

# Network Reconstruction via the Minimum Description Length Principle

Tiago P. Peixoto<sup>\*</sup>

*Inverse Complexity Lab, IT:U Interdisciplinary Transformation University Austria, Linz, Austria  
and Department of Network and Data Science, Central European University, Vienna, Austria*



(Received 12 August 2024; accepted 22 January 2025; published 20 March 2025)

A fundamental problem associated with the task of network reconstruction from dynamical or behavioral data consists in determining the most appropriate model complexity in a manner that prevents overfitting and produces an inferred network with a statistically justifiable number of edges and their weight distribution. The status quo in this context is based on  $L_1$  regularization combined with cross-validation. However, besides its high computational cost, this commonplace approach unnecessarily ties the promotion of sparsity, i.e., abundance of zero weights, with weight “shrinkage.” This combination forces a trade-off between the bias introduced by shrinkage and the network sparsity, which often results in substantial overfitting even after cross-validation. In this work, we propose an alternative nonparametric regularization scheme based on hierarchical Bayesian inference and weight quantization, which does not rely on weight shrinkage to promote sparsity. Our approach follows the minimum description length principle, and uncovers the weight distribution that allows for the most compression of the data, thus avoiding overfitting without requiring cross-validation. The latter property renders our approach substantially faster and simpler to employ, as it requires a single fit to the complete data, instead of many fits for multiple data splits and choice of regularization parameter. As a result, we have a principled and efficient inference scheme that can be used with a large variety of generative models, without requiring the number of reconstructed edges and their weight distribution to be known in advance. In a series of examples, we also demonstrate that our scheme yields systematically increased accuracy in the reconstruction of both artificial and empirical networks. We highlight the use of our method with the reconstruction of interaction networks between microbial communities from large-scale abundance samples involving on the order of  $10^4$ – $10^5$  species and demonstrate how the inferred model can be used to predict the outcome of potential interventions and tipping points in the system.

DOI: [10.1103/PhysRevX.15.011065](https://doi.org/10.1103/PhysRevX.15.011065)

Subject Areas: Complex Systems,  
Interdisciplinary Physics,  
Statistical Physics

## I. INTRODUCTION

Network reconstruction is the task of determining the unseen interactions between elements of a complex system when only data on their behavior are available—usually consisting of time-series or independent samples of their states. This task is required when it is either too difficult, costly, or simply impossible to determine the individual pairwise interactions via direct measurements. This is a common situation when we want to determine the associations between species only from their co-occurrence in population samples [1], the financial market dependencies

from the fluctuation of stock prices [2], protein structures from observed amino-acid contact maps [3,4], gene regulatory networks from expression data [5–7], neural connectivity from functional magnetic resonance imaging (fMRI), electroencephalography (EEG), micro electrode array (MEA), or calcium fluorescence imaging data [8–10], and epidemic contact tracing from infections [11], among many other scenarios.

In this work we frame network reconstruction as a statistical inference problem [12], where we assume that the observed data have been sampled from a generative statistical model that has a weighted network as part of its set of parameters. In this case, the reconstruction task consists of estimating these parameters from data. This generative framing contrasts with other inferential approaches based on, e.g., Granger causality [13] and transfer entropy [14] in important ways: (1) It allows for a principled framework that does not rely on *ad hoc* and often arbitrary confidence thresholds, (2) can reliably distinguish between direct and indirect interactions, to the extent that this information is

<sup>\*</sup>Contact author: [tiago.peixoto@it-u.at](mailto:tiago.peixoto@it-u.at)

*Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.*

obtainable from the data, (3) enables uncertainty quantification and comparison between alternative generative hypotheses, and (4) produces a generative model as a final result, which can be used not only for interpretation, but also to predict the outcomes of interventions and alternative conditions previously unseen in the data. These same features make this type of approach preferable also to noninferential heuristics such as correlation thresholding [15,16], known to suffer from substantial biases [12,17].

A central obstacle to the effective implementation of inferential network reconstruction is the need for statistical regularization: A network with  $N$  nodes can in principle have up to  $O(N^2)$  edges, and elementary maximum likelihood approaches do not allow us to distinguish between the nonexistence of an edge and the existence of an edge with a low weight magnitude, resulting in the inference of maximally dense networks, even when the true underlying network is sparse, i.e., its number of edges is far smaller than the number of possible edges. In this case, unregularized maximum likelihood would massively overfit, incorporating a substantial amount of statistical noise in its description, and thus would not meaningfully represent the true underlying system. Because of this, a robust inferential framework needs to include the ability to determine the most appropriate model complexity from data, which in this context means the right number of edges and their weight distribution, in a manner that does not overfit.

Perhaps the most common strategy to avoid overfitting is  $L_1$  regularization [18], where a penalty proportional to the sum of the weight magnitudes is subtracted from the likelihood, multiplied by a regularization parameter. This approach is a standard technique in statistics and machine learning, perhaps most well known as the central ingredient of LASSO regression [19], and has been used extensively as well for network reconstruction using multivariate Gaussian models [20–23], most famously as part of the GLASSO algorithm [24], but also for the inverse Ising model [25–31], among other problem instances.

The  $L_1$  regularization scheme has certain attractive properties: (1) It is easy to implement, (2) it results in the existence of exact zeros for the reconstructed weights, and (3) the  $L_1$  penalty function is convex. When combined with a likelihood function that is also convex (such as the one for the multivariate Gaussian [24] or the Ising model [29]), the latter property means that this scheme yields an overall convex objective function, thus enabling the use of very efficient algorithms available only for this simpler class of optimization problems. On the other hand, this scheme also has well-known significant limitations [18]: (1) The regularization parameter, which determines the final sparsity of the reconstruction, needs to be provided prior to inference, requiring the desired level of sparsity to be known beforehand, and (2) it causes “shrinkage” of the weights, i.e., the weights of the nonzero edges also decrease in magnitude, introducing a bias. The first disadvantage undermines this regularization approach in a central way,

since it requires as an *input* what it should in fact provide as an *output*: the network sparsity. Because of this, the scheme is often used together with cross-validation [24,32,33] to determine the optimal value of the regularization parameter. But not only does this significantly increase the computational cost of the approach, but in the end also typically results in substantial overfitting, as we demonstrate in this work.

Other regularization approaches also have been proposed for network reconstruction, such as  $\delta$  thresholding [34], decimation [35,36], EBIC [37], adaptive LASSO [38–40], MCP [41], SCAD [42], weight integration [43], alternative priors [44–46] with fixed number of edges [47], as well as spike-and-slab [48,49] and horseshoe [50–52] shapes. These alternatives address some limitations of  $L_1$ , in particular its detrimental shrinkage properties, but they all either still require cross-validation, rely on heuristics to determine the resulting sparsity, and/or make strong assumptions on the shape of the weight distribution.

To the best of our knowledge, there is currently no regularization scheme that is simultaneously (1) principled, (2) effective at preventing overfitting, (3) algorithmically efficient, and (4) nonparametric, in particular without requiring the network sparsity to be known in advance.

In this work we fill this gap by developing a Bayesian regularization scheme that follows the minimum description length (MDL) principle [53,54] and seeks to find the optimal sparsity and weight distribution of the edges in a manner that most compresses the observed data. Instead of weight shrinkage, our approach relies on weight quantization as a means of quantifying the information necessary to encode the inferred weights to a sufficient numerical precision, and also makes no strong assumption on the shape of the weight distribution, unlike  $L_1$  and most Bayesian approaches previously proposed for this problem [43–45,47–52]. The resulting approach does not require cross-validation and is also algorithmically efficient, especially when combined with a recent algorithm that can find the best potential edge candidates to be updated in sub-quadratic time and in parallel [55], making it applicable for the reconstruction of networks with hundreds of thousands of nodes, or even more depending on available computing resources. Combined with the fact that it does not rely on any property of the generative model other than it being conditioned on real-valued edge weights, our approach gives us a regularization scheme that can be readily employed on a broad class of inferential reconstruction problems.

This paper is divided as follows. In Sec. II we describe our inferential framework, and in Sec. III we discuss the problems with  $L_1$  regularization. In Sec. IV we present our alternative MDL framework, and in Sec. V the inference algorithm we propose to implement it, which we then evaluate with synthetic and empirical data. In Sec. VI we conduct an empirical case study where we reconstruct the large-scale network of interactions of microbial species

from empirical data and showcase how the inferred models can be used to predict the outcomes of interventions, as well as to identify “keystone” species and tipping points. We finalize in Sec. VII with a conclusion.

## II. INFERENCE FRAMEWORK

The general scenario for inferential network reconstruction consists of some data  $\mathbf{X}$  that are assumed to originate from a generative model with a likelihood

$$P(\mathbf{X}|\mathbf{W}), \quad (1)$$

where  $\mathbf{W} \in \mathbb{R}^{N \times N}$  is a symmetric matrix corresponding to the weights of an undirected graph of  $N$  nodes. In general, we are not interested in any particular *a priori* constraint on  $\mathbf{W}$ , although we usually expect it to be sparse, i.e., its number of nonzero entries scales as  $O(N)$ . In many cases, the data are represented by a  $N \times M$  matrix of  $M$  samples, with  $X_{im}$  being a value associated with node  $i$  for sample  $m$ , such that

$$P(\mathbf{X}|\mathbf{W}) = \prod_{m=1}^M P(\mathbf{x}_m|\mathbf{W}), \quad (2)$$

with  $\mathbf{x}_m$  being the  $m$ th column of  $\mathbf{X}$ . Alternatively, we may have that the network generates a Markovian time series with likelihood

$$P(\mathbf{X}|\mathbf{W}) = \prod_{m=1}^M P(\mathbf{x}_m|\mathbf{x}_{m-1}, \mathbf{W}), \quad (3)$$

given some initial state  $\mathbf{x}_0$ . Other variations are possible, but for our purposes we need only to refer to a generic posterior distribution,

$$P(\mathbf{W}|\mathbf{X}) = \frac{P(\mathbf{X}|\mathbf{W})P(\mathbf{W})}{P(\mathbf{X})}, \quad (4)$$

that we need to be able to compute up to a normalization constant. Our overall approach is designed to work independently of what kind of generative model is used, static or dynamic, with any type of likelihood, as long as it is conditioned on edge weights according to the above equation. In Appendix A we list some generative models that we will use as examples in this work, but for now it is simpler if we consider the generative model more abstractly.

Given Eq. (4) we can proceed in a variety of ways, for example, by sampling from it or computing expectations. In this work, we are interested in the maximum *a posteriori* (MAP) point estimate,

$$\hat{\mathbf{W}} = \arg \max_{\mathbf{W}} P(\mathbf{W}|\mathbf{X}) \quad (5)$$

$$= \arg \max_{\mathbf{W}} \log P(\mathbf{X}|\mathbf{W}) + \log P(\mathbf{W}), \quad (6)$$

corresponding to the most likely weighted network that generated the data. In this setting, regularization is implemented by choosing a suitable prior  $P(\mathbf{W})$ , which should, ideally, properly penalize overly complex models.

In the following we will discuss  $L_1$  regularization, and then present our alternative approach based on the MDL principle.

## III. $L_1$ REGULARIZATION, SHRINKAGE, AND OVERFITTING

The most common choice to regularize Eq. (6) is independent Laplace distributions,

$$P(\mathbf{W}|\lambda) = \prod_{i < j} \lambda e^{-\lambda |W_{ij}|} / 2, \quad (7)$$

which result in a  $L_1$  penalty to the log-likelihood:

$$\log P(\mathbf{W}|\lambda) = -\lambda \sum_{i < j} |W_{ij}| + \binom{N}{2} \log(\lambda/2). \quad (8)$$

This choice offers some attractive properties. First, the function  $\log P(\mathbf{W}|\lambda)$  is convex, and therefore if  $\log P(\mathbf{X}|\mathbf{W})$  is also convex with respect to  $\mathbf{W}$  (which is true for key problem instances such as the multivariate Gaussian or the Ising model), then Eq. (6) amounts to a convex optimization problem, which has a unique global solution with no other local maxima, and can be solved efficiently [55]. Second, this type of penalization promotes sparsity, i.e., as  $\lambda$  is increased, this will cause the values of  $W_{ij}$  to converge in succession exactly to zero at specific values of  $\lambda$ . This is in contrast to, e.g.,  $L_2$  regularization [56], where an increase in the penalty tends to merely reduce the magnitude of the weights continuously, without making them exactly zero. In this way, the value of  $\lambda$  amounts to a proxy for the density of the inferred network.

However, this choice also suffers from important shortcomings. From a modeling perspective, the prior of Eq. (7) corresponds to a maximum entropy distribution conditioned on the expected mean weight magnitude, including the zeros. This amounts to a prior that is conditioned on previous knowledge about the density of the network, encoded via the parameter  $\lambda$ . However, this information is rarely known in advance, and is almost always one of the primary objectives of inference in the first place. The likelihood of Eq. (7) is not helpful in this regard, since it is unbounded, and diverges toward  $\log P(\mathbf{W}|\lambda) \rightarrow \infty$  for  $W_{ij} = 0$  and  $\lambda \rightarrow \infty$ ; therefore we cannot infer  $\lambda$  via MAP estimation.

The most common solution employed in practice is to select  $\lambda$  via  $K$ -fold cross-validation [24,32,33], i.e., by dividing the data matrix in  $K$  disjoint subsets  $\mathbf{X} = \{\mathbf{X}_1, \dots, \mathbf{X}_K\}$ , and reconstructing the network with one subset  $k$  removed, i.e.,

$$\hat{\mathbf{W}}_k(\lambda) = \arg \max_{\mathbf{W}} \log P(\mathbf{X} \setminus \mathbf{X}_k | \mathbf{W}) + \log P(\mathbf{W} | \lambda), \quad (9)$$

and using the held-out data to compute the log-likelihood  $\mathcal{L}_k(\lambda) = \log P(\mathbf{X}_k | \hat{\mathbf{W}}_k(\lambda))$ . Finally,  $\lambda$  is chosen to maximize the mean held-out log-likelihood:

$$\hat{\lambda} = \arg \max_{\lambda} \sum_k \mathcal{L}_k(\lambda). \quad (10)$$

This standard technique is designed to prevent overfitting, since it relies on the predictive performance of the inferred model  $\hat{\mathbf{W}}_k(\lambda)$  with respect to the held-out data  $\mathbf{X}_k$ , which should deteriorate as soon as it incorporates more statistical noise. Although widespread, this solution still suffers from some key limitations. First, it increases the computational cost by at least 1 or 2 orders of magnitude, since we need to perform the optimization of Eq. (9)  $K$  times (a typical choice is  $K = 5$ ) and for many values of  $\lambda$ —usually via a combination of grid and bisection search. Furthermore, and even more importantly, this whole procedure does not sufficiently reconcile the prevention of overfitting with the promotion of sparsity, as we now demonstrate.

For our analysis it will be useful to evaluate the accuracy of a reconstruction  $\hat{\mathbf{W}}$  via the Jaccard similarity  $s(\mathbf{W}, \hat{\mathbf{W}}) \in [0, 1]$ ,

$$s(\mathbf{W}, \hat{\mathbf{W}}) = 1 - \frac{\sum_{i < j} |W_{ij} - \hat{W}_{ij}|}{\sum_{i < j} |W_{ij}| + |\hat{W}_{ij}|}, \quad (11)$$

and likewise for the binarized adjacency matrix,  $A_{ij}(\mathbf{W}) = \{1 \text{ if } |W_{ij}| > 0, \text{ else } 0\}$ , with  $s(\mathbf{A}(\mathbf{W}), \mathbf{A}(\hat{\mathbf{W}}))$ .

In Fig. 1 we consider the reconstruction of Zachary's karate club network [57], with  $N = 34$  nodes and  $E = 78$  edges serving as the nonzero couplings of a kinetic Ising model (see Appendix A), with each corresponding value of  $\mathbf{W}$  sampled from a normal distribution with mean  $N/2E = 0.22$  and standard deviation 0.01. After sampling a single time series of  $M = 1000$  transitions given a random initial state, we performed the reconstruction for a range of values of  $\lambda$ , including also a fivefold cross-validation procedure for each value. As seen in Fig. 1(a), the increase in  $\lambda$  simultaneously sparsifies and shrinks the magnitude of the inferred edge weights. The maximum of the mean held-out log-likelihood obtained from cross-validation is achieved for a value  $\lambda \approx 30.8$ , for a network with  $E = 278$  nonzero weights, shown in Fig. 1(b). The weighted and unweighted similarities with the true network are 0.83 and 0.44, respectively, affected primarily by the abundance of spurious edges incorporated into the reconstruction. If we increase instead to  $\lambda \approx 246.6$ , the similarities become 0.67 and 0.99, respectively, with  $E = 79$ —much closer to the true value—illustrating the inherent trade-off between weight magnitude preservation and sparsity. Although for this value of  $\lambda$  the binarized network is much more




FIG. 1.  $L_1$  regularization overfits when combined with cross-validation. This example considers the reconstruction of the weighted karate club network [57] ( $N = 34$  nodes and  $E = 78$  edges, weight values sampled i.i.d. from a normal distribution with mean 0.22 and standard deviation 0.01), based on  $M = 1000$  transitions from the kinetic Ising model with a random initial state, using Eqs. (6) and (7), for a range of values of the regularization strength  $\lambda$ . (a) Top to bottom: the Jaccard similarity between inferred and true weights and binarized edges, the mean held-out likelihood for a fivefold cross-validation, the number of inferred nonzero edges, and the individual values of inferred weights (with true nonzero edges shown in blue, and true zero-valued entries shown in red). The gray horizontal lines at the right-hand margin of the bottom of the panel show the true weight values. The vertical dashed lines mark the values of  $\lambda$  that maximize (b) the mean held-out likelihood and (c) the binarized Jaccard similarity. The inferred network for these two values of  $\lambda$  are shown in (b) and (c), respectively, with edge weights represented as thickness and the color representing whether it is a true (blue) or spurious (red) edge. In (a), all dashed horizontal lines, as well as the red line at the right-hand margin of the bottom panel, mark the results obtained with the MDL regularization of Sec. IV.



accurate, the average weight value is only 0.11—around half of the true value—which significantly degrades the model’s predictive performance. (We observe that without knowing the true network, it would not have been possible to determine the value of  $\lambda$  that yields the largest unweighted similarity, since neither the model likelihood nor the cross-validation procedure offer any indication in this regard.)

It is clear from this example that the  $L_1$  regularization obtained with the Laplace prior forces an unnecessary dichotomy between sparsity and low weight magnitudes. The detrimental bias of  $L_1$  to obtain regression coefficients is well known [42] and alternatives such as adaptive LASSO [38–40], MCP [41], and SCAD [42] exist which mitigate this problem, but they all involve *ad hoc* parameters that need to be determined via cross-validation, without the algorithmic attractiveness of  $L_1$ , since these alternatives are no longer convex. As a result, these alternatives are substantially more cumbersome to use.

Ideally, we should be able to determine which edges are better modeled by setting them to zero, without inherently reducing the magnitude of the nonzero ones. This is an inference problem on its own, and instead of committing beforehand to which kind of weight distribution is more appropriate, a more robust approach consists in formulating an open-ended class of distributions, and learning its precise shape *a posteriori* from the data, following the principle of maximum parsimony. This is precisely what we achieve in the next section.

#### IV. WEIGHT QUANTIZATION AND THE MINIMUM DESCRIPTION LENGTH PRINCIPLE

Our purpose is to implement a principled regularization scheme that does not rely on weight shrinkage, promotes sparsity, and obviates the need for cross-validation. Ideally, we would like a choice of prior  $P(\mathbf{W})$  that achieves this directly via the MAP estimation:

$$\hat{\mathbf{W}} = \arg \max_{\mathbf{W}} \log P(\mathbf{X}|\mathbf{W}) + \log P(\mathbf{W}). \quad (12)$$

More precisely, we want to evoke the inherent connection between Bayesian inference and data compression by interpreting the joint likelihood  $P(\mathbf{X}, \mathbf{W})$  in terms of the description length [53,54],

$$\Sigma(\mathbf{X}, \mathbf{W}) = -\log P(\mathbf{X}, \mathbf{W}) \quad (13)$$

$$= -\log P(\mathbf{X}|\mathbf{W}) - \log P(\mathbf{W}), \quad (14)$$

where the first term in the right-hand side of the last equation corresponds to the amount of information required (e.g., in bits if base 2 is used for the logarithm) to encode the data  $\mathbf{X}$  when the parameters  $\mathbf{W}$  are known, and the second term corresponds to the amount of information

required to encode the parameters  $\mathbf{W}$ . In this setting, the prior for  $\mathbf{W}$  should act as a penalty counteracting the likelihood term that prevents overly complex models from being inferred. The result of the optimization of Eq. (12) would then amount to the most compressive model—i.e., with the shortest description length—striking the ideal balance between quality of fit and model complexity, extirpating as much as possible statistical noise from the model description.

However, the above equivalence with compression is not valid when the weights  $\mathbf{W}$  are modeled as continuous values with infinite precision. In this case, their prior  $P(\mathbf{W})$  is necessarily a probability *density* function, and hence  $-\log P(\mathbf{W})$  cannot be interpreted as information content [58]. This is exactly the reason why we cannot determine the value of  $\lambda$  of the Laplace prior of Eq. (7) via MAP estimation: The most likely value is *not* necessarily the most compressive one. In order to properly quantify the model complexity, we need be to be explicit about the precision with which we describe the parameters  $\mathbf{W}$ , which means that they must be discretely distributed according to a probability *mass* function.

With the objective of precisely quantifying the model complexity, we proceed in several steps. First, we account for the sparsity of the model by introducing an auxiliary variable  $\mathbf{A}$  corresponding to the binary adjacency matrix that determines the values of  $\mathbf{W}$  that are going to be nonzero. At first, we assume  $\mathbf{A}$  is sampled from a uniform distribution conditioned on the total number of edges  $E$ ,

$$P(\mathbf{A}|E) = \frac{\delta_{\sum_{i < j} A_{ij}, E}}{\binom{N}{2}}, \quad (15)$$

and the value of  $E$  is itself sampled from a uniform prior  $P(E) = 1/[\binom{N}{2} + 1]$ —ultimately an unimportant constant added simply for completeness. This choice will allow us to benefit from sparsity, as it will lead to shorter description lengths for graphs with fewer edges. Given  $\mathbf{A}$ , we then sample the weights  $\mathbf{W}$  as

$$P(\mathbf{W}|\mathbf{A}) = \prod_{i < j} P(W_{ij})^{A_{ij}} \delta_{W_{ij}, 0}^{1-A_{ij}}, \quad (16)$$

where  $P(W_{ij})$  is the probability of a nonzero weight  $W_{ij}$ , which is the same for all  $(i, j)$ . Importantly, in order to account for precision,  $P(W_{ij})$  must be a probability mass function with a support over a discrete set of values. In particular, we may, for example, consider the weights to be the outcome of a quantization procedure,

$$W_{ij} = \Delta \left\lceil \frac{Y_{ij}}{\Delta} \right\rceil, \quad (17)$$

where  $Y_{ij}$  is some continuous auxiliary value,  $\lceil \cdot \rceil$  is the round operation, and  $\Delta$  represents the precision considered, such that we have a conditional probability mass  $P(W_{ij}|Y_{ij}, \Delta) = \delta_{W_{ij}, \Delta \lceil Y_{ij}/\Delta \rceil}$ . Note that we do not assume that the true values of  $W_{ij}$  are actually sampled in this way; instead this prior simply articulates the exact precision  $\Delta$  with which we are performing the inference.

We can then proceed by choosing a continuous distribution for the auxiliary values  $Y$ , e.g., a Laplace distribution,

$$P(Y_{ij}|\lambda) = \lambda e^{-\lambda|Y_{ij}|}/2, \quad (18)$$

in which case we obtain a probability mass,

$$\begin{aligned} P(W_{ij}|\lambda, \Delta) &= \frac{\int P(W_{ij}|Y_{ij}, \Delta) P(Y_{ij}|\lambda) dY_{ij}}{1 - \int P(0|Y_{ij}, \Delta) P(Y_{ij}|\lambda) dY_{ij}} \\ &= \begin{cases} e^{-\lambda|W_{ij}|} (e^{\lambda\Delta} - 1)/2 & \text{if } W_{ij} = \Delta \lceil W_{ij}/\Delta \rceil \\ 0 & \text{otherwise,} \end{cases} \end{aligned} \quad (19)$$

where we have excluded the value  $W_{ij} = 0$  to be consistent with our parametrization based on  $\mathbf{A}$ . Stopping at this point we would have solved two issues: (1) Via quantization we can now properly evaluate the description length of the weights and (2) we can benefit from sparsity in a manner that is independent from the weight magnitudes. However, by choosing the Laplace prior, we remain subject to weight shrinkage, since the reduction of the description length is still conditioned on their overall magnitudes—besides the discretization, Eq. (19) still amounts to  $L_1$  regularization. Therefore, our final step is to add one more level to our Bayesian hierarchy, and try to learn the weight distribution from the data, instead of assuming that the probability decays according to the Laplace distribution. For that, we assume that the discretized weights are sampled directly from an arbitrary discrete distribution,

$$P(W_{ij}|\mathbf{p}, \mathbf{z}) = \sum_{k=1}^K \delta_{W_{ij}, z_k} p_k, \quad (20)$$

where  $\mathbf{z} = \{z_k\}$  is a set of  $K$  real values representing discrete weight categories, and  $\mathbf{p} = \{p_k\}$  are their corresponding probabilities, with  $\sum_k p_k = 1$ . From this we obtain

$$P(\mathbf{W}|\mathbf{p}, \mathbf{z}, \mathbf{A}) = \prod_k p_k^{m_k} \times \prod_{i < j} \delta_{W_{ij}, 0}^{1-A_{ij}}, \quad (21)$$

with  $m_k = \sum_{i < j} A_{ij} \delta_{W_{ij}, z_k}$  being the counts of weight category  $k$ , and with a marginal distribution

$$P(\mathbf{W}|\mathbf{z}, \mathbf{A}) = \int P(\mathbf{W}|\mathbf{p}, \mathbf{z}, \mathbf{A}) P(\mathbf{p}) d\mathbf{p} \quad (22)$$

$$= \frac{\prod_k m_k!}{E!} \times \binom{K+E-1}{E}^{-1} \times \prod_{i < j} \delta_{W_{ij}, 0}^{1-A_{ij}}, \quad (23)$$

assuming a uniform prior density  $P(\mathbf{p}) = (K-1)!$  over the  $K$  simplex. The latter distribution allows for a combinatorial interpretation: It corresponds to a uniform distribution among the  $E$  edges of the weights with fixed counts  $\mathbf{m} = \{m_k\}$ , with probability  $P(\mathbf{W}|\mathbf{m}) = (\prod_k m_k! / E!) \prod_{i < j} \delta_{W_{ij}, 0}^{1-A_{ij}}$ , and the uniform distribution of the counts  $\mathbf{m}$  themselves, with probability  $P(\mathbf{m}|E) = \binom{K+E-1}{E}^{-1}$ . This interpretation allows us to fix a minor inconvenience of this model, namely that it allows for a weight category  $z_k$  to end up empty with  $m_k = 0$ , which would require us to consider category labels that never occur. Instead, we can forbid this explicitly by using a uniform prior  $P(\mathbf{m}|E) = \binom{E-1}{K-1}^{-1}$  over strictly nonzero counts  $\mathbf{m}$ , leading to a slightly modified model:

$$P(\mathbf{W}|\mathbf{z}, \mathbf{A}) = \frac{\prod_k m_k!}{E!} \times \binom{E-1}{K-1}^{-1} \times \prod_{i < j} \delta_{W_{ij}, 0}^{1-A_{ij}}. \quad (24)$$

Finally, we need a probability mass for the weight categories  $\mathbf{z}$  themselves. At this point, however, the importance of this choice is minor, since we expect  $K \ll E$ . A simple choice is to evoke the quantized Laplace once more, with

$$P(z_k|\lambda, \Delta) = \begin{cases} e^{-\lambda|z_k|} (e^{\lambda\Delta} - 1)/2 & \text{if } z_k = \Delta \lceil z_k/\Delta \rceil \\ 0 & \text{otherwise,} \end{cases} \quad (25)$$

with  $P(\mathbf{z}|\lambda, \Delta, K) = \prod_{k=1}^K P(z_k|\lambda, \Delta)$ . Although this choice implies that there will be a residual amount of shrinkage with respect to the weight categories, it will not dominate the regularization, unless the weight categories would otherwise diverge—something that can happen with certain kinds of graphical models such as the multivariate Gaussian when pairs of nodes have exactly the same values in  $\mathbf{X}$ . Because of this, we found this residual shrinkage to be beneficial in such corner cases, and otherwise we do not observe it to have a noticeable effect [59].

Lastly, we need a prior over the number of weight categories  $K$ , which we assume to be uniform in the allowed range, with  $P(K|\mathbf{A}) = (1 - \delta_{E,0})/E + \delta_{E,0}\delta_{K,0}$ . Putting it all together we have

$$\begin{aligned} P(\mathbf{W}|\lambda, \Delta) &= \sum_{\mathbf{z}, K, \mathbf{A}, E} P(\mathbf{W}|\mathbf{z}, \mathbf{A}) P(\mathbf{z}|\lambda, \Delta, K) \\ &\quad \times P(K|\mathbf{A}) P(\mathbf{A}|E) P(E) \end{aligned} \quad (26)$$

$$= \frac{\prod_k m_k! \times e^{-\lambda \sum_k |z_k|} (e^{\lambda \Delta} - 1)^K}{E! \binom{E-1}{K-1} 2^K \max(E, 1) \binom{N}{E} \left[ \binom{N}{2} + 1 \right]}, \quad (27)$$

where the remaining quantities  $\mathbf{m}$ ,  $\mathbf{z}$ ,  $K$ ,  $\mathbf{A}$ , and  $E$  in Eq. (27) should be interpreted as being functions of  $\mathbf{W}$ . Note that this model recovers the one using Eq. (19) when  $K = E$ , i.e., every weight belongs to its own category, up to an unimportant factor  $1/(E+1)$ —although we expect to obtain significantly smaller description length values with Eq. (27), since it can exploit regularities in the weight distribution and determine the most appropriate precision of the weights. We emphasize that, differently from Eq. (19), in Eq. (27) the parameter  $\Delta$  controls the precision of the weight categories  $\mathbf{z}$ , not directly of the weights themselves. The precision of the weights  $\mathbf{W}$  is primarily represented by the actual finite categories  $\mathbf{z}$ .

The model above is conditioned on two hyperparameters,  $\lambda$  and  $\Delta$ , which in principle can both be optimized, assuming a sufficiently uniform prior over them (therefore, contributing to a negligible constant to the description length). In practice, the effect of such an optimization is marginal, and it is more convenient to select a small value for  $\Delta$ , such as  $\Delta = 10^{-8}$ , which allows us to treat the weight categories as continuous for algorithmic purposes. For  $\lambda$  we use an initial value of  $\lambda = 1$ , and optimize if necessary.

We observe that we have  $-\log P(\mathbf{z}|\lambda, \Delta) = O(K\Delta)$  and  $-\log P(\mathbf{W}|\mathbf{z}, \mathbf{A}) = O(E \log E - E \log K)$ , and therefore we will typically have  $K \ll E$  when the description length is minimized, in which case the prior for  $\mathbf{z}$  will have only a minor contribution to the overall description length. Therefore, our approach is not very sensitive to alternative choices for the weight category distribution  $P(\mathbf{z}|\lambda, \Delta)$ . We observe also that the dominating term  $-\log P(\mathbf{W}|\mathbf{z}, \mathbf{A})$  from Eq. (24) is completely invariant to bijective transformations of the weights, provided the categories  $\mathbf{z}$  are transformed in the same manner. Therefore, any scale or shape dependence is driven solely by the relatively minor contribution of  $\log P(\mathbf{z}|\lambda, \Delta)$ , proportional only to the number of weight categories. Since the scale parameter  $\lambda$  can be optimized, this means that our approach does not impose any constraints on the scale and only very weak constraints on the shape of the weights (see also Ref. [59]).

Differently from the Laplace prior of Eq. (7), Eq. (27) implements a nonparametric regularization of the weights—not because it lacks parameters, but because the shape of the distribution and number of parameters adapt to what can be statistically justified from the data. In particular, the number and position of the weight categories, as well as the number of edges, will grow or shrink depending on whether they can be used to compress the data, and provide a sufficiently parsimonious network reconstruction, without overfitting.

The prior of Eq. (27) incorporates the properties we initially desired; i.e., (1) it penalizes model complexity according to the description length, (2) it exploits and

promotes sparsity, and (3) it does not rely on weight shrinkage. However, these improvements come at a cost: The prior is no longer a convex function on  $\mathbf{W}$ . Because of this, this approach requires special algorithmic considerations, which we address in the following.

## V. INFERENCE ALGORITHM

Whenever Eq. (12) corresponds to a convex optimization objective, it can be performed with relatively straightforward algorithms. The simplest one is coordinate descent [61], in which every entry in the  $\mathbf{W}$  matrix is optimized in sequence, while keeping the others fixed. This yields an algorithm with a complexity of  $O(\tau N^2)$ , provided the optimization of each entry can be done in time  $O(1)$  (e.g., via bisection search), and  $\tau$  is the number of iterations necessary for sufficient convergence. A recent significant improvement over this baseline was proposed in Ref. [55], where an algorithm was presented that solves the same optimization problem in subquadratic time, typically with a  $O(N \log^2 N)$  complexity, although the precise scaling will depend on the details of the data in general. This algorithm works as a greedy version of coordinate descent, where a subset of the entries are repeatedly proposed to be updated according to the NNDescent algorithm [62]. The latter implements a stochastic search that involves keeping a list of candidate edges incident on each node in the network as an endpoint, and updating this list by investigating the list of second neighbors in the candidate network. This algorithm typically converges in log-linear time, allowing the overall reconstruction algorithm also to converge in subquadratic time, in this way converging orders of magnitude faster than the quadratic baseline of the coordinate descent algorithm.

Although these algorithms cannot be used unmodified to our problem at hand, due to its nonconvexity and the presence of additional latent variables, we can use them as building blocks. We will consider different types of updates, according to the different types of latent variable we are considering, given our initial  $\mathbf{W}$  (initially empty with  $W_{ij} = 0$ ) and our set of weight categories  $\mathbf{z}$  (also empty initially), which are applied in arbitrary order, until the description length no longer improves.

*Edge updates.* We use the algorithm of Ref. [55] to find the  $\kappa N$  best zero-valued entries of  $\mathbf{W}$  as update candidates, usually with  $\kappa = 1$ , according to the pairwise ranking score,

$$\pi_{ij} = \max [\ln P(\mathbf{W}^+ | X), \ln P(\mathbf{W}^- | X)], \quad (28)$$

where  $\mathbf{W}^+$  is the new matrix  $\mathbf{W}$  with entry  $W_{ij}$  replaced by the smallest positive weight category in  $\mathbf{z}$ , and likewise with the smallest magnitude negative value for  $\mathbf{W}^-$ . If  $\mathbf{z}$  is currently empty, we replace the score by the gradient magnitude:

$$\pi_{ij} = \left| \frac{\partial \ln P(\mathbf{W}'|X)}{\partial W'_{ij}} \right|_{\mathbf{W}'=\mathbf{W}}. \quad (29)$$

The algorithm of Ref. [55] can obtain this set in typical time  $O(N \log^2 N)$ . We call  $\mathcal{E}$  the union of this set with the currently nonzero values of  $\mathbf{W}$ . We then visit every entry  $(i, j)$  in  $\mathcal{E}$ , and try to perform an update  $W_{ij} \rightarrow W'_{ij}$ , accepting only if it improves the objective function, chosen uniformly at random from the following options.

- (1)  $W'_{ij} \in \mathbf{z}$  chosen according to a random bisection search [63] maximizing the objective function, with each midpoint sampled uniformly at a random.
- (2)  $W'_{ij} \in \mathbb{R}$  and  $W'_{ij} \notin \mathbf{z}$  chosen according to a random bisection search in the interval of weights allowed for the problem. This will introduce a new weight category with  $\mathbf{z}' = \mathbf{z} \cup \{W'_{ij}\}$ .
- (3)  $W'_{ij} \leftarrow 0$ , if  $W_{ij} > 0$ .

*Edge moves.* Given the same set  $\mathcal{E}$  obtained as described above, we iterate over the nodes  $i$  of the network with nonzero degree, and for an incident edge  $(i, j)$  chosen uniformly at random, and an entry  $(i, u)$  chosen uniformly at random from the subset of  $\mathcal{E}$  with  $W_{iu} = 0$ , we swap their values, i.e.,  $W_{iu} \leftarrow W_{ij}$  and  $W_{ij} \leftarrow 0$ , and accept if this improves the objective function.

*Edge swaps.* We chose uniformly at random two nonzero entries in  $\mathbf{W}$ ,  $(i, j)$  and  $(u, v)$  involving four distinct nodes, and we swap their end points, i.e.,  $(W_{iv}, W_{ij}) \leftarrow (W_{ij}, W_{iv})$  and  $(W_{uj}, W_{uv}) \leftarrow (W_{uv}, W_{uj})$ , and accept only in case of improvement.

*Weight category values.* We iterate over each weight category  $z_k \in \mathbf{z}$  and perform an update  $z_k \leftarrow z'_k$ , with  $z'_k$  being the result of a bisection search over the allowed values.

*Weight category distribution.* We consider the distribution of weight categories across the nonzero entries of  $\mathbf{W}$  as a clustering problem, using the merge-split algorithm described in Ref. [64], composed of the following moves, which are accepted if it improves the quality function.

- (1) Merge. Two categories  $z_k$  and  $z_l$  are removed and merged as a single new category  $z_m$ , with a value chosen with a bisection search optimizing the objective function. The number of categories reduces by one.
- (2) Split. A single category  $z_m$  is split in two new categories  $z_k$  and  $z_l$ , chosen initially uniformly random in the range  $[z_{m-1}, z_{m+1}]$ . The entries are repeatedly moved between these categories if it improves the quality function, and the category values  $z_k$  and  $z_l$  are updated according to bisection search, until convergence. The number of categories increases by one.
- (3) Merge-split. The two steps above are combined in succession, so that the entries are redistributed in two potentially new categories, and while their total number remains the same.

The workhorse of this overall procedure is the algorithm of Ref. [55] which reduces the search for candidate updates to a subquadratic time. The remaining steps of the algorithm operate on this set with linear complexity, and therefore a single “sweep” of all latent variables can be accomplished under the same algorithmic complexity as Ref. [55], since it dominates. In practice, due to the extra latent variables and optimization steps, this algorithm finishes in time around 5–10 times longer than when using  $L_1$  regularization for a single value of  $\lambda$ , but is typically much faster than doing cross-validation and scanning over many values of  $\lambda$ . A reference C++ implementation of this algorithm is freely available as part of the `graph-tool` Python library [65].

### A. Assessment of the algorithm with synthetic data

In Fig. 2 we demonstrate the use of our algorithm on synthetic data generated from empirical networks. We simulate the kinetic and equilibrium Ising models [66] (see Appendix A) with true weights sampled i.i.d. from a normal distribution with mean  $\mu = 1/\langle k \rangle$ , with  $\langle k \rangle = 2E/N$ , and standard deviation  $\sigma = 0.01$ . We compare our approach with a prior corresponding to the true distribution, i.e.,

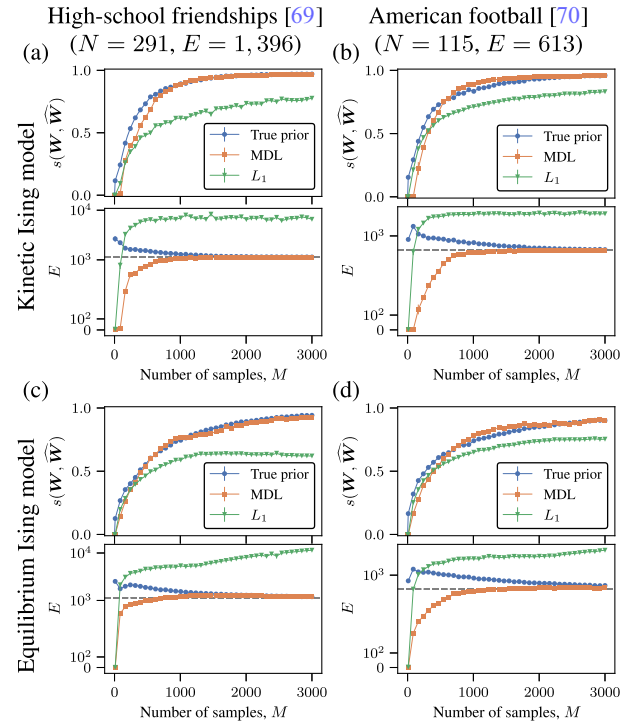


FIG. 2. Inference results for the kinetic and equilibrium Ising model for two empirical networks, as indicated in the legend, with true weights sampled as described in the text, using both the true prior, our MDL regularization, as well as  $L_1$  regularization with fivefold cross-validation. The individual panels show the Jaccard similarity  $s(\mathbf{W}, \hat{\mathbf{W}})$  between the inferred and true networks, as well as the number  $E$  of inferred nonzero edges (the dashed horizontal line shows the true value).



$$P(\mathbf{W}|\mu, \sigma, \mathbf{A}) = \prod_{i < j} \left[ \frac{e^{-(W_{ij}-\mu)^2/2\sigma^2}}{\sqrt{2\pi\sigma}} \right]^{A_{ij}} \delta_{W_{ij},0}^{1-A_{ij}}. \quad (30)$$

We observe that our MDL regularization performs very closely to the true prior, except for very few data, in which case the MDL tends to produce an empty network, whereas the true prior in fact overfits and yields a network with a number of edges larger than the true value. As we already explained, MAP estimation does not guarantee proper regularization when the priors correspond to probability densities—even if it happens to be the correct one. The overall tendency of the MDL approach is, by design, to err toward simpler models, instead of more complex ones, when the data are insufficient to provide a good accuracy.

As already discussed previously in Sec. III,  $L_1$  regularization with cross-validation performs substantially worse than our MDL approach, overfitting the data with networks that are orders of magnitude denser than they should be. Unlike using the true prior or MDL, for the examples of Fig. 2 the results obtained with  $L_1$  show no sign of converging asymptotically to true network as the data increase with  $M \rightarrow \infty$ . This behavior further highlights the general unsuitability of the  $L_1$  approach even in situations where the data are abundant. We should also mention that the results obtained with  $L_1$  took typically 1–2 orders of magnitude longer to be computed, due to the need to optimize over the parameter  $\lambda$  using a bisection search.

### B. Joint reconstruction with community detection

The use of the auxiliary sparse network  $\mathbf{A}$  and its uniform prior  $P(\mathbf{A}|\mathbf{E})$  given by Eq. (15) opens the opportunity for more structured priors that can take advantage of inferred structure for further data compression, and in this way provide improved regularization [71]. One generically applicable approach is to use the degree-corrected stochastic block model [72], here in its microcanonical formulation [73], with a likelihood

$$P(\mathbf{A}|\mathbf{b}, \mathbf{k}, \mathbf{e}) = \frac{\prod_{r < s} e_{rs}! \prod_r e_{rr}! \prod_i k_i!}{\prod_{i < j} A_{ij} \prod_i A_{ii}! \prod_r e_r!}, \quad (31)$$

where  $\mathbf{b} = \{b_i\}$  is the node partition, with  $b_i \in \{1, \dots, B\}$  being the group membership of node  $i$ ,  $\mathbf{k} = \{k_i\}$  is the degree sequence, with  $k_i$  being the degree of node  $i$ , and  $\mathbf{e} = \{e_{rs}\}$  is the group affinity matrix, with  $e_{rs}$  being the number of edges between groups  $r$  and  $s$ , or twice that if  $r = s$ . With this model, we can formulate the problem of reconstruction according to the joint posterior,

$$P(\mathbf{W}, \mathbf{A}, \mathbf{b}|\mathbf{X}) = \frac{P(\mathbf{X}|\mathbf{W})P(\mathbf{W}|\mathbf{A})P(\mathbf{A}|\mathbf{b})P(\mathbf{b})}{P(\mathbf{X})} \quad (32)$$

where the marginal distribution  $P(\mathbf{A}|\mathbf{b}) = \sum_{\mathbf{k}, \mathbf{e}} P(\mathbf{A}|\mathbf{b}, \mathbf{k}, \mathbf{e})P(\mathbf{k}|\mathbf{e})P(\mathbf{e})$  is computed using the priors described in

Ref. [73], in particular those corresponding to the hierarchical (or nested) SBM [74].

As discussed in Ref. [71], the use of this kind of structured prior can improve regularization, since it provides more opportunities to reduce the description length, whenever the underlying network is not completely random. The obtained community structure is often also useful for downstream analysis and interpretation. Algorithmically, this extension of our method is straightforward, as we need only to include merge-split partition moves [64], as done to infer the SBM.

We refer the reader to Ref. [73] for more details about the SBM formulation, and to Ref. [71] for its effect in network reconstruction.

### C. Node parameters

Besides the weighted adjacency matrix  $\mathbf{W}$ , many generative models possess additional parameters, typically values  $\boldsymbol{\theta} \in \mathbb{R}^N$  with  $\theta_i$  being a value associated with node  $i$  (e.g., a local field in the Ising model, or the mean and variance for a multivariate Gaussian). Although we could in principle include them in the diagonal of  $\mathbf{W}$ , these values often have a different interpretation and distribution, or there are more than one of them per node, and hence they need to be considered separately.

In principle, since the size of  $\boldsymbol{\theta}$  is a constant when performing network reconstruction, regularizing their inference is a relatively smaller concern. However, it would be incongruous to simply neglect this aspect in our approach. Luckily, it is a simple matter to adapt our quantized prior for  $\mathbf{W}$  into one for  $\boldsymbol{\theta}$ , the only difference being that we no longer need auxiliary variables to handle sparsity, and we can allow zeros in the discrete categories  $u_k$  for  $\boldsymbol{\theta}$ , sampled in this case as

$$P(u_k|\lambda_\theta, \Delta_\theta) = \begin{cases} e^{-\lambda_\theta|u_k|} \sinh(\lambda_\theta \Delta_\theta) & \text{if } u_k = \Delta_\theta \lceil u_k / \Delta_\theta \rceil \\ 1 - e^{-\lambda_\theta \Delta_\theta} & \text{if } u_k = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (33)$$

Proceeding as before gives us

$$P(\boldsymbol{\theta}|\lambda_\theta, \Delta_\theta) = \prod_k n_k! \times e^{-\lambda \sum_k |u_k|} \times \frac{\sinh(\lambda_\theta \Delta_\theta)^{K_\theta - \mathbb{1}_{0 \in u}} (1 - e^{-\lambda_\theta \Delta_\theta})^{\mathbb{1}_{0 \in u}}}{N! \binom{N-1}{K_\theta - 1} N}, \quad (34)$$

where  $n_k = \sum_i \delta_{\theta_i, u_k}$  are the category counts and  $K_\theta$  is the number of categories.

### D. Empirical evaluation, contrast to $L_1$ regularization, and decimation

Now we perform a comparative evaluation of our MDL regularization for empirical data where true network is not

available. We will consider the votes of all  $N = 623$  deputies in the lower house of the Brazilian congress [71], during the 2007–2011 term, as variables  $x_i \in \{-1, 0, 1\}$ , corresponding to “no,” “abstain,” and “yes,” respectively. This period corresponds to  $M = 619$  voting sessions, where all deputies cast a vote on a piece of legislation. We use a modified version of the equilibrium Ising model that allows for zero values (see Appendix A) to capture the latent coordination between individual deputies.

In Fig. 3 we show the result of the reconstruction using both our MDL scheme and  $L_1$  regularization together with fivefold cross-validation. Similarly to the karate club example of Fig. 1, the  $L_1$  inference incorporates one order of magnitude more nonzero edges ( $E = 10911$ ) when compared with the MDL method ( $E = 1540$ ). The edges obtained with  $L_1$  are mostly of low magnitude [see Fig. 3(c)], and based on our analysis for synthetic data, are much more likely to be overfitting. We can find additional evidence of this by performing cross-validation also with the MDL method, as shown in Fig. 3(d), which yields a significantly improved predictive performance of the held-out data. Both the predictive performance and description length improve when performing MDL regularization together with community detection, in accordance with Ref. [71].

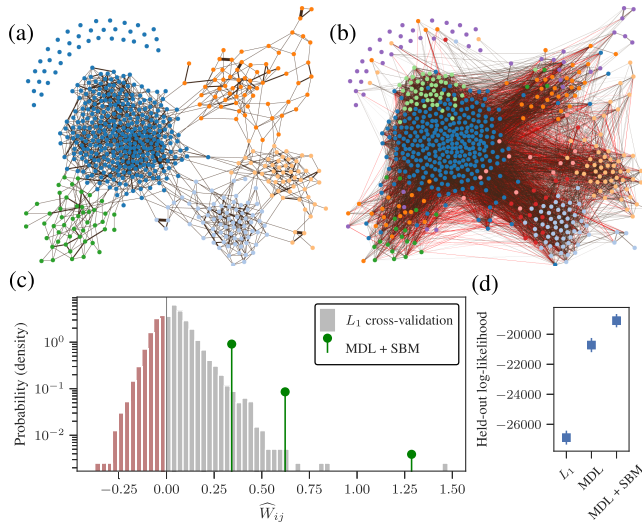


FIG. 3. Reconstructed networks of interactions between  $N = 623$  members of the lower house of the Brazilian congress [71], during the 2007–2011 term, corresponding to  $M = 619$  voting sessions. (a) Network inferred with MDL regularization, with edge weights corresponding to their thickness, and the node colors indicating the division found with the SBM incorporated into the regularization. (b) Result with  $L_1$  together with fivefold cross-validation, and negative weights shown in red. The colors indicate the group assignments found by fitting an SBM to the resulting network. (c) Weight distributions obtained with both methods and (d) the mean held-out likelihood of a fivefold cross-validation with each method, including also the MDL version without the SBM.

The interpretation of the two inferred models is fairly different. Although they share similar large-scale structures (as shown by the node colors indicating the inferred communities), the one inferred by  $L_1$  possesses weights that are often negative, indicating some amount of direct antagonism between deputies; i.e., conditioning on the “yes” vote of one deputy will increase the probability of a “no” vote of another. Instead, the model inferred by MDL possesses only positive weights, divided in three categories, indicating an overwhelming tendency toward selective agreement, instead of direct antagonism. The groups it finds also align very well with the available metadata: The largest group corresponds to the governing coalition, whereas the remaining are opposition parties forming fragmented blocks [75].

We take the opportunity to compare with another regularization scheme called “decimation,” proposed in Ref. [35]. That approach is not based on shrinkage, and instead proceeds by first performing an unregularized maximum likelihood, resulting in a full network with no zero weights, then forcing a small number of the weights with smallest magnitude to zero, and proceeding recursively on the remaining entries in the same manner, until sufficiently many entries have been set to zero. Besides the increased computational cost of inferring the weights multiple times, and starting with a full network—thus imposing an overall complexity of at least  $O(N^2)$ —this approach offers no obvious, principled criterion to determine when to stop the decimation. In Ref. [35] a heuristic stopping criterion was proposed by observing a point at which the likelihood no longer changes significantly in synthetic examples sampled from the true model. However,




FIG. 4. Decimation procedure [35] employed on the same data as Fig. 3. The bottom panel shows the maximum likelihood growing monotonically as a function of the number of nonzero edges considered during decimation, and the top panel shows the similarity (weighted and binarized) with the network inferred via MDL. The vertical dashed line corresponds to the value  $E = 24\,196$ , obtained using the stopping criterion proposed in Ref. [35].

there is no guarantee that such a signal will exist in empirical data. In Fig. 4 we show for our Brazilian congress example the model likelihood offers no clear indication of when to stop the decimation and how many edges should be inferred in the end. The Jaccard similarity with the network inferred via MDL remains low through the whole procedure, peaking at close to 1/2 for the binarized version when both networks have similar number of edges. When employing the heuristic stopping criterion proposed in Ref. [35], we obtain a network with  $E = 24196$  edges—far denser than what we obtain with MDL. Because of these limitations, we see no advantage of using decimation over our approach based on MDL.

## VI. EMPIRICAL CASE STUDY: NETWORK OF MICROBIAL INTERACTIONS

In this section we employ our regularization method for the analysis of microbial interactions [1,76–78]. We consider data from the human microbiome project (HMP) [79] and the earth microbiome project (EMP) [80], which assemble thousands of samples of microbial species abundances, collected, respectively, in various sites in the human body and in diverse habitats across the globe. More precisely, these datasets contain abundances derived from DNA sequencing, aggregated as operational taxonomic units (OTUs)—groups of individuals closely related according to their gene sequence, serving as a pragmatic proxy for individual species. The HMP dataset contains  $M = 4788$  samples involving  $N = 45\,383$  OTUs, and EMP contains  $M = 23\,323$  samples involving  $N = 126\,730$  OTUs [we used the “Silva (CR)” variant containing all samples].

The objective of our analysis is to obtain the underlying network of interactions between species from their co-occurrence in samples. To account for issues related with compositionality, i.e., the fact that the gene sequencing techniques used in these studies can provide only the relative abundance between species, instead of their absolute number, we will discard the magnitude of the read counts, and consider only the presence or absence of species, i.e., if more than zero counts have been observed for a given OTU, resulting in binary observations. We will also not attempt to model the dependence with environmental conditions and focus instead only on the mutual interactions between species. These are both substantial omissions, but our purpose is not to fully saturate the modeling requirements of this problem, but instead to demonstrate how our regularization approach can be coupled with a particular generative model to tackle a network reconstruction problem of this nature and magnitude (to the best of our knowledge, with the exception of Ref. [55], the reconstruction based on the full HMP and EMP datasets has not been performed so far).

We will model a given sample  $\mathbf{x} \in \{-1, 1\}^N$ , where  $x_i = -1$  if the species  $i$  is absent in the sample, and  $x_i = 1$  if it is present, according to the Ising model, i.e., with probability

$$P(\mathbf{x}|\mathbf{W}, \boldsymbol{\theta}) = \frac{e^{\sum_{i < j} W_{ij} x_i x_j + \sum_i \theta_i x_i}}{Z(\mathbf{W}, \boldsymbol{\theta})}, \quad (35)$$

where  $W_{ij} \in \mathbb{R}$  is the coupling strength between species  $i$  and  $j$ ,  $\theta_i$  determines the individual propensity of species  $i$  to occur in samples, independently of the other species, and  $Z(\mathbf{W}, \boldsymbol{\theta})$  is a normalization constant. A value of  $W_{ij} = 0$  indicates conditional independence between species, i.e., the presence or absence of species  $i$  does not affect the occurrence probability of species  $j$ , and vice versa, when all other species take the same value. A value  $W_{ij} > 0$  indicates cooperation, since conditioning on the presence of species  $i$  increases the probability of observing  $j$  (and vice versa), and a value  $W_{ij} < 0$  indicates antagonism, since it has the opposite effect.

We are interested in the MAP estimation,

$$\hat{\mathbf{W}}, \hat{\boldsymbol{\theta}} = \arg \max_{\mathbf{W}, \boldsymbol{\theta}} P(\mathbf{W}, \boldsymbol{\theta}|\mathbf{X}) \quad (36)$$

$$\begin{aligned} &= \arg \max_{\mathbf{W}, \boldsymbol{\theta}} \log P(\mathbf{X}|\mathbf{W}, \boldsymbol{\theta}) \\ &\quad + \log P(\mathbf{W}) + \log P(\boldsymbol{\theta}), \end{aligned} \quad (37)$$

using the MDL regularization for both  $\mathbf{W}$  and  $\boldsymbol{\theta}$  discussed previously, including also the SBM prior. (We will not include a comparison with  $L_1$ , since the cross-validation procedure it requires becomes computationally prohibitive for problems of this magnitude.)

In Fig. 5 we show the results of the reconstruction for both datasets. The inferred networks are strongly modular, composed of large node groups, connected between themselves largely by negative interactions, and a hierarchical subdivision into smaller groups (see also Fig. 6), connected between themselves by predominately positive weights. As the middle panel of Fig. 5 shows, the large-scale modular structure is strongly associated with the different body sites for the HMP network and the different habitats for EMP. This association makes it easier to interpret the negative edges between the larger groups: The different regions represent distinct microbiota, composed of specialized organisms that occupy specific ecological niches. Since we have not included dependence on environmental conditions into our model, this mutual exclusion gets encoded by negative couplings. We also observe that nodes with low degree tend to occur more seldomly (negative  $\theta_i$  values) whereas nodes with larger degrees tend to occur more frequently (positive  $\theta_i$  values).



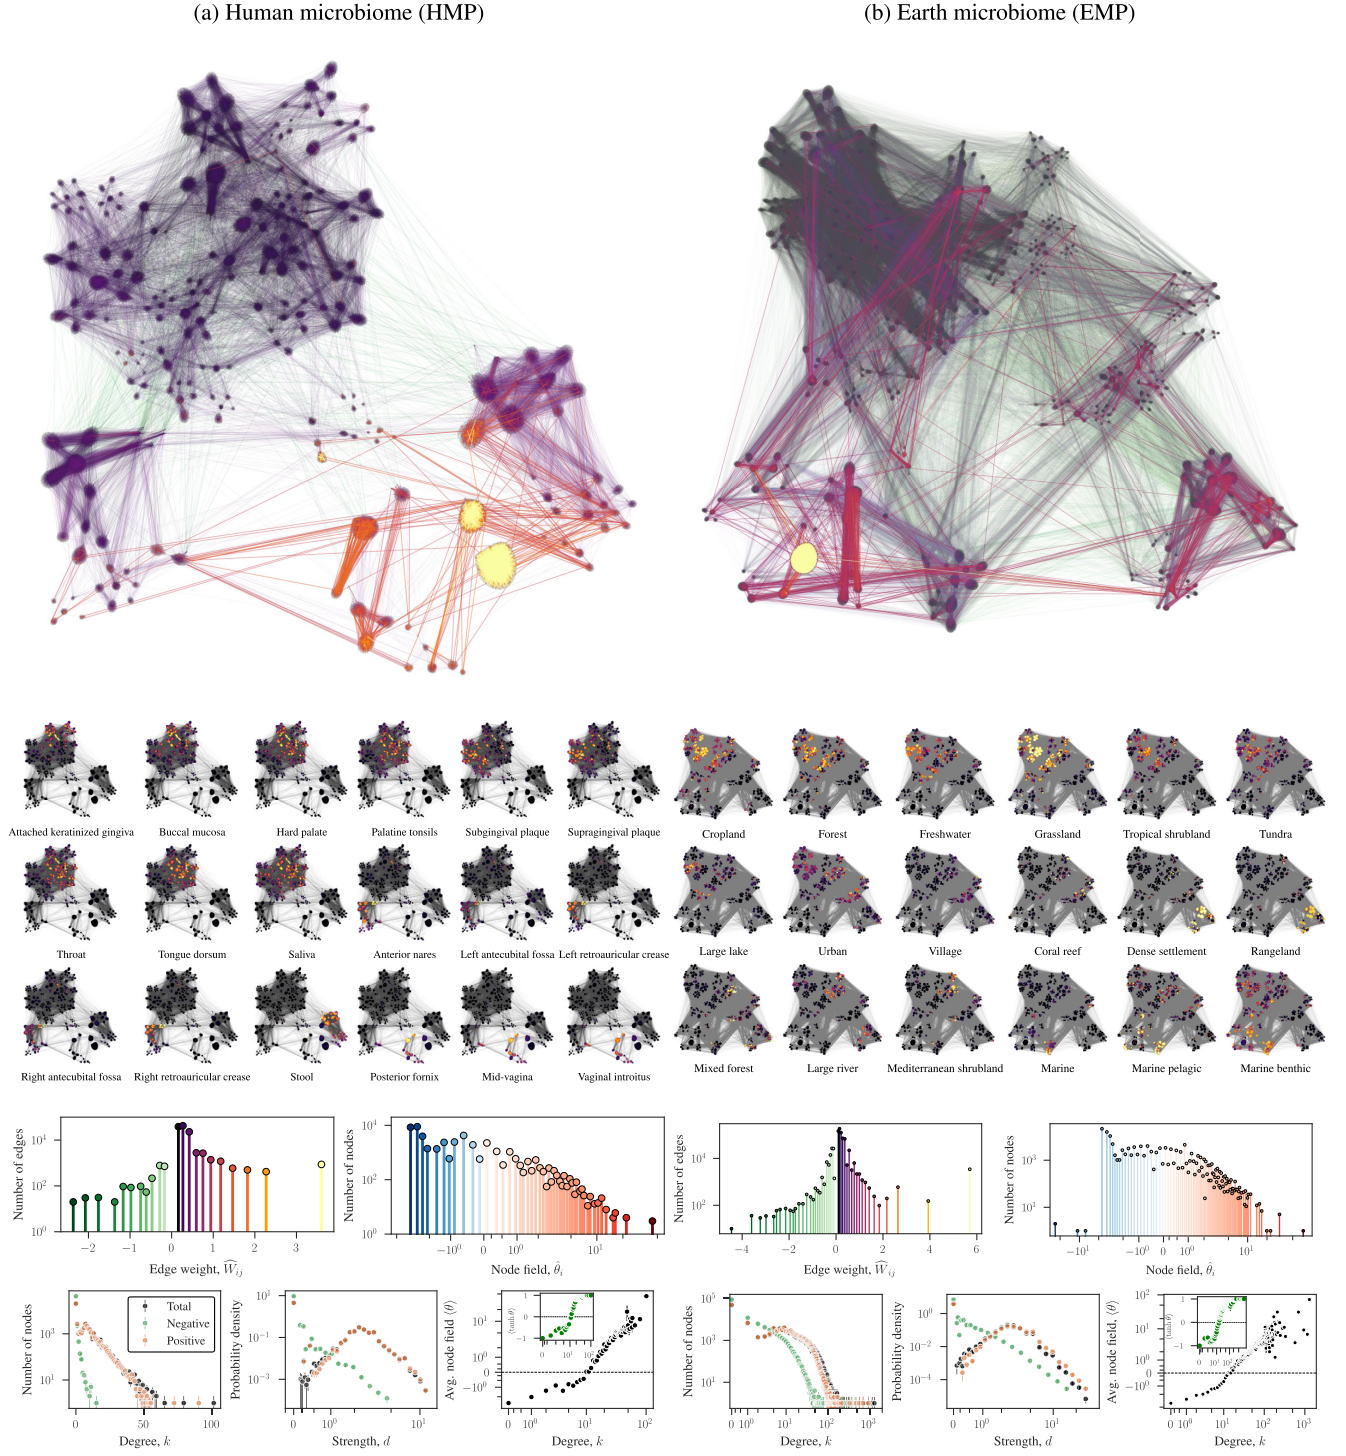


FIG. 5. Reconstructed networks for the (a) human microbiome project (HMP) with  $N = 45\,383$ ,  $M = 4\,788$ , and  $E = 122\,804$ , and (b) earth microbiome project (EMP) with  $N = 126\,730$ ,  $M = 23\,323$ , and  $E = 735\,868$ , using the MDL regularization method, together with an SBM prior. The top panels show the networks obtained, with edge weights indicated as colors. The middle panel shows the OTU counts in (a) different body sites and (b) earth habitats (the color code is the same as the “total count” panels in Fig. 9.). The bottom panel shows the edge weight distributions (with the same colors as the top panel), the node field distributions, the degree distributions (where node  $i$  has total degree  $k_i = \sum_j A_{ij}$ , as well as positive and negative degrees,  $k^+ = \sum_j A_{ij} \mathbb{1}_{W_{ij}>0}$  and  $k^- = \sum_j A_{ij} \mathbb{1}_{W_{ij}<0}$ , respectively), the node strength distributions (where node  $i$  has total strength  $d_i = \sum_j W_{ij}$ , as well as positive and negative strength,  $d_i^+ = \sum_j W_{ij} \mathbb{1}_{W_{ij}>0}$  and  $d_i^- = \sum_j W_{ij} \mathbb{1}_{W_{ij}<0}$ , respectively), and the average node fields as a function of the degrees (with the expected “magnetization” in the inset).






FIG. 6. Hierarchical community structure obtained with the nested SBM [73,74] integrated into the MDL regularization, for (a) the human microbiome project and (b) the earth microbiome project, corresponding to the same networks shown in Fig. 5. Each filled polygon represents the convex hull of the nodes that belong to the same group. Nested polygons represent nested hierarchical levels.

Strikingly, as we show in Appendix B, the groups found with the SBM that is part of our MDL regularization correlate significantly with the known taxonomic divisions of the microbial species considered, where most of the groups found contain predominately a single genus, although the same genus can be split into many inferred groups. This indicates different scales of similarity and dissimilarity between taxonomically proximal species when it comes to their interaction patterns with other species.

### A. Predicting stability and outcome of interventions

In this section we explore the fact that our reconstructed networks equip us with more than just a structural representation of the interactions between species, but it also provides a generative model that can be used to answer questions that cannot be obtained directly from the data. One particular question that is of biological significance is the extent to which an ecological system is stable to perturbations, and what are the outcomes of particular interventions. One important concept for ecology is that of “keystone” species, defined as those that have a disproportionately strong impact in its environment relative to its abundance [81,82]. In the context of microbial systems, keystone species are often associated with the impact they have once they are forcibly removed from the system, measured by the number of species that disappear as a consequence [82].

The removal of individual species can be investigated with the Ising model in the following way. First we observe that since the network structure is modular, and the weights  $\mathbf{W}$  are heterogeneous, the probability landscape for the configurations  $\mathbf{x}$  in such cases is typically multimodal, composed of many “islands” of high probability configurations separated by “valleys” of low probability

configurations—a property known as broken replica symmetry. If we consider this model to correspond to the equilibrium configuration of a dynamical system, then these islands amount to metastable configurations from which the system would eventually escape, but only after a very long time that grows exponentially with the total number of nodes. In this situation we can decompose the overall probability of a configuration  $\mathbf{x}$  as an average of all these discrete metastable states, which we will henceforth call *macrostates*, i.e.,

$$P(\mathbf{x}|\mathbf{W}, \boldsymbol{\theta}) = \sum_{\alpha} P(\mathbf{x}|\alpha)P(\alpha), \quad (38)$$

with  $P(\mathbf{x}|\alpha) = \mathbb{1}_{\mathbf{x} \in \alpha} P(\mathbf{x}|\mathbf{W}, \boldsymbol{\theta})/P(\alpha)$  being the configuration distribution when trapped in macrostate  $\alpha$ , and  $P(\alpha) = \sum_{\mathbf{x}} \mathbb{1}_{\mathbf{x} \in \alpha} P(\mathbf{x}|\mathbf{W}, \boldsymbol{\theta})$  is the asymptotic probability that this macrostate is eventually visited, if we run the underlying dynamics for an infinitely long time.

The individual macrostates  $\alpha$  have an important biological meaning, since they would correspond to the actual configurations that we would encounter in practice, as the time it would take to transition from one macrostate to another would not be experimentally observable.

We can then characterize the stability of the system precisely as its tendency to escape from one such macrostate after a small portion of the species is perturbed [see Fig. 8(a)]. In the case of a single species, we can quantify this by comparing the marginal expectation of the presence of a species  $i$  in the unperturbed system, i.e.,

$$\bar{x}_i(\alpha) = \sum_{\mathbf{x}} x_i P(\mathbf{x}|\alpha), \quad (39)$$

with the expectations obtained when another species  $j \neq i$  is forced to take value  $x_j = -1$ , i.e.,

$$\bar{x}'_i(\alpha, j) = \sum_{\mathbf{x} \setminus \{x_j\}} x_i P(\mathbf{x} \setminus \{x_j\} | x_j = -1, \mathbf{W}, \boldsymbol{\theta}, \alpha), \quad (40)$$

where in the last equation the distribution should be interpreted as the outcome of the dynamics if we initialize it with a configuration from macrostate  $\alpha$ , allowing it to potentially escape from it.

We can then quantify the magnitude of the perturbation of node  $j$  in macrostate  $\alpha$  via the expected number of additional species that disappear after it is removed:

$$z(j, \alpha) = \frac{1}{2} \sum_{i \neq j} \bar{x}_i(\alpha) - \bar{x}'_i(\alpha, j). \quad (41)$$

If the perturbation is insufficient to dislodge the system to another macrostate, we will typically have  $z(j, \alpha) \approx 0$ , otherwise the magnitude of the  $z(j, \alpha)$  will indicate the




FIG. 7. Probability distribution for the number of species lost  $z(j, \alpha)$  when a species  $j$  chosen uniformly at random is removed from the system in macrostate  $\alpha$ , averaged over many macrostates, for the Ising model inferred from the HMP data.

difference between the macrostates. We could then say that when the macrostate  $\alpha$  is encountered,  $j$  is a keystone species if  $z(j, \alpha)$  is large.

In Fig. 7 we show the distribution of  $z(j, \alpha)$  for macrostates simulated with the Ising model inferred from the HMP data, using the belief-propagation method [83], initialized with random messages. Although most perturbations are inconsequential, we observe that sometimes hundreds of species can disappear (or even appear) after a single species is removed.

In Fig. 8 we highlight two perturbations with large  $z(j, \alpha)$  values, where in one case all existing species vanish, and another where only a smaller subset remains. Importantly, the keystone species in these cases do not directly influence all species lost, instead they are affected only indirectly. In the corresponding dynamics, the move between macrostates would occur through a cascade of extinctions, where the most immediate neighbors of  $j$  would first disappear, and then the neighbors' neighbors, and so on. Notably, this kind of transition is typically irreversible; i.e., forcibly reinstating the removed species will not result in a transition in the reverse direction to the original macrostate. This kind of irreversible large-scale change that results from a small perturbation characterizes a “tipping point,” whose systematic prediction is a major challenge in ecology and climate science [84].

These predicted outcomes for the HMP data have only a tentative character, since, as we mentioned before, we are omitting important aspects from the model—such as the abundance magnitudes and environmental factors [85]—and making a strong modeling assumption with the Ising model. Nevertheless, this example already illustrates how

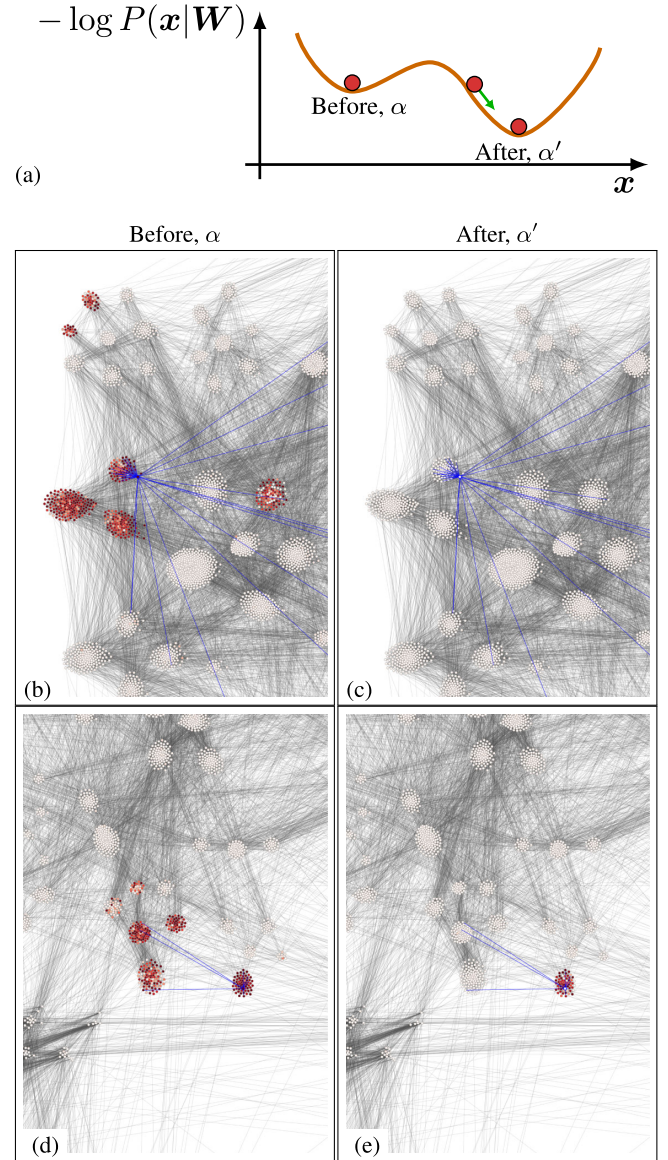


FIG. 8. Examples of single-species perturbations that cause the system to move from a macrostate  $\alpha$  to another  $\alpha'$ , as illustrated in (a), for the model inferred from the HMP data. Panels (b) and (d) show the macrostate before the perturbation, and (c) and (e) after. The node colors indicate the marginal expectations (white is absent and red is present), and the blue edges are the ones incident on the perturbed node.

even an elementary modeling assumption is sufficient to give us access to functional properties of a latent system, inferred from empirical data, which comes as a result of nonlinear emergent behavior mediated by the interactions between the species. When complemented with more realistic generative models, this kind of approach could be used to predict empirical outcomes, and be validated with experiments or additional observational data. Toward this goal, appropriate regularization is crucial, since a

model that overfits will deliver neither meaningful interpretations nor accurate predictions.

## VII. CONCLUSION

We have presented a principled nonparametric regularization scheme based on weight quantization and the MDL principle that can be used to perform network reconstruction via statistical inference in a manner that adapts to the statistical evidence available in the data. Our approach inherently produces the simplest (i.e., sparsest) network whenever the data cannot justify a denser one, and hence does not require the most appropriate number of edges to be known beforehand, and instead produces this central information as an output, as is often desired. Our approach does not rely on any specific property of the generative model being used, other than it being conditioned on a weighted adjacency matrix, and therefore it is applicable for a broad range of reconstruction problems.

Unlike  $L_1$  regularization, our method does not rely on weight shrinkage, and therefore introduces less bias in the reconstruction, including a much reduced tendency to infer spurious edges. Since it does not rely on cross-validation, it requires only a single fit to the data, instead of repeated inferences as in the case of  $L_1$ .

When combined with the subquadratic algorithm of Ref. [55], based on a stochastic second-neighbor search, our regularization scheme can be used to reconstruct networks with a large number of nodes, unlike most methods in current use, which have an algorithmic complexity that is at best quadratic on the number of nodes, typically even higher.

Our structured priors are extensible, and when coupled with inferential community detection [86] can also use latent modular network structure to improve the regularization, and as a consequence also the final reconstruction accuracy [71].

As we have demonstrated with an empirical case study on microbial interactions, our method is effective at uncovering functional latent large-scale complex systems from empirical data, which can then be used to make predictions about future behavior, the effect of interventions, and the presence of tipping points.

## ACKNOWLEDGMENTS

This work has been funded by the Vienna Science and Technology Fund (WWTF) and by the State of Lower Austria [Grant ID: 10.47379/ESS22032]. It is also in part the result of research conducted for Central European University, Private University. It was made possible by the CEU Open Access Fund.

## APPENDIX A: GENERATIVE MODELS

In our examples we use two generative models: the equilibrium Ising model [29] and the kinetic Ising model. The equilibrium Ising model is a distribution on  $N$  binary variables  $\mathbf{x} \in \{-1, 1\}^N$  given by

$$P(\mathbf{x}|\mathbf{W}, \boldsymbol{\theta}) = \frac{e^{\sum_{i<j} W_{ij}x_i x_j + \sum_i \theta_i x_i}}{Z(\mathbf{W}, \boldsymbol{\theta})}, \quad (\text{A1})$$

with  $\theta_i$  being a local field on node  $i$ , and  $Z(\mathbf{W}, \boldsymbol{\theta}) = \sum_{\mathbf{x}} e^{\sum_{i<j} W_{ij}x_i x_j + \sum_i \theta_i x_i}$  a normalization constant. Since this normalization cannot be computed analytically in closed form, we make use of the pseudolikelihood approximation [87],

$$P(\mathbf{x}|\mathbf{W}, \boldsymbol{\theta}) = \prod_i P(x_i|\mathbf{x}_{\setminus i}, \mathbf{W}, \boldsymbol{\theta}) \quad (\text{A2})$$

$$= \prod_i \frac{e^{x_i(\sum_j W_{ij}x_j + \theta_i)}}{2 \cosh(\sum_j W_{ij}x_j + \theta_i)}, \quad (\text{A3})$$

as it gives asymptotically correct results and has excellent performance in practice [29,88].

The kinetic Ising model is a Markov chain with transition probabilities conditioned on the same parameters as before, given by

$$P[\mathbf{x}(t+1)|\mathbf{x}(t), \mathbf{W}, \boldsymbol{\theta}] = \prod_i \frac{e^{x_i(t+1)(\sum_j W_{ij}x_j(t) + \theta_i)}}{2 \cosh(\sum_j W_{ij}x_j(t) + \theta_i)}. \quad (\text{A4})$$

In the case of the zero-valued Ising model with  $\mathbf{x} \in \{-1, 0, 1\}^N$ , but following the same Eq. (A1), the normalization of Eqs. (A2) and (A4) changes from  $2 \cosh(\cdot)$  to  $1 + 2 \cosh(\cdot)$ .

## APPENDIX B: SPECIES TAXONOMIES

In Fig. 9 we show the taxonomic classification of the species for the reconstructed networks of Fig. 5. The taxonomy consists of the phylum, class, order, family, and genus of each OTU. Each taxonomic category is shown as node colors, as indicated in the legend. In both cases we observe that most groups inferred with the SBM incorporated in our regularization scheme are compatible with the taxonomic division, since they tend to contain one dominating taxonomic category, at all levels—although a single category is often split into many groups, indicating structures that are not captured solely by taxonomy.



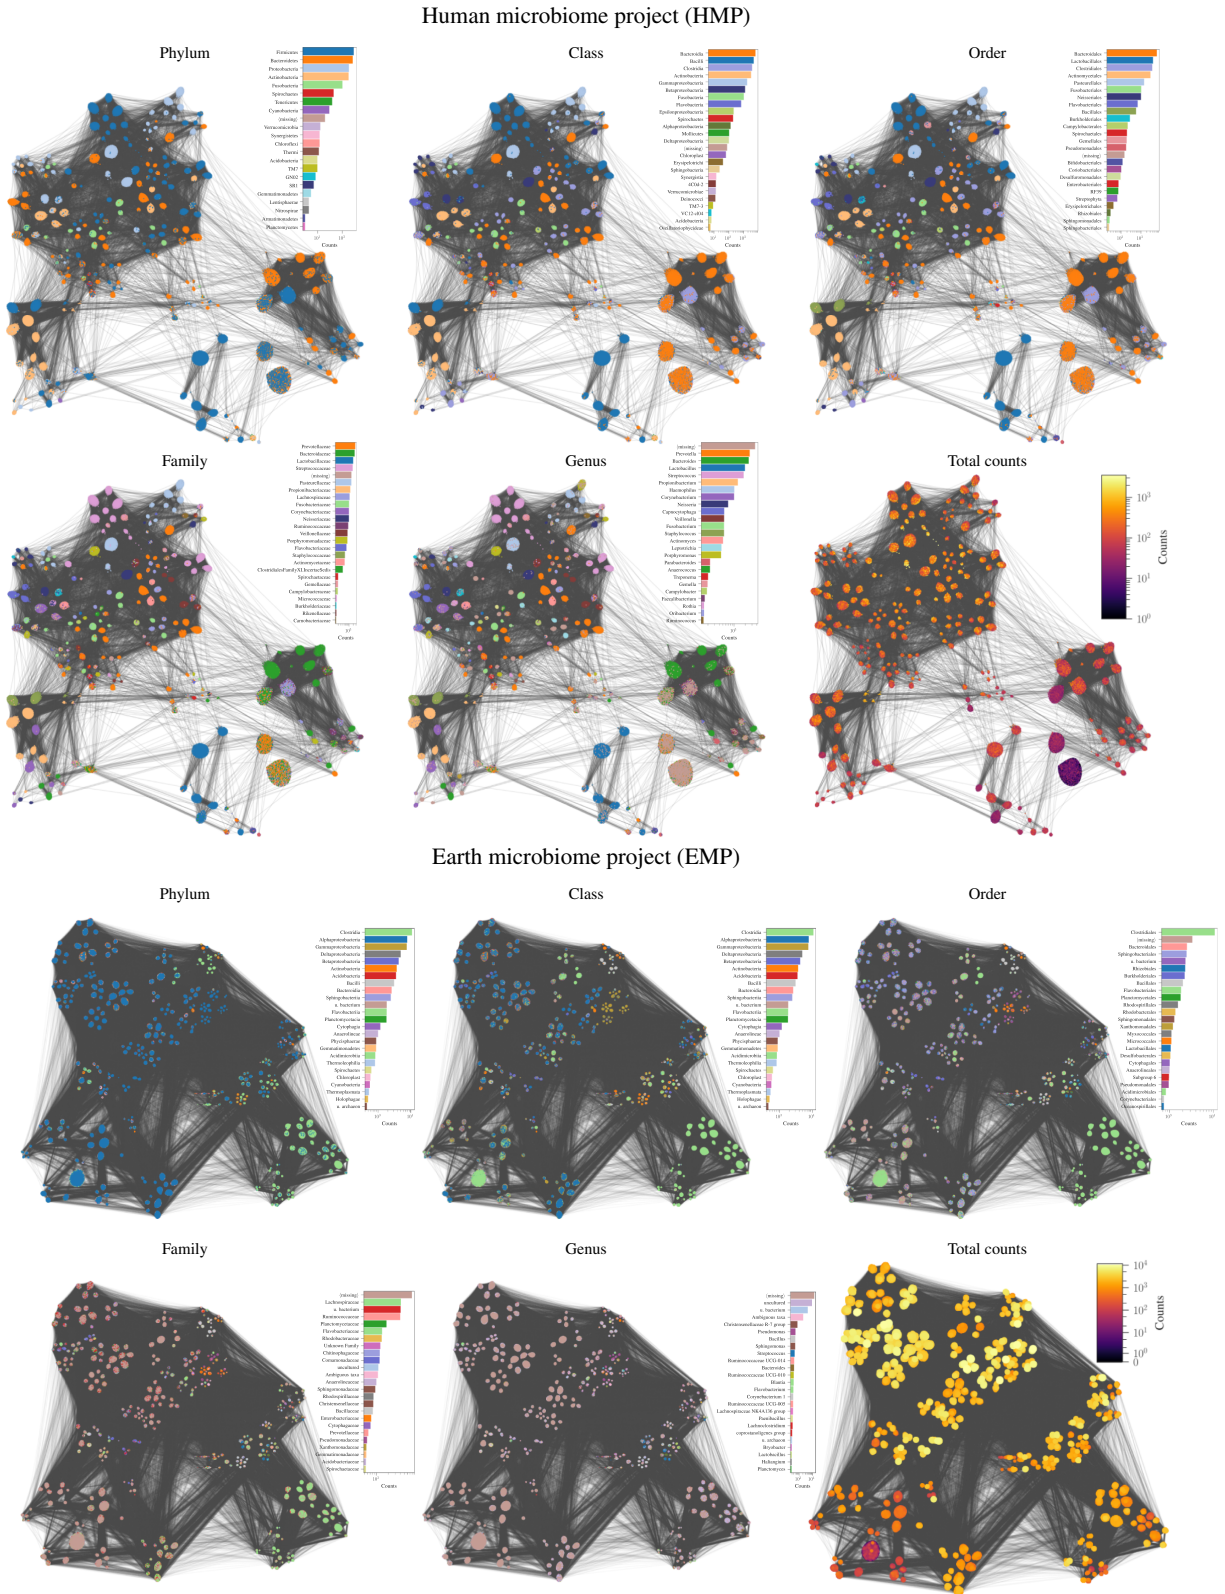


FIG. 9. Taxonomic classification of species for the human microbiome project and the earth microbiome project, overlaid on the inferred networks of Fig. 5, together with the total species count over all samples.



- [1] K. Guseva, S. Darcy, E. Simon, L. V. Alteio, A. Montesinos-Navarro, and C. Kaiser, *From diversity to complexity: Microbial networks in soils*, *Soil biology and biochemistry* **169**, 108604 (2022).
- [2] T. Bury, *A statistical physics perspective on criticality in financial markets*, *J. Stat. Mech.* (2013) P11004.
- [3] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, *Identification of direct residue contacts in protein-protein interaction by message passing*, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 67 (2009).
- [4] S. Cocco, C. Feinauer, M. Figliuzzi, R. Monasson, and M. Weigt, *Inverse statistical physics of protein sequences: A key issues review*, *Rep. Prog. Phys.* **81**, 032601 (2018).
- [5] P. D'haeseleer, S. Liang, and R. Somogyi, *Genetic network inference: From co-expression clustering to reverse engineering*, *Bioinformatics* **16**, 707 (2000).
- [6] G. Stolovitzky, D. Monroe, and A. Califano, *Dialogue on reverse-engineering assessment and methods*, *Ann. N.Y. Acad. Sci.* **1115**, 1 (2007).
- [7] V. Vinciotti, E. C. Wit, R. Jansen, E. J. C. N. de Geus, B. W. J. H. Penninx, D. I. Boomsma, and P. A. C. 't Hoen, *Consistency of biological networks inferred from micro-array and sequencing data*, *BMC Bioinf.* **17**, 254 (2016).
- [8] A. Maccione, M. Gandolfo, M. Tedesco, T. Nieuw, K. Imfeld, S. Martinoia, and B. Luca, *Experimental investigation on spontaneously active hippocampal cultures recorded by means of high-density MEAs: Analysis of the spatial resolution effects*, *Front. Neuroeng.* **3**, 1294 (2010).
- [9] J. R. Manning, R. Ranganath, K. A. Norman, and D. M. Blei, *Topographic factor analysis: A Bayesian model for inferring brain networks from neural data*, *PLoS One* **9**, e94914 (2014).
- [10] H. F. Po, A. M. Houben, A.-C. Haeb, D. R. Jenkins, E. J. Hill, H. R. Parri, J. Soriano, and D. Saad, *Inferring structure of cortical neuronal networks from firing data: A statistical physics approach*, *arXiv:2402.18788*.
- [11] Braundstein Alfredo, Ingrosso Alessandro, and Muntoni Anna Paola, *Network reconstruction from infection cascades*, *J. R. Soc. Interface* **16**, 20180844 (2019).
- [12] L. Peel, T. P. Peixoto, and M. De Domenico, *Statistical inference links data and theory in network science*, *Nat. Commun.* **13**, 6794 (2022).
- [13] J. Sun, D. Taylor, and E. M. Bollt, *Causal network inference by optimal causation entropy*, *SIAM J. Appl. Dyn. Syst.* **14**, 73 (2015).
- [14] J. G. Orlandi, O. Stetter, J. Soriano, T. Geisel, and D. Battaglia, *Transfer entropy reconstruction and labeling of neuronal connections from simulated calcium imaging*, *PLoS One* **9**, e98842 (2014).
- [15] E. Bullmore and O. Sporns, *Complex brain networks: Graph theoretical analysis of structural and functional systems*, *Nat. Rev. Neurosci.* **10**, 186 (2009).
- [16] V. L. Zogopoulos, G. Saxami, A. Malatras, K. Papadopoulos, I. Tsotra, V. A. Iconomidou, and I. Michalopoulos, *Approaches in gene coexpression analysis in eukaryotes*, *Biology* **11**, 1019 (2022).
- [17] G. T. Cantwell, Y. Liu, B. F. Maier, A. C. Schwarze, C. A. Serván, J. Snyder, and G. St-Onge, *Thresholding normally distributed data creates complex networks*, *Phys. Rev. E* **101**, 062302 (2020).
- [18] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity: The Lasso and Generalizations* (CRC Press, Boca Raton, FL, 2015).
- [19] R. Tibshirani, *Regression shrinkage and selection via the Lasso*, *J. R. Stat. Soc. Ser. B* **58**, 267 (1996).
- [20] N. Meinshausen and P. Bühlmann, *High-dimensional graphs and variable selection with the Lasso*, *Ann. Stat.* **34**, 1436 (2006).
- [21] M. Yuan and Y. Lin, *Model selection and estimation in the Gaussian graphical model*, *Biometrika* **94**, 19 (2007).
- [22] O. Banerjee, L. E. Ghaoui, and A. d'Aspremont, *Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data*, *J. Mach. Learn. Res.* **9**, 485 (2008).
- [23] L. Augugliaro, A. Abbruzzo, and V. Vinciotti,  *$\ell_1$ -penalized censored Gaussian graphical model*, *Biostatistics* **21**, e1 (2020).
- [24] J. Friedman, T. Hastie, and R. Tibshirani, *Sparse inverse covariance estimation with the graphical Lasso*, *Biostatistics* **9**, 432 (2008).
- [25] G. Bresler, E. Mossel, and A. Sly, *Reconstruction of Markov random fields from samples: Some observations and algorithms*, *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, Lecture Notes in Computer Science (Springer, Berlin, 2008), pp. 343–356.
- [26] G. Bresler, E. Mossel, and A. Sly, *Reconstruction of Markov random fields from samples: Some easy observations and algorithms*, *arXiv:0712.1402*.
- [27] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty, *High-dimensional Ising model selection using  $\ell_1$ -regularized logistic regression*, *Ann. Stat.* **38**, 1287 (2010).
- [28] G. Bresler, *Efficiently learning Ising models on arbitrary graphs*, in *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing (STOC '15)* (Association for Computing Machinery, New York, 2015), pp. 771–782.
- [29] H. C. Nguyen, R. Zecchina, and J. Berg, *Inverse statistical problems: From the inverse Ising problem to data science*, *Adv. Phys.* **66**, 197 (2017).
- [30] M. Vuffray, S. Misra, A. Lokhov, and M. Chertkov, *Interaction screening: Efficient and sample-optimal learning of Ising models*, *Advances in Neural Information Processing Systems*, Vol. 29 (Curran Associates, Inc., Red Hook, NY, 2016).
- [31] A. Y. Lokhov, M. Vuffray, S. Misra, and M. Chertkov, *Optimal structure and parameter learning of Ising models*, *Sci. Adv.* **4**, e1700791 (2018).
- [32] A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu, *Sparse permutation invariant covariance estimation*, *Electron. J. Stat.* **2**, 494 (2008).
- [33] H. Liu, K. Roeder, and L. Wasserman, *Stability approach to regularization selection (StARS) for high dimensional graphical models*, in *Proceedings of the 23rd International Conference on Neural Information Processing Systems (NIPS'10)*, Vol. 2 (Curran Associates Inc., Red Hook, NY, 2010), pp. 1432–1440.
- [34] E. Aurell and M. Ekeberg, *Inverse Ising inference using all the data*, *Phys. Rev. Lett.* **108**, 090201 (2012).

- [35] A. Decelle and F. Ricci-Tersenghi, *Pseudolikelihood decimation algorithm improving the inference of the interaction network in a general class of Ising models*, *Phys. Rev. Lett.* **112**, 070603 (2014).
- [36] S. Franz, F. Ricci-Tersenghi, and J. Rocchi, *A fast and accurate algorithm for inferring sparse Ising models via parameters activation to maximize the pseudo-likelihood*, *arXiv:1901.11325*.
- [37] R. Foygel and M. Drton, *Extended Bayesian information criteria for Gaussian graphical models*, in *Advances in Neural Information Processing Systems*, Vol. 23 (Curran Associates, Inc., Red Hook, NY, 2010).
- [38] H. Zou, *The adaptive Lasso and its Oracle properties*, *J. Am. Stat. Assoc.* **101**, 1418 (2006).
- [39] N. Meinshausen, *Relaxed Lasso*, *Computational Statistics and Data Analysis* **52**, 374 (2007).
- [40] A. Beck and M. Teboulle, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, *SIAM J. Imaging Sci.* **2**, 183 (2009).
- [41] C.-H. Zhang, *Nearly unbiased variable selection under minimax concave penalty*, *Ann. Stat.* **38**, 894 (2010).
- [42] J. Fan and R. Li, *Variable selection via nonconcave penalized likelihood and its Oracle properties*, *J. Am. Stat. Assoc.* **96**, 1348 (2001).
- [43] H. Wang, *Bayesian graphical Lasso models and efficient posterior computation*, *Bayesian Anal.* **7**, 867 (2012).
- [44] A. Mohammadi and E. C. Wit, *Bayesian structure learning in sparse Gaussian graphical models*, *Bayesian Anal.* **10**, 109 (2015).
- [45] L. Vogels, R. Mohammadi, M. Schoonhoven, and S. I. Birbil, *Bayesian structure learning in undirected Gaussian graphical models: Literature review with empirical comparison*, *arXiv:2307.02603*.
- [46] Y. Ni, V. Baladandayuthapani, M. Vannucci, and F. C. Stingo, *Bayesian graphical models for modern biological applications*, *Stat. Methods Appl.* **31**, 197 (2022).
- [47] A. Dobra, A. Lenkoski, and A. Rodriguez, *Bayesian inference for general Gaussian graphical models with application to multivariate lattice data*, *J. Am. Stat. Assoc.* **106**, 1418 (2011).
- [48] H. Wang, *Scaling it up: Stochastic search structure learning in graphical models*, *Bayesian Anal.* **10**, 351 (2015).
- [49] Z. Richard Li, T. H. McCormick, and S. J. Clark, *Bayesian joint spike-and-slab graphical Lasso*, *Proc. Mach. Learn. Res.* **97**, 3877 (2019).
- [50] C. M. Carvalho, N. G. Polson, and J. G. Scott, *Handling sparsity via the horseshoe*, in *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics* (PMLR, 2009), pp. 73–80.
- [51] Y. Li, B. A. Craig, and A. Bhadra, *The graphical horseshoe estimator for inverse covariance matrices*, *J. Comput. Graph. Stat.* **28**, 747 (2019).
- [52] Z. R. Li, T. H. McCormick, and S. J. Clark, *Using Bayesian latent Gaussian graphical models to infer symptom associations in verbal autopsies*, *Bayesian Anal.* **15**, 781 (2020).
- [53] J. Rissanen, *Information and Complexity in Statistical Modeling*, 1st ed. (Springer, New York, 2010).
- [54] P. D. Grünwald, *The Minimum Description Length Principle* (The MIT Press, Cambridge, MA, 2007).
- [55] T. P. Peixoto, *Scalable network reconstruction in subquadratic time*, *arXiv:2401.01404*.
- [56] A. N. Tikhonov, *On the stability of inverse problems*, *Proc. USSR Acad. Sci.* **39**, 195 (1943).
- [57] W. W. Zachary, *An information flow model for conflict and fission in small groups*, *J. Anthropol. Res.* **33**, 452 (1977).
- [58] Erroneously interpreting the log of a probability density  $P(\mathbf{W})$  as information could not only lead to meaningless negative values when  $P(\mathbf{W}) > 1$ , but also its value would change arbitrarily under an exchange of variables  $\mathbf{W}' = \mathbf{U}(\mathbf{W})$ , with  $\mathbf{U}(\mathbf{z})$  bijective, leading to a new probability density  $P'(\mathbf{W}') = P(\mathbf{U}^{-1}(\mathbf{W}')) \det[\nabla_{\mathbf{z}} \mathbf{U}^{-1}(\mathbf{z})|_{\mathbf{z}=\mathbf{W}'}]$ , and a new value  $-\log P'(\mathbf{W}')$  arbitrarily smaller or larger than  $\log P(\mathbf{W})$ , even though it refers to a fully equivalent model.
- [59] It is in fact easy to consider other choices for  $P(\mathbf{z})$  that eliminate residual shrinkage. A simple one is to sample the weight categories uniformly from the set of real numbers that are representable with  $q$  bits (e.g., using the IEEE 754 standard for floating-point arithmetic [60]), leading to  $P(\mathbf{z}) = \binom{2^q-1}{K}^{-1} \approx 1/2^{qK}$ . This gives virtually indistinguishable results to the choice of Eq. (25) in most examples we investigated, except, as mentioned in the main text, in situations where the weight categories would tend to diverge with  $|z_k| \rightarrow \infty$ , where we found Eq. (25) to yield a more convenient model as its residual shrinkage prevents this from happening in these edge cases.
- [60] *IEEE standard for floating-point arithmetic*, *IEEE Std 754-2008*, 1 (2008).
- [61] S. J. Wright, *Coordinate descent algorithms*, *Math. Program.* **151**, 3 (2015).
- [62] W. Dong, C. Moses, and K. Li, *Efficient  $k$ -nearest neighbor graph construction for generic similarity measures*, in *Proceedings of the 20th International Conference on World Wide Web (WWW '11)* (Association for Computing Machinery, New York, 2011), pp. 577–586.
- [63] Random bisection search will eventually succeed in finding the global minimum of a multimodal objective, hence it should in general be preferred over deterministic bisection, unless the objective function is convex.
- [64] T. P. Peixoto, *Merge-split Markov chain Monte Carlo for community detection*, *Phys. Rev. E* **102**, 012305 (2020).
- [65] T. P. Peixoto, *The graph-tool Python library*, *figshare*, 10.6084/m9.figshare.1164194, 2014, available at <https://graph-tool.skewed.de>.
- [66] To sample from the equilibrium Ising model we use MCMC with the Metropolis-Hastings criterion [67]. To determine equilibration we computed the  $\hat{R}$  of parallel chains [68], as well as the effective sample size (ESS).
- [67] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of state calculations by fast computing machines*, *J. Chem. Phys.* **21**, 1087 (1953).
- [68] A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, and P.-C. Bürkner, *Rank-normalization, folding, and localization: An improved  $\hat{R}$  for assessing convergence of MCMC (with discussion)*, *Bayesian Anal.* **16**, 667 (2021).
- [69] J. Moody, *Peer influence groups: Identifying dense clusters in large networks*, *Soc. Networks* **23**, 261 (2001).

- [70] M. Girvan and M. E. J. Newman, *Community structure in social and biological networks*, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 7821 (2002).
- [71] T. P. Peixoto, *Network reconstruction and community detection from dynamics*, *Phys. Rev. Lett.* **123**, 128301 (2019).
- [72] B. Karrer and M. E. J. Newman, *Stochastic blockmodels and community structure in networks*, *Phys. Rev. E* **83**, 016107 (2011).
- [73] T. P. Peixoto, *Nonparametric Bayesian inference of the microcanonical stochastic block model*, *Phys. Rev. E* **95**, 012317 (2017).
- [74] T. P. Peixoto, *Hierarchical block structures and high-resolution model selection in large networks*, *Phys. Rev. X* **4**, 011047 (2014).
- [75] Both methods also lead to inferences of isolated nodes. These correspond to deputies who are mostly absent from voting sessions, and more typically abstain.
- [76] Z. D. Kurtz, C. L. Müller, E. R. Miraldi, D. R. Littman, M. J. Blaser, and R. A. Bonneau, *Sparse and compositionally robust inference of microbial ecological networks*, *PLoS Comput. Biol.* **11**, e1004226 (2015).
- [77] N. Connor, A. Barberán, and A. Clauset, *Using null models to infer microbial co-occurrence networks*, *PLoS One* **12**, e0176751 (2017).
- [78] M. Layeghifard, D. M. Hwang, and D. S. Guttman, *Disentangling interactions in the microbiome: A network perspective*, *Trends Microbiol.* **25**, 217 (2017).
- [79] J. Lloyd-Price, A. Mahurkar, G. Rahnavard, J. Crabtree, J. Orvis, A. B. Hall, A. Brady, H. H. Creasy, C. McCracken, M. G. Giglio, D. McDonald, E. A. Franzosa, R. Knight, O. White, and C. Huttenhower, *Strains, functions and dynamics in the expanded Human Microbiome Project*, *Nature (London)* **550**, 61 (2017).
- [80] J. A. Gilbert, F. Meyer, J. Jansson, J. Gordon, N. Pace, J. Tiedje, R. Ley, N. Fierer, D. Field, N. Kyrpides, F.-O. Glöckner, H.-P. Klenk, K. E. Wommack, E. Glass, K. Docherty, R. Gallery, R. Stevens, and R. Knight, *The Earth Microbiome Project: Meeting report of the “1st EMP meeting on sample selection and acquisition” at Argonne National Laboratory October 6th 2010*, *Stand. Genomic Sci.* **3**, 249 (2010).
- [81] R. T. Paine, *A note on trophic complexity and community stability*, *Am. Nat.* **103**, 91 (1969).
- [82] L. S. Mills, M. E. Soulé, and D. F. Doak, *The keystone-species concept in ecology and conservation*, *BioScience* **43**, 219 (1993).
- [83] M. Mezard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, New York, 2009).
- [84] V. Dakos, B. Matthews, A. P. Hendry, J. Levine, N. Loeuille, J. Norberg, P. Nosil, M. Scheffer, and L. De Meester, *Ecosystem tipping points in an evolving world*, *Nat. Ecol. Evol.* **3**, 355 (2019).
- [85] V. Vinciotti, E. Wit, and F. Richter, *Random graphical model of microbiome interactions in related environments*, *J. Agric. Biol. Environ. Stat.* **29**, 1 (2024).
- [86] T. P. Peixoto, *Bayesian stochastic blockmodeling*, in *Advances in Network Clustering and Blockmodeling* (John Wiley & Sons, Ltd, New York, 2019), pp. 289–332.
- [87] J. Besag, *Spatial interaction and the statistical analysis of lattice systems*, *J. R. Stat. Soc. Ser. B* **36**, 192 (1974).
- [88] A. Mozeika, O. Dikmen, and J. Piili, *Consistent inference of a general model using the pseudolikelihood method*, *Phys. Rev. E* **90**, 010101(R) (2014).