

# Semisynthetic Simulation for Microbiome Data Analysis

Kris Sankaran<sup>1</sup>, Saritha Kodikara<sup>2</sup>, Jingyi Jessica Li<sup>3,4,5</sup>, and Kim-Anh Lê Cao<sup>2</sup>

<sup>1</sup>Department of Statistics, University of Wisconsin-Madison, 1300 University Ave, 53703 WI, United States

<sup>2</sup>Melbourne Integrative Genomics, School of Mathematics and Statistics, University of Melbourne, Building 184/30 Royal Parade, 3052, VIC Australia

<sup>3</sup>Department of Statistics and Data Science, University of California, Los Angeles, 520 Portola Plaza, 90095, CA, United States

<sup>4</sup>Department of Human Genetics, University of California, Los Angeles, 695 Charles E Young Dr S, 90095, CA, United States

<sup>5</sup>Department of Biostatistics, University of California, Los Angeles, 650 Charles E. Young Dr S, 90095, CA, United States

October 14, 2024

## Abstract

High-throughput sequencing data lie at the heart of modern microbiome research. Effective analysis of these data requires careful preprocessing, modeling, and interpretation to detect subtle signals and avoid spurious associations. In this review, we discuss how simulation can serve as a sandbox to test candidate approaches, creating a setting that mimics real data while providing ground truth. This is particularly valuable for power analysis, methods benchmarking, and reliability analysis. We explain the probability, multivariate analysis, and regression concepts behind modern simulators and how different implementations make trade-offs between generality, faithfulness, and controllability. Recognizing that all simulators only approximate reality, we review methods to evaluate how accurately they reflect key properties. We also present case studies demonstrating the value of simulation in differential abundance testing, dimensionality reduction, network analysis, and data integration. Code for these examples is available in an online tutorial (<https://go.wisc.edu/8994yz>) that can be easily adapted to new problem settings.

## 1 Introduction

Microbial communities play a central role in human and ecological health. Advances in sequencing technology and statistical methods have made it possible to characterize these communities at unprecedented levels of precision [37, 31, 82]. However, statistical methods have to contend with microbiome-specific challenges, like sparse read coverage, large library size variations, uncertainties in the resolved taxa, and overdispersion [71, 104, 42]. These difficulties intensify as scientific goals grow more complex. This is especially evident as microbiome data analysis shifts away from descriptive studies towards modelling that often requires experimental designs with multiple batches, longitudinal sampling, or complementary assays [29, 19, 51]. In this context, effective use of statistical methods hinges on many steps: framing precise questions, preparing suitable data, applying appropriate methods, and delivering accurate interpretations.

Given the uncertainties present in real data, guiding analysis using simulations where the ground truth is known can be tremendously helpful. Simulation has a long history in shaping microbiome data analysis. For example, McMurdie and Holmes [71] applied rarefaction to datasets simulated from a negative binomial model to clarify its impact on downstream inferences. Similarly, Kodikara et al. [52] employed an autoregressive model to assess method performance for longitudinal studies. In fact, most methodological research requires benchmarking on simulated data. Despite these strong precedents, the field is only beginning to formalize, evaluate, and share reusable simulators. In particular, recent advances have proposed simulators that leverage existing *template data* – real experimental data whose patterns the simulator should mimic –

thus simultaneously reducing development time while improving fidelity. In the remainder of this review, words in *italic font* are listed in our glossary of terms, Table 1.

This review introduces readers to emerging trends in simulation, walks through example use cases, and distills best practices. While research on simulation methods often emphasizes realism of the simulated data, their applications are typically only discussed at a high level. Here, we instead explore in-depth applications of simulated data. We first review existing packages and highlight their potential pitfalls (see Table 2). We then illustrate how researchers can get the most out of their microbiome data through simulation for various analytical tasks, ranging from effective experimental designs to data analysis strategies. In particular, this review focuses on three use cases: Power analysis, methods benchmarking, and reliability analysis.

**Power analysis** is an important step in experimental design that informs the sample size required to detect different effect sizes, thereby enabling studies to be done with minimal resources, and without compromising scientific integrity and rigor [112]. Several power calculators have been proposed for microbiome studies [50, 23, 83, 11]. However, these tools lack the flexibility of full simulators, which can be adapted to problem- and data-specific contexts. Most power calculators are limited to case-control designs and, without access to template data, they may make questionable distributional assumptions. In contrast, researchers using simulators can ensure their models reflect important properties of template data and can compare various design choices, like class imbalance and sampling times, that go beyond sample size alone.

**Benchmarking** is essential for identifying the most suitable method for a specific experimental design and data type [25]. Simulation is useful to benchmark methods when ground truth is scarce. For example, in batch effect integration, it is important to preserve biological effects, but these can be difficult to pinpoint without simulations that provide some ground truth (illustrated in the section ‘Batch effect correction’). Further, benchmarking new methods on various plausible datasets provides more insight than simply identifying the best performer in a single simulation scenario. To this end, we can train a simulator to emulate real data and then alter it to reflect multiple hypothetical scenarios. Even when ground truth exists, simulations help gauge robustness to dataset perturbations (e.g., reducing the true effect sizes), and testing many simulated scenarios is often easier than collecting multiple real datasets. Formal simulation techniques provide a more automatic approach to benchmarking and allow researchers to focus on data analysis and method selection, rather than spending time programming simulators from scratch.

**Reliability analysis** can be enabled through simulation. Many modern data analysis workflows lack sufficient theory to guide practice, and simulations offer valuable sanity checks. For example, simulation studies have found that sparse canonical correlation analysis (sCCA) can result in high false discovery rates [33, 39]. Although these studies were motivated by neuroscience, sCCA is also widely used in microbiome data integration [7, 90]. In a similar spirit, the section ‘Omics data integration’ highlights how data integration can introduce unexpected artifacts into dimensionality reduction visualizations. Hence, simulations help prevent misinterpretations that could misdirect research priorities.

This review makes the following contributions:

1. **Overview of simulation workflows:** We describe the main ingredients of modern simulation algorithms and associated assumptions. This background ensures that methods are not treated as black boxes and guides their effective use.
2. **Evaluation criteria:** We outline how to assess whether simulated data match the properties of previously observed template data, and how simulated data can inform methodological improvements.
3. **Case studies:** We present realistic case studies demonstrating how simulation can support power analysis, benchmark competing methods, and guarantee reliable conclusions.

This review is accompanied by an online tutorial (<https://go.wisc.edu/8994yz>). Each chapter of the tutorial starts with a dataset discussion, walks through the process of designing simulators, and applies the resulting models to address common questions about microbiome experimental design, method benchmarking, and result interpretation.

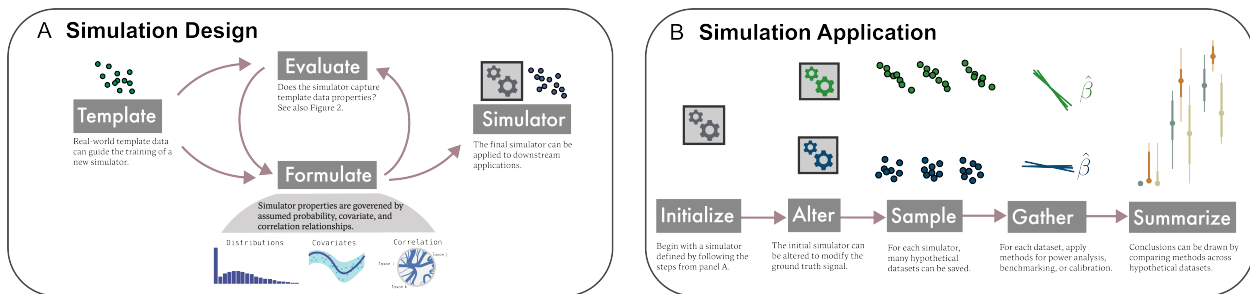


Figure 1: (A) Simulation design is an iterative process, where choices of probability distribution, experimental and biological covariates, and correlation structure can be refined according to evaluation criteria that draw attention to differences between real and template data. See Figure 2 for example criteria. (B) Given a satisfactory simulator, the same workflow applies to power analysis, benchmarking, and reliability analysis. The initial simulator can be modified to reflect hypothetical experimental or biological settings, like changes in the signal effect size or the sample size. Outputs from competing approaches can be gathered and interpreted to guide experimental design and analysis.

## 2 Simulation Workflows

Simulators vary widely, and their realism and relevance to downstream tasks is context-dependent. To help navigate this landscape, we first categorize simulators and then outline the workflow for building and refining them.

**De novo and template-based simulators.** We first note the distinction between *de novo* and template-based simulators. De novo simulators require users to specify parameters related to experimental design and distributional properties. These parameters cannot be automatically matched to real, template data. For example, in the splatter simulator [109], users can include treatment and batch effects, but these are drawn from the simulator’s internal generation mechanism. In contrast, template-based simulators first estimate the impact of sample-level *covariates* using real data. Though this does require an initial template, it can lead to more realistic synthetic data. These data are often called *semisynthetic*, reflecting the influence of the template. As many public catalogs of 16S, metagenomics, and metabolomics data are now available [78, 74], it has become easier for researchers to access relevant template data for various possible analysis. We caution that different research communities have used different terms to describe the concept of template-based simulation. For example, “semisynthetic” is used in metabolomics [105, 81] and microbiome studies [91, 69], while “reference-based” is common in single-cell genomics [62, 12]. Although de novo and template-based simulators differ in their requirements, similar factors guide their application. We will review these workflow considerations next.

**Formulation and application.** In the formulation phase (Figure 1A), a simulation model is created by adjusting parameters until the generated data pass a series of evaluation checks. In the application phase (Figure 1B), several altered versions of the simulator can be defined according to the simulation study’s goals. For example, to characterize false discovery rates, we can introduce synthetic negative control features designed to lack associations with the outcome. For each altered version of the simulator, we can generate multiple datasets, apply candidate analysis strategies, and gather summary statistics quantifying their performance. We provide practical examples of both phases in the ‘Case Studies’ section.

The formulation phase relies on concepts from probability, statistical estimation, regression, and multivariate analysis. Probabilistic models enable sampling of new data with appropriate distributional properties, *multivariate models* induce plausible associations across multiple taxa, and regression methods ensure that simulations reflect biological or experimental influences.

Table 1: A glossary of simulation terms.

Term	Definition
Covariate	A biological or experimental characteristic of a sample that is thought to impact community profiles and is used to train the simulator.
Copula	A statistical approach for modeling correlation in non-Gaussian data, often used to model associations between taxa in a community.
De Novo Simulator	A simulator whose mechanism is defined without reference to external, template data.
Factor Model	A model that induces correlation across features through latent variables. Like copulas, these are often used to induce associations between taxa.
Fit-for-Purpose	Approaches that assess simulation quality by referring to small subsets of taxa.
Global Evaluation	Approaches that assess simulation quality by referring to the entire community (all taxa).
Marginal Distribution	A probability model that reflects abundance patterns for a specific taxon.
Multivariate Model	A probability model that reflects abundance patterns across the entire community (all taxa).
Outcome-Specific	Evaluation criteria that target the simulator’s downstream application.
Reliability Analysis	The use of simulated data to calibrate interpretations of real data analysis.
Semisynthetic Data	The output from a simulator that has been designed to mimic external, template data.
Template Data	Previously gathered experimental data that can be used to train a simulator.

## 2.1 Probability distributions

The choice of probability distribution dictates many properties of the generated data. Distributions must be carefully selected, as properties that might be easy to manipulate in one distributional family might be difficult to modify in another. For example, sparsity can be more easily modified in a zero-inflated vs. ordinary negative binomial distributions. We first review univariate, then multivariate distributions.

**Univariate distributions for simulation.** For a specific taxon, we need to decide whether to simulate counts, proportions, or real numbers. Count distributions can be used to simulate amplicon sequence variant or metagenomics data before having applied any transformations. Common count distributions include the Poisson, negative binomial, and hypergeometric distributions. The Poisson distribution arises when counting events that, though individually unlikely, become common due to repeated opportunities for them to occur. For example, a stretch of DNA is unlikely to align to any given read, but given enough reads, we would expect a Poisson number of them to align. However, in real data, this often gives a poor approximation – the variance in observed counts is often higher than the Poisson can capture. Such overdispersion is more appropriately modeled using the negative binomial distribution. Finally, if the data exhibit more zeros than either Poisson or negative binomial distribution allows, then zero-inflation can be employed to introduce additional sparsity, as in zero-inflated negative binomial (ZINB) used in the single-cell RNA-seq simulator ZINB-WaVE [87]. Alternatively, presence-absence can be modeled separately from abundances, as in the simulator MIDASim [38].

When data are transformed, count distributions no longer apply. Depending on the transformation, different distributions can be used to model nonnegative, interval, or arbitrary real values. For nonnegative values, options include the truncated normal, log-normal, or variants of gamma distributions. The truncated and log-normal distributions enforce nonnegativity by either truncating or exponentiating a normal distribution. These distributions are used in microbiomeDASim [106], SparseDossa 2 [68], and CAMISIM [26]. Proportions within the interval  $[0, 1]$  can be represented by beta, Dirichlet, or generalized Gamma distribution, which include parameters for mean and concentration along the boundaries. The simulator MIDASim uses this approach to model relative abundance-transformed data. If data have been transformed



to include both positive and negative real values, then they can often be modeled with normal or Student’s T distributions, the latter being more appropriate when outliers are present. Further, such transformations can enable generative adversarial network models to perform well, as in simulators DeepMicroGen [10] and MB-GAN [88]. Continuous distributions can also be used as a preliminary sampling step for hierarchical count models. For example, the means of a count model can first be modeled as a Gamma distribution, allowing estimates for rare taxa to borrow strength from more abundant taxa. This hierarchical approach is used by metaSparsSim [79], which draws Gamma marginal distributions for individual taxa before sampling from a multivariate hypergeometric distribution for all taxa jointly.

**Multivariate distributions for simulation.** It is important that simulators generate realistic community profiles, capturing relational, multivariate structure and not just univariate, per-taxon abundances. Two common strategies for modeling these associations are (1) learning a multivariate normal covariance in an appropriately transformed space or (2) incorporating shared latent variation across features (taxa). Strategy (1) uses the multivariate normal distribution’s covariance parameter to control the relationship between features. These multivariate normal samples can then be transformed into sequencing-specific data distributions. For example, this approach is adopted by simulators using the Logistic-normal Multinomial distribution [2, 111]. This method induces correlations across features by first sampling from a correlated, multivariate normal distribution. This correlated vector is transformed into a vector of proportions using a log-ratio transformation. Given an overall sequencing depth, simulated reads are then allocated to individual taxa based on these proportions. Alternatively, a correlated multivariate normal distribution can be transformed into a multivariate distribution with known univariate marginals, a process known as *copula modeling* [44, 14]. Copula models can be used even when the marginals are drawn from different distributions. This makes copulas easily adaptable, and they are used by simulators MIDASim [38], SparseDOSSA 2 [68], and single-cell simulators scDesign2 [100] and scDesign3 [97].

In contrast, strategy (2) ties features together by assuming a latent, low-dimensional vector. The shared source of variation induces downstream correlations. This is often accomplished through variations of *factor models*. For example, ZINB-WaVE simulates entries from a model whose parameters vary according to a low-dimensional latent vector representing unobserved sample-level properties. Similarly, Latent Dirichlet Allocation assumes that samples are drawn from a multinomial distribution whose mean depends on a low-dimensional community membership vector [89]. In both univariate and multivariate settings, interpretable parameterizations can support modifications of the simulation mechanism. These modifications are helpful for validating workflows through computational negative or positive controls. For example, the mean parameter of a negative binomial model can be adjusted to reflect stronger or weaker treatment effects. In contrast, flexible machine-learning-based multivariate simulators, like MB-GAN and DeepMicroGen, can be challenging to alter in this way. Even if their simulated data are realistic, it is difficult to control them and enforce desired constraints.

## 2.2 Accounting for experimental and biological factors

How can we model differences across experimental or biological conditions? For example, a taxon’s abundance may change gradually over time or may be a marker of disease status. In this case, histograms of individual taxon abundances may reveal complex patterns, like separate peaks for disease and healthy groups, which cannot be captured by models that assume independent draws from the same probability distribution. To simulate these patterns, we can specify distributional parameters based on sample characteristics. In addition to producing more realistic data, conditioning on sample covariates enables more precise control. For example, simulators that take into account time or disease status can generate data at a denser sequence of timepoints or different sample sizes for healthy vs. disease patients, both of which can guide experimental design and method benchmarking.

Many simulators are tailored to specific experimental designs. For example, MIDASim [38] and SparseDOSSA 2 [68] support simulation from case-control designs. In both simulators, some taxa share parameters across case and control groups, representing synthetic negative controls. To create a subset of taxa that differ across groups, representing synthetic positive controls, taxa parameters can be allowed to vary. The greater the difference between parameters, the stronger the true effect. These simulators are particularly useful for

power analysis and benchmarking for differential abundance studies. Other simulators have been created specifically for longitudinal designs. For example, microbiomeDASim [106] allows taxonomic abundance to vary over time based on various plausible trends, to mimic brief disruptions or gradual development. Similarly, miaSim [27], CAMISIM [26], and DeepMicroGen [10] are designed to capture longitudinal dynamics.

Though these simulators streamline work with specific experimental designs, it can be useful to generalize and map arbitrary sample-level variables onto distributional properties. This is especially valuable in multifactorial experiments, where multiple biological or experimental characteristics can jointly influence measurements. Such generalization also allows modeling interactions or nesting between variables. For example, treatment and control groups may have divergent temporal trajectories, or treatments might have different effects across cohorts. For these applications, parameters can be linked to sample covariates using regression techniques. For example, in the scDesign3 simulator for single-cell and spatial omics data [97], mean and dispersion parameters are modeled as functions of sample characteristics using generalized additive models of location, scale, and shape (GAMLSS) [86, 98]. This flexibility allows simulated data to vary smoothly over temporal or spatial coordinates. Similarly, in the ZINB-WaVE simulator for single-cell RNA-seq data [87], sample-level covariates can influence the mean, dispersion, and zero-inflation parameters through a linear model. Beyond feature-level properties, sample covariates can modify multivariate relationships. For example, scDesign3 allows different copula covariances to be used within prespecified groups of samples.

### 3 Evaluation

Evaluating synthetic data is a crucial step in data simulation to assess whether the generated synthetic data closely mirror the statistical properties of the original data. Without proper evaluation, synthetic data may not accurately represent the underlying distribution, potentially leading to biased models or incorrect conclusions.

In evaluating synthetic data, three primary types of utility measures can be used: *fit-for-purpose measures*, *global (broad) utility measures*, and *outcome-specific (narrow) utility measures* [15] (see Figure 2). Fit-for-purpose measures provide an initial evaluation of whether the synthetic data appear reasonably close to the real data, and they can be used to improve the simulation approach [15]. In contrast, global utility measures aim to assess the multivariate characteristics of the data. Outcome-specific utility measures aim to quantify similarity in analysis results or specific model parameters between real and synthetic data.

Fit-for-purpose measures typically involve checking the univariate and bivariate distributions of observed and synthetic data. We divided these measures into two main types: univariate and bivariate evaluations (Figure 2A1-A2). In the former, we focus on whether the marginal distributions of variables in the real and synthetic data match. In the latter, we focus on whether the pairwise relationships of variables in the synthetic data resemble those in the real data, typically through bivariate distributions and correlations. We further divide evaluation criteria into graphical and quantitative measures. For example, graphical measures include side-by-side box plots or empirical cumulative density plots (univariate setting) and pairwise scatter plots or correlation heat maps (bivariate setting). Quantitative measures include the Kolmogorov–Smirnov (KS) test, the Wilcoxon Rank-Sum test, or independence tests based on kernel density estimation [17] (univariate setting) and the distance measures between correlation matrices (bivariate setting). However, while fit-for-purpose measures offer an initial assessment of synthetic data quality, they do not account for the multivariate nature of the data.

Global utility measures are built upon fit-for-purpose measures to evaluate the complex, multivariate nature of the data. We again divide these into graphical and quantitative measures (Figure 2B). For example, graphical measures include Principal Component Analysis (PCA) plots, which jointly project synthetic and real data into a two-dimensional space, or receiver operating characteristic (ROC) curves, which use binary classification to separate synthetic from real data, to determine if the synthetic and real data can be distinguished. Quantitative measures include the Friedman-Rafsky test [24] or propensity score-based [21] techniques. Global utility measures provide an aggregated similarity between simulated and real data. However, they do not guarantee that specific analyses on real and simulated data will yield similar results,

Table 2: An overview of packages that are available for microbiome community simulation. Models differ according to mechanisms they use to estimate community structure and the scope of applicability to alternative experimental designs. For applications to new environments or designs, the evaluation techniques described in the section ‘Evaluation’ can be used to compare candidates and gauge simulator fidelity.

Package Name	Probability Mechanism	Experimental Designs	Notes
CAMISIM [26]	Models taxon abundances with a log-normal distribution and does not account for featurewise correlation.	Supports separate modes for case-control and time series designs, including simulation of technical replicates.	Used in the CAMI benchmarking challenge [92, 72] and can output read-level data (Illumina, PacBio, or ONT).
DeepMicroGen [10]	Combines recurrent network and GAN architectures to model multivariate changes over time.	Supports longitudinal designs with regular spacing.	Parameters are difficult to interpret, limiting potential for control.
MB-GAN [88]	Trains a GAN to match a template dataset, capturing both marginal and multivariate structures.	Mimics the design of the template data, but cannot be adapted to alternative designs.	Primarily for Centered Log Ratio (CLR) transformed data but applicable to small sample sizes.
metaSparSim [79]	Uses a hierarchical model to sample biological composition (Gamma distribution) and technical variability (multivariate hypergeometric distribution).	Does not include covariates during parameter estimation but allows changes post-estimation to reflect experimental perturbations.	Preserves sparsity and mean-variance profiles for individual taxa, with cross-taxa correlation quality depending on Gamma parameters.
miaSim [27]	Simulates population dynamics based on various interaction models, including the generalized Lotka-Volterra model.	Supports time series designs with potential environmental interventions.	Accompanied by an interactive Shiny application but cannot be fitted to a template dataset.
microbiomeDASim [106]	Draws community profiles from a zero-inflated, truncated multivariate normal distribution.	Designed for longitudinal experiments with user-specified trends (e.g., quadratic and “hockey-stick”).	Only applicable to normalized abundance tables.
MIDASim [38]	Estimates marginals from empirical or generalized gamma models, correlating presence-absence patterns and abundances via a copula.	Supports case-control designs with varying sample sizes, sequencing depths, and taxa presence and abundance.	Scalable for both estimation and simulation, with support for template inputs and post-estimation modifications.
scDesign3 [97]	Applies GAMLSS to each feature, linking them with a Gaussian or Vine copula.	Can include covariates during estimation, enabling simulation of arbitrary designs.	Originally developed for single-cell applications, with similar models used for microbiome data [85, 110, 43].
SparseDOSSA 2 [68]	Simulates features from a zero-inflated log-normal distribution, linking them with a Gaussian copula with a regularized precision matrix.	Does not incorporate covariates during estimation, but allows post-estimation changes that use them.	Previously validated on human gut and vaginal microbiome data.
ZINB-WaVE [87]	Simulates data from a zero-inflated negative binomial model, incorporating known covariates and low-rank latent variation.	Can include covariates during estimation, enabling simulation of arbitrary designs.	Originally developed for single-cell applications, with similar models used for microbiome data [85, 110, 43].

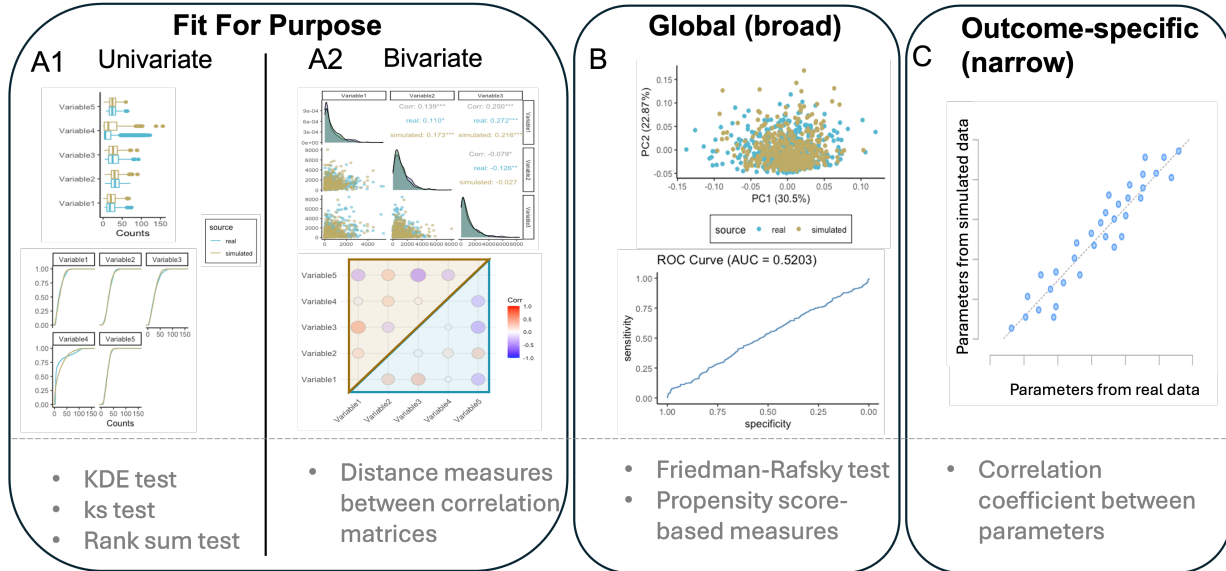


Figure 2: Utility measures for synthetic data, including visual comparison (top) and quantitative measures (bottom). (A) Fit-for-purpose, (B) global (broad), and (C) outcome-specific (narrow) utility measures. Fit-for-purpose and global utility measures both evaluate the similarity between synthetic and real data without considering any analytical objective. In contrast, outcome-specific utility assess the similarity of analysis results or specific model parameters using both real and synthetic data.

as these measures do not consider a specific analysis goal [15, 49, 20].

Outcome-specific utility measures are designed to assess the simulated data for a particular analysis goal. Since there are multiple analytical approaches, these utility measures can vary significantly. For example, if the focus is on fitting a multiple regression model between sequencing features and a biological condition of interest, then we can compare regression coefficients obtained when fitting the regression to either simulated or real data. In this situation, a scatter plot can serve as a graphical measure, and the correlation value between the parameters estimated from real and simulated data can serve as a quantitative measure (Figure 2C). However, if the objective is to perform a correlation analysis, like those which underlie microbiome network models, then the evaluation should focus on comparing the correlation matrices between the features in the real and the simulated data (for an example of this comparison, see Section ‘Power analysis for multivariate methods’). No single utility measure is universally applicable. Therefore, performing several utility measures based on the specific objective of the simulation will help modify certain aspects of the simulator, such as the selection of distributions.

## 4 Case Studies

### 4.1 Statistical power for univariate differential abundance methods

**Motivation.** Differential abundance analysis is a cornerstone of microbiome research. It is often used to highlight taxa that are associated with disease or that respond to interventions. The research community has developed various testing methods to account for specific characteristics of microbiome data, such as zero-inflation, overdispersion, library size differences, compositionality, and small sample sizes [32, 52, 38]. However, assumptions that are reasonable in one microbiome system (e.g., the gut) may not necessarily translate to others (e.g., marine). Applying a method in an inappropriate context can lead to excess false positives or reduced power. Moreover, unlike classical two-sample tests or linear models, differential

Table 3: Summary of case study datasets and analysis goals.

Study name	Number of samples	Number of features	Analysis task	Objective
ATLAS [56]	883	89	Differential abundance testing	To select the appropriate differential abundance method for a given experimental design and data type.
Type 1 Diabetes [28]	101	427	Power analysis and multivariate method	To compare complex experimental designs and multivariate analysis approaches.
American gut microbiome [70]	261	45	Benchmarking network inference	To evaluate different network inference approaches under diverse network structures.
Anaerobic digestion [6]	75	231	Batch effect correction	To give insights into the behavior of batch integration methods.
Sepsis in ICU patients [36]	57	180 bacteria, 18 fungi, 42 viruses	Omics data integration	To uncover whether different data modalities (bacteria, fungal, virus) should be analyzed together or separately.

abundance methods do not come with closed-form formulas for calculating statistical power. Therefore, semisynthetic simulation can give insights into the properties of differential abundance methods for specific data analysis applications.

We compared three common differential abundance methods commonly used in microbiome studies: DESeq2 [65], limma-voom [57] and ANCOM-BC2 [63]. ANCOM-BC2 was developed specifically for microbiome data, whereas DESeq2 and limma-voom were originally designed for RNA-seq data. The goal of this simulation is to select the most appropriate method for a specific dataset based on statistical power across sample sizes.

**Data.** We analyzed data from the ATLAS study [56]. This study profiled the gut microbiomes of  $n = 1006$  healthy adult Europeans. Notably, it discovered microbiome “tipping points,” which are unstable community composition profiles that could “tip” into more stable states. The competing stable states were related to factors like age and body mass index (BMI). We filtered down to the 89 most abundant genera in the 883 subjects with known region membership and BMI within the lean, overweight, or obese categories.

**Simulation and evaluation.** Our simulator follows scDesign3’s modeling assumptions, applying a zero-inflated negative binomial GAMLSS regression where both the mean and dispersion for each taxon can vary as a function of the BMI category. For the top ten most abundant taxa, we found that the boxplot quartiles and the observed shifts across BMI groups were captured well in the simulated data (Figure 3A1). We compared the abundance of all genera in simulated and real data for each BMI group using a kernel density-based two-sample test [16, 18]. At least 70% of genera were not statistically distinguishable (Figure 3A2). As evident by these evaluations, the simulated data were generally similar to the original data. However, for some genera with positively skewed distributions and a high number of outliers, such as *Prevotella melaninogenica et rel*, the simulation failed to capture the high degree of skewness or generate outliers, hence the significant differences between the real and simulated data returned by some of the tests.

**Data analysis results.** We simulated synthetic negative controls by re-estimating a subset of genera so that the zero-inflated negative binomial means and dispersions did not depend on BMI. These genera were declared not significant from a Wilcoxon Rank-Sum test of association with BMI ( $p$ -values corrected for multiple testing with 0.1 FDR level). Such weakly associated genera have previously been used to define pseudo-negative controls [76], whereas in our simulator, these genera are genuine negative controls. These simulation design considerations, as well as those for all other case studies, are summarized in Table 4. We

Table 4: Summary of simulation design considerations for the case studies.

Analysis task	Kept from the template	Simulator modifications	Justification
Differential abundance testing	BMI effects for the ground truth differentially abundant taxa.	Varies the proportion of ground truth differentially abundant taxa and overall sample size.	BMI effects must be removed to create synthetic negative control taxa.
Power analysis and multivariate method	T1D effects for the ground truth differentially abundant taxa; community covariance.	Varies the proportion of ground truth DA taxa and the overall sample size.	Prediction performance is known to depend on the sample size and the proportion of relevant features.
Benchmarking network inference	BMI and sequencing depth effects for all taxa.	Alters the ground truth covariance structures are substituted into the Gaussian copula.	Covariance is a common target for network inference methods, but no single edge pattern is universally applicable.
Batch effect correction	Batch and treatment effects for all taxa; community covariance.	Introduces a treatment group with a weak effect.	Overintegration is especially problematic in settings with weak signals.
Omics data integration	Sepsis and antibiotics effects for taxa in the ITS and Virome assays; community covariance.	In one setting, sepsis and antibiotics effects are removed from the 16S assay.	Dimensionality reduction plots should not introduce spurious associations.

then simulated datasets with sample sizes ranging from 50 to 1200, and applied DESeq2, limma-voom and ANCOM-BC2 using an FDR-control level of 0.1. On average, DESeq2 had better power compared to the other two approaches, and all three approaches provided valid FDR control (Figure 3B). The power increased most rapidly from 50 to 337 samples.


**Summary.** Differential abundance testing requires a delicate balance. On the one hand, parametric assumptions, like those for overdispersion or zero inflation, can improve test sensitivity. On the other hand, inappropriate assumptions can lead to invalid results. Generic benchmarking studies can help identify the appropriateness of each method for a given experimental design and data type but can only provide blanket recommendations. In contrast, we showed that simulations let us design “self-service” benchmarks where we can run our own *in silico* experiments to inform a statistical analysis workflow with more problem-specific choices. An additional benefit is that we can ask how certain changes to the data (e.g., increasing the sample size) might affect method performance.

## 4.2 Power analysis for multivariate methods

**Motivation.** Differential analysis methods can draw attention to significant taxa but can overlook important correlation structure within microbiome communities. To shed light on this more global structure, multivariate analysis methods play an essential role [59]. Theoretically characterizing the statistical efficiency of multivariate methods is an active area of research [1, 48], and practical sample size guidance typically relies on simulation [95, 35]. However, effective community-wide simulation is more challenging than what is required for differential abundance analysis, because we must pay close attention to the quality of the estimated correlations. In this case study, we explore how semisynthetic data can inform sample size calculations in an application that uses sparse partial least squares discriminant analysis (SPLS-DA) [58]. SPLS-DA is a classification method that makes use of correlations between input features to ensure stable predictions in



A1



A2

BMI Group	Lean	Overweight	Obese
Proportion of insignificant tests	0.72	0.80	0.71

B




Figure 3: Differential abundance analysis in the ATLAS study. (A1) Comparison between real and synthetic data for the ten most abundant taxa. A zero-inflated negative binomial model achieves satisfactory performance in matching quartiles of the observed data. (A2) Quantitative evaluation between real and synthetic data. Proportion of non-rejected tests from the kernel density-based global two-sample comparison test for each feature across BMI groups in real and simulated data. Insignificant test results indicate that the kernel density estimates for the given feature in each BMI group cannot distinguish between the real and simulated data. (B) Realized power and false discovery rates for the DESeq2, limma-voom and ANCOM-BC2 methods applied to the simulated data. Large samples lead to higher and less variable experimental power. FDR control is maintained on average.

small sample size settings. In the process, it computes a dimensionality reduction of the data analogous to PCA, but with the explicit goal of separating classes.

**Data.** We revisited the Type 1 Diabetes (T1D) study from Gavin et al. [28], who identified metaproteomic patterns in the microbiomes of T1D patients from a cohort of  $n = 101$  study participants. We filtered down to the 427 proteins present in at least 70% of either the T1D ( $n = 51$ ) or control ( $n = 50$ ) groups. Following the original study’s data preprocessing, we applied a centered log-ratio (CLR) transformation to the measured protein relative abundances and then used these as predictors in an SPLS-DA with T1D status as the outcome. We set the SPLS-DA hyperparameters to 5 PLS dimensions and a selection of 30 predictors. On this dataset, ten repetitions of 5-fold cross validation yields a holdout area under the receiver operating characteristic curve (AUROC) of  $0.667 \pm 0.037$ .

**Simulation and evaluation.** We fitted a Gaussian GAMLSS simulator, allowing means and variances for all proteins to depend on T1D status. Since SPLS-DA models the relationships across all proteins, we first applied a Gaussian copula using the standard sample covariance to attempt to reflect the true correlation structure. Contrary to the last section, we focus here on the quality of the simulated sample correlations rather than the univariate simulation quality (which is covered in our online tutorial section 3.3).

A simple histogram of pairwise correlations between features in simulated data showed greater variability compared to the observed data, suggesting that the simulation could be improved using a regularized covariance estimator. We therefore used the adaptive thresholding covariance matrix estimator of Cai and Liu [5] (higher thresholds apply stronger regularization, while lower thresholds reduce bias). To choose the threshold, we evaluated the correlation quality across a range of candidate values from 0.001 to 0.2 using two metrics: the KS statistic between observed and simulated histograms of pairwise correlations and the Frobenius norm between observed and simulated correlation matrices. The KS statistic formalizes our histogram check, while the Frobenius norm measures the differences between these matrices. We found that the KS statistic was minimized at 0.14, while the Frobenius error was minimized at 0.03. To balance these two metrics, we chose a threshold of 0.1. To further check the simulator, we created pairwise scatter plots (Figure 4A) and heatmaps of real and simulated sample covariance matrices (Figure 4B). The heatmaps show similar blocks of positive or negative correlation. For the pairwise scatter plots, we filtered to a subset of pairs with moderate positive correlations between 0.72 to 0.8 in the real data, as checking all 90K pairs is impossible. The scatter plots of real and simulated data overlapped well. The only notable differences were the streaks of exact zeros in the real data, a reflection of the zeros present before CLR transformation. Overall, the revised simulator with an adaptive covariance estimator seemed sufficient for a power analysis with SPLS-DA.

We next altered the fraction denoted  $\pi_1$  of nonnull proteins whose means and standard deviations depend on T1D status. We chose the nonnull proteins according to their  $p$ -values from a Wilcoxon Rank-Sum test on the original experimental data. We considered the most significant proteins as true signals in the simulator. We then applied SPLS-DA with the pre-specified hyperparameters to semisynthetic data of varying sample sizes, with  $\pi_1 \in \{\frac{10}{427}, \frac{30}{427}, \frac{50}{427}, \dots, \frac{110}{427}\}$ , giving coverage of settings where differentially abundant taxa range from rare ( $< 2.5\%$  of taxa) to relatively common ( $> 25\%$  of taxa).

**Data analysis results.** When considering 20 samples total, we found that the proportion  $\pi_1$  of true nonnull proteins has little effect on average AUROC, which is only slightly better than random (Figure 4C). Larger  $\pi_1$  was associated with less variable performance. When  $\pi_1 = \frac{110}{427}$ , performance increased quickly with sample size, plateauing at 110 samples. In contrast, for  $\pi_1 = \frac{10}{427}$ , performance increased more gradually, with room for improvement even at 200 samples. Interestingly, performance was relatively stable from  $\pi_1 = \frac{44}{427}$  to  $\frac{110}{427}$ . Altogether, this analysis suggested that if more than a few dozen proteins are thought to be associated with the outcome, then between 110 and 155 samples is sufficient. For fewer proteins, either a larger sample size or an alternative analysis strategy should be considered. A sense of the true proportion  $\pi_1$  of nonnull proteins can be gauged from the distribution of  $p$ -values in the template data, using estimators like those introduced in [99, 3, 4]. For these data, the true signal appeared to be weak. For example, Storey’s estimator returned  $\hat{\pi}_1 < 0$ .

**Summary.** We showed how semisynthetic data can help determine sample size for SPLS-DA applied to microbial proteomics. While no formulas for statistical power exist for this method, simulation allowed

us to study how both experimental (sample size) and biological (proportion of nonnull proteins) factors influence its performance. Further, we detailed the process of evaluating the simulator’s correlation structure, demonstrating how visualizations and metrics could be used to iteratively refine a model.

### 4.3 Benchmarking network inference

**Motivation.** Network models give a holistic view of interactions in microbial ecosystems. They can identify tightly connected subcommunities and keystone taxa [41, 22, 77]. Unfortunately, validating these models is difficult, since determining ground-truth edges typically requires low-throughput experimentation such as knockout studies [9]. This creates a barrier to benchmarking, both for their use in specific studies and for evaluating new methodologies. Simulation can address these issues by providing ground-truth edges.

We compared two methods: SpiecEasi [55], a graphical lasso-based algorithm tailored to compositional data, and the Ledoit-Wolf estimator [60], a covariance estimator created for high-dimensional but low-rank data. A priori, we may expect SpiecEasi to perform better on microbiome data, since it was specifically designed for this purpose. However, this comparison has not been previously reported, and it is also unclear if potentially improved accuracy justifies the additional time required to solve the SpiecEasi optimization problem. To help answer this question, we can use simulations with known covariance structures.

**Data.** We analyzed the data from the American Gut Microbiome (AGM) project, a citizen science initiative where participants submit stool samples and complete detailed diet and health surveys [70]. The study has revealed associations between survey responses and microbiome composition. We considered a subset of 261 samples available through the SpiecEasi package [54], after excluding samples with fewer than 1000 reads. We filtered down to a “core” gut microbiome [93] of the 45 taxa with an abundance of over 100 in at least 50 samples.

**Simulation and evaluation.** We fitted a zero-inflated negative binomial GAMLSS model to these data using BMI and log sequencing depth as covariates. The BMI term was included because the original AGM study [70] found a significant association between BMI and microbiome composition. The sequencing depth term allows samples with deeper sequencing to have larger means on average across all taxa. The evaluations comparing real and simulated data are available in section 3.2 of the accompanying online tutorial. Here, we focus on benchmarking estimator accuracy across several network structures beyond those observed in AGM. We define several ground-truth correlation structures representing different statistical regimes: block, banded, Duncan-Watts small-world, scale-free, and Erdős-Rényi random graph structures [80]. The block and banded covariances were defined directly, while others were derived from Gaussian graphical models with the respective structure, shown in Figure 5A. For example, the scale-free network had several hub taxa with high covariance across many neighbors, while the Erdős-Rényi network connected all pairs of taxa with equal probability. Note that while our focus here is on covariance matrix estimation, a similar study using known, structured precision matrices could also be implemented.

**Data analysis results.** Across network regimes, the Ledoit-Wolf estimator had larger bias but lower variance compared to SpiecEasi (Figure 5B). Since both estimators are regularized, they exhibited some bias towards correlations with smaller magnitudes. The Ledoit-Wolf estimator performed better in the banded and block settings, while the two approaches were comparable in other cases. All methods were challenged in the Erdős-Rényi setting, often declaring large correlations for taxa with no true relationship. Interestingly, in the block covariance scenario, SpiecEasi estimated many exact zero correlations as slightly negative. This is a consequence of SpiecEasi’s compositional assumption, which induces negative correlation across taxa. In spite of this artifact, the block covariance setting appeared to support more efficient estimation than any of the other covariance settings we considered. Thus, if we assume that a community has a block correlation structure, then fewer samples may be required. This is consistent with general statistical theory, which argues that low-dimensional block structure can greatly simplify high-dimensional covariance estimation [84].

**Summary.** We showed that simulations offer a useful lens for comparing network inference methods across diverse network structures while maintaining realistic abundance distributions. Starting from a single template dataset, we were able to simulate according to several types of ground truth correlation, enabling us to identify settings where methods are more likely to fail or succeed.

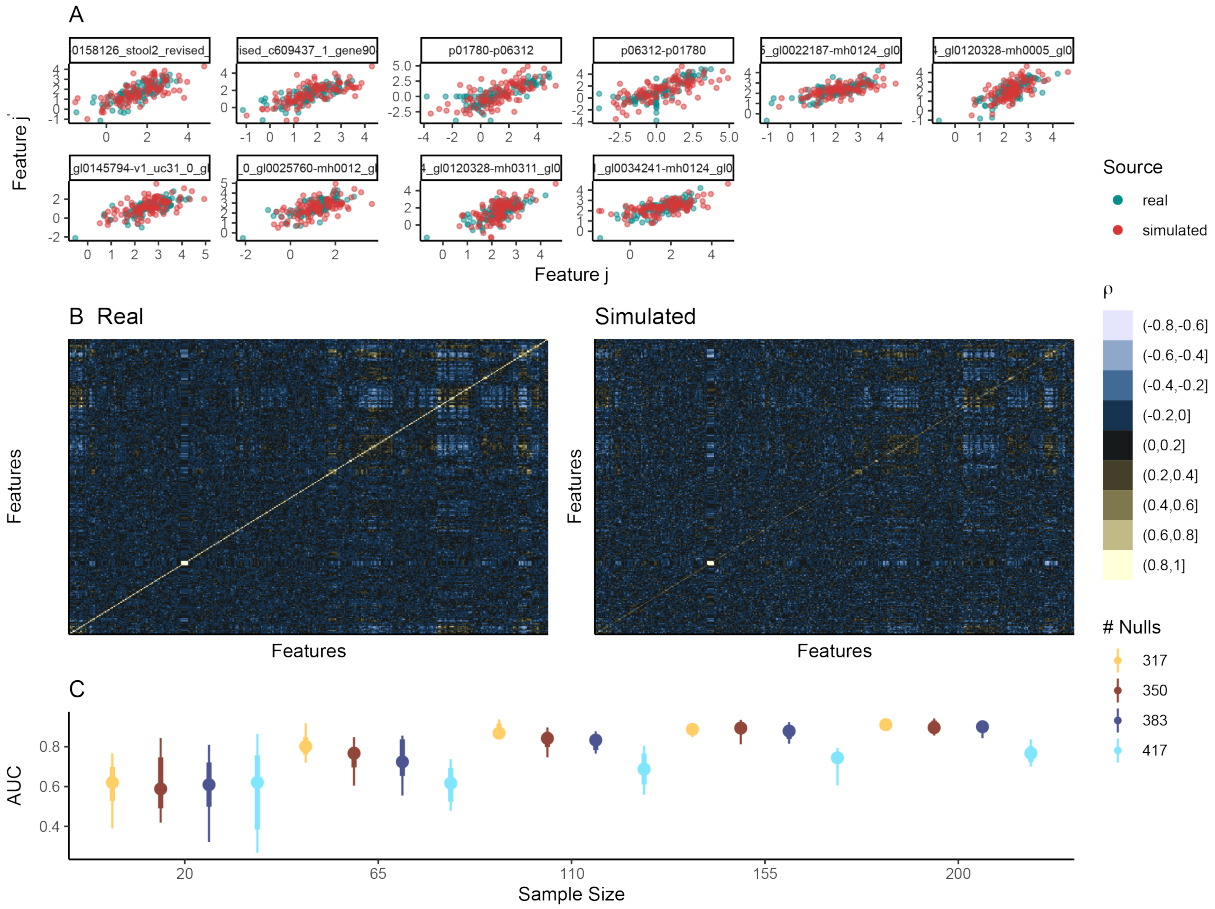


Figure 4: SPLS-DA power analysis for Type 1 Diabetes association. (A) Pairwise scatter plots for proteins with real-data correlation ranging from  $[0.72, 0.8]$ . Except for streaks at 0, a Gaussian copula appears to preserve associations between pairs of proteins. (B) A heatmap of the correlation matrices across all proteins in the real and simulated data. Blocks of positively and negatively correlated proteins appear to be preserved in the simulation. (C) Prediction accuracy of SPLS-DA across simulation settings. The  $x$  and  $y$  axes correspond to sample size and AUROC, respectively. Color corresponds to the number of null features,  $427 \times (1 - \pi_1)$ , which governs the rate of power increase as a function of sample size.

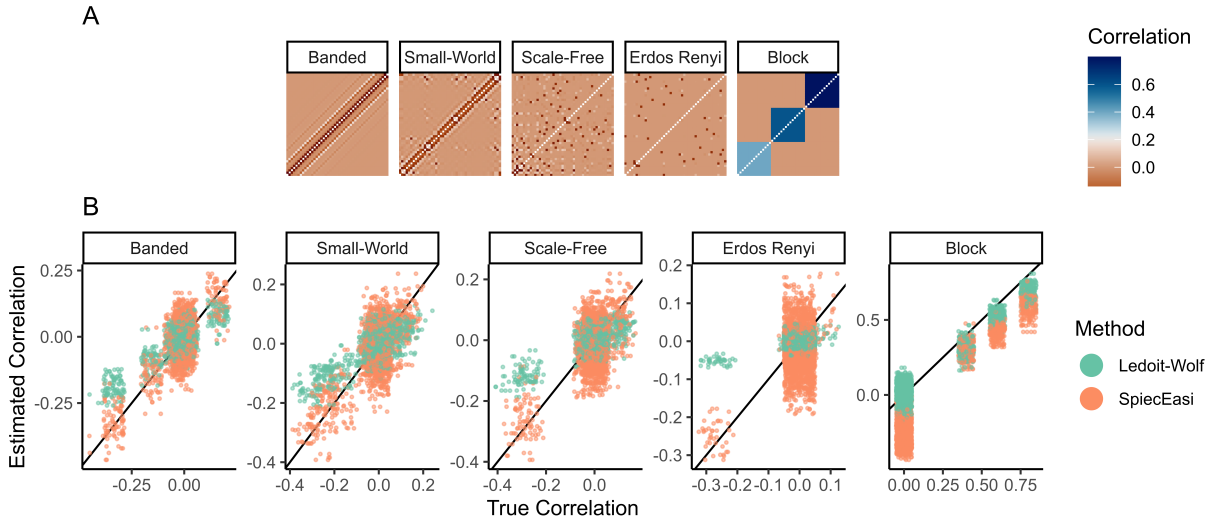


Figure 5: Benchmarking network methods for American Gut Microbiome project. (A) Ground truth correlation matrices used in the simulation study and (B) estimates made using SpiecEasi and the Ledoit-Wolf estimator. Entries  $(j, j')$  in the matrices of (A) provide the copula correlation between taxa  $j$  and  $j'$  before transformation using the zero-inflated negative binomial model. From (B), SpiecEasi tends to have lower bias, since it is generally centered along the ground truth black line, but higher variance, since its vertical spread is larger. This figure also highlights settings where estimation tends to be easier (block and banded) and more difficult (Erdős-Rényi).

#### 4.4 Batch effect correction

**Motivation.** In large microbiome studies, it is often difficult to guarantee uniform data collection and processing for all samples, leading to systematic differences across experimental groups, often referred to as batch effects [103, 66]. For example, storing samples at different temperatures might change the abundances of taxa whose marker gene sequences degrade more rapidly at some temperatures. Failing to address these batch effects can limit power by masking treatment effects. They can also compromise validity by introducing spurious differences across treatment groups [61]. Hence, a common preprocessing step in microbiome data analysis is to apply batch effect correction to standardize data across batches [30, 64].

Despite their increasingly widespread adoption in microbiome studies, batch effect correction methods must be applied carefully. Failing to account for the correction in downstream differential tests can lead to miscalibrated  $p$ -values [75] or performance estimates [96]. Further, it can be difficult to balance underintegration, where batch effects persist post-correction, against overintegration, where aggressive batch effect correction eliminates meaningful biological variation [113]. It is often unclear in advance whether these issues will arise for a given dataset or batch effect correction method. By defining ground-truth batch and biological effects, simulations give a way to evaluate batch effect correction methods and their impact on downstream inferences. Moreover, they allow quick comparison of methods across experimental designs and biological scenarios, allowing more complete evaluation than is possible in isolated benchmark datasets.

**Data.** We analyzed a study of anaerobic digestion (AD). AD is a biodegradation process underlying many bioenergy production technologies. The study Chapleur et al. [6] profiled the microbiomes from AD samples treated with phenol, a micropollutant that affects biodegradation efficiency and stability, posing challenges for large-scale industrial deployment. We use a subset of 75 samples and 231 genera discussed in Wang and Lê Cao [103], which focused on changes in community composition under two phenol concentrations. Since obtaining samples is time-consuming, the experiment was carried out over five sessions, leading to batch effects visible in the PCA plot in Figure 6A.

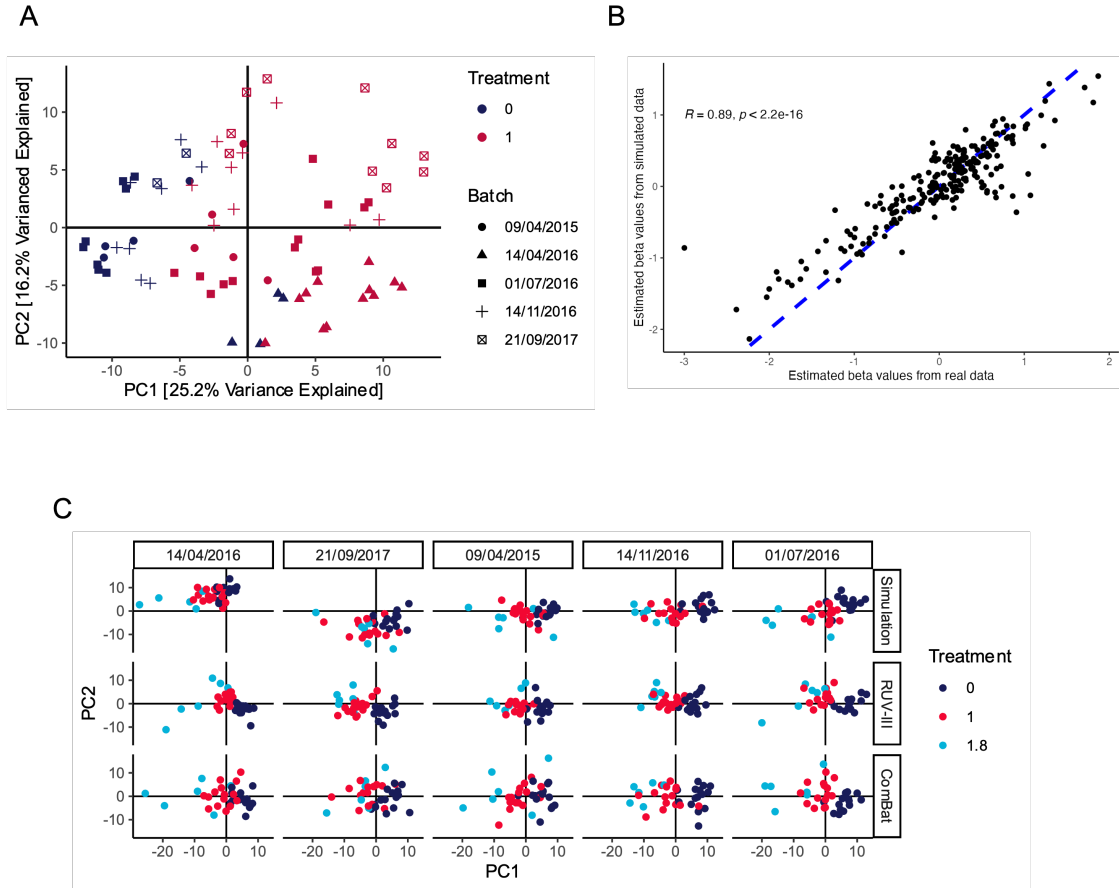


Figure 6: Assessing batch effect correction methods in an anaerobic digestion study. (A) A principal components plot of the original AD dataset reveals significant batch-to-batch variation. For example, all experimental samples generated on 14/04/2016 are shifted towards the bottom-right quadrant. Note that the technical replication structure is also visible, for example, the clusters of untreated samples from 01/07/2016 on the left-hand side. Though the two phenol concentration treatments can still be distinguished from one another, removing the batch effect can improve power. (B) Narrow evaluation using a scatterplot and correlation coefficient to compare the RUV-III batch effect regression coefficients for each taxon in real and simulated data. (C) Original and batch corrected data in the simulation with a new, less frequently sampled treatment group ( $t = 1.8$ ) introduced. Batches (columns) have been sorted from the lowest to the largest average PC1. Both RUV-III and ComBat successfully remove systematic differences between batches, but in some cases ComBat removes true differences between the  $t = 1$  and  $t = 1.8$  groups. Note that the principal components are derived separately for each row.



**Simulation and evaluation.** The purpose of this simulation experiment is to compare the performance of RUV-III [73] and ComBat [47], two batch-effect correction methods, when applied to data like those in the AD study. We used the CLR-transformed AD dataset as our simulation template. For each taxon, we applied a Gaussian GAMLSS and copula model with batch and treatment status as covariates. Given the relatively small sample size, we used an adaptive thresholding covariance matrix estimator [5] within the Gaussian copula. PCA on data simulated from this model revealed that the simulator could recapitulate the observed batch effects. We can evaluate this more precisely using a narrow utility evaluation, comparing the RUV-III batch effect regression coefficients from the real and simulated data. For each taxon, a regression coefficient was calculated using AD as the factor of interest. The regression coefficients from the simulated data closely matched those from the real data, with an overall correlation of 0.89 (Figure 6B).

Batch effect correction methods are known to be sensitive to imbalance between true biological groups [66]. We next explored whether this could pose a challenge in AD studies if we considered an additional phenol concentration treatment level. We simulated a hypothetical scenario where a stronger treatment had been applied to a small subset of samples. Each batch was assigned 15 samples at the reference concentration ( $t = 0$ ) and the previously observed comparison group ( $t = 1$ ), but only six at a new, hypothetical treatment (set to  $t = 1.8$ ). Relative to the reference, this new treatment level is anticipated to perturb the microbiome community in the same direction as the original  $t = 1$  group, but to a greater extent. Since this treatment group is less frequently sampled, there is a risk that its effect might be masked by overly aggressive batch integration.

**Data analysis results.** We compared the results of RUV-III and ComBat applied to data simulated from our altered experimental design (Figure 6C). The simulator generated data with clear batch and treatment effects (first row). Although we observed a lack of replication structure for some of the batches in Figure 6A, the simulated data preserved sufficient batch effects to evaluate batch effect correction methods. We then compared the PCA projections when RUV-III and ComBat were applied to the simulated data. Both methods centered all batches around the origin, as expected from a batch effect correction method. However, the  $t = 1.8$  group was less clearly separated from the  $t = 1$  group in the ComBat output compared to either the original simulation or the RUV-III corrected data. For a more quantitative analysis, we trained a linear discriminant analysis to predict the treatment group from the first two principal components, yielding a classification accuracy of 73.9% on the original simulation, 78.3% on the ComBat correction, and 88.3% on the RUV-III correction. These classification performance results suggest that ComBat may overintegrate imbalanced data, while RUV-III preserves more subtle treatment differences.

**Summary.** We illustrated how simulation can give insight into the behavior of candidate batch integration methods in a way that is tailored to specific data analysis contexts, with the opportunity to create new treatment scenarios not present in existing real-data benchmarks.

## 4.5 Omics data integration

**Motivation.** A single assay can only give a partial view of a microbial ecosystem. For example, 16S rRNA sequencing data characterizes bacterial community composition within a sample, but other kingdoms, metabolites, and host cells shape microbiome properties. To capture these features, different sequencing methods conducted on the same samples (e.g., ITS for fungi) or profiling techniques (e.g., NMR for metabolites) are needed. By analyzing the complementary views offered from diverse datasets, it is possible to develop a more holistic understanding of an ecosystem’s biology [90, 102, 8].

Effectively analyzing these data, however, is challenging, with no consensus on how integration strategies should be deployed across contexts. The statistical task is not simply multivariate but also multiassay, linking tables influenced by diverse biological or experimental factors. It is necessary to decide on which datasets to integrate, how each table should be normalized, and which integration method might be appropriate. Simulation can give a controlled, simplified setting within which to compare analysis strategies and can help to predict method performance under realistic biological scenarios.

Integration methods often return a dimensionality reduction plot designed to uncover shared covariation across assays, analogous to how PCA arranges samples to describe variation within a single assay. It is critical that these methods faithfully preserve between-assay and between-sample relationships. Ideally,

data integration dimensionality reduction will highlight genuine, shared axes of variation across data sources while avoiding the appearance of false associations between unrelated data sources. We design a simulation to assess the extent to which these plots can misleadingly suggest similarities in inherently “unalignable” assays [67]. This multi-assay analysis parallels the multi-batch problem discussed in the previous case study – instead of studying overintegration across batches, we consider overintegration across assays.

**Data.** We re-analyzed the data from Haak et al. [36], who studied how the gut microbiome is altered during sepsis in ICU patients. Sepsis can be triggered by microbial infection unrelated to bacteria (e.g., from the *Candida* fungus). Therefore, each sample was profiled using ITS amplicon and Virome sequencing in addition to 16S rRNA sequencing. Since sepsis is treated with antibiotics, the study included healthy patients undergoing antibiotics treatment. The study included 20 healthy controls, 23 sepsis patients, 5 healthy patients on antibiotics, and 9 ICU patients without sepsis. After applying the same filtering as [36], we obtained measurements for 180 bacterial genera, 18 fungal genera, and 42 viruses.

**Simulation and evaluation.** We considered the multiblock SPLS-DA method to integrate these class-labeled sequencing datasets [94], and simulated from the scenario where the class label was relevant only for a subset of tables. This could occur if any of the three kingdoms assayed were unrelated to sepsis or antibiotics. An ideal multiassay integration in this setting would recover the similarities in the abundance profiles across kingdoms (e.g., revealing shared clusters of samples) while avoiding introduction of false relationships with sepsis or antibiotic status. Note that multiblock SPLS-DA is the multi-assay extension of SPLS-DA (discussed in the section ‘Power analysis for multivariate methods’) and replaces the SPLS-DA objective with a weighted average of covariances across pairs of tables. We applied a Gaussian GAMLSS simulator to each assay, as they were already normalized. For the ITS and Virome assays, we allowed means and variances for each feature to depend on a categorical feature encoding both ICU and antibiotic use. For the 16S data, we used two simulators: one which conditioned taxonomic abundances on ICU/antibiotics category and one which did not. We joined all tables using a Gaussian copula with an adaptive covariance estimator to handle high-dimensionality. Therefore, in the ground truth simulation, the marginal (taxon-level) and covariance (community-level) structure from the template data are maintained. However, while the sepsis and antibiotic factors continue to influence all the ITS and Virome measurements, this relationship has been deliberately removed from one of the two simulated 16S datasets.

This design allows us to study the extent to which integration can spuriously introduce sepsis or antibiotic associations into the 16S data. Specifically, we then compared multiblock SPLS-DA output on the original and both versions of the simulated data (Figure 7).

**Data analysis results.** We find that, even after removing all association between 16S abundances and the class label, the multiblock SPLS-DA projections for the 16S block still separated by class, albeit more weakly than before. This is a consequence of the analysis’ multiblock nature: strong class associations from other tables can be artificially introduced into the 16S table. Indeed, part of the objective of the multiblock SPLS-DA algorithm is to maximize similarity in projections across all tables. In this case, in the real data, the class differences in the 16S data were much larger than those seen in this null scenario. Thus, simulation can guarantee reliable conclusions for complex integration tasks.

Further, to decide on which tables should be analyzed together, or separately, we calculated alignment measures [67] on the real data and compared them with a synthetic null where the tables are not alignable. We study the distribution of canonical correlations in both the observed and synthetic negative control data when integrating the Virome and 16S data together (Figure 8). The control was defined so that the two tables are known to have no correlation within the underlying copula model. The results showed the presence of slight, but systematic, shared variation: observed canonical correlations were consistently larger than those in the synthetic null for the top dimensions, but the increase was quite modest. This suggests that substantial variation is isolated within the tables in a way that an integrated analysis cannot capture, motivating the use of within-table analyses alongside integrative methods.

**Summary.** Analogous to using negative controls in biological experimentation, we created synthetic negative controls in a computational workflow for data integration. This allowed us to gauge evidence for potential discoveries in real data, and prioritize either simultaneous or table-specific analysis based on departures from the synthetic negative control setting.




Figure 7: Using simulation to understand properties of omics data integration methods in a transkingdom analysis. Columns correspond to three amplicon sequencing datasets for transkingdom analysis of sepsis from Haak et al. [36]. Multiblock SPLS-DA projections from real and simulated data are given in the first two rows. The bottom rows show a perturbed simulation where all associations in the 16S data have been removed. The fact that between-group differences are still visible reflects the cross-table reduction and serves as a null reference for reliability analysis.




Figure 8: Comparison between canonical correlations in real (red) and reference null (beige) samples in a transkingdom analysis. For the first twenty dimensions, the real canonical correlations are slightly larger than those in the synthetic null. These dimensions are likely to reflect true, if modest, shared variation across tables.

## 5 Discussion

We have reviewed the potential to apply simulators to microbiome study design and analysis. We have considered advances that make simulators more broadly applicable than their predecessors and have provided reproducible case studies showing how they can offer ground truth for evaluating FDR and power (sections ‘Benchmarking differential abundance methods and statistical power for univariate methods’ and ‘Power analysis for multivariate methods’), help identify regimes where competing methods differ (sections ‘Benchmarking network inference’ and ‘Batch effect correction’), and ensure valid interpretation of statistical outputs (section ‘Omics data integration’).

We have emphasized methods for evaluating simulation quality, introducing vocabulary for distinguishing between fit-for-purpose, narrow, and global metrics. These metrics can be used both to compare simulation packages, like those listed in Table 2, and to refine initial simulator formulations. Moreover, this review has identified criteria for determining a simulator’s relevance to specific applications, including distributional assumptions, ability to match a template, handling of multivariate relationships, and incorporation of experimental or biological effects. Between simulators, there are trade-offs between faithfulness, controllability, and generality of application, and the most appropriate approach will depend on the relative importance of these factors.

A key insight of modern simulation is that models can be trained on real experimental data from related contexts. A semisynthetic approach can enhance simulation quality and requires less implementation effort compared to designing a de novo mechanism. However, it is worth cautioning that even a simulator that perfectly emulates a template dataset may still be unhelpful if the template is not a good match to the motivating problem. There is a trade-off between using template data from pilot experiments, which are small but highly representative of future conditions, and large public databases, which offer more data but may not directly relate to the problem of interest.

We expect future simulators to support workflows with novel data modalities and their combinations.

For example, while we considered integration across batches and assays, their heterogeneity was mild. Developing simulators that model variation not just across batches but across cohorts, and not just for different amplicon technologies but for entirely different assay types, is a worthwhile direction for further study. Moreover, though we focused on simulating community taxonomic and metabolomic profiles, methods for simulating metagenomics reads have also been proposed [40, 46, 34], and incorporating template data into read simulation could allow systematic study of entire processing and analysis pipelines, similar to recent advances in single-cell read simulation [107].

As simulators become easier to develop and apply, the potential for reusing others' work will increase. For example, instead of creating one-off simulators for the power analysis in a grant proposal or the simulation study in a methods paper, researchers could borrow and modify existing simulator definitions and output. Curating repositories of reusable simulators, where the template data and generating mechanisms are specified, would be valuable. While we have focused on researcher-level power analysis, benchmarking, and reliability analysis, simulation also has the potential to resolve field-level controversies. For example, no consensus has emerged for the analysis of strain-level variation, with some researchers emphasizing their importance to health outcomes [108] and others cautioning against attempts to detect them with existing technology [45, 53]. Simulation can provide the ground truth and control necessary for fine-grained discussion of these issues. Indeed, it has already played an important role in a debate about the use of supervised normalization in microbiome data [101, 13]. As techniques become more powerful and accessible, computational studies using semisynthetic data will become an important part of microbiome research toolkit.

## References

- [1] Angela Andreella, Livio Fino, Bruno Scarpa, and Matteo Stocchero. Towards a power analysis for pls-based methods. 2024.
- [2] Michael B. Sohn and Hongzhe Li. A GLM-based latent variable ordination method for microbiome samples. *Biometrics*, 74(2):448–457, June 2018.
- [3] Yoav Benjamini, Abba M. Krieger, and Daniel Yekutieli. Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507, 2006.
- [4] Gilles Blanchard and Etienne Roquain. Adaptive false discovery rate control under independence and dependence. *Journal of Machine Learning Research*, 10(97):2837–2871, 2009.
- [5] Tony Cai and Weidong Liu. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*, 106(494):672–684, June 2011.
- [6] Olivier Chapleur, Céline Madigou, Raphaël Civade, Yohan Rodolphe, Laurent Mazéas, and Théodore Bouchez. Increasing concentrations of phenol progressively affect anaerobic digestion of cellulose and associated microbial communities. *Biodegradation*, 27(1):15–27, November 2015.
- [7] Jun Chen, Frederic D Bushman, James D Lewis, Gary D Wu, and Hongzhe Li. Structure-constrained sparse canonical correlation analysis with an application to microbiome data analysis. *Biostatistics*, 14(2):244–258, April 2013.
- [8] Ashwin Chetty and Ran Blekhan. Multi-omic approaches for host-microbiome data integration. *Gut Microbes*, 16, 2024.
- [9] Marc G. Chevrette, Jennifer R. Bratburd, Cameron R. Currie, and Reed M. Stubbendieck. Experimental microbiomes: Models not to scale. *mSystems*, 4(4), August 2019.
- [10] Joung Min Choi, Ming Ji, Layne T Watson, and Liqing Zhang. DeepMicroGen: a generative adversarial network-based method for longitudinal microbiome data imputation. *Bioinformatics*, 39(5), May 2023.

- [11] Thomas H. Clarke, Chris Greco, Lauren Brinkac, Karen E. Nelson, and Harinder Singh. Mpress: An r-package for accurately predicting power for comparisons of 16s rna microbiome taxa distributions including simulation by dirichlet mixture modeling. *Microorganisms*, 11(5):1166, April 2023.
- [12] Helena L. Crowell, Sarah X. Morillo Leonardo, Charlotte Soneson, and Mark D. Robinson. The shaky foundations of simulating single-cell rna sequencing data. *Genome Biology*, 24(1), March 2023.
- [13] Itay Daybog and Oren Kolodny. A computational framework for resolving the microbiome diversity conundrum. *Nature Communications*, 14(1), December 2023.
- [14] Rebecca A Deek and Hongzhe Li. Inference of microbial covariation networks using copula models with mixture margins. *Bioinformatics*, 39(7), July 2023.
- [15] Jörg Drechsler and Anna-Carolina Haensch. 30 years of synthetic data. *Statistical Science*, 39(2):221–242, 2024.
- [16] Tarn Duong. ks: Kernel density estimation and kernel discriminant analysis for multivariate data inr. *Journal of Statistical Software*, 21(7), 2007.
- [17] Tarn Duong, Bruno Goud, and Kristine Schauer. Closed-form density-based framework for automatic detection of cellular morphology changes. *Proceedings of the National Academy of Sciences*, 109(22):8382–8387, 2012.
- [18] Tarn Duong, Bruno Goud, and Kristine Schauer. Closed-form density-based framework for automatic detection of cellular morphology changes. *Proceedings of the National Academy of Sciences*, 109(22):8382–8387, May 2012.
- [19] Claire Duvallet, Sean M. Gibbons, Thomas Gurry, Rafael A. Irizarry, and Eric J. Alm. Meta-analysis of gut microbiome studies identifies disease-specific and shared responses. *Nature Communications*, 8(1), December 2017.
- [20] Khaled El Emam, Lucy Mosquera, Xi Fang, and Alaa El-Hussuna. An evaluation of the replicability of analyses using synthetic health data. *Scientific Reports*, 14(1):6978, 2024.
- [21] Khaled El Emam, Lucy Mosquera, and Richard Hoptroff. *Practical synthetic data generation: balancing privacy and the broad availability of data*. O’Reilly Media, 2020.
- [22] Marco Fabbri, Daniel Scicchitano, Marco Candela, Silvia Turrone, and Simone Rampelli. Connect the dots: sketching out microbiome interactions through networking approaches. *Microbiome Res. Rep.*, 2(4):25, July 2023.
- [23] Tahsin Ferdous, Lai Jiang, Irina Dinu, Julie Groizeleau, Anita L. Kozyrskyj, Celia M.T. Greenwood, and Marie-Claire Arrieta. The rise to power of the microbiome: power and sample size calculation for microbiome studies. *Mucosal Immunology*, 15(6):1060–1070, November 2022.
- [24] Jerome H. Friedman and Lawrence C. Rafsky. Graph-theoretic measures of multivariate association and prediction. *The Annals of Statistics*, 11(2), June 1983.
- [25] Sarah Friedrich and Tim Friede. On the role of benchmarking data sets and simulations in method comparison studies. *Biometrical Journal*, 66(1):2200212, 2024.
- [26] Adrian Fritz, Peter Hofmann, Stephan Majda, Eik Dahms, Johannes Dröge, Jessika Fiedler, Till R Lesker, Peter Belmann, Matthew Z DeMaere, Aaron E Darling, Alexander Sczyrba, Andreas Bremges, and Alice C McHardy. CAMISIM: simulating metagenomes and microbial communities. *Microbiome*, 7(1):17, February 2019.



- [27] Yu Gao, Yağmur Şimşek, Emma Gheysen, Tuomas Borman, Yi Li, Leo Lahti, Karoline Faust, and Daniel Rios Garza. miasim: an R/Bioconductor package to easily simulate microbial community dynamics. *Methods Ecol. Evol.*, May 2023.
- [28] Patrick G. Gavin, Jane A. Mullaney, Dorothy Loo, Kim-Anh Lê Cao, Peter A. Gottlieb, Michelle M. Hill, Danny Zipris, and Emma E. Hamilton-Williams. Intestinal metaproteomics reveals host-microbiota interactions in subjects at risk for type 1 diabetes. *Diabetes Care*, 41(10):2178–2186, August 2018.
- [29] Georg K. Gerber. *Longitudinal Microbiome Data Analysis*, page 97–111. Elsevier, 2015.
- [30] Sean M Gibbons, Claire Duvall, and Eric J Alm. Correcting for batch effects in case-control microbiome studies. *PLoS Comput. Biol.*, 14(4):e1006102, April 2018.
- [31] Jack A Gilbert, Martin J Blaser, J Gregory Caporaso, Janet K Jansson, Susan V Lynch, and Rob Knight. Current understanding of the human microbiome. *Nat. Med.*, 24(4):392–400, April 2018.
- [32] Gregory B Gloor, Jean M Macklaim, Vera Pawlowsky-Glahn, and Juan J Egozcue. Microbiome datasets are compositional: and this is not optional. *Frontiers in microbiology*, 8:2224, 2017.
- [33] Alexej Gossmann, Pascal Zille, Vince Calhoun, and Yu-Ping Wang. FDR-corrected sparse canonical correlation analysis with applications to imaging genomics. *IEEE Trans. Med. Imaging*, 37(8):1761–1774, August 2018.
- [34] Hadrien Gourel, Oskar Karlsson-Lindsjö, Juliette Hayer, and Erik Bongcam-Rudloff. Simulating illumina metagenomic data with InSilicoSeq. *Bioinformatics*, 35(3):521–522, February 2019.
- [35] Rosember Guerra-Urzola, Katrijn Van Deun, Juan C. Vera, and Klaas Sijtsma. A guide for sparse pca: Model comparison and applications. *Psychometrika*, 86(4):893–919, June 2021.
- [36] Bastiaan W. Haak, Ricard Argelaguet, Cormac M. Kinsella, Robert F. J. Kullberg, Jacqueline M. Lankelma, Martin Deijs, Michelle Klein, Maarten F. Jebbink, Floor Hugenholtz, Sarantos Kostidis, Martin Giera, Theodorus B. M. Hakvoort, Wouter J. de Jonge, Marcus J. Schultz, Tom van Gool, Tom van der Poll, Willem M. de Vos, Lia M. van der Hoek, and W. Joost Wiersinga. Integrative transkingdom analysis of the gut microbiome in antibiotic perturbation and critical illness. *mSystems*, 6(2):10.1128/msystems.01148–20, 2021.
- [37] Jo Handelsman. Metagenomics: application of genomics to uncultured microorganisms. *Microbiol. Mol. Biol. Rev.*, 68(4):669–685, December 2004.
- [38] Mengyu He, Ni Zhao, and Glen A Satten. MIDASim: a fast and simple simulator for realistic microbiome data. *bioRxiv.org*, March 2024.
- [39] Wenxing Hu, Dongdong Lin, Shaolong Cao, Jingyu Liu, Jiayu Chen, Vince D Calhoun, and Yu-Ping Wang. Adaptive sparse multiple canonical correlation analysis with application to imaging (epi)genomics study of schizophrenia. *IEEE Trans. Biomed. Eng.*, 65(2):390–399, February 2018.
- [40] Ben Jia, Liming Xuan, Kaiye Cai, Zhiqiang Hu, Liangxiao Ma, and Chaochun Wei. NeSSM: A next-generation sequencing simulator for metagenomics. *PLoS One*, 8(10):e75448, October 2013.
- [41] Duo Jiang, Courtney R Armour, Chenxiao Hu, Meng Mei, Chuan Tian, Thomas J Sharpton, and Yuan Jiang. Microbiome multi-omics network analysis: Statistical considerations, limitations, and opportunities. *Front. Genet.*, 10:995, November 2019.
- [42] Ruochen Jiang, Wei Vivian Li, and Jingyi Jessica Li. mbimpute: an accurate and robust imputation method for microbiome data. *Genome Biology*, 22(1), June 2021.

- [43] Shuang Jiang, Guanghua Xiao, Andrew Y Koh, Jiwoong Kim, Qiwei Li, and Xiaowei Zhan. A bayesian zero-inflated negative binomial regression model for the integrative analysis of microbiome data. *Bio-statistics*, 22(3):522–540, December 2019.
- [44] Harry Joe. *Dependence modeling with copulas*. Chapman & Hall/CRC Monographs on Statistics and Applied Probability. CRC Press, London, England, January 2023.
- [45] Jethro S Johnson, Daniel J Spakowicz, Bo-Young Hong, Lauren M Petersen, Patrick Demkowicz, Lei Chen, Shana R Leopold, Blake M Hanson, Hanako O Agresta, Mark Gerstein, Erica Sodergren, and George M Weinstock. Evaluation of 16S rRNA gene sequencing for species and strain-level microbiome analysis. *Nat. Commun.*, 10(1):5029, November 2019.
- [46] Stephen Johnson, Brett Trost, Jeffrey R Long, Vanessa Pittet, and Anthony Kusalik. A better sequence-read simulator program for metagenomics. *BMC Bioinformatics*, 15 Suppl 9:S14, September 2014.
- [47] W E Johnson, C Li, and A Rabinovic. Adjusting batch effects in microarray expression data using empirical bayes methods. *Biostatistics*, 8(1):118–127, January 2007.
- [48] Iain M. Johnstone and Debashis Paul. Pca in high dimensions: An orientation. *Proceedings of the IEEE*, 106(8):1277–1292, 2018.
- [49] Alan F Karr, Christine N Kohonen, Anna Oganian, Jerome P Reiter, and Ashish P Sanil. A framework for evaluating the utility of data altered to protect confidentiality. *The American Statistician*, 60(3):224–232, 2006.
- [50] Brendan J. Kelly, Robert Gross, Kyle Bittinger, Scott Sherrill-Mix, James D. Lewis, Ronald G. Collman, Frederic D. Bushman, and Hongzhe Li. Power and sample-size estimation for microbiome studies using pairwise distances and permanova. *Bioinformatics*, 31(15):2461–2468, April 2015.
- [51] Rob Knight, Alison Vrbancac, Bryn C. Taylor, Alexander Aksenov, Chris Callewaert, Justine Debelius, Antonio Gonzalez, Tomasz Kosciolk, Laura-Isobel McCall, Daniel McDonald, Alexey V. Melnik, James T. Morton, Jose Navas, Robert A. Quinn, Jon G. Sanders, Austin D. Swafford, Luke R. Thompson, Anupriya Tripathi, Zhenjiang Z. Xu, Jesse R. Zaneveld, Qiyun Zhu, J. Gregory Caporaso, and Pieter C. Dorrestein. Best practices for analysing microbiomes. *Nature Reviews Microbiology*, 16(7):410–422, May 2018.
- [52] Saritha Kodikara, Susan Ellul, and Kim-Anh Lê Cao. Statistical challenges in longitudinal microbiome data analysis. *Briefings in Bioinformatics*, 23(4), July 2022.
- [53] M Senthil Kumar, Eric V Slud, Christine Hehnly, Lijun Zhang, James Broach, Rafael A Irizarry, Steven J Schiff, and Joseph N Paulson. Differential richness inference for 16S rRNA marker gene surveys. *Genome Biol.*, 23(1):166, August 2022.
- [54] Zachary Kurtz, Christian Mueller, Emily Miraldi, and Richard Bonneau. *SpiecEasi: Sparse Inverse Covariance for Ecological Statistical Inference*, 2024. R package version 1.1.3, commit 5f396da85baa114b31c13d9744c05387a1b04c23.
- [55] Zachary D. Kurtz, Christian L. Müller, Emily R. Miraldi, Dan R. Littman, Martin J. Blaser, and Richard A. Bonneau. Sparse and compositionally robust inference of microbial ecological networks. *PLOS Computational Biology*, 11(5):e1004226, May 2015.
- [56] Leo Lahti, Jarkko Salojärvi, Anne Salonen, Marten Scheffer, and Willem M De Vos. Tipping elements in the human intestinal ecosystem. *Nature communications*, 5(1):4344, 2014.
- [57] Charity W Law, Yunshun Chen, Wei Shi, and Gordon K Smyth. voom: Precision weights unlock linear model analysis tools for rna-seq read counts. *Genome biology*, 15:1–17, 2014.

- [58] Kim-Anh Lê Cao, Simon Boitard, and Philippe Besse. Sparse pls discriminant analysis: biologically relevant feature selection and graphical displays for multiclass problems. *BMC Bioinformatics*, 12(1), June 2011.
- [59] Kim-Anh Lê Cao and Zoe Marie Welham. *Multivariate data integration using R: methods and applications with the mixOmics package*. Chapman and Hall/CRC, 2021.
- [60] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, February 2004.
- [61] Jeffrey T Leek, Robert B Scharpf, Héctor Corrada Bravo, David Simcha, Benjamin Langmead, W Evan Johnson, Donald Geman, Keith Baggerly, and Rafael A Irizarry. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat. Rev. Genet.*, 11(10):733–739, October 2010.
- [62] Hechen Li, Ziqi Zhang, Michael Squires, Xi Chen, and Xiuwei Zhang. scmultisim: simulation of multi-modality single cell data guided by cell-cell interactions and gene regulatory networks. March 2023.
- [63] Huang Lin and Shyamal Das Peddada. Analysis of compositions of microbiomes with bias correction. *Nature communications*, 11(1):3514, 2020.
- [64] Wodan Ling, Jiuyao Lu, Ni Zhao, Anju Lulla, Anna M Plantinga, Weijia Fu, Angela Zhang, Hongjiao Liu, Hoseung Song, Zhigang Li, Jun Chen, Timothy W Randolph, Wei Li A Koay, James R White, Lenore J Launer, Anthony A Fodor, Katie A Meyer, and Michael C Wu. Batch effects removal for microbiome data via conditional quantile regression. *Nat. Commun.*, 13(1):5418, September 2022.
- [65] Michael I Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for rna-seq data with *deseq2*. *Genome biology*, 15:1–21, 2014.
- [66] Malte D Luecken, M Büttner, K Chaichoompu, A Danese, M Interlandi, M F Mueller, D C Strobl, L Zappia, M Dugas, M Colomé-Tatché, and Fabian J Theis. Benchmarking atlas-level data integration in single-cell genomics. *Nat. Methods*, 19(1):41–50, January 2022.
- [67] Rong Ma, Eric D. Sun, David Donoho, and James Zou. Principled and interpretable alignability testing and integration of single-cell data. *Proceedings of the National Academy of Sciences*, 121(10), February 2024.
- [68] Siyuan Ma, Boyu Ren, Himel Mallick, Yo Sup Moon, Emma Schwager, Sagun Maharjan, Timothy L. Tickle, Yiren Lu, Rachel N. Carmody, Eric A. Franzosa, Lucas Janson, and Curtis Huttenhower. A statistical model for describing and simulating microbial community profiles. *PLOS Computational Biology*, 17(9):e1008913, September 2021.
- [69] Venkata Suhas Maringanti, Vanni Bucci, and Georg K. Gerber. Mditre: Scalable and interpretable machine learning for predicting host status from temporal microbiome dynamics. *mSystems*, 7(5), October 2022.
- [70] Daniel McDonald, Embriette Hyde, Justine W Debelius, James T Morton, Antonio Gonzalez, Gail Ackermann, Alexander A Aksenov, Bahar Behsaz, Cairiona Brennan, Yingfeng Chen, Lindsay De-Right Goldasich, Pieter C Dorrestein, Robert R Dunn, Ashkaan K Fahimipour, James Gaffney, Jack A Gilbert, Grant Gogul, Jessica L Green, Philip Hugenholtz, Greg Humphrey, Curtis Huttenhower, Matthew A Jackson, Stefan Janssen, Dilip V Jeste, Lingjing Jiang, Scott T Kelley, Dan Knights, Tomasz Kosciolk, Joshua Ladau, Jeff Leach, Clarisse Marotz, Dmitry Meleshko, Alexey V Melnik, Jessica L Metcalf, Hosein Mohimani, Emmanuel Montassier, Jose Navas-Molina, Tanya T Nguyen, Shyamal Peddada, Pavel Pevzner, Katherine S Pollard, Gholamali Rahnavard, Adam Robbins-Pianka, Naseer Sangwan, Joshua Shorenstein, Larry Smarr, Se Jin Song, Timothy Spector, Austin D Swafford, Varykina G Thackray, Luke R Thompson, Anupriya Tripathi, Yoshiki Vázquez-Baeza, Alison Vrbanc,

Paul Wischmeyer, Elaine Wolfe, Qiyun Zhu, Rob Knight, Allison E Mann, Amnon Amir, Angel Frazier, Cameron Martino, Carlito Lebrilla, Catherine Lozupone, Cecil M Lewis, Jr, Charles Raison, Chi Zhang, Christian L Lauber, Christina Warinner, Christopher A Lowry, Chris Callewaert, Cinnamon Bloss, Dana Willner, Daniela Domingos Galzerani, David J Gonzalez, David A Mills, Deepak Chopra, Dirk Gevers, Donna Berg-Lyons, Dorothy D Sears, Doug Wendel, Elijah Lovelace, Emily Pierce, Emily TerAvest, Evan Bolyen, Frederic D Bushman, Gary D Wu, George M Church, Gordon Saxe, Hanna D Holscher, Ivo Ugrina, J Bruce German, J Gregory Caporaso, Jacob M Wozniak, Jacqueline Kerr, Jacques Ravel, James D Lewis, Jan S Suchodolski, Janet K Jansson, Jarrad T Hampton-Marcell, Jason Bobe, Jeroen Raes, John H Chase, Jonathan A Eisen, Jonathan Monk, Jose C Clemente, Joseph Petrosino, Julia Goodrich, Julia Gauglitz, Julian Jacobs, Karsten Zengler, Kelly S Swanson, Kim Lewis, Kris Mayer, Kyle Bittinger, Lindsay Dillon, Livia S Zaramela, Lynn M Schriml, Maria G Dominguez-Bello, Marta M Jankowska, Martin Blaser, Meg Pirrung, Michael Minson, Mike Kurisu, Nadim Ajami, Neil R Gottel, Nicholas Chia, Noah Fierer, Owen White, Patrice D Cani, Pawel Gajer, Philip Strandwitz, Purna Kashyap, Rachel Dutton, Rachel S Park, Ramnik J Xavier, Robert H Mills, Rosa Krajmalnik-Brown, Ruth Ley, Sarah M Owens, Scott Klemmer, Sébastien Matamoros, Siavash Mirarab, Stephanie Moorman, Susan Holmes, Tara Schwartz, Tifani W Eshoo-Anton, Tim Vigers, Vineet Pandey, Will Van Treuren, Xin Fang, Zhenjiang Zech Xu, Alan Jarmusch, Justin Geier, Nicolai Reeve, Ricardo Silva, Evguenia Kopylova, Dominic Nguyen, Karenina Sanders, Rodolfo Antonio Salido Benitez, Arthur Cole Heale, Max Abramson, Jérôme Waldispühl, Alexander Butyaev, Chris Drogaris, Elena Nazarova, Madeleine Ball, Beau Gunderson, and The American Gut Consortium. American gut: An open platform for citizen science microbiome research. *mSystems*, 3(3), June 2018.

- [71] Paul J. McMurdie and Susan Holmes. Waste not, want not: Why rarefying microbiome data is inadmissible. *PLoS Computational Biology*, 10(4):e1003531, April 2014.
- [72] Fernando Meyer, Adrian Fritz, Zhi-Luo Deng, David Koslicki, Till Robin Lesker, Alexey Gurevich, Gary Robertson, Mohammed Alser, Dmitry Antipov, Francesco Beghini, Denis Bertrand, Jaqueline J. Brito, C. Titus Brown, Jan Buchmann, Aydin Buluç, Bo Chen, Rayan Chikhi, Philip T. L. C. Clausen, Alexandru Cristian, Piotr Wojciech Dabrowski, Aaron E. Darling, Rob Egan, Eleazar Eskin, Evangelos Georganas, Eugene Goltsman, Melissa A. Gray, Lars Hestbjerg Hansen, Steven Hofmeyr, Pingqin Huang, Luiz Irber, Huijue Jia, Tue Sparholt Jørgensen, Silas D. Kieser, Terje Klemetsen, Axel Kola, Mikhail Kolmogorov, Anton Korobeynikov, Jason Kwan, Nathan LaPierre, Claire Lemaitre, Chenhao Li, Antoine Limasset, Fabio Malcher-Miranda, Serghei Mangul, Vanessa R. Marcelino, Camille Marchet, Pierre Marijon, Dmitry Meleshko, Daniel R. Mende, Alessio Milanese, Niranjana Nagarajan, Jakob Nissen, Sergey Nurk, Leonid Oliner, Lucas Paoli, Pierre Peterlongo, Vitor C. Piro, Jacob S. Porter, Simon Rasmussen, Evan R. Rees, Knut Reinert, Bernhard Renard, Espen Mikal Robertsen, Gail L. Rosen, Hans-Joachim Ruscheweyh, Varuni Sarwal, Nicola Segata, Enrico Seiler, Lizhen Shi, Fengzhu Sun, Shinichi Sunagawa, Søren Johannes Sørensen, Ashleigh Thomas, Chengxuan Tong, Mirko Trajkovski, Julien Tremblay, Gherman Urtskiy, Riccardo Vicedomini, Zhengyang Wang, Ziyi Wang, Zhong Wang, Andrew Warren, Nils Peder Willassen, Katherine Yelick, Ronghui You, Georg Zeller, Zhengqiao Zhao, Shanfeng Zhu, Jie Zhu, Ruben Garrido-Oter, Petra Gastmeier, Stephane Haquard, Susanne Häußler, Ariane Khaledi, Friederike Maechler, Fantin Mesny, Simona Radutoiu, Paul Schulze-Lefert, Nathiana Smit, Till Strowig, Andreas Bremges, Alexander Sczyrba, and Alice Carolyn McHardy. Critical assessment of metagenome interpretation: the second round of challenges. *Nature Methods*, 19(4):429–440, April 2022.
- [73] Ramyar Molania, Johann A Gagnon-Bartsch, Alexander Dobrovic, and Terence P Speed. A new normalization for nanostring ncounter gene expression data. *Nucleic Acids Res.*, 47(12):6073–6083, July 2019.
- [74] Efrat Muller, Yadid M Algavi, and Elhanan Borenstein. The gut microbiome-metabolome dataset collection: a curated resource for integrative meta-analysis. *NPJ Biofilms Microbiomes*, 8(1):79, October 2022.

- [75] Vegard Nygaard, Einar Andreas Rødland, and Eivind Hovig. Methods that remove batch effects while retaining group differences may lead to exaggerated confidence in downstream analyses. *Biostatistics*, 17(1):29–39, August 2015.
- [76] Michael Olbrich, Axel Künstner, and Hauke Busch. MBECS: Microbiome batch effects correction suite. *BMC Bioinformatics*, 24(1):182, May 2023.
- [77] Samuel Ozminkowski and Claudia Solís-Lemus. Identifying microbial drivers in biological phenotypes with a bayesian network regression model. *Ecology and Evolution*, 14(5), May 2024.
- [78] Edoardo Pasoli, Lucas Schiffer, Paolo Manghi, Audrey Renson, Valerie Obenchain, Duy Tin Truong, Francesco Beghini, Faizan Malik, Marcel Ramos, Jennifer B Dowd, Curtis Huttenhower, Martin Morgan, Nicola Segata, and Levi Waldron. Accessible, curated metagenomic data through ExperimentHub. *Nat. Methods*, 14(11):1023–1024, November 2017.
- [79] Ilaria Patuzzi, Giacomo Baruzzo, Carmen Losasso, Antonia Ricci, and Barbara Di Camillo. metaSPARSim: a 16S rRNA gene sequencing count data simulator. *BMC Bioinformatics*, 20(Suppl 9):416, November 2019.
- [80] Carel F. W. Peeters, Anders Ellern Bilgrau, and Wessel N. van Wieringen. rags2ridges: A one-stop- $\ell_2$ -shop for graphical modeling of high-dimensional precision matrices. *Journal of Statistical Software*, 102(4):1–32, 2022.
- [81] Vasileios C. Pezoulas, Dimitrios I. Zaridis, Eugenia Mylona, Christos Androutsos, Kosmas Apostolidis, Nikolaos S. Tachos, and Dimitrios I. Fotiadis. Synthetic data generation methods in healthcare: A review on open-source tools and methods. *Computational and Structural Biotechnology Journal*, 23:2892–2910, December 2024.
- [82] Yishay Pinto and Ami S. Bhatt. Sequencing-based analysis of microbiomes. *Nature Reviews Genetics*, June 2024.
- [83] Gibraan Rahman, Daniel McDonald, Antonio Gonzalez, Yoshiki Vázquez-Baeza, Lingjing Jiang, Climent Casals-Pascual, Daniel Hakim, Amanda Hazel Dilmore, Brent Nowinski, Shyamal Peddada, and Rob Knight. Determination of effect sizes for power analysis for microbiome studies using large microbiome databases. *Genes*, 14(6):1239, June 2023.
- [84] Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing l1-penalized log-determinant divergence. *Electron. J. Stat.*, 5(none):935–980, January 2011.
- [85] Boyu Ren, Sergio Bacallado, Stefano Favaro, Tommi Vatanen, Curtis Huttenhower, and Lorenzo Trippa. Bayesian mixed effects models for zero-inflated compositions in microbiome data analysis. *The Annals of Applied Statistics*, 14(1), March 2020.
- [86] R A Rigby and D M Stasinopoulos. Generalized additive models for location, scale and shape. *J. R. Stat. Soc. Ser. C. Appl. Stat.*, 54(3):507–554, June 2005.
- [87] Davide Risso, Fanny Perraudeau, Svetlana Gribkova, Sandrine Dudoit, and Jean-Philippe Vert. A general and flexible method for signal extraction from single-cell rna-seq data. *Nature Communications*, 9(1), January 2018.
- [88] Ruichen Rong, Shuang Jiang, Lin Xu, Guanghua Xiao, Yang Xie, Dajiang J Liu, Qiwei Li, and Xiaowei Zhan. MB-GAN: Microbiome simulation via generative adversarial network. *Gigascience*, 10(2), February 2021.
- [89] Kris Sankaran and Susan P Holmes. Latent variable modeling for the microbiome. *Biostatistics*, 20(4):599–614, October 2019.

- [90] Kris Sankaran and Susan P Holmes. Multitable methods for microbiome data integration. *Front. Genet.*, 10:627, August 2019.
- [91] Musfiqur Sazal, Kalai Mathee, Daniel Ruiz-Perez, Trevor Cickovski, and Giri Narasimhan. Inferring directional relationships in microbial communities using signed bayesian networks. *BMC Genomics*, 21(S6), December 2020.
- [92] Alexander Sczyrba, Peter Hofmann, Peter Belmann, David Koslicki, Stefan Janssen, Johannes Dröge, Ivan Gregor, Stephan Majda, Jessika Fiedler, Eik Dahms, Andreas Bremges, Adrian Fritz, Ruben Garrido-Oter, Tue Sparholt Jørgensen, Nicole Shapiro, Philip D Blood, Alexey Gurevich, Yang Bai, Dmitrij Turaev, Matthew Z DeMaere, Rayan Chikhi, Niranjan Nagarajan, Christopher Quince, Fernando Meyer, Monika Balvočiūtė, Lars Hestbjerg Hansen, Søren J Sørensen, Burton K H Chia, Bertrand Denis, Jeff L Froula, Zhong Wang, Robert Egan, Dongwan Don Kang, Jeffrey J Cook, Charles Deltel, Michael Beckstette, Claire Lemaitre, Pierre Peterlongo, Guillaume Rizk, Dominique Lavenier, Yu-Wei Wu, Steven W Singer, Chirag Jain, Marc Strous, Heiner Klingenberg, Peter Meinicke, Michael D Barton, Thomas Lingner, Hsin-Hung Lin, Yu-Chieh Liao, Genivaldo Gueiros Z Silva, Daniel A Cuevas, Robert A Edwards, Surya Saha, Vitor C Piro, Bernhard Y Renard, Mihai Pop, Hans-Peter Klenk, Markus Göker, Nikos C Kyrpides, Tanja Woyke, Julia A Vorholt, Paul Schulze-Lefert, Edward M Rubin, Aaron E Darling, Thomas Rattei, and Alice C McHardy. Critical assessment of metagenome interpretation—a benchmark of metagenomics software. *Nature Methods*, 14(11):1063–1071, October 2017.
- [93] Ashley Shade and Jo Handelsman. Beyond the venn diagram: the hunt for a core microbiome. *Environmental microbiology*, 14(1):4–12, 2012.
- [94] Amrit Singh, Casey P Shannon, Benoît Gautier, Florian Rohart, Michaël Vacher, Scott J Tebbutt, and Kim-Anh Lê Cao. Diablo: an integrative approach for identifying key molecular drivers from multi-omics assays. *Bioinformatics*, 35(17):3055–3062, January 2019.
- [95] John R. Skalski, Shelby M. Richins, and Richard L. Townsend. A statistical test and sample size recommendations for comparing community composition following pca. *PLOS ONE*, 13(10):e0206033, October 2018.
- [96] Charlotte Soneson, Sarah Gerster, and Mauro Delorenzi. Batch effect confounding leads to strong bias in performance estimates obtained by cross-validation. *PLoS ONE*, 9(6):e100335, June 2014.
- [97] Dongyuan Song, Qingyang Wang, Guanao Yan, Tianyang Liu, Tianyi Sun, and Jingyi Jessica Li. scdesign3 generates realistic in silico data for multimodal single-cell and spatial omics. *Nature Biotechnology*, 42(2):247–252, May 2023.
- [98] D Mikis Stasinopoulos and Robert A Rigby. Generalized additive models for location scale and shape (GAMLSS) inr. *J. Stat. Softw.*, 23(7), 2007.
- [99] John D. Storey, Jonathan E. Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(1):187–205, December 2003.
- [100] Tianyi Sun, Dongyuan Song, Wei Vivian Li, and Jingyi Jessica Li. scdesign2: a transparent simulator that generates high-fidelity single-cell gene expression count data with gene correlations captured. *Genome biology*, 22(1):163, 2021.
- [101] Gerry Tonkin-Hill. GitHub - gtonkinhill/TCGA\_analysis — github.com. [https://github.com/gtonkinhill/TCGA\\_analysis](https://github.com/gtonkinhill/TCGA_analysis), 2023. [Accessed 21-06-2024].
- [102] Mireia Valles-Colomer, Cristina Menni, Sarah E. Berry, Ana M. Valdes, Tim D. Spector, and N. Segata. Cardiometabolic health, diet and the gut microbiome: a meta-omics perspective. *Nature Medicine*, 29:551–561, 2023.



- [103] Yiwen Wang and Kim-Anh Lê Cao. Managing batch effects in microbiome data. *Briefings in Bioinformatics*, 21(6):1954–1970, November 2019.
- [104] Sophie Weiss, Zhenjiang Zech Xu, Shyamal Peddada, Amnon Amir, Kyle Bittinger, Antonio Gonzalez, Catherine Lozupone, Jesse R. Zaneveld, Yoshiki Vázquez-Baeza, Amanda Birmingham, Embriette R. Hyde, and Rob Knight. Normalization and microbial differential abundance strategies depend upon data characteristics. *Microbiome*, 5(1), March 2017.
- [105] Cecilia Wieder, Rachel P. J. Lai, and Timothy M. D. Ebbels. Single sample pathway analysis in metabolomics: performance evaluation and application. *BMC Bioinformatics*, 23(1), November 2022.
- [106] Justin Williams, Hector Corrada Bravo, Jennifer Tom, and Joseph Nathaniel Paulson. micro-biomedasim: Simulating longitudinal differential abundance for microbiome data. *F1000Research*, 8:1769, February 2020.
- [107] Guanao Yan, Dongyuan Song, and Jingyi Jessica Li. screadsim: a single-cell rna-seq and atac-seq read simulator. *Nature Communications*, 14(1), November 2023.
- [108] Yan Yan, Long H Nguyen, Eric A Franzosa, and Curtis Huttenhower. Strain-level epidemiology of microbial communities and the human microbiome. *Genome Med.*, 12(1):71, August 2020.
- [109] Luke Zappia, Belinda Phipson, and Alicia Oshlack. Splatter: simulation of single-cell RNA sequencing data. *Genome Biol.*, 18(1), December 2017.
- [110] Yanyan Zeng, Jing Li, Chaochun Wei, Hongyu Zhao, and Tao Wang. mbdnoise: microbiome data denoising using zero-inflated probabilistic principal components analysis. *Genome Biology*, 23(1), April 2022.
- [111] Yanyan Zeng, Daolin Pang, Hongyu Zhao, and Tao Wang. A zero-inflated logistic normal multinomial model for extracting microbial compositions. *J. Am. Stat. Assoc.*, pages 1–31, February 2022.
- [112] Yiran Zhang, Rita Hedo, Anna Rivera, Rudolph Rull, Sabrina Richardson, and Xin M Tu. Post hoc power analysis: is it an informative and meaningful analysis? *General psychiatry*, 32(4), 2019.
- [113] Zhaojun Zhang, Divij Mathew, Tristan Lim, Kaishu Mason, Clara Morral Martinez, Sijia Huang, E John Wherry, Katalin Susztak, Andy J Minn, Zongming Ma, and Nancy R Zhang. Signal recovery in single cell batch integration. *bioRxiv.org*, September 2023.