



# Identifying Keystone Species in the Human Gut Microbiome from Metagenomic Timeseries Using Sparse Linear Regression

Charles K. Fisher, Pankaj Mehta\*

Department of Physics, Boston University, Boston, Massachusetts, United States of America

## Abstract

Human associated microbial communities exert tremendous influence over human health and disease. With modern metagenomic sequencing methods it is now possible to follow the relative abundance of microbes in a community over time. These microbial communities exhibit rich ecological dynamics and an important goal of microbial ecology is to infer the ecological interactions between species directly from sequence data. Any algorithm for inferring ecological interactions must overcome three major obstacles: 1) a correlation between the abundances of two species does not imply that those species are interacting, 2) the sum constraint on the relative abundances obtained from metagenomic studies makes it difficult to infer the parameters in timeseries models, and 3) errors due to experimental uncertainty, or mis-assignment of sequencing reads into operational taxonomic units, bias inferences of species interactions due to a statistical problem called “errors-in-variables”. Here we introduce an approach, Learning Interactions from Microbial Time Series (LIMITS), that overcomes these obstacles. LIMITS uses sparse linear regression with bootstrap aggregation to infer a discrete-time Lotka-Volterra model for microbial dynamics. We tested LIMITS on synthetic data and showed that it could reliably infer the topology of the inter-species ecological interactions. We then used LIMITS to characterize the species interactions in the gut microbiomes of two individuals and found that the interaction networks varied significantly between individuals. Furthermore, we found that the interaction networks of the two individuals are dominated by distinct “keystone species”, *Bacteroides fragilis* and *Bacteroides stercosis*, that have a disproportionate influence on the structure of the gut microbiome even though they are only found in moderate abundance. Based on our results, we hypothesize that the abundances of certain keystone species may be responsible for individuality in the human gut microbiome.

**Citation:** Fisher CK, Mehta P (2014) Identifying Keystone Species in the Human Gut Microbiome from Metagenomic Timeseries Using Sparse Linear Regression. PLoS ONE 9(7): e102451. doi:10.1371/journal.pone.0102451

**Editor:** Bryan A. White, University of Illinois, United States of America

**Received:** March 10, 2014; **Accepted:** June 17, 2014; **Published:** July 23, 2014

**Copyright:** © 2014 Fisher, Mehta. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Funding:** PM and CKF were supported by a Sloan Research Fellowship (to PM), National Institutes of Health Grants K25GM086909 (to PM), and internal funding from Boston University. 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.

\* Email: pankajm@bu.edu

## Introduction

Metagenomic sequencing technologies have revolutionized the study of the human-associated microbial consortia making up the human microbiome. Sequencing methods now allow researchers to estimate the relative abundance of the species in a community without having to culture individual species [1–3]. These studies have shown that microbial cells vastly outnumber human cells in the body, and that symbiotic microbial communities are important contributors to human health [1]. For example, a recent study by Ridaura et al [4] demonstrated that transplants of gut microbial consortia are sufficient to induce obesity in previously lean mice or to promote weight loss in previously obese mice, suggesting an intriguing hypothesis that the composition of the gut microbiome may also contribute to obesity in humans. Many other studies have found significant links between the composition of human-associated microbial consortia and diseases including cancer and autism spectrum disorder [4–10]. Despite the recent revelations highlighting the importance of the microbiome to human health, relatively little is known about the ecological structure and dynamics of these microbial communities.

A microbial community consists of a vast number of species, all of which must compete for space and resources. In addition to competition, there are also many symbiotic interactions where certain species benefit from the presence of other microbial species. For example, a small molecule that is secreted by one species can be metabolized by another [11]. These species interactions provide a window with which to view the ecology of a microbial community, and allow one to make predictions about the effect of perturbations on a population [12]. For example, removing a species that engages in mutualistic interactions may diminish the abundance of other species that depend on it for survival. Given their utility for understanding the ecology of a community, there is tremendous interest in developing techniques to infer interactions between species from metagenomic data [12–14].

There are two approaches to inferring dependencies between microbial species from metagenomic studies: cross-sectional analysis, and timeseries analysis [12–16]. Cross-sectional studies pool samples of the relative abundances of the microbial species in a particular environment (e.g. the gut) from multiple individuals and utilize correlations in the relative abundances as proxies for

effective interactions between species. By contrast, timeseries analysis follows the relative abundances of the microbial species in a particular environment, for a single individual, over time and utilizes dynamical modeling (e.g. ordinary differential equations) to understand dependencies between species.

Any methods for making reliable inferences about species interactions from metagenomic studies must overcome three major obstacles. First, as shown below, a correlation between the abundances of two species does not imply that those species are interacting. Second, metagenomic methods measure the relative, not absolute, abundances of the microbial species in a community. This makes it difficult to infer the parameters in timeseries models. Finally, errors due to experimental measurement errors and/or mis-assignment of sequencing reads into operational taxonomic units (OTUs), bias inferences of species interactions due to a statistical problem called “errors-in-variables” [17]. We will show that each of these obstacles can be overcome using a new method we call LIMITS (Learning Interactions from Microbial Time Series). LIMITS obtains a reliable estimate for the topology of the directed species interaction network by employing sparse linear regression with bootstrap aggregation (“Bagging”) to learn the species interactions in a discrete-time Lotka-Volterra (dLV) model of population dynamics from a time series of relative species abundances [18,19].

## Results

### Correlation does not imply interaction

Many previous works use the correlation between the relative abundances of two microbial species in an environment (e.g. the gut) as a proxy for how much the species interact. In particular, a high degree of correlation between the abundances of two species is often taken as a proxy for a strong mutualistic interaction, and large anti-correlations, as indicative of a strong competitive interactions. Using correlations as a proxy for interactions suffers from several drawbacks. First, there are important subtleties involved in calculating correlations between species from relative abundances, but previous studies have presented algorithms (e.g. SparCC) to mitigate these problems [14]. More importantly, the abundances of two species may be correlated even if those species do not directly interact. For example, if species A directly interacts with species B, and species B directly interacts with species C, the abundances of species A and C are likely to be correlated even though they do not directly interact. Finally, since correlation matrices are necessarily symmetric, all interactions learned using correlations must also be symmetric.

The problems with using correlations in species abundances as proxies for species interactions can be illustrated with a simple numerical simulation. We used the dLV model (Eq. 2) to simulate timeseries of the absolute abundances of 10 species for 1000 timesteps, starting from 100 different initial conditions, for two arbitrary species interaction matrices (see Materials and Methods). See Figure S1 for example time series. Figure 1 compares the Pearson correlation matrices calculated from the absolute species abundances obtained from the dLV simulations and the true interaction matrices. It is clear from Fig. 1 that there is no obvious relationship between the correlations and the species interactions. That is, correlations in species abundances are actually very poor proxies for species interactions.

In general, the relationship between the interaction coefficients ( $c_{ij}$ ) and the correlations in the species abundances is described by a complicated non-linear function that is difficult to compute or utilize. This can be seen by linearizing the dynamics of  $\ln x_i(t)$  (Eq. 4) around their equilibrium values,  $\ln \langle x_i(t) \rangle$ . This stochastic

process is a first-order autoregressive model described by

$$\ln x(t+1) = \omega + J \ln x(t) + \zeta(t), \quad (1)$$

where  $\omega_i = -\sum_j c_{ij} \langle x_j \rangle$ , and  $J_{ij} = \delta_{ij} + c_{ij} \langle x_j \rangle$  is the Jacobian obtained by linearizing around the equilibrium species abundances, and  $\zeta(t) = \ln \eta(t)$ . If  $V$  is the covariance matrix with elements  $V_{ij} = \langle (\ln x_i(t) - \langle \ln x_i(t) \rangle)(\ln x_j(t) - \langle \ln x_j(t) \rangle) \rangle$  and  $\Sigma$  is the covariance matrix of the Gaussian noise  $\zeta(t)$ , then  $\text{vec}(V) = (I_{n^2} - J \otimes J)^{-1} \text{vec}(\Sigma)$ , where  $\otimes$  is the Kronecker product and  $\text{vec}$  is the matrix vectorization operator [20]. Thus, the interaction matrix is related to the covariance matrix by complex relation even in the linear regime of the dynamics. For this reason, it is very difficult, if not impossible, to determine the interaction coefficients using only knowledge of the correlations and equilibrium abundances.

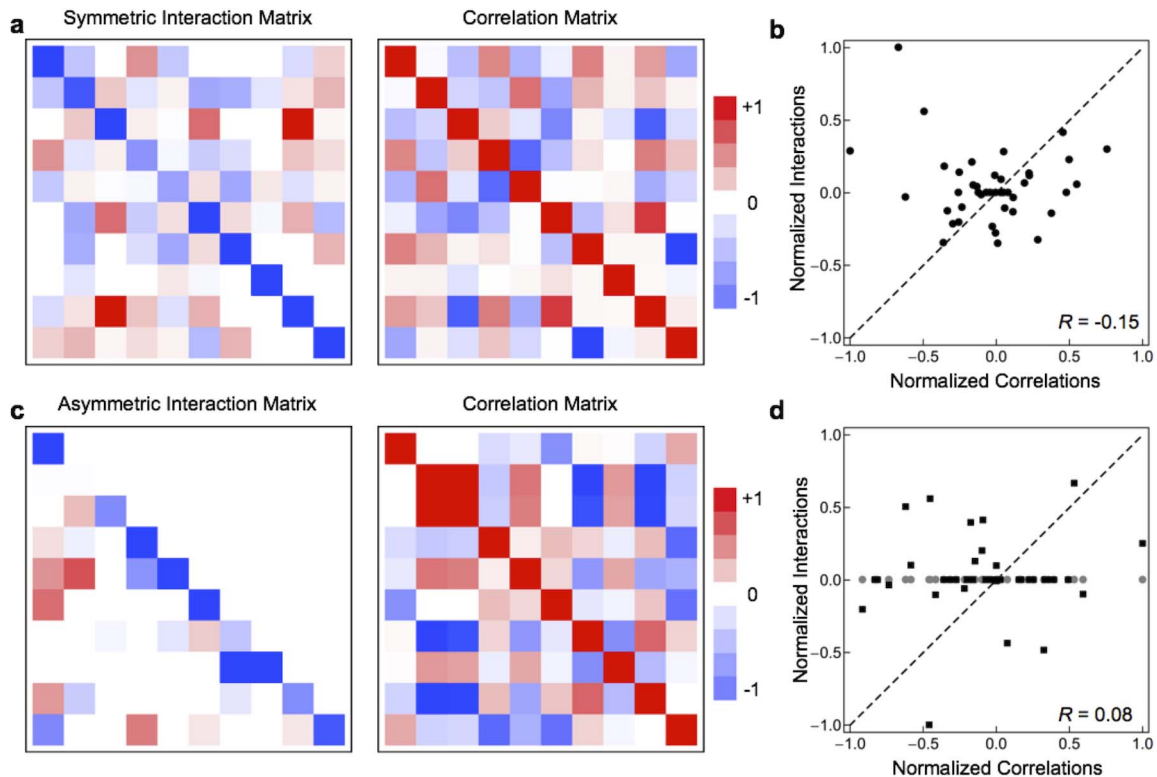
Cross-sectional studies that pool data across individuals and utilize the correlations between the abundances of different species as proxies for species dependencies are especially affected by this problem. This suggests that time-series data is likely to be more suited for inferring ecological interactions than cross-sectional data.

### Timeseries inference with relative species abundances

Even though the species interaction coefficients cannot be inferred from the correlations in species abundances, it is possible to reliably infer the interaction matrix using timeseries models. To do so, one utilizes a discrete time Lotka-Volterra Model (dLV) that relates the abundance of species  $i$  at a time  $t+1$  ( $x_i(t+1)$ ) to the abundances of *all* the species in the ecosystem at a time  $t$  ( $\vec{x} = \{x_1(t), \dots, x_N(t)\}$ ). These interactions are encoded in the dLV through a set of interaction coefficients,  $c_{ij}$ , that describe the influence species  $j$  has on the abundance of species  $i$  [19], and inferring these interaction coefficients is the major goal of this work. The effect of species  $j$  on species  $i$  can be beneficial ( $c_{ij} > 0$ ), competitive ( $c_{ij} < 0$ ), or the two species may not interact ( $c_{ij} = 0$ ).

As shown in the Materials and Methods, given a time-series of the absolute abundances of the microbes in an ecosystem, one can learn the interaction coefficients by performing a linear regression of  $\ln x_i(t+1) - \ln x_i(t)$  against  $\vec{x}(t) - \langle \vec{x} \rangle$ , where  $\langle x \rangle$  is the vector of the equilibrium abundances. It is important to note that each of these linear regressions can be performed independently for species  $i=1, \dots, N$ . In the following, we assume that the population dynamics are stable and that the equilibrium  $\langle x \rangle$  (or  $\langle \vec{x} \rangle$ ) can be estimated by taking the median species abundances over the time series.

Recall that most modern metagenomic techniques can only measure the relative abundances of microbes, not absolute abundances. This introduces additional complications into the problem of inferring species interactions using timeseries data. Although it is straight forward to infer species interactions by applying linear regression to a timeseries of absolute abundances, it is not *a priori* clear that linear regression still works when applied to a timeseries of the relative abundances. An important technical problem that arises when using relative abundances is that the design matrix for the regression is singular because of the sum constraint on relative abundances of species ( $\sum_i \tilde{x}_i(t) = 1$ ). As a result, there is no unique solution to the ordinary least squares problem applied to timeseries of relative species abundances. Nevertheless, the design matrix can be made to be invertible if one, or more, of the species are not included as variables in the regression.



**Figure 1. There is no simple relation between interaction coefficients and correlations in abundance.** a) A symmetric interaction matrix and the corresponding correlation matrix. b) There is no relation between the interaction parameters and the correlations in abundance for the symmetric interaction matrix. c) An asymmetric interaction matrix and the corresponding correlation matrix. d) There is no relation between the interaction parameters and the correlations in abundance for the asymmetric interaction matrix. Points from above the diagonal in the interaction matrix are gray circles, whereas points from below the diagonal are black squares. In a and c, matrix elements have been scaled so that the smallest negative element is  $-1$ , the largest positive element is  $+1$ , and all elements retain their sign. In b and d, interaction coefficients were scaled so that the largest element by absolute value has  $|c_{ij}|=1$ . doi:10.1371/journal.pone.0102451.g001

This insight motivates the use of a forward stepwise regression for selecting the covariate species that explain the changes in abundance of species  $i$ . In such a procedure, interactions and species are added sequentially to the regression as long as they improve the predictive power of model (see Fig. 2a). Since the design matrix now only contains a sub-set of all possible species, it is never singular and the linear regression problem is well-defined. Furthermore, the goal of forward selection is to include only the strongest, most important species interactions in the model. Therefore, the resulting interaction networks are sparse and, hence, easily interpretable.

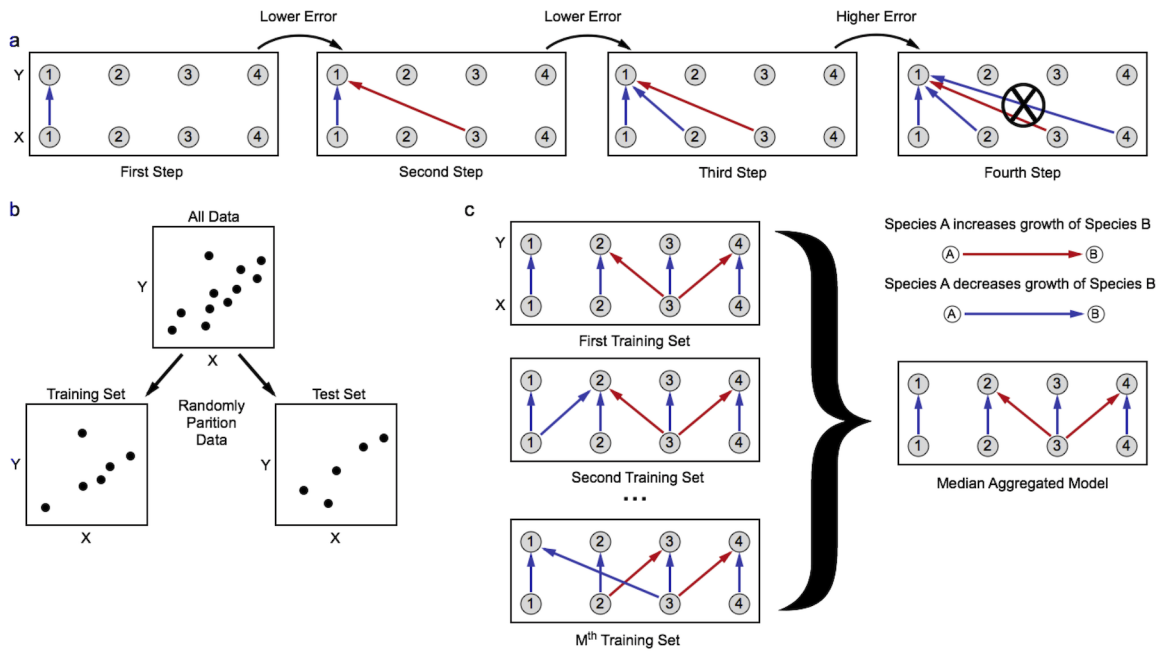
The procedure for forward stepwise regression is illustrated in Fig. 2a. We know that each species must interact with itself, so  $c_{ij}$  is the only interaction coefficient allowed to be nonzero in the first iteration. In each subsequent iteration, one additional interaction  $c_{ij}$  is included in the model by scanning over all other species and choosing the one that produces the lowest error at predicting a test dataset. This is repeated as long as the prediction error decreases by a pre-specified percentage that controls the sparsity of the model. A larger required improvement in prediction results in more sparse solutions. Note that the threshold defining the required improvement in prediction is a relative improvement of one model relative to another (expressed as a percent), not a measure of the absolute error in out of sample prediction.

Forward stepwise regression is a greedy algorithm, which results in a well-known instability [18]. This instability can be ‘cured’ using a method called bootstrap aggregation, or ‘Bagging’

(Fig. 2b,c) [18]. To bag forward stepwise regression, the data are randomly partitioned into a training set used for the regression and a test set used for evaluating the prediction error. The required improvement in prediction is a percentage that refers to how much the mean squared error evaluated on the test dataset must decrease in order to include an additional variable in the regression. The random partitioning of the data into training and test sets is repeated many times, each one resulting in a different estimate for the interaction matrix. The classical approach to Bagging calls for averaging these different estimates, but this destroys the sparsity of the solution. For this reason, we use the median of the estimates instead of averages. This still greatly improves the stability of the inferred interaction matrix but preserves its sparsity. We call our algorithm LIMITS (Learning Interactions from Microbial Time Series).

Figure 3 presents the results from applying LIMITS to infer the same interaction matrices discussed in Figure 1. The data consist of either absolute or relative abundances from timeseries with 500 timesteps and 10 different initial conditions. The sample sizes are quite large, but not as large as used for the calculation of the correlations. The inferred parameters match the true interaction coefficients very accurately – the smallest  $R^2$  between the inferred and true parameters is 0.82 – for both the symmetric (Fig. 3a–b) and asymmetric (Fig. 3c–d) interaction matrices using either absolute or relative species abundances.

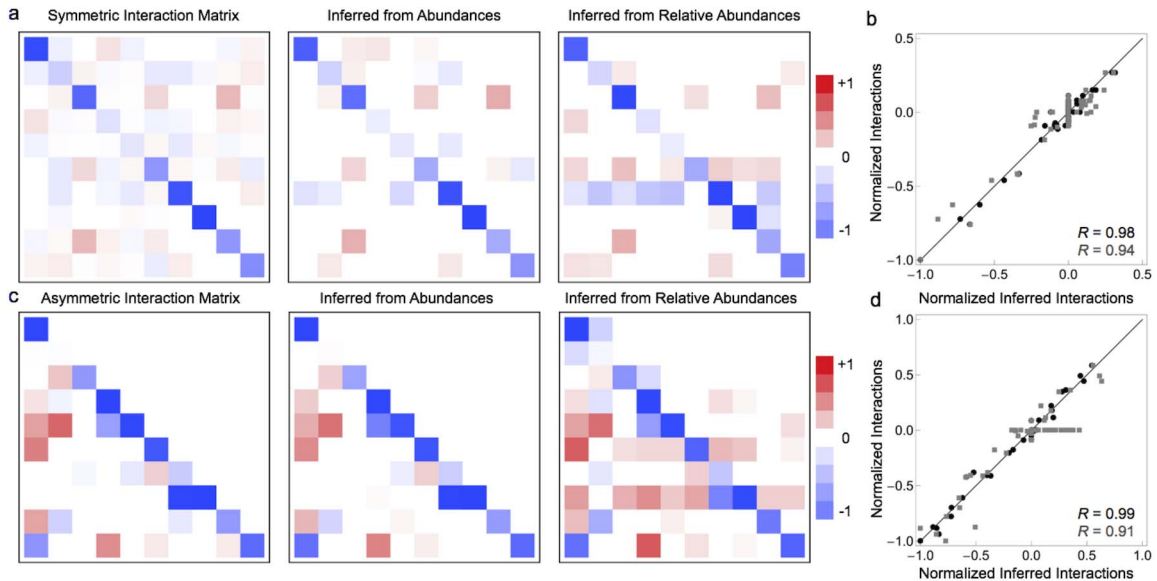
To ensure that the exceptional performance of our sparse linear regression approach to inferring species interactions was not a



**Figure 2. Schematic illustrating forward stepwise regression and median bootstrap aggregating.** a) In forward stepwise regression, interactions are added to the model one at a time as long as including the additional covariate lowers the prediction error by a pre-defined threshold. b) The prediction error used for variable selection is evaluated by randomly partitioning the data into a training set used for the regression and a test set used to evaluate the prediction error. c) Multiple models are built by repeatedly applying forward stepwise regression to random partitions of the data, each containing half the data points. The models are aggregated, or “bagged”, by taking the median, which improves the stability of the fit while preserving the sparsity of the inferred interactions. doi:10.1371/journal.pone.0102451.g002

fluke due to a particular choice of interaction matrices, we calculated the correlation between the true and inferred parameters for many randomly generated interaction matrices (see Fig. 4

and Materials and Methods). Note that the average magnitude of  $c_{ij}$  relative to the noise in these simulations is fixed, and that these simulations have no measurement errors (see next section for more



**Figure 3. Example fits of interaction parameters using sparse linear regression.** a) A symmetric interaction matrix (left), the corresponding matrix inferred from absolute abundance data (middle), and the corresponding matrix inferred from relative abundance data (right). b) There is good agreement between the true and inferred interactions, from both absolute (black) and relative (gray) abundances, for the symmetric interaction matrix. c) An asymmetric interaction matrix (left), the corresponding matrix inferred from absolute abundance data (middle), and the corresponding matrix inferred from relative abundance data (right). d) There is good agreement between the true and inferred interactions, from both absolute (black) and relative (gray) abundances, for the asymmetric interaction matrix. The prediction error threshold was set to 5% in for all fits. doi:10.1371/journal.pone.0102451.g003

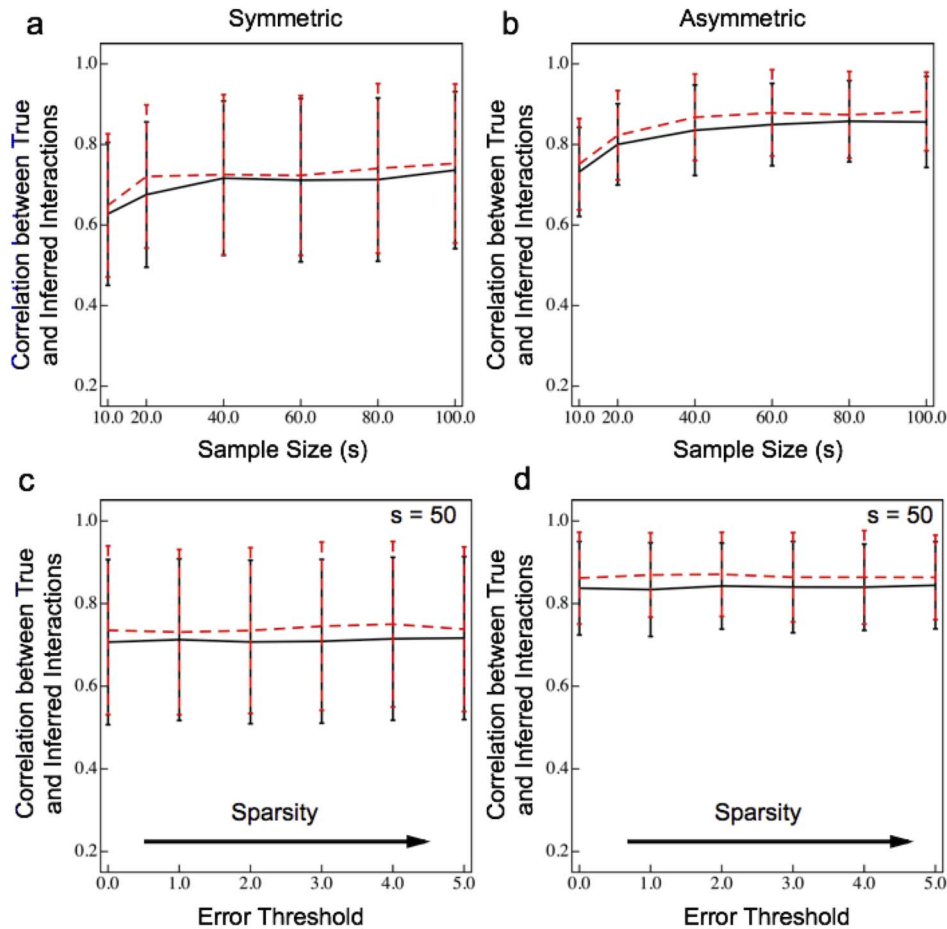
on this). The performance of the algorithm obviously depends on sample size, but LIMITS generally perform admirably at inferring ecological interactions (see Figs. 4a–b for symmetric and asymmetric matrices, respectively). Furthermore, Figs. 4c–d show that our results do *not* strongly depend on the required improvement in prediction over most reasonable choices of threshold (0–5%). These results demonstrate that in the absence of measurement noise, LIMITS can successfully learn the interaction parameters from both absolute and relative abundances.

### Inferring interactions in the presence of measurement errors

Up to this point, our analyses have ignored the impact of “measurement noise” on the inferred species interactions. There are two important sources of measurement noise in metagenomic data. The first source is experimental noise introduced by sequencing errors. The second, and perhaps larger source of noise, is the mis-classification of sequencing reads into operational taxonomic units (OTUs). Most metagenomic studies rely on the sequencing of 16S rRNA to estimate species composition and diversity in a community. These 16S sequences are binned into

groups, or OTUs, that contain sequences with a predetermined degree of similarity. By comparing the sequences in an OTU to known sequences in an annotated database, it is often possible to assign OTUs to particular species or strains. In general, this is an extremely difficult bioinformatics problem [21] and is likely to be a significant source of measurement errors. Thus, any algorithm for inferring species interactions must be robust to measurement errors.

At first glance, it is tempting to assume that measurement noise, which we assume is multiplicative, simply adds to the stochastic ( $\ln \eta_i(t)$ ) term that acts on the dependent variable ( $\ln x_i(t+1) - \ln x_i(t)$ ) and, therefore, should have little impact on the inferred interactions (see Methods). However, as we discuss below, this is not the case since the  $x_i(t)$  also act as the independent variables in the regression. Standard regression techniques assume that the independent variables are known exactly, and violation of this assumption results in biased parameter estimates even for asymptotically large sample sizes [17]. For example, in the simplest case of a 1-dimensional regression  $Y = \alpha + \beta X$  the estimator  $\hat{\beta}$  is always less than the real  $\beta$ , i.e.  $\hat{\beta} = \text{COV}(X, Y) / \text{VAR}(X) \leq \beta$  with equality only if there is no measurement error on  $X$ . The bias induced by using noisy



**Figure 4. Performance of sparse linear regression as a function of sample size and the prediction error threshold.** a) Performance on absolute (red) and relative (black) abundances as a function of sample size for symmetric interaction matrices. b) Performance on absolute (red) and relative (black) abundances as a function of sample size for asymmetric interaction matrices. c) Performance on absolute (red) and relative (black) abundances as a function of the out-of-bag error threshold for symmetric interaction matrices. d) Performance on absolute (red) and relative (black) abundances as a function of the out-of-bag error threshold for symmetric interaction matrices. Error bars correspond to  $\pm$  one standard deviation, and lines connect the means.

doi:10.1371/journal.pone.0102451.g004

independent variables in regression is known as the “errors-in-variables” problem in the statistical literature [17]. Analysis of the errors-in-variables bias for multivariate regression is more complicated than the 1-dimensional example; nevertheless, it can be stated quite generally that the interaction parameters inferred in the presence of significant measurement errors will be incorrect. The most reliable method for mitigating the errors-in-variables bias is to measure additional data on some “instrumental variables” that provide information on the true values of the relative species abundances [17]. Unfortunately, in most cases we often do not have access to such additional data.

Although the errors-in-variables bias cannot be eliminated, the topology of the interaction network can still be reliably inferred using our sparse linear regression approach even when the measurements of the relative species abundances are very noisy. Knowledge of which interactions are beneficial ( $c_{ij} > 0$ ), competitive ( $c_{ij} < 0$ ), or zero ( $c_{ij} = 0$ ) defines the topology of the interaction network. In the following simulations, we focus only on whether a given interaction is zero ( $c_{ij} = 0$ ) or not ( $c_{ij} \neq 0$ ) because we found that errors in the signs of the interactions were rare (Fig. S2). The accuracy of the interaction topologies inferred from noisy relative abundance data were assessed using simulations with randomly chosen (asymmetric) interaction matrices. We computed the specificity – the fraction of species pairs correctly identified as non-interacting – and the sensitivity – the fraction of species pairs correctly identified as interacting – of the inferred topologies (Fig. 5) [22]. Figure 5a shows that LIMITS produced specificities between 60% and 80% for different choices of the prediction error threshold, and that the performance was relatively insensitive to multiplicative measurement noise up to, and even beyond, 10%. By contrast, Fig. 5b shows that applying forward variable selection to the entire dataset, i.e. without Bagging, produced results that were very sensitive to the choice of the prediction error threshold and even produced specificities as low as 0%. The sensitivities for detecting interacting species with (Fig. 5c) or without (Fig. 5d) Bagging are both quite good, ranging between 70% and 80% for measurement noise up to 10%. These results demonstrate that both the forward stepwise regression and the median bootstrap aggregation are crucial components of the LIMITS algorithm. Moreover, LIMITS reliably infers the topology of the species interaction network even when there are significant errors-in-variables.

### Keystone species in the human gut microbiome

Emboldened by the success of our algorithm on synthetic data, we applied LIMITS to infer the species interactions in the gut microbiomes of two individuals. The data from Caporaso et al [3] were obtained from the MGRAST database [23] in a pre-processed form; i.e. as relative species abundances instead of raw sequencing data. These data consisted of approximately a half-year of daily sampled relative species abundances for individual (a) and a full year of daily sampled relative species abundances for individual (b). Note that an interaction matrix has of order  $n^2$  free parameters for  $n$  species. Also, Figures 4 and S2 show that the performance of LIMITS plateaus around  $n^2$  data points. Thus, in general, we require at least  $n^2$  timepoints in order to infer the interactions from  $n$  species. The number of available timepoints was  $O(100)$ , so we considered only the 10 most abundant species from individuals (a) and (b). Because the most abundant species were not entirely the same in individuals (a) and (b), we studied 14 species obtained by taking the union of the 10 most abundant species from individual (a) with the 10 most abundant species from individual (b). The 14 species are listed in Fig. 6.

The species interaction network of the gut microbiome of individual (a) (shown in Fig. 6a) is dominated by the species *Bacteroides fragilis* even though it is found only in moderate abundances. *Bacteroides fragilis* has 6 outgoing interactions in individual (a), in contrast to the other 13 species that have 0–3 outgoing interactions. The species interaction network of the gut microbiome of individual (b) (shown in Fig. 6b) is also dominated by a single species, *Bacteroides stercosis*, which is also found only in moderate abundances. *Bacteroides stercosis* has 4 outgoing interactions in individual (b), in contrast to the other 13 species that have 0–2 outgoing interactions. In addition, many of the interactions involving *Bacteroides fragilis* and *Bacteroides stercosis* are beneficial interactions. Based on these results, we refer to *Bacteroides fragilis* and *Bacteroides stercosis* as “keystone species” of the human gut microbiome because these two species exert tremendous influence on the structure of the microbial communities, even though they have lower median abundances than some other species.

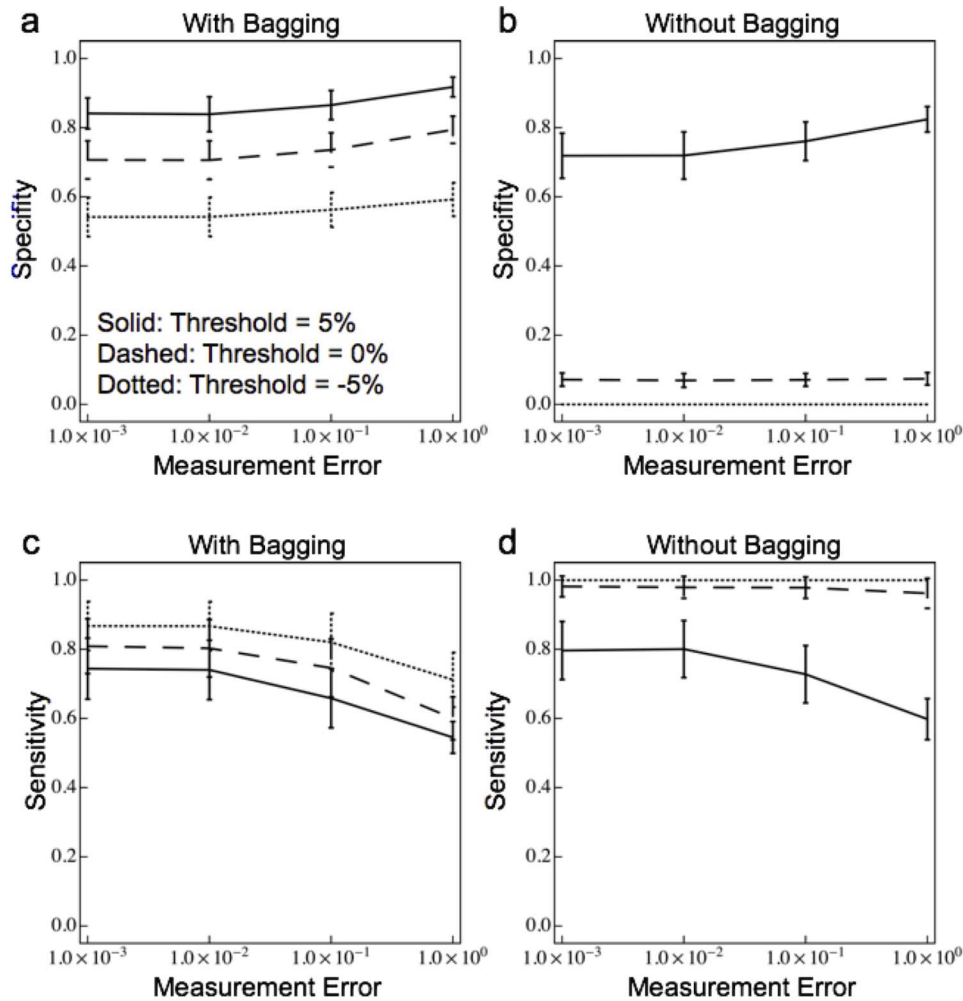
Additionally, we observed that the species interaction topology of the gut microbiome of individual (a) differs substantially from the species interaction topology of the gut microbiome of individual (b), as is clear from Fig. 6. In individual (a), *Bacteroides fragilis* is much more abundant than *Bacteroides stercosis* and, in turn, it is *Bacteroides fragilis* that dominates the interaction network of individual (a). Likewise, in individual (b), *Bacteroides stercosis* is much more abundant than *Bacteroides fragilis* and, in turn, it is *Bacteroides stercosis* that dominates the interaction network of individual (b). This observation motivates us to propose an intriguing hypothesis, that the abundances of certain keystone species are responsible for individuality of the human gut microbiome. Of course, much more data, from a larger population, will be required to confirm or reject this hypothesis.

### Discussion


Metagenomic methods are providing an unprecedented window into the composition and structure of microbial communities. They are revolutionizing our knowledge of microbial ecology and highlight the important roles played by the human microbiome in health and disease. Nevertheless, it is important to carefully consider the tools used to analyze these data and to address their associated challenges. We have highlighted three major obstacles that must be addressed by any study designed to use metagenomic data to analyze species interactions: 1) a correlation between the abundances of two species does not imply that those species are interacting, 2) the sum constraint on the relative abundances obtained from metagenomic studies makes it difficult to infer the parameters in timeseries models, and 3) errors due to experimental uncertainty, or mis-assignment of sequencing reads into operational taxonomic units (OTUs), bias inferences of species interactions due to a statistical problem called “errors-in-variables”.

To overcome these obstacles, we have introduced a novel algorithm, LIMITS, for inferring species interaction coefficients that combines sparse linear regression with bootstrap aggregation (Bagging). Our method provides reliable estimates for the topology of the species interaction network even when faced with significant measurement noise. The interaction networks constructed using our approach are sparse, including only the strongest ecological interactions. Regularizing the inference of the interaction network by favoring sparse solutions has the benefit that the results are easily interpretable, enabling the identification of keystone species with many important interactions. Furthermore, our work suggests that it is difficult to learn species interactions from cross-sectional





**Figure 5. Sensitivity and specificity of predicted interactions as a function of measurement error for bagged and unbagged models.** Specificity refers to the fraction of species pairs correctly identified as non-interacting, while sensitivity refers to the fraction of species pairs correctly identified as interaction. Both measures range from 0 (poor performance) to 1 (good performance). a) Specificity of sparse linear regression with Bagging as a function of measurement error for different prediction error thresholds. b) Specificity of sparse linear regression trained on the entire data set without Bagging as a function of measurement error for different prediction thresholds. c) Sensitivity of sparse linear regression with Bagging as a function of measurement error for different prediction error thresholds. d) Sensitivity of sparse linear regression trained on the entire data set without Bagging as a function of measurement error for different prediction error thresholds. Notice that without bagging, model performance is extremely sensitive to choice of the threshold for the required improvement in prediction for adding new interactions. doi:10.1371/journal.pone.0102451.g005



**Figure 6. Interaction topologies of abundant species in the guts of two individuals.** The size of a node denotes the median relative species abundance, beneficial interactions are shown as solid red arrows, and competitive interactions are shown as dashed blue arrows. In individual a) species 4 *Bacteroides fragilis* acts as a keystone species with 6 outgoing interactions, compared to a median number of outgoing interactions of 1. In individual b) species 5 *Bacteroides stercoris* acts as a keystone species with 4 outgoing interactions, compared a median number of outgoing interactions of 1. The 14 species included in the model were obtained by taking the union of the top 10 most abundant species from individuals a and b. The required improvement in prediction was set to 3%, graphs obtained using other prediction thresholds are shown in the Supporting Information. doi:10.1371/journal.pone.0102451.g006

studies that pool samples of the relative abundances of the microbial species from multiple individuals. This highlights the importance of collecting extended time-series data for understanding microbial ecological dynamics.

We applied LIMITS to time-series data to infer ecological interaction networks of two individuals and found that the interaction networks are dominated by distinct keystone species. This motivated us to propose a hypothesis: that the abundances of certain keystone species are responsible for individuality of the human gut microbiome. While more data will be required to confirm or reject this hypothesis, it is intriguing to examine its potential consequences for the human microbiome. The keystone species hypothesis implies that even small perturbations to an environment can have a large impact on the composition of its resident microbial consortia if those perturbations affect a small number of important “keystone” species. Moreover, relatively small differences in individual diets, or minor differences in the interaction between the host immune system and the gut microbiota, that affect keystone species may be sufficient to organize gut microbial consortia into distinct types of communities, or “enterotypes” [24,25].

Our analysis identified the closely related species *Bacteroides fragilis* and *Bacteroides stercosis* as potential keystone species of the gut microbiome [26]. Previous studies have suggested that the abundance of *Bacteroides fragilis* modulates the levels of several metabolites and, in turn, the composition of the gut microbiome in a mouse model of gastrointestinal abnormalities associated with autism spectrum disorder [9]. Abundance of both *Bacteroides fragilis* and *Bacteroides stercosis* are associated with an increased risk of colon cancer [5,6,10], and previous authors even suggest that *Bacteroides fragilis* acts as a critical “keystone pathogen” in the development of the disease [8]. Classical ecological models of species interaction demonstrate that the manner of the interaction between two species is not solely a function of their identity, but is highly dependent on the environment in which the interaction takes place [27,28]. Increased abundance of *Bacteroides* species is associated with high fat diets, including typical Western diets with a high consumption of red meat that are associated with increased cancer risk [5]. It is possible that *Bacteroides fragilis* and *Bacteroides stercosis* act as keystone species in individuals consuming high fat diets due to their ability to convert bile into metabolites that are used by other members of the microbial community [5,9].

The keystone species hypothesis can be experimentally tested by perturbing the abundance of individual species in a microbial consortium and observing the effect on the composition of the community. Our prediction is that most perturbations will have little impact on the overall structure of the microbial community, but perturbations applied to a small number of keystone species will have a large impact on the structure of the community. Due to ethical concerns, it is difficult to envision a direct experimental test of the keystone species hypothesis in human microbiota and, therefore, to test our specific predictions in regards to the keystone species *Bacteroides fragilis* and *Bacteroides stercosis*. Nevertheless, experimental tests could be performed in animal models, or even in culture if a large enough microbial consortia can be assembled.

## Materials and Methods

### Discrete Time Lotka-Volterra Model for absolute abundances

Metagenomic sequencing methods have made it possible to follow the time evolution of a microbial population by determining the relative abundances of the species in a community in discrete

intervals (e.g. one day). Given the discrete nature of these data, it is most sensible to use a discrete-time model of population dynamics. The discrete-time Lotka-Volterra (dLV) model of population dynamics (sometimes called the Ricker model) relates the abundance of species  $i$  at a time  $t + \delta t$  ( $x_i(t + \delta t)$ ) to the abundances of *all* the species in the ecosystem at a time  $t$  ( $\vec{x} = \{x_1(t) \dots x_N(t)\}$ ). These interactions are encoded in the dLV through a set of interaction coefficients,  $c_{ij}$ , that describe the influence species  $j$  has on the abundance of species  $i$  [19], and inferring these interaction coefficients is the major goal of this work. The abundance of a species can also change due to environmental and demographic stochastic effects. The dLV can be generalized to include stochasticity by including a log-normally distributed multiplicative noise,  $\eta_i(t)$ . Specifically, the dynamics are modeled by the equations

$$x_i(t + \delta t) = \eta_i(t)x_i(t) \exp\left(\delta t \sum_j c_{ij}(x_j(t) - \langle x_j \rangle)\right), \quad (2)$$

where  $\langle x_j \rangle$  is equilibrium abundance of species  $j$  and is set by the carrying capacity of the environment. In writing these equations in this form, we have assumed that in the absence of noise the dLV equations have a unique steady-state solution with the abundances given by  $\langle x_j \rangle$ . Notice that in the absence of multiplicative noise and the limit  $\delta t \rightarrow 0$ , the dLV model reduces to the usual Lotka-Volterra differential equations:

$$\frac{1}{x_i(t)} \frac{dx_i(t)}{dt} = \sum_j c_{ij}(x_j(t) - \langle x_j \rangle). \quad (3)$$

In what follows, without loss of generality, we set  $\delta t = 1$ . This is equivalent to measuring time in units of  $\delta t$ . To fit microbial data, it is actually helpful to work with the logarithm of Equation (2). Furthermore, we assume that the sampling time and the update time are both equal to 1. Taking the logarithm of Eq. 2 yields,

$$\ln x_i(t + 1) - \ln x_i(t) = \zeta_i(t) + \sum_j c_{ij}(x_j(t) - \langle x_j \rangle), \quad (4)$$

where by construction  $\zeta_i(t) = \ln \eta_i(t)$  is a normally distributed variable. This logarithmic form of the dLV model is especially convenient for inferring species interactions from time series of species abundances because the inference problem reduces to standard linear regression (as discussed below).

Thus, far we have assumed that it is possible to directly measure the absolute abundances  $x_i(t)$ . However, in practice, metagenomic sequencing studies typically provide relative abundances  $\tilde{x}_i = Z^{-1}x_i$ , where  $Z = \sum_j x_j$ . Provided that the number of species is large ( $N \gg 1$ ), the fluctuations in the total population size  $Z(t) = \sum_i x_i(t)$  around its mean value  $\langle Z \rangle = \sum_i \langle x_i \rangle$  will be small. In this case, the dynamics of the relative abundances ( $\tilde{x}_i(t)$ ) are well-described by a modified dLV model:

$$\tilde{x}_i(t + 1) \approx \eta_i(t)\tilde{x}_i(t) \exp\left(\sum_j \tilde{c}_{ij}(\tilde{x}_j(t) - \langle \tilde{x}_j \rangle)\right), \quad (5)$$

where we have defined the new interaction coefficients  $\tilde{c}_{ij} = \langle Z \rangle c_{ij}$  which are related to the true interaction coefficients  $c_{ij}$  by the average population size. Thus, relative species abundance data can be modeled using the dLV model, but the interaction coefficients are known only up to an arbitrary multiplicative constant.



However, as discussed in the main text, the design matrix for relative species abundances is singular so simple linear regression fails.

In all of the simulations discussed in the main text the stochasticity was set to  $\ln \eta_i(t) \sim \mathcal{N}(0, 0.1)$ .

### Linear Regression

Suppose we are given data consisting of the absolute (or relative) abundances of the species in a population of  $N$  species in the form of timeseries of length  $T$  starting from  $M$  initial conditions. We infer each row ( $\vec{c}_i = \{c_{ij}\}_{j=1}^N$ ) of the interaction coefficient matrix ( $C$ ) separately. The equilibrium population  $\langle x \rangle$  (or  $\langle \tilde{x} \rangle$ ) is assumed to correspond to the population median. The design matrix is an  $M(T-1) \times N$  matrix with rows  $X_I = \{x_1^{(k)}(t) - \langle x_1 \rangle, \dots, x_N^{(k)}(t) - \langle x_N \rangle\}$ . The data vector has length  $M(T-1)$  and is given by  $\vec{v}_i = \{\ln x_i^{(k)}(1) - \ln x_i^{(k)}(0), \dots, \ln x_i^{(k)}(T) - \ln x_i^{(k)}(T-1)\}$ . Note that any of the timepoints with  $x_i^{(k)}(t) = 0$  were left out of the regression because the logarithm of zero is undefined. The least squares estimate for the interaction coefficients is  $\hat{c}_i = X^+ \vec{v}_i$ , where the  $+$  denotes the pseudo-inverse.

### Outline of the LIMITS Algorithm

Here, we present a high-level outline of the LIMITS algorithm (see Fig. 2). An implementation of the LIMITS algorithm written in Mathematica (Wolfram Research, Inc.) is available in the Supporting Information and on the author's webpage at <http://physics.bu.edu/~pankajm>. Since all of the regressions are performed independently for each species, we will only describe the algorithm for inferring a row of the interaction matrix ( $\vec{c}_i = \{c_{ij}\}_{j=1}^N$ ). One simply loops over  $i$  to obtain the full interaction matrix. Moreover, bootstrap aggregating simply involves performing the whole procedure  $L$  times, thereby constructing multiple estimates  $\vec{c}_i^{(1)}, \dots, \vec{c}_i^{(L)}$  and taking their median. Thus, the only thing that takes some effort to explain is how to construct one of estimate of  $\vec{c}_i$ .

1. First, we randomly partition the data into a training set and a test set, each containing half the data points.
2. A set of active coefficients is initialized to ACTIVE =  $\{c_{ii}\}$  and a set of inactive coefficients is initialized to INACTIVE =  $\{c_{ij}\}_{j \neq i}$ .
3. A linear regression including only species  $i$  is performed on the training set, and the inferred coefficient is used to calculate a prediction error (called ERROR) for the test dataset.
4. For each coefficient  $c$  in INACTIVE create TEST = ACTIVE  $\cup \{c\}$  and perform a linear regression using the coefficients in TEST against the training dataset.
5. Next, the inferred coefficients are used to calculate the prediction errors for the test dataset. The particular species  $j$  with the smallest prediction error is retained, and we call this error ERROR( $j$ ).
6. If  $100 \times (\text{ERROR} - \text{ERROR}(j)) / \text{ERROR}$  is greater than a pre-specified required improvement in prediction then we set ERROR = ERROR( $j$ ), ACTIVE = ACTIVE  $\cup \{c_{ij}\}$ , and  $c_{ij}$

### References

1. Turnbaugh PJ, Ley RE, Hamady M, Fraser-Liggett CM, Knight R, et al. (2007) The human microbiome project. *Nature* 449: 804–810.
2. Yatsunenko T, Rey FE, Manary MJ, Trehan I, Dominguez-Bello MG, et al. (2012) Human gut microbiome viewed across age and geography. *Nature* 486: 222–227.
3. Caporaso JG, Lauber CL, Costello EK, Berg-Lyons D, Gonzalez A, et al. (2011) Moving pictures of the human microbiome. *Genome Biol* 12: R50.
4. Ridaura VK, Faith JJ, Rey FE, Cheng J, Duncan AE, et al. (2013) Gut microbiota from twins discordant for obesity modulate metabolism in mice. *Science* 341: 1241214.

is deleted from INACTIVE, otherwise we terminate the loop and return an estimate for the interactions  $\vec{c}_i = \{c_{ij}\}_{j=1}^N$ .

### Generating Test Interaction Matrices

Interaction matrices (of size  $10 \times 10$  for 10 species) were randomly generated to ensure that they were sparse, and that they resulted in a stable internal equilibrium. First, the equilibrium abundances were drawn from a lognormal distribution with median 1 and scale parameter 0.1. The diagonal elements of the matrix were randomly drawn from a uniform distribution over the interval  $c_{ii} \sim (-1.9/\bar{x}_i, -0.1/\bar{x}_i)$ , where  $\bar{x}_i$  is the randomly chosen equilibrium abundance of species  $i$ . Next, random  $c_{ij}$  parameters were added to the matrix one at a time as long as the Lotka-Volterra system remained stable until 10 off diagonal interactions had been added. The algorithm we used for generating random interaction matrices is provided in Code S1.

### Supporting Information

**Code S1 Code for the LIMITS algorithm and an example implemented in Mathematica (Wolfram Research, Inc.) Version 8.0.**

(NB)

**Figure S1 Example Lotka-Volterra Time Series.** Representative simulations of Lotka-Volterra time series with a) a symmetric interaction matrix and b) an asymmetric interaction matrix.

(TIFF)

**Figure S2 Errors in Interaction Topology as a Function of Sample Size.** The errors in inferred interaction topologies as a function of sample size for a) a symmetric interaction matrices and b) an asymmetric interaction matrices. The black line represents the true positive rate, the dashed blue line represents the false negative rate, and the red line represents the rate of sign errors.

(TIFF)

**Figure S3 Interaction topologies of abundant species in the guts of two individuals using different prediction error thresholds.** The size of a node denotes the median relative species abundance, beneficial interactions are shown as solid red arrows, and competitive interactions are shown as dashed blue arrows. The 14 species included in the model were obtained by taking the union of the top 10 most abundant species from individuals a and b.

(TIFF)

### Acknowledgments

We would like to acknowledge useful conversations with Sara Collins and Daniel Segré and thank Alex Lang and Javad Noorbakhsh for comments on the manuscript.

### Author Contributions

Conceived and designed the experiments: CKF PM. Analyzed the data: CKF. Wrote the paper: CKF PM.

5. Moore W, Moore LH (1995) Intestinal oras of populations that have a high risk of colon cancer. *Applied and Environmental Microbiology* 61: 3202–3207.
6. Ahn J, Sinha R, Pei Z, Dominianni C, Wu J, et al. (2013) Human gut microbiome and risk for colorectal cancer. *Journal of the National Cancer Institute* 105: 1907–1911.
7. Turnbaugh PJ, Hamady M, Yatsunenko T, Cantarel BL, Duncan A, et al. (2008) A core gut microbiome in obese and lean twins. *Nature* 457: 480–484.
8. Hajishengallis G, Darveau RP, Curtis MA (2012) The keystone-pathogen hypothesis. *Nature Reviews Microbiology* 10: 717–725.
9. Hsiao EY, McBride SW, Hsien S, Sharon G, Hyde ER, et al. (2013) Microbiota modulate behavioral and physiological abnormalities associated with neurodevelopmental disorders. *Cell* 155: 1451–1463.
10. Wu S, Rhee KJ, Albesiano E, Rabizadeh S, Wu X, et al. (2009) A human colonic commensal promotes colon tumorigenesis via activation of t helper type 17 t cell responses. *Nature medicine* 15: 1016–1022.
11. Klitgord N, Segrè D (2010) Environments that induce synthetic microbial ecosystems. *PLoS computational biology* 6: e1001002.
12. Faust K, Raes J (2012) Microbial interactions: from networks to models. *Nature Reviews Microbiology* 10: 538–550.
13. Stein RR, Bucci V, Toussaint NC, Buffie CG, Ratsch G, et al. (2013) Ecological modeling from timeseries inference: Insight into dynamics and stability of intestinal microbiota. *PLoS computational biology* 9: e1003388.
14. Friedman J, Alm EJ (2012) Inferring correlation networks from genomic survey data. *PLoS Computational Biology* 8: e1002687.
15. Ruan Q, Dutta D, Schwalbach MS, Steele JA, Fuhrman JA, et al. (2006) Local similarity analysis reveals unique associations among marine bacterioplankton species and environmental factors. *Bioinformatics* 22: 2532–2538.
16. Xia LC, Ai D, Cram J, Fuhrman JA, Sun F (2013) Efficient statistical significance approximation for local similarity analysis of high-throughput time series data. *Bioinformatics* 29: 230–237.
17. Fuller WA (1980) Properties of some estimators for the errors-in-variables model. *The Annals of Statistics*: 407–422.
18. Breiman L (1996) Bagging predictors. *Machine learning* 24: 123–140.
19. Hofbauer J, Hutson V, Jansen W (1987) Coexistence for systems governed by difference equations of lotka-volterra type. *Journal of Mathematical Biology* 25: 553–570.
20. Enders W (2008) *Applied econometric time series*. John Wiley & Sons.
21. Wooley JC, Godzik A, Friedberg I (2010) A primer on metagenomics. *PLoS computational biology* 6: e1000667.
22. Bishop CM, Nasrabadi NM (2006) *Pattern recognition and machine learning*, volume 1. springer New York.
23. Glass EM, Wilkening J, Wilke A, Antonopoulos D, Meyer F (2010) Using the metagenomics rast server (mg-rast) for analyzing shotgun metagenomes. *Cold Spring Harbor Protocols* 2010: pdb-prot5368.
24. Arumugam M, Raes J, Pelletier E, Le Paslier D, Yamada T, et al. (2011) Enterotypes of the human gut microbiome. *Nature* 473: 174–180.
25. Turnbaugh PJ, Ridaura VK, Faith JJ, Rey FE, Knight R, et al. (2009) The effect of diet on the human gut microbiome: a metagenomic analysis in humanized gnotobiotic mice. *Science translational medicine* 1: 6ra14.
26. Shah H, Collins M (1989) Proposal to restrict the genus *bacteroides* (castellani and chalmers) to *bacteroides fragilis* and closely related species. *International journal of systematic bacteriology* 39: 85–87.
27. MacArthur R (1970) Species packing and competitive equilibrium for many species. *Theoretical population biology* 1: 1–11.
28. Chesson P, Kuang JJ (2008) The interaction between predation and competition. *Nature* 456: 235–238.