A Bayesian zero-inflated negative binomial regression model for the integrative analysis of microbiome data

SHUANG JIANG

Department of Statistical Science, Southern Methodist University, Dallas, TX 75275, USA

GUANGHUA XIAO

Quantitative Biomedical Research Center, University of Texas Southwestern Medical Center, Dallas, TX 75390, USA

ANDREW Y. KOH

*Department of Pediatrics, University of Texas Southwestern Medical Center, Dallas, TX 75390, USA and
Department of Microbiology, University of Texas Southwestern Medical Center, Dallas, TX 75390, USA*

JIWOONG KIM

Quantitative Biomedical Research Center, University of Texas Southwestern Medical Center, Dallas, TX, 75390, USA

QIWEI LI[†]

Department of Mathematical Sciences, The University of Texas at Dallas, Richardson, TX 75080, USA
qiwei.li@utdallas.edu

XIAOWEI ZHAN^{*,†}

Quantitative Biomedical Research Center, University of Texas Southwestern Medical Center, Dallas, TX 75390, USA
xiaowei.zhan@utsouthwestern.edu

SUMMARY

Microbiome omics approaches can reveal intriguing relationships between the human microbiome and certain disease states. Along with identification of specific bacteria taxa associated with diseases, recent scientific advancements provide mounting evidence that metabolism, genetics, and environmental factors can all modulate these microbial effects. However, the current methods for integrating microbiome data and other covariates are severely lacking. Hence, we present an integrative Bayesian zero-inflated negative binomial regression model that can both distinguish differentially abundant taxa with distinct phenotypes and quantify covariate-taxa effects. Our model demonstrates good performance using simulated data. Furthermore, we successfully integrated microbiome taxonomies and metabolomics in two real microbiome datasets to provide biologically interpretable findings. In all, we proposed a novel integrative Bayesian regression model that features bacterial differential abundance analysis and microbiome-covariate effects quantifications, which makes it suitable for general microbiome studies.

*To whom correspondence should be addressed.

[†]Qiwei Li and Xiaowei Zhan contributed equally to this work. Both are considered as the corresponding authors.

Keywords: Bayesian regression; Count data; Feature selection; Integrative analysis; Microbiome;

1. INTRODUCTION

The human microbiome is estimated to contain 3.0×10^{13} bacteria (Sender *and others*, 2016) and 3.3×10^6 microbial genes (Qin *and others*, 2010). Microbial communities have a profound impact on human health (Ursell *and others*, 2012). Recently, microbiome studies have identified disease-associated bacteria taxa in type 2 diabetes (Karlsson *and others*, 2013), liver cirrhosis (Qin *and others*, 2014), inflammatory bowel disease (Halfvarson *and others*, 2017), and melanoma patients responsive to cancer immunotherapy (Frankel *and others*, 2017). An increasing number of research projects continue to systematically investigate the role of the microbiome in human diseases (Integrative, 2014).

While innovations in next-generation sequencing technology continue to shape the next steps in the microbiome field, the statistical methods used in microbiome research have not kept pace. For instance, metagenomic shotgun sequencing (MSS) generates a massive amount of sequence reads that can provide species or isolate level taxonomic resolution (Segata *and others*, 2012). The subsequent comparative statistical analysis assesses whether specific species are associated with a phenotypic state or experimental condition.

Upon surveying commonly used statistical approaches, one method focuses on the comparison of multi-taxa (Chen *and others*, 2012; Kelly *and others*, 2015; Zhao *and others*, 2015; Wu *and others*, 2016), frequently termed the microbiome community. However, those approaches do not aim to identify differentially abundant species—making clinical interpretation, mechanistic insights, and biological validations difficult. Another approach interrogates each individual bacteria taxa for different groups or conditions. For example, La Rosa *and others* (2015) utilizes a Wilcoxon rank-sum test or Kruskal–Wallis test for groupwise comparisons on microbiome compositional data. Recently, methods developed for RNA-seq data have been adapted to microbiome studies, e.g. the negative-binomial regression model in DESeq2 (Love *and others*, 2014) and overdispersed Poisson model in edgeR (Robinson *and others*, 2010). These methods, however, are not optimized for microbiome datasets.

Microbial abundance can be affected by covariates, such as metabolites, antibiotics, and host genetics. These confounding variables need to be adjusted for more accurate differential abundance analysis. Ultimately, there may be a clinical need to quantify the associations between microbiome and clinical confounders (Kinross *and others*, 2011; Maier *and others*, 2018; Zhu *and others*, 2018). One common approach is to calculate pairwise correlations between all taxa and covariates (Li *and others*, 2008), but this method may be significantly underpowered. Other model-based methods (Chen and Li, 2013; Wadsworth *and others*, 2017) have been proposed to detect covariate-taxa associations, but the taxon-outcome associations have been ignored. Recently, Li *and others* (2018) developed a multivariate zero-inflated logistic normal model to quantify the associations between microbiome abundances and multiple factors (e.g. disease risk factors or health outcomes) based on microbiome compositional data instead of the count data.

Here, we propose a Bayesian integrative model to analyze microbiome count data. Our model jointly identifies differentially abundant taxa among multiple groups and simultaneously quantifies the taxon-covariate associations. Our modeling construction includes several advantages. First, it characterizes the over-dispersion and zero-inflation frequently observed in microbiome count data by introducing a zero-inflated negative binomial (ZINB) model. Second, it models the heterogeneity from different sequencing depths, covariate effects, and group effects via a log-linear regression framework on the ZINB mean components. Last, we propose two feature selection processes to simultaneously detect differentially abundant taxa and estimate the covariate-taxa associations using the *spike-and-slab* priors. We compute Bayesian posterior probabilities for these correlated features and provide the Bayesian false discovery rate (FDR). Extensive and thorough use of simulated data demonstrates that our model largely improved performance when compared with existing methods. We present two applications of real microbiome

datasets with various covariate sets. Biological interpretations of our results confirm those of previous studies and offer a more comprehensive understanding of the underlying mechanism in disease etiology.

The article is organized as follows. In Section 2, we introduce the integrative hierarchical mixture model and the prior formulations. Section 3 supplies a brief discussion of the Markov chain Monte Carlo (MCMC) algorithm and the resulting posterior inference. In Section 4, we evaluate model performance on simulated data through a comparison study. We investigate the covariate association in Section 5. Two real-data analyses are shown in Section 6. Our conclusions are presented in Section 7.

2. HIERARCHICAL MODEL

Our model starts with a high-dimensional count matrix where each entry represents the count of sequence reads belonging to a taxonomy such as bacterial species. Specifically, we denote $\mathbf{Y}_{n \times p}$ (usually $n \ll p$) be a microbial abundance matrix, with $y_{ij} \in \mathbb{N}, i = 1, \dots, n, j = 1, \dots, p$ representing the observed count of the i -th sample and j -th taxon out of the total n samples and p taxa (features). Note that the proposed model can also be applied to an operational taxonomic unit (OTU) count table obtained via 16S metagenomic approaches. For an OTU table, each feature would be a taxonomic unit of a bacteria species or genus depending on the sequence similarity threshold (e.g. 97%). We also denote a covariate matrix $\mathbf{X}_{n \times R}$ where each entry x_{ir} represents the measurement of the r -th covariate on the i -th sample. The graphical formulation of the proposed model is summarized in [Figures S1 and S2 in the supplementary material](#) available at *Biostatistics* online.

2.1. Count generating process

In practice, the microbial abundance matrix $\mathbf{Y}$ is characterized by an inflated amount of zeros, resulting from insufficient sampling depth. Meanwhile, the abundance matrix usually consists of extremely large counts. Based on these two facts, we assume that each count is sampled from a ZINB distribution so as to simultaneously account for both zero-inflation and over-dispersion presented in $\mathbf{Y}$:

$$y_{ij} | \pi, \lambda_{ij}, \phi_j \sim \pi I(y_{ij} = 0) + (1 - \pi) \text{NB}(y_{ij}; \lambda_{ij}, \phi_j), \quad (2.1)$$

where $\pi \in [0, 1]$ represents the weight of generating extra zeros, $I(\cdot)$ is an indicator function, and $\text{NB}(y; \lambda, \phi)$ denotes a negative binomial distribution for random variable y with the expectation λ and dispersion $1/\phi$. Under this parameterization, the variance of y is $\lambda + \lambda^2/\phi$. A small value of ϕ allows modeling of extra variation. Note that increasing ϕ towards infinity yields a Poisson distribution with both expectation and variance equal to λ . We assume a Gamma prior $\text{Ga}(a_\phi, b_\phi)$ for the dispersion parameter ϕ .

An equivalent way to model this count generating process is to introduce a latent binary variable r_{ij} , such that

$$y_{ij} | r_{ij}, \lambda_{ij}, \phi_j \begin{cases} \sim \text{NB}(\lambda_{ij}, \phi_j) & \text{if } r_{ij} = 0 \\ = 0 & \text{if } r_{ij} = 1 \end{cases}, \quad (2.2)$$

where r_{ij} is from a Bernoulli distribution with parameter π , i.e. $\sim \text{Bernoulli}(\pi)$. We further impose $\pi \sim \text{Beta}(a_\pi, b_\pi)$, which leads to a Beta-Bernoulli prior for r_{ij} with expectation $a_\pi/(a_\pi + b_\pi)$.

2.2. Integrative modeling with feature selection

Microbiome count data are characterized by high variability in the number of reads among samples from different groups (due to distinct biological conditions), or even the same group (due to uneven sequencing

depths). To accommodate this setting, we parameterize the mean parameter λ_{ij} of the negative binomial distribution as the multiplicative effects of two positive random effects: (i) the size factor s_i reflects how the sequencing depth affects counts across all taxa observed in the i -th sample; (ii) the normalized abundance α_{ij} for the j -th taxon in the i -th sample once the sample-specific variability has been accounted for. Our goal is to find a subset of p taxa that enables us to discriminate the n samples from K distinct groups. We introduce a binary latent vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)$, with $\gamma_j = 1$ indicating that the j -th taxon has significantly differential abundances among the K groups, and $\gamma_j = 0$ otherwise. Therefore, conditional on $r_{ij} = 0$, we reparameterize the negative binomial kernel of Equation (2.2) as follows:

$$y_{ij}|r_{ij} = 0, \gamma_j, s_i, \alpha_{ijk}, \alpha_{ij0} \sim \begin{cases} \text{NB}(y_{ij}; s_i \alpha_{ij0}, \phi_j) & \text{if } \gamma_j = 0 \\ \text{NB}(y_{ij}; s_i \alpha_{ijk}, \phi_j) & \text{if } \gamma_j = 1 \text{ and } z_i = k \end{cases}. \quad (2.3)$$

Here, z_i is the sample allocation indicator. Collectively, $\mathbf{z}_{n \times 1} = (z_1, z_2, \dots, z_n)^T$ indicates the membership for each sample, where $z_i = k, k \in \{1, \dots, K\}$ reveals that the i -th sample belongs to the k -th group. s_i is the size factor of the i -th sample, which can be estimated from the data (see Section 2.3). We assume an independent Bernoulli prior $\gamma_j \sim \text{Bernoulli}(\omega)$ for each taxon j , and further impose a beta hyperprior on ω to formulate a Beta-Bernoulli prior, i.e. $\omega \sim \text{Beta}(a_\omega, b_\omega)$. The choice of a_ω and b_ω incorporates the prior belief that a certain percentage of taxa are discriminatory.

We further specify a log-link function to integrate the covariates into the modeling construction for each normalized abundance:

$$\begin{cases} \log \alpha_{ij0} &= \mu_{0j} + \mathbf{x}_i \boldsymbol{\beta}_j^T & \text{if } \gamma_j = 0 \\ \log \alpha_{ijk} &= \mu_{0j} + \mu_{kj} + \mathbf{x}_i \boldsymbol{\beta}_j^T & \text{if } \gamma_j = 1 \text{ and } z_i = k \end{cases}, \quad (2.4)$$

where μ_{0j} is a feature-specific baseline parameter for taxon j . Note that $\exp(\mu_{0j})$'s can be considered as scaling factors adjusting for feature-specific levels across all samples. The group-specific parameter μ_{kj} captures the baseline shift between the k -th group and the reference group. We set $\mu_{kj} = 0$ if the k -th group is the reference group to avoid identifiability problems arising from the sum of the components. $\mathbf{x}_i$, the i -th row of covariate matrix $\mathbf{X}$, contains all the covariate measurements for sample i . Here, $\boldsymbol{\beta}_j$ is a 1-by- R vector, with each element β_{rj} modeling the global effect of the r -th covariate on the observed counts for the j -th taxon. In practice, not all of the covariates are related to the abundance of a taxon. Therefore, we allow different sets of covariates to affect different taxa by specifying a *spike-and-slab* prior (Brown and others, 1998; Ishwaran and others, 2005) as $\beta_{rj} \sim (1 - \delta_{rj})I(\beta_{rj} = 0) + \delta_{rj}N(0, \sigma_{\beta j}^2)$, where $\delta_{rj} = 1$ indicates the r -th covariate is associated with the normalized abundance for the j -th feature, and $\delta_{rj} = 0$ otherwise. This modeling approach allows us to identify significant covariate-taxa associations, via the selection of the nonzero β_{rj} coefficients, for all discriminatory and non-discriminatory taxonomic features. We complete the model by setting $\mu_{0j} \sim N(0, \sigma_{0j}^2)$, $\mu_{kj} \sim N(0, \sigma_{\mu j}^2)$, and $\delta_{rj} \sim \text{Beta-Bernoulli}(a_p, b_p)$. Letting $\sigma_{0j}^2 = 10^2$ for all j yields a vague prior for the feature-specific baseline parameter. An inverse-gamma (IG) hyperprior $\text{IG}(a, b)$ is shared by $\sigma_{\mu j}^2$ and $\sigma_{\beta j}^2$.

2.3. Size factor estimation

The parameterization of the negative binomial mean, as shown in Equation (2.3), is a product of the size factor and the normalized abundance. It is typical to normalize the size factor first to ensure model identifiability. Hence, the plug-in estimator (equivalent to a point-mass prior) of s_i is adopted to facilitate the inference based on the normalized abundance α_{ij} as shown in Equation (2.4). The plug-in estimators can be calculated from the observed count matrix $\mathbf{Y}$. There have been a number of proposals to estimate

the size factors in the context of RNA-seq data analyses. Both [Witten \(2011\)](#) and [Li and others \(2017\)](#) conducted a comprehensive literature review. However, the assumptions of many existing methods for RNA-seq are likely not appropriate for highly diverse microbial environments ([Weiss and others, 2017](#)). A so-called cumulative sum scaling (CSS) method has been developed by [Paulson and others \(2013\)](#) as $\hat{s}_i^{\text{CSS}} \propto \sum_{j=1}^p y_{ij} I(y_{ij} \leq q_i^{\text{CSS}})$, where the default value of q_i^{CSS} is 50. CSS can be viewed as an adaptive extension of [Bullard and others \(2010\)](#), and it is better suited for microbiome data. Moreover, a new normalization method named geometric mean of pairwise ratios (GMPR) has been proposed by [Chen and others \(2018\)](#), aiming to handle the zero-inflated sequencing data. GMPR calculates the size factor s_i based on the median count ratio of nonzero counts between the i -th sample and the remaining samples. It has been shown to be robust to differential and outlier OTUs. Combining this with some constraints such as $\sum_{i=1}^n \log \hat{s}_i = 0$ (i.e. $\prod_{i=1}^n \hat{s}_i = 1$), we are able to obtain a set of identifiable values. In this article, both CSS and GMPR are considered.

3. MODEL FITTING AND POSTERIOR INFERENCE

Our model space consists of $(\mathbf{R}, \boldsymbol{\phi}, \boldsymbol{\mu}_0, \mathbf{M}, \mathbf{B}, \boldsymbol{\gamma}, \boldsymbol{\Delta}, \omega, \pi)$ with the extra zero indicators $\mathbf{R} = (r_{ij}, i = 1, \dots, n, j = 1, \dots, p)$, the dispersion parameters $\boldsymbol{\phi} = (\phi_j, j = 1, \dots, p)$, the feature-specific baselines $\boldsymbol{\mu}_0 = (\mu_{0j}, j = 1, \dots, p)$, the group-specific baselines $\mathbf{M} = (\mu_{kj}, k = 1, \dots, K, j = 1, \dots, p)$, the covariate effects $\mathbf{B} = (\beta_{rj}, r = 1, \dots, R, j = 1, \dots, p)$, the discriminatory taxa indicators $\boldsymbol{\gamma} = (\gamma_j, j = 1, \dots, p)$, and the association indicators $\boldsymbol{\Delta} = (\delta_{rj}, r = 1, \dots, R, j = 1, \dots, p)$. We explore the posterior distribution via a MCMC algorithm based on stochastic search variable selection with within-model updates ([Savitsky and Vannucci, 2010](#)). Full details can be found in the [supplementary material](#) available at *Biostatistics* online.

We are interested in distinguishing taxa that are differentially abundant among different groups, via $\boldsymbol{\gamma}$, as well as their associations with covariates, via $\boldsymbol{\Delta}$. One way to summarize the posterior distributions of these binary parameters is via the marginal posterior probability of inclusion (PPI). Suppose $t = 1, \dots, T$ index the MCMC iterations after burn-in. Then PPI of each γ_j and δ_{rj} can be written as $\text{PPI}(\gamma_j) = \frac{1}{T} \sum_{t=1}^T \gamma_j^{(t)}$ and $\text{PPI}(\delta_{rj}) = \frac{1}{T} \sum_{t=1}^T \delta_{rj}^{(t)}$, respectively. Subsequently, important features and covariates can be selected based on a given PPI threshold. Following [Newton and others \(2004\)](#), we choose a threshold that controls the Bayesian FDR. Specifically, we solve the following equations to determine the thresholds: $\text{FDR}_{\boldsymbol{\gamma}}(c_{\boldsymbol{\gamma}}) = \frac{\sum_{j=1}^p (1 - \text{PPI}(\gamma_j)) I(1 - \text{PPI}(\gamma_j) < c_{\boldsymbol{\gamma}})}{\sum_{j=1}^p I(1 - \text{PPI}(\gamma_j) < c_{\boldsymbol{\gamma}})}$, $\text{FDR}_{\boldsymbol{\Delta}}(c_{\boldsymbol{\Delta}}) = \frac{\sum_{r=1}^R \sum_{j=1}^p (1 - \text{PPI}(\delta_{rj})) I(1 - \text{PPI}(\delta_{rj}) < c_{\boldsymbol{\Delta}})}{\sum_{r=1}^R \sum_{j=1}^p I(1 - \text{PPI}(\delta_{rj}) < c_{\boldsymbol{\Delta}})}$, where $I(\cdot)$ is an indicator function. A well-accepted setting is to set both $\text{FDR}_{\boldsymbol{\gamma}}$ and $\text{FDR}_{\boldsymbol{\Delta}}$ equal to 0.05, which corresponds to an expected FDR of 5%.

4. SIMULATION STUDY

In this section, we evaluated the proposed model using simulated data. In particular, we considered two methods (CSS and GMPR) introduced in Section 2.3 for estimating the size factor s_i 's. We also compared our model with other existing methods described in the prior microbiome studies. In order to mimic metagenome sequencing data from real data applications (Section 6), we chose the parameters as follows: we set n samples for $K = 2$ groups with balanced group size $n_1 = n_2 = n/2$. We chose a large number of candidate features by setting the number of taxa $p = 300$, and randomly selected 20 true discriminant features to evaluate our model performance. Each row of $\mathbf{Y}$, denoted as $\mathbf{y}_i$, was generated from a Dirichlet-multinomial distribution as described in [Wadsworth and others \(2017\)](#). For $i = 1, \dots, n$, we let $\mathbf{y}_i \sim \text{Multinomial}(N_i, \boldsymbol{\pi}_i)$ with the row sum $N_i \sim \text{Discrete Uniform}(2 \times 10^7, 6 \times 10^7)$ and $\boldsymbol{\pi}_i = (\pi_{i1}, \dots, \pi_{ip}) \sim \text{Dirichlet}(\mathbf{a}_i)$. We further incorporated the feature and covariate effects through $\mathbf{a}_i = (a_{i1}, \dots, a_{ip})$ by setting $a_{ij} = \exp(a_{ij}^*)$ with $a_{ij}^* \sim \text{Normal}(\mu_{0j} + \mu_{kj} + \mathbf{x}_i \boldsymbol{\beta}_j^T, \sigma_e^2)$. Here, a larger value

of σ_e^2 corresponded to a higher noise level. Compared with Equation (2.3), this data generating process is different from the assumption of the proposed model. We set $\mu_{0j} \sim \text{Uniform}(8, 10)$, $\mu_{1j} = 0$ for all j and $\mu_{2j} = \pm 2$ for all selected discriminating features and 0 otherwise. Then for the covariate effects, we first obtained the covariate matrix $X_{n \times R}$ by sampling each row $\mathbf{x}_i$ from the covariate matrix of the liver cirrhosis study in Section 6.1 (with $n = 237$ and $R = 7$). In particular, we sampled $n/2$ covariate records from healthy and disease groups, respectively. For each taxon j , we then randomly selected $m \in \{0, 2, 4, 6\}$ out of R covariates and let the corresponding $\beta_{rj} \sim \pm \text{Uniform}(0.5, 1)$ while setting the rest $\beta_{rj} = 0$. Lastly, we randomly set $\pi_0 n p$ counts in $\mathbf{Y}$ to be zeros to mimic the zero-inflation in the real data. To summarize, we varied the following settings in order to comprehensively examine the model performance: (i) sample size per group $n/2 = 10$ or 30 ; (ii) noise level $\sigma_e = 0.5, 1.0$, or 1.5 ; 3) zero proportion $\pi_0 = 30\%, 40\%$, or 70% . In the main text, we present the results obtained from the simulated datasets that $n/2 = 30$, $\sigma_e^2 = 1$, and $\pi_0 = 40\%$, and the remaining results can be found in [Section S2 in the supplementary material](#) available at *Biostatistics* online.

The hyperparameters were specified using the following default settings. For the binary variables with Beta-Bernoulli priors $\gamma_j \sim \text{Beta-Bernoulli}(a_\omega, b_\omega)$, $\delta_{rj} \sim \text{Beta-Bernoulli}(a_p, b_p)$ and $r_{ij} \sim \text{Beta-Bernoulli}(a_\pi, b_\pi)$, we set $a_\omega = 0.2$, $b_\omega = 1.8$, $a_p = 0.4$, and $b_p = 0.6$, which means that 10% of the taxa are expected to be discriminant features, and 20% of the covariate coefficients to be nonzero. We chose $a_\pi = b_\pi = 1$ assuming that about half of the zeros are truly missing. For the dispersion parameter with $\text{Ga}(a_\phi, b_\phi)$ prior, we set $a_\phi = 1$, $b_\phi = 0.01$ to obtain a vague gamma prior with mean of 100 and variance of 10 000. Next, we specified a flat prior $\text{IG}(a = 2, b = 10)$ for the variance term $\sigma_{\mu j}^2$ and $\sigma_{\beta j}^2$. The sensitivity analysis reported in the [Section S3.2 in the supplementary material](#) available at *Biostatistics* online contains more details on the choice of a and b . When implementing our model on a dataset, we ran four independent chains with different starting points where each feature or covariate was randomly initialized to have $\gamma_j = 1$ or 0 , $\delta_{rj} = 1$ or 0 . We set 20 000 iterations as the default and discarded the first half as burn-in. To assess the concordance between four chains, we looked at all the pairwise correlation coefficients between the marginal PPI of $\boldsymbol{\gamma}$ and $\boldsymbol{\Delta}$. As mentioned by [Stingo and others \(2013\)](#), high values of correlation suggest that MCMC chains are run for a satisfactory number of iterations. After ensuring convergence, we assessed our model performance based on the averaged result over four chains.

Our goal was to identify the discriminating features (e.g. taxa) and the significant feature-covariate associations (i.e. all nonzero γ_j and δ_{rj} in our model). We thus obtained the PPI for all γ_j and δ_{rj} , and visualized the accuracy in feature selection using the receiving operating characteristic (ROC) curve. We also computed the false positive rate when all feature-covariate associations were zero. We further considered two types of competitors for model comparison. The first type, similar to the proposed model, can simultaneously identify discriminating features and detect the feature-covariate associations. Here, we compared with the multivariate zero-inflated logistic-normal (MZILN) regression model proposed by [Li and others \(2018\)](#). The MZILN model treats the sample allocation vector as an observed covariate for each sample. Therefore, we combined the group label with other observed covariates to create a new covariate matrix, and the MZILN model gave a regularized estimation of the regression coefficient between each feature and covariate. The selected discriminating features and feature-covariate associations corresponded to the nonzero coefficient estimations. The second type of method achieves the same goal in two separate stages. The first stage consists of four methods to select discriminating features based on p -values, including the Wilcoxon rank-sum test (Wilcoxon test) and three differential expression analysis methods implemented by the R packages *metagenomeSeq* ([Paulson and others, 2013](#)), *edgeR* ([Robinson and others, 2010](#)), and *limma* ([Ritchie and others, 2015](#)). Specifically, *metagenomeSeq* assumes a zero-inflated Gaussian model, *edgeR* models count data using a negative binomial distribution, and *limma* adopts a linear model for the log-transformed count data. Then, the discriminating features were selected to be those with BH ([Benjamini and Hochberg, 1995](#)) adjusted p -values smaller than 0.05. To make a head to head comparison in the first stage, we also included a simplified version of the ZINB model by excluding the

covariate term $\mathbf{x}_j\beta_j^T$ in Equation (2.4). In the second stage, we considered the following feature selection strategies for each p -value based method. They are: (i) correlation test, (ii) lasso regression, (iii) random forest, and (iv) multivariate linear regression. We centered the selected discriminating features by group, and the rest across all samples. For the correlation test, the Pearson correlation coefficients were calculated between the log scaled compositional data and the covariate measurements for each outcome group. Next, a Fisher z-transformation (Fisher, 1915) was applied to obtain the p -values for testing the significance of correlation. For lasso regression, we calculated the true positive rates and the false positive rates with respect to a range of lasso penalty. For the last two, we fitted a random forest model or a multivariate linear regression model between each feature and the covariate matrix X , which yielded variable importance measures or p -values. In all, we have four choices in the first stage {Wilcoxon test, metagenomeSeq, edgeR, limma} and four choices in the second stage {correlation test, lasso regression, random forest, multivariate linear regression}, with $4 \times 4 = 16$ choices in total. For clear visualization of the result, we excluded limma in the second stage due to its relatively inferior performance in the first stage. We also dropped random forest and linear regression since they showed similar performance as the lasso regression. Besides, all the p -values generated using different methods were adjusted using the BH method to control the FDR.

For each of the four scenarios, Figure 1 compares the model performance through the averaged ROC curve over 100 simulated datasets. We also include the area under curve (AUC) for each approach. For detecting discriminating features, the proposed method consistently shows high AUC (> 0.98) across all scenarios, and similar results for capturing the feature-covariate associations ($\text{AUC} > 0.90$). Moreover, the proposed method maintains a low FDR even when all β_{ij} are 0. The correlation-based method shows low false positive rates in the case where the true number of contributing covariate is 0 but has low power when {2, 4, 6} out of 7 covariates have nonzero contribution. In addition, the proposed model achieves the highest true positive fraction under a fixed small value of FDR in all scenarios, with the MZILN model and metagenomeSeq performing the second and third best in estimating the discriminating feature indicator γ (shown in the left column of Figure 1, Figures S3–S5 in the supplementary material available at *Biostatistics* online). We also noticed that the MZILN model could not outperform the two-stage methods in estimating the feature-covariate association indicator Δ . The above conclusions hold for using either CSS or GMPR to estimate the plug-in size factors. To test if our model is robust to the choice of size factor estimation methods, we further conducted a sensitivity analysis in Section S3.1 in the supplementary material available at *Biostatistics* online. The result, as shown in Figure S6 in the supplementary material available at *Biostatistics* online, suggests that our model is considerably robust to the choice of different normalization methods, while CSS and GMPR have a marginal performance improvement. Furthermore, we reach the same conclusion with varying group sizes, log-scale noise levels, and zero proportions. In particular, the proposed ZINB model is robust to a larger amount of extra zeros. Either decreasing the group size or increasing the noise level hampers the performance of all the methods. Nevertheless, the ZINB model still consistently outperforms the alternative approaches in estimating γ and Δ . Results are summarized in Figures S3–S5 in the supplementary material available at *Biostatistics* online.

5. FEATURE-COVARIATE ASSOCIATION ANALYSIS

In this section, we demonstrate that our model can estimate the association between a taxonomic feature and a covariate by adjusting for the remaining confounders. As a comparison, current approaches rely on correlation analysis between the pairwise microbiome and covariates. Specifically, those analyses converted each observed taxonomic count to a fraction (or termed percentages, intensities) by sample. Next the Pearson correlation coefficients were calculated between the log scaled fractions and the covariate measurements for each outcome group. Lastly, a Fisher z-transformation (Fisher, 1915) was applied to obtain the p -values for testing the significance of correlation.

Fig. 1. Averaged ROC curves for the discriminating feature indicator γ (left) and the feature-covariate association indicator Δ (right) with respect to different numbers of nonzero covariate coefficients, i.e. (a) 0, (b) 2, (c) 4, and (d) 6 out of 7, over 100 replicates in each scenario.

*The correlation-based methods showed low false positive rates in the case where there is no truly nonzero covariate coefficients.

Our model constructs a regression framework to quantify the relationship between the normalized abundance α_{ijk} and covariates through the Equation (2.4). Based on Equation (2.4), given a feature j and a covariate r of interest, we first normalized the observed abundance using CSS and performed logarithmic transformation. Next, to calculate $x_{ir}\hat{\beta}_{rj}$ and the group shift $\hat{\mu}_{kj}$, we subtracted the estimated feature-specific influence $\hat{\mu}_{0j}$ and other covariates' impact $\sum_{r' \neq r} x_{ir'}\hat{\beta}_{r'j}$ from the transformed abundance. Lastly, we could evaluate whether our model provided a reasonable estimation ($\hat{\beta}_{rj}$) of the feature-covariate association between covariate r and the normalized and adjusted observations of feature j .

Here, we demonstrated the advantages of the proposed model in estimating the feature-covariate association over the correlation-based method through simulation. For each feature, we randomly selected four out of seven covariates to have nonzero linear effects on the latent abundances, and generated a simulated dataset following the description in Section 4. We kept the same prior and algorithm settings to obtain the estimations for all parameters of interest. We chose a 5% Bayesian FDR for estimating Δ . Among all feature-covariate combinations, the proposed model achieved sensitivity and specificity rates of 82.9% and 86.7%, respectively. We randomly chose several pairs of feature and covariate and compared the proposed method and the correlation-based method. Figure 2 displays the results of 2 examples, where the true values of δ_{rj} were 1. The two dashed lines in Figure 2a or b have the same slope of $\hat{\beta}_{rj}$ as our estimated covariate effect. Both plots suggest that the proposed model is able to capture the feature-covariate relationship. Notice that we did not adjust for the group-specific effect. Hence the differences between two dashed lines represents the group-specific parameter $\hat{\mu}_{kj}$. These results illustrated the advantages in simultaneously detecting the discriminating features and quantifying the feature-covariate associations. Furthermore, we also validated that the proposed model had correctly captured the direction of covariate effects in both cases. Figure 2b and d show the results from the correlation-based model. The slope of each dashed line represents the Pearson correlation coefficient. However, there was no significant result as all p -values were greater than 0.05, suggesting that the correlation test can be underpowered. The correlation-based method failed to isolate the covariate of interest from the confounders, and it might suggest a wrong direction of covariate effect, as shown in Figure 2d.

6. REAL-DATA ANALYSIS

We applied the proposed model on two real-data sets: one with hundreds of samples and the other with only 24 samples. Compared with the analysis methods used in the original publications, our model demonstrates better performance in detecting differentially abundant bacteria. In addition, our model supports adjusting for biologically meaningful covariates. When adjusting for the metabolic pathway quantities (or metabolites through metabolomics technology) as covariates, our model estimates the association between taxa and metabolism-related functions (or metabolites).

6.1. Liver cirrhosis dataset

Cirrhosis is a late-stage condition of scarring or fibrosis of the liver caused by liver disease such as hepatitis B, hepatitis C, and non-alcoholic fatty liver disease (Abubakar and others, 2015). The liver is connected to the gastrointestinal tract via the hepatic portal and bile secretion systems. Interestingly, distinct gut microbiota signatures have been associated with both early-stage liver diseases and end-stage liver cirrhosis (Garcia-Tsao and Wiest, 2004; Yan and others, 2011; Benten and Wiest, 2012). We applied our model on a gut microbiome dataset from a liver cirrhosis study carried out by Qin and others (2014). All metagenome sequenced samples were available from the NCBI Short Read Archive and the curated microbial abundance matrix was accessible from ExperimentHub (Pasolli and others, 2017).

The full dataset includes 237 samples with their observed microbial abundance matrix Y profiled from the gut microbiome at the species taxonomic level. The study has two patient groups, including 114

Fig. 2. Feature-covariate association analysis: comparison of the results given by the proposed method ((a) and (b)) and correlation-based method ((c) and (d)) from the simulated dataset, where the two features shown (randomly selected for illustration) were truly discriminating with the covariate effect $\beta_{1,68} > 0$ and $\beta_{1,127} < 0$ by simulation. The proposed method provided a reasonable estimation ($\hat{\beta}_{ij}$) of the feature-covariate association.

healthy controls and 123 liver cirrhosis patients. We filtered out the taxa with extremely low abundance before the analysis as suggested in [Wadsworth and others \(2017\)](#). We obtained 528 taxa that had at least two observed counts in both groups for further analysis. As for the covariate information, we used MetaCyc, a collection of microbial pathways and enzymes involved in metabolism for an extensive amount of organisms ([Caspi and others, 2007](#)). We incorporated the 529 MetaCyc pathway measurements for 237 individuals in the study, and reduced the high correlation among the pathways by average linkage clustering on their correlation matrix ([Wadsworth and others, 2017](#)). Specifically, we kept the pathway with the largest fold-change between two groups in each cluster, and decided the number of clusters such that the correlations between the resulting pathways were less than 0.5. Logarithmic transformation and normalization (zero mean and unit variance) were applied to the selected covariates to ensure the zero mean and unit standard deviation. After the pre-processing step, we had seven covariates representing metabolic functions.

In [Qin and others \(2014\)](#), differential analyses were based on the Wilcoxon test and the p -values were corrected by the BH method. Although a stringent threshold of significance level (0.0001) was used, the

authors discovered 79 differentially abundant species and had to restrictively report the 30 top candidates in each group. Figure 5a is the cladogram of the discriminating taxa selected by different methods, with blue dots representing the results by *Qin and others* (2014) and red dots reported by the ZINB model. As suggested in our simulation study, these results may reflect a high FDR as covariate effects were not factored in the analysis. In addition, the Wilcoxon test cannot account for the pathway effects and thus the associations between bacteria and metabolic pathways were not identified.

We applied the proposed Bayesian ZINB model to simultaneously analyze the microbial abundance matrix of bacteria and their metabolic pathway abundance. We set a similar hyperparameter setting as discussed in Section 4 by first specifying $a_\mu = a_\beta = 2$ and $b_\mu = b_\beta = 10$. Next, we set a_ω , b_ω , a_p , b_p , a_ϕ , b_ϕ to be the same as their default values discussed in Section 4. We ran four independent Markov chains with different starting points. Each chain had 40 000 iterations with the first half discarded as burn-in. We checked the convergence visually and calculated the pairwise Pearson correlation for PPIs, which ranged from 0.988 to 0.994 for γ 's and from 0.982 to 0.989 for δ 's. These concluded highly consistent results. Figure 4a shows the PPIs for all 528 taxa, where the dashed line represents the threshold corresponding to an expected FDR of 0.05. We identified 19 differentially expressed taxa, the majority of which are more abundant in the liver cirrhosis group.

Figure 4c shows the posterior mean of μ_{2j} for all identified discriminating taxa, and Table S1 in the supplementary material available at *Biostatistics* online contains all the detailed parameter estimations for those taxa. Interestingly, two clear taxonomic branches are distinguished by our model (as indicated by red dots, Figure 5a): the genera *Veillonella* and *Streptococcus*, both of which can originate from the oral cavity. Of note, oral commensal bacteria are able to colonize the distal intestinal tract in liver cirrhosis patients (*Qin and others*, 2014), probably due to bile acid changes. *Veillonella* spp. and *Streptococcus* spp. have been identified as more abundant in patients with primary biliary cholangitis (*Tang and others*, 2018), another hepatic disorder which shares pathophysiologic features with liver cirrhosis (*Ridlon and others*, 2013). Figure 3a and b show the identified associations between microbiota and metabolic pathways. For example, L-alanine biosynthesis (PWY0-1061) is positively correlated with *Veillonella*. Alanine is a gluconeogenesis precursors in liver metabolism, and increased alanine is thought to induce pyruvate kinase in *Veillonella*. Thus, this connection between alanine synthesis and *Veillonella* is intriguing and potentially novel, and biologic validation experiments might offer further clarification.

6.2. Metastatic melanoma dataset

The proposed Bayesian ZINB model can perform integrative analysis of microbiome taxonomic data and other omics datasets. In this section, we applied this model to simultaneously analyze microbiome and metabolomics data from a study of advanced stage melanoma patients receiving immune checkpoint inhibitor therapy (ICT) (*Frankel and others*, 2017). The data were collected using MSS and unbiased shotgun metabolomics. Here, we aim at identifying unique microbiome taxonomic and metabolomic signatures in those patients who responded favorably to ICT.

A subset of patients in this study ($n = 24$) were treated with ipilimumab and nivolumab (IN), a combination therapy that has been shown to be more efficacious than therapy with anti-PD1 or anti-CTLA4 therapy alone. Sixteen patients responded to treatment and eight patients had progression. We performed quality control steps on MSS reads and profiled them using MetaPhlAn *Segata and others* (2012) as described in *Frankel and others* (2017). We filtered out taxa with at most one observation in either patient group, which left $p = 248$ taxa from species to kingdom level. For the same fecal samples, we performed metabolomics profiling and quantified 1901 patients' metabolite compounds as the covariate matrix X . We are interested in statistically assessing how the biochemical volumes between patient groups are associated with bacteria burden or quantities. We adopted the same strategy mentioned in Section 6.1 to reduce the correlation between covariates, which resulted in a 24×9 matrix as the covariate matrix of X .

Fig. 3. Real-data analysis: heatmap showing the effect from covariates, the MetaCyc pathway abundances, in two studies. (a and b) We use a liver cirrhosis dataset and show the effect between covariate effects and all microbiome, or differential abundant microbiome, respectively. (c and d) We use metastatic melanoma dataset and show the effect between covariate effects and all microbiome, or differential abundant microbiome, respectively).

In the model fitting stage, for prior specification, we used $a_p = 0.2$, $b_p = 1.8$ to obtain a sparser covariate effect due to the small sample size, and it suggested about 10% of taxa-covariate associations were significant. We kept the same default setting for the rest of the hyperparameters. Next, we ran four

Fig. 4. Real-data analysis: plots for γ PPI and credible interval. The horizontal dashed line in the PPI plot represents the threshold controlling the Bayesian false discovery rate < 0.05 . All taxa whose PPI pass the threshold are included in (c) and (d), where each horizontal bar is the 95% credible interval for μ_{2j} (group-specific parameter) with posterior mean shown in circle. Each arrow in (a) and (b) points out the taxon with largest absolute value of μ_{2j} in one patient group as shown in c and d.

independent chains with different starting points, and discarded the first half of 40 000 iterations for each chain. Although the small sample size ($n = 24$) posed challenges for parameter estimation, the results showed high pairwise Pearson correlations of PPIs for γ (ranging from 0.989 to 0.992) and δ (ranging from 0.927 to 0.953). Figure 4b shows the PPIs for all taxa, and Figure 4d illustrates the posterior means of the selected taxa. Table S2 in the supplementary material available at *Biostatistics* online includes detailed parameter estimations of the taxa in Figure 4d.

Our model jointly identified differentially abundant taxa and revealed the microbiome–metabolite associations. First, among all seven taxa identified, it is of specific interest to investigate the responder-enriched taxon *Bifidobacterium* (genus level), *Bifidobacteriaceae* (family level). *Bifidobacterium*, nesting within *Bifidobacteriaceae*, is a genus of gram-positive, nonmotile, often branched anaerobic bacteria (Schell and others, 2002). *Bifidobacteria* are one of the major genera of bacteria that make up the gastrointestinal tract microbiota in mammals. This result about *Bifidobacterium* is supported by recent melanoma studies. Sivan and others (2015) compared melanoma growth in mice harboring specific microbiota, and used sequencing of the 16S ribosomal RNA to identify *Bifidobacterium* as associated with the antitumor effects. They also found that oral administration of *Bifidobacterium* augmented ICT efficacy. Moreover, Matson and others (2018) detected significant association between several species from *Bifidobacterium* with patients’

Fig. 5. Real-data analysis: cladograms of the identified discriminating taxa (shown in dots). Red dots: taxa found by the proposed model; blue dots: taxa found by methods reported in the original studies. Each arrow in (a and b) points out the taxon with the largest absolute value of μ_{2j} (group-specific parameter) in one patient group, as shown in Figure 4c and d.

outcomes in an immunotherapy treatment study for metastatic melanoma. Both studies showed consistent direction of effect, as did our model. The responder-enriched taxon *Bifidobacterium* were estimated to negatively correlate with 2-oxoarginine and 2-hydroxypalmitate in Figure 4d. The suppression of these fatty-acid metabolites may induce better cancer treatment as they were shown to have the oncogenic signaling role in cancer cells (Louie and others, 2013).

7. CONCLUSION

In this article, we presented a Bayesian ZINB model for analysis of high-throughput sequencing microbiome data. Our method is novel in simultaneously incorporating the effect from measurable genetic covariates and identifying differentially abundant taxa for multiple patient groups in one statistical

framework. This allows for integrative analysis of microbiome data and other omics data. Our method is flexible, as it allows for identification and estimation of the association between covariates and each taxon's abundance. These results could potentially guide clinical decisions for precision shaping of the microbiome, although results would need to be validated in preclinical models first. In addition, our method is computationally efficient in posterior inferences. We implemented the MCMC algorithm to analyze the data from two MSS studies with results readily available in minutes.

In real-data analysis, the identified differentially abundant taxa by our model are often cluttered in the same phylogenetic branch. These results are achieved without imposing the phylogenetic structures in the model. This highlights that the results from our model are biologically interpretable and thus capable of guiding further biological mechanism studies. Our results on the metastatic melanoma study uncover novel relationships between taxa and metabolites which merit further experimental investigation.

The framework of our model allows for several extensions. For example, the current method supports two phenotype groups. If there are multiple groups (e.g. the intermediate phenotypes), the current model can incorporate group-specific parameters while holding the other parameters unchanged (e.g. the normalized microbiome abundance can be inferred in the same way). Then the same posterior inferences can be applied. The proposed model based on a regression framework considers the microbiome normalized abundance as the response and integrates the omics data, e.g. metabolite compounds, as predictors. Similarly, [Lloyd-Price and others \(2019\)](#) used the microbial abundance as the response and a type of genomics data (i.e. single-nucleotide polymorphism) as predictors to identify several inflammatory bowel disease associated host–microbial interactions. Both methods focus on the omics effect on microbial abundance. However, the interaction between the microbiome and the host is bidirectional. Therefore, it is worthwhile to consider using the microbial features as predictors to investigate their modulations on any biological process with quantitative omics measurements. For instance, [Richards and others \(2019\)](#) explored how microbial abundances induced changes in chromatin accessibility and transcription factor binding of host genetics. Another interesting extension would be to analyze correlated covariates such as longitudinal clinical measurements ([Zhang and others, 2017](#)).

8. SOFTWARE

An R package `IntegrativeBayes` is available on GitHub:

<https://github.com/shuangj00/IntegrativeBayes>. All the simulated datasets and two real datasets presented in Section 6 are available on figshare:

<https://figshare.com/projects/IntegrativeBayesZINB/57980>.

SUPPLEMENTARY MATERIAL

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

ACKNOWLEDGMENTS

We thank Jessie Norris for her comments on the manuscript. We thank the authors of [Li and others \(2018\)](#) for sharing the R codes to evaluate the MZILN model. We thank the editor and two reviewers for their comments that led to improvements in the manuscript. *Conflict of Interest:* Dr Andrew Koh is a consultant for Merck and Saol Therapeutics. No conflict of interest reported by the rest authors.

FUNDING

This work is supported by the following fundings: NIH 5P30CA142543, NIH 5P50CA070907, NIH 5R01AI123163, NIH 1U01CA213338-01A1, NIH R01AI123163, NIH 5R01GM126479, NIH

5R01HG008983. O'Donnell Brain Institute Pilot Award. The Roberta I. and Norman L. Pollock Fund. The Centers for Disease Control/National Center for Emerging and Zoonotic Infectious Diseases, SBIR Phase II 200-2017-95643.

REFERENCES

- ABUBAKAR, II, TILLMANN, T. AND BANERJEE, A. (2015). Global, regional, and national age-sex specific all-cause and cause-specific mortality for 240 causes of death, 1990-2013: a systematic analysis for the Global Burden of Disease study 2013. *Lancet* **385**, 117–171.
- 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 (Methodological)* **57**, 289–300.
- BENTEN, D. AND WIEST, R. (2012). Gut microbiome and intestinal barrier failure—the Achilles heel in hepatology? *Journal of Hepatology* **56**, 1221–1223.
- BROWN, P. J., VANNUCCI, M. AND FEARN, T. (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society. Series B (Methodological)* **60**, 627–641.
- BULLARD, J. H., PURDOM, E., HANSEN, K. D. AND DUDOIT, S. (2010). Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments. *BMC Bioinformatics* **11**, 94.
- CASPI, R., FOERSTER, H., FULCHER, C. A., KAIPA, P., KRUMMENACKER, M., LATENDRESSE, M., PALEY, S., RHEE, S. Y., SHEARER, A. G., TISSIER, C. and others. (2007). The MetaCyc Database of metabolic pathways and enzymes and the BioCyc collection of Pathway/Genome Databases. *Nucleic Acids Research* **36**, D623–D631.
- CHEN, J., BITTINGER, K., CHARLSON, E. S., HOFFMANN, C., LEWIS, J., WU, G. D., COLLMAN, R. G., BUSHMAN, F. D. AND LI, H. (2012). Associating microbiome composition with environmental covariates using generalized UniFrac distances. *Bioinformatics* **28**, 2106–2113.
- CHEN, J. AND LI, H. (2013). Variable selection for sparse Dirichlet-multinomial regression with an application to microbiome data analysis. *Annals of Applied Statistics* **7**, 418–442.
- CHEN, L., REEVE, J., ZHANG, L., HUANG, S., WANG, X. AND CHEN, J. (2018). GMPR: a robust normalization method for zero-inflated count data with application to microbiome sequencing data. *PeerJ* **6**, e4600.
- FISHER, R. A. (1915). Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika* **10**, 507–521.
- FRANKEL, A. E., COUGHLIN, L. A., KIM, J., FROELICH, T. W., XIE, Y., FRENKEL, E. P. AND KOH, A. Y. (2017). Metagenomic shotgun sequencing and unbiased metabolomic profiling identify specific human gut microbiota and metabolites associated with immune checkpoint therapy efficacy in melanoma patients. *Neoplasia* **19**, 848–855.
- GARCIA-TSAO, G. AND WIEST, R. (2004). Gut microflora in the pathogenesis of the complications of cirrhosis. *Best Practice & Research: Clinical Gastroenterology* **18**, 353–372.
- HALFVARSON, J., BRISLAWN, C. J., LAMENDELLA, R., VÁZQUEZ-BAEZA, Y., WALTERS, W. A., BRAMER, L. M., D'AMATO, M., BONFIGLIO, F., McDONALD, D., GONZALEZ, A. and others. (2017). Dynamics of the human gut microbiome in inflammatory bowel disease. *Nature Microbiology* **2**, 17004.
- INTEGRATIVE, HMP. (2014). The Integrative Human Microbiome Project: dynamic analysis of microbiome-host omics profiles during periods of human health and disease. *Cell Host & Microbe* **16**, 276.
- ISHWARAN, H., SUNIL RAO, J. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics* **33**, 730–773.
- KARLSSON, F. H., TREMAROLI, V., NOOKAEW, I., BERGSTRÖM, G., BEHRE, C. J., FAGERBERG, B., NIELSEN, J. AND BÄCKHED, F. (2013). Gut metagenome in European women with normal, impaired and diabetic glucose control. *Nature* **498**, 99.

- KELLY, B. J., GROSS, R., BITTINGER, K., SHERRILL-MIX, S., LEWIS, J. D., COLLMAN, R. G., BUSHMAN, F. D. AND LI, H. (2015). Power and sample-size estimation for microbiome studies using pairwise distances and PERMANOVA. *Bioinformatics* **31**, 2461–2468.
- KINROSS, J. M., DARZI, A. W. AND NICHOLSON, J. K. (2011). Gut microbiome-host interactions in health and disease. *Genome Medicine* **3**, 14.
- LA ROSA, P. S., ZHOU, Y., SODERGREN, E., WEINSTOCK, G. AND SHANNON, W. D. (2015). Hypothesis testing of metagenomic data. In: Izard, J. and Rivera, M. C. (editors), *Metagenomics for Microbiology*. Cambridge, Massachusetts: Elsevier, pp. 81–96.
- LI, M., WANG, B., ZHANG, M., RANTALAINEN, M., WANG, S., ZHOU, H., ZHANG, Y., SHEN, J., PANG, X., ZHANG, M. and others. (2008). Symbiotic gut microbes modulate human metabolic phenotypes. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 2117–2122.
- LI, Q., GUINDANI, M., REICH, B. J., BONDELL, H. D. AND VANNUCCI, M. (2017). A Bayesian mixture model for clustering and selection of feature occurrence rates under mean constraints. *Statistical Analysis and Data Mining* **10**, 393–409.
- LI, Z., LEE, K., KARAGAS, M. R., MADAN, J. C., HOEN, A. G., OMALLEY, A. J. AND LI, H. (2018). Conditional regression based on a multivariate zero-inflated logistic-normal model for microbiome relative abundance data. *Statistics in Biosciences* **10**, 587–608.
- LLOYD-PRICE, J., ARZE, C., ANANTHAKRISHNAN, A. N., SCHIRMER, M., AVILA-PACHECO, J., POON, T. W., ANDREWS, E., AJAMI, N. J., BONHAM, K. S., BRISLAWN, C. J. and others. (2019). Multi-omics of the gut microbial ecosystem in inflammatory bowel diseases. *Nature* **569**, 655.
- LOUIE, S. M., ROBERTS, L. S., MULVIHILL, M. M., LUO, K. AND NOMURA, D. K. (2013). Cancer cells incorporate and remodel exogenous palmitate into structural and oncogenic signaling lipids. *Biochimica et Biophysica Acta* **1831**, 1566–1572.
- LOVE, M. I., HUBER, W. AND ANDERS, S. (2014). Moderated estimation of fold change and dispersion for RNA-Seq data with DESeq2. *Genome Biology* **15**, 550.
- MAIER, L., PRUTEANU, M., KUHN, M., ZELLER, G., TELZEROW, A., ANDERSON, E. E., BROCHADO, A. R., FERNANDEZ, K. C., DOSE, H., MORI, H. and others. (2018). Extensive impact of non-antibiotic drugs on human gut bacteria. *Nature* **555**, 623.
- MATSON, V., FESSLER, J., BAO, R., CHONGSUWAT, T., ZHA, Y., ALEGRE, M.-L., LUKE, J. J. AND GAJEWSKI, T. F. (2018). The commensal microbiome is associated with anti-PD-1 efficacy in metastatic melanoma patients. *Science* **359**, 104–108.
- NEWTON, M. A., NOUEIRY, A., SARKAR, D. AND AHLQUIST, P. (2004). Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics* **5**, 155–176.
- PASOLLI, E., SCHIFFER, L., MANGHI, P., RENSON, A., OBENCHAIN, V., TRUONG, D. T., BEGHINI, F., MALIK, F., RAMOS, M., DOWD, J. B. and others. (2017). Accessible, curated metagenomic data through Experimenthub. *Nature Methods* **14**, 1023.
- PAULSON, J. N., STINE, O. C., BRAVO, H. C. AND POP, M. (2013). Differential abundance analysis for microbial marker-gene surveys. *Nature Methods* **10**, 1200.
- QIN, J., LI, R., RAES, J., ARUMUGAM, M., BURGDORF, K. S., MANICHANH, C., NIELSEN, T., PONS, N., LEVENEZ, F., YAMADA, T. and others. (2010). A human gut microbial gene catalogue established by metagenomic sequencing. *Nature* **464**, 59.
- QIN, N., YANG, F., LI, A., PRIFTI, E., CHEN, Y., SHAO, L., GUO, J., LE CHATELIER, E., YAO, J., WU, L. and others. (2014). Alterations of the human gut microbiome in liver cirrhosis. *Nature* **513**, 59.

- RICHARDS, A. L., MUEHLBAUER, A. L., ALAZIZI, A., BURNS, M. B., FINDLEY, A., MESSINA, F., GOULD, T. J., CASCARDO, C., PIQUE-REGI, R., BLEKHMEN, R. *and others.* (2019). Gut microbiota has a widespread and modifiable effect on host gene regulation. *mSystems* **4**, e00323-18.
- RIDLON, J. M., ALVES, J. M., HYLEMON, P. B. AND BAJAJ, J. S. (2013). Cirrhosis, bile acids and gut microbiota: unraveling a complex relationship. *Gut Microbes* **4**, 382–387.
- RITCHIE, M. E., PHIPSON, B., WU, D., HU, Y., LAW, C. W., SHI, W. AND SMYTH, G. K. (2015). limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Research* **43**, e47.
- ROBINSON, M. D., MCCARTHY, D. J. AND SMYTH, G. K. (2010). edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* **26**, 139–140.
- SAVITSKY, T. AND VANNUCCI, M. (2010). Spiked Dirichlet process priors for Gaussian process models. *Journal of Probability and Statistics* **2010**, Article ID 201489.
- SCHILL, M. A., KARMIRANTZOU, M., SNEL, B., VILANOVA, D., BERGER, B., PESSI, G., ZWAHLN, M.-C., DESIERE, F., BORK, P., DELLEY, M. *and others.* (2002). The genome sequence of *Bifidobacterium longum* reflects its adaptation to the human gastrointestinal tract. *Proceedings of the National Academy of Sciences, India. Section B* **99**, 14422–14427.
- SEGATA, N., WALDRON, L., BALLARINI, A., NARASIMHAN, V., JOUSSON, O. AND HUTTENHOWER, C. (2012). Metagenomic microbial community profiling using unique clade-specific marker genes. *Nature Methods* **9**, 811.
- SENDER, R., FUCHS, S. AND MILO, R. (2016). Revised estimates for the number of human and bacteria cells in the body. *PLoS Biology* **14**, e1002533.
- SIVAN, A., CORRALES, L., HUBERT, N., WILLIAMS, J. B., AQUINO-MICHAELS, K., EARLEY, Z. M., BENYAMIN, F. W., LEI, Y. M., JABRI, B., ALEGRE, M.-L. *and others.* (2015). Commensal *Bifidobacterium* promotes antitumor immunity and facilitates anti-PD-L1 efficacy. *Science*, 350, 1084–1089.
- Stingo, F. C., Guindani, M., Vannucci, M. and Calhoun, V. D. (2013). AN INTEGRATIVE BAYESIAN MODELING APPROACH TO IMAGING GENETICS. *Journal of the American Statistical Association* **108**, 876–891.
- TANG, R., WEI, Y., LI, Y., CHEN, W., CHEN, H., WANG, Q., YANG, F., MIAO, Q., XIAO, X., ZHANG, H. *and others.* (2018). Gut microbial profile is altered in primary biliary cholangitis and partially restored after UDCA therapy. *Gut* **67**, 534–541.
- URSELL, L. K., METCALF, J. L., PARFREY, L. W. AND KNIGHT, R. (2012). Defining the human microbiome. *Nutrition Reviews* **70**, S38–S44.
- 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.
- WEISS, S., XU, Z. Z., PEDDADA, S., AMIR, A., BITTINGER, K., GONZALEZ, A., LOZUPONE, C., ZANEVELD, J. R., VÁZQUEZ-BAEZA, Y., BIRMINGHAM, A. *and others.* (2017). Normalization and microbial differential abundance strategies depend upon data characteristics. *Microbiome* **5**, 27.
- WITTEN, D. M. (2011). Classification and clustering of sequencing data using a Poisson model. *The Annals of Applied Statistics* **5**, 2493–2518.
- WU, C., CHEN, J., KIM, J. AND PAN, W. (2016). An adaptive association test for microbiome data. *Genome Medicine* **8**, 56.
- YAN, A. W., FOUTS, D. E., BRANDL, J., STÄRKEL, P., TORRALBA, M., SCHOTT, E., TSUKAMOTO, H., NELSON, K. E., BRENNER, D. A. AND SCHNABL, B. (2011). Enteric dysbiosis associated with a mouse model of alcoholic liver disease. *Hepatology* **53**, 96–105.
- ZHANG, X., MALLICK, H., TANG, Z., ZHANG, L., CUI, X., BENSON, A. K. AND YI, N. (2017). Negative binomial mixed models for analyzing microbiome count data. *BMC Bioinformatics* **18**, 4.

- ZHAO, N., CHEN, J., CARROLL, I. M., RINGEL-KULKA, T., EPSTEIN, M. P., ZHOU, H., ZHOU, J. J., RINGEL, Y., LI, H. AND WU, M. C. (2015). Testing in microbiome-profiling studies with MiRKAT, the microbiome regression-based kernel association test. *American Journal of Human Genetics* **96**, 797–807.
- ZHU, W., WINTER, M. G., BYNDLOSS, M. X., SPIGA, L., DUERKOP, B. A., HUGHES, E. R., BÜTTNER, L., DE LIMA ROMÃO, E., BEHRENDT, C. L., LOPEZ, C. A. *and others*. (2018). Precision editing of the gut microbiota ameliorates colitis. *Nature* **553**, 208.

[Received December 18, 2018; revised October 07, 2019; accepted for publication October 09, 2019]