but in this form is rarely used in real situations. By allowing the parameter θ and the weights w_k to depend on the regressor, the model becomes extremely flexible and prone to overfit the data. Therefore, the models considered are typically used as special cases of the model displayed in (1). We prefer the NCoRM mixture model to the DDP mixture model if we can identify subpopulations with different levels of response and which are associated with different regressor values. Clearly, this is the case in simulated datasets I and II but also simulated dataset III when jumps are introduced ($c = 1$). On the other hand, the DDP works well if there is a wide range of responses in each subpopulation.

5 Conclusions

Normalized compound random measures are a large class of dependent nonparametric processes. The jumps of the processes are expressed as the product of a jump from a Lévy process and a random variable. This allows the dependence of the nonparametric processes to be modelled through the dependence in the random variables. In this paper, we have developed Markov chain Monte Carlo methods to estimate nonparametric mixture models where the mixing measure is given a normalized compound random measure prior with a wide-range of dependences between the underlying random variables. The NCoRM approach could be generalized to allow the jump locations to depend on regressors and the MCMC method could be simply extended. The examples illustrate priors constructed using linear models and Gaussian processes. Other types of dependence could be included such as time series models, spatial models or hierarchical models. The MCMC methods are efficient and depend on approximating the tail mass integral of a Lévy process. Examples of appropriate approximations are given for the most popular classes of Lévy processes used in Bayesian nonparametrics.

A simulation study illustrates that the NCoRM can provide better out-of-sample predictive performance than probit stick-breaking process mixtures in a range of simulated data sets and can outperform DDP mixtures if the mean of the responses does not vary smoothly.

Supplementary Material

Appendix of “Modelling and computation using NCoRM mixtures for density regression” (DOI: [10.1214/17-BA1072SUPP](https://doi.org/10.1214/17-BA1072SUPP); .pdf).

References

- Andrieu and Vihola (2016). “Establishing some order amongst exact approximations of MCMCs.” *Annals of Applied Probability*, 26: 2661–2696. [MR3563190](#). doi: <https://doi.org/10.1214/15-AAP1158>. 905
- Andrieu, C. and Roberts, G. O. (2009). “The pseudo-marginal approach for efficient Monte Carlo computations.” *Annals of Statistics*, 37: 697–725. [MR2502648](#). doi: <https://doi.org/10.1214/07-AOS574>. 898, 900, 902
- Arbel, J. and Prünster, I. (2016). “A moment-matching Ferguson & Klass algorithm.” *Statistics and Computing*, 27: 3–17. 906
- Atchadé, Y. F. and Rosenthal, J. S. (2005). “On Adaptive Markov Chain Monte Carlo Algorithms.” *Bernoulli*, 11: 815–828.
- Brix, A. (1999). “Generalised gamma measures and shot-noise Cox processes.” *Advances in Applied Probability*, 31: 929–953. 906
- Chen, C., Rao, V. A., Buntine, W., and Teh, Y. W. (2013). “Dependent Normalized Random Measures.” In *Proceedings of the International Conference on Machine Learning*. 898
- De Iorio, M., Müller, P., Rosner, G. L., and MacEachern, S. N. (2004). “An ANOVA model for dependent random measures.” *Journal of the American Statistical Association*, 99: 205–215. 907, 910
- Doucet, A., Pitt, M., Deligiannidis, G., and Kohn, R. (2015). “Efficient Implementation of Markov chain Monte Carlo when Using an Unbiased Likelihood Estimator.” *Biometrika*, 102: 295–313. 905, 906
- Dunson, D. B. (2010). “Nonparametric Bayes applications to biostatistics.” In *Bayesian Nonparametrics*, 223–273. Cambridge University Press. [MR2730665](#). 897
- Favaro, S. and Teh, Y. W. (2013). “MCMC for Normalized Random Measure Mixture Models.” *Statistical Science*, 28: 335–359. 902, 904
- Fearnhead, P., Papaspiliopoulos, O., Roberts, G. O., and Stuart, A. (2010). “Random-weight particle filtering of continuous time processes.” *Journal of the Royal Statistical Society, Series B*, 74: 497–512. 903

- Foti, N. and Williamson, S. (2012). “Slice sampling normalized kernel-weighted completely random measure mixture models.” In *Advances in Neural Information Processing Systems 25*, 2240–2248. 898
- Griffin, J. E. (2011). “The Ornstein-Uhlenback Dirichlet process and other time-varying processes for Bayesian nonparametric inference.” *Journal of Statistical Planning and Inference*, 141: 3648–3664. 898
- Griffin, J. E., Kolossiatos, M., and Steel, M. F. J. (2013). “Comparing Distributions By Using Dependent Normalized Random-Measure Mixtures.” *Journal of the Royal Statistical Society, Series B*, 75: 499–529. MR3065477. doi: <https://doi.org/10.1111/rssb.12002>. 898
- Griffin, J. and Leisen, F. (2017). “Appendix of “Modelling and computation using NCoRM mixtures for density regression”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/17-BA1072SUPP>. 902
- Griffin, J. E. and Leisen, F. (2017). “Compound Random Measures and their use in Bayesian nonparametrics.” *Journal of the Royal Statistical Society, Series B*, 79: 525–545. 898, 899, 900, 901
- Griffin, J. E. and Walker, S. G. (2011). “Posterior simulation of normalized random measure mixtures.” *Journal of Computational and Graphical Statistics*, 20: 241–259. 904
- James, L. F., Lijoi, A., and Prünster, I. (2009). “Posterior Analysis for Normalized Random Measures with Independent Increments.” *Scandinavian Journal of Statistics*, 36: 76–97. 902
- Leisen, F. and Lijoi, A. (2011). “Vectors of Poisson-Dirichlet processes.” *Journal of Multivariate Analysis*, 102: 482–495. 898
- Leisen, F., Lijoi, A., and Spano, D. (2013). “A Vector of Dirichlet processes.” *Electronic Journal of Statistics*, 7: 62–90. 898
- Lijoi, A. and Nipoti, B. (2014). “A class of hazard rate mixtures for combining survival data from different experiments.” *Journal of the American Statistical Association*, 109: 802–814. 898
- Lijoi, A., Nipoti, B., and Prünster, I. (2014a). “Bayesian inference with dependent normalized completely random measures.” *Bernoulli*, 20: 1260–1291. 898
- Lijoi, A., Nipoti, B., and Prünster, I. (2014b). “Dependent mixture models: clustering and borrowing information.” *Computational Statistics and Data Analysis*, 71: 417–433. 898
- Lijoi, A. and Prünster, I. (2010). “Models beyond the Dirichlet Process.” In *Bayesian Nonparametrics*, 80–136. Cambridge University Press. 902
- Lyne, A.-M., Girolami, M., Strathmann, H., Simpson, D., and Atchade, Y. (2015). “On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods.” *Statistical Science*, 30: 443–467. MR3432836. doi: <https://doi.org/10.1214/15-STS523>. 903

- MacEachern, S. N. (1999). “Dependent nonparametric processes.” In *ASA Proceedings of the Section on Bayesian Statistical Science*. 897
- 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. 907
- Neal, R. M. (2000). “Markov chain sampling methods for Dirichlet process mixture models.” *Journal of Computational and Graphical Statistics*, 9: 249–265. 902
- Papaspiliopoulos, O. (2009). “A methodological framework for Monte Carlo probabilistic inference for diffusion processes.” In *Bayesian Time Series Models*, 91–112. Cambridge University Press. 903
- Ranganath, R. and Blei, D. M. (2017). “Correlated Random Measures.” *Journal of the American Statistical Association*, to appear. 898
- Regazzini, E., Lijoi, A., and Prünster, I. (2003). “Distributional results for means of normalized random measures with independent increments.” *Annals of Statistics*, 31: 560–585. 897
- Rhee, C.-H. and Glynn, P. W. (2015). “Unbiased estimation with square root convergence for SDE models.” *Operations Research*, 63: 1026–1043. 903
- Rodriguez, A. and Dunson, D. B. (2011). “Nonparametric Bayesian models through probit stick-breaking processes.” *Bayesian Analysis*, 6: 145–178. 911
- Sethuraman, J. (1994). “A constructive definition of Dirichlet priors.” *Statistica Sinica*, 4: 639–650. 897
- Silverman, B. W. (1985). “Some aspects of the spline smoothing approach to nonparametric curve fitting.” *Journal of the Royal Statistical Society, Series B*, 47: 1–52. 910
- Todeschini, A. and Caron, F. (2016). “Exchangeable random measures for sparse and modular graphs with overlapping communities.” ArXiv: 1602.0211. 898
- Yu, Y. and Meng, X.-L. (2011). “To Center or Not to Center: That is Not the Question – An Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Efficiency.” *Journal of Computational and Graphical Statistics*, 20: 531–570.
- Zhu, W. and Leisen, F. (2015). “A multivariate extension of a vector of Poisson-Dirichlet processes.” *Journal of Nonparametric Statistics*, 27: 89–105. 898

Acknowledgments

The authors would like to acknowledge helpful comments from two referees, an associate editor and an editor. Fabrizio Leisen’s research was supported by the European Community’s Seventh Framework Programme [FP7/2007-2013] under grant agreement no: 630677.