RESEARCH ARTICLE

A unified framework for unconstrained and constrained ordination of microbiome read count data

Stijn Hawinkel[1]*, Frederiek-Maarten Kerckhof[2], Luc Bijnens[3,4], Olivier Thas[1,4,5]

1 Department of Data Analysis and Mathematical Modelling, Ghent University, Ghent, Belgium, 2 Center for Microbial Ecology and Technology, Ghent University, Ghent, Belgium, 3 Quantitative Sciences, Janssen Pharmaceutical companies of Johnson and Johnson, Beerse, Belgium, 4 Center for Statistics, Hasselt University, Hasselt, Belgium, 5 National Institute for Applied Statistics Research Australia (NIASRA), University of Wollongong, Wollongong, Australia

* stijn.hawinkel@ugent.be

Abstract

Explorative visualization techniques provide a first summary of microbiome read count datasets through dimension reduction. A plethora of dimension reduction methods exists, but many of them focus primarily on sample ordination, failing to elucidate the role of the bacterial species. Moreover, implicit but often unrealistic assumptions underlying these methods fail to account for overdispersion and differences in sequencing depth, which are two typical characteristics of sequencing data. We combine log-linear models with a dispersion estimation algorithm and flexible response function modelling into a framework for unconstrained and constrained ordination. The method is able to cope with differences in dispersion between taxa and varying sequencing depths, to yield meaningful biological patterns. Moreover, it can correct for observed technical confounders, whereas other methods are adversely affected by these artefacts. Unlike distance-based ordination methods, the assumptions underlying our method are stated explicitly and can be verified using simple diagnostics. The combination of unconstrained and constrained ordination in the same framework is unique in the field and facilitates microbiome data exploration. We illustrate the advantages of our method on simulated and real datasets, while pointing out flaws in existing methods. The algorithms for fitting and plotting are available in the R-package *RCM*.

Introduction

Explorative visualization is a key first step in the analysis of high-dimensional ecological datasets. It provides insights into the strongest patterns in the dataset, unbiased by the researcher's prior beliefs. It can also help to formulate new hypotheses to be tested in a subsequent study. Nowadays, microbiological communities are characterized by sequencing either marker genes or the entire metagenome of a sample, and attributing the sequences to their matching operational taxonomic units (OTUs), species or other phylogenetic levels. Throughout this paper,

we will refer to the lowest level to which the reads are attributed as *taxa*. Sample-specific variables, such as patient baseline characteristics or environmental conditions, can also be recorded. Microbiome sequencing datasets typically contain information on thousands of microbial taxa, whereas the number of samples and sample-specific variables is usually in the order of tens to hundreds. These data are thus *high-dimensional*, and require a dimension reduction before visualization. Apart from the biological variability, also the measurement procedure including the DNA-extraction, amplification and sequencing steps, introduces additional variability and technical artefacts, such as differences in sequencing depth [1]. At best, data visualization methods must be insensitive to this technical noise, while accurately capturing the biological signal. The first aim of such a dimension reduction is to optimally represent (dis)similarities between samples in an *ordination*: samples that are similar in high dimensional space should also be represented close together in a two or three dimensional visualization. A second aim is to elucidate which taxa drive the (dis)similarities between samples by assigning taxon scores. These taxon scores indicate how strongly the different taxa differ in abundance between the samples. A final objective might be to identify which sample-specific variables can explain the (dis)similarities in taxa composition between samples. Over the last years, methods that attempt to visualize variability in a dataset (unconstrained ordination) and methods that explore the role of sample-specific variables in shaping the community (constrained ordination), have evolved independently.

A popular ordination method for the microbiome is *principal coordinates analysis* (PCoA) [2], also known as *metric* multidimensional scaling [3]. First, the data analyst chooses a particular distance measure, which is calculated for every pair of samples in the high-dimensional space. Next, samples are represented in two dimensions such that their pairwise Euclidean distances approximate their corresponding distances in high dimensional space as closely as possible. However, no matter how well motivated the choice of distance measure for a particular application, the contribution of the individual taxa to the separation between the samples is lost in the distance calculation; see Fig 1A. One exception is PCoA with Euclidean distances, which is equivalent to Principal Components Analysis and which does directly yield taxon scores. However, most often dedicated ecological distance measures are used, such that taxon scores have to be added to the PCoA plots as weighted sample scores [4], but these scores do not reflect their contributions to the distance measures. Moreover, distance-based approaches have been shown to be affected by differences in dispersion [5] and library size [6, 7] between the samples.

Correspondence analysis (CA) [8] is a classical statistical method for the exploration of contingency tables, which allows for quantification of taxon contributions to the sample ordination. Canonical correspondence analysis (CCA) [9] even allows restricting the sample ordination to be explained by sample-specific variables (see Fig 2A). This technique thus allows for unconstrained (CA) and constrained (CCA) analysis in the same framework, which greatly enhances their use for researchers. Correspondence analysis relies on residuals for capturing the discrepancy between observed counts and the counts expected in case of identical taxa composition in all samples (sample homogeneity). It implicitly assumes a certain mean-variance relationship for normalization of these residuals. However, a residual-based approach is not well adapted to skewed data, and its mean-variance assumption is too rigid to account for the overdispersion which is typically encountered in sequencing data [5]. Moreover, both CA and CCA implicitly assume unimodal response functions, i.e. for each taxon the expected abundance shows a bell-shaped functional relationship with a *score*. This score may be latent (CA) or observed (CCA), and represents the value of a particular sample along a *gradient* of e.g. environmental conditions. CCA makes strong assumptions on the shape of these taxon response functions [9, 10].

Fig 1. Unconstrained ordination methods. (A): Principal coordinates (PCoA) sample ordination with Bray-Curtis dissimilarity on relative abundances of the Turnbaugh mice dataset. Taxon scores were added as weighted sample scores. Coloured symbols represent mice, percentages on the axes indicate fraction of eigenvalue to the sum of all eigenvalues. Only the six taxa with taxon scores furthest from the origin are plotted. (B): Biplot of the unconstrained RC(M) ordination of the same dataset. Arrows represent taxa, the ratios of the ψ parameters reflect the relative importance of the corresponding dimensions. Only the six taxa with strongest departure from homogeneity are shown for clarity. The sample ordination is similar to PCoA, but the RC(M) method proposes a more principled approach to identifying the taxa that contribute most to the separation of the samples. LF/PP: low fat, plantpolysaccharide rich.
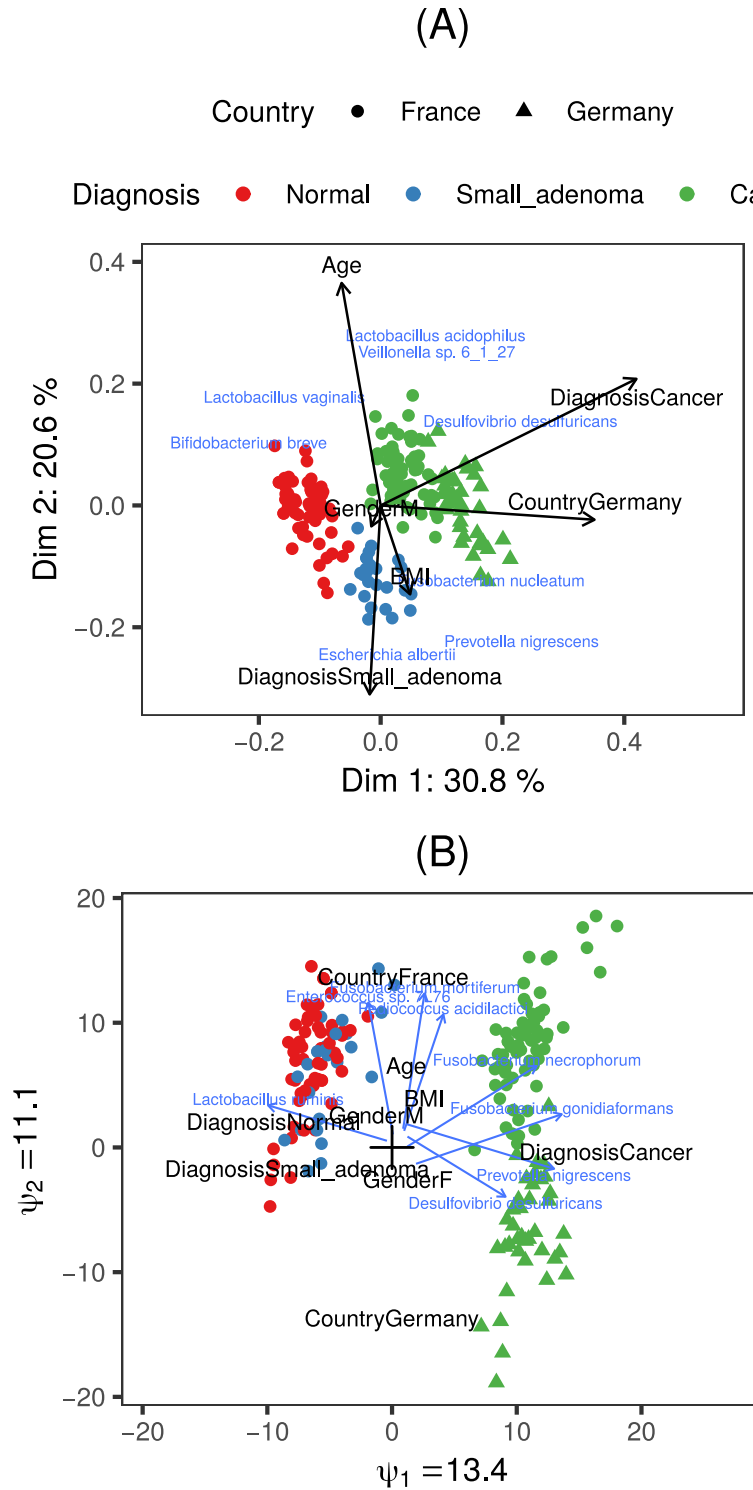
Fig 2. Constrained ordination methods. (A): Triplot of canonical correspondence analysis (CCA) of the Zeller data. Dots represent samples, the taxon labels indicate the location of the peaks of the taxon response functions under strict assumptions. For clarity, only the eight taxa with peaks furthest from the origin are shown. Percentages along the axes indicate fractions of total inertia (departure from sample homogeneity) explained by the dimension. Arrows depict the contribution of the variables to the environmental gradient. (B): Triplot of the constrained ordination of the same dataset by the RC(M) method with linear response functions. Arrows represent taxon response functions, and labels represent variables constituting the environmental gradient. The ratio of the ψ parameters reflects the relative

importance of the corresponding dimensions. Only the eight taxa that react most strongly to the environmental gradients (the longest arrows) are shown. Two Fusobacterium species are among the taxa most sensitive to the environmental gradient, and are more abundant in cancer patients than in the others, which is in accordance with the findings of [11].

Recently, new data visualization methods for sequence count data have been proposed that aim to account for their compositionality [12]. Compositional data are constrained to a constant sum that is unrelated to their composition (e.g. the library size for sequencing data). As a result, only the proportions of the components (e.g. taxa) are meaningful, and an increase in proportion (relative abundance) of one taxon automatically entails a decrease in proportion of some other taxon or taxa. These visualization methods take the compositional nature of the data into account by working on log-ratios of relative abundances, and allow to visualize the role of the taxa in the ordination. However, since sequence count tables have very high zero count frequencies, the addition of pseudocounts prior to the log-ratio transformation is needed to avoid logarithms of zero or division by zero. The choice of the pseudocount is arbitrary and can strongly affect the eventual ordination [13]. In addition, normalizing to relative abundances and using ratios, discards the information on the variance of the counts that is contained in the library size and taxon abundance [14]. As a result, these methods fail to account for heteroscedasticity, and can be distorted by technical artefacts such as differences in library size.

Over the last years, row-column interaction models for unconstrained ordination of ecological data have gained traction. Their main idea is that a statistical model is defined for the count table, and that within this model a small number of sample-taxon interaction terms is estimated. These interaction terms summarize the dataset in low dimension and can be used for plotting purposes. [15–18]. One such method is *gomms* [17]. However, it assumes inappropriate distributions with a common dispersion parameter for all taxa and does not plot the taxon scores. In ecology, a similar branch of models, referred to as *latent variable models*, has recently gained popularity. Unlike the original row-column interaction models [19], latent variable models consider the sample scores as random effects and make prior distributional assumptions on them. This renders the fitting procedure computationally intensive, without providing a clear improvement to the ordination plot as compared to fixed effects models [15, 20, 21]. Latent variable models have also been developed from a finite mixture perspective, in which samples and taxa are assigned to a small number of latent clusters. The drawback of this approach is that it lacks the liberty of assigning unique scores to all samples and taxa, such that the final ordination does not provide a comprehensive overview of the variability of the dataset [16].

As the preceding examples illustrate, a rich literature exists on ordination of ecological data, but few methods bridge the gap between unconstrained and constrained ordination. Correspondence analysis [8, 9] is a rare exception, but it is too restrictive for sequence count data. Existing row-column interaction methods [15, 17] and compositional data analysis have no counterpart for constrained analysis [12], whereas distance based methods have to resort to inefficient two-step approaches [22]. On the other hand, many methods for constrained ordination focus on the estimation of either the gradient or the response curve. As a result, they do not produce comprehensive triplots which simultaneously show the relationships between samples, taxa and sample-specific variables [10, 23, 24].

Upon combining ideas of log-linear analysis of contingency tables [18, 19], dispersion estimation for sequencing data [25] and flexible response function estimation [10, 24], we present a row-column interaction model tailored to the visualization of the strongest signals in a

microbiome count dataset. Being based on a statistical regression model, like other model-based approaches, our method has the flexibility to correct for observed confounders such as sequencing center or technology, and to adequately deal with the mean-variance relationships of sequencing data. Our method integrates unconstrained and constrained ordination into the same framework, which simplifies the workflow of microbiome data exploration. Our fitting algorithm is simpler, faster and more stable than that of other model-based ordination methods. It is implemented in R [26] in the form of the *RCM* package, which enables the creation of annotated graphs of the ordinations. Unlike distance-based ordination methods, the underlying assumptions of our method are explicitly stated and can be verified through simple diagnostic plots. Still, it is important to note that the RC(M) method cannot be used for statistical inference, but is meant only for explorative visualization.

Comparisons of ordination methods have mainly focused on sample ordination, either from the viewpoint of ordination along a gradient [5, 27–31] or clustering [6, 14], but their conclusions are not in accordance. They rely mainly on simulated data based on gradients with hypothesized response functions [27–30, 32], and on clusters of samples with similar compositions [5, 30, 32] or on real datasets with supposedly known gradients or clusters [5, 30–33]. Few studies pay attention to the role of the taxa in the ordination, but none of them does so in a quantitative way [5, 32, 34, 35]. Here we present a simulation study that evaluates sample ordination as well as identification of taxa that contribute to the separation of the samples.

Materials and methods

Real data analyses were run on a Dell laptop, and simulations were run on a server with 12 cores and on the high performance computing facilities of VSC (the Flemish Supercomputer Center). All analyses were run with the R programming language versions 3.3.1, 3.4.3 and 3.5.1 [26]. All R-code used for the publication is available in the S1 File. The code for fitting and plotting the RC(M) models can be found in the R-package *RCM*, which can be installed from https://github.com/CenterForStatistics-UGent/RCM.

Datasets

The Human Microbiome Project (HMP, V1-3 region of the 16S rRNA gene) [36] and the American Gut Project (AGP) [37] provide microbiome count datasets of healthy human volunteers. Data from two studies on the colorectal microbiome of cancer patients, referred to as the Zeller data [11] and the Kostic data [38] are also included. Furthermore, a study on several generations of gnotobiotic mice, referred to as the Turnbaugh data [39], provides non-human microbiome data. A study on microbes in cooling water provides data from a non-mammalian source, referred to as the Props data [40]. All datasets are available in the S2 File.

Simulation study

Simulations were set up by assuming a particular count distribution, for which the parameters were estimated from a real dataset. Parameter values for the taxa and samples were then sampled from this pool of realistic parameter estimates for every Monte Carlo simulation. We chose the negative binomial, zero-inflated negative binomial and Dirichlet multinomial as count distributions. The Dirichlet multinomial distribution generates even higher zero frequencies than observed in microbiome data [41], but it was included because of its common use in microbiome science [42]. Parameter values were obtained as follows. Library sizes were randomly sampled from a pool of observed library sizes of the HMP datasets. The taxon-wise mean abundance and dispersion parameters from the negative binomial distribution were

estimated by maximum likelihood from the mid-vagina, stool and tongue dorsum samples from the HMP and from the AGP data. The overdispersion parameter of the Dirichlet multinomial was estimated from the AGP dataset using the method of moments. The mixing proportions of the zero-inflated negative binomial were estimated by maximum likelihood from the AGP data. Datasets were generated with 60 samples and 1000 taxa.

Two sets of scenarios were considered. In a first set, no biological signal was introduced. The first scenario consisted in simulating data with the negative binomial distribution such that in each of four groups of 15 samples, the sampled library sizes were multiplied with a constant: 0.2, 1, 5 and 10 for the four groups. This generates technical variability that should not be picked up by the ordination method. The second scenario was similar, but now the sampled taxon-wise dispersions were multiplied by 0.2, 1, 2 and 5 for the four groups. The second set of scenarios were designed to represent different types of biological signal that should be detected and visualized by the ordination method. Counts were also generated for 4 equally sized groups of samples, but with different taxa compositions.

In the first scenario, which will be referred to as NB, initially one taxa composition was sampled for all the groups. This composition was then altered for every group separately by multiplying a random sample of 10% of the taxon abundances by a fold change of 5 so as to make them differentially abundant (DA). Counts were generated with the negative binomial distribution. The second setting, referred to as NB(cor), was identical to the first, except that counts were generated with between-taxon correlations. These taxon correlation networks were estimated by SpiecEasi [43] on the mid-vagina, stool and tongue dorsum datasets of the HMP and on the AGP data. A correlation network was sampled for every Monte Carlo instance. The third scenario, referred to as NB(phy), was also similar to NB, only now a random phylogenetic tree was created for every dataset. Next, the tree was divided into 20 clusters of related taxa, and differential abundance was introduced in one of the clusters with a fold change of 5. This assures that the DA taxa are phylogenetically related, similar to the second approach in [44]. The fourth simulation scenario, which will be referred to as DM, used the Dirichlet multinomial distribution, for which DA is introduced as for the NB scenario. The fifth scenario, referred to as ZINB, was again similar to the NB setup, but used the zero-inflated negative binomial distribution. The DA is introduced only in the count part of the distribution. Further details and additional simulation scenarios can be found in Section 3.1 of the S1 Appendix. Except for the data generated with the Dirichlet multinomial distribution, the data generated in this way are not *compositional*, as they do not obey a sum constraint. However, any analysis method that incorporates an estimate of the sequencing depth will implicitly treat the data as compositional.

Competitor ordination methods.   As competitor ordination methods we include (1) detrended correspondence analysis (DCA), (2) ordination through PCoA with (a) Bray-Curtis dissimilarities on absolute abundances (Bray-Curtis-Abs), rarefied absolute abundances (Bray-Curtis-rare), relative abundances (Bray-Curtis) and log-transformed abundances (after adding a pseudocount of 1) (Bray-Curtis-Log), with (b) Jensen-Shannon divergence (JSD), with (c) unweighted and weighted UniFrac distances (UniFrac and w-UniFrac), and (3) ordination through non-metric multidimensional scaling with Bray-Curtis dissimilarities on relative abundances (Bray-Curtis-NMDS) and (4) DPCoA using *phyloseq* [45]. Correspondence analysis approximating the Pearson's chi-squared (CApearson), the contingency ratio (CAcontRat) and the chi-squared distance (CAchisq) was implemented according to [46]. The ordination based on the Hellinger distance (Hellinger) follows [47]. Compositional data analysis (CoDa) biplots follow [12]. The *gomms* R-package was used to run the GOMMS ordination method [17] and the *tsne* R-package for the t-SNE method [48]. The *gllvm* method augmented with the negative binomial distribution was fitted with the *gllvm* R-package [49]. All methods were

applied to count matrices trimmed for taxa below a prevalence threshold of 5% or with a total count lower than 10% of the number of samples. Ordinations in three dimensions were requested.

Performance metrics. The results of all ordinations on the simulated datasets were evaluated for separation of the sample clusters through silhouettes [50] and through a pseudo F-statistic [33, 51]. The contribution of the taxa to the correct separation of the samples is quantified by the "taxon ratio". This metric is based on the average inner product of the DA taxon scores and the sample scores (see next section for a definition of the scores) of samples in which the taxa are known to be differentially abundant. This yields a measure of how much these DA taxa contribute to the separation of the samples. The mean inner product of the non-DA taxon scores with the same sample scores should be small for an ordination method that can discriminate between DA and non-DA taxa. The ratio of the former to the latter mean inner product is the taxon ratio. It is used as a measure of method performance in terms of taxon identification. Evidently, these performance metrics can only be calculated when the underlying truth is known, e.g. in simulations, but not for real data. Finally, also the correlations of the sample scores with the observed library size are calculated. These summary measures allow a quick evaluation of all simulation runs, but inevitably high values for these measures do not always correspond to an aesthetically pleasing biplot. As a result these measures should only be used for the comparison of different methods in relative terms.

Results

The RC(M) model

The unconstrained RC(M) method and biplots. A typical microbiome count dataset is represented as an n×p count table **X** for n samples and p taxa. An n×d matrix of sample-specific variables **Q** (the metadata) can also be available; categorical variables are represented by 0/1 dummy variables. In the unconstrained Row-Column interaction model of dimension M (RC(M)), the expected count of taxon j in sample i is modelled as

$$\log(E(X_{ij})) = u_i + v_j + \sum_{m=1}^{M} \psi_m r_{im} s_{jm}, \tag{1}$$

in which $u_i + v_j$ represents the independence model (note that we refer to the model as "RC (M)", and to the R-package as "*RCM*"). The independence model describes the expected counts under the assumption of equal taxa composition in all samples (sample homogeneity). In the current context, $\exp(u_i)$ is a measure of sequencing depth of sample i, and $\exp(v_j)$ is the mean relative abundance of taxon j. The factor r_{im} is a sample score that captures departure from homogeneity in sample i in dimension m, and s_{jm} is a taxon score for taxon j in dimension m. Because the sample and taxon scores are normalized for identifiability (see Section 2.1.5 of the S1 Appendix), the parameter ψ_m is a measure of overall strength of departure from homogeneity in dimension m. The constant M is the number of dimensions of the ordination, which is usually 2 or 3, as this is the number of dimensions that can be plotted. This mean model is augmented with a negative binomial count distribution for X_{ij}, which captures the high variance and high zero frequency in microbiome count data [5, 14]. The term $\sum_{m=1}^{M} \psi_m r_{im} s_{jm}$ in Eq 1 can be used to make interpretable biplots for visualizing departures from homogeneity. In 2D one can plot $\psi_1 r_{i1}$ versus $\psi_2 r_{i2}$ to obtain a sample ordination plot. Samples close together on this plot depart similarly from homogeneity and thus have similar taxa compositions (see Fig 1B). To reveal the role of the individual taxa in this ordination, we add the p taxon scores s_{j1} versus s_{j2} as arrows on the same plot. The orthogonal projection of

(s_{j1}, s_{j2}) on $(\psi_1 r_{i1}, \psi_2 r_{i2})$ gives $\sum_{m=1}^{2} \psi_m r_{im} s_{jm}$, which quantifies the deviation of taxon j in sample i from sample homogeneity; see Eq 1.

Loosely speaking, taxa have a higher expected abundance in samples for which the sample dots and taxon arrows lie at the same side of the origin, and a lower expected abundance if they lie at opposite sides. See Section 2 of the S1 Appendix for a detailed description of the estimating algorithm and the construction of biplots, Section 4 for real data examples.

Finally, it is important to note that the RC(M) model in all its forms is overparametrized. To allow for model identifiability, restrictions are imposed on some of its parameters (see Section 2.1.5 of the S1 Appendix). This also implies that the RC(M) model fit is no longer a full maximum-likelihood solution, and classical statistical inference, such as hypothesis testing and confidence intervals, are not available. The RC(M) method should thus only be used for data exploration.

Conditioning in the RC(M)-model. Technical sample-specific variables such as batch effects and sequencing center and technology are known to affect the observed counts [52]. When these confounding variables are known, they can be included in the RC(M) model. This effectively filters out their effect, similar to conditioning in correspondence analysis [53] and latent variable models [54, 55]. With **G** an n×k confounder matrix (a subset of **Q**), model (1) is extended to

$$\log(E(X_{ij})) = \overbrace{u_i + v_j}^{\text{Independence model}} + \underbrace{\sum_{l=1}^{k} \zeta_{jl} g_{il}}_{\text{Extended null model}} + \overbrace{\sum_{m=1}^{M} \psi_m r_{im} s_{jm}}^{\text{Biological signal}}. \tag{2}$$

In this model, ζ_{jl} is a parameter such that the interaction term $\zeta_{jl} g_{il}$ captures the departure from homogeneity of taxon j in sample i due to variable l. As a result, the biological signal term $\sum_{m=1}^{M} \psi_m r_{im} s_{jm}$ is free of the effect of the confounding variables. This is illustrated in Fig 3. Details can be found in Section 2.1.4 of the S1 Appendix. Conditioning on observed confounders can be applied in the unconstrained as well as in the constrained RC(M) model (see next section).

The constrained RC(M) model. The idea of a constrained ordination is to visualize the variability in the dataset that can be explained by sample-specific variables [9, 10]. Constrained ordination is traditionally performed by finding an environmental gradient $\boldsymbol{\alpha}_m$ for every dimension m. Let $\mathbf{c}_i$ represent the i^{th} row of **C** (a subset of **Q**, excluding **G**) containing the sample-specific variables for which one wishes to investigate the effect on the taxa composition. The environmental gradient then defines an environmental score $h_{im} = \boldsymbol{\alpha}_m^t \mathbf{c}_i$ for every sample i. This h_{im} can be seen as an equivalent of the row score r_{im}, but constrained to be a linear combination of sample-specific variables. Each taxon j is allowed to react to this environmental score in a different way through taxon-specific response functions $f_{jm}(h_{im})$. The constrained RC(M) model then becomes

$$\log(E(X_{ij})) = u_i + v_j + \sum_{m=1}^{M} \psi_m f_{jm}(\boldsymbol{\alpha}_m^t \mathbf{c}_i), \tag{3}$$

in which u_i, v_j and ψ_m play the same role as in models 1 and 2. The difference with the classical gradient analysis methods is that we use the response functions to model the *departure from homogeneity*. In this way, our method automatically accounts for differences in sequencing depth and taxon abundance. The environmental gradient $\boldsymbol{\alpha}_m$ is estimated by maximizing the likelihood ratio between a model with the taxon-specific response functions f_{jm} of model 3,

Fig 3. Effect of conditioning on unconstrained RC(M) ordination. (A): Unconstrained RC(M) sample ordination of the anterior nares samples of the HMP dataset without conditioning. (B): Ordination of the same sample, but after conditioning on the main sequencing center (Washington University genome center (WUGC), J. Craig Venter Institute (JCVI), Baylor College of Medicine (BCM) and Broad Institute (BI)). The ratio of the ψ parameters reflects the relative importance of the corresponding dimensions.

https://doi.org/10.1371/journal.pone.0205474.g003

and a model with a common response function, $f_m = f_{1m} = f_{2m} = \cdots = f_{pm}$, for all taxa. This encourages maximal niche separation between the taxa [10]. The correct shape of the response function has been the subject of theoretical debate [18, 23, 56], but it evidently depends on the taxon, as well as on the available sample-specific variables and their observed values. For easy interpretability we propose to use linear response functions $f_{jm}(h_{im}) = \beta_{0jm} + \beta_{1jm} h_{im}$, analogous to redundancy analysis [57] (see Section 2.1.5 of the Supplementary material for details of the estimation procedure). These response functions can easily be represented in two dimensions by an arrow originating in $(-\frac{\beta_{0j1}}{\beta_{1j1}}, -\frac{\beta_{0j2}}{\beta_{1j2}})$, with slope $\frac{\beta_{1j2}}{\beta_{1j1}}$ and magnitude proportional to $\sqrt{\beta_{1j1}^2 + \beta_{1j2}^2}$. The origin of the arrow then corresponds to the values of the environmental scores, (h_{i1}, h_{i2}), at which the taxon j does not depart from homogeneity in the first two dimensions. The direction and magnitude of the arrow indicate to which sample-specific variables the taxa abundances respond most strongly, and in which samples the departure from homogeneity is largest. See Fig 2B for an example of such an ordination. The (approximate) validity of the linearity assumption can be verified through diagnostic plots (see Fig 4 and Section 4.4.3 in the S1 Appendix).

A more flexible approach to modelling the taxa niches is provided by non-parametric estimation of the response functions with generalized additive models (GAMs) [58], similar to [24]. It provides possibly improved constrained sample ordination and gradient estimation, but also allows the researcher to study the way the taxa react to the environment with less prejudice. Fig 5 shows that different taxa can react entirely differently (and non-linearly) to changes in their environment. Quadratic response functions are frequently used implicitly [9] or explicitly [10, 59] to model unimodal response functions; they are also implemented in the

Fig 4. Diagnostic plots for the constrained RC(M) model with linear response functions on the Zeller data. (A) Triplot with samples coloured by deviance. No clusters of samples with high deviance are visible, which would have pointed to a group of poorly fit samples. (B) Residual plot in function of the first environmental gradient. A clear increase in positive deviance residuals is visible towards for positive environmental scores, which points to a violation of the linearity assumption. (C) Triplot with samples coloured by their influence on the parameter for the "Cancer" level of the diagnosis variable. On the right side of the plot, one sample with a strong negative and one with a strong positive influence on the parameter estimate are visible. These samples may deserve further scrutiny.

RCM R-package. They are, however, harder to depict in a triplot than linear response functions, while still providing less flexibility than non-parametrically estimated response functions. Moreover, for some taxa the estimated parameters of quadratic response functions may make the response curve convex rather than concave [60].

Diagnostic tools for the RC(M) ordination. Almost all ordination methods come with certain assumptions, but they are rarely explicitly mentioned, let alone checked by the user. The advantage of model-based approaches such as the RC(M) model, is that they explicitly state their assumptions, and allow them to be checked [15, 55]. Deviance residuals are a standard diagnostic tool in generalized linear models [61], and can be used to detect taxa and samples that poorly fit the model, or to detect misspecification of the response function. Influence functions can help to identify samples or taxa with a dominant role in shaping the final ordination [62]. Both of these diagnostic plots are available in the *RCM* package and can point researchers to outlying and possibly interesting samples and taxa that deserve further scrutiny (see Section 2.4 of the S1 Appendix for examples).

Simulation study

No-Signal Simulations. Fig 6 shows the pseudo F-statistics for the no-signal simulations with the negative binomial distribution. Since sequencing depths are assumed to be unrelated to the biological composition of a sample [12, 14], they should not affect the sample ordinations by, for example, forming clusters of samples with similar library sizes. Many methods appear to be insensitive to library size variability (as can be seen from their very small pseudo F-statistics),

Fig 5. RC(M) ordination with nonparametric response functions. One-dimensional triplot of the first dimension of the constrained RC(M) ordination with non-parametrically estimated response functions of the Zeller data. Coloured lines represent taxon response functions. The horizontal dotted line represents the expected taxon abundances under sample homogeneity. Only the eight taxa that react most strongly to changes in the environmental score are shown for clarity. Black labels show the variables constituting the gradient and vertical dashes at the bottom represent the sample scores. The horizontal positions of the variable labels with respect to the vertical dashed line at zero indicate how much they contribute to the environmental gradient; the vertical stacking is only for readability.

except the ordinations based on Hellinger distances, PCoA with Bray-Curtis dissimilarities on absolute and logged abundances, *gllvm* and the compositional data analysis (CoDa). The latter method's sensitivity to the library sizes can also be seen in S1 Fig, where the correlations between the sample scores and the library sizes for the first three dimensions are shown. It has been noted before that distance-based methods are sensitive to differences in dispersion between different sample groups [5, 17]. Our simulations confirm that all PCoA methods investigated, as well as CoDa, Hellinger distance, *gllvm* and our RC(M) method tend to cluster samples with the same dispersion levels together, even when all samples have equal taxa compositions (see Fig 6).

Biological Signal Simulations. As shown in Fig 7, the biological signal is best detected with the RC(M) method (large silhouette and pseudo-F values) and RC(M) succeeds best in
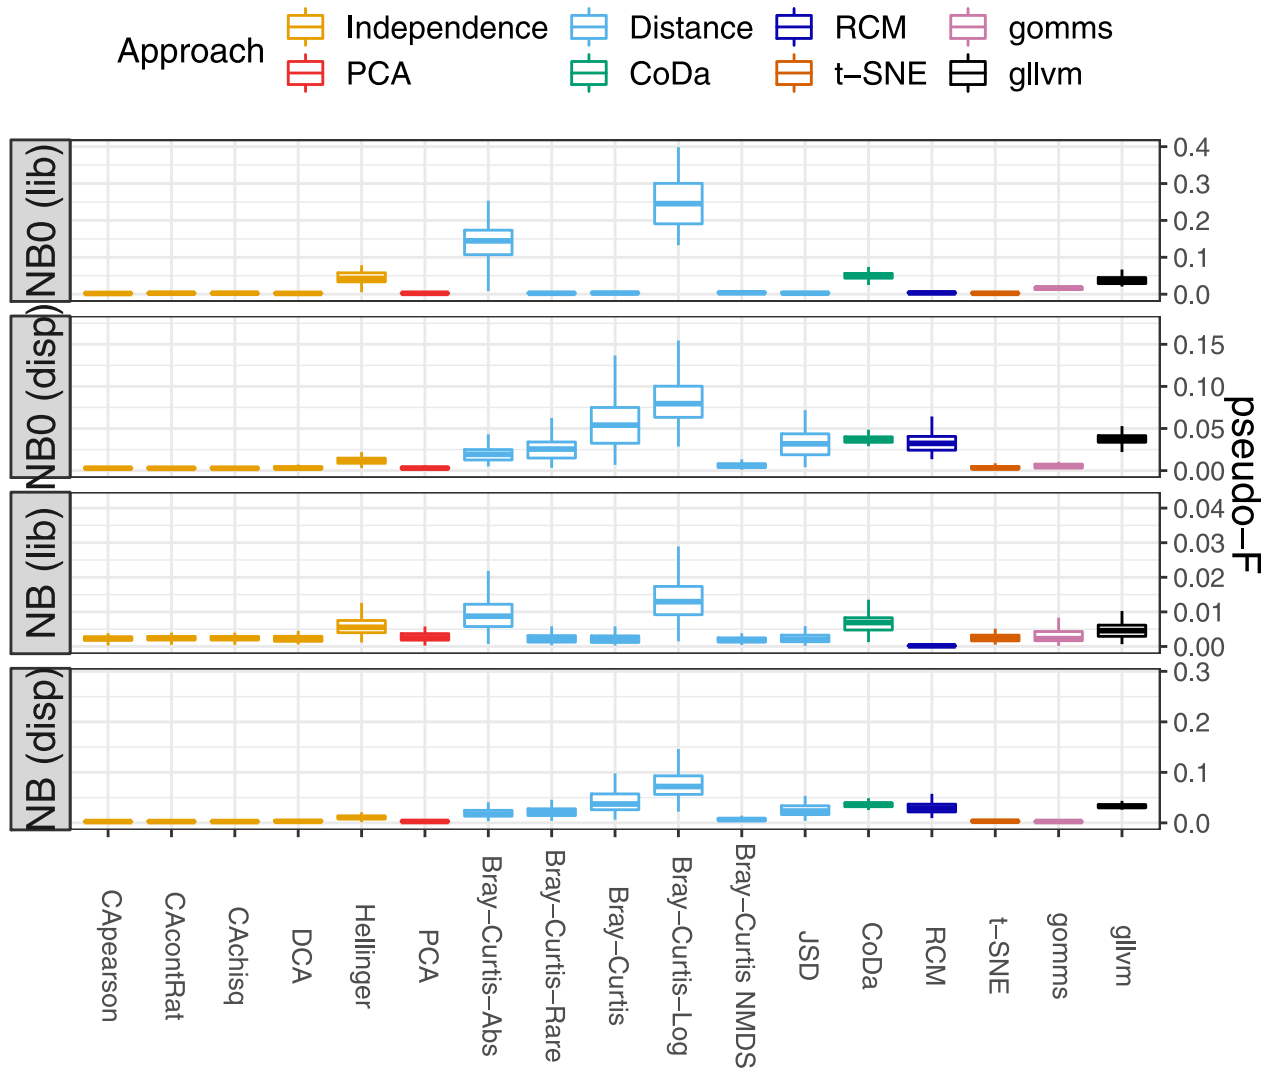
Fig 6. Results of simulations without signal. Boxplots of the pseudo-F statistic for sample clustering (y-axis) for several ordination methods (x-axis) for 100 parametric simulation runs. All samples have the same mean taxon composition, but four groups of samples differ in mean library sizes or mean dispersions. See Section Competitor ordination methods for the meaning of the abbreviations. As clustering according to library size or dispersion is undesirable, a small pseudo-F value is preferred. **Top**: Four groups with differences in library sizes. **Bottom**: Four groups with differences in dispersions. See S1 Appendix for details.

https://doi.org/10.1371/journal.pone.0205474.g006

identifying the driving taxa (large taxon ratio). This holds for all scenarios, except for data generated by the Dirichlet multinomial (DM) distribution. Also, detrended correspondence analysis (DCA) and PCA are good at detecting the important taxa. Note that the variability of the silhouette and especially of the pseudo F-value is seen to increase with their mean for all methods under study. This positive mean-variance relation is a known property of non-central F-distributions. More results, with similar conclusions, can be found in Section 3 of the S1 Appendix.

Discussion

Unconstrained and constrained ordination techniques that are currently employed in microbial ecology rely mainly on eigenvalues/eigenvectors and singular value decompositions.

Fig 7. Results of biological signal simulations. Boxplots of the silhouette (top), pseudo-F statistic (center) and taxon ratio (bottom) for several ordination methods (x-axis) over 100 parametric simulation runs. See Section Competitor ordination methods for the meaning of the abbreviations. 10% of the taxa were made differentially abundant in each of 4 sample groups, with a fold change of 5. As there are true differences in composition between the groups, a large pseudo-F value is preferred. Columns correspond to the simulation scenario: negative binomial (NB) (cor: data generation with taxon correlation, phy: phylogenetically correlated taxa were made differentially abundant), Dirichlet multinomial (DM) and zero-inflated negative binomial (ZINB). See S1 Appendix for details.

https://doi.org/10.1371/journal.pone.0205474.g007

Although having the advantage of computational efficiency, they are too rigid to deal with some of the more peculiar aspects of microbial amplicon sequencing data. For instance, sequencing depths varying between samples and taxon-wise overdispersions are two characteristics of microbiome data that may distort ordinations [5, 15]. One possible reason why these flaws received little attention, is because the assumptions underlying these ordination methods are rarely stated explicitly, and hence they are almost never checked. Researchers in microbial ecology should become more aware of assumptions and limitations of the ordination methods. Ordination methods developed for ecological data with directly observed species counts may no longer be valid for sequencing data, because sequence counts are only a proxy of abundance and the biological and technical variability show specific characteristics. Dimension reduction for plotting inevitably entails information loss, but using ordination methods that are inappropriate for the data type may yield misleading results. Another reason for the wide use of distance-based approaches may be the computational speed of their underlying matrix calculations. Yet on modern computers, certain simple, model-based methods can also be fitted within reasonable time spans.

Distance-based methods are currently very popular ordination methods in microbiomics. However, by calculating distances between samples, the information on which taxa discriminate the samples is discarded. As a result, distance based methods cannot directly identify which taxa drive the differences between samples, limiting their use for data exploration.

Compositional data analysis (CoDa) analyzes ratios between taxon counts rather than the counts themselves. Although sequencing data often should be treated as compositional indeed, these methods ignore the count origin and the associated heteroscedasticity. As a result, the sample scores of their ordinations correlate strongly with the library sizes, which are considered as technical artefacts. This is highly problematic for the interpretation of their ordination diagrams. Especially in datasets with a low signal-to-noise ratio, differences in library sizes, rather than biological signal, may be depicted in the ordination graphs. Because of the common association of library sizes with sample-specific variables, this may incorrectly confirm the researcher's prior beliefs in differences in microbiome composition, whereas actually, none exist.

Despite their slightly longer computation times (about one minute per dataset with our *RCM* package), ordination methods based on count regression models are more flexible to deal with these issues, and have gained popularity over the recent years. Model-based ordination methods can include an offset to account for varying sequencing depths, and can be easily augmented with skewed count distributions with taxon-wise parameters to address heteroscedasticity. Furthermore, they can condition out the effect of other confounding variables. The main idea is that interaction terms between samples and taxa capture departures from equal taxa composition in the samples. These interaction terms can then be plotted to visualize the strongest signal in the dataset. These strongest signals need not necessarily come from the most abundant taxa. Since an explicit mean model is stated, standard diagnostic tools can be employed to assess model assumptions. Moreover, outlying or influential observations can be identified, which can reveal useful information to researchers.

Latent variable models for ordination of ecological data have been developed over the recent years. They share many of the advantages of row-column interaction models, such as explicit model statement and the option of conditioning on baseline covariates. However, the inclusion of latent sample variables as random effects in a Bayesian framework greatly increases the computational burden. The inclusion of random effects does not improve the explorative ordination plots, such that simpler fixed effects may be preferable. If statistical inference were the goal, then random effects would be preferred. Also, latent variable models have no counterpart for constrained ordination.

Just as row-column interaction models, correspondence analysis tries to represent departures from sample homogeneity in few dimensions. Still, for skewed and overdispersed data, an additive model for departure from equal sample composition is inappropriate and produces ordination plots dominated by outliers. A multiplicative model as employed in the RC(M) model is more appropriate for these data.

The performance of ordination methods can be assessed quantitatively through simulations. Our comprehensive simulation study confirms a good performance of the RC(M) method, both in terms of sample separation as in the identification of taxa that contribute to these separations. The RC(M) method is not sensitive to library size variation, but, just as many other ordination methods, it is somewhat sensitive to differences in dispersions between samples.

We believe the potential of row-column interaction models is underemployed in the analysis of all types of high-dimensional data, despite the availability of contemporary fitting algorithms and computing power. However, given the reasonably good performance of CoDa techniques in our simulations, a combination of model-based approaches that correctly model the mean-variance structure, and models that account for compositionality would probably further improve visualization methods for the microbiome. Also, extending the RC(M) method to allow for significance testing (e.g. through permutations as in [63]) would be an interesting avenue for future research.

Constrained ordinations include sample-specific variables in the visualization. Despite a very rich theoretical foundation, they are less frequently employed in the microbial ecology practice. We combined the row-column interaction model with flexible response modeling using linear response functions as well as non-parametrically estimated response functions. Linear response functions yield easily interpretable triplots, and the linearity assumption can be verified using diagnostic plots. Non-parametrically estimated response function allow maximal flexibility in modelling the taxon niches. Our method uniquely combines unconstrained and constrained ordination into the same framework for fitting and plotting, which greatly facilitates comprehensive exploration of microbiome datasets.

Our methods for visualization of microbiome data are implemented in the R-package *RCM* (available at http://github.com/CenterForStatistics-UGent/RCM). The package comes with a custom-written fitting algorithm for the RC(M) model as well as several ready-to-use plotting functions.

Supporting information

S1 Appendix. A detailed discussion of the RC(M) method, with illustrations on real datasets. Further, a detailed description of the setup and results of the simulation study, followed by a list of software versions.
(PDF)

S1 File. Auxiliary R-code. All R-code for making the graphs shown in the publication, along with the code for the simulation study.
(GZ)

S2 File. Data. All datasets used in this publication.
(GZ)

S1 Fig. Correlations of library sizes and row scores. Boxplots with the correlation of sample scores with observed library sizes (y-axis) for different ordination methods (x-axis). Side panels indicate the different parametric simulation scenarios, see Section Simulation study for an

explanation of the codes used. Top panels show the dimension of the sample score.
(EPS)

Acknowledgments

Author Contributions

Conceptualization: Stijn Hawinkel, Olivier Thas.

Formal analysis: Stijn Hawinkel.

Methodology: Stijn Hawinkel.

Software: Stijn Hawinkel, Frederiek-Maarten Kerckhof.

Supervision: Olivier Thas.

Visualization: Stijn Hawinkel.

Writing – original draft: Stijn Hawinkel, Olivier Thas.

Writing – review & editing: Frederiek-Maarten Kerckhof, Luc Bijnens, Olivier Thas.

References

1. Stackebrandt E, Goebel BM. Taxonomic Note: A Place for DNA-DNA Reassociation and 16S rRNA Sequence Analysis in the Present Species Definition in Bacteriology. International Journal of Systematic and Evolutionary Microbiology. 1994; 44(4):846–849. https://doi.org/10.1099/00207713-44-4-846

2. Gower JC. In: Principal Coordinates Analysis. John Wiley & Sons, Ltd; 2005.

3. Richardson. Multidimensional Psychophysics. Psychological Bulletin. 1938; 35:659–660.

4. Oksanen, J, Blanchet, FG, Friendly, M, Kindt, R, Legendre, P, McGlinn, D, et al. vegan: Community Ecology Package; 2017. Available from: https://CRAN.R-project.org/package=vegan.

5. Warton DI, Wright ST, Wang Y. Distance-based multivariate analyses confound location and dispersion effects. Methods in Ecology and Evolution. 2012; 3(1):89–101. https://doi.org/10.1111/j.2041-210X.2011.00127.x

6. Weiss S, Xu ZZ, Peddada S, Amir A, Bittinger K, Gonzalez A, et al. Normalization and microbial differential abundance strategies depend upon data characteristics. Microbiome. 2017; 5(27).

7. Wong RG, Wu JR, Gloor GB. Expanding the UniFrac Toolbox. PLOS ONE. 2016; 11(9):1–20. https://doi.org/10.1371/journal.pone.0161196

8. Benzecri JP. L'analyse des données. Population. 1975; 30(6):1190.

9. ter Braak CJF. Canonical Correspondence Analysis: A New Eigenvector Technique for Multivariate Direct Gradient Analysis. Ecology. 1986; 67(5):1167–1179. https://doi.org/10.2307/1938672

10. Zhu M, Hastie T, Walther G. Constrained ordination analysis with flexible response functions. Ecological Modelling. 2005; 187:524–536. https://doi.org/10.1016/j.ecolmodel.2005.01.049

11. Zeller G, Tap J, Voigt AY, Sunagawa S, Kultima JR, Costea PI, et al. Potential of fecal microbiota for early-stage detection of colorectal cancer. Mol Syst Biol. 2014; 10(766). https://doi.org/10.15252/msb.20145645 PMID: 25432777

12. Gloor GB, Reid G. Compositional analysis: A valid approach to analyze microbiome high-throughput sequencing data. Can J Microbiol. 2016; 62(8):692–703. https://doi.org/10.1139/cjm-2015-0821 PMID: 27314511

13. Costea PI, Zeller G, Sunagawa S, Bork P. A fair comparison. Nature Methods. 2014; 11:359. https://doi.org/10.1038/nmeth.2897 PMID: 24681719

14. McMurdie PJ, Holmes S. Waste Not, Want Not: Why Rarefying Microbiome Data Is Inadmissible. PLoS Comput Biol. 2014; 10(4):e1003531. https://doi.org/10.1371/journal.pcbi.1003531 PMID: 24699258

15. Hui FKC, Taskinen S, Pledger S, Foster SD, Warton DI. Model-based approaches to unconstrained ordination. Methods in Ecology and Evolution. 2015; 6(4):399–411. https://doi.org/10.1111/2041-210X.12236

16. Pledger S, Arnold R. Multivariate methods using mixtures: Correspondence analysis, scaling and pattern-detection. Computational Statistics & Data Analysis. 2014; 71(C):241–261. https://doi.org/10.1016/j.csda.2013.05.013

17. Sohn MB, Li H. A GLM-based latent variable ordination method for microbiome samples. Biometrics. 2017; p. e–pub ahead of print.

18. Yee TW, Hadi AF. Row–column interaction models, with an R implementation. Computational Statistics. 2014; 29(6):1427–1445. https://doi.org/10.1007/s00180-014-0499-9

19. Goodman L. Simple Models for the Analysis of Association in Cross-Classifications Having Ordered Categories. 1979; 74:537–552.

20. Xu L, Paterson AD, Xu W. Bayesian latent variable models for hierarchical clustered count outcomes with repeated measures in microbiome studies. Genetic Epidemiology. 2017; 41(3):221–232. https://doi.org/10.1002/gepi.22031 PMID: 28111783

21. Hui FKC, Warton DI, Ormerod JT, Haapaniemi V, Taskinen S. Variational Approximations for Generalized Linear Latent Variable Models. Journal of Computational and Graphical Statistics. 2017; 26(1):35–43. https://doi.org/10.1080/10618600.2016.1164708

22. Anderson MJ, Willis TJ. Canonical analysis of principal coordinates: A useful method of constrained ordination for ecology. Ecology. 2003; 84(2):511–525. https://doi.org/10.1890/0012-9658(2003)084%5B0511:CAOPCA%5D2.0.CO;2

23. ter Braak CJF, Prentice IC. A Theory of Gradient Analysis. 1988; 18(Supplement—C):271–317.

24. Yee TW. Constrained additive ordination. Ecology. 2006; 87(1):203–213. https://doi.org/10.1890/05-0283 PMID: 16634311

25. Robinson MD, Smyth GK. Moderated statistical tests for assessing differences in tag abundance. Bioinformatics. 2007; 23(21):2881–2887. https://doi.org/10.1093/bioinformatics/btm453 PMID: 17881408

26. R Core Team. R: A Language and Environment for Statistical Computing; 2015. Available from: http://www.R-project.org/.

27. Minchin P. An Evaluation of the Relative Robustness of Techniques for Ecological Ordination. 1987; 69:89–107.

28. Faith DP, Minchin P, Belbin L. Compositional dissimilarity as a robust measure of ecological distance. 1987; 69:57–68.

29. Legendre P, Gallagher ED. Ecologically meaningful transformations for ordination of species data. Oecologia. 2001; 129(2):271–280. https://doi.org/10.1007/s004420100716 PMID: 28547606

30. Kuczynski J, Liu Z, Lozupone C, McDonald D, Fierer N, Knight R. Microbial community resemblance methods differ in their ability to detect biologically relevant patterns. Nat Methods. 2010; 7(10):813–819. https://doi.org/10.1038/nmeth.1499 PMID: 20818378

31. Ruokolainen L, Salo K. Differences in performance of four ordination methods on a complex vegetation dataset. Science. 2006; 43:269–275.

32. Fukuyama J, McMurdie PJ, Dethlefsen L, Relman DA, Holmes S. Comparisons of distance methods for combining covariates and abundances in microbiome studies. Pac Symp Biocomput. 2012; p. 213–224. PMID: 22174277

33. Schmidt TSB, Rodrigues JFM, von Mering C. A family of interaction-adjusted indices of community similarity. The Isme Journal. 2016; 11(3):791–807. https://doi.org/10.1038/ismej.2016.139 PMID: 27935587

34. Dray S, Pavoine S, de Cárcer DA. Considering external information to improve the phylogenetic comparison of microbial communities: A new approach based on constrained Double Principal Coordinates Analysis (cDPCoA). Molecular Ecology Resources. 2015; 15(2):242–249. https://doi.org/10.1111/1755-0998.12300 PMID: 24974884

35. Clarke K. Nonparametric Multivariate Analyses of Changes in Community Structure. 1993; 18:117–143.

36. Peterson J, Garges S, Giovanni M, McInnes P, Wang L, Schloss JA, et al. The NIH Human Microbiome Project. Genome Res. 2009; 19(12):2317–2323. https://doi.org/10.1101/gr.096651.109 PMID: 19819907

37. AmericanGut org. The American gut project. https://githubcom/biocore/American-Gut/blob/master/data/AG/AG_100nttxt. 2015.

38. Kostic AD, Gevers D, Pedamallu CS, Michaud M, Duke F, Earl AM, et al. Genomic analysis identifies association of Fusobacterium with colorectal carcinoma. Genome Res. 2012; 22(2):292–298. https://doi.org/10.1101/gr.126573.111 PMID: 22009990

39. Turnbaugh PJ, Ridaura VK, Faith JJ, Rey FE, Knight R, Gordon JI. The Effect of Diet on the Human Gut Microbiome: A Metagenomic Analysis in Humanized Gnotobiotic Mice. Sci Transl Med. 2009; 1 (6):6ra14–6ra14. https://doi.org/10.1126/scitranslmed.3000322 PMID: 20368178

40. Props R, Kerckhof FM, Rubbens P, De Vrieze J, Hernandez Sanabria E, Waegeman W, et al. Absolute quantification of microbial taxon abundances. The ISME Journal. 2016; 11:584–587. https://doi.org/10.1038/ismej.2016.117 PMID: 27612291

41. Hawinkel S, Mattiello F, Bijnens L, Thas O. A broken promise: Microbiome differential abundance methods do not control the false discovery rate. Briefings in Bioinformatics. 2017; p. bbx104.

42. La Rosa PS, Brooks JP, Deych E, Boone EL, Edwards DJ, Wang Q, et al. Hypothesis Testing and Power Calculations for Taxonomic-Based Human Microbiome Data. PLoS ONE. 2012; 7(12):e52078. https://doi.org/10.1371/journal.pone.0052078 PMID: 23284876

43. Kurtz ZD, Müller CL, Miraldi ER, Littman DR, Blaser MJ, Bonneau RA. Sparse and Compositionally Robust Inference of Microbial Ecological Networks. PLoS Comput Biol. 2015; 11(5):e1004226. https://doi.org/10.1371/journal.pcbi.1004226 PMID: 25950956

44. Chen J, Bittinger K, Charlson ES, Hoffmann C, Lewis J, Wu GD, et al. Associating microbiome composition with environmental covariates using generalized UniFrac distances. Bioinformatics. 2012; 28 (16):2106–2113. https://doi.org/10.1093/bioinformatics/bts342 PMID: 22711789

45. McMurdie PJ, Holmes S. phyloseq: An R package for reproducible interactive analysis and graphics of microbiome census data. PLoS ONE. 2013; 8(4):1–11. https://doi.org/10.1371/journal.pone.0061217

46. Gower J, Lubbe S, le Roux N. Understanding Biplots. vol. 1; 2011. https://doi.org/10.1002/9780470973196

47. Rao CR. A review of canonical coordinates and an alternative to correspondence analysis using Hellinger distance. Qüestiió. 1995; 19(1-3):23–63.

48. van der Maaten L, Hinton G. Visualizing data using t-SNE. Journal of Machine Learning Research. 2008; 9(Nov):2579–2605.

49. Niku J, Brooks W, Herliansyah R, Hui FKC, Taskinen S, Warton DI. Gllvm: Generalized Linear Latent Variable Models; 2018. Available from: https://CRAN.R-project.org/package=gllvm.

50. Rousseeuw PJ. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. Journal of Computational and Applied Mathematics. 1987; 20(Supplement—C):53–65. https://doi.org/10.1016/0377-0427(87)90125-7

51. Anderson MJ. A new method for non-parametric multivariate analysis of variance. Austral Ecology. 2001; 26(1):32–46. https://doi.org/10.1111/j.1442-9993.2001.01070.pp.x

52. Hiergeist A, Reischl U, Gessner A. Multicenter quality assessment of 16S ribosomal DNA-sequencing for microbiome analyses reveals high inter-center variability. International Journal of Medical Microbiology. 2016; 306(5):334–342. https://doi.org/10.1016/j.ijmm.2016.03.005 PMID: 27052158

53. Legendre P, Legendre LFJ. Numerical Ecology. Developments in Environmental Modelling. Elsevier Science; 2012.

54. Niku J, Warton DI, Hui FKC, Taskinen S. Generalized Linear Latent Variable Models for Multivariate Count and Biomass Data in Ecology. Journal of Agricultural, Biological and Environmental Statistics. 2017; 22(4):498–522. https://doi.org/10.1007/s13253-017-0304-7

55. Warton DI, Blanchet FG, O'Hara RB, Ovaskainen O, Taskinen S, Walker SC, et al. So Many Variables: Joint Modeling in Community Ecology. Trends in Ecology & Evolution. 2015; 30(12):766–779. https://doi.org/10.1016/j.tree.2015.09.007

56. Macarthur R, Levins R. The Limiting Similarity, Convergence, and Divergence of Coexisting Species. The American Naturalist. 1967; 101(921):377–385. https://doi.org/10.1086/282505

57. van den Wollenberg AL. Redundancy analysis an alternative for canonical correlation analysis. Psychometrika. 1977; 42(2):207–219. https://doi.org/10.1007/BF02294050

58. Hastie T, Tibshirani R. Generalized Additive Models. Statistical Science. 1986; 1(3):297–310. https://doi.org/10.1214/ss/1177013609

59. Yee TW. A new technique for maximum-likelihood canonical gaussian ordination. Ecological Monographs. 2004; 74(4):685–701. https://doi.org/10.1890/03-0078

60. Zhang Y, Thas O. Constrained Ordination Analysis with Enrichment of Bell-Shaped Response Functions. PLOS ONE. 2016; 11(4):1–21. https://doi.org/10.1371/journal.pone.0154079

61. McCullagh P, Nelder JA. Generalized Linear Models, Second Edition. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis; 1989.

62. Hampel FR, Ronchetti EM, Rousseeuw PJ, Stahel WA. Robust Statistics: The Approach Based on Influence Functions. vol. 07. John Wiley & Sons, Inc.; 2011.

63. Anderson MJ. In: Permutational Multivariate Analysis of Variance (PERMANOVA). American Cancer Society; 2017. p. 1–15. Available from: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118445112.stat07841.