

Taxonomies of networks from community structure

Jukka-Pekka Onnela,^{1,2,3,4,*} Daniel J. Fenn,^{4,5,*} Stephen Reid,³ Mason A. Porter,^{4,6} Peter J. Mucha,⁷
 Mark D. Fricker,^{4,8} and Nick S. Jones^{3,4,9,10}

¹*Department of Biostatistics, Harvard School of Public Health, Boston, Massachusetts 02115, USA*

²*Department of Health Care Policy, Harvard Medical School, Boston, Massachusetts 02115, USA*

³*Department of Physics, University of Oxford, Oxford OX1 3PU, United Kingdom*

⁴*CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, United Kingdom*

⁵*Mathematical and Computational Finance Group, University of Oxford, Oxford OX1 3LB, United Kingdom*

⁶*Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute, University of Oxford, OX1 3LB, United Kingdom*

⁷*Carolina Center for Interdisciplinary Applied Mathematics, Department of Mathematics and Institute for Advanced Materials, Nanoscience & Technology, University of North Carolina, Chapel Hill, North Carolina 27599, USA*

⁸*Department of Plant Sciences, University of Oxford, South Parks Road, Oxford OX1 3RB, United Kingdom*

⁹*Oxford Centre for Integrative Systems Biology, Department of Biochemistry, University of Oxford, Oxford OX1 3QU, United Kingdom*

¹⁰*Department of Mathematics, Imperial College, London SW7 2AZ, United Kingdom*

(Received 30 November 2011; published 10 September 2012)

The study of networks has become a substantial interdisciplinary endeavor that encompasses myriad disciplines in the natural, social, and information sciences. Here we introduce a framework for constructing taxonomies of networks based on their structural similarities. These networks can arise from any of numerous sources: They can be empirical or synthetic, they can arise from multiple realizations of a single process (either empirical or synthetic), they can represent entirely different systems in different disciplines, etc. Because mesoscopic properties of networks are hypothesized to be important for network function, we base our comparisons on summaries of network community structures. Although we use a specific method for uncovering network communities, much of the introduced framework is independent of that choice. After introducing the framework, we apply it to construct a taxonomy for 746 networks and demonstrate that our approach usefully identifies similar networks. We also construct taxonomies within individual categories of networks, and we thereby expose nontrivial structure. For example, we create taxonomies for similarity networks constructed from both political voting data and financial data. We also construct network taxonomies to compare the social structures of 100 Facebook networks and the growth structures produced by different types of fungi.

DOI: [10.1103/PhysRevE.86.036104](https://doi.org/10.1103/PhysRevE.86.036104)

PACS number(s): 89.75.Hc

I. INTRODUCTION

Although there is a long tradition of scholarship on networks, the last two decades have witnessed substantial advances in network science due to developments in physics, mathematics, computer science, sociology, and numerous other disciplines [1,2]. Given that the questions asked by researchers in different fields can be surprisingly similar, it would be useful to be able to highlight similarities in network structures across disciplines in a systematic way. One way to approach this is to formulate a suitable means of comparing networks and to use this means to develop taxonomies of networks. Such taxonomies have the potential to facilitate the identification of problems from different disciplines that might be approached similarly in terms of both empirical analyses and theoretical modeling. For example, if a biological network representing covariation of neural activity in different regions of the brain is demonstrated to be structurally similar to a financial network representing correlations of stock returns, then certain types of edge thresholding methods or structural null models might be applicable to both situations. (In this paper, we use the terms “edge” and “link” interchangeably.)

From a historical perspective, classification of objects has often been central to the progress of science, as demonstrated

by the periodic table of elements in chemistry and phylogenetic trees of organisms in biology [3]. It is plausible that an organization of networks has the potential to shed light on mechanisms for generating networks, reveal how an unknown network should be treated once one has discerned its position in a taxonomy, or help identify a network family’s anomalous members. Further potential applications of network taxonomies include unsupervised study of multiple realizations of a given model process (e.g., characterizing the similarities and differences of many different networks drawn from the Erdős-Rényi random graph model using the same parameter values), examination of multiple empirical networks with known similar origins or generative processes, and the detection of anomalous changes in temporally ordered series of networks. In this paper, we develop a framework for the creation of network taxonomies [4]. In so doing, we develop the requisite diagnostic tools and discuss several case studies that suggest how our methodology can help illuminate relationships both between and within families of networks.

In aiming to construct taxonomies of networks, one has to consider the scales at which one wants to compare differences in network structures. Much research has focused on extremes—either microscopic (e.g., node degree) or macroscopic (e.g., mean geodesic distance) properties—and numerous researchers have, for example, reported that many empirical networks possess heavy-tailed degree distributions or the small-world property [1,5]. Given the ubiquity of such

*These authors contributed equally to this work.

findings, it is clear that more nuanced approaches are needed to make useful comparisons between networks. Indeed, interpretations of microscopic and macroscopic approaches often implicitly assume that networks are homogeneous and ignore “mesoscopic” structures in networks. To overcome some of these limitations, earlier work has focused on the statistics of small, *a priori* specified modules called “motifs” [6,7], role-to-role connectivity profiles of nodes [8], the isolation of statistically significant structures called “backbones” [9], interrelations of network modules [10], examination of the number of nodes located within “shells” [11], and the self-similarity of networks as characterized by fractal exponents [12]. The taxonomic framework that we develop in the present paper builds on the idea of examining network modules by computing community structures [13,14], as was also done in the work of Ref. [15], and we subsequently compare signatures derived from community structure across networks. Importantly, although we use a specific method to identify network communities, much of the introduced framework is independent of that choice. Consequently, our comparative framework can accommodate a large variety of community detection schemes.

The remainder of this paper is organized as follows. First, we discuss the detection of communities in networks in order to find coherent groups of nodes that are densely connected to each other. We then introduce *mesoscopic response functions* (MRFs), which allow us to probe how the community structure of a network changes as a function of a resolution parameter that determines network scales of interest. We then illustrate MRFs using several examples of networks and compare the MRFs for several well-known generative models of networks. We use MRFs to develop a means to measure distance between a pair of networks, and we use this comparative measure to cluster networks and thereby develop taxonomies. Using 746 networks from numerous different fields, we construct a taxonomy of these networks. We then construct taxonomies of networks within fields using several case studies: voting in the United States (US) Senate, voting in the United Nations General Assembly, Facebook networks at US universities, fungal networks, and networks of stock returns in the New York Stock Exchange. In each example, we expose structure that either is illuminating or can be checked against information from an external source (e.g., previously published investigations). This suggests that our method for comparing networks is capturing important similarities and differences. We conclude with a brief summary and discussion of our results. In addition, we provide further details in the appendices and Supplemental Material [16]. Among other topics, we examine the robustness of the obtained taxonomies, address some computational issues, tabulate some of the basic properties of the networks that we study, and provide references for the network data sources used in this study.

II. MULTIREOLUTION COMMUNITY DETECTION

Our approach is based on network *community structure* [13,14]. A community consists of a set of nodes for which there are more edges (or, in the case of weighted networks, a larger total edge weight) connecting the nodes in the set than

what would be expected by chance. The algorithmic detection of communities is a particularly active area of network science, in part because communities are thought to be related to functional units in many networks and in part because they can strongly influence dynamical processes that operate on networks [13,14].

In this paper, we detect communities using the multiresolution Potts method [13,14,17], which is a generalization of modularity optimization [13,14,17–20]. (Modularity optimization is one of the most popular approaches for detecting communities.) Given a network adjacency matrix \mathbf{A} with elements A_{ij} , we find communities by minimizing the Hamiltonian of the infinite-range N -state Potts spin glass

$$\begin{aligned} \mathcal{H}(\lambda) &= - \sum_{i \neq j} J_{ij}(\lambda) \delta(C_i, C_j) \\ &= - \sum_{i \neq j} (A_{ij} - \lambda P_{ij}) \delta(C_i, C_j), \end{aligned} \quad (1)$$

where C_i indicates the community (state) of node (spin) i , the quantity λ is a resolution parameter, and $\mathbf{J}(\lambda)$ is the coupling matrix with entries $J_{ij}(\lambda)$ representing the interaction strength between node i and node j in the Potts Hamiltonian. We use the (undirected-network) null model $P_{ij} = k_i k_j / (2m)$, where k_i denotes the strength (total edge weight) of node i and m is the total edge weight in the network [18]. By tuning the resolution parameter λ , we can detect communities at multiple scales of a network. Our particular choice of J_{ij} implies that we are optimizing modularity (with the addition of the resolution parameter) [13,14].

To compare networks, we create profiles of summary statistics that characterize the community structure of each network at different mesoscopic scales. We also study a wide variety of networks that contain different numbers of nodes and edges. (We enumerate the networks that we consider in Table II of the Supplemental Material.) To ensure that we can compare the profiles for different networks, we sweep the resolution parameter λ from a minimum value Λ_{\min} to a maximum value Λ_{\max} (discussed in detail below). We define these quantities separately for each network such that the number of communities η into which a network is partitioned is 1 at Λ_{\min} and is equal to the total number of nodes N at Λ_{\max} . In other words, one can think of λ as a parameter that controls the fragmentation of a network into communities.

To find the minimum and maximum resolution-parameter values, consider the interactions in Eq. (1). An interaction is called *ferromagnetic* when $J_{ij} > 0$ and *antiferromagnetic* when $J_{ij} < 0$. For each pair of nodes i and j , we find the resolution $\lambda = \Lambda_{ij}$ at which the interaction J_{ij} is neutral [i.e., $J_{ij}(\Lambda_{ij}) = 0$], leading to $\Lambda_{ij} = A_{ij}/P_{ij}$. We thereby identify two special resolutions:

$$\Lambda_{\min} = \max_{ij} \{ \Lambda_{ij} | \eta(\lambda) = 1 \}, \quad (2)$$

$$\Lambda_{\max} = \max_{ij} \{ \Lambda_{ij} \} + \epsilon, \quad (3)$$

where $\epsilon > 0$ is any small number (we use $\epsilon = 10^{-6}$ in the present paper). The resolution Λ_{\min} is the largest Λ_{ij} value for which community detection yields a single community; note that this need not be the minimum nonzero value of Λ_{ij} . Including the small number ϵ in the definition of Λ_{\max} ensures

that all edges are antiferromagnetic at resolution $\lambda = \Lambda_{\max}$ and thereby forces each node into its own community.

III. MESOSCOPIC RESPONSE FUNCTIONS (MRFS)

To describe how a network disintegrates into communities as the value of λ is increased from Λ_{\min} to Λ_{\max} [see Fig. 1(a) for a schematic], one needs to select summary statistics. There are many possible ways to summarize such a disintegration process, and we focus on three diagnostics that characterize fundamental properties of network communities.

First, we use the value of the Hamiltonian $\mathcal{H}(\lambda)$ (1), which is a scalar quantity closely related to network modularity and quantifies the energy of the system [13,14]. Second, we calculate a partition entropy $S(\lambda)$ to characterize the community size distribution. To do this, let n_k denote the number of nodes in community k and define $p_k = n_k/N$ to be the probability to choose a node from community k uniformly at random. This yields a (Shannon) partition entropy of $S(\lambda) = -\sum_{k=1}^{\eta(\lambda)} p_k \log p_k$, which quantifies the disorder in the associated community size distribution. Third, we use the number of communities $\eta(\lambda)$.

Because we need to normalize \mathcal{H} , S , and η to compare them across networks, we define an *effective energy*

$$\mathcal{H}_{\text{eff}}(\lambda) = \frac{\mathcal{H}(\lambda) - \mathcal{H}_{\min}}{\mathcal{H}_{\max} - \mathcal{H}_{\min}} = 1 - \frac{\mathcal{H}(\lambda)}{\mathcal{H}_{\min}}, \quad (4)$$

where $\mathcal{H}_{\min} = \mathcal{H}(\Lambda_{\min})$ and $\mathcal{H}_{\max} = \mathcal{H}(\Lambda_{\max})$; an *effective entropy*

$$S_{\text{eff}}(\lambda) = \frac{S(\lambda) - S_{\min}}{S_{\max} - S_{\min}} = \frac{S(\lambda)}{\log N}, \quad (5)$$

where $S_{\min} = S(\Lambda_{\min})$ and $S_{\max} = S(\Lambda_{\max})$; and an *effective number of communities*

$$\eta_{\text{eff}}(\lambda) = \frac{\eta(\lambda) - \eta_{\min}}{\eta_{\max} - \eta_{\min}} = \frac{\eta(\lambda) - 1}{N - 1}, \quad (6)$$

where $\eta_{\min} = \eta(\Lambda_{\min}) = 1$ and $\eta_{\max} = \eta(\Lambda_{\max}) = N$.

Some networks contain a small number of entries Λ_{ij} that are orders of magnitude larger than most other entries. For example, in the network of Facebook friendships at Caltech [21,22], 98% of the Λ_{ij} entries are less than 100, but 0.02% of them are larger than 8000. These large Λ_{ij} values arise when two low-strength nodes become connected. Using the null model $P_{ij} = k_i k_j / (2m)$, the interaction between two nodes i and j becomes antiferromagnetic when $\lambda > A_{ij} / P_{ij} = 2m A_{ij} / (k_i k_j)$. If a network has a large total edge weight but both i and j have small strengths compared to other nodes in the network, then λ needs to be large to make the interaction antiferromagnetic. In prior studies, network community structure has been investigated at different mesoscopic scales by considering plots of various diagnostics as a function of the resolution parameter λ [13,14,17]. In the present example, such plots would be dominated by interactions that require large resolution-parameter values to become antiferromagnetic. To overcome this issue, we define the *effective fraction of antiferromagnetic edges*

$$\xi = \xi(\lambda) = \frac{\ell^A(\lambda) - \ell^A(\Lambda_{\min})}{\ell^A(\Lambda_{\max}) - \ell^A(\Lambda_{\min})} \in [0, 1], \quad (7)$$

where $\ell^A(\lambda)$ is the total number of antiferromagnetic interactions for the given value of λ . In other words, it is the number of Λ_{ij} elements that are smaller than λ . Thus, $\ell^A(\Lambda_{\min})$ is the largest number of antiferromagnetic interactions for which a network still forms a single community, and the effective number of antiferromagnetic interactions $\xi(\lambda)$ is the number of antiferromagnetic interactions (normalized to the unit interval) in excess of $\ell^A(\Lambda_{\min})$. The function $\xi(\lambda)$ increases monotonically in λ .

Sweeping λ from Λ_{\min} to Λ_{\max} corresponds to sweeping the value of ξ from 0 to 1. (One can think of λ as a continuous variable and ξ as a discrete variable that changes with events.) As we perform such sweeping for a given network, the number of communities increases from $\eta(\xi = 0) = 1$ to $\eta(\xi = 1) = N$ and yields a vector $[\mathcal{H}_{\text{eff}}(\xi), S_{\text{eff}}(\xi), \eta_{\text{eff}}(\xi)]$ whose components we call the *mesoscopic response functions* (MRFs) of that network. (We also sometimes refer to the vector itself as an MRF.) Because $\mathcal{H}_{\text{eff}} \in [0, 1]$, $S_{\text{eff}} \in [0, 1]$, $\eta_{\text{eff}} \in [0, 1]$, and $\xi \in [0, 1]$ for every network, we can compare the MRFs across networks and use them to identify groups of networks with similar mesoscopic structures. In Fig. 1(b), we show the Zachary Karate Club network [23] for different values of

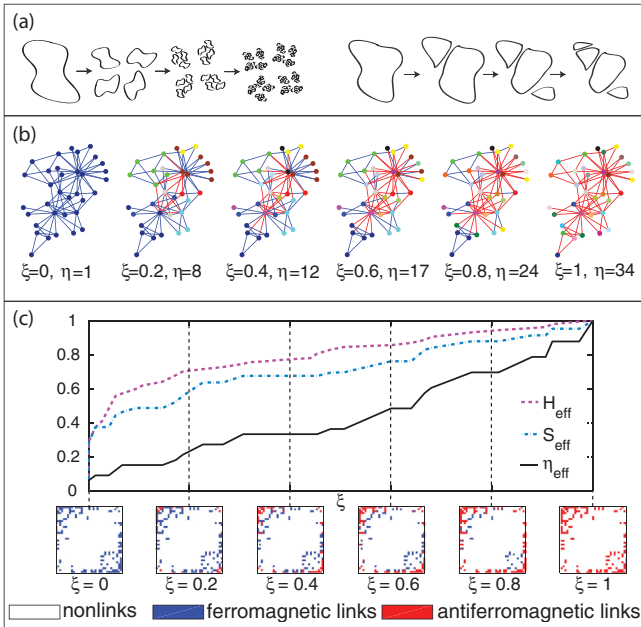


FIG. 1. (Color online) (a) Schematic of some of the ways that a network can break up into communities as the value of λ (or ξ) is increased. (b) Zachary Karate Club network [23] for different values of the effective fraction of antiferromagnetic edges ξ . All interactions are either ferromagnetic or antiferromagnetic; i.e., for the values of ξ that we used, there are no neutral interactions. We color edges in blue if the corresponding interactions are ferromagnetic, and we color them in red if the interactions are antiferromagnetic. We color the nodes based on community affiliation. (c) The \mathcal{H}_{eff} , S_{eff} , and η_{eff} MRFs, and the interaction matrix \mathbf{J} for different values of ξ . We color elements of the interaction matrix by depicting the absence of an edge in white, ferromagnetic edges in blue (dark gray), and antiferromagnetic edges in red (light gray).




FIG. 2. (Color online) Values of the three mesoscopic response functions $[H_{\text{eff}}(\xi), S_{\text{eff}}(\xi), \eta_{\text{eff}}(\xi)]$ as a function of $\xi \in [0, 1]$. By construction, the MRFs start from the bottom front corner $[H_{\text{eff}}(\xi = 0), S_{\text{eff}}(\xi = 0), \eta_{\text{eff}}(\xi = 0)]$ and end at the top back corner $[H_{\text{eff}}(\xi = 1), S_{\text{eff}}(\xi = 1), \eta_{\text{eff}}(\xi = 1)]$. We show a schematic representation of $[H(\xi), S(\xi), \eta(\xi)]$ for two networks; one is in blue (solid curve) and the other is in red (dashed curve). The surface represents the mean value of $\eta_{\text{eff}}(\xi)$ for binned values of $[H_{\text{eff}}(\xi), S_{\text{eff}}(\xi)]$ computed using the 189 networks identified in the right column of Table I.

ξ . As more edges become antiferromagnetic, the network fragments into smaller communities, and panel (c) shows the corresponding MRFs. In Fig. 2, we plot the coordinates of the three-dimensional MRF vector against each other. In Fig. 3, we show example MRFs for several other networks.

Although minimizing Eq. (1) is an NP-hard problem [24] and \mathcal{H} possesses a complicated landscape of local optima for many networks [25], there exist numerous good computational heuristics that make finding a nearly optimal




FIG. 3. (Color online) Example scalar mesoscopic response functions (MRFs). The curves show \mathcal{H}_{eff} (magenta, dashed), S_{eff} (blue, dash-dotted), and η_{eff} (black, solid) as a function of the effective fraction of antiferromagnetic edges ξ for the following networks: (a) New York Stock Exchange (NYSE), 1980–1999 [27]; (b) fractal (10,2,8) [29]; (c) Biogrid *D. melanogaster* [35]; (d) Garfield scientometrics citations [36]; (e) United Kingdom House of Commons voting, 2001–2005 [30]; and (f) roll-call voting of 108th United States House of Representatives [31–34].

partition of the network into communities at a given resolution computationally tractable [13,14]. Thus far, we have reported results that we obtained by optimizing modularity using the locally greedy Louvain algorithm [26] because its speed is important for studying large networks. We have compared the results that we report in the present work to those obtained from optimizing modularity using spectral and simulated-annealing algorithms, and we obtained similar MRFs and taxonomies for them (see Appendix B for more details).

IV. EXAMPLES OF MRFs

The shapes of the MRFs summarize many factors—including the fraction of possible edges in a network that are actually present, the relative weights of inter- versus intra-community edges, the edge weights compared with the expected edge weights in the null model, the number of edges that need to become antiferromagnetic for a community to fragment, and the way in which the communities fragment (e.g., whether a community splits in half or a single node leaves a community when a particular edge becomes antiferromagnetic). To understand the effects of some of these factors on the shapes of the MRFs, we consider some examples.

Of particular interest are plateaus in the η_{eff} and S_{eff} curves that are accompanied by large increases in \mathcal{H}_{eff} . As illustrated in Fig. 3(a), the New York Stock Exchange (NYSE) network from 1980 to 1999 [27] provides a good example of this behavior. This network is an example of a so-called *similarity network*. We use this label to describe networks that have been constructed by starting from some node-level quantity or attribute and then defining the edges based on some form of similarity or correlation measure between each pair of nodes. Similarity networks tend to be complete (or almost complete) and weighted networks, except when they have been deliberately thresholded. In this particular example, each node represents a stock, and the strength of the edge connecting stocks i and j is linear in the Pearson correlation between the daily logarithmic returns of the stocks. (See Sec. IX E for more details.) Plateaus imply that as the resolution λ is increased (leading to an increase in \mathcal{H}_{eff}), the communities remain unchanged even though the number and strength of antiferromagnetic interactions increase. As λ is increased and more interactions become antiferromagnetic, there is an increased energy incentive for communities to break up. Community partitions in such plateaus tend to be robust and have the potential to represent interesting structures [13,14,17,28].

In Fig. 3(b), we show MRFs for a “fractal” network [29], which demonstrates that plateaus in the η_{eff} and S_{eff} curves need not be accompanied by significant changes in \mathcal{H}_{eff} . Such plateaus can be explained by considering the distribution of Λ_{ij} values. If several interactions have identical values of Λ_{ij} , then the interactions all become antiferromagnetic at exactly the same resolution value. This leads to a significant increase in the effective fraction of antiferromagnetic edges ξ but only a small change in \mathcal{H}_{eff} . If these interactions do not result in additional communities, then we obtain plateaus in the η_{eff} and S_{eff} curves.

To demonstrate other qualitative behaviors, we show the MRFs for the Biogrid *Drosophila melanogaster* network and the Garfield Scientometrics citation network in Fig. 3(c) and Fig. 3(d), respectively. A common feature in these MRFs is the sharp initial increase in the curves that results from the networks initially breaking into two communities.

Another family of networks, which we will discuss in more detail in our case studies, are political voting networks. These voting networks are also similarity networks: We have constructed these networks so that an edge between two nodes indicates the level of agreement on votes between two entities, and each edge takes a value between 0 and 1. In Fig. 3(e), we show the MRFs for the voting network of the United Kingdom House of Commons during the period 2001–2005 [30]; in Fig. 3(f), we show the MRFs for the roll-call voting network for the 108th (2003–2004) United States House of Representatives [31–34]. In both cases, we observe that sharp increases in \mathcal{H}_{eff} can be accompanied by only small changes in η_{eff} and S_{eff} . To see how this can arise, we again consider the distribution of Λ_{ij} values. If there is a large difference between consecutive Λ_{ij} values, then a large increase in λ is needed to increase ξ ; this results in a large change in \mathcal{H}_{eff} . However, the change in η_{eff} is small because this results only in a single additional antiferromagnetic interaction.

V. COMPARING NETWORK MODELS

To provide further insights into MRFs, we consider Erdős-Rényi (ER) [37], Barabási-Albert (BA) [38], and Watts-Strogatz (WS) [39] networks. These network models are stochastic, and there is a large ensemble of possible network realizations for each choice of parameter values in these models. However, even with the ensuing structural variation, networks generated by a given one of these three models exhibit similar properties at mesoscopic and macroscopic scales, so we expect MRFs for different realizations of a given model to be similar. In Fig. 4, we compare the MRFs for




FIG. 4. (Color online) MRFs for 1000 realizations of Erdős-Rényi (ER), Barabási-Albert (BA), and Watts-Strogatz (WS) networks. Each network has $N = 1000$ nodes and mean degree $\langle k \rangle = 10$. For each value of ξ , the upper curves show the maximum values of \mathcal{H}_{eff} (top row), S_{eff} (middle row), and η_{eff} (bottom row) for all networks in the ensemble; the lower curves show the corresponding minimum value, and the dashed curves show the corresponding mean.




FIG. 5. (Color online) Upper panels: MRFs for Watts-Strogatz networks for different values of the rewiring probability p . Each network has $N = 1000$ nodes and mean degree $\langle k \rangle = 10$. Lower panels: Distribution of Λ_{ij} values for each network. The MRFs for $p = 1$ are indistinguishable from those of an Erdős-Rényi network with $N = 1000$ and $\langle k \rangle = 10$.

1000 realizations of each model for networks with $N = 1000$ nodes and mean degree $\langle k \rangle = 10$. For the WS networks, we set the edge rewiring probability at $p = 0.1$. As illustrated in Fig. 4, we obtain a narrow range of possible MRFs for fixed parameter values. This comparison illustrates that the MRF profiles of the three different models are distinctive. In addition, for each model, there is little variation in the shapes of the MRFs for different network realizations that use the same parameter values.

It is also instructive to consider variation in MRF shapes for a particular network model for different parameter values. We focus on WS networks because they illuminate the effect of the distribution of Λ_{ij} values on the shapes of the MRFs. In Fig. 5, we show MRFs for WS networks for different values of the edge rewiring probability p . (We continue using $N = 1000$ and $\langle k \rangle = 10$.) We also show the distribution of Λ_{ij} values for each network.

For small rewiring probabilities, the MRFs have lots of steps. As with prior examples, we can see how this feature arises by considering the distribution of Λ_{ij} values. When the rewiring probability is small, many nodes possess the same degree, which results in the presence of many interactions with identical Λ_{ij} values (see the bottom left panel of Fig. 5). Because several interactions have identical Λ_{ij} values, these interactions all become antiferromagnetic at exactly the same resolution-parameter value, so the behavior of MRFs changes only for a small number of ξ values. As the rewiring probability p is increased, the degree and Λ_{ij} distributions become more heterogeneous, which leads to smoother MRFs. For a rewiring probability of $p = 1$, the WS network is an ER network.

VI. MEASURING DISTANCE BETWEEN NETWORKS

In the framework that we have introduced in this paper, comparing two networks at the mesoscopic level amounts to characterizing the differences in behavior of the corresponding MRFs. To quantify such differences, we define a distance between two networks with respect to one of the summary statistics as the area between the corresponding MRFs. For example, the distance between two networks i and j with respect to the effective energy \mathcal{H}_{eff} is given by

$$d_{ij}^{\mathcal{H}} = \int_0^1 |\mathcal{H}_{\text{eff}}^i(\xi) - \mathcal{H}_{\text{eff}}^j(\xi)| d\xi. \quad (8)$$

TABLE I. Network categories, the total number of networks assigned to each category, and the number of networks from each category included in the taxonomy in Fig. 6. For the full taxonomy that uses all 746 networks, see Fig. 1 of the Supplemental Material.

Category	All networks	Taxonomy networks
Political: voting	285	23
Facebook	100	15
Fungal	65	12
Synthetic	58	0
Financial	54	6
Metabolic	43	15
Social	26	26
Political: cosponsorship	26	26
Other	23	0
Protein interaction	22	22
Political: committee	16	16
Brain	12	12
Language	8	8
Collaboration	8	8
Total	746	189

For the effective entropy and effective number of communities, the distances are given by $d_{ij}^S = \int_0^1 |S_{\text{eff}}^i(\xi) - S_{\text{eff}}^j(\xi)| d\xi$ and $d_{ij}^\eta = \int_0^1 |\eta_{\text{eff}}^i(\xi) - \eta_{\text{eff}}^j(\xi)| d\xi$, respectively.

We represent the resulting three sets of distances (computed for each pair of networks from the 746 networks that we consider; see Table I) in matrix form as $\mathbf{D}^{\mathcal{H}}$, \mathbf{D}^S , and \mathbf{D}^η . These distance measures have several desirable properties. First, they compare MRFs across all network scales (i.e., for all values of ξ); second, each distance is bounded between 0 and 1; third, the distances are easy to interpret, as each of them corresponds to the geometric area between (a certain dimension of) a pair of MRFs; and finally, we find *a posteriori* that these distances can be used to cluster networks accurately (see the discussions below).

We have computed MRFs for the energy \mathcal{H} , entropy S , and number of communities η , but we can proceed similarly with any desired summary statistic. If two diagnostics provide similar information, then one of them can be excluded without significant loss of information. We checked whether the summary statistics were sufficiently different—for the set of networks that we considered—for it to be worthwhile to include all of them by calculating the Pearson correlation coefficient between their corresponding distance measures. The correlations between the pairs of distances are $r(d_{ij}^{\mathcal{H}}, d_{ij}^S) \doteq 0.36$, $r(d_{ij}^{\mathcal{H}}, d_{ij}^\eta) \doteq 0.24$, and $r(d_{ij}^S, d_{ij}^\eta) \doteq 0.58$. These correlations are not sufficiently high to justify excluding any of the summary statistics.

In the interest of parsimony—and given the nonvanishing correlations between the distance measures—we reduce the number of distance measures using principal component analysis (PCA) [40]. Starting with \mathcal{N} networks, we create a $\frac{1}{2}\mathcal{N}(\mathcal{N} - 1) \times 3$ matrix in which each column corresponds to the vector representation of the upper triangle of one of the distance matrices $\mathbf{D}^{\mathcal{H}}$, \mathbf{D}^S , \mathbf{D}^η , and we perform a PCA on this matrix. We then define a distance matrix \mathbf{D}^p with elements $d_{ij}^p = w_{\mathcal{H}}d_{ij}^{\mathcal{H}} + w_S d_{ij}^S + w_\eta d_{ij}^\eta$, where the weights are the coefficients for the first principal component, and

we normalize the sum of squared coefficients to unity. The coefficients are $w_{\mathcal{H}} \doteq 0.24$, $w_S \doteq 0.79$, and $w_\eta \doteq 0.57$. The first component accounts for about 69% of the variance, so the distance matrix \mathbf{D}^p provides a reasonable single-variable projection of the distance matrices $\mathbf{D}^{\mathcal{H}}$, \mathbf{D}^S , and \mathbf{D}^η .

It is important that the distance measures for comparing networks are robust to small perturbations in network structure. Because many of the networks that we study are constructed empirically, they might contain false positives and false negatives. In other words, the networks might falsely identify a relationship where none exists, and they also might fail to identify an existing relationship. Consequently, the topology and edge weights of an observed network might be slightly different than those of the actual underlying network. To test the robustness of our distance measures to such observational errors, we recalculate the MRFs for a subset of relatively small unweighted networks in which, for each network, we rewire a number of edges corresponding to a given percentage of the total number of edges (5%, 10%, 20%, 50%, or 100%). See Appendix A for more details. (We study networks with up to 1000 nodes and consider only a subset of 25 networks because of the computational costs of rewiring a large number of networks multiple times; however, we have performed the same investigation for five different subsets of 25 networks and obtained similar results. We list the networks in each subset in Table I of the Supplemental Material.) We investigate two rewiring mechanisms: one in which the degree distribution is maintained, where we also ensure after each rewiring that the network consists of a single connected component; and another in which the only constraint is that the network continues to consist of a single connected component after each edge rewiring [41]. We find in both cases that the structures of the block-diagonalized distance matrices for the 25 networks (see Figs. 14 and 15 in Appendix A) are robust to random perturbations of the networks, thereby suggesting that our MRF distance measures are not sensitive to small structural perturbations.

VII. CLUSTERING NETWORKS

We assign each of the 746 networks to a category based on its type (see Table I). Due to the varying availability of different types of network data, the networks that we consider are not distributed evenly across these categories. Many of the networks are either different temporal snapshots of the same system or different realizations of the same type of network. To have a more balanced distribution across the different categories, we focus on 189 of the 746 networks. We include only categories for which we have eight or more networks, and we select a subset of networks (uniformly at random) from each of the larger categories. We also exclude all synthetic networks. See Sec. IV of the Supplemental Material for the list of networks that we consider and Fig. 1 in Sec. II of the Supplemental Material for a dendrogram showing a taxonomy that we constructed using all 746 networks.

Our primary reason for assigning each network to a category is to use such an external categorization to help assess the quality of taxonomies produced by the unsupervised MRF clustering. For each way of computing distance, we construct a dendrogram for the set of networks using average linkage




FIG. 6. (Color) Taxonomy for 189 networks (described in the text). We construct the dendrogram (tree) using the distance matrix \mathbf{D}^p and average linkage clustering. We order the leaves of the dendrogram to minimize the distances between adjacent nodes (using the MATLAB option `optimalleaforder`) and color the leaves to indicate the type of network.

clustering, which is an agglomerative hierarchical clustering technique [13,42,43]. In Fig. 6, we show a dendrogram obtained from the distance matrix \mathbf{D}^p . The colored rectangle underneath each leaf indicates the network category. Contiguous blocks of color demonstrate that networks from the same category have been grouped together using the MRF clustering method, and the presence of such contiguous color blocks is an indication of the success of the MRF clustering scheme.

The assignment of a network to one of these categories is, of course, to some extent subjective, as several of the networks could belong to more than one category. For example, we could categorize the network of jazz musicians [44] as either a collaboration network or a social network. The initial selection of network categories is also somewhat subjective. One could argue that if one has a social network category, then it is not necessary to have a collaboration network category as well because a collaboration network is a type of social network. We have attempted to maintain a balance between having too many categories and having too few of them. When such ambiguities have arisen, we have systematically chosen the most specific of the relevant categories (e.g., we placed the jazz musician network in the category of collaboration networks rather than in the category of social networks). Obviously, our methodology does not depend on such choices.

VIII. TAXONOMIES OF EMPIRICAL NETWORKS

All of the networks in some categories appear in blocks of adjacent leaves in the dendrogram in Fig. 6. For example, there is a cluster of political voting networks at the far left of the dendrogram. This cluster includes voting networks from the US Senate, the US House of Representatives, the UK House of Commons, and the United Nations General Assembly (UNGA). The clustering of these voting networks suggests that there are some common features in the network representations of the different legislative bodies. We also obtain blocks that consist of all political committee networks and all metabolic networks.

There are also several categories for which all except one or two networks cluster into a contiguous block. For example, all

but two of the fungal networks appear in the same block, and all but one of the Facebook networks are clustered together. The isolated Facebook network is the Caltech network, which is the smallest network of this type and which appears in a group next to that containing all of the other Facebook networks. We also remark that the social organization revealed by the community structure of the Caltech Facebook network has been shown to be different from those of the other Facebook networks [21,22].

Networks of certain categories do not appear in near-contiguous blocks. For example, protein interaction networks appear in several clusters. These networks represent interactions within several different organisms, so we would not expect all of them to be clustered together. Moreover, the data that we employed include examples of protein interaction networks for the same organism in which the interactions were identified using different experimental techniques, and these networks do not cluster together. This supports previous work suggesting that the properties of protein interaction networks are very sensitive to the experimental procedure used to identify the interactions [45,46]. Social networks are also distributed throughout the dendrogram. This is unsurprising given the extremely broad nature of the category, which includes networks of very different sizes with edges representing a diverse range of social interactions. The leftmost outlying social network is the network of Marvel comic book characters [47], which is arguably an atypical social network.

The grouping (and, to some extent, the nongrouping) of networks by category suggests that the PCA distance matrix \mathbf{D}^p between MRFs of different networks produces a sensible taxonomy. It is important to ask, however, whether a simpler approach based on a single network diagnostic, such as edge density, can be comparably successful at constructing a taxonomy. In Appendix D, we demonstrate using some well-known diagnostics that this does not appear to be the case, as the diagnostics we tried were unable to reproduce or explain the classifications that we produced using the MRFs.

In order to compare the aggregate shapes of the MRFs across categories, we show the bounds of the \mathcal{H}_{eff} , \mathcal{S}_{eff} , and η_{eff} curves for each category in Fig. 7. We again consider all empirical network categories with at least eight networks in them. This illustrates that the MRFs for some classes of networks (such as political cosponsorship and metabolic networks) are very similar to each other, whereas there are large variations in the MRFs for other categories (such as social and protein interaction networks). The variety of different MRFs for the social and protein interaction networks is consistent with the fact that their constituent networks are scattered throughout the dendrogram in Fig. 6.

IX. CASE STUDIES

We now consider several case studies, in which we generate taxonomies using multiple realizations of particular types of networks or multiple time snapshots of particular networks. This enables us to compare these networks and (in some cases) illustrate possible connections between network function and mesoscopic network structure.




FIG. 7. (Color online) MRFs for all of the network categories containing at least eight networks (see Table I). At each value of ξ , the upper curve shows the maximum value of \mathcal{H}_{eff} (magenta, left panel in each category), S_{eff} (blue, center panels), and η_{eff} (black, right panels) for all networks in the category and the lower curve shows the minimum value. The dashed curves show the corresponding mean MRFs.

A. Voting in the United States Senate

Our first example deals with roll-call voting in the US Senate [31–34,48]. Establishing a taxonomy of networks detailing the voting similarities of individual legislators complements previous studies of these data, and it facilitates the comparison of voting similarity networks across time. We consider Congresses 1–110, which cover the period 1789–2008. As in Ref. [34], we construct networks from the roll-call data [31,32] for each two-year Congress such that the adjacency matrix element $A_{ij} \in [0,1]$ represents the number of times Senators i and j voted the same way on a bill (either both in favor of it or both against it) divided by the total number of bills on which both of them voted. Following the approach of Ref. [32], we consider only “nonunanimous” roll-call votes, which are defined as votes in which at least 3% of the Senators were in the minority.

Much research on the US Congress has been devoted to the ebb and flow of partisan polarization over time and the influence of parties on roll-call voting [33,34]. In highly polarized legislatures, representatives tend to vote along party lines, so there are strong similarities in the voting patterns of members of the same party and strong differences between members of different parties. In contrast, during periods of low polarization, the party lines become blurred. The notion of partisan polarization can be used to help understand the taxonomy of Senates in Fig. 8, in which we consider two measures of polarization. The first measure uses DW-Nominate scores (a multidimensional scaling technique commonly used in political science [32,33]), where the extent of polarization is given by the absolute value of the difference between the mean first-dimension DW-Nominate scores for members of one party and the same mean for members of the other party [31–33]. In particular, we use the simplest such measure of polarization, called MPR polarization, which assumes a competitive two-party system and hence cannot be calculated prior to the 46th Senate. The second measure that we consider is the maximum modularity Q over partitions of

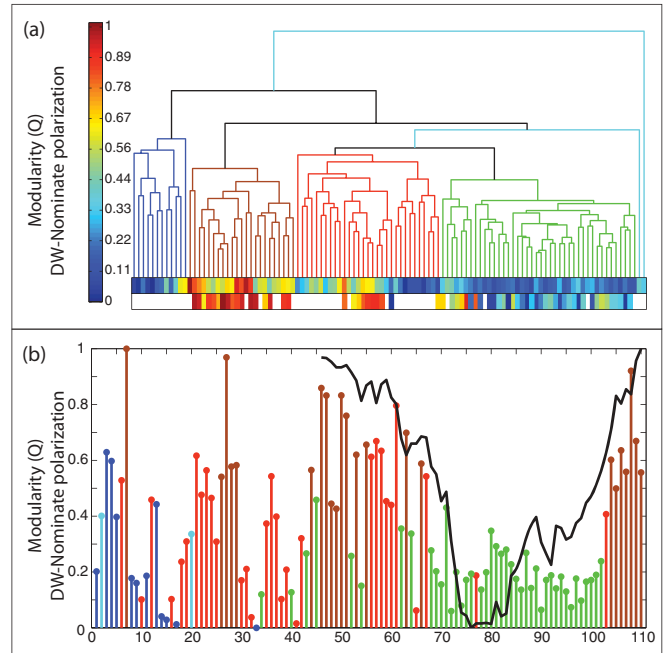


FIG. 8. (Color) (a) Dendrogram for Senate roll-call voting networks for the 1st–110th Congresses. Each leaf in the dendrogram represents a single Senate. The two horizontal color bars below the dendrograms indicate polarization measured in terms of optimized modularity (upper bar) and DW-Nominate scores (lower bar). We color the branches in the dendrogram corresponding to periods of similar polarization. (b) Polarization of the US Senate as a function of time, which we label using the Congress number. The height of each stem indicates the level of polarization measured using optimized modularity, and the color of each stem gives the cluster membership of each Senate in (a). The black curve shows the DW-Nominate polarization. Note that we have normalized both measures to lie in the interval $[0,1]$.

a network. It was shown recently that Q is a good measure of polarization even for Congresses without clear party divisions [34]. Modularity is given in terms of the energy \mathcal{H} in Eq. (1) by $Q = -\mathcal{H}(\lambda = 1)/(2m)$.

In Fig. 8(a), we include bars under the dendrograms to represent the two polarization measures, both of which have been normalized to lie in the interval $[0,1]$. The bars demonstrate that Senates with similar levels of polarization (measured in terms of both DW-Nominate scores and optimized modularity values) are usually assigned to the same group, suggesting that our MRF clustering technique groups Senates based on the polarization of roll-call votes. We have also colored dendrogram groups according to their mean levels of polarization using optimized modularity, where the brown group in the dendrogram corresponds to the most polarized Senates and the blue group corresponds to the least polarized Senates. We chose the specific number of groups by inspection of the dendrogram. Although one ought to expect similarity in the results from the modularity-based measure of polarization and the MRF clustering, it is important to stress that the MRF clustering method is based on different principles; modularity attempts to quantify the extent to which a given network is “modular,” whereas the MRF clustering explicitly




FIG. 9. (Color online) Comparison of the (low-polarization) 85th Senate and the (high-polarization) 108th Senate. The panels show (a) the \mathcal{H}_{eff} MRFs and (b) the cumulative distributions of Λ_{ij} values.

compares the differences in modular structures between any two networks at all scales.

In Fig. 8(a), we also show the clusters that we obtained for the Senate. They closely match the different periods of polarization that have been identified using optimized modularity and DW-Nominate [34]. The cluster with the highest mean polarization (shown in brown) consists of Senates 7, 26–29, 44, 46–51, 53, 55, 66, and 104–110. The 104th–110th Congresses correspond to a period of extremely high polarization following the 1994 “Republican Revolution,” in which the Republican party earned majority status in the House of Representatives for the first time in more than 40 years [31,33,34]. The cluster with the second highest mean polarization (shown in red) includes several contiguous blocks of Senates, such as those from Congresses 21–25, 35–39, and 56–61. The 21st–25th Congresses (1829–1839) corresponded to a period of partisan conflict between supporters of John Quincy Adams and Andrew Jackson; it lasted until the emergence of the Whigs and the Democratic Party in the 25th Congress [34,49]. The American Civil War started during the 37th Congress, and a third party known as the Populist Party was strong during the 56th–58th Congresses.

The main differences between different clusters occur in the \mathcal{H}_{eff} response functions. For the most polarized Senates, there is a sharp shoulder in the \mathcal{H}_{eff} MRF that becomes less pronounced as the polarization decreases. We illustrate this in Fig. 9, in which we compare the \mathcal{H}_{eff} MRFs for the (low-polarization) 85th and (high-polarization) 108th Senates. The shoulder in the \mathcal{H}_{eff} curve for the 108th Senate is very pronounced, which can be explained by considering the distribution of Λ_{ij} values. The 108th Senate has a bimodal Λ_{ij} distribution that contains a trough at $\Lambda_{ij} = 1$. Recall that $\Lambda_{ij} = A_{ij}/P_{ij}$, so Λ_{ij} compares the observed voting similarity A_{ij} of legislators i and j with the similarity $P_{ij} = k_i k_j / (2m)$ expected from random voting. If $\Lambda_{ij} < 1$, legislators i and j vote differently more frequently than expected (with respect to the chosen null model); if $\Lambda_{ij} > 1$, they vote more similarly than expected. Therefore, the peaks in the Λ_{ij} distribution above and below 1 correspond, respectively, to intraparty and interparty voting blocs. In a Senate with low polarization, legislators from different parties often vote in the same manner, so the values of Λ_{ij} no longer separate two distinct types of behavior.

We also examined roll-call voting networks in the US House of Representatives and found many similar features as the ones

that we have presented for the US Senate. For example, the highly polarized 104th–110th Congresses, which followed the “Republican Revolution,” appear in the same cluster for both the House and Senate. We also observed some differences in the clusters for the two chambers. For example, the 78th–102nd Senates all appear in the same cluster. For the House, however, Congresses 80, 88, 89, and 98–102 are not in the same cluster as the other Congresses between 78 and 102; instead, they are assigned to a cluster that also includes the 26th–28th Houses. This was a particularly eventful period: The 25th Congress saw the emergence of the Whigs and the Democratic Party, and the abolitionist movement was also prevalent (e.g., the *Amistad* seizure occurred in 1839 during the 26th Congress).

B. Voting in the United Nations General Assembly

The United Nations General Assembly (UNGA) is one of the principal organs of the United Nations (UN), and it is the only part of the UN in which all member nations have equal representation. Although most resolutions are neither legally nor practically enforceable because the General Assembly lacks enforcement powers on most issues, it is the only forum in which a large number of states meet and vote regularly on international issues. It also provides an interesting point of comparison with roll-call voting in the US Congress, as the level of agreement on UN resolutions tends to be much higher than that in the Senate and House [50].

We study voting for the 1st–63rd sessions (covering the period 1946–2008), where each session corresponds to a year [51]. For each session, we define an adjacency matrix \mathbf{A} whose elements A_{ij} represent the number of times countries i and j voted in the same manner in a session (i.e., the sum of the number of times both countries voted yea on the same resolution, both countries voted nay on the same resolution, or both countries abstained from voting on the same resolution) divided by the total number of resolutions on which the UNGA voted in a session. The matrix \mathbf{A} , with elements $A_{ij} \in [0, 1]$, thereby represents a (similarity) network of weighted edges between countries.

We cluster UNGA sessions by comparing MRFs for the corresponding voting networks. In Fig. 10, we plot a




FIG. 10. (Color online) Dendrogram for the United Nations General Assembly resolution voting network for the 1st–63rd sessions (excluding the 19th session), covering the period 1946–2008. Each leaf in the dendrogram represents a single session. In the main text, we discuss the coloring of groups of dendrogram branches.

dendrogram of the UNGA sessions and highlight some of the clusters, which correspond to notable periods in the recent history of international relations. The red cluster in the middle of the dendrogram consists of all post-Cold War sessions (1992–2008) except 1995. This group forms a larger cluster with some UNGA sessions from the 1970s and a cluster consisting of 1946, 1948, and 1950. These last three sessions (which are shown in magenta) are all noteworthy: 1946 was the first session of the UNGA, the Universal Declaration of Human Rights was introduced during the 1948 session, and the “Uniting for Peace” resolution was passed during the 1950 session. At the rightmost part of the dendrogram, we color in black a group that consists of all sessions from 1979 to 1991 (excluding 1980). The beginning of this period marked the end of détente between the Soviet Union and the United States following the former’s invasion of Afghanistan at the end of 1979, and the end of this period saw the end of the Cold War. The large blue cluster in the leftmost part of the dendrogram consists primarily of sessions from before 1971 (though it also includes the sessions in 1977 and 1995).

C. Facebook

We now consider Facebook networks for 100 US universities [21,22]. The nodes in each network represent users of the Facebook social networking site, and the unweighted edges represent reciprocated “friendships” between users at a single-time snapshot in September 2005. We consider only edges between Facebook users at the same university, as this allows us to compare the structure of the networks at the different institutions. These networks represent complete data sets obtained directly from Facebook. In contrast to the previous examples, we are not comparing snapshots of the same network at different times but are instead comparing multiple realizations of the same type of network that have evolved independently. Such real-world ensembles of network data are rare, and constructing a taxonomy will hopefully allow us to compare and contrast the social organization at these institutions.

In Fig. 11, we show the dendrogram for Facebook networks that we produced by comparing MRFs. The two color bars below the dendrogram indicate (top) the number of nodes N in each network and (bottom) the fraction of possible edges

f_e that are present (i.e., the edge density). The Facebook networks range in size from 762 to 41,536 nodes, and the edge density varies from 0.2% to 6%. In contrast to previous examples, we observe in this case that two simple network properties appear to explain most of the observed clustering of the networks. An important feature of this example is that the \mathcal{H}_{eff} , S_{eff} , and η_{eff} MRFs are each very similar in shape and lie in a narrow range across all 100 institutions (see Fig. 7). Such extreme similarity is remarkable—as one can see in Fig. 7, this contrasts starkly with most of the other examples that we examine—and it suggests that all of the Facebook networks have very similar mesoscopic structural features. If one also considers demographic information, then one can find interesting differences between the networks [21,22], but the structural similarity indicated by the MRFs is striking.

D. Fungi

We also examined fungal mycelial networks extracted from time series of digitized images of colony growth. In these undirected, planar, weighted networks, the nodes represent hyphal tips, branch points, or anastomoses (hyphal fusions), and the edges represent the interconnecting hyphal cords weighted by their conductivity [52–54]. For comparison, we also digitized weighted networks of the acellular slime mold *Physarum polycephalum* [55]. Fungal networks look like trees but contain additional edges (known as *cross-links*) that generate cycles.

As shown in Fig. 12(a), we find using our method that replicate networks from different species at comparable time points are grouped together. Furthermore, the aggregate clustering pattern reflects increasing levels of cross-linking that are characteristic of different species, as illustrated in Fig. 12(b); this ranges from the low levels in *Resinicium bicolor* to intermediate levels in *Phanerochaete velutina* and highly cross-linked networks formed by *Phallus impudicus*. By constructing a dendrogram for only one species but including data from repeated experiments and over time [see Fig. 12(c)], we observe a progression from trees at early developmental times to an increasingly cross-linked network later in mycelium growth [52,56]. In early growth, the developmental stage appears to dominate the clustering pattern, as networks from different replicates but of similar age are grouped together. At later times, however, the fungal networks exhibit a high aggregate level of similarity, and the fine-grained clustering predominantly reflects the subtle changes in structure evolving within each replicate.

E. New York Stock Exchange

As our final example, we consider a set of stock-return correlation networks for the New York Stock Exchange (NYSE), which is the largest stock exchange in the world (as measured by the aggregate US dollar value of the securities listed on it). Each node represents a stock, and the strength of the edge connecting stocks i and j is linear in the Pearson product-moment correlation coefficient between the daily logarithmic returns of the stocks [27]. We consider $N = 100$ stocks during the time period 1985–2008 and construct two networks for each year: one for the first half and one for the




FIG. 11. (Color online) Dendrogram for 100 Facebook networks of US universities at a single-time snapshot in September 2005. We order the leaves of the dendrogram to minimize the distances between adjacent nodes. The color bars below the dendrogram indicate (top) the number of nodes N in the networks and (bottom) the edge density e_d .

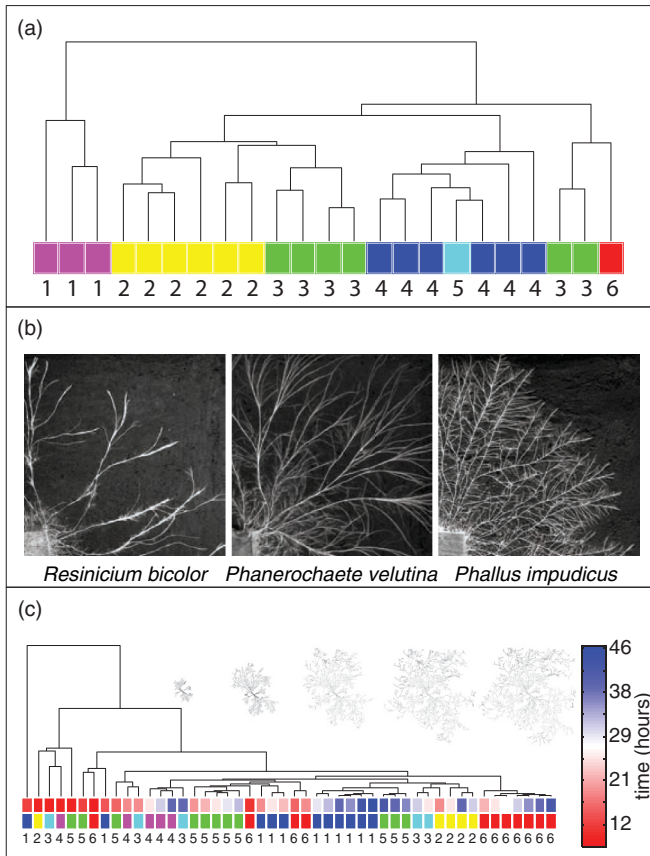


FIG. 12. (Color) (a) Dendrogram of networks for six different species of *Saprotrophic basidiomycetes* and the slime mold *Physarum polycephalum*. Each leaf represents a replicate experiment. The colors and numbers correspond to the following species: (1) *Resinicium bicolor*, (2) *Physarum polycephalum*, (3) *Phallus impudicus*, (4) *Phanerochaete velutina*, (5) *Stropharia caerulea*, and (6) *Agrocybe gibberosa*. (b) Images illustrating the network structure of the different species [53]. (c) Dendrogram of network development in six replicate time series of *Phanerochaete velutina*. We color the leaves by time, and the color bar underneath the leaves indicates experiment number (1, . . . , 6). In the inset, we show extracted networks that illustrate the transition from simple branching trees to increasing levels of interconnection (i.e., cross-linking) with time.

second half. This yields a sequence of (almost) fully connected, weighted adjacency matrices whose elements quantify the similarity of two stocks (normalized to the unit interval for each time window).

We show the dendrogram for the NYSE networks in Fig. 13. The first division of these networks classifies them into two groups (which we have colored in blue and red). The red cluster appears to correspond to periods of market turmoil, including the networks for the second half of 1987 (including the Black Monday crash of October 1987), all of 2000–2002 (including and following the bursting of the dot-com bubble), and the second half of 2007 and all of 2008 (including the credit and liquidity crisis). The value of the NYSE composite index, which measures the aggregate performance of all common stocks listed on the NYSE [57], supports our hypothesis that the red cluster is associated with periods of market turmoil. Indeed, the networks in the red cluster correspond (with one




FIG. 13. (Color online) Dendrogram for 48 NYSE networks during the period 1985–2008. For each year, we include a network both for its first half (H1) and for its second half (H2) [27]. The split of the dendrogram into two clusters (a blue group on the left and a red group on the right) appears to correspond to volatility. Leaf color indicates mean daily volatility of the composite index.

or two exceptions) to the periods of high volatility of the composite index (see Fig. 13), where we define the mean volatility during a time period as the mean absolute value of the daily return of the composite index.

X. CONCLUSIONS

We have developed an approach that facilitates the comparison of diverse networks by summarizing network community structure using what we call mesoscopic response functions (MRFs). We have demonstrated how this approach can be used to group networks both across categories and within categories. Our work builds on prior research on network community structure, which has focused predominantly on algorithmic detection of communities rather than on subsequently using communities for applications (such as comparing sets of networks).

The development of algorithmic methods to detect communities is frequently motivated by the idea that community structure in a network representing a system has some bearing on the function of that system. If different networks perform different functions—and if their functions are constrained, at least in part, by their mesoscopic structure—then it should be possible in principle to derive a functional classification of networks based on community structure. Although this has mostly been presented as a presumption in the existing literature, it is actually an empirically testable hypothesis. Indeed, we have shown in the present paper that one can systematically exploit mesoscopic structure to obtain useful comparisons of networks. This allows one to derive taxonomies for networks, and these taxonomies also appear to have correspondence with functional similarities. We observed that networks that were not grouped with other members of the same class appeared to be unusual in some respects, and we also demonstrated that we could detect historically noted financial and political changes from time-ordered sequences of networks.

We believe that our framework has the potential to aid in the exploration and exploitation of similarities in network structures across both network types and disciplinary boundaries.

ACKNOWLEDGMENTS

We are particularly grateful to A. D’Angelo and Facebook, A. Merdzanovic, M. E. J. Newman, E. Voeten, and Voteview [31] (maintained by K. Poole and H. Rosenthal) for data. We thank A. Lewis, G. Villar, S. Agarwal, D. Smith, and the anonymous referees for useful suggestions. We adapted code originally written by R. Guimerà for the simulated-annealing algorithm and code by R. Lambiotte for the Louvain algorithm. This research was supported in part by the Fulbright Program (J.P.O.) and NIH grant P01 AG-031093 (J.P.O.), the James S. McDonnell Foundation (M.A.P.: 220020177), NSF (P.J.M.: DMS-0645369), and EPSRC and BBSRC (N.S.J.: EP/I005765/1, EP/I005986/1, and BBD0201901).

APPENDIX A: ROBUSTNESS OF CLUSTERING

To examine the robustness of our clustering to false positives (false links) and false negatives (false nonlinks), we consider two network rewiring mechanisms, and we apply the rewiring to each network among 25 of the 746 networks in Table II in the Supplemental Material. We highlight these 25 networks in that table. (In the Supplemental Material, we also consider five additional subsets of 25 networks.) The first step in the procedure is to randomly rewire a number of edges corresponding to a given percentage (5%, 10%, 20%, 50%, or 100%) of the total number of edges in each network, subject to the constraints that we preserve the networks’ degree distributions and the fact that each network consists of a single connected component [58]. (That is, such a rewiring of a number of edges equal to $x\%$ of the L edges in a network amounts, in practice, to performing $\lceil xL \rceil$ rewiring steps; the same edge can be rewired multiple times.) Second, we randomly rewire a given number of the edges subject only to the constraint that each rewired network still consists of a single component.

Because we are perturbing the original networks, we focus on the distance matrices \mathbf{D}^t , \mathbf{D}^s , and \mathbf{D}^n , because they can be calculated directly for each network. We consider 25 of the 746 original networks of varying sizes and edge densities; we highlight these networks in bold in Table II of the Supplemental Material. In Fig. 14, the first column shows the matrices for the original networks. (Note that the node orderings for \mathbf{D}^t , \mathbf{D}^s , and \mathbf{D}^n are not necessarily the same in Fig. 14 because of the block diagonalization of the matrices.) The subsequent columns show the mean-distance matrices as increasing numbers of edges are rewired with the degree distributions preserved; for a given row, the node ordering in each column is fixed. The elements of the mean-distance matrices for the randomizations are given by the mean pairwise distances between networks, where the mean is calculated over all possible pairs between 10 perturbations of each network.

We now describe in more detail the procedure that we use to construct the mean-distance matrices. Let A and B represent two different (unperturbed) networks and let the sequences A_1, A_2, \dots, A_{10} and B_1, B_2, \dots, B_{10} represent 10 realizations of the perturbation process (e.g., at the 5% level) for the networks. We then calculate the mean of the ensuing $10 \times 10 = 100$ pairwise distance values. Based on visual inspection of Fig. 14, the matrices for the first few columns for




FIG. 14. (Color online) The first column shows the block-diagonalized distance matrices \mathbf{D}^t (top row), \mathbf{D}^s (middle row), and \mathbf{D}^n (bottom row) for the 25 networks listed in bold in Table II of the Supplemental Material. The other columns show block-diagonalized mean-distance matrices following randomizations of the original networks in which a given percentage of edges are rewired and the degree distributions of the networks are preserved. (We also constrain each rewired network to consist of a single connected component.) The components of the mean-distance matrices are given by the mean pairwise distances between the randomized networks’ MRFs (see the main text for a more detailed explanation). The ordering of the nodes in the plots is fixed for each row.

all of the distances are fairly similar to the original distance matrices. This suggests some notion of robustness in our clustering technique. We study only 25 networks because of the computational costs of rewiring a large number of networks multiple times; however, we have performed the same investigation for five additional subsets of 25 networks and obtained similar results. We list the networks in each subset of 25 in Table I in the Supplemental Material.

To carry out a more thorough randomization of each network, we now rewire every edge in each network 10 times on average. In Fig. 15, we show the \mathbf{D}^t , \mathbf{D}^s , and \mathbf{D}^n mean-distance matrices for this number of rewirings. We again calculate the mean distance using the method described in the previous paragraph. The first column again shows the distance matrices for the original networks. The second and third columns, respectively, show the distance matrices for randomizations in which the degree distributions are preserved and destroyed. The node orderings of the matrices in the second and third columns are again the same as the orderings for the matrix of the first column of the corresponding row. The second column in Fig. 15 demonstrates that some block structure remains in the distance matrices when the degree distributions are preserved. The third column shows that much of this structure is destroyed (though some block structure is still visible) when the degree distributions are not preserved. When the networks are “fully randomized” in this way—with the only constraint being that each rewired network must consist of a single connected component—one is in effect producing random graphs. These random graphs, however, have common properties, such as number of nodes and edge density.

APPENDIX B: COMPUTATIONAL HEURISTICS

1. Robustness of network MRFs

We detected all communities in the main text using the locally greedy Louvain algorithm [26]; however, several




FIG. 15. (Color online) Block-diagonalized distance matrices \mathbf{D}^H (top row), \mathbf{D}^S (middle row), and \mathbf{D}^η (bottom row) for the 25 networks listed in bold in Table II of the Supplemental Material. The first column shows the distance matrices for the original networks. The second column shows the mean-distance matrices following randomizations of the original networks in which 10 times the total number of edges in the networks have been rewired such that the degree distributions are preserved and the rewired networks each consist of a single connected component. The third column shows the mean-distance matrices following randomizations of the original networks in which 10 times the total number of edges in the networks have been rewired but only the fact that the networks consist of single connected components is preserved (i.e., the degree distributions are not preserved). The components of the mean-distance matrices for the randomizations are given by the mean pairwise distances between the networks.

alternative heuristics exist, so we now investigate whether the choice of heuristic has any effect on the results. In Ref. [25], Good *et al.* demonstrated that there can be extreme near-degeneracies in the energy function. In particular, there can be an exponential (or larger) number of low-energy (i.e., high-modularity) partitions. Given this, it is unsurprising that different energy-optimization heuristics can yield very different partitions for the same network. Good *et al.* suggested that the reason for such behavior is that different heuristics sample different regions of the energy landscape. Because of the potential sensitivity of results to the choice of heuristic, one should treat individual partitions by particular heuristics with caution. However, one can have more confidence in the validity of the partitions if different heuristics produce similar results. Here we compare the results for the Louvain algorithm [26] with those for a spectral algorithm [19] and simulated annealing [59].

In Fig. 16, we show MRFs for three networks calculated using Louvain [26], spectral [19], and simulated-annealing [59] algorithms. For all three networks, the three algorithms agree very closely on the shapes of the \mathcal{H} , S , and η MRFs. The MRFs are most similar for the roll-call voting network of the 102nd US Senate [32–34], and the \mathcal{H} MRF is almost identical for the three heuristics. In general, we observe the largest differences




FIG. 16. (Color online) Comparison of the MRFs produced using spectral [19], Louvain [26], and simulated-annealing [59] optimization heuristics. We show the MRFs for (a) the Zachary Karate Club network [23], (b) the roll-call voting network of the 102nd US Senate [32–34], and (c) the Garfield small-world citation network [36].

in the shapes of the MRFs when considering those produced using the spectral algorithm versus those produced using the other two algorithms. The energy values that we obtained using this particular spectral algorithm (which is a bipartitioning method) tended to be larger—i.e., less optimal—than those that we obtained from the Louvain and simulated-annealing algorithms. Moreover, the Louvain and simulated-annealing algorithms are much more popular than spectral algorithms in investigations of community structure [14] (and life is short), so we only compare results using the Louvain and simulated-annealing algorithms for the remainder of this appendix.

2. Robustness of resulting network taxonomies

Although Fig. 16 shows good agreement between the shapes of the MRFs that we obtain from the different computational heuristics, we nevertheless check that the small differences that do occur do not have a significant effect on the resulting network taxonomy. Because of the computational cost of detecting communities using simulated annealing, we investigate the effect on the taxonomy using a subset consisting of 25 small networks. We highlight all of the networks that we consider with an asterisk in Table II of the Supplemental Material. (The largest network that we include is the cat brain cortical/thalamic network [60], which has 1,170 nodes.) Indeed, MRFs for small networks tend to be much noisier than those for large networks—see, for example, Fig. 16(a), which shows the MRFs for the 34-node Zachary Karate Club network—so we expect that differences between algorithms are likely to be more pronounced for small networks.

In Fig. 17, we show dendrograms obtained using the Louvain and simulated-annealing modularity optimization algorithms for a subset of 15 networks selected uniformly at random from the 25-network subset that we mark using asterisks in Table II of the Supplemental Material. On visual

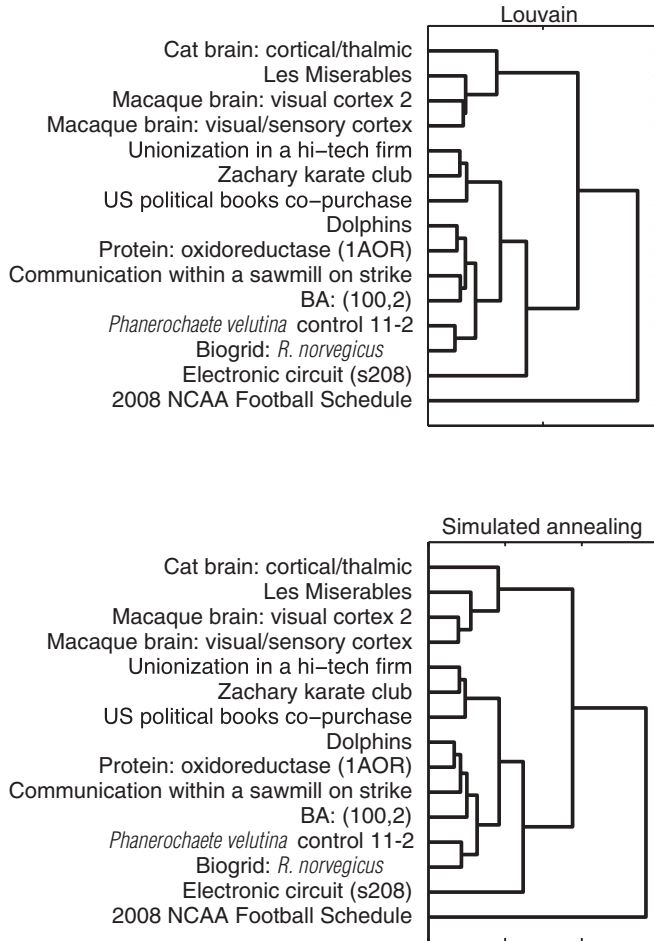


FIG. 17. Comparison of the dendrograms produced using a Louvain algorithm (top panel) and simulated annealing (bottom panel) for a subset of 15 networks. The only difference between the two dendrograms is the order in which the “Communication within a sawmill on strike” and the “BA: (100,2)” networks cluster together and the distances at which the other networks become assigned to the same cluster. The names of the networks are as they appear in Table II in the Supplemental Material.

inspection, the dendrograms appear to be very similar, as there are only a few small differences in the heights at which leaves and clusters combine. To quantify the similarity between a pair of dendrograms with underlying distance matrices denoted \mathbf{s} and \mathbf{t} , we define a correlation coefficient φ as

$$\varphi = \frac{\sum_{i < j} (s_{ij} - \bar{s})(t_{ij} - \bar{t})}{\sqrt{[\sum_{i < j} (s_{ij} - \bar{s})^2][\sum_{i < j} (t_{ij} - \bar{t})^2]}}, \quad (\text{B1})$$

where \bar{s} is the mean of the distances s_{ij} and \bar{t} is the mean of the distances t_{ij} . A pair of dendrograms derived from identical distance matrices have a correlation coefficient of $\varphi = 1$. The correlation for the example dendrograms shown in Fig. 17 is 0.997. To get a better sense of the extent of this correlation, we compare the observed correlations with those obtained for randomized dendrograms. To make the comparison, we first produce a distribution of correlation coefficients φ between pairs of dendrograms drawn from a large number of empirical (unrandomized) dendrograms




FIG. 18. (Color online) Comparison of the distributions of correlation coefficients between empirical Louvain dendrograms and empirical (red, hollow) and randomized (blue, solid) simulated-annealing dendrograms. See the text for details.

produced by the Louvain and simulated-annealing algorithms. Because of the computational costs of calculating the MRFs from the simulated-annealing algorithm, we only consider the subset of 25 networks identified in Appendix B 1. We select 15 networks uniformly at random from this subset of 25 networks and generate two dendrograms similar to those in Fig. 17: One corresponds to the distance matrix produced by the Louvain algorithm, and the other corresponds to the distance matrix produced by simulated annealing. We then calculate the correlation coefficient between the two distance matrices. We repeat this process 10,000 times to obtain 10,000 correlation coefficients, whose distribution we show using the hollow red histogram in Fig. 18. This procedure makes it possible to compare a large number of dendrograms at the computational cost of calculating simulated annealing MRFs for a total of 25 networks, which we highlight with asterisks in Table II of the Supplemental Material.

We compare this observed distribution of correlation coefficients to a randomized reference. We focus on the correlation between empirical Louvain dendrograms (i.e., empirical dendrograms that result from distance matrices produced by the Louvain method) and randomized simulated-annealing dendrograms (i.e., dendrograms that result from distance matrices produced by the simulated-annealing algorithm and are then randomized). We proceed as follows: For each of the 10,000 dendrogram pairs that we assembled from 15-network subsets, we create 100 randomizations of the simulated-annealing dendrogram, and we then calculate the correlation coefficient between each of these randomized dendrograms and the corresponding empirical Louvain dendrogram. The resulting distribution from 10,000 repetitions is the solid blue histogram in Fig. 18. To randomize the simulated-annealing dendrogram, we used the double-permutation procedure described in Refs. [61,62]. This procedure has two steps. First, we randomize the distances at which the different clusters are combined. For example, consider an unrandomized dendrogram in which clusters A and B are combined at a distance of 0.45, and clusters C and D are combined at a distance of 0.65; after the randomization, A and B might be combined at a distance of 0.65, and C and D might be combined at a distance of 0.45. Second, we randomize the networks corresponding to each leaf in the dendrogram. This two-step randomization procedure maintains the underlying distances and the topology of the dendrogram.

As mentioned above, we show the distributions of correlation coefficients between empirical Louvain dendrograms and both empirical (unrandomized) and randomized simulated-annealing dendrograms in Fig. 18. The correlation is clearly much higher for the empirical case, as there is only a very slight overlap in the tails of the two distributions. The correlation between the Louvain and simulated-annealing dendrograms is greater than 0.99 for about 63% of the studied dendrograms.

APPENDIX C: DIAGNOSTIC FOR ASSESSING THE CLUSTERING FROM DIFFERENT DISTANCE MEASURES

An examination of the leaf colors of the dendrogram in Fig. 7 illustrates that the employed distance measure groups networks from a variety of categories—including political voting networks, political committee networks, Facebook networks, metabolic networks, and fungal networks. A visual comparison provides a reasonable starting point for assessing the effectiveness of different distance measures at clustering networks. To quantify how effectively each distance matrix (\mathbf{D}^H , \mathbf{D}^S , \mathbf{D}^n , and \mathbf{D}^p) clusters networks of the same type, we introduce a *clustering diagnostic*, which we denote by $\alpha(h)$ and will explain shortly. Because the assignment of networks to categories is subjective and because some of the categories include networks of very different types, it would be inappropriate to assess the effectiveness of a distance measure based on how well it clusters networks in very broad categories. We thus focus our examination on narrower categories whose constituent networks are clustered fairly tightly in Fig. 7. This includes the following eight categories of networks: Facebook, metabolic, political cosponsorship, political committee, political voting, financial, brain, and fungal.

The value of the clustering diagnostic $\alpha(h)$ depends on where one “cuts” the dendrograms. We start by constructing a dendrogram for each of the four distance matrices \mathbf{D}^H , \mathbf{D}^S , \mathbf{D}^n , and \mathbf{D}^p . Performing a horizontal cut through a dendrogram at a given height h splits the dendrogram into multiple disconnected clusters (see Fig. 17). For each such cluster, we calculate the proportion of networks from a particular category that are contained in it. For example, if a cut produces three clusters and if we consider the Facebook category, then we might find that one cluster contains two tenths of the Facebook networks, a second cluster has three tenths of those networks, and the third cluster has the remaining half of those networks. We calculate these membership fractions for each network category and for each cluster. We then identify, for each category, what we call the *plurality cluster*, which we define as the cluster that includes the largest fraction of networks from that category. In the above example, the third cluster is the plurality cluster for the Facebook category. Our diagnostic $\alpha(h)$ is then defined by adding across all eight categories the fraction of networks in the plurality clusters:

$$\alpha(h) = \sum_{j=1}^8 \gamma_j(h), \quad (\text{C1})$$

where $\gamma_j(h)$ is the plurality fraction for the j th category of networks for the given cut at height h of the dendrogram.




FIG. 19. (Color online) Comparison of the effectiveness of the employed distance measures at clustering networks of the same category. As discussed in the main text, we quantify this using the clustering diagnostic $\alpha(h)$. We calculate dendrograms from four distance matrices (\mathbf{D}^H , \mathbf{D}^S , \mathbf{D}^n , and \mathbf{D}^p) and compare the resulting values of $\alpha(h)$ for different sets of categories. (a) The value of the clustering diagnostic $\alpha(h)$ as a function of dendrogram cut level h (i.e., where the dendrogram is split to clusters) for the following eight categories of networks: Facebook, metabolic, political cosponsorship, political committee, political voting, financial, brain, and fungal. (b) The value of $\alpha(h)$ for the largest five of the above eight categories (Facebook, metabolic, political cosponsorship, political committee, and political voting) and (c) for the smallest five of the above eight categories (Facebook, metabolic, financial, brain, and fungal). The maximum possible value of $\alpha(h)$ in each panel is equal to the number of categories considered in each panel. The values of $\alpha(h)$ obtained using the PCA-distance matrix \mathbf{D}^p (gray solid curve) are usually higher than those obtained using the other three distance measures. This suggests that PCA distance is the most effective of the four employed clustering measures.

We perform similar calculations for each level of the dendrogram and use the resulting values of $\alpha(h)$ to assess the effectiveness of the different distance measures at clustering the networks. For example, at the root of the dendrogram, all of the networks are in a single cluster, so the maximum fraction of networks in the same cluster is one for every network category. Given the above choice of eight categories, this yields $\alpha = 8$. However, as one considers lower levels of the dendrogram, the clusters break up more and more, so the fraction of networks in the plurality cluster in each category typically decreases. Effective distance measures ought to result in relatively high values for $\alpha(h)$ for many values of h .

In Fig. 19, we compare the values of $\alpha(h)$ at each level of the dendrogram for \mathbf{D}^H , \mathbf{D}^S , \mathbf{D}^n , and \mathbf{D}^p . For each of the different subsets of networks and for most of the dendrogram levels, the PCA distance matrix \mathbf{D}^p is the most effective of the employed distance measures at clustering networks of the same category. This agrees with our visual assessment (i.e., our identification of contiguous blocks of color) of the different measures.

APPENDIX D: USING SIMPLE CHARACTERISTICS TO CLUSTER NETWORKS

We established in Section VIII that the PCA distance matrix \mathbf{D}^p between MRFs can produce sensible network taxonomies, and we now consider briefly whether the observed taxonomies can be explained using simple summary statistics. We consider




FIG. 20. (Color online) Taxonomy for 189 networks. We constructed the dendrogram using the distance matrix \mathbf{D}^p and average linkage clustering. We order the leaves of the dendrogram to minimize the distances between adjacent nodes, and we color the leaves to indicate the type of network. The three color bars below the dendrogram indicate whether the network corresponding to each leaf is weighted or unweighted (top), the number of nodes N in the networks (middle), and the edge density e_d (bottom).

only a few specific properties, though, of course, there are myriad other network diagnostics that one might consider.

Perhaps the three simplest properties of an undirected network are the following: (1) whether it has weighted or unweighted edges; (2) the number of nodes N ; and (3) the

edge density $e_d = 2L/[N(N - 1)]$ (where L is the number of edges, which we distinguish from the total edge weight m in weighted networks). The top colored row in Fig. 20 indicates that many of the weighted networks are clustered together at the far left of the dendrogram. However, there are also weighted networks scattered throughout the dendrogram, so whether a network is weighted or unweighted does not explain the observed classification. The third colored row provides a clearer explanation for the cluster of networks at the left: These are not simply weighted networks, as they are in fact similarity networks, so nearly all possible edges are present and they have weights indicating connection strengths. However, this property alone cannot explain the observed classification, as several of the weighted networks containing nearly all possible edges do not appear at the far left of the dendrogram. In fact, there are many clusters in the dendrogram that contain networks with very different fractions of possible edges. The total number of nodes, shown by the second colored row in the figure, again explains some of the clustering, as networks with similar numbers of nodes are clustered together in some regions of the dendrogram. However, there are also numerous examples in which networks with the same number of nodes appear in different clusters. Therefore, none of these three simple network diagnostics can explain the observed classification by itself.

-
- [1] M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).
- [2] M. E. J. Newman, *Phys. Today* **61**, 33 (2008).
- [3] K. Börner, *Atlas of Science: Visualizing What We Know* (MIT Press, Cambridge, MA, 2010).
- [4] The code used for constructing network taxonomies in this paper is available at www.jponnela.com.
- [5] R. Albert and A.-L. Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
- [6] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, *Science* **298**, 824 (2002).
- [7] R. Milo, S. Itzkovitz, N. Kashtan, R. Levittand, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, and U. Alon, *Science* **303**, 1538 (2004).
- [8] R. Guimerà, M. Sales-Pardo, and L. A. N. Amaral, *Nature Phys.* **3**, 63 (2007).
- [9] M. A. Serrano, M. Boguná, and A. Vespignani, *Proc. Natl. Acad. Sci. USA* **106**, 6483 (2009).
- [10] A. Arenas, J. Borge-Holthoefer, A. Gómez, and G. Zamora-López, *New J. Phys.* **12**, 053009 (2010).
- [11] J. Bagrow, E. Bollt, J. Skufca, and D. Ben-Avraham, *Europhys. Lett.* **81**, 68004 (2008).
- [12] C. Song, S. Havlin, and H. A. Makse, *Nature* **433**, 392 (2005).
- [13] M. A. Porter, J.-P. Onnela, and P. J. Mucha, *Not. Am. Math. Soc.* **56**, 1082 (2009).
- [14] S. Fortunato, *Phys. Rep.* **486**, 75 (2010).
- [15] A. Lancichinetti, M. Kivela, J. Saramaki, and S. Fortunato, *PLoS ONE* **5**, e11976 (2010).
- [16] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.86.036104> for a taxonomy for all 746 networks, details on the models used to generate synthetic networks, and a summary of properties and literature references for all networks.
- [17] J. Reichardt and S. Bornholdt, *Phys. Rev. E* **74**, 016110 (2006).
- [18] M. E. J. Newman and M. Girvan, *Phys. Rev. E* **69**, 026113 (2004).
- [19] M. E. J. Newman, *Proc. Nat. Acad. Sci. USA* **103**, 8577 (2006).
- [20] J. Reichardt and S. Bornholdt, *Phys. Rev. Lett.* **93**, 218701 (2004).
- [21] A. L. Traud, E. D. Kelsic, P. J. Mucha, and M. A. Porter, *SIAM Rev.* **53**, 526 (2011).
- [22] A. L. Traud, P. J. Mucha, and M. A. Porter, *Physica A* **391**, 4165 (2012).
- [23] W. W. Zachary, *J. Anthropol. Res.* **33**, 452 (1977).
- [24] U. Brandes, D. Delling, M. Gaertler, R. Görke, M. Hoefer, Z. Nikoloski, and D. Wagner, *IEEE Trans. Knowl. Data Eng.* **20**, 172 (2008).
- [25] B. H. Good, Y. A. de Montjoye, and A. Clauset, *Phys. Rev. E* **81**, 046106 (2010).
- [26] V. D. Blondel, J. Guillaume, R. Lambiotte, and E. Lefebvre, *J. Stat. Mech.* (2008) P10008.
- [27] J.-P. Onnela, A. Chakraborti, K. Kaski, J. Kertész, and A. Kanto, *Phys. Rev. E* **68**, 056110 (2003).
- [28] A. Arenas, A. Fernandez, and S. Gómez, *New J. Phys.* **10**, 053039 (2008).
- [29] O. Sporns, *BioSystems* **85**, 55 (2006).
- [30] D. Firth and A. Spirling, working paper (2003), <http://www2.warwick.ac.uk/fac/sci/statistics/staff/academic/firth/software/tapir/firth-spirling.pdf>.
- [31] K. T. Poole (2009), data available at <http://voteview.com>.
- [32] K. T. Poole and H. Rosenthal, *Congress: A Political-Economic History of Roll Call Voting* (Oxford University Press, Oxford, 1997).

- [33] N. M. McCarty, K. T. Poole, and H. Rosenthal, *Polarized America: The Dance of Ideology and Unequal Riches* (MIT Press, Cambridge, MA, 2007).
- [34] A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter, arXiv:0907.3509 (2012).
- [35] C. Stark, B. J. Breitkreutz, T. Reguly, L. Boucher, A. Breitkreutz, and M. Tyers, *Nucleic Acids Res.* **34**, D535 (2006).
- [36] E. Garfield, I. H. Sher, and R. J. Torpie, *The Use of Citation Data in Writing the History of Science* (Institute for Scientific Information, Philadelphia, 1964).
- [37] P. Erdős and A. Rényi, *Publ. Math.* **6**, 290 (1959).
- [38] A.-L. Barabási and R. Albert, *Science* **286**, 509 (1999).
- [39] D. J. Watts and S. H. Strogatz, *Nature (London)* **393**, 440 (1998).
- [40] I. T. Jolliffe, *Principal Component Analysis* (Springer-Verlag, New York, 1986).
- [41] S. Maslov and K. Sneppen, *Science* **296**, 910 (2002).
- [42] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification* (Wiley, New York, 2001).
- [43] The three most common linkage-clustering algorithms are single, average, and complete linkage clustering; they join clusters of objects based on the smallest, mean, and largest distance between objects in clusters, respectively. We used the cophenetic correlation to compare how well each algorithm preserves the pairwise distances between the networks [63]. If we let t_{ij} represent the distance in a dendrogram between networks i and j , then for a dendrogram constructed from the PCA-distance matrix \mathbf{D}^p with elements d_{ij}^p , the cophenetic correlation is defined as
$$\frac{\sum_{i < j} (d_{ij}^p - \bar{d}^p)(t_{ij} - \bar{t})}{\sqrt{\left[\sum_{i < j} (d_{ij}^p - \bar{d}^p)^2 \right] \left[\sum_{i < j} (t_{ij} - \bar{t})^2 \right]}}$$
, where \bar{d}^p is the mean value of the lower-triangular elements of the PCA distance matrix \mathbf{D}^p and \bar{t} is the mean distance between the leaves of the dendrogram. We obtain the following cophenetic correlations for the different linkage-clustering algorithms: 0.65 (single), 0.78 (average), and 0.62 (complete). Hence, dendrograms constructed using average linkage clustering preserve the distances in \mathbf{D}^p better than those constructed using the other two clustering techniques. We thus use average linkage clustering to construct all of the dendrograms in this paper.
- [44] P. Gleiser and L. Danon, *Adv. Complex Sys.* **6**, 565 (2003).
- [45] L. Hakes, J. W. Pinney, D. L. Robertson, and S. C. Lovell, *Nat. Biotechnol.* **26**, 69 (2008).
- [46] H. Yu *et al.*, *Science* **322**, 104 (2008).
- [47] R. Alberich, J. Miro-Julia, and F. Rossello, arXiv:cond-mat/0202174 (2002).
- [48] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, *Science* **328**, 876 (2010).
- [49] S. Kernell, G. C. Jacobson, and T. Kousser, *The Logic of American Politics*, 4th ed. (CQ Press, Washington, DC, 2009).
- [50] K. T. Macon, P. J. Mucha, and M. A. Porter, *Physica A* **391**, 343 (2012).
- [51] We excluded the 19th session because only one resolution was considered by the UNGA that year.
- [52] D. P. Bebber, J. Hynes, P. R. Darrah, L. Boddy, and M. D. Fricker, *Proc. Roy. Soc. B* **274**, 2307 (2007).
- [53] M. Fricker, L. Boddy, and D. Bebber, in *The Mycota. Vol VIII, Biology of the Fungal Cell*, edited by R. J. Howard and N. A. R. Gow (Springer, Berlin, 2007), pp. 307–328.
- [54] M. D. Fricker, L. Boddy, T. Nakagaki, and D. P. Bebber, in *Adaptive Networks: Theory, Models and Applications*, edited by T. Gross and H. Sayama (Springer, Berlin, 2009), pp. 51–70.
- [55] A. Tero, S. Takagi, T. Saigusa, K. Ito, D. P. Bebber, M. D. Fricker, K. Yumiki, R. Kobayashi, and T. Nakagaki, *Science* **327**, 439 (2010).
- [56] L. L. M. Heaton, E. López, P. K. Maini, M. D. Fricker, and N. S. Jones, *Proc. Roy. Soc. B* **277**, 3265 (2010).
- [57] See <http://www.nyse.com/>
- [58] S. Maslov and K. Sneppen, *Science* **296**, 910 (2002).
- [59] R. Guimerà, M. Sales-Pardo, and L. A. N. Amaral, *Phys. Rev. E* **70**, 025101 (2004).
- [60] J. W. Scannell, G. A. P. C. Burns, C. C. Hilgetag, M. A. O’Neil, and M. P. Young, *Cereb. Cortex* **9**, 277 (1999).
- [61] F.-J. Lapointe and P. Legendre, *J. Classif.* **12**, 265 (1995).
- [62] F.-J. Lapointe and P. Legendre, *Syst. Zool.* **39**, 1 (1990).
- [63] R. R. Sokal and F. J. Rohlf, *Taxon* **11**, 33 (1962).