

RESEARCH ARTICLE

# Sparse and Compositionally Robust Inference of Microbial Ecological Networks

Zachary D. Kurtz<sup>1</sup> , Christian L. Müller<sup>2,3</sup> , Emily R. Miraldi<sup>1,2,3</sup> , Dan R. Littman<sup>1</sup>, Martin J. Blaser<sup>1</sup>, Richard A. Bonneau<sup>2,3,4</sup> \*

**1** Departments of Microbiology and Medicine, New York University School of Medicine, New York, New York, United States of America, **2** Department of Biology, Center for Genomics and Systems Biology, New York University, New York, New York, United States of America, **3** Courant Institute of Mathematical Sciences, New York University, New York, New York, United States of America, **4** Simons Center for Data Analysis, Simons Foundation, New York, New York, United States of America

 These authors contributed equally to this work.

\* [rb133@nyu.edu](mailto:rb133@nyu.edu)



 OPEN ACCESS

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

**Editor:** Christian von Mering, University of Zurich and Swiss Institute of Bioinformatics, SWITZERLAND

**Received:** August 19, 2014

**Accepted:** March 2, 2015

**Published:** May 7, 2015

**Copyright:** © 2015 Kurtz et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All code and data are available from <http://bonneaulab.bio.nyu.edu/>

**Funding:** This work was supported by the National Institutes of Health grants T32AI007180-30 (ZDK, ERM), R01 DK103358-01 (DRL, RAB) and R01 GM63270 (MJB) and the Simons Foundation (CLM, ERM, RAB). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing Interests:** The authors have declared that no competing interests exist.

## Abstract

16S ribosomal RNA (rRNA) gene and other environmental sequencing techniques provide snapshots of microbial communities, revealing phylogeny and the abundances of microbial populations across diverse ecosystems. While changes in microbial community structure are demonstrably associated with certain environmental conditions (from metabolic and immunological health in mammals to ecological stability in soils and oceans), identification of underlying mechanisms requires new statistical tools, as these datasets present several technical challenges. First, the abundances of microbial operational taxonomic units (OTUs) from amplicon-based datasets are compositional. Counts are normalized to the total number of counts in the sample. Thus, microbial abundances are not independent, and traditional statistical metrics (e.g., correlation) for the detection of OTU-OTU relationships can lead to spurious results. Secondly, microbial sequencing-based studies typically measure hundreds of OTUs on only tens to hundreds of samples; thus, inference of OTU-OTU association networks is severely under-powered, and additional information (or assumptions) are required for accurate inference. Here, we present SPIEC-EASI (**SP**arse **I**nverse **C**ovariance Estimation for **E**cological **A**ssociation Inference), a statistical method for the inference of microbial ecological networks from amplicon sequencing datasets that addresses both of these issues. SPIEC-EASI combines data transformations developed for compositional data analysis with a graphical model inference framework that assumes the underlying ecological association network is sparse. To reconstruct the network, SPIEC-EASI relies on algorithms for sparse neighborhood and inverse covariance selection. To provide a synthetic benchmark in the absence of an experimentally validated gold-standard network, SPIEC-EASI is accompanied by a set of computational tools to generate OTU count data from a set of diverse underlying network topologies. SPIEC-EASI outperforms state-of-the-art methods to recover edges and network properties on synthetic data under a variety of scenarios. SPIEC-EASI also reproducibly predicts previously unknown microbial associations using data from the American Gut project.

## Author Summary

Genomic survey of microbes by 16S rRNA gene sequencing and metagenomics has inspired appreciation for the role of complex communities in diverse ecosystems. However, due to the unique properties of community composition data, standard data analysis tools are likely to produce statistical artifacts. For a typical experiment studying microbial ecosystems these artifacts can lead to erroneous conclusions about patterns of associations between microbial taxa. We developed a new procedure that seeks to infer ecological associations between microbial populations, by 1) taking advantage of the proportionality invariance of relative abundance data and 2) making assumptions about the underlying network structure when the number of taxa in the dataset is larger than the number of sampled communities. Additionally, we employed a novel tool to generate biologically plausible synthetic data and objectively benchmark current association inference tools. Finally, we tested our procedures on a large-scale 16S rRNA gene sequencing dataset sampled from the human gut.

## Introduction

Low-cost metagenomic and amplicon-based sequencing promises to make the resolution of complex interactions between microbial populations and their surrounding environment a routine component of observational ecology and experimental biology. Indeed, large-scale data collection efforts (such as Earth Microbiome Project [1], the Human Microbiome Project [2], and the American Gut Project [3]) bring an ever-increasing number of samples from soil, marine and animal-associated microbiota to the public domain. Recent research efforts in ecology, statistics, and computational biology have been aimed at reliably inferring novel biological insights and testable hypotheses from population abundances and phylogenies. Classic objectives in community ecology include, (i) the accurate estimation of the number of taxa (observed and unobserved) from microbial studies [4] and, related to that, (ii) the estimation of community diversity within and across different habitats from the modeled population counts [5]. Moreover, some microbial compositions appear to form distinct clusters, leading to the concept of enterotypes, or ecological steady states in the gut [6], but their existence has not been established with certainty [7]. Another aim of recent studies is the elucidation of connections between microbes and environmental or host covariates. Examples include a novel statistical regression framework for relating microbiome compositions and covariates in the context of nutrient intake [8], observations that microbiome compositions strongly correlate with disease status in new-onset Crohn's disease [9], and the connections between helminth infection and the microbiome diversity [10].

One goal of microbiome studies is the accurate inference of microbial ecological interactions from population-level data [11]. 'Interactions' are inferred by detecting significant (typically non-directional) associations between sampled populations, e.g., by measuring frequency of co-occurrence [12, 13]. Microbiota are measured by profiling variable regions of bacterial 16S rRNA gene sequences. These regions are amplified, sequenced, and the resulting reads are then grouped into common Operational Taxonomic Units (OTUs) and quantified, with OTU counts serving as a proxy to the underlying microbial populations' abundances. Knowledge of interaction networks (here, a measure of microbial association) provides a foundation to predictively model the interplay between environment and microbial populations. A recent example is the successful construction of a dynamic differential equation model to describe the primary succession of intestinal microbiota in mice [14]. A commonly used tool to infer a

network is correlation analysis; that is computing Pearson's correlation coefficient among all pairs of OTU samples, and an interaction between microbes is assumed when the absolute value of the correlation coefficient is sufficiently high [15, 16].

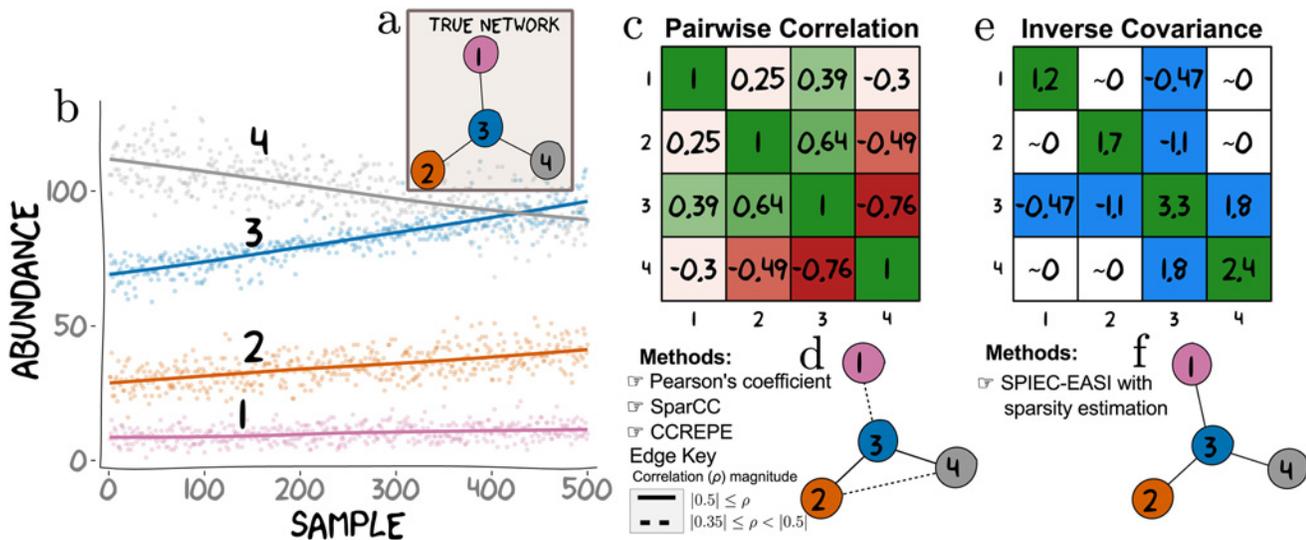
However, applying traditional correlation analysis to amplicon surveys of microbial population data is likely to yield spurious results [9, 17]. To limit experimental biases due to sampling depth, OTU count data is typically transformed by normalizing each OTU count to the total sum of counts in the sample. Thus, communities of microbial relative abundances, termed compositions, are not independent, and classical correlation analysis may fail [18]. Recent methods such as Sparse Correlations for Compositional data (SparCC) [17] and Compositionally Corrected by REnormalization and PErmutation (CCREPE) [9, 11, 19] are designed to account for these compositional biases and represent the state of the art in the field. Yet, it is not clear that correlation is the proper measure of association. For example, correlations can arise between OTUs that are indirectly connected in an ecological network (we expand on this point below).

Dimensionality poses another challenge to statistical analysis of microbiome studies, as the number of measured OTUs  $p$  is on the order of hundreds to thousands whereas the number of samples  $n$  generally ranges from tens to hundreds. This implies that any meaningful interaction inference scheme must operate in the underdetermined data regime ( $p > n$ ), which is viable only if additional assumptions about the interaction network can be made. As technological developments lead to greater sequencing depths, new computational methods that address the ( $p > n$ ) challenge will become increasingly important.

In the present work, we present a novel strategy to infer networks from (potentially high-dimensional) community composition data. We introduce SPIEC-EASI (**S**Parse **I**nvers**E** **C**ovariance **E**stimation for **E**cological **A**sociation **I**nference, pronounced *speakeasy*), a new statistical method for the inference of microbial ecological networks and generation of realistic synthetic data. SPIEC-EASI inference comprises two steps: First, a transformation from the field of compositional data analysis is applied to the OTU data. Second, SPIEC-EASI estimates the interaction graph from the transformed data using one of two methods: (i) neighborhood selection [20, 21] and (ii) sparse inverse covariance selection [22, 23]. Unlike empirical correlation or covariance estimation, used in SparCC and CCREPE, our pipeline seeks to infer an underlying graphical model using the concept of conditional independence.

Informally, two nodes (e.g. OTUs) are conditionally independent if, given the state (e.g. abundance) of all other nodes in the network, neither node provides additional information about the state of the other. A link between any two nodes in the graphical model implies that the OTU abundances are not conditionally independent and that there is a (linear) relationship between them that cannot be better explained by an alternate network wiring. In this way, our method avoids detection of correlated but indirectly connected OTUs, thus ensuring parsimony of the resulting network model (for more detail, see [Materials and Methods](#) and [Fig 1](#)). This model is an undirected graph where links between nodes represent signed associations between OTUs. The use of graphical models has gained considerable popularity in network biology [24–26] and, more recently, in structural biology [27], particularly to correct for transitive correlations in protein structure prediction [28].

To properly benchmark our inference scheme and compare its performance with other state-of-the-art schemes [9, 17], SPIEC-EASI is accompanied by a synthetic data generation routine, which generates realistic synthetic OTU data from networks with diverse topologies. This is significant because, to date, (i) no experimentally verified set of “gold-standard” microbial interactions exists, (ii) previous synthetic benchmark data do not accurately reflect the actual properties of microbiome data [17], and (iii) theoretical and empirical work from high-



**Fig 1. Conditional independence vs correlation analysis for a toy dataset.** In an ecosystem, the abundance of any OTU is potentially dependent on the abundances of other OTUs in the ecological network. Here, we simulate abundances from a network where OTU 3 directly influences (via some set of biological mechanisms) the abundances of OTUs 1, 2 and 4 (a). The inference goal here is to recover the underlying network from the simulated data. b) Absolute abundances of these four OTUs were drawn from a negative-binomial distribution across 500 samples according to the true network (as described in the Methods section). c) Computing all pairwise Pearson correlation yields a symmetric matrix showing patterns of association (positive correlations are green and negative are red). We thresholded entries of the correlation matrix to generate relevance networks. d) A threshold at  $\rho \geq |0.35|$  (represented by dashed and solid edges) results in a network in which OTU 3 is connected to all other OTUs with an additional connection between OTU 2 and OTU 4. A more stringent threshold at  $\rho \geq |0.5|$ , results in a sparser relevance network (notably missing the edge between OTU 3 and OTU 1), and is represented in d by solid edges only. Importantly, no single threshold recovers the true underlying hub topology. e) The inverse sample covariance matrix yields a symmetric matrix where entries are approximately zero if the corresponding OTU pairs are conditionally independent. The network (f) inferred from the non-zero entries (colored in blue in e) identifies the correct hub network. Thus, it is possible to choose a threshold for the sample inverse covariance that faithfully recovers the true network. Such a threshold is not guaranteed to exist for correlation or covariance (the metric used by SparCC and CCREPE). Intuitively, this is because simultaneous direct connections can induce strong correlations between nodes that do not have direct relationships (e.g. OTU 2-4). Conversely, weak correlations can arise between directly connected nodes (e.g. OTU 1-3). Although correlation is a useful measure of association in many contexts, it is a pairwise metric and therefore limited in a multivariate setting. On the other hand, SPIEC-EASI's estimate of entries in the inverse covariance matrix depend on the conditional states of all available nodes. This feature helps SPIEC-EASI avoid detection of indirect network interactions.

doi:10.1371/journal.pcbi.1004226.g001

dimensional statistics [29–31] suggests that network topology can strongly impact network recovery and performance and thus must be considered in the design of synthetic datasets.

We show that SPIEC-EASI is a scalable inference engine that (i) yields superior performance with respect to state-of-the-art methods in terms of interaction recovery and network features in a diverse set of realistic synthetic benchmark scenarios, (ii) provides the most stable and reproducible network when applied to real data, and (iii) reliably estimates an invertible covariance matrix which can be used for additional downstream statistical analysis. In agreement with statistical theory [29], inference on the synthetic datasets demonstrates that the degree distribution of the underlying network has the largest effect on performance, and this effect is observed across all methods tested. SPIEC-EASI network inference applied to actual data from the American Gut Project (AGP) shows (i) that clusters of strongly connected components are likely to contain OTUs with common family membership and (ii) that actual gut microbial networks are likely composites of archetypical network topologies. In the Materials and Methods section, we present statistical and computational aspects of SPIEC-EASI. We then benchmark SPIEC-EASI, comparing it to current inference schemes using synthetic data. We then apply SPIEC-EASI to measurements available from the AGP database. The SPIEC-EASI pipeline is implemented in the R package [SpiecEasi] freely available at <https://github.com/zdk123/SpiecEasi>. All presented numerical data is available at <http://bonneaulab.bio.nyu.edu/>.

## Materials and Methods

SPIEC-EASI comprises both an inference and a synthetic data generation module. Fig 2 summarizes the key components of the pipeline. In this section, we introduce all statistical and computational aspects of the inference scheme and then describe our approach for generating realistic synthetic datasets.

### Data processing and transformation of standard OTU count data

For this discussion, a table of OTU count data, typical output of 16S rRNA gene sequencing data curation pipelines (e.g., mothur [32], QIIME [33]) are given. The OTU data are stored in a matrix  $W \in \mathbb{N}_0^{n \times p}$  where  $w^{(j)} = [w_1^{(j)}, w_2^{(j)}, \dots, w_p^{(j)}]$  denotes the  $p$ -dimensional row vector of OTU counts from the  $j^{th}$  sample,  $j = 1, \dots, n$ , with total cumulative count  $m^{(j)} = \sum_{i=1}^p w_i^{(j)}$ ;  $\mathbb{N}_0$  denotes the set of natural numbers  $\{0, 1, 2, \dots\}$ . As described above, to account for sampling biases, microbiome data is typically transformed by normalizing the raw count data  $w^{(j)}$  with respect to the total count  $m^{(j)}$  of the sample [10]. We thus arrive at vectors of relative abundances or compositions  $x^{(j)} = [\frac{w_1^{(j)}}{m^{(j)}}, \frac{w_2^{(j)}}{m^{(j)}}, \dots, \frac{w_p^{(j)}}{m^{(j)}}]$  for sample  $j$ . Due to this normalization OTU abundances are no longer independent, and the sample space of this  $p$ -part composition  $x^{(j)}$  is not the unconstrained Euclidean space but the  $p$ -dimensional unit simplex  $\mathbb{S}^p \doteq \{x \mid x_i > 0, \sum_{i=1}^p x_i = 1\}$ . Thus, OTU compositions from  $n$  samples are constrained to lie in the unit simplex,  $X \in \mathbb{S}^{n \times p}$ . This restriction of the data to the simplex prohibits the application of standard statistical analysis techniques, such as linear regression or empirical covariance estimation. Covariance matrices of compositional data exhibit, for instance, a negative bias due to closure effects.

Major advances in the statistical analysis of compositional data were achieved by Aitchison in the 1980's [18, 34]. Rather than considering compositions in the simplex, Aitchison proposed log-ratios,  $\log_{x_j}^{x_i}$ , as a basis for studying compositional data. The simple equivalence  $\log_{x_j}^{x_i} = \log \frac{w_i/m}{w_j/m} = \log \frac{w_i}{w_j}$  implies that statistical inferences drawn from analysis of log-ratios of compositions are equivalent to those that could be drawn from the log-ratios of the unobserved absolute measurements, also termed the *basis*.

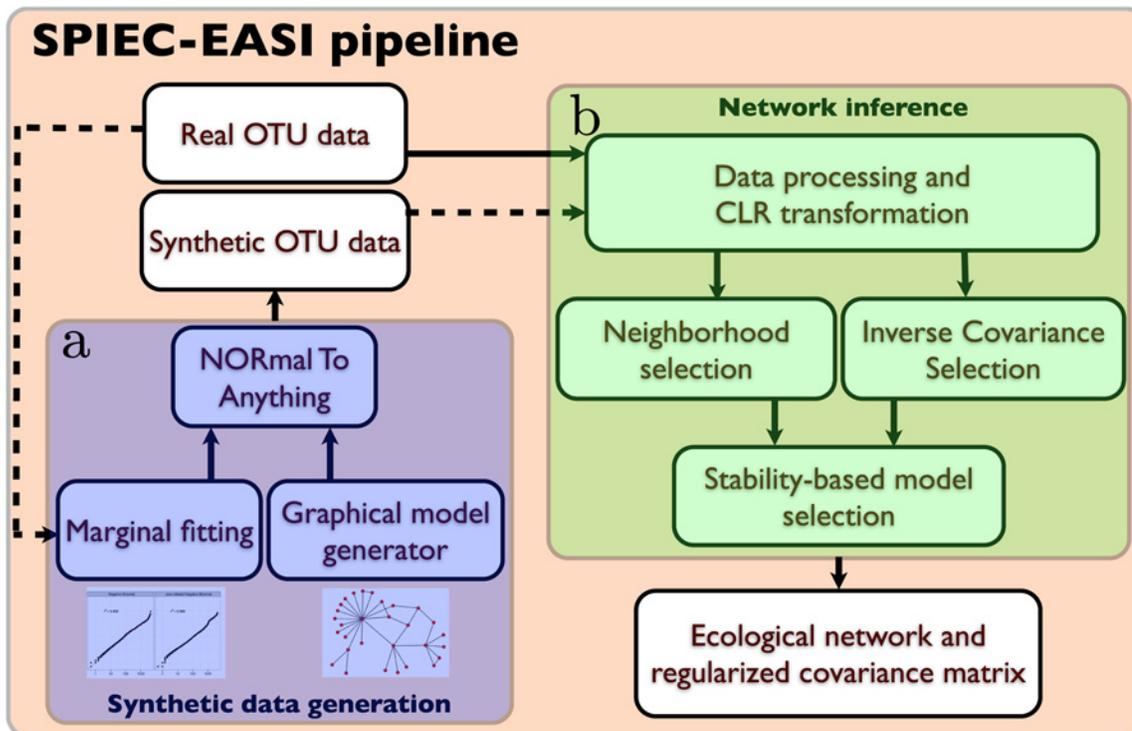
Aitchison also proposed several statistically equivalent log-ratio transformations to remove the unit-sum constraint of compositional data [18]. Here we apply the centered log-ratio (clr) transform:

$$z \doteq \text{clr}(x) = [\log(x_1/g(x)), \dots, \log(x_p/g(x))] = [\log(w_1/g(w)), \dots, \log(w_p/g(w))] \quad (1)$$

where  $g(x) = \left[ \prod_{i=1}^p x_i \right]^{1/p}$  is the geometric mean of the composition vector. The clr transform is symmetric and isometric with respect to the component parts. The resulting vector  $z$  is constrained to a zero sum. The clr transform maps the data from the unit simplex to a  $p - 1$ -dimensional Euclidean space, and the corresponding population covariance matrix  $\Gamma = \text{Cov}[\text{clr}(X)] \in \mathbb{R}^{(p-1) \times (p-1)}$  is also singular [18]. The covariance matrix  $\Gamma$  is related to the population covariance of the log-transformed absolute abundances  $\Omega = \text{Cov}[\log W]$  via the relationship [34]:

$$\Gamma = G\Omega G \quad (2)$$

where  $G = I_p - \frac{1}{p}J$ ,  $I_p$  is the  $p$ -dimensional identity matrix, and  $J = [j_1, j_2, \dots, j_i, \dots, j_p]$ ,  $j_i = [1, 1, \dots, 1]$  the  $p$ -dimensional all-ones vector. For high-dimensional data,  $p > 0$ , the matrix



**Fig 2. Workflow of the SPIEC-EASI pipeline.** The SPIEC-EASI pipeline consists of two independent parts for **a**) synthetic data generation and **b**) network inference. **a**) Synthetic data generation requires an OTU count table and a user-selected network topology. Internally, the parameters of a statistical distribution (the zero-inflated Negative binomial model is suggested) are fit to the OTU marginals of the real data, and are combined with the randomly-generated network in the Normal to Anything (NORTA) approach to generate correlated count data. **b**) Network inference proceeds in three stages on synthetic or real OTU count data: First, data is pre-processed and centered log-ratio (CLR) transformed to ensure compositional robustness. Next, the user selects one of two graphical model inference procedures: 1) Neighborhood selection (the MB method) or 2) inverse covariance selection (the glasso method). SPIEC-EASI network inference assumes that the underlying network is sparse. We infer the correct model sparseness by the Stability Approach to Regularization Selection (StARS), which involves random subsampling of the dataset to find a network with low variability in the selected set of edges. SPIEC-EASI outputs include an ecological network (from the non-zero entries of the inverse covariance matrix) and an invertible covariance matrix. If the network was inferred from synthetic data, it can be compared with the input network to assess inference quality.

doi:10.1371/journal.pcbi.1004226.g002

$G$  is close to the identity matrix, and thus we can assume that a finite sample estimator  $\hat{\Gamma}$  of  $\Gamma$  serves as a good approximation of  $\hat{\Omega}$ . This approximation serves as the basis of our network inference scheme. Finally, because real-world OTU data often contain samples with a zero count for low-abundance OTUs, we add a unit pseudo count to the original count data to avoid numerical problems with the clr transform.

### Inference of microbial associations from microbial abundance datasets

Our key objective is to learn a network of pairwise taxon-taxon associations (putative interactions) from clr-transformed microbiome compositions  $Z \in \mathbb{R}^{n \times p}$ . We represent the network as an undirected, weighted graph  $\mathcal{G} = (V, E)$ , where the vertex set  $V = \{v_1, \dots, v_p\}$  represents the  $p$  taxa (e.g., OTUs) and the edge set  $E \subset V \times V$  the possible associations among them. Our formal approach is to learn a probabilistic graphical model [35] (i) that is consistent with the observed data and (ii) for which the (unknown) graph  $\mathcal{G}$  encodes the conditional dependence structure between the random variables (in our case, the observed taxa). Graphical models over undirected graphs (also known as Markov networks or Markov Random Fields) have a straightforward distributional interpretation when the data are drawn from a probability distribution  $\pi(x)$  that

belongs to an exponential family [36, 37]. For example, when the data are drawn from a multivariate normal distribution  $\pi(x) = \mathcal{N}(x|\mu, \Sigma)$  with mean  $\mu$  and covariance  $\Sigma$ , the non-zero elements of the off-diagonal entries of the inverse covariance matrix  $\Theta = \Sigma^{-1}$ , also termed the precision matrix, defines the adjacency matrix of the graph  $\mathcal{G}$  and thus describes the factorization of the normal distribution into conditionally dependent components [35]. Conversely, if and only if an entry in  $\Theta$ :  $\Theta_{i,j} = 0$ , then the two variables are conditionally independent, and there is no edge between  $v_i$  and  $v_j$  in  $\mathcal{G}$ . We seek to estimate the inverse covariance matrix from the data, thereby inferring associations based on conditional independence. This is fundamentally distinct from SparCC and CCREPE (see S1 Table), which essentially estimate pairwise correlations (though other pairwise metrics could be considered for CCREPE). We highlight this key difference in Fig 1. For an intuitive introduction to graphical models in the context of biological networks see Bühlmann *et. al*, 2014 [38].

Inferring the exact underlying graph structure in the presence of a finite amount of samples is, in general, intractable. However, two types of statistical inference procedures have been useful in high-dimensional statistics due to their provable performance guarantees under assumptions about the sample size  $n$ , dimensionality  $p$ , underlying graph properties, and the generating distribution [29, 39]. The first approach, neighborhood selection [20, 39], aims at reconstructing the graph on a node-by-node basis where, for each node, a penalized regression problem is solved. The second approach is the penalized maximum likelihood method [22, 23], where the entire graph is reconstructed by solving a global optimization problem, the so-called covariance selection problem [40]. The key advantages of these approaches are that (i) their underlying inference procedures can be formulated as convex (and hence tractable) optimization problems, and (ii) they are applicable even in the underdetermined regime ( $p > n$ ), provided that certain structural assumptions about the underlying graph hold. One assumption is that the true underlying graph is reasonably sparse, e.g., that the number of taxon-taxon associations scales linearly with the number of measured taxa.

**Graphical model inference.** The SPIEC-EASI pipeline comprises two types of inference schemes, neighborhood and covariance selection. The neighborhood selection framework, introduced by Meinshausen and Bühlmann [20] and thus often referred as the MB method, tackles graph inference by solving  $p$  regularized linear regression problems, leading to local conditional independence structure predictions for each node. Let us denote the  $i^{th}$  column of the data matrix  $Z$  by  $Z^i \in \mathbb{R}^n$  and the remaining columns by  $Z^{-i} \in \mathbb{R}^{n \times p-1}$ . For each node  $v_i$ , we solve the following convex problem:

$$\hat{\beta}^{i,\lambda} = \operatorname{argmin}_{\beta \in \mathbb{R}^{p-1}} \left( \frac{1}{n} \| Z^i - Z^{-i} \beta \|^2 + \lambda \| \beta \|_1 \right), \tag{3}$$

where  $\|a\|_1 = \sum_{i=1}^{p-1} |a_i|$  denotes the L1 norm, and  $\lambda \geq 0$  is a scalar tuning parameter. This so-called LASSO problem [41] aims at balancing the least-square fit and the number of necessary predictors (the non-zero components  $\beta_j$  of  $\beta$ ) by tuning the  $\lambda$  parameter. We define the local neighborhood of a node  $v_i$  as  $N_i^\lambda = \{j \subset \{1, \dots, p\} \setminus i : \hat{\beta}_j^{i,\lambda} \neq 0\}$ . The final edge set  $E$  of  $\mathcal{G}$  can be defined via the intersection or the union operation of the local neighborhoods. An edge  $e_{i,j}$  between node  $v_i$  and  $v_j$  exists if  $j \in N_i^\lambda \cap i \in N_j^\lambda$  or  $j \in N_i^\lambda \cup i \in N_j^\lambda$ . For edges in the set  $j \in N_i^\lambda \cap i \in N_j^\lambda$ , the edge weights,  $e_{i,j}$  and  $e_{j,i}$ , are estimated using the average of the two corresponding  $\beta$  entries. From a theoretical point of view, both edge selection choices are asymptotically consistent under certain technical assumptions [20]. The choice of the  $\lambda$  parameter controls the sparsity of the local neighborhood, which requires tuning [42]. We present our parameter selection strategy at the end of this section.

The second inference approach, (inverse) covariance selection, relies on the following penalized maximum likelihood approach. In the standard Gaussian setting, the related convex optimization problem reads:

$$\hat{\Theta} = \underset{\Theta \in \text{PD}}{\operatorname{argmin}} (-\log \det(\Theta) + \operatorname{tr}(\Theta \hat{\Sigma}) + \lambda \|\Theta\|_1), \quad (4)$$

where PD denotes the set of symmetric positive definite matrices  $\{A: x^T A x > 0, \forall x \in \mathbb{R}^p\}$ ,  $\hat{\Sigma}$  the empirical covariance estimate,  $\|\cdot\|_1$  the element-wise L1 norm, and  $\lambda \geq 0$  a scalar tuning parameter. For  $\lambda = 0$ , the expression is identical to the maximum likelihood estimate of a normal distribution  $\mathcal{N}(x|0, \Sigma)$ . For non-zero  $\lambda$ , the objective function (also referred as the graphical Lasso [22]) encourages sparsity of the underlying precision matrix  $\Theta$ . The non-zero, off-diagonal entries in  $\Theta$  define the adjacency matrix of the interaction graph  $\mathcal{G}$  which, similar to MB, depends on the proper choice of the penalty parameter  $\lambda$ . Originally, this estimator was shown to have theoretical guarantees on consistency and recovery only under normality assumptions [43]. However, recent theoretical [29, 44] work shows that distributional assumptions can be considerably relaxed, and the estimator is applicable to a larger class of problems, including inference on discrete (count) data. In addition, nonparametric approaches, such as sparse additive models, can be used to “gaussianize” the data prior to network inference [45]. We thus propose the following estimator for inferring microbial ecological associations. Given clr-transformed OTU data  $Z \in \mathbb{R}^{n \times p}$ , we propose the modified optimization problem:

$$\hat{\Omega}^{-1} = \underset{\Omega^{-1} \in \text{PD}}{\operatorname{argmin}} (-\log \det(\Omega^{-1}) + \operatorname{tr}(\Omega^{-1} \hat{\Gamma}) + \lambda \|\Omega^{-1}\|_1), \quad (5)$$

where  $\hat{\Gamma}$  is the empirical covariance estimate of  $Z$ , and  $\Omega^{-1}$  is the inverse covariance (or precision matrix) of the underlying (unknown) basis. As stated above,  $\hat{\Gamma}$  will be a good approximation for the basis covariance matrix  $\hat{\Omega}$  because  $p \gg n > 0$ . The resulting solution is constrained to the set of PD matrices, ensuring that the penalized estimator has full rank  $p$ . The non-zero off-diagonal entries of the estimated matrix  $\Omega^{-1}$  define the inferred network  $\mathcal{G}$ , and their values are the signed edge weights of the graph. To reduce the variance of the estimate, the covariance matrix  $\hat{\Gamma}$  can also be replaced by the empirical correlation matrix  $\hat{R} = D \hat{\Gamma} D$ , where  $D$  is a diagonal matrix that contains the inverse of the estimated element-wise standard deviations.

The covariance selection approach has two advantages over the neighborhood selection framework. First, we obtain unique weights associated with each edge in the network. No averaging or subsequent edge selection is necessary. Second, the covariance selection framework provides invertible precision and covariance matrix estimates that can be used in further downstream microbiome analysis tasks, such as regression and discriminant analysis [10].

**Model selection.** For both neighborhood and covariance selection, the tuning parameter  $\lambda \in [0, \lambda_{\max}]$  controls the sparsity of the final model. Rather than inferring a single graphical model, both methods produce a  $\lambda$ -dependent solution path with the complete and the empty graph as extreme networks. A number of model selection criteria, such as Bayesian Information Criteria [46] and resampling schemes [47], have been used. Here we use a popular model selection scheme known as the Stability Approach to Regularization Selection (StARS) [48]. This method repeatedly takes random subsamples (80% in the standard setting) of the data and estimates the entire graph solution path based on this subsample. For each subsample, the  $\lambda$ -dependent incidence frequencies of individual edges are retained, and a measure of overall edge stability is calculated. StARS selects the  $\lambda$  value at which subsampled non-empty graphs are the least variable (most stable) in terms of edge incidences. For the selected graph, the

observed edge frequencies indicate the reproducibility, and likely the predictive power, and are used to rank edges according to confidence.

**Theoretical and computational aspects.** Learning microbial graphical models with neighborhood or inverse covariance selection schemes has important theoretical and practical advantages over current methods. A wealth of theoretical results are available that characterize conditions for asymptotic and finite sample guarantees for the estimated networks [20, 29, 39, 43, 46]. Under certain model assumptions, the number of samples  $n$  necessary to infer the true topology of the graph in the neighborhood selection framework is known to scale as  $n = O(d^3 \log(p))$ , where  $d$  is the maximum vertex (or node) degree of the underlying graph (i.e. the maximum size of any local neighborhood). Additional assumptions on the sample covariance matrices reduce the scaling to  $n = O(d^2 \log(p))$  [39]. This implies that graph recovery and precision matrix estimation is indeed possible even in the  $p \gg n$  regime, and that the underlying graph topology strongly impacts edge recovery. The latter observation means that, even if the number of interactions  $e$  is constant, graphs with large hub nodes, perhaps representing keystone species in microbial networks, or, more generally, scale-free graphs with a few highly connected nodes, will be more difficult to recover than networks with evenly distributed neighborhoods. In addition to these theoretical results, a second advantage is that well-established, efficient, and scalable implementations are available to infer microbial ecological networks from OTU data in practice. Thus, SPIEC-EASI methods will efficiently scale as microbiome dataset dimensions grow (e.g., due to technological advances that increase the number of OTUs detected per sample). The SPIEC-EASI inference engine relies on the R package *huge* [49], which includes algorithms to solve neighborhood and covariance selection problems [20, 22], as well as the StARS model selection.

## Generation of synthetic microbial abundance datasets

Estimating the absolute and comparative performance of network inference schemes from biological data remains a fundamental challenge in biology. In the context of gene regulatory network inference, recent community-wide efforts, such as the DREAM (Dialogue for Reverse Engineering Assessments and Methods) Challenges (<http://www.the-dream-project.org/>), have considerably advanced our understanding about feasibility, accuracy, and applicability of a large number of developed methods. In the DREAM challenges, both real data from “gold standard” regulatory networks (e.g., networks where the true topology is known from independent experimental evidence) and realistic in-silico data (using, e.g., the GeneNetWeaver pipeline [50]) are included. In the context of microbiome data and microbial ecological networks, neither a gold standard nor a realistic synthetic data generator exist. SPIEC-EASI is accompanied by a set of computational tools that allow the generation of realistic synthetic OTU data. As outlined in Fig 2, real taxa count data serve as input to SPIEC-EASI’s synthetic data generation pipeline. The pipeline enables one to: (i) fit the marginal distributions of the count data to a parametric statistical model and (ii) specify the underlying graphical model architecture (e.g., scale-free).

**The NorTA approach.** The parametric statistical model and network topology are then combined in the ‘Normal To Anything’ (NorTA) [51] approach to generate synthetic OTU data that resemble real measurements from microbial communities but with user-specified network topologies. NorTA [51] is an approximate technique to generate arbitrary continuous and discrete multivariate distributions, given (1) a target correlation structure  $R$  with entries  $\rho_{i,j}$  and (2) a target univariate marginal distribution  $U_i$ . To achieve this task, NorTA relies on normal-copula functions [51–53]. A  $n \times p$  matrix of data  $U$  is sampled from a normal

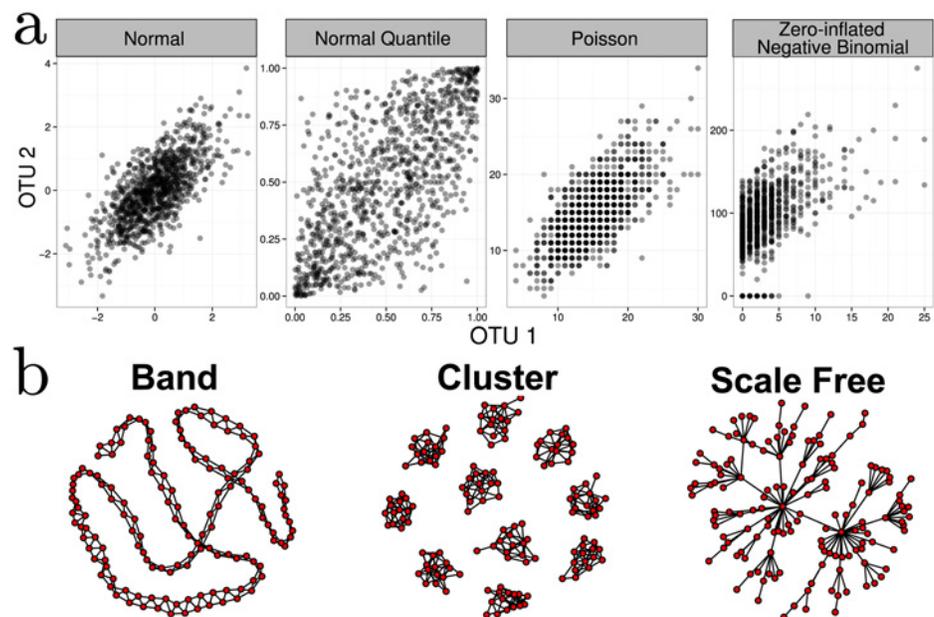
distribution with zero mean and a  $p \times p$  correlation matrix  $R_N$ . For each marginal  $U_i$ , the Normal cumulative distribution function (CDF) is transformed to the target distribution via its inverse CDF. For any target distribution  $P$  with CDF  $\Xi$ , we can thus generate multivariate correlated data via

$$U_{P_i} = \Xi^{-1}(\Phi(U_{N_i})), \tag{6}$$

where  $U_N \sim \mathcal{N}(0, R_N)$  and  $\Phi(U) = \int_{-\infty}^U \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{u^2}{2\sigma^2}} du$ , the CDF of a univariate normal. In Fig 3a, we illustrate this process for bivariate Poisson and negative binomial data ( $n = 1000$  and correlation  $\rho_{ij} = 0.7$ ).

In the original NorTA approach, an element-wise monotone transformation  $c_U(\cdot)$  with  $R_N = c_U(R)$  is applied to account for slight differences in correlation structure between normal and target distribution samples [51]. Here we neglect this transformation step because we observe that the log-transformed data from exponential count distributions, such as the Poisson and Negative binomial, are already close to  $R$ , provided that the mean is greater than one, particularly when the counts data are log-transformed (S1 Fig). In practice, SPIEC-EASI relies on routines from base R and VGAM packages [54, 55] to estimate the uniform quantiles of the normal data and to fit the desired CDF with estimated parameters.

**Fitting marginal distribution to real OTU data.** Prior to fitting marginal distributions to real data, several commonly used pre-processing steps are applied. For any given OTU abundance table of size  $n \times K$ , we first select  $p$  non-zero columns. To account for experimental



**Fig 3. a) Bivariate illustration of the NorTA approach.** First normal data, incorporating the target correlation structure, is generated. Uniform data are then generated for each margin via the normal density function. These is then converted to an arbitrary marginal distribution (Poisson and Zero-inflated Negative Binomial shown as examples) via its quantile function. To generate realistic synthetic data, parameters for these margins are fit to real data. **b) Examples of band-like, cluster, and scale-free network topologies**

doi:10.1371/journal.pcbi.1004226.g003

differences in sample sequencing, we then normalize samples to a median sequencing depth by multiplying all counts by the ratio of minimum desirable sampling depth to the total sum of counts in that sample and rounding to the nearest whole number, which is preferable to rarefaction [56]. These filtered and sequencing-depth-normalized data serve as the marginal counts, which are fit to a parametric distribution  $U_i$  and used as input to the NorTA approach. The concrete target marginal distribution depends on the actual microbiome dataset. For gut microbiome data (e.g. from HMP or APG), the zero-inflated Negative Binomial (ziNB) distribution is a good choice, as it accounts for both overdispersion [56, 57] and the preponderance of zero-count data points in microbial count datasets. The fitting procedure is done within a maximum likelihood framework. The corresponding optimization problem is solved with the Quasi-Newton methods with box-constraints, as implemented in the `optim` function in R [54]. In [S1 Text](#), we use quantile-quantile plots to compare ziNB to several other candidate distributions (e.g., lognormal, Poisson, NB) and show that ziNB has superior fit.

**Generation of network topologies and correlation matrices.** Under normality assumption, the non-zero pattern of the precision matrix corresponds to the adjacency matrix of the underlying undirected graph. We use this property to generate target covariance (correlation) matrices originating from different graph topologies. The pipeline to generate a network structure for simulated data proceeds in three steps: (i) Generate an undirected graph, in the form of an adjacency matrix, with a desired topology and sparsity, (ii) convert the adjacency matrix to a positive-definite precision matrix by assigning positive and negative edge weights and appropriate diagonal entries, and (iii) invert  $\Theta$  and convert the resulting covariance matrix  $\Sigma$  to a correlation matrix ( $R = D\Sigma D$ , where,  $D$  is a diagonal matrix with diagonal entries  $1/\sqrt{\sigma_i}$ ).

Among many potential graph structures, we focus on three representative network structures: band-like, cluster, and scale-free graphs (see [Fig 3b](#) for graphical examples). Maximum network degree strongly impacts network recovery, and thus our choice of network topologies spans a range of maximum degrees (band < cluster < scale-free). In addition, cluster and scale-free lend themselves to hypothetical ecological scenarios. Cluster graphs may be seen as archetypal models for microbial communities that populate different disjoint niches (clusters) and have only few associations across niches. Scale-free graphs, ubiquitous in many other facets of network biology (such as gene regulatory, protein-protein and social networks), serve as a baseline model for a microbial community that comprises (1) a few “keystone” species (hub nodes with many partners) that are essential for coordinating/stabilizing the community and (2) many dependent species that are sparsely connected to each other. The sparsity of the networks is controlled by the number of edges,  $e < p(p - 1)/2$ , in the graph. The topologies are generated according to the following algorithms, starting with an empty  $p \times p$  adjacency matrix:

1. **Band:** A band-type network consists of a chain of nodes that connect only their nearest neighbors. Let  $e = e_{used} + e_{available}$ , the number of edges already used and available, respectively. Fill the next available off-diagonal vector with edges if and only if  $e_{available} \geq$  number of elements in the off-diagonal.
2. **Cluster:** A cluster network comprises  $h$  independent groups of randomly connected nodes. For given  $p$  and  $e$  we divide the set of nodes into  $h$  components of (approximately) identical size and set the number of edges in each component to  $e_{comp} = e/h$ . For each component, we generate a random (Erdős-Renyi) graph for which we randomly assign an edge between two nodes in the cluster with probability  $p = \frac{e_{comp}}{h(h-1)/2}$ .
3. **Scale-free:** The distribution of degrees, the number of edges per node, in a scale-free network is described by a power law, implying that the central node or nodes (potentially

keystone species in an ecological network) have proportionally more connections. We use the standard preferential attachment scheme [58] until  $p - 1$  edges are exhausted.

After generation of these standard adjacency matrices, we randomly remove or add edges until the adjacency matrix has exactly  $e$  edges. All schemes generate symmetric adjacency matrices that describes a graph, with entries of 1 if an edge exists and 0 otherwise.

From the adjacency matrices, we generate precision matrices by uniformly sampling non-zero entries  $\Theta_{i,j} \in [-\Theta_{\max}, -\Theta_{\min}] \cup [\Theta_{\min}, \Theta_{\max}]$ , where  $\Theta_{\min}, \Theta_{\max} > 0$  are model parameters and describe the strength of the conditional dependence among the nodes. To ensure that the precision matrix is positive definite with tunable condition number  $\kappa = \text{cond}(\Theta)$ , we scale the diagonal entries  $\Theta_{i,i}$  by a constant  $c$  using binary search. The precision matrix  $\Theta$  is then converted to a correlation matrix  $R$  to be used as input to the NorTA approach.

## Results

### Network inference on synthetic microbiome data

Given that no large-scale experimentally validated microbial ecological network exists, we use SPIEC-EASI's data generator capabilities to synthesize data whose OTU count distributions faithfully resemble microbiome count data. By varying parameters known to influence network recovery (network topology, association strength, sample number) and quantifying performance on resulting networks, we rigorously assess SPIEC-EASI inference relative to state-of-the-art inference methods, SparCC [17] and CCREPE [9], as well as standard Pearson correlation.

**Benchmark setup.** We modeled the synthetic datasets on American Gut Project data using SPIEC-EASI's data generation module. The count data, accessed February, 2014 at [www.microbio.me/qiime](http://www.microbio.me/qiime) and available at <https://github.com/zdk123/SpiecEasi/tree/master/inst/extdata>, come from two sampling rounds and comprise several thousand OTUs. Round 1 data contains  $n_1 = 304$ , and Round 2 data contains  $n_2 = 254$  samples. As filtering steps, OTUs were removed from the input data if present in fewer than 37% of the samples, while samples were removed if total sequencing depth fell below the 1st quartile (10,800 sequence reads). Thus, we arrived at a total of  $p = 205$  distinct OTUs. We also generated smaller-dimensional datasets ( $p = 68$ ) with fewer zero counts by requiring that OTUs be present in  $> 60\%$  of the samples. We used Round 1 data and fit the  $n_1$  count histograms to a ziNB distribution (for justification of this, see [S1 Text](#)). The empirical effective number  $n_{\text{eff}}$  is 13.5 for  $p = 205$  and 7.5 for  $p = 68$  data. The resulting parametrized marginal distributions served as input to NorTA.

As described above, network topology is expected to influence network recovery; thus, we consider the three previously described topologies (band-like, cluster, and scale-free) as representative microbial networks. We hypothesize that any method that successfully infers the sets of associations underlying these archetypal networks from synthetic datasets is likely to also perform well in the context of true microbiomes, whose underlying network architecture is unknown but expected to be a mixture of these network types. For all networks, we fix the total number of edges  $e$  to the respective number of OTUs  $p$ , and we analyze a medium ( $p = 68$ ) and a high-dimensional scenario ( $p = 205$ ). Microbial association strength is controlled by the range of values in off-diagonal entries in the precision matrices  $\Theta$  and the condition number  $\kappa = \text{cond}(\Theta)$ . We use  $\Theta_{\min} = 2$  and  $\Theta_{\max} = 3$  with either condition number  $\kappa = 10$  or 100. In this setting,  $\kappa$  controls the spread of the absolute correlation values (and thus the strength of indirect associations) present in the synthetic data. The relationship between condition number and distribution of correlation is illustrated in [S3 Fig](#). For each network type and size, we generate 20 distinct instances. For each instance, we then use the NorTA approach to generate a

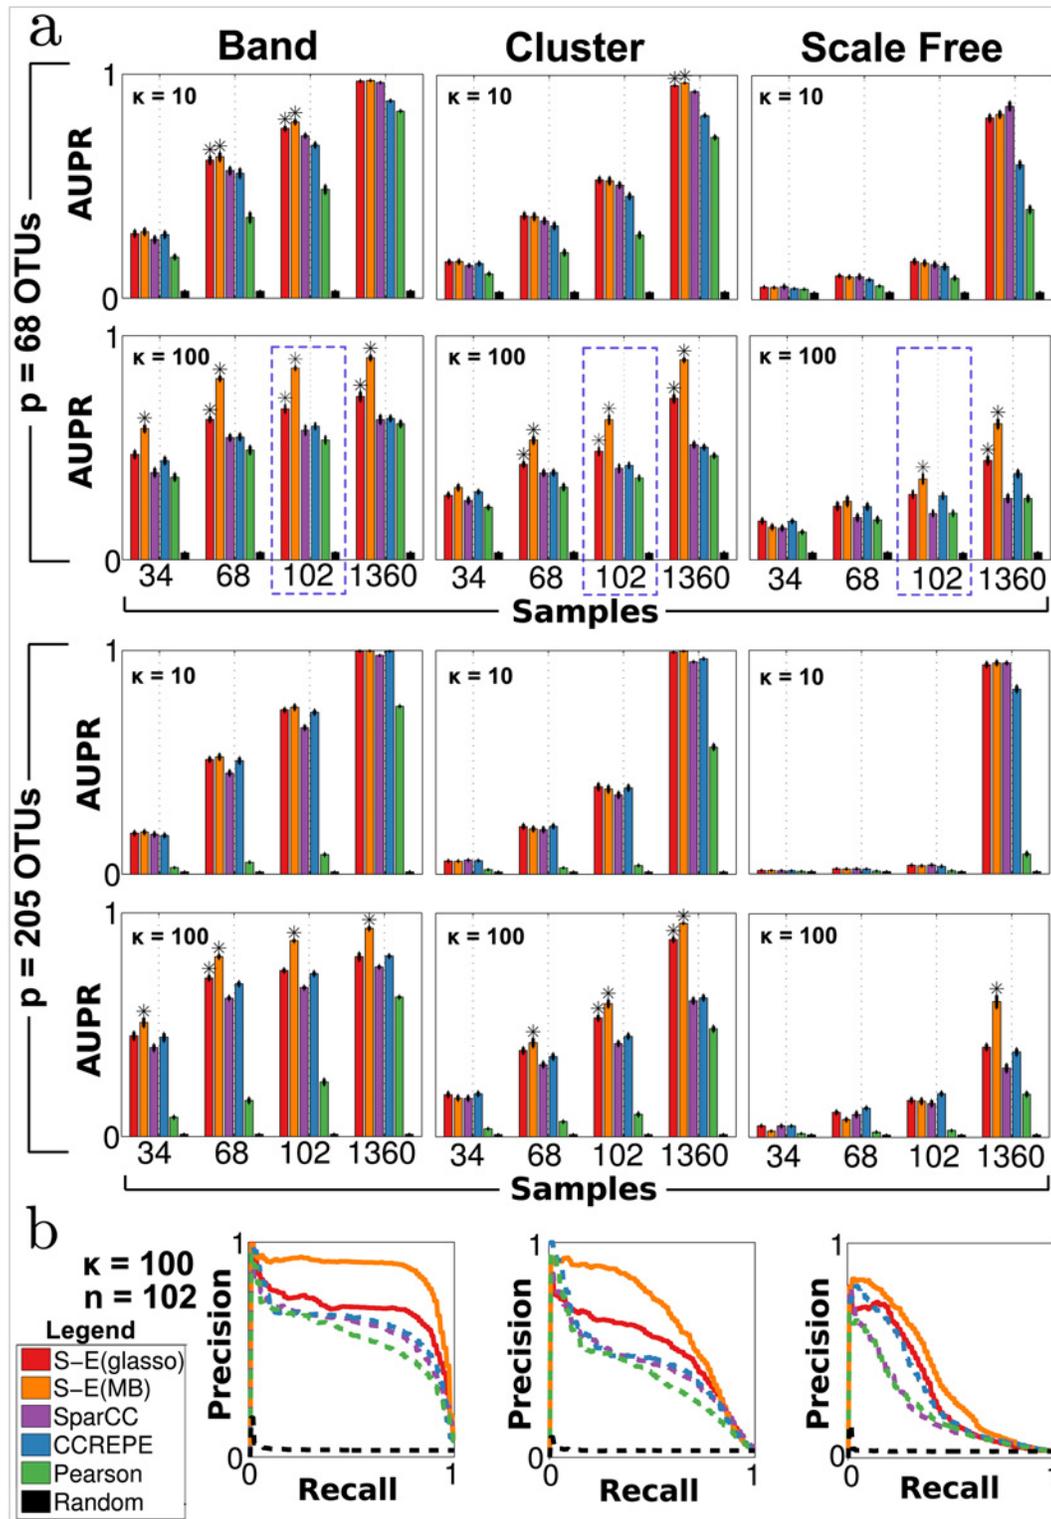
maximum of  $n = 1360$  synthetic microbial count data samples. In [S2 Fig](#), we highlight the fidelity between the target microbiome dataset and the synthetic datasets, especially in terms of dataset dimensions and OTU count distributions. To assess the effect of sample size on network recovery, we test methods on a range of sample sizes:  $n = 34, 68, 102, 1360$ .

We compare SPIEC-EASI's covariance selection method (referred to as S-E(glasso)) and neighborhood selection method (referred to as S-E(MB)) to SparCC and CCREPE, methods which were also designed to be robust to compositional artifacts. As a baseline reference, we also compared all methods to Pearson correlation, which is neither compositionally robust nor appropriate for estimating correlation in the under-determined regime. Both of these methods, however, infer interactions from correlations and do not consider the concept of conditional independence. We improved the runtime of the original SparCC implementation (available at <https://bitbucket.org/yonatanf/sparcc>) and include the updated code in our SPIEC-EASI package. The original SparCC package also includes a small benchmark test case available at <https://bitbucket.org/yonatanf/sparcc/src/9a1142c179f7/example>). The recovery performance results for S-E(MB), SparCC, and CCREPE for this test case can be found in [S9 Fig](#). The CCREPE implementation is downloaded from <http://www.bioconductor.org/packages/release/bioc/html/ccrepe.html>.

**Recovery of microbial networks.** To quantify each method's ability to recover the true underlying association network, we evaluated performance in terms of precision-recall (P-R) curves and area under P-R curves (AUPR). For each method, we ranked edge predictions according to confidence. For SparCC, CCREPE and Pearson correlation, edge predictions were ranked according to p-value. SPIEC-EASI edge predictions were ranked according to edge stability, inferred by StARS model selection step at the most stable tuning parameter  $\lambda_{\text{StARS}}$ . [Fig 4](#) summarizes methods' performance on 960 independent synthetic datasets for a total of 48 conditions (4 samples sizes  $\times$  2 conditions numbers  $\times$  3 network topologies  $\times$  2 dimensions).

We observe the following key trends. First, the performance of all methods improves with increasing sample size. Under certain scenarios, even near-perfect recovery ( $\text{AUPR} \approx 1$ ) is possible in the large sample limit ( $n = 1360$ ). Second, all methods show a clear dependence on the network topology. Best performance is achieved for band graphs, followed by cluster and scale-free graphs. These results are consistent with theoretical results [29], which show that the maximum node degree  $d$  reduces the probability of correctly inferring network edges for fixed sample size (scale-free networks have highest maximum degree, followed by cluster and then band.) Third, for most scenarios, the SPIEC-EASI methods, particularly S-E(MB), perform as well or significantly better than all control methods. Standard Pearson correlation is outperformed by all methods that take the compositional nature of the data into account. Forth, in the large sample limit ( $n = 1360$ ), S-E(MB) is the only method that recovers a significant portion of edges under all tested scenarios (particularly scale-free networks). Additionally, we observe ([S10 Fig](#)) that SPIEC-EASI performs well on synthetic data generated by the SparCC benchmark [17, 59].

The present results suggest that complete network recovery is likely an unrealistic goal for microbiome studies, given that most studies have at most hundreds of samples. In addition, the P-R curves are based on ranking predicted edges. To generate a final network, confidence-based criteria must be applied to select a final set of edges for network inclusion, and, to date, no optimal selection process exists. Nonetheless, if we focus on the set of high-confidence interactions (i.e. the top-ranked entries in the edge list), we see that S-E methods, particularly S-E (MB), can achieve very high precision for all network types (see [Fig 4](#) for a representative P-R curve).



**Fig 4. Precision-recall performance on synthetic datasets.** a) Red = S-E(glasso), orange = S-E(MB), purple = SparCC, blue = CCREPE, green = Pearson correlation, black = random. Area under precision-recall (AUPR) vs. number of samples  $n$  for different  $\kappa$  values are depicted. Bars represent average over 20 synthetic datasets, and error bars represent standard error. Asterisks denote conditions under which SPIEC-EASI methods had significantly higher AUPR relative to all other control methods ( $P < 0.05$  for all one-sided T tests). b) Representative precision-recall curves for  $p = 68$ ,  $n = 102$ ,  $\kappa = 100$ ; solid and dashed lines denote SPIEC-EASI and control methods, respectively.

doi:10.1371/journal.pcbi.1004226.g004

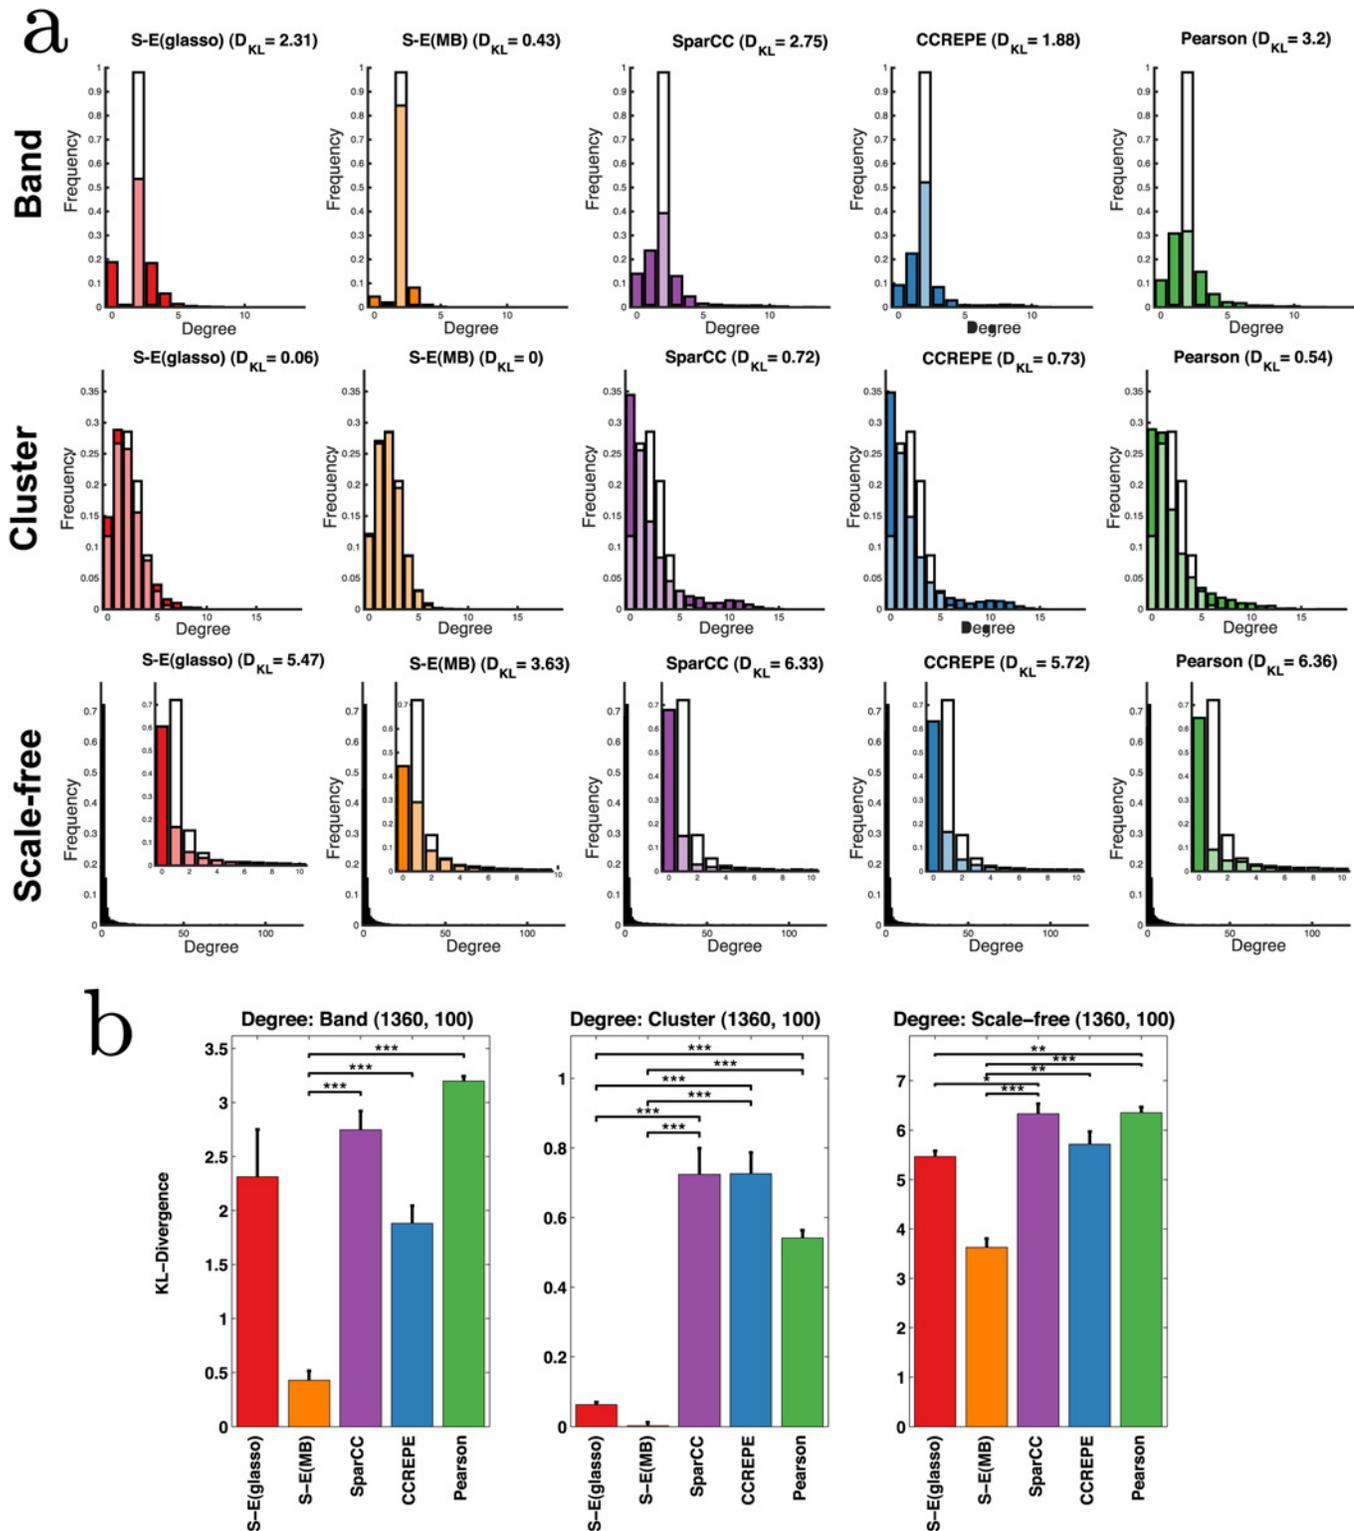
Overall, these results suggest that S-E methods outperform current state-of-the-art methods in terms of network recovery under most tested scenarios, with S-E(MB) showing superior performance over S-E(glasso).

**Recovery of global network properties.** Accurate recovery of global network properties (e.g., degree distribution, number of connected components, shortest path lengths) from taxa abundance data would help define the underlying topology of the ecological network (e.g., cluster versus scale-free) [15, 58]. At this point, little is known about the underlying topology of microbial ecological networks, but, as elaborated in the Discussion, such information could be incorporated as a constraint into SPIEC-EASI's inference methods, thereby further improving prediction. Thus, we tested how well SPIEC-EASI and other methods recover global network properties from the synthetic datasets and evaluate whether these methods might be able to provide insight into global network architecture, perhaps even in the underdetermined regime, where the prediction of individual edges is less accurate. To control for the disparate means by which individual methods ranked edge confidences (e.g., stability for the SPIEC-EASI methods, estimation of p-values for SparCC, Pearson and CREPE), for each synthetic dataset and method, final networks were generated by selecting the top 205 predicted edges and comparing to the true synthetic network topologies for  $p = 205$  OTUs.

We first consider (node) degree distributions, where node degree is defined as the number of edges each node has. In Fig 5, we show the empirical degree distribution and the underlying ground truth for all methods and networks types,  $n = 1360$  and  $\kappa = 100$ . Scale-free networks are characterized by exponential degree distributions, in which few nodes (e.g., hubs and, in our context, potential keystone taxa) have very high degree (e.g., interact with other taxa), while most nodes/taxa have few interactions. In contrast, nodes in cluster networks have relatively even degree, which depends on cluster size. In the ecological context, cluster networks would be consistent with niche communities that share few interactions with microbiota outside of one's niche community; this structure is also reflected in degree distributions. Using Kullback-Leibler (KL) divergence to measure the dissimilarity between methods' predicted degree distributions and the true degree distribution we see that S-E(MB) outperforms all other methods in recovering degree distributions (Fig 5). This performance improvement also holds for smaller samples sizes.

Another common topological feature is betweenness centrality, which, similar to degree, betweenness centrality can be used to gauge the relative importance of a node (e.g., taxon) to the (ecological) network. Betweenness centrality, as the fraction of shortest paths between all other nodes in the network that contain the given node, highlights central nodes. The distribution of nodes' betweenness centrality provides information about the network architecture (S4 Fig). Specifically, scale-free networks are expected to have a few nodes with very high betweenness centrality that connect most other nodes to each other; in scale-free networks, betweenness centrality can approach unity. For cluster networks, the maximum betweenness centrality is limited by the total number of independent clusters. In band networks, similar to scale-free, all nodes are connected; however, the degree is fixed and so the betweenness centrality distribution is roughly uniform from zero to one. For smaller sample sizes  $n < 1360$ , no method dominates. However, for the largest sample size,  $n = 1360$ , S-E(MB) is again significantly better than all other methods for five out of six conditions with the exception of scale-free networks  $\kappa = 10$ , where SparCC recovery is best (S5 Fig).

We next consider distributions over graph geodesic distances. The geodesic distance is the length of the shortest path between two nodes. Given the existence of highly connected hubs in scale-free networks, geodesic distances for scale-free networks tend to be short, a feature that is described as the "small-world" property. Thus, the geodesic distributions in the scale-free network are a lot smaller than for the band and cluster networks (S6 Fig). In recovery of geodesic



**Fig 5. a)** Predicted degree distributions (colored) are overlaid with the true degree distribution (white) for  $n = 1360$  samples,  $p = 205$  OTUs,  $\kappa = 100$ . Lighter shades correspond to regions of overlap between predicted and true distributions. Dissimilarity between the distributions is measured by KL divergence,  $D_{KL}$ . **b)** Bars represent the average  $D_{KL}$  over three independent sets of synthetic datasets (7 datasets per set); error bars represent standard error. Divergences were compared between S-E and control methods using one-sided T-tests; \*\*\*, \*\*, \* correspond to  $P < 0.001$ ,  $0.01$ , and  $0.05$ .

doi:10.1371/journal.pcbi.1004226.g005

distance distributions, S-E(MB) performs equivalently or significantly better than all other methods for scale-free networks as well as band graphs across all conditions. For cluster networks, the other methods generally outperform SPIEC-EASI methods for smaller sample sizes ( $n < 1360$ ). In the large sample limit  $n = 1360$ , S-E(MB) has significantly better recovery of geodesic distance distributions relative to all control methods, even for cluster graphs (S7 Fig).

Finally, we analyzed the number and size of connected components in the inferred graphs. While all synthetic band and scale-free synthetic networks form a single connected component containing all nodes, cluster networks have a variable number of connected components. In terms of cluster number recovery, all methods predicted too many connected components. Overall, S-E methods had lower error rates for band and scale-free networks over all sample sizes. For high sample number ( $n = 1360$ ), S-E(MB) had significantly better recover of cluster size across all network types (S8 Fig), with nearly perfect recovery for cluster graphs ( $D_{KL} = 0$ , S9 Fig).

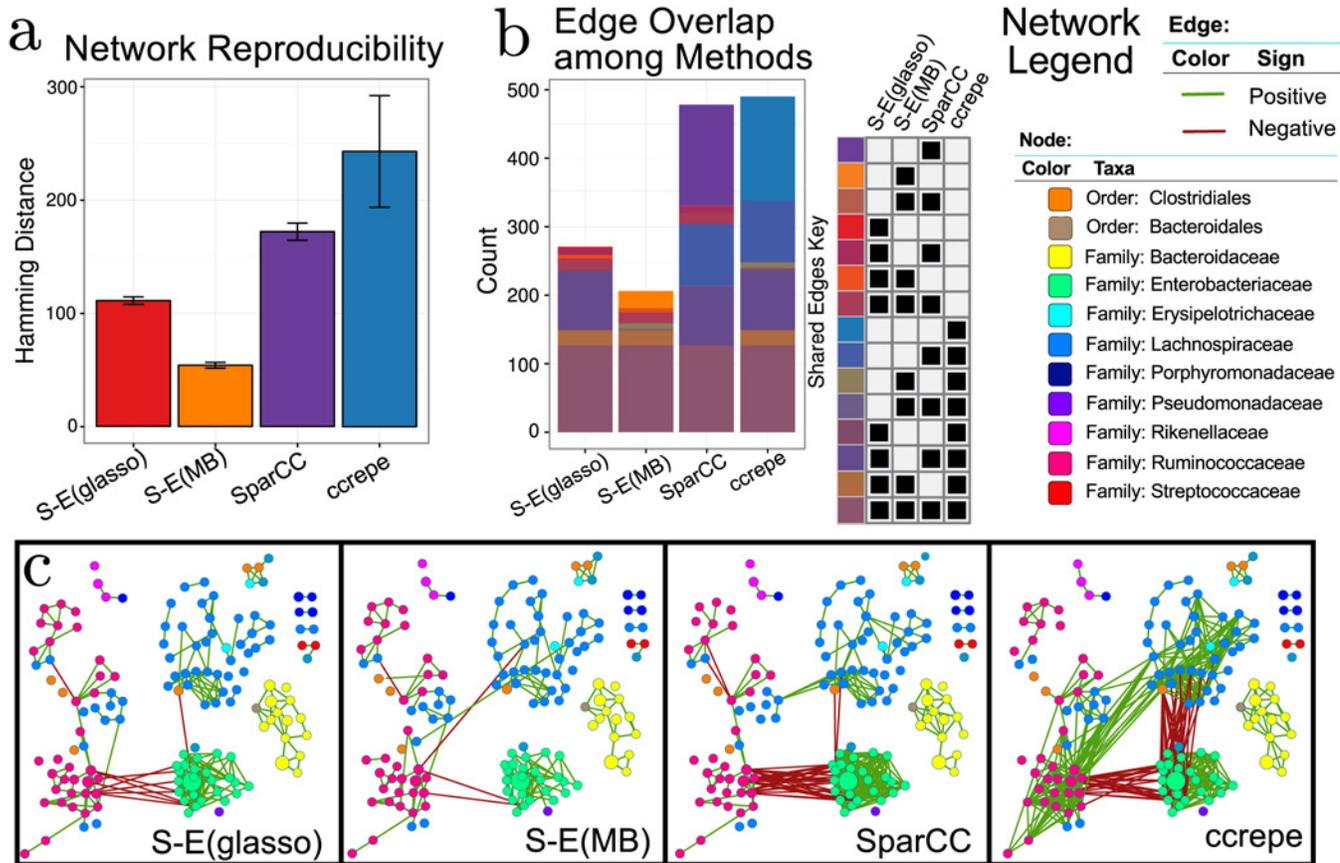
## Inference of the American Gut network

Thus far we have used the  $n_1 = 304$  first-round AGP samples as a means to construct realistic synthetic microbiome data sets with SPIEC-EASI's data generation module. In this section, we apply SPIEC-EASI inference methods to construct ecological association networks from the AGP data directly. To do this, we first filter out rare OTUs by selecting only the top 205 OTUs (to match the dimensionality of the synthetic data) in the combined AGP dataset (by frequency of presence) before adding a pseudo-count and total-sum normalization. Although there is no independent means to assess the accuracy of these hypothetical networks, we can assess their reproducibility and consistency. For each method, we first infer a single representative network of taxon-taxon interactions from Round 1 AGP abundance data. For SPIEC-EASI, the StARS model selection approach is used to select the final model network. For SparCC, we use a threshold  $\rho_t = 0.35$  to construct a relevance network from the SparCC-inferred correlation matrix; i.e. an edge between nodes  $v_i, v_j$  is present in the SparCC network if  $|\rho_{i,j}| > \rho_t$  [17]. Similarly, we use a q-value cut-off of  $10^{-24}$  to create an interaction network from CCREPE-corrected significance scores of Pearson's correlation coefficient [9]. For each method, we thus arrive at a reference network that can be considered the hypothetical gold standard. We then use the  $n_2 = 254$  Round 2 AGP samples as an independent test set and learn a new model network from these data alone. We measure consistency between the two network models by computing the Hamming distance between the reference and new network models, i.e., the difference between the upper triangular part of the two adjacency matrices. For the present data, the Hamming distance can vary between  $p(p-1)/2 = 20910$  (no edges in common) and a minimum of 0 for identical networks. Confidence intervals for Hamming distances can be obtained by combining Round 1 and 2 samples into a unified dataset, repeatedly subsampling these data into two disjoint groups of size  $n_1$  and  $n_2$ , and repeating the entire inference procedure.

Fig 6a shows network reproducibility for SPIEC-EASI methods, SparCC, and CCREPE. The S-E(MB) has smallest the Hamming distance, followed by S-E(glasso), SparCC, and CCREPE. In S-E(MB), the edge disagreement is roughly 50 with very small error bars. At the other extreme, CCREPE edge disagreement is 250 edges and highly variable.

These numerical experiments clearly demonstrate that SPIEC-EASI networks are more reproducible than other current methods.

Finally, we use each inference method to construct a candidate American Gut microbiome association network from the unified dataset of size  $n_1 + n_2 = 558$  (Fig 6c). We analyze the differences between the reconstructed networks by quantifying the number of unique and shared predicted edges (Fig 6b). All four methods agree on a core network that consists of 127 edges.



**Fig 6. a) Network reproducibility for inference methods (see main text for details).** Bars represent mean Hamming distance, and errorbars are 95% confidence intervals. **b) Visualization of edge overlap between networks inferred with SPIEC-EASI, SparCC, and CCREPE.** **c) Network visualizations with OTU nodes colored by Family lineage (or Order, when the Family of the OTU is unknown), edges are colored by sign (positive: green, negative: red), and the node diameter proportional to the geometric mean of that OTU's relative abundance.**

doi:10.1371/journal.pcbi.1004226.g006

These edges are mostly found within OTUs of the same taxonomic group. This phenomenon, termed assortativity, has also been observed in other microbial network studies [11]. Assortativity is one of the most salient features of the AGP-derived networks, and, for all networks, the assortativity coefficients for each network are close to unity (e.g., maximum assortativity, S11 Fig). The SparCC network comprises about twice as many edges as the SPIEC-EASI networks. SparCC infers 147 distinct edges; these additional edges correspond to negative associations between OTUs of Ruminococcaceae (genus *Faecalibacterium*) and Enterobacteriaceae families (various genera) and a dense web of correlations within Enterobacteriaceae OTUs. Similarly, CCREPE identified 152 edges uniquely, with many negative edges between Enterobacteriaceae and Lachnospiraceae (genera: *Blautia*, *Roseburia* and unknown); additionally, CCREPE uniquely predicted positive edges between the Lachnospiraceae and Ruminococcaceae (genus: *Faecalibacterium*). Both SPIEC-EASI methods produce relatively sparse networks by comparison. S-E(glasso) infers a total of 271 total edges (with one unique edge), and S-E(MB) infers 206 edges with 25 unique edges. In scale with edge predictions, both CCREPE and SparCC infer networks with large maximum degree (33 and 30, respectively), while the S-E(MB) and S-E(glasso) networks have a maximum degree of sixteen and eight, respectively (S11 Fig). However, even though CCREPE and SparCC predict a similar number of total edges, the global

network properties are distinct. CCREPE predicts a higher maximum betweenness centrality and a larger number of nodes in the largest connected component (100).

In summary, analysis of the AGP networks suggests that the SPIEC-EASI inference schemes construct more reproducible taxon-taxon interactions than SparCC and CCREPE and infer considerably sparser model networks than the other two methods. These observations may be explained as follows: SparCC and CCREPE aim to recover correlation networks, which contain both direct edges as well as indirect (e.g., spurious) edges (due to correlation alone). SparCC and CCREPE may recover indirect edges less robustly than direct edges, an explanation that would be consistent with the Hamming distance reproducibility analysis. In addition, all methods' resulting networks suggest that the topology of the American Gut association network cannot be attributed to a specific network class. Instead, these networks are a mixture of band, scale-free, and cluster network type, and they exhibit high phylogenetic assortativity within highly connected components.

## Discussion

Inferring interactions among different microbial species within a community and understanding their influence on the environment is of central importance in ecology and medicine [19, 60]. An ever increasing number of recent amplicon-based sequencing studies have uncovered strong correlations between microbial community composition and environment in diverse and highly relevant domains of life [1, 9, 10, 61–63]. These studies alone underscore the need to understand how the microbial communities adapt, develop, and interact with the environment [5]. Elucidation of species interactions in microbial communities across different environments remains, however, a formidable challenge. Foremost, available high-throughput experimental data are compositional in nature, overdispersed, and usually underdetermined with respect to statistical inference. In addition, for most microbes few to no ecological interactions are known, thus the ecological interaction network must be constructed *de novo*, in the absence of guiding assumptions and a set of “gold standard” interactions for validation.

To overcome both challenges, we present SPIEC-EASI (Sparse InversE Covariance Estimation for Ecological Association Inference), a computational framework that includes statistical methods for the inference of microbial ecological interactions from 16S rRNA gene sequencing datasets and a sophisticated synthetic microbiome data generator with controllable underlying species interaction topology. SPIEC-EASI's inference engine includes two well-known graphical model estimators, neighborhood selection [20] and sparse inverse covariance selection [22, 23, 46] that are extended by compositionally robust data transformations for application to the specific context of microbial abundance data.

The synthetic data pipeline was used to generate realistic-looking gut microbiome datasets for a controlled benchmark of SPIEC-EASI's inference performance relative to two state-of-the-art methods, SparCC [17] and CCREPE [9]. We showed that neighborhood selection (S-E (MB)) outperforms SparCC and CCREPE in terms of recovery of taxon-taxon interactions and global network topology features under almost all tested benchmark scenarios, while covariance selection (S-E(glasso)) performs competitively with and sometimes better than SparCC and CCREPE.

Through our simulation study, we demonstrate that several other factors, in addition to total number of samples, affect network recovery. Foremost and in agreement with theoretical results from high-dimensional statistics [29, 30, 39], network topology has a significant impact, as network recovery performance is nearly doubled from scale-free to cluster to band (Fig 4) for fixed sample size, number of taxa, and condition number. We also demonstrated dependence to strength of direct interactions (and thus strength of correlations) within a given

network. Our simulation study provides the community with rough guidelines for requisite sample sizes, given state-of-the-art network inference and basic assumptions about the underlying network. This is of obvious importance to experimental design and the estimation of statistical power. Here, we used the synthetic data pipeline to generate datasets characteristic of the gut microbiome. However, the SPIEC-EASI data generator is generic and therefore enables researchers to generate synthetic datasets that resemble microbiome samples in terms of taxa dispersion and marginal distributions from their field of research, such as soil or sea water ecosystems [1].

Our application study on real American Gut Project (AGP) data revealed that inference with SPIEC-EASI produced more consistent and sparser interaction networks than SparCC and CCREPE. In addition, our AGP network analysis revealed several biologically relevant observations. Specifically, we observed that OTUs were more likely to interact with phylogenetically related OTUs (Fig 6c and S11 Fig). In addition, our gut microbial interaction networks appear to be a composite of network types, as we find evidence for scale-free, band-like, and cluster subnetworks.

An important advantage of neighborhood and covariance selection as underlying inference frameworks is their ability to include prior knowledge about the underlying data or network structure from independent scientific studies in a principled manner. For example, in the neighborhood selection scheme, the standard LASSO approach can be augmented by a group penalty [64] that takes into account *a priori* known group structure. The assortativity observed in our gut microbial interaction networks suggests that such a grouping of OTUs based on phylogenetic relationship might improve inference. Moreover, if verified species interactions are available for a certain microbial contexts, this knowledge can be included in covariance and neighborhood selection by relaxing the penalty term on these interactions. This strategy has already been fruitfully applied to inference of similarly high-dimensional transcriptional regulatory networks [65]. Finally, in agreement with theoretical and empirical work in high-dimensional statistics, our synthetic benchmark results confirmed that networks with scale-free structures elude accurate inference even if the underlying network is globally sparse. Recent modified neighborhood [30] and covariance selection [31] schemes improve recovery of scale-free networks and can be conveniently included into SPIEC-EASI.

Finally, although the main focus of this work is inference of microbial interaction networks, estimation of the regularized inverse covariance matrix with S-E(glasso) will be key to addressing several other important questions arising from microbiome studies. For example, statistical methods to infer which taxa are responsive to design factors in 16S gene amplicon studies is an active area of research. Most methods test each taxon independently one-at-a-time (see [56] and references therein) even though taxa are actually highly correlated and thought to ecologically interact. Inference of taxa responses from 16S rRNA gene sequencing datasets could be improved by modeling this correlation structure through incorporation of the inverse covariance matrix into the statistical model [66].

Other, more complex questions are motivated by a desire to understand why and how ecosystems evolve with time. In the dynamic modeling setting, association networks have already been successfully used as an underlying structure to fit a differential-equation-based model of gut microbiome development in mice [14]. Thus, association networks provide the underlying topology for dynamic models, which can be used to develop hypotheses about how the ecosystems might respond to specific perturbations [5].

In conclusion, SPIEC-EASI is an improvement over state-of-the-art methods for inference of microbial ecological networks from microbiome composition datasets. We demonstrate this through rigorous benchmarking with synthetic networks and also through application to a true biological dataset. In addition, the LASSO underpinnings of the SPIEC-EASI inference

methods provide a flexible and principled mathematical framework to incorporate additional information about microbial ecological association networks as it becomes available, thereby improving prediction. We anticipate that SPIEC-EASI network inference will serve as a backbone for more sophisticated modeling endeavors, engendering new hypotheses and predictions of relevance to environmental ecology and medicine.

## Supporting Information

**S1 Text. Comparative marginal fits to American Gut Project count data.** A discussion of the methods used to simulate biologically-plausible correlated count data, and some results comparing different generative distributions.

(PDF)

**S1 Table. Method comparison.** Table to compare some of the features of SPIEC-EASI, SparCC, CCREPE and Pearson's correlation coefficient.

(PDF)

**S1 Fig. Relationship between target and empirical correlation after NorTA transformation.**

The recovery of empirical Pearson correlations generated from the NORITA process, using zero-inflated Negative Binomial as a model (x-axis) versus the input multivariate Normal empirical correlations (upper panels) or inverse correlations (lower panels) on untransformed counts (a) or log-transformed counts (b). Simulated data are with  $p = 205$  OTUs,  $n = 20,000$  samples with 10 replicates on each plot.

(PDF)

**S2 Fig. Heatmaps of AGP and synthetic data.** Visual comparison of American Gut Project data with synthetically generated datasets. Heatmaps show  $\log_{10}(\text{counts})$ . For all network types (band, graph, scale-free), synthetic datasets are consistent with real datasets in terms of number of OTUs, number of samples, and OTU count abundances across samples.

(PDF)

**S3 Fig. Effect of Condition Number on Correlations Distribution.** The condition of a Precision matrix is the ratio of the largest to smallest eigenvalue/singular value. The relationship between condition number and correlation distribution for the synthetic networks; increasing condition number corresponds to increasing the strength of correlations in the network.

(PDF)

**S4 Fig. Betweenness Centrality for Synthetic Networks.** Examples of betweenness centrality distributions for each network type (in white) overlaid with the distribution predicted by S-E (MB) (in orange) for  $\kappa = 100$ ,  $n = 1360$  samples,  $p = 205$  OTUs.

(PDF)

**S5 Fig. Recovery Performance of Betweenness Centrality.** Performance results for betweenness centrality (red = S-E(glasso) orange = S-E(MB), purple = SparCC, blue = CCREPE, green = Pearson). Bars represent the average KL-divergence over three independent sets of synthetic datasets (7 datasets per set); error bars represent standard error. Asterisks indicate that an S-E method had significantly better recovery of the true betweenness centrality distributions ( $p < 0.05$  for one-sided T tests in comparison to each control method).

(PDF)

**S6 Fig. Geodesic Distance for Synthetic Networks.** Examples of geodesic distance distributions for each network type (in white) overlaid with the distribution predicted by S-E(MB) (in orange) for  $\kappa = 100$ ,  $n = 1360$  samples,  $p = 205$  OTUs (PDF)

**S7 Fig. Recovery Performance of Geodesic Distance.** Performance results for geodesic distance distributions (red = S-E(glasso), orange = S-E(MB), purple = SparCC, blue = CCREPE, green = Pearson). Bars represent the average KL-divergence over three independent sets of synthetic datasets (7 datasets per set); error bars represent standard error. Asterisks indicate that an S-E method had significantly better recovery of the true geodesic distance distributions ( $p < 0.05$  for one-sided T tests in comparison to each control method). (PDF)

**S8 Fig. Cluster Sizes for Synthetic Networks.** Examples of cluster size distributions for each network type (in white) overlaid with the distribution predicted by S-E(MB) (in orange), for  $\kappa = 100$ ,  $n = 1360$  samples,  $p = 205$  OTUs. (PDF)

**S9 Fig. Recovery Performance of Cluster Sizes.** Performance results for cluster size distributions (red = S-E(glasso), orange = S-E(MB), purple = SparCC, blue = CCREPE, green = Pearson). Bars represent the average KL-divergence over three independent sets of synthetic datasets (7 datasets per set); error bars represent standard error. Asterisks indicate that an S-E method had significantly better recovery of the true cluster size distributions ( $P < 0.05$  for one-sided T tests in comparison to each control method). (PDF)

**S10 Fig. Synthetic test case from SparCC.** Using the example data provided by the SparCC package, we inferred networks from SparCC (threshold correlation at  $\pm 0.35$ ), S-E(MB) and CCREPE (thresholding q-value at  $5 * 10^{-19}$ ). In this setting, SparCC and SPIEC-EASI correctly recover four true edges, including the association sign. (PDF)

**S11 Fig. Assortativity coefficients in inferred American Gut networks.** Network assortativity coefficients at the Phyla and Family level for each of the four inference methods. Assortativity is a measure of the tendency for nodes to be connected with nodes of the same taxonomic class. (PDF)

## Acknowledgments

We would like to thank Eric Alm and Jonathan Friedman for helpful discussions.

## Author Contributions

Conceived and designed the experiments: ZDK CLM ERM RAB. Performed the experiments: ZDK CLM ERM. Analyzed the data: ZDK CLM ERM RAB. Contributed reagents/materials/analysis tools: ZDK CLM ERM. Wrote the paper: ZDK CLM ERM DRL MJB RAB.

## References

1. Gilbert J, Meyer F, Jansson J, Gordon J, Pace N, et al. (2010) The earth microbiome project: Meeting report of the "1st emp meeting on sample selection and acquisition" at argonne national laboratory october 6th 2010. Standards in Genomic Sciences 3.

2. Turnbaugh PJ, Ley RE, Hamady M, Fraser-Liggett C, Knight R, et al. (2007) The human microbiome project: exploring the microbial part of ourselves in a changing world. *Nature* 449: 804. doi: [10.1038/nature06244](https://doi.org/10.1038/nature06244) PMID: [17943116](https://pubmed.ncbi.nlm.nih.gov/17943116/)
3. AmGut. The american gut project. <http://humanfoodproject.com/american-gut/>. Accessed: 2014-01-30.
4. Bunge J, Willis A, Walsh F (2014) Estimating the number of species in microbial diversity studies. *Annual Review of Statistics and Its Application* 1: 427–445. doi: [10.1146/annurev-statistics-022513-115654](https://doi.org/10.1146/annurev-statistics-022513-115654)
5. Foster JA, Krone SM, Forney LJ (2008) Application of ecological network theory to the human microbiome. *Interdisciplinary perspectives on infectious diseases* 2008: 839501. doi: [10.1155/2008/839501](https://doi.org/10.1155/2008/839501) PMID: [19259330](https://pubmed.ncbi.nlm.nih.gov/19259330/)
6. Arumugam M, Raes J, Pelletier E, Le Paslier D, Yamada T, et al. (2011) Enterotypes of the human gut microbiome. *Nature* 473: 174–180. doi: [10.1038/nature09944](https://doi.org/10.1038/nature09944) PMID: [21508958](https://pubmed.ncbi.nlm.nih.gov/21508958/)
7. Koren O, Knights D, Gonzalez A, Waldron L, Segata N, et al. (2013) A guide to enterotypes across the human body: meta-analysis of microbial community structures in human microbiome datasets. *PLoS computational biology* 9: e1002863. doi: [10.1371/journal.pcbi.1002863](https://doi.org/10.1371/journal.pcbi.1002863) PMID: [23326225](https://pubmed.ncbi.nlm.nih.gov/23326225/)
8. Chen J, Li H (2013) Variable selection for sparse dirichlet-multinomial regression with an application to microbiome data analysis. *The Annals of Applied Statistics* 7: 418–442. doi: [10.1214/12-AOAS592](https://doi.org/10.1214/12-AOAS592)
9. Gevers D, Kugathasan S, Denson LA, Vázquez-Baeza Y, Van Treuren W, et al. (2014) The treatment-naive microbiome in new-onset crohn’s disease. *Cell Host—Microbe* 15: 382–392. doi: [10.1016/j.chom.2014.02.005](https://doi.org/10.1016/j.chom.2014.02.005) PMID: [24629344](https://pubmed.ncbi.nlm.nih.gov/24629344/)
10. Lee SC, Tang MS, Lim YAL, Choy SH, Kurtz ZD, et al. (2014) Helminth colonization is associated with increased diversity of the gut microbiota. *PLoS Negl Trop Dis* 8: e2880. doi: [10.1371/journal.pntd.0002880](https://doi.org/10.1371/journal.pntd.0002880) PMID: [24851867](https://pubmed.ncbi.nlm.nih.gov/24851867/)
11. Faust K, Sathirapongsasuti JF, Izard J, Segata N, Gevers D, et al. (2012) Microbial Co-occurrence Relationships in the Human Microbiome. *PLoS Computational Biology* 8: e1002606. doi: [10.1371/journal.pcbi.1002606](https://doi.org/10.1371/journal.pcbi.1002606) PMID: [22807668](https://pubmed.ncbi.nlm.nih.gov/22807668/)
12. Fuhrman JA, Steele JA (2008) Community structure of marine bacterioplankton: Patterns, networks, and relationships to function. In: *Aquatic Microbial Ecology*. volume 53, pp. 69–81. doi: [10.3354/ame01222](https://doi.org/10.3354/ame01222)
13. Barberán A, Bates ST, Casamayor EO, Fierer N (2012) Using network analysis to explore cooccurrence patterns in soil microbial communities. *The ISME journal* 6: 343–51. doi: [10.1038/ismej.2011.119](https://doi.org/10.1038/ismej.2011.119) PMID: [21900968](https://pubmed.ncbi.nlm.nih.gov/21900968/)
14. Marino S, Baxter NT, Huffnagle GB, Petrosino JF, Schloss PD (2014) Mathematical modeling of primary succession of murine intestinal microbiota. *Proceedings of the National Academy of Sciences of the United States of America* 111: 439–44. doi: [10.1073/pnas.1311322111](https://doi.org/10.1073/pnas.1311322111) PMID: [24367073](https://pubmed.ncbi.nlm.nih.gov/24367073/)
15. Deng Y, Jiang YH, Yang Y, He Z, Luo F, et al. (2012) Molecular ecological network analyses. *BMC Bioinformatics* 13: 113. doi: [10.1186/1471-2105-13-113](https://doi.org/10.1186/1471-2105-13-113) PMID: [22646978](https://pubmed.ncbi.nlm.nih.gov/22646978/)
16. Steele JA, Countway PD, Xia L, Vigil PD, Beman JM, et al. (2011) Marine bacterial, archaeal and protistan association networks reveal ecological linkages. *The ISME journal* 5: 1414–25. doi: [10.1038/ismej.2011.24](https://doi.org/10.1038/ismej.2011.24) PMID: [21430787](https://pubmed.ncbi.nlm.nih.gov/21430787/)
17. Friedman J, Alm EJ (2012) Inferring correlation networks from genomic survey data. *PLoS computational biology* 8: e1002687. doi: [10.1371/journal.pcbi.1002687](https://doi.org/10.1371/journal.pcbi.1002687) PMID: [23028285](https://pubmed.ncbi.nlm.nih.gov/23028285/)
18. Aitchison J (1981) A new approach to null correlations of proportions. *Mathematical Geology* 13: 175–189. doi: [10.1007/BF01031393](https://doi.org/10.1007/BF01031393)
19. Faust K, Raes J (2012) Microbial interactions: from networks to models. *Nat Rev Micro* 10: 538–550. doi: [10.1038/nrmicro2832](https://doi.org/10.1038/nrmicro2832)
20. Meinshausen N, Bühlmann P (2006) High Dimensional Graphs and Variable Selection with the Lasso. *The Annals of Statistics* 34: 1436–1462. doi: [10.1214/009053606000000281](https://doi.org/10.1214/009053606000000281)
21. Bonneau R, Reiss D, Shannon P, Facciotti M, Hood L, et al. (2006) The inferelator: an algorithm for learning parsimonious regulatory networks from systems-biology data sets de novo. *Genome Biology* 7: R36. doi: [10.1186/gb-2006-7-5-r36](https://doi.org/10.1186/gb-2006-7-5-r36) PMID: [16686963](https://pubmed.ncbi.nlm.nih.gov/16686963/)
22. Friedman J, Hastie T, Tibshirani R (2008) Sparse inverse covariance estimation with the graphical lasso. *Biostatistics (Oxford, England)* 9: 432–441. doi: [10.1093/biostatistics/kxm045](https://doi.org/10.1093/biostatistics/kxm045)
23. Banerjee O, El Ghaoui L, d’Aspremont A (2008) Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine . . .* 9: 485–516.
24. Wille A, Zimmermann P, Vranová E, Fürholz A, Laule O, et al. (2004) Sparse graphical Gaussian modeling of the isoprenoid gene network in *Arabidopsis thaliana*. *Genome Biol* 5: R92. doi: [10.1186/gb-2004-5-11-r92](https://doi.org/10.1186/gb-2004-5-11-r92) PMID: [15535868](https://pubmed.ncbi.nlm.nih.gov/15535868/)

25. Friedman N (2004) Inferring Cellular Networks Using Probabilistic Graphical Models. *Science* 303: 799–805. doi: [10.1126/science.1094068](https://doi.org/10.1126/science.1094068) PMID: [14764868](https://pubmed.ncbi.nlm.nih.gov/14764868/)
26. Bonneau R (2008) Learning biological networks: from modules to dynamics. *Nature chemical biology* 4: 658–64. doi: [10.1038/nchembio.122](https://doi.org/10.1038/nchembio.122) PMID: [18936750](https://pubmed.ncbi.nlm.nih.gov/18936750/)
27. Jones DT, Buchan DWA, Cozzetto D, Pontil M (2011) Psicov: Precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics*.
28. Marks DS, Hopf TA, Sander C (2012) Protein structure prediction from sequence variation. *Nature biotechnology* 30: 1072–80. doi: [10.1038/nbt.2419](https://doi.org/10.1038/nbt.2419) PMID: [23138306](https://pubmed.ncbi.nlm.nih.gov/23138306/)
29. Ravikumar P, Wainwright MJ, Raskutti G, Yu B (2011) High-dimensional covariance estimation by minimizing L1-penalized log-determinant divergence. *Electronic Journal of Statistics* 5: 935–980. doi: [10.1214/11-EJS631](https://doi.org/10.1214/11-EJS631)
30. Tandon R, Ravikumar P (2014) Learning Graphs with a Few Hubs. In: *Proceedings of The 31st International Conference on Machine Learning*. pp. 602–610. URL <http://jmlr.org/proceedings/papers/v32/taandon14.html>.
31. Liu Q, Ihler AT (2011) Learning scale free networks by reweighted l1 regularization. In: *AISTATS*. pp. 40–48.
32. Schloss PD, Westcott SL, Ryabin T, Hall JR, Hartmann M, et al. (2009) Introducing mothur: Open-source, platform-independent, community-supported software for describing and comparing microbial communities. *Applied and Environmental Microbiology* 75: 7537–7541. doi: [10.1128/AEM.01541-09](https://doi.org/10.1128/AEM.01541-09) PMID: [19801464](https://pubmed.ncbi.nlm.nih.gov/19801464/)
33. La Rosa PS, Brooks JP, Deych E, Boone EL, Edwards DJ, et al. (2012) Hypothesis testing and power calculations for taxonomic-based human microbiome data. *PloS one* 7: e52078. doi: [10.1371/journal.pone.0052078](https://doi.org/10.1371/journal.pone.0052078) PMID: [23284876](https://pubmed.ncbi.nlm.nih.gov/23284876/)
34. Aitchison J (1986) *The statistical analysis of compositional data*. London; New York: Chapman and Hall.
35. Koller D, Friedman N (2009) *Probabilistic graphical models: principles and techniques*. MIT press.
36. Lauritzen SL (1996) *Graphical models*. Oxford University Press.
37. Wainwright MJ, Jordan MI (2008) Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* 1: 1–305. doi: [10.1561/2200000001](https://doi.org/10.1561/2200000001)
38. Bühlmann P, Kalisch M, Meier L (2014) High-dimensional statistics with a view toward applications in biology. *Annual Review of Statistics and Its Application* 1: 255–278. doi: [10.1146/annurev-statistics-022513-115545](https://doi.org/10.1146/annurev-statistics-022513-115545)
39. Ravikumar P, Wainwright MJ, Lafferty JD, et al. (2010) High-dimensional ising model selection using l1-regularized logistic regression. *The Annals of Statistics* 38: 1287–1319. doi: [10.1214/09-AOS691](https://doi.org/10.1214/09-AOS691)
40. Dempster AP (1972) Covariance selection. *Biometrics*: 157–175. doi: [10.2307/2528966](https://doi.org/10.2307/2528966)
41. Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)* 58: 267–288.
42. Lederer J, Müller CL (2015) Don't fall for tuning parameters: Tuning-free variable selection in high dimensions with the TREX. In: *AAAI Conference on Artificial Intelligence*.
43. Lam C, Fan J (2009) Sparsistency and rates of convergence in large covariance matrix estimation. *The Annals of Statistics* 37: 4254–4278. doi: [10.1214/09-AOS720](https://doi.org/10.1214/09-AOS720)
44. Loh PL, Wainwright MJ (2013) Structure estimation for discrete graphical models: Generalized covariance matrices and their inverses. *The Annals of Statistics* 41: 3022–3049. doi: [10.1214/13-AOS1162](https://doi.org/10.1214/13-AOS1162)
45. Liu H, Lafferty J, Wasserman L (2009) The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *J Mach Learn Res* 10: 2295–2328.
46. Yuan M, Lin Y (2007) Model selection and estimation in the Gaussian graphical model. *Biometrika* 94: 19–35. doi: [10.1093/biomet/asm018](https://doi.org/10.1093/biomet/asm018)
47. Hastie T, Tibshirani R, Friedman J, Hastie T, Friedman J, et al. (2009) *The Elements of Statistical Learning*. Springer.
48. Liu H, Roeder K, Wasserman L (2010) Stability approach to regularization selection (stars) for high dimensional graphical models. *Proceedings of the Twenty-Third Annual Conference on Neural Information Processing Systems (NIPS)*: 1–14.
49. Zhao T, Liu H, Roeder K (2012) The huge package for high-dimensional undirected graph estimation in R. *The Journal of Machine Learning Research* 13: 1059–1062.
50. Schaffter T, Marbach D, Floreano D (2011) GeneNetWeaver: in silico benchmark generation and performance profiling of network inference methods. *Bioinformatics (Oxford, England)* 27: 2263–2270. doi: [10.1093/bioinformatics/btr373](https://doi.org/10.1093/bioinformatics/btr373)

51. Nelsen RB (1999) An introduction to copulas. Springer Series in Statistics. Springer.
52. Madsen L, Dalthorp D (2007) Simulating correlated count data. *Environmental and Ecological Statistics* 14: 129–148. doi: [10.1007/s10651-007-0008-1](https://doi.org/10.1007/s10651-007-0008-1)
53. Cario MC, Nelson BL (1997) Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. *Industrial Engineering*: 1–19.
54. R Development Core Team (2011). R: A Language and Environment for Statistical Computing. URL <http://www.r-project.org>.
55. Yee TW (2007) The VGAM Package for Categorical Data Analysis. *Journal of Statistical Software* 32: 1–34.
56. McMurdie PJ, Holmes S (2014) Waste Not, Want Not: Why Rarefying Microbiome Data Is Inadmissible. *PLoS Computational Biology* 10: e1003531. doi: [10.1371/journal.pcbi.1003531](https://doi.org/10.1371/journal.pcbi.1003531) PMID: [24699258](https://pubmed.ncbi.nlm.nih.gov/24699258/)
57. La Rosa PS, Brooks JP, Deych E, Boone EL, Edwards DJ, et al. (2012) Hypothesis testing and power calculations for taxonomic-based human microbiome data. *PloS one* 7: e52078. doi: [10.1371/journal.pone.0052078](https://doi.org/10.1371/journal.pone.0052078) PMID: [23284876](https://pubmed.ncbi.nlm.nih.gov/23284876/)
58. Barabási AL, Albert R (1999) Emergence of Scaling in Random Networks. *Science* 286: 509–512. doi: [10.1126/science.286.5439.509](https://doi.org/10.1126/science.286.5439.509) PMID: [10521342](https://pubmed.ncbi.nlm.nih.gov/10521342/)
59. Friedman J. SparCC package and data. <https://bitbucket.org/yonatanf/sparcc/src/9a1142c179f7/example>. Accessed: 2014-11-05.
60. Longman RS, Yang Y, Diehl GE, Kim SV, Littman DR (2013) Microbiota: Host interactions in mucosal homeostasis and systemic autoimmunity. In: Cold Spring Harbor symposia on quantitative biology. Cold Spring Harbor Laboratory Press, volume 78, pp. 193–201.
61. HMP Consortium (2012) Structure, function and diversity of the healthy human microbiome. *Nature* 486: 207–214. doi: [10.1038/nature11234](https://doi.org/10.1038/nature11234) PMID: [22699609](https://pubmed.ncbi.nlm.nih.gov/22699609/)
62. de Vos WM, de Vos EA (2012) Role of the intestinal microbiome in health and disease: from correlation to causation. *Nutrition Reviews* 70: S45–S56. doi: [10.1111/j.1753-4887.2012.00505.x](https://doi.org/10.1111/j.1753-4887.2012.00505.x) PMID: [22861807](https://pubmed.ncbi.nlm.nih.gov/22861807/)
63. Scher JU, Sczesnak A, Longman RS, Segata N, Ubeda C, et al. (2013) Expansion of intestinal prevotella copri correlates with enhanced susceptibility to arthritis. *eLife* 2. doi: [10.7554/eLife.01202](https://doi.org/10.7554/eLife.01202) PMID: [24192039](https://pubmed.ncbi.nlm.nih.gov/24192039/)
64. Yuan M, Lin Y (2006) Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68: 49–67. doi: [10.1111/j.1467-9868.2005.00532.x](https://doi.org/10.1111/j.1467-9868.2005.00532.x)
65. Greenfield A, Hafemeister C, Bonneau R (2013) Robust data-driven incorporation of prior knowledge into the inference of dynamic regulatory networks. *Bioinformatics* 29: 1060–1067. doi: [10.1093/bioinformatics/btt099](https://doi.org/10.1093/bioinformatics/btt099) PMID: [23525069](https://pubmed.ncbi.nlm.nih.gov/23525069/)
66. Lin W, Shi P, Feng R, Li H (2014) Variable selection in regression with compositional covariates. *Biometrika* accepted.