A Bayesian Mixture Model for Clustering and Selection of Feature Occurrence Rates under Mean Constraints

Qiwei Li¹, Michele Guindani², Brian J. Reich³, Howard D. Bondell[†]
and

Marina Vannucci^{*4}

Mon May 22 21:12:46 2017

In this paper, we consider the problem of modeling a matrix of count data, where multiple features are observed as counts over a number of samples. Due to the nature of the data generating mechanism, such data are often characterized by a high number of zeros and overdispersion. In order to take into account the skewness and heterogeneity of the data, some type of normalization and regularization is necessary for conducting inference on the occurrences of features across samples. We propose a zero-inflated Poisson mixture modeling framework that incorporates a model-based normalization through prior distributions with mean constraints, as well as a feature selection mechanism, which allows us to identify a parsimonious set of discriminatory features, and simultaneously cluster the samples into homogenous groups. We show how our approach improves on the accuracy of the clustering with respect to more standard approaches for the analysis of count data, by means of a simulation study and an application to a *bag-of-words* benchmark data set, where the features are represented by the frequencies of occurrence of each word.

Keywords: count data; Bayesian nonparametrics; Poisson mixture; feature selection; text analysis.

1 Introduction

¹ Department of Statistics, Rice University

² Department of Statistics, University of California, Irvine

³ Department of Statistics, North Carolina State University

⁴ Corresponding author, marina@rice.edu

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: [10.1002/sam.11350](https://doi.org/10.1002/sam.11350)

Modern data science often involves data sets where the features of interest are measured as counts. For example, text documents are typically summarized by their word frequencies, with the size of the dictionary determining the number of features [?, see e.g.]Griffiths2004, Airolidi2006. In those applications, clustering and feature selection techniques are often employed to provide low-dimensional summaries of the data and investigate the topical content of a text [3, 4]. Similarly, in ecological survey data [?, see e.g.]Gee1985, species counts are observed on a relatively large number of sites, with the objective of characterizing the different ecosystems across the sites. In biology, sequence data – e.g., SAGE data, RNA-Seq and microbiome data [6, 7, 8] – are represented as a matrix of short sequence tags (e.g. taxonomic units in a microbiome experiment) and corresponding observed reads for several samples. Investigators are typically interested in discovering tags whose abundances significantly differ across samples.

One common characteristic of those studies is that the observed frequency of a feature depends on the sampling effort, or – for the analysis of text documents – the document length. This usually results in datasets characterized by two distinctive attributes. On the one hand, the datasets contain a high percentage of zero counts. A zero count can either indicate a missing trait in the population or be due to a limited sample. Furthermore, the datasets are highly variable, both with respect to the total number of counts per sample and to the total number of counts per feature. Thus, the observed distributions of counts are typically skewed and overdispersed, since a large number of features are recorded at low frequencies whereas a few features are recorded very frequently. The amount of overdispersion may also vary sample to sample [9, 10].

In order to take into account the skewness and heterogeneity of the data, some type of normalization and regularization is often necessary for conducting inference on the features' occurrence rates [11, 12, 13, 14]. For example, [9] proposes a Poisson log linear model, where the Poisson intensities are robustly estimated after a normalization step which takes into account the total number of reads observed for each sample and each feature. More recently, [4] have proposed a Hierarchical Poisson modeling framework where the word occurrence rate is moderated by the document's length, and a low-dimensional representation of the data is achieved by means of a set of latent (mixture) components and sparsity inducing shrinkage priors.

In this paper, we propose a novel regularization approach for estimating occurrence rates in zero-inflated Poisson mixture models. Zero-inflated Poisson models have often been employed to fit count data characterized by overdispersion and a high number of zeros, both in the econometrics and the statistics literature [15, 16, 17]. Our proposal is characterized by priors that employ "soft" constraints on expected values of both samples' and features' scaling parameters, in order to normalize the information content of each sample. This is in contrast to typical approaches that use "hard" constraints as data-dependent plug-in estimates based on

the observed counts, see for example [14], [18], [7], [19], and [9], which *a-priori* condition the inference over unknown parameters on the observed counts. Such plug-in estimates can be regarded as maximum likelihood estimators in multiple stage approaches and somewhat akin to Empirical Bayes methods, therefore relying on implicit assumptions of exchangeability of the observations, which may not be always justified in practice and can introduce bias in the estimation of posterior uncertainties [20, 21]. In addition, “soft” constraints based on a moderate amount of information have been shown to produce more flexible and less biased estimates [22]. Our approach is similar in spirit to the recent proposal by [23], who impose a stochastic constraint on the quantile functions of infinite mixtures, in order to avoid identifiability restrictions in Bayesian nonparametric approaches for quantile regression.

We further regularize the estimation problem by allowing priors that enable feature selection and discriminate samples into clusters. In the analysis of a text document, for example, our approach allows to identify a subset of discriminatory features (words) which are exclusive to particular topics, identified from clusters of documents, automatically balancing the influence of frequency and exclusivity of words across samples by virtue of our model-based inference [24, 4]. More specifically, to achieve feature selection we introduce latent discriminatory variables, similarly as in [25] and [26], who proposed the use of latent indicators for variable selection in the context of finite mixtures of Gaussian distributions. Samples are then clustered on the basis of the similarity of the resulting vectors of discriminatory features by means of an infinite Dirichlet process mixture (DPM). Dirichlet process mixtures have been often employed in Bayesian modeling for clustering purposes, since they allow to estimate the number of clusters (i.e. the number of mixing components) directly from the data [27]. Alternatively, one could consider a finite mixture model, and use reversible jump Markov chain Monte Carlo (RJMCMC) to determine the number of components, at the expense of increased computational cost [28, 25].

In recent Bayesian nonparametric literature, it has become common to employ an Indian Buffet Process (IBP) characterization to identify non-zero elements in a general matrix of n samples $\times p$ (latent) features [29]. For example, IBP Compound Dirichlet process models [30] and Beta-negative binomial processes [31] have been introduced for the analysis of count data in topic modeling. Differently than those, the constrained Poisson mixtures (CoPoM) approach we propose is aimed at normalizing the data through the use of mean constraints, and clustering the available samples based on the entire subset of selected observed features, rather than assigning single elements of the $n \times p$ matrix to either one of the mixture components (e.g. a topic) or none.

By means of a simulation study, we show how our approach improves on the accuracy of the clustering performance with respect to more standard approaches for the analysis of count data. We then present an application to a *bag-of-words* benchmark data set, where the features are represented by the frequencies of occurrence of each word.

The rest of the paper is organized as follows. In Section 2 we introduce the zero-inflated hierarchical mixture model and discuss the prior formulations. In Section 3 we briefly describe the MCMC algorithm and discuss the resulting posterior inference. In Section 4 we illustrate the performance of our method on simulated data and then present an application to document clustering, based on a *bag-of-words* benchmark data set. Section 5 concludes the paper with a discussion of the modeling choices, namely the use of a Poisson likelihood versus alternatives, and with future research directions.

2 Zero-inflated Poisson mixture model with feature selection

We consider a $n \times p$ matrix of counts, x_{ij} , $i = 1, \dots, n$, $j = 1, \dots, p$, observed on a set of p features and n samples. We assume that a large number of the counts is zero, either because the feature is truly missing in a subset of the population or due to limitations of the sampling effort. Thus, we start by considering a zero-inflated Poisson mixture model, i.e. a mixture model where we constrain one of the kernels to be degenerate at zero,

$$x_{ij} \sim \pi \delta_0(x_{ij}) + (1 - \pi) \int \text{Poi}(x_{ij}; \lambda) G(d\lambda), \quad (1)$$

where $\pi \in [0, 1]$, $\text{Poi}(x; \lambda)$ denotes a Poisson distribution for the random variable x , with expectation λ , $\delta_c(\cdot)$ indicates a point mass distribution on $c \in \mathbb{R}$, and $G(\cdot)$ denotes a general mixing distribution, which we use to model the overdispersion of the data. Note that if $G(\cdot)$ is Gamma distributed, then (1) defines a zero-inflated Negative Binomial [17]. Alternatively, we can write (1) also by introducing latent indicator variables $r_{ij} \sim \text{Bern}(\pi)$, such that if $r_{ij} = 1$ then $x_{ij} = 0$, whereas if $r_{ij} = 0$ then $x_{ij} \sim \int \text{Poi}(x_{ij}; \lambda) G(d\lambda)$.

2.1 Feature selection

We envision that only some of the features are relevant to discriminate the n samples into distinct clusters. In particular, we postulate the existence of a latent binary vector $\gamma = (\gamma_1, \dots, \gamma_p)$, with $\gamma_j = 1$ if the counts associated to feature j are relevant to discriminate among the n samples, and $\gamma_j = 0$ otherwise. Let $p_\gamma = \sum_{j=1}^p \gamma_j$ denote the number of discriminatory features and, correspondingly, let $p - p_\gamma$ indicate the number of non-informative features. Our model formulation assumes that any count, say x_{ij} , that maps to a discriminatory feature is drawn from a zero-inflated Poisson distribution with intensity parameter λ_{ij} , while counts that map to non-discriminatory features are drawn from a “null” model, which can be characterized as a zero-inflated Poisson distribution with intensity

parameter λ_{ij}^0 . In other words, our representation allows for features to be discriminatory, although their count may be zero in some samples. Conditioning on $r_{ij} = 0$, we can write our model as follows,

$$x_{ij} | r_{ij} = 0, \gamma_j, \lambda_{ij} \underset{ind}{\sim} \begin{cases} \text{Poi}(x_{ij}; \lambda_{ij}) & \text{if } \gamma_j = 1 \\ \text{Poi}(x_{ij}; \lambda_{ij}^0) & \text{if } \gamma_j = 0, \end{cases} \quad (2)$$

whereas if $r_{ij} = 1$ we assume $x_{ij} = 0$, irrespectively of γ_j , with $i = 1, \dots, n$ and $j = 1, \dots, p$. A common choice for the prior on the vector γ is to assume independent Bernoulli distributions on the individual components, with a common hyperparameter ω , i.e. $\gamma_j | \omega \sim \text{Bern}(\omega)$, which is equivalent to a Binomial prior on the number of discriminatory features, that is $p_\gamma | \omega \sim \text{Bin}(p, \omega)$. The hyperparameter ω can be elicited as the proportion of features expected *a-priori* to be in the discriminatory set. This prior assumption can be further relaxed by formulating a $\text{Be}(a_\omega, b_\omega)$ hyperprior on ω , which leads to a beta-binomial prior for p_γ with expectation $pa_\omega/(a_\omega + b_\omega)$. A vague prior on ω can be obtained as in [32] by imposing the constraint $a_\omega + b_\omega = 2$, corresponding to a prior effective sample size of 2, with some desirable mean percentage of inclusion. Analogously, we assume $\pi \sim \text{Be}(a_\pi, b_\pi)$ as a prior on the zero-inflation weight.

In order to account for the variability observed across samples and across features, we further parametrize the intensity parameters of the Poisson distributions as the multiplicative effect of three random effects: a) a scaling factor capturing how the sampling effort affects sample-specific occurrences across all features, denoted by s_i ; b) a scaling factor capturing feature-specific levels across all samples, denoted by g_j (e.g. a common usage preposition in the analysis of text documents); and c) a term capturing the occurrence rate for each count, once all the previous global effects have been accounted for, denoted by d_{ij} . Specifically, we parametrize the Poisson intensity in (2) as $\lambda_{ij} = s_i g_j d_{ij}$, if $\gamma_j = 1$, whereas we set $\lambda_{ij}^0 = s_i g_j d_0$ if $\gamma_j = 0$ in the "null" model. Here, d_0 is a parameter capturing homogenous background noise across samples and features. Since background noise should be the result of the variability captured by the scaling factors s_i and g_j , and also to ensure the identifiability of all model parameters, we assume $d_0 = 1$, so that the significance of each feature is completely revealed by the values of the random effects d_{ij} across the samples.

Multiplicative characterizations of the Poisson intensity parameter of model (2) are typical both in the frequentist as well as in the Bayesian literature to account for latent

heterogeneity and overdispersion in count data [?, e.g. see]the latter for examples of this specification in spatial statistics]Witten2011, Cameron2013,Airoidi2016,Banerjee2004. To simplify both the prior specification and the computational algorithms, it is sometimes convenient to reparametrize model (2) by using the logarithmic transformations $\tilde{\lambda}_{ij} = \log\{\lambda_{ij}\}$, and, consequently, $\tilde{s}_i = \log\{s_i\}$, $\tilde{g}_j = \log\{g_j\}$ and $\tilde{d}_{ij} = \log\{d_{ij}\}$, such that $\tilde{\lambda}_{ij} = \tilde{s}_i + \tilde{g}_j + \tilde{d}_{ij}$. Under this reparametrization, the base component is characterized by $\tilde{d}_0 = 0$. In Section 2.2, we describe a regularizing prior specification for the scaling factors that allows flexible modeling of count data, and also avoids some limiting assumptions commonly made when estimating those models.

2.2 Estimation of scaling factors via mean constraints

In Poisson models with a multiplicative intensity λ_{ij} as defined in Section 2.1, the inferential interest is often limited to the estimation of the occurrence rates d_{ij} , whereas scaling factors are often normalized in order to regularize the inferential problem and ensure the identifiability of the relevant parameters d_{ij} . One typical choice is to estimate the scaling factors by means of plug-in estimators based on the observed counts. For example, in the context of the estimation of RNA-seq abundances, [7], [19], and [9] fix $\hat{s}_i = \sum_{j=1}^p x_{ij} / \sum_{i=1}^n \sum_{j=1}^p x_{ij}$, so that $\sum_{i=1}^n s_i = 1$. Similarly, [14] propose $\hat{s}_i = m_i / \sum_{i=1}^n m_i$, where m_i is the median of the distribution of the ratios of the counts for observation i to their geometric mean. As a further example, [18] propose taking $\hat{s}_i = q_i / \sum_{i=1}^n q_i$, with q_i the 75th percentile of the counts for observation i . Many of the above examples further fix $\hat{g}_j = \sum_{i=1}^n x_{ij}$. While convenient, the use of plug-in estimates for estimating s_i and g_j has noticeable drawbacks, since *a-priori* they condition the inference over unknown parameters on the observed data counts. Indeed, such plug-in estimates can be regarded as maximum likelihood estimators in multiple stage approaches and somewhat akin to Empirical Bayes methods, therefore relying on implicit assumptions of exchangeability of the observations, which may not be always justified in practice and can introduce bias in the estimation of posterior uncertainties [20, 21].

One of the contributions of this manuscript is to provide an alternative normalization approach, through the use of priors for the vectors $\mathbf{s} = (s_1, \dots, s_n)$ and $\mathbf{g} = (g_1, \dots, g_p)$, that respectively capture the sampling and feature heterogeneity in the count data, and avoid fixing those values *a-priori*. A simple choice would be to use conjugate priors, $s_i \sim \text{Ga}(a_s, b_s)$ and

$g_j \sim \text{Ga}(a_g, b_g)$. However, the Poisson intensity in (2) depends on the product $s_i g_j$, and further constraints are necessary to allow identifiability of the parameters and simultaneous inference on the occurrence rates d_{ij} . In the context of quantile regression, [23] impose a stochastic constraint on the quantile functions of infinite mixtures to ensure identifiability of the parameter estimates. Here, we consider priors on s_i and g_j that impose normalizing constraints on the expected values, but still provide a flexible estimate of the posterior densities.

For computational convenience, we consider the logarithmic transformation $\tilde{s}_i = \log\{s_i\}$ and $\tilde{g}_j = \log\{g_j\}$, so that $s_i g_j = \exp\{\tilde{s}_i + \tilde{g}_j\}$. We assume that the priors for $\tilde{s}_i$ and $\tilde{g}_j$ are mixture distributions,

$$\tilde{s}_i | \cdot \sim \sum_{m=1}^M \phi_m^s f_m^s(\tilde{s}_i | \cdot), \quad \text{and} \quad \tilde{g}_j | \cdot \sim \sum_{l=1}^L \phi_l^g f_l^g(\tilde{g}_j | \cdot), \quad (3)$$

with $0 \leq \phi_m^s, \phi_l^g \leq 1$, $\sum_{m=1}^M \phi_m^s = \sum_{l=1}^L \phi_l^g = 1$, and M, L are positive integers. The use of mixture distributions allows flexible estimation of the posterior density of $\tilde{s}_i$ and $\tilde{g}_j$. We further build each mixture so to satisfy the desired constraint, i.e. $E[\tilde{s}_i] = c_s$ and $E[\tilde{g}_j] = c_g$, where c_s and c_g are some fixed values. [34] demonstrates that any distribution with a mean constraint can be generated by an infinite sum of two-components mixture distributions, where the two-component mixtures are constrained to have the required expected value. Therefore, we assume that each $f_m^s(\cdot)$ and $f_l^g(\cdot)$ in (3) are themselves a two-component Gaussian mixture, as

$$\begin{aligned} f_m^s(\tilde{s}_i | t_m, \eta_m) &= t_m \text{N}(\eta_m, \sigma_s^2) + (1-t_m) \text{N}\left(\frac{c_s - t_m \eta_m}{1-t_m}, \sigma_s^2\right), \\ f_l^g(\tilde{g}_j | q_l, \mu_l) &= q_l \text{N}(\mu_l, \sigma_g^2) + (1-q_l) \text{N}\left(\frac{c_g - q_l \mu_l}{1-q_l}, \sigma_g^2\right), \end{aligned} \quad (4)$$

with $0 \leq t_m, q_l \leq 1$. It is immediate to check that the densities in (4) satisfy the desired constraint. Of course, one could consider other prior formulations satisfying that constraint. However, the proposed mixture-of-mixtures prior is attractive because it allows flexibility in the estimation of the unknown s_i and g_j 's, by spanning a wide class of distributions, e.g. skewed and multi-modal densities.

Furthermore, if $M = L = \infty$, then (3) can be interpreted as Bayesian nonparametric infinite mixtures. In particular, Dirichlet process mixture models have been extensively used in recent literature for flexible density estimation, both for continuous and discrete data [?, see, e.g.]trippa2011, Kyung2011, taddy2012. The Dirichlet process assumes that the mixing

distribution can be written as a discrete random measure, $F(\cdot) = \sum_{k=1}^{\infty} \phi_k \delta_{\theta_k^*}(\cdot)$, where the weights ϕ_k are defined by the [38] stick-breaking construction, i.e. $\phi_1 = V_1$, $\phi_k = V_k \prod_{u=1}^{k-1} (1 - V_u)$, $V_k \sim \text{Be}(1, \alpha)$, $k = 1, 2, \dots$, and the atoms $\theta_k^* \sim F_0$, with F_0 a baseline parametric model describing the prior expectation of the Dirichlet process. In symbols, we write $F \sim \text{DP}(\alpha, F_0)$. The concentration (or mass) parameter α provides a measure of the precision of the random measure around the baseline parametric model [?, see, for details on the Dirichlet process,]]Hjortal2010, Mullerbook2015. We note that theoretical results on large support and consistency of models based on discrete kernels have not been discussed in the literature. Indeed, given the discreteness of the support on the natural numbers, some technical issues make the derivation of such results more complex than for mixtures of continuous distributions.

We conclude this section by specifying the distributions of the hyper-parameters in (??). More specifically, we assume that the two mixtures are characterized by $\eta_m \sim \text{N}(0, \tau_\eta)$, $t_m \sim \text{Be}(a_t, b_t)$, and $\mu_l \sim \text{N}(0, \tau_\mu)$, $q_l \sim \text{Be}(a_q, b_q)$, whereas ϕ_m^s and ϕ_l^g are obtained according to the stick-breaking construction. Since the aim of this specification is simply to achieve an automatic normalization of the scaling factors, we further assume $\sigma_s^2 = \sigma_g^2 = 1$.

2.3 Clustering selected features via Dirichlet process mixtures

In the analysis of count data across multiple samples, one common objective is to characterize and cluster the observed samples into homogenous groups on the basis of the estimated features' occurrence rates. In this section we provide a method for clustering the n samples, based on the set of selected discriminatory features from Section 2.1. We use the superscript (γ) to index the set of discriminatory features, characterized by $\gamma_j = 1$ in (2). Similarly, (γ^c) indicates the set of non-discriminatory features, characterized by $\gamma_j = 0$ and the "null" Poisson distribution with intensity parameter $\lambda_{ij}^0 = s_i g_j d_0$. Thus, each data sample is represented by the $1 \times p$ vector $\mathbf{x}_i$ of observations $\{x_{ij}\}$, with $\mathbf{x}_i^{(\gamma)}$ and $\mathbf{x}_i^{(\gamma^c)}$ indicating the subsets of features corresponding to $\gamma_j = 1$ and $\gamma_j = 0$, respectively. We assume that samples can be clustered based on the subset of selected features $\mathbf{x}_i^{(\gamma)}$, by means of a zero-inflated infinite mixture of Poisson distributed components. For that purpose, we introduce an auxiliary set of clustering allocation variables, $\mathbf{z} = \{z_1, \dots, z_n\}$, defined so that $z_i = k$ if and only if the vector of observations $\mathbf{x}_i^{(\gamma)}$ belongs to cluster k , for some integer $k \geq 1$. Then, we can characterize the likelihood for the selected features as

$$x_{ij}^{(\gamma)} | z_i = k, r_{ij}, s_i, g_j \sim r_{ij} \overset{ind}{\delta_0}(x_{ij}) + (1 - r_{ij}) \text{Poi}(x_{ij}; \lambda_{ijk}^*) \quad (5)$$

where $\lambda_{ijk}^* = s_i g_j d_{kj}^*$, with $d_{kj}^* \sim \text{Ga}(a, b)$ being a non-negative occurrence rate for feature j , common to all non-zero observations assigned to the k -th cluster. We further assume that each $z_i \sim \sum_{k=1}^{\infty} w_k \delta_{\{k\}}$, where the weights w_k are defined by the [38] stick-breaking construction, so that (5) effectively defines a zero-inflated conjugate Dirichlet process mixture model. This modeling framework allows us to cluster the samples based on the vectors of selected features, by allowing $d_{ij} = d_{jk}^*$ when $r_{ij} = 0$, for some integer k , while at the same time the number of clusters is estimated as a by-product of the usual posterior inference. The dimension of the component-specific vector $\mathbf{d}_k^{*(\gamma)} = \{d_{k1}^*, \dots, d_{kp_\gamma}^*\}$ depends on the outcome of the feature selection procedure. As a matter of fact, the joint prior probability of a given allocation of the n samples into K groups is given by

$$p(z_1, \dots, z_n) = \frac{\alpha^K \Gamma(\alpha) \prod_{k=1}^K \Gamma(n_k)}{\Gamma(\alpha + n)}, \quad (6)$$

that describes the Ewens distribution [40, 41, 42]. Here, α is the concentration parameter of the Dirichlet process mixture model, which defines the expected number of clusters as $E(k) \approx \alpha \log(n)$ [41]. As $\alpha \rightarrow 0$, the number of clusters goes to 1, while for $\alpha \rightarrow \infty$ the number of clusters goes to n . We complete our prior specification by placing a $\text{Ga}(a_\alpha, b_\alpha)$ hyperprior on α as in [43]. In the absence of prior information, we suggest choosing the values of a_α and b_α to obtain a fair degree of support for $\alpha \approx 1$.

The proposed constrained Poisson mixture model is summarized in Figure 1. Our zero-inflated Poisson mixture model provides a more flexible framework for density estimation of overdispersed count data with respect to widely used Negative Binomial models, which are indeed a special case of our framework.

Mixture model
Likelihood:
$x_{ij} \gamma_j = 1, z_i, s_i, g_j, r_{ij} \overset{ind}{\sim} r_{ij} \delta_0(x_{ij})$

$x_{ij} \mid \gamma_j = 0, s_i, g_j, r_{ij} \overset{ind}{\sim} r_{ij} \delta_0(x_{ij}) +$
Zero-inflation latent indicator prior:
$r_{ij} \sim \text{Bern}(\pi), \quad \pi \sim \text{Be}(a_\pi, \ell$
Feature selection prior:
$\gamma_j \sim \text{Bern}(\omega), \quad \omega \sim \text{Be}(a_\omega, \ell$
Mixing distribution for selected discriminatory features:
$z_i \sim \sum_{k=1}^{\infty} w_k \delta_{\{k\}}$
$\mathbf{d}_k^{*(\gamma)} \sim \prod_{\{j: \gamma_j=1\}} \text{Ga}(a, b)$
$w_k = V_k \prod_{u=1}^{k-1} (1 - V_u), V_k \sim \text{Be}(1, \ell$
$\alpha \sim \text{Ga}(a_\alpha, b_\alpha)$
Priors on subject- and feature-specific scaling factors:

$\tilde{s}_i v_i, \varepsilon_i, \mathbf{t}, \boldsymbol{\eta} \sim \sum_{m=1}^M \phi_m^s \left[t_m \text{N}(\eta_m, \sigma$
$\phi_m^s = V_m \prod_{u=1}^{m-1} (1 - V_u), V_m \sim \text{Be}(a_m$
$\eta_m \sim \text{N}(0, \tau_\eta)$
$t_m \sim \text{Be}(a_t, b_t)$
$\tilde{g}_j \xi_j, \psi_j, \mathbf{q}, \boldsymbol{\mu} \sim \sum_{l=1}^L \phi_l^g \left[q_l \text{N}(\mu_l,$
$\phi_l^g = V_l \prod_{u=1}^{l-1} (1 - V_u), V_l \sim \text{Be}(a_l, b$
$\mu_l \sim \text{N}(0, \tau_\mu)$
$q_l \sim \text{Be}(a_q, b_q)$
Fixed
hyperparameters:
$a, b, a_\alpha, b_\alpha, a_\omega, b_\omega,$

Table 1: Hierarchical formulation of the proposed constrained Poisson mixture (CoPoM) model.

3 Model fitting

We now briefly describe the MCMC algorithm for posterior inference. Our inferential strategy allows to simultaneously infer group structure in the samples while identifying the

discriminatory features.

3.1 MCMC algorithm

Our primary interest lies in the identification of the discriminatory features, via the vector γ , and the estimation of the sample clustering allocations, via the vector $\mathbf{z}$. For this, we design a Markov chain Monte Carlo (MCMC) algorithm based on Metropolis search variable selection algorithms [44, 45] and Gibbs sampling methods for Dirichlet process mixture models [46]. We also sample the sample-specific and feature-specific scaling factors, $\mathbf{s}$ and $\mathbf{g}$. We give full details of our MCMC algorithm in the Appendix and report here a brief description of the most relevant updates.

Update of γ : This is done via an *add-delete-swap* algorithm. In this approach, a new candidate vector, say γ^{new} , is generated by randomly choosing between two types of moves. For the add/delete move, we select at random one of the elements in the current vector, say γ^{old} , and change its value from 0 to 1, or vice versa. For the swap move, we select two elements in γ^{old} with different inclusion status and swap their values. Then, the Metropolis-Hastings ratio can be written as

$$m_{MH} = \frac{p(\gamma^{new} | \mathbf{z}, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}) J(\gamma^{old} | \gamma^{new})}{p(\gamma^{old} | \mathbf{z}, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}) J(\gamma^{new} | \gamma^{old})},$$

where $J(\cdot | \cdot)$ indicates the proposal probability distribution for the selected move. The move is accepted with probability $\min(1, m_{MH})$. We should notice that the feature selection and the cluster structure are determined simultaneously in the MCMC algorithm. Therefore, to improve mixing, it is necessary to allow the selection to stabilize for any visited cluster configuration. As suggested in [47], we repeat the above Metropolis step $E = 20$ times within each iteration. In the applications of this paper, no improvement in the MCMC performance was noticed beyond this value.

Update of $\mathbf{z}$: Since we have assumed a conjugate baseline parametric distribution, $\prod_{\{j: \gamma_j=1\}} \text{Ga}(a, b)$, we can integrate analytically over the cluster-specific parameters $\mathbf{d}_k^{*(\gamma)}$ and directly sample the cluster assignment indicators of the selected features, $\mathbf{z}$, according to Algorithm 3 of [46]. More specifically, the Gibbs sampler iteratively samples the full conditionals,

$$p(z_i | \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}) = \begin{cases} \frac{n_{k,-i}}{n-1+\alpha} f(\mathbf{x}_i | z_i = k, \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}_{-i}) \\ 0.8in \text{for } z_i = k, k = 1, \dots, K_{-i}, \\ \frac{\alpha}{n-1+\alpha} f(\mathbf{x}_i | \gamma, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_i) \quad \text{for } z_i = K_{-i} + 1, \end{cases} \quad (7)$$

where $\mathbf{z}_{-i}$ denotes all the elements in $\mathbf{z}$ excluding the i -th one, $n_{k,-i}$ is the size of cluster k in $\mathbf{z}_{-i}$, and K_{-i} is the number of unique values in $\mathbf{z}_{-i}$. Note that $f(\mathbf{x}_i | z_i = k, \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}_{-i})$ is the integrated likelihood, with updated $\mathbf{d}_k^{*(\gamma)}$ based on its prior and all observations except the i -th one. See details in the Appendix.

Update of $\mathbf{s}$ and $\mathbf{g}$: The prior distribution for the scaling factors $\mathbf{s}$ and $\mathbf{g}$ is a Dirichlet process mixture on the log-transformed values $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{g}}$. Since the mixture of log-normal distribution that we have assumed as baseline measure in the Dirichlet process is not conjugate to the Poisson likelihood, we consider a finite truncation of the Dirichlet process, so that M and L in (3) are large but finite [48]. Such a choice allows to simplify computations considerably, and still achieves flexible estimate of posterior densities, such as those commonly obtained in a Bayesian nonparametric framework. We employ Metropolis-Hastings steps for all $\tilde{s}_i$'s and $\tilde{g}_j$'s.

Update of $\mathbf{R}$: The full conditional for the zero-inflation latent indicators takes into account that we only need to update those r_{ij} 's that correspond to zero counts. For positive counts, necessarily $r_{ij} = 0$. We use a Metropolis-Hasting within Gibbs sampling step for updating r_{ij} and π , after sampling d_{kj}^* for those features corresponding to $\gamma_j = 0$. Details are given in the Appendix.

3.2 Posterior inference

We obtain posterior inference on the parameters by post-processing of the MCMC samples after burn-in. We start by obtaining a probabilistic assessment of the cluster allocations by analyzing the MCMC samples of $\mathbf{z}$. One way to summarize the posterior distribution of $\mathbf{z}$ is via the *maximum-a-posteriori* (MAP) estimate that can be calculated as

$$\begin{aligned}\hat{\mathbf{z}}_{\text{MAP}} &= \operatorname{argmax}_{1 \leq b \leq B} p(\mathbf{z}^{(b)} | \gamma^{(b)}, \tilde{\mathbf{s}}^{(b)}, \tilde{\mathbf{g}}^{(b)}, \mathbf{R}^{(b)}, \mathbf{X}) \\ &= \operatorname{argmax}_{1 \leq b \leq B} p(\mathbf{z}) \prod_{i=1}^n f(\mathbf{x}_i | z_i^{(b)}, \mathbf{z}_{-i}^{(b)}, \gamma^{(b)}, \tilde{\mathbf{s}}^{(b)}, \tilde{\mathbf{g}}^{(b)}, \mathbf{R}^{(b)}, \mathbf{X}_{-i}),\end{aligned}$$

with $b = 1, \dots, B$ indicating the MCMC iterations, after burn-in, and where the marginal posterior probability that sample i is allocated to cluster k can be calculated through Equation (??). An alternative estimate relies on the computation of a matrix of posterior pairwise probabilities of co-clustering, i.e. the probabilities that observation i and observation i' are assigned to the same cluster, $p_{ii'} = p(z_i = z_{i'} | \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X})$, as suggested by [49], among others. These probabilities can be estimated by computing empirical frequencies of co-clustering based on the MCMC samples, resulting in an $n \times n$ symmetric pairwise probability matrix (PPM). Then, a point estimate for the cluster memberships, $\hat{\mathbf{z}}_{\text{PPM}}$, is

obtained by minimizing the sum of squared deviations of its association matrix from the PPM, i.e.

$$\hat{\mathbf{z}}_{\text{PPM}} = \underset{\mathbf{z}}{\operatorname{argmin}} \sum_{i < i'} [I(z_i = z_{i'}) - p_{ii'}]^2.$$

As for feature selection, the MAP estimate of γ can be obtained by enumerating all visited MCMC samples $\gamma^{(b)}$ and then considering the set of features that maximizes the posterior density. Alternatively, we can estimate the marginal posterior probability of inclusion (PPI) of single features as the proportion of MCMC iterations, after burn-in, in which the corresponding γ_j were equal to 1, that is $\text{PPI}(j) = \sum_{u=1}^B (\gamma_j^{(b)} | \mathbf{z}^{(b)}, \tilde{\mathbf{s}}^{(b)}, \tilde{\mathbf{g}}_j^{(b)}, \mathbf{r}_j^{(b)}, \mathbf{x}_j) / B$. A point estimate of γ is then obtained by identifying those PPI values that exceed a given threshold. The optimal threshold is typically chosen based on a decision theoretic criterion, e.g. to maximize power under a constraint on the number of false positives [50, 51].

4 Applications

We first explore performances on simulated data and then show results on a *bag-of-words* benchmark data set. We also demonstrate the superiority of our constrained Poisson mixture (CoPoM) model over other widely adopted methods for the analysis of overdispersed count data.

4.1 Simulation study

Data were generated with $n = 30$ samples and $K = 3$ clusters. We then simulated observations from a zero-inflated mixture of Poisson distributions, assuming three mixture components and $p_\gamma = 50$ discriminating features,

$$x_{ij} \sim 0.25\delta_0(x_{ij}) + 0.75\{I(1 \leq i \leq 6)\text{Poi}(s_i g_j d_{1j}^*) + I(7 \leq i \leq 21)\text{Poi}(s_i g_j d_{2j}^*) + I(22 \leq i \leq 30)\text{Poi}(s_i g_j d_{3j}^*)\},$$

where the first six observations were drawn from the first distribution, the next 15 from the second and the last nine from the third distribution. We added 950 noisy features, which we generated from a zero-inflated Poisson model, $x_{ij} \sim 0.25\delta_0(x_{ij}) + 0.75\text{Poi}(s_i g_j)$. We simulated the s_i 's as independent and identically distributed $s_i \sim \text{U}(0.5, 1.5)$ and the g_j 's as $g_j \sim \text{Exp}(1/3)$. For the mixture components, we simulated $\tilde{d}_{kj}^*, k=1,2,3$ from a standard Normal distribution $\text{N}(0, \sigma^2)$ with $\sigma^2 = 1$.

As for hyperparameter settings, we used the following default settings. We set the hyperparameters that control the base distribution of the mixture DP prior to $a = b = 1$, which leads to a Gamma distribution with mean and variance equal to 1, and $a_\alpha = b_\alpha = 1$, setting the

expectation and variance of the concentration parameter α to 1. As for the Beta prior on the feature selection parameter ω , we set $a_\omega = 0.2, b_\omega = 1.8$, resulting in the proportion of features expected *a-priori* to discriminate the different groups to be $a_\omega/(a_\omega + b_\omega) = 10\%$. For the priors on $\tilde{s}_i$'s and $\tilde{g}_j$'s, we use the following default settings: $M = n/2 = 15, L = p/2 = 500, c_s = c_g = 0, \sigma_s = \sigma_g = 1, \tau_\eta = \tau_\mu = 1, a_m = b_m = 1, a_l = b_l = 1, a_t = b_t = 1$, and $a_q = b_q = 1$. Results we report below were obtained by running one MCMC chain with 10,000 iteration, discarding the first 1,000 as burn-in. We started the chain from a model with 2 randomly chosen γ_j 's set to 1 and with each observation assigned to a different cluster.

We first describe posterior inference on the relevant parameters as a result of our normalization and regularization approach. Figure 1 shows the heatmap of the pairwise posterior probabilities, $p(z_i = z_{i'} | \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X})$, of allocating observations i and i' to the same cluster, after burn-in. It is evident from the map that the inspection of the highest posterior allocation probabilities allows to reconstruct the true allocation structure quite well.

Figure 1: Simulation study: Heatmap of the pairwise posterior probabilities $p(z_i = z_{i'} | \cdot)$.

As for the feature selection, Figure 2 shows the marginal posterior probability of inclusion (PPI) of each feature $p(\gamma_j = 1 | \mathbf{z}, \tilde{\mathbf{s}}, \mathbf{g}_j, \mathbf{r}_{\cdot j}, \mathbf{x}_{\cdot j})$, after burn-in. The red dots indicate the truly discriminatory features. A threshold of 0.5 on the marginal probabilities results in a median model that includes 18 features, all of which are included in the set of discriminatory features used in the data generation process. Figure 3 shows the estimated marginal posterior distribution of the parameter d_{ij} , for two of the discriminating features. These were obtained by post-MCMC compositional sampling of the parameters as in [10]. Notice that the distributions are skewed, with high probability mass on low counts and small mass on extreme values, confirming that our modeling approach is able to capture the overdispersion characterizing the simulated data.

Figure 2: Simulation study: Marginal posterior probabilities of inclusion $p(\gamma_j = 1 | \cdot)$, with the red dots indicating the truly discriminatory features.

Figure 3: Simulation study: Estimated marginal posterior distribution of $E[d_{ij}]$ for two of the discriminatory features, obtained by post-MCMC compositional sampling. The distributions are skewed, with high probability mass on low counts and small mass on extreme values, capturing the overdispersion of the simulated data.

In order to further quantify the accuracy of our algorithm, and to compare its performances with other methods available in the literature, we looked at results under different scenarios in terms of sample size, n , number of noisy features, $p - p_\gamma$, and cluster dispersion, σ . In particular, we considered $\sigma^2 = 1, 0.5$, $n = 21, 99$ (equally distributed among the $K = 3$ mixture components) and $p = 200, 1000$ (with $p_\gamma = 50$ discriminating features). For each of these $2 \times 2 \times 2 = 8$ scenarios, we simulated 50 data sets and ran our MCMC algorithm with the default settings describe above and ran our MCMC algorithm with the default settings described above.

Table 2: Simulation study with data from a zero-inflated Poisson mixture: Adjusted Rand index (ARI) values, averaged over 50 replicates, achieved by the R packages edgeR, PoiClaClu and by our CoPoM method under different simulated scenarios. Standard deviations are indicated in parentheses.

Scenario			Competing methods		
n	p	σ^2	edgeR	PoiClaClu	CoPoM (PPM)
99	200	1.00	0.412 (0.129)	0.550 (0.162)	0.894 (0.129)
99	200	0.50	0.257 (0.100)	0.227 (0.095)	0.895 (0.116)
99	1000	1.00	0.307 (0.125)	0.179 (0.107)	0.925 (0.122)
99	1000	0.50	0.127 (0.103)	0.043 (0.029)	0.907 (0.106)
21	200	1.00	0.492 (0.194)	0.516 (0.210)	0.788 (0.120)
21	200	0.50	0.287 (0.155)	0.280 (0.169)	0.744 (0.128)
21	1000	1.00	0.240	0.134	0.716

			(0.131)	(0.107)	(0.081)
21	1000	0.50	0.135 (0.095)	0.101 (0.068)	0.670 (0.130)

For the analysis of the clustering results, we quantified performance via the adjusted Rand index [?, ARI,] Hubert1985, a variant of the Rand index [53], on the basis of the vector of point estimates $\hat{\mathbf{z}}_{\text{PPM}} = (\hat{z}_1^{\text{PPM}}, \dots, \hat{z}_n^{\text{PPM}})$. Let $A = \sum_{i>i'} I(z_i = z_{i'}) I(\hat{z}_i^{\text{PPM}} = \hat{z}_{i'}^{\text{PPM}})$ be the number of pairs of observations that belong to the same group in both the true clustering and the estimated one; $B = \sum_{i>i'} I(z_i = z_{i'}) I(\hat{z}_i^{\text{PPM}} \neq \hat{z}_{i'}^{\text{PPM}})$ be the number of pairs which belong to the same group in the true clustering but different groups in the estimated one; $C = \sum_{i>i'} I(z_i \neq z_{i'}) I(\hat{z}_i^{\text{PPM}} = \hat{z}_{i'}^{\text{PPM}})$ be the number of pairs in different groups in the true partition but assigned to the same group in the estimated one; $D = \sum_{i>i'} I(z_i \neq z_{i'}) I(\hat{z}_i^{\text{PPM}} \neq \hat{z}_{i'}^{\text{PPM}})$, the number of pairs assigned to different groups in both the truth and the estimate. Then the RI is defined as

$$RI = \frac{A + D}{A + B + C + D},$$

and the ARI as

$$ARI = \frac{\binom{n}{2} (A + D) - [(A + B)(A + C) + (C + D)(B + D)]}{\binom{n}{2} - [(A + B)(A + C) + (C + D)(B + D)]}.$$

The Rand index yields values between 0 and 1, while the ARI can yield negative values [54]. The larger the index, the more accurate the clustering result.

Table 3: Simulation study with data from a Poisson mixture: Adjusted Rand index (ARI) values, averaged over 50 replicates, achieved by the R packages edgeR, PoiClaClu and by our CoPoM method under different simulated scenarios. Standard deviations are indicated in parentheses.

Scenario			Competing methods		
n	p	σ^2	edgeR	PoiClaClu	CoPoM (PPM)
99	200	1.00	0.590 (0.133)	1.000 (0.000)	1.000 (0.000)
99	200	0.50	0.437	0.998	1.000

			(0.106)	(0.007)	(0.000)
99	1000	1.00	0.394 (0.111)	0.999 (0.006)	1.000 (0.000)
99	1000	0.50	0.232 (0.082)	0.947 (0.067)	0.991 (0.061)
21	200	1.00	0.701 (0.167)	1.000 (0.000)	1.000 (0.000)
21	200	0.50	0.556 (0.138)	0.983 (0.059)	0.997 (0.021)
21	1000	1.00	0.477 (0.143)	0.987 (0.058)	1.000 (0.000)

21	1000	0.50	0.312 (0.150)	0.863 (0.190)	1.000 (0.000)
----	------	------	------------------	------------------	------------------

For comparison, we selected two commonly employed estimation methods for count data, based on Poisson mixtures and implemented in the freely available R packages edgeR [13] and PoiClaClu [9]. Both edgeR and PoiClaClu incorporate plug-in estimates of scaling factors for normalization purposes: edgeR considers a negative binomial likelihood, whereas PoiClaClu aims at clustering count data by using a regularized Poisson log linear model. Unlike our modeling approach, which allows to directly estimate the cluster assignments, through the latent $\hat{z}$, those methods do not provide individual allocation estimates, but rather yield a dissimilarity matrix that can be transformed into a tree via hierarchical clustering. In order to make the comparison with our CoPoM model feasible, we considered those estimates that achieved the maximum ARI values. Table 1 reports results on clustering performances of all methods in terms of ARI values, averaged over the 50 replicates, under the different simulated scenarios. For our method, we report the results obtained by using the PPM estimates for cluster allocation. The MAP estimates performed similarly (not shown). In all replicates and scenarios considered, we generated data from a zero-inflated Poisson model, fixing $\pi = 0.25$. The percentage of observed zeros was around 44%. Results show that CoPoM consistently outperforms competing methods in terms of clustering accuracy. This is to be expected, since the competing methods do not incorporate the variable selection, and one might expect the inclusion of noisy features to mask the recovery of the true clustering structure. Table 3 shows the performances of our algorithm when data were generated from a Poisson mixture with three components (i.e. $\pi = 0$). Also in this setting, our zero-inflated Poisson model performs favorably or similarly with respect to the other methods in all cases. The widely used edgeR method shows the worst performances, especially for weakening signal strength, i.e. decreasing σ or increasing p , whereas our CoPoM method and PoiClaClu show the best performances.

Table 4: Simulation study: Average false positive rates (FPRs), true positive rates (TPRs), and areas under ROC curves (AUCs), achieved by the CoPoM model under different simulated scenarios, for different values of the threshold, τ , on the PPIs.

	CoPoM (PPI)
--	-------------

Scenario			$\tau = 0.2$		$\tau = 0.4$		$\tau = 0.6$		
n	p	σ^2	FPR	TPR	FPR	TPR	FPR	TPR	
99	200	1.00	0.075	0.888	0.019	0.827	0.005	0.790	
99	200	0.50	0.057	0.803	0.014	0.724	0.004	0.667	
99	1000	1.00	0.008	0.803	0.002	0.760	0.001	0.724	

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

99	1000	0.50	0.006	0.677	0.001	0.625	0.000	0.580	
21	200	1.00	0.354	0.872	0.110	0.700	0.033	0.576	
21	200	0.50	0.269	0.706	0.073	0.505	0.020	0.389	
21	1000	1.00	0.054	0.579	0.014	0.476	0.004	0.414	
21	1000	0.50	0.044	0.431	0.011	0.327	0.003	0.266	

Since our CoPoM model also allows the selection of a subset of discriminatory features, we quantified its performances in terms of averaged false positive rate (FPR) and true positive rate (TPR) achieved under the different simulated scenarios, for different values of the threshold on the PPIs. Results are reported in Table 4, and the corresponding receiver operating characteristic (ROC) curves are shown in Figure 4. Each sub-figure, corresponding to different values of n and p , shows that the estimate becomes more accurate with increasing separation between the clusters, captured by the between-cluster variability parameter σ .

Figure 4: Simulation study: Receiver operating characteristic (ROC) curves, for different values of the threshold on the PPIs, obtained by CoPoM under different simulated scenarios. The bold curve corresponds to $\sigma^2 = 1$, whereas the dashed curve corresponds to $\sigma^2 = 0.5$.

Unlike the edgeR and PoiClaClu methods, which use plug-in estimates, our model formulation is characterized by priors on the samples' and features' scaling parameters that impose soft constraints on the expected values. To appreciate the effect of such prior formulation, we looked at the frequentist coverages of our estimates and compared those to an implementation of our model that employs simple conjugate Gamma priors on the s_i 's and g_j 's parameters. For example, for the scenario with $n = 21, p = 200, \sigma^2 = 1$, the average frequentist coverages of the 95% credible intervals obtained with our model with constrained priors, calculated over 100 simulated datasets, were .91, .92, .92, for the s_i 's, g_j 's and $s_i g_j$'s parameters, respectively, while those obtained with unconstrained Gamma priors were 0, .001, .8, for the s_i 's, g_j 's and $s_i g_j$'s parameters, respectively.

We conclude this section by conducting a sensitivity analysis on the prior specification of our method and reporting results for different hyperparameter values of the priors for ω and α . In particular, we considered four scenarios: $(a_\omega = 0.02, b_\omega = 1.98)$, $(a_\omega = 1, b_\omega = 1)$, $(a_\alpha = 1, b_\alpha = 10)$, and $(a_\alpha = 1, b_\alpha = 0.1)$, while setting all the other hyperparameter to default values. Table 4

shows that clustering and feature selection performance, in terms of the adjusted Rand index and the AUC, are fairly robust to the choice of the hyperparameters. When no prior information is available, we suggest to use the default choice $a_\omega = 2p_{\text{prior}}/p$, $b_\omega = 2 - a_\omega$, and $a_\alpha = b_\alpha = 1$.

Table 5: Simulation study: Sensitivity analysis on the values of the hyperparameters of the priors on α and π

Scenario				Performance	
a_ω	b_ω	a_α	b_α	ARI	AUC
0.02	1.98	1.00	1.00	0.498	0.879
0.20	1.80	1.00	0.10	0.384	0.864
0.20	1.80	1.00	1.00	0.562	0.846
0.20	1.80	1.00	10.0	0.370	0.833
1.00	1.00	1.00	1.00	0.452	0.869

4.2 Application to the analysis of *bag-of-words* data

Bag-of-words data sets report the frequencies of occurrence of each word in a text document. Clustering of documents on the basis of a subset of relevant words is one of the most common tasks when analyzing *bag-of-words* data. Here we illustrate the performance of the CoPoM approach using the widely employed *Brown Corpus* [55]. In linguistics, a corpus defines a large and structured collection of texts, which are often used for conducting statistical analyses and tests of hypotheses about a linguistic variety or particular characteristics of a language. The Brown Corpus consists of 500 samples, distributed across 15 genres. Each sample begins at a random sentence-boundary in the article and continues up to the first sentence boundary after 2,000 words. The total vocabulary is about 50,000 words and half of them occur equal or less than once in the corpus. Thus, the data are typically characterized by an excessive number of zeros and overdispersion.

Figure 5: Brown Corpus data set: Topleft: trace plot of the number of clusters K ; Topright: histogram of the number of clusters K ; Bottomleft: trace plot of the number of selected variables p_γ ; Bottomright: marginal posterior probabilities of inclusion $p(\gamma_j = 1 | \cdot)$, with the red dots indicating the truly discriminatory features.

For the application of this paper, we selected a subset of the Brown Corpus, composed of five sports reportage, three society reportage, seven spot news, and seven editorials, whose length range from 2,200 to 2,374 words, for a total of 22 texts and 8,826 features. The data is quite sparse, as about 90% of the counts are zero.

Table 6: Brown Corpus data set: Point estimate $\hat{\mathbf{z}}_{\text{PPM}}$ of cluster membership obtained from the pairwise probability matrix of co-clustering. See Section 3.2 for details.

	Sports						Society				S
$\hat{\mathbf{z}}_{\text{PPM}}$	1	1	1	1	1	2	2	2	3	3	

We report results obtained by running the MCMC algorithm described in Section 3.1 with the same default settings used in the simulations. Figure 5 (top) shows the posterior inference on the number of clusters, more specifically the MCMC trace plot across iterations (top-left) and the resulting estimate of the posterior distribution of K , after burn-in (top-right). Figure 5 (bottom) reports the results of the feature selection. More specifically, Figure 5 (bottom-left) reports the trace plot for the total number of

included features, whereas Figure 5 (bottom-right) shows the estimated marginal posterior probabilities of inclusion (PPIs) of each single feature, $p(\gamma_j = 1 | \mathbf{z}, \tilde{\mathbf{s}}, \mathbf{g}_j, \mathbf{r}_j, \mathbf{x}_j)$, after burn-in. A threshold of 0.5 on the marginal probabilities results in a median model that includes 203 features (2.3% of the total). Finally, Figure 6 shows the heatmap of the pairwise posterior probabilities of co-clustering, $p(z_i = z_{i'} | \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X})$. Table 5 summarizes the posterior sample allocations $\hat{\mathbf{z}}$, based on the PPM estimate, and the corresponding ARI value. Our results suggest that the proposed CoPoM model is able to roughly recover distinctive features of the 4 broad categories.

Figure 6: Brown Corpus data set: Heatmap of the pairwise posterior probabilities $p(z_i = z_{i'} | \cdot)$.

For comparison, we already pointed out that the R packages edgeR and PoiClaClu do not provide a single point estimate, $\hat{\mathbf{z}}$, but yield a dissimilarity matrix that can be used as input in a hierarchical clustering algorithm. Within our approach, the squared Euclidean distance between each pair of observations, based on the selected subset of discriminatory features, can be defined as

$$d(\mathbf{x}_i^{(\gamma)}, \mathbf{x}_{i'}^{(\gamma)}) = \sqrt{\sum_{\{j: \gamma_j = 1\}} (d_{z_i j}^* - d_{z_{i'} j}^*)^2},$$

and an estimate of d_{kj}^* can be computed as

$$\hat{d}_{kj}^* = \frac{\sum_{\{i: \hat{z}_i = k, r_{ij} = 0\}} x_{ij}}{\hat{g}_j \sum_{\{i: \hat{z}_i = k, r_{ij} = 0\}} \hat{s}_i},$$

with $\hat{\mathbf{z}}_{\text{PPM}}$ and $\hat{\mathbf{s}}$ and $\hat{\mathbf{g}}$ the MAP estimates of the parameters. Note that the distance between each pair of observations within the same cluster is zero. Figure 7 shows the dendrogram of the hierarchical cluster analysis for the CoPoM model and the two

R packages edgeR and PoiClaClu [9], employing the Euclidean distance and the complete agglomeration method. Our CoPoM model and Witten's methods appear to be the most efficient to separate the four genres.

Figure 7: Bag-of-words data: The hierarchical clustering dendrograms using the Euclidean distance measure and the complete agglomeration method achieved by the CoPoM model, PoiClaClu and edgeR, respectively.

5 Discussion

In this paper, we have introduced a zero-inflated Constrained Poisson Mixture (CoPoM) model for the analysis of count data, where multiple features are observed as counts over a number of samples. We assume that the data are characterized by an excessive number of zeros and by overdispersion, that is a large number of features do not get frequently observed, whereas a few features are characterized by large occurrence rates in one or more samples. We assume that the amount of overdispersion may also vary from sample to sample. Our constrained Poisson mixture (CoPoM) approach proposes to regularize the estimation of the feature occurrence rates by means of a model-based normalization approach, which employs prior distributions with mean constraints on some of the model parameters. It also incorporates a feature selection mechanism, for the simultaneous identification of a parsimonious set of discriminatory features and group structures in the samples.

When applied to simulated data, the CoPoM model has shown improved accuracy of the clustering performance with respect to more standard approaches. We have also presented an application to a *bag-of-words* benchmark data set, where the data matrix is represented by the frequencies of word occurrences in multiple documents. We have shown the good performance of our Poisson mixture model with mean constraints in terms of feature selection and the clustering of the samples into larger topical groups.

Our modeling framework enables discrimination of samples based on a subset of selected features. Alternative feature selection approaches are often used to describe overlapping clustering of (latent) features across samples, [56, 57, 58], with

resulting post-hoc interpretation of the subset of non-exclusive pairings. Furthermore, several models have been proposed recently for the analysis of overdispersed count data as an alternative to Poisson mixtures. For example, in the Bayesian nonparametric literature, rounded Gaussian kernel models have been introduced as a more flexible and robust choice for analyzing both underdispersed and overdispersed count data [59, 60]. To date rounded Gaussian kernel methods have been developed mostly for density estimation. In our manuscript, instead, we are concerned with the estimation and comparison of features' occurrence rates across multiple samples, and feature selection, which require some degree of regularization due to the nature of the data generating mechanism. Future work will aim at exploring how our inferential aims can be comprised into the domain of those more recent approaches.

Although here we have focused on the analysis of text documents, our methodology is quite general. In particular, the class of prior distributions with mean constraints that we propose can be successfully employed on other types of high-dimensional count data sets, such as those encountered in ecology, genomics and spatial statistics, where the normalization of Poisson intensities is commonly employed to account for the overdispersion and heterogeneity observed across samples and across features. Future work will explore how the approaches presented here can be extended to the analysis of multivariate vectors of dependent count data, observed spatially or longitudinally in time.

References [url@samestyle](mailto:author@samestyle)

- [1] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Sciences*, vol. 101, no. Suppl 1, pp. 5228–5235, 2004.
- [2] E. M. Airoldi, A. G. Anderson, S. E. Fienberg, and K. K. Skinner, "Who wrote Ronald Reagan's radio addresses?" *Bayesian Analysis*, vol. 1, no. 2, pp. 289–319, 2006.

- [3] E. M. Airoidi, E. A. Erosheva, S. E. Fienberg, C. Joutard, T. Love, and S. Shringarpure, "Reconceptualizing the classification of PNAS articles," *Proceedings of the National Academy of Sciences*, vol. 107, no. 49, pp. 20 899–20 904, 2010.
- [4] E. M. Airoidi and J. M. Bischof, "A regularization scheme on word occurrence rates that improves estimation and interpretation of topical content," *Journal of the American Statistical Association*, no. To appear, 2016.
- [5] J. M. Gee, R. M. Warwick, M. Schaanning, J. A. Berge, and W. G. Ambrose Jr, "Effects of organic enrichment on meiofaunal abundance and community structure in sublittoral soft sediments," *Journal of Experimental Marine Biology and Ecology*, vol. 91, no. 3, pp. 247–262, 1985.
- [6] S. Blackshaw, S. Harpavat, J. Trimarchi, L. Cai, H. Huang, W. P. Kuo, G. Weber, K. Lee, R. E. Fraioli, and S.-H. Cho, "Genomic analysis of mouse retinal development," *PLoS Biology*, vol. 2, no. 9, p. e247, 2004.
- [7] J. C. Marioni, C. E. Mason, S. M. Mane, M. Stephens, and Y. Gilad, "RNA-Seq: an assessment of technical reproducibility and comparison with gene expression arrays," *Genome Research*, vol. 18, no. 9, pp. 1509–1517, 2008.
- [8] J. Chen and H. Li, "Variable selection for sparse dirichlet-multinomial regression with an application to microbiome data analysis," *The Annals of Applied Statistics*, vol. 7, no. 1, 2013.
- [9] D. M. Witten, "Classification and clustering of sequencing data using a Poisson model," *The Annals of Applied Statistics*,

vol. 5, no. 4, pp. 2493–2518, 2011.

[10] M. Guindani, N. Sepúlveda, C. D. Paulino, and P. Müller, “A Bayesian semiparametric approach for the differential analysis of sequence counts data,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 63, no. 3, pp. 385–404, 2014.

[11] E. M. Airoldi, D. Blei, E. A. Erosheva, and S. E. Fienberg, *Handbook of Mixed Membership Models and Their Applications*. 1em plus 0.5em minus 0.4em CRC Press, 2014.

[12] M. D. Robinson and A. Oshlack, “A scaling normalization method for differential expression analysis of RNA-Seq data,” *Genome Biology*, vol. 11, no. 3, p. R25, 2010.

[13] M. D. Robinson, D. J. McCarthy, and G. K. Smyth, “edgeR: a Bioconductor package for differential expression analysis of digital gene expression data,” *Bioinformatics*, vol. 26, no. 1, pp. 139–140, 2010.

[14] S. Anders and W. Huber, “Differential expression analysis for sequence count data,” *Genome Biology*, vol. 11, no. 10, p. R106, 2010.

[15] D. Lambert, “Zero-inflated poisson regression, with an application to defects in manufacturing,” *Technometrics*, vol. 34, no. 1–14, 1992.

- [16] A. C. Cameron and P. K. Trivedi, *Microeconometrics: methods and applications*. 1em plus 0.5em minus 0.4em Cambridge University Press, 2005.
- [17] —, *Regression Analysis of Count Data*. 1em plus 0.5em minus 0.4em Cambridge University Press, 2013, vol. 53.
- [18] J. H. Bullard, E. Purdom, K. D. Hansen, and S. Dudoit, “Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments,” *BMC Bioinformatics*, vol. 11, no. 1, p. 94, 2010.
- [19] A. Mortazavi, B. A. Williams, K. McCue, L. Schaeffer, and B. Wold, “Mapping and quantifying mammalian transcriptomes by RNA-Seq,” *Nature Methods*, vol. 5, no. 7, pp. 621–628, 2008.
- [20] C. N. Morris, “Parametric empirical Bayes inference: theory and applications,” *Journal of the American Statistical Association*, vol. 78, no. 381, pp. 47–55, 1983.
- [21] A. Gelman, “Objections to Bayesian statistics,” *Bayesian Analysis*, vol. 3, no. 3, pp. 445–449, 2008.
- [22] P. Gustafson, “On model expansion, model contraction, identifiability, and prior information: two illustrative scenarios involving mismeasured variables (with discussion),” *Statistical Science*, vol. 20, pp. 111–140, 2005.

- [23] B. J. Reich, H. D. Bondell, and H. J. Wang, "Flexible Bayesian quantile regression for independent and clustered data," *Biostatistics*, vol. 11, no. 2, pp. 337–352, 2010.
- [24] M. E. Roberts, B. M. Stewart, and E. M. Airolidi, "A model of text for experimentation in the social sciences," *Journal of the American Statistical Association*, vol. 111, pp. 988–1003, 2016.
- [25] M. G. Tadesse, N. Sha, and M. Vannucci, "Bayesian variable selection in clustering high-dimensional data," *Journal of the American Statistical Association*, vol. 100, no. 470, pp. 602–617, 2005.
- [26] A. E. Raftery and N. Dean, "Variable selection for model-based clustering," *Journal of the American Statistical Association*, vol. 101, no. 473, pp. 168–178, 2006.
- [27] P. Müller, F. A. Quintana, A. Jara, and T. Hanson, *Bayesian Nonparametric Data Analysis*. 1em plus 0.5em minus 0.4em Springer, 2015.
- [28] S. Richardson and P. J. Green, "On Bayesian analysis of mixtures with an unknown number of components (with discussion)," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 59, no. 4, pp. 731–792, 1997.
- [29] T. L. Griffiths and Z. Ghahramani, "The Indian Buffet process: An introduction and review," *The Journal of Machine Learning Research*, vol. 12, pp. 1185–1224, 2011.

- [30] S. Williamson, C. Wang, K. A. Heller, and D. M. Blei, "The IBP compound Dirichlet process and its application to focused topic modeling," in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 1151–1158.
- [31] M. Zhou, "Beta-Negative Binomial Process and Exchangeable? Random Partitions for Mixed-Membership Modeling," in *Advances in Neural Information Processing Systems*, 2014, pp. 3455–3463.
- [32] P. Brown, M. Vannucci, and T. Fearn, "Bayesian wavelength selection in multicomponent analysis," *Journal of Chemometrics*, vol. 12(3), pp. 173–182, 1998.
- [33] S. Banerjee, B. P. Carlin, and A. E. Gelfand, *Hierarchical Modeling and Analysis for Spatial Data*. 1em plus 0.5em minus 0.4em CRC Press, 2014.
- [34] P. Hoff, "Nonparametric estimation of convex models via mixture," *Annals of Statistics*, vol. 31, pp. 174–200, 2003.
- [35] L. Trippa and G. Parmigiani, "False discovery rates in somatic mutation studies of cancer," *The Annals of Applied Statistics*, pp. 1360–1378, 2011.
- [36] M. Kyung, J. Gill, and G. Casella, "Sampling schemes for generalized linear Dirichlet process random effects models," *Statistical Methods & Applications*, vol. 20, no. 3, pp. 259–290, 2011.

- [37] M. A. Taddy and A. Kottas, "Mixture modeling for Marked Poisson processes," *Bayesian Analysis*, vol. 7, no. 2, pp. 335–362, 2012.
- [38] J. Sethuraman, "A constructive definition of Dirichlet priors," *Statistica Sinica*, vol. 4, pp. 639–650, 1994.
- [39] N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker, *Bayesian Nonparametrics*. 1em plus 0.5em minus 0.4em Cambridge University Press, 2010, vol. 28.
- [40] W. J. Ewens, "The sampling theory of selectively neutral alleles," *Theoretical Population Biology*, vol. 3, no. 1, pp. 87–112, 1972.
- [41] C. E. Antoniak, "Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems," *The Annals of Statistics*, vol. 2, no. 6, pp. 1152–1174, 1974.
- [42] H. Crane *et al.*, "The ubiquitous ewens sampling formula," *Statistical Science*, vol. 31, no. 1, pp. 1–19, 2016.
- [43] M. D. Escobar and M. West, "Bayesian density estimation and inference using mixtures," *Journal of the American Statistical Association*, vol. 90, no. 430, pp. 577–588, 1995.

[44] E. I. George and R. E. McCulloch, "Approaches for Bayesian variable selection," *Statistica Sinica*, vol. 7, no. 2, pp. 339–373, 1997.

[45] P. J. Brown, M. Vannucci, and T. Fearn, "Multivariate Bayesian variable selection and prediction," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 60, no. 3, pp. 627–641, 1998.

[46] R. M. Neal, "Markov chain sampling methods for Dirichlet process mixture models," *Journal of Computational and Graphical Statistics*, vol. 9, no. 2, pp. 249–265, 2000.

[47] S. Kim, M. G. Tadesse, and M. Vannucci, "Variable selection in clustering via Dirichlet process mixture models," *Biometrika*, vol. 93, no. 4, pp. 877–893, 2006.

[48] H. Ishwaran and L. F. James, "Gibbs sampling methods for stick-breaking priors," *Journal of the American Statistical Association*, vol. 96, no. 453, pp. 161–173, 2001.

[49] D. B. Dahl, "Model-based clustering for expression data via a Dirichlet process mixture model," *Bayesian Inference for Gene Expression and Proteomics*, pp. 201–218, 2006.

[50] P. Müller, G. Parmigiani, C. Robert, and J. Rousseau, "Optimal sample size for multiple testing: the case of gene expression microarrays," *Journal of the American Statistical Association*, vol. 99, no. 468, pp. 990–1001, 2004.

- [51] M. Guindani, P. Müller, and S. Zhang, "A Bayesian discovery procedure," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 71, no. 5, pp. 905–925, 2009.
- [52] L. Hubert and P. Arabie, "Comparing partitions," *Journal of Classification*, vol. 2, no. 1, pp. 193–218, 1985.
- [53] W. M. Rand, "Objective criteria for the evaluation of clustering methods," *Journal of the American Statistical Association*, vol. 66, no. 336, pp. 846–850, 1971.
- [54] J. M. Santos and M. Embrechts, "On the use of the adjusted Rand index as a metric for evaluating supervised classification," in *Artificial Neural Networks–ICANN 2009*. 1em plus 0.5em minus 0.4em Springer, 2009, pp. 175–184.
- [55] H. Kucera and N. Francis, *Computational Analysis of Present-day American English*. 1em plus 0.5em minus 0.4em Brown University Press, 1967.
- [56] T. Broderick, J. Pitman, and M. I. Jordan, "Feature allocations, probability functions, and paintboxes," *Bayesian Analysis*, vol. 8, no. 4, pp. 801–836, 12 2013.
- [57] J. Lee, P. Müller, K. Gulukota, and Y. Ji, "A bayesian feature allocation model for tumor heterogeneity," *Ann. Appl. Stat.*, vol. 9, no. 2, pp. 621–639, 06 2015.

- [58] J. Lee, P. Müller, P., S. Sengupta, K. Gulukota, and Y. Ji, “Bayesian inference for intratumour heterogeneity in mutations and copy number variation,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 65, no. 4, pp. 547–563, 2016.
- [59] A. Canale and D. B. Dunson, “Bayesian kernel mixtures for counts,” *Journal of the American Statistical Association*, vol. 106, no. 496, pp. 1528–1539, 2011.
- [60] A. Canale and I. Prünster, “Robustifying Bayesian nonparametric mixtures for count data,” *Biometrics*, vol. 73(1), pp. 174–184, 2017.

Appendix: Details of the MCMC algorithm

We start by writing the marginal likelihood for each sample i , $i = 1, \dots, n$, after integrating out the parameters $\mathbf{d}_{kj}^*$.

$$\begin{aligned}
 & f(\mathbf{x}_i | \gamma, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_i) \\
 &= \int f(\mathbf{x}_i | z_i = k, \gamma, \mathbf{d}_k^*, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_i) p(\mathbf{d}_k^*) d\mathbf{d}_k^* \\
 &= \int \prod_{\{j: \gamma_j=1, r_{ij}=0\}} \text{Poi}(x_{ij}; s_i g_j d_{kj}^*) \prod_{\{j: \gamma_j=0, r_{ij}=0\}} \text{Poi}(x_{ij}; s_i g_j) \prod_{\{j: \gamma_j=1, r_{ij}=0\}} \text{Ga}(d_{kj}^*; a, b) d\mathbf{d}_{kj}^* \\
 &= \prod_{\{j: r_{ij}=0\}} \frac{(s_i g_j)^{x_{ij}}}{x_{ij}!} \exp \left\{ -s_i \sum_{\{j: \gamma_j=0, r_{ij}=0\}} g_j \right\} \prod_{\{j: \gamma_j=1, r_{ij}=0\}} \frac{b^a}{\Gamma(a)} \frac{\Gamma(a + x_{ij})}{(b + s_i g_j)^{a+x_{ij}}}.
 \end{aligned}$$

where $s_i = \exp\{\tilde{s}_i\}$, $g_j = \exp\{\tilde{g}_j\}$, and $d_{kj}^* = \exp\{\tilde{d}_{kj}^*\}$

In terms of each feature j , $j = 1, \dots, p$, we write

$$f(\mathbf{x}_j | \mathbf{z}, \gamma_j = 0, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{r}_j) = \prod_{\{i: r_{ij}=0\}} \frac{(s_i g_j)^{x_{ij}}}{x_{ij}!} \exp \left\{ -g_j \sum_{\{i: r_{ij}=0\}} s_i \right\}$$

and

$$f(\mathbf{x}_j | \mathbf{z}, \gamma_j = 1, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{r}_j) = \prod_{\{i: r_{ij}=0\}} \frac{(s_i g_j)^{x_{ij}}}{x_{ij}!} \left(\frac{b^a}{\Gamma(a)} \right)^K$$

$$\times \prod_{k=1}^K \frac{\Gamma(a + \sum_{\{i: z_i=k, r_{ij}=0\}} x_{ij})}{(b + g_j \sum_{\{i: z_i=k, r_{ij}=0\}} s_i)^{a + \sum_{\{i: z_i=k, r_{ij}=0\}} x_{ij}}}.$$

At each MCMC iteration, we perform the following steps:

Update of the set of discriminatory features γ : We randomly perform an *add-delete-swap* step. We repeat this step 20 times to ensure validation of the feature selection for each given cluster assignment. The general Hasting ratio can be written as

$$\begin{aligned} r &= \frac{p(\gamma^* | \mathbf{z}, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}_j, \mathbf{R}, \omega, \mathbf{X})}{p(\gamma^{(b-1)} | \mathbf{z}, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}_j, \mathbf{R}, \omega, \mathbf{X})} \frac{J(\gamma^{(b-1)} | \gamma^*)}{J(\gamma^* | \gamma^{(b-1)})} \\ &= \frac{f(\mathbf{X} | \mathbf{z}, \gamma^*, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}_j, \mathbf{R})}{f(\mathbf{X} | \mathbf{z}, \gamma^{(b-1)}, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}_j, \mathbf{R})} \frac{p(\gamma^* | \omega)}{p(\gamma^{(b-1)} | \omega)} \frac{J(\gamma^{(b-1)} | \gamma^*)}{J(\gamma^* | \gamma^{(b-1)})}. \end{aligned}$$

For both add/delete and swap steps, the proposal density ratio equals to 1.

More specifically, for the add-delete step, we randomly sample a feature $j, j = 1, \dots, p$ and propose to change the value of γ_j with probability $\min(1, m_{\text{MH}})$, where

$$\begin{aligned}
m_{MH}^{\text{add}} &= \frac{\left(\frac{b^a}{\Gamma(a)}\right)^K \prod_{k=1}^K \frac{\Gamma(a + \sum_{i:z_i=k, r_{ij}=0} x_{ij})}{(b + g_j \sum_{i:z_i=k, r_{ij}=0} s_i)^{a + \sum_{i:z_i=k, r_{ij}=0} x_{ij}}}{\exp\left(-g_j \sum_{\{i:r_{ij}=0\}} s_i\right)} \frac{\omega}{1-\omega}, \\
m_{MH}^{\text{del}} &= \frac{\exp\left(-g_j \sum_{\{i:r_{ij}=0\}} s_i\right)}{\left(\frac{b^a}{\Gamma(a)}\right)^K \prod_{k=1}^K \frac{\Gamma(a + \sum_{i:z_i=k, r_{ij}=0} x_{ij})}{(b + g_j \sum_{i:z_i=k, r_{ij}=0} s_i)^{a + \sum_{i:z_i=k, r_{ij}=0} x_{ij}}} \frac{1-\omega}{\omega}.
\end{aligned}$$

For the swap step, we randomly swap two features if applicable, i.e. change the value of a feature γ_{j_1} from 1 to 0, and, correspondingly, of another currently non-discriminatory feature j_2 , i.e. change the value of γ_{j_2} from 0 to 1, with probability $\min(1, r_{MH}^{\text{swp}})$, where

$$\begin{aligned}
m_{MH}^{swp} = & \frac{\exp\left(-g_{j_1} \sum_{\{i:r_{ij}=0\}} s_i\right)}{\Gamma(a + \sum_{\{i:z_i=k, r_{ij}=0\}} x_{ij_1})} \\
& \prod_{k=1}^K \frac{1}{(b + g_{j_1} \sum_{\{i:z_i=k, r_{ij}=0\}} s_i)^{a + \sum_{\{i:z_i=k, r_{ij}=0\}} x_{ij_1}}} \\
& \prod_{k=1}^K \frac{\Gamma(a + \sum_{\{i:z_i=k, r_{ij}=0\}} x_{ij_2})}{(b + g_{j_2} \sum_{\{i:z_i=k, r_{ij}=0\}} s_i)^{a + \sum_{\{i:z_i=k, r_{ij}=0\}} x_{ij_2}}} \\
& \times \frac{1}{\exp\left(-g_{j_2} \sum_{\{i:r_{ij}=0\}} s_i\right)}.
\end{aligned}$$

Finally, we update the hyperparameter ω for inclusion probability,

$$\omega | \gamma, a_\omega, b_\omega \sim \text{Be}(a_\omega + p_\gamma, b_\omega + n - p_\gamma).$$

Update of the cluster allocation for the selected features $\mathbf{z}$: At each iteration, we perform a Gibbs sampling to update z_i sequentially from observation 1 to n according to Algorithm 3 in [46]. In order to do this, we first need to integrate out the parameter $\mathbf{d}_{k\cdot}^*$ based on the its prior and all observations except itself, i.e. all of $\mathbf{x}_{i'}$ for which $i' \neq i$.

$$\begin{aligned}
& f(\mathbf{x}_{i\cdot} | z_i = k, \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}_{-i\cdot}) \\
& = \int f(\mathbf{x}_{i\cdot} | z_i = k, \gamma, \mathbf{d}_{k\cdot}^*, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_{i\cdot}) p(\mathbf{d}_{k\cdot}^* | \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}_{-i}, \tilde{\mathbf{g}}, \mathbf{R}_{-i\cdot}, \mathbf{X}_{-i\cdot}) d\mathbf{d}_{k\cdot}^*.
\end{aligned}$$

$$\begin{aligned}
&= \int \prod_{\{j:\gamma_j=1, r_{ij}=0\}} \text{Poi}(x_{ij}; s_i g_j d_{kj}^*) \prod_{\{j:\gamma_j=0, r_{ij}=0\}} \text{Poi}(x_{ij}; s_i g_j) \\
&\quad \prod_{\{j:\gamma_j=1, r_{ij}=0\}} \text{Ga}\left(d_{kj}^*; a + \sum_{\{i': z_{i'}=k, r_{i'j}=0, i' \neq i\}} x_{i'j}, b + g_j \sum_{\{i': z_{i'}=k, r_{i'j}=0, i' \neq i\}} s_{i'}\right) dd_{kj}^* \\
&= \prod_{\{j:r_{ij}=0\}} \frac{(s_i g_j)^{x_{ij}}}{x_{ij}!} \exp\left\{-s_i \sum_{\{j:\gamma_j=0, r_{ij}=0\}} g_j\right\} \\
&\quad \prod_{\{j:\gamma_j=1, r_{ij}=0\}} \frac{\left(b + g_j \sum_{\{i': z_{i'}=k, r_{i'j}=0, i' \neq i\}} s_{i'}\right)^{a+} \sum_{\{i': z_{i'}=k, r_{i'j}=0, i' \neq i\}} x_{i'j}}{\Gamma(a + \sum_{\{i': z_{i'}=k, r_{i'j}=0, i' \neq i\}} x_{i'j})} \\
&\quad \prod_{\{j:\gamma_j=1, r_{ij}=0\}} \frac{\Gamma\left(a + \sum_{\{i': z_{i'}=k, r_{i'j}=0\}} x_{i'j}\right)}{\left(b + g_j \sum_{\{i': z_{i'}=k, r_{i'j}=0\}} s_{i'}\right)^{a+} \sum_{\{i': z_{i'}=k, r_{i'j}=0\}} x_{i'j}}.
\end{aligned}$$

Then, we propose to either:

- form a new cluster (change the value of z_i to $K_{-i} + 1$) with probability $p_{K_{-i}+1}$, where

$$p_{K_{-i}+1} = \frac{\frac{\alpha}{n-1+\alpha} f(\mathbf{x}_i | \gamma, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_i)}{\frac{\alpha}{n-1+\alpha} f(\mathbf{x}_i | \gamma, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_i) + \sum_{k=1}^K \frac{n_{k,-i}}{n-1+\alpha} f(\mathbf{x}_i | z_i = k, \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}_{-i})}.$$

- allocate the observation to any of the existing clusters, say cluster $k \in \{1, \dots, K_{-i} + 1\}$, with probability p_k , where

$$p_k = \frac{\frac{n_{k,-i}}{n-1+\alpha} f(\mathbf{x}_i | z_i = k, \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}_{-i})}{\frac{\alpha}{n-1+\alpha} f(\mathbf{x}_i | \gamma, \tilde{s}_i, \tilde{\mathbf{g}}, \mathbf{r}_i) + \sum_{k=1}^K \frac{n_{k,-i}}{n-1+\alpha} f(\mathbf{x}_i | z_i = k, \mathbf{z}_{-i}, \gamma, \tilde{\mathbf{s}}, \tilde{\mathbf{g}}, \mathbf{R}, \mathbf{X}_{-i})}.$$

Finally, we update the concentration parameter α , by following the algorithm in [43], that is we generate an auxiliary variable $\eta_\alpha | \alpha \sim \text{Be}(\alpha+1, n)$ and then we sample α from a mixture of two gamma densities,

$$\alpha | \eta_\alpha, \mathbf{z} \sim \pi_\eta \text{Ga}(a_\alpha + K, b - \log(\eta_\alpha)) + (1 - \pi_\eta) \text{Ga}(a_\alpha + K - 1, b_\alpha - \log(\eta_\alpha)),$$

with the weights π_η defined by $\pi_\eta / (1 - \pi_\eta) = (a_\alpha + K - 1) / (n(b_\alpha - \log(\eta_\alpha)))$.

Update of the scaling factors: We can rewrite the prior distribution in (3)–(??) by introducing latent auxiliary variables, that specify how the $\tilde{s}_i$ and $\tilde{g}_j$ are assigned to any of the inner and outer mixture components. More specifically, we can introduce a $n \times 1$ vector of assignment indicators, ν , with $\nu_i = m$ indicating that $\tilde{s}_i$ is a sample from $f_m^s(\tilde{s}_i | t_m, \eta_m)$. The weights ϕ_m^s determine the probability of each value $\nu_i = m$, with $m = 1, 2, \dots$. Correspondingly, we can consider a $n \times 1$ vector ε of binary elements ε_i , where $\varepsilon_i = 1$ indicates that, given $\nu_i = m$, $\tilde{s}_i$ is drawn from a $N(\eta_m, \sigma_s^2)$ with probability t_m , and $\varepsilon_i = 0$ indicates that $\tilde{s}_i = 0$ is drawn from the right component of $f_m^s(\tilde{s}_i | t_m, \eta_m)$, i.e. $N\left(\frac{c_s - t_m \eta_m}{1 - t_m}, \sigma_s^2\right)$, with probability $1 - t_m$. Similarly,

we can introduce a $p \times 1$ vector ξ , with $\xi_j = l$ indicating that $\tilde{g}_j$ is sampled from $f_l^g(\tilde{g}_j | q_l, \mu_l)$, $l = 1, 2, \dots$, and $\phi_l^g = p(\xi = l)$. Correspondingly, given the assignments obtained in the vector ξ , we can define a $p \times 1$ vector ψ of binary elements ψ_j , where $\psi_j = 1$ indicates that, given $\xi_j = l$, then $\tilde{g}_j$ is drawn from $N(\mu_l, \sigma_g^2)$, whereas $\psi_j = 0$ indicates that $\tilde{g}_j$ is from $N\left(\frac{c_g - q_l \mu_l}{1 - q_l}, \sigma_g^2\right)$. Thus, the prior model (3)–(??) can be rewritten as

$$\tilde{s}_i | v_i, \varepsilon_i, \mathbf{t}, \eta \sim N\left(\varepsilon_i \eta_{v_i} + (1 - \varepsilon_i) \frac{c_s - t_{v_i} \eta_{v_i}}{1 - t_{v_i}}, \sigma_s^2\right) \quad (8)$$

and

$$\tilde{g}_j | \xi_j, \psi_j, \mathbf{q}, \mu \sim N\left(\psi_j \mu_{\xi_j} + (1 - \psi_j) \frac{c_g - q_{\xi_j} \mu_{\xi_j}}{1 - q_{\xi_j}}, \sigma_g^2\right), \quad (9)$$

where $\mathbf{t}$, η , $\mathbf{q}$, and μ denote the vectors of t_m , η_m , q_l , and μ_l , respectively. Therefore, the update of the sample and feature specific scaling factors s_i and g_j can proceed as follows, after logarithmic transformation:

a) Update of the $\tilde{s}_i$'s: We perform Metropolis sampling to update $\tilde{s}_i$, where $\tilde{s}_i = \log\{s_i\}$, sequentially from observation 1 to n . We propose a new $\tilde{s}_i^*$ from $N(\tilde{s}_i^{(b-1)}, \tau_s^2)$ and accept it with probability $\min(1, m_{\text{MH}})$, where

$$\begin{aligned} m_{\text{MH}} &= \frac{p(\tilde{s}_i^* | \gamma, \tilde{\mathbf{g}}, \mathbf{r}_i, v_i, \varepsilon_i, \mathbf{t}, \eta, \mathbf{x}_i)}{p(\tilde{s}_i^{(b-1)} | \gamma, \tilde{\mathbf{g}}, \mathbf{r}_i, v_i, \varepsilon_i, \mathbf{t}, \eta, \mathbf{x}_i)} \frac{J(\tilde{s}_i^{(b-1)} | \tilde{s}_i^*)}{J(\tilde{s}_i^* | \tilde{s}_i^{(b-1)})} \\ &= \frac{f(\mathbf{x}_i | \gamma, \tilde{s}_i^*, \tilde{\mathbf{g}}, \mathbf{r}_i)}{f(\mathbf{x}_i | \gamma, \tilde{s}_i^{(b-1)}, \tilde{\mathbf{g}}, \mathbf{r}_i)} \frac{p(\tilde{s}_i^* | v_i, \varepsilon_i, \mathbf{t}, \eta)}{p(\tilde{s}_i^{(b-1)} | v_i, \varepsilon_i, \mathbf{t}, \eta)} \end{aligned}$$

$$\begin{aligned}
&= \frac{s_i^* \sum_{\{j:r_{ij}=0\}} x_{ij} \exp\left(-s_i^* \sum_{\{j:\gamma_j=0, r_{ij}=0\}} g_j\right)}{s_i^{(b-1)} \sum_{\{j:r_{ij}=0\}} x_{ij} \exp\left(-s_i^{(b-1)} \sum_{\{j:\gamma_j=0, r_{ij}=0\}} g_j\right)} \\
&\times \frac{\prod_{j:\gamma_j=1, r_{ij}=0} (b + s_i^* g_j)^{-a-x_{ij}}}{\prod_{j:\gamma_j=1, r_{ij}=0} (b + s_i^{(b-1)} g_j)^{-a-x_{ij}}} \\
&\times \frac{\mathcal{N}\left(\tilde{s}_i^*; \varepsilon_i \eta_{v_i} + (1-\varepsilon_i) \frac{c_s - t_{v_i} \eta_{v_i}}{1-t_{v_i}}, \sigma_s^2\right)}{\mathcal{N}\left(\tilde{s}_i^{(b-1)}; \varepsilon_i \eta_{v_i} + (1-\varepsilon_i) \frac{c_s - t_{v_i} \eta_{v_i}}{1-t_{v_i}}, \sigma_s^2\right)}.
\end{aligned}$$

Since v_i , ε_i , $\mathbf{t}$, and $\boldsymbol{\eta}$ have conjugate full conditionals, we use Gibbs sampling to update them one after another

- Gibbs sampling for updating $v_i, i = 1, \dots, n$:

$$\pi(v_i = l | v_{-i}, \varepsilon_i, \mathbf{t}, \boldsymbol{\eta}, \tilde{s}_i) \propto \phi_m^s \mathcal{N}\left(\tilde{s}_i; \varepsilon_i \eta_m + (1-\varepsilon_i) \frac{c_s - t_m \eta_m}{1-t_m}, \sigma_s^2\right).$$

- Gibbs sampling for updating $\varepsilon_i, i = 1, \dots, n$:

$$\pi(\varepsilon_i | v_i = m, \varepsilon_{-i}, \mathbf{t}, \eta, \tilde{s}_i) \propto \begin{cases} (1-t_m) \mathcal{N}\left(\tilde{s}_i; \frac{c_s - t_m \eta_m}{1-t_m}, \sigma_s^2\right) & \text{if } \varepsilon_i = 0 \\ t_m \mathcal{N}(\tilde{s}_i; \eta_m, \sigma_s^2) & \text{if } \varepsilon_i = 1 \end{cases}.$$

- Gibbs sampling for updating $t_m, m = 1, \dots, M$:

$$t_m | v, \varepsilon \sim \text{Be}\left(a_t + \sum_{i=1}^n I(v_i = m) I(\varepsilon_i = 1), b_t + \sum_{i=1}^n I(v_i = m) I(\varepsilon_i = 0)\right).$$

- Gibbs sampling for updating $\eta_m, m = 1, \dots, M$:

$$\eta_m | v, \varepsilon, \mathbf{t}, \tilde{\mathbf{s}} \sim \mathcal{N}\left(\frac{c_m / \sigma_s^2}{e_m / \sigma_s^2 + 1 / \tau_\eta^2}, \frac{1}{e_m / \sigma_s^2 + 1 / \tau_\eta^2}\right),$$

where $c_m = \sum_{\{i: v_i = m, \varepsilon_i = 1\}} \tilde{s}_i - \frac{t_m}{1-t_m} \sum_{\{i: v_i = m, \varepsilon_i = 0\}} \left(\tilde{s}_i - \frac{c_s}{1-t_m}\right)$ and $e_m = \sum_{i=1}^n I(v_i = m) I(\varepsilon_i = 1) + \sum_{\{i: v_i = m, \varepsilon_i = 0\}} \left(\frac{t_m}{1-t_m}\right)^2$.

- Gibbs Sampling for updating $\phi_m^s, m = 1, \dots, M$ by stick-breaking process [48]:

$$\begin{aligned} \phi_1^s &= v_1, \\ \phi_2^s &= (1-v_1)v_2, \\ &\vdots \\ \phi_M^s &= (1-v_1) \cdots (1-v_{M-1})v_M, \end{aligned}$$

where $v_m | v \sim \text{Be}\left(a_m + \sum_{i=1}^n I(v_i = m), b_m + \sum_{i=1}^n I(v_i > m)\right).$

b) Update of the g_j 's: We perform Metropolis sampling to update $\tilde{g}_j$, where $\tilde{g}_j = \log\{g_j\}$, sequentially from feature

1 to p . We propose a new $\tilde{g}_j^*$ from $N(\tilde{g}_j^{(b-1)}, \tau_g^2)$ and accept it with probability $\min(1, m_{MH})$, where

$$\begin{aligned} m_{MH} &= \frac{p(\tilde{g}_j^* | \mathbf{z}, \gamma_j, \tilde{\mathbf{s}}, \mathbf{r}_j, \xi_j, \psi_j, \mathbf{q}, \mu, \mathbf{x}_{\cdot j})}{p(\tilde{g}_j^{(b-1)} | \mathbf{z}, \gamma_j, \tilde{\mathbf{s}}, \mathbf{r}_j, \xi_j, \psi_j, \mathbf{q}, \mu, \mathbf{x}_{\cdot j})} \frac{J(\tilde{g}_j^{(b-1)} | \tilde{g}_j^*)}{J(\tilde{g}_j^* | \tilde{g}_j^{(b-1)})} \\ &= \frac{f(\mathbf{x}_{\cdot j} | \mathbf{z}, \gamma_j, \tilde{\mathbf{s}}, \tilde{g}_j^*, \mathbf{r}_j)}{f(\mathbf{x}_{\cdot j} | \mathbf{z}, \gamma_j, \tilde{\mathbf{s}}, \tilde{g}_j^{(b-1)}, \mathbf{r}_j)} \frac{p(\tilde{g}_j^* | \xi_j, \psi_j, \mathbf{q}, \mu)}{p(\tilde{g}_j^{(b-1)} | \xi_j, \psi_j, \mathbf{q}, \mu)} \end{aligned}$$

$$= \left\{ \begin{array}{ll} \frac{g_j^* \sum_{\{i:r_{ij}=0\}} x_{ij} \exp \left\{ -g_j^* \sum_{\{i:r_{ij}=0\}} s_i \right\}}{g_j^{(b-1)} \sum_{\{i:r_{ij}=0\}} x_{ij} \exp \left\{ -g_j^{(b-1)} \sum_{\{i:r_{ij}=0\}} s_i \right\}} \times \frac{N \left(\tilde{g}_j^*; \psi_j \mu_{\xi_j} + (1-\psi_j) \frac{c_g - q_{\xi_j} \mu_{\xi_j}}{1-q_{\xi_j}}, \sigma_g^2 \right)}{N \left(\tilde{g}_j^{(b-1)}; \psi_j \mu_{\xi_j} + (1-\psi_j) \frac{c_g - q_{\xi_j} \mu_{\xi_j}}{1-q_{\xi_j}}, \sigma_g^2 \right)} & \text{if } \gamma_j = 0 \\ \frac{g_j^* \sum_{\{i:r_{ij}=0\}} x_{ij} \prod_{k=1}^K (b + g_j^* \sum_{\{i:z_i=k, r_{ij}=0\}} s_i)^{-a - \sum_{\{i:z_i=k, r_{ij}=0\}} x_{ij}}}{g_j^{(b-1)} \sum_{\{i:r_{ij}=0\}} x_{ij} \prod_{k=1}^K (b + g_j^{(b-1)} \sum_{\{i:z_i=k, r_{ij}=0\}} s_i)^{-a - \sum_{\{i:z_i=k, r_{ij}=0\}} x_{ij}}} \times \frac{N \left(\tilde{g}_j^*; \psi_j \mu_{\xi_j} + (1-\psi_j) \frac{c_g - q_{\xi_j} \mu_{\xi_j}}{1-q_{\xi_j}}, \sigma_g^2 \right)}{N \left(\tilde{g}_j^{(b-1)}; \psi_j \mu_{\xi_j} + (1-\psi_j) \frac{c_g - q_{\xi_j} \mu_{\xi_j}}{1-q_{\xi_j}}, \sigma_g^2 \right)} & \text{if } \gamma_j = 1 \end{array} \right. .$$

Since ξ , ψ , $\mathbf{q}$, and μ have conjugate full conditionals, we use Gibbs sampling to update them one after another

- Gibbs sampling for updating $\xi_j, j = 1, \dots, p$:

$$p(\xi_j = l | \xi_{-j}, \psi_j, \mathbf{q}, \mu, \tilde{\mathbf{g}}_j) \propto \phi_l^g \mathcal{N}\left(\tilde{\mathbf{g}}_j; \psi_j \mu_l + (1 - \psi_j) \frac{c_g - q_l \mu_l}{1 - q_l}, \sigma_g^2\right).$$

- Gibbs sampling for updating $\psi_j, j = 1, \dots, p$:

$$\pi(\psi_j | \xi_j = l, \psi_{-j}, \mathbf{q}, \mu, \tilde{\mathbf{g}}_j) \propto \begin{cases} (1 - q_l) \mathcal{N}\left(\tilde{\mathbf{g}}_j; \frac{c_g - q_l \mu_l}{1 - q_l}, \sigma_g^2\right) & \text{if } \psi_j = 0 \\ q_l \mathcal{N}(\tilde{\mathbf{g}}_j; \mu_l, \sigma_g^2) & \text{if } \psi_j = 1 \end{cases}.$$

- Gibbs sampling for updating $q_l, l = 1, \dots, L$:

$$q_l | \xi, \psi \sim \text{Be}\left(a_q + \sum_{j=1}^p I(\xi_j = l) I(\psi_j = 1), b_q + \sum_{j=1}^p I(\xi_j = l) I(\psi_j = 0)\right).$$

- Gibbs sampling for updating $\mu_l, l = 1, \dots, L$:

$$\mu_l | \xi, \psi, \mathbf{q}, \tilde{\mathbf{g}} \sim \mathcal{N}\left(\frac{c_l / \sigma_g^2}{e_l / \sigma_g^2 + 1 / \tau_\mu^2}, \frac{1}{e_l / \sigma_g^2 + 1 / \tau_\mu^2}\right),$$

where $c_l = \sum_{\{j: \xi_j = l, \psi_j = 1\}} \tilde{\mathbf{g}}_j - \frac{q_l}{1 - q_l} \sum_{\{j: \xi_j = l, \psi_j = 0\}} \left(\tilde{\mathbf{g}}_j - \frac{c_g}{1 - q_l}\right)$ and $e_l = \sum_{j=1}^p I(\xi_j = l) I(\psi_j = 1) + \sum_{\{j: \xi_j = l, \psi_j = 0\}} \left(\frac{q_l}{1 - q_l}\right)^2$.

- Gibbs Sampling for updating $\phi_l^g, l = 1, \dots, L$ by stick-breaking process [48]:

$$\begin{aligned} \phi_1^g &= v_1, \\ \phi_2^g &= (1 - v_1) v_2, \end{aligned}$$

$\vdots$
 $\phi_L^g = (1 - v_1) \cdots (1 - v_{L-1}) v_L,$
 where $v_l | \xi \sim \text{Be}\left(1 + \sum_{j=1}^p I(\xi_j = l), 1 + \sum_{j=1}^p I(\xi_j > l)\right).$

Update of zero-inflation latent indicator r_{ij} : The full conditionals of the r_{ij} 's can be obtained after considering that we need to consider only those cases for which $x_{ij} = 0$, and

$$\begin{aligned}
 p(r_{ij} | x_{ij} = 0, z_i = k, \gamma_j, \tilde{s}_i, \tilde{g}_j, d_{kj}^*, \pi) &\propto f(x_{ij} = 0 | z_i = k, \gamma_j, \tilde{s}_i, \tilde{g}_j, d_{kj}^*, r_{ij}) p(r_{ij} | \pi) \\
 &= (\exp\{-s_i g_j d_{kj}^{*\gamma_j}\})^{1-r_{ij}} \pi^{r_{ij}} (1-\pi)^{1-r_{ij}} \\
 &= \pi^{r_{ij}} [(1-\pi) \exp\{-s_i g_j d_{kj}^{*\gamma_j}\}]^{1-r_{ij}}.
 \end{aligned}$$

In order to sample from the above full conditional, we proceed with a Metropolis Hasting within Gibbs approach, where within each iteration we first propose a value d_{kj}^* for which $\gamma_j = 0$ by sampling

$$d_{kj}^* | \cdot \sim \text{Ga}\left(a + \sum_{\{i: z_i = k, r_{ij} = 0\}} x_{ij}, b + g_j \sum_{\{i: z_i = k, r_{ij} = 0\}} s_i\right),$$

and then propose a value for r_{ij} by sampling

$$p(r_{ij} | x_{ij} = 0, z_i = k, \gamma_j, \tilde{s}_i, \tilde{g}_j, d_{kj}^*, \pi) = \begin{cases} \frac{\pi^{r_{ij}} [(1-\pi) e^{-s_i g_j}]^{1-r_{ij}}}{\pi + (1-\pi) \exp\{-s_i g_j\}} & \text{if } \gamma_j = 0 \\ \frac{\pi^{r_{ij}} [(1-\pi) \exp\{-s_i g_j d_{kj}^*\}]^{1-r_{ij}}}{\pi + (1-\pi) \exp\{-s_i g_j d_{kj}^*\}} & \text{if } \gamma_j = 1. \end{cases}$$

After the Metropolis Hasting step, we use a Gibbs sampling step to update π , as

$$\pi \mid \mathbf{R}, a_\pi, b_\pi \sim \text{Be} \left(a_\pi + \sum_{i=1}^n \sum_{j=1}^p r_{ij}, b_\pi + \sum_{i=1}^n \sum_{j=1}^p I(x_{ij} = 0) - \sum_{i=1}^n \sum_{j=1}^p r_{ij} \right).$$

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Li, Q;Guindani, M;Reich, BJ;Bondell, HD;Vannucci, M

Title:

A Bayesian mixture model for clustering and selection of feature occurrence rates under mean constraints

Date:

2017-12-01

Citation:

Li, Q., Guindani, M., Reich, B. J., Bondell, H. D. & Vannucci, M. (2017). A Bayesian mixture model for clustering and selection of feature occurrence rates under mean constraints. STATISTICAL ANALYSIS AND DATA MINING, 10 (6), pp.393-409. <https://doi.org/10.1002/sam.11350>.

Persistent Link:

<http://hdl.handle.net/11343/293010>