



Model Selection for Maternal Hypertensive Disorders with Symmetric Hierarchical Dirichlet Processes

Beatrice Franzolini, Antonio Lijoi, Igor Prünster

No. 672

April 2022

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Model Selection for Maternal Hypertensive Disorders with Symmetric Hierarchical Dirichlet Processes

Beatrice Franzolini¹ Antonio Lijoi² Igor Prünster²

¹ Singapore Institute for Clinical Sciences (SICS),

Agency for Science, Technology and Research (A*STAR), Singapore

² Department of Decision Sciences and BIDSa, Bocconi University Milan,
and Collegio Carlo Alberto, Torino, Italy.

Abstract

Hypertensive disorders of pregnancy occur in about 10% of pregnant women around the world. Though there is evidence that hypertension impacts maternal cardiac functions, the relation between hypertension and cardiac dysfunctions is only partially understood. The study of this relationship can be framed as a joint inferential problem on multiple populations, each corresponding to a different hypertensive disorder diagnosis, that combines multivariate information provided by a collection of cardiac function indexes. A Bayesian nonparametric approach seems particularly suited for this setup and we demonstrate it on a dataset consisting of transthoracic echocardiography results of a cohort of Indian pregnant women. We are able to perform model selection, provide density estimates of cardiac function indexes and a latent clustering of patients: these readily interpretable inferential outputs allow to single out modified cardiac functions in hypertensive patients compared to healthy subjects and progressively increased alterations with the severity of the disorder. The analysis is based on a Bayesian nonparametric model that relies on a novel hierarchical structure, called symmetric hierarchical Dirichlet process. This is suitably designed so that the mean parameters are identified and used for model selection across populations, a penalization for multiplicity is enforced, and the presence of unobserved relevant factors is investigated through a latent clustering of subjects. Posterior inference relies on a suitable Markov Chain Monte Carlo algorithm and the model behaviour is also showcased on simulated data.

Keywords— Bayesian nonparametrics, clustering populations, Dirichlet process, hierarchical partitions, hierarchical process, hypertensive disorders of pregnancy, model based clustering

1 Introduction

Hypertensive disorders of pregnancy are a class of high blood pressure disorders that occur during the second half of pregnancy, which include gestational hypertension, preeclampsia and severe preeclampsia. They are characterized by a diastolic blood pressure higher than 90 mm Hg and/or a systolic blood pressure higher than 140 mm Hg and they are

often accompanied by proteinuria. These disorders affect about 10% of pregnant women around the world, with preeclampsia occurring in 2–8% of all pregnancies (Timokhina et al., 2019). These disorders represent one of the leading causes of maternal and fetal morbidity and mortality, contributing to 7–8% of maternal death worldwide (Dolea and AbouZahr, 2003; Shah et al., 2009; McClure et al., 2009). The World Health Organization estimates that the incidence of preeclampsia is seven times higher in developing countries than in developed countries. However, the occurrence of these diseases appears under-reported in low and middle income countries, implying that the true incidence is unknown (Igberase and Ebeigbe, 2006; Malik et al., 2019). While there is evidence that hypertensive disorders of pregnancy are related with the development of cardiac dysfunctions both in the mother and in the child (Bellamy et al., 2007; Davis et al., 2012; Ambrožic et al., 2020; Garcia-Gonzalez et al., 2020; Aksu et al., 2021; deMartelly et al., 2021), there is no common agreement on the relation between the severity of hypertension and cardiac dysfunction (Tatapudi and Pasumathy, 2017b) and echocardiography is not included in baseline evaluation of hypertensive disorders of pregnancy. Further investigations on these disorders are needed, especially for developing countries, where women often give birth at a younger age with respect to developed countries.

The goal of this work is to detect which cardiac function is altered and under which hypertensive disorders by relying on a principled Bayesian nonparametric approach. An interesting case-control study to explore the relation between cardiac dysfunction and hypertensive disorders is provided by Tatapudi and Pasumathy (2017a), where the measures of ten different cardiac function indexes were recorded in four groups of pregnant women in India. Groups of women are characterized by different hypertensive disorder diagnoses, that are naturally ordered based on the severity of the diagnosed disorder: healthy (C), gestational hypertension (G), mild preeclampsia (M) and severe preeclampsia (S). Hypertensive diagnoses are used as identifiers for what we call populations of patients and we refer to cardiac function indexes also with the term response variables. For each response variable we want to determine a partition of the four populations of patients. This amounts to identifying similarities between different hypertensive disorders, with respect to each cardiac index. Supposing, for instance, that the selected partition assigns all the populations to the same cluster, one can conclude that no alteration is shown for the corresponding cardiac index across different hypertensive diseases.

Our goal of identifying a partition of the four patients’ populations for each of the ten responses can be rephrased as a problem of multiple model selection: we want to select the most plausible partition for each cardiac index. Frequentist hypothesis testing does not allow to deal with more than two populations in a straightforward way, and pairwise comparisons may lead to conflicting conclusions. Conversely, a Bayesian approach yields the posterior distribution on the space of partitions, which can be used for simultaneous comparisons. Moreover, the presence of $M = 10$ jointly tested cardiac indexes requires to perform model selection repeatedly ten times. Once again, a Bayesian approach seems

to be preferred, because, as observed for instance by [Scott and Berger \(2006\)](#), it does not require the introduction of a penalty term for multiple comparison, thanks to the prior distribution build-in penalty.

Here we design a Bayesian nonparametric model, that is tailored to deal with both a collection of ordered populations and the multivariate information of the response variables, while preserving the typical flexibility of nonparametric models and producing easily interpretable results. When applied to the dataset on transthoracic echocardiography results for a cohort of Indian pregnant women in [Section 5](#), our model effectively identifies modified cardiac functions in hypertensive patients compared to healthy subjects and progressively increased alterations with the severity of the disorder, in addition to other more subtle findings. The observed data $X_{i,j,m}$ represent the measurement of the m -th response variable (cardiac index) on the i -th individual (pregnant woman) in the j -th population (hypertensive disorder) and, as in standard univariate ANOVA models, they are assumed to be partially exchangeable across disorders. This means that for every $m \in \{1, \dots, M\}$, the law of $((X_{i,1,m})_{i \geq 1}, \dots, (X_{i,J,m})_{i \geq 1})$ is invariant with respect to permutations within each sequence of random variables, namely for any positive integers n_1, \dots, n_J

$$((X_{i,1,m})_{i=1}^{n_1}, \dots, (X_{i,J,m})_{i=1}^{n_J}) \stackrel{d}{=} ((X_{\sigma_1(i),1,m})_{i=1}^{n_1}, \dots, (X_{\sigma_J(i),J,m})_{i=1}^{n_J})$$

for all permutations σ_j of $(1, \dots, n_j)$, with $j = 1, \dots, J$. This is a natural generalization of exchangeability to tackle heterogeneous data and, by de Finetti's representation theorem, it amounts to assuming the existence of a collection of (possibly dependent) random probability measures $\{\pi_{j,m} : j = 1, \dots, J \ m = 1, \dots, M\}$ such that

$$X_{i,j,m} \mid \pi_{j,m} \stackrel{\text{iid}}{\sim} \pi_{j,m} \quad i = 1, \dots, n_j$$

Hence, for any two populations $j \neq j'$, homogeneity corresponds to $\pi_{j,m} = \pi_{j',m}$ (almost surely). However, a reliable assessment of this type of homogeneity is troublesome when having just few patients per diagnosis, as it happens in the mild preeclampsia subsample. Without relying on simplifying parametric assumptions, a small sub-sample size may not be sufficiently informative to infer equality of entire unknown distributions. To overcome this issue, without introducing parametric assumptions, we resort to an alternative weaker notion of homogeneity between populations j and j' : we only require the conditional means of the two populations to (almost surely) coincide

$$\mathbb{E}(X_{i,j,m} \mid \pi_{j,m}) = \mathbb{E}(X_{i,j',m} \mid \pi_{j',m}). \quad (1)$$

According to this definition, the detection of heterogeneities in cardiac function indexes amounts to inferring which cardiac indexes have means that differ across diagnoses, as it is done in standard parametric ANOVA models. Besides clustering populations according to [\(1\)](#), it is also of interest to cluster patients, both within and across different groups,

once the effect of the specific hypertensive disorder is taken into account. This task may be achieved by assuming a model that decomposes the observations as

$$X_{i,j,m} = \theta_{j,m} + \varepsilon_{i,j,m} \quad \varepsilon_{i,j,m} | (\xi_{i,j,m}, \sigma_{i,j,m}^2) \stackrel{\text{ind}}{\sim} \text{N}(\xi_{i,j,m}, \sigma_{i,j,m}^2) \quad (2)$$

and the $\xi_{i,j,m}$ have a symmetric distribution around the origin, in order to ensure $E(\xi_{i,j,m}) = 0$. In view of this decomposition, we will let $\theta_{j,m}$ govern the clustering of populations while the $(\xi_{i,j,m}, \sigma_{i,j,m}^2)$'s determine the clustering of individuals, namely patients, after removing the effect of the specific hypertensive disorder. In order to pursue this, for each cardiac index m , we will specify a hierarchical process prior for $(\xi_{i,j,m}, \sigma_{i,j,m}^2)$ that is suited to infer the clustering structure both within and across different hypertensive disorders for a specific cardiac index. In particular, we will deploy a novel instance of hierarchical Dirichlet process, introduced in [Teh et al. \(2006\)](#), that we name *symmetric*, to highlight its centering in 0.

Early examples of Bayesian nonparametric models for ANOVA can be found in [Cifarelli and Regazzini \(1978\)](#) and [Muliere and Petrone \(1978\)](#), while the first popular proposal, due to [De Iorio et al. \(2004\)](#), uses the dependent Dirichlet process (DDP) ([MacEachern, 2000](#)) and is therefore termed ANOVA-DDP. This model is mainly tailored to estimate populations' probability distributions, while we draw inferences over clusters of populations' means and obtain estimates of the unknown distributions as a by-product. Moreover, the ANOVA-DDP of [De Iorio et al. \(2004\)](#) was not introduced as a model selection procedure. A popular Bayesian nonparametric model, that does cluster populations and can be used for model selection, is the nested Dirichlet process of [Rodriguez, Dunson, and Gelfand \(2008\)](#). As shown in [Camerlenghi et al. \(2019a\)](#), such a prior is biased towards homogeneity, in the sense that even a single tie between populations j and j' , namely $X_{i,j,m} = X_{i',j',m}$ for some i and i' , entails $\pi_{j,m} = \pi_{j',m}$ (almost surely). In order to overcome such a drawback, a novel class of nested, and more flexible, priors has been proposed in [Camerlenghi et al. \(2019a\)](#). See also [Soriano and Ma \(2017\)](#) for related work. Interesting alternatives that extend the analysis to more than two populations can be found in [Christensen and Ma \(2020\)](#), [Lijoi, Prünster, and Rebaudo \(2022\)](#) and in [Beraha, Guglielmi, and Quintana \(2021\)](#). Another similar proposal is the one by [Gutiérrez et al. \(2019\)](#), whose model identifies differences over cases' distributions and the control group. These models imply that two populations belong to the same cluster if they share the entire distribution. However, as already mentioned, distribution-based clustering is not ideal when dealing with scenarios as the one of hypertensive dataset. Further evidence will be provided in [Section 5.1](#), through simulation studies. In addition, note that all these contributions deal with only one response variable and would need to be suitably generalized to fit the setup of this paper. As far as the contributions treating multiple response variables are concerned, uses of nonparametric priors for multiple testing can be found, for instance, in [Gopalan and Berry \(1998\)](#), [Do, Müller, and Tang \(2005\)](#), [Dahl and Newton \(2007\)](#), [Guindani, Müller, and Zhang \(2009\)](#), [Martin and Tokdar \(2012\)](#) and

more recently in [Cipolli, Hanson, and McLain \(2016\)](#), who propose an approximate finite Pólya tree multiple testing procedure to compare two-samples' locations, and in [Denti et al. \(2020\)](#). However, in all these contributions, models are developed directly over summaries of the original data (e.g. averages, z-scores) and, as such, do not allow to draw any inference on the entire distributions and clusters of subjects.

The outline of the paper is as follows. In [Section 2](#) we introduce the model, which makes use of an original hierarchical prior structure for symmetric distributions ([Section 2.2](#)). In [Section 3](#) we derive the prior law of the random partitions induced by the model, key ingredient for the Gibbs sampling scheme devised in [Section 4](#). In [Section 5](#), we first present a series of simulation studies that highlight the behaviour of the model before applying it to obtain our results on cardiac dysfunction in hypertensive disorders. [Section 6](#) contains some concluding remarks. As Supplementary Material we provide the datasets and Python codes, some further background material and details about the derivation of the posterior sampling scheme as well as additional simulation studies and results on the application, including an analysis of prior sensitivity.

2 The Bayesian nonparametric model

The use of discrete nonparametric priors for Bayesian model-based clustering has become standard practice. The Dirichlet process (DP) ([Ferguson, 1973](#)) is the most popular instance, and clustering is typically addressed by resorting to a mixture model, which with our data structure amounts to

$$X_{i,j,m} | \psi_{i,j,m} \stackrel{\text{ind}}{\sim} k(X_{i,j,m}; \psi_{i,j,m}), \quad \psi_{i,j,m} | \tilde{p}_{j,m} \stackrel{\text{ind}}{\sim} \tilde{p}_{j,m}$$

for $m = 1, \dots, M$, $j = 1, \dots, J$ and $i = 1, \dots, n_j$. Here $k(\cdot; \cdot)$ is some kernel and the $\tilde{p}_{j,m}$'s are discrete random probability measures. Hence, the $\psi_{i,j,m}$'s may exhibit ties. The model specification for $\tilde{p}_{j,m}$ will be tailored to address the following goals: (i) cluster the J probability distributions based on their means; (ii) cluster the observations $X_{i,j,m}$ according to the ties induced on the $\psi_{i,j,m}$'s by the $\tilde{p}_{j,m}$'s for a given fixed j and across different j 's. These two issues will be targeted separately: we first design a clustering scheme for the populations, through the specification of a prior on the means of the $X_{i,j,m}$'s and, then, we cluster the data using a hierarchical DP having a specific invariance structure that is ideally suited to the application at hand.

2.1 The prior on disease-specific locations

As a model for the observations we consider a nonparametric mixture of Gaussian distributions specified as

$$X_{i,j,m} | (\boldsymbol{\theta}_m, \boldsymbol{\xi}_m, \boldsymbol{\sigma}_m^2) \stackrel{\text{ind}}{\sim} \mathcal{N}(\theta_{j,m} + \xi_{i,j,m}, \sigma_{i,j,m}^2) \quad (3)$$

where $\boldsymbol{\theta}_m = (\theta_{1,m}, \dots, \theta_{J,m})$, $\boldsymbol{\xi}_m = (\xi_{1,1,m}, \dots, \xi_{1,n_1,m}, \xi_{2,1,m}, \dots, \xi_{n_J,J,m})$, with a similar definition for the vector $\boldsymbol{\sigma}_m^2$, and $\mathcal{N}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2 . The assumption in (3) clearly reflects (2). Moreover, in order to account for the two levels of clustering we are interested in, we will assume that

$$(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M) \sim P, \quad (\xi_{i,j,m}, \sigma_{i,j,m}^2) | \tilde{q}_{j,m} \stackrel{\text{iid}}{\sim} \tilde{q}_{j,m} \quad (i = 1, \dots, n_j) \quad (4)$$

where $\tilde{q}_{1,m}, \dots, \tilde{q}_{J,m}$ are discrete random probability measures independent from $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$. Thus, the likelihood corresponds to

$$\prod_{m=1}^M \prod_{j=1}^J \prod_{i=1}^{n_j} \frac{1}{\sigma_{i,j,m}} \varphi\left(\frac{x_{i,j,m} - \theta_{j,m} - \xi_{i,j,m}}{\sigma_{i,j,m}}\right) \tilde{q}_{j,m}(\mathrm{d}\xi_{i,j,m}, \mathrm{d}\sigma_{i,j,m}) \quad (5)$$

with φ denoting the standard Gaussian density. Relevant inferences can be carried out if one is able to marginalize this expression with respect to both $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$ and $(\tilde{q}_{1,m}, \dots, \tilde{q}_{J,m})$ for each $m = 1, \dots, M$.

This specification allows to address the model selection problem in the following way. If \mathcal{M}^m stands for the space of all partitions of the J populations for the m -th cardiac function index, then $\mathcal{M}^m = \{M_b^m : b = 1, \dots, \text{card}(\mathcal{P}_J)\}$ where \mathcal{P}_J is the collection of all possible partitions of $[J] = \{1, \dots, J\}$. In our specific case, $J = 4$ and $\text{card}(\mathcal{P}_J) = 15$, thus we have 15 competing models per cardiac index. Each competing model corresponds to a specific partition in \mathcal{M}^m . In particular, the partition arises from ties between the population specific means in $\boldsymbol{\theta}_m$ and, hence, the distribution P in (4) needs to associate positive probabilities to ties between the parameters within the vector $\boldsymbol{\theta}_m$, for each $m = 1, \dots, M$.

Let us start considering as distribution P a well-known effective clustering prior, i.e. a mixture of DPs in the spirit of [Antoniak \(1974\)](#), namely

$$\begin{aligned} \theta_{j,m} | \tilde{p}_m &\stackrel{\text{iid}}{\sim} \tilde{p}_m & j = 1, \dots, J \\ \tilde{p}_m | \omega &\stackrel{\text{iid}}{\sim} \text{DP}(\omega, G_m) & m = 1, \dots, M \\ \omega &\sim p_\omega \end{aligned} \quad (6)$$

where $\text{DP}(\omega, G_m)$ denotes the DP with concentration parameter ω and non-atomic baseline probability measure G_m and p_ω is a probability measure on \mathbb{R}^+ . The discreteness of the DP implies the presence (with positive probability) of ties within the vector of locations $\boldsymbol{\theta}_m$ associated to a certain cardiac index m , as desired. The ties give rise to a random partition: as shown in [Antoniak \(1974\)](#), the probability of observing a specific partition of the elements in $\boldsymbol{\theta}_m$ consisting of $k \leq J$ distinct values with respective frequencies n_1, \dots, n_k coincides with

$$\Pi_k^{(J)}(n_1, \dots, n_k) = \frac{\omega^k}{(\omega)_J} \prod_{i=1}^k (n_i - 1)! \quad (7)$$

where $(\omega)_J = \Gamma(\omega + J)/\Gamma(\omega)$. The use of a shared concentration parameter over (7) to address multiple model selection has been already successfully employed in Moser, Rodríguez and Lofland (2021), where they cluster parameters in a probit model. When there is no pre-experimental information available on competing partitions, the use of (7) as a prior for model selection has some relevant benefits. Indeed, it induces borrowing of strength across diagnoses and, being ω random, it generates borrowing of information also across cardiac indexes, thus improving the Bayesian learning mechanism. These two features can also be given a frequentist interpretation in terms of desirable penalties. As a matter of fact, the procedure penalizes for the multiplicity of the model selections that are performed. The penalty has to be meant in the following way: while J and/or M increase, the prior odds change in favor of less complex models. For more details on this, see Scott and Berger (2010). Summing up, the mixture of DPs automatically induces a prior distribution on $\{\mathcal{M}^m : m = 1, \dots, M\}$, that arises from (7) combined with the prior p_ω on ω , and it presents desirable properties for model selection that can be interpreted either in terms of borrowing of information or in terms of penalties.

However, in the analysis of hypertensive disorders, some prior information on competing models is available, and this is not yet incorporated in (7). In fact, as already mentioned, there is a natural order of the diagnoses, which is given by the severity of the disorders, i.e. C, G, M, S. Partitions that do not comply with this ordering, e.g. $\{\{C, S\}\{G\}, \{M\}\}$, should be excluded from the support of the prior. Thus, we consider a prior over \mathcal{M}^m that associates zero probability to partitions that do not respect the natural order of the diagnoses and a probability proportional to that in (7) for the remaining partitions, i.e.

$$\mathbb{P}(M_b^m \mid \omega) \propto \begin{cases} \Pi_k^{(J)}(n_1, \dots, n_k) & \text{if } M_b^m \text{ is compatible with the natural order} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

This amounts to a distribution P for $(\theta_1, \dots, \theta_M)$ given by

$$\begin{aligned} (\theta_{1,m}, \dots, \theta_{J,m}) \mid \omega &\stackrel{\text{ind}}{\sim} P_{\omega, G_m} & m = 1, \dots, M \\ \omega &\sim p_\omega \end{aligned} \quad (9)$$

where P_{ω, G_m} is the distribution obtained sampling a partition according to (8) and associating to each cluster a unique value sampled from G_m . Using (9) as a prior for the disease-specific locations, we preserve the desirable properties of the mixture of DPs mentioned before, while incorporating prior information on the severity of the diseases.

As detailed in the next section, we further define random probability measures $\tilde{q}_{j,m}$ that satisfy the symmetry condition

$$\tilde{q}_{j,m}(A \times B) = \tilde{q}_{j,m}((-A) \times B) \quad \text{a.s.} \quad (10)$$

for any A and B . This condition ensures that the parameters $\theta_{j,m}$, for $j = 1, \dots, J$ and $m = 1, \dots, M$, in (3) are identified, namely $\mathbb{E}(X_{i,j,m} \mid \boldsymbol{\theta}_m, \tilde{q}_{j,m}) = \theta_{j,m}$ with probability one. This identifiability property is crucial to make inference over the location parameters $\boldsymbol{\theta}_m$'s. Similar model specifications for discrete exchangeable data have been proposed and studied in Dalal (1979b), Doss (1984), Diaconis and Freedman (1986) and Ghosal, Ghosh, and Ramamoorthi (1999), of which (5) represents a generalization to density functions and partially exchangeable data.

2.2 The prior for the error terms

While the clustering of populations is governed by (8), we use a mixture of hierarchical discrete processes for the error terms. This has the advantage of modeling the clustering of the observations, both within and across different samples, once the disease-specific effects are account for. This clustering structure allows to model heterogeneity across patients in a much more realistic way with respect to standard ANOVA models based on assumption of normality. Cardiac indexes may be influenced by a number of factors that are not directly observed in the study, such as pre-existing conditions (Hall, George, and Granger, 2011) and psychosocial factors (Pedersen et al., 2017). These unobserved relevant factors may be shared across patients with the same or a different diagnosis and may also result in outliers. To take into account this latent heterogeneity of the data, we introduce the hierarchical symmetric DP that satisfies the symmetry condition in (10) and, moreover, allows to model heterogeneous data similarly to the hugely popular hierarchical DP (Teh et al., 2006).

The basic building block of the proposed prior is the invariant Dirichlet process, which was introduced for a single population ($J = 1$) in an exchangeable framework by Dalal (1979a). Such a modification of the DP satisfies a symmetry condition, in the sense that it is a random probability measure that is invariant with respect to a chosen group of transformations \mathcal{G} . A more formal definition and detailed description of the invariant DP can be found in Section A of the Supplement. For our purposes it is enough to consider the specific case of the symmetric Dirichlet process, which can be constructed through a symmetrization of a Dirichlet process. Consider a non-atomic probability measure P_0 on \mathbb{R} and let $\tilde{Q}_0 \sim \text{DP}(\alpha, P_0)$. If

$$\tilde{Q}(A) = \frac{\tilde{Q}_0(A) + \tilde{Q}_0(-A)}{2} \quad \forall A \in \mathcal{B}(\mathbb{R}) \quad (11)$$

where $-A = \{x \in \mathbb{R} : -x \in A\}$, then \tilde{Q} is symmetric about 0 (almost surely) and termed symmetric DP, in symbols $\tilde{Q} \sim \text{s-DP}(\alpha, P_0)$. For convenience and without loss of generality, we assume that P_0 is symmetric: this implies that P_0 is the expected value of \tilde{Q} making it an interpretable parameter. The random probability measure \tilde{Q} is the basic building block of the hierarchical process that we use to model the heterogeneity

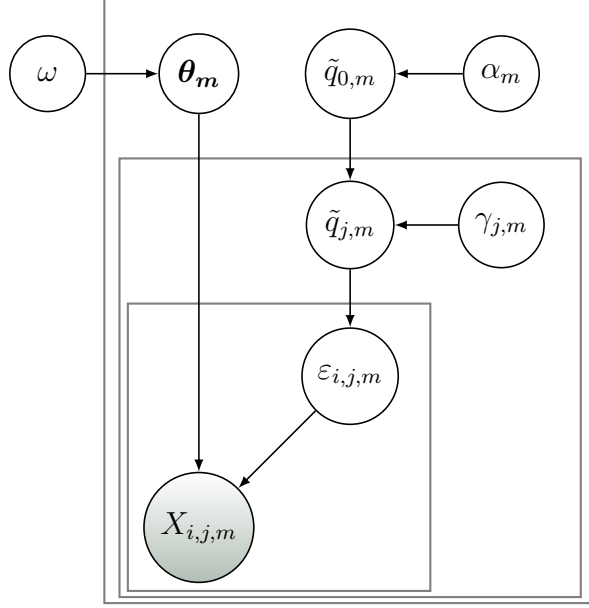


Figure 1: Graphical representation of the model. Each node represents a random variable and each rectangle denotes conditional i.i.d. replications of the model within the rectangle.

of the error terms across different populations, $j = 1, \dots, J$, in such a way that clusters identified by the unique values can be shared within and across populations. This prior is termed *symmetric hierarchical Dirichlet process* (s-HDP) and described as

$$\begin{aligned} \tilde{q}_{j,m} \mid \gamma_{j,m}, \tilde{q}_{0,m} &\stackrel{\text{ind}}{\sim} \text{s-DP}(\gamma_{j,m}, \tilde{q}_{0,m}) \\ \tilde{q}_{0,m} \mid \alpha_m &\stackrel{\text{ind}}{\sim} \text{s-DP}(\alpha_m, P_{0,m}) \end{aligned} \quad (12)$$

where $\gamma_{j,m}$ and α_m are positive parameters and $P_{0,m}$ is a non-atomic probability distribution symmetric about 0. We use the notation $(\tilde{q}_{1,m}, \dots, \tilde{q}_{J,m}) \sim \text{s-HDP}(\boldsymbol{\gamma}_m, \alpha_m, P_{0,m})$, where $\boldsymbol{\gamma}_m = (\gamma_{1,m}, \dots, \gamma_{J,m})$. This definition clearly ensures the validity of (10). A graphical model representation of the over-all proposed model is displayed in Figure 1.

Still referring to the decomposition of the observations into disease-specific locations and an error term, i.e. $X_{i,j,m} = \theta_{j,m} + \varepsilon_{i,j,m}$, it turns out that the $\varepsilon_{i,j,m}$'s are from a symmetric hierarchical DP mixture (s-HDP mixture) with a normal kernel. Hence, the patients' clusters are identified through the $\varepsilon_{i,j,m}$, which, according to (3), are conditionally independent from a $\mathcal{N}(\xi_{i,j,m}, \sigma_{i,j,m}^2)$ given $(\xi_{i,j,m}, \sigma_{i,j,m}^2)$. The choice of the specific invariant DP is aimed at ensuring that $\mathbb{E}(\varepsilon_{i,j,m} \mid \tilde{q}_{j,m}) = 0$. The clusters identified by the s-HDP mixture can be interpreted as representing common unobserved factors across patients, once the disease-specific locations have been accounted for. Indeed, for any pair of patients, we may consider the decomposition $X_{i,j,m} - X_{i',j',m} = \Delta_{\theta}^{(m)} + \Delta_{\xi}^{(m)} + (e_{i,j,m} - e_{i',j',m})$ where $\Delta_{\theta}^{(m)} = \theta_{j,m} - \theta_{j',m}$, $\Delta_{\xi}^{(m)} = \xi_{i,j,m} - \xi_{i',j',m}$ and $e_{i,j,m}$ and $e_{i',j',m}$ are independent and normally distributed random variables with zero mean and variances $\sigma_{i,j,m}^2$ and $\sigma_{i',j',m}^2$, respectively.

Hence, patients' clustering reflects the residual heterogeneity that is not captured by

the disease-specific component $\Delta_\theta^{(m)}$ and are related to the subject-specific locations $\Delta_\xi^{(m)}$ and to the zero-mean error component $(e_{i,j,m} - e_{i',j',m})$. In view of this interpretation, using a s-HDP mixture over error terms offers a three-fold advantage. Firstly, the presence of clearly separated clusters of patients within and across populations will indicate the presence of unobserved relevant factors which affect the cardiac response variables. Secondly, single patients with very low probabilities of co-clustering with all other subjects will have to be interpreted as outliers. Finally, the estimated clustering structure can also be used to check whether the relative effect of a certain disease (with respect to another) is fully explained by the corresponding $\Delta_\theta^{(m)}$. To clarify this last point consider two diseases: if the posterior co-clustering probabilities among patients sharing the same disease are different between the two populations, this will indicate that different diagnoses not only have an influence on disease-specific locations (which is measured by $\Delta_\theta^{(m)}$), but they also have an impact on the shape of the distribution of the corresponding cardiac index. More details on this can be found in Section D of the Supplement.

3 Marginal distributions and random partitions

As emphasized in the previous sections, ties among the $\theta_{j,m}$'s and the $(\xi_{i,j,m}, \sigma_{i,j,m}^2)$'s are relevant for inferring the clustering structure both among the populations (hypertensive diseases) and among the individual units (patients). Indeed, for each m (cardiac index) they induce a random partition that emerges as a composition of two partitions generated respectively by the prior in (9) and the s-HDP. The laws of these random partitions are not only crucial to understand the clustering mechanism, but also necessary in order to derive posterior sampling schemes. In this section such a law is derived and used to compute the predictive distributions that, jointly with the likelihood, determine the full conditionals of the Gibbs sampler devised in Section 4. To reduce the notational burden, in this and the following section, we remove the dependence of observations and parameters on the specific response variable m and denote with $\phi_{i,j}$ the pair $(\xi_{i,j}, \sigma_{i,j}^2)$ and with ϕ the collection $(\phi_{1,1}, \dots, \phi_{1,n_1}, \phi_{2,1}, \dots, \phi_{n_J,J})$.

Conditionally on ω , the law of the partition in (8) leads to the following predictive distribution for the disease-specific locations

$$\theta_j \mid \omega, \theta_1, \dots, \theta_{j-1} \sim a_j(\omega, \theta_1, \dots, \theta_{j-1}) \delta_{\theta_{j-1}} + [1 - a_j(\omega, \theta_1, \dots, \theta_{j-1})] G$$

where

$$a_j(\omega, \theta_1, \dots, \theta_{j-1}) = \frac{\sum_{(*)} \Pi_k^{(J)}(n_1, \dots, n_k)}{\sum_{(\Delta_j)} \Pi_k^{(J)}(n_1, \dots, n_k)} \quad (13)$$

where the sum at the denominator runs over the set of partitions consistent with the one generated by $(\theta_1, \dots, \theta_{j-1})$ and the one at the numerator runs over a subset of those

partitions where one further has $\theta_j = \theta_{j-1}$. For $j = 4$, the predictive equals

$$\theta_4 \mid \omega, \theta_1, \theta_2, \theta_3 \begin{cases} \frac{3}{\omega+3}\delta_{\theta_3} + \frac{\omega}{\omega+3}G & \text{if } \theta_1 = \theta_2 = \theta_3 \\ \frac{2}{\omega+2}\delta_{\theta_3} + \frac{\omega}{\omega+2}G & \text{if } \theta_1 \neq \theta_2 = \theta_3 \\ \frac{1}{\omega+1}\delta_{\theta_3} + \frac{\omega}{\omega+1}G & \text{otherwise} \end{cases}$$

Explicit expressions for the function a , for $j = 1, 2, 3$, can be easily computed using (13) and (8) and are provided in Section B of the Supplement.

Moving to second-level partitions induced by the s-HDP, we recall that the key concept for studying random partitions on multi-sample data is the *partially exchangeable partition probability function* (pEPPF). See, e.g., Lijoi, Nipoti, and Prünster (2014) and Camerlenghi et al. (2019b). The pEPPF returns the probability of a specific multi-sample partition and represents the appropriate generalization of the well-known single-sample EPPF, which in the DP case corresponds to (7). Discreteness of the s-HDP $(\tilde{q}_1, \dots, \tilde{q}_m)$ in (12) induces a partition of the elements of ϕ into equivalence classes identified by the distinct values. Taking into account the underlying partially exchangeable structure, such a random partition is characterized by the pEPPF

$$\tilde{\Pi}_k^{(N)}(\mathbf{n}_1, \dots, \mathbf{n}_J) = \mathbb{E} \left(\int_{\Phi^k} \prod_{j=1}^J \prod_{h=1}^k \tilde{q}_{j,m}^{n_{j,h}}(d\phi_i) \right) \quad (14)$$

where $\mathbf{n}_j = (n_{j,1}, \dots, n_{j,k})$ are non-negative integers, for any $j = 1, \dots, J$, such that $n_{j,h}$ is the number of elements in ϕ corresponding to population j and belonging to cluster h . Thus $\sum_{j=1}^J n_{j,h} \geq 1$ for any $h = 1, \dots, k$, $\sum_{h=1}^k n_{j,h} = n_j$ and $\sum_{h=1}^k \sum_{j=1}^J n_{j,h} = N$. The determination of probability distributions of this type is challenging and only recently the first explicit instances have appeared in the literature. See e.g., Lijoi et al. (2014), Camerlenghi et al. (2019a) and Camerlenghi et al. (2019b). With respect to the hierarchical case considered in Camerlenghi et al. (2019b), the main difference is that here we have to take into account the specific structure (11) of the $\tilde{q}_{j,m}$. The almost sure symmetry of the process generates a natural random matching between sets in the induced partition. Therefore, instead of studying the marginal law in (14), we derive the joint law of the partition and of the random matching. Formally, consider a specific partition $\{A_1^+, A_1^-, \dots, A_k^+, A_k^-\}$ of ϕ , such that, for $h = 1, \dots, k$, all the elements in A_h^+ belong to $\mathbb{R}^+ \times \mathbb{R}^+$, all the elements in A_h^- belong to $\mathbb{R}^- \times \mathbb{R}^+$ and, if $\phi_{i,j} \in A_h^+$ and $\phi_{i',j'} \in A_h^-$, then the element-wise absolute values of $\phi_{i,j}$ and $\phi_{i',j'}$ are equal. Denote with $n_{j,h}^+$ the number of elements in $A_h^+ \cap \{\phi_{i,j}, i = 1, \dots, n_j\}$ and with $n_{j,h}^-$ the number of elements in $A_h^- \cap \{\phi_{i,j}, i = 1, \dots, n_j\}$. The probability of observing $\{A_1^+, A_1^-, \dots, A_k^+, A_k^-\}$ is

$$\tilde{\Pi}_k^{(N)}(\mathbf{n}_1^+, \mathbf{n}_1^-, \dots, \mathbf{n}_J^+, \mathbf{n}_J^-) = \mathbb{E} \left(\int_{\Phi^k} \prod_{j=1}^J \prod_{h=1}^k \tilde{q}_{j,m}^{n_{j,h}^+ + n_{j,h}^-}(d\phi) \right) \quad (15)$$




Figure 2: Chinese restaurant franchise representation of the symmetric hierarchical DP for $J = 2$ populations. Each circle represents a table.

with $\mathbf{n}_j^+ = (n_{j,1}^+, \dots, n_{j,k}^+)$. As for the determination of (15), a more intuitive understanding may be gained if one considers its corresponding Chinese restaurant franchise (CRF) metaphor, which displays a variation of both the standard Chinese restaurant franchise of Teh et al. (2006) and the skewed Chinese restaurant process of Iglesias, Orellana, and Quintana (2009). Figure 2 provides a graphical representation. The scheme is as follows: there are J restaurants sharing the same menu and the customers are identified by their choice of $\phi_{i,j}$ but, unlike in the usual CRF, at each table two *symmetric dishes* are served. Denote with $\phi_{t,j}^* = (\xi_{t,j}^*, \sigma_{t,j}^{2*})$ and $-\phi_{t,j}^* = (-\xi_{t,j}^*, \sigma_{t,j}^{2*})$ the two dishes served at table t in restaurant j , with $\phi_h^{**} = (\xi_h^{**}, \sigma_h^{**2})$ and $-\phi_h^{**} = (-\xi_h^{**}, \sigma_h^{**2})$ the h -th pair of dishes in the menu and with $n_{j,h}^+$ and $n_{j,h}^-$ the number of customers in restaurant j eating dish ϕ_h^{**} and $-\phi_h^{**}$, respectively. This means that two options are available to a customer entering restaurant j : she/he will either sit at an already occupied table, with probability proportional to the number of customers at that table or will sit at a new table with probability proportional to the concentration parameter γ_j . In the former case, the customer will choose the dish $\phi_{t,j}^*$ with probability $1/2$ and $-\phi_{t,j}^*$ otherwise. In the latter case, the customer will eat a dish served at another table of the franchise with probability proportional to half the number of tables that serve that dish, or will make a new order with probability proportional to the concentration parameter α . In view of this scheme, the probability in (15) turns out to be

$$\tilde{\Pi}_k^{(N)}(\mathbf{n}_1^+, \dots, \mathbf{n}_J^-) = 2^{-N} \bar{\Pi}_k^{(N)}(\mathbf{n}_1^+ + \mathbf{n}_1^-, \dots, \mathbf{n}_J^+ + \mathbf{n}_J^-)$$

and $\bar{\Pi}_k^{(N)}$ on the right-hand-side is the pEPPF of the hierarchical DP derived in Camerlenghi et al. (2019b), namely

$$\bar{\Pi}_k^{(N)}(\mathbf{n}_1, \dots, \mathbf{n}_k) = \left(\prod_{j=1}^J \frac{\prod_{i=1}^k (\gamma_j)_{n_{j,i}}}{(\gamma_j)_{n_j}} \right) \sum_{\ell} \frac{\alpha^k}{(\alpha)_{|\ell|}} \prod_{h=1}^k (\ell_{\bullet,h} - 1)! \prod_{j=1}^J P(K_{n_{j,h}} = \ell_{j,h})$$

where each sums runs over all $\ell_{j,h}$ in $\{1, \dots, n_{j,h}\}$, if $n_{j,h} \geq 1$, and equals 1 if $n_{j,h} = 0$,

whereas $\ell_{\bullet,h} = \sum_{j=1}^J \ell_{j,h}$ and $|\ell| = \sum_{j=1}^J \sum_{h=1}^k \ell_{j,h}$. Note that the latent variable $\ell_{j,h}$ is the number of tables in restaurant j serving the h -th pair of dishes. Moreover, $K_{n_{j,h}}$ is a random variable denoting the number of distinct clusters, out of $n_{j,h}$ observations generated by a DP with parameter γ_j and diffuse baseline P_0 and it is well-known that

$$\mathbb{P}(K_{n_{j,h}} = \ell_{j,h}) = \frac{\gamma_j^{\ell_{j,h}}}{(\gamma_j)_{n_{j,h}}} |\mathfrak{s}(n_{j,h}, \ell_{j,h})|$$

where $|\mathfrak{s}(n_{j,h}, \ell_{j,h})|$ is the signless Stirling number of the first kind. In view of this, one can deduce the predictive distribution

$$\begin{aligned} \mathbb{P}(\phi_{n_j+1,j} \in \cdot | \phi) &= \frac{\gamma_j}{i-1+\gamma_j} \sum_{\ell} \frac{\alpha}{|\ell|+\alpha} \pi(\ell | \phi) P_0(\cdot) \\ &+ \sum_{h=1}^k \left[\frac{n_{j,h}^+ + n_{j,h}^-}{n_j + \gamma_j} + \frac{\gamma_j}{n_j + \gamma_j} \sum_{\ell} \frac{\ell_{\bullet,h}}{|\ell|+\alpha} \pi(\ell | \phi) \right] \left(\frac{\delta_{\phi_h^{**}}(\cdot) + \delta_{-\phi_h^{**}}(\cdot)}{2} \right) \end{aligned}$$

where

$$\pi(\ell | \phi) \propto \frac{\alpha^k}{(\alpha)_{|\ell|}} \prod_{h=1}^k (\ell_{\bullet,h} - 1)! \prod_{j=1}^J \frac{\gamma_j^{\ell_{j,h}}}{(\gamma_j)_{n_{j,h}^+ + n_{j,h}^-}} |\mathfrak{s}(n_{j,h}^+ + n_{j,h}^-, \ell_{j,h})| \mathbb{1}_{\{1, \dots, n_{j,h}^+ + n_{j,h}^-\}}(\ell_{j,h})$$

is the posterior distribution of the latent variables $\ell_{j,h}$'s, and $\mathbb{1}_A$ is the indicator function of set A .

4 Posterior inference

The findings of the previous section are the key ingredients to perform posterior inference with a marginal Gibbs sampler. The output of the sampler is structured into three levels: the first produces posterior probabilities on partitions of disease-specific locations; the second generates density estimates; the third provides clusters of patients. For notational simplicity, we omit the dependence on m , except for the description of the sampling step that generates ω . Recall that $\theta = (\theta_1, \dots, \theta_J)$ and $\phi = \{(\phi_{1,j}, \dots, \phi_{n_j,j}) : j = 1, \dots, J\}$, with $\phi_{i,j} = (\xi_{i,j}, \sigma_{i,j}^2)$. The target distribution of the sampler is the joint distribution of θ , ϕ and ω conditionally on the observed data \mathbf{X} .

Sampling ϕ . In view of the CRF representation of the s-HDP, $t_{i,j}$ stands for the label of the table where the i -th customer in restaurant j sits and $h_{t,j}$ for the dish label served at table t in restaurant j and with \mathbf{t} and \mathbf{h} we denote the corresponding arrays. Moreover, define the assignment variable $s_{i,j} = \mathbb{1}(\phi_{i,j} = \phi_{t_{i,j},j}^*) - \mathbb{1}(\phi_{i,j} = -\phi_{t_{i,j},j}^*)$ and \mathbf{s} is the corresponding arrays. In order to generate ϕ , we need to sample

- (i) $(t_{i,j}, s_{i,j})$ for $i = 1, \dots, n_j$ and $j = 1, \dots, J$;
- (ii) $h_{t,j}$ for $t \in \mathbf{t}$ and $j = 1, \dots, J$;

(ii) ϕ_h^{**} for $h \in \mathbf{h}$.

Note that, using the latent allocation indicators in \mathbf{t} and \mathbf{h} , the sampling scheme is more efficient than sampling directly from the full conditional of each $\phi_{i,j}$, since the algorithm can update more than one parameter simultaneously (Neal, 2000). Define $\varepsilon_{i,j} = X_{i,j} - \theta_j$ and denote with $h(\varepsilon_{i,j}|\phi^*)$ the conditional normal density of $\varepsilon_{i,j}$ given $\phi^* = (\xi^*, \sigma^{2*})$, while the marginal density is

$$\bar{h}(\varepsilon_{i,j}) = \int h(\varepsilon_{i,j}|\phi)P_0(d\phi)$$

To sample $(t_{i,j}, s_{i,j})$ from their joint full conditional, we first sample $t_{i,j}$ from

$$P(t_{i,j} = t \mid \mathbf{t}^{-(i,j)}, \mathbf{h}^{-(i,j)}, \phi^{*-(i,j)}, \phi^{**-(i,j)}, \varepsilon_{i,j}) \propto \begin{cases} n_{t,j}^{-(i,j)} p_{\text{old}}(\varepsilon_{i,j}|\phi_{t,j}^*) & \text{if } t \in \mathbf{t}^{-(i,j)} \\ \gamma_j p_{\text{new}}(\varepsilon_{i,j}|\phi^{**-(i,j)}) & \text{if } t = t^{\text{new}} \end{cases}$$

where $\mathbf{t}^{-(i,j)}$, $\mathbf{h}^{-(i,j)}$, $\phi^{*-(i,j)}$, $\phi^{**-(i,j)}$ coincide with the arrays \mathbf{t} , \mathbf{h} , ϕ^* , ϕ^{**} after having removed the entries corresponding to the i -th customer in restaurant j . Moreover

$$p_{\text{old}}(\varepsilon_{i,j}|\phi_{t,j}^*) = \frac{1}{2}h(\varepsilon_{i,j}|\phi_{t,j}^*) + \frac{1}{2}h(\varepsilon_{i,j} - \phi_{t,j}^*)$$

and

$$p_{\text{new}}(\varepsilon_{i,j}|\phi^{**-(i,j)}) = \sum_{h=1}^{k^{-(i,j)}} \frac{\ell_{\bullet,h}}{|\ell| + \alpha} \left\{ \frac{1}{2}h(\varepsilon_{i,j}|\phi_h^{**}) + \frac{1}{2}h(\varepsilon_{i,j} - \phi_h^{**}) \right\} + \frac{\alpha}{|\ell| + \alpha} \bar{h}(\varepsilon_{i,j})$$

Then we sample $s_{i,j}$ from its full conditional

$$p(s_{i,j} = s \mid \phi^*, t_{i,j}, \varepsilon_{i,j}) \propto \begin{cases} h(\varepsilon_{i,j}|\phi_{t_{i,j}}^*) & \text{if } s = 1 \\ h(\varepsilon_{i,j} - \phi_{t_{i,j}}^*) & \text{if } s = -1 \end{cases}$$

The conditional distribution of $h_{t,j}$ is

$$p(h_{t,j} = h \mid \mathbf{t}, \mathbf{h}^{-(t,j)}, \phi^{**-(t,j)}, \mathbf{s}, \boldsymbol{\varepsilon}) \propto \begin{cases} \ell_{\bullet,h}^{-(t,j)} \prod_{\{(i,j): t_{i,j}=t\}} h(s_{i,j} \varepsilon_{i,j}|\phi_h) & \text{if } h \in \mathbf{h}^{-(t,j)} \\ \alpha \int \prod_{\{(i,j): t_{i,j}=t\}} h(s_{i,j} \varepsilon_{i,j}|\phi) P_0(d\phi) & \text{if } h = h^{\text{new}} \end{cases}$$

Finally, when P_0 is conjugate with respect to the Gaussian kernel, the full conditional distribution of ϕ_h^{**} is obtained in closed form as posterior distribution of a Gaussian model, using as observations the collection $\{(s_{i,j} \varepsilon_{i,j}) : h_{t_{i,j},j} = h\}$.

Sampling $\boldsymbol{\theta}$. When sampling the disease-specific location parameters, one can rely on a Chinese restaurant process restricted to those partitions that are consistent with the ordering of the diseases. Thus, in order to generate $\boldsymbol{\theta}$, we first sample the labels

$\mathbf{t}_\theta = \{t_1, \dots, t_J\}$, where t_j is the label of the table where the j -th customer sits. Then, we sample the dish θ_t^* associated to table t for all $t \in \mathbf{t}_\theta$. If $z_{i,j} = X_{i,j} - \xi_{i,j}$, the conditional density of $\mathbf{z}_j = (z_{1,j}, \dots, z_{n_j,j})$ associated to the location parameter θ^* , given $\boldsymbol{\sigma}_j = (\sigma_{1,j}, \dots, \sigma_{n_j,j})$, is

$$f_{\theta^*}(\mathbf{z}_j | \boldsymbol{\sigma}_j) = \frac{1}{\sqrt{2\pi} \prod_{i=1}^{n_j} \sigma_{i,j}} \exp \left\{ -\frac{1}{2} \sum_{i=1}^{n_j} \frac{(z_{i,j} - \theta^*)^2}{\sigma_{i,j}^2} \right\}$$

Under the prior in (9), the full conditional distribution of \mathbf{t}_θ is provided by

$$p(t_j = t \mid t_1, \dots, t_{j-1}, \theta_{j-1}, \mathbf{z}_j, \boldsymbol{\sigma}_j) \propto \begin{cases} a(\omega, \theta_1, \dots, \theta_{j-1}) f_{\theta_{j-1}}(\mathbf{z}_j | \boldsymbol{\sigma}_j) & \text{if } t = t_j \\ [1 - a(\omega, \theta_1, \dots, \theta_{j-1})] \int f_\theta(\mathbf{z}_j | \boldsymbol{\sigma}_j) G(d\theta) & \text{if } t = t^{\text{new}} \\ 0 & \text{otherwise} \end{cases}$$

Finally, when G is conjugate with respect to the Gaussian kernel, the full conditional distribution of θ_t^* , given $\{\mathbf{z}_j : t_j = t\}$, is obtained in closed form using conjugacy of the Normal-Normal model.

Sampling the concentration parameter. Finally, the concentration parameter ω can be sampled through an importance sampling step using as importance distribution the prior p_ω over ω . Denoting with M_m the selected partition for $\boldsymbol{\theta}_m$ and with T_m the number of clusters in M_m , we have

$$p(\omega \mid M_m : m = 1, \dots, M) \propto p_\omega(\omega) \frac{\omega^{\sum_{m=1}^M T_m - M}}{(\omega + 2)^M (\omega^2 + \omega + 3)^M}.$$

5 Results

5.1 Simulation studies

We perform a series of simulation studies with two main goals. First, we aim to highlight the drawbacks of clustering based on the entire distribution with respect to our proposal in the context of small sample sizes. Second, we check the model's ability of detecting the presence of underlying relevant factors in the sense described in Section 2.2.

To accomplish the first goal, we compare the results obtained using our model with the nested Dirichlet process (NDP) (Rodriguez et al., 2008), arguably the most popular Bayesian model to cluster populations. Mimicking the real hypertensive dataset, we simulate data for 4 samples, ideally corresponding to four diseases, with respective sample sizes of 50, 19, 9 and 22, which correspond to the sample sizes of the real data investigated in Section 5.2. Since the NDP does not allow to treat jointly multiple response variables,

we consider only one response variable to ensure a fair comparison. The observations are sampled from the following distributions and 100 simulation studies are performed.

$$\begin{aligned} X_{i,1} &\stackrel{iid}{\sim} 0.5\mathcal{N}(0, 0.5) + 0.5\mathcal{N}(2, 0.5) && \text{for } i = 1, \dots, n_1 \\ X_{i,2} &\stackrel{iid}{\sim} 0.5\mathcal{N}(2, 0.5) + 0.5\mathcal{N}(4, 0.5) && \text{for } i = 1, \dots, n_2 \\ X_{i,3} &\stackrel{iid}{\sim} 0.5\mathcal{N}(4, 0.5) + 0.5\mathcal{N}(6, 0.5) && \text{for } i = 1, \dots, n_3 \\ X_{i,4} &\stackrel{iid}{\sim} 0.5\mathcal{N}(6, 0.5) + 0.5\mathcal{N}(8, 0.5) && \text{for } i = 1, \dots, n_4 \end{aligned}$$

Note that the true data generating process corresponds to samples from distinct distributions with pairwise sharing of a mixture component. Alternative scenarios are considered in the additional simulation studies that can be found in Section D of the Supplement.

The implementation of the NDP was carried out through the marginal sampling scheme proposed in [Zuanetti et al. \(2018\)](#), which is suitably extended to accommodate hyperpriors on the concentration parameters of the NDP. To simplify the choice of the hyperparameters, as suggested by [Gelman et al. \(2013, p. 535 and p. 551–554\)](#) we estimate both models over standardized data. For our model, we set $G_m = \mathcal{N}(0, 1)$ and $P_{0,m} = \text{NIG}(\mu = 0, \tau = 1, \alpha = 2, \beta = 4)$. Here, $\text{NIG}(\mu, \tau, \alpha, \beta)$ indicates a normal inverse gamma distribution. The base distribution for the NDP is $\text{NIG}(\mu = 0, \tau = 0.01, \alpha = 3, \beta = 3)$, as in [Rodriguez et al. \(2008\)](#). Finally, we use gamma priors with shape 3 and rate 3 for all concentration parameters, which is a common choice. For each simulation study, we perform 10,000 iterations of the MCMC algorithms with the first 5,000 used as burn-in.

Table 1 displays summaries of the results on population clustering, darker rows correspond to partitions that are not consistent with the natural ordering of the diseases. The true clustering structure is given by the finest partition. As already observed in [Rodriguez et al. \(2008\)](#), the NDP tends to identify fewer, rather than more clusters, due to the presence of small sample sizes. Using the *maximum a posteriori* estimate, our model correctly identifies the partition in 99 out of 100 simulation studies and a partition with three elements or more in 100 out of 100 simulation studies. The same counts for the NDP are, respectively, 0 out of 100 and 21 out of 100. Analogous conclusions can be drawn looking at posterior probability averages and medians across the 100 simulation studies (see Table 1) leaving no doubt about the model to be preferred under this scenario.

Finally, we randomly select three simulation studies among the 100 to better understand the performance in estimating the other model parameters. Here we comment on one of the studies, the other two leading to similar results are reported in Section D.1.1 of the Supplement. Figure 3a shows point estimates and credible intervals for the population-specific location parameters $\theta_1, \theta_2, \theta_3, \theta_4$. The true means belong to the 95% credible intervals.

Moreover, it turns out that the model is able to detect the presence of two clusters of subjects leading to a posterior distribution for the number of clusters that is rather


Table 1: Simulation studies summaries.

Partitions	sHDP			NDP		
	MAP count	Average post. prob.	Median post. prob.	MAP count	Average post. prob.	Median post. prob.
$\{1,2,3,4\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2,3,4\}$	0	0.000	0.000	2	0.020	0.000
$\{1,2\}\{3,4\}$	0	0.000	0.000	72	0.695	0.860
$\{1,3,4\}\{2\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2\}\{3,4\}$	0	0.027	0.007	3	0.035	0.000
$\{1,2,3\}\{4\}$	0	0.000	0.000	5	0.061	0.000
$\{1,4\}\{2,3\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2,3\}\{4\}$	1	0.054	0.015	0	0.014	0.000
$\{1,3\}\{2,4\}$	0	0.000	0.000	0	0.000	0.000
$\{1,2,4\}\{3\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2,4\}\{3\}$	0	0.000	0.000	0	0.000	0.000
$\{1,2\}\{3\}\{4\}$	0	0.004	0.000	18	0.175	0.032
$\{1,3\}\{2\}\{4\}$	0	0.000	0.000	0	0.000	0.000
$\{1,4\}\{2\}\{3\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2\}\{3\}\{4\}$	99	0.915	0.954	0	0.000	0.000


concentrated on the true value, see Figure 3b–3d. Moreover, the point estimate for the subject partition, obtained minimizing the Binder loss function, also contains two clusters, proving the ability of the model to detect the underlying relevant factor. In Section D of the Supplement, a number of additional simulation studies are conducted, both using alternative specifications over the disorder-specific parameters and different data generating mechanisms: the results highlight a good performance of the model, which appears also able to detect outliers, to highlight non-location effects of the disorders and to produce reliable outputs even under deviation from symmetry.

5.2 Impact of hypertensive disorders on maternal cardiac dysfunction


Our analysis is based on the dataset of [Tatapudi and Pasumorthy \(2017a\)](#), which can be obtained from <https://data.mendeley.com/datasets/d72zr4xggx/1>. The dataset contains observations for 10 cardiac function measurements collected through a prospective case-control study on women in the third semester of pregnancy divided in $n_1 = 50$ control cases (C), $n_2 = 19$ patients with gestational hypertension (G), $n_3 = 9$ patients with mild preeclampsia (M) and $n_4 = 22$ patients with severe preeclampsia (S). The cases are women admitted from 2012 to 2014 to King George Hospital in Visakhapatnam, India. The healthy sample is composed by normotensive pregnant women. All women with hypertension were on antihypertensive treatment with oral Labetalol or Nifedipine. Women with severe hypertension were treated with either oral nifedipine and parenteral




(a) 95% credible intervals for population-specific locations



(b) Number of second-level clusters.



(c) Co-clustering.



(d) Co-clustering.

Figure 3: Panel (a): Mean point estimates and 95% credible intervals for the four populations, vertical lines correspond to true values. Panel (b): Posterior distribution on the number of second-level clusters. Panels (c) and (d): heatmaps of second-level clustering, darker colors correspond to higher probability of co-clustering; in (c) patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities and in (d) patients are reordered based on co-clustering probabilities.

labetalol or a combination. For more details on the dataset, we refer to [Tatapudi and Pasumarthu \(2017b\)](#). The prior specification is the same as in the previous section. Sections E.2 and E.3 of the Supplement contain a prior-sensitivity analysis and show rather robust results w.r.t. different prior specifications. Inference is based on 10,000 MCMC iterations with the first half used as burn-in.

Table 2 displays the posterior distributions for the partitions of unknown disease-specific means along with the corresponding entropy measurements, that can be used as measures of uncertainty. First note that if one takes also the ordering among distinct disease-specific locations into account, the posterior partition probabilities are, as desired, concentrated on specific orders of the associated unique values for all ten cardiac indexes. For instance, we have $\mathbb{P}(\{\theta_{C,CI} = \theta_{G,CI} = \theta_{M,CI}\} \{\theta_{S,CI}\} \mid X) = \mathbb{P}(\theta_{C,CI} = \theta_{G,CI} = \theta_{M,CI} > \theta_{S,CI} \mid X) = 0.463$. The ordered partitions with the highest posterior probability

Table 2: Posterior probabilities over partitions of means. Maximum a posteriori probabilities are in **bold**.

partitions	CI	CWI	LVMI	IVST	LVPW	EF	FS	EW	AW	E/A
{C,G,M,S}	0.021	0.000	0.000	0.000	0.000	0.365	0.303	0.096	0.000	0.000
{C}{G,M,S}	0.002	0.546	0.001	0.083	0.016	0.078	0.190	0.021	0.036	0.000
{C,G}{M,S}	0.002	0.000	0.001	0.000	0.000	0.037	0.038	0.072	0.076	0.049
{C,M,S}{G}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G}{M,S}	0.001	0.139	0.001	0.019	0.024	0.028	0.078	0.042	0.232	0.055
{C,G,M}{S}	0.463	0.000	0.595	0.000	0.000	0.276	0.045	0.498	0.020	0.002
{C,S}{G,M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G,M}{S}	0.146	0.099	0.188	0.551	0.672	0.074	0.164	0.092	0.260	0.033
{C,M}{G,S}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,G,S}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G,S}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,G}{M}{S}	0.233	0.000	0.107	0.000	0.000	0.083	0.062	0.114	0.091	0.371
{C,M}{G}{S}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,S}{G}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G}{M}{S}	0.133	0.216	0.108	0.347	0.288	0.060	0.121	0.065	0.287	0.491
$\sum \log_{15} (p_i^{-p_i})$	0.501	0.430	0.415	0.361	0.289	0.632	0.688	0.598	0.613	0.424




Figure 4: 95% credible intervals for population-specific locations for CI and CWI

are displayed in Table 3.

Considering the posterior probabilities summarized in Table 2 and in Table 3, we find that the cardiac index (CI) is reduced in severe preeclampsia compared to all other patients, indicating reduced myocardial contractility in the presence of the most severe disorder. The cardiac work index (CWI) is a good indicator to distinguish between cases and control, but not among cases. The left ventricular mass index (LVMI) is increased in severe preeclampsia patients compared to other pregnant women, indicating ventricular remodelling. While inter ventricular septal thickness (IVST) and left ventricular posterior wall thickness (LVPW) differ both between cases and controls and between severe preeclampsia and other disorders, indicating a progressive increase in the indexes with the severity of the disorder. The posterior probabilities associated to indexes of systolic function such as ejection fraction (EF) and fraction shortening (FS) are relatively concentrated on the partition of complete homogeneity, letting us to conclude that no differences

Table 3: Posterior probabilities over ordered partitions of means.

cardiac index	ordered partition with highest posterior probability	posterior prob
CI	$\{\text{C}, \text{G}, \text{M}\} > \{\text{S}\}$	0.463
CWI	$\{\text{C}\} < \{\text{G}, \text{M}, \text{S}\}$	0.546
LVMl	$\{\text{C}, \text{G}, \text{M}\} < \{\text{S}\}$	0.595
IVST	$\{\text{C}\} < \{\text{G}, \text{M}\} < \{\text{S}\}$	0.548
LVPW	$\{\text{C}\} < \{\text{G}, \text{M}\} < \{\text{S}\}$	0.671
EF	$\{\text{C}, \text{G}, \text{M}, \text{S}\}$	0.365
FS	$\{\text{C}, \text{G}, \text{M}, \text{S}\}$	0.303
EW	$\{\text{C}, \text{G}, \text{M}\} > \{\text{S}\}$	0.497
AW	$\{\text{C}\}_i \{\text{G}, \text{M}\} < \{\text{S}\}$	0.256
E/A	$\{\text{C}\} > \{\text{G}\} > \{\text{M}\} > \{\text{S}\}$	0.466

are present among patients. As for the parameters of the diastolic function, the posterior distribution for the E-wave indicator identifies a modified index in severe preeclampsia patients, while the mean E/A ratio indicates a decreasing diastolic function with the severity of the disorder. The posterior for the A-wave index is actually concentrated on three distinct partitions, leaving a relatively high uncertainty regarding the modifications of the index. However, considering jointly the three partitions with the highest posterior probability, differences are detected between control and cases with a total posterior probability equal to 0.779. Figure 4 shows point estimates and credible intervals for disorder-specific location parameters for the first two cardiac indexes. Analogous plots for all cardiac indexes can be found in Section E.1 of the Supplement.

Table 4 shows the results obtained using the prior in (7), instead of (8). We remark that for all ten cardiac indexes, the posterior associates negligible probabilities to partitions that are in contrast with the natural order of the diagnoses. This is particularly reassuring in that the model, even without imposing such an order a priori, is able to single it out systematically across cardiac indexes. Moreover, we observe how the partitions identified by MAP are the same of Table 2 for all cardiac index except AW. However, even under this alternative prior, the A-wave index is concentrated on the same three distinct partitions leading to the conclusion that there exists a difference between cases and control.

As far as prediction and second-level clustering are concerned, Figure 5 displays the density estimates and the heatmap of co-clustering probabilities between pairs of patients for the E/A ratio and LVMl. Figure 5b shows that co-clustering probabilities are similar within and across diagnoses, indicating that the effect of the diseases on the distribution of the cardiac index is mostly explained through shifts between disease-specific locations. Moreover Figure 5b suggests the presence of three outliers that have low probability of co-clustering with all the other subjects and that would be ignored by the model using

Table 4: Posterior probabilities over partitions of means. Maximum a posteriori probabilities are in **bold**.

partitions	CI	CWI	LVMI	IVST	LVPW	EF	FS	EW	AW	E/A
$\{\text{C,G,M,S}\}$	0.019	0.000	0.000	0.000	0.000	0.332	0.247	0.078	0.000	0.000
$\{\text{C}\}\{\text{G,M,S}\}$	0.002	0.643	0.001	0.114	0.031	0.065	0.130	0.048	0.080	0.000
$\{\text{C,G}\}\{\text{M,S}\}$	0.004	0.000	0.003	0.000	0.000	0.044	0.019	0.152	0.073	0.103
$\{\text{C,M,S}\}\{\text{G}\}$	0.004	0.000	0.000	0.000	0.000	0.037	0.105	0.013	0.000	0.000
$\{\text{C}\}\{\text{G}\}\{\text{M,S}\}$	0.002	0.065	0.002	0.047	0.078	0.027	0.036	0.063	0.424	0.167
$\{\text{C,G,M}\}\{\text{S}\}$	0.316	0.000	0.527	0.000	0.000	0.178	0.032	0.288	0.002	0.000
$\{\text{C,S}\}\{\text{G,M}\}$	0.023	0.000	0.000	0.000	0.000	0.019	0.103	0.006	0.000	0.000
$\{\text{C}\}\{\text{G,M}\}\{\text{S}\}$	0.173	0.089	0.124	0.472	0.594	0.033	0.054	0.064	0.140	0.042
$\{\text{C,M}\}\{\text{G,S}\}$	0.002	0.000	0.001	0.003	0.000	0.044	0.031	0.017	0.000	0.000
$\{\text{C,G,S}\}\{\text{M}\}$	0.018	0.000	0.000	0.000	0.000	0.061	0.067	0.016	0.000	0.000
$\{\text{C}\}\{\text{G,S}\}\{\text{M}\}$	0.005	0.163	0.001	0.095	0.006	0.028	0.040	0.015	0.016	0.000
$\{\text{C,G}\}\{\text{M}\}\{\text{S}\}$	0.213	0.000	0.124	0.000	0.000	0.052	0.014	0.121	0.036	0.241
$\{\text{C,M}\}\{\text{G}\}\{\text{S}\}$	0.074	0.000	0.137	0.003	0.000	0.041	0.022	0.055	0.001	0.000
$\{\text{C,S}\}\{\text{G}\}\{\text{M}\}$	0.014	0.000	0.000	0.000	0.000	0.011	0.067	0.004	0.000	0.000
$\{\text{C}\}\{\text{G}\}\{\text{M}\}\{\text{S}\}$	0.133	0.040	0.079	0.265	0.291	0.029	0.033	0.059	0.229	0.448
$\sum \log_{15} (p_i^{-p_i})$	0.687	0.407	0.509	0.501	0.371	0.828	0.886	0.823	0.582	0.505

a more traditional ANOVA structure. On the other hand, Figure 5e shows a slightly different pattern for co-clustering probabilities in the fourth square, which suggests that the heterogeneity between severe preeclampsia patients and the other patients is not entirely explained by shifts in disease-specific locations. Finally, Figure 5f suggests the presence of an underlying relevant factor. The corresponding figures for all ten response variables are reported in Section E.1 of the Supplement and can be used for prediction and for a graphical analysis aimed at controlling the presence of underlying relevant factors, outliers and differences across diseases distinct from shifts between disease-specific locations.

Our results are coherent with almost all of the findings in Tatapudi and Pasumarthu (2017b), where results were obtained through a series of independent frequentist tests. However, importantly, we are able to provide more insights thanks to the simultaneous comparison approach and the latent clustering of subjects. For instance, considering the response LVMI, Tatapudi and Pasumarthu (2017b) detected a significant increase in cases compared to controls and an increase in severe preeclampsia compared to gestational hypertensive and mild preeclampsia patients. Such results do not clarify whether a modification exists between the control group and gestational hypertensive patients or between the latter and mild preeclampsia patients. Moreover, in contrast to our analysis, their results do not provide any information concerning the presence of underlying common factors, outliers or distributional effects (different from shifts in locations).

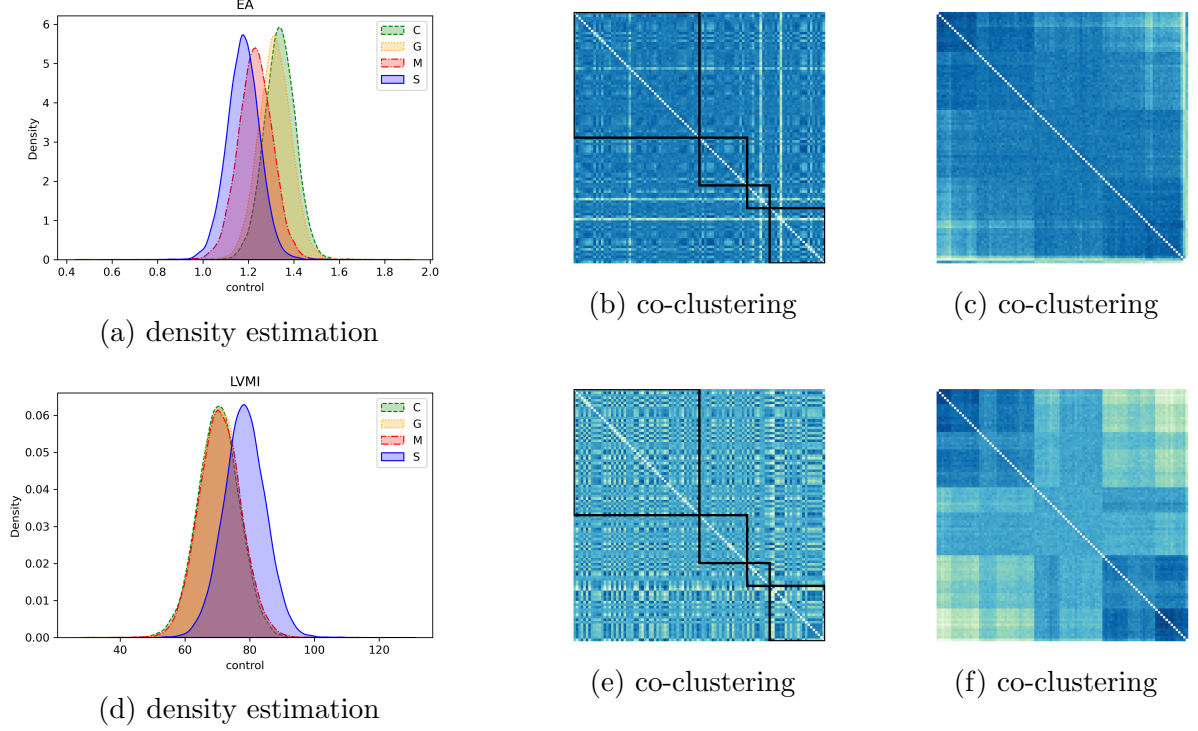


Figure 5: Panels (a) and (d): density estimates. Panels (b)–(c) and (e)–(f): heatmaps of the posterior probabilities of co-clustering; in (b) and (e) patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities; in (c) and (f) patients are reordered based on co-clustering probabilities.

6 Concluding remarks

We designed a Bayesian nonparametric model to detect clusters of hypertensive disorders over different cardiac function indexes and found modified cardiac functions in hypertensive patients compared to healthy subjects as well as progressively increased alterations with the severity of the disorder. The proposed model has application potential also beyond the considered setup when the goal is to cluster populations according to multivariate information: it borrows strength across response variables, preserves the flexibility intrinsic to nonparametric models, and correctly detects partitions of populations even in presence of small sample sizes, when alternative distribution-based clustering models tend to underestimate the number of clusters. The key component of the model is the s-HDP, a hierarchical nonparametric structure for the error terms that offers flexibility and serves as a tool to investigate the presence of unobserved factors, outliers and effects other than changes in locations. Interesting extensions of the model include generalizations to other types of invariances in order to accommodate identifiability in generalized linear models, for instance in presence of count data and a log link function, as well as generalizations to other types of processes, beyond the Dirichlet process.

Acknowledgements

Most of the paper was completed while B. Franzolini was a Ph.D. student at the Bocconi University, Milan. A. Lijoi and I. Prünster are partially supported by MIUR, PRIN Project 2015SNS29B.

References

- Aksu, E., B. Cuglan, A. Tok, E. Celik, A. Doganer, A. Sokmen, and G. Sokmen (2021). Cardiac electrical and structural alterations in preeclampsia. *The Journal of Maternal-Fetal & Neonatal Medicine*, 1–10.
- Ambrožic, J., M. Lucovnik, K. Prokšelj, J. Toplišek, and M. Cvijic (2020). Dynamic changes in cardiac function before and early postdelivery in women with severe preeclampsia. *Journal of hypertension* 38, 1367–1374.
- Antoniak, C. E. (1974). Mixtures of Dirichlet processes with applications to bayesian nonparametric problems. *Annals of Statistics* 2, 1152–1174.
- Bellamy, L., J.-P. Casas, A. D. Hingorani, and D. J. Williams (2007). Pre-eclampsia and risk of cardiovascular disease and cancer in later life: systematic review and meta-analysis. *BMJ* 335, 974.
- Beraha, M., A. Guglielmi, and F. A. Quintana (2021). The semi-hierarchical Dirichlet process and its application to clustering homogeneous distributions. *Bayesian Analysis* 16, 1187–1219.
- Camerlenghi, F., D. B. Dunson, A. Lijoi, I. Prünster, and A. Rodriguez (2019a). Latent nested nonparametric priors (with discussion). *Bayesian Analysis* 14, 1303–1356.
- Camerlenghi, F., A. Lijoi, P. Orbanz, and I. Prünster (2019b). Distribution theory for hierarchical processes. *Annals of Statistics* 47, 67–92.
- Christensen, J. and L. Ma (2020). A Bayesian hierarchical model for related densities by using Pólya trees. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82, 127–153.
- Cifarelli, D. M. and E. Regazzini (1978). Problemi statistici non parametrici in condizioni di scambiabilità parziale e impiego di medie associative. *Quaderni Istituto Matematica Finanziaria, Torino*.
- Cipolli, W., T. Hanson, and A. C. McLain (2016). Bayesian nonparametric multiple testing. *Computational Statistics & Data Analysis* 101, 64–79.

- Dahl, D. B. and M. A. Newton (2007). Multiple hypothesis testing by clustering treatment effects. *Journal of the American Statistical Association* 102, 517–526.
- Dalal, S. (1979a). Dirichlet invariant processes and applications to nonparametric estimation of symmetric distribution functions. *Stochastic Processes and their Applications* 9, 99–107.
- Dalal, S. R. (1979b). Nonparametric and robust Bayes estimation of location. *Optimizing methods in statistics* 9, 141–166.
- Davis, E. F., M. Lazdam, A. J. Lewandowski, S. A. Worton, B. Kelly, Y. Kenworthy, et al. (2012). Cardiovascular risk factors in children and young adults born to preeclamptic pregnancies: a systematic review. *Pediatrics* 129, 1552–1561.
- De Iorio, M., P. Müller, G. L. Rosner, and S. N. MacEachern (2004). An ANOVA model for dependent random measures. *Journal of the American Statistical Association* 99, 205–215.
- deMartelly, V. A., J. Dreixler, A. Tung, A. Mueller, S. Heimberger, A. A. Fazal, H. Naseem, R. Lang, E. Kruse, M. Yamat, et al. (2021). Long-term postpartum cardiac function and its association with preeclampsia. *Journal of the American Heart Association* 10, e018526.
- Denti, F., M. Guindani, F. Leisen, A. Lijoi, W. D. Wadsworth, and M. Vannucci (2020). Two-group Poisson-Dirichlet mixtures for multiple testing. *Biometrics* 76, 10.1111/biom.13314.
- Diaconis, P. and D. Freedman (1986). On inconsistent Bayes estimates of location. *Annals of Statistics* 14, 68–87.
- Do, K.-A., P. Müller, and F. Tang (2005). A Bayesian mixture model for differential gene expression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 54, 627–644.
- Dolea, C. and C. AbouZahr (2003). Global burden of hypertensive disorders of pregnancy in the year 2000. Technical report, GBD 2000 Working Paper, World Health Organization, Geneva.
- Doss, H. (1984). Bayesian estimation in the symmetric location problem. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 68, 127–147.
- Escobar, M. D. and M. West (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* 90, 577–588.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Annals of Statistics* 1, 209–230.

- Garcia-Gonzalez, C., G. Georgiopoulos, S. A. Azim, F. Macaya, N. Kametas, P. Nihoyannopoulos, K. H. Nicolaidis, and M. Charakida (2020). Maternal cardiac assessment at 35 to 37 weeks improves prediction of development of preeclampsia. *Hypertension* 76, 514–522.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013). *Bayesian data analysis*. CRC press.
- Ghosal, S., J. K. Ghosh, and R. Ramamoorthi (1999). Consistent semiparametric Bayesian inference about a location parameter. *Journal of Statistical Planning and Inference* 77, 181–193.
- Gopalan, R. and D. A. Berry (1998). Bayesian multiple comparisons using Dirichlet process priors. *Journal of the American Statistical Association* 93, 1130–1139.
- Guindani, M., P. Müller, and S. Zhang (2009). A Bayesian discovery procedure. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71, 905–925.
- Gutiérrez, L., A. F. Barrientos, J. González, and D. Taylor-Rodríguez (2019). A Bayesian nonparametric multiple testing procedure for comparing several treatments against a control. *Bayesian Analysis* 14, 649–675.
- Hall, M. E., E.M. George, J. P. Granger (2018). The heart during pregnancy. *Revista Española de Cardiología (English Edition)* 64, 1045–1050.
- Igberase, G. and P. Ebeigbe (2006). Eclampsia: ten-years of experience in a rural tertiary hospital in the Niger delta, Nigeria. *Journal of Obstetrics and Gynaecology* 26, 414–417.
- Iglesias, P. L., Y. Orellana, and F. A. Quintana (2009). Nonparametric Bayesian modelling using skewed Dirichlet processes. *Journal of Statistical Planning and Inference* 139, 1203–1214.
- Lijoi, A., B. Nipoti, and I. Prünster (2014). Bayesian inference with dependent normalized completely random measures. *Bernoulli* 20, 1260–1291.
- Lijoi, A., I. Prünster, and G. Rebaudo (2022). Flexible clustering via hidden hierarchical Dirichlet priors. *Scandinavian Journal of Statistics*, forthcoming.
- MacEachern, S. N. (2000). Dependent Dirichlet processes. *Technical Report, Department of Statistics, The Ohio State University*.
- Malik, A., B. Jee, and S. K. Gupta (2019). Preeclampsia: disease biology and burden, its management strategies with reference to India. *Pregnancy Hypertension* 15, 23–31.
- Martin, R. and S. T. Tokdar (2012). A nonparametric empirical Bayes framework for large-scale multiple testing. *Biostatistics* 13, 427–439.

- McClure, E. M., S. Saleem, O. Pasha, and R. L. Goldenberg (2009). Stillbirth in developing countries: a review of causes, risk factors and prevention strategies. *The Journal of Maternal-fetal & Neonatal Medicine* 22, 183–190.
- Moser, S., A. Rodríguez, and C. L. Lofland (2021). Multiple ideal points: Revealed preferences in different domains. *Political Analysis* 29, 139–166.
- Muliere, P. and S. Petrone (1978). A Bayesian predictive approach to sequential search for an optimal dose: parametric and nonparametric models. *Journal of the Italian Statistical Society* 2, 349–364.
- Neal, R. M. (2000). Markov chain sampling methods for Dirichlet process mixture models. *Journal of Computational and Graphical Statistics* 9, 249–265.
- Pedersen, S.S., R. Von Känel, P.J. Tully, and J. Denollet(2017). Psychosocial perspectives in cardiovascular disease. *European Journal of Preventive Cardiology* 24, 108–115.
- Rodriguez, A., D. B. Dunson, and A. E. Gelfand (2008). The nested Dirichlet process. *Journal of the American Statistical Association* 103, 1131–1154.
- Scott, J. G. and J. O. Berger (2006). An exploration of aspects of Bayesian multiple testing. *Journal of Statistical Planning and Inference* 136, 2144–2162.
- Scott, J. G. and J. O. Berger (2010). Bayes and empirical-bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*, 2587–2619.
- Shah, A., B. Fawole, J. M. M’Imunya, F. Amokrane, I. Nafiou, J.-J. Wolomby, et al. (2009). Cesarean delivery outcomes from the WHO global survey on maternal and perinatal health in Africa. *International Journal of Gynecology & Obstetrics* 107, 191–197.
- Soriano, J. and L. Ma (2017). Probabilistic multi-resolution scanning for two-sample differences. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79, 547–572.
- Tatapudi, R. and L. R. Pasumathy (2017a). Data for: Maternal cardiac function in gestational hypertension, mild and severe preeclampsia and normal pregnancy: A comparative study. <https://data.mendeley.com/datasets/d72zr4xggx/1>. DOI:10.17632/d72zr4xggx.1 Licensed under a Creative Commons Attribution 4.0 International licence.
- Tatapudi, R. and L. R. Pasumathy (2017b). Maternal cardiac function in gestational hypertension, mild and severe preeclampsia and normal pregnancy: A comparative study. *Pregnancy Hypertension* 10, 238–241.

- Teh, Y. W., M. I. Jordan, M. J. Beal, and D. M. Blei (2006). Hierarchical Dirichlet processes. *Journal of the American Statistical Association* 101, 1566–1581.
- Timokhina, E., T. Kuzmina, A. Strizhakov, E. Pitskhelauri, I. Ignatko, and V. Belousova (2019). Maternal cardiac function after normal delivery, preeclampsia, and eclampsia: A prospective study. *Journal of Pregnancy* 2019, 2090–2727.
- Zuanetti, D. A., P. Müller, Y. Zhu, S. Yang, and Y. Ji (2018). Clustering distributions with the marginalized nested Dirichlet process. *Biometrics* 74, 584–594.

Supplement to “Model Selection for Maternal Hypertensive Disorders with Symmetric Hierarchical Dirichlet Processes”

A. A review of the invariant Dirichlet process

We provide a brief review of the invariant Dirichlet process (IDP), introduced by Dalal (1979a), which serves as building block for the proposed model. After recalling the definition, we present two representations of the process: the first is the analog of the stick-breaking construction of the Dirichlet Process (DP), whereas the second is an extension of the generalized Pólya urn scheme of Blackwell and MacQueen (1973).

Let (E, \mathcal{E}) be any measurable Euclidean space and $\mathcal{G} = \{g_1, \dots, g_L\}$ be a finite group of measurable transformations on (E, \mathcal{E}) .

Definition A.1 (Invariant Probability). A probability measure P_0 on (E, \mathcal{E}) is a \mathcal{G} -invariant probability distribution, if $P_0(A) = P_0(g_l(A))$, for any $A \in \mathcal{E}$ and $l = 1, \dots, L$.

Definition A.2 (Invariant Random Probability). A random probability \tilde{p} on (E, \mathcal{E}) is said \mathcal{G} -invariant, if it is almost surely \mathcal{G} -invariant.

Definition A.3 (Invariant Partition). A measurable partition A_1, A_2, \dots, A_K of E is a \mathcal{G} -invariant partition, if $A_k = g_l(A_k)$, $\forall k = 1, \dots, K$ and $\forall l = 1, \dots, L$.

Definition A.4 (Invariant Dirichlet Process). A random probability \tilde{p} is an IDP with group of transformations \mathcal{G} , if

1. \tilde{p} is almost surely \mathcal{G} -invariant

$$\tilde{p}(A) = \tilde{p}(g_l(A)) \quad \text{for } l = 1, \dots, L \quad \text{a.s.}$$

2. there exists a \mathcal{G} -invariant probability distribution P_0 on (E, \mathcal{E}) and $\alpha \in \mathbb{R}^+$, such that for any $k \in \mathbb{N}$ and any \mathcal{G} -invariant measurable partition A_1, \dots, A_k

$$(\tilde{p}(A_1), \dots, \tilde{p}(A_k)) \sim D_{k-1}(\alpha P_0(A_1), \dots, \alpha P_0(A_k))$$

where α is the concentration parameter and P_0 the baseline probability measure.

The notation $\tilde{p} \sim \text{IDP}(\alpha, P_0, \mathcal{G})$ indicates that the random measure \tilde{p} is distributed according to an IDP. Note that, if $\tilde{p} \sim \text{DP}(\alpha, P_0)$, then \tilde{p} is not an IDP, since it is not an invariant random probability. Vice versa, if $\tilde{p} \sim \text{IDP}(\alpha, P_0, \mathcal{G})$, then \tilde{p} is not a DP, since its finite dimensional distributions over non \mathcal{G} -invariant partitions are not Dirichlet. However there is a strong relationship between the two processes as shown in Dalal (1979a).

Theorem A.5 (Dalal, 1979a). Let $\tilde{q} \sim \text{DP}(\alpha, P_0)$ and $\tilde{p} \sim \text{IDP}(\alpha, P_0, \mathcal{G})$. Define

$$q^*(\cdot) = \frac{1}{L} \sum_{l=1}^L \tilde{q}(g_l(\cdot \dots))$$

then

$$\tilde{p} \stackrel{d}{=} q^*$$

Tiwari (1988) provided also a constructive definition for the IDP, which is the analogue of the stick-breaking representation of Sethuraman (1994) for the DP.

Proposition A.6 (Tiwari, 1988). If $\tilde{p} \sim \text{IDP}(\alpha, P_0, \mathcal{G})$, then

$$\tilde{p} = \sum_{h=1}^{\infty} \pi_h \sum_{l=1}^L \delta_{g_l(\phi_h^*)}$$

$$\text{with} \quad \pi_h = \frac{\pi'_h}{L} \prod_{r=1}^{h-1} (1 - \pi'_r) \quad \pi'_r \sim \text{Beta}(1, \alpha) \quad \phi_h^* \stackrel{iid}{\sim} P_0.$$

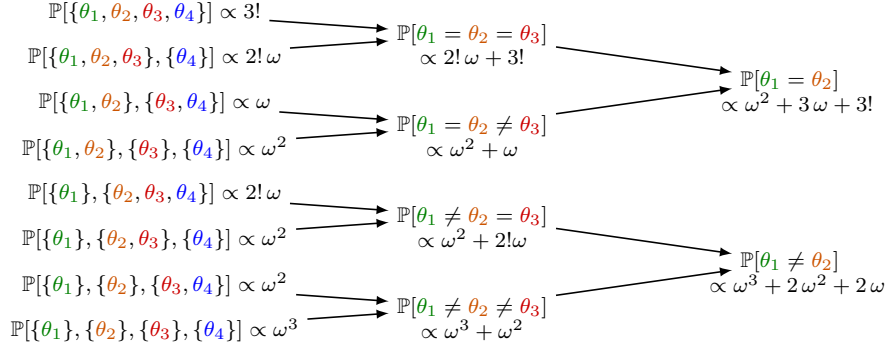


Fig 1: A priori partitions' probabilities joint (on the left) and marginals

Moreover, if $\phi_i \mid \tilde{Q} \stackrel{iid}{\sim} \tilde{Q}$ with $\tilde{Q} \sim \text{IDP}(\alpha, P_0, \mathcal{G})$, by integrating out \tilde{Q} , we get the corresponding generalized Pólya urn representation for the process

$$\phi_1 \sim P_0$$

$$\phi_i \mid \phi_1, \dots, \phi_{i-1} \sim \sum_{j=1}^{i-1} \frac{1}{i-1+\alpha} \left(\frac{1}{L} \sum_{l=1}^L \delta_{g_l(\phi_j)} \right) + \frac{\alpha}{i-1+\alpha} P_0$$

For more details about IDPs we refer to Dalal (1979a), Dalal (1979b), Hannum and Hollander (1983), Doss (1984), Diaconis and Freedman (1986), Tiwari (1988), Ferguson, Phadia, and Tiwari (1992) and Ghosal et al. (1999).

B. Predictive distribution for disease-specific locations

We recall that the prior over partitions is given by

$$\mathbb{P}(M_b^m \mid \omega) \propto \begin{cases} \omega^{k-1} \prod_{i=1}^k (n_i - 1)! & \text{if } M_b^m \text{ is compatible with the natural order} \\ 0 & \text{otherwise} \end{cases}$$

where k is the number of distinct clusters according to the partition M_b^m and n_1, \dots, n_k are the clusters' frequencies. Thus, being $J = 4$, one obtains the probabilities in Figure 1, starting from which it is possible to compute the joint distribution of $(\theta_{1,m}, \dots, \theta_{4,m})$ conditional on ω

$$\theta_{1,m} \mid \omega \sim G_m$$

$$\theta_{2,m} \mid \theta_{1,m}, \omega \sim \frac{\omega^2 + 3\omega + 6}{(\omega + 2)(\omega^2 + \omega + 3)} \delta_{\theta_{1,m}} + \frac{\omega^3 + 2\omega^2 + 2\omega}{(\omega + 2)(\omega^2 + \omega + 3)} G_m$$

$$\theta_{3,m} \mid \theta_{1,m}, \theta_{2,m}, \omega \sim \begin{cases} \frac{2\omega+6}{\omega^2+3\omega+6} \delta_{\theta_{2,m}} + \frac{\omega^2+\omega}{\omega^2+3\omega+6} G_m & \text{if } \theta_{1,m} = \theta_{2,m} \\ \frac{\omega+2}{\omega^2+2\omega+2} \delta_{\theta_{2,m}} + \frac{\omega^2+\omega}{\omega^2+2\omega+2} G_m & \text{if } \theta_{1,m} \neq \theta_{2,m} \end{cases}$$

$$\theta_{4,m} \mid \theta_{1,m}, \theta_{2,m}, \theta_{3,m}, \omega \sim \begin{cases} \frac{3}{\omega+3} \delta_{\theta_{3,m}} + \frac{\omega}{\omega+3} G_m & \text{if } \theta_{1,m} = \theta_{2,m} = \theta_{3,m} \\ \frac{2}{\omega+2} \delta_{\theta_{3,m}} + \frac{\omega}{\omega+2} G_m & \text{if } \theta_{1,m} \neq \theta_{2,m} = \theta_{3,m} \\ \frac{1}{\omega+1} \delta_{\theta_{3,m}} + \frac{\omega}{\omega+1} G_m & \text{otherwise} \end{cases}$$

C. Alternative priors over disorder-specific locations

For comparison purposes and prior sensitivity analysis we consider also two alternative priors over the disorder-specific locations: a uniform prior, which does not penalize multiplicity but incorporates the prior information on

the severity of disorders, and a mixture of Dirichlet processes (DPs), which penalizes for multiplicity but does not reflect prior information.

C.1. Uniform prior

The uniform prior is obtained associating zero-probability to nonsensical partitions and a uniform prior over the remaining, i.e.

$$\mathbb{P}(M_b^m) \propto \begin{cases} \frac{1}{8} & \text{if } M_b^m \text{ is compatible with the natural order} \\ 0 & \text{otherwise} \end{cases}$$

The predictive distributions are

$$\theta_j | \theta_1, \dots, \theta_{j-1} \sim \frac{1}{2} \delta_{\theta_{j-1}} + \frac{1}{2} G$$

and the full conditional distribution of t_j is

$$p(t_j = t | \mathbf{t}_\theta^{(-j)}, \boldsymbol{\theta}^{(-j)}, \mathbf{z}_j, \boldsymbol{\sigma}_j) \propto \begin{cases} f_{\theta_{j-1}}(\mathbf{z}_j | \boldsymbol{\sigma}_j) & \text{if } t = t_{j-1} \\ \int f_\theta(\mathbf{z}_j | \boldsymbol{\sigma}_j) G(d\theta) & \text{if } t = t^{\text{new}} \end{cases}$$

Note that with this prior there is not a common concentration parameter and therefore there is no borrowing of information across cardiac indexes as well as no Occam's razor effect.

C.2. Mixture of DPs prior

Using mixtures of DPs as prior, the locations $(\theta_1, \dots, \theta_J)$, conditionally on ω , are from a DP and the law of the partition (described in Section 2.1 of the main paper) yields the well-known predictive distributions

$$\theta_j | \omega, \theta_1, \dots, \theta_{j-1} \sim \sum_{t=1}^{T_{j-1}} \frac{n_t}{j-1+\omega} \delta_{\theta_t^*} + \frac{\omega}{j-1+\omega} G$$

with T_{j-1} the number of distinct values θ_t^* in $(\theta_1, \dots, \theta_{j-1})$ and $n_t = \text{card}\{i \in \{1, \dots, j-1\} : \theta_i = \theta_t^*\}$. From this, one easily deduces that the conditional prior odds against two populations sharing the same location is

$$\frac{\mathbb{P}(\theta_{j,m} \neq \theta_{j',m} | \omega)}{\mathbb{P}(\theta_{j,m} = \theta_{j',m} | \omega)} = \frac{\Pi_2^{(2)}(1, 1)}{\Pi_1^{(2)}(2)} = \omega$$

Under the mixture of DP prior, the full conditional distribution of t_j is

$$p(t_j = t | \mathbf{t}_\theta^{(-j)}, \boldsymbol{\theta}^{*(-j)}, \mathbf{z}_j, \boldsymbol{\sigma}_j) \propto \begin{cases} n_t^{-j} f_{\theta_t^*}(\mathbf{z}_j | \boldsymbol{\sigma}_j) & \text{if } t \in \mathbf{t}^{(-j)} \\ \omega \int f_\theta(\mathbf{z}_j | \boldsymbol{\sigma}_j) G(d\theta) & \text{if } t = t^{\text{new}} \end{cases}$$

where $\mathbf{t}^{(-j)} = \{t_{j'} : j' \neq j\}$, $\boldsymbol{\theta}^{*(-j)} = \{\theta_t^* : t \in \mathbf{t}^{(-j)}\}$ and n_t^{-j} denotes the number of customers already allocated to table t , after removing the j -th customer.

Moreover, if the prior p_ω for the concentration parameter is chosen to be gamma with shape a and rate b , the full conditional for the parameter ω can be obtained by generalizing the result for a single mixture of DPs in [Escobar and West \(1995\)](#), as follows. Denote with T_m the number of distinct values of $\boldsymbol{\theta}_m = \{\theta_{1,m}, \dots, \theta_{d,m}\}$, for $m = 1, \dots, M$ and note that ω depends on the data only through T_1, \dots, T_M . The full conditional distribution of ω is:

$$\begin{aligned} p(\omega | T_1, \dots, T_M) &\propto p_\omega(\omega) \cdot \prod_{m=1}^M p(T_m | \omega) \\ &\propto p_\omega(\omega) \cdot \prod_{m=1}^M \left[c_d(T_m) d! \omega^{T_m} \frac{\Gamma(\omega)}{\Gamma(\omega + d)} \right] \end{aligned}$$

where $p_\omega(\omega)$ is the prior density of ω and $c_d(T_m) = p(T_m \mid \omega = 1)$. Therefore

$$p(\omega \mid T_1, \dots, T_m) \propto p_\omega(\omega) \cdot \omega^{\sum_m T_m - M} (\omega + d)^M \prod_{m=1}^M \left[\int_0^1 u^\omega (1-u)^{d-1} du \right]$$

Defining M auxiliary random variables u_m for $m = 1, \dots, M$ such that $u_m \mid \omega \stackrel{iid}{\sim} \text{Beta}(\omega + 1, d)$. It follows that

$$p(\omega, u_1, \dots, u_M \mid T_1, \dots, T_m) \propto p_\omega(\omega) \cdot \omega^{\sum_m T_m - M} (\omega + d)^M \prod_{m=1}^M u_m^\omega \prod_{m=1}^M (1 - u_m)^{d-1}$$

Finally if $p_\omega \equiv \text{Gamma}(a, b)$, then

$$\begin{aligned} p(\omega \mid u_1, \dots, u_M, T_1, \dots, T_m) &\propto \omega^{a + \sum_m T_m - M - 1} (\omega + d)^M \exp \left\{ -\omega \left(b - \sum_{m=1}^M \log(u_m) \right) \right\} \\ &\propto \sum_{v=0}^M \binom{M}{v} \frac{d^v \Gamma \left(a + \sum_{m=1}^M T_m - v \right)}{\left(b - \sum_{m=1}^M \log(u_m) \right)^{a + \sum_{m=1}^M T_m - v}} \times \text{Gamma} \left(a + \sum_{m=1}^M T_m - v, b - \sum_{m=1}^M \log(u_m) \right) \end{aligned}$$

so that the conditional distribution of ω is a mixture of $M + 1$ Gamma distributions. The sampling of ω becomes

- (i) Sample u_m , for $m = 1, \dots, M$, independently from a $\text{Beta}(\omega + 1, J)$, where J is the number of populations.
- (ii) Sample v_ω from

$$p(v_\omega = v \mid u_1, \dots, u_M) = \binom{M}{v} d^v \Gamma \left(a + \sum_{m=1}^M T_m - v \right) \left(b - \sum_{m=1}^M \log(u_m) \right)^v$$

for $v \in \{0, \dots, M\}$, where T_m is the number of distinct values in θ_m , for $m = 1, \dots, M$.

- (iii) Sample ω from $\text{Gamma} \left(a + \sum_{m=1}^M T_m - v, b - \sum_{m=1}^M \log(u_m) \right)$.

D. Additional simulations studies

This section provides additional results obtained from simulation studies. It is divided in four sub-sections based on the data generating process (DGP) used to simulate observations.


D.1. Generating mechanism with underlying relevant factor

The DGP used in this sub-section is the same considered in Section 5.1 of the main paper, i.e.


$$\begin{aligned} X_{i,1} &\stackrel{iid}{\sim} 0.5\mathcal{N}(0, 0.5) + 0.5\mathcal{N}(2, 0.5) && \text{for } i = 1, \dots, n_1 \\ X_{i,2} &\stackrel{iid}{\sim} 0.5\mathcal{N}(2, 0.5) + 0.5\mathcal{N}(4, 0.5) && \text{for } i = 1, \dots, n_2 \\ X_{i,3} &\stackrel{iid}{\sim} 0.5\mathcal{N}(4, 0.5) + 0.5\mathcal{N}(6, 0.5) && \text{for } i = 1, \dots, n_3 \\ X_{i,4} &\stackrel{iid}{\sim} 0.5\mathcal{N}(6, 0.5) + 0.5\mathcal{N}(8, 0.5) && \text{for } i = 1, \dots, n_4 \end{aligned}$$

D.1.1. Inferential results from two additional randomly selected studies


In Figures 2 and 3 below, we display the plots regarding the inference for two additional randomly selected simulation studies among the 100 of Section 5.1. Like for the simulation study already discussed in Section 5.1, the true means belong to the 95% credible intervals and the model correctly identifies the two clusters.




(a) Inference on location parameters



(b) Number of second-level clusters.




(c) Co-clustering.




(d) Co-clustering.


Fig 2: Panel (a): Results of the 37th simulation study. Mean point estimates and 95% credible intervals for the four populations, vertical lines correspond to true values. Panel (b): Posterior distribution on the number of second-level clusters. Panels (c) and (d): heatmaps of second level clustering, darker colors correspond to higher probability of co-clustering; in (c) patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities and in (d) patients are reordered based on co-clustering probabilities.




(a) Inference on location parameters



(b) Number of second-level clusters.



(c) Co-clustering.



(d) Co-clustering.

Fig 3: Results of the 9th simulation study. Panel (a): Mean point estimates and 95% credible intervals for the four populations, vertical lines correspond to true values. Panel (b): Posterior distribution on the number of second-level clusters. Panels (c) and (d): heatmaps of second level clustering, darker colors correspond to higher probability of co-clustering; in (c) patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities and in (d) patients are reordered based on co-clustering probabilities.

D.1.2. Inference results with alternative priors

Below we display the results obtained on the simulated data of Section 5.1 of the main paper using the alternative priors described in Section C.

TABLE 1
Simulation studies summaries.

	sHDP-with mixture of DPs			sHDP-with unifor prior		
Partitions	MAP count	Average post. prob.	Median post. prob.	MAP count	Average post. prob.	Median post. prob.
$\{1,2,3,4\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2,3,4\}$	0	0.000	0.000	0	0.000	0.000
$\{1,2\}\{3,4\}$	0	0.000	0.000	0	0.000	0.000
$\{1,3,4\}\{2\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2\}\{3,4\}$	5	0.083	0.022	0	0.030	0.009
$\{1,2,3\}\{4\}$	0	0.000	0.000	0	0.001	0.000
$\{1,4\}\{2,3\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2,3\}\{4\}$	2	0.056	0.012	1	0.051	0.014
$\{1,3\}\{2,4\}$	0	0.000	0.000	0	0.000	0.000
$\{1,2,4\}\{3\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2,4\}\{3\}$	0	0.000	0.000	0	0.000	0.000
$\{1,2\}\{3\}\{4\}$	0	0.002	0.000	0	0.003	0.000
$\{1,3\}\{2\}\{4\}$	0	0.000	0.000	0	0.000	0.000
$\{1,4\}\{2\}\{3\}$	0	0.000	0.000	0	0.000	0.000
$\{1\}\{2\}\{3\}\{4\}$	93	0.859	0.918	99	0.916	0.956

Both models perform better than the NDP, whose results are in Table 1 of the main paper, confirming the advantages of location-based clustering in presence of small sample sizes, when compared to distribution-based clustering. Moreover, sHDP-with mixture of DPs has a slightly worse performance with respect to our main proposal; this was expected, since the corresponding prior incorporates less information and ignores the natural order of the four populations.

D.2. Generating mechanism with outliers

We present here a simulation study with a twofold goal: (1) compare again the location-based clustering approach of our proposal with the distribution-based clustering approach of the nested Dirichlet process (NDP) under a different DGP; (2) study the performance of our model in presence of outliers. The simulated data have been sampled according to the following DGP


$$\begin{aligned}
 \text{DGP 1: } & X_{i,1} \stackrel{iid}{\sim} \mathcal{N}(0, 0.5) & \text{for } i = 1, \dots, n_1 - 1 \\
 & X_{n_1,1} \sim \mathcal{N}(4, 0.5) \\
 & X_{i,2} \stackrel{iid}{\sim} \mathcal{N}(1, 0.5) & \text{for } i = 1, \dots, n_2 \\
 & X_{i,3} \stackrel{iid}{\sim} \mathcal{N}(1, 0.5) & \text{for } i = 1, \dots, n_3 \\
 & X_{i,4} \stackrel{iid}{\sim} \mathcal{N}(2, 0.5) & \text{for } i = 1, \dots, n_4
 \end{aligned}$$

Thus, the true partition is $\{1\}, \{2, 3\}, \{4\}$. Moreover, there is one outlier in the first sample.


TABLE 2
Posterior probabilities over the space of partitions.

	s-HDP	NDP
$\{1,2,3,4\}$	0	0
$\{1\}\{2,3,4\}$	0	0
$\{1,2\}\{3,4\}$	0	0
$\{1,3,4\}\{2\}$	0	0
$\{1\}\{2\}\{3,4\}$	0	0
$\{1,2,3\}\{4\}$	0.013	0.980
$\{1,4\}\{2,3\}$	0	0
$\{1\}\{2,3\}\{4\}$	0.771	0.010
$\{1,3\}\{2,4\}$	0	0
$\{1,2,4\}\{3\}$	0	0
$\{1\}\{2,4\}\{3\}$	0	0
$\{1,2\}\{3\}\{4\}$	0.006	0.020
$\{1,3\}\{2\}\{4\}$	0	0
$\{1,4\}\{2\}\{3\}$	0	0
$\{1\}\{2\}\{3\}\{4\}$	0.210	0

Table 2 displays the posterior probabilities obtained using our model (s-HDP) and the NDP. Our model largely outperforms the competitor.



(a) Co-clustering.



(b) Co-clustering.

Fig 4: Posterior similarity matrices for the simulation study under DGP 1. In (a) patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities; in (b) patients are reordered based on co-clustering probabilities.

Figure 4 shows the posterior co-clustering probabilities obtained in the simulations study. Our proposal is able to correctly identify the outlier.

D.3. Generating mechanism with non-location effects

We present a simulation study with a twofold goal: (1) compare again the location-based clustering approach of our proposal with the distribution-based clustering approach of the nested Dirichlet process (NDP) under a different DGP; (2) study the performance of our model for the case of heterogeneity between populations not being fully explained by shift in locations. The simulated data have been sampled according to the following DGP.


$$\begin{aligned}
 \text{DGP 2: } X_{i,1} &\stackrel{iid}{\sim} 0.5\mathcal{N}(-1, 0.5) + 0.5\mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_1 \\
 X_{i,2} &\stackrel{iid}{\sim} \mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_2 \\
 X_{i,3} &\stackrel{iid}{\sim} \mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_3 \\
 X_{i,4} &\stackrel{iid}{\sim} \mathcal{N}(2, 0.5) && \text{for } i = 1, \dots, n_4
 \end{aligned}$$

Thus, the true partition is $\{1\}, \{2, 3\}, \{4\}$. Moreover, the relative effect of the first population w.r.t. the others is not fully explained by the shift of the location, since the whole distribution is different and not only the mean.


TABLE 3
Posterior probabilities over the space of partitions.

	s-HDP	NDP
$\{1, 2, 3, 4\}$	0	0
$\{1\}\{2, 3, 4\}$	0.001	0
$\{1, 2\}\{3, 4\}$	0	0
$\{1, 3, 4\}\{2\}$	0	0
$\{1\}\{2\}\{3, 4\}$	0.001	0
$\{1, 2, 3\}\{4\}$	0.058	0.98
$\{1, 4\}\{2, 3\}$	0	0
$\{1\}\{2, 3\}\{4\}$	0.706	0.01
$\{1, 3\}\{2, 4\}$	0	0
$\{1, 2, 4\}\{3\}$	0	0
$\{1\}\{2, 4\}\{3\}$	0	0
$\{1, 2\}\{3\}\{4\}$	0.019	0.02
$\{1, 3\}\{2\}\{4\}$	0	0
$\{1, 4\}\{2\}\{3\}$	0	0
$\{1\}\{2\}\{3\}\{4\}$	0.214	0

Table 3 displays the posterior probabilities obtained using our model (s-HDP) and the NDP. Our model largely outperforms the competitor.



(a) Co-clustering.



(b) Co-clustering.

Fig 5: Posterior similarity matrices for the simulation study under DGP 2. In (a) patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities; in (b) patients are reordered based on co-clustering probabilities.

Figure 5 shows the posterior co-clustering probabilities. Our proposal is able to correctly identify the non-location effect (see Figure 5(a)).

D.4. Simulation studies under non-symmetric data generating process

Here we report three simulation studies to check the performance of the model in presence of deviations from symmetry. The simulated data have been sampled according to the following DGPs.


$$\begin{aligned}
 \text{DGP 3: } & X_{i,1} \stackrel{iid}{\sim} \mathcal{N}(0, 0.5) && \text{for } i = 1, \dots, n_1 \\
 & X_{i,2} \stackrel{iid}{\sim} \text{Gamma}(3, 3) && \text{for } i = 1, \dots, n_2 \\
 & X_{i,3} \stackrel{iid}{\sim} \mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_3 \\
 & X_{i,4} \stackrel{iid}{\sim} \mathcal{N}(2, 0.5) && \text{for } i = 1, \dots, n_4 \\
 \\
 \text{DGP 4: } & X_{i,1} \stackrel{iid}{\sim} 0.7\mathcal{N}(-1, 0.5) + 0.3\mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_1 \\
 & X_{i,2} \stackrel{iid}{\sim} \mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_2 \\
 & X_{i,3} \stackrel{iid}{\sim} \mathcal{N}(1, 0.5) && \text{for } i = 1, \dots, n_3 \\
 & X_{i,4} \stackrel{iid}{\sim} \mathcal{N}(2, 0.5) && \text{for } i = 1, \dots, n_4 \\
 \\
 \text{DGP 5: } & X_{i,1} \stackrel{iid}{\sim} \text{Gamma}(10, 10) && \text{for } i = 1, \dots, n_1 \\
 & X_{i,2} \stackrel{iid}{\sim} \text{Gamma}(10, 10) && \text{for } i = 1, \dots, n_2 \\
 & X_{i,3} \stackrel{iid}{\sim} \text{Gamma}(10, 10) && \text{for } i = 1, \dots, n_3 \\
 & X_{i,4} \stackrel{iid}{\sim} 0.5\mathcal{N}(0, 0.5) + 0.5\mathcal{N}(2, 0.5) && \text{for } i = 1, \dots, n_4
 \end{aligned}$$

Under all DGPs the model is misspecified due to lack of symmetry in one or more populations. Under DGP 3 and DGP 4 the true partition is $\{1\}, \{2, 3\}, \{4\}$, while under DGP 5 it is $\{1, 2, 3, 4\}$. In DGP 3 the second population differs from the others also in distribution (what we called non-location effect), the same is true for the first and the fourth populations respectively under DGP 4 and DGP 5. Table 4 shows that the model is able to detect the right clustering of the population-specific locations under all three DGPs. Moreover, Figure 6 shows co-clustering probabilities that differ in correspondence of the populations affected by non-location effects, more or less evidently based on the DGP used to generate the data. These results under misspecification are reassuring: the model appears


robust in estimating the partitions of the locations and, moreover, the different within-population patterns of co-clustering probabilities still highlight heterogeneities different than shifts in population-specific locations.

TABLE 4
Posterior probabilities over the space of partitions.


	DGP 3	DGP 4	DGP 5
$\{1,2,3,4\}$	0	0.001	0.494
$\{1\}\{2,3,4\}$	0	0	0.023
$\{1,2\}\{3,4\}$	0	0	0.014
$\{1,3,4\}\{2\}$	0	0	0
$\{1\}\{2\}\{3,4\}$	0	0	0.004
$\{1,2,3\}\{4\}$	0.016	0	0.375
$\{1,4\}\{2,3\}$	0	0	0
$\{1\}\{2,3\}\{4\}$	0.736	0.788	0.047
$\{1,3\}\{2,4\}$	0	0	0
$\{1,2,4\}\{3\}$	0	0	0
$\{1\}\{2,4\}\{3\}$	0	0	0
$\{1,2\}\{3\}\{4\}$	0.015	0	0.030
$\{1,3\}\{2\}\{4\}$	0	0	0
$\{1,4\}\{2\}\{3\}$	0	0	0
$\{1\}\{2\}\{3\}\{4\}$	0.232	0.211	0.012



(a) Co-clustering DGP 3.



(b) Co-clustering DGP 4.




(c) Co-clustering DGP 5.

Fig 6: Posterior similarity matrices under DGP 3-4-5. Patients are ordered based on the diagnosis and the four black squares highlight the within-sample probabilities.


E. Hypertensive dataset

E.1. Additional results


The figures below report the density estimates, the heatmaps of co-clustering probabilities between pairs of patients and population-specific credible intervals for all ten response variables.



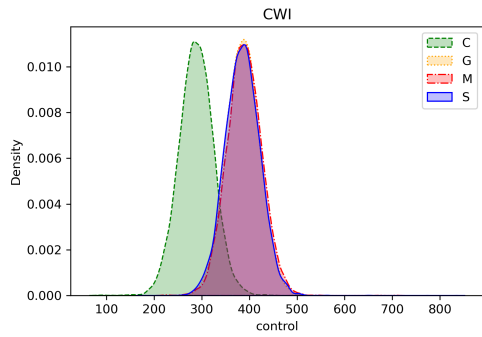
(a) density estimation




(b) co-clustering




(c) co-clustering



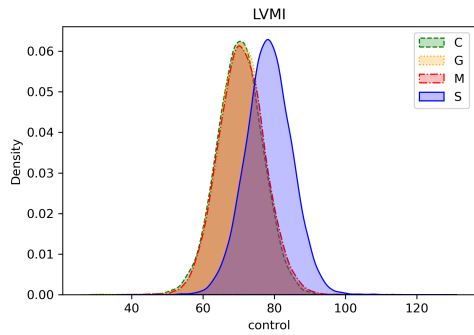
(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering




(a) density estimation




(b) co-clustering




(c) co-clustering





(a) density estimation











(b) co-clustering



(c) co-clustering





E.2. Prior sensitivity to hyperpriorparameters

Here we verify the robustness of the model w.r.t. different specifications of the hyperparameters. We consider two alternative specifications for the hyperparameters, which differ from the one employed in Section 5.2.

Prior specification 1: $G_m = N(0, 1)$; $P_{0,m} \equiv \text{NIG}(\mu = 0, \tau = 0.01, \alpha = 3, \beta = 3)$; all concentration parameters have prior equal to $\text{Gamma}(3, 3)$


Prior specification 2: $G_m = N(0, 2)$; $P_{0,m} \equiv \text{NIG}(\mu = 0, \tau = 1, \alpha = 2, \beta = 4)$; all concentration parameters have prior equal to $\text{Gamma}(0.1, 0.1)$

The model turns out to be rather robust w.r.t. the choice of the hyperparameters, leading to the same selected models for all cardiac indexes under all considered specifications. The detailed results are in the following tables and figures, which report the posterior over partitions of locations, the density estimates, and the posterior similarity matrices for the last cardiac index.


TABLE 5

Posterior probabilities over partitions of means, using prior specification 1. Maximum a posteriori probabilities are in **bold**.


partitions	CI	CWI	LVMI	IVST	LVPW	EF	FS	EW	AW	E/A
$\{C, G, M, S\}$	0.018	0.000	0.000	0.000	0.000	0.371	0.276	0.100	0.000	0.000
$\{C\}\{G, M, S\}$	0.002	0.526	0.001	0.086	0.015	0.068	0.207	0.025	0.028	0.000
$\{C, G\}\{M, S\}$	0.002	0.000	0.000	0.000	0.000	0.038	0.035	0.058	0.072	0.045
$\{C, M, S\}\{G\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C\}\{G\}\{M, S\}$	0.001	0.139	0.000	0.021	0.023	0.025	0.087	0.034	0.244	0.054
$\{C, G, M\}\{S\}$	0.436	0.000	0.612	0.000	0.000	0.279	0.04	0.499	0.007	0.001
$\{C, S\}\{G, M\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C\}\{G, M\}\{S\}$	0.157	0.100	0.180	0.542	0.678	0.073	0.172	0.103	0.265	0.026
$\{C, M\}\{G, S\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C, G, S\}\{M\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C\}\{G, S\}\{M\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C, G\}\{M\}\{S\}$	0.252	0.000	0.092	0.000	0.000	0.081	0.054	0.113	0.087	0.361
$\{C, M\}\{G\}\{S\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C, S\}\{G\}\{M\}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\{C\}\{G\}\{M\}\{S\}$	0.131	0.234	0.116	0.351	0.284	0.066	0.130	0.068	0.295	0.513



(a) density estimation



(b) co-clustering




(c) co-clustering


TABLE 6

Posterior probabilities over partitions of means, using prior specification 2. Maximum a posteriori probabilities are in **bold**.


partitions	CI	CWI	LVMI	IVST	LVPW	EF	FS	EW	AW	E/A
{C,G,M,S}	0.023	0.000	0.000	0.000	0.000	0.341	0.281	0.109	0.000	0.000
{C}{G,M,S}	0.002	0.484	0.000	0.097	0.055	0.079	0.199	0.028	0.029	0.000
{C,G}{M,S}	0.002	0.000	0.001	0.000	0.000	0.036	0.029	0.042	0.074	0.058
{C,M,S}{G}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G}{M,S}	0.001	0.134	0.001	0.022	0.028	0.033	0.090	0.033	0.238	0.068
{C,G,M}{S}	0.408	0.000	0.585	0.000	0.000	0.254	0.036	0.494	0.014	0.001
{C,S}{G,M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G,M}{S}	0.145	0.111	0.184	0.530	0.643	0.077	0.172	0.105	0.254	0.019
{C,M}{G,S}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,G,S}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G,S}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,G}{M}{S}	0.247	0.000	0.097	0.000	0.000	0.089	0.050	0.106	0.076	0.346
{C,M}{G}{S}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,S}{G}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G}{M}{S}	0.172	0.270	0.131	0.351	0.274	0.091	0.144	0.084	0.315	0.508



(a) density estimation



(b) co-clustering



(c) co-clustering


E.3. s-HDP with uniform prior estimates on the Hypertensive Dataset

Here we report the results on the real dataset, obtained with the s-HDP with independent uniform priors on disease-specific locations, described in Section C.1. This prior induces independence between different cardiac indexes and no borrowing of information (i.e. penalization for multiplicity) is applied. Moreover, compared to the priors used in Section 5 of the main paper, here the prior associates higher probability to finer partitions and, thus, does not apply a Ockham's-razor penalty, resulting in a different MAP for the EF.


TABLE 7

Posterior probabilities over partitions obtained through independent uniform priors. Maximum a posteriori probabilities are in **bold**.


partitions	CI	CWI	LVMI	IVST	LVPW	EF	FS	EW	AW	E/A
{C,G,M,S}	0.009	0.000	0.000	0.000	0.000	0.248	0.216	0.047	0.000	0.000
{C}{G,M,S}	0.002	0.568	0.001	0.084	0.014	0.078	0.205	0.027	0.039	0.000
{C,G}{M,S}	0.003	0.000	0.002	0.000	0.000	0.082	0.079	0.160	0.102	0.055
{C,M,S}{G}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G}{M,S}	0.001	0.143	0.001	0.024	0.029	0.027	0.087	0.041	0.262	0.064
{C,G,M}{S}	0.376	0.000	0.555	0.000	0.000	0.324	0.060	0.422	0.005	0.002
{C,S}{G,M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G,M}{S}	0.157	0.115	0.188	0.614	0.730	0.078	0.189	0.096	0.304	0.045
{C,M}{G,S}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,G,S}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G,S}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,G}{M}{S}	0.353	0.000	0.173	0.000	0.000	0.125	0.088	0.162	0.087	0.378
{C,M}{G}{S}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C,S}{G}{M}	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
{C}{G}{M}{S}	0.099	0.174	0.079	0.278	0.227	0.039	0.077	0.045	0.201	0.457
$\sum \log_{15} (p_i^{-p_i})$	0.493	0.426	0.432	0.352	0.269	0.664	0.725	0.624	0.603	0.448



(a) density estimation



(b) co-clustering



(c) co-clustering

E.4. NDP estimates on the Hypertensive Dataset

Finally we display the results obtained with ten independent NDPs used on the real dataset. As expected, the NDP tends to identify coarser partitions. Moreover, the independence between cardiac indexes of the NDP approach leads to more concentrated posterior probabilities, because no borrowing of information (i.e. penalization for multiplicity) is applied.

TABLE 8
Posterior probabilities over partitions obtained through independent NDPs. Maximum a posteriori probabilities are in **bold**.

partitions	CI	CWI	LVMI	IVST	LVPW	EF	FS	EW	AW	E/A
{C,G,M,S}	0.117	0.000	0.000	0.000	0.000	0.613	0.394	0.116	0.000	0.000
{C}{G,M,S}	0.004	0.999	0.001	0.696	0.663	0.047	0.099	0.049	0.313	0.000
{C,G}{M,S}	0.010	0.000	0.014	0.000	0.001	0.027	0.035	0.206	0.043	0.768
{C,M,S}{G}	0.013	0.000	0.000	0.000	0.000	0.040	0.067	0.051	0.000	0.000
{C}{G}{M,S}	0.001	0.000	0.001	0.013	0.163	0.005	0.015	0.088	0.468	0.013
{C,G,M}{S}	0.552	0.000	0.906	0.000	0.000	0.103	0.091	0.154	0.002	0.000
{C,S}{G,M}	0.070	0.000	0.000	0.000	0.000	0.025	0.069	0.029	0.000	0.000
{C}{G,M}{S}	0.077	0.001	0.010	0.207	0.136	0.010	0.032	0.050	0.093	0.006
{C,M}{G,S}	0.009	0.000	0.003	0.023	0.000	0.035	0.045	0.017	0.001	0.000
{C,G,S}{M}	0.047	0.000	0.000	0.000	0.000	0.068	0.081	0.073	0.000	0.000
{C}{G,S}{M}	0.003	0.001	0.000	0.052	0.027	0.007	0.022	0.012	0.030	0.000
{C,G}{M}{S}	0.065	0.000	0.047	0.000	0.000	0.011	0.016	0.071	0.007	0.208
{C,M}{G}{S}	0.025	0.000	0.017	0.004	0.000	0.007	0.017	0.033	0.001	0.000
{C,S}{G}{M}	0.006	0.000	0.000	0.000	0.000	0.006	0.014	0.023	0.000	0.000
{C}{G}{M}{S}	0.007	0.000	0.002	0.007	0.012	0.001	0.007	0.032	0.044	0.006
$\sum \log_{15} \left(p_i^{-p_i} \right)$	0.603	0.016	0.167	0.349	0.368	0.567	0.785	0.898	0.509	0.239

References

- Blackwell, D. and J. B. MacQueen (1973). Ferguson distributions via Pólya urn schemes. *Annals of Statistics* 1, 353–355.
- Dalal, S. (1979a). Dirichlet invariant processes and applications to nonparametric estimation of symmetric distribution functions. *Stochastic Processes and their Applications* 9, 99–107.
- Dalal, S. R. (1979b). Nonparametric and robust Bayes estimation of location. *Optimizing methods in statistics* 9, 141–166.
- Diaconis, P. and D. Freedman (1986). On inconsistent Bayes estimates of location. *Annals of Statistics* 14, 68–87.
- Doss, H. (1984). Bayesian estimation in the symmetric location problem. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 68, 127–147.
- Escobar, M. D. and M. West (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* 90, 577–588.
- Ferguson, T. S., E. G. Phadia, and R. C. Tiwari (1992). Bayesian nonparametric inference. *Lecture Notes-Monograph Series* 17, 127–150.
- Ghosal, S., J. K. Ghosh, and R. Ramamoorthi (1999). Consistent semiparametric Bayesian inference about a location parameter. *Journal of Statistical Planning and Inference* 77, 181–193.
- Hannum, R. and M. Hollander (1983). Robustness of Ferguson’s Bayes estimator of a distribution function. *The Annals of Statistics*, 632–639.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* 4, 639–650.
- Tiwari, R. C. (1988). Convergence of Dirichlet invariant measures and the limits of Bayes estimates. *Communications in Statistics-Theory and Methods* 17, 375–393.