



HHS Public Access

Author manuscript

J Am Stat Assoc. Author manuscript; available in PMC 2018 February 08.

Published in final edited form as:

J Am Stat Assoc. 2017 ; 112(520): 1430–1442. doi:10.1080/01621459.2017.1288631.

Bayesian Nonparametric Ordination for the Analysis of Microbial Communities

Boyu Ren, PhD,

Department of Biostatistics, Harvard University, Boston, MA 02115

Sergio Bacallado,

University academic staff, Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridgeshire CB2 1TN, UK

Stefano Favaro,

Associate Professor, University of Torino and Collegio Carlo, C.so Unione Sovietica 218/bis, Torino 10134, Italy

Susan Holmes, and

Professor, Department of Statistics, Stanford University, Sequoia Hall, Stanford, CA 94305

Lorenzo Trippa

Assistant Professor, Department of Biostatistics, Harvard University, Boston, MA 02115

Abstract

Human microbiome studies use sequencing technologies to measure the abundance of bacterial species or Operational Taxonomic Units (OTUs) in samples of biological material. Typically the data are organized in contingency tables with OTU counts across heterogeneous biological samples. In the microbial ecology community, ordination methods are frequently used to investigate latent factors or clusters that capture and describe variations of OTU counts across biological samples. It remains important to evaluate how uncertainty in estimates of each biological sample's microbial distribution propagates to ordination analyses, including visualization of clusters and projections of biological samples on low dimensional spaces. We propose a Bayesian analysis for dependent distributions to endow frequently used ordinations with estimates of uncertainty. A Bayesian nonparametric prior for dependent normalized random measures is constructed, which is marginally equivalent to the normalized generalized Gamma process, a well-known prior for nonparametric analyses. In our prior, the dependence and similarity between microbial distributions is represented by latent factors that concentrate in a low dimensional space. We use a shrinkage prior to tune the dimensionality of the latent factors. The resulting posterior samples of model parameters can be used to evaluate uncertainty in analyses

Correspondence to: Boyu Ren; Lorenzo Trippa.

SUPPLEMENTARY MATERIAL

Supplementary document This file contains the proofs for our propositions, details about the fast algorithm we described in subsection 3.1, the PCoA plots for the simulated data in Figure 3(f), Global Patterns dataset, and Ravel's vaginal microbiome dataset, and a section to illustrate the computation time and convergence diagnosis of the MCMC sampler. (.pdf file)

R code for simulation study and applications The MCMC sampler is formatted into a R package and the code to reproduce Figure 3–5 can be found in the vignette of the package (DirFactor.zip). We also include the code for the algorithm described in subsection 3.1 and the code for supplementary Figure S2–S6 (Misc.zip). (zip file)

routinely applied in microbiome studies. Specifically, by combining them with multivariate data analysis techniques we can visualize credible regions in ecological ordination plots. The characteristics of the proposed model are illustrated through a simulation study and applications in two microbiome datasets.

Keywords

Dependent Dirichlet processes; Bayesian factor analysis; Uncertainty of ordination; Microbiome data analysis

1 Introduction

Next generation sequencing (NGS) has transformed the study of microbial ecology. Through the availability of cheap efficient amplification and sequencing, marker genes such as 16S rRNA are used to provide inventories of bacteria in many different environments. For instance soil and waste water microbiota have been inventoried (DeSantis et al., 2006) as well as the human body (Dethlefsen et al., 2007). NGS also enables researchers to describe the *metagenome* by computing counts of DNA reads and matching them to the genes present in various environments.

Over the last ten years, numerous studies have shown the effects of environmental and clinical factors on the bacterial communities of the human microbiome. These studies enhance our understanding of how the microbiome is involved in obesity (Turnbaugh et al., 2009), Crohn's disease (Quince et al., 2013), or diabetes (Kostic et al., 2015). Studies are currently underway to improve our understanding of the effects of antibiotics (Dethlefsen and Relman, 2011), pregnancy (DiGiulio et al., 2015), and other perturbations to the human microbiome.

Common microbial ecology pipelines either start by grouping the 16S rRNA sequences into known Operational Taxonomic Units (OTUs) or taxa as done in Caporaso et al. (2010), or denoising and grouping the reads into more refined strains sometimes referred to as oligotypes, phylotypes, or ribosomal variants (RSV) (Rosen et al., 2012; Eren et al., 2014; Callahan et al., 2016). We will call all types of groupings OTUs to maintain consistency. In all cases the data are analyzed in the form of contingency tables of read counts per sample for the different OTUs, as exemplified in Table 1. Associated to these contingency tables are clinical and environmental covariates such as time, treatment, and patients' BMI, information collected on the same biological samples or environments. These are sometimes misnamed "metadata"; this contiguous information is usually fundamental in the analyses. The data are often assembled in multi-type structures, for instance `phyloseq` (McMurdie and Holmes, 2013) uses lists (S4 classes) to capture all the different aspects of the data at once.

Currently bioinformaticians and statisticians analyze the preprocessed microbiome data using linear ordination methods such as Correspondence Analysis (CA), Canonical or Constrained Correspondence Analysis (CCA), and Multidimensional Scaling (MDS) (Caporaso et al., 2010; Oksanen et al., 2015; McMurdie and Holmes, 2013). Distance-based

ordination methods use measures of between-sample or Beta diversity, such as the Unifrac distance (Lozupone and Knight, 2005). These analyses can reveal clustering of biological samples or taxa, or meaningful ecological or clinical gradients in the community structure of the bacteria. Clustering, when it occurs indicates a latent variable which is discrete, whereas gradients correspond to latent continuous variables. Following these exploratory stages, confirmatory analyses can include differential abundance testing (McMurdie and Holmes, 2014), two-sample tests for Beta diversity scores (Anderson et al., 2006), ANOVA permutation tests in CCA (Oksanen et al., 2015), or tests based on generalized linear models that include adjustment for multiple confounders (Paulson et al., 2013).

The interaction between these tasks can be problematic. In particular, the uncertainty in the estimation of OTUs' prevalence is often not propagated to subsequent steps (Peiffer et al., 2013). Moreover, unequal sequencing depths generate variations of the number of OTUs with zero counts across biological samples. Finally, the hypotheses tested in the inferential step are often formulated after significant exploration of the data and are sensitive to earlier choices in data preprocessing.

These issues motivate a Bayesian approach that enables us to integrate the steps of the analytical pipeline. Holmes et al. (2012); La Rosa et al. (2012); Ding and Schloss (2014) have suggested the use of a simple Dirichlet-Multinomial model for these data; however, in those analyses the multinomial probabilities for each biological sample are independent in the prior and posterior, which fails to capture underlying relationships between biological samples. The simple Dirichlet-Multinomial model is also not able to account for strong positive correlations (high co-occurrences (Faust et al., 2012)) or negative correlations (checker board effect (Koenig et al., 2011)) that can exist between different species (Gorvitovskaia et al., 2016).

We propose a Bayesian procedure, which jointly models the read counts from different OTUs and sample-specific latent multinomial distributions, allowing for correlations between OTUs. The prior assigned to these multinomial probabilities is highly flexible, such that the analysis learns the dependence structure from the data, rather than constraining it *a priori*. The method can deal with uncertainty coherently, provides model-based visualizations of the data, and is extensible to describe the effects of observed clinical and environmental covariates.

Bayesian analysis with Dirichlet priors is a convenient starting point for microbiome data, since the OTU distributions are inherently discrete. Moreover, Bayesian nonparametric priors for discrete distributions, suitable for an unbounded number of OTUs, have been the topic of intense research in recent years. General classes of priors such as normalized random measures have been developed, and their properties in relation to classical estimators of species diversity are well-understood (Ferguson, 1973; Lijoi and Prünster, 2010). The problem of modeling dependent distributions has also been extensively studied since the proposal of the Dependent Dirichlet Process (MacEachern, 2000) by Müller et al. (2004), Rodríguez et al. (2009), and Griffin et al. (2013)).

In this paper, we try to capture the variation in the composition of microbial communities as a result of a group of unobserved samples' characteristics. With this goal we introduce a model which expresses the dependence between OTUs abundances in different environments through vectors embedded in a low dimensional space. Our model has aspects in common with nonparametric priors for dependent distributions, including a generalized Dirichlet type marginal prior on each distribution, but is also similar in spirit to the multivariate methods currently employed in the microbial ecology community. Namely, it allows us to visualize the relationship between biological samples through low dimensional projections.

The paper is organized as follows. Section 2 describes a prior for dependent microbial distributions, first constructing the marginal prior of a single discrete distribution through manipulation of a Gaussian process and then extending this to multiple correlated distributions. The extension is achieved through a set of continuous latent factors, one for each biological sample, whose prior has been frequently used in Bayesian factor analyses. Section 3 derives an MCMC sampling algorithm for posterior inference and a fast algorithm to estimate biological samples' similarity. Section 4 discusses a method for visualizing the uncertainty in ordinations through conjoint analysis. Section 5 contains analyses of simulated data, which serve to demonstrate desirable properties of the method, followed by applications to real microbiome data in Section 6. Section 7 discusses potential improvement and concludes. The code for implementing the analyses discussed in this article is included in the Supplementary Materials.

2 Probability Model

In Table 1, we illustrate an example of a typical OTU table with 10 biological samples, where half are healthy subjects, and half are Inflammatory Bowel disease (IBD) patients. This contingency table is a subset of the data in Morgan et al. (2012) and records the observed frequencies of five most abundant genus level OTUs in all biological samples based on 16S rRNA sequencing results. Let Z_i be the i th observed OTU (e.g. Z_1 is Bacteroides) and n_{ij} be the observed frequency of OTU Z_i in biological sample j . As an example, $n_{11} = 1822$ is the observed frequency of Bacteroides in the biological sample Ctrl1. We will denote an OTU table as $(n_{ij})_{i \in \mathcal{I}, j \in \mathcal{J}}$, where \mathcal{I} is the number of observed OTUs and \mathcal{J} the number of biological samples.

For the biological sample j , we will assume the vector (n_{1j}, \dots, n_{Ij}) follows a multinomial distribution, noting that our analysis extends easily to the case in which the total count

$\sum_{i=1}^I n_{ij}$ is a Poisson random variable. The unobserved multinomial probabilities of OTUs present in biological sample j determine the distribution of the frequencies n_{ij} . These probabilities form a discrete probability measure, which we call a microbial distribution, on the space \mathcal{Z} of all OTUs.

We denote this discrete measure as \mathcal{P}^j and $\mathcal{P}^j(\{Z_i\})$ gives the probability of sampling Z_i from biological sample j . If we consider all J biological samples, we expect there will be variation in the respective \mathcal{P}^j s. This variation usually can be explained by specific characteristics of the biological sample. For instance, in Table 1, we can see the empirical multinomial probability of Enterococcus is higher in healthy controls than in IBD patients on average.

This variation has been discovered in prior publication (Morgan et al., 2012) and is attributed to the IBD status. Microbiome studies aim to elucidate the characteristics that explain these types of variations.

Our method focuses on modeling the distributions P^j 's and the variations among them. For biological samples labelled in $\mathcal{J} = \{1, \dots, J\}$, we assume they have the same infinite set of OTUs $Z_1, Z_2, \dots \in \mathcal{Z}$. We let the number of OTUs present in a biological sample be infinity to make our model nonparametric in consideration of the fact that there might be an unknown number of OTUs that are not observed in the experiment. We specify the probability mass assigned to a group of OTUs $A \subset \mathcal{Z}$ as

$$P^j(A) = M^j(A) / M^j(\mathcal{Z}),$$

$$M^j(A) = \sum_{i=1}^{\infty} (Z_i \in A) \sigma_i \langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+2}, \quad (1)$$

where $\sigma_i \in (0, 1)$, $\mathbf{X}_i, \mathbf{Y}^j \in \mathbb{R}^m$, (\cdot) is the indicator function, and $x^+ = x \times (x > 0)$. In addition, $\langle \cdot, \cdot \rangle$ is the standard inner product in \mathbb{R}^m .

In this model specification, σ_i is related to the average abundance of OTU i across all biological samples. When σ_i is large, the average probability mass assigned to OTU Z_i will also be large. We refer to \mathbf{X}_i and \mathbf{Y}^j as OTU vector and biological sample vectors respectively. The variation of the P^j 's is determined by the vectors \mathbf{Y}^j , which can be treated as latent characteristics of the biological samples that associate with microbial composition; for example, an unobserved feature of the subject's diet, such as vegetarianism, could affect the abundance of certain OTUs. We assume there are m such characteristics, and the l th component in \mathbf{Y}^j is the measurement of the l th latent characteristic in biological sample j . The vector \mathbf{X}_i denotes the effects of each of the m latent characteristics on the abundance of the OTU Z_i . Therefore \mathbf{X}_i has m entries.

In subsection 2.1 we consider a single microbial distribution P^j with fixed parameter \mathbf{Y}^j and define a prior on $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots)$ and $(\mathbf{X}_i)_{i \geq 1}$ which makes P^j a Dirichlet process (Ferguson, 1973). The degree of similarity between the discrete distributions $\{P^j; j \in \mathcal{J}\}$ is summarized by the Gram matrix $(\phi(j, j')) = \langle \mathbf{Y}^j, \mathbf{Y}^{j'} \rangle; j, j' \in \mathcal{J}$. Subsection 2.2 discusses the interpretation of this matrix. Subsection 2.3 proposes a prior for the parameters $\{\mathbf{Y}^j, j \in \mathcal{J}\}$ which has been previously used in Bayesian factor analysis, and which has the effect of shrinking the dimensionality of the Gram matrix $(\phi(j, j'))$ and is used to infer the number of latent characteristics m . The parameters $\{\mathbf{Y}^j, j \in \mathcal{J}\}$ or $(\phi(j, j'))$ can be used to visualize and understand variations of microbial distributions across biological samples.

2.1 Construction of a Dirichlet Process

The prior on $\sigma = (\sigma_1, \sigma_1, \dots)$ is the distribution of ordered points $(\sigma_j > \sigma_{j+1})$ in a Poisson process on $(0, 1)$ with intensity

$$\nu(\sigma) = \alpha \sigma^{-1} (1 - \sigma)^{-1/2}, \quad (2)$$

where $\alpha > 0$ is a concentration parameter. Denote the index of component of \mathbf{Y}^j and \mathbf{X}_i as l . Fix j , and let $\mathbf{Y}^j = (Y_{l,j}, l \leq m)$ be a fixed vector in \mathbb{R}^m such that $\langle \mathbf{Y}^j, \mathbf{Y}^j \rangle = 1$. We let $\mathbf{X}_i = (X_{l,i}, l \leq m)$ be a random vector for $i = 1, 2, \dots$ and $X_{l,i}$ be independent and $N(0, 1)$ a priori for $l = 1, 2, \dots, m$ and $i = 1, 2, \dots$. Finally, let G be a nonatomic probability measure on the measurable space $(\mathcal{Z}, \mathcal{F})$, where \mathcal{F} is the sigma-algebra on \mathcal{Z} , and Z_1, Z_2, \dots is a sequence of independent random variables with distribution G . We claim that the probability distribution P^j defined in Equation (1) is a Dirichlet Process with base measure G .

We note that the point process σ defines an infinite sequence of positive numbers, the products $\langle \mathbf{X}_i, \mathbf{Y}^j \rangle$, $i = 1, 2, \dots$, are independent Gaussian $N(0, 1)$ variables, and that the intensity ν satisfies the inequality $\int_0^1 \sigma d\nu < \infty$. These facts directly imply that with probability 1, $0 < M^j(A) < \infty$ when $G(A) > 0$. It also follows that for any sequence of disjoint sets $A_1, A_2, \dots \in \mathcal{F}$ the corresponding random variables $M^j(A_i)$'s are independent. In different words, M^j is a completely random measure (Kingman, 1967). The marginal Lévy intensity can be factorized as $\mu_M(ds) \times G(dz)$, where

$$\mu_M(ds) \propto \int_0^1 \nu(\sigma) \left(\frac{1}{\sigma}\right)^{1/2} s^{-1/2} \exp\left(-\frac{s}{2\sigma}\right) d\sigma ds \propto \frac{\exp(-s/2)}{s} ds, \quad \text{for } s \in (0, \infty).$$

The above expression shows that M^j is a Gamma process. We recall that the Lévy intensity of a Gamma process is proportional to the map $s \mapsto \exp(-c \times s) \times s^{-1}$, where c is a positive scale parameter. In Ferguson (1973) it is shown that a Dirichlet process can be defined by normalizing a Gamma process. It directly follows that P^j is a Dirichlet Process with base measure G .

Remark—Our construction can be extended to a wider class of normalized random measures (James, 2002; Regazzini et al., 2003) by changing the intensity ν that defines the Poisson process σ . If we set

$$\nu(\sigma) = \alpha \sigma^{-1-\beta} (1 - \sigma)^{-1/2+\beta},$$

$\beta \in [0, 1)$, in our definition of M^j , then the Lévy intensity of the random measure in (1) becomes proportional to

$$s^{-1-\beta} \exp(-s/2).$$

In this case the Lévy intensity indicates that \mathcal{M}^j is a generalized Gamma process (Brix, 1999). We recall that by normalizing this class one obtains normalized generalized Gamma processes (Lijoi et al., 2007), which include the Dirichlet process and the normalized Inverse Gaussian process (Lijoi et al., 2005) as special cases.

A few comments capture the relation between our definition of $P^j(A)$ in (1) and alternative definitions of the Dirichlet Process. If we normalize h independent $\text{Gamma}(a/h, 1/2)$ variables, we obtain a vector with $\text{Dirichlet}(a/h, \dots, a/h)$ distribution. To interpret our construction we can note that, when $a/h < 1/2$, each of the $\text{Gamma}(a/h, 1/2)$ components can be obtained by multiplying a $\text{Beta}(a/h, 1/2 - a/h)$ variable and an independent $\text{Gamma}(1/2, 1/2)$. The distribution of the $\langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+2}$ variables in (1) is in fact a mixture with a $\text{Gamma}(1/2, 1/2)$ component and a point mass at zero. Finally if we let h increase to ∞ , the law of the ordered $\text{Beta}(a/h, 1/2 - a/h)$ converges weakly to the law of ordered points of a Poisson point process on $(0, 1)$ with intensity ν (see Supplementary Document S1).

2.2 Dependent Dirichlet Processes

We use the representation for Dirichlet processes from Equation (1) to define a family of dependent Dirichlet processes labelled by a general index set \mathcal{J} . The dependency structure of this family is related to $(\phi(j, j') = \langle \mathbf{Y}^j, \mathbf{Y}^{j'} \rangle)_{j, j' \in \mathcal{J}}$. Geometrically $\phi(j, j')$ is the cosine of the angle between \mathbf{Y}^j and $\mathbf{Y}^{j'}$. The dependent Dirichlet processes is defined by setting

$$P^j(A) = \frac{\sum_i (Z_i \in A) \times \sigma_i \langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+2}}{\sum_i \sigma_i \langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+2}}, \quad \forall j \in \mathcal{J}, \quad (3)$$

for every $A \in \mathcal{A}$. Here the sequence (Z_1, Z_2, \dots) and the array $(\mathbf{X}_1, \mathbf{X}_2, \dots)$, as in Section 2.1, contain independent and identically distributed random variables, while σ is our Poisson process on the unit interval defined in (2). We will use the notation $Q_{i,j} = \langle \mathbf{X}_i, \mathbf{Y}^j \rangle$. This construction has an interpretable dependency structure between the P^j 's that we state in the next proposition.

Proposition 1—There exists a real function $\eta: [0, 1] \rightarrow [0, 1]$ such that the correlation between $P^j(A)$ and $P^{j'}(A)$ is equal to $\eta(\phi(j, j'))$ for every A that satisfies $G(A) > 0$. In different words, the correlation between $P^j(A)$ and $P^{j'}(A)$ does not depend on the specific measurable set A , it is a function of the angle defined by \mathbf{Y}^j and $\mathbf{Y}^{j'}$.

The proof is in the Supplementary Document S2. The first panel of Figure 1 shows a simulation of P^j 's. In this figure $\mathcal{J} = \{1, 2, 3, 4\}$. When $\phi(j, j')$, the cosine of the angle between two vectors \mathbf{Y}^j and $\mathbf{Y}^{j'}$, corresponding to distinct biological samples j and j' , decreases to -1 the random measures tend to concentrate on two disjoint sets. The second panel shows the function η that maps the $\phi(j, j')$'s into the correlations $\text{corr}(P^j(A), P^{j'}(A)) = \eta(\phi(j, j'))$. As expected the correlation increases with $\phi(j, j')$.

We want to point out that the construction in (3) extends easily to the setting where we are given any positive semi-definite kernel $\phi: \mathcal{J} \times \mathcal{J} \rightarrow (-1, 1)$ capturing the similarity between biological samples labelled by \mathcal{J} . Mercer's theorem (Mercer, 1909) guarantees the kernel is represented by the inner product in an \mathcal{L}^2 space, whose elements are infinite-dimensional analogues of the vectors \mathbf{Y}^j . The analysis presented in this section is unchanged in this general setting.

The next proposition provides mild conditions that guarantee a large support for the dependent Dirichlet processes that we defined.

Proposition 2—Consider a collection of probability measures $(F_j, j = 1, \dots, J)$ on \mathcal{Z} and a positive definite kernel ϕ . Assume that $\mathcal{J} = \{1, \dots, J\}$ and the support of G coincides with \mathcal{Z} . The prior distribution in (3) assigns strictly positive probability to the neighborhood $\{(F_1^i, \dots, F_J^i) : |\int f_i dF_j^i - \int f_i dF_j| < \varepsilon, i=1, \dots, L, j=1, \dots, J\}$, where $\varepsilon > 0$ and $f_i, i = 1, \dots, L$, are bounded continuous functions.

In what follows we will replace the constraint $\langle \mathbf{Y}^j, \mathbf{Y}^j \rangle = 1$ with the requirement $\langle \mathbf{Y}^j, \mathbf{Y}^j \rangle < \infty$. The two constraints are equivalent for our purpose, because we normalize

$$M^j(\cdot) = \sum_i (Z_i \in \cdot) \times \sigma_i \langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+2} \text{ and } \langle \mathbf{Y}^j, \mathbf{Y}^j \rangle \text{ can be viewed as a scale parameter.}$$

2.3 Prior on biological sample parameters

This subsection deals with the task of estimating the parameters $\mathbf{Y}^j, j \in \mathcal{J} = \{1, \dots, J\}$, that capture most of the variability observed when comparing J biological samples with different OTU counts. We define a joint prior on these factors which makes them concentrate on a low dimensional space; equivalently, the prior tends to shrink the nuclear norm of the Gram matrix $(\phi(j_1, j_2))_{j_1, j_2 \in \mathcal{J}}$. The problem of estimating low dimensional factor loadings or a low-rank covariance matrix is common in Bayesian factor analysis, and the prior defined below has been used in this area of research.

The parameters \mathbf{Y}^j can be interpreted as key characteristics of the biological samples that affect the relative abundance of OTUs. As in factor analysis, it is difficult to interpret these parameters unambiguously (Press and Shigemasa, 1989; Rowe, 2002); however, the angles between their directions have a clear interpretation. As observed in Figure 1, if the kernel $\phi(j_1, j_2) \approx \sqrt{\phi(j_1, j_1)\phi(j_2, j_2)}$, the two microbial distributions p_{j_1} and p_{j_2} will be very similar. If $\phi(j_1, j_2) \approx 0$, then there will be little correlation between OTUs' abundances in the two samples. If $\phi(j_1, j_2) \approx -\sqrt{\phi(j_1, j_1)\phi(j_2, j_2)}$, then the two microbial distributions are concentrated on disjoint sets. This interpretation suggests Principal component analysis (PCA) of the Gram matrix $(\phi(j_1, j_2))_{j_1, j_2 \in \mathcal{J}}$ as a useful exploratory data analysis technique.

It is common in factor analysis to restrict the dimensionality of factor loadings. In our model, this is accomplished by assuming \mathbf{Y}^j to be in \mathbb{R}^m and adding an error term ε in the definition of $Q_{i,j}$, the OTU-specific latent weights,

$$Q_{i,j} = \langle \mathbf{X}_i, \mathbf{Y}^j \rangle + \varepsilon_{i,j}, \quad (4)$$

where the $\varepsilon_{i,j}$ are independent standard normal variables. Recall that each sample-specific random distribution P^j is obtained by normalizing the random variables $\sigma_i(Q_{i,j}^+)^2$. If we denote the covariance matrix of $(Q_{i,1}, \dots, Q_{i,j})$ as Σ , this factor model specification indicates $\Sigma = \mathbf{Y}^T \mathbf{Y} + \mathbf{I}$ conditioning on \mathbf{Y} , where \mathbf{I} is the identity matrix and $\mathbf{Y} = (\mathbf{Y}^1, \dots, \mathbf{Y}^J)$. As a result, the correlation matrix \mathbf{S} induced by Σ only depends on \mathbf{Y} .

In most applications the dimensionality m is unknown. Several approaches to estimate m have been proposed (Lopes and West, 2004; Lee and Song, 2002; Lucas et al., 2006; Carvalho et al., 2008; Ando, 2009). However, most of them involve either calculation of Bayes Factors or complex MCMC algorithms. Instead we use a normal shrinkage prior proposed by Bhattacharya and Dunson (2011). This prior includes an infinite sequence of factors ($m = \infty$), but the variability captured by this sequence of latent factors rapidly decreases to zero. A key advantage of the model is that it does not require the user to choose the number of factors. The prior is designed to replace direct selection of m with the shrinkage toward zero of the unnecessary latent factors. In addition, this prior is nearly conjugate, which simplifies computations. The prior is defined as follows,

$$\gamma_l \sim \text{Gamma}(a_l, 1), \quad \gamma'_{l,j} \sim \text{Gamma}(v/2, v/2), \quad Y_{l,j} | \gamma \sim N \left(0, (\gamma'_{l,j})^{-1} \prod_{k \leq l} \gamma_k^{-1} \right), \quad l \geq 1, j \in \mathcal{J}, \quad (5)$$

where the random variables $\gamma = (\gamma_b, \gamma'_{l,j}; l, j \geq 1)$ are independent and, conditionally on these variables, the $Y_{l,j}$'s are independent.

When $a_l > 1$, the shrinkage strength *a priori* increases with the index l , and therefore the variability captured by each latent factor tends to decrease with l . We refer to Bhattacharya and Dunson (2011) for a detailed analysis of the prior in (5). In practice, the assumption of infinitely many factors is replaced for data analysis and posterior computations by a finite and sufficiently large number m of factors. The choice of m is based on computational considerations. It is desirable that posterior variability of the last components ($l \sim m$) of the factor model in (4) is negligible. This prior model is conditionally conjugate when paired with the dependent Dirichlet processes prior in subsection 2.2, a relevant and convenient characteristic for posterior simulations. We summarize the full model with a plate diagram, shown in Figure 2.

3 Posterior Analysis

Given an exchangeable sequence W_1, \dots, W_n from $P^j = M^j \times M^j(\mathcal{Z})^{-1}$ as defined in subsection 2.1, we can rewrite the likelihood function using variable augmentation as in James et al. (2009),

$$\prod_{i=1}^n P^j(\{W_i\}) = \int_0^\infty \frac{\exp[-M^j(\mathcal{Z})T] \times T^{n-1}}{\Gamma(n)} \prod_{i=1}^I M^j(\{W_i^*\})^{n_i} dT \quad (6)$$

Here W_1^*, \dots, W_I^* is the list of distinct values in (W_1, \dots, W_n) and n_1, \dots, n_I are the occurrences in (W_1, \dots, W_n) , so that $\sum_{i=1}^I n_i = n$. We use expression (6) to specify an algorithm that allows us to infer microbial abundances P^1, \dots, P^J in J biological samples.

We proceed, similarly to Muliere and Tardella (1998) and Ishwaran and James (2001), using truncated versions of the processes in subsection 2.2. We replace $\sigma = \{\sigma_i, i \geq 1\}$ with a finite number I of independent Beta($e_i, 1/2 - e_i$) points in $(0, 1)$. Supplementary Document S1 shows that when I diverges, and $e_i = \alpha/I$, this finite dimensional version converges weakly to the process in (2). Each point σ_i is paired with a multivariate normal $\mathbf{Q}_i = (Q_{i,1}, \dots, Q_{i,J})$ with mean zero and covariance Σ . The distribution of $M_{i,j} = \sigma_i(Q_{i,j}^+)^2$ is a mixture of a point mass at zero and a Gamma distribution. In this section \mathbf{Q} and σ are finite dimensional, and the normalized vectors P^j , which assign random probabilities to I OTUs in J biological samples, are proportional to $(M_{1,j}, \dots, M_{I,j}), j = 1, \dots, J$. Note that P^j conditional on $(Q_{1,j} > 0), \dots, (Q_{I,j} > 0)$ follows a Dirichlet distribution with parameters proportional to $(Q_{1,j} > 0), \dots, (Q_{I,j} > 0)$.

The algorithm is based on iterative sampling from the full conditional distributions. We first provide a description assuming that Σ is known. We then extend the description to allow sampling under the shrinkage prior in Section 2.3 and to infer Σ .

With I OTUs and J biological samples, the typical dataset is $\mathbf{n} = (\mathbf{n}_1, \dots, \mathbf{n}_J)$, where $\mathbf{n}_j = (n_{1,j}, \dots, n_{I,j})$ and $n_{i,j}$ is the absolute frequency of the i th OTU in the j th biological sample. We use the notation $n^j = \sum_{i=1}^I n_{i,j}$, $n_i = \sum_{j=1}^J n_{i,j}$, $\sigma = (\sigma_1, \dots, \sigma_I)$, $\mathbf{Y} = (\mathbf{Y}^j, j = 1, \dots, J)$ and $\mathbf{Q} = (Q_{i,j}, 1 \leq i \leq I, 1 \leq j \leq J)$. By using the representation in (6) we introduce the latent random variables $\mathbf{T} = (T_1, \dots, T_I)$ and rewrite the posterior distribution of (σ, \mathbf{Q}) :

$$p(\sigma, \mathbf{Q} | \mathbf{n}) \propto \left(\prod_{j=1}^J \prod_{i=1}^I (\sigma_i Q_{i,j}^+)^{n_{i,j}} \right) \times \prod_{j=1}^J \left(\sum_{i=1}^I \sigma_i Q_{i,j}^+ \right)^{-n^j} \times \pi(\sigma, \mathbf{Q}) \quad (7)$$

$$\propto \int \pi(\boldsymbol{\sigma}, \mathbf{Q}) \prod_{j=1}^J \left\{ \left(\prod_{i=1}^I (\sigma_i Q_{i,j}^{+2})^{n_{i,j}} \right) \frac{T_j^{n_j-1} \exp\left(-T_j \sum_i \sigma_i Q_{i,j}^{+2}\right)}{\Gamma(n_j)} \right\} d\mathbf{T}, \quad (8)$$

where π is the prior. In order to obtain approximate $(\boldsymbol{\sigma}, \mathbf{Q})$ sampling we specify a Gibbs sampler for $(\boldsymbol{\sigma}, \mathbf{Q}, \mathbf{T})$ with target distribution

$$p(\boldsymbol{\sigma}, \mathbf{Q}, \mathbf{T} | \mathbf{n}) \propto \pi(\boldsymbol{\sigma}, \mathbf{Q}) \prod_{j=1}^J \left\{ \left(\prod_{i=1}^I (\sigma_i Q_{i,j}^{+2})^{n_{i,j}} \right) \frac{T_j^{n_j-1} \exp\left(-T_j \sum_i \sigma_i Q_{i,j}^{+2}\right)}{\Gamma(n_j)} \right\}. \quad (9)$$

The sampler iterates the following steps:

[Step 1] Sample T_j independently, one for each biological sample $j = 1, \dots, J$,

$$T_j | \mathbf{Q}, \boldsymbol{\sigma}, \mathbf{n} \sim \text{Gamma}(n_j, \sum_i \sigma_i Q_{i,j}^{+2}).$$

[Step 2] Sample \mathbf{Q}_i independently, one for each OTU $i = 1, \dots, I$. The conditional density of $\mathbf{Q}_i = (Q_{i,1} \dots Q_{i,J})$ given $\boldsymbol{\sigma}, \mathbf{T}, \mathbf{n}$ is log-concave, and the random vectors $\mathbf{Q}_i, i = 1, \dots, I$, given $\boldsymbol{\sigma}, \mathbf{T}, \mathbf{n}$ are conditionally independent.

We simulate, for $j = 1, \dots, J$, from

$$p(Q_{i,j} | \mathbf{Q}_{i,-j}, \boldsymbol{\sigma}, \mathbf{T}, \mathbf{n}) \propto Q_{i,j}^{+2n_{i,j}} \times \exp\left(-T_j \sigma_i Q_{i,j}^{+2}\right) \times \exp\left(-\frac{(Q_{i,j} - \mu_{i,j})^2}{2s_j^2}\right), \quad (10)$$

where $\mathbf{Q}_{i,-j} = (Q_{i,1}, \dots, Q_{i,j-1}, Q_{i,j+1}, \dots, Q_{i,J})$, $\mu_{i,j} = E[Q_{i,j} | \mathbf{Q}_{i,-j}]$, $s_j^2 = \text{var}[Q_{i,j} | \mathbf{Q}_{i,-j}]$, with the proviso $0^0 = 1$. Since \mathbf{Q}_i is a multivariate normal, both $\mu_{i,j}$ and s_j have simple closed form expressions.

When $n_{i,j} = 0$ the density in (10) reduces to a mixture of truncated normals:

$$(1 - p_1) N\left(Q_{i,j}; \frac{\mu_{i,j}}{\Delta_{i,j}}, \frac{s_j^2}{\Delta_{i,j}}\right) I(Q_{i,j} > 0) + p_1 N(Q_{i,j}; \mu_{i,j}, s_j^2) I(Q_{i,j} \leq 0),$$

$$p_1 = \frac{\Phi(0; \mu_{i,j}, s_j^2) N\left(0; \frac{\mu_{i,j}}{\Delta_{i,j}}, \frac{s_j^2}{\Delta_{i,j}}\right)}{\Phi(0; \mu_{i,j}, s_j^2) N\left(0; \frac{\mu_{i,j}}{\Delta_{i,j}}, \frac{s_j^2}{\Delta_{i,j}}\right) + N(0; \mu_{i,j}, s_j^2) \left(1 - \Phi\left(0; \frac{\mu_{i,j}}{\Delta_{i,j}}, \frac{s_j^2}{\Delta_{i,j}}\right)\right)},$$

and $\Delta_{i,j}=1+2\sigma_i T_j s_j^2$. Here $\mathcal{N}(\cdot; \mu, s^2)$ and $\Phi(\cdot; \mu, s^2)$ are the density and cumulative density functions of a normal variable with mean μ and variance s^2 .

When $n_{ij} > 0$ the density $p(Q_{i,j}|\mathbf{Q}_{i,-j}, \boldsymbol{\sigma}, \mathbf{T}, \mathbf{n})$ remains log-concave, and the support becomes $(0, +\infty)$. We update $Q_{i,j}$ using a Metropolis-Hastings step with proposal identical to the Laplace approximation $N(\hat{\mu}_{i,j}, \hat{s}_{i,j}^2)$ of the density in (10),

$$\hat{\mu}_{i,j} = \frac{\mu_{i,j}/s_j^2 + \sqrt{\mu_{i,j}^2/s_j^4 + 8n_{i,j}(2\sigma_i T_j + 1/s_j^2)}}{2(2\sigma_i T_j + 1/s_j^2)}, \quad \hat{s}_{i,j}^2 = \left(\frac{2n_{i,j}}{\hat{\mu}_{i,j}^2} + 2T_j \sigma_i + \frac{1}{s_j^2} \right)^{-1}. \quad (11)$$

Here $\hat{\mu}_{i,j}$ maximizes the density (10), and $\hat{s}_{i,j}^2$ is obtained from the second derivative of the log-density at $\hat{\mu}_{i,j}$. We found the approximation accurate. In Supplementary Document S4 we provide bounds of the total variation distance between the target (10) and the approximation (11). When n_{ij} increases, the bound of the total variation decreases to zero. See also Figure S1 in the Supplementary Document.

[Step 3] Sample σ_i independently, one for each OTU $i = 1, \dots, I$, from the density

$p(\sigma_i|\mathbf{Q}, \mathbf{T}, \mathbf{n}) \propto \pi(\sigma_i) \sigma_i^{n_i} \exp(-\sigma_i \sum_{j=1}^J T_j Q_{i,j}^{+2})$. The σ_i 's are a priori independent Beta($a/I, 1/2 - a/I$) variables. We use piecewise constant bounds for

$\sigma \rightarrow \exp(-\sigma_i \sum_{j=1}^J T_j Q_{i,j}^{+2})$, $\sigma \in [0, 1]$ and an accept/reject step to sample from $p(\sigma_i|\mathbf{Q}, \mathbf{T}, \mathbf{n})$.

We now consider inference on Σ using the prior on \mathbf{Y} in subsection 2.3. The goal is to generate approximate samples of \mathbf{Y} from the posterior. We exploit the identity of the conditional distributions of \mathbf{Y} given $(\boldsymbol{\sigma}, \mathbf{T}, \mathbf{Q}, \mathbf{n})$ and \mathbf{Q} . In order to sample \mathbf{Y} from the posterior we can therefore directly apply the MCMC transitions in Bhattacharya and Dunson (2011), with \mathbf{Q} replacing the observable variables in their work.

3.1 Self-consistent estimates of biological samples' similarity

We discuss an EM-type algorithm to estimate the correlation matrix \mathbf{S} of the vectors $(Q_{i,1}, \dots, Q_{i,J})$, $i = 1, \dots, I$. Under our construction in subsection 2.3, we interpret \mathbf{S} as the normalized version of Gram matrix $(\hat{\phi}(j, j'))_{j, j' \in \mathcal{J}}$ between biological samples. In this subsection we describe an alternative estimating procedure, distinct from the Gibbs sampler, which does not require tuning of the prior probability model. The algorithm can be used for MCMC initialization and for exploratory data analyses. It assumes that the observed OTU abundances are representative of the microbial distributions, i.e. $\mathbf{P}^j = (n_{1,j}/n^j, \dots, n_{I,j}/n^j)$. Under this assumption, for each biological sample j ,

$$\sigma_i Q_{i,j}^{+2} \times (n_{i,j} > 0) \propto n_{i,j}, i=1, \dots, I, \text{ and } Q_{i,j} \leq 0 \text{ when } n_{i,j}=0. \quad (12)$$

For σ_i , $i = 1, \dots, I$, we use a moment estimate $\hat{\sigma}_i = (1/J) \sum_j (n_{i,j} / \sum_{i \neq i'} n_{i',j})$. The procedure uses these estimates and at iteration $t + 1$ generates the following results:

[Expectation] Impute repeatedly \mathbf{Q} , $\ell = 1, \dots, D$ times, consistently with the constraints (12) and using a $N(0, \Sigma_\ell)$ joint distribution. Here Σ_ℓ is the estimate of Σ , the covariance matrix of $(Q_{i,1}, \dots, Q_{i,j})$, after the t -th iteration. For each replicate $\ell = 1, \dots, D$, we fix $Q_{i,j}^\ell$ for all (i, j) pairs with strictly positive $n_{i,j}$ counts at $\sqrt{n_{i,j} / \hat{\sigma}_i}$ and sample jointly, conditional on these values, negative $Q_{i,j}^\ell$ values for the remaining (i, j) pairs with $n_{i,j} = 0$. We use these $Q_{i,j}^\ell$ values to approximate $\mathcal{L}(\Sigma)$, the full data log-likelihood, our target function as in any other EM algorithm. **[Maximization]** Set Σ_{t+1} equal to the empirical covariance matrix of the $(Q_{i,1}^\ell, \dots, Q_{i,j}^\ell)$ vectors, thus maximizing the $\mathcal{L}(\Sigma)$ approximation.

We iterate until convergence of Σ_t . Then, after the last iteration, the inferred covariance matrix of $(Q_{i,1}, \dots, Q_{i,j})$ directly identifies an estimate of \mathbf{S} . We evaluated the algorithm using in-silico datasets from the simulation study in Section 5. Overall it generates estimates that are slightly less accurate compared to posterior estimation based on MCMC simulations. We use the datasets considered in Figure 3(a), with number of factors fixed at three and n' at 100,000, for a representative example. In this case the average RV-coefficient between the true \mathbf{S} and the estimated matrix is 0.93 for the EM-type algorithm and 0.95 for posterior simulations. In our work the described procedure reduced the computing time to approximately 10% compared to the Gibbs sampler. More details on this procedure are provided in the Supplementary Document S5.

4 Visualizing uncertainty in ordination plots

Ordination methods such as Multidimensional Scaling of ecological distances or Canonical Correspondence Analysis are central in microbiome research. Given posterior samples of the model parameters, we use a procedure to plot credible regions in visualizations such as Fig 3(f). The methods that we consider here are all related to PCA and use the normalized Gram matrix \mathbf{S} between biological samples. We recall that in our model \mathbf{S} is the correlation matrix of $(Q_{i,1}, \dots, Q_{i,j})$. Based on a single posterior instance of \mathbf{S} , we can visualize biological samples in a lower dimensional space through PCA, with each biological sample projected once. Naively, one could think that simply overlaying projections of the principal component loadings generated from different posterior samples of \mathbf{S} on the same graph would show the variability of the projections. However, these super-impositions could be spurious if we carry out PCA for each \mathbf{S} sample separately. One possible problem is principal component (PC) switching, when two PCs have similar eigenvalues. Another problem is the ambiguity of signs in PCA, which would lead to random signs of the loadings that result in symmetric groups of projections of the same biological sample at different sides of the axes. More generally PCA projections from different posterior samples of \mathbf{S} are difficult to compare, as the different lower dimensional spaces are not aligned.

We alternatively identify a consensus lower dimensional space for all posterior samples of \mathbf{S} (Escoufier, 1973; Lavit et al., 1994; Abdi et al., 2005). We list the three main steps used to visualize the variability of \mathbf{S} .

1. Identify a normalized Gram matrix \mathbf{S}_0 that best summarizes K posterior samples of normalized Gram matrix $\mathbf{S}_1, \dots, \mathbf{S}_K$. One simple criterion is to minimize L_2 loss element-wise. This leads to $\mathbf{S}_0 = (\sum_i \mathbf{S}_i) / K$. Alternatively, we can define \mathbf{S}_0 as the normalized Gram matrix that maximizes similarity with $\mathbf{S}_1, \dots, \mathbf{S}_K$. One possible similarity metric between two symmetric square matrices \mathbf{A} and \mathbf{B} is the RV-coefficient (Robert and Escoufier, 1976),

$$RV(\mathbf{A}, \mathbf{B}) = \text{Tr}(\mathbf{AB}) / \sqrt{\text{Tr}(\mathbf{AA})\text{Tr}(\mathbf{BB})}$$
. We refer to Holmes (2008) for a discussion on RV-coefficients.

2. Identify the lower dimensional consensus space V based on \mathbf{S}_0 . Assume we want $\dim(V) = 2$; the basis of V will be the orthonormal eigenvectors \mathbf{v}_1 and \mathbf{v}_2 of \mathbf{S}_0 corresponding to the largest eigenvalues λ_1 and λ_2 . The configuration of all biological samples in V is visualized by projecting rows of \mathbf{S}_0 onto

$$V: (\psi_1^0, \psi_2^0) = \mathbf{S}_0(\mathbf{v}_1 \lambda_1^{-1/2}, \mathbf{v}_2 \lambda_2^{-1/2})$$
. As in a standard PCA, this configuration best approximates the normalized Gram matrix in the L_2 sense:

sence:
$$(\psi_1^0, \psi_2^0) = \text{argmin}_{\langle \psi_1, \psi_2 \rangle = 0} \|\mathbf{S}_0 - (\psi_1, \psi_2)(\psi_1, \psi_2)'\|^2$$
.

3. Project the rows of posterior sample \mathbf{S}_k onto V by
$$(\psi_1^k, \psi_2^k) = \mathbf{S}_k(\mathbf{v}_1 \lambda_1^{-1/2}, \mathbf{v}_2 \lambda_2^{-1/2})$$
. Over laying all the ψ^k displays uncertainty of \mathbf{S} in the same linear subspace. Posterior variability of the biological samples' projections is visualized in V by plotting each row of the matrices (ψ_1^k, ψ_2^k) , $k = 1, \dots, K$, in the same figure. A contour plot is produced for each biological sample (see for example Fig 3(f)) to facilitate visualization of the posterior variability of its position in the consensus space V .

5 Simulation Study

In this section, we evaluate the procedure described in Section 3 and explore whether the shrinkage prior allows us to infer the number of factors and the normalized Gram matrix between biological samples \mathbf{S} . We also consider the estimates $E(\hat{P}^j | \mathbf{h})$ obtained with our joint model, one for each biological sample j , and compare their precision with the empirical estimator. Throughout the section, we assumed the number of factors is $m = 10$ when running the posterior simulations.

We first defined a scenario with distributions P^j generated from the prior (1), with $I = 68$ OTUs and $J = 22$ biological samples. The true number of factors is m_0 , and for biological samples $j = 1, \dots, m_0/2$, the vector $\mathbf{Y}^j = (Y_{I,j}, 1 \leq I \leq m_0)$ has elements $I = m_0/2 + 1, \dots, m_0$ equal to zero, while symmetrically, for $j = J/2 + 1, \dots, J$, the vectors \mathbf{Y}^j have the elements $I = 1, \dots, m_0/2$ equal to zero. The underlying normalized Gram matrix \mathbf{S} is therefore block-diagonal. After generating the distributions P^j , we sampled with fixed total counts (n^j) per

biological sample $n^j = 1,000$. We produced 50 replicates with $m_0 = 3, 6, \text{ and } 9$. In our simulations the non-zero components $Y_{i,j}$'s are independent standard normal.

We use PCA-type summaries for the posterior samples of \mathbf{Y} generated from $p(\mathbf{Y}|\mathbf{n})$. Computations are based on the $J \times J$ normalized Gram matrix \mathbf{S} . At each MCMC iteration we generate approximate samples \mathbf{Y} from the posterior, compute \mathbf{S} by normalizing the Gram matrix $\mathbf{Y}'\mathbf{Y}$, and operate standard spectral decomposition on \mathbf{S} . This allows us to estimate the ranked eigenvalues, i.e. the principal components' variance of our \mathbf{Q} latent vectors (after normalization), by averaging over the MCMC iterations. Figure 3(a) shows the variability captured by the first 10 principal components, with the box-plots illustrating posterior means' variability across our 50 replicates. The proportion of variability associated to each *principal component* decreases rapidly after the *true* number of factors $m_0 = 3, 6, 9$. This suggests that the shrinkage model (Bhattacharya and Dunson, 2011) tends to produce posterior distributions for our \mathbf{Y} latent variables that concentrates around a linear subspace.

Figure 3(c) illustrates the accuracy of the estimated normalized Gram matrix $\hat{\mathbf{S}}$ with n^j equal to 1,000, 10,000, and 100,000. We estimated the unknown $J \times J$ normalized Gram matrix \mathbf{S} with the posterior mean of the normalized Gram matrix, which we approximate by averaging over MCMC iterations. We summarized the accuracy using the RV coefficient between $\hat{\mathbf{S}}$ and \mathbf{S} , see Robert and Escoufier (1976) for a discussion on this metric. The box-plots illustrate variability of estimates' accuracy across 50 simulation replicates. As expected, when the total counts per sample increases from 10,000 to 100,000, we only observe limited gain in accuracy. Indeed the overall number of observed OTUs with positive counts per biological sample remains comparable, with expected values equal to 30 and 33 when the total counts per biological sample are fixed at 10,000 and 100,000 respectively. We also note that when m_0 increases, the accuracy decreases.

We investigate interpretability of our model by using distributions \mathcal{P}^j generated from a probability model that slightly differs from the prior. More precisely, the i th random weight in \mathcal{P}^j , conditionally on \mathbf{Y} and \mathbf{X} , is defined proportional to a monotone function of $\langle \mathbf{X}_i, \mathbf{Y}^j \rangle^+$. We considered for example

$$P^j(A) = \frac{\sum_i \sigma_i \langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+a} I_{Z_i}(A)}{\sum_i \sigma_i \langle \mathbf{X}_i, \mathbf{Y}^j \rangle^{+a}}, \quad a > 0. \quad (13)$$

When the monotone function is quadratic the probability model becomes identical to our prior. In Figure 3(b) and Figure 3(d) we used model (13) with $a = 1$ to generate datasets. We repeated the same simulation study summarized in the previous paragraphs.

We evaluated the effectiveness of borrowing information across biological samples for estimating the vectors \mathcal{P}^j . The accuracy metric that we used is the total variation distance. We compared the Bayesian estimator $E(\mathcal{P}^j|\mathbf{n})$ and the empirical estimator $\hat{\mathcal{P}}^j$ which assigns mass $n_{i,j}/n^j$ to the i^{th} OTUs. The advantage of pooling information varies with the similarity between biological samples. To reflect this, we generated \mathcal{P}^j with non-zero components of \mathbf{Y}

sampled from a zero mean multivariate normal with $cov(Y_{l,j}, Y_{l,j'})$ equal to θ . We considered the case when β^j is generated either from our prior or model (13) with $a = 0.5, 1, 3$. In addition, we considered $\theta = 0.5, 0.75, 0.95$, $I = 68$, $J = 22$, and $m_0 = 3$, while n^j varies from 10 to 100.

The results are summarized in Figure 3(e) which shows the average difference in total variation, contrasting the Bayesian and empirical estimators. The results, both when the model is correctly specified, and when mis-specified, quantify the advantages in using a joint Bayesian model.

We complete this section with one illustration of the method in Section 4. We simulate a dataset with two clusters by generating $Y_{l,j}$ for $l = 1, \dots, m_0$ from $\mathcal{N}(-3, 1)$ when $j = 1, \dots, J/2$ and from $\mathcal{N}(3, 1)$ when $j = J/2 + 1, \dots, J$. All $Y_{l,j}$ are different from zero. We expected a low n^j to be sufficient for detecting the clusters. We sampled β^j from the prior and set $J = 22$, $I = 68$, $m = 3$, and $n^j = 100$. The PC plot and the biological sample specific credible regions are shown in Figure 3(f). In the PC plot the two clusters are illustrated with different colors. In this simulation exercise the posterior credible regions leave little ambiguity both on the presence of clusters and also on samples-specific cluster membership. To compare this with the Principal Coordinates Analysis (PCoA) method used in microbiome studies, we plot the ordination results using PCoA based on the Bray-Curtis dissimilarity matrix derived from the empirical microbial distributions (See Figure S3). We can see that the PCoA point estimate is similar to the centroids identified by the proposed Bayesian ordination method.

6 Application to microbiome datasets

In this section, we apply our Bayesian analysis to two microbiome datasets. We show that our method gives results that are consistent with previous studies, and we show our novel visualization of uncertainty in ordination plots. We start with the Global Patterns data (Caporaso et al., 2011) where human-derived and environmental biological samples are included. We then considered data on the vaginal microbiome (Ravel et al., 2011).

6.1 Global Patterns dataset

The Global Patterns dataset includes 26 biological samples derived from both human and environmental specimens. There are a total of 19,216 OTUs, and the average total counts per biological sample is larger than 100,000. We collapsed all taxa OTUs to the genus level—a standard operation in microbiome studies—and yielded 996 distinct genera. We treated these genera as OTUs' and fit our model to this collapsed dataset. We ran one MCMC chain for 50,000 iterations and recorded posterior samples every 10 iterations.

We first performed a cluster analysis of biological samples based on their microbial compositions. For each posterior sample of the model parameters, we computed β^j for $j = 1, \dots, J$ and calculated the Bray-Curtis dissimilarity matrix between biological samples. We then clustered the biological samples using this dissimilarity matrix with Partitioning Among Medoids (PAM) (Tibshirani et al., 2002). By averaging over the MCMC iterations for the clustering results from each dissimilarity matrix, we obtained the posterior probability of two biological samples being clustered together. Figure 4(a) illustrates the

clustering probabilities. We can see that biological samples belonging to a specific specimen type are tightly clustered together while different specimens tend to define separate clusters. This is consistent with the conclusion in Caporaso et al. (2011), where the authors suggest, that within specimen microbiome variations are limited when compared to variations across specimen types. We also observed that biological samples from the skin are clustered with those from the tongue. This is to some extent an expected result, because both specimens are derived from humans, and because the skin microbiome has often OTUs frequencies comparable to other body sites (Grice and Segre, 2011).

We then visualized the biological samples using ordination plots and applying the method described in Section 4. We fixed the dimension of the consensus space V at three. We plotted all biological samples' projections onto V along with contours to visualize their posterior variability. The results are shown in Figure 4(b–d). We observe a clear separation between human-derived (tongue, skin, and feces) biological samples and biological samples from free environments. This separation is mostly identified by the first two compromise axes. The third axis defines a saline/non-saline samples separation. Biological samples derived from saline environment (e.g. Ocean) are well separated when projected on this axis from those derived from non-saline environment (e.g. Creek freshwater). We observed small 95% credible regions for all biological samples projections. This low level of uncertainty captured by the small credible regions in Figure 4(b–d) is mainly explained by the large total counts n^j for all biological samples. Finally, to compare the ordination results to those given by standard methods used in microbiome studies, we generated ordination results using PCoA. Figure S4 shows that the relative positions of different types of biological samples in PCoA plots and in the Bayesian ordination plots are similar.

6.2 The Vaginal Microbiome

We also consider a dataset previously presented in Ravel et al. (2011) which contains a larger number of biological samples (900) and a simpler bacterial community structure. These biological samples are derived from 54 healthy women. Multiple biological samples are taken from each individual, ranging from one to 32 biological samples per individual. Each woman has been classified, before our microbiome sequencing data were generated, into vaginal community state subtypes (CST). This dataset contains only species level taxonomic information, and we filtered OTUs by occurrence. We only retain species with more than five reads in at least 10% of biological samples. This filtering resulted in 31 distinct OTUs. We ran one MCMC chain with 50,000 iterations.

We performed the same analyses as in the previous subsection. The results are shown in Figure 5. Clustering probabilities indicate strong within CST similarity (panel a). There is one exception, CST IV-A samples, in some cases, presenting low levels of similarities when compared to each other and tend to cluster with CST I, CST III, and CST IV-B samples. This is because CST IV-A is characterized as a highly heterogeneous subtype (Ravel et al., 2011). The ordination plots are consistent with the discoveries in Ravel et al. (2011). A tetrahedron shape is recovered, and CST I, II, III, IV-B occupy the four vertices. CST II is well separated from other CSTs by the third axis. This pattern is similar to the one observed in the plots generated using PCoA (Figure S5). We also observed a sub-clustering in CST II

which has not been detected and discussed in Ravel et al. (2011). This difference in the results can be due to distinct clustering metrics in the analyses.

Note that there are two biological samples with large credible regions, indicating high uncertainty of the corresponding positions. This uncertainty propagates on their cluster membership. Both biological samples have small total counts compared to the others. The lack of precision when using biological samples with small sequencing depth leads to high uncertainty in ordination and classification. It is therefore important to account for uncertainty in the validation of subgroups biological differences—in our case CST subtypes—based on microbiome profiling. Our example suggests also the importance of uncertainty summaries when microbiome profiles are used to classify samples. Uncertainty summaries allow us to retain all samples, including those with low counts, without the risk of overinterpreting the estimated locations and projections. This also argues for the retention of raw counts in microbiome studies (McMurdie and Holmes, 2014). By using raw counts, we can evaluate the uncertainty of our estimates and exploit the information and statistical power carried by the full dataset; whereas if we downsample the data we lose information and increase uncertainty on the projections.

It is ubiquitous to have biological samples with relevant differences in their total counts, and in some cases the number of OTUs and the total number of reads can be comparable. In this cases, the empirical estimates of microbial distributions are not reliable, and an assessment of the uncertainty is necessary for downstream analyses. The two biological samples with low total counts in the vaginal microbiome dataset are examples. For biological samples with a scarce amount of data our model provides measures of uncertainty and allows uncertainty visualizations with ordination plots.

7 Conclusion

We propose a joint model for multinomial sampling of OTUs in multiple biological samples. We apply a prior from Bayesian factor analysis to estimate the similarity between biological samples, which is summarized by a Gram matrix. Simulation studies give evidence that this parameter is recovered by the Bayes estimate, and in particular, the inherent dimensionality of the latent factors is effectively learned from the data. The simulation also demonstrated that the analysis yields more accurate estimates of the microbial distributions by borrowing information across biological samples.

In addition, we provide a robust method to visualize the uncertainty in ecological ordinations, furnishing each point in the plot with a credible region. Two published microbiome datasets were analyzed, and the results are consistent with previous findings. The second analysis demonstrates that the level of uncertainty can vary across biological samples due to differences in sampling depth, which underlines the importance of modeling multinomial sampling variations coherently. We believe our analysis will mitigate artifacts arising from rarefaction, thresholding of rare species, and other preprocessing steps.

There are several directions for development which are not explored here. We highlight the possibility of incorporating prior knowledge about the biological samples, such as the

subject or group identifier in a clinical study. To achieve this, we can augment the latent factors \mathbf{Y}^j by a vector of covariates $(b_1 w_1^j, \dots, b_p w_p^j)$ whose coefficients b could be given a normal prior, for example. The posterior distribution of the coefficients could be used to infer the magnitude of covariates' effects. A less straightforward extension involves moving away from the assumption of *a priori* exchangeability between OTUs to include prior information about phylogenetic or functional relationships between them. In our present analysis, these relationships are not taken into account in the definition of the prior for microbial distributions.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

B. Ren is supported by National Science Foundation under Grant No. DMS-1042785. S. Favaro is supported by the European Research Council (ERC) through StG N-BNP 306406. L. Trippa has been supported by the Claudia Adams Barr Program in Innovative Basic Cancer Research. S. Holmes was supported by the NIH grant R01AI112401. We thank Persi Diaconis, Kris Sankaran and Lan Huong Nguyen for helpful suggestions and improvements.

Funding was provided by the National Science Foundation, European Research Council, Claudia Adams Barr Program in Innovative Basic Cancer Research and the National Institutes of Health.

References

- Abdi, H., O'Toole, AJ., Valentin, D., Edelman, B. Computer Vision and Pattern Recognition-Workshops, 2005 CVPR Workshops IEEE Computer Society Conference on. IEEE; 2005. Distatis: The analysis of multiple distance matrices; p. 42-42.
- Anderson MJ, Ellingsen KE, McArdle BH. Multivariate dispersion as a measure of beta diversity. *Ecology Letters*. 2006; 9(6):683–693. [PubMed: 16706913]
- Ando T. Bayesian factor analysis with fat-tailed factors and its exact marginal likelihood. *Journal of Multivariate Analysis*. 2009; 100(8):1717–1726.
- Bhattacharya A, Dunson DB. Sparse Bayesian infinite factor models. *Biometrika*. 2011; 98(2):291. [PubMed: 23049129]
- Brix A. Generalized gamma measures and shot-noise cox processes. *Advances in Applied Probability*. 1999:929–953.
- Callahan BJ, McMurdie PJ, Rosen MJ, Han AW, Johnson AJA, Holmes SP. Dada2: High-resolution sample inference from illumina amplicon data. *Nature methods*. 2016; 13(7):581–583. [PubMed: 27214047]
- Caporaso JG, Kuczynski J, Stombaugh J, Bittinger K, Bushman FD, Costello EK, Fierer N, Pena AG, Goodrich JK, Gordon JI, Knight R. Qiime allows analysis of high-throughput community sequencing data. *Nature methods*. 2010; 7(5):335–336. [PubMed: 20383131]
- Caporaso JG, Lauber CL, Walters WA, Berg-Lyons D, Lozupone CA, Turnbaugh PJ, Fierer N, Knight R. Global patterns of 16s rRNA diversity at a depth of millions of sequences per sample. *Proceedings of the National Academy of Sciences*. 2011; 108(Supplement 1):4516–4522.
- Carvalho CM, Chang J, Lucas JE, Nevins JR, Wang Q, West M. High-dimensional sparse factor modeling: applications in gene expression genomics. *Journal of the American Statistical Association*. 2008; 103(484)
- DeSantis TZ, Hugenholtz P, Larsen N, Rojas M, Brodie EL, Keller K, Huber T, Dalevi D, Hu P, Andersen GL. Greengenes, a chimera-checked 16s rRNA gene database and workbench compatible with arb. *Applied and environmental microbiology*. 2006; 72(7):5069–5072. [PubMed: 16820507]

- Dethlefsen L, McFall-Ngai M, Relman DA. An ecological and evolutionary perspective on human–microbe mutualism and disease. *Nature*. 2007; 449(7164):811–818. [PubMed: 17943117]
- Dethlefsen L, Relman DA. Incomplete recovery and individualized responses of the human distal gut microbiota to repeated antibiotic perturbation. *Proceedings of the National Academy of Sciences*. 2011; 108(Supplement 1):4554–4561.
- DiGiulio D, Callahan BJ, McMurdie PJ, Costello EK, Lyell DJ, Robaczewska A, Sun CL, Goltsman DSA, Wong RJ, Shaw G, Stevenson DK, Holmes S, R DA. Temporal and spatial variation of the human microbiota during pregnancy. 2015 to appear.
- Ding T, Schloss PD. Dynamics and associations of microbial community types across the human body. *Nature*. 2014; 509(7500):357. [PubMed: 24739969]
- Eren AM, Borisy GG, Huse SM, Welch JLM. Oligotyping analysis of the human oral microbiome. *Proceedings of the National Academy of Sciences*. 2014; 111(28):E2875–E2884.
- Escoufier Y. Le traitement des variables vectorielles. *Biometrics*. 1973:751–760.
- Faust K, Sathirapongsasuti JF, Izard J, Segata N, Gevers D, Raes J, Huttenhower C. Microbial co-occurrence relationships in the human microbiome. *PLoS Comput Biol*. 2012; 8(7):e1002606. [PubMed: 22807668]
- Ferguson TS. A Bayesian analysis of some nonparametric problems. *The annals of statistics*. 1973:209–230.
- Gorvitovskaia A, Holmes SP, Huse SM. Interpreting prevotella and bacteroides as biomarkers of diet and lifestyle. *Microbiome*. 2016; 4(1):1. [PubMed: 26739322]
- Grice EA, Segre JA. The skin microbiome. *Nature Reviews Microbiology*. 2011; 9(4):244–253. [PubMed: 21407241]
- Griffin JE, Kolossiatis M, Steel MFJ. Comparing distributions by using dependent normalized random-measure mixtures. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2013; 75(3):499–529.
- Holmes I, Harris K, Quince C. Dirichlet multinomial mixtures: generative models for microbial metagenomics. *PloS one*. 2012; 7(2):e30126. [PubMed: 22319561]
- Holmes, S. Probability and statistics: Essays in honor of David A Freedman. Institute of Mathematical Statistics; 2008. Multivariate data analysis: the french way; p. 219-233.
- Ishwaran H, James LF. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*. 2001; 96(453):161–173.
- James LF. Poisson process partition calculus with applications to exchangeable models and bayesian nonparametrics. 2002 arXiv preprint math/0205093.
- James LF, Lijoi A, Prünster I. Posterior analysis for normalized random measures with independent increments. *Scandinavian Journal of Statistics*. 2009; 36(1):76–97.
- Kingman J. Completely random measures. *Pacific Journal of Mathematics*. 1967; 21(1):59–78.
- Koenig JE, Spor A, Scalfone N, Fricker AD, Stombaugh J, Knight R, Angenent LT, Ley RE. Succession of microbial consortia in the developing infant gut microbiome. *Proceedings of the National Academy of Sciences*. 2011; 108(Supplement 1):4578–4585.
- Kostic AD, Gevers D, Siljander H, Vatanen T, Hyötyläinen T, Hämäläinen A, Peet A, Tillmann V, Pöhö P, Mattila I. The dynamics of the human infant gut microbiome in development and in progression toward type 1 diabetes. *Cell host & microbe*. 2015; 17(2):260–273. [PubMed: 25662751]
- La Rosa PS, Brooks JP, Deych E, Boone EL, Edwards DJ, Wang Q, Sodergren E, Weinstock G, Shannon WD. Hypothesis testing and power calculations for taxonomic-based human microbiome data. *PloS one*. 2012; 7(12):e52078. [PubMed: 23284876]
- Lavit C, Escoufier Y, Sabatier R, Traissac P. The ACT (statis method). *Computational Statistics & Data Analysis*. 1994; 18(1):97–119.
- Lee S, Song X. Bayesian selection on the number of factors in a factor analysis model. *Behaviormetrika*. 2002; 29(1):23–39.
- Lijoi A, Mena RH, Prünster I. Hierarchical mixture modeling with normalized inverse-gaussian priors. *Journal of the American Statistical Association*. 2005; 100(472):1278–1291.

- Lijoi A, Mena RH, Prünster I. Controlling the reinforcement in Bayesian nonparametric mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2007; 69(4): 715–740.
- Lijoi, A., Prünster, I. Models beyond the dirichlet process. In: Hjort, NL, Holmes, C, Müller, P., Walker, SG., editors. *Bayesian nonparametrics*. Cambridge University Press; 2010. p. 80-136. Chapter 3
- Lopes HF, West M. Bayesian model assessment in factor analysis. *Statistica Sinica*. 2004; 14(1):41–68.
- Lozupone C, Knight R. Unifrac: a new phylogenetic method for comparing microbial communities. *Applied and environmental microbiology*. 2005; 71(12):8228–8235. [PubMed: 16332807]
- Lucas J, Carvalho C, Wang Q, Bild A, Nevins JR, West M. Sparse statistical modelling in gene expression genomics. *Bayesian Inference for Gene Expression and Proteomics*. 2006; 1
- MacEachern, SN. Unpublished manuscript. Department of Statistics, The Ohio State University; 2000. Dependent dirichlet processes.
- McMurdie PJ, Holmes S. phyloseq: an r package for reproducible interactive analysis and graphics of microbiome census data. *PLOS one*. 2013; 8(4):e61217. [PubMed: 23630581]
- McMurdie PJ, Holmes S. Waste not, want not: why rarefying microbiome data is inadmissible. *PLoS Comput Biol*. 2014; 10(4):e1003531. [PubMed: 24699258]
- Mercer J. Functions of positive and negative type, and their connection with the theory of integral equations. *Philosophical transactions of the royal society of London Series A, containing papers of a mathematical or physical character*. 1909:415–446.
- Morgan XC, Tickle TL, Sokol H, Gevers D, Devaney KL, Ward DV, Reyes JA, Shah SA, LeLeiko N, Snapper SB, et al. Dysfunction of the intestinal microbiome in inflammatory bowel disease and treatment. *Genome biology*. 2012; 13(9):1.
- Muliere P, Tardella L. Approximating distributions of random functionals of ferguson-dirichlet priors. *Canadian Journal of Statistics*. 1998; 26(2):283–297.
- Müller P, Quintana F, Rosner G. A method for combining inference across related nonparametric Bayesian models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2004; 66(3):735–749.
- Oksanen J, Blanchet FG, Kindt R, Legendre P, Minchin PR, O’Hara RB, Simpson GL, Solymos P, Stevens MHH, Wagner H. *vegan: Community Ecology Package*. 2015 Nov.
- Paulson JN, Stine OC, Bravo HC, Pop M. Differential abundance analysis for microbial marker-gene surveys. *Nature methods*. 2013; 10(12):1200–1202. [PubMed: 24076764]
- Peiffer JA, Spor A, Koren O, Jin Z, Tringe SG, Dangl JL, Buckler ES, Ley R. Diversity and heritability of the maize rhizosphere microbiome under field conditions. *Proceedings of the National Academy of Sciences*. 2013; 110(16):6548–6553.
- Press, SJ., Shigemasu, K. *Contributions to probability and statistics*. Springer; 1989. Bayesian inference in factor analysis; p. 271-287.
- Quince C, Lundin EE, Andreasson AN, Greco D, Rafter J, Talley NJ, Agreus L, Andersson AF, Engstrand L, D’Amato M. The impact of crohn’s disease genes on healthy human gut microbiota: a pilot study. *Gut*. 2013:952–954.
- Ravel J, Gajer P, Abdo Z, Schneider GM, K SSK, McCulle SL, Karlebach S, Gorle R, Russell J, Tacket CO, Brotman RM. Vaginal microbiome of reproductive-age women. *Proceedings of the National Academy of Sciences*. 2011; 108(Supplement 1):4680–4687.
- Regazzini E, Lijoi A, Prünster I. Distributional results for means of normalized random measures with independent increments. *Annals of Statistics*. 2003:560–585.
- Robert P, Escoufier Y. A unifying tool for linear multivariate statistical methods: the RV-coefficient. *Applied statistics*. 1976:257–265.
- Rodríguez A, Dunson DB, Gelfand AE. Bayesian nonparametric functional data analysis through density estimation. *Biometrika*. 2009; 96(1):149–162. [PubMed: 19262739]
- Rosen MJ, Callahan BJ, Fisher DS, Holmes S. Denoising pcr-amplified metagenome data. *BMC bioinformatics*. 2012; 13(1):283. [PubMed: 23113967]
- Rowe DB. *Multivariate Bayesian statistics: models for source separation and signal unmixing* CRC Press. 2002

- Tibshirani R, Hastie T, Narasimhan B, Chu G. Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences*. 2002; 99(10): 6567–6572.
- Turnbaugh PJ, Hamady M, Yatsunenko T, Cantarel BL, Duncan A, Ley RE, Sogin ML, Jones WJ, Roe BA, Affourtit JP, Egholm M, Henrissat B, Heath AC, Knight R, Gordon JI. A core gut microbiome in obese and lean twins. *Nature*. 2009 Jan; 457(7228):480–484. [PubMed: 19043404]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

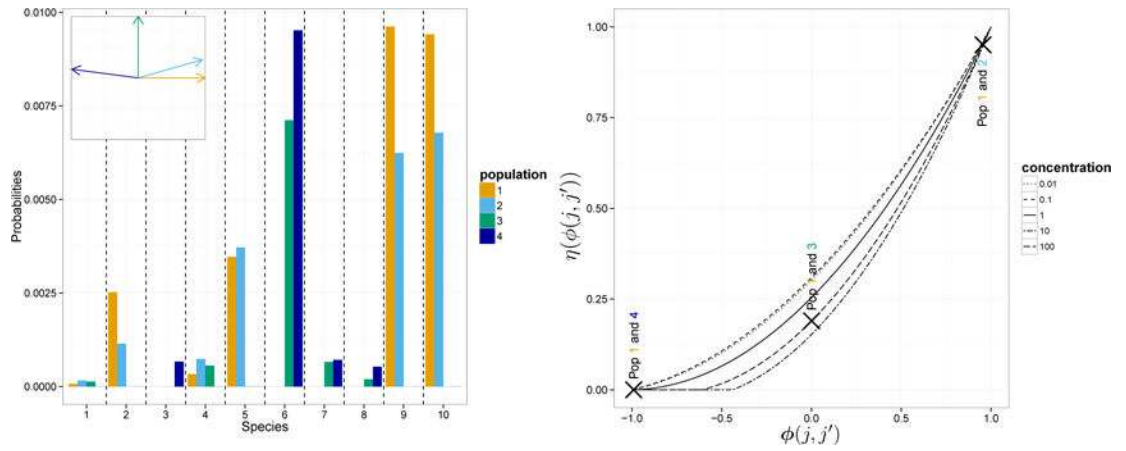


Figure 1.




Figure 2. Plate diagram. We include the factor model for the latent variables $Q_{i,j}$ as well as the matrix S . Nodes encompassed by a rectangle are defined over the range of indices indicated at the corner of the rectangle, and the connections shown within the rectangle are between nodes with the same index. We use j to index biological samples, i to index microbial species and l to index the components of latent factors.

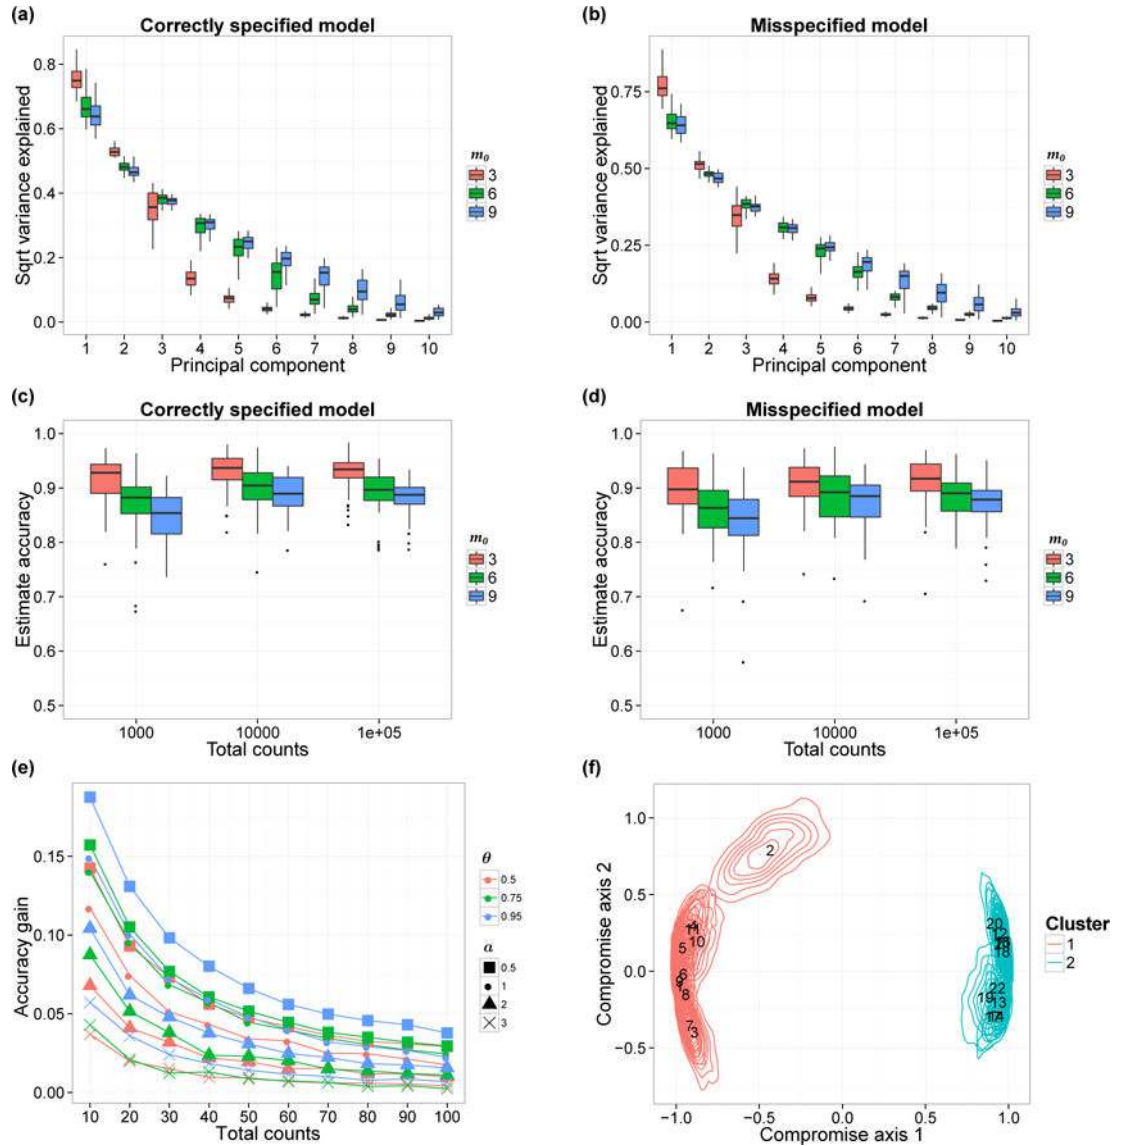


Figure 3.

(a–b) Estimated proportion of variability captured by the first 10 PCs. Each box-plot here shows the variability of the estimated proportion across 50 simulation replicates. We show the results when the data are generated from the prior (Panel a) and from the model in (13) with $a = 1$ (Panel b). (c–d) Accuracy of the correlation matrix estimates $\hat{\Sigma}$. The box-plots show the variability of the accuracy in 50 simulation replicates, with data generated from the prior (Panel c) and from model (13) with $a = 1$ (Panel d). We vary the true number of factors m_0 (colors) and n^j and show the corresponding accuracy variations. (e) Comparison between Bayesian estimates of the underlying microbial distributions \hat{P}^j and the empirical estimates. We consider the average total variation difference, averaging across all J biological samples. Each curve shows the relationship between n^j and average accuracy gain. We set $m_0 = 3$ and the parameter a varies from 0.5 to 3 (shapes). The similarity parameter θ is equal to 0.5, 0.75 or 0.95 (colors). (f) PCoA plot with confidence regions. We visualize the confidence regions using the method in Section 4. Each contour illustrates the uncertainty of a single biological

sample's position. Colors indicate cluster membership and annotated numbers are biological samples' IDs.

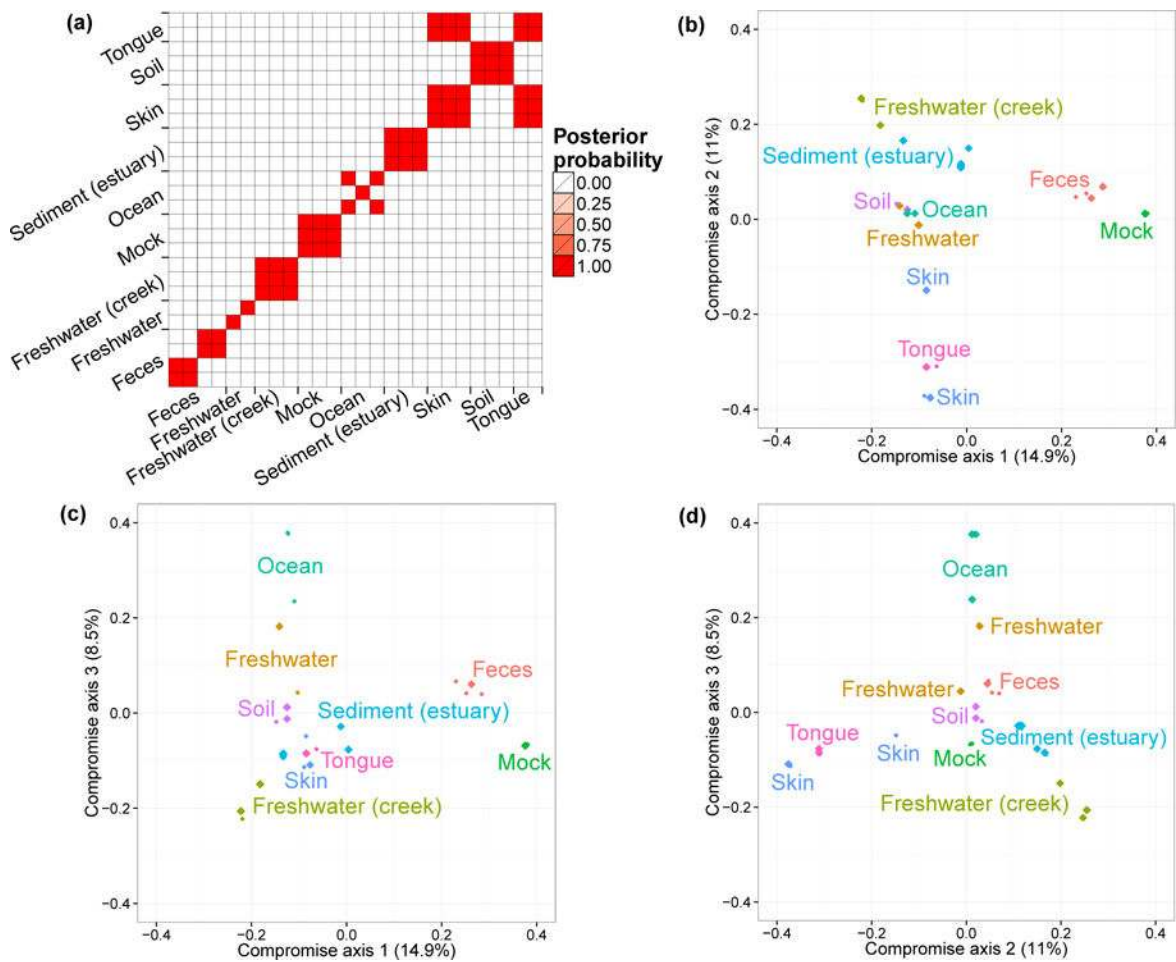


Figure 4.

(a) Posterior Probability of each pair of biological samples (j, j') being clustered together. The labels on axes indicate the environment of origin for each biological sample. (b–d) Ordination plots of biological samples and 95% posterior credible regions. We illustrate the first three compromise axes with three panels. Panel (b) plots projections on the first and second axes. Panel (c) plots projections on the first and third axes. Panel (d) plots projections on the second and third axes. The percentages on the three axes are the ratios of the corresponding S_0 eigenvalues and the trace of the matrix. The credible regions for some biological samples are so small that appears as single points. Colors and annotated text indicate the environments.

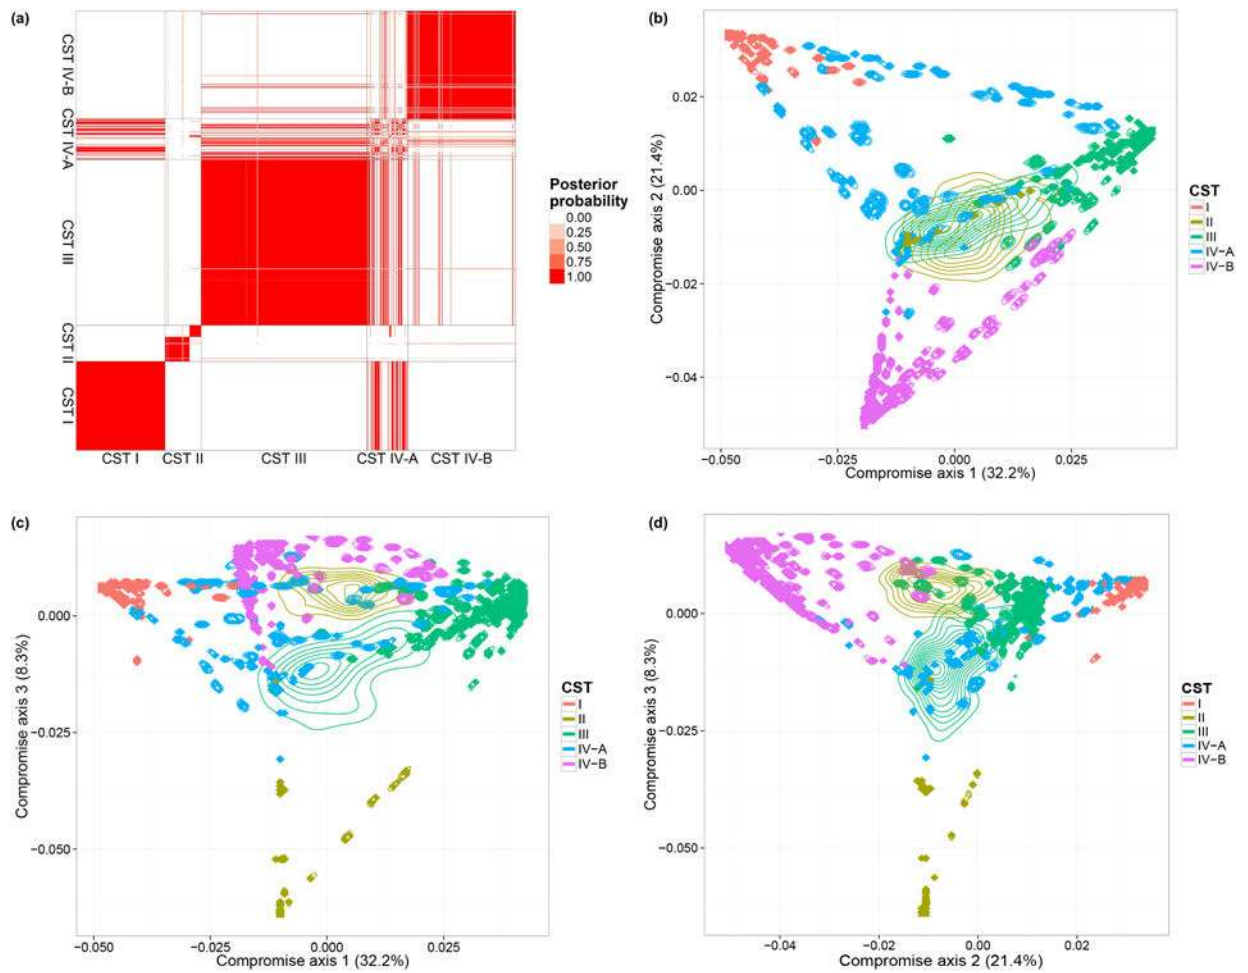


Figure 5.

(a) Posterior Probability of each pair of biological samples (j, j') being clustered together. The labels on axes indicate the CST for each biological sample. (b–d) Ordination plots of biological samples and posterior credible regions. We illustrate the first three compromise axes with three panels. The percentages on the three axes are the ratios of the corresponding S_0 eigenvalues and the trace of the matrix. Colors and indicate CSTs.

Table 1

An example of OTU table derived from data published in Morgan et al. (2012).

OTU	Ctrl1	Ctrl2	Ctrl3	Ctrl4	Ctrl5	IBD1	IBD2	IBD3	IBD4	IBD5
Bacteroidesa	1822	913	147	2988	4616	172	3516	657	550	1423
Bifidobacterium	0	162	0	0	84	0	85	1927	0	286
Collinsella	1359	0	0	206	0	327	0	0	160	122
Enterococcus	621	0	0	3	40	0	0	0	0	0
Streptococcus	75	139	2161	110	97	1820	85	58	5	294