

Weight thresholding on complex networks

Xiaoran Yan,^{1,*} Lucas G. S. Jeub,² Alessandro Flammini,^{1,2} Filippo Radicchi,² and Santo Fortunato^{1,2}

¹Indiana University Network Science Institute (IUNI), 1001 E SR 45/46 Bypass, Bloomington, Indiana 47408, USA

²School of Informatics, Computing and Engineering, Indiana University, 700 N Woodlawn Ave, Bloomington, Indiana 47408, USA



(Received 21 June 2018; published 8 October 2018)

Weight thresholding is a simple technique that aims at reducing the number of edges in weighted networks that are otherwise too dense for the application of standard graph-theoretical methods. We show that the group structure of real weighted networks is very robust under weight thresholding, as it is maintained even when most of the edges are removed. This appears to be related to the correlation between topology and weight that characterizes real networks. On the other hand, the behavior of other properties is generally system dependent.

DOI: [10.1103/PhysRevE.98.042304](https://doi.org/10.1103/PhysRevE.98.042304)

I. INTRODUCTION

Many real networks have weighted edges [1], representing the intensity of the interaction between pairs of vertices. Also, some weighted networks, e.g., financial [2] and brain networks [3], have a high density of edges. Analyzing very dense graphs with tools of network science is often impossible, unless some preprocessing technique is applied to reduce the number of connections. Different recipes of edge pruning, or *graph sparsification*, have been proposed in recent years [4–10].

In practical applications such as the preprocessing of data about brain, financial, and biological networks [2,11,12], weight thresholding is the most popular approach to sparsification. It consists in removing all edges with weight below a given threshold. Ideally, one would like to eliminate as many edges as possible without drastically altering key features of the original system. A recent study of functional brain networks investigates how the graph changes as a function of the threshold value [13], finding that conventional network properties are usually disrupted early on by the pruning procedure. Further, many standard measures do not behave smoothly under progressive edge removal, and hence are not reliable measures to assess the effective change to the system structure induced by the removal of edges.

By analyzing several synthetic and real weighted networks, we show that, while local and global network features are often quickly lost under weight thresholding, the procedure does not alter the mesoscopic organization of the network [14]: groups (e.g., communities) survive even when most of the edges are removed.

In addition, we introduce a measure, the *minimum absolute spectral similarity* (MASS), that estimates the variation of spectral properties of the graph when edges are removed. Spectral properties are theoretically related to group structures in general [15–17]. The MASS is stable under weight thresholding for many real networks, making it a potential

alternative to expensive community detection algorithms for testing the robustness of communities under thresholding.

II. METHODS

Let us consider a weighted and undirected graph G composed of N vertices and M edges. Edges have positive weights, and the graph topology is described by the symmetric weight matrix W , where the generic element $W_{uv} = W_{vu} > 0$ if there is a weighted edge between nodes u and v , while $W_{uv} = W_{vu} = 0$ otherwise. Weight thresholding removes all edges with weight lower than a threshold value. This means that the resulting graph \tilde{G} has a thresholded weight matrix \tilde{W} , whose generic element $\tilde{W}_{uv} = \tilde{W}_{vu} = W_{uv}$ if $W_{uv} \geq \theta$, and $\tilde{W}_{uv} = \tilde{W}_{vu} = 0$ otherwise. The thresholded graph \tilde{G} is therefore a subgraph of G with the same number of nodes.

We first examined synthetic networks generated by the Lancichinetti, Fortunato, and Radicchi (LFR) benchmark [18]. We then extend the analysis to several real networks from different domains: structural brain networks [19], the world trade network [20], the airline network [21], and the coauthorship network of faculty of Indiana University. We treat all networks as undirected, weighted graphs with positive edge weights. We show the variation of several graph properties as a function of the fraction of removed edges. The properties we have chosen include all network-level measures from Ref. [13] as well as mesoscopic structure measures:

(1) *Characteristic path length (CPL)*, the average length of all shortest paths connecting pairs of vertices of the network. Once the network becomes disconnected, CPL is not defined, and we set its value to 0.

(2) *Global efficiency*, the average inverse distance between all pairs of vertices of the network [22]. Global efficiency remains well defined after network disconnection, as disconnected vertex pairs simply have a inverse distance of 0.

(3) *Transitivity*, or global clustering coefficient, is the ratio of triangles to triplets in the network, where a triplet is a motif consisting of one vertex and two links incident to the vertex.

(4) *Community structure*, which we detect with two methods. The first is based on modularity maximization [23] via the

*Corresponding author: Xiaoran.a.yan@gmail.com

Louvain algorithm [24,25]. The second method uses k -means spectral clustering [15] with a constant k equal to the one found with the Louvain algorithm.

(5) *Core-periphery structure*, the vertices of which we detect using the method introduced for the weighted coreness measure [26].

(6) *DeltaCon*, a metric indicating the similarity between the original graph and the thresholded one based on graph diffusion properties [27].

For all three partitioning algorithms, i.e., Louvain, k -means, and coreness, we measure the similarity between the partitions of the sparsified graph \tilde{G} and those of G using the *adjusted mutual information* (AMI) [28]. To overcome the randomness of the partitioning algorithms, we sample 100 different partition outcomes from the original graph, 10 from the sparsified graph, and use the maximum AMI between any pair. These numbers are picked so that the same algorithm returns consistent results (AMI very close to 1) on independent runs on the original graph.

For global efficiency we take the ratio between the value of the measure on the sparsified graph and the corresponding value on the initial graph. CPL grows with edge removal, and we therefore normalize it by its largest value before disconnection. Since DeltaCon scales directly with weights, we normalize each matrix entry such that it reaches 0 on empty graphs. This way all our measures are confined in the interval $[0,1]$, and their trends can be compared (transitivity naturally varies in this range). We also plot the relative size of the largest connected component to keep track of splits of the network during edge removal. More details of how these graph properties are calculated are given in Appendix A.

We also add another measure, capturing the variation of spectral properties of the graph. To define this measure we recall that the Laplacian of G is defined as $L = D - W$, where D is the diagonal matrix of the weighted degrees (strengths), with entries $D_u = \sum_v W_{uv}$. Similarly, for the sparsified graph \tilde{G} , $\tilde{L} = \tilde{D} - \tilde{W}$. The Laplacian has the spectral decomposition

$$L = D - W = V \Lambda V^{-1},$$

where the columns of the matrix V are the eigenvectors v_1, v_2, \dots, v_N of the Laplacian, and the entries in the diagonal matrix Λ are the corresponding eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$.

The difference between the sparsified Laplacian \tilde{L} and the original L can be quantified by the *minimum relative spectral similarity* (MRSS) [29],

$$\sigma_{\min}^R = \min_{x^T L x} \frac{x^T \tilde{L} x}{x^T L x}, \quad (1)$$

where $x^T L x$ is the Laplacian quadratic form and x any N -dimensional real vector. The MRSS is a direct adaptation of relative spectral bounds. Intuitively, the input vector x determines the “direction” along which we measure the change of the graph, and by taking the minimum we consider the worst case scenario. However, the value of MRSS drops to zero as soon as \tilde{G} becomes disconnected. Because of this mathematical degeneracy, it is also numerically unstable for many optimization algorithms.

To overcome this issue, we instead propose the *absolute spectral similarity* with respect to the input vector x , defined as

$$\sigma(x) = 1 - \frac{x^T \Delta L x}{x^T \lambda_N x}, \quad (2)$$

where λ_N is the largest eigenvalue of the original graph Laplacian, and $\Delta L = \Delta D - \Delta W$ is the graph Laplacian of the *difference graph* ΔG , whose vertices are the same as in G , while the edges are the ones removed by the chosen sparsification procedure. Without loss of generality, we consider only unit length input vectors, $|x| = 1$. In Appendix B, we prove that weight thresholding optimizes the expected value of $\sigma(x)$, if we assume the entries of the vector x are independent identically distributed random variables.

Since the input vector is variable, we again consider the worst case scenario and use the *minimum absolute spectral similarity* (MASS),

$$\sigma_{\min} = \min_{|x|=1} \left(1 - \frac{x^T \Delta L x}{\lambda_N} \right) = \frac{\lambda_N - \lambda_N^\Delta}{\lambda_N}, \quad (3)$$

where λ_N^Δ is the largest eigenvalue of the difference Laplacian ΔL .

MASS is consistent with our intuition that disconnecting a few peripheral vertices, while a bigger change than removing redundant edges, should have a small impact on the organization of the system. Many local and global network features become ill-defined as soon as the network disconnects, and hence they are not reliable measures to assess the effect of thresholding. Mesoscopic properties, like communities, on the other hand, remain meaningful even after the network becomes disconnected, and a reliable measure should be robust in such situations.

A major advantage of MASS is its numerical stability and computational efficiency (see Appendix C). Defined as the ratio of the largest eigenvalues of two graph Laplacians, it can be computed using standard numerical libraries [30–32]. (Readers can find our Matlab implementation at <https://github.com/IU-AMBITION/MASS>.) In spite of its simplicity, MASS satisfies a series of theoretical properties (see Appendix E).

III. RESULTS

To illustrate the changes of the above graph properties under weight thresholding, we first compare a synthetic weighted network with its binary counterpart; see Fig. 1. The binary network with 1000 vertices and eight planted communities is generated using the LFR benchmark [18] (the parameters of the LFR benchmarks used here are listed in Appendix A), which produces realistic networks by capturing power-law distributions of both degree and community size. Since all edges have the same weight, the thresholding is done by removing a fraction of randomly selected edges [in Fig. 1(a), all curves are averaged over 10 random realizations]. While the network remains well connected even after we remove 50% of edges, the AMI scores relative to the mesoscopic group structures drop already when a small fraction of edges is removed. The MASS curve decays even faster,

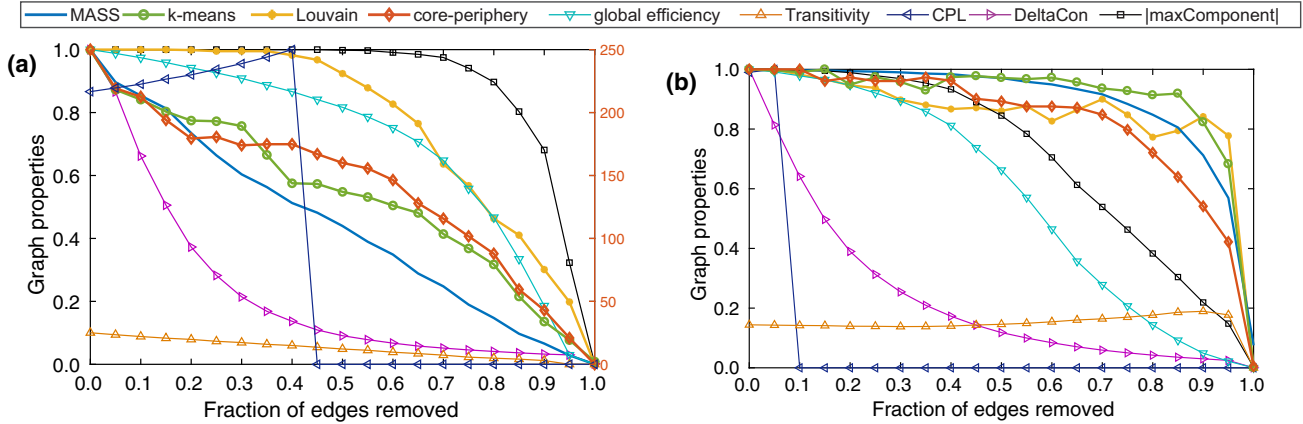


FIG. 1. Variation of MASS and other variables under weight thresholding, for synthetic binary networks generated by the LFR benchmark (a) and a weighted counterpart with noisy correlation between edge weight and the degrees of its endpoints (b). The “|maxComponent|” curve represents the fraction of vertices remaining in the largest connected component.

following a diagonal line, which suggests that it captures changes in spectral properties that are not directly related to the mesoscopic structure of the network.

In many real weighted networks, there is a power law relation between the degree of a vertex and its strength (i.e., the total weight carried by the edges adjacent to the vertex) [1,33]. We first set the weight of an edge to be proportional to the product of the degrees of its endpoints, changing the binary network into a weighted one. To account for the presence of noise we added a uniform error in the range of $[-w, w]$ on top of each weighted edge, with w being the corresponding edge weight. The effect of thresholding on these networks is shown in Fig. 1(b). As a result, the network’s group structures are now robust under weight thresholding, despite the fact that a substantial fraction of vertices become disconnected at an early stage. In this realistic weighted network, the MASS curve now has a quite similar trend as the curves describing the variation of group structure. On the other hand the CPL becomes ill-defined as soon as the network is disconnected. DeltaCon also drops much earlier compared with other measures, as diffusion is heavily affected by network disconnection.

Let us now discuss the analysis of the real networks. We list their basic properties in Table I. All four networks have positive weight-degree correlation coefficient, as it often happens in real weighted networks [1]. Another observation is that they also have high values of the weighted coreness measure [26] (except for structural brain networks), which means weakly connected peripheral vertices will quickly become disconnected.

The structural brain network is built from diffusion weighted imaging MRI scans of 40 experiment participants [19]. Each personal network has 234 vertices representing brain regions. On average, they have 7862 weighted edges representing the fiber density connecting couples of regions. Here we threshold each network individually and considered the population average of each graph property for our analysis.

Figure 2(a) shows that the network is robust when the light edges are removed: it takes the removal of a substantial fraction of edges to get appreciable changes in all measures. In particular, the MASS is basically unaffected until more

than half of the edges are deleted, and it follows qualitatively the trend of the similarity (AMI) of the mesoscopic group partitions.

The airline network is constructed from public data on flights between major airports around the world, with the edge weight representing the number of flights as well as the capacity of the plane operating each flight [21]. Small airports have only weak connections to the rest of the system, due to the limited traffic they handle. On the other hand, major hubs have some of the strongest connections, leading to a strong core-periphery structure with 0.5848 weight-degree correlation. The system gets quickly disconnected under weight thresholding, and we see big drops in other graph properties early on, including CPL, DeltaCon, and global efficiency, as shown in Fig. 2(b). Community and core-periphery structure, on the other hand, remain fairly stable against edge removal, and the MASS again follows a similar trend as the group partition similarity (AMI) curves.

The world trade network is constructed from economic trading data between 250 countries. We aggregate the directed edges into 18 389 undirected edges representing bidirectional trade volumes [20]. This network is an extreme example of strong core-periphery structure with close to 1 weight-degree correlation (0.9539) and coreness measure (0.8373). As a result [see Fig. 2(c)], it can be sparsified very aggressively without large variations in all measures except for CPL and transitivity. All three mesoscopic group structures, k -means, Louvain, and core periphery, have very stable AMI scores. The MASS values also reflect the same trend.

The coauthor network captures the academic collaborations between institutions revealed through papers authored by Indiana University faculty. It is built from Thomson Reuters Web of Science data (Web of Knowledge version 5 [34]) from 2008 to 2013. Like the airline network, it has a strong core-periphery structure with 0.6247 weight-degree correlation. However, its peripheral vertices get disconnected at a much slower rate. Again, the mesoscopic group structure remains robust under thresholding, captured by k -means, Louvain, core-periphery, and the MASS curves [see Fig. 2(d)]. In contrast, CPL, DeltaCon, transitivity, and global efficiency again demonstrate very different patterns.

TABLE I. Networks studied in this paper and their properties.

Name	No. of samples	No. of vertices	No. of edges	No. of k (Louvain)	No. of weight degree correlation	No. of coreness
Structural brain networks	40	234	4046 (mean)	9.9 (mean)	0.2603 (mean)	0.3009 (mean)
World trade network	1	250	18 389	3	0.9539	0.8373
Airline network	1	3253	18 997	20	0.5848	0.6776
Coauthorship network	1	2855	75 058	7	0.6247	0.8223

We conclude that group structure, including community and core-periphery structures, is a very robust feature that survives even when most edges are removed. The MASS is also quite consistent in capturing group structure across these real-world networks. The variation of the measure is rather smooth and barely affected by the disconnection of small subgraphs, making it a stable and efficient measure for evaluating thresholding effects. We remark that the empirically observed robustness of the group structure does not hold if the weight-degree correlation is destroyed (see Appendix D).

While we currently lack a full theoretical understanding of the relationship between MASS and community structure, here we provide some mathematical justification. Matrix perturbation theory studies the change of graph spectrum of a real symmetric matrix A by “perturbing” it. In our context, if we treat the original Laplacian matrix L as A and the difference Laplacian ΔL as the perturbation, the classical

Weyl’s theorem directly relates to MASS,

$$\frac{|\lambda_i - \tilde{\lambda}_i|}{\lambda_N} \leq \frac{|\lambda_N^\Delta|}{\lambda_N} = 1 - \sigma_{\min}, \quad 1 \leq i \leq N.$$

We now consider a generalized version of the Davis-Kahan theorem on subspaces,

$$\|\sin \angle(V_k, \tilde{V}_k)\|_F \leq \frac{2\sqrt{k}|\lambda_N^\Delta|}{\delta_k}, \quad 1 \leq k \leq N,$$

where $\|A\|_F$ represents the Frobenius norm of the matrix A , V_k, \tilde{V}_k are the k -dimensional subspaces spanned by the eigenbasis v_1, v_2, \dots, v_k and $\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_k$, respectively, $\delta_k = \lambda_{k+1} - \lambda_k$, and $\sin \angle(V_k, \tilde{V}_k)$ is a diagonal matrix whose entries are the sines of the angles between the corresponding eigenvectors in the two eigenbasis. If σ_{\min} approaches 1, or equivalently λ_N^Δ approaches 0, we then have a tight upper bound on the rotation of the eigenbasis v_1, v_2, \dots, v_k . According to spectral graph theory, these smaller eigenvectors play a

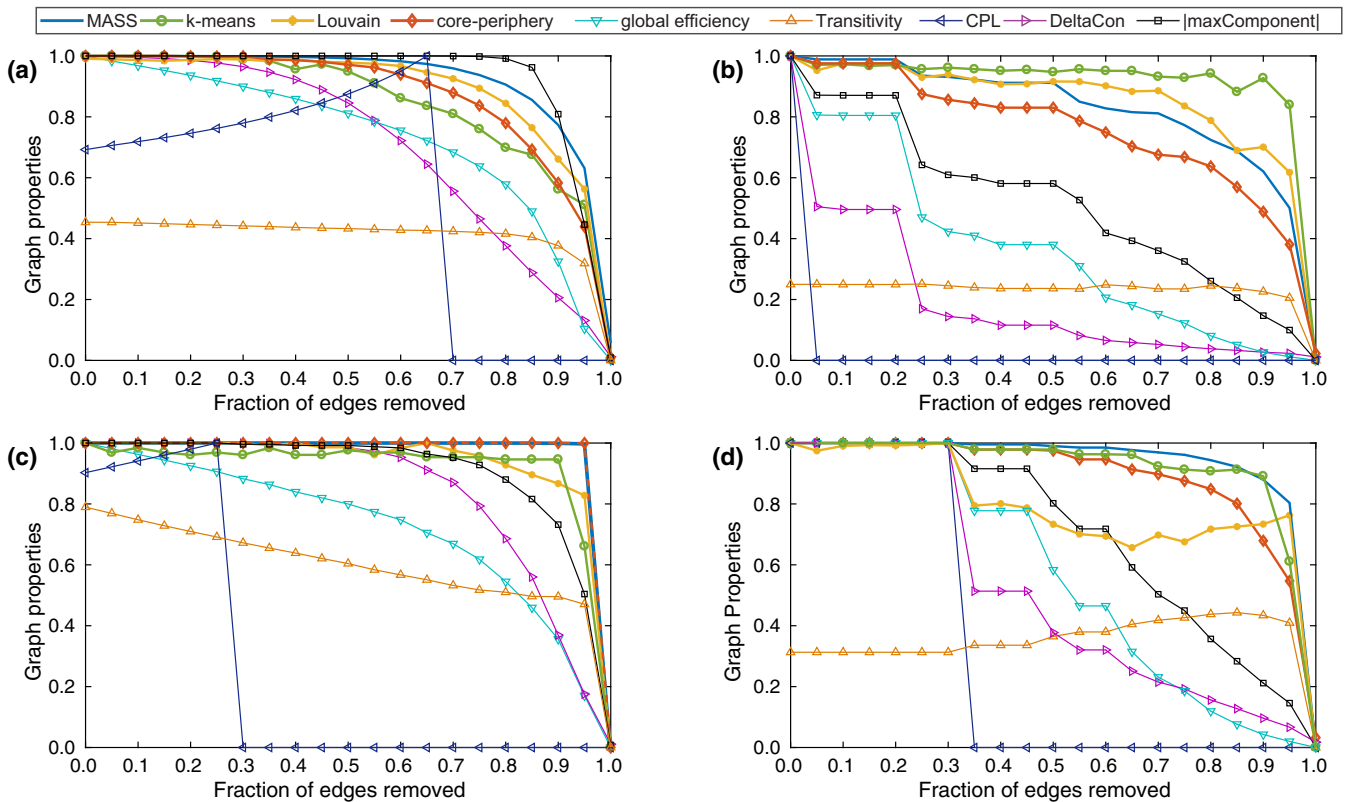


FIG. 2. Variation of MASS and other variables under weight thresholding, for structural brain networks (a), the global airline network (b), the world trade network (c), and the coauthorship network (d). The “|maxComponent|” curve represents of the fraction of vertices remaining in the largest connected component.

fundamental role in defining community structures [15,16]. Preservation of MASS therefore guarantees a one-sided upper bound on the change to community structure.

IV. CONCLUSIONS

We have carried out a detailed analysis of weight thresholding on weighted networks. In general, it appears that group structure is fairly robust under weight thresholding, in contrast to other features. We found that this is due to the peculiar correlation between weight and degree that is commonly observed in real networks, according to which large weights are more likely to be carried by links attached to high degree vertices.

We have also introduced a measure, the minimum absolute spectral similarity (MASS), to estimate the effect that sparsification procedures have on spectral features of the network. In case studies above we have seen that MASS behaves similarly to traditional group structure measures when there is correlation between weight and degree.

This work deals with weight thresholding, but the analysis can be easily repeated with more sophisticated graph sparsification methods. In the future, we plan to investigate more closely the relationship between MASS and group structure, as well as the role played by the weight-degree correlation.

ACKNOWLEDGMENTS

F.R. acknowledges financial support through National Science Foundation (CMMI-1552487) and U.S. Army Research Office (W911NF-16-1-0104).

APPENDIX A: ALGORITHMIC DETAILS OF GRAPH PROPERTIES

Here we provide more details of how the graph properties are defined and calculated. All experiments are conducted under Matlab version R2017a.

Characteristic path length (CPL) is the average length of all shortest paths connecting pairs of vertices of the network. Once the network becomes disconnected, CPL is not defined, and we set its value to 0.

Transitivity, or global clustering coefficient, is the ratio of triangles to triplets in the network, where a triplet is a motif consisting of one vertex and two links incident to the vertex.

Global efficiency is defined as $\frac{2}{N(N-1)} \sum_{u,v} \frac{1}{d_{uv}}$, where d_{uv} represents the shortest path between the vertices u and v [22]. Notice here if we define $d_{uv} = \infty$ for unreachable vertex pairs, global efficiency remains robust under network disconnections. The Matlab code for CPL, transitivity, and global efficiency are provided by the Brain Connectivity Toolbox [35].

Community structure: The specific Louvain implementation we used is Ref. [36] with all default parameter settings, whereas the k -means spectral clustering algorithm follows the pseudocode in Ref. [15] with a normalized graph Laplacian. The constant k that is set to be the planted ground truth in synthetic experiments. For real-world networks, we set it to be the same stable k value found by the Louvain algorithm (for all four networks Louvain was able to find stable k).

Core-periphery structure: The partition of core and periphery vertices uses the Matlab code provide by the authors of Ref. [26]. The algorithm is based on the following definition of the weighted coreness measure:

$$Q_C = \frac{1}{Z} \left[\sum_{u,v \in C_c} (W_{uv} - \bar{W}) - \sum_{u,v \in C_p} (W_{uv} - \bar{W}) \right]. \quad (\text{A1})$$

In Eq. (A1), C_c and C_p represents the bisection of the network into the core and periphery subsets, \bar{W} is the average edge weight, and $Z = \sum_{u,v} |W_{uv} - \bar{W}|$ is the normalizing constant, so that $0 \leq Q_C \leq 1$. The idea is that if there is a core-periphery structure, there are many heavy edges joining pairs of vertices of the core C_c and many light edges joining pairs of vertices of the periphery, yielding a value of Q_C appreciably larger than 0. Therefore, by maximizing Q_C^{\max} over all possible bisections of the network, we also get the C_c and C_p for our experiments.

We measure the similarity between the partitions of the sparsified graph \tilde{G} and those of G using the adjusted mutual information (AMI) [28], for all three partitioning algorithms: Louvain, k -means, and coreness. Notice here that the AMI is taken after removal of single isolated vertices, because they do not constitute meaningful communities. To overcome the randomness of the partitioning algorithms, we sample 100 different partition outcomes from the original graph, 10 from the sparsified graph, and use the maximum AMI between any pair. These numbers are picked so that the same algorithm returns consistent results (AMI very close to 1) on independent runs on the original graph. The resulting curve is thus an upper bound of any individual pairs.

DeltaCon [27] is a general graph similarity metric based on diffusion results on graphs. It aggregates the affinities between all pairs of vertices using a rooted Euclidean distance (RootED),

$$\text{RootED}(S_G, S_{\tilde{G}}) = \sqrt{\sum_{u,v} \left(\sqrt{S_G^{uv}} - \sqrt{S_{\tilde{G}}^{uv}} \right)^2},$$

where the vertex affinity matrices S_G and $S_{\tilde{G}}$ are calculated by distributed diffusion processes around each vertex. The affinity matrix involves the calculation of personalized PageRank [37], which is intuitively captured by the following matrix power series:

$$S_G \approx [I - \epsilon W]^{-1} = I + \epsilon W + \epsilon^2 W^2 + \epsilon^3 W^3 + \dots$$

Here $0 \leq \epsilon \leq 1$ represents the decay factor of diffusion over longer distances. Under the default settings, with $\epsilon = \frac{1}{\max_u D_u}$, longer range diffusion decays quickly, and DeltaCon thus puts a stronger emphasis on local structure. We use the Matlab code provided by the authors of Ref. [27], which uses a fast Belief Propagation approximation of personalized PageRank. Since DeltaCon scales directly with weights, we normalize the matrix by multiplying a constant so that DeltaCon reaches 0 on empty graphs.

For the LFR benchmark [18], we used the binary version [38]. We generated multiple synthetic networks until we get one with eight planted communities. The parameters of the LFR benchmark are listed in Table II.

TABLE II. Parameters of the LFR benchmark.

No. of vertices	1000
Average degree	20
Maximum degree	100
Minus exponent for the degree sequence	2
Minus exponent for the community size distribution	1
Minimum community size	10
Maximum community size	800

APPENDIX B: OPTIMALITY OF LINEAR THRESHOLD UNDER EXPECTED SPECTRAL SIMILARITY

In (2), we defined the *absolute spectral similarity* as a function of the input vector x . Besides the worst case σ_{\min} , we can also define an average case similarity measure,

$$\sigma_{\text{exp}} = E_{x \sim P(x)} \left[1 - \frac{x^T \Delta L x}{\lambda_N} \right] = 1 - \frac{E_{x \sim P(x)} [x^T \Delta L x]}{\lambda_N},$$

where the input vectors are drawn from a distribution $P(x)$. If we assume the entries of the vector x are independent identically distributed random variables, we have

$$\begin{aligned} \sigma_{\text{exp}} &= 1 - \frac{\sum_{u \in V, v \in V} \Delta W_{uv} E_{x_u, x_v \sim P(x)} [(x_u - x_v)^2]}{\lambda_N}, \\ &\propto 1 - \frac{C}{\lambda_N} \sum_{u \in V, v \in V} (W_{uv} - \tilde{W}_{uv}), \end{aligned} \quad (\text{B1})$$

where we have used linearity of expectation and the fact that $C = E_{x_u, x_v \sim P(x)} [(x_u - x_v)^2]$ is independent of u and v because the entries of x are assumed to be independent and identically distributed. Hence, the *expected spectral similarity* σ_{exp} simply becomes the maximum when total edge weights are kept as much as possible. This completes the proof that weight thresholding optimizes the expected value of $\sigma(x)$ with independent identically distributed x entries.

APPENDIX C: COMPUTATIONAL EFFICIENCY OF MASS

Recall that the MASS measure is defined as

$$\sigma_{\min} = \min_{\forall x, |x|=1} \left(1 - \frac{x^T \Delta L x}{\lambda_N} \right) = 1 - \frac{\lambda_N^\Delta}{\lambda_N}, \quad (\text{C1})$$

where λ_N^Δ is the largest eigenvalue of the difference Laplacian ΔL .

As a general spectral measure, MASS automatically captures important mesoscopic structures in the data. We suggest users of MASS take all types of group structure into account. However, if the application really concerns the community structures, additional validation can be done relatively easily. The Davis-Kahan theorem (see the main text) provides a theoretical connection between the largest eigenvalue of the difference Laplacian and the smallest eigenvectors associated with community structures. Empirically, we can also consider the *average rotational angle* (ARA) of the corresponding eigenvectors,

$$\eta = E_{2 \leq i \leq k} [\cos \angle(v_i, \tilde{v}_i)], \quad (\text{C2})$$

where $\cos \angle(v_i, \tilde{v}_i)$ represents the cosine of the angle between the respective eigenvectors of G and \tilde{G} corresponding to their i th smallest eigenvalue. The integer k is the number of relevant communities and can be selected based on spectral graph theory, the Louvain method, or domain knowledge if it is available.

A major advantage of the formulation in Eq. (C1) is its numerical stability and computational efficiency. Designing efficient and stable algorithms for finding the eigenvalues of a matrix is one of the most important problems in numerical analysis. The state-of-the-art iterative solvers in popular numerical packages today are inherently more stable for larger eigenvalues, including the current implementation of Matlab which we use for this work [30–32].

With only computing the largest eigenvalues of two graph Laplacians, the MASS measure is therefore among the most efficient and stable spectral properties. To demonstrate its computational efficiency, we compare the running time of computing MASS [Eqs. (3)] with those of traditional community detection algorithms, as well as DeltaCon and efficiency measures in Table III.

APPENDIX D: RESULTS ON NETWORKS WITH RANDOMIZED EDGE WEIGHT

To demonstrate the effect of weight thresholding on networks with no weight-degree correlation, we rerun the experiment on the synthetic and world trade network with randomized edge weight. In both cases, MASS and mesoscopic structures fall quickly as the edges are removed (Fig. 3).

APPENDIX E: THEORETICAL PROPERTIES OF MASS

In Ref. [27] the authors proposed several theoretical axioms for graph similarity measures. Here we adapt them to the specialized task of comparing G with its thresholded subgraph \tilde{G} . We interpret the axioms as desired properties that good subgraph similarity metrics (SSM) must satisfy. We first introduce the full list of axioms with interpretations:

- (1) *Zero identity*: The SSM returns 0 if \tilde{G} is an empty graph, 1 if \tilde{G} is identical to the original graph.
- (2) *Monotonicity*: The SSM (nonstrictly) monotonically decreases or increases as we threshold out more edges.
- (3) *Robustness*: The SSM will not drop to zero from relative large values by removing a single edge, even if the graph becomes disconnected (unless it becomes an empty graph).
- (4) *Submodularity*: Removing the same set of edges has a greater impact on the subgraph similarity measure for smaller graphs.
- (5) *Weight awareness*: When thresholding out a single edge, the greater the edge weight, the greater the impact on the subgraph similarity measure.
- (6) *Structure awareness*: The SSM suffers from a greater impact if thresholding creates disconnected components.

Next, we demonstrate that the proposed MASS measure satisfies these axioms. Recall that we define MASS as

$$\sigma_{\min}(\tilde{G}, G) = 1 - \frac{\lambda_N^\Delta}{\lambda_N},$$

TABLE III. Computational time (in seconds) of graph properties.^a

Measures	MASS	DeltaCon	Global efficiency	k-means	Louvain
World trade network	0.181	0.457	154.5	1.619	1.355
Airline network	13.60	767.2	157.0	37.63	112.9

^aAll measures are taken 21 times, in correspondence to the thresholds [0, 0.05, 0.1, 0.15, . . . , 0.95, 1]. The experiment is conducted under Matlab version R2017a. Code packages are provided by the authors of Refs. [27,35,36].

where λ_N^A is the largest eigenvalue of the Laplacian of the difference graph $\Delta G = \{V, \Delta E = E - \tilde{E}, \Delta W = W - \tilde{W}\}$.

The first axiom requires that the subgraph similarity measure returns 0 for a completely thresholded graph and 1 for the original graph.

Property 1 (Zero Identity): $\sigma_{\min}(\emptyset, G) = 0$ and $\sigma_{\min}(G, G) = 1$, where \emptyset represents an empty graph $\emptyset = \{V, \emptyset, 0_{N,N}\}$.

Proof. The Zero property is trivially satisfied as $\Delta G = G$ for $\tilde{G} = \emptyset$ and thus $\lambda_N^A = \lambda_N$. The Identity property holds because $\Delta G = \emptyset$ for $\tilde{G} = G$. According to spectral graph theory, all eigenvalues of the Laplacian matrix are 0 for an empty graph with N connected components, and we have $\lambda_N^A = 0$. ■

The second axiom requires the subgraph similarity measure to be monotonically decreasing as we threshold out more and more edges.

Property 2 (Monotonicity): $\sigma_{\min}(A, G) \leq \sigma_{\min}(B, G)$ if A is a subgraph of B , where A, B are both subgraphs of the original graph G .

Proof. By complement, we know that ΔB is a subgraph of ΔA . Assume that $\lambda_N^{\Delta B} = v^T \Delta L_B v$, where v is the corresponding unit length eigenvector. Because of the monotonicity of the Laplacian quadratic form, we have $v^T \Delta L_A v \geq \lambda_N^{\Delta B}$. We also have $\lambda_N^{\Delta A} \geq v^T \Delta L_A v \geq \lambda_N^{\Delta B}$. Therefore $\sigma_{\min}(A, G) \leq \sigma_{\min}(B, G)$. ■

Monotonicity alone does not prevent degeneracy when the subgraph becomes disconnected. The third axiom thus states that the subgraph similarity measure will not drop to zero from relative large values by removing an arbitrary edge. A general

proof for weighted graphs is difficult to formulate. Here we focus on simple graphs.

Property 3 (Robustness): Let A be a subgraph of B by removing an arbitrary edge, where A, B are both subgraphs of the original graph G , and all graphs are unweighted. If $\sigma_{\min}(B, G) > 0.5$, we have $\sigma_{\min}(A, G) > 0$.

Proof. Since $\sigma_{\min}(B, G) = 1 - \frac{\lambda_N^{\Delta B}}{\lambda_N} > 0.5$, we have $2\lambda_N^{\Delta B} < \lambda_N$. The largest eigenvalue of the Laplacian of ΔB is bounded on both sides by

$$\max_i d_i^{\Delta B} + 1 \leq \lambda_N^{\Delta B} \leq \max_{i,j \in \Delta B} (d_i^{\Delta B} + d_j^{\Delta B}),$$

where $d_i^{\Delta B}$ denotes the degree of vertex i in ΔB . Without loss of generality, we assume that ΔA is ΔB plus the edge (u, v) . If (u, v) becomes the new maximizer for the upper bound of the Laplacian of ΔA , we have

$$\lambda_N^{\Delta A} \leq d_u^{\Delta A} + d_v^{\Delta A} = d_u^{\Delta B} + d_v^{\Delta B} + 2 \leq 2\lambda_N^{\Delta B} < \lambda_N.$$

If the maximizer is $(u', v') \neq (u, v)$, we instead get

$$\lambda_N^{\Delta A} \leq d_{u'}^{\Delta A} + d_{v'}^{\Delta A} \leq d_{u'}^{\Delta B} + d_{v'}^{\Delta B} + 1 < 2\lambda_N^{\Delta B} < \lambda_N.$$

Therefore, we always have

$$\sigma_{\min}(A, G) = 1 - \frac{\lambda_N^{\Delta A}}{\lambda_N} > 0.$$

Monotonicity and smoothness concerns different thresholding on the same graph G . We can similarly derive submodularity for the same thresholding on different graphs. In other

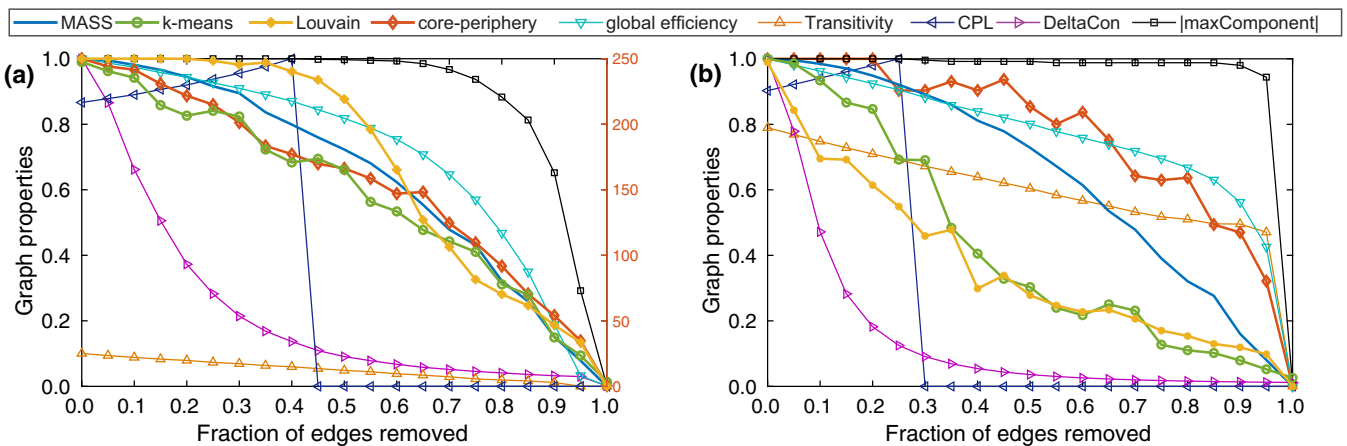


FIG. 3. Variation of MASS and other variables under weight thresholding, for a synthetic network generated by the LFR benchmark (a) and the world trade network (b) with randomized weight and 0 weight-degree correlation. The “|maxComponent|” curve represents the fraction of vertices remaining in the largest connected component.

words, removing the same set of edges has a greater impact on the similarity measure for smaller graphs.

Property 4 (Submodularity): Let A be a subgraph of B . For any common thresholding on both graphs such that $\Delta A = \Delta B \in A$, we have $\sigma_{\min}(\tilde{A}, A) \leq \sigma_{\min}(\tilde{B}, B)$.

Proof. Assume $\lambda_N^A = v^T L_A v$, where v is the corresponding unit length eigenvector. Because of the monotonicity of the Laplacian quadratic form, we have $v^T L_B v \geq \lambda_N^A$. We also have $\lambda_N^B \geq v^T L_B v \geq \lambda_N^A$. Since $\Delta A = \Delta B$, we get $\lambda_N^{\Delta A} = \lambda_N^{\Delta B}$. Therefore,

$$\sigma_{\min}(\tilde{A}, A) = 1 - \frac{\lambda_N^{\Delta A}}{\lambda_N^A} \leq 1 - \frac{\lambda_N^{\Delta B}}{\lambda_N^B} = \sigma_{\min}(\tilde{B}, B).$$

■

The fifth axiom asserts that when thresholding out a single edge, the greater the edge weight, the greater the impact on the similarity measure.

Property 5 (Weight Awareness): Let A, B be different subgraphs of the original graph G by removing a single edge, with the edge $(u, v) \notin A$ (but $\in B$), and $(u', v') \notin B$ (but $\in A$). If the edge weights follow $W_{uv} \geq W_{u'v'}$, we have $\sigma_{\min}(A, G) \leq \sigma_{\min}(B, G)$.

Proof. By complement, we know that ΔA and ΔB both consist of a single edge and $W_{uv}^{\Delta A} \geq W_{u'v'}^{\Delta B}$. The largest eigenvalue of the Laplacian of a single-edge graph is simply $\lambda_N^{\Delta A} = 2W_{uv} \geq 2W_{u'v'} = \lambda_N^{\Delta B}$, and we thus have $\sigma_{\min}(A, G) = 1 - \frac{2W_{uv}}{\lambda_N} \leq 1 - \frac{2W_{u'v'}}{\lambda_N} = \sigma_{\min}(B, G)$. ■

The last property provides an important structural constraint. If thresholding creates disconnected components, it

should have a greater impact on the similarity measure. Axiomatizing this property in its most general form is difficult, we thus focus on an intuitive special case: the unweighted Barbell graph.

Property 6 (Structure Awareness): Let graph G be an unweighted graph with two nonoverlapping stars of equal size $N/2$ connected by a single edge (u, v) , where u and v are the two center vertices. Assume that A, B are subgraphs of G , which differs only by swapping a single edge, with the edge $(u, v) \notin A$ (but $\in B$), and $(u', v') \notin B$ (but $\in A$). Then we have $\sigma_{\min}(A, G) \leq \sigma_{\min}(B, G)$.

Proof. By complement, we know that ΔA is composed of two connected stars while ΔB consists of two disconnected stars. Without loss of generality, let us assume that $|\Delta A| = |\Delta B| = k$, and that the two stars in ΔB have sizes $k_1 \leq k_2$, with $k = k_1 + k_2$. Since the largest eigenvalue of Laplacian of an unweighted k -star is exactly k , and the spectrum of the Laplacian of a disconnected graph is simply the union of those of its components, we have $\lambda_N^{\Delta B} = k_2$.

Without loss of generality, assume that the bridge edge (u, v) in ΔA has its endpoint u in the bigger component, which is of size k_2 . Vertex u therefore has a degree that is at least k_2 [It will equal $k_2 + 1$ if (u', v') is in the opposite star component.] Vertex u and its neighbors thus form a $(k_2 + 1)$ -star subgraph of ΔA . Since the largest eigenvalue of Laplacian of an unweighted $(k_2 + 1)$ -star is exactly $k_2 + 1$, by the monotonicity of the Laplacian quadratic form, we have $\lambda_N^{\Delta A} \geq k_2 + 1 \geq \lambda_N^{\Delta B}$. Therefore, $\sigma_{\min}(A, G) \leq \sigma_{\min}(B, G)$. ■

[1] A. Barrat, M. Barthelemy, R. Pastor-Satorras, and A. Vespignani, The architecture of complex weighted networks, *Proc. Natl. Acad. Sci. USA* **101**, 3747 (2004).

[2] A. Namaki, A. H. Shirazi, R. Raei, and G. R. Jafari, Network analysis of a financial market based on genuine correlation and threshold method, *Physica A (Amsterdam)* **390**, 3835 (2011).

[3] E. Bullmore and O. Sporns, Complex brain networks: Graph theoretical analysis of structural and functional systems, *Nat. Rev. Neurosci.* **10**, 186 (2009).

[4] M. Tumminello, T. Aste, T. Di Matteo, and R. N. Mantegna, A tool for filtering information in complex systems, *Proc. Natl. Acad. Sci. USA* **102**, 10421 (2005).

[5] M. Ángeles Serrano, M. Boguná, and A. Vespignani, Extracting the multiscale backbone of complex weighted networks, *Proc. Natl. Acad. Sci. USA* **106**, 6483 (2009).

[6] F. Radicchi, J. J. Ramasco, and S. Fortunato, Information filtering in complex weighted networks, *Phys. Rev. E* **83**, 046101 (2011).

[7] D. A. Spielman and N. Srivastava, Graph sparsification by effective resistances, *SIAM J. Comput.* **40**, 1913 (2011).

[8] N. Dianati, Unwinding the hairball graph: Pruning algorithms for weighted complex networks, *Phys. Rev. E* **93**, 012304 (2016).

[9] M. Coscia and F. M. H. Neffke, Network Backboning with Noisy Data, in *2017 IEEE 33rd International Conference on Data Engineering (ICDE)* (IEEE, San Diego, CA, 2017), pp. 425–436.

[10] T. Kobayashi, T. Takaguchi, and A. Barrat, The structured backbone of temporal social ties, [arXiv:1804.08828](https://arxiv.org/abs/1804.08828) [physics.soc-ph] (2018).

[11] M.-E. Lynall, D. S. Bassett, R. Kerwin, P. J. McKenna, M. Kitzbichler, U. Müller, and E. Bullmore, Functional connectivity and brain networks in schizophrenia, *J. Neurosci.* **30**, 9477 (2010).

[12] S. Allesina, A. Bodini, and C. Bondavalli, Secondary extinctions in ecological networks: Bottlenecks unveiled, *Ecol. Model.* **194**, 150 (2006).

[13] K. A. Garrison, D. Scheinost, E. S. Finn, X. Shen, and R. T. Constable, The (in)stability of functional brain network measures across thresholds, *NeuroImage* **118**, 651 (2015).

[14] S. Fortunato, Community detection in graphs, *Phys. Rep.* **486**, 75 (2010).

[15] U. Luxburg, A tutorial on spectral clustering, *Stat. Comput.* **17**, 395 (2007).

[16] S. Sarkar, S. Chawla, P. A. Robinson, and S. Fortunato, Eigenvector dynamics under perturbation of modular networks, *Phys. Rev. E* **93**, 062312 (2016).

[17] B. Iriarte, Largest Eigenvalue of the Laplacian matrix: Its eigenspace and transitive orientations, *SIAM J. Discrete Math.* **30**, 2146 (2016).

[18] A. Lancichinetti, S. Fortunato, and F. Radicchi, Benchmark graphs for testing community detection algorithms, *Phys. Rev. E* **78**, 046110 (2008).

[19] R. F. Betzel, A. Griffa, A. Avena-Koenigsberger, J. Goñi, J.-P. Thiran, P. Hagmann, and O. Sporns, Multi-scale community

- organization of the human structural connectome and its relationship with resting-state functional connectivity, *Netw. Sci.* **1**, 353 (2013).
- [20] L. De Benedictis and L. Tajoli, The world trade network, *World Econ.* **34**, 1417 (2011).
- [21] J. Patokallio, Openflights data, <http://openflights.org/> (2009).
- [22] V. Latora and M. Marchiori, Efficient Behavior of Small-World Networks, *Phys. Rev. Lett.* **87**, 198701 (2001).
- [23] M. E. J. Newman, Modularity and community structure in networks, *Proc. Natl. Acad. Sci. USA* **103**, 8577 (2006).
- [24] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, Fast unfolding of communities in large networks, *J. Stat. Mech.: Theory Exp.* (2008) P10008.
- [25] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, Community structure in time-dependent, multiscale, and multiplex networks, *Science* **328**, 876 (2010).
- [26] M. Rubinov, R. J. F. Ypma, C. Watson, and E. T. Bullmore, Wiring cost and topological participation of the mouse brain connectome, *Proc. Natl. Acad. Sci. USA* **112**, 10032 (2015).
- [27] D. Koutra, N. Shah, J. T. Vogelstein, B. Gallagher, and C. Faloutsos, DeltaCon: Principled massive-graph similarity function with attribution, *ACM Trans. Knowl. Discov. Data* **10**, 1 (2016).
- [28] N. X. Vinh, J. Epps, and J. Bailey, Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance, *J. Machine Learn. Res.* **11**, 2837 (2010).
- [29] D. A. Spielman and S.-H. Teng, Spectral sparsification of graphs, *SIAM J. Comput.* **40**, 981 (2011).
- [30] D. C. Sorensen, Implicit application of polynomial filters in ak-step Arnoldi method, *SIAM J. Matrix Anal. Appl.* **13**, 357 (1992).
- [31] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods* (SIAM, Philadelphia, PA, 1998).
- [32] G. W. Stewart, A Krylov–Schur algorithm for large eigenproblems, *SIAM J. Matrix Anal. Appl.* **23**, 601 (2001).
- [33] A. Lancichinetti and S. Fortunato, Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities, *Phys. Rev. E* **80**, 016118 (2009).
- [34] Clarivate Analytics, Web of Science data (Thomson Reuters), <http://iuni.iu.edu/resources/web-of-science> (2015).
- [35] M. Rubinov and O. Sporns, Complex network measures of brain connectivity: Uses and interpretations, *NeuroImage* **52**, 1059 (2010).
- [36] L. G. S. Jeub, M. Bazzi, I. S. Jutla, and P. J. Mucha, A generalized Louvain method for community detection implemented in MATLAB, <http://netwiki.amath.unc.edu/GenLouvain> (2011–2017).
- [37] T. H. Haveliwala, Topic-sensitive pagerank: A context-sensitive ranking algorithm for web search, *IEEE Trans. Knowledge Data Eng.* **15**, 784 (2003).
- [38] Binary version of the LFR benchmark as described in [18], https://sites.google.com/site/andrealancichinetti/files/binary_networks.tar.gz.