



BIROn - Birkbeck Institutional Research Online

Papageorgiou, Georgios (2019) Bayesian density regression for discrete outcomes. Australian and New Zealand Journal of Statistics 61 (3), pp. 336-359. ISSN 1467-842X.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/28540/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html> or alternatively contact lib-eprints@bbk.ac.uk.

Bayesian density regression for discrete outcomes

Georgios Papageorgiou
Department of Economics, Mathematics and Statistics
Birkbeck, University of London, UK
g.papageorgiou@bbk.ac.uk

August 13, 2019

Abstract

We develop Bayesian models for density regression with emphasis on discrete outcomes. The problem of density regression is approached by considering methods for multivariate density estimation of mixed scale variables, and obtaining conditional densities from the multivariate ones. The approach to multivariate mixed scale outcome density estimation that we describe represents discrete variables, either responses or covariates, as discretised versions of continuous latent variables. We present and compare several models for obtaining these thresholds in the challenging context of count data analysis where the response may be over- and/or under-dispersed in some of the regions of the covariate space. We utilise a nonparametric mixture of multivariate Gaussians to model the directly observed and the latent continuous variables. The paper presents a Markov chain Monte Carlo algorithm for posterior sampling, sufficient conditions for weak consistency, and illustrations on density, mean and quantile regression utilizing simulated and real datasets.

Keywords: Dirichlet process mixtures; joint models; Kullback-Leibler property; latent variables; over-dispersion; under-dispersion

1 Introduction

We consider methods for Bayesian density regression, with special attention to discrete responses. The basic objective is to estimate the conditional probability mass function of a response y given a vector of covariates \mathbf{x} , $f(y|\mathbf{x})$. This allows us to study how the distribution of the response changes with covariates, but of course from the estimated conditional density, other quantities of interest, such as the conditional mean, median or other quantiles, may be obtained. We approach the problem of conditional density estimation by considering methods for multivariate density estimation $f(y, \mathbf{x})$ of mixed scale outcomes. From the multivariate density we can obtain the conditional density using $f(y|\mathbf{x}) = f(y, \mathbf{x}) / \int f(y, \mathbf{x}) dy$.

Mixture models provide a very popular approach to density estimation. The general form of a mixture model for the joint density of (y, \mathbf{x}) is given by

$$f_P(y, \mathbf{x}) = \int_{\Theta} k(y, \mathbf{x}; \boldsymbol{\theta}) dP(\boldsymbol{\theta}),$$

where $k(\cdot, \cdot)$ is a probability kernel characterised by parameter $\boldsymbol{\theta}$ and $P(\cdot)$ is a probability measure on the parameter space Θ . In a Bayesian setting, the model specification is completed by specifying a prior distribution for the mixing measure $P(\cdot)$.

Here we adopt a nonparametric approach in which the prior on $P(\cdot)$ is taken to be a Dirichlet process (DP) (Ferguson, 1973), resulting in a so-called DP mixture model (DPMM). Due to the discreteness of the DP, the DPMM for the joint density can be expressed as $f_P(y, \mathbf{x}) = \sum_{h=1}^{\infty} \pi_h k(y, \mathbf{x}; \boldsymbol{\theta}_h)$. DPMMs were introduced by Lo (1984) and became popular after the work of Escobar & West (1995) and Müller, Erkanli & West (1996). They allow for flexible modelling of densities without having to specify the number of mixture components while at the same time achieving a balance between over- and under-fitting.

The implied DPMM for the conditional takes the form of a predictor-dependent mixture $f_P(y|\mathbf{x}) = \sum_{h=1}^{\infty} \pi_h(\mathbf{x}) k(y|\mathbf{x}; \boldsymbol{\theta}_h)$, where $\pi_h(\mathbf{x}) = \pi_h g(\mathbf{x}; \boldsymbol{\theta}_h) / f_P(\mathbf{x})$, showing how a flexible model for density regression is induced by the DPMM for the joint density. Furthermore, the mixture formulation allows for complex relationships between \mathbf{x} and functionals of the conditional to be captured. For instance, the implied model for the conditional mean takes the form $E(Y|\mathbf{x}) = \sum_{h=1}^{\infty} \pi_h(\mathbf{x}) E(Y|\mathbf{x}; \boldsymbol{\theta}_h)$, which is the approach to nonparametric regression that was introduced by Müller, Erkanli & West (1996). Related modelling approaches, which can also be used to analyse discrete data, include those of Shahbaba & Neal (2009), Taddy & Kottas (2010), Dunson & Bhattacharya (2011) and Hannah, Blei & Powell (2011).

An alternative formulation directly models the conditional $f(y|\mathbf{x})$ as a predictor dependent DPMM. Dunson, Pillai & Park (2007) and Chung & Dunson (2009) provided such model formulations for continuous responses. Directly modelling the conditional distribution can be advantageous over modelling the joint when the covariate vector \mathbf{x} is of high dimension, as the joint model requires estimation of the distribution of the high dimensional vector $(y, \mathbf{x}^\top)^\top$. A possible remedy to this problem of the joint model is to decompose the overall dependence among variables into clusters by assuming conditional independence, thereby reducing the problem of estimating a high dimensional distribution to the problem of estimating many univariate ones. On the other hand, the advantages of modelling the joint distribution include the simplicity by which a predictor-dependent mixture model is formulated. More importantly, modelling the joint density provides a direct mechanism for dealing with missing data, as the modelled joint $f(y, \mathbf{x})$ can be used for imputing missing responses and/or covariates, under a missing at random assumption (Kunihama, Halpern & Herring, 2019). Further, the joint formulation, under a conditional independence assumption, can be used to model a variety of types of covariates, such as functional data and images (Dunson & Bhattacharya, 2011).

Discretization of the latent continuous variables into the observed discrete ones requires the specification of either fixed cut-points or of models for the cut-points. There is a growing literature on DPMMs that utilise latent variables and fixed cut-points. We note the work of Kottas, Müller & Quintana (2005) and DeYoreo & Kottas (2018) who focus on ordinal data, DeYoreo & Kottas (2015) who present regression models for binary outcomes, Canale & Dunson (2015) who treat the problem of mixed-scale density estimation from both a theoretical and an applied perspective, and Norets & Pelenis (2012) who also present theory and applications, but based on finite mixture models. However, models with fixed cut-points have yet to be extended to include covariates that are fixed by design, such as binary treatment allocation variables in clinical trial settings or the offset term in count regression settings. Hence, in this paper, we work with modelled cut-points and we introduce covariates that are fixed by design through the modelled cut-points, as was also done by Papageorgiou, Richardson & Best (2015). We consider several models for specifying these cut-points in the context of count data analysis and we offer a comparison of their performance in a simulation study.

Specific models that we consider for the specification of the cut-points, within a DPMM, include the Poisson, negative binomial and generalised Poisson kernels. Whereas the negative binomial can model over-dispersed counts, the generalised Poisson can model both over- and under-dispersed counts. As was mentioned by Canale & Dunson (2011), DPMMs that utilise the Poisson (or negative binomial) kernel, even though they seem very flexible on the surface, in reality are not, as they are not able to adequately model under-dispersed counts. This is our motivation for considering the three kernels mentioned above.

Importantly, the paper provides a theoretical examination of the properties of the proposed model by presenting sufficient conditions for attaining weak consistency. To do so, the paper utilises the theorem of Schwartz (1965, Theorem 6.1) (see also Ghosh & Ramamoorthi (2003)) and the work of Wu & Ghosal (2008). The conditions that we provide are also sufficient for many special cases of the model presented here, namely DPMM with product mixture kernels that have been utilised multiple times in the literature, see e.g. Taddy & Kottas (2010) and Dunson & Bhattacharya (2011). Furthermore, by utilizing the conditions that we present here as a starting point, one could derive sufficient conditions for weak consistency of DP mixtures of generalised linear models. Surprisingly, this topic that has received very little attention in the literature. To the best of the author’s knowledge, the only other paper that considers it is that of Hannah, Blei & Powell (2011).

The remainder of this paper is arranged as follows. Section 2 presents the methodology and Section 3 presents a brief description of the MCMC algorithm we have implemented and methods for posterior inference, with most of the details presented in Appendix 8. Section 4 presents sufficient conditions for weak consistency, with the proof deferred to Appendix 9. Sections 5 and 6 present results from a simulation study and an applications to a real dataset. The paper concludes with a brief discussion.

2 Methodology

2.1 Model specification

Let Y denote a discrete response with support on (a subset of) the non-negative integers and let $\mathbf{X} = (\mathbf{X}_d^\top, \mathbf{X}_c^\top)^\top$ denote a vector of p mixed scale covariates, where \mathbf{X}_d is a vector of p_d discrete variables and \mathbf{X}_c is a vector of p_c continuous variables, $p_d + p_c = p$. Our goal is to jointly model $\mathbf{Z} = (Y, \mathbf{X}^\top)^\top$ as a draw from an unknown density $f \in \mathcal{F}$ with respect to an appropriate measure, where \mathcal{F} denotes the set of all such densities. The next few paragraphs describe how a prior Π on \mathcal{F} is induced.

Discrete variables, either responses or covariates, are assumed to be discretised or rounded versions of continuous latent variables (Muthen, 1984). Our presentation below concerns generic discrete and continuous latent variables denoted by Z and Z^* respectively. Observed and latent variables are connected according to the rule

$$Z = z \text{ if and only if } z^* \in R(z) = (c_{z-1}, c_z], z = 0, 1, 2, \dots,$$

where $R(z)$ is an interval on the real line with bounds given by: $c_{-1} = -\infty$, and for $q \geq 0$, $c_q = c_q(\boldsymbol{\lambda}) = \Phi^{-1}\{F(q; \boldsymbol{\lambda})\}$. Here $\Phi(\cdot)$ is the cumulative distribution function of a standard normal variable, and $F(\cdot; \boldsymbol{\lambda})$ denotes an appropriate cumulative distribution function. Further, latent variables are assumed to be independent draws from a $N(0, 1)$ distribution, where the mean and variance are restricted to be zero and one respectively as they are non-identifiable by the data. It is easy to see that with this specification the marginal distribution of Z is $F(z; \boldsymbol{\lambda})$:

$$\Pr(Z \leq z) = \Pr\{Z^* < \Phi^{-1}[F(z; \boldsymbol{\lambda})]\} = \Pr[\Phi(Z^*) < F(z; \boldsymbol{\lambda})] = F(z; \boldsymbol{\lambda}), \quad (1)$$

where the last equality follows because $\Phi(Z^*)$ has a uniform distribution on the unit interval (see e.g. Robert & Casella (2005)).

Generally, equation (1) is satisfied if one assumes $c_q = G^{-1}[F(q; \boldsymbol{\lambda})]$ and that the latent variables are independent draws from G , where G is a continuous cumulative distribution function. Common choices for the density of the latent variables Z^* include the Student’s t, Weibull, lognormal and gamma. In this paper, we focus on the case where $G = \Phi$ and we examine alternatives for function $F(\cdot; \boldsymbol{\lambda})$. The next few paragraphs discuss choices for $F(\cdot; \boldsymbol{\lambda})$ for count and binomial data.

Count data: For modelling counts, we may take $F(\cdot; \boldsymbol{\lambda})$ to be the distribution function of a $\text{Poisson}(H\xi_1)$ variable. Here $\boldsymbol{\lambda} = (\xi_1, H)^\top$, where $\xi_1 > 0$ denotes the Poisson rate and H the offset term. Hence, the associated probability mass function (pmf) is given by $\Pr(Z = z; \xi_1) = \exp(-H\xi_1)(H\xi_1)^z/z!$, with implied mean and variance both equal to $H\xi_1$. Further, to account for potential over-dispersion in the counts, we may take $F(\cdot; \boldsymbol{\lambda})$ to be the negative binomial distribution function. With this choice, the vector $\boldsymbol{\lambda} = (\xi_1, \xi_2, H)^\top$ includes the offset term H and two unknown parameters $\boldsymbol{\xi} = (\xi_1, \xi_2)^\top$ that allow for extra flexibility compared to the flexibility afforded by the single parameter Poisson distribution. The pmf is given by

$$\Pr(Z = z; \xi_1, \xi_2) = \frac{\Gamma(z + \xi_1)}{\Gamma(\xi_1)\Gamma(z + 1)} \left(\frac{\xi_2}{H + \xi_2} \right)^{\xi_1} \left(\frac{H}{H + \xi_2} \right)^z,$$

where $\xi_1 > 0$ and $\xi_2 > 0$, and it implies a mean and a variance of $H\xi_1/\xi_2$ and $H\xi_1/\xi_2(1 + H/\xi_2)$, respectively. Both of these choices, however, are quite restrictive as they require the variance to be equal or greater than the mean. Under-dispersion cannot be modelled in a satisfactory way even with nonparametric mixtures of Poisson or negative binomial pmfs. For this reason, we also consider the generalised Poisson (Consul & Famoye, 1992) distribution function that allows for both over- and under-dispersion relative to the Poisson. The pmf is given by

$$\Pr(Z = z; \xi_1, \xi_2) = \xi_1[\xi_1 + (\xi_2 - 1)z]^{z-1} \xi_2^{-z} \exp\{-[\xi_1 + (\xi_2 - 1)z]/\xi_2\}/z!, \quad (2)$$

where $\xi_1 > 0$ and $\xi_2 > 0$. It may be shown that the pmf implies that $E(Z) = \xi_1$ and $\text{var}(Z) = \xi_2^2\xi_1$. Hence, the distribution is over-dispersed when $\xi_2 > 1$, under-dispersed when $\xi_2 < 1$ and it reduces to the Poisson pmf when $\xi_2 = 1$. When $\xi_2 < 1$ an upper bound m is set on the counts, where m is the largest integer for which $\xi_1 + (\xi_2 - 1)m > 0$, so that the pmf remains positive on its support. In this case the normalizing constant of the pmf in (2) needs to be computed because the sum of probabilities $\sum_{z=0}^m \Pr(z; \xi_1, \xi_2)$ is not necessarily equal to one. Lastly, it is straight forward to include an offset term H in this pmf, by replacing ξ_1 by $H\xi_1$.

Binomial data: For modelling binomial data, we may take $F(\cdot; \boldsymbol{\lambda})$ to be the distribution function of a $\text{Binomial}(H, \xi_1)$ variable, with ξ_1 denoting the success probability and H the number of trials. A more flexible approach would be to take $F(\cdot; \boldsymbol{\lambda})$ to be the beta-binomial distribution function, where $\boldsymbol{\lambda} = (\xi_1, \xi_2, H)^\top$ includes two unknown parameters. The associated pmf is given by

$$\Pr(Z = z; \xi_1, \xi_2) = \binom{H}{z} \frac{\text{Beta}(z + \xi_1, H - z + \xi_2)}{\text{Beta}(\xi_1, \xi_2)},$$

where $\xi_1 > 0$ and $\xi_2 > 0$.

Binary data: The special case of binomial data with $H = 1$ can equivalently be treated as

$$\begin{aligned} Z = 0 & \text{ if and only if } z^* < 0, \\ z^* & \sim N(\mu_z^*, 1), \end{aligned} \quad (3)$$

which is the approach of our preference as it allows for simpler posterior sampling.

We induce a prior Π on \mathcal{F} by assuming a nonparametric mixture model for $\mathbf{Z} = (Y, \mathbf{X}_d^\top, \mathbf{X}_c^\top)^\top$,

$$f_P(\mathbf{z}) = \int_{\Theta} k(\mathbf{z}; \boldsymbol{\theta}) dP(\boldsymbol{\theta}) \quad (4)$$

that utilises a parametric kernel $k(\mathbf{z}; \boldsymbol{\theta})$ and a nonparametric model for the random mixing distribution P on Θ .

The kernel $k(\mathbf{z}; \boldsymbol{\theta})$ is obtained by assuming a q -dimensional Gaussian for the continuous observed and latent variables, $\mathbf{Z}^* = (Y^*, \mathbf{X}_d^{*\top}, \mathbf{X}_c^\top)^\top$, and integrating out the latent variables

$$k(\mathbf{z}; \boldsymbol{\theta}) = \int_{R(y)} \int_{R(x_d)} \mathbb{N}(\mathbf{z}^*; \boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*) d\mathbf{x}_d^* dy^*, \quad (5)$$

where $q = 1 + p$ and $\boldsymbol{\theta} = (\boldsymbol{\xi}, \boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$ denotes the kernel parameters. Due to the non-identifiability of the location and scale parameters of the distribution of the latent variables, with the exception of the location parameter in (3), the mean $\boldsymbol{\mu}^*$ and covariance $\boldsymbol{\Sigma}^*$ are of the form

$$\boldsymbol{\mu}^* = \begin{pmatrix} \mathbf{0} \\ \boldsymbol{\mu} \end{pmatrix}, \quad \boldsymbol{\Sigma}^* = \begin{bmatrix} \mathbf{C} & \boldsymbol{\nu}^\top \\ \boldsymbol{\nu} & \boldsymbol{\Sigma} \end{bmatrix}, \quad (6)$$

where \mathbf{C} is the covariance matrix of the latent continuous variables and has diagonal elements equal to one i.e. it is a correlation matrix. Further, $\boldsymbol{\Sigma}$ is the unrestricted covariance matrix of the directly observed continuous variables. Specific examples are provided later in the paper and they concern: (i) a count response and a continuous covariate—see (11); and (ii) a count response and a binary and a continuous covariate—see (12). Below we consider some special cases, where the kernel can be simplified, and draw connections to the literature.

As a first special case we consider the so-called product kernel, obtained when $\boldsymbol{\nu} = \mathbf{0}$ and $\mathbf{C} = \mathbf{I}$, where \mathbf{I} is the identity matrix. The choice $\mathbf{C} = \mathbf{I}$ implies that the discrete variables are conditionally independent and $\boldsymbol{\nu} = \mathbf{0}$ implies that discrete and continuous variables are conditionally independent. Further, assuming that $\boldsymbol{\Sigma}$ is diagonal implies that the continuous variables are conditionally independent. Within a Bayesian nonparametric framework, such kernels have been utilised by Taddy & Kottas (2010) and Dunson & Bhattacharya (2011).

As another special case we consider the scenario where all p covariates are continuous. Here the joint mean is $\boldsymbol{\mu}^* = (\mu_y, \boldsymbol{\mu}_x^\top)^\top$ and the joint covariance $\boldsymbol{\Sigma}^*$ has submatrix $\mathbf{C} = 1$, $\boldsymbol{\nu}$ a p -dimensional vector, and $\boldsymbol{\Sigma}$ a $p \times p$ positive definite matrix. The kernel in (5) may be written as

$$k(\mathbf{z}; \boldsymbol{\theta}) = \int_{R(y)} \mathbb{N}(\mathbf{z}^*; \boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*) dy^* = \mathbb{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) \int_{R(y)} \mathbb{N}(y^*; m^*, v^*) dy^*, \quad (7)$$

where $m^* = \mathbb{E}(y^* | \mathbf{x}) = \mu_y + \boldsymbol{\nu}^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})$ and $v^* = \text{var}(y^* | \mathbf{x}) = 1 - \boldsymbol{\nu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\nu}$.

When the response is binary, (7) becomes

$$k(Y = 1, \mathbf{x}; \boldsymbol{\theta}) = \mathbb{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) \Phi(m^* / \sqrt{v^*}),$$

implying a probit regression model for the conditional probability of success. A Bayesian non-parametric approach for binary regression based on this kernel has been developed by DeYoreo & Kottas (2015).

2.2 Prior specification

Following a Bayesian nonparametric approach, we assign to the unknown mixing distribution $P(\cdot)$ a Dirichlet process (DP) prior (Ferguson, 1973). A DP prior is characterised by two parameters: a total mass or concentration parameter α and a base distribution P_0 over the parameter space. According to the stick-breaking representation (Sethuraman, 1994)

$$P(\cdot) = \sum_{h=1}^{\infty} \pi_h \delta_{\boldsymbol{\theta}_h}(\cdot),$$

which when combined with (4) leads to the following DPMM for $\mathbf{z} = (y, \mathbf{x}_d^\top, \mathbf{x}_c^\top)^\top$

$$f_P(\mathbf{z}) = \sum_{h=1}^{\infty} \pi_h k(\mathbf{z}; \boldsymbol{\theta}_h). \quad (8)$$

In the above countable mixture, the weights $\pi_h, h \geq 1$, are constructed by the so-called stick-breaking process: $\pi_1 = v_1$, and for $l \geq 2$, $\pi_l = v_l \prod_{h=1}^{l-1} (1 - v_h)$, where $v_k, k \geq 1$, are independent draws from a Beta(1, α) distribution. Further, the atoms $\boldsymbol{\theta}_h, h \geq 1$, are obtained as independent draws from the base distribution P_0 , which consists of three independent priors for the elements of $\boldsymbol{\theta}_h = (\boldsymbol{\xi}_h, \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h^*)$, $h \geq 1$, described next.

Firstly, the priors on the set of parameters $\{\boldsymbol{\xi}_h\}$ depend on the choice of the function $F(\cdot)$ in (1). For all functions, we take these priors to be very close to uninformative. Table 1 provides a summary. For the rate ξ_1 of the Poisson distribution, we take the prior to be Gamma(1, 0.1), a gamma distribution with mean 10 and variance 100. The same prior is taken for the ξ_1 and ξ_2 parameters of the negative binomial distribution function. Further, the mean parameter of the generalised Poisson distribution, ξ_1 , is given the same gamma prior, while the dispersion parameter, ξ_2 , is given a normal prior with mean and variance equal to 1, and truncated from below at 0.05. In addition, the binomial probability of success is given a uniform prior, while the two parameters of the beta-binomial distribution are given vague gamma priors.

Table 1: Prior distributions on the pmf parameters.

| pmf | Prior distribution |
|--|--|
| 1. Poisson(ξ_1) | $\xi_1 \sim \text{Gamma}(1, 0.1)$ |
| 2. Negative binomial(ξ_1, ξ_2) | $\xi_1, \xi_2 \sim \text{Gamma}(1, 0.1)$ |
| 3. Generalised Poisson(ξ_1, ξ_2) | $\xi_1 \sim \text{Gamma}(1, 0.1)$ $\xi_2 \sim \text{N}(1, 1)\mathbf{1}(\xi_2 > 0.05)$ |
| 4. Binomial(ξ_1) | $\xi_1 \sim \text{Beta}(1, 1)$ |
| 5. Beta-binomial(ξ_1, ξ_2) | $\xi_1, \xi_2 \sim \text{Gamma}(1, 0.1)$ |

Secondly, the prior on $\boldsymbol{\mu}_h$, the non-zero part of $\boldsymbol{\mu}_h^*$, is taken to be multivariate normal $\boldsymbol{\mu}_h \sim \text{N}(\mathbf{d}, \mathbf{D})$. The mean \mathbf{d} is taken to be equal to the centre of the dataset. Specifically, the part of \mathbf{d} that corresponds to continuous variables is taken to be equal to the sample mean while the part that corresponds to binary variables is taken to be a transformation of the observed sample proportion. Let p^* denote an observed sample proportion. The corresponding prior mean is taken to be $-\Phi^{-1}(1 - p^*)$, which along with the mechanism in (3) implies a prior proportion equal to p^* . Further, the covariance matrix \mathbf{D} is taken to be diagonal. Its elements that correspond to variances of continuous variables are set equal to a small multiple (here taken to be 1/8) of the square of the observed data range (Richardson & Green, 1997), while the elements that correspond to binary variables are set equal to a constant (here taken to be 5).

Lastly, the prior distribution assigned to the restricted covariance matrices $\boldsymbol{\Sigma}_h^*, h \geq 1$, in (6) is specified by utilizing the methods of Zhang, Boscardin & Belin (2006) and Barnard, McCulloch & Meng (2000): we add into the model variance parameters that are non-identifiable by the data and then separate identifiable from non-identifiable parameters. The starting point is a Wishart($\mathbf{E}_h; \eta, \mathbf{H}$) prior for unrestricted $q \times q$ covariance matrices $\mathbf{E}_h, h \geq 1$:

$$p(\mathbf{E}_h; \eta, \mathbf{H}) \propto |\mathbf{E}_h|^{(\eta-q-1)/2} \exp[\text{tr}(-\mathbf{H}^{-1}\mathbf{E}_h/2)],$$

where

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_{11} & \mathbf{H}_{12} \\ \mathbf{H}_{12}^\top & \mathbf{H}_{22} \end{bmatrix},$$

where \mathbf{H}_{11} is a $(1 + p_d) \times (1 + p_d)$ correlation matrix, \mathbf{H}_{22} is a $p_c \times p_c$ unrestricted covariance matrix, and \mathbf{H}_{12} is a $(1 + p_d) \times p_c$ matrix of covariances.

We decompose $\mathbf{E}_h = \mathbf{D}_h^{1/2} \boldsymbol{\Sigma}_h^* \mathbf{D}_h^{1/2}$, where $\mathbf{D}_h = \text{Diag}(d_{h,1}^2, \dots, d_{h,1+p_d}^2, 1, \dots, 1)$ is a diagonal matrix of $(1 + p_d)$ non identifiable variance parameters and p_c ones that correspond to identifiable variances, and $\boldsymbol{\Sigma}_h^*$ is a covariance matrix that satisfies the restrictions imposed by the data. The Jacobian of this transformation is $J(\mathbf{E}_h \rightarrow \mathbf{D}_h, \boldsymbol{\Sigma}_h^*) = \prod_{j=1}^{(1+p_d)} d_{hj}^{q-1} = |\mathbf{D}_h|^{(q-1)/2}$. Hence, we obtain the following prior for the pairs $(\mathbf{D}_h, \boldsymbol{\Sigma}_h^*), h \geq 1$:

$$p(\mathbf{D}_h, \boldsymbol{\Sigma}_h^*; \eta, \mathbf{H}) \propto |\mathbf{E}_h|^{(\eta-q-1)/2} \exp[\text{tr}(-\mathbf{H}^{-1} \mathbf{E}_h/2)] J(\mathbf{E}_h \rightarrow \mathbf{D}_h, \boldsymbol{\Sigma}_h^*).$$

In our analyses we take η to be equal to $q + 2$ and \mathbf{H} to be diagonal, with sub-matrix \mathbf{H}_{11} equal to the identity matrix and with sub-matrix \mathbf{H}_{22} having entries equal to a small multiple ($1/8$) of the square of the observed data range.

3 Posterior sampling and inference

3.1 MCMC sampler

The main tools that we utilise in developing an MCMC sampler are the ‘blocked’ approach of Ishwaran & James (2001) and adaptive Metropolis algorithms (Roberts & Rosenthal, 2009) to achieve optimal scaling (Roberts & Rosenthal, 2001).

We start by truncating the countable mixture in (8) to include T components and by introducing an allocation variable δ . The model is now written in the following equivalent way

$$\begin{aligned} \mathbf{z} | \boldsymbol{\theta}, \delta = l &\sim k(\mathbf{z}; \boldsymbol{\theta}_l), \\ \Pr(\delta = l | \alpha) &= \pi_l, l = 1, 2, \dots, T. \end{aligned}$$

The likelihood associated with n independent and identically distributed observations $\mathbf{z}_i, i = 1, 2, \dots, n$, can be written as

$$\ell(\boldsymbol{\theta}, \alpha; \mathbf{z}, \boldsymbol{\delta}) = \ell(\boldsymbol{\theta}, \alpha; \mathbf{z}_i, \delta_i = l_i, i = 1, \dots, n) = \prod_i k(\mathbf{z}_i; \boldsymbol{\theta}_{l_i}) \pi_{l_i}.$$

Recall that \mathbf{z}_i consist of a discrete response y_i , p_d discrete covariates $\mathbf{x}_{d,i} = (x_{d,i,1}, \dots, x_{d,i,p_d})^\top$ and p_c continuous covariates $\mathbf{x}_{c,i} = (x_{c,i,1}, \dots, x_{c,i,p_c})^\top, i = 1, 2, \dots, n$. Augmenting the likelihood with the latent continuous variables \mathbf{y}_i^* and $\mathbf{x}_{d,i}^* = (x_{d,i,1}^*, \dots, x_{d,i,p_d}^*)^\top$, we obtain

$$\begin{aligned} \ell(\boldsymbol{\theta}, \alpha; \mathbf{z}, \boldsymbol{\delta}, \mathbf{y}^*, \mathbf{x}_d^*) &= \\ \prod_i \left\{ \mathbf{1}(y_i^* \in R(y_i)) \left[\prod_{m=1}^{p_d} \mathbf{1}(x_{d,i,m}^* \in R(x_{d,i,m})) \right] \mathbf{N}(y_i^*, \mathbf{x}_{d,i}^*, \mathbf{x}_{c,i}; \boldsymbol{\mu}_{l_i}^*, \boldsymbol{\Sigma}_{l_i}^*) \pi_{l_i} \right\}, \end{aligned}$$

from which the full posterior follows

$$\pi(\boldsymbol{\theta}, \boldsymbol{\delta}, \mathbf{y}^*, \mathbf{x}_d^*, \alpha | \mathbf{y}, \mathbf{x}_d, \mathbf{x}_c) \propto \ell(\boldsymbol{\theta}, \alpha; \mathbf{z}, \boldsymbol{\delta}, \mathbf{y}^*, \mathbf{x}_d^*) p_0(\boldsymbol{\theta}, \alpha).$$

3.2 An exact algorithm

Truncation of the mixture density in (8) can be avoided by implementing a slice sampler (Walker, 2007; Papaspiliopoulos, 2008). This requires augmenting the likelihood with uniform random variables $U_i \sim \text{U}(0, 1), i = 1, \dots, n$, such that the complete likelihood becomes

$$\prod_i \left\{ \mathbf{1}(y_i^* \in R(y_i)) \left[\prod_{m=1}^{p_d} \mathbf{1}(x_{d,i,m}^* \in R(x_{d,i,m})) \right] \text{N}(y_i^*, \mathbf{x}_{d,i}^*, \mathbf{x}_{c,i}; \boldsymbol{\mu}_{l_i}^*, \boldsymbol{\Sigma}_{l_i}^*) \mathbf{1}(u_i < \pi_{l_i}) \right\}.$$

Details on the updating steps of the MCMC algorithm are provided in Appendix 8. We note that throughout the paper we utilise the truncated sampler as truncation allows sampling from the full posterior distribution, including sampling from the posterior of the random mixing distribution P , which in turn allows for proper uncertainty quantification for density estimates. The same strategy was utilised by DeYoreo & Kottas (2018).

3.3 Posterior inference

Recall that $\mathbf{z} = (y, \mathbf{x}_d^\top, \mathbf{x}_c^\top)^\top$ and $\mathbf{z}^* = (y^*, \mathbf{x}_d^{*\top}, \mathbf{x}_c^\top)^\top$. Then the model for \mathbf{z} , truncated to include T components, is expressed as

$$f_P(\mathbf{z}) = \sum_{h=1}^T \pi_h k(\mathbf{z}; \boldsymbol{\theta}_h) = \sum_{h=1}^T \pi_h \int_{R(y)} \int_{R(x_d)} \text{N}(\mathbf{z}^*; \boldsymbol{\mu}_h^*, \boldsymbol{\Sigma}_h^*) dy^* d\mathbf{x}_d^*. \quad (9)$$

Further, let $\boldsymbol{\mu}_h^*$ and $\boldsymbol{\Sigma}_h^*$ be partitioned as follows

$$\boldsymbol{\mu}_h^* = \begin{pmatrix} \boldsymbol{\mu}_{h,d} \\ \boldsymbol{\mu}_{h,x_c} \end{pmatrix}, \quad \boldsymbol{\Sigma}_h^* = \begin{bmatrix} \boldsymbol{\Sigma}_{h,d,d} & \boldsymbol{\Sigma}_{h,d,x_c} \\ \boldsymbol{\Sigma}_{h,x_c,d} & \boldsymbol{\Sigma}_{h,x_c,x_c} \end{bmatrix},$$

where subscript d denotes the continuous latent variables underlying discrete ones $(y^*, \mathbf{x}_d^{*\top})^\top$. Then (9) may be expressed as

$$\sum_{h=1}^T \pi_h \text{N}(\mathbf{x}_c; \boldsymbol{\mu}_{h,x_c}, \boldsymbol{\Sigma}_{h,x_c,x_c}) \int_{R(y)} \int_{R(x_d)} \text{N}(y^*, \mathbf{x}_d^* | \mathbf{x}_c; \boldsymbol{\mu}_{h,d,c}, \boldsymbol{\Sigma}_{h,d,c}) dy^* d\mathbf{x}_d^*, \quad (10)$$

where $\boldsymbol{\mu}_{h,d,c} = \boldsymbol{\mu}_{h,d} + \boldsymbol{\Sigma}_{h,d,x_c} \boldsymbol{\Sigma}_{h,x_c,x_c}^{-1} (\mathbf{x}_c - \boldsymbol{\mu}_{h,x_c})$ and $\boldsymbol{\Sigma}_{h,d,c} = \boldsymbol{\Sigma}_{h,d,d} - \boldsymbol{\Sigma}_{h,d,x_c} \boldsymbol{\Sigma}_{h,x_c,x_c}^{-1} \boldsymbol{\Sigma}_{h,x_c,d}$.

Further, to obtain an expression for the conditional $f_P(y|\mathbf{x})$, first let $\mathbf{x} = (\mathbf{x}_d^\top, \mathbf{x}_c^\top)^\top$ and $\mathbf{x}^* = (\mathbf{x}_d^{*\top}, \mathbf{x}_c^\top)^\top$. In addition, let $\boldsymbol{\mu}_{h,x}^*$ and $\boldsymbol{\Sigma}_{h,xx}^*$ denote the mean and covariance of \mathbf{x}^* , which we partition as

$$\boldsymbol{\mu}_{h,x}^* = \begin{pmatrix} \boldsymbol{\mu}_{h,x_d} \\ \boldsymbol{\mu}_{h,x_c} \end{pmatrix}, \quad \boldsymbol{\Sigma}_{h,xx}^* = \begin{bmatrix} \boldsymbol{\Sigma}_{h,x_d,x_d} & \boldsymbol{\Sigma}_{h,x_d,x_c} \\ \boldsymbol{\Sigma}_{h,x_c,x_d} & \boldsymbol{\Sigma}_{h,x_c,x_c} \end{bmatrix}.$$

Now, the conditional $f_P(y|\mathbf{x})$ may be expressed as

$$f_P(y|\mathbf{x}) = \frac{f_P(\mathbf{z})}{f_P(\mathbf{x})} = \frac{\sum_{h=1}^T \pi_h k(\mathbf{z}; \boldsymbol{\theta}_h)}{\sum_{h=1}^T \pi_h k(\mathbf{x}; \boldsymbol{\theta}_h)},$$

where $k(\mathbf{x}; \boldsymbol{\theta}_h) = \int_{R(x_d)} \text{N}(\mathbf{x}^*; \boldsymbol{\mu}_{h,x}^*, \boldsymbol{\Sigma}_{h,xx}^*) d\mathbf{x}_d^*$. Lastly, utilizing a similar factorization as in (10), we may write

$$k(\mathbf{x}; \boldsymbol{\theta}_h) = \text{N}(\mathbf{x}_c; \boldsymbol{\mu}_{h,x_c}, \boldsymbol{\Sigma}_{h,x_c,x_c}) \int_{R(x_d)} \text{N}(\mathbf{x}_d^* | \mathbf{x}_c; \boldsymbol{\mu}_{h,x_d,c}, \boldsymbol{\Sigma}_{h,x_d,c}) d\mathbf{x}_d^*$$

where $\boldsymbol{\mu}_{h,x_d,c} = \boldsymbol{\mu}_{h,x_d} + \boldsymbol{\Sigma}_{h,x_d x_c} \boldsymbol{\Sigma}_{h,x_c x_c}^{-1} (\boldsymbol{x}_c - \boldsymbol{\mu}_{h,x_c})$ and $\boldsymbol{\Sigma}_{h,x_d,c} = \boldsymbol{\Sigma}_{h,x_d x_d} - \boldsymbol{\Sigma}_{h,x_d x_c} \boldsymbol{\Sigma}_{h,x_c x_c}^{-1} \boldsymbol{\Sigma}_{h,x_c x_d}$. Hence, we find that $f_P(y|\boldsymbol{x})$ can be expressed as

$$\frac{\sum_{h=1}^T \pi_h \mathcal{N}(\boldsymbol{x}_c; \boldsymbol{\mu}_{h,x_c}, \boldsymbol{\Sigma}_{h,x_c x_c}) \int_{R(y)} \int_{R(x_d)} \mathcal{N}(y^*, \boldsymbol{x}_d^* | \boldsymbol{x}_c; \boldsymbol{\mu}_{h,d,c}, \boldsymbol{\Sigma}_{h,d,c}) dy^* d\boldsymbol{x}_d^*}{\sum_{h=1}^T \pi_h \mathcal{N}(\boldsymbol{x}_c; \boldsymbol{\mu}_{h,x_c}, \boldsymbol{\Sigma}_{h,x_c x_c}) \int_{R(x_d)} \mathcal{N}(\boldsymbol{x}_d^* | \boldsymbol{x}_c; \boldsymbol{\mu}_{h,x_d,c}, \boldsymbol{\Sigma}_{h,x_d,c}) d\boldsymbol{x}_d^*}.$$

Of interest is the quantity

$$f(y|\boldsymbol{x}) = \int \frac{f_P(\boldsymbol{z})}{f_P(\boldsymbol{x})} dP.$$

Given samples from the posterior of P , denoted by $P_s, s = 1, \dots, S$, $f(y|\boldsymbol{x})$ will be approximated by

$$f(y|\boldsymbol{x}) = S^{-1} \sum_{s=1}^S f_{P_s}(y|\boldsymbol{x}).$$

For each sampled conditional $f_{P_s}(y|\boldsymbol{x})$ we calculate all functionals of interest and thereby obtain posteriors for these functionals. This process is carried out on a grid of \boldsymbol{x} values, enabling inference about the dependence of the conditional $f(y|\boldsymbol{x})$ and its functionals on \boldsymbol{x} .

4 Weak Consistency

Here we provide sufficient conditions under which the proposed mixture model for the joint density attains weak posterior consistency at the true distribution f_0 . Weak consistency refers to the property of the posterior distribution to concentrate in regions of \mathcal{F} that are close to the true distribution f_0 in the weak topology sense. To formalise this concept we next provide some basic definitions, followed by two important theorems based on which we establish weak consistency for the proposed model.

Recall that \mathcal{F} denotes the space of mixed scale densities with respect to some suitable measure. A weak neighbourhood of $f_0 \in \mathcal{F}$ of radius ϵ is defined as

$$V_\epsilon(f_0) = \left\{ f \in \mathcal{F} : \left| \int \phi_i f - \int \phi_i f_0 \right| < \epsilon, i = 1, \dots, m \right\},$$

where $\phi_i, i = 1, \dots, m$, are bounded, continuous functions.

Further recall that Π is a prior on \mathcal{F} and let $\boldsymbol{Z}_1, \dots, \boldsymbol{Z}_n$ be i.i.d. with common density f_0 . The posterior probability of $A \subset \mathcal{F}$ is given by

$$\Pi(A | \boldsymbol{Z}_1, \dots, \boldsymbol{Z}_n) = \frac{\int_A \prod_{i=1}^n f(\boldsymbol{Z}_i) \Pi(df)}{\int_{\mathcal{F}} \prod_{i=1}^n f(\boldsymbol{Z}_i) \Pi(df)}.$$

A prior Π is said to be weakly consistent at f_0 if

$$\Pi(V_\epsilon | \boldsymbol{Z}_1, \dots, \boldsymbol{Z}_n) \rightarrow 1$$

for all weak neighbourhoods V_ϵ of f_0 , with \mathcal{P}_{f_0} probability 1.

The Kullback-Leibler (KL) neighbourhood of $f_0 \in \mathcal{F}$ of radius ϵ is defined as follows

$$K_\epsilon(f_0) = \left\{ f \in \mathcal{F} : \int f_0 \log(f_0/f) < \epsilon \right\}.$$

A density f_0 is said to be in the KL support of Π if $\Pi(K_\epsilon(f_0)) > 0$ for all $\epsilon > 0$.

To establish weak consistency we utilise the following theorem of Schwartz (1965).

Theorem: If f_0 is in the KL support of Π , then the posterior is weakly consistent at f_0 .

To prove the KL property for the mixture model described in this paper, we utilise Theorem 1 of Wu & Ghosal (2008), stated below.

Theorem: Let $f_0(y, \mathbf{x})$ denote the true density, Π^* the prior on $\mathcal{M}(\Theta)$, the space of probability measures on Θ , and Π the prior induced on \mathcal{F} . If for any $\epsilon > 0$, there exists mixing distribution P_ϵ and a $\mathcal{W} \subset \mathcal{M}(\Theta)$ with $\Pi^*(\mathcal{W}) > 0$, such that

$$A_1: \sum_y \int f_0(y, \mathbf{x}) \log[f_0(y, \mathbf{x})/f_{P_\epsilon}(y, \mathbf{x})]d\mathbf{x} < \epsilon,$$

$$A_2: \sum_y \int f_0(y, \mathbf{x}) \log[f_{P_\epsilon}(y, \mathbf{x})/f_P(y, \mathbf{x})]d\mathbf{x} < \epsilon \text{ for every } P \in \mathcal{W},$$

then $f_0 \in \text{KL}(\Pi)$.

The main result on the weak consistency of the proposed model is stated in the following lemma. It is based on a special case of the overall model, namely the case where there is a discrete response and p continuous covariates, and hence the kernel is the one that appears in (7).

Lemma: Under the following conditions we may establish that the prior defined in (4) and (5) satisfies conditions A_1 and A_2 of the theorem of Wu & Ghosal (2008):

C_1 : the true density $f_0(y, \mathbf{x})$ can be expressed as $f_0(y, \mathbf{x}) = \int K(y; \boldsymbol{\xi})f_0^*(\boldsymbol{\xi}, \mathbf{x})d\boldsymbol{\xi}$ and it is compactly supported,

C_2 : the density f_0^* is continuous, compactly supported and it satisfies $0 < f_0^*(\mathbf{z}) < M$ almost everywhere,

C_3 : there exists an m such that $\inf_{\|\mathbf{x}\| \leq m} f_0^*(\mathbf{x}|\boldsymbol{\xi}) \geq c$ for all $\boldsymbol{\xi}$,

C_4 : for \mathbf{x} such that $\|\mathbf{x}\| > m$, $f_0^*(\boldsymbol{\xi}, \mathbf{x}) < c$. Further, f_0^* is decreasing as $\|\mathbf{x}\|$ increases more than m ,

C_5 : $|\sum_y \int_x f_0(y, \mathbf{x}) \log f_0(y, \mathbf{x})d\mathbf{x}| < \infty$.

Hence, $f_0 \in \text{KL}(\Pi)$ and by the theorem of Schwartz (1965), the posterior is weakly consistent at f_0 . The proof, which is provided in Appendix 9, clarifies why these five conditions are needed.

5 Simulation study

Here we present results from a simulation study. The key aim is to provide insights into the effect of the kernel choice on posterior inference for a regression surface and other functionals of the conditional pmfs.

The data-generating mechanism consists of a continuous covariate $X \sim U(0, 11)$ and a count response Y , which has the following conditional mean function

$$\mu_x = E(Y|X = x) = 1 + \sin(\pi x/5) + x/4.$$

The mechanism from which we generate the responses is

$$Y|X = x \sim \begin{cases} \text{Poisson}(H\mu_x), & x < 3, \\ \text{Poisson}(H\mu_x|\epsilon_1|), & 3 < x < 6, \\ \text{round}(H\mu_x + \epsilon_2), & 6 < x < 9, \\ \text{Poisson}(H\mu_x|\epsilon_3|), & x > 9, \end{cases}$$

where H denotes the offset term, here generated from $H \sim U(10, 30)$. Furthermore, $\text{round}(\cdot)$ is a function that rounds its argument to the closest integer and $\epsilon_i, i = 1, 2, 3$, are normally distributed random errors: $\epsilon_1 \sim N(1, 0.15^2)$, $\epsilon_2 \sim N(0, 2^2)$, and $\epsilon_3 \sim N(1, 0.30^2)$. Because the mechanism places positive probability on negative realizations y , we take the realised response to be $\max(0, y)$. The mechanism generates responses that are Poisson distributed over the range $x < 3$, ‘mildly’ over-dispersed over $3 < x < 6$, under-dispersed over $6 < x < 9$, and ‘severely’ over-dispersed relative to the Poisson distribution over $x > 9$. We take the sample size to be $n = 500$. A simulated dataset is shown in Figure 1.

For each simulated dataset we fit a model of the form

$$f(y_i^*, x_i) = \sum_{h=1}^{\infty} \pi_h N \left(\begin{bmatrix} 0 \\ \mu_h \end{bmatrix}, \begin{bmatrix} 1.0 & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix} \right), \quad (11)$$

where y_i^* denotes the latent continuous variable underlying the count response $y_i, i = 1, \dots, n$. Observed and latent responses are connected by $Y_i = y_i$ if and only if $c_{y_i-1} < y_i^* \leq c_{y_i}$, where $c_{y_i} = \Phi^{-1}[F(y_i; H, \boldsymbol{\xi})]$. We consider three choices for the function $F(,)$ that appears in the definition of the cut-points. These are the Poisson, negative binomial (NB) and generalised Poisson (GP) distribution functions. We note that the data generating mechanism is not nested within any of the models we fit.

Results presented are based on 20 replicate datasets. For each dataset and for each choice of $F(,)$ we obtained 40,000 posterior samples of which we discarded the first 20,000 as burn-in. Of the remaining 20,000 samples, we retained one every 5. Furthermore, during posterior simulation we obtained samples of conditional pmfs $f(Y|X = x)$ for 25 values of x equally spaced between 0.1 and 10.90 and for offset term H equal to the mean of the sampled offset terms. For each sampled conditional pmf, we calculated the mean and 25th and 75th percentiles. We compared sampled and true values of these functionals by calculating the medians of the posteriors of the total squared errors $\{\sum_{c=1}^{25} (p_c - p_c^{r,s})^2 : s = 1, \dots, 4000, r = 1, \dots, 20\}$, where p_c denotes the true value of the parameter of interest, here the mean, 25th and 75th percentile, of the c th conditional pmf, $c = 1, \dots, 25$, and $p_c^{r,s}$ denotes the s th sampled value of the parameter of interest, $s = 1, \dots, 4000$, when fitting the model to the r th replicate dataset, $r = 1, \dots, 20$.

Table 2: Simulation study results: posterior medians of total squared errors for the three functionals of interest and the three kernels.

| | Mean | Q ₂₅ | Q ₇₅ |
|---------------------|-------|-----------------|-----------------|
| Poisson | 0.762 | 0.796 | 1.587 |
| Negative binomial | 0.513 | 0.862 | 1.428 |
| Generalised Poisson | 0.577 | 0.681 | 1.211 |

In Table 2 we present the total errors. Concerning estimation of the mean surface, the model that utilises the NB distribution function performs the best, reducing the total errors of the models that utilise the Poisson and GP distribution functions by 33% ($= 1 - 0.513/0.762$) and 11% respectively. Concerning estimation of the quantiles, the best performance is achieved utilizing the GP function. Specifically, estimation of the first quantile under the DPMM with the GP function is improved by 14% and 20% relative to the DPMMs that utilise the Poisson and NB distribution functions, while estimation of the third quantile is improved by 24% and 15%, respectively.

In Table 3 results are presented in more detail. Although the more detailed results are not always clear-cut, there are some general observations that can be made. Firstly, over the range $3 < x < 6$, where the response is over-dispersed relative to the Poisson, the DPMMs with the NB and GP distribution functions perform better than the DPMM with the Poisson distribution function. The same is true also

Table 3: Simulation study results: posterior medians of total squared errors for the three functionals of interest and the three kernels.

| | $x < 3$ | | | $3 < x < 6$ | | | $6 < x < 9$ | | |
|----|---------|-----------------|-----------------|-------------|-----------------|-----------------|-------------|-----------------|-----------------|
| | Mean | Q ₂₅ | Q ₇₅ | Mean | Q ₂₅ | Q ₇₅ | Mean | Q ₂₅ | Q ₇₅ |
| P | 0.272 | 0.238 | 0.257 | 0.081 | 0.066 | 0.194 | 0.086 | 0.125 | 0.196 |
| NB | 0.143 | 0.114 | 0.161 | 0.056 | 0.053 | 0.119 | 0.038 | 0.304 | 0.297 |
| GP | 0.150 | 0.142 | 0.183 | 0.054 | 0.049 | 0.110 | 0.037 | 0.101 | 0.106 |

| | $x > 9$ | | |
|----|---------|-----------------|-----------------|
| | Mean | Q ₂₅ | Q ₇₅ |
| P | 0.267 | 0.295 | 0.869 |
| NB | 0.236 | 0.234 | 0.755 |
| GP | 0.289 | 0.314 | 0.763 |

over the range $x < 3$ where the response is Poisson distributed. Secondly, estimation is substantially improved under the DPMM with the GP kernel over the region of the covariate space where the response is under-dispersed, $6 < x < 9$. Thirdly, over the region $x > 9$, the DPMM with the NB distribution function does better than the DPMMs with the Poisson and GP distribution functions. It is a bit surprising that over $x > 9$, the DPMM with the Poisson distribution function does better than that with the GP distribution function for estimating the mean and first quantile functions, although the differences in estimation of the mean are mostly due to the results concerning estimation of conditional pmf $f(Y|X = x)$ for $x = 10.9$, which is at the edge of the covariate space.

Further, in Figure 1 we present a simulated dataset along with plots of the estimated mean and 25th and 75th percentile curves utilizing the GP distribution function. We can see that the model fits well over all regions of the covariate space.

Figure 2 presents estimated conditional pmfs $f(y|x)$, and 95% credible intervals, for three values of covariate x , namely $x = 2.3, 3.7, 6.3$, to show model performance over regions of the covariate space where the response is equi-, over- and under-dispersed relative to the Poisson. The three rows of the figure correspond to the three values of x and the three columns to the models with the Poisson, NB and GP distribution functions. In the first row, where the response is Poisson distributed, we see that all models fit well. In the second row, where the response is over-dispersed, we see that the model that utilises the Poisson distribution function cannot adopt to the thicker tails. Lastly, in the third row, we see that only the model that utilises the GP function can adopt to the under-dispersion.

6 Application

We present an application that we adopt from Bailey et al. (2009) who examined the impact of commercial fishing on deep-sea fish populations in the northeast Atlantic. The dataset, that is available in Hilbe (2014), includes $n = 147$ observations on scientific trawls that were made in two distinct time periods, from 1977 to 1989 and from 1997 to 2002, at depths from 0.8 to 4.8 Km. The response variable y is ‘fish abundance’, a count of the number of fish caught in each of the 147 trawls. With each trawl, there is an associated offset term H that is calculated based on the size of the swept area (Km²). The two explanatory variables are x_1 , a time period indicator, and x_2 , the average depth (Km) of the trawl. The main interest here is on the effect of the time periods that are thought to reflect the effect of the development of commercial fishing. The period 1977–1989 (with 97 observations), which we refer to as




Figure 1: Simulation study results obtained from the DPMM with the generalised Poisson distribution function. The three curves show the results concerning the mean and 25th and 75th percentiles.

the ‘early’ period, is before and during the development of commercial fishing, and the period 1997–2002 (with 50 observations), which we refer to as the ‘late’ period, is considered post-commercial fishing. The dataset is displayed in Figure 3.

Commercial fishing is limited to depths of less than approximately 1.6 Km. Hence, it may be reasonable to expect fish abundance in waters deeper than 1.6 Km to be unaffected by commercial fishing. However, it is also possible that the effects of commercial fishing are transmitted to the deeper waters. In fact, Bailey et al. (2009) concluded that fish abundance reduced significantly between the two periods at all depths between 0.8 and 2.5 Km. One of the two explanations they considered was that these reductions were due to commercial fishing and its effects cascading to deeper waters. Therefore, we find it interesting to re-analyse the dataset using the methods we have described.

We let $\mathbf{z}_i = (y_i, x_{i1}, x_{i2})^\top, i = 1, \dots, n$, denote the i th vector of observed response and covariates, where the count of fish caught y_i is associated with offset term H_i , the size of the swept area. The model we fit to $\mathbf{z}_i, i = 1, \dots, n$, is a DP mixture of the form

$$f_P(y_i, x_{i1}, x_{i2}) = \sum_{h=1}^T \pi_h \int_{R(y_i)} \int_{R(x_{i1})} \mathbf{N}(\mathbf{z}_i^*; \boldsymbol{\mu}_h^*, \boldsymbol{\Sigma}_h^*) dy_i^* d\mathbf{x}_{i1}^*,$$

where $\mathbf{z}_i^* = (y_i^*, x_{i1}^*, x_{i2}^*)^\top$ is the vector of the latent and directly observed continuous variables. Further, $\mathbf{N}(\mathbf{z}_i^*; \boldsymbol{\mu}_h^*, \boldsymbol{\Sigma}_h^*)$ denotes a trivariate Gaussian density with one restriction on the mean vector and two on the covariance matrix

$$\boldsymbol{\mu}_h^* = \begin{bmatrix} 0 \\ \mu_{1h} \\ \mu_{2h} \end{bmatrix}, \quad \boldsymbol{\Sigma}_h^* = \begin{bmatrix} 1.0 & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & 1.0 & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_{33} \end{bmatrix}. \quad (12)$$

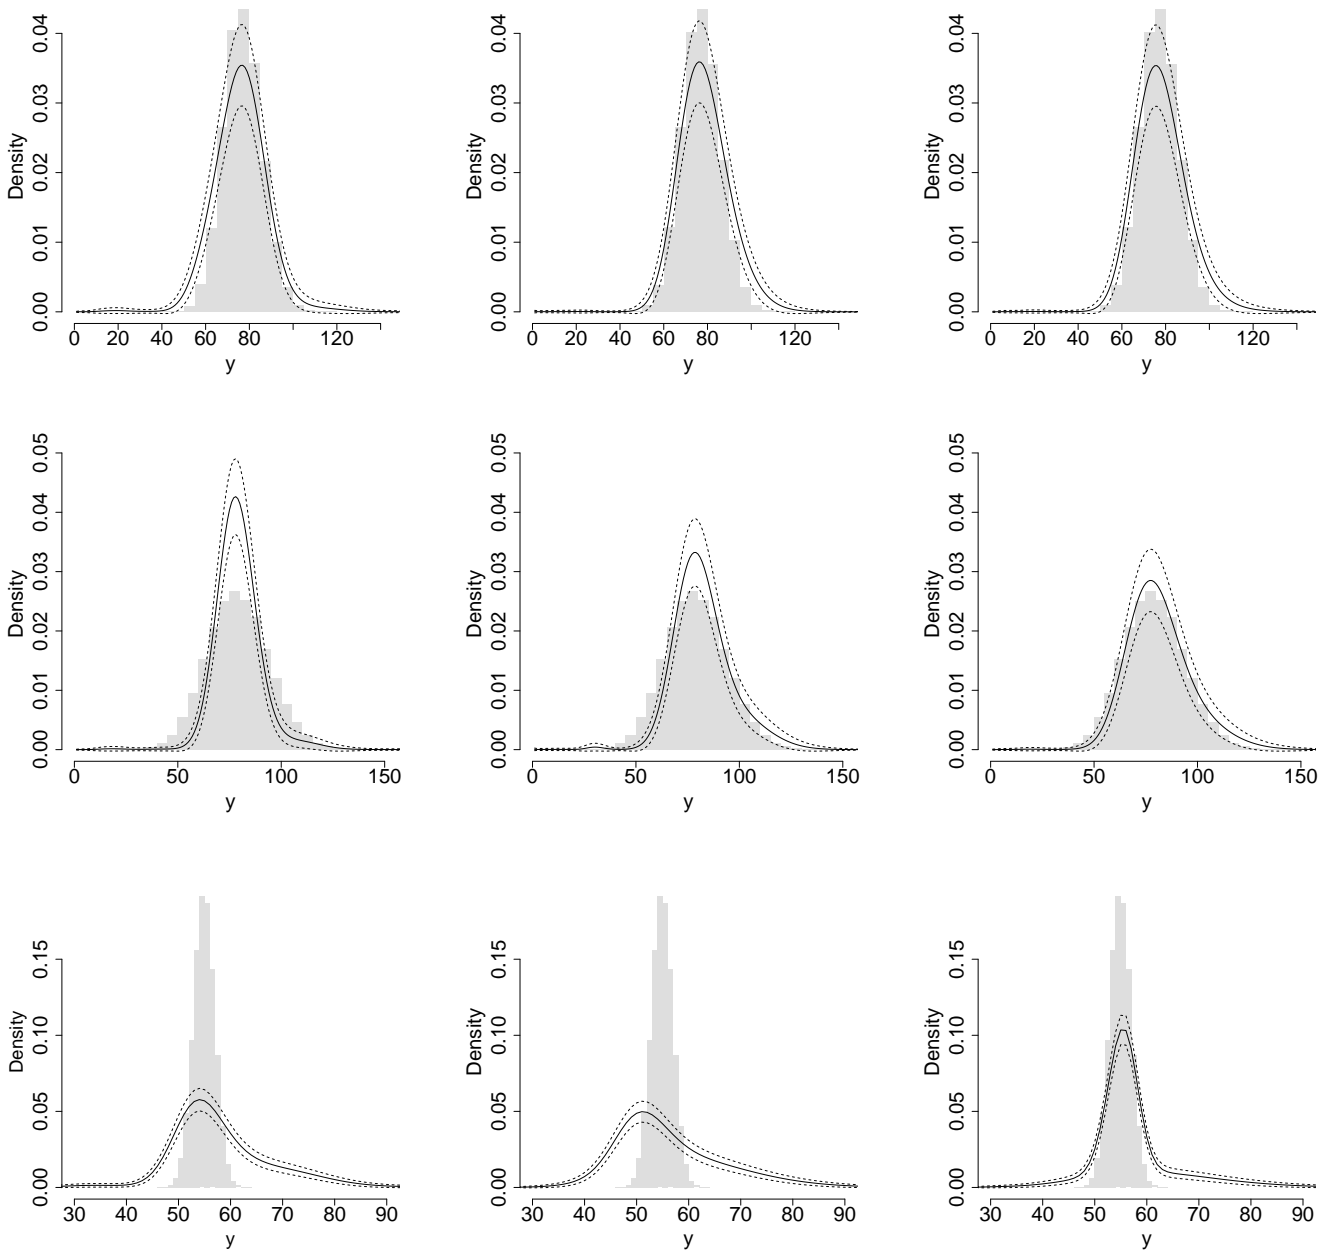


Figure 2: Estimated conditional pmfs $f(y|x)$ and 95% credible intervals. The histograms represent data generated from the true model. The rows correspond to $x = 2.3, 3.7, 6.3$ where the response is equi-, over- and under-dispersed relative to the Poisson. The three columns correspond to the models with Poisson, NB and GP distribution functions.




Figure 3: A scatter-plot of the fishing data. Standardised responses (Y/H) plotted against depth (x_2) with different symbols for the two time periods (x_1).

The two rules for connecting observed and latent variables are as follows

$$Y_i = y_i \text{ if and only if } c_{y_{i-1}} < y_i^* \leq c_{y_i}, \text{ where } c_{y_i} = \Phi^{-1}[F(y_i; H_i, \boldsymbol{\xi})],$$

$$X_{i1} = 0 \text{ if and only if } x_{i1}^* < 0,$$

where $F(\cdot, \cdot)$, that appears in the definition of the cut-points, is taken to be the distribution function of a negative binomial random variable. This choice was guided by the presence of over-dispersion and lack of under-dispersion in the response variable over the predictor space.

To fit this model, we ran the MCMC algorithm for 60,000 iterations, and retained one sample every five, after discarding the first 20,000 as burn-in. Let $f_{P_s}(y_i, x_{i1}, x_{i2})$ denote the s th sampled joint density, $s = 1, \dots, 8,000$. From it, we can compute the s th sampled conditional pmf $f_{P_s}(y_i | x_{i1}, x_{i2})$ that describes the possible values and associated probabilities for the count variable y_i , for the given values of the covariates, x_{i1}, x_{i2} , and the given value of the offset term H_i , the size of the swept area. We sampled conditional pmfs for the early ($x_1 = 0$) and late ($x_1 = 1$) periods, in combination with 25 equally spaced depths (x_2), ranging from 0.90 to 4.75 Km, and with offset term (H) fixed at the value of mean observed offset.

This procedure enables inference about how the shape of the conditional pmf changes with covariates, which is what we refer to as pmf regression. Further, for each sampled conditional $f_{P_s}(y_i | x_{i1}, x_{i2})$, we can calculate all functionals of interest. Here, we are interested in the median of the conditional pmf, and on how it changes with covariates, which is what we refer to as median regression. Other functionals, such as the mode or quantiles of interest, can also be computed. This is an important feature of the approach presented here: whereas traditional methods, such as generalised linear models, allow us to examine only how the mean changes with covariates, the current method allows for more detailed examination of the effects of the covariates on the response distribution.

Results, in terms of standardised responses Y/H , are presented in Figure 4. First, Figure 4 (a) displays the results for median regression along with 90% pointwise credible intervals. With solid lines are the

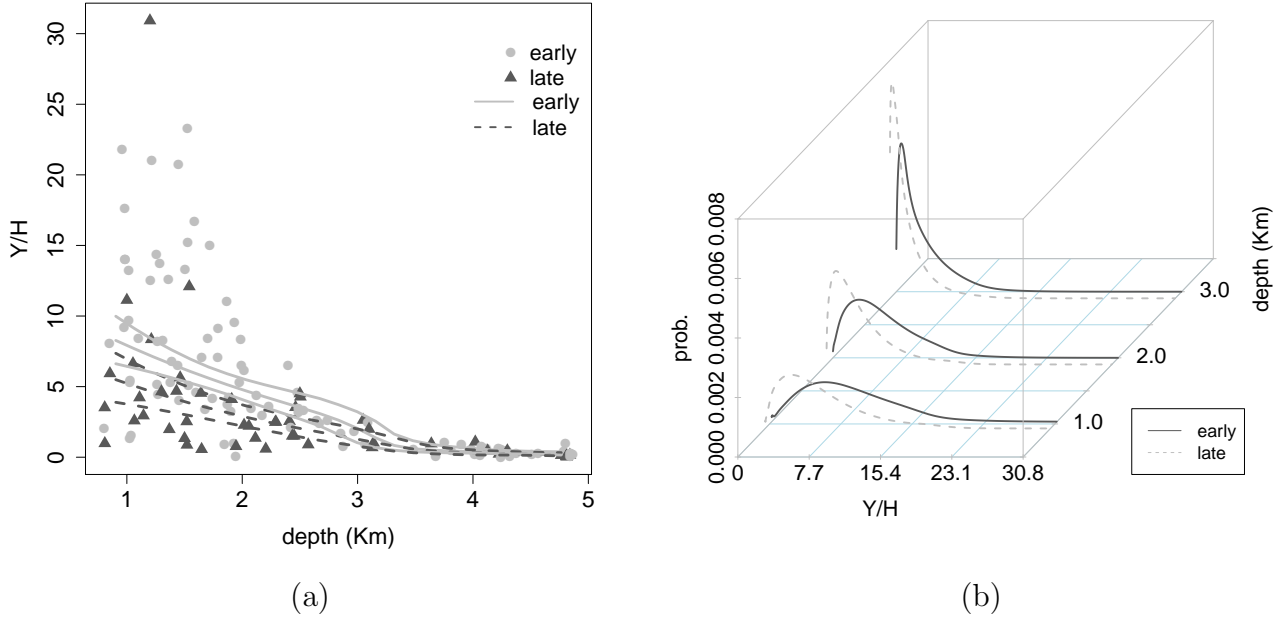


Figure 4: Results from the analysis of the fishing data. Figure (a) displays results on median regression along with 90% credible intervals. Figure (b) displays results on pmf regression.

results for the early period and with dashed lines the results for the late period. Clearly, median fish abundance decreases with depth for both periods. The median abundance for the late period is below that of the early period for all depths up to about 3.5 Km. Credible intervals do not overlap for depths between about 1.3 and 2.1 Km. Second, Figure 4 (b) displays the results for pmf regression. With solid lines are the estimated pmfs for the early period and with dashed lines the ones for the late period. At depth 1.0 Km, the pmf of the late period gives higher probability to smaller rates i.e. to smaller counts associated with offset term equal to unity. As the depth increases, the two pmfs give higher probability to smaller numbers, while the estimated pmf of the late period continues to give higher probability to smaller numbers.

7 Discussion

We have developed Bayesian models for pmf regression with emphasis on count responses. The method represents discrete variables as continuous latent variables that have been discretised and utilises Dirichlet process mixtures of Gaussians to model the joint density of the observed and latent continuous variables. The joint density forms the basis for carrying out inference on the conditional densities and its functionals.

The assumed mechanism by which latent continuous variables become observed discrete ones utilises cut-points that are expressed as $\Phi[F(\cdot)]$, where $\Phi(\cdot)$ is the distribution function of a standard normal variable and $F(\cdot)$ is an appropriate distribution function, the choice of which is made by the data analyst, depending on the needs of the particular data analysis that is being carried out. We have considered and evaluated several functions $F(\cdot)$. We have shown utilizing simulated and real data how flexible the proposed model is and the diverse types of Bayesian inferences one can obtain by utilizing this model, including pmf, mean and quantile regression. Another attractive feature of the current model is the ease by which missing data can be handled under a missing at random mechanism. See e.g. Dunson &

Bhattacharya (2011) for further details.

The method we have proposed can be computationally expensive. There are several parts of the MCMC algorithm that can contribute to that. First, obtaining samples of discrete distributions $f(y|\mathbf{x})$ over a set of covariate values \mathbf{x} , to enable inference about the dependence of the pmf on the covariates, is very computationally intensive, especially when a natural upper bound on the values of y is not present. There are additional features in the model that can make it computationally intensive. These are the restrictions that are placed on the model for the latent variables, the zero mean and unit variance, which are, of course, also present in the DP mixtures of Gaussians. Furthermore, the numerical integration over the unobserved latent variables can also be numerically intensive, depending on the number of discrete variables present in the model. Lastly, it is also computationally demanding to handle the possibly high dimensional covariance and precision matrices which are present in the model for the joint density. The MCMC algorithms for fitting the presented models, with any choice of the discussed functions $F(\cdot)$, are available in the R package `BNSP` (Papageorgiou, 2019).

8 Appendix I

Our MCMC sampler proceeds as follows

1. Update $v_h \sim \text{Beta}(n_h + 1, n - \sum_{l=1}^h n_h + \alpha)$, where n_h is the number of observations allocated in the h th cluster.
2. The joint posterior of $(\mathbf{D}_h, \mathbf{\Sigma}_h^*)$ is given by

$$p(\mathbf{D}_h, \mathbf{\Sigma}_h^* | \dots) \propto |\mathbf{D}_h|^{\eta/2-1} |\mathbf{\Sigma}_h^*|^{(\eta-q-1-n_h)/2} \exp[-\text{tr}(\mathbf{H}^{-1} \mathbf{E}_h + \mathbf{\Sigma}_h^{*-1} \mathbf{S}_h)/2],$$

where $\mathbf{S}_h = \sum_{i:\delta_i=h} (\mathbf{z}_i^* - \boldsymbol{\mu}_h^*)(\mathbf{z}_i^* - \boldsymbol{\mu}_h^*)^\top$.

As the above is a non-standard density, sampling from it requires a Metropolis-Hastings step. We take the proposal density to be $\mathbf{E}_h^{(p)} \sim \text{Wishart}(\mathbf{E}_h^{(p)}; \psi, \mathbf{E}_h^{(t)}/\psi)$, where $\mathbf{E}_h^{(t)} = \mathbf{D}_h^{(t)1/2} \mathbf{\Sigma}_h^{*(t)} \mathbf{D}_h^{(t)1/2}$ and $\mathbf{D}_h^{(t)}, \mathbf{\Sigma}_h^{*(t)}$ are realizations from the previous iteration. Proposed values for $\mathbf{D}_h, \mathbf{\Sigma}_h^*$ are obtained by decomposing $\mathbf{E}_h^{(p)} = \mathbf{D}_h^{(p)1/2} \mathbf{\Sigma}_h^{*(p)} \mathbf{D}_h^{(p)1/2}$ and they are accepted with probability

$$\alpha = \min \left\{ \frac{p(\mathbf{D}_h^{(p)}, \mathbf{\Sigma}_h^{*(p)} | \dots) t(\mathbf{D}_h^{(t)}, \mathbf{\Sigma}_h^{*(t)} | \mathbf{D}_h^{(p)}, \mathbf{\Sigma}_h^{*(p)})}{p(\mathbf{D}_h^{(t)}, \mathbf{\Sigma}_h^{*(t)} | \dots) t(\mathbf{D}_h^{(p)}, \mathbf{\Sigma}_h^{*(p)} | \mathbf{D}_h^{(t)}, \mathbf{\Sigma}_h^{*(t)})}, 1 \right\},$$

where the proposal density is given by $t(\mathbf{D}_h^{(p)}, \mathbf{\Sigma}_h^{*(p)} | \mathbf{D}_h^{(t)}, \mathbf{\Sigma}_h^{*(t)}) = \text{Wishart}(\mathbf{E}_h^{(p)}; \psi, \mathbf{E}_h^{(t)}/\psi) J(\mathbf{E}_h^{(p)} \rightarrow \mathbf{D}_h^{(p)}, \mathbf{\Sigma}_h^{*(p)})$. The free parameter ψ is chosen adaptively (Roberts & Rosenthal, 2009) so as to achieve an acceptance ratio of about 20–25% (Roberts & Rosenthal, 2001). Here the acceptance ratio that we adjust with parameter ψ is the average acceptance ratio over the non-empty clusters.

3. Next we describe how the means $\boldsymbol{\mu}_h, h \geq 1$ are updated. Recall that the q -dimensional mean $\boldsymbol{\mu}_h^*$ is restricted to have some of its elements equal to zero, see (6). These are the means that correspond to the latent variables underlying the discrete (but not the binary) variables. Denote these by \mathbf{z}_1 and their dimension by p_1 . Further, the unrestricted elements of $\boldsymbol{\mu}_h^*$, denoted by $\boldsymbol{\mu}_h$, correspond to the means of the latent variables underlying the binary and continuous variables. Denote these by \mathbf{z}_2 and their dimension by p_2 , hence $p_1 + p_2 = q$. Writing the joint pdf of $(\mathbf{z}_1^\top, \mathbf{z}_2^\top)^\top$ as

$$(\mathbf{z}_1^\top, \mathbf{z}_2^\top)^\top | (\boldsymbol{\mu}_h^*, \mathbf{\Sigma}_h^*) \sim \text{N} \left(\boldsymbol{\mu}_h^* = \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\mu}_h \end{bmatrix}, \mathbf{\Sigma}_h^* = \begin{bmatrix} \boldsymbol{\Sigma}_{11h} & \boldsymbol{\Sigma}_{12h} \\ \boldsymbol{\Sigma}_{21h} & \boldsymbol{\Sigma}_{22h} \end{bmatrix} \right),$$

it is easy to see that $\boldsymbol{\mu}_h$ are updated from $f(\boldsymbol{\mu}_h | \dots) \propto$

$$p_0(\boldsymbol{\mu}_h) \prod_{\{i:l_i=h\}} \text{N}(\mathbf{z}_{2i}; \boldsymbol{\mu}_h + \boldsymbol{\Sigma}_{21h} \boldsymbol{\Sigma}_{11h}^{-1} \mathbf{z}_{1i}, \boldsymbol{\Sigma}_{22h} - \boldsymbol{\Sigma}_{21h} \boldsymbol{\Sigma}_{11h}^{-1} \boldsymbol{\Sigma}_{12h}). \quad (13)$$

With prior $\boldsymbol{\mu}_h \sim \text{N}(\mathbf{d}, \mathbf{D})$, the updating distribution is

$$\text{N}\left\{(n_h \mathbf{W}_h^{-1} + \mathbf{D}^{-1})^{-1} \left[\mathbf{W}_h^{-1} \sum_{l_i=h} (\mathbf{z}_{2i} - \boldsymbol{\Sigma}_{21h} \boldsymbol{\Sigma}_{11h}^{-1} \mathbf{z}_{1i}) + \mathbf{D}^{-1} \mathbf{d} \right], \right. \\ \left. (n_h \mathbf{W}_h^{-1} + \mathbf{D}^{-1})^{-1} \right\},$$

where $\mathbf{W}_h = \boldsymbol{\Sigma}_{22h} - \boldsymbol{\Sigma}_{21h} \boldsymbol{\Sigma}_{11h}^{-1} \boldsymbol{\Sigma}_{12h}$.

4. We describe the step for updating $\boldsymbol{\xi}_h$, $h \geq 1$, assuming that the dataset consists of binary and continuous covariates and a count or a binomial response variable. We update $\boldsymbol{\xi}_h$, $h \geq 1$, from the marginal posterior having integrated out y^* , the latent variable that corresponds to the response

$$p(\boldsymbol{\xi}_h | \dots) \propto \prod_{\{i:\delta_i=h\}} \left(\Phi\{[c_{y_i}(\boldsymbol{\lambda}_i) - E_i^*]/\text{sd}_i^*\} - \Phi\{[c_{y_i-1}(\boldsymbol{\lambda}_i) - E_i^*]/\text{sd}_i^*\} \right) p_0(\boldsymbol{\xi}_h),$$

where $E_i^* = \text{E}(y_i^* | \mathbf{x}_{d,i}^*, \mathbf{x}_{c,i})$ is the conditional expectation and $\text{sd}_i^* = \text{sd}(y_i^* | \mathbf{x}_{d,i}^*, \mathbf{x}_{c,i})$ is the conditional standard deviation. They are obtained using standard theory on multivariate normal densities, as was done in (13). Further, priors $p_0(\boldsymbol{\xi}_h)$ are defined in Table 1.

For all cases, updating requires a Metropolis-Hastings step. We provide details next.

- (a) For Poisson mixtures, the proposed value $\xi_h^{(p)}$ is obtained from $\xi_h^{(p)} \sim \text{Gamma}(\tau \xi_h^{(t)} \xi_h^{(t)}, \tau \xi_h^{(t)})$, that is from a gamma distribution with mean $\xi_h^{(t)}$ and variance $1/\tau$, where $\xi_h^{(t)}$ denotes the current value. The acceptance probability is given by

$$\min\left\{1, \frac{p(\xi_h^{(p)} | \dots) \text{Gamma}(\xi_h^{(t)}; \tau \xi_h^{(p)} \xi_h^{(p)}, \tau \xi_h^{(p)})}{p(\xi_h^{(t)} | \dots) \text{Gamma}(\xi_h^{(p)}; \tau \xi_h^{(t)} \xi_h^{(t)}, \tau \xi_h^{(t)})}\right\}. \quad (14)$$

Here τ is introduced as a free parameter which is adjusted adaptively (Roberts & Rosenthal, 2009) in order to achieve an acceptance ratio of about 20–25% (Roberts & Rosenthal, 2001). The acceptance ratio that we adjust with parameter τ is the average acceptance ratio over the non-empty clusters.

- (b) For binomial mixtures, the proposed value $\xi_h^{(p)}$ is obtained from $\xi_h^{(p)} \sim \text{Beta}(\xi_h^{(p)}; a^{(t)}, b^{(t)})$, where $b^{(t)} = \xi_h^{(t)} - 1 + \xi_h^{(t)}(1 - \xi_h^{(t)})^2 \tau$ and $a^{(t)} = b^{(t)} \xi_h^{(t)} / (1 - \xi_h^{(t)})$ that define a beta distribution with mean $\xi_h^{(t)}$ and variance $1/\tau$. The expression of the acceptance probability follows along the same lines as (14) and hence omitted.
- (c) For negative binomial mixtures, parameter vector $\boldsymbol{\xi}_h = (\xi_{1h}, \xi_{2h})^\top$ is updated in a single step. Proposed values for the elements of $\boldsymbol{\xi}_h^{(p)} = (\xi_{1h}^{(p)}, \xi_{2h}^{(p)})^\top$ are obtained from independent gamma distributions similar in form to the gamma distribution shown in part 4a for Poisson mixtures. We utilise a common tuning parameter τ in the two gamma proposal distributions.
- (d) For beta-binomial mixtures, the elements of $\boldsymbol{\xi}_h = (\xi_{1h}, \xi_{2h})^\top$ are also updated in a single step. Proposed values for $\boldsymbol{\xi}_h^{(p)} = (\xi_{1h}^{(p)}, \xi_{2h}^{(p)})^\top$ are obtained from independent gamma distributions with a single tuning parameter, as was done in part 4c.

(e) For generalised Poisson mixtures, parameter vector $\boldsymbol{\xi}_h = (\xi_{1h}, \xi_{2h})^\top$ is updated in two steps utilizing two tuning parameters, τ_1 and τ_2 . Proposed values for the mean ξ_{1h} are obtained from $\xi_{1h}^{(p)} \sim \text{Gamma}(\tau_1 \xi_h^{(t)}, \tau_1 \xi_h^{(t)})$ and those for the variance ξ_{2h} from $\xi_{2h}^{(p)} \sim \text{N}(\xi_{2h}^{(t)}, 1/\tau_2)$, a normal distribution centered at the previous realization and with variance $1/\tau_2$.

5. We impute the latent variables $y_i^*, \mathbf{x}_{d,i}^* = (x_{d,i,1}^*, \dots, x_{d,i,p_d}^*)^\top, i = 1, \dots, n$, from the conditional

$$\begin{aligned} (y_i^*, \mathbf{x}_{d,i}^{*\top})^\top | \mathbf{x}_{c,i} &\sim \text{N}(y_i^*, \mathbf{x}_{d,i}^* | \mathbf{x}_{c,i}; \boldsymbol{\mu}_{h,d,c,i}, \boldsymbol{\Sigma}_{h,d,c}) \mathbf{1}[y_i^* \in R(y_i)] \\ &\times \prod_{m=1}^{p_d} \mathbf{1}[x_{d,i,m}^* \in R(x_{d,i,m})], \end{aligned}$$

where $\boldsymbol{\mu}_{h,d,c,i}$ and $\boldsymbol{\Sigma}_{h,d,c}$ were defined after (10). The imputation utilises the algorithm of Robert (2009) according to which imputation is done one variable at a time given all other ones. Here with subscript h we denote the cluster in which sampling unit i is allocated.

6. We update the cluster allocation variables $\delta_i, i = 1, 2, \dots, n$, according to probabilities $\text{Pr}(\delta_i = h)$ obtained from the marginalised posterior

$$k_{ih} \text{N}(\mathbf{x}_{c,i}; \boldsymbol{\mu}_{2h}, \boldsymbol{\Sigma}_{22h}^*) \int \dots \int \text{N}(y_i^*, \mathbf{x}_{d,i}^* | \mathbf{x}_{c,i}; \boldsymbol{\mu}_{h,d,c,i}, \boldsymbol{\Sigma}_{h,d,c}) dy_i^* d\mathbf{x}_{d,i}^*,$$

where $k_{ih} = \mathbf{1}(u_i < \pi_h)$ for the slice sampler while $k_{ih} = \pi_h$ for the truncated sampler.

7. Label switching moves (Papaspiliopoulos & Roberts, 2008):

- (a) Choose randomly two nonempty clusters, a and b say, and propose to exchange their labels. The acceptance probability of this move is $\min\{1, (\pi_b/\pi_a)^{n_a - n_b}\}$. If the proposed move is accepted, we exchange allocation variables and cluster specific parameters.
- (b) Choose randomly a cluster, a say, and propose to exchange the labels of clusters a and $a + 1$, and at the same time propose to exchange v_a with v_{a+1} . Cluster a is chosen randomly among clusters labelled $1, \dots, n^* - 1$, where n^* is the nonempty cluster with the largest label. The acceptance probability of this move is $\min\{1, (1 - v_{a+1})^{n_a} / (1 - v_a)^{n_{a+1}}\}$, and if it is accepted, we exchange allocation variables and cluster specific parameters.

8. We update concentration parameter α using the method described by Escobar & West (1995). Assuming a $\text{Gamma}(\alpha|a, b)$ prior (mean = a/b), the posterior is expressed as a mixture of two gamma distributions:

$$\alpha \sim \pi_\eta \text{Gamma}(a + k, b - \log(\eta)) + (1 - \pi_\eta) \text{Gamma}(a + k - 1, b - \log(\eta)), \quad (15)$$

where k is the number of non-empty clusters, $\pi_\eta = (a + k - 1) / \{a + k - 1 + n[b - \log(\eta)]\}$ and

$$\eta | \alpha, k \sim \text{Beta}(\alpha + 1, n). \quad (16)$$

Hence the algorithm proceeds as follows: with α and k fixed at their current values, we sample η from (16). Then, based on the same k and the value of η , we sample a new α value from (15).

9 Appendix II

We start by constructing a density f_{p_ϵ} such that $\sum_y \int_{\mathbf{x}} f_0(y, \mathbf{x}) \log[f_0(y, \mathbf{x})/f_{p_\epsilon}(y, \mathbf{x})] d\mathbf{x} < \epsilon$ for any $\epsilon > 0$.
Let

$$dP_m(\boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\nu}, \boldsymbol{\Sigma}) = \delta_{\boldsymbol{\nu}}(\mathbf{0}) \delta_{\boldsymbol{\Sigma}}(h_m^2 \mathbf{I}) f_0^*(\boldsymbol{\xi}, \boldsymbol{\mu}),$$

where $h_m = m^{-\eta}$ for some $\eta > 0$ and $\delta_x(A) = \mathbf{1}(x \in A)$.

Hence, by utilizing the kernel in (7), we may write

$$f_{p_m}(y, \mathbf{x}) = \int k(\mathbf{z}; \boldsymbol{\theta}) dP_m(\boldsymbol{\theta}) = \int \int \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, h_m^2 \mathbf{I}) K(y; \boldsymbol{\xi}) f_0^*(\boldsymbol{\xi}, \boldsymbol{\mu}) d\boldsymbol{\mu} d\boldsymbol{\xi},$$

where $K(y; \boldsymbol{\xi})$ is the chosen model e.g. the Poisson, negative binomial or generalised Poisson pmf for count data.

Now, by utilizing the transformation $\mathbf{a} = (\mathbf{x} - \boldsymbol{\mu})/h_m$, we obtain

$$f_{p_m}(y, \mathbf{x}) = \int \int \mathbf{N}(\mathbf{a}; \mathbf{0}, \mathbf{I}) K(y; \boldsymbol{\xi}) f_0^*(\boldsymbol{\xi}, \mathbf{x} - h_m \mathbf{a}) d\mathbf{a} d\boldsymbol{\xi}.$$

By the continuity of f_0^* , we have that $f_0^*(\boldsymbol{\xi}, \mathbf{x} - h_m \mathbf{a}) \rightarrow f_0^*(\boldsymbol{\xi}, \mathbf{x})$ as $h_m \rightarrow 0$. Further, recalling that by condition C_2 , f_0^* is bounded, by the dominated convergence theorem we have that

$$f_{p_m}(y, \mathbf{x}) \rightarrow \int K(y; \boldsymbol{\xi}) f_0^*(\boldsymbol{\xi}, \mathbf{x}) d\boldsymbol{\xi} = f_0(y, \mathbf{x}),$$

where the last equality follows from condition C_1 . Therefore, as $m \rightarrow \infty$, $\log[f_0(y, \mathbf{x})/f_{p_m}(y, \mathbf{x})] \rightarrow 0$ for all y and \mathbf{x} . To show that

$$\sum_y \int_{\mathbf{x}} f_0(y, \mathbf{x}) \log[f_0(y, \mathbf{x})/f_{p_m}(y, \mathbf{x})] d\mathbf{x} \rightarrow 0, \quad (17)$$

we need to find a function that dominates $|\log[f_0(y, \mathbf{x})/f_{p_m}(y, \mathbf{x})]|$ and that is f_0 -integrable.

To this end, first observe that due to condition C_2 , $f_{p_m}(y, \mathbf{x})$ is bounded from above by

$$f_{p_m}(y, \mathbf{x}) \leq M \int \int \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, h_m^2 \mathbf{I}) K(y; \boldsymbol{\xi}) d\boldsymbol{\mu} d\boldsymbol{\xi} \leq M.$$

It follows that

$$\log \frac{f_0(y, \mathbf{x})}{f_{p_m}(y, \mathbf{x})} \geq \log \frac{f_0(y, \mathbf{x})}{M}. \quad (18)$$

Further, for $\|\mathbf{x}\| > m$ and any y , from C_1 and C_4 , we have

$$\begin{aligned} f_{p_m}(y, \mathbf{x}) &\geq \int \int_{\|\boldsymbol{\mu}\| < \|\mathbf{x}\|} \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, h_m^2 \mathbf{I}) K(y; \boldsymbol{\xi}) f_0^*(\boldsymbol{\xi}, \boldsymbol{\mu}) d\boldsymbol{\mu} d\boldsymbol{\xi} \\ &\geq f_0(y, \mathbf{x}) \int_{\|\boldsymbol{\mu}\| < \|\mathbf{x}\|} \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, h_m^2 \mathbf{I}) d\boldsymbol{\mu} \geq f_0(y, \mathbf{x})/3, \end{aligned} \quad (19)$$

where the last inequality follows by a suitable choice of h_m^2 .

Furthermore, for $\|\mathbf{x}\| \leq m$ and any y , from C_1 and C_3 , we have

$$\begin{aligned} f_{p_m}(y, \mathbf{x}) &\geq \int \int_{\|\boldsymbol{\mu}\| < m} \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, h_m^2 \mathbf{I}) K(y; \boldsymbol{\xi}) f_0^*(\boldsymbol{\mu} | \boldsymbol{\xi}) f_0^*(\boldsymbol{\xi}) d\boldsymbol{\mu} d\boldsymbol{\xi} \\ &\geq c f_0(y) \int_{\|\boldsymbol{\mu}\| < m} \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, h_m^2 \mathbf{I}) d\boldsymbol{\mu} \geq c f_0(y)/3. \end{aligned} \quad (20)$$

Combining (19) and (20)

$$f_{p_m}(y, \mathbf{x}) \geq \begin{cases} f_0(y, \mathbf{x})/3, & \|\mathbf{x}\| > m, \\ c f_0(y)/3, & \|\mathbf{x}\| \leq m, \end{cases}$$

from which follows that

$$\log \frac{f_0(y, \mathbf{x})}{f_{p_m}(y, \mathbf{x})} \leq \xi(y, \mathbf{x}) = \begin{cases} \log 3, & \|\mathbf{x}\| > m, \\ \log\{[3f_0(y, \mathbf{x})]/[c f_0(y)]\}, & \|\mathbf{x}\| \leq m. \end{cases} \quad (21)$$

Now, from (18) and (21)

$$\left| \log \frac{f_0(y, \mathbf{x})}{f_{p_m}(y, \mathbf{x})} \right| \leq \max \left\{ \xi(y, \mathbf{x}), \left| \log \frac{f_0(y, \mathbf{x})}{M} \right| \right\},$$

where the right-hand side is f_0 -integrable due to C_5 , and it follows that (17) holds.

For any given $\epsilon > 0$, condition A_1 is satisfied by f_{P_m} with suitable choice of m . Hence, we take $f_{P_\epsilon} = f_{P_m}$.

Further, to show that condition A_2 is satisfied, observe that

$$c = \inf_{\mathbf{z} \in \mathcal{Z}} \inf_{\boldsymbol{\theta} \in \Theta} k(\mathbf{z}; \boldsymbol{\theta}) > 0, \quad (22)$$

where \mathcal{Z} denotes the sample space.

In addition, note that the family of maps $\{\boldsymbol{\theta} \rightarrow k(\mathbf{z}; \boldsymbol{\theta}) : \mathbf{z} \in \mathcal{Z}\}$ is uniformly equicontinuous over compact space Θ . To see this, write

$$\begin{aligned} k(\mathbf{z}; \boldsymbol{\theta}) &= \int_{R(y)} \mathbf{N}(\mathbf{z}^*; \boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*) dy^* = \\ &\mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) \int_{R(y)} \mathbf{N}(y^*; m^*, v^*) dy^* = \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) p(y; \mathbf{x}, \boldsymbol{\theta}), \end{aligned}$$

where m^* and v^* were defined below (7). Now, $|k(\mathbf{z}; \boldsymbol{\theta}) - k(\mathbf{z}; \boldsymbol{\theta}')|$ is expressed as

$$|\mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) p(y; \mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) - \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}', \boldsymbol{\Sigma}') p(y; \mathbf{x}, \boldsymbol{\xi}', \boldsymbol{\mu}', \boldsymbol{\Sigma}')| \quad (23)$$

$$\begin{aligned} &\leq p(y; \mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) |\mathbf{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) - \mathbf{N}(\mathbf{x}; \boldsymbol{\mu}', \boldsymbol{\Sigma}')| + \\ &\mathbf{N}(\mathbf{x}; \boldsymbol{\mu}', \boldsymbol{\Sigma}') |p(y; \mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) - p(y; \mathbf{x}, \boldsymbol{\xi}', \boldsymbol{\mu}, \boldsymbol{\Sigma})| + \\ &\mathbf{N}(\mathbf{x}; \boldsymbol{\mu}', \boldsymbol{\Sigma}') |p(y; \mathbf{x}, \boldsymbol{\xi}', \boldsymbol{\mu}, \boldsymbol{\Sigma}) - p(y; \mathbf{x}, \boldsymbol{\xi}', \boldsymbol{\mu}', \boldsymbol{\Sigma}')|. \end{aligned} \quad (24)$$

Due to the equicontinuity of the multivariate normal pdf (Wu & Ghosal, 2008; Canale & De Blasi, 2017) the first and last terms in the right-hand side (24) can be made arbitrarily small for all $\mathbf{z} \in \mathcal{Z}$.

Furthermore, the middle term can be made arbitrarily small because of the following expression for the difference of the probabilities

$$|p(y; \mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) - p(y; \mathbf{x}, \boldsymbol{\xi}', \boldsymbol{\mu}, \boldsymbol{\Sigma})| = \left| \int_{\Phi^{-1}(F\{y-1; \boldsymbol{\xi}\})}^{\Phi^{-1}(F\{y; \boldsymbol{\xi}\})} \mathbf{N}(y^*; m^*, v^*) dy^* - \int_{\Phi^{-1}(F\{y-1; \boldsymbol{\xi}'\})}^{\Phi^{-1}(F\{y; \boldsymbol{\xi}'\})} \mathbf{N}(y^*; m^*, v^*) dy^* \right|,$$

and the equicontinuity of the cut-point function $\Phi^{-1}(F\{y; \boldsymbol{\xi}\})$, viewed as a function of $\boldsymbol{\xi}$.

Hence, for any $\delta > 0$, there exist $\mathbf{z}_1, \dots, \mathbf{z}_m$, such that for any $\mathbf{z} \in \mathcal{Z}$

$$\sup_{\boldsymbol{\theta} \in \Theta} |k(\mathbf{z}; \boldsymbol{\theta}) - k(\mathbf{z}_i; \boldsymbol{\theta})|, \quad (25)$$

for some $\mathbf{z}_i, i = 1, \dots, m$.

Let

$$\mathcal{U} = \{P : \left| \int_{\Theta} k(\mathbf{z}_i; \boldsymbol{\theta}) dP_m(\boldsymbol{\theta}) - \int_{\Theta} k(\mathbf{z}_i; \boldsymbol{\theta}) dP(\boldsymbol{\theta}) \right| < c\delta, i = 1, \dots, m\}.$$

It follows that \mathcal{U} is a weak neighbourhood of P_m with $\Pi(\mathcal{U}) > 0$.

Now, for some $P \in \mathcal{U}$ and any $\mathbf{z} \in \mathcal{Z}$ we have that

$$\begin{aligned} |f_p(\mathbf{z}) - f_{p_m}(\mathbf{z})| &= \left| \int k(\mathbf{z}; \boldsymbol{\theta}) dP(\boldsymbol{\theta}) - \int k(\mathbf{z}; \boldsymbol{\theta}) dP_m(\boldsymbol{\theta}) \right| \\ &= \left| \int k(\mathbf{z}; \boldsymbol{\theta}) dP(\boldsymbol{\theta}) \pm \int k(\mathbf{z}_i; \boldsymbol{\theta}) dP(\boldsymbol{\theta}) \right. \\ &\quad \left. \pm \int k(\mathbf{z}_i; \boldsymbol{\theta}) dP_m(\boldsymbol{\theta}) - \int k(\mathbf{z}; \boldsymbol{\theta}) dP_m(\boldsymbol{\theta}) \right|, \end{aligned} \quad (26)$$

where \mathbf{z}_i is chosen from (25) and with $\pm a$ we mean add and subtract a .

It follows that the expression (26) is $\leq 3c\delta$. Further, recalling (22), from which follows that $f_{p_m}(\mathbf{z}) = \int k(\mathbf{z}; \boldsymbol{\theta}) dP_m(\boldsymbol{\theta}) > c$, and dividing both sides of (26) by f_{p_m} , we obtain

$$\left| \frac{f_{p_m}(\boldsymbol{\theta})}{f_p(\boldsymbol{\theta})} - 1 \right| \leq \frac{3\delta}{1 - 3\delta}.$$

Hence, condition A_2 is satisfied for any $P \in \mathcal{U}$ as

$$\sum_y \int f_0(y, \mathbf{x}) \log \frac{f_{P_m}(y, \mathbf{x})}{f_P(y, \mathbf{x})} d\mathbf{x} < \frac{3\delta}{1 - 3\delta}.$$

This completes the proof.

References

- BAILEY, D., COLLINS, M., GORDON, J., ZUUR, A. & PRIEDE, I. (2009). Long-term changes in deep-water fish populations in the northeast atlantic: a deeper reaching effect of fisheries? *Proceedings of the Royal Society of London B: Biological Sciences* **276**, 1965–1969.
- BARNARD, J., MCCULLOCH, R. & MENG, X.L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* **10**, 1281–1311.
- CANALE, A. & DE BLASI, P. (2017). Posterior asymptotics of nonparametric location-scale mixtures for multivariate density estimation. *Bernoulli* **23**, 379–404.
- CANALE, A. & DUNSON, D.B. (2011). Bayesian kernel mixtures for counts. *Journal of the American Statistical Association* **106**, 1528–1539.
- CANALE, A. & DUNSON, D.B. (2015). Bayesian multivariate mixed-scale density estimation. *Statistics and its Interface* **8**, 195–201.
- CHUNG, Y. & DUNSON, D.B. (2009). Nonparametric bayes conditional distribution modeling with variable selection. *Journal of the American Statistical Association* **104**, 1646–1660.
- CONSUL, P.C. & FAMOYE, F. (1992). Generalized Poisson regression model. *Communications in Statistics - Theory and Methods* **21**, 89–109.
- DEYOREO, M. & KOTTAS, A. (2015). A fully nonparametric modelling approach to binary regression. *Bayesian Analysis* **10**, 821–847.
- DEYOREO, M. & KOTTAS, A. (2018). Bayesian nonparametric modeling for multivariate ordinal regression. *Journal of Computational and Graphical Statistics* **27**, 71–84.
- DUNSON, D.B. & BHATTACHARYA, A. (2011). Nonparametric bayes regression and classification through mixtures of product kernels. In *Bayesian Statistics 9, Proceedings of the Ninth Valencia International Conference on Bayesian Statistics*, eds. J. Bernardo, M. Bayarri, J. Berger, D. A.P., D. Heckerman, A. Smith & M. West. Oxford University Press, pp. 145–164.
- DUNSON, D.B., PILLAI, N. & PARK, J.H. (2007). Bayesian density regression. *Journal of the Royal Statistical Society: Series B* **69**, 163–183.
- ESCOBAR, M.D. & 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.
- GHOSH, J.K. & RAMAMOORTHY, R. (2003). *Bayesian Nonparametrics*. New York: Springer-Verlag.
- HANNAH, L.A., BLEI, D.M. & POWELL, W.B. (2011). Dirichlet process mixtures of generalized linear models. *Journal of Machine Learning Research* **12**, 1923–1953.
- HILBE, J.M. (2014). *COUNT: Functions, data and code for count data*. URL <http://CRAN.R-project.org/package=COUNT>. R package version 1.3.2.

- ISHWARAN, H. & JAMES, L. (2001). Gibbs sampling methods for stick breaking priors. *Journal of the American Statistical Association* **96**, 161–173.
- KOTTAS, A., MÜLLER, P. & QUINTANA, F. (2005). Nonparametric Bayesian modeling for multivariate ordinal data. *Journal of Computational and Graphical Statistics* **14**, 610–625.
- KUNIHAMA, T., HALPERN, C.T. & HERRING, A.H. (2019). Non-parametric Bayes models for mixed scale longitudinal surveys. *Journal of the Royal Statistical Society: Series C* .
- LO, A.Y. (1984). On a class of Bayesian nonparametric estimates: I. Density estimates. *The Annals of Statistics* **12**, 351–357.
- MÜLLER, P., ERKANLI, A. & WEST, M. (1996). Bayesian curve fitting using multivariate normal mixtures. *Biometrika* **83**, 67–79.
- MUTHEN, B. (1984). A general structural equation model with dichotomous, ordered categorical, and continuous latent variable indicators. *Psychometrika* **49**, 115–132.
- NORETS, A. & PELENIS, J. (2012). Bayesian modeling of joint and conditional distributions. *Journal of Econometrics* **168**, 332–346.
- PAPAGEORGIU, G. (2019). *BNSP: Bayesian Non- And Semi-Parametric Model Fitting*. URL <https://CRAN.R-project.org/package=BNSP>. R package version 2.1.0.
- PAPAGEORGIU, G., RICHARDSON, S. & BEST, N. (2015). Bayesian non-parametric models for spatially indexed data of mixed type. *Journal of the Royal Statistical Society: Series B* **77**, 973–999.
- PAPASPILIOPOULOS, O. (2008). A note on posterior sampling from Dirichlet mixture models. *Technical report*, University of Warwick.
- PAPASPILIOPOULOS, O. & ROBERTS, G.O. (2008). Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models. *Biometrika* **95**, 169–186.
- RICHARDSON, S. & GREEN, P. (1997). On Bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society: Series B* **59**, 731–792.
- ROBERT, C.P. (2009). Simulation of truncated normal variables. *Statistics and Computing* **5**, 121–125.
- ROBERT, C.P. & CASELLA, G. (2005). *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- ROBERTS, G.O. & ROSENTHAL, J.S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science* **16**, 351–367.
- ROBERTS, G.O. & ROSENTHAL, J.S. (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics* **18**, 349–367.
- SCHWARTZ, L. (1965). On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **4**, 10–26.
- SETHURAMAN, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* **4**, 639–650.

- SHAHBABA, B. & NEAL, R.M. (2009). Nonlinear models using Dirichlet process mixtures. *Journal of Machine Learning Research* **10**, 1829–1850.
- TADDY, M.A. & KOTTAS, A. (2010). A Bayesian nonparametric approach to inference for quantile regression. *Journal of Business & Economic Statistics* **28**, 357–369.
- WALKER, S.G. (2007). Sampling the Dirichlet mixture model with slices. *Communications in Statistics - Simulation and Computation* **36**, 45–54.
- WU, Y. & GHOSAL, S. (2008). Kullback Leibler property of kernel mixture priors in Bayesian density estimation. *Electronic Journal of Statistics* **2**, 298–331.
- ZHANG, X., BOSCARDIN, J.W. & BELIN, T.R. (2006). Sampling correlation matrices in Bayesian models with correlated latent variables. *Journal of Computational & Graphical Statistics* **15**, 880–896.