

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

MCMC Computations for Bayesian Mixture Models Using Repulsive Point Processes

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

MCMC Computations for Bayesian Mixture Models Using Repulsive Point Processes / Beraha, Mario; Argiento, Raffaele; Møller, Jesper; Guglielmi, Alessandra. - In: JOURNAL OF COMPUTATIONAL AND GRAPHICAL STATISTICS. - ISSN 1061-8600. - ELETTRONICO. - 31:2(2022), pp. 422-435. [10.1080/10618600.2021.2000424]

Availability:

This version is available at: <https://hdl.handle.net/11585/850159> since: 2022-01-31

Published:

DOI: <http://doi.org/10.1080/10618600.2021.2000424>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

**Mario Beraha, Raffaele Argiento, Jesper Møller & Alessandra Guglielmi (2022)
MCMC Computations for Bayesian Mixture Models Using Repulsive Point Processes,
Journal of Computational and Graphical Statistics, 31:2, 422-435**

The final published version is available online at
<https://dx.doi.org/10.1080/10618600.2021.2000424>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

MCMC computations for Bayesian mixture models using repulsive point processes

Mario Beraha^{*†}, Raffaele Argiento[‡], Jesper Møller[§], Alessandra Guglielmi^{*}

18 April 2021

Abstract

Repulsive mixture models have recently gained popularity for Bayesian cluster detection. Compared to more traditional mixture models, repulsive mixture models produce a smaller number of well separated clusters. The most commonly used methods for posterior inference either require to fix a priori the number of components or are based on reversible jump MCMC computation. We present a general framework for mixture models, when the prior of the ‘cluster centres’ is a finite repulsive point process depending on a hyperparameter, specified by a density which may depend on an intractable normalizing constant. By investigating the posterior characterization of this class of mixture models, we derive a MCMC algorithm which avoids the well-known difficulties associated to reversible jump MCMC computation. In particular, we use an ancillary variable method, which eliminates the problem of having intractable normalizing constants in the Hastings ratio. The ancillary variable method relies on a perfect simulation algorithm, and we demonstrate this is fast because the number of components is typically small. In several simulation studies and an application on sociological data, we illustrate the advantage of our new methodology over existing methods, and we compare the use of a determinantal or a repulsive Gibbs point process prior model.

Keywords: birth-death Metropolis-Hastings algorithm, cluster estimation, pairwise interaction point process, intractable normalizing constant, normalized infinitely divisible distribution, perfect simulation.

1 Introduction

Mixture models are useful when partitioning observations $\mathbf{y} = (y_1, \dots, y_n)$ into groups/clusters as well as when approximating densities that are not otherwise modelled by standard parametric distributions; see Fruhwirth-Schnatter et al. (2019) and the references therein. Our approach to finite mixture models extends that in Argiento and De Iorio (2019), but the focus in this paper will be on prior specification and Bayesian MCMC computations when the aim is cluster

^{*}Department of Mathematics, Politecnico di Milano, Milano, Italy

[†]Department of Computer Science, Università di Bologna, Bologna, Italy

[‡]Department of Statistical Sciences, Università Cattolica del Sacro Cuore, Milano, Milano, Italy

[§]Department of Mathematical Sciences, Aalborg University, Aalborg, Denmark

Jesper Møller was supported by the ‘Danish Council for Independent Research – Natural Sciences’ grant DFF – 7014-00074 ‘Statistics for point processes in space and beyond’

detection. Our objective is partly to present a general framework for mixture models based on repulsive point process priors for ‘cluster centres’, arguing why this is useful, and partly to derive a MCMC algorithm which avoids the well-known difficulties associated to reversible jump MCMC computation. In several simulation studies and an application on sociological data, we illustrate the advantage of our new methodology over existing methods, and we compare the use of the different repulsive point process priors. Moreover, when introducing a hyperparameter in such priors, we demonstrate that perfect simulation is fast in connection to a useful ancillary variable method.

1.1 Setting

For specificity, assume each $y_i \in \mathbb{R}^q$ with $q \geq 1$. It will always be obvious from the context whether we consider y_i (and other variables considered later on) as a random variable, a realization, or an argument of a function. Denote $\{k(\cdot | \tau) : \tau \in \Theta\}$ a parametric family of densities (with respect to Lebesgue measure on \mathbb{R}^q or counting measure on a countable subset of \mathbb{R}^q), where the specification of the parameter space Θ is application dependent. We refer to this parametric family as the kernel (of the mixture model). We consider each y_i to follow a mixture of these densities: Let $\boldsymbol{\tau} = (\tau_1, \dots, \tau_m)$ specify m densities where $\tau_h \in \Theta$ for $h = 1, \dots, m$, and let $\boldsymbol{w} = (w_1, \dots, w_m)$ specify weights $w_h > 0$ with $\sum_{h=1}^m w_h = 1$. We assume that $\boldsymbol{\tau}, \boldsymbol{w}$, and $m \leq +\infty$ are random. The case where m is a fixed positive integer may be considered as a special case. Note that m is determined by $\boldsymbol{\tau}$ as well as by \boldsymbol{w} . Conditioned on $(\boldsymbol{w}, \boldsymbol{\tau})$, the observations are assumed to be independent identically distributed (iid) with a distribution given by the following mixture density:

$$y_i \stackrel{\text{iid}}{\sim} \sum_{h=1}^m w_h k(\cdot | \tau_h), \quad i = 1, \dots, n. \quad (1)$$

The densities $k(\cdot | \tau_h)$, $h = 1, \dots, m$ are usually referred to as the ‘components’ of the mixture. In this context, cluster detection means estimating allocation parameters $\boldsymbol{c} = (c_1, \dots, c_n) \in \{1, \dots, m\}^n$ where the sets $\{y_i : c_i = h\}$, $h = 1, \dots, m$ are the clusters. The number of clusters in the mixture model is the number of allocated components in (1), i.e., the number of unique values in (c_1, \dots, c_n) .

We make prior assumptions as follows. To control the number of clusters, m is random and finite; the case $m = +\infty$ would be relevant for nonparametric inference (Müller and Mitra, 2013), but this context is not addressed in this paper. Only when $m < +\infty$ is not fixed, it can be consistently estimated, cf. Argiento and De Iorio (2019) and Miller and Harrison (2018). We let $\tau_h = (\mu_h, \gamma_h)$, thinking of μ_h as a continuous random parameter in \mathbb{R}^q which specifies a ‘cluster centre’ of cluster h , and of γ_h as a positive random parameter ($q = 1$) or a continuous covariance matrix ($q \geq 2$) (or, in simple settings, a fixed positive number) which specifies the amount of dispersion of the data points in cluster h (for example, $k(\cdot | \tau_h)$ could be a normal density with mean μ_h and variance γ_h). To make posterior inference more robust, we add a hyperparameter ξ to the prior distribution of (μ_1, \dots, μ_m) . Furthermore, since the mixture density in (1) does not depend on the order of the components, we can assume that

- (a) the conditional marginal prior density $p(\mu_1, \dots, \mu_m | \xi, m)$ is exchangeable,

that is, for any fixed integer $m \geq 1$, it is invariant under permutations of μ_1, \dots, μ_m . Note that $\boldsymbol{\mu} = \{\mu_1, \dots, \mu_m\}$ is then a finite point process, specifying both the random number m

of components and the locations of the cluster centres. Finally, a priori we make conditional independence assumptions: Conditioned on (ξ, m) , we have that

- (b) (w_1, \dots, w_m) , (μ_1, \dots, μ_m) , and $(\gamma_1, \dots, \gamma_m)$ are a priori independent,
- (c) given m , the conditional marginal prior distribution of (w_1, \dots, w_m) does not depend on ξ ,
- (d) the γ_h 's are iid, with a prior distribution which does not depend on (ξ, m) ,

and conditioned on $(\xi, m, w_1, \dots, w_m, \mu_1, \dots, \mu_m, \gamma_1, \dots, \gamma_m)$, we have that

- (e) the c_i 's are iid with a prior distribution given by $P(c_i = h | \mathbf{w}) = w_h$.

Hence, the random parameter here consists of $(\xi, \{\mu_1, \dots, \mu_m\}, w_1, \dots, w_m, \gamma_1, \dots, \gamma_m, c_1, \dots, c_n)$. By Bayes' theorem, using the generic notation $p(\cdot)$ for a density and $p(\cdot | \cdot)$ for a conditional density, the posterior density becomes

$$p(\xi, \{\mu_1, \dots, \mu_m\}, w_1, \dots, w_m, \gamma_1, \dots, \gamma_m, c_1, \dots, c_n | y_1, \dots, y_n) \propto p(\xi)p(m | \xi)p(\mu_1, \dots, \mu_m | \xi, m)p(w_1, \dots, w_m | m) \left[\prod_{h=1}^m p(\gamma_h) \right] \left[\prod_{i=1}^n w_{c_i} k(y_i | (\mu_{c_i}, \gamma_{c_i})) \right]. \quad (2)$$

The dominating measure for (2) is given in Section 4 which contains measure theoretical details; see Section 2 for further prior specifications. In brief, the prior specification of $\boldsymbol{\mu}$ and \mathbf{w} requires particular attention, whilst for the prior specification of the remaining parameters we use a standard setting, following Fraley and Raftery (2007).

1.2 Previous work on repulsive mixture models

The most used mixture models assume the μ_h 's are iid and independent of m , cf. Fruhwirth-Schnatter et al. (2019). This assumption, although convenient for mathematical tractability, is often an oversimplification and might produce misleading results in producing too many clusters. This motivated Petralia et al. (2012), Xu et al. (2016), Fúquene et al. (2019), Quinlan et al. (2020), Bianchini et al. (2020), and Xie and Xu (2019) to explicitly define prior models with repulsion between the locations, thereby obtaining well separated components.

In Petralia et al. (2012), Fúquene et al. (2019), and Quinlan et al. (2020), m is finite and fixed, but, as mentioned before, this cannot guarantee posterior consistency of the number of components. However, Xu et al. (2016), Bianchini et al. (2020), and Xie and Xu (2019) assumed m to be finite and random. In particular, Xu et al. (2016) and Bianchini et al. (2020) dealt with determinantal point process (DPP) priors for $\boldsymbol{\mu}$. A DPP density has to be approximated as described in Lavancier et al. (2015) where the calculation will increase exponentially fast as the dimension q increases, or as described in Bardenet and Titsias (2015) but at the price that the model parameters are hard to interpret. Xie and Xu (2019) had no hyperparameter ξ in their repulsive mixture model, which used as a prior for $\boldsymbol{\mu}$ conditioned on m , a tempered repulsive pairwise interaction point process density of the form

$$p(\mu_1, \dots, \mu_m | m) = \frac{1}{Z_m} \left[\prod_{i=1}^m \phi_1(\mu_i) \right] \left[\prod_{1 \leq i < j \leq m} \phi_2(\|\mu_i - \mu_j\|)^{1/m} \right]$$

with respect to m -fold Lebesgue measure on \mathbb{R}^q . Here, $\|\cdot\|$ denotes usual distance, ϕ_1 is a non-negative function, $0 \leq \phi_2 \leq 1$ is a non-decreasing function (this implies repulsiveness), and Z_m is the normalizing constant. Note that if $\phi_2(\cdot) = 1$, then μ_1, \dots, μ_m are iid and independent of m . Apart from this case, Z_m is intractable and has to be approximated by numerical methods, a non-trivial task which limits both efficiency and feasibility as q increases.

As far as posterior simulation is concerned, Xu et al. (2016) and Bianchini et al. (2020) proposed to simulate $(\mathbf{w}, \boldsymbol{\tau})$ using a reversible jump MCMC algorithm, cf. Green (1995). At every iteration of this algorithm, either a split move (in which one component is killed and two new ones are created, hence increasing the dimension by one), or a combine move (in which two components are merged into a single one, hence decreasing the dimension by one) is proposed. As discussed in Green (2010), Richardson and Green (1997), and Dellaportas and Papageorgiou (2006), in order to obtain good mixing properties of the reversible jump MCMC algorithm, it is crucial to define appropriate proposal distributions that generate the new values in the split move. In general, this is a complex task that depends heavily on the kernel under consideration.

Similarly to how Miller and Harrison (2018) studied a classical mixture model, Xie and Xu (2019) considered in the observation model (1) to marginalize with respect to a prior of $(\mathbf{w}, \boldsymbol{\tau})$ and derived a ‘marginal MCMC algorithm’. However, although this algorithm compared with the reversible jump MCMC algorithm has smaller auto-correlations for the number of clusters, it requires the calculation of the normalizing constants Z_1, Z_2, \dots up to some truncation, and inference is limited to the number of clusters and the posterior mean of the mixture density.

1.3 Our contribution and outline

We discuss a general framework for mixture models based on repulsive point process priors for ‘cluster centres’ $\boldsymbol{\mu}$ and derive a new MCMC algorithm avoiding the problem with reversible jump MCMC computation.

Our first contribution is the proposal of the prior of $\boldsymbol{\mu}$ conditioned on ξ , cf. item (a) in Section 1.1: We consider a general setting with a repulsive finite point process density, including the case of a DPP (any DPP except the special case of a Poisson process is repulsive) or a density specified by an unnormalized density, e.g. a pairwise interaction point process density, which involves a normalizing constant Z_ξ which in general (except the special case of a Poisson process) is intractable. As a particular simple example of a pairwise interaction point process, we assume a Strauss process (defined later in Section 2.1). Note that the prior distributions for $\boldsymbol{\mu}$ in all the papers cited in Section 1.2 can all be considered as special cases of our prior for $\boldsymbol{\mu}$. Notice also that Z_ξ will never appear in our posterior simulation algorithm.

The second contribution is the algorithm for posterior simulation from our model. This contribution builds upon Argiento and De Iorio (2019) and is mainly based on two assumptions, namely \mathbf{w} and $\boldsymbol{\mu}$ are chosen a priori independent and the mixture weights \mathbf{w} are defined by *normalization* of iid infinitely divisible random variables, i.e. \mathbf{w} follows a *normalized infinitely divisible* distribution (Favaro et al., 2011). Indeed, Argiento and De Iorio (2019) introduced the class of *normalized independent point processes mixture* models and showed that this class can be framed into the Bayesian nonparametric context. In this way, several ideas and algorithms developed in the nonparametric literature for normalized random measures with independent increments (NRMI – see Regazzini et al., 2003) can be adapted to the finite dimensional case. Here we extend Argiento and De Iorio (2019) building a Metropolis-within Gibbs sampler, referred to as *conditional Gibbs sampler* in the Bayesian nonparametric literature, see Papaspiliopou-

los and Roberts (2008). In particular, we relax the usual assumption of μ_1, \dots, μ_m being iid and independent of m , still being able to propose a transformation of $\boldsymbol{\mu}$ into allocated cluster centres $\boldsymbol{\mu}^{(a)} = \{\mu_{c_i} : i = 1, \dots, n\}$ and non-allocated cluster centres $\boldsymbol{\mu}^{(na)} = \boldsymbol{\mu} \setminus \boldsymbol{\mu}^{(a)}$. This allow us to simulate from the full conditional of $\boldsymbol{\mu}$ without resorting to the split and combine moves of the reversible jump MCMC algorithm as used in Xu et al. (2016) and Bianchini et al. (2020). In fact posterior updates of $\boldsymbol{\mu}^{(a)}$ becomes easy and when updating $\boldsymbol{\mu}^{(na)}$ we use the Metropolis-Hasting birth-death algorithm in Geyer and Møller (1994). The Metropolis-Hasting birth-death algorithm has the advantage that the choice of the kernel does not impact on the acceptance rate of the algorithm.

We impose the hyperprior on ξ , the parameter in the repulsive point process prior controlling the intensity of the point process, to make posterior inference more robust, cf. Section 1.1, unlike previous literature (apart from Bianchini et al., 2020). When making posterior updates of ξ in our MCMC algorithm, if the prior density for $\boldsymbol{\mu}$ conditioned on ξ has an intractable normalizing constant Z_ξ , we get rid of Z_ξ mentioned by using the single exchange algorithm in Murray et al. (2006) coming from the ancillary variable algorithm in Møller et al. (2006). These algorithms require perfect simulation of an ancillary variable following the same distribution as $\boldsymbol{\mu}$ conditioned on ξ . Interestingly, perfect simulation is feasible in our context because m will typically be small (in our examples it is effectively always less than 10).

The remainder of this paper is organized as follows. Sections 2 and 3 specify our further prior assumptions on the cluster centers $\boldsymbol{\mu}$ and the mixture weights \boldsymbol{w} , respectively. Section 4 derive the posterior density, using the useful superposition of $\boldsymbol{\mu}$ mentioned above, and provides the technical details needed when dealing with point process densities (we aim at keeping this as simple as possible). Section 5 details our Metropolis-within-Gibbs sampler for posterior simulation. Sections 6.1 and 6.2 discuss prior elicitation when the prior for $\boldsymbol{\mu}$ is the Strauss process and the DPP (conditioned on $\{m \geq 1\}$, see Section 2). Section 7 presents various simulation studies comparing posterior inference and MCMC mixing obtained using reversible jump or our Metropolis-within-Gibbs sampler, and using a DPP, a Strauss process, or a non-repulsive prior for $\boldsymbol{\mu}$. Furthermore, an application to a sociological data set is discussed in Section 8. The article concludes with a discussion in Section 9. In the Appendix we provide practical details on our Metropolis-within-Gibbs sampler, collects additional simulation studies, including an illustration on the advantages of using a Strauss process over a DPP as prior for $\boldsymbol{\mu}$, and discusses possible extensions. The code for posterior simulation has been implemented in C++ and linked to Python.

2 Prior specification of $\boldsymbol{\mu}$

For the prior specification of $\boldsymbol{\mu}$, introduced in item (a) in Section 1.1, which is the first original contribution of our work, a few technical details are needed to start. Let $\Omega = \cup_{m=0}^{\infty} \Omega_m$ denote the space of all finite subsets (point configurations) of \mathbb{R}^q , where Ω_m denotes the space of all finite subsets of cardinality m , with $\Omega_0 = \{\emptyset\}$, where \emptyset denotes the empty point configuration; although we cannot have 0 groups, it becomes convenient in Section 4 to include Ω_0 into the definition of Ω . We equip each Ω_m with the smallest σ -algebra making the mapping of pairwise distinct $(\mu_1, \dots, \mu_m) \in \mathbb{R}^{qm}$ into $\{\mu_1, \dots, \mu_m\} \in \Omega_m$ measurable. The σ -algebra on Ω is the smallest σ -algebra containing the union of the σ -algebras on each Ω_m . Then $\boldsymbol{\mu}$ is absolutely continuous with respect to a measure on Ω which, with an abuse of notation, is denoted $d\boldsymbol{\mu}$ and

defined as follows. For sets $B = \cup_{m=0}^{\infty} B_m$ with $B_m \subseteq \Omega_m$,

$$\int_B d\boldsymbol{\mu} = \sum_{m=0}^{\infty} \frac{1}{m!} \int_{B_m} d\boldsymbol{\mu}_m,$$

where the notation means the following. For $m = 0$, we interpret the term in the sum as $[\emptyset \in A]$. We set $\boldsymbol{\mu}_m = \{\mu_1, \dots, \mu_m\}$ and, with an abuse of notation, write $d\boldsymbol{\mu}_m$ for Lebesgue measure $d\mu_1 \cdots d\mu_m$ on \mathbb{R}^{qm} . Further, we write $\int_{B_m} d\boldsymbol{\mu}_m$ for $\int_{\mathbb{R}^{qm}} \mathbb{I}[\boldsymbol{\mu} \in B_m] d\boldsymbol{\mu}_m$, where $\mathbb{I}[\cdot]$ denotes the indicator function. Then, conditioned on ξ , the density of $\boldsymbol{\mu}$ with respect to $d\boldsymbol{\mu}$ is given by

$$p(\boldsymbol{\mu} | \xi) = p(m | \xi) p(\mu_1, \dots, \mu_m | \xi, m), \quad \boldsymbol{\mu} = \{\mu_1, \dots, \mu_m\} \in \Omega, \quad m \geq 1,$$

setting $p(\mu_1, \dots, \mu_m | \xi, m) = 0$ if $m = 0$. This means that we consider the prior process prior restricted to the event that $\boldsymbol{\mu}$ is non-empty.

2.1 Repulsive pairwise-interaction point process priors

When incorporating repulsiveness in the prior density $p(\boldsymbol{\mu} | \xi)$, we suggest a repulsive pairwise-interaction point process. This is a popular class of models in statistical physics and spatial statistics; see Møller and Waagepetersen (2004) and the references therein. The repulsive pairwise-interaction density is of the form

$$p(\boldsymbol{\mu} | \xi) = \frac{1}{Z_\xi} \left[\prod_{h=1}^m \phi_1(\mu_h | \xi) \right] \left[\prod_{1 \leq i < j \leq m} \phi_2(\|\mu_i - \mu_j\| | \xi) \right], \quad (3)$$

where $\phi_1(\cdot | \xi) \geq 0$ is an integrable function, $0 \leq \phi_2(\cdot | \xi) \leq 1$ is a non-decreasing function, and Z_ξ is a normalizing constant. Note that $Z_\xi < +\infty$, but in general Z_ξ is intractable. An exception is the special case $\phi_2(\cdot | \xi) = 1$ (a Poisson process with intensity function $\phi_1(\cdot | \xi)$ and conditioned on not being empty), where $Z_\xi = 1 - \exp(-\int \phi_1(\mu_h | \xi) d\mu_h)$.

For simplicity and specificity, in Sections 7-8, we follow Bianchini et al. (2020) in letting ξ be a positive random variable and using an empirical Bayesian approach with

$$\phi_1(\mu_h | \xi) = \xi \mathbb{I}[\mu_h \in R], \quad (4)$$

where $R \subset \mathbb{R}^q$ is the smallest rectangular region containing the data \mathbf{y} and with sides parallel to the usual axes in \mathbb{R}^q (they advocate the use of this choice over other more complicated situations).

The simplest non-trivial case is a Strauss prior,

$$\phi_2(r | \xi) = \alpha^{\mathbb{I}[r \leq \delta]}, \quad (5)$$

so that ξ enters only in the expression of ϕ_1 . Here, $\delta > 0$ is a fixed parameter, specifying the range of interaction, and $0 \leq \alpha \leq 1$ is a fixed interaction parameter. Note that, if $\alpha = 0$, we set $0^0 = 1$ and obtain a so-called hard core point process. If $\alpha = 1$, we obtain a model with no interaction which is like a Poisson process except that we condition on that $\boldsymbol{\mu}$ is non-empty.

2.2 Repulsive priors specified by an unnormalized density

In the following we consider a general prior model given by

$$p(\boldsymbol{\mu} | \xi) = \frac{1}{Z_\xi} g(\boldsymbol{\mu} | \xi), \quad (6)$$

where $g(\cdot | \xi)$ is a so-called unnormalized density, meaning that $g(\cdot | \xi)$ is a non-negative measurable function such that the normalizing constant Z_ξ is finite. Note that by assumption $g(\emptyset | \xi) = 0$. Specific examples of (6) can be found in Møller and Waagepetersen (2004) and the references therein. In our simulation study and application example (Sections 7-8) we focus on the Strauss prior and a specific DPP prior given below, but considering (6) is useful in order to give a general exposition of our methodology.

To describe interaction in the general model (6), one possibility is to assume that for any $\boldsymbol{\mu} \in \Omega$ and $\mu^* \in \mathbb{R}^q \setminus \boldsymbol{\mu}$ we have $g(\boldsymbol{\mu} \cup \{\mu^*\} | \xi) > 0 \Rightarrow g(\boldsymbol{\mu} | \xi) > 0$, and then consider the so-called Papangelou conditional intensity defined by

$$\lambda(\mu^*, \boldsymbol{\mu} | \xi) := g(\boldsymbol{\mu} \cup \{\mu^*\} | \xi) / g(\boldsymbol{\mu} | \xi)$$

(taking $0/0 := 0$). Then we have repulsiveness if $\lambda(\mu^*, \boldsymbol{\mu} | \psi)$ is a non-increasing function of $\boldsymbol{\mu}$, that is, $\lambda(\mu^*, \boldsymbol{\mu} | \psi) \geq \lambda(\mu^*, \boldsymbol{\mu} \cup \{\mu'\} | \psi)$ for any $\mu' \in \mathbb{R}^q \setminus \boldsymbol{\mu} \cup \{\mu^*\}$, where inequality can not be replaced by an identity. Clearly, this is true for (3) when $\phi_2(\cdot | \xi) \neq 1$.

2.3 Determinantal point process priors

A DPP density (conditioned on that the DPP is non-empty) is a special case of (6) but with repulsion characterized in another way than above (Hough et al., 2009; Lavancier et al., 2015; Biscio et al., 2016; Møller and O'Reilly, 2021). To work with a DPP density, we consider a compact region $R \subset \mathbb{R}^q$ with $\int_R dx > 0$, and a complex covariance function $C : R \times R \mapsto \mathbb{C}$ with a spectral representation

$$C(x, x' | \xi) = \sum_{i=1}^{\infty} \lambda_i \varphi_i(x) \overline{\varphi_i(x')}, \quad x, x' \in R, \quad (7)$$

where the φ_i 's form an orthonormal basis for the $L^2(R)$ -space of complex functions defined on R , each $\lambda_i \geq 0$, and $\sum_{i=1}^{\infty} \lambda_i < +\infty$. Then existence of the DPP is equivalent to that all $\lambda_i \leq 1$, cf. Macchi (1975). Note that we suppress in the notation that the eigenvalues λ_i 's and the eigenfunctions φ_i 's may depend on ξ .

A special case of a DPP occurs when C is a projection of finite rank m , let us say

$$C(x, x' | \xi) = \sum_{i=1}^m \varphi_i(x) \overline{\varphi_i(x')}, \quad x, x' \in R.$$

From (7) we obtain a density

$$p(\mu_1, \dots, \mu_m | \xi) = \det\{C(\mu_h, \mu'_h)\}_{h,h'=1,\dots,m} \quad \text{for } \mu_1, \dots, \mu_m \in R, \quad (8)$$

where $\det\{C(\mu_h, \mu'_h)\}_{h,h'=1,\dots,m}$ is the determinant of the $m \times m$ matrix $\{C(\mu_h, \mu'_h)\}_{h,h'=1,\dots,m}$. A point process with density (8) is called a projection DPP with kernel C . Note that it consists of exactly m points in R .

The general construction of a DPP is given by introducing a random projection

$$K(x, x' | \xi) = \sum_{i=1}^{\infty} B_i \varphi_i(x) \overline{\varphi_i(x')}, \quad (9)$$

where B_1, B_2, \dots are independent Bernoulli variables with means $\lambda_1, \lambda_2, \dots$, respectively. Then a DPP with kernel C is a finite point process on R which conditioned on B_1, B_2, \dots is a projection DPP with kernel K ; it can be shown that the distribution of this DPP depends only on C , cf. Hough et al. (2006). Note that $\sum_{i=1}^{\infty} B_i$ is the random number of points. In particular, assuming all $\lambda_i < 1$ and defining C' as C in (7) but with each λ_i replaced by $\lambda'_i = \lambda_i / (1 - \lambda_i)$, the DPP has unnormalized density

$$g(\boldsymbol{\mu} | \xi) = \det\{C'(\mu_h, \mu_{h'})\}_{h, h'=1, \dots, m} \quad \text{for pairwise distinct } \boldsymbol{\mu} = \{\mu_1, \dots, \mu_m\} \subset R, m \geq 1. \quad (10)$$

Most DPP densities are specified as in (10) with the kernel coming from a parametric family of (often real) covariance functions with all eigenvalues < 1 , see Lavancier et al. (2015). The advantage of using such models is that we can avoid including the Bernoulli variables as ancillary variables in the posterior, whilst the problem is to find a spectral representation. Note that when we condition on that the DPP is non-empty, the normalizing constant is given by

$$Z_\xi = \prod_{i=1}^{\infty} (1 - \lambda_i)^{-1} - 1. \quad (11)$$

For our purpose it is easiest to let R be rectangular and use a spectral approach with Fourier basis functions for the eigenfunctions, cf. Lavancier et al. (2015). In Sections 7-8, we follow Bianchini et al. (2020) in making this choice of eigenfunctions and letting $R = [-\frac{1}{2}, \frac{1}{2}]^q$ and

$$C(x, x' | \xi) = \sum_{j \in \mathbb{Z}^q} \lambda_j \cos(2\pi j \cdot (x - x')), \quad (12)$$

where \mathbb{Z} is the set of integers, \cdot denotes the usual inner product on \mathbb{R}^q , and $\lambda_j = \chi(j)$ is specified by the spectral density χ of the power exponential spectral model from Lavancier et al. (2015). Specifically,

$$\lambda_j = \xi \frac{\alpha^q \Gamma(q/2 + 1)}{\pi^{q/2} \Gamma(q/\beta + 1)} \exp(-\|\alpha j\|^\nu), \quad (13)$$

where $\Gamma(\cdot)$ is the gamma function, α and β are fixed positive parameters, and λ_j depends on the parameter $\xi > 0$ so that $\lambda_j \leq 1$ and $\sum_{j \in \mathbb{Z}^q} \lambda_j < \infty$. For details, see Lavancier et al. (2015), noting that ξ is the intensity of the DPP if we do not condition on that the DPP is non-empty.

When dealing with computations, in the sum of (12) and in the corresponding product $\prod_{j \in \mathbb{Z}^q} \dots$ for the normalizing constant, cf. (11), we may replace the infinite lattice \mathbb{Z}^q with a finite set, which is most naturally given by $\{-N, -N+1, \dots, 0, \dots, N-1, N\}^q$, where $N > 0$ is an integer. Then $m \leq (2N+1)^q$; in Bianchini et al. (2020), $N = 50$ for $q = 1, 2$. Bardenet and Titsias (2015) suggested an alternative approach, which does not require the spectral approach used above but specifies the DPP density directly by (10) and exploits certain bounds for the product in (11). However, it is then harder to interpret the parameters, and in particular to work with an intensity parameter.

In order to fix the values of α and β , we could follow Lavancier et al. (2015) who proposed to approximate some summaries such as the pair correlation function which depends only on (α, β) . Instead, in Section 6, we discuss an empirical Bayesian approach to select hyperparameters and hyperpriors for both the Strauss process given by (3)-(5) and the DPP given by (12)-(13).

3 Normalized infinite divisible prior for the weights

When deriving full conditional distributions for our Metropolis-within-Gibbs sampler given in Section 5, it becomes convenient to introduce ancillary variables t and u as specified below.

Conditioned on m , let $\mathbf{s} = (s_1, \dots, s_m)$ consists of iid positive continuous random variables, with the distribution of each s_h not depending on m , and with \mathbf{s} independent of $(\xi, \boldsymbol{\mu}, \boldsymbol{\gamma})$. Set $t = \sum_{h=1}^m s_h$ and $\mathbf{w} = (s_1/t, \dots, s_m/t)$, so \mathbf{s} and (\mathbf{w}, t) are in a one-to-one correspondence. In particular, in Sections 7-8, we assume each s_h follows a gamma distribution, in which case our model can be referred to as a *finite Dirichlet mixture model with repulsive locations*. We point out that the idea of building the weights \mathbf{w} by normalization not only has computational advantages – as discussed in Section 1.3 – but it also allows us to embed the model into the large class of mixtures obtained by normalization of finite point processes (Argiento and De Iorio, 2019). This latter class, to be defined, requires only the distribution of s_h 's to be infinitely divisible, and it is the finite-dimensional counterpart of the normalized random measures with independent increments, which has been thoroughly investigated in the last two decades in the Bayesian nonparametric literature (see, for instance, Regazzini et al., 2003; James et al., 2009; Lijoi and Prünster, 2010). The weights \mathbf{w} resulting from a finite normalization have distribution on the simplex that is denominated as *normalized infinite divisible* following Favaro et al. (2011). It is worth underlining that our Metropolis-within-Gibbs sampler, cf. Section 5, works also for normalized infinite divisible priors different from the Dirichlet distribution, such as those introduced in Argiento and De Iorio (2019).

One of the advantages of building the distribution of the weights \mathbf{w} by normalization is that computations are easier. The main idea is to consider a gamma random variable v with scale parameter one and shape parameter n , which is independent of $(\xi, \boldsymbol{\mu}, \boldsymbol{\gamma}, \mathbf{s}, c)$. Then, we set the ancillary variable $u := v/t$. It is immediate to show that u is well defined, i.e., it has a density with respect to Lebesgue measure given by

$$p(u) = \frac{u^{n-1}}{\Gamma(n)} \int_0^\infty t^n e^{-ut} p(t) dt$$

where $p(\cdot)$ in the integral is the density function of t . We show below (see (18)) that, conditioned on u , the full conditional of the unnormalized weights s_h 's factorize (i.e., the weights are conditionally independent), so that simulation will be drastically simplified. We notice that the *trick* of introducing the ancillary variable u is familiar in the context of normalized completely random measure as mixing measures for mixture models. It was studied in James et al. (2009) in the infinite dimensional case and largely exploited by Argiento and De Iorio (2019) and Argiento et al. (2016) in the finite dimensional setting.

4 Posterior distribution and a useful decomposition of $\boldsymbol{\mu}$

To specify the posterior obtained by considering all parameters introduced so far, including (\mathbf{s}, t, u) , we first notice that the dominating measure implicitly used in (2) leads to a new dominating measure ν given as follows. Let Ξ and Γ denote the spaces where ξ and each γ_h take values, respectively, equipped with some appropriate σ -algebras and measures $d\xi$ and $d\gamma_h$ (typically Borel σ -algebras and Lebesgue measures). For $m = 1, 2, \dots$, set $\mathbf{s}_m = (s_1, \dots, s_m)$ and $\boldsymbol{\gamma}_m = (\gamma_1, \dots, \gamma_m)$, let $d\mathbf{s}_m$ denote Lebesgue measure on \mathbb{R}_+^m , let $d\boldsymbol{\gamma}_m$ denote the product measure $\prod_{h=1}^m d\gamma_h$, and consider arbitrary measurable subsets $A \subseteq \Xi$, $B_m \subseteq \Omega_m$, $C_m \subseteq \mathbb{R}_+^m$,

$D_m \subseteq \Gamma^m$, $E_m \subseteq \{1, \dots, m\}^n$, and $F \subseteq \mathbb{R}_+$ (we still let the σ -algebra of Ω_m be induced by the Borel σ -algebra of \mathbb{R}^{qm} and the mapping $\mathbb{R}^{qm} \ni (\mu_1, \dots, \mu_m) \mapsto \{\mu_1, \dots, \mu_m\} \in \Omega_m$ with μ_1, \dots, μ_m pairwise distinct). Then the measure $d\boldsymbol{\mu}$ together with the other reference measures lead to

$$\begin{aligned} \nu(A \times \{\cup_{m=0}^{\infty} B_m \times C_m \times D_m \times E_m\} \times F) \\ = \int_A d\xi \sum_{m=0}^{\infty} \frac{1}{m!} \int_{B_m} d\boldsymbol{\mu}_m \int_{C_m} d\mathbf{s}_m \int_{D_m} d\boldsymbol{\gamma}_m \sum_{c_1, \dots, c_n=1}^m \mathbb{I}[c \in E_m] \int_F du. \end{aligned} \quad (14)$$

The posterior density of the new parameter $(\xi, \boldsymbol{\mu}, \mathbf{s}, \boldsymbol{\gamma}, \mathbf{c}, u)$ with respect to ν is then

$$p(\xi, \boldsymbol{\mu}, \mathbf{s}, \boldsymbol{\gamma}, \mathbf{c}, u | \mathbf{y}) \propto p(\xi) p(\boldsymbol{\mu} | \xi) \left[\prod_{h=1}^m p(\gamma_h) p(s_h) \right] p(u | t) \frac{1}{t^n} \left[\prod_{i=1}^n s_{c_i} k(y_i | (\mu_{c_i}, \gamma_{c_i})) \right]. \quad (15)$$

In the algorithm for posterior simulation presented in Section 5, we find it useful to split $\boldsymbol{\mu}$ into those cluster centres which are used to allocate the data, and those which are not, that is, $\boldsymbol{\mu}^{(a)} = \{\mu_{c_1}, \dots, \mu_{c_n}\}$ and $\boldsymbol{\mu}^{(na)} = \boldsymbol{\mu} \setminus \boldsymbol{\mu}^{(a)}$. For the points of these processes we use the notation $\boldsymbol{\mu}^{(a)} = \{\mu_1^{(a)}, \dots, \mu_k^{(a)}\}$ and $\boldsymbol{\mu}^{(na)} = \{\mu_1^{(na)}, \dots, \mu_\ell^{(na)}\}$. Note that $1 \leq k \leq m$, $\ell \geq 0$, and the product measure $d\boldsymbol{\mu} \times d\boldsymbol{\mu}$ on $\Omega \times \Omega$ lifted by the map $(\mathbf{x}, \mathbf{z}) \mapsto \mathbf{x} \cup \mathbf{z}$ results in the measure $d\boldsymbol{\mu}$. Hence, $(\boldsymbol{\mu}^{(a)}, \boldsymbol{\mu}^{(na)})$ conditioned on ξ has density

$$p(\boldsymbol{\mu}^{(a)}, \boldsymbol{\mu}^{(na)} | \xi) = p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi)$$

with respect to the product measure $d\boldsymbol{\mu}^{(a)} \times d\boldsymbol{\mu}^{(na)}$ (thinking of the measures $d\boldsymbol{\mu}, d\boldsymbol{\mu}^{(a)}, d\boldsymbol{\mu}^{(na)}$ as being identical but of course not thinking of $\boldsymbol{\mu}, \boldsymbol{\mu}^{(a)}, \boldsymbol{\mu}^{(na)}$ as being identical).

Obviously, $(\boldsymbol{\mu}, \mathbf{s}, \boldsymbol{\gamma}, \mathbf{c})$ and $(\boldsymbol{\mu}^{(a)}, \mathbf{s}^{(a)}, \boldsymbol{\gamma}^{(a)}, \boldsymbol{\mu}^{(na)}, \mathbf{s}^{(na)}, \boldsymbol{\gamma}^{(na)}, \mathbf{c})$ are in a one-to-one correspondence, and the cardinalities of the point processes $\boldsymbol{\mu}^{(a)}$ and $\boldsymbol{\mu}^{(na)}$ satisfy $1 \leq k < +\infty$ and $0 \leq \ell < +\infty$. Finally, setting $n_h = \#\{i : c_i = h\}$, we obtain from (14) and (15) the posterior density

$$\begin{aligned} p(\xi, \boldsymbol{\mu}^{(a)}, \mathbf{s}^{(a)}, \boldsymbol{\gamma}^{(a)}, \mathbf{c}, \boldsymbol{\mu}^{(na)}, \mathbf{s}^{(na)}, \boldsymbol{\gamma}^{(na)}, u | \mathbf{y}) \propto \\ p(\xi) p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi) \left[\prod_{h=1}^k p(\gamma_h^{(a)}) p(s_h^{(a)}) (s_h^{(a)})^{n_h} \prod_{i:c_i=h} k(y_i | (\mu_h^{(a)}, \gamma_h^{(a)})) \right] \\ \times \left[\prod_{h=1}^{\ell} p(\gamma_h^{(na)}) p(s_h^{(na)}) \right] p(u | t) \frac{1}{t^n} \end{aligned} \quad (16)$$

with respect to a new dominating measure defined by (using an obvious notation)

$$\begin{aligned} \nu'(A \times \left\{ \cup_{k=1}^{\infty} B_k^{(a)} \times C_k^{(a)} \times D_k^{(a)} \times E_k^{(a)} \right\} \times \left\{ \cup_{\ell=0}^{\infty} B_\ell^{(na)} \times C_\ell^{(na)} \times D_\ell^{(na)} \right\} \times F) \\ = \int_A d\xi \sum_{k=1}^{\infty} \frac{1}{k!} \int_{B_k^{(a)}} d\boldsymbol{\mu}_k^{(a)} \int_{C_k^{(a)}} d\mathbf{s}_k^{(a)} \int_{D_k^{(a)}} d\boldsymbol{\gamma}_k^{(a)} \sum_{\substack{c_1, \dots, c_n=1: \\ \#\{c_1, \dots, c_n\}=k}}^k \mathbb{I}[c \in E_k^{(a)}] \\ \times \sum_{\ell=0}^{\infty} \frac{1}{\ell!} \int_{B_\ell^{(na)}} d\boldsymbol{\mu}_\ell^{(na)} \int_{C_\ell^{(na)}} d\mathbf{s}_\ell^{(na)} \int_{D_\ell^{(na)}} d\boldsymbol{\gamma}_\ell^{(na)} \int_F du. \end{aligned} \quad (17)$$

Without introducing the ancillary variable u , that is, when leaving out the term $p(u|t)$ in (16)), it becomes difficult to derive the full conditionals for the allocated and non-allocated variables $\boldsymbol{\mu}^{(a)}, \mathbf{s}^{(a)}, \boldsymbol{\gamma}^{(a)}, \boldsymbol{\mu}^{(na)}, \mathbf{s}^{(na)}, \boldsymbol{\gamma}^{(na)}$. This is due to the term $1/t^n$ in (16), noting that $t = \sum_{h=1}^k s_h^{(a)} + \sum_{h=1}^\ell s_h^{(na)}$, which makes it impossible to factorize according to the allocated and non-allocated variables. When including u we obtain that

$$p(u|t) \frac{1}{t^n} = \frac{u^{n-1}}{(n-1)!} \exp(-ut) t^n \frac{1}{t^n} = \frac{u^{n-1}}{(n-1)!} \left[\prod_{h=1}^k \exp(-us_h^{(a)}) \right] \left[\prod_{h=1}^\ell \exp(-us_h^{(na)}) \right], \quad (18)$$

which does not depend on t . Using (16) and (18), a factorization is obtained which is useful for the Metropolis-within-Gibbs sampler described in the following section.

5 Algorithm for posterior simulation

5.1 Metropolis-within-Gibbs sampler

In our Metropolis-within-Gibbs sampler for simulating from the posterior (16), a single iteration is given by updating from full conditionals for five blocks of variables as specified by the first line in the following steps (A)-(E). Note that we use the notation $p(\cdot|\cdots)$ to indicate that we consider a variable or collection of variables \cdot given the remaining variables \cdots (including the data \mathbf{y}).

- (A) Update the non-allocated variables $(\boldsymbol{\mu}^{(na)}, \mathbf{s}^{(na)}, \boldsymbol{\gamma}^{(na)})$ from their full conditional as given by the following steps (i)-(iii), noting the following. Since the cardinality of each of the vectors $\mathbf{s}^{(na)}$ and $\boldsymbol{\gamma}^{(na)}$ agrees with the cardinality of $\boldsymbol{\mu}^{(na)}$, it is of paramount importance to resort to a *collapsed* Gibbs sampler. Therefore, in (i) we sample $\boldsymbol{\mu}^{(na)}$ from the conditional density obtained by integrating out $(\mathbf{s}^{(na)}, \boldsymbol{\gamma}^{(na)})$, and then in (ii)-(iii) we sample $\mathbf{s}^{(na)}$ and $\boldsymbol{\gamma}^{(na)}$ from their respective full conditionals, hence knowing the cardinality ℓ of $\boldsymbol{\mu}^{(na)}$.

- (i) Sample from the conditional density obtained by integrating out $(\mathbf{s}^{(na)}, \boldsymbol{\gamma}^{(na)})$ and given by

$$p(\boldsymbol{\mu}^{(na)} | \xi, \boldsymbol{\mu}^{(a)}, \mathbf{s}^{(a)}, \boldsymbol{\gamma}^{(a)}, \mathbf{c}, u, \mathbf{y}) \propto p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi) \psi(u)^\ell \quad (19)$$

with respect to $d\boldsymbol{\mu}^{(na)}$. Here, $\psi(u)$ denotes the Laplace transform of the density $p(s_h)$, and we can get rid of the last term $\psi(u)^\ell$, since ℓ is the cardinality of $\boldsymbol{\mu}^{(na)}$. In Section 5.2 we verify (19) and give details for simulation from (19). If, after this update, $\ell = 0$, the following two items (ii) and (iii) are skipped.

- (ii) Sample $\mathbf{s}^{(na)}$ from its full conditional,

$$p(\mathbf{s}^{(na)} | \cdots) \propto \prod_{h=1}^\ell p(s_h^{(na)}) \exp(-us_h^{(na)}).$$

That is, sample independently ℓ values from the exponential tilting of the prior density. Depending on the specific choice of $p(s_h^{(na)})$, this can be done exactly or requires a Metropolis-Hastings step.

(iii) Sample $\gamma^{(na)}$ from its full conditional,

$$p(\gamma^{(na)} | \dots) \propto \prod_{h=1}^{\ell} p(\gamma_h^{(na)}).$$

That is, sample independently ℓ values from the prior density $p(\gamma_h^{(na)})$.

(B) Update the allocated variables $(\mu^{(a)}, \mathbf{s}^{(a)}, \gamma^{(na)})$:

(i) Sample $\mu^{(a)}$ from its full conditional,

$$p(\mu^{(a)} | \dots) \propto p(\mu^{(a)} \cup \mu^{(na)} | \xi) \prod_{h=1}^k \left[\prod_{i:c_i=h} k(y_i | (\mu_h^{(a)}, \gamma_h^{(a)})) \right],$$

where by (6) we can replace $p(\mu^{(a)} \cup \mu^{(na)} | \xi)$ by $g(\mu^{(a)} \cup \mu^{(na)} | \xi)$. We do this by updating each of $\mu_h^{(a)}$ from

$$p(\mu_h^{(a)} | \dots) \propto g(\mu_h^{(a)} \cup \{\mu^{(a)} \setminus \{\mu_h^{(a)}\}\} \cup \mu^{(na)}) \prod_{i:c_i=h} k(y_i | (\mu_h^{(a)}, \gamma_h^{(a)})).$$

Appendix Section A.1 discusses how to construct a proposal density for sampling from $p(\mu_h^{(a)} | \dots)$ using a Metropolis-Hastings step.

(ii) Sample $\mathbf{s}^{(a)}$ from its full conditional,

$$p(\mathbf{s}^{(a)} | \dots) \propto \prod_{h=1}^k (s_h^{(a)})^{n_h} e^{-us_h^{(a)}} p(s_h^{(a)}).$$

Here, the $s_h^{(a)}$'s are independent conditional to everything else, so they can be updated individually using a Metropolis-Hastings step.

(iii) Sample $\gamma^{(a)}$ from its full conditional,

$$p(\gamma^{(a)} | \dots) \propto \prod_{h=1}^k p(\gamma_h^{(a)}) \prod_{i:c_i=h} k(y_i | (\mu_h^{(a)}, \gamma_h^{(a)})).$$

Unless $p(\gamma_h^{(a)})$ and $k(y_i | (\mu_h^{(a)}, \gamma_h^{(a)}))$ are conjugate, we apply a Metropolis step for the $\gamma_h^{(a)}$'s.

Since in this step we have conditioned with respect to \mathbf{c} too, k denotes the number of clusters and is fixed.

(C) Sample each c_i from its full conditional, which is a discrete distribution over $1, \dots, k + \ell$ given by

$$\begin{aligned} p(c_i = h | \dots) &\propto s_h^{(a)} k(y_i | (\mu_h^{(a)}, \gamma_h^{(a)})), & h = 1, \dots, k, \\ p(c_i = k + h | \dots) &\propto s_h^{(na)} k(y_i | (\mu_h^{(na)}, \gamma_h^{(na)})), & h = 1, \dots, \ell. \end{aligned}$$

After this, with a positive probability it may happen that $c_i > k$ for some i 's, so that some non-allocated components have become allocated, and some allocated components have become non-allocated. Then a simple relabelling of $(\mu^{(a)}, \mathbf{s}^{(a)}, \gamma^{(a)}, \mu^{(na)}, \mathbf{s}^{(na)}, \gamma^{(na)})$ and \mathbf{c} is needed, so that \mathbf{c} takes values in $\{1, \dots, k\}^n$.

(D) Sample ξ from its full conditional,

$$p(\xi | \dots) \propto p(\xi) p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi).$$

This requires a Metropolis-Hastings step, which is not straightforward when Z_ξ in (6) is not expressible in closed form, e.g. in the case of a repulsive pairwise interaction point process. Details on how this issue is overcome are given in Section 5.3.

(E) Sample u from its full conditional, which is just a gamma distribution with shape parameter n and inverse scale t .

5.2 Updating the non-allocated variables

This section provides the remaining details of step (A)(i). By (16),

$$\begin{aligned} p(\boldsymbol{\mu}^{(na)} | \xi, \boldsymbol{\mu}^{(a)}, \boldsymbol{S}^{(a)}, \boldsymbol{\gamma}^{(a)}, \boldsymbol{c}, u, \boldsymbol{y}) &= \int \int p(\boldsymbol{\mu}^{(na)}, \boldsymbol{s}^{(na)}, \boldsymbol{\gamma}^{(na)} | \xi, \boldsymbol{\mu}^{(a)}, \boldsymbol{S}^{(a)}, \boldsymbol{\gamma}^{(a)}, \boldsymbol{c}, u, \boldsymbol{y}) d\boldsymbol{s}^{(na)} d\boldsymbol{\gamma}^{(na)} \\ &\propto \int \int p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi) \left[\prod_{h=1}^{\ell} p(\gamma_h^{(na)}) p(s_h^{(na)}) \right] p(u | t) \frac{1}{t^n} d\boldsymbol{s}^{(na)} d\boldsymbol{\gamma}^{(na)} \\ &\propto \int p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi) \left[\prod_{h=1}^{\ell} \exp(-u s_h^{(na)}) p(s_h^{(na)}) \right] d\boldsymbol{s}^{(na)} \quad (20) \\ &= p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi) \psi(u)^\ell \quad (21) \end{aligned}$$

where (20) follows by integrating over $\gamma_h^{(na)}$ and using (18), and (21) by applying the definition of $\psi(u)$. This verifies (19).

Note that (19) identifies an unnormalized density for $\boldsymbol{\mu}^{(na)}$ with respect to $d\boldsymbol{\mu}^{(na)}$. In our examples, the unnormalized density in (21) will be hereditary, that is, $p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}^{(na)} | \xi) > 0$ implies $p(\boldsymbol{\mu}^{(a)} \cup \boldsymbol{\mu}'^{(na)} | \xi) > 0$ whenever $\boldsymbol{\mu}'^{(na)}$ consists of one more point than $\boldsymbol{\mu}^{(na)}$. Moreover, in our examples, this density is defined on a compact set, and so we can easily employ the birth-death Metropolis-Hastings algorithm in Geyer and Møller (1994). Specifically, we use Algorithm 11.3 in Møller and Waagepetersen (2004).

5.3 Sampling the hyperparameter ξ

When the density $p(\boldsymbol{\mu} | \xi)$ is expressible in close form, a standard Metropolis-Hastings move can be employed to update ξ from its full conditional. However, when Z_ξ is intractable, it is a doubly-intractable problem, since a ratio of unknown normalizing constants appears in the Hastings ratio. In fact, if $p(\xi'; \xi | \dots)$ is a proposal density for the Metropolis-Hastings step for the full conditional of ξ , then the acceptance ratio amounts to

$$\alpha(\xi'; \xi | \dots) = \frac{p(\xi') g(\boldsymbol{\mu} | \xi') p(\xi; \xi') \frac{Z_\xi}{Z_{\xi'}}}{p(\xi) g(\boldsymbol{\mu} | \xi) p(\xi'; \xi) \frac{Z_{\xi'}}{Z_\xi}}, \quad (22)$$

which is intractable due to the term $Z_\xi/Z_{\xi'}$. To overcome this issue, we can use the exchange algorithm described in Murray et al. (2006) and inspired by the *single auxiliary variable method* proposed by Møller et al. (2006). For further details, see Appendix A.2. This algorithm requires

generating an ancillary variable following the same distribution of $\boldsymbol{\mu}$ given ξ' . To this end, we employ the dominated coupling from the past algorithm in Kendall and Møller (2000).

In previous literature, the use of the ancillary variable algorithms in Møller et al. (2006) and Murray et al. (2006) has been limited because of the high computational cost associated to perfect simulation. In contrast, in our examples perfect simulation is fast. As an example, approximating the density of a DPP with $N = 50$ in dimension $q = 2$, as done in Bianchini et al. (2020), is around 25 times more expensive than running a perfect simulation from a Strauss process with parameters and prior for ξ chosen as in Section 6.1; for further comparisons, see Appendix B.1. The perfect simulation step is very fast since m (the number of components in the mixture model) is typically small (in our examples it is always less than 10). However, when dealing with applications with a very large number of clusters, such as topic modeling, where the number of clusters is usually between 50 and 100, cf. Blei et al. (2003), we expect perfect simulation to be potentially a bottleneck and limit the use of the exchange algorithm. Although not investigated here, in such cases we could avoid perfect simulation by replacing the exact exchange algorithm of Murray et al. (2006) with asymptotically exact algorithms that should offer smaller computational cost; see for instance Lyne et al. (2015) and Liang et al. (2016).

6 Prior elicitation

In this section, we discuss prior elicitation and how to set hyperparameters when the prior for $\boldsymbol{\mu}$ is the Strauss process or the DPP with power exponential spectral density.

6.1 Prior elicitation for the Strauss process

Consider the Strauss process prior given by (3)-(5). In addition to the parameter ξ which controls the intensity, the process depends on two parameters $\alpha \in [0, 1]$ and $\delta > 0$ which control repulsiveness and the range of interaction, respectively. Initially we investigated cases where α and δ were random, but then our simulated datasets yielded a large number of clusters a posteriori. Moreover, when fitting mixtures of Gaussian densities to data generated from heavy-tailed distributions, as also discussed at the beginning of Section 7, in general better density estimates were obtained when using a larger (i.e., larger than the true value) number of components in the mixture model. For this reason, we obtained a posteriori values of α and δ that induced less repulsive behaviors than desired. Therefore, we suggest below to fix α and δ via an empirical Bayes procedure, and let only ξ be random.

We propose to estimate α and δ as follows. Denote by $p(r)$ the kernel density estimate of the empirical distribution of the pairwise distances between observations; in all the examples, we have obtained such an estimate using the default bandwidth selection procedure in Python's `scipy` package. Since δ should be large enough to induce repulsion of redundant clusters, but not too large to affect density estimation severely, we suggest to fix δ as the smallest local minimum point of $p(r)$, that is, $\delta = \min_{r>0} \{r : r \text{ is a local minimum for } p(r)\}$. Further, α should be small enough to encourage separation between the allocated means. Consider, for example, the case with two clusters $\{y_i : c_i = h'\}$ and $\{y_i : c_i = h''\}$ where $0 < \|\mu_{h'} - \mu_{h''}\| \leq \delta$ but the distances from $\mu_{h'}$ and $\mu_{h''}$ to all the other μ_h 's are greater than δ . Then, by (15), the full conditional of $\mu_{h'}$ has density

$$p(\mu_{h'} \mid \cdots) \propto \alpha \prod_{i: c_i = h'} k(y_i \mid \mu_{h'}, \gamma_{h'}).$$

Now, the point of using a repulsive prior is that if the cardinality of cluster h' , that is $\#\{i : c_i = h'\}$, is small, the repulsiveness should prevail on the within-cluster likelihood: That is, regardless of how well the value of parameter $\mu_{h'}$ ‘fits’ data in cluster h' , the full conditional density associated to that value should be small because $\mu_{h'}$ is near to the cluster center $\mu_{h''}$. A rough estimate of α can be obtained by assuming $\alpha = \exp(-n^* \log(k_s))$. Here, n^* represents the minimum cluster size needed to balance the repulsive behavior induced by the prior, while k_s represents a ‘guess’ of $k(\cdot | \cdot)$ in a small cluster. In our experiments, we assumed that clusters with less than 5% of the data should be considered small and thus we fixed $n^* = n/20$. Further, we fixed $\log(k_s) = 1$ so that this term did not affect the definition of α . In addition, preliminary sensitivity analysis on α led us to conclude that posterior inference is robust.

Finally, we assume that ξ is random. An upper bound for the expected number of points in μ is $\xi|R|$, and given an upper bound M_{\max} on the expected number of components, we assume the prior for ξ to be uniform over the interval $(|R|^{-1}, M_{\max}|R|^{-1})$. Since the number of clusters is smaller than the number of components, M_{\max} is an upper bound for both, to be fixed in each application according to prior belief.

6.2 Prior elicitation for the power exponential spectral DPP model

For the DPP defined on \mathbb{R}^q by the spectral density χ used in (13), existence is ensured if $0 < \alpha \leq \alpha_{\max}$, where

$$(\alpha_{\max})^q = \frac{\pi^{q/2} \Gamma(q/\beta + 1)}{\xi \Gamma(q/2 + 1)},$$

cf. Lavancier et al. (2015). So we let $\alpha = s \alpha_{\max}$ with $0 < s < 1$ (as specified below), which implies existence of the DPP restricted to any compact subset of \mathbb{R}^q . Note that the DPP density given by (12)-(13) refers to the case $R = [-1/2, 1/2]^q$, and a simple rescaling is needed in the density expression when we fix R to be the smallest rectangle containing all the observations, cf. Lavancier et al. (2015).

Recall that ξ is the expected number of points in μ . We let a priori ξ be uniformly distributed over $[1, M_{\max}]$, where M_{\max} is fixed (as in the case of the Strauss process, cf. Section 6.1). As noted in Lavancier et al. (2015), the parameters (s, β) are harder to interpret via (13). In our examples, we fix $s = 0.5$ and perform sensitivity analysis on β , concluding that inference is robust.

7 Simulation studies

In this section, we compare the reversible jump algorithm in Bianchini et al. (2020) to our Metropolis-within-Gibbs sampler presented in Section 5, and show the advantages of repulsive mixtures over non-repulsive ones. We refer to our Metropolis-within-Gibbs sampler as the ‘M-w-G sampler’ and the reversible jump algorithm as ‘RJ’. In Appendix B, we illustrate the advantages of using a Strauss process over a DPP as prior for μ and provide further simulations when the dimension q or the number of components m increase. In particular, we conclude that the computational cost of posterior inference under the DPP grows exponentially with data dimension q , whilst the computational cost associated to the Strauss process is almost constant as data dimension increases, and that posterior summaries obtained under the DPP and Strauss process are almost identical. This motivates the use of the Strauss process as a prior for μ .

In this section, we study posterior inference in *misspecified* settings, i.e., when the generating process does not coincide with the model used to fit data; for a formal definition of misspecification, see Kleijn et al. (2006). In misspecified settings, there is a trade-off between the accuracy of the density and number of clusters estimation recovered by the mixture model, cf. Guha et al. (2019). This indicates that more accurate density estimates correspond to overestimated number of clusters and vice-versa. In fact, to recover the shape of non-Gaussian data, several Gaussian components (with similar values of the mean parameters) are needed. We expect that the repulsiveness induced by the prior for μ favours cluster over density estimation.

We consider two simulation scenarios, the first one is as in Miller and Dunson (2019), where the authors generated iid data y_1, \dots, y_n using a two-step procedure as follows. First, a mixture density f_0 with m_0 components is selected. Second, a random density \tilde{f} is drawn from a Dirichlet process mixture, with base measure given by f_0 . Specifically,

$$\begin{aligned} y_1, \dots, y_n \mid P &\stackrel{\text{iid}}{\sim} \tilde{f}(\cdot) = \int \mathcal{N}(\cdot \mid \theta, 0.25^2) P(d\theta), \\ P &\sim DP(af_0), \quad f_0 = \sum_{h=1}^{m_0} w_{0h} \mathcal{N}(\mu_{0h}, \sigma_{0h}^2), \end{aligned} \tag{23}$$

where $DP(af_0)$ denotes the Dirichlet process with total mass parameter a and centering probability measure induced by f_0 . We fix $a = 500$, $m_0 = 4$, $\mathbf{w}_0 = (0.25, 0.25, 0.3, 0.2)$, $\mu_0 = (-3.5, 3, 0, 6)$, and $\sigma_0 = (0.8, 0.5, 0.4, 0.5)$. Following Miller and Dunson (2019), we interpret the data generating density \tilde{f} as a perturbation of the ‘true’ density f_0 , so the goal is to recover f_0 and m_0 .

The second simulation scenario considers draws from the following mixture of two components:

$$y_1, \dots, y_n \stackrel{\text{iid}}{\sim} 0.5 t_q(1, \mu_0, \Sigma_0) + 0.5 MSN_q(\omega, \mu_1, \sigma_1). \tag{24}$$

Here $t_q(1, \mu_0, \Sigma_0)$ denotes the density of a q -dimensional Student distribution with one degree of freedom, location μ_0 , and scale matrix Σ_0 . Furthermore, $MSN_d(\omega, \mu_1, \sigma_1)$ denotes the density of a q -dimensional random vector, where each entry is drawn independently from a skew normal distribution with mean $\mu_1 + \omega \sigma_1 \sqrt{2/\pi}$, being μ_1 the location parameter and ω the scale parameter of the skew normal distribution. The dimension q and the other parameters in (24) will be specified later.

For both scenarios, the kernel $k(\cdot \mid \cdot)$ in (1) is either the univariate or the multivariate Gaussian density. In addition to the prior assumptions (a)-(e) in Section 1.1, we let a priori $(w_1, \dots, w_m) = (s_1/t, \dots, s_m/t)$, with \mathbf{s} as in Section 3, where each s_h follows a gamma distribution with shape and scale equal to one, and assume each γ_h to be either inverse-gamma distributed (if data are univariate) or inverse-Wishart distribution (for multivariate data), with hyperparameters fixed as in Fraley and Raftery (2007). Finally, unless otherwise stated, parameters of the Strauss point process or the DPP are chosen as discussed in Sections 6.1-6.2. In particular, we fix $M_{\max} = 30$.

7.1 Monitoring MCMC mixing

In this section, data are given by 500 observations simulated in accordance to (23). The marginal prior for μ is the DPP specified in Section 6.2, where in order to identify the effect of the algorithm on posterior inference, we keep the intensity parameter ξ fixed. Furthermore, we

Params.		RJ		M-w-G sampler		
ξ	β	ESS	$\mathbb{E}[m \mid \text{data}]$	ESS	$\mathbb{E}[m \mid \text{data}]$	$\mathbb{E}[k \mid \text{data}]$
4	10	90.63	4.33	8201.41	4.01	4.00
4	2.5	62.46	4.402	3735.80	4.01	4.00
4	25	83.32	4.44	2971.05	4.02	4.00

Table 1: Summary of the MCMC simulations for the reversible jump algorithm (RJ) in Bianchini et al. (2020) and our Metropolis-within-Gibbs sampler (M-w-G sampler). ESS denotes the effective sample size out of the 10,000 MCMC samples.

consider three possible values for the hyperparameters ξ and β in the DPP prior, cf. Table 1. For each choice of hyperparameter values, we ran both MCMC algorithms (M-w-G sampler and RJ) for 20,000 iterations, discarding the first 10,000 as burn-in and without thinning the chain. In order to compare the results, we consider the effective sample size (ESS) of the number of components in the mixture (m in our notation) as well as its autocorrelation.

Table 1 reports, for different combination of hyperparameters, the posterior expected value of m as well as the effective sample sizes for m obtained by the two algorithms. Since in our M-w-G sampler the number of clusters can be smaller than m , the table also shows the posterior expected value of k (the number of allocated components/clusters). Figure 1 shows for both algorithms trace plots and autocorrelation plots for m when $\xi = 4$ and $\beta = 10$ (first row of Table 1). Note that both algorithms offer good estimates of the number of components in the mixture. However, the performance of our M-w-G sampler is superior to the RJ algorithm in all the settings of hyperparameters we tested: Our M-w-G sampler generally produces a (much) higher effective sample size and overall better mixing of the chains.

7.2 Comparison with DPM and FM

We focus on the differences between repulsive and non-repulsive mixtures using two further simulations. For the class of non-repulsive mixtures, we consider (i) the finite mixture models (FM) of Gaussian densities in Argiento and De Iorio (2019) and Miller and Harrison (2018), and (ii) the Dirichlet Process Mixture (DPM) of Gaussian densities.

Both FM and DPM require the choice of a base measure P_0 that we fix as the normal-inverse-Wishart distribution (or the normal-inverse-gamma distribution in the univariate case). Hyperparameters are fixed according to Fraley and Raftery (2007) to provide a weakly informative prior. Moreover, the concentration parameter in the Dirichlet process is fixed to one, and for the FM model we consider as prior for m the shifted Poisson distribution (with support $\{1, 2, \dots\}$) so that the prior mean of the components is equal to four if the data generating process is (23) and to two if the data generating process is (24). Finally, for our model, we assume the Strauss process prior for μ .

Posterior simulation from the FM model was carried out using the R package `AntMAN`¹, while for the DPM we used the R package `BNPMix` (Corradin et al., 2020). For all three models, we ran the MCMC algorithm for 100,000 iterations, discarding the first 50,000 as a burn-in and keeping one of every ten iterations, so that in each case the final sample size is 5,000.

In the first simulation study, data are given by 400 simulated observations from (23). Figure 2 shows the true data generating density, together with Bayesian mixture density estimates

¹available at <https://github.com/bbodin/AntMAN>

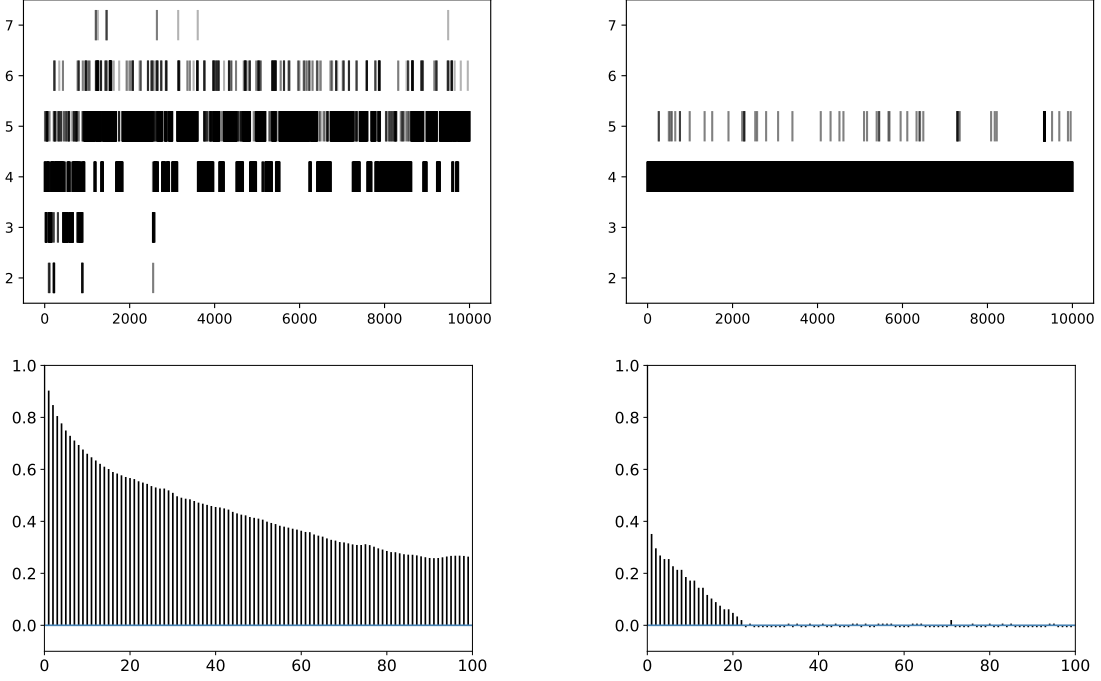


Figure 1: Trace plots (top) and autocorrelations (bottom) of m when $\xi = 4$ and $\beta = 10$. Left: RJ. Right: M-w-G sampler.

obtained by our model and the DPM (left), as well as the distribution of the number of clusters under the three models (right). Here, by Bayesian density estimate we always mean the posterior expectation of the mixture density evaluated on a fixed grid of points. As expected, under this misspecified setting, the use of non-repulsive mixture models overestimate the number of clusters. For instance, to recover the shape of the leftmost *bell* of the data generating density in Figure 2, several Gaussian components (with close cluster centers) are needed. Our model instead, due to the repulsiveness induced by the prior on μ , ‘correctly’ identifies four clusters.

For the second simulation study, we simulated 500 observations from (24) in each of the cases $q = 1$ and $q = 5$, where we fixed $\mu_0 = (-5, \dots, -5)$, $\Sigma_0 = I_q$, $\omega = 2$, $\mu_1 = 5$, and $\sigma_1 = 1$. Figure 3 reports density estimates when $q = 1$ (left) and the posterior distribution of the number of clusters for the three models (right) when $q = 1, 5$. Note that, among the three models, our is the one that gives highest posterior probability to the true value $k = 2$. When $q = 5$, DPM assigns the highest probability to three clusters. Appendix B.2 contains a comparison of the cluster estimates under the three models considered when $q = 1$, and we conclude that the repulsive mixture model is the one that better recovers the true clustering of the data in this example.

8 Teenager problematic behavior dataset

In this section, we apply our model, with the Strauss process prior for μ , to a dataset consisting of $n = 6504$ observations coming from the Wave 1 data of the National Longitudinal Study of Adolescent to Adult Health, which is available at <http://www.icpsr.umich.edu/icpsrweb/>




Figure 2: Posterior inference based on data simulated from (23). To the left, Bayesian mixture density estimates under the Strauss process and the DPM priors for μ , together with the true mixture density which has four components. To the right, posterior distributions of the number of allocated components under the Strauss process, FM, and DPM priors for μ .

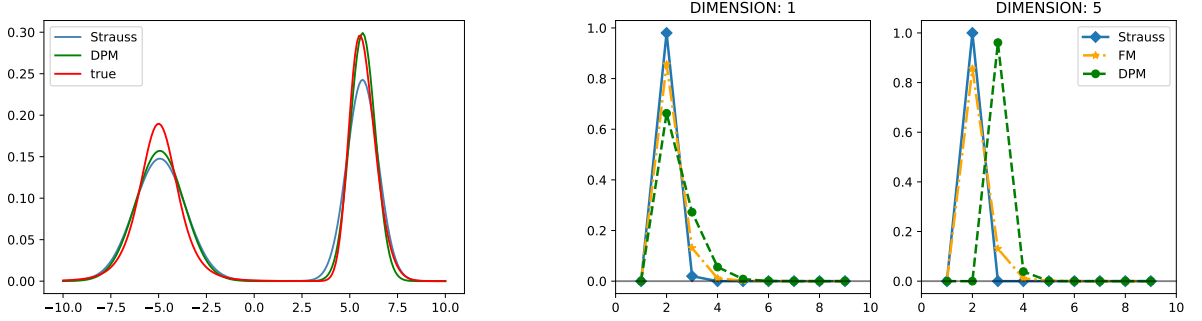


Figure 3: Posterior inference based on data simulated from (24). To the left, when $q = 1$, Bayesian mixture density estimates under the Strauss process and the DPM priors for μ , together with the true mixture density which has two components. To the right, when $q = 1, 5$, posterior distributions of the number of allocated components under the Strauss process, FM, and DPM priors for μ .

ICPSR/studies/21600. The data were also considered in Collins and Lanza (2009) and Li et al. (2018).

The dataset corresponds to six survey items pertaining problematic behaviors in teenagers, so that for the i 'th teenager, $y_i = (y_{i1}, \dots, y_{i6}) \in \{0, 1\}^6$ is a binary vector, where $y_{ij} = 1$ means a positive answer to entry j . The six entries correspond to (i) 'lied to parents', (ii) 'loud/rowdy/unruly in a public place', (iii) 'damaged property', (iv) 'stolen from a store', (v) 'stolen something worth less than 50 dollars', and (vi) 'taken part in a group fight'.

We let the kernel in (1) be given by

$$k(y | \mu_h) = \prod_{j=1}^6 \mu_{hj}^{y_j} (1 - \mu_{hj})^{1-y_j}, \quad y = (y_1, \dots, y_6) \in \{0, 1\}^6, \quad (25)$$

so that the six entries in $y_i = (y_{i1}, \dots, y_{i6})$ are conditionally independent binary random variables with success probability vector $\mu_h = (\mu_{h1}, \dots, \mu_{h6})$. Note that there is no parameter γ ,




Figure 4: Estimated cluster centres $\hat{\mu}_h^{(a)}$ (in blue) and empirical estimates μ_h^{emp} (in orange) when the clusters are sorted according to cluster sizes as given by the estimated weights \hat{w}_h , $h = 1, \dots, 5$ (specified at the top of each plot). The clusters are also labelled as (1), \dots , (5) (specified at the bottom of each plot).

and the probability vector (μ_1, \dots, μ_m) belongs to $R = [0, 1]^6$. The mixture model with kernel (25) is known as a *latent class model*.

As the prior for $\boldsymbol{\mu}$, we assume the Strauss process on R with parameters $\delta = 0.4$, $\alpha = e^{-n^*}$ (with $n^* = 50$), and a uniform prior on ξ with $M_{\max} = 30$, cf. Section 6.1. In this context, we may consider the μ_h 's as cluster centres/locations, where repulsion among the μ_h 's is meant to favor identification of the clusters.

We ran our posterior simulation algorithm for 20,000 iterations, after discarding other 20,000 iterations as burn-in and saving one of every ten iterations. So the final sample is of size $M = 2,000$, and we denote μ_h^j the value of μ_h at iteration $j = 1, \dots, M$. Below we summarize our findings for the cluster centres and compare to what was obtained in Li et al. (2018), where the authors used a finite mixture model with the same kernel (25) as ours, but fixed the number of clusters to be equal to four.

We obtained $P(k = 5 | \text{data}) \approx 1$. As usually done in Bayesian mixture modelling, we computed a point estimate of the latent partition of the data (as given by the unknown c_i 's) by selecting, among the partitions visited during the MCMC iterations, the minimum point of the Binder loss function with equal misclassification cost, cf. Binder (1978). Then, we evaluated the weights in each cluster by $\hat{w}_h = \#\hat{C}_h/n$, $h = 1, \dots, 5$, where \hat{C}_h is the estimated index set of data in cluster h . Furthermore, as in Molitor et al. (2010), we estimated the cluster centres by

$$\hat{\mu}_h^{(a)} = \frac{1}{M} \sum_{j=1}^M \frac{1}{\#\hat{C}_h} \sum_{i \in \hat{C}_h} \mu_{c_i}^j, \quad h = 1, \dots, 5.$$

Figure 4 shows these estimates, together with the empirical frequencies in each cluster as given by

$$\mu_h^{\text{emp}} = \frac{1}{\#\hat{C}_h} \sum_{i \in \hat{C}_h} y_i, \quad h = 1, \dots, 5.$$

Note that in Figure 4, the estimated clusters are labeled (1), \dots , (5) and ordered by the estimated weights.




Figure 5: Pairwise distances between the estimated $\hat{\mu}_h^{(a)}$, $h = 1, \dots, 5$.

The following interpretation of the clusters is consistent with the one given in Li et al. (2018): Figure 4 shows that cluster (1) accounts for 59% of the data and groups teenagers with few problematic behaviours, since all estimated and empirical cluster centers in the leftmost panel in Figure 4 are small. Further, cluster (2) groups 18% of the subjects and describes minor problematic behaviours (relating to the first and second survey items). Finally, clusters (3), (4), and (5) represent smaller groups of teenagers who are truly problematic, as their tendency to commit small crimes (cluster 4) or fights (clusters 3 and 5) is very high.

In Figure 4, there are discrepancies between the empirical frequencies and our estimates, see for instance the estimates of μ_{h1} and μ_{h2} in cluster (2) and of μ_{h3} and μ_{h6} in cluster (5). These discrepancies can be explained by the use of the repulsive prior, which encourages separation among clusters.

Moreover, Figure 5 shows the pairwise Euclidean distances among the estimated $\hat{\mu}_h^{(a)}$, $h = 1, \dots, 5$. Here, the smallest distance is around 0.41, which is close to the value of δ (which we fixed to be equal to 0.4). Note that $\hat{\mu}_5^{(a)}$ is very close to both $\hat{\mu}_2^{(a)}$ and $\hat{\mu}_3^{(a)}$; and $\hat{\mu}_1^{(a)}$ and $\hat{\mu}_2^{(a)}$ are close as well.

Finally, we performed posterior inference with $\delta = 0.5$ and $n^* = 100$ to induce more separation. In this case, our inference gave four estimated clusters, in agreement with Li et al. (2018). However, compared to Figure 4, the estimated cluster centres $\hat{\mu}_h^{(a)}$ were then further more different from the empirical frequencies μ_h^{emp} . As noticed in Section 7, this trade-off between density versus cluster estimation accuracy is not surprising.

9 Discussion

In this work we have contributed to the fast-growing literature on repulsive mixture models. A main contribution is the introduction of a unifying framework which encompasses previously proposed repulsive mixtures as special cases. In our setting, a repulsive point process is assumed as prior for ‘cluster centres’ of the parametric kernel densities, thus making it more likely having a small number of well separated clusters in the mixture model. In particular, we have showed the usefulness of the Strauss process prior, which is a simple example of a repulsive pairwise interaction point process.

By studying posterior characterization of the repulsive point process, we were able to derive

a Metropolis-within-Gibbs sampler that avoids the arduous choice of problem-specific reversible jump proposals (Xu et al., 2016; Bianchini et al., 2020) and the computationally expensive evaluation of infinite summations and integrals over the parameters space (Xie and Xu, 2019) as seen in previous work. When deriving the posterior distribution of the repulsive point process prior, we extended the approach in Argiento and De Iorio (2019) but framing our model within the class of normalized point processes mixture models.

Our MCMC algorithm can also handle cases when the point process density involves an intractable normalizing constant, which has not been considered in the previous literature. In particular, we used an ancillary variable method which eliminates the problem of having a ratio of normalizing constants in the Hastings ratio when making posterior simulations for full conditional of the hyperparameter. Since our mixture model is parsimonious (i.e., the number of components is typically small), the ancillary variable method relying on a perfect simulation algorithm is fast.

We tested our approach by extensive simulation studies, comparing it to the reversible jump approach of Xu et al. (2016) and Bianchini et al. (2020), where we concluded that our Metropolis-within-Gibbs (M-w-G) sampler has better mixing. Our M-w-G sampler scales well with data dimension and this feature was particularly evident when we assumed the Strauss process as a prior for the cluster centers. Furthermore, since repulsive mixture models encourage a small number of well separated components, thus controlling the computational cost, our algorithm was shown to scale well with sample size too.

Finally, we illustrated the advantages of repulsive mixtures against the popular Dirichlet process mixtures and finite mixtures. We concluded that repulsive mixtures are especially useful when the model is misspecified.

Several further extensions are possible. Beyond mixture models for cluster detection, feature allocation problems and regression settings could be considered. Further, adapting our approach to hierarchical and nested settings, where multiple groups of data are present, could be of interest. Finally, extensions of our model to handle extremely high dimension data are also of interest, for instance in the field of genomics, where a repulsive prior would help in deriving interpretable results characterized by few and well separated clusters.

A Further details on the Metropolis-within-Gibbs sampler used for posterior simulation

This section provides additional details for the Metropolis-within-Gibbs (M-w-G) sampler in Section 5.

A.1 The choice of the proposal distribution

For most choices of the point process density $p(\boldsymbol{\mu} | \xi)$ and the mixture kernel $k(\cdot | \cdot)$, the update of the allocated means $\mu_h^{(a)}$ requires sampling from an unnormalized distribution, which we do via a Metropolis-Hastings step. As proposal distribution we use a mixture of two normal distributions with means equal to the current value of $\mu_h^{(a)}$ but with different variances so that

$$p(\mu'; \mu_h^{(a)}) = \kappa \mathcal{N}(\mu' | \mu_h^{(a)}, \underline{\sigma}^2 I) + (1 - \kappa) \mathcal{N}(\mu' | \mu_h^{(a)}, \bar{\sigma}^2 I), \quad (26)$$

where $\kappa = 0.9$, $\underline{\sigma} = 0.1$, and $\bar{\sigma} = 1.5$ when $q = 1, 2$ and $\bar{\sigma} = 1.5q$ when $q > 2$. The intuition that led us to consider such a proposal is as follows, where for ease of notation we drop the superscript (a) when considering a current value of $\mu_h^{(a)}$, denoted μ_1 , and another cluster centre μ_2 . Suppose that μ_1 and μ_2 are close and far from the remaining points in $\boldsymbol{\mu}$. If the number of observations allocated to μ_1 is small, we want a proposal distribution $p(\mu'_1; \mu_1)$ that gives significant mass to values that are far from μ_2 , so that, given the repulsiveness of the point process, this proposal is likely to be accepted. This is the case when we sample from the second component of (26) (in fact, if μ'_1 is far from μ_1 , with sufficiently large probability it is far from μ_2 as well). On the other hand, if the number of observations allocated to μ_1 is large, we want a proposal that gives significant mass to a neighborhood of the current value of μ_1 , to get a precise fit of the data. This is what happens if we sample from the first component of (26).

For the second component in (26), instead of fixing $\bar{\sigma}$ as we do, an alternative is to exploit the properties of $g(\cdot | \xi)$ as follows. Suppose we condition on sampling from $\mathcal{N}(\mu'_1 | \mu_1, \bar{\sigma}^2 I)$ in (26). Then $\|\mu'_1 - \mu_1\|^2 / \bar{\sigma}^2 \sim \chi^2(q)$, the chi-squared distribution with q degrees of freedom. Considering the Strauss density, a possibility is to fix $\bar{\sigma}$ to give sufficiently high mass to values of μ'_1 that are outside the range of interaction of μ_1 , i.e., such that $P(\|\mu'_1 - \mu_1\|^2 > \delta) > p_0$ for some fixed p_0 , with the intuition that this gives a positive probability to μ'_1 being distant at least δ also from μ_2 . Considering the DPP density instead, the same argument holds but replacing δ with the range of correlation r_0 , cf. Lavancier et al. (2015). That is, (12) implies that C is of the form $C(\mu_1, \mu_2) = C_0(r)$ with $r = \|\mu_1 - \mu_2\|$, and defining the corresponding correlation function $R(r) = C_0(r)/C_0(0)$, r_0 is chosen such that $R(r)$ is effectively zero.

A.2 The exchange algorithm and perfect simulation

With the same notation as Section 5.3, the exchange algorithm (Murray et al., 2006) consists of the following steps:

1. Propose $\xi' \sim p(\xi'; \xi)$.
2. Generate an auxiliary variable $\boldsymbol{\mu}^{\text{aux}} \sim g(\boldsymbol{\mu} | \xi') / Z_{\xi'} \propto g(\boldsymbol{\mu} | \xi')$.
3. Accept ξ' with probability $\min\{1, \alpha^*\}$ where

$$\alpha^* \equiv \alpha^*(\xi; \xi' | \dots) = \frac{p(\xi')g(\boldsymbol{\mu} | \xi')p(\xi; \xi')}{p(\xi)g(\boldsymbol{\mu} | \xi)p(\xi'; \xi)} \times \frac{g(\boldsymbol{\mu}^{\text{aux}} | \xi)}{g(\boldsymbol{\mu}^{\text{aux}} | \xi')}.$$

Comparing α^* to the acceptance ratio in (22), note that the ratio $Z_{\xi}/Z_{\xi'}$ has been replaced by a ratio of unnormalized densities, evaluated in the auxiliary variable $\boldsymbol{\mu}^{\text{aux}}$. The main difficulty is sampling $\boldsymbol{\mu}^{\text{aux}}$, which must follow the distribution of $\boldsymbol{\mu}$ given ξ' . To this end, we employ the stochastic dominated coupling from the past algorithm in Kendall and Møller (2000), which extends the coupling from the past algorithm in Propp and Wilson (1996) to uncountable partially ordered spaces. Specifically, we employed in our code Algorithm 11.7 in Møller and Waagepetersen (2004).

B Additional simulation studies

In addition to the simulation studies in Section 7, below we discuss different aspects of the M-w-G sampler and posterior inference.




Figure 6: Per-iteration run-times as a function of data dimension q in case of DPP (with truncation levels $N = 5$ or 10) and Strauss process priors for $\boldsymbol{\mu}$.

B.1 Comparison of run-times and posterior inference when using DPP and Strauss process priors

For $q = 1, 2, \dots, 5$, we simulated $n = 200$ observations from (24) with $\mu_0 = (-5, \dots, -5)$, $\Sigma_0 = I_q$, $\omega = 1$, $\mu_1 = 5$, and $\sigma_1 = 1$. Then we applied our M-w-G sampler when the marginal prior for $\boldsymbol{\mu}$ is either the DPP or the Strauss process, with hyperparameters as in Section 6. Here, we considered two truncation levels for the approximation of the DPP density in (12), namely $N = 5$ and $N = 10$ (for comparison, Bianchini et al. (2020) suggested $N = 50$ when $q = 1$).

Figure 6 shows the per-iteration run-times of the M-w-G sampler as a function of the dimension q under either the DPP or Strauss process prior for $\boldsymbol{\mu}$. For each value of N , the computational cost associated to the DPP grows exponentially fast as the dimension q increases, unlike in the case of the Strauss process. In fact, the unnormalized density of the Strauss process is almost immediate to compute, and since the Strauss prior is quite informative on the number of components, cf. Section 6.1, the perfect simulation algorithm (see Section 5.3) does not impact significantly on the computational cost. Although not appreciable from Figure 6, the computational cost of our algorithm increases significantly with data dimension q also when we consider the Strauss process; in this case, the per-iteration computational cost goes from 0.0016 sec when $q = 1$ to 0.07 sec when $q = 5$, i.e., it increases by a factor of roughly 50.

As a further comparison, we simulated 500 univariate observations from model (23) and made again posterior computations under the Strauss process or the DPP prior for $\boldsymbol{\mu}$, where for the DPP density we fixed $\beta = 10$ (corresponding to the highest ESS in Table 1). For both cases of prior models, we ran the M-w-G sampler for 100,000 iterations discarding the first 50,000 as a burn-in and keeping one every ten iterations, for a final sample size of 5,000. Figure 7 shows the true data generating density, together with Bayesian mixture density estimates and posterior distributions of the number of clusters under the two point process priors. Note that the two density estimates, as well as the two posterior distributions of the number of clusters, overlap almost perfectly. The Strauss process seems a good choice to model the prior of $\boldsymbol{\mu}$ since it, for a much smaller computational cost, provides same posterior summaries as the DPP.

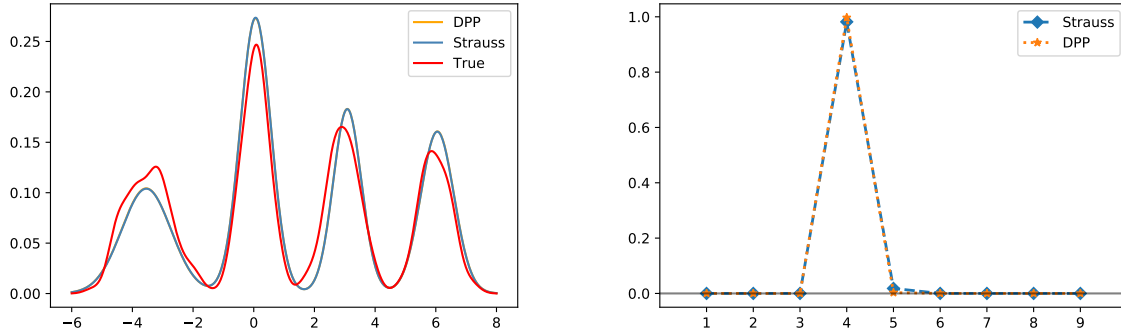


Figure 7: Bayesian mixture density estimates (left) and posterior distributions of the number of clusters (right) under the Strauss process (blue lines) and DPP (orange lines) priors for μ , together with the true mixture density which has four components. The orange lines overlap almost perfectly with the blue lines so that they are hardly visible.




Figure 8: Posterior similarity matrices and ARI scores under the three models for the mixture of the univariate t and skew-normal distributions discussed in Section 7.2. The colors are on a logit scale to highlight differences around one.

B.2 Accuracy of cluster estimates

Figure 8 shows the posterior similarity matrices and the Adjusted Rand Index (ARI) scores for the univariate mixture of t and skew-normal distribution discussed in Section 7.2. The ARI is computed from the cluster labels \mathbf{c} at each iteration of the MCMC chain as a measure of similarity between the estimated clusters and the true cluster. It is bounded by 1 and the larger value it assumes, the more similar is the estimated cluster to the true one. We report the posterior mean of the ARI \pm one standard deviation on top of each posterior similarity matrix in Figure 8. The difference in the posterior similarity matrices is not so pronounced, but our repulsive mixture model gives the best ARI.

B.3 The effect of the number of clusters

We consider how the number of clusters affects the performance of our M-w-G sampler. When μ is distributed as the Strauss process, at every step of the MCMC algorithm a perfect simulation of μ is required. The perfect simulation algorithm we use has a finite but random computational




Figure 9: Locations of the true data generating process (left) and run-time comparison (right). The plot of the locations should be intended as follows: for $m = 5$ only the points labelled accordingly are considered, for $m = 9$ the points labeled as $m = 5$ and $m = 9$ are considered and so on. The run-times (in seconds) over 50 independently simulated datasets for each value of m are denoted by dots, we also report the median times as diamonds with a dashed line connecting them.

cost and, as argued in Section 5.3, it might become infeasible for a large number of clusters. On the other hand, when μ is a DPP, the approximation of its density requires computing the determinant of the matrix C' in (10), which scales cubically with m . Furthermore, for the specific DPP considered in (12) computing C' requires the evaluation of $O(N^q m^2)$ inner products.

We generated $n = 500$ observations from a mixture of $m = 5, 9, 17, 25$ bivariate Gaussian densities, with locations given in Figure 9 (left), equal covariance matrices given by $0.5I_2$, and with equal mixture weights. We compared the run-times (in seconds) required to complete 200 iterations with our M-w-G sampler when μ is distributed either as the Strauss or the determinantal point process. Prior hyperparameters are fixed as in Section 6 (with $M_{\max} = 5m$) and Section 7. For the DPP, we considered two truncation levels of the spectral density, $N = 10, 50$. For each choice of m we generated 50 independent datasets and for the 200 M-w-G sampler iterations we used fixed and different independent random seeds.

In Figure 9 (right) for each m the run-times over the 50 independent datasets are denoted by dots, the median times by diamonds, and the median times are connected by a dashed line. We see that the DPP with $N = 50$ is the most computationally demanding model for all values of m . When $m = 5, 9$, the Strauss process is significantly faster (up to 10 times faster) than the DPP with $N = 10$; instead, when $m = 17$, they have comparable computational costs. When $m = 25$, the perfect simulation algorithm starts to become more demanding; for example, the computational cost for the Strauss process is almost twice the one for the DPP with $N = 10$.

B.4 The effect of the data dimension

Below we compare our repulsive mixture model, the finite mixture model (FM) in Argiento and De Iorio (2019), and the Dirichlet process mixture model (DPM). See Section 7.2 for further details on how posterior inference is performed under the different models. In particular, we fix the hyper-parameters according to Sections 6 and 7.

		$q = 5$	$q = 10$	$q = 15$	$q = 20$	$q = 25$	$q = 30$
Strauss	ARI	1.0	1.0	1.0	1.0	1.0	1.0
	ESS	240.3	250.1	0.0	0.0	0.0	0.0
	$\mathbb{E}[k \mid \text{data}]$	2.01	2.005	2.0	2.0	2.0	2.0
FM	ARI	1.0	1.0	1.0	1.0	1.0	1.0
	ESS	7.4	0.0	0.0	0.0	0.0	0.0
	$\mathbb{E}[k \mid \text{data}]$	2.01	2.0	2.0	2.0	2.0	2.0
DPM	ARI	1.0	0.0	0.0	0.0	0.0	0.0
	ESS	0.0	0.0	0.0	0.0	0.0	0.0
	$\mathbb{E}[k \mid \text{data}]$	2.00	1.0	1.0	1.0	1.0	1.0

Table 2: Adjusted Rand Index (ARI), effective sample size for the chain of the number of clusters k and posterior mean of k under the repulsive mixture model (Strauss), the non repulsive finite mixture model (FM) and the Dirichlet process mixture model (DPM).

If the sample size n is not significantly larger than the data dimension q , the use of commonly employed MCMC algorithms may be problematic for the following reasons. If $k(\cdot \mid \cdot)$ is a multivariate Gaussian density with non-zero correlations, the number of parameters to be estimated is much larger than n . Further, the curse of dimensionality, common to all clustering problems (Kriegel et al., 2009), implies a poor mixing of the algorithms. In addition to that, when considering a repulsive mixture model, things might be further complicated by either the need of perfect simulation to update possible hyperparameters ξ (when μ follows the Strauss process) or the computation of the spectral density (when μ follows the DPP given by (12)) which becomes prohibitive even for moderate values of q , as shown in Figure 6. Therefore, below we consider only the Strauss process and perform a simulation to assess the performance of repulsive versus non-repulsive mixtures when $q = 2, 5, 10, 15, 20, 25, 30$ increases. Moreover, we simulated $n = 200$ observations from

$$y_i \stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(-5/\sqrt{q}\mathbf{1}_q, I_q) + 0.5\mathcal{N}(5/\sqrt{q}\mathbf{1}_q, I_q)$$

where $\mathbf{1}_q$ denotes the vector in \mathbb{R}^q with elements all equal to one.

Table 2 reports posterior summaries as q increases for the three models. MCMC chains were run for 11,000 iterations discarding the first 10,000 as burn-in, so that the effective sample size must be referred to a total number of MCMC iterations equal to 1,000. It is clear from Table 2 that as q increases, the mixing of the chains becomes progressively worse for all the models. In particular, the table shows the effective sample size (ESS) for the three cases: For our repulsive mixture model, the number of clusters k is constant for all the MCMC iterations when $q \geq 15$, and so ESS is zero; for FM, the ESS is zero when $q \geq 10$; and for DPM the ESS is zero for all values of q . The difference in the ARI scores is simply explained by the different strategy of initialization of the different software we ran: In our code for the M-w-G sampler, observations are initially randomly subdivided into 10 clusters; in the package **AntMAN**, which we used to fit the FM model, one cluster per observation is created; in the package **BNPMix**, used to fit the DPM model, all observations are initially allocated to one single cluster. In the latter case, the proposal of a new cluster is never accepted. Using our software or the package **AntMAN** instead, after a few MCMC iterations the observations are (correctly) partitioned into $k = 2$ clusters and no additional cluster is ever created.

Considering the effective sample size of k , Table 2 shows that repulsive mixture models might offer an advantage over non-repulsive mixture models when $q \leq 10$. We believe that

the poor performance of FM and DPM is due to prior assumptions for the following reason. Note that both models assume that the parameters $\{(\mu_h, \gamma_h)\}_h$ are a priori iid and normal inverse-Wishart distributed, with $\mathbb{E}[\mu_h] = 0$. Thus, as q increases, the multivariate Gaussian distribution becomes more and more concentrated around the mean, due to the so-called curse of dimensionality, so that proposing a new value for μ_h from the prior that is near to any of the observations becomes less likely. Instead, when considering a Strauss point process as prior for $\boldsymbol{\mu}$, the proposed means are not concentrated around the origin, which led to a better mixing when $q = 5, 10$. When $q \geq 15$, we believe that the volume of the rectangle containing all observations becomes so large that also the repulsive mixture models suffer from the curse of dimensionality.

Perfect simulation is not a bottleneck here, as the number of points in the Strauss process is small. However, in one of several independent simulations, an unlucky initialization led to a large value of m in the first few iterations. As a consequence, the perfect simulation algorithm took longer to coalesce and indeed caused an out-of-memory problem on a 32 GB laptop.

Finally, when $q \rightarrow +\infty$, Chandra et al. (2020) show how the posterior distribution under non repulsive mixture models either assigns all the observations to the same cluster or each observation to a separate cluster. These authors propose to consider mixtures in a latent space to overcome such issue, similarly to Ghahramani and Hinton (1996). Extensions of latent mixture models to account for repulsiveness are currently being investigated; see Ghilotti (2021).

C Removing the rectangular support assumption

Often we have assumed that the points of $\boldsymbol{\mu}$ have support given by a rectangular set R : For the theory in Sections 2–5, we made that assumption only for specificity and simplicity; in Section 7, we considered Gaussian mixture models and determined the rectangle R from the observations; while in Section 8, we considered the multivariate Bernoulli kernel and $R = [0, 1]^q$. Apart from the case of a DPP prior, it is often easy to modify everything without assuming R is rectangular and even compactness of R may not be needed. In fact, the birth-death Metropolis-Hastings algorithm, which we always use to simulate the non-allocated process $\boldsymbol{\mu}^{(na)}$, can be specified in a very general setting, see Geyer and Møller (1994). On the other hand, for a DPP prior, compactness of R is needed when specifying a DPP density with respect to $d\boldsymbol{\mu}$, and R needs to be a rectangle in order to use the spectral approach discussed in Lavancier et al. (2015). Recently, Poinas and Lavancier (2021) proposed a novel approximation of a general DPP density that does not require R to be rectangular (but still requires R is bounded).

References

- Argiento, R., Bianchini, I., and Guglielmi, A. (2016). “Posterior sampling from ε -approximation of normalized completely random measure mixtures.” *Electronic Journal of Statistics*, 10(2), 3516–3547.
- Argiento, R. and De Iorio, M. (2019). “Is infinity that far? A Bayesian nonparametric perspective of finite mixture models.” Technical report, available at arXiv:1904.09733.
- Bardenet, R. and Titsias, M. (2015). “Inference for determinantal point processes without spectral knowledge.” In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Gar-

- nett, R. (eds.), *Advances in Neural Information Processing Systems 28*, 3393–3401. Curran Associates, Inc.
- Bianchini, I., Guglielmi, A., and Quintana, F. A. (2020). “Determinantal point process mixtures via spectral density approach.” *Bayesian Analysis*, 15, 187–214.
- Binder, D. A. (1978). “Bayesian cluster analysis.” *Biometrika*, 65(1), 31–38.
- Biscio, C. A. N., Lavancier, F., et al. (2016). “Quantifying repulsiveness of determinantal point processes.” *Bernoulli*, 22(4), 2001–2028.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). “Latent Dirichlet allocation.” *Journal of Machine Learning Research*, 3(Jan), 993–1022.
- Chandra, N. K., Canale, A., and Dunson, D. B. (2020). “Escaping the curse of dimensionality in Bayesian model based clustering.” Technical report, available at arXiv:2006.02700.
- Collins, L. M. and Lanza, S. T. (2009). *Latent Class and Latent Transition Analysis: With Applications in the Social, Behavioral, and Health Sciences*. John Wiley & Sons, New York.
- Corradin, R., Canale, A., and Nipoti, B. (2020). *BNPmix: Bayesian Nonparametric Mixture Models*. R package version 0.2.7.
- Dellaportas, P. and Papageorgiou, I. (2006). “Multivariate mixtures of normals with unknown number of components.” *Statistics and Computing*, 16(1), 57–68.
- Favaro, S., Hadjicharalambous, G., and Prünster, I. (2011). “On a class of distributions on the simplex.” *Journal of Statistical Planning and Inference*, 141(9), 2987–3004.
- Fraley, C. and Raftery, A. E. (2007). “Bayesian regularization for normal mixture estimation and model-based clustering.” *Journal of Classification*, 24(2), 155–181.
- Fruhwirth-Schnatter, S., Celeux, G., and Robert, C. P. (2019). *Handbook of Mixture Analysis*. Chapman and Hall/CRC, New York.
- Fúquene, J., Steel, M., and Rossell, D. (2019). “On choosing mixture components via non-local priors.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(5), 809–837.
- Geyer, C. J. and Møller, J. (1994). “Simulation procedures and likelihood inference for spatial point processes.” *Scandinavian Journal of Statistics*, 21, 359–373.
- Ghahramani, Z. and Hinton, G. E. (1996). “The EM algorithm for mixtures of factor analyzers.” Technical report, CRG-TR-96-1, University of Toronto.
- Ghilotti, L. (2021). “Bayesian clustering of high-dimensional data via latent repulsive mixtures.” Master’s thesis, Politecnico di Milano. Available upon request.
- Green, P. J. (1995). “Reversible jump Markov chain Monte Carlo computation and Bayesian model determination.” *Biometrika*, 82(4), 711–732.
- (2010). “Trans-dimensional Markov chain Monte Carlo.” In Green, P. J., Hjort, N. L., and Richardson, S. (eds.), *Highly Structured Stochastic Systems*, 179–198. Oxford University Press, Oxford U.K.

- Guha, A., Ho, N., and Nguyen, X. (2019). “On posterior contraction of parameters and interpretability in Bayesian mixture modeling.” Technical report, available at arXiv:1901.05078.
- Hough, J. B., Krishnapur, M., Peres, Y., and Virág, B. (2006). “Determinantal processes and independence.” *Probability Surveys*, 3, 206–229.
- Hough, J. B., Krishnapur, M., Peres, Y., and Virág, B. (2009). *Zeros of Gaussian Analytic Functions and Determinantal Point Processes*. Providence: American Mathematical Society.
- James, L. F., Lijoi, A., and Prünster, I. (2009). “Posterior analysis for normalized random measures with independent increments.” *Scandinavian Journal of Statistics*, 36(1), 76–97.
- Kendall, W. S. and Møller, J. (2000). “Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes.” *Advances in Applied Probability*, 32, 844–865.
- Kleijn, B. J., van der Vaart, A. W., et al. (2006). “Misspecification in infinite-dimensional Bayesian statistics.” *The Annals of Statistics*, 34(2), 837–877.
- Kriegel, H.-P., Kröger, P., and Zimek, A. (2009). “Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering.” *ACM transactions on knowledge discovery from data*, 3(1), 1–58.
- Lavancier, F., Møller, J., and Rubak, E. (2015). “Determinantal point process models and statistical inference.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(4), 853–877.
- Li, Y., Lord-Bessen, J., Shiyko, M., and Loeb, R. (2018). “Bayesian latent class analysis tutorial.” *Multivariate Behavioral Research*, 53(3), 430–451.
- Liang, F., Jin, I. H., Song, Q., and Liu, J. S. (2016). “An adaptive exchange algorithm for sampling from distributions with intractable normalizing constants.” *Journal of the American Statistical Association*, 111(513), 377–393.
- Lijoi, A. and Prünster, I. (2010). “Models beyond the Dirichlet process.” In Hjort, N., Holmes, C., Müller, P., and Walker, S. (eds.), *Bayesian Nonparametrics*, 80–136. Cambridge University Press, Cambridge.
- Lyne, A.-M., Girolami, M., Atchadé, Y., Strathmann, H., Simpson, D., et al. (2015). “On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods.” *Statistical Science*, 30(4), 443–467.
- Macchi, O. (1975). “The coincidence approach to stochastic point processes.” *Advances in Applied Probability*, 7, 83–122.
- Miller, J. W. and Dunson, D. B. (2019). “Robust Bayesian inference via coarsening.” *Journal of the American Statistical Association*, 114(527), 1113–1125.
- Miller, J. W. and Harrison, M. T. (2018). “Mixture models with a prior on the number of components.” *Journal of the American Statistical Association*, 113(521), 340–356.
- Molitor, J., Papathomas, M., Jerrett, M., and Richardson, S. (2010). “Bayesian profile regression with an application to the National Survey of Children’s Health.” *Biostatistics*, 11(3), 484–498.

- Møller, J. and O'Reilly, E. (2021). “Couplings for determinantal point processes and their reduced Palm distributions with a view to quantifying repulsiveness.” *Advances in Applied Probability* (to appear). Available at arXiv:1806.07347.
- Møller, J., Pettitt, A. N., Reeves, R., and Berthelsen, K. K. (2006). “An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants.” *Biometrika*, 93(2), 451–458.
- Møller, J. and Waagepetersen, R. P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall/CRC, Boca Raton.
- Müller, P. and Mitra, R. (2013). “Bayesian nonparametric inference – why and how.” *Bayesian Analysis*, 8, 269–302.
- Murray, I., Ghahramani, Z., and MacKay, D. J. C. (2006). “MCMC for doubly-intractable distributions.” In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, UAI’06, 359–366. Arlington, Virginia, USA: AUAI Press.
- Papaspiliopoulos, O. and Roberts, G. O. (2008). “Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models.” *Biometrika*, 95(1), 169–186.
- Petralia, F., Rao, V., and Dunson, D. B. (2012). “Repulsive mixtures.” In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 25*, 1889–1897. Curran Associates, Inc.
- Poinas, A. and Lavancier, F. (2021). “Asymptotic approximation of the likelihood of stationary determinantal point processes.” Technical report, available at arXiv:2103.02310.
- Propp, J. G. and Wilson, D. B. (1996). “Exact sampling with coupled Markov chains and applications to statistical mechanics.” *Random Structures & Algorithms*, 9(1-2), 223–252.
- Quinlan, J. J., Quintana, F. A., and Page, G. L. (2020). “Parsimonious hierarchical modeling using repulsive distributions.” *Test* (to appear).
- Regazzini, E., Lijoi, A., and Prünster, I. (2003). “Distributional results for means of normalized random measures with independent increments.” *The Annals of Statistics*, 31(2), 560–585.
- Richardson, S. and Green, P. J. (1997). “On Bayesian analysis of mixtures with an unknown number of components (with discussion).” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(4), 731–792.
- Xie, F. and Xu, Y. (2019). “Bayesian repulsive gaussian mixture model.” *Journal of the American Statistical Association*, 187–203.
- Xu, Y., Müller, P., and Telesca, D. (2016). “Bayesian inference for latent biologic structure with determinantal point processes (DPP).” *Biometrics*, 72(3), 955–964.