Bayesian variable selection for multivariate zero-inflated models: Application to microbiome count data

KYU HA LEE*

The Forsyth Institute, 245 First Street, Cambridge, MA 02142, USA and Department of Oral Health Policy and Epidemiology, Harvard School of Dental Medicine, Boston, MA 02115, USA
klee@forsyth.org

BRENT A. COULL

Department of Biostatistics, Harvard T.H. Chan School of Public Health, 665 Huntington Avenue, Boston, MA 02115, USA

ANNA-BARBARA MOSCICKI

Department of Pediatrics, David Geffen School of Medicine, University of California Los Angeles, Los Angeles, CA 10833, USA

BRUCE J. PASTER

The Forsyth Institute, 245 First Street, Cambridge, MA 02142, USA and Department of Oral Medicine, Infection, and Immunity, Harvard School of Dental Medicine, Boston, MA 02115, USA

JACQUELINE R. STARR

The Forsyth Institute, 245 First Street, Cambridge, MA 02142, USA and Department of Oral Health Policy and Epidemiology, Harvard School of Dental Medicine, Boston, MA 02115, USA

SUMMARY

Microorganisms play critical roles in human health and disease. They live in diverse communities in which they interact synergistically or antagonistically. Thus for estimating microbial associations with clinical covariates, such as treatment effects, joint (multivariate) statistical models are preferred. Multivariate models allow one to estimate and exploit complex interdependencies among multiple taxa, yielding more powerful tests of exposure or treatment effects than application of taxon-specific univariate analyses. Analysis of microbial count data also requires special attention because data commonly exhibit zero inflation, i.e., more zeros than expected from a standard count distribution. To meet these needs, we developed a Bayesian variable selection model for multivariate count data with excess zeros that incorporates information on the covariance structure of the outcomes (counts for multiple taxa), while estimating associations with the mean levels of these outcomes. Though there has been much work on zero-inflated models for longitudinal data, little attention has been given to high-dimensional multivariate zero-inflated data modeled via a general correlation structure. Through simulation, we compared performance of the

*To whom correspondence should be addressed.

proposed method to that of existing univariate approaches, for both the binary (“excess zero”) and count parts of the model. When outcomes were correlated the proposed variable selection method maintained type I error while boosting the ability to identify true associations in the binary component of the model. For the count part of the model, in some scenarios the univariate method had higher power than the multivariate approach. This higher power was at a cost of a highly inflated false discovery rate not observed with the proposed multivariate method. We applied the approach to oral microbiome data from the Pediatric HIV/AIDS Cohort Oral Health Study and identified five (of 44) species associated with HIV infection.

Keywords: Bayesian variable selection; Markov chain Monte Carlo; Microbiome sequencing data; Multivariate analysis; Zero-inflated models.

1. INTRODUCTION

The human microbiome plays a critical role in maintaining health and causing both acute and chronic disease. Microbes live in communities in which multiple species establish synergistic and antagonistic relationships (Pflughoeft and Versalovic, 2012). These interactions allow some species to thrive and keep others in check. The complex biological dependencies among taxa demand statistical methods that account for and exploit this interdependence. There are valid and powerful methods for jointly analyzing microbiome sequence data as predictors of health outcomes, but there are fewer methodological options for analyzing microbiome community data as a set of joint endpoints. We specifically address three challenges that commonly arise in analysis of microbiome sequencing data as responses (dependent variables): excess zeros, interdependence of the endpoints, and the need for outcome-specific covariate selection.

First, in most human microbiome studies, a large proportion of microbial taxa is absent in the majority of subjects, leading to many more zero counts for each taxon than expected on the basis of a Poisson, negative binomial, or Dirichlet-multinomial distribution (e.g., see [Supplementary Material A](#) available at *Biostatistics* online.) (Chen and Li, 2013). Application of a conventional linear model that uses untransformed or logarithmic-transformed counts is inappropriate for zero-inflated data (Loeys and others, 2012; Xu and others, 2015). An intuitive approach to analyzing zero-inflated count data is to view the data as arising from an underlying zero-inflated distribution, which is a mixture of a point mass at zero and a count distribution, such as Poisson (Lambert, 1992).

Second, as mentioned above, microbiome sequencing data are typically multivariate (joint response) count data sampled from communities of interdependent species. Naïve application of a univariate, taxon-by-taxon approach implicitly assumes that counts of each taxon are uncorrelated. Although one could control for the type I error in this approach, this generally results in loss of power (Breiman and Friedman, 1997; La Rosa and others, 2012). Multivariate non-parametric methods that compare bacterial community composition between two groups are available (Mantel, 1967; Mantel and Valand, 1970; Anderson, 2001); these are generally less powerful than regression methods and often do not quantify the magnitude of group differences. One approach for joint modeling of multivariate microbial sequence count data is Dirichlet-multinomial regression (Holmes and others, 2012; La Rosa and others, 2012; Chen and Li, 2013; Wadsworth and others, 2017). However, the Dirichlet-multinomial model imposes restrictions that may misrepresent features of multivariate taxa count data distributions. For example, despite that relationships among microbial species can be either positively or negatively correlated, the dependence between Dirichlet variates are always negatively correlated (Aitchison and Ho, 1989; Li, 2015).

Multivariate zero-inflated regression models can address both excess zeros as well as interdependent responses. Such methods that have been developed to date have been scaled to model only a small number of interdependent count endpoints, which include bivariate (Arab and others, 2012; Fox, 2013) and trivariate (Li and others, 1999) zero-inflated Poisson models. In some cases, existing methods have incorporated a restrictive covariance structure among outcomes, which may not always be appropriate. Specific examples of such restrictions include zero-inflated models for longitudinal data only with variance

components (with no covariance components) (Lee and others, 2006; Hall, 2000; Long and others, 2015), models with dependence structures specific to spatial-temporal data (Earnest and others, 2007; Fernandes and others, 2009; Wang and others, 2015a), and models including latent factors that can induce only positive correlations among outcomes (Neelon and Chung, 2017).

A third impediment to developing and applying a multivariate analysis technique to microbiome data is that due to having more than one endpoint, there is a large number of potential covariate-endpoint associations to be modeled. It is well recognized that variable selection helps improve prediction accuracy and reduce the cost of measurement and storage of future data. The need for variable selection techniques is well appreciated for high-dimensional covariate data and may be less well known in the context of multiple outcomes. Although there exist variable selection methods for multivariate normal (MVN) (Brown and others, 1998; Lee and others, 2017) or multinomial responses (Wadsworth and others, 2017), we know of no such technique applied to methods for multivariate zero-inflated outcomes.

We have developed multivariate zero-inflated regression models by relaxing requirements regarding the covariance structure and incorporating a Bayesian variable selection approach. The proposed methods can be used to identify zero-inflated count outcomes associated with a set of covariates while accounting for the covariance structure of the outcomes. Since it is implausible that all outcomes are relevant to the same subset of covariates, we enable the proposed model to perform outcome-specific variable selection, i.e., to identify exposures or treatments associated with particular outcomes, in this case microbial taxa. Spike-and-slab approaches have been widely used for Bayesian variable selection (George and McCulloch, 1993, 1997), including for multivariate linear regression problems (Brown and others, 1998; Lee and others, 2017). In this work, we extend the spike-and-slab approach to the context of multivariate zero-inflated data.

We use the newly developed model to analyze data from the Pediatric HIV/AIDS Cohort Oral Health Study (PHACS). PHACS is an ongoing prospective cohort study at 15 US clinical sites, designed to assess the health effects of HIV infection and antiretroviral therapy on youth perinatally infected with HIV (PHIV) compared with exposed but uninfected (PHEU) youth (Alperen and others, 2014; Tassiopoulos and others, 2016; Starr and others, 2018). The data analyzed were from a cross-sectional study focused on oral health and the oral microbiome (Moscicki and others, 2016; Ryder and others, 2017). All participants were exposed to HIV perinatally, the period when they became HIV infected if at all. Emerging from the womb, it is likely that they began acquiring their oral microbiota at birth and should have had oral microbiota similar to those of adults by 3 years of age (Mueller and others, 2015; Perez-Muñoz and others, 2017). Thus, if there is a causal association, the oral bacterial sequences we measured at 10–22 years of age more likely resulted from perinatal HIV infection rather than the reverse. This is why we treat taxa's counts as endpoints and HIV as the exposure. The goals of this analysis are (i) to identify taxa associated with HIV infection; and (ii) to estimate and test the association of HIV infection with counts of the identified taxa.

The remainder of this article is organized as follows. Section 2 describes the proposed Bayesian framework, including model formulation and specification of prior distributions. Section 3 describes results from simulation studies conducted to compare the operating performance of the proposed variable selection approach versus an existing univariate method. Section 4 describes results from the PHACS data analysis. In Section 5, we further discuss the method and results.

2. METHODS

In this section, we describe the proposed multivariate zero-inflated model, present prior distributions for model parameters and the variable selection strategy, discuss interpretation of the regression parameters, and summarize the computational scheme and implementation (see [Supplementary Material B](#) available at *Biostatistics* online. for a summary of model parameters and C for implementation details).

2.1. Model formulations

Suppose that count outcomes Y_{ij} are observed for taxon $j=1, \dots, q$ and subject $i=1, \dots, n$. We use an approach that assumes that Y_{ij} follow a multivariate zero-inflated Poisson (MZIP) distribution, which is a mixture of a Poisson distribution and point mass distribution at zero ($\mathcal{I}_0$):

$$\begin{aligned} Y_{ij} &\sim \text{Poisson}(\lambda_{ij}), & \text{if } U_{ij} = 1, \\ &\sim \mathcal{I}_0, & \text{if } U_{ij} = 0, \end{aligned} \quad (2.1)$$

where U_{ij} is an unobservable indicator for the excess zeros for taxon j in subject i , and λ_{ij} is the mean of the Poisson distribution. The model implies that some zeros occur through a Poisson process whereas others represent the impossibility for a given taxon to be observed in some subjects. In practice, regression analysis based on the MZIP model proceeds by placing structure on λ_{ij} and $P[U_{ij} = 1]$, specifically as a function of the covariates and random effects. Toward this, let $\mathbf{x}_i$ and $\mathbf{z}_i$ be a vector of p_x and p_z covariates for the i^{th} subject that will be considered in the model for λ_{ij} and $P[U_{ij} = 1]$, respectively. With this formulation it is not necessary for the presence and the count of a taxon to depend on the same set of covariates.

For the count (Poisson) model part, we consider the following general modeling specification:

$$\log(\lambda_{ij}) = \beta_{0j} + \mathbf{x}_i^\top \boldsymbol{\beta}_j + \log(\xi_i) + V_{ij}, \quad (2.2)$$

where $\boldsymbol{\beta}_0 = (\beta_{01}, \dots, \beta_{0q})^\top$ are the outcome-specific intercepts and $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jp_x})^\top$ are the outcome-specific vectors of fixed-effect regression parameters. The random effects $\mathbf{V}_i = (V_{i1}, \dots, V_{iq})^\top$ characterize the unobserved characteristics that are associated with the mean count for taxon j in subject i and account for within-subject correlations. The term $\log(\xi_i)$ is included as an offset variable for settings in which one is interested in the incidence density λ_{ij}/ξ_i . For application to genetic sequence counts, setting ξ_i to the total number of sequencing reads accounts for individual variation in sequencing depth.

To account for the dependency structure in the binary part of the model, we adopt a multivariate probit model (Ashford and Sowden, 1970). Letting $U_{ij} = I(w_{ij} \geq 0)$, with indicator function $I(\cdot)$, we consider a MVN distribution for the latent variable $\mathbf{w}_i = (w_{i1}, \dots, w_{iq})^\top$, with location vector $\boldsymbol{\alpha}_0 + A^\top \mathbf{z}_i$ and variance-covariance matrix R . Here, $\boldsymbol{\alpha}_0 = (\alpha_{01}, \dots, \alpha_{0q})^\top$ are the intercepts and A is the $p_z \times q$ coefficient matrix whose columns are $\boldsymbol{\alpha}_j = (\alpha_{j1}, \dots, \alpha_{jp_z})^\top$, $j=1, \dots, q$. Then the probability density function of U_{ij} is given by

$$\mathbf{w}_i | \boldsymbol{\alpha}_0, A, R, \mathbf{z}_i \sim \text{MVN}_q(\boldsymbol{\alpha}_0 + A^\top \mathbf{z}_i, R). \quad (2.3)$$

R is restricted to be a correlation matrix to ensure identifiability of all model parameters (Chib and Greenberg, 1998; Liu, 2001). R measures the dependence between U_{ij} and $U_{ij'}$ by using the correlations among the elements of the vector $\mathbf{w}_i$. Let $\bar{\mathbf{Y}}$, $\bar{\mathbf{w}}$, $\bar{\mathbf{V}}$, and $\bar{\boldsymbol{\phi}}$ denote the collections of $\mathbf{Y}_i$, $\mathbf{w}_i$, $\mathbf{V}_i$, and $\boldsymbol{\phi}_i$, respectively, across all subjects. Let B be the $p_x \times q$ coefficient matrix whose columns $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jp_x})^\top$, $j=1, \dots, q$. Combining (2.1), (2.2), and (2.3), the augmented data likelihood function, as a function of the unknown parameters $\Theta = \{\boldsymbol{\alpha}_0, A, \boldsymbol{\beta}_0, B, \bar{\mathbf{V}}, R\}$, is:

$$\begin{aligned} L(\bar{\mathbf{Y}}, \bar{\mathbf{w}} | \Theta) &= \prod_{i=1}^n \prod_{j=1}^q \left[I(w_{ij} \geq 0) \frac{\exp \{y_{ij} \log(\lambda_{ij}) - \lambda_{ij}\}}{y_{ij}!} + I(w_{ij} < 0) I(y_{ij} = 0) \right] \\ &\times \prod_{i=1}^n (2\pi)^{-q/2} |R|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\mathbf{w}_i - \boldsymbol{\alpha}_0 - A^\top \mathbf{z}_i)^\top R^{-1} (\mathbf{w}_i - \boldsymbol{\alpha}_0 - A^\top \mathbf{z}_i) \right\}. \end{aligned} \quad (2.4)$$

2.2. Prior specification and covariance structure

We complete the Bayesian formulation of the proposed framework by specifying prior distributions for the unknown parameters. To facilitate outcome-specific variable selection, we adopt spike-and-slab priors for the regression parameters in both parts of the proposed model. Such a prior has been widely used in the context of Bayesian stochastic search variable selection (George and McCulloch, 1993, 1997). Specifically, we assign the following priors for β_j and α_j :

$$\begin{aligned}\beta_{j,k}|\gamma_{j,k}, \sigma_{\beta,k}^2 &\sim \gamma_{j,k}\text{Normal}(0, v_{\beta,j}^2 \sigma_{\beta,k}^2) + (1 - \gamma_{j,k})\mathcal{I}_0 \text{ and} \\ \alpha_{j,l}|\delta_{j,l}, \sigma_{\alpha,l}^2 &\sim \delta_{j,l}\text{Normal}(0, v_{\alpha,j}^2 \sigma_{\alpha,l}^2) + (1 - \delta_{j,l})\mathcal{I}_0,\end{aligned}\quad (2.5)$$

where $\gamma_k = (\gamma_{1,k}, \dots, \gamma_{q,k})^\top$ and $\delta_l = (\delta_{1,l}, \dots, \delta_{q,l})^\top$, for $k = 1, \dots, p_x$, $l = 1, \dots, p_z$, are vectors of binary latent variables indicating the membership of each regression parameter to one of the mixture components in (2.5). $v_{\beta,j}^2$ and $v_{\alpha,j}^2$ are the hyperparameters to specify. The k th covariate is considered to be associated with mean counts for the j -th outcome if the data support $\gamma_{j,k} = 1$ over $\gamma_{j,k} = 0$. A similar interpretation holds for $\delta_{j,k}$. We use independent Bernoulli hyperpriors for $\gamma_{j,k}$ and $\delta_{j,l}$ with inclusion probability $\omega_{\beta,k}$ and $\omega_{\alpha,l}$, respectively.

As outlined in Section 2.1, the dependence among mean bacterial counts of multiple taxa is accounted for by the distribution of random effects V_i . Specifically, we structure this dependence through a hierarchical Poisson-logNormal model (Aitchison and Ho, 1989; Fox, 2013), which corresponds to a MVN prior, $V_i|\Sigma_V \sim \text{MVN}(\mathbf{0}, \Sigma_V)$. We use an unstructured covariance pattern for Σ_V and R , thus imposing no specific structure for the dependence among outcomes. Under the unstructured model, we adopt a conjugate hyperprior, inverse-Wishart(Ψ_0, ρ_0), for Σ_V . For R , we use $\pi(R) \propto |R|^{-(q+1)/2}$, which is the prior for a correlation matrix based on Jeffreys' prior distribution for the variance-covariance matrix (Box and Tiao, 2011). In addition, we assign $\alpha_0 \sim \text{MVN}(\mu_{\alpha_0}, \sigma_{\alpha_0}^2 R)$, where μ_{α_0} is the hyperparameter to be specified. We discuss the desirable properties of the prior distributions for R and α_0 in Supplementary Material C available at *Biostatistics* online (Chib and Greenberg, 1998; Liu and Sun, 2000). We assign $\text{MVN}(\mu_{\beta_0}, \sigma_{\beta_0}^2 I_q)$ for the intercepts β_0 , where μ_{β_0} is the hyperparameter to be specified, and I_q is the $q \times q$ identity matrix. Lastly, we use the conjugate hyperpriors inverse-Gamma($a_{\beta,k}, b_{\beta,k}$), inverse-Gamma($a_{\alpha,l}, b_{\alpha,l}$), inverse-Gamma(a_{β_0}, b_{β_0}) and inverse-Gamma($a_{\alpha_0}, b_{\alpha_0}$), for $\sigma_{\beta,k}^2, \sigma_{\alpha,l}^2, \sigma_{\beta_0}^2$, and $\sigma_{\alpha_0}^2$, respectively.

2.3. Induced marginal incident density ratio

Because the MZIP model is a mixture model, the regression parameters, $\beta_{j,k}$, have latent interpretations: $\beta_{j,k}$ represents the change in log mean count of taxon j associated with a one-unit increase in the covariate k , in a susceptible sub-population (Preisser and others, 2012). The relationship between $E[Y_{ij}|V_{ij}]$ and the parameters from the proposed MZIP model is given by

$$E[Y_{ij}|V_{ij}] = \lambda_{ij}P(U_{ij} = 1) = \exp(\beta_{0j} + \mathbf{x}_i^\top \beta_j + \log(\xi_i) + V_{ij}) \Phi(\alpha_{0j} + \mathbf{z}_i^\top \alpha_j),$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. For models with log-links and normally distributed random effects such as the one we propose, it is straightforward to marginalize the conditional expectation over the random effects distribution (Long and others, 2015). Under a model with $\mathbf{x}_i = \mathbf{z}_i$, the ratio of means for a one-unit increase in covariate k is given by

$$\text{IDR}(\mathbf{x}_{i,(-k)}) = \frac{E[Y_{ij}|x_{ik} = x + 1, \mathbf{x}_{i,(-k)}]}{E[Y_{ij}|x_{ik} = x, \mathbf{x}_{i,(-k)}]} = \frac{E[E[Y_{ij}|x_{ik} = x + 1, \mathbf{x}_{i,(-k)}, V_{ij}]]}{E[E[Y_{ij}|x_{ik} = x, \mathbf{x}_{i,(-k)}, V_{ij}]]}$$

$$= e^{\beta_{jk}} \frac{\Phi(\alpha_{0j} + \mathbf{x}_{i,(-k)}^\top \boldsymbol{\alpha}_{j,(-k)} + (x+1)\alpha_{j,k})}{\Phi(\alpha_{0j} + \mathbf{x}_{i,(-k)}^\top \boldsymbol{\alpha}_{j,(-k)} + x\alpha_{j,k})}, \quad (2.6)$$

where $\mathbf{x}_{i,(-k)}$ and $\boldsymbol{\alpha}_{j,(-k)}$ are the vectors $\mathbf{x}_i$ and $\boldsymbol{\alpha}_j$, respectively, with the k -th element removed. Although we obtain $\text{IDR}(\mathbf{x}_{i,(-k)})$ by marginalizing $E[Y_{ij}|V_{ij}, U_{ij} = 1]$ over the latent mixture distribution and the distribution of random effect V_{ij} , the quantity still depends on the values of other covariates $\mathbf{x}_{i,(-k)}$, which can be addressed several ways (Preisser and others, 2012). For continuous covariates, one could obtain a covariate-adjusted $\text{IDR}(\bar{\mathbf{x}}_{(-k)})$ by either (i) inserting mean values of covariates $\mathbf{x}_{i,(-k)}$, or (ii) assuming specific covariate distributions and marginalizing the quantity over these distributions. For discrete covariates, one could (iii) empirically marginalize the $\text{IDR}(\mathbf{x}_{i,(-k)})$ over the observed distribution of covariates, or (iv) present multiple different values for the $\text{IDR}(\mathbf{x}_{i,(-k)})$, one for each category defined by unique covariate profiles.

As has been well described, estimating the variance of measures such as the $\text{IDR}(\mathbf{x}_{i,(-k)})$ in a frequentist framework would require an extra statistical technique such as bootstrap resampling (Albert and others, 2014). An advantage of the Bayesian paradigm is that estimation of uncertainty for $\text{IDR}(\mathbf{x}_{i,(-k)})$ follows directly from the variance of its posterior distribution, estimated by evaluating its expression at each scan of the Markov chain Monte Carlo (MCMC) scheme.

2.4. Markov chain Monte Carlo

We perform estimation and inference for the proposed model by using a Gibbs sampling algorithm to generate samples from the posterior distribution. In the corresponding MCMC scheme, parameters are updated either by exploiting conjugacies inherent in the model structure or by using a Metropolis–Hastings step. However, MCMC is far from straightforward because the joint posterior distribution under the proposed framework involves (i) the unobserved multivariate latent variables, $\mathbf{w}_i$; (ii) the augmented data likelihood function based upon the latent mixture distribution; (iii) spike-and-slab mixture priors for the regression parameters; (iv) a high-dimensional parameter space due to the unstructured pattern for Σ_V and R ; and (v) restrictions on the correlation parameters in R . Therefore, we develop an efficient MCMC sampling scheme based on a data augmentation algorithm (Tanner and Wong, 1987) in which the computational challenge of high-dimensionality is avoided by iterating between an “imputation step,” in which values of the unobserved latent variables $\mathbf{w}_i$ are imputed and updated, and a “posterior step,” in which the model parameters are updated. In the posterior step, we used a parameter expanded data augmentation method (Liu, 2001) to update R for computational efficiency (see [Supplementary Material C](#) available at *Biostatistics* online for detailed description of the proposed computation scheme). We have developed a series of core functions in C to improve the computation speed, for which we provide the algorithm in the `mBvs` package for R (<https://cran.r-project.org/web/packages/mBvs>). As an example of computational time, it takes approximately 3 min to generate 10 000 MCMC scans on a 2.5 GHz Intel Core i7 MacBook Pro for the analysis of the PHACS data ($n=254$, $q=44$, $p_x=p_z=2$).

3. SIMULATION STUDIES

We evaluated the performance of the proposed method on simulated data. We generated data sets under six scenarios with varying outcome correlation structures and association patterns between a covariate and the vector of outcomes in the two model parts.

3.1. Set-up and data generation

We generated samples of size $n = 300$ with $q = 20$ outcomes and $p_x = p_z = 1$ covariate under the proposed model given in (2.1), (2.2), and (2.3). The covariate was generated from $\text{Normal}(0, 2)$ and the intercepts set to $\boldsymbol{\alpha}_0 = \mathbf{1}_q$ and $\boldsymbol{\beta}_0 = 5 \cdot \mathbf{1}_q$. In Scenarios I–III and VI, we varied the scale and sign of the

association between the covariate and outcomes in the two model parts by setting $B = A = [0.05, 0.10, 0.15, 0.20, 0.25, -0.05, -0.10, -0.15, -0.20, -0.25, \mathbf{0}_{10}^\top]$. In Scenario IV, the covariate was associated with outcomes in only one of the two model parts: $B = [0.05, 0.10, 0.15, 0.20, 0.25, \mathbf{0}_{15}^\top]$ and $A = [\mathbf{0}_5, 0.05, 0.10, 0.15, 0.20, 0.25, \mathbf{0}_{10}^\top]$. We considered the *null* case in Scenario V by setting all elements of B and A to zero. We set each variance-covariance matrix Σ_V in the count part of the model and R in the binary part of the model to a correlation matrix with an exchangeable structure with correlation c_1 within the block of the first ten outcomes, an exchangeable structure with correlation c_2 within the block of the second 10 outcomes, and a common cross-block correlation of c_3 for pairs of outcomes from different blocks. In Scenarios I, IV, and V, the outcomes associated with the covariate were highly correlated and outcomes unassociated with the covariate only moderately correlated, $(c_1, c_2, c_3) = (0.70, 0.30, 0.20)$. In Scenario II, the outcomes associated with the covariate were moderately correlated and the remaining outcomes weakly correlated, $(c_1, c_2, c_3) = (0.40, 0.05, 0.10)$. In Scenario III, the outcomes associated with the covariate were weakly correlated and those unassociated with the covariate highly correlated, $(c_1, c_2, c_3) = (0.20, 0.70, 0.30)$. In Scenario VI, each outcome is assumed to follow a univariate zero-inflated Poisson (UZIP) distribution, indicating $(c_1, c_2, c_3) = (0, 0, 0)$ (independence), and $\text{diag}(\Sigma_V) = \text{diag}(R) = \mathbf{0}$ (no overdispersion).

3.2. Analyses and specification of hyperparameters

We fit the proposed MZIP model to 600 simulated data sets, 100 under each of the six scenarios. For comparison purposes in each data set, we also implemented UZIP regression with and without the lasso penalty by using the `psc1` (Zeileis and others, 2008) and `mpath` (Wang and others, 2015b) R packages, respectively. As outlined in Section 2.2, the proposed Bayesian framework requires the specification of several hyperparameters. For the intercepts β_0 and α_0 and their variance components $\sigma_{\alpha_0}^2$ and $\sigma_{\beta_0}^2$, we assigned non-informative priors by setting $(\mu_{\beta_0}, a_{\beta_0}, b_{\beta_0}) = (\mu_{\alpha_0}, a_{\alpha_0}, b_{\alpha_0}) = (\mathbf{0}_q, 0.7, 0.7)$. For the regression parameters, β and α , and their variance components, σ_α^2 and σ_β^2 , the hyperparameters $(v_{\beta,j}, a_{\beta,k}, b_{\beta,k}) = (v_{\alpha,j}, a_{\alpha,l}, b_{\alpha,l})$, $k = 1, \dots, p_x$, $l = 1, \dots, p_z$, were set to $(10, 0.7, 0.7)$ to make the corresponding priors fairly non-informative. The hyperparameters $\omega_{\beta,k} = \omega_{\alpha,l}$ were set to either 0.1 or 0.5, implying 0.1 or 0.5 *a priori* probability, respectively, of each covariate to be selected as associated with each outcome. Finally, we set $(\Psi_0, \rho_0) = (\psi I_q, q + \psi + 1)$ with $\psi = 3$, corresponding to a prior distribution of Σ_V centered at I_q and with variance of the diagonal elements equal to 2.0. We ran each MCMC chain for 1 million iterations with the first half taken as burn-in.

For the proposed Bayesian MZIP model, we perform variable selection based on the marginal posterior distribution of variable selection indicators, $\gamma_{j,1}$ and $\delta_{j,1}$. Here, we applied a marginal posterior probability cutoff of 0.5. Between the two univariate approaches implemented, in initial simulation studies the penalized UZIP model tended to select the covariate for all outcomes in both model parts when the outcomes were correlated. For this reason, and because it would be one typical practice, for the UZIP regression analyses we performed variable selection by applying 95% confidence intervals with a false discovery rate (FDR)-controlling procedure (Benjamini and Hochberg, 1995) to account for multiple comparisons.

We assessed performance of the variable selection feature of the model by calculating four quantities based on the true positives (TP; the number of outcomes associated with the covariate and selected into the model) and false positives (FP; the number of outcomes unrelated to the covariate and mistakenly selected into the model), where τ is the number of outcomes that are truly associated with the covariate: the true positive rate, $\text{TPR} = \text{TP}/\tau$, the false positive rate, $\text{FPR} = \text{FP}/(20 - \tau)$, the positive predictive value, $\text{PPV} = \text{TP}/(\text{TP} + \text{FP})$, and the negative predictive value, $\text{NPV} = (20 - \tau - \text{FP})/(20 - \text{TP} - \text{FP})$.

3.3. Results

We focus the presentation of results in this section on the MZIP model with $\omega_{\beta,1} = \omega_{\alpha,1} = 0.1$. This is to demonstrate the improvement gained by the proposed multivariate approach over an analogous univariate

Table 1. Four operating characteristics* (%) and the number of outcomes selected to be associated with the covariate (q_{sel}) for the univariate zero-inflated Poisson (UZIP)[†] and the proposed multivariate zero-inflated Poisson (MZIP)[‡] models across six simulation scenarios described in Section 3.1. The number of outcomes generated (q) and the number of outcomes that are truly associated with the covariate (τ) for each scenario are presented under the corresponding scenario number

Scenario (q, τ)		UZIP				MZIP							
		Binary		Count		$(\omega_\alpha = \omega_\beta = 0.1)$				$(\omega_\alpha = \omega_\beta = 0.5)$			
						Binary		Count		Binary		Count	
Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)		
I (20, 10)	TPR	54.3	(9.7)	99.1	(2.9)	67.4	(8.4)	82.5	(5.5)	78.4	(8.1)	87.4	(5.5)
	FPR	0.1	(1.0)	86.0	(11.1)	0.3	(1.8)	0.8	(2.8)	4.6	(6.5)	3.1	(5.5)
	PPV	99.8	(1.7)	53.7	(3.4)	99.6	(2.4)	99.1	(3.0)	95.0	(6.8)	96.9	(5.3)
	NPV	68.9	(4.7)	94.4	(17.0)	75.6	(4.9)	85.2	(4.2)	81.9	(5.6)	88.7	(4.4)
	q_{sel}	5.4	(1.0)	18.5	(1.2)	6.8	(0.8)	8.3	(0.6)	8.3	(1.1)	9.1	(0.8)
II (20, 10)	TPR	53.7	(9.9)	99.3	(2.6)	60.8	(9.6)	75.2	(8.0)	73.0	(8.8)	84.1	(7.0)
	FPR	0.4	(1.9)	87.1	(10.0)	0.5	(2.1)	0.6	(2.4)	4.0	(5.8)	2.7	(5.1)
	PPV	99.5	(2.8)	53.4	(2.9)	99.4	(2.8)	99.3	(2.8)	95.3	(6.7)	97.2	(5.4)
	NPV	68.6	(4.6)	96.4	(14.2)	72.1	(4.9)	80.4	(5.3)	78.4	(5.6)	86.2	(5.5)
	q_{sel}	5.4	(1.0)	18.6	(1.1)	6.1	(1.0)	7.6	(0.8)	7.7	(1.1)	8.7	(0.8)
III (20, 10)	TPR	56.0	(10.2)	99.1	(2.8)	60.0	(10.0)	73.5	(8.3)	74.3	(9.2)	82.6	(7.9)
	FPR	0.2	(1.5)	88.7	(11.1)	0.2	(1.5)	0.1	(1.0)	3.9	(5.8)	1.8	(4.1)
	PPV	99.6	(2.5)	53.0	(3.3)	99.6	(2.4)	99.9	(1.2)	95.2	(7.1)	98.1	(4.3)
	NPV	69.8	(5.0)	92.3	(23.2)	71.7	(5.1)	79.4	(5.2)	79.2	(6.3)	85.4	(5.9)
	q_{sel}	5.6	(1.0)	18.8	(1.2)	6.0	(1.0)	7.4	(0.8)	7.8	(0.9)	8.4	(1.0)
IV (20, 5)	TPR	56.8	(16.3)	98.8	(4.8)	63.0	(17.4)	82.6	(11.2)	79.0	(15.1)	88.6	(10.7)
	FPR	0.1	(0.7)	86.8	(8.6)	0.4	(1.6)	0.3	(1.3)	3.3	(4.6)	2.5	(4.2)
	PPV	99.8	(2.0)	27.6	(2.5)	98.5	(6.0)	99.3	(3.6)	90.3	(12.8)	93.3	(10.8)
	NPV	87.6	(4.2)	95.9	(17.1)	89.2	(4.6)	94.6	(3.4)	93.4	(4.5)	96.3	(3.4)
	q_{sel}	2.9	(0.8)	18.0	(1.3)	3.2	(0.9)	4.2	(0.6)	4.5	(1.0)	4.8	(0.8)
V [§] (20, 0)	TPR	—	—	—	—	—	—	—	—	—	—	—	—
	FPR	0.1	(0.5)	86.8	(8.5)	0.1	(0.7)	0.1	(0.7)	1.9	(3.4)	1.1	(2.3)
	PPV	—	—	—	—	—	—	—	—	—	—	—	—
	NPV	100.0	(0.0)	100.0	(0.0)	100.0	(0.0)	100.0	(0.0)	100.0	(0.0)	100.0	(0.0)
	q_{sel}	0.0	(0.1)	17.4	(1.7)	0.0	(0.1)	0.0	(0.1)	0.4	(0.7)	0.2	(0.5)
VI (20, 10)	TPR	58.1	(10.7)	100.0	(0.0)	57.7	(10.8)	100.0	(0.0)	72.2	(9.6)	100.0	(0.0)
	FPR	0.0	(0.0)	0.0	(0.0)	0.1	(1.2)	0.0	(0.0)	4.7	(6.3)	0.1	(1.1)
	PPV	100.0	(0.0)	100.0	(0.0)	99.8	(1.9)	100.0	(0.0)	94.4	(7.3)	99.9	(1.0)
	NPV	70.9	(5.5)	100.0	(0.0)	70.6	(5.4)	100.0	(0.0)	77.8	(6.1)	100.0	(0.0)
	q_{sel}	5.8	(1.1)	10.0	(0.0)	5.8	(1.1)	10.0	(0.0)	7.7	(1.1)	10.0	(0.1)

Note: Throughout the mean and standard deviation (SD) values are based on results from 100 simulated datasets.

* TPR=TP/ τ , FPR=FP/(20 - τ), PPV=TP/(TP+FP), NPV=(20 - τ - FP)/(20 - TP+FP), where TP is the number of outcomes associated with the covariate and selected into the model, FP is the number of outcomes unrelated to the covariate and mistakenly selected into the model.

[†] For variable selection in UZIP analyses, we used 95% confidence intervals with a false discovery rate controlling procedure.

[‡] For variable selection in MZIP models, we applied a marginal posterior probability cutoff of 0.5 for both γ and δ . The hyperparameters $\omega_\alpha = \omega_\beta$ are set to either 0.1 or 0.5

[§] Since there is no outcomes that are truly associated with the covariate in Scenario V, TPR and PPV are not presented.

method, while implementing the fairest comparison to the univariate method. When the values of the overall prior inclusion probabilities ($\omega_{\beta,1}$ and $\omega_{\alpha,1}$) increased from 0.1 to 0.5, the MZIP model tended to select one to two more variables on average, yielding higher TPR and NPV but also a bit higher FPR and lower PPV in both parts of the model (Table 1). When the outcome variables were uncorrelated (Scenario VI), the variable selection capability for the MZIP model with $\omega_{\beta,1}=\omega_{\alpha,1}=0.1$ was almost the same as that of UZIP model.

Across scenarios in which the outcomes were correlated (I–III), the binary part of the MZIP model was more sensitive than in the UZIP approach (Table 1), for example, with TPR = 61% versus 54% in Scenario II, respectively. The TPR in UZIP models was insensitive to the strength of correlation among outcomes, whereas the MZIP TPR increased to 67% in Scenario I, in which there was a stronger correlation among outcomes associated with the covariate. Both the UZIP and the MZIP methods successfully identified the covariate associations for the four outcomes with the largest effect sizes ($|\alpha_{j,1}| \geq 0.20$; Table 2). However, the MZIP model performed much better at detecting smaller-magnitude associations: when $|\alpha_{j,1}| = 0.10$ (outcomes 2 and 7), associations were correctly included in 40% and 20% of MZIP and UZIP models, respectively; the corresponding inclusion rates were 95% and 60% when $|\alpha_{j,1}| = 0.15$ (outcomes 3 and 8).

The multivariate approach yielded much more substantial improvement for the count part of model when outcomes were correlated (Table 1). Even controlling the FDR, the univariate approach generally exhibited inflated type I error, as high as 86% across Scenarios I–V. In contrast, across all scenarios the MZIP model had a low probability of false discovery (FPR < 1%) while also exhibiting high TPR that ranged from 73% to 83% in Scenarios I–IV. The relatively poor performance of the UZIP method is not due to bias, since the estimated association for outcomes unassociated with the covariate ($\beta_{j,1}$, $j=11, \dots, 20$) were very close to zero (Table 2). However, the medians of the asymptotic standard errors (SE) for $\hat{\beta}_{j,1}$ were 20 times smaller than the empirical standard deviations of $\hat{\beta}_{j,1}$ (rang between 0.036 and 0.054) (Table 2), was also observed in Scenarios II–V (Supplementary Material D available at *Biostatistics* online). Thus, the univariate approach appears to perform poorly in estimation of the SE for count model parameters when outcomes are correlated. Consequently, the estimated confidence intervals are too narrow.

In the *null case* (Scenario V), both approaches successfully excluded the covariate for all outcomes for the binary part of the model even when outcomes were strongly correlated; the UZIP model exhibited a high false discovery rate (87%) for the count part of the model.

We ran extensive additional simulations to explore other factors (detailed in Table E.1 and E.2 in Supplementary Material E available at *Biostatistics* online), including a larger number of outcomes ($q = 50$), a lower signal density (4~5%), a smaller sample size ($n = 150$) and negative correlations. Briefly, the results were similar to those described above, with the MZIP performing much better for variable selection and the UZIP exhibiting inflated type I error. We also compared the proposed MZIP to a univariate non-parametric method, the Wilcoxon rank sum test. Although type I error was well controlled in the univariate non-parametric method, substantially higher power was achieved by the MZIP.

To summarize, compared with the univariate approach, the proposed multivariate method improved upon the UZIP's performance for the binary part of the model by maintaining type I error while boosting the ability to identify true associations under the simulated settings. For the count part of the model, there were some scenarios in which the power of UZIP was higher than with the MZIP approach. This higher power of UZIP was at a cost of a highly inflated false discovery rate, whereas the MZIP FPR was < 1%. Performance of the MZIP model was enhanced by increasing the prior inclusion probabilities, $\omega_{\beta,k} = \omega_{\alpha,l} = 0.5$. The TPR then exceeded 80% in all non-null scenarios and FPR remained < 4% across all scenarios.

Table 2. *Estimated regression parameters and inclusion probabilities for the univariate zero-inflated Poisson (UZIP)[†], * and the proposed multivariate zero-inflated Poisson (MZIP)[‡] models for Scenario I*

<i>j</i>	True $\alpha_{j,1}, \beta_{j,1}$	UZIP				MZIP [†]			
		Binary		Count		Binary		Count	
		$\alpha_{j,1}$ Est (SE)	$m_{\alpha,j}$	$\beta_{j,1}$ Est (SE)	$m_{\beta,j}$	$\alpha_{j,1} \delta_{j,1}=1$ PM (SD)	$\delta_{j,1}$ PM	$\beta_{j,1} \gamma_{j,1}=1$ PM (SD)	$\gamma_{j,1}$ PM
1	0.05	0.05 (0.04)	0.02	0.042 (0.002)	0.95	0.06 (0.04)	0.04	0.03 (0.02)	0.03
2	0.10	0.11 (0.04)	0.14	0.088 (0.002)	1.00	0.11 (0.04)	0.43	0.09 (0.02)	1.00
3	0.15	0.15 (0.05)	0.63	0.142 (0.002)	1.00	0.15 (0.04)	0.95	0.15 (0.02)	1.00
4	0.20	0.20 (0.05)	0.91	0.191 (0.002)	1.00	0.20 (0.04)	1.00	0.19 (0.02)	1.00
5	0.25	0.25 (0.05)	0.99	0.233 (0.002)	1.00	0.24 (0.04)	1.00	0.24 (0.02)	1.00
6	-0.05	-0.04 (0.04)	0.01	-0.046 (0.002)	0.97	-0.06 (0.04)	0.04	-0.06 (0.02)	0.23
7	-0.10	-0.11 (0.04)	0.25	-0.104 (0.002)	0.99	-0.11 (0.04)	0.36	-0.11 (0.02)	1.00
8	-0.15	-0.14 (0.05)	0.55	-0.143 (0.002)	1.00	-0.15 (0.04)	0.95	-0.15 (0.02)	1.00
9	-0.20	-0.20 (0.05)	0.95	-0.194 (0.002)	1.00	-0.20 (0.04)	1.00	-0.20 (0.02)	1.00
10	-0.25	-0.25 (0.05)	0.99	-0.251 (0.002)	1.00	-0.24 (0.04)	1.00	-0.25 (0.02)	1.00
11	0.00	-0.01 (0.04)	0.00	0.001 (0.002)	0.82	-0.01 (0.04)	0.01	0.00 (0.01)	0.00
12	0.00	0.02 (0.04)	0.00	0.001 (0.002)	0.84	0.02 (0.04)	0.02	0.00 (0.01)	0.00
13	0.00	-0.01 (0.04)	0.00	-0.004 (0.002)	0.89	-0.00 (0.04)	0.01	-0.00 (0.01)	0.00
14	0.00	-0.00 (0.04)	0.00	-0.002 (0.002)	0.82	-0.00 (0.04)	0.01	-0.00 (0.01)	0.00
15	0.00	0.00 (0.04)	0.00	-0.000 (0.002)	0.93	-0.00 (0.04)	0.01	-0.00 (0.01)	0.01
16	0.00	-0.00 (0.04)	0.00	-0.004 (0.002)	0.80	-0.00 (0.04)	0.01	-0.00 (0.01)	0.00
17	0.00	0.00 (0.04)	0.00	-0.005 (0.002)	0.85	0.00 (0.04)	0.01	-0.00 (0.01)	0.00
18	0.00	0.01 (0.04)	0.00	0.007 (0.002)	0.96	0.00 (0.04)	0.01	-0.00 (0.01)	0.00
19	0.00	0.00 (0.04)	0.01	-0.003 (0.002)	0.85	-0.00 (0.04)	0.01	0.00 (0.01)	0.00
20	0.00	-0.01 (0.04)	0.00	0.002 (0.002)	0.83	-0.00 (0.04)	0.01	-0.00 (0.01)	0.00

Note: Throughout values are based on results from 100 simulated datasets.

* The medians of the maximum likelihood estimate (Est) and SE of $\alpha_{j,1}$ and $\beta_{j,1}$, the proportion that the covariate is selected for the j^{th} outcome ($m_{\alpha,j}, m_{\beta,j}$) are calculated.

[†] The empirical standard deviations of Est($\beta_{j,1}$) range between 0.036 and 0.054. (not presented in the table)

[‡] The medians of the posterior means (PM) and posterior standard deviation (SD) of $\alpha_{j,1}$ and $\beta_{j,1}$ (conditioning on $\delta_{j,1} = 1$ and $\gamma_{j,1} = 1$, respectively), the medians of the posterior means of $\delta_{j,1}$ and $\gamma_{j,1}$ (marginal posterior probability of inclusion) are computed. The hyper-parameters $\omega_{\alpha} = \omega_{\beta}$ are set to 0.1.

4. Application to pediatric HIV/AIDS cohort study data

The proposed MZIP method was originally motivated by research into whether caries-associated bacteria differ in PHIV and PHEU youth (Moscicki and others, 2016; Ryder and others, 2017). The 254 subjects were age 10–22 years at the time of an oral health examination done from September 2012 to January 2014. Subgingival dental plaque samples were collected at two preselected sites and excluded if participants had antibiotic exposure in the previous 3 months. DNA was isolated from plaque specimens and 16S rDNA sequenced (Caporaso and others, 2011; Gomes and others, 2015). The sequencing reads were trimmed, filtered, and grouped using the DADA2 pipeline, and reads matched to the curated Human Oral Microbiome Database (99.9% of reads matched to the species or genus level) (Dewhurst and others, 2010). Each subject had a count (number of sequencing reads) for each taxon identified in the study.

4.1. Analysis details and prior specifications

We focused our analysis on $q = 44$ taxa: 14 known caries-associated species (Aas and others, 2008) and any additional species that were highly correlated with them in this dataset (Figure 1a). HIV infection status ($x_{i1}=z_{i1}$; 0, uninfected; 1, infected) was the covariate of primary interest, with adjustment for participants' age ($x_{i2}=z_{i2}$) without performing variable selection on it (i.e., it was “forced” into the model). To account for sequencing depth variation across samples, the total number of sequencing reads was included as an offset in the count model.

We fit the proposed MZIP model and the UZIP model to PHACS data. For the Bayesian MZIP approach, we set the hyperparameters, $(\mu_{\beta_0}, \mu_{\alpha_0}, a_{\beta_0}, b_{\beta_0}, a_{\alpha_0}, b_{\alpha_0}, v_{\beta,j}, v_{\alpha,j}, a_{\beta,k}, b_{\beta,k}, a_{\alpha,l}, b_{\alpha,l}, \Psi_0, \rho_0)$, to the same values as in Section 3. Since performance of the MZIP was improved when the prior inclusion probabilities were increased from 0.1 to 0.5 in the simulation studies, we set $\omega_{\beta,k} = \omega_{\alpha,l} = 0.5$. We ran two independent MCMC chains for 2 million iterations, each with the first half taken as burn-in. We assessed convergence of the MCMC sampler by plotting traces of the MCMC scans for each parameter. A visual assessment of convergence to the stationary distribution was carried out by overlaying plots for the two MCMC chains.

We also calculated the posterior median with 95% credible intervals for the marginal IDR for MZIP (described in Section 2.3). We age-adjusted the estimate by using the mean value, 16 years.

4.2. Results

With a marginal posterior probability cutoff of 0.5, the MZIP method identified three species (*Fusobacterium periodonticum*, *Lachnoanaerobaculum orale*, and *Veillonella* genus NOI) for which counts were associated with HIV infection among “susceptible” individuals, and it selected two species (*Actinomyces graevenitzi* and *Leptotrichia* sp oral taxon 215) for which the probability of participants being susceptible to these two species was associated with HIV infection (Figure 2). In contrast, the frequentist UZIP analysis with 95% confidence intervals identified 39 species whose levels were associated with HIV infection in the susceptible population, including the three associations identified by MZIP. The UZIP analysis identified no species for which the probability of being “susceptible” to that species was associated with HIV infection. Based on the methods' relative performance in the simulation study, it is plausible that the UZIP's lack of accounting for outcomes' correlation patterns, which are complex (Figure 1a), grossly inflated the type I error rate in the count model and may also have decreased sensitivity of the zero model.

Indeed, comparing the estimated associations between HIV infection and the five taxa selected based on the MZIP model (Table 3), the uncertainty associated with the count part of UZIP model was much smaller than that generated by the MZIP method. As with the simulation, this may have accounted for the detection of 39 associations, many of which are presumably false positive associations. Because of how we calculated the IDR estimate, it has an age-specific interpretation. For example, 16-year-old youth, PHIV youth had 34% (95% credible interval 2%, 47%) and 48% (95% credible interval 14%, 76%) lower abundance of *F. periodonticum* and *L. orale* species, respectively, compared with PHEU youth.

The proposed MZIP model captures within-subject dependence among multiple outcomes via two correlation components, R and Σ_V . The dependence patterns arising from the empirical correlations appear to reflect, strongly, the correlation structure predicted by the binary component of the proposed MZIP (Figure 1). This implies that in these data, the presence of taxa is more structured than their counts. This result is not attributable to smoothing of the empirical correlation from having added one to every count, because the results were not sensitive to changes of this value to 0.5, 0.1, and 0.01. The model also provides an opportunity to quantify and compare the contribution of zero inflation versus other sources of overdispersion to microbiome taxon abundance (see [Supplementary Material F](#) available at *Biostatistics* online for posterior estimates of $\Sigma_{Vj,j}$ and $\alpha_{0,j}$).

Fig. 1. Observed and estimated correlations among counts of 44 microbial species in the Pediatrics HIV/AIDS Cohort Study (PHACS): (a) empirical correlations calculated by $\text{cor}(\log(y + 1))$; (b) Posterior median of R ; (c) Posterior median of R_v , calculated based on posterior samples of Σ_v from the proposed multivariate zero-inflated Poisson (MZIP) model fit.

Fig. 2. Analysis of Pediatric HIV/AIDS Cohort Study (PHACS) data: marginal posterior inclusion probabilities for the HIV status covariate in relation to excess zeros and counts of 44 microbial species as estimated via the proposed multivariate zero-inflated Poisson (MZIP) model. We adjust for participants age, a potential confounder, but do not perform variable selection on it. “NOI” stands for “not otherwise identified.”

Table 3. *Estimated regression parameters and inclusion probabilities for the five species identified as associated with HIV infection by using a multivariate zero-inflated Poisson (MZIP)* and the univariate zero-inflated Poisson (UZIP)† models*

	Binary		Count		IDR‡
	$\alpha_{j,1} \delta_{j,1} = 1$ PM (95% CI)	IP	$\beta_{j,1} \gamma_{j,1} = 1$ PM (95% CI)	IP	PM (95% CI)
<i>Actinomyces graevenitzi</i>	-0.39 (-0.67 to -0.09)	0.54	0.25 (-0.22 to 0.72)	0.11	0.96 (0.57–1.08)
<i>Fusobacterium periodonticum</i>	0.16 (-0.60 to 0.48)	0.07	-0.43 (-0.66 to -0.14)	0.93	0.66 (0.53–0.98)
MZIP <i>Lachnoanaerobaculum orale</i>	0.10 (-0.17 to 0.33)	0.03	-0.67 (-1.15 to -0.25)	0.95	0.52 (0.24–0.86)
<i>Leptotrichia</i> sp oral taxon 215	-0.44 (-0.75 to -0.17)	0.63	-0.04 (-0.19 to 0.20)	0.02	0.86 (0.74–1.00)
<i>Veillonella</i> genus NOI§	-0.04 (-0.60 to 0.45)	0.05	-0.37 (-0.60 to -0.12)	0.88	0.71 (0.57–1.00)
	$\alpha_{j,l}$ Est (95% CI _F)	$\beta_{j,k}$ Est (95% CI _F)			
<i>Actinomyces graevenitzi</i>	-0.27 (-0.89 to 0.35)	-0.46 (-0.69 to -0.23)			
<i>Fusobacterium periodonticum</i>	-0.08 (-0.93 to 0.77)	-0.80 (-0.84 to -0.75)			
UZIP <i>Lachnoanaerobaculum orale</i>	0.31 (-0.38 to 0.99)	-2.58 (-2.79 to -2.36)			
<i>Leptotrichia</i> sp oral taxon 215	-0.41 (-1.04 to 0.21)	-0.26 (-0.41 to -0.11)			
<i>Veillonella</i> genus NOI	0.20 (-1.02 to 1.41)	-0.24 (-0.27 to -0.21)			

* In the MZIP model, the posterior median (PM) and 95% credible interval (CI) of regression parameters and incidence density ratio (IDR) are computed.

† In the UZIP model, the maximum likelihood estimates (Est) and 95% confidence intervals (CI_F) of regression parameters are computed.

‡ Adjusted for individuals with age of 16 years.

§ Not otherwise identified.

Note: Inclusion probability (IP) greater than 0.5 is highlighted in bold.

5. DISCUSSION

We have described the development of a new Bayesian variable selection method that addresses challenges arising in the analysis of microbiome sequencing data: excess zero counts and high-dimensional outcomes with a complex association structure. Applying the proposed multivariate approach led to the identification of two species for which the probability of being susceptible to those species was associated with HIV infection; these associations did not meet FDR thresholds when the existing univariate approach was applied. In addition, based on the estimated induced marginalized IDR under the proposed model, another two species were less abundant in HIV-infected youth aged 16 years compared with PHEU youth of the same age.

One might question how realistic are these analyses when they are adjusted for only one confounder, age. Some reassurance might be provided from the observation that in univariate analyses that included

additional confounders (e.g., sex, race, and dental visit in previous year as a marker of oral hygiene), inference was not greatly altered compared with models including only HIV status and age. We are continuing to study performance in a range of datasets, including more complete confounder adjustment. We are also working to scale up the proposed method in the number of endpoints, as discussed further below, and also in the number of covariates, both of which are required for integrated omics analyses.

The simulation study demonstrated superior performance of the proposed MZIP approach over the existing UZIP method when outcomes were correlated. The sample size was small enough that asymptotic assumptions under the frequentist-based UZIP model did not hold. This affected estimation of the asymptotic variance of the regression parameters for the count model, which was not a limitation for the proposed multivariate Bayesian approach. This difference in performance is primarily because (i) for small data settings, estimation is generally more stable with Bayesian approaches, which exploit information from both the observed data likelihood and prior distribution; and (ii) the MZIP method uses information not only on the mean model but also from the structure of covariance among outcomes.

We used a multivariate probit model for the binary part of the MZIP mixture model. An alternative is to assume a multivariate logistic distribution for $\mathbf{w}_i$ (O'Brien and Dunson, 2004), for which posterior computation can be facilitated based on a data augmentation algorithm (Albert and Chib, 1993). However, initial numerical studies using the latter approach resulted in prohibitively slow mixing of the MCMC algorithms due to sparseness of data, even for data with $q = 10$ and assuming an unstructured covariance pattern. This is because the multivariate logistic model specification requires the estimation of n more latent parameters than does the multivariate probit model. Thus, the multivariate probit model in the MZIP proved to be much more computationally tractable.

We have presented the model in its most general form that allows the importance of each covariate, as well as the correlation structure among the multivariate outcomes, to vary across the binary and the count components of the model. This gives the user maximal flexibility and provides evidence on how a covariate is associated with each response, i.e., with more zeros or higher counts. The question arises whether the complexity of the model is necessary or whether simpler models should suffice. Analysis of the motivating data suggests that different correlation structures were needed in this case. It is difficult to provide a general answer to this question until we have had more opportunity to apply it a range of datasets and compare results with those obtained in simpler models. One would not be able to make this comparison if the most general model is not available as a basis for comparison. Yet, simpler models might well be useful in other datasets.

Thus, the software implementation of the proposed approach offers more parsimonious versions of the model, simplified by imposing additional restrictions regarding the model parameters. For example, the two model parts can be forced to have one common variable selection indicator by setting $\gamma_{j,k} = \delta_{j,l}$. In practice, such a restriction might facilitate implementation by providing a single vector of variable selection indicators, i.e., one list as to which species are associated with each covariate. A different assumption is that both model parts share the same covariance pattern ($R = R_V$), which will greatly reduce the number of parameters to be estimated and thus the computational complexity in the MCMC algorithm, especially for data with large q . In our initial analysis, fitting this restricted model to the PHACS data yielded unreliable estimates of R , because the assumption that $R = R_V$ is violated in these data (Figure 1). Again, the restricted model may serve well in other datasets. Therefore, we made available the algorithms to implement both types of simpler MZIP models as options in the `mBVS` package.

There are several ways the proposed framework could be extended. First, marginalized zero-inflated models have recently been developed so that inference can be made on the marginal mean of the sampled population via a set of unified regression coefficients (Long and others, 2015; Tabb and others, 2016). The unified regression coefficients have better interpretability, as the marginalized models do not require the additional steps described in Section 2.3 to address the dependence of parameter values on $\mathbf{x}_{i,(-k)}$. In some applications, it may be more appropriate to interpret the two sets of regression coefficients separately, yet

there also may be other applications for which interpretability of regression parameters would be enhanced by adopting a marginalized model within the proposed multivariate Bayesian variable selection method. Marginalization may also provide more stable model fitting. Second, although we focused data analysis on a preselected subset of species in the application, often microbiologists' goal is to perform whole-community oral microbiome analysis, which generally involves several hundred taxa. We are currently working to address the computational issues arising from an even higher-dimensional parameter space. Because the complexity mainly results from the flexible unstructured covariance model, we propose to scale up the proposed MZIP method by adopting alternative correlation structures that can flexibly accommodate potentially complicated patterns among hundreds of taxa. A final possibility is to study the interaction between the marginal posterior probability cutoff and the prior inclusion probability in controlling the FDR at the desired level under the proposed model.

In conclusion, the proposed framework gives researchers valid and powerful statistical tools to overcome major methodological barriers in microbiome sequencing data analysis. Beyond the study of the human microbiome, the methods, software, and guidance from simulation studies in this work will be useful in any field requiring analysis of multivariate zero-inflated count data.

6. SOFTWARE

R-package `mBvs` contains codes to implement proposed Bayesian framework described in the article. The package is currently available in CRAN (<https://cran.r-project.org/web/packages/mBvs>).

7. SUPPLEMENTARY MATERIAL

Supplementary material is available online at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

The authors thank the associate editor and four anonymous reviewers for their constructive comments and suggestions. Portions of this research were conducted on the O2 High Performance Compute Cluster, supported by the Research Computing Group, at Harvard Medical School. *Conflict of Interest*: None declared.

FUNDING

National Institutes of Health (P01CA134294, P30ES000002, U01HD052104, U01HD052102, R21DE026872, and R03DE027486).

REFERENCES

- AAS, J. A., GRIFFEN, A. L., DARDIS, S. R., LEE, A. M., OLSEN, I., DEWHIRST, F. E., LEYS, E. J. AND PASTER, B. J. (2008). Bacteria of dental caries in primary and permanent teeth in children and young adults. *Journal of Clinical Microbiology* **46**, 1407–1417.
- AITCHISON, J. AND HO, C.H. (1989). The multivariate Poisson-log normal distribution. *Biometrika* **76**, 643–653.
- ALBERT, J. AND CHIB, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* **88**, 669–679.
- ALBERT, J., WANG, W. AND NELSON, S. (2014). Estimating overall exposure effects for zero-inflated regression models with application to dental caries. *Statistical Methods in Medical Research* **23**, 257–278.

- ALPEREN, J., BRUMMEL, S., TASSIOPOULOS, K., MELLINS, C., KACANEK, D., SMITH, R., SEAGE, G. AND MOSCICKI, A. (2014). Prevalence of and risk factors for substance use among perinatally human immunodeficiency virus-infected and perinatally exposed but uninfected youth. *Journal of Adolescent Health* **54**, 341–349.
- ANDERSON, M. J. (2001). A new method for non-parametric multivariate analysis of variance. *Austral Ecology* **26**, 32–46.
- ARAB, A., HOLAN, S., WIKLE, C. AND WILDHABER, M. (2012). Semiparametric bivariate zero-inflated Poisson models with application to studies of abundance for multiple species. *Environmetrics* **23**, 183–196.
- ASHFORD, J. AND SOWDEN, R. (1970). Multi-variate probit analysis. *Biometrics* **26**, 535–546.
- BENJAMINI, Y. AND HOCHBERG, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B* **57**, 289–300.
- BOX, G. AND TIAO, G. (2011). *Bayesian Inference in Statistical Analysis*. Hoboken, NJ, USA: John Wiley & Sons.
- BREIMAN, L. AND FRIEDMAN, J. (1997). Predicting multivariate responses in multiple linear regression. *Journal of the Royal Statistical Society: Series B* **59**, 3–54.
- BROWN, P., VANNUCCI, M. AND FEARN, T. (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society: Series B* **60**, 627–641.
- CAPORASO, J., LAUBER, C., WALTERS, W., BERG-LYONS, D., LOZUPONE, C., TURNBAUGH, P., FIERER, N. AND KNIGHT, R. (2011). Global patterns of 16S rRNA diversity at a depth of millions of sequences per sample. *Proceedings of the National Academy of Sciences of the United States of America* **108**, 4516–4522.
- CHEN, J. AND LI, H. (2013). Variable selection for sparse Dirichlet-multinomial regression with an application to microbiome data analysis. *The Annals of Applied Statistics* **7**, 418–442.
- CHIB, S. AND GREENBERG, E. (1998). Analysis of multivariate probit models. *Biometrika* **85**, 347–361.
- DEWHIRST, F. E., CHEN, T., IZARD, J., PASTER, B. J., TANNER, A.C.R., YU, W.-H., LAKSHMANAN, A. AND WADE, W. G. (2010). The human oral microbiome. *Journal of Bacteriology* **192**, 5002–5017.
- EARNEST, A., MORGAN, G., MENGENSEN, K., RYAN, L., SUMMERHAYES, R. AND BEARD, J. (2007). Evaluating the effect of neighbourhood weight matrices on smoothing properties of conditional autoregressive (car) models. *International Journal of Health Geographics* **6**, 54.
- FERNANDES, M. V. M., SCHMIDT, A. M. AND MIGON, H. S. (2009). Modelling zero-inflated spatio-temporal processes. *Statistical Modelling* **9**, 3–25.
- FOX, J. (2013). Multivariate zero-inflated modeling with latent predictors: modeling feedback behavior. *Computational Statistics & Data Analysis* **68**, 361–374.
- GEORGE, E. I. AND MCCULLOCH, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association* **88**, 881–889.
- GEORGE, E.I. AND MCCULLOCH, R.E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica* **7**, 339–374.
- GOMES, B., BERBER, V., KOKARAS, A., CHEN, T. AND PASTER, B. (2015). Microbiomes of endodontic-periodontal lesions before and after chemomechanical preparation. *Journal of Endodontics* **41**, 1975–1984.
- HALL, D. B. (2000). Zero-inflated Poisson and binomial regression with random effects: a case study. *Biometrics* **56**, 1030–1039.
- HOLMES, I., HARRIS, K. AND QUINCE, C. (2012). Dirichlet multinomial mixtures: generative models for microbial metagenomics. *PLoS One* **7**, e30126.
- LA ROSA, P. S., BROOKS, J. P., DEYCH, E., BOONE, E. L., EDWARDS, D. J., WANG, Q., SODERGREN, E., WEINSTOCK, G. AND SHANNON, W. D. (2012). Hypothesis testing and power calculations for taxonomic-based human microbiome data. *PLoS One* **7**, e52078.

- LAMBERT, D. (1992). Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics* **34**, 1–14.
- LEE, A. H., WANG, K., SCOTT, J. A., YAU, K. K. W. AND MCLACHLAN, G. J. (2006). Multi-level zero-inflated Poisson regression modelling of correlated count data with excess zeros. *Statistical Methods in Medical Research* **15**, 47–61.
- LEE, K. H., TADESSE, M. G., BACCARELLI, A. A., SCHWARTZ, J. AND COULL, B. A. (2017). Multivariate Bayesian variable selection exploiting dependence structure among outcomes: application to air pollution effects on DNA methylation. *Biometrics* **73**, 232–241.
- LI, C., LU, J., PARK, J., KIM, K., BRINKLEY, P. A. AND PETERSON, J. P. (1999). Multivariate zero-inflated Poisson models and their applications. *Technometrics* **41**, 29–38.
- LI, H. (2015). Microbiome, metagenomics, and high-dimensional compositional data analysis. *Annual Review of Statistics and Its Application* **2**, 73–94.
- LIU, C. (2001). Bayesian analysis of multivariate probit models—discussion on the art of data augmentation. *Journal of Computational and Graphical Statistics* **10**, 75–81.
- LIU, C. AND SUN, D. X. (2000). Analysis of interval-censored data from fractionated experiments using covariance adjustment. *Technometrics* **42**, 353–365.
- LOEYS, T., MOERKERKE, B., DE SMET, O. AND BUYSSE, A. (2012). The analysis of zero-inflated count data: Beyond zero-inflated Poisson regression. *British Journal of Mathematical and Statistical Psychology* **65**, 163–180.
- LONG, D. L., PREISSER, J. S., HERRING, A. H. AND GOLIN, C. E. (2015). A marginalized zero-inflated Poisson regression model with random effects. *Journal of the Royal Statistical Society: Series C* **64**, 815–830.
- MANTEL, N. (1967). The detection of disease clustering and a generalized regression approach. *Cancer Research* **27** (2 Part 1), 209–220.
- MANTEL, N. AND VALAND, R. S. (1970). A technique of nonparametric multivariate analysis. *Biometrics* **26**, 547–558.
- MOSCICKI, A.-B., YAO, T.-J., RYDER, M. I., RUSSELL, J. S., DOMINY, S. S., PATEL, K., MCKENNA, M., VAN DYKE, R. B., SEAGE III, G. R., HAZRA, R. and others. (2016). The burden of oral disease among perinatally HIV-infected and HIV-exposed uninfected youth. *PLoS One* **11**, e0156459.
- MUELLER, N. T., BAKACS, E., COMBELICK, J., GRIGORYAN, Z. AND DOMINGUEZ-BELLO, M. G. (2015). The infant microbiome development: mom matters. *Trends in Molecular Medicine* **21**, 109–117.
- NEELON, B. AND CHUNG, D. (2017). The LZIP: a Bayesian latent factor model for correlated zero-inflated counts. *Biometrics* **73**, 185–196.
- O'BRIEN, S. M. AND DUNSON, D. B. (2004). Bayesian multivariate logistic regression. *Biometrics* **60**, 739–746.
- PEREZ-MUÑOZ, M. E., ARRIETA, MARIE-CLAIRE, RAMER-TAIT, AMANDA E. AND WALTER, JENS. (2017). A critical assessment of the “sterile womb” and “in utero colonization” hypotheses: implications for research on the pioneer infant microbiome. *Microbiome* **5**, 48.
- PFLUGHOEFT, K. J. AND VERSALOVIC, J. (2012). Human microbiome in health and disease. *Annual Review of Pathology: Mechanisms of Disease* **7**, 99–122.
- PREISSER, J. S., STAMM, J. W., LONG, D. L. AND KINCADE, M. E. (2012). Review and recommendations for zero-inflated count regression modeling of dental caries indices in epidemiological studies. *Caries Research* **46**, 413–423.
- RYDER, M. I., YAO, T.-J., RUSSELL, J. S., MOSCICKI, A.-B. AND SHIBOSKI, C. H. (2017). Prevalence of periodontal diseases in a multicenter cohort of perinatally HIV-infected and HIV-exposed and uninfected youth. *Journal of Clinical Periodontology* **44**, 2–12.

- STARR, J. R., HUANG, Y., LEE, K. H., MURPHY, C. M., MOSCICKI, A.-B., SHIBOSKI, C. H., RYDER, M. I., YAO, T.-J., FALLER, L. L., VAN DYKE, R. B. *and others.* (2018). Oral microbiota in youth with perinatally acquired HIV infection. *Microbiome* **6**, 100.
- TABB, L. P., TCHETGEN, E. J., WELLENIUS, G. A. AND COULL, B. A. (2016). Marginalized zero-altered models for longitudinal count data. *Statistics in Biosciences* **8**, 181–203.
- TANNER, M. A. AND WONG, W. H. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* **82**, 528–540.
- TASSIOPOULOS, K., PATEL, K., ALPEREN, J., KACANEK, D., ELLIS, A., BERMAN, C., ALLISON, S. M., HAZRA, R., BARR, E., CANTOS, K. *and others.* (2016). Following young people with perinatal HIV infection from adolescence into adulthood: the protocol for PHACS AMP Up, a prospective cohort study. *BMJ Open* **6**, e011396.
- WADSWORTH, W. D., ARGIENTO, R., GUINDANI, M., GALLOWAY-PENA, J., SHELBURNE, S. A. AND VANNUCCI, M. (2017). An integrative Bayesian Dirichlet-multinomial regression model for the analysis of taxonomic abundances in microbiome data. *BMC Bioinformatics* **18**, 94.
- WANG, X., CHEN, M., KUO, R. C. AND DEY, D. K. (2015a). Bayesian spatial-temporal modeling of ecological zero-inflated count data. *Statistica Sinica* **25**, 189–204.
- WANG, Z., MA, S. AND WANG, C. (2015b). Variable selection for zero-inflated and overdispersed data with application to health care demand in Germany. *Biometrical Journal* **57**, 867–884.
- XU, L., PATERSON, A. D., TURPIN, W. AND XU, W. (2015). Assessment and selection of competing models for zero-inflated microbiome data. *PLoS One* **10**, e0129606.
- ZEILEIS, A., KLEIBER, C. AND JACKMAN, S. (2008). Regression models for count data in R. *Journal of Statistical Software* **27**, 1–25.

[Received October 27, 2017; revised September 6, 2018; accepted for publication September 9, 2018]