

Bayesian clustering of multiple zero-inflated outcomes

Beatrice Franzolini¹, Andrea Cremaschi², Willem van den Boom³, and Maria De Iorio⁴

^{1,2,4}Singapore Institute for Clinical Sciences, Agency for Science, Technology and Research

^{3,4}Yong Loo Lin School of Medicine, National University of Singapore

⁴Department of Statistical Science, University College London

Abstract

Several applications involving counts present a large proportion of zeros (excess-of-zeros data). A popular model for such data is the Hurdle model, which explicitly models the probability of a zero count, while assuming a sampling distribution on the positive integers. We consider data from multiple count processes. In this context, it is of interest to study the patterns of counts and cluster the subjects accordingly. We introduce a novel Bayesian nonparametric approach to cluster multiple, possibly related, zero-inflated processes. We propose a joint model for zero-inflated counts, specifying a Hurdle model for each process with a shifted Negative Binomial sampling distribution. Conditionally on the model parameters, the different processes are assumed independent, leading to a substantial reduction in the number of parameters as compared to traditional multivariate approaches. The subject-specific probabilities of zero-inflation and the parameters of the sampling distribution are flexibly modelled via an *enriched* finite mixture with random number of components. This induces a two-level clustering of the subjects based on the zero/non-zero patterns (outer clustering) and on the sampling distribution (inner clustering). Posterior inference is performed through tailored MCMC schemes. We demonstrate the proposed approach on an application involving the use of the messaging service WhatsApp.

Keywords – conditional algorithm, excess-of-zeros data, enriched priors, Hurdle model, finite mixtures, marginal algorithm, nested clustering

1 Introduction

Count data presenting excess of zeros are commonly encountered in applications. These can arise in several settings, such as healthcare, medicine, or sociology. In this scenario, the observations carry structural information about the data-generating process, i.e., an *inflation* of zeros. The analysis of zero-inflated data requires the specification of models beyond standard count distributions, such as Poisson or Negative Binomial. Commonly used models are the Zero-Inflated (Lambert, 1992), the hurdle (Mullahy, 1986) and the zero-altered (Heilbron, 1994) models. The first class assumes the existence of a probability mass at zero and a distribution over $\mathbb{N}_0 = \{0, 1, 2, \dots\}$. This type of models explicitly differentiates between the zeros originating from a common underlying process, such as the utilisation of a service, described by the sampling distribution on \mathbb{N}_0 , and those arising from a structural phenomenon, such as the ineligibility to use the service, which are modelled by the point mass. Very popular Zero-Inflated models are the Zero-Inflated Poisson (ZIP)

and the Zero-Inflated Negative Binomial (ZINB) models, where the sampling distribution is chosen to be a Poisson and a Negative Binomial, respectively. These models allow for inflation in the number of zeros and departures from standard distributional assumptions on the moments of the sampling distribution. For instance, the ZIP model allows the mean and the variance of the distribution to be different from each other (as opposed to a standard Poisson distribution), while the ZINB additionally captures overdispersion in the data.

Hurdle models are a very popular choice of distributions for modelling zero-inflated counts. Differently from the Zero-Inflated ones, these models handle zeros and positive observations separately, assuming on the latter a sampling distribution with support on $\mathbb{N} = \mathbb{N}_0 \setminus \{0\}$. Thus the distribution of the count data is given by:

$$\mathbb{P}(Y_i = y_i) = \begin{cases} (1 - p_i), & y_i = 0 \\ p_i g(y_i | \boldsymbol{\mu}_i), & y_i > 0 \end{cases} \quad (1)$$

where p_i and g now capture two distinct features of the data. Hurdle models present appealing features that can make them preferable to Zero-Inflated models. Firstly, hurdle distributions allow for both inflation and deflation of zero-counts. Indeed, under a Zero-Inflated model, the probability of observing a zero is always greater than the corresponding probability under the sampling distribution, thus making it impossible to capture deflation in the number of zeros (Min and Agresti, 2005). Secondly, and more importantly for our work, the probability of zero counts in hurdle models is independent of the parameters controlling the distribution of non-zero counts. This feature improves interpretability and facilitates parameter estimation. Note that the Zero-Altered model proposed by Heilbron (1994) is a modified hurdle model in which the two parts are connected by specifying a direct link between the model parameters.

Univariate models for zero-inflated data can be extended to multivariate settings, where several variables presenting excess of zeros are recorded, e.g. in applications involving questionnaires or microbiome data analysis. In this context, a multivariate extension of the ZIP model has been proposed by Li et al. (1999), through a finite mixture with ZIP marginals. In this construction, the number of parameters increases linearly as the number d of zero-inflated processes increases, as the total number of parameters is $3d + 2$. See also Liu and Tian (2015), Liu et al. (2019) and, Tian et al. (2018) for simplified versions of the previous construction involving a smaller number of parameters and better distributional properties.

In a Bayesian parametric setting, Fox (2013) proposes the joint modelling of two related zero-inflated outcomes. Their strategy is based on the ZIP model, with the same Bernoulli component to capture the extra zeros for both processes. Correlation between subject-specific outcomes is accounted for through the specification of a joint random effect distribution for the parameters governing the sampling distribution of the two processes. Alternatively, Lee et al. (2020) model the binary variables indicating whether an observation is positive or not via a multivariate probit model (Chib and Greenberg, 1998; García-Zattera et al., 2007). In this approach, the vectors of latent continuous variables characterising the multivariate probit are modelled jointly assuming a random unstructured correlation matrix describing their dependence.

In several applications, knowledge relative to the grouping of the subjects is also available, thus providing additional information that can be exploited in the model (Choo-Wosoba et al., 2018). Moreover, the clustering structure can be estimated by assuming a prior distribution on the partition of the subjects, e.g. via the popular Dirichlet process (Li et al., 2017) or a mixture with random number of components as proposed in Hu et al. (2022).

In the context of Bayesian semiparametric approaches, Shuler et al. (2021) propose to model multivariate zero-inflated count data by linking different Dirichlet Process mixtures of ZINB models through the use of the popular dependent Dirichlet process (MacEachern, 1999). In particular, the probability of zeros and the sampling distribution are modelled via two distinct single-p DDP, where the location parameters of the mixture depend on a categorical covariate. The proposed approach yields flexible estimation of the partition of the subjects, although it does not allow for sharing of information a priori between the two components of the ZINB model, thus yielding two separate clustering structures. A different semiparametric approach is proposed by Arab et al. (2012), which exploits the multivariate ZIP construction of Li et al. (1999) to model bivariate count data, but the proportion of zeros and the intensity of the sampling distribution are modelled through the introduction of spline regression terms. The spline approach is flexible and computationally tractable when d is small. For larger dimensions, this model would induce a non-trivial computational burden.

The focus of this work is clustering of individuals based on multiple, possibly related, zero-inflated processes. To this end, we propose a Bayesian approach for joint modelling of zero-inflated count data, based on finite mixtures with random number of components. In particular, we specify a hurdle model for each process with a shifted Negative Binomial sampling distribution on the positive integers. Let n denote the sample size and d the number of processes under study. The subject-specific probabilities of zero-inflation p_{ij} for the i -th individual and the j -th process, $i = 1, \dots, n$, $j = 1, \dots, d$, and the parameter vector of the sampling distribution μ_{ij} are flexibly modelled via an *enriched* mixture with random number of components, borrowing ideas from the Bayesian nonparametric literature on the Dirichlet process. One of the main novelties of our work is to combine a recent representation of finite mixture models with random number of components presented in Argiento and De Iorio (2022) with a finite extension of the enriched nonparametric prior proposed by Wade et al. (2011) to achieve a two-level clustering of the subjects, where at the *outer* level individuals are clustered based on the pattern of zero/non-zero observations, while within each outer cluster they are grouped at a finer level (which we refer to as *inner* level) according to the distribution of the non-zero counts. Figure 1 provides an illustration of the nested clustering structure.

Enriched priors in Bayesian nonparametrics generalise concepts developed by Consonni and Veronese (2001), who propose a general methodology for the con-

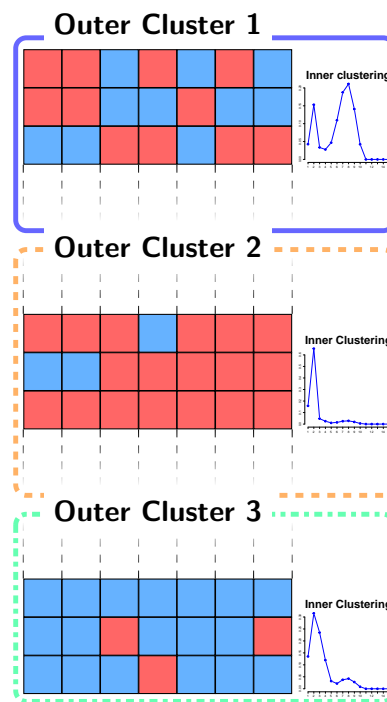


Figure 1: Example of two-level clustering induced by the enriched mixture with random number of components. The observations are first clustered based on their zero/non-zero patterns, indicated in the figure in blue and red, respectively. Within each outer cluster, subjects are grouped based on the sampling distribution of the non-zero observations. The inner clustering structure is here depicted via a multimodal discrete distribution, representing a finite mixture.

struction of enriched conjugate families for the parametric natural exponential families. The idea underlying this approach is to decompose the joint prior distribution for a vector of parameters indexing a multivariate exponential family into tractable conditional distributions. In particular, distributions belonging to the multivariate natural exponential family satisfy the *conditional reducibility* property, which allows reparameterising the distribution in terms of a parameter vector, whose components are variation and likelihood independent. Then, it is possible to construct an enriched standard conjugate family on the parameter vector, closed under i.i.d. sampling, which leads to the breaking down of the global inference procedure into several independent subcomponents. Such parameterisation achieves greater flexibility in prior specification relative to the standard conjugate one, while still allowing for efficient computations (see, for example, [Consonni et al., 2004](#)). An example of this class of parametric priors is the enriched Dirichlet distribution ([Connor and Mosimann, 1969](#)).

In a Bayesian nonparametric framework, [Wade et al. \(2011\)](#) first propose an enrichment of the Dirichlet process ([Ferguson, 1973](#)) that is more flexible with respect to the precision parameter but still conjugate, by defining a joint random probability measure on the measurable product space $(\mathcal{X}, \mathcal{Y})$ in terms of the marginal and conditional distributions, P_X and $P_{Y|X}$, and assigning independent Dirichlet process priors to each of these terms. The enriched Dirichlet process enables a nested clustering structure that is particularly appealing in our setting and allows for a finer control of the dependence structure between X and Y . This construction has been employed also in nonparametric regression problems to model the joint distribution of the response and the covariates ([Wade et al., 2014](#); [Gadd et al., 2020](#)), as well as in longitudinal data analysis ([Zeldow et al., 2021](#)) and causal inference ([Roy et al., 2018](#)). Recently, [Rigon et al. \(2022\)](#) propose the enriched Pitman–Yor process which leads to a more robust clustering estimation.

In this work, we consider the joint distribution of d zero-inflated process, where the d -dimensional vectors of probabilities (p_{i1}, \dots, p_{id}) correspond to X , while the parameters of the sampling distributions μ_{ij} correspond Y . The enrichment of the prior is achieved by modelling both P_X and $P_{Y|X}$ through a mixture with random number of components (see, for instance, [Miller and Harrison, 2018](#)). We exploit the recent construction by [Argiento and De Iorio \(2022\)](#) based on Normalised Independent Finite Point Processes, which allows for a wider choice of prior distributions for the unnormalised weights of the mixture. Therefore, the proposed model offers more flexibility, while preserving computational tractability.

The motivating application for the proposed model is the analysis of multiple count data collected from a questionnaire on the frequency of use of the messaging service WhatsApp ([ClinicalTrials.gov, 2021](#)). In particular, the questionnaire concerns the sharing of COVID-19-related information via WhatsApp messages, either directly or by forwarding, over the course of a week. For each subject, responses to the same seven questions are recorded over seven consecutive days, providing information on a subject’s WhatsApp use (see Table S1 in Supplementary Material). In this set-up, the multiple count processes correspond to the seven questions, all of which displaying an excess of zeros (see Figure S2 in Supplementary Material).

The manuscript is organised as follows. Section 2 introduces a novel enriched prior process for multiple zero-inflated outcomes, while Section 3 describes the Markov chain Monte Carlo (MCMC) algorithm designed for posterior inference. We demonstrate the model on the WhatsApp application in Section 4. We conclude the paper in Section 5.

2 The model

2.1 Likelihood

Let Y_{ij} be the count of subject $i = 1, \dots, n$ for outcome $j = 1, \dots, d$ and let $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{id})$ be the d -dimensional vector of observations for subject i . To take into account the zero-inflated nature of the data, we assume for each outcome j a hurdle model. Each observed count Y_{ij} is equal to zero with probability $1 - p_{ij}$, while with probability p_{ij} it is distributed according to a probability mass function (pmf) $g(\cdot | \boldsymbol{\mu}_{ij})$ with support on \mathbb{N} . Assuming conditional independence among responses, the likelihood for a subject is given by:

$$\mathbb{P}(\mathbf{Y}_i = \mathbf{y}_i | \mathbf{p}_i, \boldsymbol{\mu}_i) = \prod_{j=1}^d f(y_{ij} | p_{ij}, \boldsymbol{\mu}_{ij}) \quad f(y | p, \boldsymbol{\mu}) = \begin{cases} 1 - p, & y = 0 \\ p g(y | \boldsymbol{\mu}), & y > 0 \end{cases} \quad (2)$$

with $\mathbf{p}_i = (p_{i1}, \dots, p_{id}) \in (0, 1)^d$, $\boldsymbol{\mu}_i = (\boldsymbol{\mu}_{i1}, \dots, \boldsymbol{\mu}_{id})$, $i = 1, \dots, n$. In what follows, we set g to be a shifted Negative Binomial distribution with parameters $\boldsymbol{\mu}_{ij} = (r_{ij}, \theta_{ij})$ and pmf:

$$g(y | r_{ij}, \theta_{ij}) = \frac{(y + r_{ij} - 2)!}{(r_{ij} - 1)!(y - 1)!} \theta_{ij}^{y-1} (1 - \theta_{ij})^{r_{ij}}, \quad y \in \mathbb{N} \quad (3)$$

where $r_{ij} \in \mathbb{N}$ and $\theta_{ij} \in (0, 1)$, for $i = 1, \dots, n$ and $j = 1, \dots, d$. Different parametric choices for g are possible (e.g. a shifted Poisson), or even nonparametric alternatives could be employed. Note that the conditional independence assumption among the multiple processes leads to a significant reduction in the number of parameters as compared to multivariate zero-inflated models.

2.2 Enriched finite mixture model

In this work, we propose an *enriched* extension of the Normalised Independent Finite Point Process (Norm-IFPP) of [Argiento and De Iorio \(2022\)](#) and specify a joint prior for $(\mathbf{p}_i, \boldsymbol{\mu}_i)$ as conditionally dependent processes. This allows us to account for interindividual heterogeneity, overdispersion, and outliers and induces data-driven nested clustering of the observations. Each subject is first assigned to an *outer* cluster, and then clustered again at an *inner* level, providing increased interpretability. Differently from previous work on Bayesian nonparametric enriched processes, we opt for a finite mixture with random number of components, where the weights are obtained through the normalisation of a finite point process. Finite mixture models with random number of components have received increasing attention in the last years (see, for example, [Malsiner-Walli et al., 2016](#); [Miller and Harrison, 2018](#)). The representation of [Argiento and De Iorio \(2022\)](#) allows for the specification of a wide range of distributions for the weights and simultaneously leads to effective and widely applicable MCMC schemes on which [Algorithms 1](#) and [2](#) are based. More specifically, they show that a finite mixture model is equivalent to a realisation of a stochastic process with random dimension and infinite-dimensional support, leading to flexible distributions for the weights of the mixture given by the normalisation of a finite point process. We thus employ this approach as it allows for efficient computations via a conditional algorithm, as compared to labour-intensive reversible jump algorithms common in mixture models. An alternative efficient conditional sampler for mixtures with a random number of components is the recently proposed telescopic sampler by [Frühwirth-Schnatter et al. \(2021\)](#).

In the proposed framework, the observations are assumed to be sampled from a mixture with an inner and an outer component. As kernel of the mixture, we assume

the hurdle model in (2), which distinguishes between the probabilities of being non-zero \mathbf{p}_i and the parameters of the sampling distribution $(\mathbf{r}_i, \boldsymbol{\theta}_i)$. The components of the outer mixture are determined by different probabilities of non-zero outcomes, denoted with $\mathbf{p}_m^* = (p_{m1}^*, \dots, p_{md}^*)$, for $m = 1, \dots, M$, with M the number of outer mixture components. The components of the inner mixtures are characterised by distinct parameters of the sampling distribution, denoted with $\mathbf{r}_{ms}^* = (r_{ms1}^*, \dots, r_{msd}^*)$ and $\boldsymbol{\theta}_{ms}^* = (\theta_{ms1}^*, \dots, \theta_{msd}^*)$, for $s = 1, \dots, S_m$ and $m = 1, \dots, M$, where S_m is the number of mixture components within the m -th outer mixture component. Letting $\boldsymbol{\psi}_{msj}^* = (p_{mj}^*, r_{msj}^*, \theta_{msj}^*)$ and $\boldsymbol{\psi}_{ms}^* = (\boldsymbol{\psi}_{ms1}^*, \dots, \boldsymbol{\psi}_{msd}^*)$, the mixture model is as follows:

$$\begin{aligned}
\mathbf{Y}_i \mid \{\boldsymbol{\psi}_{ms}^*\}, \mathbf{w}, \{\mathbf{q}_m\} &\stackrel{\text{iid}}{\sim} \underbrace{\sum_{m=1}^M w_m}_{\text{outer level}} \underbrace{\sum_{s=1}^{S_m} q_{ms}}_{\text{inner level}} \prod_{j=1}^d f(y_{ij} \mid \boldsymbol{\psi}_{msj}^*) \\
\mathbf{q}_m = (q_{m1}, \dots, q_{mS_m}) \mid S_m &\sim \text{Dirichlet}_{S_m}(\gamma_S, \dots, \gamma_S) \\
\mathbf{w} = (w_1, \dots, w_M) \mid M &\sim \text{Dirichlet}_M(\gamma_M, \dots, \gamma_M) \\
\mathbf{p}_m^* &\stackrel{\text{iid}}{\sim} \prod_{j=1}^d \text{Beta}(\alpha, \beta) \\
\mathbf{r}_{ms}^* &\stackrel{\text{iid}}{\sim} \prod_{j=1}^d \text{Geometric}(\zeta) \\
\boldsymbol{\theta}_{ms}^* &\stackrel{\text{iid}}{\sim} \prod_{j=1}^d \text{Beta}(\eta, \lambda) \\
S_1, \dots, S_M \mid M &\stackrel{\text{iid}}{\sim} \text{Poi}_0(\Lambda_S) \\
M &\sim \text{Poi}_0(\Lambda_M)
\end{aligned} \tag{4}$$

where the kernel $f(y_{ij} \mid \boldsymbol{\psi}_{msj}^*)$ is defined via conditionally independent hurdle models in (2)–(3). Here $\text{Dirichlet}_M(\gamma_M, \dots, \gamma_M)$ denotes the symmetric Dirichlet distribution defined on the $(M - 1)$ -dimensional simplex with mean $1/M$, which is the distribution of the normalised mixture weights. $\text{Beta}(\alpha, \beta)$ indicates the Beta distribution with mean $\alpha/(\alpha + \beta)$ and variance $\alpha\beta/((\alpha + \beta)^2(\alpha + \beta + 1))$, $\text{Geometric}(\zeta)$ the Geometric distribution with mean $1/\zeta$, and $\text{Poi}_0(\Lambda)$ the shifted Poisson distribution, such that if $X \sim \text{Poi}_0(\Lambda)$ then $X - 1$ has a Poisson distribution with mean Λ . Moreover, M and S_m , for $m = 1, \dots, M$, indicate the random number of components at the outer and inner level of the enriched Norm-IFPP, respectively.

The outer mixture is a mixture of multivariate Bernoulli distributions, and coincides with the widely-used Latent Class model (Lazarsfeld and Henry, 1968). Moreover, being conditionally independent of the actual values of the non-zero observations, it offers further computation advantages as shown in Section 3.

Model (4) induces a partition of the subject indices $\{1, \dots, n\}$ at an outer and an inner level. Let c_i and z_i , for $i = 1, \dots, n$, denote the allocation variables which indicate to which component of the mixture each subject is assigned to at the outer and inner level, respectively. When two subjects, i and l , are assigned to the same component of the outer level mixture, then the probabilities of observing a zero for the two subjects are the same, $\mathbf{p}_i = \mathbf{p}_l$, and the two subjects are assigned to the same cluster, i.e. $c_i = c_l$. Moreover, if the two subjects are also assigned to the same component of the inner level mixture, we have $z_i = z_l$ and $\boldsymbol{\mu}_i = \boldsymbol{\mu}_l$ (with obviously $c_i = c_l$). However, the vectors of parameters $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_l$ characterising

the sampling distribution might be different even when $c_i = c_l$ and, consequently, the two subjects might be assigned to different clusters at the inner level. This is reflected in the components of the vectors of parameters $(\mathbf{p}_i, \boldsymbol{\mu}_i)$ and $(\mathbf{p}_l, \boldsymbol{\mu}_l)$, which might share only the component corresponding to the probability of zero outcomes or both components.

Using allocation variables, the conditional dependence structure between outer and inner levels is the following. Let

$$\tilde{Y}_{ij} = \begin{cases} 1 & \text{if } Y_{ij} > 0 \\ 0 & \text{if } Y_{ij} = 0 \end{cases} \quad (5)$$

$\tilde{\mathbf{Y}}_i = (\tilde{Y}_{i1}, \dots, \tilde{Y}_{id})$, $\mathcal{C}_m = \{i : c_i = m\}$, and $\mathcal{C}_{ms} = \{i : c_i = m, z_i = s\}$.
Outer mixture:

$$\begin{aligned} \tilde{\mathbf{Y}}_i | \mathbf{p}_i &\sim \prod_{j=1}^d p_{ij}^{\tilde{y}_{ij}} (1 - p_{ij})^{1 - \tilde{y}_{ij}}, \quad \tilde{y}_{ij} \in \{0, 1\} \\ \mathbf{p}_i &= \mathbf{p}_{c_i}^* \\ \mathbf{p}_1^*, \dots, \mathbf{p}_M^* | M &\stackrel{\text{iid}}{\sim} \prod_{j=1}^d \text{Beta}(\alpha, \beta) \\ \Pr(c_i = m) &\propto \Gamma_m, \quad m = 1, \dots, M \\ \Gamma_1, \dots, \Gamma_M &\stackrel{\text{iid}}{\sim} \text{Gamma}(\gamma_M, 1) \\ M &\sim \text{Poi}_0(\Lambda_M) \end{aligned} \quad (6)$$

Inner mixture:

$$\begin{aligned} \mathbf{Y}_i | M, c_i = m, \mathbf{p}_m^*, \mathbf{r}_{mi}, \boldsymbol{\theta}_{mi} &\sim \prod_{j=1}^d f(y_{ij} | p_{mj}^*, r_{mij}, \theta_{mij}) \\ (\mathbf{r}_{mi}, \boldsymbol{\theta}_{mi}) &= (\mathbf{r}_{mz_i}^*, \boldsymbol{\theta}_{mz_i}^*) \\ \mathbf{r}_{m1}^*, \dots, \mathbf{r}_{mS_m}^* | S_m &\stackrel{\text{iid}}{\sim} \prod_{j=1}^d \text{Geometric}(\zeta) \\ \boldsymbol{\theta}_{m1}^*, \dots, \boldsymbol{\theta}_{mS_m}^* | S_m &\stackrel{\text{iid}}{\sim} \prod_{j=1}^d \text{Beta}(\eta, \lambda) \\ \Pr(z_i = s | c_i = m) &\propto \Delta_{ms}, \quad i \in \mathcal{C}_m, \quad s = 1, \dots, S_m \\ \Delta_{m1}, \dots, \Delta_{mS_m} &\stackrel{\text{iid}}{\sim} \text{Gamma}(\gamma_S, 1) \\ S_1, \dots, S_M | M &\stackrel{\text{iid}}{\sim} \text{Poi}_0(\Lambda_S) \end{aligned} \quad (7)$$

where, as before, we denote with \mathbf{p}_m^* , \mathbf{r}_{ms}^* and $\boldsymbol{\theta}_{ms}^*$ the component-specific parameters, which are assumed a priori independent and $\text{Gamma}(\alpha, \beta)$ is the Gamma distribution with mean α/β and variance α/β^2 . The choice of Gamma distribution for the unnormalised weight of the mixture leads to the standard Dirichlet distribution for the normalised weights. In this setting, the computations are greatly simplified by the introduction of a latent variable, conditionally on which the unnormalised weights are independent. See [Argiento and De Iorio \(2022\)](#) for details. Note that the inner mixture is here defined conditionally on the probabilities $p_{m,j}$ of being zero and not on $\tilde{\mathbf{Y}}_i$. Thus, while conditioning on $p_{m,j}$, \mathbf{Y}_i is still allowed to present zero entries. Finally, we highlight that representations (4) and (6)-(7) are equivalent.

3 Inference

Posterior inference can be performed through both a conditional and a marginal algorithm, derived by extending the algorithms by [Argiento and De Iorio \(2022\)](#) to the enriched set-up. The conditional algorithm is described in [Algorithm 1](#), while in [Algorithm 2](#) we present the marginal one.

The conditional algorithm is very flexible and allows for different prior distributions on the weights of the two mixtures as well as on M and S_m (see [Argiento and De Iorio, 2022](#), for details). In [Algorithm 2](#), we use the notation q_M and q_S to denote the prior on M and S_m , respectively, and we set them both equal to a shifted Poisson for the application in [Section 4](#). Furthermore, h_{out} and h_{in} denote the prior distribution on the unnormalised weights (in our case Gamma distributions) of the outer and inner mixture, respectively, $\psi_{\text{out}}(u)$ and $\psi_{\text{in}}(u)$ denote the corresponding Laplace transforms of h_{out} and h_{in} (in our case $\psi_{\text{out}}(u) = (u + 1)^{-\gamma_M}$ and $\psi_{\text{in}}(u) = (u + 1)^{-\gamma_S}$).

To implement the marginal algorithm, we need to derive the marginal likelihood of the data, conditionally on cluster membership. The likelihood in [Eq. \(4\)](#) can be written as:

$$\prod_{i=1}^n \prod_{j=1}^d \left\{ (1 - p_{ij})^{1 - \tilde{y}_{ij}} p_{ij}^{\tilde{y}_{ij}} \left\{ \frac{(y_{ij} + r_{ij} - 2)!}{(r_{ij} - 1)! (y_{ij} - 1)!} \theta_{ij}^{y_{ij} - 1} (1 - \theta_{ij})^{r_{ij}} \right\}^{\tilde{y}_{ij}} \right\} \quad (8)$$

Recall that c_i and z_i denote the labels of the clusters to which the i -th subject belongs to in the outer and the inner clustering, respectively. The marginal likelihood of the data conditionally on the cluster allocation is obtained marginalising with respect to the prior distributions defined in [\(6\)](#) and [\(7\)](#). For a vector of counts y , we obtain:

$$\begin{aligned} \mathcal{M}(y \mid \mathbf{c}, \mathbf{z}) &= \prod_{j=1}^d \left\{ \prod_{m=1}^K \left\{ \mathcal{M}_{\text{Bern}}(y_{j\mathcal{C}_m}^*) \prod_{s=1}^{K_m} \mathcal{M}_{\text{NB}}(y_{j\mathcal{C}_{ms}}^*) \right\} \right\} \\ \mathcal{M}_{\text{Bern}}(y) &= \frac{B(\alpha + n^1, \beta + n^0)}{B(\alpha, \beta)} \\ \mathcal{M}_{\text{NB}}(y) &= \sum_{r=1}^{+\infty} \left\{ \frac{B(\eta + \sum_i (y_i - 1) \tilde{y}_i, \lambda + r \sum_i \tilde{y}_i)}{B(\eta, \lambda)} \prod_i \left(\frac{(y_i + r - 2)!}{(r - 1)! (y_i - 1)!} \right)^{\tilde{y}_i} (1 - \zeta)^{r-1} \zeta \right\} \end{aligned}$$

where $\mathcal{C}_m = \{i : c_i = m\}$, $\mathcal{C}_{ms} = \{i : c_i = m, z_i = s\}$, $y_{j\mathcal{C}_m}^*$ is the vector of observations y_{ij} such that $c_i = m$, for $j = 1, \dots, d$. Similarly, $y_{j\mathcal{C}_{ms}}^*$ is the vector of observations y_{ij} such that $c_i = m$ and $z_i = s$. Moreover, $B(\cdot, \cdot)$ denotes the Beta function, $n^1 = \sum_i \tilde{y}_i$, $n^0 = \sum_i (1 - \tilde{y}_i)$, \tilde{y}_i is defined as in [Eq. \(5\)](#) and the last two summations run over the elements of the vector \tilde{y} . Here K and K_m are the numbers of clusters at the outer and inner level, respectively. Note that by cluster we mean an occupied component (i.e. a mixture component to which at least one observation has been assigned), with $K \leq M$ and $K_m \leq S_m$, $m = 1, \dots, M$.

When implementing the marginal algorithm, after updating the latent variables \bar{U} and U_m , we could add an extra step involving a shuffle of the nested partition structure as suggested by [Wade et al. \(2014\)](#) to improve mixing. More details and an empirical comparison of the two algorithms are provided in [Section S3](#) of [Supplementary Material](#).

4 Application to WhatsApp use during COVID-19

4.1 Data description and preprocessing

We apply our model to a dataset on WhatsApp use during COVID-19 ([ClinicalTrials.gov, 2021](#)). The data consist of a questionnaire filled out by participants living

Algorithm 1 Conditional algorithm

Input: $(y_{ij})_{ij}$ and parameter initialisation

Output: posterior distribution of cluster allocation and other parameters

for i in $1:n$ **do**

Sample c_i and z_i from

$$\mathbb{P}[c_i = m, z_i = s \mid \text{rest}] \propto \Gamma_m \Delta_{ms} \prod_{j=1}^d f(y_{ij} \mid \mathbf{p}_m^*, \mathbf{r}_{ms}^*, \boldsymbol{\theta}_{ms}^*)$$

end for

Compute K , the number of allocated components at the outer level

Relabel the outer level clusters so that the first K components of the mixture are allocated

Sample the latent variable \bar{u} from Gamma $(n, \sum_{m=1}^M \Gamma_m)$

Set $M = K + x$, where

$$x \sim q_x \quad q_x \propto \frac{(x + K)!}{x!} \psi_{\text{out}}(\bar{u})^x q_M(K + x) \quad \text{for } x = 0, 1, \dots$$

Sample the unnormalised weights of the outer measure from

$$\mathbb{P}[\Gamma_m \in d\omega \mid \text{rest}] \propto \omega^{n_m} e^{-\omega \bar{u}} h_{\text{out}}(\omega) d\omega \quad \text{for } m = 1, \dots, M$$

where n_m is the cardinality of outer level cluster m and $n_m = 0$ for $m > K$

for m in $1:K$ **do**

Sample \mathbf{p}_m^* from the full conditional.

Compute the number K_m of allocated components at the inner level

Relabel the inner level clusters so that the first K_m components are allocated

Sample the latent variable u_m from Gamma $(n_m, \sum_{s=1}^{S_m} \Delta_{ms})$

Set $S_m = K_m + x$ where

$$x \sim q_x \quad q_x \propto \frac{(x + K_m)!}{x!} \psi_{\text{in}}(u_m)^x q_S(K_m + x) \quad \text{for } x = 0, 1, \dots$$

Sample the unnormalised weights of the m -th inner mixture from

$$\mathbb{P}[\Delta_{ms} \in dq \mid \text{rest}] \propto q^{n_{ms}} e^{-\omega u_m} h_{\text{in}}(q) dq \quad \text{for } s = 1, \dots, S_m$$

where n_{ms} is the cardinality of inner level cluster s and $n_{ms} = 0$ for $s > K_m$

for s in $1:K_m$ **do**

Sample $(\mathbf{r}_{ms}^*, \boldsymbol{\theta}_{ms}^*)$ from the full conditional

end for

for s in $(K_m + 1):S_m$ **do**

Sample \mathbf{r}_{ms}^* from the prior

Sample $\boldsymbol{\theta}_{ms}^*$ from the prior

end for

end for

for m in $(K + 1):M$ **do**

Sample \mathbf{p}_m^* and S_m from the prior

for s in $1:S_m$ **do**

Sample Δ_{ms} from the prior

Sample \mathbf{r}_{ms}^* from the prior

Sample $\boldsymbol{\theta}_{ms}^*$ from the prior

end for

end for

Algorithm 2 Marginal algorithm

Input: $(y_{ij})_{ij}$ and parameter initialisation

Output: posterior distribution of cluster allocation and other posterior summaries

for i in $1:n$ **do**
 Sample c_i

$\mathbb{P}[c_i = m \mid \mathbf{c}^{-(i)}, \mathbf{z}^{-(i)}, \bar{U}, U_1, \dots, U_K]$

$$\propto \begin{cases} \left(n_m^{-(i)} + \gamma_M \right) \prod_{j=1}^d \frac{\mathcal{M}_{\text{Bern}} \left(y_{j\mathcal{C}_m^{+(i)}}^* \right)}{\mathcal{M}_{\text{Bern}} \left(y_{j\mathcal{C}_m^{-(i)}}^* \right)} \left(\frac{n_{ms}^{-(i)} + \gamma_S}{L_m^{-(i)}} \prod_{j=1}^d \frac{\mathcal{M}_{\text{NB}} \left(y_{j\mathcal{C}_m^{+(i)}}^* \right)}{\mathcal{M}_{\text{NB}} \left(y_{j\mathcal{C}_m^{-(i)}}^* \right)} \right) & \text{if } m = m_{\text{old}} \\ \quad + \frac{L_m^{-(i)} - n_m^{-(i)} - \gamma_S}{L_m^{-(i)}} \prod_{j=1}^d \mathcal{M}_{\text{NB}} \left(y_{ij} \right) & \\ \frac{\Lambda_M + (K^{-(i)} + 1)(\bar{u} + 1)^{\gamma_M}}{\Lambda_M + K^{-(i)}(\bar{u} + 1)^{\gamma_M}} \frac{\Lambda_M \gamma_M}{(\bar{u} + 1)^{\gamma_M}} \prod_{j=1}^d \mathcal{M}_{\text{Bern}} \left(y_{ij} \right) \mathcal{M}_{\text{NB}} \left(y_{ij} \right) & \text{otherwise} \end{cases}$$

where $n_m^{-(i)}$ and $n_{ms}^{-(i)}$ are the cardinalities of outer and inner clusters after removing the i -th observation, $\mathcal{C}_m^{-(i)} = \mathcal{C}_m \setminus \{i\}$ and $\mathcal{C}_m^{+(i)} = \mathcal{C}_m \cup \{i\}$, and similarly for $\mathcal{C}_{ms}^{+(i)}$, $\mathcal{C}_{ms}^{-(i)}$, $K^{-(i)}$, and $K_m^{-(i)}$. Here the subscript ‘old’ denotes an existing (occupied) cluster and

$$L_m^{-(i)} = \frac{\Lambda_S + (K_m^{-(i)} + 1)(u_m + 1)^{\gamma_S}}{\Lambda_S + K_m^{-(i)}(u_m + 1)^{\gamma_S}} \frac{\Lambda_S \gamma_S}{(u_m + 1)^{\gamma_S}} + n_m^{-(i)} + \gamma_S$$

Sample z_i

$\mathbb{P}[z_i = s \mid \mathbf{c}, \mathbf{z}^{-(i)}, U_m]$

$$\propto \begin{cases} \left(n_{ms}^{-(i)} + \gamma_S \right) \prod_{j=1}^d \frac{\mathcal{M}_{\text{NB}} \left(y_{j\mathcal{C}_{ms}^{+(i)}}^* \right)}{\mathcal{M}_{\text{NB}} \left(y_{j\mathcal{C}_{ms}^{-(i)}}^* \right)} & \text{if } s = s_{\text{old}} \\ \frac{\Lambda_S + (K_m^{-(i)} + 1)(u_m + 1)^{\gamma_S}}{\Lambda_S + K_m^{-(i)}(u_m + 1)^{\gamma_S}} \frac{\Lambda_S \gamma_S}{(u_m + 1)^{\gamma_S}} \prod_{j=1}^d \mathcal{M}_{\text{NB}} \left(y_{ij} \right) & \text{otherwise} \end{cases}$$

Note that when a subject i is assigned to a new outer cluster, then the full conditional distribution of z_i is degenerate and a new auxiliary variable U_m has to be sampled before moving to the next subject $i + 1$.

end for

Sample the latent variables \bar{U} and U_1, \dots, U_K from their full conditional:

$$\mathbb{P}[\bar{U} \in d\bar{u} \mid \text{rest}] \propto \left(\frac{\Lambda_M}{(\bar{u} + 1)^{\gamma_M}} + K \right) \exp \left\{ \frac{\Lambda_M}{(\bar{u} + 1)^{\gamma_M}} \right\} \frac{\bar{u}^{n-1}}{(\bar{u} + 1)^{n+K\gamma_M}}, \quad \bar{u} > 0$$

$$\mathbb{P}[U_m \in du_m \mid \text{rest}] \propto \left(\frac{\Lambda_S}{(u_m + 1)^{\gamma_S}} + K_m \right) \exp \left\{ \frac{\Lambda_S}{(u_m + 1)^{\gamma_S}} \right\} \frac{(u_m)^{n_m-1}}{(u_m + 1)^{n_m+K_m\gamma_S}}, \quad u_m > 0$$

in India. Each subject answers the same $d = 7$ questions for $T = 7$ consecutive days on the number of ($j = 1$) COVID-19 messages forwarded, ($j = 2$) WhatsApp groups to which COVID-19 messages were forwarded, ($j = 3$) people to whom COVID-19 messages were forwarded, ($j = 4$) unique forwarded messages received in personal

chats, ($j = 5$) people from whom forwarded messages were received, ($j = 6$) personal chats that discussed COVID-19, ($j = 7$) WhatsApp groups that mentioned COVID-19. Table S1 in Supplementary Material provides the list of the questions, as well as a brief description. In what follows, the first replicate ($t = 1$) corresponds to Sunday for all subjects, $t = 2$ to Monday, up to $T = 7$ corresponding to Saturday. The questionnaire responses were collected in June and July 2021, during India’s infection wave of the Delta variant of the SARS-CoV-2 virus that causes coronavirus disease 2019 (COVID-19).

From the initial 1156 respondents, we remove two subjects for which no answers are available, resulting in a final sample size of $n = 1154$. Moreover, 19% of the observations are missing. We also treat counts higher than 400, which are very rare (7 observations out of 56 546), as missing data as they are very far from the range of the majority of the data. We handle missing data using a two-step procedure. Firstly, whenever possible, we recover missing zeros using deterministic imputation based on respondent’s answers to other sections of the questionnaire. For instance, if the answer to the question “did you send any message of this kind today?” is “no” and there is a missing value for the question “how many?”, we can reasonably assume that the answer to the latter question is zero. In this way, we can recover 0.5% of the missing observations. Secondly, the remaining missing values are imputed using random forest imputation (as implemented in the R package `mice`, van Buuren and Groothuis-Oudshoorn, 2011). In Section S2 of Supplementary Material, we provide more details on the data imputation technique and we present an empirical study to quantify the impact of data imputation on the results presented in the next section. Figure S2 of Supplementary Material displays the data after imputation.

To account for the fact that T repeated observations are available for each subject and process, we need to slightly modify model (4). We do so by assuming that the different time points are independent of each other, so that repeated observations can be straightforwardly included into the proposed model. Let Y_{ijt} denote the count for the i -th subject and the j -th process at time t , $i = 1, \dots, n$, $j = 1, \dots, d$ and $t = 1, \dots, T$. We assume that Y_{ijt} are conditionally independent, given the parameters of the model. Thus, the likelihood contribution of each subject i is given by $\prod_{t=1}^T \prod_{j=1}^d f(y_{ijt} | \psi_{msj}^*)$. It must be highlighted that we are clustering individuals based on the pattern of all their observations, at each time point t and for each process j .

Finally we note that, thanks to the probabilistic structure of the hurdle model for zero-inflated data, \mathbf{p}_i and the sampling distribution $g(\cdot | \boldsymbol{\mu}_i)$ reflect two distinct features of the respondents’ behaviour: \mathbf{p}_i represents the probability of engaging in some COVID-19 related WhatsApp activity, while $g(\cdot | \boldsymbol{\mu}_i)$ captures the behaviour of those subjects who have actually engaged in the activity.

4.2 Results

Posterior inference is performed through the conditional algorithm described in Algorithm 1. We run the algorithm for 15 000 MCMC iterations, discarding the first 5000 as burn-in.

Figure 2 shows that, at the outer level, the posterior distributions of the number of both components and clusters present a mode at the value three.

As point estimate of the cluster allocation, we report the configuration that minimises the posterior expectation of Binder’s loss function (Binder, 1978) under equal misclassification costs, which is a common choice in the applied Bayesian nonparametrics literature (Lau and Green, 2007). Briefly, this expectation of the loss measures the difference for all possible pairs of subjects between the posterior probability of co-clustering and the estimated cluster allocation. We refer to the





Figure 2: Posterior distribution of the number of outer mixture components M (left panel) and clusters K , i.e., number of occupied components to which at least one observation is assigned (right panel).



(a) Outer clusters relative frequencies

(b) Outer level Bernoulli parameters p_{mj}^*

Figure 3: Relative frequency of the outer clusters (left panel) and the posterior means of the cluster-specific probabilities of a non-zero count p_{mj}^* (right panel) corresponding to the posterior estimate of the clustering allocation obtained by minimising Binder’s loss function.

resulting cluster allocation as the Binder estimate.

The Binder estimate of the outer clustering contains three clusters, whose characteristics are summarised in Figures 3 and 4. The largest cluster corresponds to WhatsApp users who on most days report a zero count for all $d = 7$ questions. The individuals in the other two clusters use WhatsApp more frequently when it comes to forwarding COVID-19 messages ($j = 1, 2$), receiving forwarded messages ($j = 3, 4, 5$) and having COVID-19 mentioned in their WhatsApp groups ($j = 7$). The main feature distinguishing Cluster 2 from Cluster 3 in terms of probabilities p_i of non-zero counts is that on most days Cluster 2, unlike Cluster 3, discusses COVID-19 also in personal chats (question $j = 6$).

Figures 5 and 6 display the main characteristics of the inner clusters. We are interested in the posterior distribution of the number of the inner clusters per outer cluster, as well as the inner clustering within each outer cluster. To this end, we run the MCMC algorithm fixing the outer cluster allocation to its Binder estimate, thus obtaining the conditional posterior distribution of the inner clustering. The results reveal substantial variability in the distribution of non-zero counts within outer Clusters 1 and 2 (see Figure 5, bottom panel). The majority of counts in outer Cluster 1 are zero, leaving little variation in the counts for the inner clustering. As most individuals present zero counts (for most processes) at an inner cluster level, it becomes difficult to detect specific patterns as it is also evident from the fact that many co-clustering probabilities are in the range 0.3-0.6 (see Figure 6). Notably, around a quarter of the individuals in outer Cluster 2, as captured by its inner Cluster 2, forward COVID-19 messages to many more people (question $j = 3$) than




Figure 4: Estimated pmfs for the seven questions within each outer cluster (conditionally on the counts being positive) corresponding to the posterior estimate of the clustering allocation obtained by minimising Binder’s loss function. Shaded areas represent the 95% credible intervals.

subjects in inner Cluster 1 of outer Cluster 2. Figure 4 also supports the fact that outer Cluster 2 engages with WhatsApp in a much more persistent manner than the other outer clusters. These results highlight that a sizeable minority of WhatsApp users has a relatively large propensity to spread COVID-19 messages during a critical phase of the pandemic. This is in line with a similar survey in Singapore (Tan et al., 2021) and findings on “superspreaders” on other social media.

5 Conclusion

In this work, we propose a Bayesian model for multiple zero-inflated count data, building on the well-established hurdle model and exploiting the flexibility of finite mixture models with random number of components. The main contribution of this work is the construction of an *enriched* finite mixture with random number of components, which allows for two level (nested) clustering of the subjects based on their pattern of counts across different processes. This structure enhances interpretability of the results and has the potential to better capture important features of the data. We design a conditional and a marginal MCMC sampling scheme to perform posterior inference. The proposed methodology has wide applicability, since excess-of-zeros count data arise in many fields. Our motivating application involves answers to a questionnaire on the use of WhatsApp in India during the COVID-19 pandemic. Our analysis identifies a two-level clustering of the subjects: the outer cluster allocation reflects daily probabilities of engaging in different WhatsApp activities, while the inner level informs on the number of messages conditionally on the fact that the subject is indeed receiving/sending messages on WhatsApp. Any two subjects are clustered together if they show a similar pattern across the multiple responses. We find three different well-distinguished respondent behaviours corresponding to the three outer clusters: (i) subjects with low probability of daily utilisation; (ii) subjects with high probability of sending/receiving all types of mes-




Figure 5: Posterior distribution of the number of inner clusters per outer cluster (top panel), relative frequency of the inner clusters corresponding to the Binder estimate of the inner cluster allocation (central panel), cluster-specific empirical means of the counts (bottom panel). For outer Cluster 1, the latter is only shown for the four largest inner clusters for visualisation purposes. Results are obtained conditionally on the Binder estimate of the outer clustering.

sages and (iii) subjects with high probability for all considered messages except for non-forwarded messages in personal chats. Interestingly, the inner level clustering and the outer cluster specific estimates of the sampling distribution g highlight similarities between the outer Clusters 1 and 3, where subjects tend to send/receive fewer messages compared to outer Cluster 2. Moreover, we are able to identify those subjects with high propensity to spread COVID-19 messages during the critical phase of the pandemic and for these subjects we do not find notable differences in terms of types of messages sent or received. Our results are in line with existing literature on the topic. Future work involves the development of more complex clustering hierarchies and techniques able to identify processes that most inform the clustering structure.




Figure 6: Heatmaps of the posterior co-clustering probabilities for the inner clusters per outer cluster. Results are obtained conditionally on the Binder estimate of the outer cluster allocation. Observations are reordered based on the co-clustering probability profiles, through hierarchical clustering.

Funding: This work was partially supported by the NUS Centre for Trusted Internet and Community [grant number CTIC-RP-20-09].

Acknowledgements: We thank Dr. Jean Liu and the Synergy Lab at Yale-NUS College for providing the data.

References

- Arab, A., S. H. Holan, C. K. Wikle, and M. L. Wildhaber (2012). Semiparametric bivariate zero-inflated Poisson models with application to studies of abundance for multiple species. *Environmetrics* 23(2), 183–196.
- Argiento, R. and M. De Iorio (2022). Is infinity that far? A Bayesian nonparametric perspective of finite mixture models. *The Annals of Statistics forthcoming*. arXiv:1904.09733v1.
- Binder, D. A. (1978). Bayesian cluster analysis. *Biometrika* 65(1), 31–38.
- Chib, S. and E. Greenberg (1998). Analysis of multivariate probit models. *Biometrika* 85(2), 347–361.
- Choo-Wosoba, H., J. Gaskins, S. Levy, and S. Datta (2018). A Bayesian approach for analyzing zero-inflated clustered count data with dispersion. *Statistics in Medicine* 37(5), 801–812.
- ClinicalTrials.gov (2021). WhatsApp in India during the COVID-19 pandemic. Identifier NCT04918849. U.S. National Library of Medicine. Available from <https://clinicaltrials.gov/ct2/show/NCT04918849>.
- Connor, R. J. and J. E. Mosimann (1969). Concepts of independence for proportions with a generalization of the Dirichlet distribution. *Journal of the American Statistical Association* 64(325), 194–206.
- Consonni, G. and P. Veronese (2001). Conditionally reducible natural exponential families and enriched conjugate priors. *Scandinavian Journal of Statistics* 28(2), 377–406.
- Consonni, G., P. Veronese, and E. Gutiérrez-Pena (2004). Reference priors for exponential families with simple quadratic variance function. *Journal of Multivariate Analysis* 88(2), 335–364.

- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* 1(2), 209–230.
- Fox, J.-P. (2013). Multivariate zero-inflated modeling with latent predictors: Modeling feedback behavior. *Computational statistics & data analysis* 68, 361–374.
- Frühwirth-Schnatter, S., G. Malsiner-Walli, and B. Grün (2021). Generalized mixtures of finite mixtures and telescoping sampling. *Bayesian Analysis* 16(4), 1279–1307.
- Gadd, C., S. Wade, and A. Boukouvalas (2020, 26–28 Aug). Enriched mixtures of generalised Gaussian process experts. In S. Chiappa and R. Calandra (Eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, Volume 108 of *Proceedings of Machine Learning Research*, pp. 3144–3154. PMLR.
- García-Zattera, M. J., A. Jara, E. Lesaffre, and D. Declerck (2007). Conditional independence of multivariate binary data with an application in caries research. *Computational statistics & data analysis* 51(6), 3223–3234.
- Heilbron, D. C. (1994). Zero-altered and other regression models for count data with added zeros. *Biometrical Journal* 36(5), 531–547.
- Hu, G., H.-C. Yang, Y. Xue, and D. K. Dey (2022). Zero-inflated Poisson model with clustered regression coefficients: Application to heterogeneity learning of field goal attempts of professional basketball players. *Canadian Journal of Statistics advance online publication*.
- Lambert, D. (1992). Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics* 34(1), 1–14.
- Lau, J. W. and P. J. Green (2007). Bayesian model-based clustering procedures. *Journal of Computational and Graphical Statistics* 16(3), 526–558.
- Lazarsfeld, P. F. and N. W. Henry (1968). *Latent Structure Analysis*. Houghton Mifflin.
- Lee, K. H., B. A. Coull, A.-B. Moscicki, B. J. Paster, and J. R. Starr (2020). Bayesian variable selection for multivariate zero-inflated models: Application to microbiome count data. *Biostatistics* 21(3), 499–517.
- Li, C.-S., J.-C. Lu, J. Park, K. Kim, P. A. Brinkley, and J. P. Peterson (1999). Multivariate zero-inflated Poisson models and their applications. *Technometrics* 41(1), 29–38.
- Li, Q., M. Guindani, B. J. Reich, H. D. Bondell, and M. Vannucci (2017). A Bayesian mixture model for clustering and selection of feature occurrence rates under mean constraints. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 10(6), 393–409.
- Liu, Y. and G.-L. Tian (2015). Type I multivariate zero-inflated Poisson distribution with applications. *Computational Statistics & Data Analysis* 83, 200–222.
- Liu, Y., G.-L. Tian, M.-L. Tang, and K. C. Yuen (2019). A new multivariate zero-adjusted Poisson model with applications to biomedicine. *Biometrical Journal* 61(6), 1340–1370.

- MacEachern, S. N. (1999). Dependent nonparametric processes. In *ASA proceedings of the section on Bayesian statistical science*, Volume 1, pp. 50–55. Alexandria, Virginia. Virginia: American Statistical Association; 1999.
- Malsiner-Walli, G., S. Frühwirth-Schnatter, and B. Grün (2016). Model-based clustering based on sparse finite Gaussian mixtures. *Statistics and Computing* 26(1), 303–324.
- Miller, J. W. and M. T. Harrison (2018). Mixture models with a prior on the number of components. *Journal of the American Statistical Association* 113(521), 340–356.
- Min, Y. and A. Agresti (2005). Random effect models for repeated measures of zero-inflated count data. *Statistical Modelling* 5(1), 1–19.
- Mullahy, J. (1986). Specification and testing of some modified count data models. *Journal of Econometrics* 33(3), 341–365.
- Rigon, T., B. Scarpa, and S. Petrone (2022). Enriched Pitman-Yor processes. arXiv:2003.12200v2.
- Roy, J., K. J. Lum, B. Zeldow, J. D. Dworkin, V. L. Re III, and M. J. Daniels (2018). Bayesian nonparametric generative models for causal inference with missing at random covariates. *Biometrics* 74(4), 1193–1202.
- Shuler, K., S. Verbanic, I. A. Chen, and J. Lee (2021). A bayesian nonparametric analysis for zero-inflated multivariate count data with application to microbiome study. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 70(4), 961–979.
- Tan, E. Y., R. R. Wee, Y. E. Saw, K. J. Heng, J. W. Chin, E. M. Tong, and J. C. Liu (2021). Tracking private WhatsApp discourse about COVID-19 in Singapore: Longitudinal infodemiology study. *Journal of Medical Internet Research* 23(12), e34218.
- Tian, G.-L., Y. Liu, M.-L. Tang, and X. Jiang (2018). Type I multivariate zero-truncated/adjusted Poisson distributions with applications. *Journal of Computational and Applied Mathematics* 344, 132–153.
- van Buuren, S. and K. Groothuis-Oudshoorn (2011). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software* 45(3), 1–67.
- Wade, S., D. B. Dunson, S. Petrone, and L. Trippa (2014). Improving prediction from Dirichlet process mixtures via enrichment. *The Journal of Machine Learning Research* 15(1), 1041–1071.
- Wade, S., S. Mongelluzzo, and S. Petrone (2011). An enriched conjugate prior for Bayesian nonparametric inference. *Bayesian Analysis* 6(3), 359–385.
- Zeldow, B., J. Flory, A. Stephens-Shields, M. Raebel, and J. A. Roy (2021). Functional clustering methods for longitudinal data with application to electronic health records. *Statistical Methods in Medical Research* 30(3), 655–670.

Supplementary material for
“Bayesian clustering of multiple zero-inflated outcomes”
by

Beatrice Franzolini, Andrea Cremaschi, Willem van den Boom and Maria De Iorio

S1 Questionnaire questions

Table 1: Items in the questionnaire on WhatsApp activity. Each question corresponds to a process in the model.

| Index j | Behaviour | Type of Message | Question |
|-----------|-------------|------------------|--|
| 1 | Sender | Forwarded | How many different COVID-19 messages did you forward today? |
| 2 | Sender | Forwarded | How many different WhatsApp groups did you forward COVID-19 messages to today? |
| 3 | Sender | Forwarded | How many different people did you forward COVID-19 messages to today? |
| 4 | Recipient | Forwarded | How many unique forwarded messages did you receive in your personal chats? |
| 5 | Recipient | Forwarded | How many different people did you receive forwarded messages from today? |
| 6 | Participant | Personal comment | How many of your personal chats discussed COVID-19 today? |
| 7 | Recipient | Both | How many of your WhatsApp groups mentioned COVID-19 today? |

S2 Data imputation

Missing data are imputed via multivariate imputation by chained equations as implemented in the R package `mice` [van Buuren and Groothuis-Oudshoorn \(2011\)](#). The method can be summarised as follows: for each variable presenting missing entries a conditional distribution is estimated given all other variables. To this end, random forests are employed. Missing values are then imputed from such distribution. Imputation is iterated five times conditionally on previous values. More details on the procedure can be found in [van Buuren and Groothuis-Oudshoorn \(2011\)](#).

In [Table 2](#) we compare results presented in the main manuscript with those obtained from other four imputed datasets. The table contains the following posterior inference summaries: maximum a posteriori (MAP) estimates of the number of outer clusters (MAP-est K) and of the total number of inner clusters (MAP-est $\sum_m K_m$); Hellinger distance (H-dist) and Jensen-Shannon Divergence (JSD) between the posterior distribution of K and $\sum_m K_m$ from the analysis presented in the main manuscript and the distribution obtained from each replicate; two adjusted rand indexes between the outer partition in the main manuscript and the outer partition obtained from each of the replication dataset. We estimate the partition by (i) minimising the Binder loss estimate of the outer clustering (Adj Rand index Binder) (ii) applying the clustering method of Medvedovic (Adj Rand index Medvedovic), as implemented in the R package `mcclust`.

The Hellinger distance and the Jensen-Shannon divergence take values between 0 and 1, with higher values corresponding to more dissimilarities between the two distributions. The adjusted rand index varies between 0 and 1, with higher values indicating higher similarities between clustering allocations.

Table 2: Impact of data imputation.

| | Main Analysis | Replica1 | Replica 2 | Replica 3 | Replica 4 |
|-----------------------|---------------|----------|-----------|-----------|-----------|
| MAP-est K | 3 | 3 | 3 | 3 | 3 |
| MAP-est $\sum_m K_m$ | 13 | 13 | 10 | 11 | 13 |
| H-dist K | - | 0.309 | 0.425 | 0.269 | 0.307 |
| JSD K | - | 0.101 | 0.189 | 0.081 | 0.098 |
| H-dist $\sum_m K_m$ | - | 0.299 | 0.453 | 0.323 | 0.285 |
| JSD $\sum_m K_m$ | - | 0.113 | 0.249 | 0.128 | 0.106 |
| Adj Rand index Binder | - | 0.732 | 0.679 | 0.720 | 0.781 |
| Adj Rand index Medv. | - | 0.749 | 0.722 | 0.749 | 0.807 |

S3 Marginal algorithm

We compare the performance of the two sampling schemes presented in Section 3 (i.e., the conditional algorithm and the marginal algorithm) on the dataset on WhatsApp use during COVID-19.

We perform posterior inference employing both MCMC schemes and compare them in terms of effective sample size per iteration (ESS/#iter) and integrated autocorrelation time (IAT) computed with the R packages `coda` and `LaplacesDemon`, respectively. ESS/#iter and IAT are computed for the number of outer clusters, the number of inner clusters, and the log-likelihood $\log \mathcal{M}(y | \mathbf{c}, \mathbf{z})$. Results for the conditional algorithm are based on 10000 iterations. Results for the marginal algorithm are obtained with 1000 iterations.

Table 3 highlights that the marginal algorithm, as expected, leads to higher effective sample sizes per iteration as a consequence of the smaller number of parameters to sample and the lower dependence across them. This ultimately allows for a better mixing of the chain.

Table 3: Comparison between conditional and marginal algorithm.

| | Log-likelihood | | Num. of outer clusters | | Num. of inner clusters | |
|-----------|----------------|-------|------------------------|--------|------------------------|-------|
| | Cond. | Marg. | Cond. | Marg. | Cond. | Marg. |
| ESS/#iter | 0.013 | 0.208 | 0.005 | 0.076 | 0.006 | 0.142 |
| IAT | 81.266 | 4.937 | 215.0 | 15.634 | 126.80 | 8.011 |

However, one iteration of the marginal algorithm is associated with a higher computational cost mainly due to the evaluation of the integrals in $\mathcal{M}_{\text{Bern}}$ and \mathcal{M}_{NB} . In particular, the computation of \mathcal{M}_{NB} requires approximating an infinite sum. Here, we employed a Monte Carlo approximation based on 100 samples. Alternatively, we could rely on numerical approximation or introduce an auxiliary variable to be sampled from its full conditional.

Additionally, the variability of the posterior distributions obtained with the marginal algorithm depends on the availability of an efficient mechanism for sampling the auxiliary variables corresponding to the inner mixtures U_1, \dots, U_M from their full conditional. Here, given the importance of U_1, \dots, U_M in determining the full conditional of the allocating variables \mathbf{c} and \mathbf{z} , we sample the auxiliary variables any time a new inner cluster is created, in addition to sampling them at the end of the `for` cycle as described in Section 3.

While the marginal algorithm can be used to derive a point estimate of the predictive distribution for a new observation, it does not provide appropriate credible intervals for either outer or inner mixtures. Thus, in terms of uncertainty quantification, it can be used only for deriving credible balls for the clustering structure.

On the contrary, the conditional algorithm provides proper uncertainty quantification for both the mixtures and the clustering, while it can also be straightforwardly implemented.

S4 Data

Figure 7: Bar plots of the frequencies for the $d = 7$ responses across the seven days of the week. For visualisation purposes, plots contain only counts smaller than 20.

