

Bayesian graphical compositional regression for microbiome data

Jialiang Mao Yuhao Chen Li Ma

Duke University, Durham, NC

Abstract

An important task in microbiome studies is to test the existence of and give characterization to differences in the microbiome composition across groups of samples. Important challenges of this problem include the large within-group heterogeneities among samples and the existence of potential confounding variables that, when ignored, increase the chance of false discoveries and reduce the power for identifying true differences. We propose a probabilistic framework to overcome these issues by combining three ideas: (i) a phylogenetic tree-based decomposition of the cross-group comparison problem into a series of local tests, (ii) a graphical model that links the local tests to allow information sharing across taxa, and (iii) a Bayesian testing strategy that incorporates covariates and integrates out the within-group variation, avoiding potentially unstable point estimates. With the proposed method, we analyze the American Gut data to compare the gut microbiome composition of groups of participants with different dietary habits. Our analysis shows that (i) the frequency of consuming fruit, seafood, vegetable and whole grain are closely related to the gut microbiome composition and (ii) the conclusion of the analysis can change drastically when different sets of relevant covariates are adjusted, indicating the necessity of carefully selecting and including possible confounders in the analysis when comparing microbiome compositions with data from observational studies.

1 Introduction

The human microbiome is the community of numerous microbes that inhabit the human body. Understanding the microbiome can provide insights into various aspects of human health. For example, diseases such as obesity and Type 2 diabetes have been shown to be related to the gut microbiome (??). Next generation sequencing technologies provide ways of profiling the microbiome. This is typically achieved through either shotgun sequencing on the entire genomes of microbes, or through targeting a signature gene—the 16S ribosomal RNA (rRNA) gene that provides barcodes of species identity. The 16S rRNA gene of the bacteria in the samples is sequenced and the sequences are clustered into operational taxonomic units (OTUs) using preprocessing pipelines such as QIIME (?). Traditionally, OTUs are used as representatives of the species at 97% similarity level and their counts (defined as the counts of sequences in each OTU cluster) form the basis of analyzing the composition of the human microbiome (?). More recently, ? proposed to use amplicon sequence variants (ASVs) to achieve more precise characterization of species. Our methodology herein applies to both OTUs and ASVs though we will use “OTUs” to refer to the unit of species.

One important task of microbiome studies is to compare the composition of the microbial community of groups of subjects (????). Various methods have been proposed for comparing two groups of data samples and their underlying distributions, ranging from the classic *t*-test to numerous recently developed methods (see for example ?, ? and ?). However, these generic approaches are either inapplicable or severely underpowered when the samples are OTU counts. In this paper, we analyze the American Gut data (?) to compare the gut microbiome composition of groups of people with different dietary habits. The American Gut data includes samples from a large population and provides abundant sample-specific covariate information. Previous analyses of the American Gut data that compare groups of microbiome samples can be improved since they either focus on certain summary statistics

of the microbiome compositions (?), or disregard the observational nature of the American Gut project by ignoring possible confounders (?). Our approach aims at addressing these limitations.

There are several features of the OTU counts that make this two-sample-problem challenging: (i) high dimensionality: the number of OTUs in the study is often large, (ii) “rare biosphere” (?) or sparsity: the total OTU count of a sample is dominated by a few OTUs, with others having counts closed or equal to zero, (iii) complex covariance structure: the correlations among counts of different OTUs are complicated, and (iv) large overdispersions: the counts of samples in a same group often show large within-group heterogeneities. Fortunately, along with the OTU counts, one can construct a phylogenetic tree that encodes the evolutionary relationships of these OTUs and collect abundant covariate information about the host participant of each sample. In this paper, we propose a testing method, called Bayesian graphical compositional regression (BGCR), that utilizes such information to tackle the aforementioned challenges.

BGCR aims to effectively account for the specific features of microbiome data by combining three techniques—(i) a decomposition of the multinomial likelihood along the phylogenetic tree, (ii) graphical models, and (iii) the Bayesian testing framework. Broadly speaking, our approach falls into a stream of works based on the Dirichlet-multinomial model (DM) (?) that puts a Dirichlet prior on the multinomial parameters to resolve the large overdispersions in OTU counts. Recent developments along the line use the phylogenetic tree—which summarizes the evolutionary relationship among the OTUs and thus serves as a proxy to their functional relationship—to enrich the model construction. Specifically, ? adopt the Dirichlet tree multinomial model (DTM) in which OTU counts are aggregated along the phylogenetic tree and a DM model is introduced to characterize how the OTU counts on each internal node of the tree are distributed into its child nodes. BGCR adopts a similar decomposition to transform the original testing problem into multiple testing involving a

collection of node-specific tests along the phylogenetic tree.

Instead of treating the node-specific tests independently, we use a bottom-up graphical structure to introduce dependency among them. Our motivation is that some of these tests may have poor statistical power due to the limited number of samples and relevant OTU counts. The dependency structure allows effective information sharing among the node-specific problems, thereby improving the power of the tests. Our choice of a bottom-up autoregressive specification comes from the observation that compositional differences are often observed to cluster into “chains” along the phylogenetic tree (?). Such chains can arise in two ways: either due to a structural constraint of DTM—a cross-group difference at any particular node in the phylogenetic tree induces (weaker) differences in its ancestors, or as a result of the functional relatedness of OTUs in the phylogenetic tree.

The decomposition scheme and the graphical structure are merged with the Bayesian testing framework to give fast and interpretable inference through a suite of computational techniques including numerical integration and message passing. On the one hand, exact posterior summaries are available through a recursive algorithm that does not require Monte Carlo simulations; on the other, the testing results are given a full probabilistic characterization with the uncertainty quantified in a coherent manner. In addition, we show that the Bayesian test results in substantially improved power, likely due to the way the within-group variability is dealt with—integrated out rather than estimated as in existing approaches.

Besides the phylogenetic tree, covariate information is also used to assist the modeling of the microbiome composition in recent works focusing on prediction (?) and selection of the covariates with significant associations to the OTU counts (?????). In contrast, little attention has been paid to the cross-group comparison scenario, where incorporating covariates is even more crucial since many microbiome studies are observational and the effects of unadjusted confounders could invalidate the testing results by creating false positives. As a starting point, BGCR allows a node-specific regression adjustment for important covariates.

This adjustment can be incorporated in the inference mechanism with little extra computational burden and is able to reduce the number of false positives effectively, as we will illustrate in our numerical examples. In our analysis of the American Gut data, we show that when a set of possible confounders are ignored, the testing results can be misleading.

In summary, in comparison to existing approaches, the contribution of BGCR are three-fold: (i) it allows a Bayesian framework for testing cross-group differences in microbiome composition along a phylogenetic structure; (ii) it provides a principled probabilistic modeling framework for borrowing information across taxa, thereby enhancing the power for detecting cross-group differences; and (iii) it uses a principled way to incorporate additional covariates in the Bayesian testing setting while maintaining the computational efficiency of the proposed framework.

To close the introduction, we connect BGCR to some relevant references in the DM literature stream. Using DM, ? proposed a generalized Wald-type test statistic based on method-of-moments estimates of the Dirichlet parameters. Due to the large number of OTUs, this test has large degrees of freedom and is usually underpowered. One simple attempt to alleviate this problem is to aggregate OTUs from the same genus and study their composition at the genus level as suggested by ?. However, one could also aggregate the OTUs to other levels in the hierarchy of biological classification, such as the family or order level, yet different levels of aggregation can result in inconsistent testing results (?). Similar to BGCR, ? proposed the PhyloScan test, which also decomposes the original problem to a series of node-specific tests along the phylogenetic tree and tries to incorporate dependencies among these tests. BGCR and PhyloScan differ in the way they deal with these underlying dependencies. PhyloScan introduces dependencies through a scanning procedure by adding up the test statistics in triplets of neighboring nodes while BGCR considers a probabilistic graphical model that describes the dependencies in a generative manner. The latter allows more flexible borrowing of information beyond neighboring triplets along with

a fully probabilistic interpretation.

In [Section 2](#), we describe our method for testing and characterizing the cross-group differences of OTU compositions. In [Section 3](#), four representative simulation scenarios are considered to evaluate the performance of the proposed method. An application of the method to the American Gut data is shown in [??](#). [??](#) concludes.

2 Method

2.1 Data and background

In this section, we set up some notations and briefly review DM and DTM for OTU compositions. The microbiome dataset we work with contains three parts: OTU counts in each sample, a phylogenetic tree over the OTUs, and a set of covariates for each sample.

OTU counts. Consider a microbiome dataset with OTU counts of two groups of subjects on K OTUs denoted by $\Omega = \{\text{OTU}_1, \text{OTU}_2, \dots, \text{OTU}_K\} = \{\omega_1, \omega_2, \dots, \omega_K\}$. Let n_i be the number of samples in group i , $i = 0, 1$. For the j -th sample in group i , $j = 1, \dots, n_i$, let $\mathbf{y}_{ij} = (y_{ij1}, \dots, y_{ijK})$ be the vector of its OTU counts, where y_{ijl} is the number of the l -th OTU in this sample for $l = 1, \dots, K$. Moreover, let $N_{ij} = \sum_{l=1}^K y_{ijl}$ be the total number of OTU counts in that sample. For simplicity, let \mathbf{Y} be the OTU counts of all the samples. [Table 1](#) illustrate the OTU counts for a certain group of samples.

Phylogenetic tree. Let \mathcal{T} be a rooted full binary phylogenetic tree that describes the evolutionary relations of the K OTUs. Let \mathcal{I} denote the set of its internal nodes (i.e., non-leaf nodes). We denote each node A of \mathcal{T} by the set of its descendant OTUs. For example, $A = \{\omega_l\}$, $1 \leq l \leq K$ represents a leaf of \mathcal{T} that contains a single OTU ω_l ; $A = \Omega$ is the root containing all the OTUs. Since the tree is full, each internal node has exactly two children. For $A \in \mathcal{I}$, let A_l, A_r be the left and right children of A . If $A \neq \Omega$, let A_p be its parent and A_s be its sibling (i.e., the node in \mathcal{T} that has the same parent as A). If A is a

leaf, for the j -th sample in group i , we let $y_{ij}(A) = y_{ijl}$ be the count of the OTU that A represents; if $A \in \mathcal{I}$, we recursively define the count in A to be the aggregation of the counts in its two children: $y_{ij}(A) = y_{ij}(A_l) + y_{ij}(A_r)$. Equivalently, $y_{ij}(A) = \sum_{\{l: \omega_l \in A\}} y_{ijl}$. **Figure 1** shows an example of a phylogenetic tree over 6 OTUs. In this example, $A = \{\omega_3, \omega_4, \omega_5, \omega_6\}$, $y_{ij}(A) = y_{ij3} + y_{ij4} + y_{ij5} + y_{ij6}$.

id	ω_1	ω_2	\cdots	ω_K	Sum
1	y_{i11}	y_{i12}	\cdots	y_{i1K}	N_{i1}
2	y_{i21}	y_{i22}	\cdots	y_{i2K}	N_{i2}
\vdots	\vdots	\vdots	\ddots	\vdots	\vdots
n_i	y_{in_i1}	y_{in_i2}	\cdots	y_{in_iK}	N_{in_i}

Table 1: OTU counts for the i -th group.

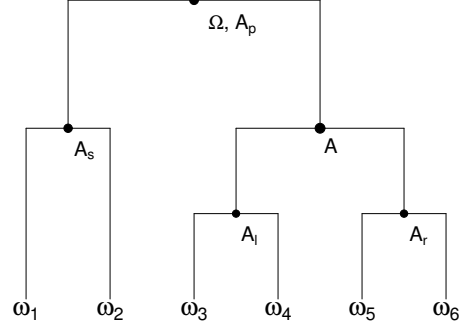


Figure 1: A phylogenetic tree over 6 OTUs.

Covariates. Microbiome studies often collect covariate information about the samples. For example, these covariates may include the demographic information of the participants, their dietary habits, and histories of disease. For the j -th sample in group i , let $\mathbf{x}_{ij} = (1, x_{ij1}, \dots, x_{ijp})' \in \mathbb{R}^{p+1}$ denote the p covariates that we want to include in the analysis. Moreover, let $z_{ij} \in \{0, 1\}$ be its group indicator.

? used the Dirichlet-multinomial model (DM) to account for the within-group heterogeneity among the OTU counts:

$$\begin{aligned} \mathbf{y}_{ij} \mid N_{ij}, \boldsymbol{\pi}_{ij} &\stackrel{\text{ind}}{\sim} \text{Multinomial}(N_{ij}, \boldsymbol{\pi}_{ij}) \\ \boldsymbol{\pi}_{ij} \mid \boldsymbol{\pi}_i, \nu &\stackrel{\text{i.i.d.}}{\sim} \text{Dirichlet}(\nu \boldsymbol{\pi}_i) \end{aligned} \tag{2.1}$$

for $i = 0, 1$, where $\boldsymbol{\pi}_i = (\pi_{i1}, \dots, \pi_{iK})$ with $\sum_{l=1}^K \pi_{il} = 1$, and $\nu > 0$ a dispersion parameter that controls the within-group variability. A limitation of DM is its induced correlations among OTUs—that they are independent up to the summation constraint (?), which is

unrealistic since the OTUs have inherent and complicated relationships with each other, partly reflected in their evolutionary history summarized in a phylogenetic tree.

we propose to adopt the Dirichlet-tree multinomial model (DTM) by directly incorporating the phylogenetic tree in the modeling of the OTU counts. Specifically, DTM models how the counts in $A \in \mathcal{I}$ are distributed to its two children by a beta-binomial model:

$$\begin{aligned} y_{ij}(A_l) \mid y_{ij}(A), \theta_{ij}(A) &\stackrel{\text{i.i.d.}}{\sim} \text{Binomial}(y_{ij}(A), \theta_{ij}(A)) \\ \theta_{ij}(A) \mid \theta_i(A), \nu(A) &\stackrel{\text{i.i.d.}}{\sim} \text{Beta}(\theta_i(A)\nu(A), (1 - \theta_i(A))\nu(A)), \end{aligned} \quad (2.2)$$

where $\theta_{ij}(A)$ denotes the proportion of counts in node A that are distributed to its left child A_l , $\theta_i(A)$ the mean of the $\theta_{ij}(A)$'s in group i , and $\nu(A)$ a precision or dispersion parameter that controls the variability of $\theta_{ij}(A)$ around the group mean, for $i = 1, 2, j = 1, \dots, n_i$. Let $\mathcal{L}_{BB}(\theta_i(A), \nu(A) \mid y_{ij}(A_l), y_{ij}(A_r))$ denote the likelihood of the parameters in (2.2) obtained by integrating out $\theta_{ij}(A)$, which with slightly more general notations can be written as

$$\mathcal{L}_{BB}(\theta, \nu \mid y_1, y_2) = \begin{cases} B(\theta\nu + y_1, (1 - \theta)\nu + y_2) / B(\theta\nu, (1 - \theta)\nu), & \text{if } \nu < \infty \\ \theta^{y_1} (1 - \theta)^{y_2}, & \text{if } \nu = \infty \end{cases} \quad (2.3)$$

for $y_1, y_2 \in \{0, 1, 2, \dots\}$, $\theta \in [0, 1]$ and $\nu \in (0, \infty]$, where $B(\cdot, \cdot)$ is the Beta function. It can be shown that the DM likelihood for each sample can be factorized into a series of beta-binomial likelihoods along the tree

$$\mathcal{L}_{DM}(\boldsymbol{\pi}_i, \nu \mid \mathbf{y}_{ij}) = \prod_{A \in \mathcal{I}} \mathcal{L}_{BB}(\theta_i(A), \nu_i(A) \mid y_{ij}(A_l), y_{ij}(A_r)) \quad (2.4)$$

provided that $\theta_i(A) = \pi_i(A_l) / \pi_i(A)$ and $\nu_i(A) = \nu \pi_i(A)$, where $\pi_i(A) = \sum_{l \in A} \pi_{il}$ for $A \in \mathcal{I}$ (??). In this sense, DTM could be seen as a generalization of DM.

2.2 Bayesian compositional regression for microbiome data

The likelihood factorization (2.4) suggests a “divide-and-conquer” strategy to perform inference on $\boldsymbol{\pi}_i$. Specifically, inference on $\boldsymbol{\pi}_i$ could be achieved equivalently by doing inference on $\{\theta_i(A) : A \in \mathcal{I}\}$, $i = 0, 1$. In this section, we use this strategy to transform the original problem of comparing microbial compositions of two groups of sample into a series of tests on $\{\theta_i(A) : A \in \mathcal{I}\}$ in a more general situation that allows adjustment for covariates.

For each $A \in \mathcal{T}$, we modify the beta-binomial specification in DTM to allow regression adjustments for covariates:

$$\begin{aligned} y_{ij}(A_l) \mid y_{ij}(A), \theta_{ij}(A) &\stackrel{\text{ind}}{\sim} \text{Binomial}(y_{ij}(A), \theta_{ij}(A)) \\ \theta_{ij}(A) \mid \mathbf{x}_{ij}, z_{ij}, \nu(A) &\sim \text{Beta}(\theta_{\mathbf{x}_{ij}, z_{ij}}(A)\nu(A), (1 - \theta_{\mathbf{x}_{ij}, z_{ij}}(A))\nu(A)) \\ g(\theta_{\mathbf{x}_{ij}, z_{ij}}(A)) &= \mathbf{x}_{ij}^\top \boldsymbol{\beta}(A) + z_{ij} \gamma(A), \end{aligned} \tag{2.5}$$

where $g : [0, 1] \rightarrow \mathbb{R}$ is a link function, such as the logit link $g(x) = \log(x/(1 - x))$ for $x \in (0, 1)$. $\boldsymbol{\beta}(A) \in \mathbb{R}^{p+1}$ and $\gamma(A) \in \mathbb{R}$ are the unknown parameters of the local model on A . We refer to this model as the Bayesian compositional regression (BCR). Under BCR, testing the original null that there is no cross-group difference in the microbial composition is equivalent to jointly testing a set of local hypotheses on all $A \in \mathcal{I}$:

$$H_0(A) : \gamma(A) = 0 \quad \text{vs} \quad H_1(A) : \gamma(A) \neq 0. \tag{2.6}$$

Instead of modeling the OTU counts around some “group-specific centroid” as in DM and DTM, we model them around some “covariate-specific centroid” with the grouping information included as a special “covariate”. Recent literature has shown that many covariates have associations and possible effects on the microbiome composition (??????). In the testing scenario, some of these covariates may also be correlated with the grouping variable. Ignor-

ing such confounding can reduce the power for identifying differences, or worse yet, lead to false positives. The regression adjustments in (2.5) block out the possible confounders and give us an opportunity to improve the accuracy of the testing results and reduce the chance of false discoveries.

We shall take a Bayesian approach to the local testing problem (2.6) for three reasons: (1) it provides a natural way to deal with the dispersion $\nu(A)$, essentially a nuisance parameter, through integration instead of potentially unstable points estimates; (2) it allows the introduction of dependencies among the local hypotheses in a generative manner as shown in Section 2.3; and (3) it gives a fully probabilistic characterization of the uncertainty.

To this end, we first need to specify priors for $\beta(A)$ and for $\gamma(A)$ under $H_1(A)$. A simple choice is to put independent normal priors on the elements of $\beta(A)$ and $\gamma(A)$. Note that using vague proper priors on the model specific parameters $\gamma(A)$ is often problematic in testing and would cause the so-called “Bartlett’s Paradox” (?). Alternatively, various principles for constructing “objective” priors for the coefficients can be adopted. For example, generalizations of the g -prior and the mixture of g -priors (?) for linear regressions to GLMs can be employed here, see for example ? and ?. Specifically, one could apply the local information metric (LIM) g -prior as suggested by ? on $\gamma(A)$. In our setting, the difference in dimensions between the parameter space under the alternative and the null is only one; therefore, putting normal prior $N(0, \sigma_\gamma^2(A))$ with reasonably large variance (such as $\sigma_\gamma^2(A) = 10$) would give reasonable results. For the “common” parameters $\beta(A)$, we adopt independent normal priors $N(0, \sigma_\beta^2(A))$ with large $\sigma_\beta(A)$ on its elements. For example, when the covariates are standardized (rescaled to have a mean of zero and a standard deviation of one), $N(0, 16)$ covers most probable values of $\beta(A)$ in many applications.

Without further knowledge about the nuisance dispersion parameter $\nu(A)$, we put a prior $G_A(\nu)$ on it and integrate it out. This is key to the improvement gained by applying a Bayesian testing scheme compared to its frequentist counterparts that rely on a point

estimator of $\nu(A)$ (Section 3). It is necessary for $G_A(\nu)$ to have a large support to allow various levels of dispersions. We hence take $\log_{10} \nu(A) \sim \text{Unif}(-1, 4)$ that covers a wide range of dispersion levels but does not put too much prior mass on unreasonably large or small values (?). Other choices with unbounded support such as Gamma priors on $\nu(A)$ can also be adopted.

To perform posterior inference, we introduce the state indicator of the test on node A for $A \in \mathcal{I}$:

$$S(A) = \begin{cases} 0, & H_0(A) \text{ is true,} \\ 1, & H_1(A) \text{ is true.} \end{cases} \quad (2.7)$$

We assume for now that the $S(A)$'s are independent and let $\Pr(S(A) = 1) = \rho(A)$ *a priori*. The testing problem on A now becomes an inference problem on $S(A)$. Let $M_s(A)$ be the marginal likelihood under $H_s(A)$, $s \in \{0, 1\}$. We have

$$M_s(A) = C \iiint \prod_{i=1}^2 \prod_{j=1}^{n_i} \mathcal{L}_{BB}(g^{-1}(\mathbf{x}_{ij}^\top \boldsymbol{\beta} + z_{ij}\gamma), \nu \mid y_{ij}(A_l), y_{ij}(A_r)) dF_{s,A}(\boldsymbol{\beta}) dF_{s,A}(\gamma) dG_A(\nu),$$

where C is a constant with respect to the parameters, $F_{s,A}(\boldsymbol{\beta}), F_{s,A}(\gamma)$ are the priors for $\boldsymbol{\beta}(A)$ and $\gamma(A)$ under $H_s(A)$, $\mathcal{L}_{BB}(\cdot)$ the beta-binomial marginal likelihood defined in (2.3), and $G_A(\nu)$ the prior on $\nu(A)$. We give details on the computational strategy for evaluating $M_s(A)$ in Online Supplementary Materials A. Given $M_s(A)$, the posterior of $S(A)$ is

$$S(A) \mid \tilde{\rho}(A) \sim \text{Bernoulli}(\tilde{\rho}(A)), \text{ where } \tilde{\rho}(A) = \frac{\rho(A)M_1(A)}{(1 - \rho(A))M_0(A) + \rho(A)M_1(A)}.$$

We shall refer to $\tilde{\rho}(A) = \Pr(S(A) = 1 \mid \mathbf{Y})$ as the *posterior marginal alternative probability* (PMAP) on node A , denoted also as $\text{PMAP}(A)$. The PMAP is large when the evidence against the local null is strong. The set of PMAPs along \mathcal{T} can be used to test the original null and pinpoint the differences. We postpone the details on using PMAPs for making such

decisions based on multiple testing considerations to Section 2.5.

We end this section with three further comments on our model specification in BCR. Firstly, we assume that the two groups share a common dispersion parameter $\nu(A)$ to simplify the computation. If this assumption is not likely to hold in the data, group-specific dispersions can be adopted. Secondly, BCR does not automatically select covariates. This is not restrictive in the testing scenario since our primary goal is not to identify the best predictive model for the microbiome composition or to estimate the “effect size” of specific covariates on the microbiome composition, but to preclude them from introducing bias to the analysis. In practice, all the suspect confounders can be included in the model even some of them may turn out to have no confounding effects provided that the number of covariates is not too large. We note that one must practice caution in adopting covariate selection in this context. Traditional statistical variable selection is based on the predictive ability of the covariates, and can lead to inappropriate elimination of confounders and result in false positives. We provide more discussion and a simple numerical illustration in Online Supplementary Materials C. Thirdly, in the regression setup, we assume that the regression coefficients $\beta(A)$ are the same for both groups. This “common slope” assumption essentially serves as an identifiability constraint, without which the null hypothesis is ill-defined—the two groups will be identical only for a specific combination of covariate values but not others, and consequently it would not make sense to test whether two groups of samples involving different covariate values are “identical” in microbiome composition.

2.3 Bayesian graphical compositional regression

In BCR, the test on node A is performed based only on the empirical evidence in A . Let $\mathcal{S} = \{S(A) : A \in \mathcal{I}\}$, the collection of state indicators of all internal nodes in \mathcal{T} . Elements in \mathcal{S} are independent *a priori*. However, this independence assumption disregards the inherent relations of the $S(A)$ ’s—that they are naturally linked by the phylogenetic tree. Introducing

suitable dependency structures on \mathcal{S} can enhance inference because (1) cross-group differences often occur along the phylogenetic tree in a clustered manner, forming chains, which is pointed out in ? and will be confirmed in our data analysis, and (2) from a statistical perspective, some local tests may involve only few counts and their power is thus limited, especially when the nodes are close to the leaves of \mathcal{T} . Introducing a dependency structure on \mathcal{S} allows these nodes to borrow information from each other and thus increases the power of the tests. The dependency structure we introduce should satisfy two desiderata. On the one hand, it should allow flexible information sharing among nodes that are close on the phylogenetic tree. On the other hand, it should be simple enough to keep the posterior inference tractable.

Recall that the node-specific beta-binomial models in BCR are naturally linked together by the phylogenetic tree, which provides a proxy to their functional relationship. This suggests that we can use the phylogenetic information when imposing dependency structures on \mathcal{S} . In particular, we consider a “bottom-up” auto-regressive structure such that if the signal (cross-group difference) is present at a certain node $A \in \mathcal{T}$, there is a chance for it to be “carried upwards” to A_p . Formally, for $A \in \mathcal{I}$, we let

$$\begin{aligned} \text{logit}[\Pr(S(A) = 1 \mid S(A_l), S(A_r))] &= \alpha(A) + \tau(A) \cdot \mathbb{1}_{[S(A_l)+S(A_r) \geq 1]} \\ &+ \kappa(A) \cdot \mathbb{1}_{[S(A_l)+S(A_r)=2]}, \end{aligned} \tag{2.8}$$

where $\tau(A), \kappa(A) \geq 0$ together describe how likely signals at A_l or A_r are passed upwards along the tree. If A_l or A_r (or both) is a leaf node, we set $S(A_l)$ or $S(A_r)$ to be zero since there is no test performed on that node. This model can be embedded in a family of models specified by the conditional probabilities

$$\Pr(S(A) = s \mid S(A_l) = s_l, S(A_r) = s_r) = \rho_{s_l s_r, s}(A), \tag{2.9}$$

where $s, s_l, s_r \in \{0, 1\}$ and $\rho_{s_l s_r, s} \in [0, 1]$. Specifically, the auto-regressive model forces $\rho_{01, s}(A) = \rho_{10, s}(A)$ since usually there is no specific prior information to differentiate between the two children of a node. Equivalently, the conditional probability of $S(A)$ given $S(A_l), S(A_r)$ can be characterized by a transition matrix

$$\boldsymbol{\rho}(A) = \begin{pmatrix} \rho_{00,0}(A) & \rho_{00,1}(A) \\ \rho_{01,0}(A) & \rho_{01,1}(A) \\ \rho_{10,0}(A) & \rho_{10,1}(A) \\ \rho_{11,0}(A) & \rho_{11,1}(A) \end{pmatrix}$$

where the elements in each row of $\boldsymbol{\rho}(A)$ sum up to 1. We denote $\boldsymbol{\rho} = \{\boldsymbol{\rho}(A) : A \in \mathcal{I}\}$ and defer the specification of elements in $\boldsymbol{\rho}$ to Section 2.5. The set of auto-regressive models with transition probabilities $\boldsymbol{\rho}$ induces a joint distribution $F_{\mathcal{S}}$ on \mathcal{S} . Incorporating this distribution into BCR leads to a hierarchical formulation of our model for each $A \in \mathcal{I}$:

$$\begin{aligned} S(A) \mid \boldsymbol{\rho} &\sim F_{\mathcal{S}}, \quad \nu(A) \mid G_A \stackrel{\text{ind}}{\sim} G_A \\ \boldsymbol{\beta}(A) &\stackrel{\text{ind}}{\sim} \text{N}_{p+1}(\mathbf{0}, \sigma_{\boldsymbol{\beta}}^2(A) \cdot I_{p+1}) \\ \gamma(A) &\stackrel{\text{ind}}{\sim} \mathbb{1}_{[S(A)=0]} \cdot \delta_0 + \mathbb{1}_{[S(A)=1]} \cdot \text{N}(0, \sigma_{\gamma}^2(A)) \\ g(\theta_{\mathbf{x}_{ij}, z_{ij}}(A)) &= \mathbf{x}_{ij}^{\top} \boldsymbol{\beta}(A) + z_{ij} \gamma(A) \\ \theta_{ij}(A) \mid \mathbf{x}_{ij}, z_{ij}, \nu(A) &\sim \text{Beta}(\theta_{\mathbf{x}_{ij}, z_{ij}}(A) \nu(A), (1 - \theta_{\mathbf{x}_{ij}, z_{ij}}(A)) \nu(A)) \\ y_{ij}(A_l) \mid y_{ij}(A), \theta_{ij}(A) &\stackrel{\text{ind}}{\sim} \text{Binomial}(y_{ij}(A), \theta_{ij}(A)), \quad i = 1, 2, j = 1, \dots, n_i, \end{aligned} \tag{2.10}$$

where I_{p+1} is the $(p+1)$ -dimensional identity matrix, $\mathbb{1}_{[\cdot]}$ the indicator function, and δ_0 a point mass at zero. We shall refer to this model as the Bayesian graphical compositional regression (BGCR).

An alternative way to introduce dependencies among elements in \mathcal{S} that also incorporate the phylogenetic information is by utilizing the “top-town” Markov tree model (MT) (??),

which instead specifies $\Pr(S(A) \mid S(A_p))$ for $A \in \mathcal{T} \setminus \{\Omega\}$. However, the “explaining away” effect (?) of the auto-regressive model is important for microbiome data. MT pushes signals downwards along the tree to both children of a node and implies cooccurrence of signals in the sibling nodes, which is not the typical situation in the microbiome context. Instead, the signals in microbiome studies often form chains, with only one of the two children nodes share the signal with the parent node. This is the primary reason for adopting a “bottom-up” autoregressive model.

2.4 Inference under BGCR

As in BCR, the marginal posteriors $\{\Pr(S(A) = 1 \mid \mathbf{Y}) : A \in \mathcal{I}\}$ play a pivotal role in the posterior inference under BGCR. We next show that these quantities can be calculated exactly up to the approximation to $M_s(A)$ ’s, without entailing Monte Carlo simulations.

The joint distribution introduced by BGCR on \mathcal{S} can be represented by a Bayesian network, i.e., a directed acyclic graph (DAG) $\mathcal{G}(V_{\mathcal{S}}, B_{\mathcal{S}})$ with nodes $V_{\mathcal{S}} = \{S(A) : A \in \mathcal{I}\}$ and edges $B_{\mathcal{S}} = \{S(A) \rightarrow S(A_p) : A \in \mathcal{I} \setminus \{\Omega\}\}$. Moralizing $\mathcal{G}(V_{\mathcal{S}}, B_{\mathcal{S}})$ yields an undirected graph $\mathcal{U}(V_{\mathcal{S}}, E_{\mathcal{S}})$ from which the conditional dependencies in \mathcal{S} can be read off directly (?). Specifically, the vertices of \mathcal{U} are the same of those of \mathcal{G} ; there is an undirected edge between $S(A)$ and $S(A')$ in \mathcal{U} if A and $A' \in \mathcal{I}$ are siblings or if they form a parent-child pair: $E_{\mathcal{S}} = \{\{S(A), S(A')\} : A' = A_p \text{ or } A' = A_s, A \in \mathcal{I} \setminus \{\Omega\}\}$. Moreover, we refer to a chain of edges as a “path”. Two nodes $S(A)$ and $S(A')$ in \mathcal{U} are conditionally independent given a set of nodes $v = \{S(A_1), \dots, S(A_m)\} \subset V_{\mathcal{S}}$ if and only if any path connecting $S(A)$ and $S(A')$ passes some elements in v . **Figure 2** (i), (ii), and (iii) respectively give illustrations of a fictional phylogenetic tree, the corresponding Bayesian network \mathcal{G} on \mathcal{S} , and the moralized undirected graph \mathcal{U} . The structure of \mathcal{U} induced by BGCR has two implications on posterior computations. On the one hand, \mathcal{U} contains many loops, which at first glance would substantially complicate the posterior and the inference algorithm. If \mathcal{U}

were a tree-type graph without loops, as in MT, simple and exact inference can be achieved by a forward-backward algorithm. On the other hand, all the loops in \mathcal{U} include only three nodes and the dependencies among $S(A)$'s are local— $S(A)$ is independent of other elements in \mathcal{S} once conditioning on $S(A_p)$, $S(A_s)$, $S(A_l)$ and $S(A_r)$, suggesting that \mathcal{U} is not “too far” from a tree-type graph that allows exact inference. Intuitively, if we treat each loop as a unit, \mathcal{U} can be viewed as a “tree of loops” on which exact inference can now be performed.

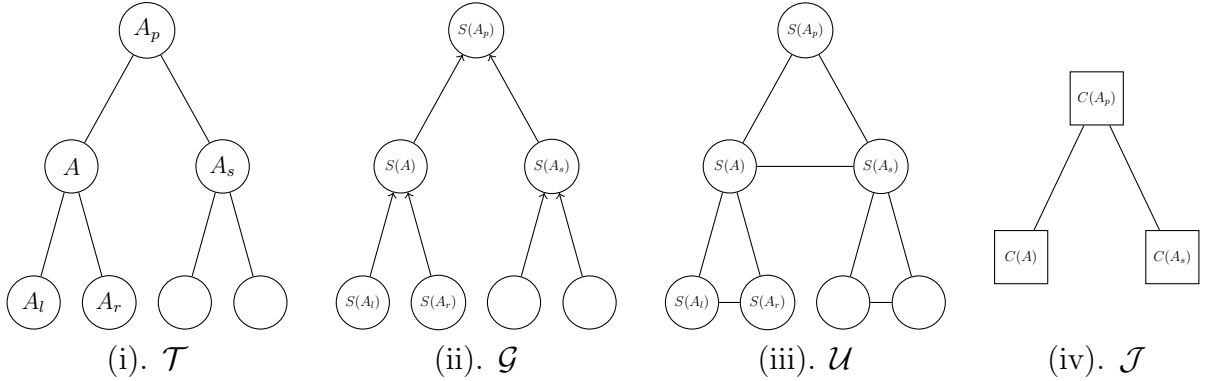


Figure 2: (i). A fictional phylogenetic tree (the leaves of \mathcal{T} are not shown); (ii). The “bottom-up” Bayesian network \mathcal{G} on \mathcal{S} corresponds to \mathcal{T} ; (iii). The undirected graph after moralizing \mathcal{G} ; (iv). The clique tree derived from \mathcal{U} .

Formally, for $S(A) \in V_{\mathcal{S}}$, we call $C(A) = (S(A), S(A_l), S(A_r))$ the *clique* of $S(A)$, or simply the *clique* of A . The conditional dependencies among $\mathcal{C} = \{C(A) : A \in \mathcal{I}\}$ are inherited from \mathcal{U} and can be represented by an undirected graph $\mathcal{J}(V_{\mathcal{C}}, E_{\mathcal{C}})$ in a similar manner, where $V_{\mathcal{C}} = \mathcal{C}$, $E_{\mathcal{C}} = \{\{C(A), C(A')\} : A' = A_p, A \in \mathcal{I} \setminus \{\Omega\}\}$. See Figure 2 (iv) for an illustration. It is easy to see that \mathcal{J} contains no loop and that any two nodes in \mathcal{J} are connected by a unique path. Therefore, \mathcal{J} is an undirected tree. We refer to \mathcal{J} as the *clique tree* and perform inference thereon. Note that $C(A)$ can take on eight values, denoted by $\mathbf{c} = \{c_1, c_2, \dots, c_8\}$ where $c_i = (c_{i1}, c_{i2}, c_{i3}) = (\lfloor \frac{i-1}{4} \rfloor \bmod 2, \lfloor \frac{i-1}{2} \rfloor \bmod 2, (i-1) \bmod 2)$ (each element is the corresponding digit of the binary representation of i).

For $A \in \mathcal{I}$, without loss of generality, let A be the left child of A_p . Eq. (2.9) induces the

top-down transition probabilities on the clique tree \mathcal{J} :

$$\begin{aligned}\Pr(C(A) = c_{i'} \mid C(A_p) = c_i) &= \Pr((S(A), S(A_l), S(A_r)) = c_{i'} \mid (S(A_p), S(A), S(A_s)) = c_i) \\ &= \Pr((S(A_l), S(A_r)) = (c_{i'2}, c_{i'3}) \mid S(A) = c_i),\end{aligned}\tag{2.11}$$

for $C(A) \in \mathcal{C}$, where $i, i' = 1, 2, \dots, 8$. As with $\boldsymbol{\rho}(A)$, these transition probabilities could be organized into an 8×8 matrix $\boldsymbol{\xi}(A)$ with the i' -th element in row i defined by $\xi_{ii'}(A) = \Pr(C(A) = c_{i'} \mid C(A_p) = c_i)$. Each row of this matrix represents the conditional distribution of $C(A)$ when $C(A_p)$ takes one of the eight possible values in \mathcal{C} . Since $\Omega \in \mathcal{T}$ does not have a parent, we let each row of $\boldsymbol{\xi}(\Omega)$ be the induced marginals of $C(\Omega)$.

In addition, for $A \in \mathcal{I}$, let $\mathcal{T}(A)$ be the subtree of \mathcal{T} with A as the root. Recall that for the j -th sample in group i , $y_{ij}(A)$ denotes the OTU counts in node A . Let $\mathbf{y}_{ij}(A)$ denote all the OTU counts that fall into the subtree $\mathcal{T}(A)$. Note that the difference between $y_{ij}(A)$ and $\mathbf{y}_{ij}(A)$ is important: $y_{ij}(A)$ contains only the counts in A while $\mathbf{y}_{ij}(A)$ contains the set of $y_{ij}(A_d)$'s for all $A_d \in \mathcal{T}(A)$. Let $\mathbf{Y}(A) = \{\mathbf{y}_{ij}(A) : i = 1, 2, j = 1, \dots, n_i\}$ be the counts in $\mathcal{T}(A)$ from all the samples. Bayesian inference on $S(A)$ relies on evaluating the posterior transition probability matrix $\tilde{\boldsymbol{\xi}}(A)$ for each $C(A) \in \mathcal{C}$. We next describe how these posteriors can be calculated with a recursive algorithm. Specifically, the root of the clique tree, $C(\Omega)$, first collect information iteratively from other nodes. After all the information is collected, it is distributed downwards along the tree to finish the information sharing.

Information collection. For $i = 1, 2, \dots, 8$, let $\phi_i(A) = \Pr(\mathbf{Y}(A) \mid C(A_p) = c_i)$ be the prior predictive distribution of $\mathbf{Y}(A)$ evaluated at the observed counts given $C(A_p) = c_i$.

When $A \in \mathcal{I} \setminus \{\Omega\}$, it is given by

$$\begin{aligned}
\phi_i(A) &:= \Pr(\mathbf{Y}(A) \mid C(A_p) = c_i) \\
&= \sum_{1 \leq i' \leq 8} \Pr(\mathbf{Y}(A) \mid C(A) = c_{i'}) \Pr(C(A) = c_{i'} \mid C(A_p) = c_i) \\
&= \sum_{1 \leq i' \leq 8} \xi_{ii'}(A) M_{c_{i'1}}(A) \Pr(\mathbf{Y}(A_l) \mid C(A) = c_{i'}) \Pr(\mathbf{Y}(A_r) \mid C(A) = c_{i'}) \\
&= \sum_{1 \leq i' \leq 8} \xi_{ii'}(A) M_{c_{i'1}}(A) \phi_{i'}(A_l) \phi_{i'}(A_r).
\end{aligned} \tag{2.12}$$

When A has no children in \mathcal{I} , by definition $\phi_i(A) = 1$ for all i . When $A = \Omega$, $C(A)$ does not have a parent in the clique tree. To simplify the notation, we can introduce an imaginary node Ω_p that serves as Ω 's parent. Without loss of generality, we let Ω be the left child of Ω_p and let Ω_s with $\phi_i(\Omega_s) = 1$ be Ω 's ‘‘imaginary’’ sibling. By setting $\Pr(C(\Omega) = c_{i'} \mid C(\Omega_p) = c_i)$ to the corresponding marginals of $C(\Omega)$ for $c_i, c_{i'} \in \mathbf{c}$, we can keep the formulation in (2.12).

As a byproduct of the inference algorithm, $\phi_1(\Omega)$ gives the marginal likelihood of the data, which can be used to guide the selection of hyper-parameters (Section 2.6). Moreover, since the marginal likelihood is available, variable selection can be achieved by putting spike-and-slab priors on the regression coefficients (details are in Online Supplementary Materials C).

Information distribution. By Bayes' theorem, the posterior transition probability

$$\begin{aligned}
\tilde{\xi}_{ii'}(A) &:= \Pr(C(A) = c_{i'} \mid C(A_p) = c_i, \mathbf{Y}) \\
&= \frac{\Pr(C(A) = c_{i'}, \mathbf{Y}(\Omega) \mid C(A_p) = c_i)}{\Pr(\mathbf{Y}(\Omega) \mid C(A_p) = c_i)} \\
&= \frac{\Pr(C(A) = c_{i'} \mid C(A_p) = c_i) \Pr(\mathbf{Y}(A) \mid C(A) = c_{i'})}{\Pr(\mathbf{Y}(A) \mid C(A_p) = c_i)} \\
&= \frac{\xi_{ii'}(A) M_{c_{i'1}}(A) \phi_{i'}(A_l) \phi_{i'}(A_r)}{\phi_i(A)}.
\end{aligned} \tag{2.13}$$

Let $\tilde{\xi}(A)$ be the matrix of the posterior transition probabilities of $C(A)$ given $C(A_p)$. Based on $\{\tilde{\xi}(A) : A \in \mathcal{I}\}$, it is easy to compute the PMAP on each node $A \in \mathcal{I}$. Specifically,

starting from the root of the tree, each row of $\tilde{\xi}(\Omega)$ represents the posterior marginals of the clique $C(\Omega)$. Thus

$$\Pr(C(\Omega) = c_{i'} \mid \mathbf{Y}) = \sum_{1 \leq i \leq 8} \tilde{\xi}_{ii'}(\Omega)/8 \quad (2.14)$$

for $i' = 1, 2, \dots, 8$. By marginalization, we can get $\text{PMAP}(\Omega)$. The PMAPs on other nodes can be computed by induction. For example, given the clique marginals of $C(A_p)$, we can compute the clique marginals of $C(A)$

$$\Pr(C(A) = c_{i'} \mid \mathbf{Y}) = \sum_{1 \leq i \leq 8} \Pr(C(A_p) = c_i \mid \mathbf{Y}) \tilde{\xi}_{ii'}(A) \quad (2.15)$$

for $i' = 1, 2, \dots, 8$, from which the PMAP on A can be obtained by marginalization. We can also compute the *posterior joint alternative probability* (PJAP) that captures the empirical evidence against the global null. Formally,

$$\text{PJAP} = 1 - \Pr(S(A) = 0, A \in \mathcal{T} \mid \mathbf{Y}) = 1 - \prod_{A \in \mathcal{I}} \tilde{\xi}_{11}(A). \quad (2.16)$$

Note that the above steps are a variant of the junction tree algorithm (?) that efficiently calculate the marginals in a graphical model. Directly applying the standard junction tree algorithm on \mathcal{J} outputs the PMAPs, however, it does not allow efficient computation of the PJAP. Algorithm 1 summarizes the entire inference recipe for testing the existence of cross-sample differences under BGCR.

2.5 Decision making

The PJAP and the PMAPs provide the bases of making decisions about the original and the node-specific hypotheses along the tree. For the original testing problem, we reject the null if $\text{PJAP} > 0.5$, which corresponds to the Bayes optimal decision rule under the simple 0-1 loss. Similarly, we reject $H_0(A)$ and call A a significant node if $\text{PMAP}(A) > L$ for some threshold

Algorithm 1 BGCR for comparing microbiome composition

Construct the clique tree \mathcal{J} . ▷ Preprocessing

for A in \mathcal{I} **do**

 Compute the marginal likelihoods of the local test on A .

 Compute the prior transition matrix $\boldsymbol{\xi}(A)$ on the clique of A .

 Compute $D(A)$ — the depth of A defined as the number of edges from A to the root of \mathcal{T} .

end for

for d in $\max_{A \in \mathcal{I}}\{D(A)\} : 0$ **do** ▷ Recursive information collection

for A with $D(A) = d$ **do**

if $C(A)$ has no children in \mathcal{J} **then**

 Let $\phi_i(A_l) = \phi_i(A_r) = 1, i = 1, 2, \dots, 8$.

else

 Compute $\phi_i(A) = \sum_{1 \leq i' \leq 8} \boldsymbol{\xi}_{ii'}(A) M_{c_{i'1}}(A) \phi_{i'}(A_l) \phi_{i'}(A_r), \quad i = 1, 2, \dots, 8$.

end if

end for

end for

for A in \mathcal{I} **do** ▷ Information distribution

 Compute $\tilde{\boldsymbol{\xi}}_{ii'}(A) = \frac{\boldsymbol{\xi}_{ii'}(A) M_{c_{i'1}}(A) \phi_{i'}(A_l) \phi_{i'}(A_r)}{\phi_i(A)}, \quad i, i' = 1, 2, \dots, 8$.

end for

Compute the PJAP. ▷ Global information summary

for d in $0 : \max_{A \in \mathcal{I}}\{D(A)\}$ **do** ▷ Local information summary

for A with $D(A) = d$ **do**

if $A = \Omega$ **then**

 Let $\Pr(C(A_p) = c_i \mid \mathbf{Y}) = 1/8, \quad i = 1, 2, \dots, 8$.

else

 Compute $\Pr(C(A) = c_{i'} \mid \mathbf{Y}) = \sum_{1 \leq i \leq 8} \Pr(C(A_p) = c_i \mid \mathbf{Y}) \tilde{\boldsymbol{\xi}}_{ii'}(A), \quad i' = 1, 2, \dots, 8$.

 Compute PMAP(A) by marginalization.

end if

end for

end for

Report nodes with $\text{PMAP}(A) > L$ for some threshold L . ▷ Decision making

Reject the original null if $\text{PJAP} > 0.5$.

$0 < L < 1$. For example, $L = 0.5$ is recommended by ?. Details on decision making can be found in Online Supplementary Materials B. It is worth noting that by reporting the significant nodes and marking them on the phylogenetic tree (or simply mark the PMAPs along the tree without explicitly providing a decision), we have a natural way to characterize and visualize the cross-group differences that offers more insights than merely providing a decision about a test on the original null. For example, it sheds light on the set of microbes most relevant to the cross-group differences.

2.6 Specifications of hyper-parameters

In BGCR, we need to specify $\alpha(A)$, $\tau(A)$ and $\kappa(A)$ in (2.8). For simplicity, we let α, τ, κ be global parameters that are the same for all $A \in \mathcal{I}$.

Choice of α . The *prior joint alternative probability* (PrJAP) can be written as

$$\text{PrJAP} = 1 - \left(\frac{1}{1 + \exp(\alpha)} \right)^{|\mathcal{I}|},$$

which is a monotone function of α . We can then choose $\alpha = \alpha_0$ such that the PrJAP is at a desired level such as 0.5.

Choice of κ . τ and κ together control the “stickiness” of the signals, that is, how likely the “chaining” pattern will occur along the phylogenetic tree. To encapsulate the “explaining away” effect, we set $\kappa = 0$ and assume that the signal at A could be “explained away” by the signal at one of its children. Another choice for κ is $\kappa = \tau$, which imposes additive effect of signals at the sibling nodes on the parent.

Choice of τ . Given $\alpha = \alpha_0, \kappa = 0$, we use an empirical Bayes procedure to choose the value of τ . Given τ , the marginal likelihood is given by $\phi_1(\Omega)$ as defined in (2.12). We maximize $\phi_1(\Omega)$ over τ to get $\hat{\tau}$ as the prior choice for τ . In search of τ , we focus on the interval $[0, \tau_{max})$, where τ_{max} is some predetermined upper bound on τ . One way to choose

τ_{max} is based on the *prior marginal alternative probability* (PrMAP) on each node, which is a monotonically increasing function of τ . Therefore, the total number of nodes with PrMAP exceeding the threshold L also increases with τ . We can preselect an upper bound of the total prior expected number of true alternatives and solve for τ_{max} accordingly. In practice, when the number of OTUs is not very large (≤ 150), we find that setting $\tau_{max} = 6$ usually works well (this corresponds to setting the prior expected sum of PrMAPs equal to 2 with 100 total OTUs). When $\tau = 0$, BGCR incorporates BCR as a special case.

3 Numerical examples

3.1 Evaluating the performance of BGCR

We first carry out three simulation studies to illustrate the performance of BGCR. For these simulations, several synthetic datasets are obtained based on the July 29, 2016 version of the fecal data of the American Gut Project (?), which collected microbiome samples from different body sites of a large number of participants and offers publicly accessible datasets that can be downloaded from <http://americangut.org>. The full dataset we use contains counts of 27774 OTUs from 8327 fecal samples. Each sample is taken from a unique participant individual; 455 covariates describing various aspects of the subjects are collected, such as their demographical information and dietary behaviors. Moreover, a rooted full binary phylogenetic tree on all OTUs is available.

Although the number of OTUs in the study is large, many cells in the OTU table contain zero or very few counts. For illustration purpose, we focus on the top 50, 75, and 100 OTUs with the largest overall counts across samples. Using different numbers of OTUs allow us to evaluate the inference of increasing dimensionality on the performance of different tests. Note that even with these filtered top OTUs, the data still demonstrate the four challenging features of the OTU counts listed in [Section 1](#). We further narrow down the samples to 561

middle-aged (people in their 30s, 40s and 50s) male Caucasian participants from the west census region to reduce the large variation across different simulation rounds (we will use the full dataset in our later data analysis). Reducing the sample size in the simulation is also necessary due to the speed of the competitors. In each simulation, we randomly divide the data into two roughly equal-sized groups to create the data under the null. The data under the alternative is generated under three scenarios:

- I. **Cross-group difference exists at a single OTU.** This scenario is also considered in ?. In each round of the simulation, we randomly select an OTU and increase its count in the second group by a given percentage p . This induces cross-group differences on the parent node of the selected OTU. In this case, the differences are considered local on the tree. We consider $K = 50, 75$, and 100 with $p = 100\%, 150\%$, and 250% (To rule out obvious simulations with either too weak or too strong signals, we place additional constraints in the random selection. Specifically, we only select OTUs with sample means in the middle 80% range of all the K OTU sample means).
- II. **Cross-group difference exists at multiple OTUs.** Similar to the first scenario, *eight* OTUs are randomly selected and their counts are increased by a given percentage p of the subjects in the second group to create cross-group differences that are more global on the tree. We consider $K = 50, 75$, and 100 with $p = 50\%, 100\%$, and 150% .
- III. **Cross-group difference exists at a chain of nodes in the phylogenetic tree.** Cross-group differences of OTU compositions often cluster into chains as shown in ?? and ?. To see how the graphical structure in BGCR helps increase the power of the test in this situation, we consider a case in which cross-group differences are present in a fixed chain of nodes in the phylogenetic tree. In particular, we focus on the top 100 OTUs ($K = 100$) with their phylogenetic tree shown in [Figure 3](#). Consider a chain of three nodes in the tree as shown in red in [Figure 3](#), we create cross-group

differences at these nodes by systematically modifying the counts of their descendant OTUs. Specifically, we increase the counts of OTU 1, 2 and 3 (marked in **Figure 3**) of all the subjects in the second group by $0.33p$, $0.67p$ and p percent, respectively. We consider the case when $p = 75\%$, 100% , and 125% .

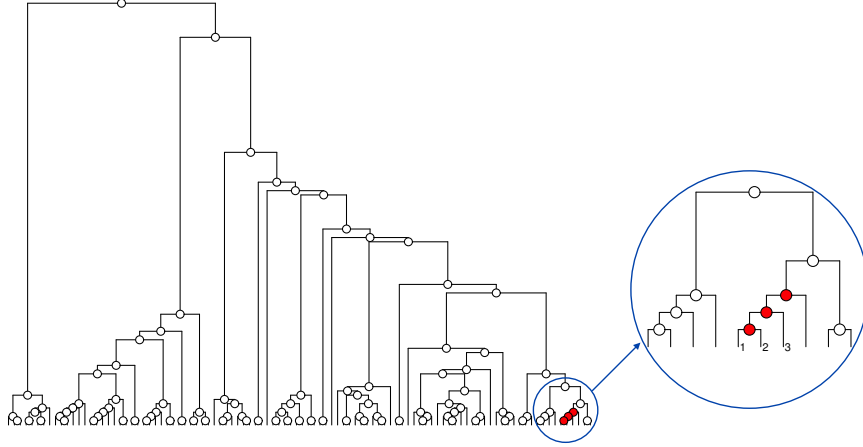


Figure 3: The phylogenetic tree of the top 100 OTUs with the largest overall counts. Each circle represents an internal node; each tip represents an OTU. The chain of nodes with designed cross-group differences in simulation III is marked red. The zoom-in plot gives a better view of these nodes and their descendant OTUs.

In each scenario, we carry out 3000 rounds of simulations. For BGCR, we let $\gamma(A) \sim N(0, 10)$ and $\log_{10} \nu(A) \sim \text{Unif}(-1, 4)$ independently for $A \in \mathcal{I}$. Hyperparameters are chosen according to Section 2.6 with $\text{PrJAP} = 0.5$, $\kappa = 0$, and $\tau_{max} = 6$. We compare BGCR with DTM using the maximum of the single node statistic (DTM-1), using the maximum of the triplet statistic (DTM-3) (?) and the DM test (?).

The ROC plots under the three simulation scenarios with $K = 100$ are shown in ??. Similar results under Scenario II and III with $K = 50$ and $K = 75$ and shown in Figure S1 and Figure S2 in online supplementary materials. Overall, BGCR outperforms the competitors in all these scenarios. In general, the three methods that incorporate phylogenetic information

perform better than the DM test. In Scenario I and II, ?? and Figure S1 show that DTM-1 and DTM-3 work better when K is small. Compared with BGCR, these tests perform worse with the increase of K . This is due to the inherent features of the microbiome data. When K is small, the samples typically contain large counts of many OTUs included in the study and are likely to be described by the DM-type models. When K becomes larger, the OTU table gets sparser, making it harder for DTM-1 and DTM-3 to estimate the right dispersion parameters. In contrast, the dispersion parameters are integrated out rather than estimated in BGCR.

To illustrate how borrowing information among neighboring nodes helps increase the power of the node-specific tests, we compare BGCR with BCR in Scenario III. Under the null, the estimated γ by the empirical Bayes procedure in BGCR are zero in all but a few of the simulations (Figure S3 left). Therefore, BGCR essentially degenerates to BCR and does not increase the chance of false positives (Figure S3 middle and right). Under the alternative, by introducing positive dependencies among the local tests, BGCR obtains larger PJAPs and thus provides a better chance to correctly reject the null. To see how this would influence the decision procedure, we show the proportion of tests that reject the null under different decision thresholds L in Figure S4 in online supplementary materials. When $0.5 < L < 0.9$, the gain by introducing the dependencies among tests is remarkable. Figure S5 in online supplementary materials shows the estimated γ 's under the alternatives. When the cross-group difference increases, the chaining pattern gets stronger, causing the estimated γ to increase. Figure S6 in online supplementary materials shows the corresponding PJAPs of the two models under comparison.

3.2 Visualizing the cross-group differences with BGCR

In many applications, characterizing and visualizing the cross-group difference is often of more interests than merely testing its existence. As mentioned in Section 2.5, this can

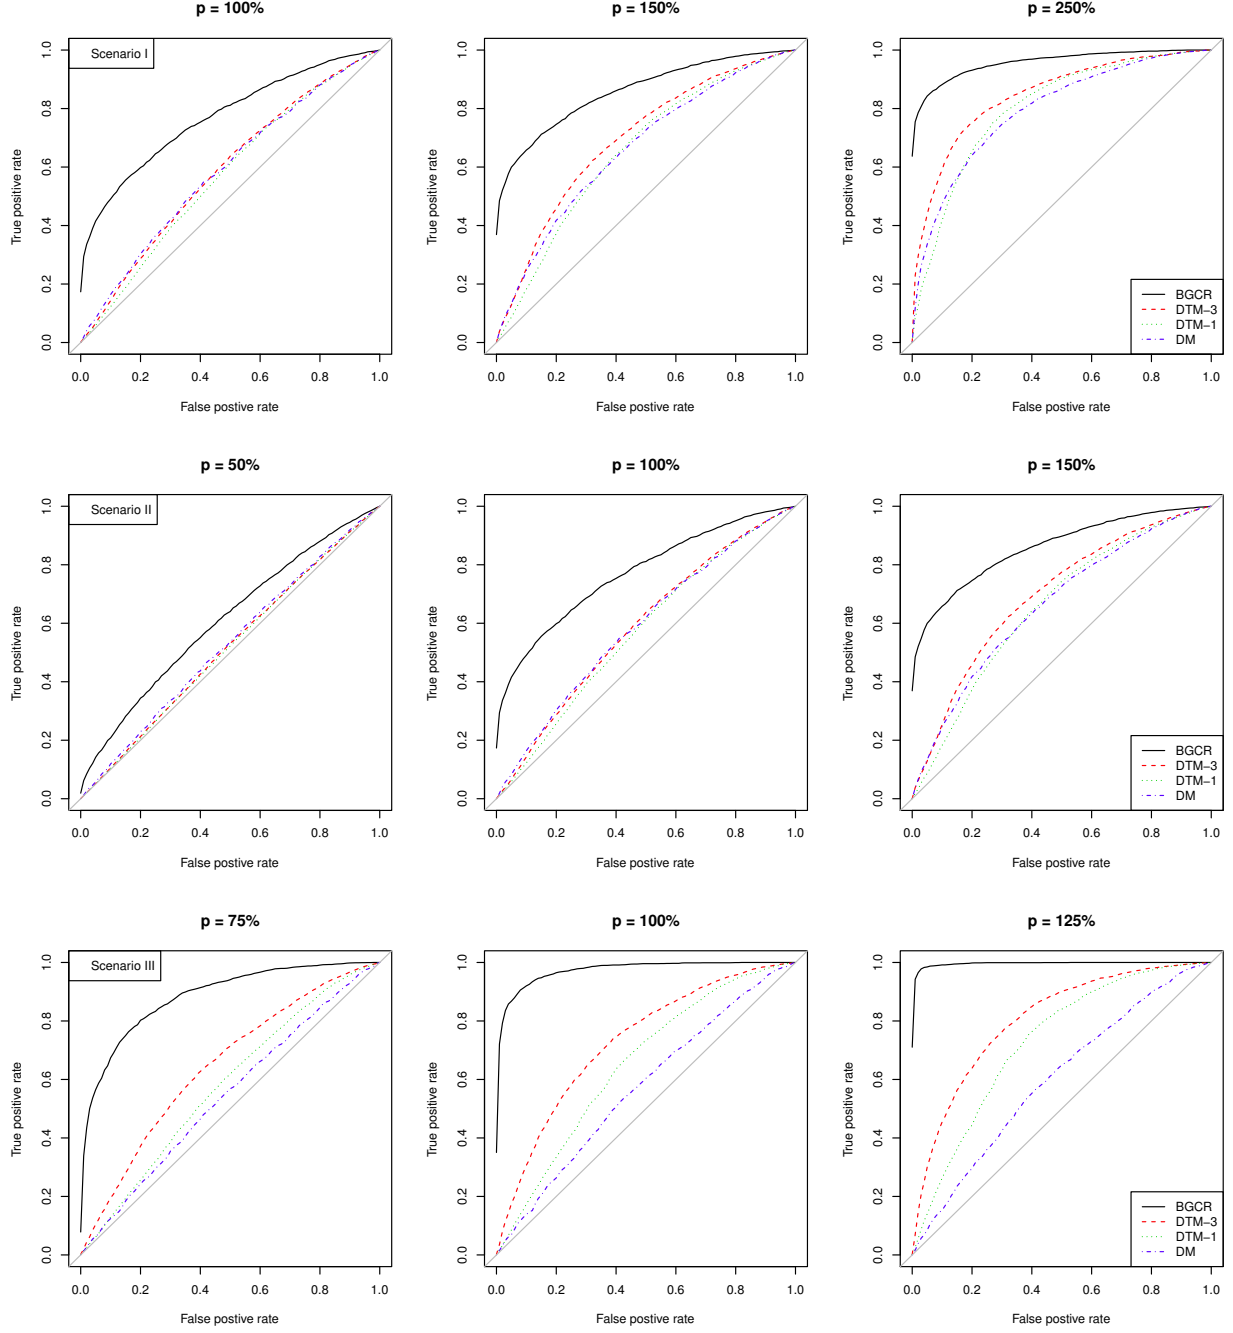


Figure 4: ROC curves for Scenario I (top row), II (middle row) and III (bottom row) with $K = 100$. The columns are indicated by the percent of counts increased in the second group (p).

be achieved by marking the PMAPs along the phylogenetic tree. As an example, consider a specific round of simulation in Scenario III. In this simulation, $\text{PJAP}_{\text{BGCR}} = 0.910$, $\text{PJAP}_{\text{BCR}} = 0.651$, $\hat{\gamma} = 5.521$. ?? plots the PMAPs obtained by the two models along the phylogenetic tree (only shows the relevant part of the tree). It is clear that BGCR benefits from the graphical structure and reveals the designed chaining pattern of the cross-group differences (see Figure 3).

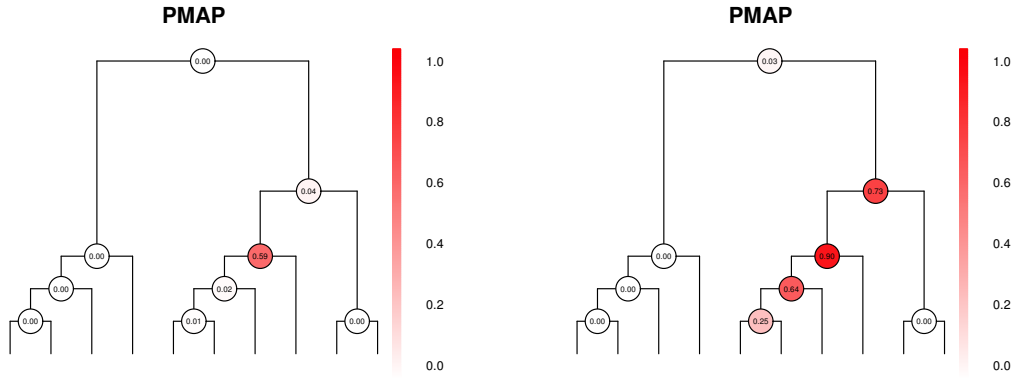


Figure 5: PMAPs from BCR and BGCR under a specific simulation case in Scenario III. Only nodes in the zoom-in subtree in Figure 3 are shown since the PMAPs at other nodes are essentially zero. Left: PMAPs obtained by BCR; Right: PMAPs obtained by BGCR. The number in each node denotes its PMAP. We also color the nodes by their PMAPs to give a better visualization.

3.3 Adjusting for covariates

In practice, OTU compositions can depend on various factors. When the data are gathered from observational studies, unadjusted covariates can lead to false positives. Even when the data do arise from randomized experiments, incorporating relevant covariates can improve the statistical power for identifying cross-group differences. We illustrate with another simulation scenario the necessity of incorporating confounders in the testing procedure:

IV. Cross-group difference exists at a single OTU with an unbalanced confounder. We start from the same American Gut dataset as in the previous simulation scenarios. Now we randomly select 200 *male* and 200 *female* Caucasians, all middle aged and from the west region (the reason we take a random sample instead of using the full dataset is the large computational burden brought by the competitors). To purposefully make gender an unbalanced confounder, we select an OTU (OTU ‘4352657’, denoted by ω_c) and increase the counts of this OTU for all male participants by 175%. We then randomly select 160 males and 40 females into the first group and put the rest into the second group to get the data under the null. To simulate data under the alternative, we select another OTU (OTU ‘4481131’, denoted by ω_s) and increase the count of this OTU in the second group by 175%.

We carry out 750 simulations with $K = 50$ and compare BGCR with DTM-3 and DM (we do not simulate 3000 rounds or consider $K = 75$ or 100 since some competitors are too slow). ?? shows the histograms of the statistics these methods used for decision making under the null—PJAP for BGCR, p -values (or equivalently, $1 - p$ -values) for DTM-3 and DM. From ?? (ii)-(iv), it is clear that all three methods fail in controlling false positives without adjusting for the “gender” of the participants. After incorporating “gender” in each node-specific regression, BGCR is able to keep the PJAPs under the null at a reasonable level (?? (i)).

Adjusting for confounders is also important in the characterization of the cross-group differences. For example, consider a specific round of simulation in Scenario IV. ?? marks the PMAPs reported by BGCR on the phylogenetic tree under the null and the alternative, with or without “gender” included. Under the null, ignoring “gender” leads to a false discovery at the parent of ω_c and also a false rejection of the hypothesis that no cross-group difference exists at threshold 0.95. Under the alternative, the PJAPs reported by BGCR are closed to 1 with or without adjusting for “gender”, suggesting that incorporating “gender” in the model

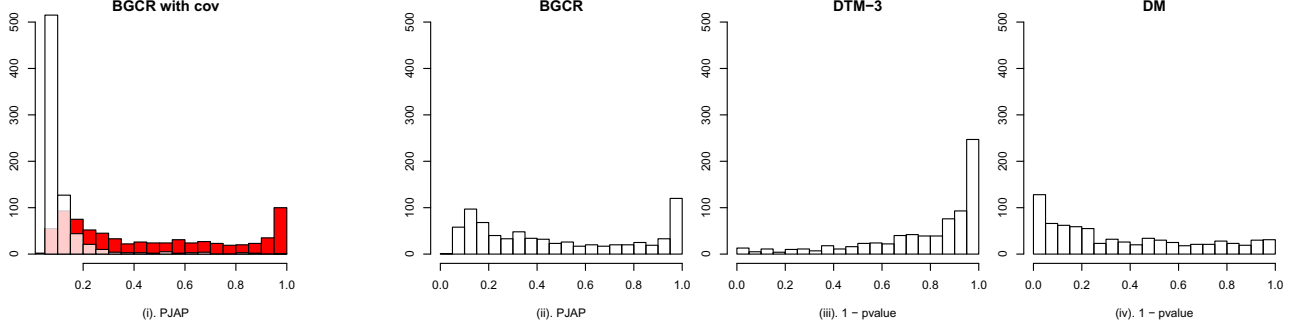


Figure 6: Histograms of the statistics used for decision making under the null. Plot (i) shows the histogram of PJAPs under the null (white) and under the alternative (red) with “gender” adjusted. The other three plots present the histograms of the decision-making statistics of the three models under comparison under the null without adjusting for “gender”.

is likely to keep the decision about the original hypothesis unchanged. However, failing to adjust for “gender” contaminates the pattern of the cross-group differences (?? (iii)-(iv)).

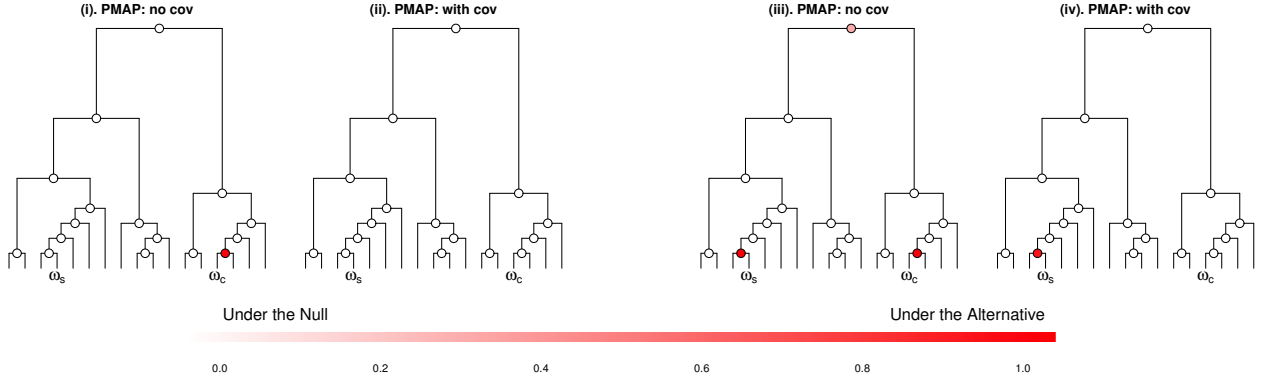


Figure 7: PMAPs of a specific simulation in Scenario IV. (i) and (ii): PMAPs under the null; (iii) and (iv): PMAPs under the alternative (only a subtree is shown since PMAPs on other nodes are essentially zero). The nodes are colored by their PMAPs. Under the alternative, true cross-group differences are only present at the parent of ω_s . A false discovery is made at the parent of ω_c when “gender” is not adjusted for.

4 Application to the American Gut project

4.1 Cross-group comparison of OTU compositions

In this section, we apply BGCR to the same dataset from the American Gut project as in [Section 3](#). The American Gut project collects many binary or ordered categorical covariates of the participants that can be used to divide them into groups. In particular, we focus on comparing the OTU compositions of groups of participants categorized by their dietary habits, which can be characterized by their frequencies of consuming different kinds of food. We consider eight types of food: *alcohol*, *diary*, *fruit*, *meat/egg*, *seafood*, *sweets*, *vegetable* and *whole grain*. For each kind of food, the dataset records each participant’s frequency of consumption per week as a categorical variable with five categories: *Never*, *Occasionally (1-2 times per week)*, *Rarely (less than once per week)*, *Regularly (3-5 times per week)* and *Daily*. We dichotomize the categories based on whether the food is consumed less than 3 times per week and use this binary variable to divide the participants into two groups.

In our analysis, we only use the top 100 OTUs with the largest overall counts across samples. This choice is made mainly for quality control. For OTUs containing nil or only very few counts over all samples, their counts are largely influenced by sequencing errors and are highly unreliable. Using these OTUs is likely to contaminate the analysis by bringing in pure noises that are hard to quantify. We note that prescreening the OTUs to drop those who do not have substantial counts in any samples is a common recommended practice in the literature (see for example ??).

Besides diet, there are other possible factors that may influence the OTU composition. Including these factors in the model could help reduce the chance of false discoveries and improve the power of the inference. In our analysis, we consider five such covariates: *the participant’s biological sex*, *has the participant been diagnosed with diabetes*, *does the participant have IBD (inflammatory bowel disease)*, *has the participant used antibiotic in the past year*,

and does the participant consume probiotic less than 3 times per week. All five covariates are binary. In addition to adjusting for these non-dietary covariates when investigating a certain kind of food, dietary habits of other foods are also natural candidates for confounders to be included in the analysis, especially when they are unbalanced across the two groups. For example, ?? shows the unbalancedness of vegetable and grain consumption of participants across the two groups defined by fruit consumption. Since the covariate information in the study is collected by asking the participants to take a non-compulsory questionnaire, some covariates are missing for certain participants. For simplicity, we ignore all participants with missing values in any variable mentioned above and perform a complete-case analysis. Moreover, we restrict our samples to participants with age in range $20 \sim 69$ and BMI $18.5 \sim 30$ to disregard potential outliers. The resulting sample size in each comparison is shown in ??.

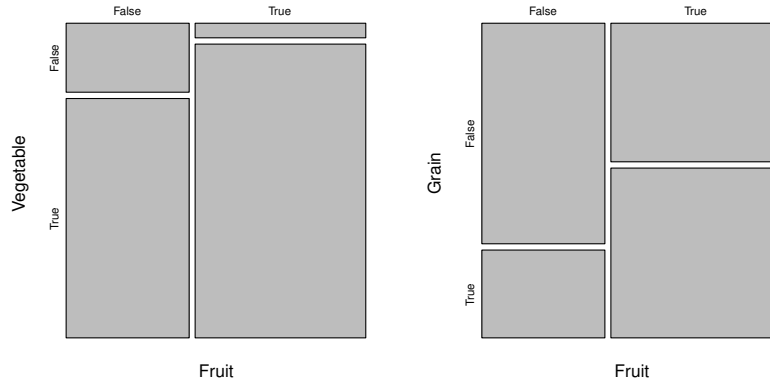


Figure 8: Unbalancedness of vegetable and grain consumption in the two groups defined by fruit consumption. The area of each block represents the proportion of the counts in the corresponding cell of the 2×2 contingency table.

For each type of food, we use BGCR to compare the group-level OTU composition across the two resulting groups. In each comparison, we consider BGCR with different sets of covariates adjusted—first no covariate, then the five non-dietary covariates, and finally the non-dietary covariates as well as dietary habits on other foods. The $(1 - \text{PJAP})$'s reported

Diet	Group size		1-PJAP		
	False	True	No cov	Non-diet cov	All cov ($\hat{\tau}$)
Alcohol	2099	995	0.134	0.345	0.555 (0)
Diary	1733	1361	1.27×10^{-2}	6.97×10^{-2}	0.517 (0)
Fruit	1296	1798	6.86×10^{-7}	1.29×10^{-8}	0.259 (3.93)
Meat/Egg	798	2296	0.867	0.857	0.757 (0)
Seafood	2637	457	1.99×10^{-2}	9.11×10^{-3}	0.166 (0)
Sweets	2005	1089	0.200	0.347	0.912 (0)
Vegetable	375	2719	$< 1.00 \times 10^{-8}$	$< 1.00 \times 10^{-8}$	$< 1.00 \times 10^{-8}$ (5.13)
Whole grain	1735	1359	$< 1.00 \times 10^{-8}$	$< 1.00 \times 10^{-8}$	3.85×10^{-7} (5.63)

Table 2: (1-PJAP)’s for testing OTU compositions of different dietary habits. In the last column, red cells indicate rejections of the global null under threshold $L = 0.5$. The estimated $\hat{\tau}$ when all the covariates are adjusted is shown in the parentheses in the last column.

by BGCR under each comparison are shown in ??.

When all the covariates are adjusted, the test rejects the global null at threshold $L = 0.5$ when the grouping food is *fruit*, *seafood*, *vegetable*, or *whole grain*. This is consistent with the findings that dietary habits are closely related to gut microbiome compositions in general (????).

Our analysis also shows that adjusting for relevant covariates, which is missing in previous analyses, can significantly change the testing results and can reverse the testing decision on the null hypothesis. For example, when no covariate or only non-dietary covariates are adjusted, the null of no cross-group difference is likely to be rejected for *diary* and *sweets*, which leads to the same conclusion as in ?. However, once other dietary covariates are included, these two diets tend to have no significant association with the microbiome composition. Typically, the posterior probability of the null (1 - PJAP) increases when more covariates are incorporated in the model. This reflects the complex nature of the microbial community that it is associated with various factors (?).

Similar phenomena can be seen on the more refined PMAP scale. For the four comparisons that reject the null, we visualize the cross-group differences by marking the PMAPs on

the phylogenetic tree (??). Similar plots of PMAPs under BGCR with no covariate or only non-dietary covariates can be found in online supplementary materials (Figure S7 and Figure S8). These plots show that adjusting for relevant covariates greatly changes the characterization of the cross-group differences. Many significant nodes reported by BGCR with some covariates unadjusted turn out to be false positives when all the covariates are included.

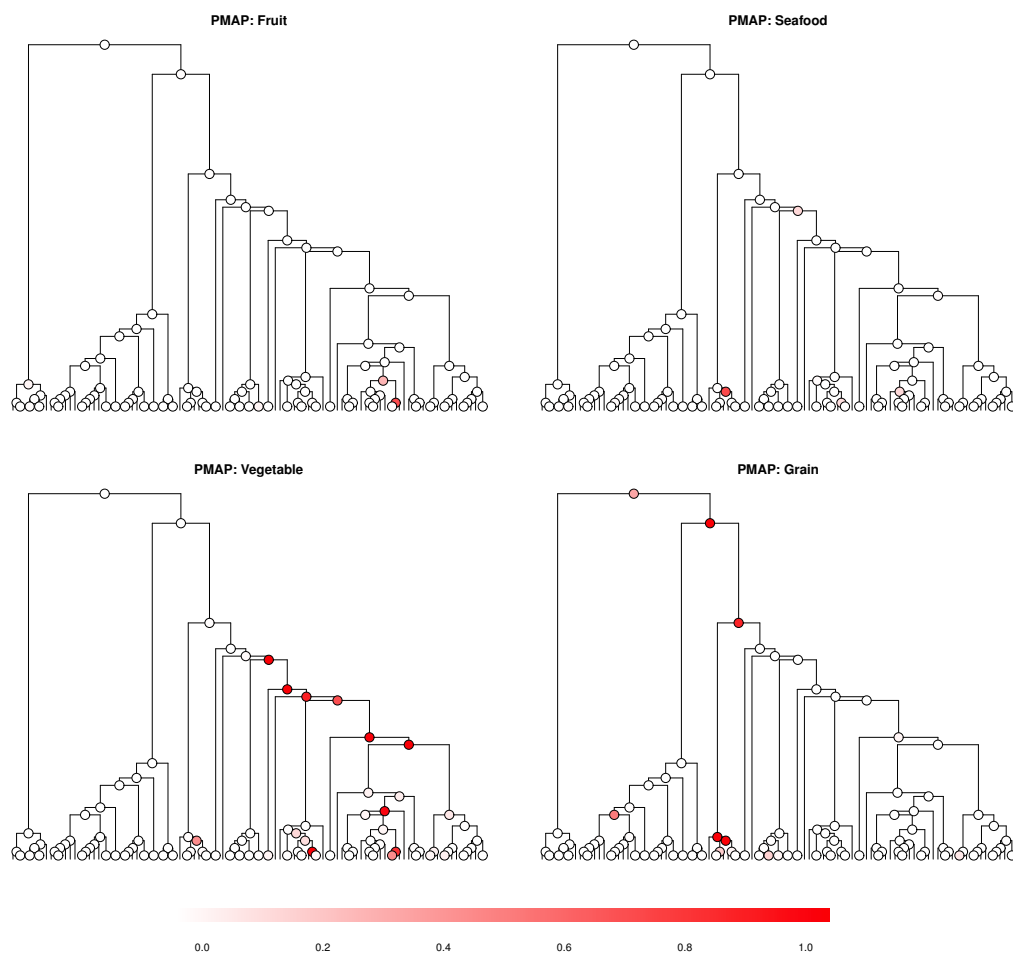


Figure 9: PMAPs for the four comparisons that reject the global null. The nodes are colored by PMAPs reported by BGCR with both non-dietary covariates and dietary covariates adjusted.

4.2 BGCR vs BCR

In this section, we evaluate the performance of BGCR compared with BCR in our data application. In BGCR, the global parameter τ controls the level of dependencies among the node-specific tests. We note in Section 2.6 that when $\tau = 0$, BGCR degenerates to BCR. This nested feature together with our empirical Bayes strategy to estimate τ allow BGCR to decide whether the dependencies are necessary based on the data—when the estimated $\hat{\tau} = 0$, BGCR simply performs independent tests without introducing the graphical structure.

Figure 1 shows $\hat{\tau}$ in each comparison when BGCR with all covariates is used. In all four comparisons that the null is accepted at threshold $L = 0.5$, $\hat{\tau} = 0$. In these cases, no node is significant and the “signals” are trivially independent. For *fruit*, *vegetable* and *whole grain*, $\hat{\tau} > 0$, suggesting that cross-group differences might cluster into chains. To see whether it is necessary to introduce the graphical structure among the node-specific tests, one can test $H_0^g : \tau = 0$ vs $H_1^g : \tau \neq 0$. For example, when the prior $\tau \sim \text{Uniform}(0, 6)$ is used, the Bayes factor comparing H_1^g and H_0^g for *fruit*, *vegetable*, *seafood* and *whole grain* are 1.022, 0.830, 138.699 and 6.922, respectively. For *vegetable* and *whole grain*, the cross-group differences cluster into chains (Figure 1) and thus it is necessary to introduce the graphical structure. We can also look at the posterior on τ in the four comparisons that reject the null as shown in Figure 2. For *vegetable* and *whole grain*, the signals demonstrate a strong chaining pattern and the posteriors on τ concentrate to some positive values. For *fruit* and *seafood*, the cross-group differences are local to certain nodes. Therefore, not much information about the dependencies of the node-specific tests is provided by the data and the posteriors on τ are very similar to the prior.

5 Concluding remarks

We have introduced a Bayesian framework for testing the existence of cross-group differences in the OTU composition. The original testing problem is transformed into a series

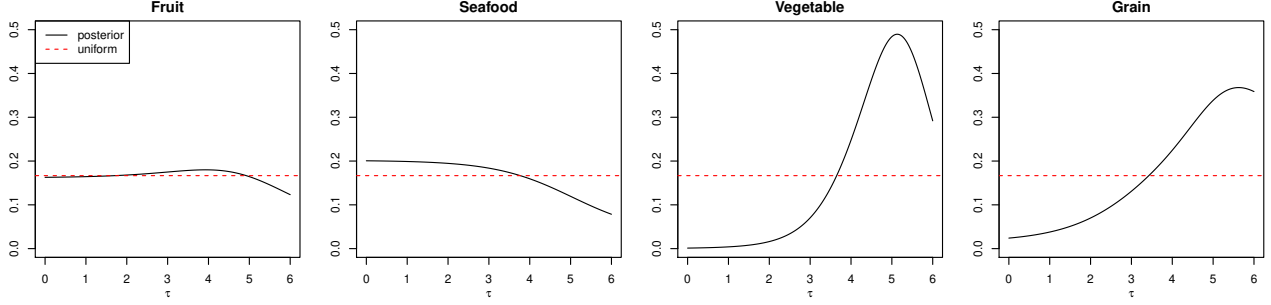


Figure 10: Posteriors on τ with the uniform prior. Shown for the four comparisons that reject the null at threshold $L = 0.5$.

of dependent node-specific tests linked together by the phylogenetic tree. A key feature of the Bayesian formulation is that the dispersion parameters are integrated out, providing more robust inference than the methods that use point estimates of the within-group variation. We use numerical integration to evaluate the Bayes factors, circumventing Monte Carlo simulations to give fast inference. By introducing dependencies among the local tests, information sharing is allowed among neighboring nodes. This further improves the power of the test when the cross-group differences cluster into chains along the tree, which is often observed in practice. We derive an exact message passing algorithm that is scalable to the size of the tree to carry out inference under this information borrowing.

Because microbiome data are often collected by observational studies, it is important to adjust for possible confounders in testing. This is crucial in reducing the chance of false discoveries and the residual overdispersion. Our model achieves this goal by incorporating the covariates via a regression model, which adds little computational burdens to the Bayesian testing framework. A relevant problem is to select covariates in the testing scenario. Generally, this must be done with caution since incorporating standard variable selection procedures can substantially affect or even invalidate the meaning of the testing result on the two-group difference due to possible confounding. Readers may refer to Online Supplementary Materials C for more details on this along with some numerical examples.

Besides testing, our model gives a full probabilistic characterization of the cross-group differences that can be naturally visualized. Note that although we focus on two group comparisons in this paper, the proposed method can be generalized to multiple group settings with little extra effort.

Software

The DTM methods and the DM test are implemented with the R packages `PhyloScan` and `HMP`. An R package for BGCR is freely available at <https://github.com/MaStatLab/BGCR> and will be soon available on CRAN. Source code for the numerical examples are available at https://github.com/MaStatLab/BGCR_analysis.

Acknowledgment

LM’s research is partly supported by NSF grants DMS-1612889, DMS-1749789, and a Google Faculty Research Award. Part of this work was completed when JM was supported by a Duke Forge Fellowship.

References

- Barbieri, M. M., J. O. Berger, et al. (2004). Optimal predictive model selection. *The annals of statistics* 32(3), 870–897.
- Berger, J. O., L. R. Pericchi, J. Ghosh, T. Samanta, F. De Santis, J. Berger, and L. Pericchi (2001). Objective bayesian methods for model selection: introduction and comparison. *Lecture Notes-Monograph Series*, 135–207.
- Callahan, B. J., P. J. McMurdie, and S. P. Holmes (2017). Exact sequence variants should replace operational taxonomic units in marker-gene data analysis. *The ISME journal* 11(12), 2639.

- Caporaso, J. G., J. Kuczynski, J. Stombaugh, K. Bittinger, F. D. Bushman, E. K. Costello, N. Fierer, A. G. Peña, J. K. Goodrich, J. I. Gordon, et al. (2010). Qiime allows analysis of high-throughput community sequencing data. *Nature methods* 7(5), 335–336.
- Chen, H. and J. H. Friedman (2017). A new graph-based two-sample test for multivariate and object data. *Journal of the American statistical association* 112(517), 397–409.
- Chen, J. and H. Li (2013). Variable selection for sparse dirichlet-multinomial regression with an application to microbiome data analysis. *The annals of applied statistics* 7(1).
- Crouse, M. S., R. D. Nowak, and R. G. Baraniuk (1998). Wavelet-based statistical signal processing using hidden markov models. *IEEE Transactions on signal processing* 46(4), 886–902.
- David, L. A., C. F. Maurice, R. N. Carmody, D. B. Gootenberg, J. E. Button, B. E. Wolfe, A. V. Ling, A. S. Devlin, Y. Varma, M. A. Fischbach, et al. (2014). Diet rapidly and reproducibly alters the human gut microbiome. *Nature* 505(7484), 559–563.
- Dennis, S. Y. (1991). On the hyper-dirichlet type 1 and hyper-liouville distributions. *Communications in Statistics-Theory and Methods* 20(12), 4069–4081.
- Dennis, S. Y. (1996). A bayesian analysis of tree-structured statistical decision problems. *Journal of statistical planning and inference* 53(3), 323–344.
- George, E. I. and R. E. McCulloch (1997). Approaches for bayesian variable selection. *Statistica sinica*, 339–373.
- Goedert, J. J., X. Hua, G. Yu, and J. Shi (2014). Diversity and composition of the adult fecal microbiome associated with history of cesarean birth or appendectomy: analysis of the american gut project. *EBioMedicine* 1(2-3), 167–172.
- Graf, D., R. Di Cagno, F. Fåk, H. J. Flint, M. Nyman, M. Saarela, and B. Watzl (2015). Contribution of diet to the composition of the human gut microbiota. *Microbial ecology in health and disease* 26(1), 26164.
- Grantham, N. S., B. J. Reich, E. T. Borer, and K. Gross (2017). Mimix: a bayesian

- mixed-effects model for microbiome data from designed experiments. *arXiv preprint arXiv:1703.07747*.
- Held, L., D. S. Bové, I. Gravestock, et al. (2015). Approximate bayesian model selection with the deviance statistic. *Statistical Science* 30(2), 242–257.
- Hildebrandt, M. A., C. Hoffmann, S. A. Sherrill-Mix, S. A. Keilbaugh, M. Hamady, Y.-Y. Chen, R. Knight, R. S. Ahima, F. Bushman, and G. D. Wu (2009). High-fat diet determines the composition of the murine gut microbiome independently of obesity. *Gastroenterology* 137(5), 1716–1724.
- Holmes, C. C., F. Caron, J. E. Griffin, D. A. Stephens, et al. (2015). Two-sample bayesian nonparametric hypothesis testing. *Bayesian Analysis* 10(2), 297–320.
- Koller, D. and N. Friedman (2009). *Probabilistic graphical models: principles and techniques*. MIT press.
- La Rosa, P. S., J. P. Brooks, E. Deych, E. L. Boone, D. J. Edwards, Q. Wang, E. Sodergren, G. Weinstock, and W. D. Shannon (2012). Hypothesis testing and power calculations for taxonomic-based human microbiome data. *PloS one* 7(12), e52078.
- Lauritzen, S. L. and D. J. Spiegelhalter (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 157–224.
- Li, H. (2015). Microbiome, metagenomics, and high-dimensional compositional data analysis. *Annual Review of Statistics and Its Application* 2, 73–94.
- Li, Y. and M. A. Clyde (2015). Mixtures of g-priors in generalized linear models. *arXiv preprint arXiv:1503.06913*.
- Liang, F., R. Paulo, G. Molina, M. A. Clyde, and J. O. Berger (2008). Mixtures of g priors for bayesian variable selection. *Journal of the American Statistical Association* 103(481), 410–423.
- Ma, L. and J. Soriano (2018). Analysis of distributional variation through graphical multi-

- scale beta-binomial models. *Journal of Computational and Graphical Statistics*, 1–13.
- McDonald, D., M. Hornig, C. Lozupone, J. Debelius, J. A. Gilbert, and R. Knight (2015). Towards large-cohort comparative studies to define the factors influencing the gut microbial community structure of ASD patients. *Microbial ecology in health and disease* 26.
- Müller, P., G. Parmigiani, and K. Rice (2006). Fdr and bayesian multiple comparisons rules.
- Müller, P., G. Parmigiani, C. Robert, and J. Rousseau (2004). Optimal sample size for multiple testing: the case of gene expression microarrays. *Journal of the American Statistical Association* 99(468), 990–1001.
- Qin, J., Y. Li, Z. Cai, S. Li, J. Zhu, F. Zhang, S. Liang, W. Zhang, Y. Guan, D. Shen, et al. (2012). A metagenome-wide association study of gut microbiota in type 2 diabetes. *Nature* 490(7418), 55–60.
- Ren, B., S. Bacallado, S. Favaro, T. Vatanen, C. Huttenhower, and L. Trippa (2017). Bayesian nonparametric mixed effects models in microbiome data analysis. *arXiv preprint arXiv:1711.01241*.
- Rothschild, D., O. Weissbrod, E. Barkan, A. Kurilshikov, T. Korem, D. Zeevi, P. I. Costea, A. Godneva, I. N. Kalka, N. Bar, et al. (2018). Environment dominates over host genetics in shaping human gut microbiota. *Nature* 555(7695), 210.
- Singh, R. K., H.-W. Chang, D. Yan, K. M. Lee, D. Ucmak, K. Wong, M. Abrouk, B. Farahnik, M. Nakamura, T. H. Zhu, et al. (2017). Influence of diet on the gut microbiome and implications for human health. *Journal of translational medicine* 15(1), 73.
- Sogin, M. L., H. G. Morrison, J. A. Huber, D. M. Welch, S. M. Huse, P. R. Neal, J. M. Arrieta, and G. J. Herndl (2006). Microbial diversity in the deep sea and the underexplored “rare biosphere”. *Proceedings of the National Academy of Sciences* 103(32), 12115–12120.
- Soriano, J. and L. Ma (2017). Probabilistic multi-resolution scanning for two-sample differences. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(2), 547–572.

- Tang, Y., L. Ma, D. L. Nicolae, et al. (2018). A phylogenetic scan test on a dirichlet-tree multinomial model for microbiome data. *The Annals of Applied Statistics* 12(1), 1–26.
- Tang, Y. and D. L. Nicolae (2017). Mixed effect dirichlet-tree multinomial for longitudinal microbiome data and weight prediction. *arXiv preprint arXiv:1706.06380*.
- Turnbaugh, P. J., R. E. Ley, M. A. Mahowald, V. Magrini, E. R. Mardis, and J. I. Gordon (2006). An obesity-associated gut microbiome with increased capacity for energy harvest. *nature* 444(7122), 1027–131.
- Wadsworth, W. D., R. Argiento, M. Guindani, J. Galloway-Pena, S. A. Shelburne, and M. Vannucci (2017). An integrative bayesian dirichlet-multinomial regression model for the analysis of taxonomic abundances in microbiome data. *BMC bioinformatics* 18(1), 94.
- Wang, T. and H. Zhao (2017). A dirichlet-tree multinomial regression model for associating dietary nutrients with gut microorganisms. *Biometrics*.
- Wellman, M. P. and M. Henrion (1993). Explaining ‘explaining away’. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3), 287–292.
- Wu, G. D., J. Chen, C. Hoffmann, K. Bittinger, Y.-Y. Chen, S. A. Keilbaugh, M. Bewtra, D. Knights, W. A. Walters, R. Knight, et al. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science* 334(6052), 105–108.
- Xia, F., J. Chen, W. K. Fung, and H. Li (2013). A logistic normal multinomial regression model for microbiome compositional data analysis. *Biometrics* 69(4), 1053–1063.

Supplementary Materials

A Computational strategies

The crux of doing inference with BCR is to calculate the marginal likelihoods $M_s(A)$. We use the computational strategies in ? to compute the integrals. Specifically, for fixed ν , the inner integrals on the regression parameters are evaluated based on Laplace approximation; then the outer integral on ν is calculated with finite Riemann approximation.

With a bit abuse of the notations, in this section, we use $\beta(A)$ to denote the ‘active’ regression coefficient in (2.10). That is, under the null, $\beta(A)$ is just the $\beta(A)$ in (2.10); under the alternative, $\beta(A)$ denotes $(\beta(A)^\top, \gamma(A)^\top)^\top$. \mathbf{x}_{ij} is redefined to be the ‘active’ covariates in the same sense. With these notations, we have $g(\theta_{\mathbf{x}_{ij}}(A)) = \mathbf{x}_{ij}^\top \beta(A)$ and $\theta_{ij}(A) \mid \mathbf{x}_{ij}, \nu(A) \sim \text{Beta}(\theta_{\mathbf{x}_{ij}}(A)\nu(A), (1 - \theta_{\mathbf{x}_{ij}}(A))\nu(A))$ for each local beta-binomial model.

The computational strategies are the same under both hypotheses. Let $\pi_A(\beta)$ be the prior density of $\beta(A)$ under either hypothesis. For fixed ν in the support of $G_A(\nu)$, by the Laplace approximation, the inner integral is

$$\begin{aligned} L_\nu(A) &= \int \prod_{i=1}^2 \prod_{j=1}^{n_i} \mathcal{L}_{BB}(g^{-1}(\mathbf{x}_{ij}^\top \beta), \nu \mid y_{ij}(A_l), y_{ij}(A_r)) \pi_A(\beta) d\beta \\ &= \int \exp \left\{ \sum_{i,j} \log \mathcal{L}_{BB}(g^{-1}(\mathbf{x}_{ij}^\top \beta), \nu \mid y_{ij}(A_l), y_{ij}(A_r)) + \log \pi_A(\beta) \right\} d\beta \\ &= \int \exp \{h_\nu(\beta)\} d\beta \\ &\approx \exp \left\{ h_\nu(\hat{\beta}_\nu) \right\} \cdot (2\pi)^{d/2} \cdot | - H_\nu(\hat{\beta}_\nu) |^{-1/2} \\ &= \hat{L}_\nu(A) \end{aligned}$$

where $h_\nu(\beta) = \sum_{i,j} \log \mathcal{L}_{BB}(g^{-1}(\mathbf{x}_{ij}^\top \beta), \nu \mid y_{ij}(A_l), y_{ij}(A_r)) + \log \pi_A(\beta)$, $\hat{\beta}_\nu$ is the maximizer of $h_\nu(\beta)$, $H_\nu(\hat{\beta}_\nu)$ is the Hessian matrix of $h_\nu(\beta)$ at $\beta = \hat{\beta}_\nu$. d is the degrees of freedom of β , which is $(p+2)$ under the alternative and $(p+1)$ under the null. We describe a Newton-

Raphson algorithm to solve for $\hat{\beta}_\nu$ below. The log-likelihood function is strictly log-concave and the Newton-Raphson method generally converges after only a few iterations. Finally, to get $M_s(A)$, we compute the outer integral on ν , $\int L_\nu(A) dG_A(\nu)$, with finite Riemann approximations. Specifically, after calculating $\hat{L}_\nu(A)$ at a grid of ν 's: $\nu_1, \nu_2, \dots, \nu_M$, we have

$$\int L_\nu(A) dG_A(\nu) \approx \sum_{m=2}^M \hat{L}_{\nu_m}(A) (G_A(\nu_m) - G_A(\nu_{m-1})).$$

Newton-Raphson for $\hat{\beta}_\nu$. In this subsection, we shall fix our attention on a specific node A and suppress the ‘(A)’ in the notations. Moreover, we let $y_{ij}(A_l) = t_{ij}$ for simplicity and express the local Beta-Binomial regression model on A as

$$t_{ij} \sim \text{Binomial}(y_{ij}, \theta_{ij}), \quad \theta_{ij} \sim \text{Beta}(\theta_{\mathbf{x}_{ij}}\nu, (1 - \theta_{\mathbf{x}_{ij}})\nu), \quad \text{and} \quad g(\theta_{\mathbf{x}}) = \mathbf{x}^\top \boldsymbol{\beta}.$$

The contribution to the log marginal likelihood from the j -th observation in group i is

$$\begin{aligned} l_{ij} &= \log \mathcal{L}_{BB}(g^{-1}(\mathbf{x}_{ij}^\top \boldsymbol{\beta}), \nu \mid t_{ij}, y_{ij} - t_{ij}) \\ &= \log \Gamma(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) + \log \Gamma((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \log \Gamma(\nu + y_{ij}) \\ &\quad - \log \Gamma(\theta_{\mathbf{x}_{ij}}\nu) - \log \Gamma((1 - \theta_{\mathbf{x}_{ij}})\nu) + \log \Gamma(\nu). \end{aligned}$$

Taking the first derivative w.r.t. $\boldsymbol{\beta}$,

$$\frac{\partial l_{ij}}{\partial \boldsymbol{\beta}} = \frac{\partial l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}} \cdot \frac{\partial \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}} \cdot \frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}}$$

where $\eta_{ij} = \mathbf{x}_{ij}^\top \boldsymbol{\beta}$. Now with ϕ denoting the digamma function,

$$\frac{\partial l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}} = \nu [\phi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) - \phi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \phi(\theta_{\mathbf{x}_{ij}}\nu) + \phi((1 - \theta_{\mathbf{x}_{ij}})\nu)].$$

With the logit link, $\theta_{\mathbf{x}_{ij}} = g^{-1}(\eta_{ij}) = 1/(1 + e^{-\eta_{ij}})$, and

$$\frac{\partial \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}} = (g^{-1})'(\eta_{ij}) = \theta_{\mathbf{x}_{ij}}(1 - \theta_{\mathbf{x}_{ij}}).$$

Thus

$$\frac{\partial l_{ij}}{\partial \boldsymbol{\beta}} = \nu \theta_{\mathbf{x}_{ij}}(1 - \theta_{\mathbf{x}_{ij}}) [\phi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) - \phi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \phi(\theta_{\mathbf{x}_{ij}}\nu) + \phi((1 - \theta_{\mathbf{x}_{ij}})\nu)] \mathbf{x}_{ij}.$$

The second derivative of l_{ij} w.r.t. $\boldsymbol{\beta}$ is

$$\frac{\partial^2 l_{ij}}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^\top} = \frac{\partial^2 l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}^2} \cdot \left(\frac{\partial \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}} \right)^2 \cdot \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}} \right) \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}^\top} \right) + \frac{\partial l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}} \cdot \frac{\partial^2 \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}^2} \cdot \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}} \right) \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}^\top} \right) + \frac{\partial l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}} \cdot \frac{\partial \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}} \frac{\partial^2 \eta_{ij}}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^\top}.$$

The third term on the right-hand side is equal to zero. With ψ being the trigamma function,

$$\frac{\partial^2 l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}^2} = \nu^2 [\psi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) + \psi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \psi(\theta_{\mathbf{x}_{ij}}\nu) - \psi((1 - \theta_{\mathbf{x}_{ij}})\nu)].$$

Thus the first term is

$$\begin{aligned} & \frac{\partial^2 l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}^2} \cdot \left(\frac{\partial \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}} \right)^2 \cdot \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}} \right) \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}} \right)^\top \\ &= \nu^2 [\psi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) + \psi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \psi(\theta_{\mathbf{x}_{ij}}\nu) - \psi((1 - \theta_{\mathbf{x}_{ij}})\nu)] \theta_{\mathbf{x}_{ij}}^2 (1 - \theta_{\mathbf{x}_{ij}})^2 \mathbf{x}_{ij} \mathbf{x}_{ij}^\top. \end{aligned}$$

The second term, which has expectation zero, is

$$\begin{aligned} & \frac{\partial l_{ij}}{\partial \theta_{\mathbf{x}_{ij}}} \cdot \frac{\partial^2 \theta_{\mathbf{x}_{ij}}}{\partial \eta_{ij}^2} \cdot \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}} \right) \left(\frac{\partial \eta_{ij}}{\partial \boldsymbol{\beta}} \right)^\top \\ &= \nu [\phi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) - \phi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \phi(\theta_{\mathbf{x}_{ij}}\nu) + \phi((1 - \theta_{\mathbf{x}_{ij}})\nu)] \theta_{\mathbf{x}_{ij}} (1 - \theta_{\mathbf{x}_{ij}}) (1 - 2\theta_{\mathbf{x}_{ij}}) \mathbf{x}_{ij} \mathbf{x}_{ij}^\top. \end{aligned}$$

For each $i = 1, 2, j = 1, 2, \dots, n_i$, let

$$\begin{aligned} a_{ij} &= \phi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) - \phi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \phi(\theta_{\mathbf{x}_{ij}}\nu) + \phi((1 - \theta_{\mathbf{x}_{ij}})\nu) \\ b_{ij} &= \psi(\theta_{\mathbf{x}_{ij}}\nu + t_{ij}) + \psi((1 - \theta_{\mathbf{x}_{ij}})\nu + y_{ij} - t_{ij}) - \psi(\theta_{\mathbf{x}_{ij}}\nu) - \psi((1 - \theta_{\mathbf{x}_{ij}})\nu). \end{aligned}$$

Since the total log likelihood is $l = \sum_{i,j} l_{ij}$,

$$\frac{\partial l}{\partial \boldsymbol{\beta}} = \nu \sum_{i,j} a_{ij} \theta_{\mathbf{x}_{ij}} (1 - \theta_{\mathbf{x}_{ij}}) \mathbf{x}_{ij} = \nu \mathbf{X}^\top W_1 \mathbf{z},$$

where the rows of \mathbf{X} are \mathbf{x}_{ij}^\top , $W_1 = \text{diag}(a_{ij})$ and $\mathbf{z} = (\theta_{\mathbf{x}_{11}}(1 - \theta_{\mathbf{x}_{11}}), \dots, \theta_{\mathbf{x}_{1n_1}}(1 - \theta_{\mathbf{x}_{1n_1}}), \dots, \theta_{\mathbf{x}_{2n_2}}(1 - \theta_{\mathbf{x}_{2n_2}}))^\top$. The rows of \mathbf{X} , W_1 and the elements of \mathbf{z} are ordered first by j and then i .

$$\begin{aligned} \frac{\partial^2 l}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^\top} &= \nu^2 \sum_{i,j} b_{ij} \theta_{\mathbf{x}_{ij}}^2 (1 - \theta_{\mathbf{x}_{ij}})^2 \mathbf{x}_{ij} \mathbf{x}_{ij}^\top + \nu \sum_{i,j} a_{ij} \theta_{\mathbf{x}_{ij}} (1 - \theta_{\mathbf{x}_{ij}}) (1 - 2\theta_{\mathbf{x}_{ij}}) \mathbf{x}_{ij} \mathbf{x}_{ij}^\top \\ &= -\nu \mathbf{X}^\top W_2 \mathbf{X}, \end{aligned}$$

where $W_2 = -\text{diag}(\nu b_{ij} \theta_{\mathbf{x}_{ij}}^2 (1 - \theta_{\mathbf{x}_{ij}})^2 + a_{ij} \theta_{\mathbf{x}_{ij}} (1 - \theta_{\mathbf{x}_{ij}}) (1 - 2\theta_{\mathbf{x}_{ij}}))$. The columns of W_2 is also ordered first by j and then by i .

When applying Laplace approximation to evaluate the marginal likelihood for a fixed ν ,

$$L_\nu = \int \exp\{l(\boldsymbol{\beta}) + \log \pi(\boldsymbol{\beta})\} d\boldsymbol{\beta},$$

where π is the prior on $\boldsymbol{\beta}$. For example, with $\pi(\boldsymbol{\beta})$ is the independent normal $N(0, \sigma_k^2)$ on the k -th element of $\boldsymbol{\beta}$, let $h_\nu(\boldsymbol{\beta}) = l(\boldsymbol{\beta}) + \log \pi(\boldsymbol{\beta}) = l(\boldsymbol{\beta}) - \boldsymbol{\beta}^\top \Sigma^{-1} \boldsymbol{\beta} / 2$, where $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$, we have

$$\begin{aligned} \frac{\partial h_\nu(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} &= \frac{\partial l}{\partial \boldsymbol{\beta}} - \Sigma^{-1} \boldsymbol{\beta} \\ \frac{\partial^2 h_\nu(\boldsymbol{\beta})}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^\top} &= \frac{\partial^2 l}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^\top} - \Sigma^{-1}. \end{aligned}$$

Hence the Newton-Raphson step for solving the MLE of $\boldsymbol{\beta}$ given ν is given by

$$\hat{\boldsymbol{\beta}}^{(t+1)} = \hat{\boldsymbol{\beta}}^{(t)} + \left(\mathbf{X}^\top W_2^{(t)} \mathbf{X} + \Sigma^{-1}/\nu \right)^{-1} \left(\mathbf{X}^\top W_1^{(t)} \mathbf{z}^{(t)} - \Sigma^{-1} \hat{\boldsymbol{\beta}}^{(t)}/\nu \right).$$

Under the alternative, suppose that $\pi(\boldsymbol{\beta}) = \pi(\boldsymbol{\beta}_1)\pi(\gamma)$, where $\boldsymbol{\beta}_1$ are the coefficients for the covariates and γ for the group indicator. Instead of using independent normal prior on γ , the LIM g -prior (?) could be adopted. Using the independent normal prior for $\boldsymbol{\beta}_1$, we have

$$\begin{aligned} h_\nu(\boldsymbol{\beta}) &= l(\boldsymbol{\beta}) + \log \pi(\boldsymbol{\beta}) \\ &= l(\boldsymbol{\beta}) - \boldsymbol{\beta}_1^\top \Sigma^{-1} \boldsymbol{\beta}_1 / 2 - g^{-1} \mathcal{J}_\nu(\hat{\gamma}) \gamma^2 / 2 \\ &= l(\boldsymbol{\beta}) - \boldsymbol{\beta}_1^\top \Sigma^{-1} \boldsymbol{\beta}_1 / 2 - g^{-1} \nu (\mathbf{X}^\top \hat{W}_2 \mathbf{X})_2 \gamma / 2 \end{aligned}$$

where $(\mathbf{X}^\top \hat{W}_2 \mathbf{X})_2$ denote the block of the Hessian matrix corresponding to γ . Therefore,

$$\begin{aligned} \frac{\partial h_\nu(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} &= \nu \mathbf{X}^\top W_1 \mathbf{z} - \begin{pmatrix} \Sigma^{-1} \boldsymbol{\beta}_1 \\ g^{-1} \nu (\mathbf{X}^\top \hat{W}_2 \mathbf{X})_2 \gamma \end{pmatrix} \\ \frac{\partial^2 h_\nu(\boldsymbol{\beta})}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^\top} &= -\nu \mathbf{X}^\top W_2 \mathbf{X} - \begin{pmatrix} \Sigma^{-1} & \mathbf{0} \\ \mathbf{0} & g^{-1} \nu (\mathbf{X}^\top \hat{W}_2 \mathbf{X})_2 \end{pmatrix}. \end{aligned}$$

The resulting NR update is

$$\begin{aligned} \hat{\boldsymbol{\beta}}^{(t+1)} &= \hat{\boldsymbol{\beta}}^{(t)} + \left(\mathbf{X}^\top W_2^{(t)} \mathbf{X} + \begin{pmatrix} \Sigma^{-1}/\nu & \mathbf{0} \\ \mathbf{0} & g^{-1} (\mathbf{X}^\top W_2^{(t)} \mathbf{X})_2 \end{pmatrix} \right)^{-1} \\ &\quad \times \left(\mathbf{X}^\top W_1^{(t)} \mathbf{z}^{(t)} - \begin{pmatrix} \hat{\Sigma}^{-1} \boldsymbol{\beta}_1^{(t)}/\nu \\ g^{-1} (\mathbf{X}^\top W_2^{(t)} \mathbf{X})_2 \hat{\gamma}^{(t)} \end{pmatrix} \right). \end{aligned}$$

B More on decision making

We first consider the original hypothesis that there is no cross-group difference. Taking a decision theoretic perspective, let $d(\mathbf{y}) \in \{0, 1\}$ be some decision rule, with $d(\mathbf{y}) = 1$ corresponding to the rejection of the global null that there are no cross-group differences in the OTU composition. When the loss function is

$$L(d(\mathbf{y}), c) = c \cdot \mathbb{1}_{[H_0 \text{ is true}]}d(\mathbf{y}) + (1 - c) \cdot \mathbb{1}_{[H_1 \text{ is true}]}(1 - d(\mathbf{y}))$$

for some $0 \leq c \leq 1$, one can show that the Bayes optimal decision rule is $d(\mathbf{y}) = \mathbb{1}_{[\text{PJAP} > c]}$. In particular, when $c = 0.5$, this gives the optimal decision under the simple 0-1 loss.

The decision on reporting the significant nodes is essentially a multiple testing problem. One way to address this problem is to use loss functions specified with the false positives and false negatives (?). For example, let $d_i(\mathbf{y}) \in \{0, 1\}$ be the decision rule on the i -th node; again, $d_i(\mathbf{y}) = 1$ corresponds to the rejection of the node-specific null. Let FD and FN denote the number of false positives and false negatives. The posterior expectation of FD and FN are

$$\begin{aligned}\overline{\text{FD}} &= \sum (1 - \text{PMAP}_i) \times d_i(\mathbf{y}), \\ \overline{\text{FN}} &= \sum \text{PMAP}_i \times (1 - d_i(\mathbf{y})).\end{aligned}$$

It can be shown that under the loss $L(d(\mathbf{y}), t) = t \times \text{FD} + \text{FN}$, the Bayes optimal decision rule, which minimizes the posterior expected loss $\bar{L}(d(\mathbf{y}), t) = t \times \overline{\text{FD}} + \overline{\text{FN}}$ has the form $d_i(\mathbf{y}) = \mathbb{1}_{[\text{PMAP}_i > c']}$ with the optimal threshold $c' = t/(t + 1)$, $t \geq 0$ (?). In our application, we use $c' = 0.5$ that corresponds to $t = 1$ which is also recommended by ? from a Bayesian model choice perspective. Note that one can also consider loss functions that directly take into account the dependency among the hypotheses being tested. In our framework, such dependency is incorporated only through the probability model, not in the decision theoretic part.

C Covariate selection

As we noted in Section 2.4, covariate selection is achievable in BGCR by putting a spike-and-slab prior on the regression coefficients (?). For example, let $r_l \stackrel{\text{ind}}{\sim} \text{Bernoulli}(q_l)$ where $q_l \in (0, 1)$, $l = 2, \dots, p + 1$. For $A \in \mathcal{I}$, we can modify the prior on $\beta(A)$ to be

$$\beta_l(A) \stackrel{\text{ind}}{\sim} (1 - r_l)\delta_0 + r_l N(0, \sigma_l^2(A)), \quad l = 2, \dots, p + 1, \quad (\text{S1})$$

where δ_0 is a point mass at zero, $\sigma_l^2(A)$'s are chosen for $N(0, \sigma_l^2(A))$ to cover all reasonable values of $\beta_l(A)$ while not supporting unreasonable values of $\beta_l(A)$.

Let $\mathbf{r} = (r_2, \dots, r_{p+1}) \in \{0, 1\}^p$. The independent Bernoulli priors on r_l induce the following prior on \mathbf{r}

$$\pi(\mathbf{r}) = \prod_{l=2}^{p+1} q_l^{r_l} (1 - q_l)^{1-r_l}.$$

Conditioning on \mathbf{r} , the marginal likelihood of the data, $\phi_1(\Omega \mid \mathbf{r})$, is available as a byproduct of the BGCR inference algorithm (Section 2.4). When the number of covariates is not too large, this allows us to get the posterior of \mathbf{r} by Bayes theorem:

$$\pi(\mathbf{r} \mid \mathbf{Y}) \propto \pi(\mathbf{r}) \phi_1(\Omega \mid \mathbf{r}).$$

We modify our simulation scenario IV in Section ?? to give a simple illustration of the covariate selection procedure. Consider the data simulated under the alternative, in which the counts of OTU ‘4481131’ (ω_s) are increased by 175% in the second group. Instead of using “gender” as a confounder, we generate two covariates for each sample:

$$x_{ij2} \stackrel{\text{iid}}{\sim} N(0, 1), \quad x_{ij3} \stackrel{\text{iid}}{\sim} N(0, 1).$$

Suppose that the first covariate is relevant to the counts of a specific OTU while the second

covariate has nothing to do with the OTU counts. Specifically, we increase the counts of OTU ‘4352657’ (ω_c) in the j -th sample in group i by $(x_{ij2} \times 175\%)$ (when this value is less than -1 , we set the count to zero). We note that due to the large variation in OTU counts, the signal injected on ω_c is quite weak.

Consider a specific round of simulation. We let $r_2 = r_3 = 0.5$, $\sigma_l^2(A) = 10$ for $l = 2, 3$ and fit BGCR with the prior in (??). ?? summaries the posterior probabilities of the four possible models. In comparison, each model has equal prior probabilities. Therefore, the important variable is correctly identified.

Covariate in the model	None	2	3	2 and 3
Posterior probability	0.223	0.320	0.186	0.270

Table S1: Posterior probabilities of different models (no confounding).

Although a covariate selection procedure can be incorporated in BGCR, one must proceed with caution since this can substantially affect or even invalidate the meaning of the testing result on the two-group difference. To see this intuitively, consider the following simplistic but representative scenario. Suppose there is a (close-to) perfect confounding covariate which explains virtually all the difference across the two groups. Once this covariate is included in the model then there is no remaining cross-group difference and the two-group comparison will not favor the alternative. However, including the covariate into the model may not improve the fit to the observed data in any substantive manner as its effect is largely overlapping with that of the intercept (i.e., the group label). Consequently, statistical model selection strategies, both Bayesian or frequentist, would very likely to exclude this covariate from the model. This would lead to a significant testing result on the two-group differences. As a simple illustration, in the previous example, suppose instead we have

$$x_{1j2} \stackrel{\text{iid}}{\sim} N(0, 1), \quad x_{2j2} \stackrel{\text{iid}}{\sim} N(2, 1).$$

In this case, the first covariate is a confounding variable. ?? summaries the posterior probabilities of the four possible models. Due to the strong confounding effect, the first covariate is excluded from the model, which would lead to false positives in the testing scenario.

Covariate in the model	None	2	3	2 and 3
Posterior probability	0.599	0.001	0.400	≈ 0

Table S2: Posterior probabilities of different models (with confounding).

D Additional figures

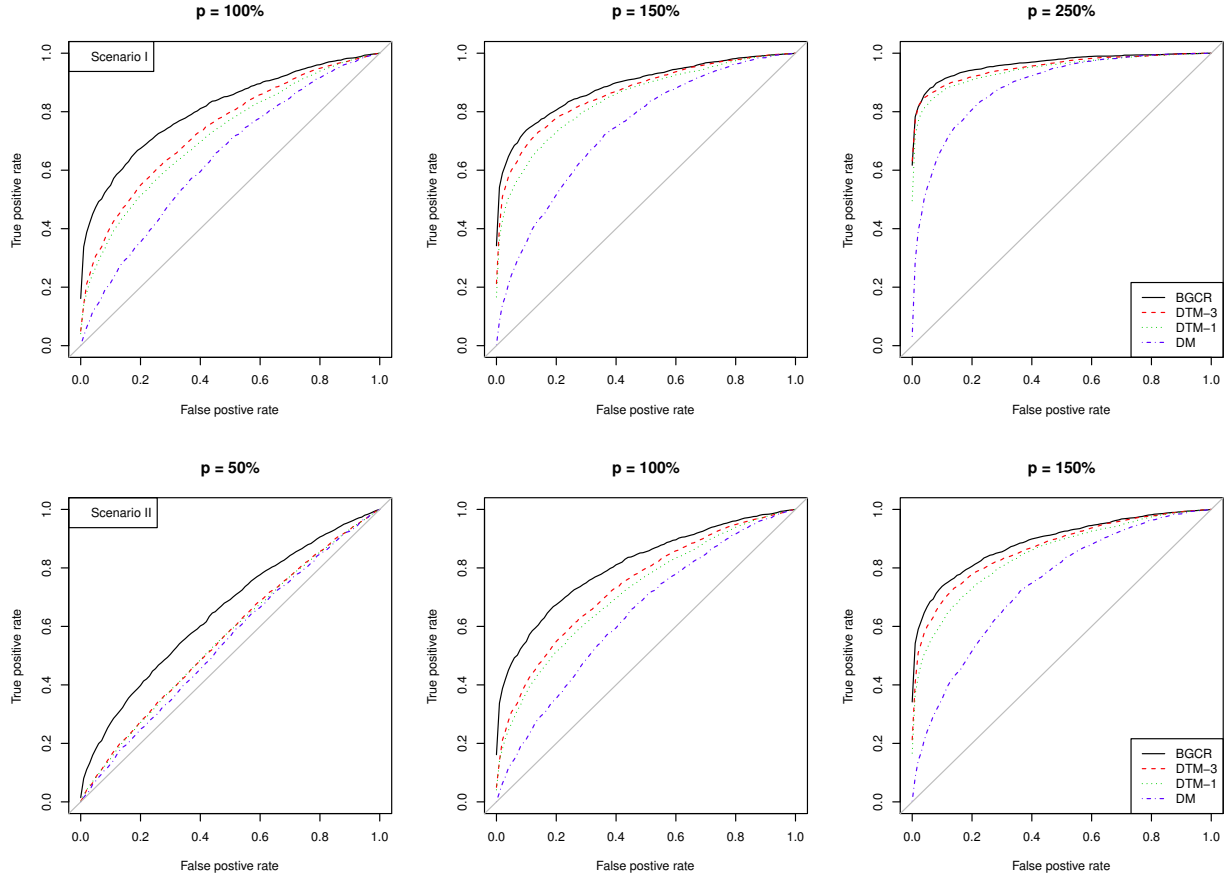


Figure S1: ROC curves for Scenario I and II with $K = 50$. The columns are indicated by the percent of count increased in the second group (p).

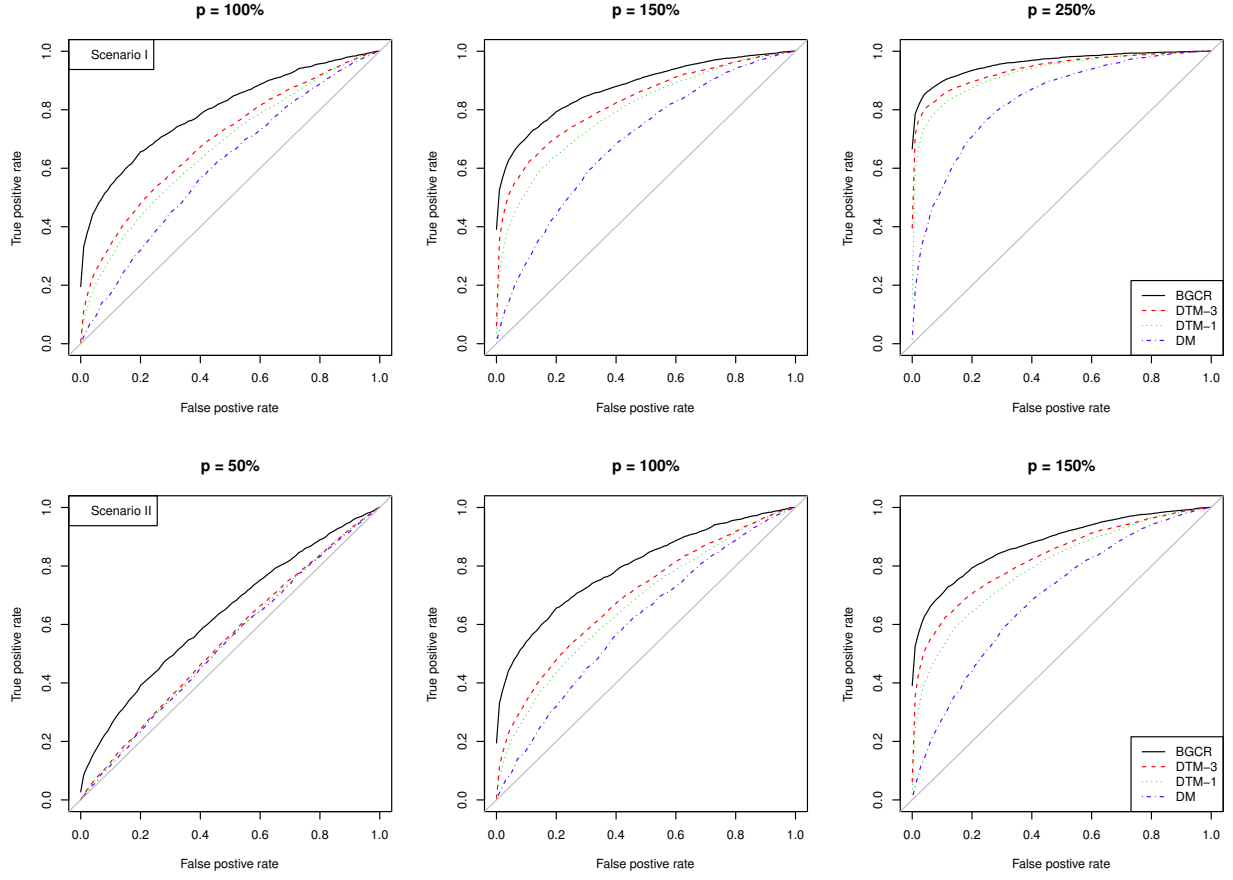


Figure S2: ROC curves for Scenario I and II with $K = 75$. The columns are indicated by the percent of count increased in the second group (p).

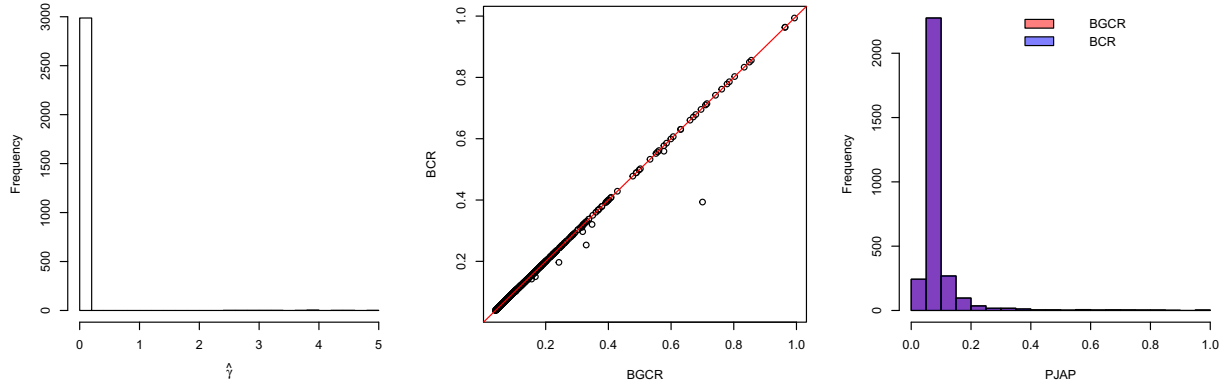


Figure S3: BGCR vs BCR under the null in scenario 3. Left: Histogram of the estimated γ in BGCR; Middle: PJAPs of BGCR vs BCR; Right: Histograms of the PJAPs of BGCR and BCR.

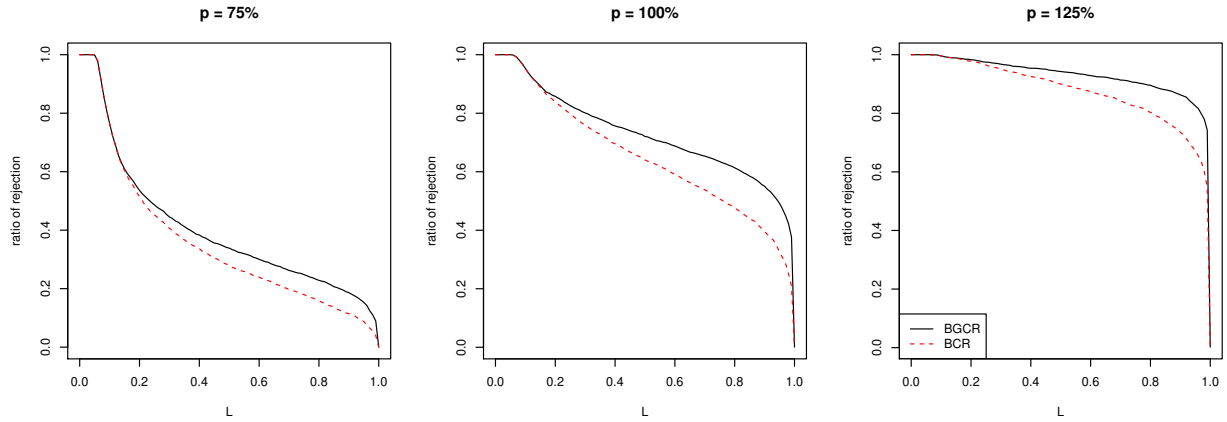


Figure S4: Ratio of rejection under the alternatives in Scenario III. The columns are indicated by the percent of count increased in the second group (p).

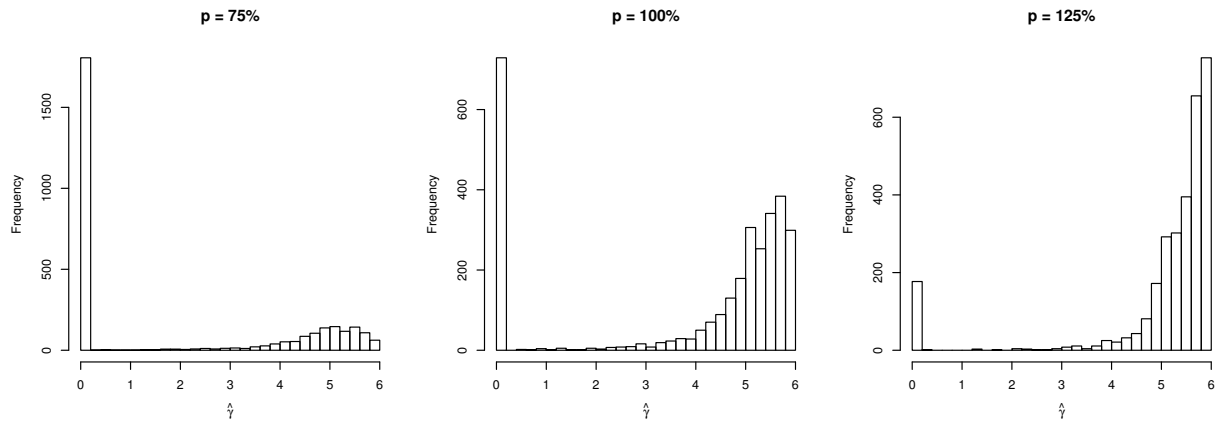


Figure S5: Estimated γ under the alternatives in scenario 3.

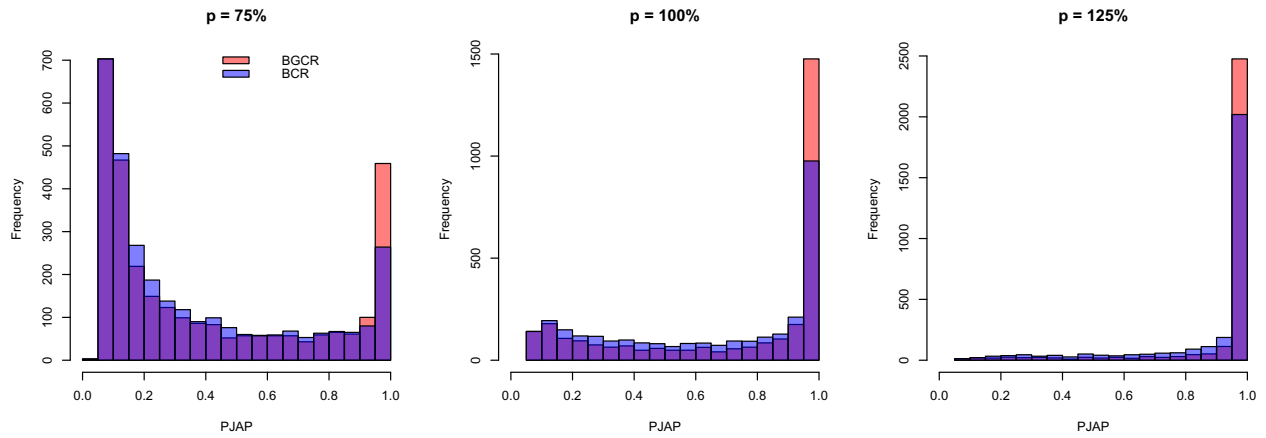


Figure S6: Histograms of the PJAPs under the alternatives in scenario 3.

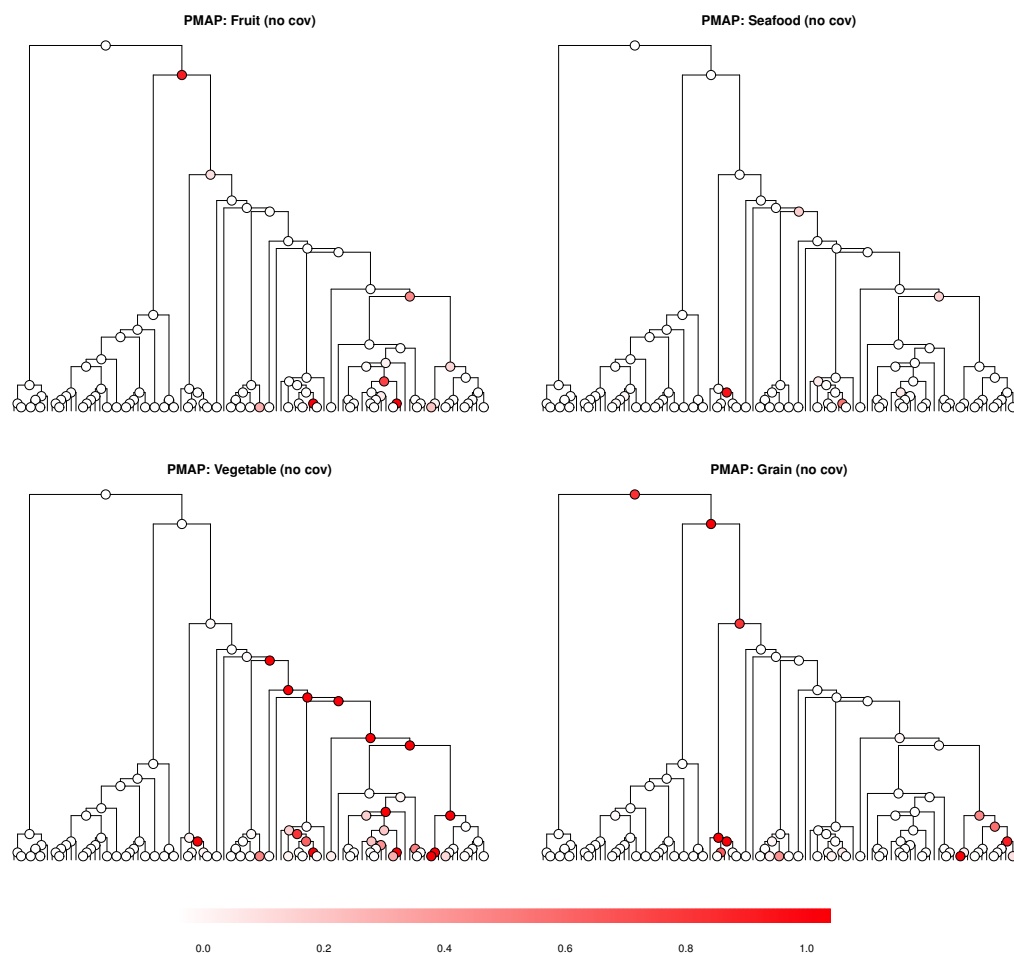


Figure S7: PMAPs for the four comparisons that reject the global null. The nodes are colored by PMAPs reported by BGCR with no covariate adjusted.

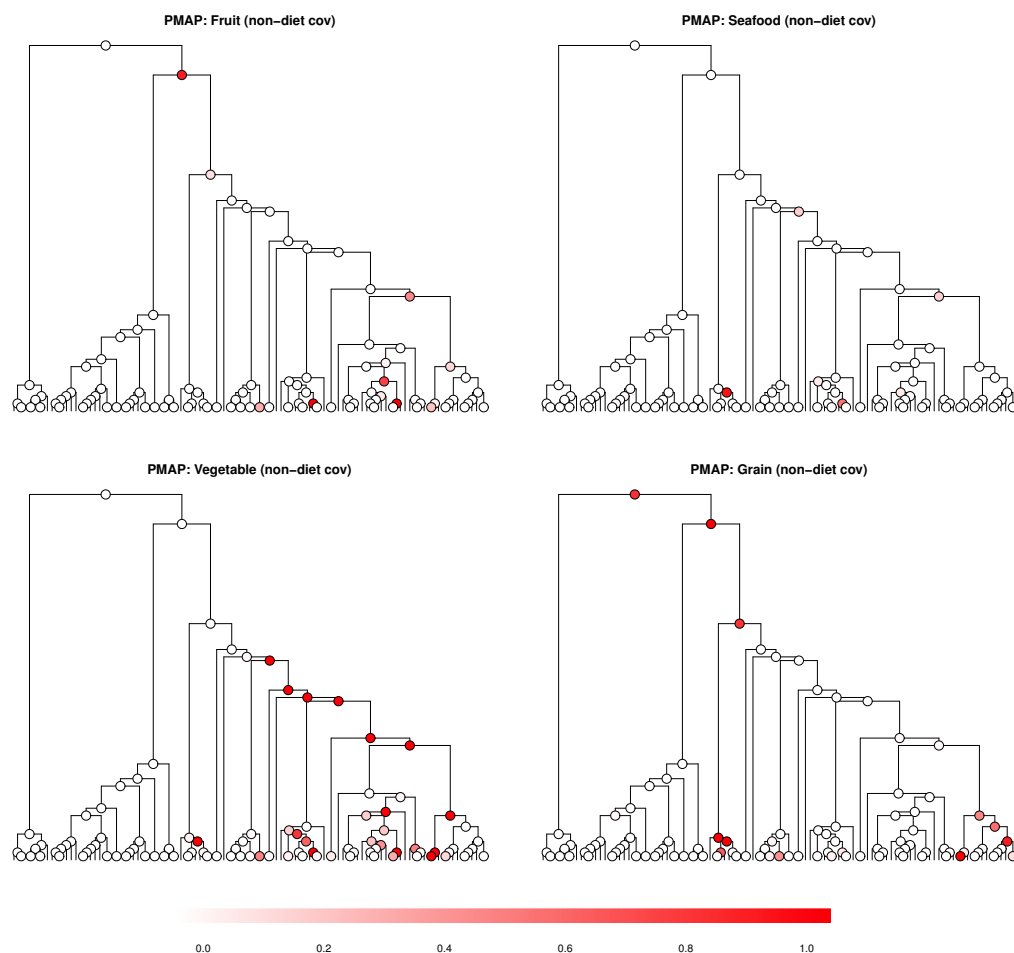


Figure S8: PMAPs for the four comparisons that reject the global null. The nodes are colored by PMAPs reported by BGCR with only non-dietary covariates adjusted.

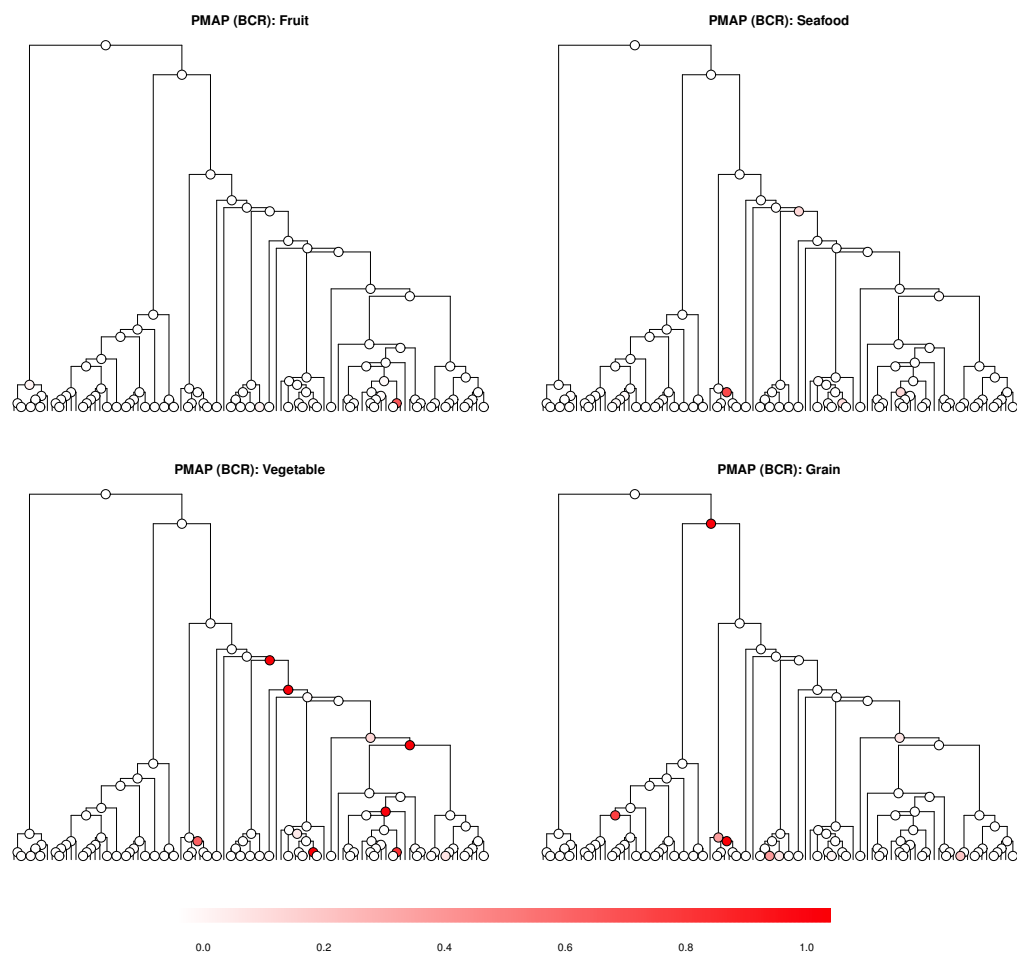


Figure S9: PMAPs for the four comparisons that reject the global null. The nodes are colored by PMAPs reported by BCR with both non-dietary covariates and dietary covariates adjusted.