

RESEARCH

Non-backtracking Cycles: Length Spectrum Theory and Graph Mining Applications

Leo Torres^{1*}, Pablo Suárez-Serrato^{2,3} and Tina Eliassi-Rad^{1,4}

*Correspondence: leo@leotrs.com

¹ Network Science Institute,
Northeastern University, 360
Huntington Ave, 02115-5000,
Boston, MA, USA
Full list of author information is
available at the end of the article

Abstract

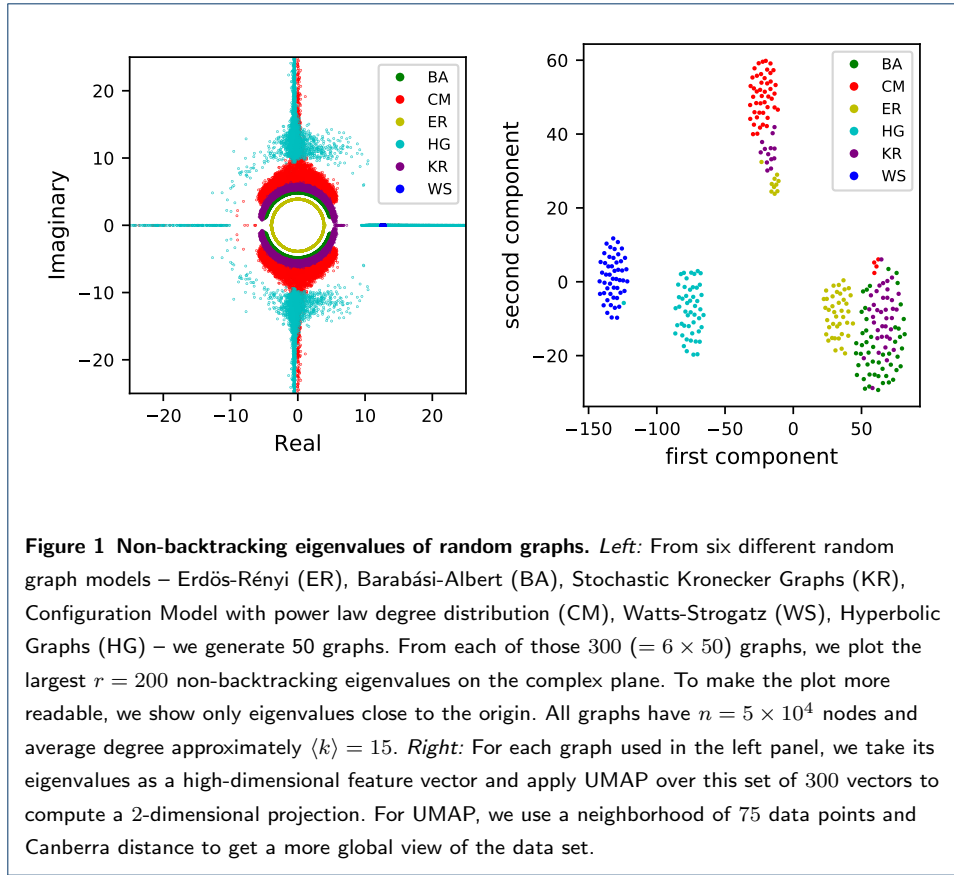
Graph distance and graph embedding are two fundamental tasks in graph mining. For graph distance, determining the structural dissimilarity between networks is an ill-defined problem, as there is no canonical way to compare two networks. Indeed, many of the existing approaches for network comparison differ in their heuristics, efficiency, interpretability, and theoretical soundness. Thus, having a notion of distance that is built on theoretically robust first principles and that is interpretable with respect to features ubiquitous in complex networks would allow for a meaningful comparison between different networks. For graph embedding, many of the popular methods are stochastic and depend on black-box models such as deep networks. Regardless of their high performance, this makes their results difficult to analyze which hinders their usefulness in the development of a coherent theory of complex networks. Here we rely on the theory of the *length spectrum* function from algebraic topology, and its relationship to the non-backtracking cycles of a graph, in order to introduce two new techniques: Non-Backtracking Spectral Distance (NBD) for measuring the distance between undirected, unweighted graphs, and Non-Backtracking Embedding Dimensions (NBED) for finding a graph embedding in low-dimensional space. Both techniques are interpretable in terms of features of complex networks such as presence of hubs, triangles, and communities. We showcase the ability of NBD to discriminate between networks in both real and synthetic data sets, as well as the potential of NBED to perform anomaly detection. By taking a topological interpretation of non-backtracking cycles, this work presents a novel application of topological data analysis to the study of complex networks.

Keywords: graph distance; graph embedding; algebraic topology; length spectrum

1 Introduction

As the network science literature continues to expand and scientists compile more examples of real life networked data sets coming from an ever growing range of domains [1, 2], there is a need to develop methods to compare complex networks, both within and across domains. Many such graph distance measures have been proposed [3, 4, 5, 6, 7, 8, 9, 10, 11, 12], though they vary in the features they use for comparison, their interpretability in terms of structural features of complex networks, computational costs, as well as in the discriminatory power of the resulting distance. This reflects the fact that complex networks represent a wide variety of systems whose structure and dynamics are difficult to encapsulate in a single distance score. For the purpose of providing a principled, interpretable, efficient and effective notion of distance, we turn to the *length spectrum* function. The length spectrum function can be defined on a broad class of metric spaces that includes Riemannian manifolds and graphs. The discriminatory power of the length spectrum is well known in other contexts: it can distinguish certain one-dimensional metric spaces up to isometry [13], and it determines the Laplacian spectrum in the case of closed hyperbolic surfaces [14]. However, it is not clear if this discriminatory power is also present in the case of complex networks. Accordingly, we present a study on the following question: is the length spectrum function useful for the comparison of complex networks?

We answer this question in the positive by introducing the Non-Backtracking Spectral Distance (NBD): a principled, interpretable, efficient, and effective measure that quantifies the distance between two undirected, unweighted networks. NBD has several desirable properties. First, NBD is based on the theory of the length spectrum and the set of non-backtracking cycles of a graph (a non-backtracking cycle is a closed walk that does not retrace any edges immediately after traversing them); these provide the theoretical background of our method. Second, NBD is interpretable in terms of features of complex networks such as existence of hubs and triangles. This helps in the interpretation and visualization of distance scores yielded by NBD. Third, NBD is more computationally efficient than other comparable methods. Indeed, NBD depends only on the computation of a few of the largest eigenvalues of the non-backtracking matrix of a graph, which requires only slightly more computational time than the spectral decomposition of the adjacency matrix.



Fourth, we have extensive empirical evidence that demonstrates the effectiveness of NBD at distinguishing real and synthetic networks.

NBD is intimately related to the eigenvalues of the non-backtracking matrix, which serve as a proxy for the length spectrum of a graph (see Section 4.1). Motivated by the usefulness of the eigenvalues of the non-backtracking matrix, we then turn our focus to its eigenvectors. We discuss the potential of using the eigenvectors as a graph embedding technique, which we refer to as Non-Backtracking Embedding Dimensions (or NBED for short). This technique computes a low dimensional embedding of each directed edge in a graph. We show that the peculiar visualizations yielded by NBED are particularly apt at providing a rich visualization of the embedded network, and use the patterns found therein to perform anomaly detection on the Enron email corpus [15].

The contributions of this paper are as follows. By highlighting the topological interpretation of the non-backtracking cycles of a graph, we propose Non-Backtracking Spectral Distance (NBD) as a novel graph distance method, and

discuss the potential of Non-Backtracking Embedding Dimensions (NBED) as a promising embedding technique. Along the way, we present an efficient algorithm to compute the non-backtracking matrix, and discuss the data visualization capabilities of its complex eigenvalues (see Figure 1) and eigenvectors (Figure 12). With this work we also release publicly available source code that implements NBD, NBED and the related matrix computations [16].

The rest of this paper is structured as follows. Section 2 provides necessary background information on the length spectrum, non-backtracking cycles, and the non-backtracking matrix. In Section 3 we discuss previous investigations related to ours. Section 4 explains the connection between these objects, as well as a discussion of the properties of the non-backtracking matrix that make it relevant for the study of complex networks. It also presents our efficient algorithm to compute it. Section 5 presents our distance method NBD and provides experimental evidence of its performance by comparing it to other distance techniques. In Section 6 we discuss our embedding method NBED based on the eigenvectors of the non-backtracking matrix and provide extensive visual analysis of the resulting edge embeddings, as well as mention its shortcomings and necessary future lines of study. We conclude in Section 8 with a discussion of limitations and future work.

2 Background

Here we introduce two theoretical constructions: the length spectrum of a metric space and the set of non-backtracking cycles of a graph. Our analysis pivots on the fact that the latter is intimately related to a particular subset of the (domain of the) former. For a list of all symbols and abbreviations used in this work, see Table 1.


2.1 Length Spectrum

Consider a metric space X and a point $p \in X$. A closed curve that goes through p is called a *loop*, and p is called its *basepoint*. Two loops are *homotopic* to one another relative to a given basepoint when there exists a continuous transformation from one to the other that fixes the basepoint. The *fundamental group* of X with basepoint p is denoted by $\pi_1(X, p)$ and is defined as the first homotopy group of X , i.e., the set of all loops in X that go through p , under the equivalence relation of based homotopy. An element of the fundamental group is called a *homotopy class*. If a homotopy class contains the loop c , we denote it as $[c]$. The group operation is

Symbol	Definition
$\pi_1(X, p)$	the fundamental group of X with basepoint p
$[c]$	homotopy class of loop c
$\mathcal{L}, \mathcal{L}'$	length spectrum function, relaxed length spectrum
$Conv(X)$	if X is a graph, $Conv(X)$ is its 2-core
$G = (V, E)$	an undirected graph with node set V and edge set E
n, m	number of nodes and number of edges of a graph G
e, e^{-1}	a directed edge $e = u \rightarrow v$ and its inverse $e^{-1} = v \rightarrow u$
NBC	non-backtracking cycle, in which no edge is followed by its inverse
B	$2m \times 2m$ non-backtracking matrix of a graph
B'	$2n \times 2n$ matrix whose eigenvalues are the same as those of B , save for ± 1
$\lambda_k = a_k + ib_k$	k -th largest eigenvalue, in magnitude, of B
$Re(\lambda), Im(\lambda)$	real and imaginary parts of the complex number λ
P, Q	$n \times 2m$ directed incidence matrices of a graph
p_k	the fraction of a graphs' nodes with degree k
γ	the degree exponent in the case when $p_k \sim k^{-\gamma}$
$\langle k^i \rangle$	i -th moment of the degree distribution p_k
$\deg(u)$	degree of node u
$nnz(A)$	number of non-zero elements of binary matrix A
$NBD(G, H)$	non-backtracking distance between graphs G and H
$d(\{\lambda_k\}, \{\mu_k\})$	distance between eigenvalues λ_k and eigenvalues μ_k
r	number of eigenvalues computed from one graph
$\rho = \sqrt{\lambda_1}$	magnitude threshold for eigenvalue computation
r_0	number of eigenvalues whose magnitude is greater than ρ
σ	spread parameter of the RBF kernel

Table 1 Notation used in this work.

the product of two homotopy classes, defined as the concatenation of two of their representative loops and written multiplicatively as $[c_1][c_2]$ whenever $[c_1]$ and $[c_2]$ are elements of $\pi_1(X, p)$. Since c_1 starts at the same node at which c_2 ends, they can be traversed one after the other, which uniquely defines another based loop. A *conjugacy class* of the fundamental group is a set C of homotopy classes such that for any two classes $[c_1], [c_2] \in C$ there exists a third class $[g]$ that satisfies $[c_1] = [g^{-1}][c_2][g]$. Here, the inverse of the loop g is defined as the same loop traversed in reverse. Closed curves without a distinguished basepoint are called *free loops*, and they correspond to conjugacy classes of $\pi_1(X, p)$. For an example of conjugacy classes of the fundamental group of a graph G , see Figure 2. A well-known fact of homotopy theory is that if X is path-connected (i.e., if any two points can be connected through a continuous curve in X) then $\pi_1(X, p)$ is unique up to isomorphism, regardless of basepoint p . In the present work we only consider connected graphs, hence, we simply write $\pi_1(X)$ when there is no ambiguity. For more on homotopy, we refer the interested reader to [17, 18].



The length spectrum \mathcal{L} is a function on $\pi_1(X)$ which assigns to each homotopy class of loops the infimum length among all of the representatives in its conjugacy class.^[1] Note, importantly, that the definition of length of a homotopy class considers the length of those loops not only in the homotopy class itself, but in all other conjugate classes. In the case of compact geodesic spaces, such as finite metric graphs, this infimum is always achieved. For a finite graph where each edge has length one, the value of \mathcal{L} on a homotopy class then equals the number of edges contained in the length-minimizing loop. That is, for a graph $G = (V, E)$, $v \in V$, if $[c] \in \pi_1(G, v)$ and c achieves the minimum length k in all classes conjugate to $[c]$, we define $\mathcal{L}([c]) = k$.

Our interest in the length spectrum is supported by the two following facts. First, graphs are *aspherical*. More precisely, the underlying topological space of any graph is aspherical, i.e., all of its homotopy groups of dimension greater than one are trivial.^[2] Therefore, it is logical to study the only non-trivial homotopy group, the fundamental group $\pi_1(G)$. Second, Constantine and Lafont [13] showed that the

^[1]The definition presented here is also known as *marked* length spectrum. An alternative definition of the (unmarked) length spectrum does not depend on π_1 ; see for example [14].

^[2]This follows from G being homotopy equivalent to a bouquet of k circles, where k is the rank of the fundamental group of G . The universal covering of a bouquet of circles is contractible, which is equivalent to the space being aspherical. (See [18].)

length spectrum of a graph determines a certain subset of it up to isomorphism. Thus, we aim to determine when two graphs are close to each other by comparing their length spectra relying on the main theorem of [13]. For completeness, we briefly mention it here; it is known as *marked length spectrum rigidity*.

Theorem 1 (Constantine-Lafont [13]) *For a metric space X define $\text{Conv}(X)$ as the minimal set to which X retracts by deformation. Let X_1, X_2 be a pair of compact, non-contractible, geodesic spaces of topological dimension one. If the marked length spectra of X_1 and X_2 are the same, then $\text{Conv}(X_1)$ is isometric to $\text{Conv}(X_2)$.*

When G_1, G_2 are graphs, $\text{Conv}(G_i), i = 1, 2$, corresponds to the subgraph resulting from iteratively removing nodes of degree 1 from G_i , which is precisely the 2-core of G_i [19]. Equivalently, the 2-core of G is the maximal subgraph in which all vertices have degree at least 2. Thus, Theorem 1 states that *when two graphs have the same length spectrum, their 2-cores are isomorphic*.

Given these results, it is natural to use the length spectrum as the basis of a measure of graph distance. Concretely, given two graphs, we aim to efficiently quantify how far their 2-cores are from being isomorphic by measuring the distance between their length spectra. In Section 4.1, we explain our approach at implementing a computationally feasible solution for this problem.

2.2 Non-Backtracking Cycles

Consider an undirected, unweighted graph $G = (V, E)$ and suppose $|E| = m$. For an undirected edge $e = \{u, v\} \in E$, we write $u \rightarrow v$ and $v \rightarrow u$ for its two possible orientations, and define \vec{E} as the set of all $2m$ directed edges. Given a directed edge $e = u \rightarrow v \in \vec{E}$, define e^{-1} as the same edge traversed in the inverse order, $e^{-1} = v \rightarrow u$. A *path* in G is a sequence of directed edges $e_1 e_2 \dots e_k$ such that if $e_i = u_i \rightarrow v_i$ then $v_i = u_{i+1}$ for $i = 1, \dots, k-1$. Here, k is called *length* of the path. A path is *closed* if $v_k = u_1$. A path is called *non-backtracking* when an edge is never followed by its inverse, $e_{i+1} \neq e_i^{-1}$ for all i . A closed path is also called a

cycle. A cycle is a *non-backtracking cycle* (or NBC for short) when it is a closed non-backtracking path and, in addition, $e_k \neq e_1^{-1}$.^[3]

The *non-backtracking matrix* B is a $2m \times 2m$ matrix whose rows and columns are indexed by directed edges. For two edges $u \rightarrow v$ and $k \rightarrow l$, B is given by

$$B_{k \rightarrow l, u \rightarrow v} = \delta_{vk}(1 - \delta_{ul}), \quad (1)$$

where δ_{uv} is the Kronecker delta. Thus, there is a 1 in the entry indexed by row $k \rightarrow l$ and column $u \rightarrow v$ when $u \neq l$ and $v = k$, and a 0 otherwise. Intuitively, one can interpret the B matrix as the (unnormalized) transition matrix of a random walker that does not perform backtracks: the entry at row $k \rightarrow l$ and column $u \rightarrow v$ is positive if and only if a walker can move from node u to node v (which equals node k) and then to l , without going back to u .

Importantly, NBCs are topologically relevant because backtracking edges are homotopically trivial, that is, the length-2 cycle $u \rightarrow v, v \rightarrow u$ can always be contracted to a point [20]. Indeed, in Figure 2 we removed edges are precisely those that form backtracks.

Observe that the matrix B tracks each pair of incident edges that do not comprise a backtrack. As a consequence we have the following result.

Lemma 1 $B_{k \rightarrow l, i \rightarrow j}^p$ equals the number of non-backtracking paths of length $p + 1$ that start with edge $i \rightarrow j$ and end with edge $k \rightarrow l$. Moreover, $\text{tr}(B^p)$ is proportional to the number of closed non-backtracking paths (i.e., NBCs) of length p in G .

Proof The first fact is reminiscent of the well-known fact that A_{uv}^p gives the number of paths that start at node u and end at node v , for any $u, v \in G$, where A is the adjacency matrix of G . The proof is the same as the proof of this latter fact, so we omit the details for brevity. The only difference is that B^p gives the number of non-backtracking paths of length $p + 1$, while A^p gives the number of paths of length p . To see why, notice that if the entry $B_{k \rightarrow l, i \rightarrow j}$ is positive, then the length-2 path $i \rightarrow j \rightarrow l$ must exist and it cannot contain any backtracks. Accordingly, if $B_{k \rightarrow l, i \rightarrow j}^2$

^[3]Other authors reserve the use of the term *cycle* for special cases of closed paths such as the set of *simple cycles*, which are closed paths that do not intersect each other. In this work we use *cycle* and *closed path* interchangeably.

is positive then the length-3 path $i \rightarrow j \rightarrow k \rightarrow l$ must be non-backtracking. It remains to show that $\text{tr}(B^p)$ is proportional to the number of NBCs of length k , and not the number of NBCs of length $p + 1$ as one may think from the previous paragraph. Suppose $a \rightarrow b \rightarrow c \rightarrow a$ is a valid closed non-backtracking path, that is, the nodes a, b, c form a triangle. In this case, $B_{c \rightarrow a, a \rightarrow b}^2$ must be positive, as must be $B_{a \rightarrow b, a \rightarrow b}^3$. The latter represents the existence of some node c such that the path $a \rightarrow b \rightarrow c \rightarrow a \rightarrow b$ exists and is non-backtracking, or, equivalently, that $a \rightarrow b \rightarrow c \rightarrow a$ is a non-backtracking cycle. Note that $B_{b \rightarrow c, b \rightarrow c}^3$ and $B_{c \rightarrow a, c \rightarrow a}^3$ are also positive and all signal the existence of the same triangle a, b, c . Thus the sum of all diagonal elements $\text{tr}(B^3)$ counts every triangle exactly six times, since there are three edges each with two possible orientations. A proof by induction shows that the above argument holds for any p . \square

The second claim of Lemma 1 will be fundamental in our later exposition. We prefer to count NBCs of length p using $\text{tr}(B^p)$, as opposed to the entries of B^{p-1} , because of the well-known fact that $\text{tr}(B^p) = \sum_i \lambda_i^p$, for any square matrix B , and where each λ_i is an eigenvalue of B [21]. Also of importance is the fact that B is not symmetric, and hence its eigenvalues are, in general, complex numbers. Since B is a real matrix, all of its eigenvalues come in conjugated pairs: if $a + ib$ is an eigenvalue then so is $a - ib$, where i is the imaginary unit. This implies that we can write the above equation as $\text{tr}(B^p) = \sum_i \lambda_i^p = \sum_i \text{Re}(\lambda_i^p)$, where $\text{Re}(z)$ is the real part of a complex number, since the imaginary part of conjugated eigenvalues cancel out. In the rest of this work we refer to B 's eigenvalues as the *non-backtracking eigenvalues*.

If one is interested in B 's eigenvalues rather than in the matrix itself, one may use the so-called Ihara determinant formula [22, 23], which states that the eigenvalues of B different than ± 1 are also the eigenvalues of the $2n \times 2n$ block matrix

$$B' = \begin{pmatrix} A & I - D \\ I & 0 \end{pmatrix} \quad (2)$$

where A is the adjacency matrix, D is the diagonal matrix with the node degrees, and I is the $n \times n$ identity matrix.

3 Related work

Hashimoto [22] discussed the non-backtracking cycles of a graph (and the associated non-backtracking matrix) in relation to the theory of Zeta functions in graphs. Terras [20] explained the relationship between non-backtracking cycles and the free homotopy classes of a graph. More recently, the non-backtracking matrix has been applied to diverse applications such as node centrality [24, 25] and community detection [26, 27, 28], and to the data mining tasks of clustering [29] and embedding [30]. In particular, the application to community detection is of special interest because it was shown that the non-backtracking matrix performs better at spectral clustering than the Laplacian matrix in some cases [26]. Hence, there is recent interest in describing the eigenvalue distribution of the non-backtracking matrix in models such as the Erdős-Rényi random graph and the stochastic block model [27, 31]. Our work differs from other applied treatments of the non-backtracking matrix in that we arrive at its eigenvalues from first principles, as a relaxed version of the length spectrum. Concretely, we use the eigenvalues to compare graphs because the spectral moments of the non-backtracking matrix describe certain aspects of the length spectrum (see Section 4). The spectral moments of other matrices (e.g., adjacency and Laplacian matrices) also describe structural features of networks [32, 33].

Many distance methods have been proposed recently [3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. This proliferation is due to the fact that there is no definitive way of comparing two graphs, especially complex networks since they present a vast diversity of structural features. The methods that are most related to ours fall in two categories: combinatorial enumeration or estimation of different subgraphs (such as motifs or shortest paths), or those that depend on spectral properties of some matrix representation of the network, such as the method **Lap** from Section 5.3.1.

The topic of embedding has also seen a sharp increase of activity in recent years, motivated by the ubiquity of machine learning and data mining applications to structured data sets; see [34, 35] for recent reviews. Our method NBED differs from most others in that it yields an edge embedding that differentiates the two possible orientations of an edge. Our focus is in providing an interpretable visual analysis of the resulting embedding, and on its application to anomaly detection.

4 Operationalizing the length spectrum

We want to quantify the dissimilarity between two graphs by measuring the dissimilarity between their length spectra. To do this, there are two main obstacles to overcome: (i) computing the exact length spectrum of a given graph is not computationally feasible as it depends on enumerating the length of every NBC, and (ii) it is not immediately clear how to compare two length spectra functions that come from two distinct graphs, because they are defined on disjoint domains (the fundamental groups of the graphs)^[4]. In order to overcome these obstacles, we propose the use of what we call *the relaxed length spectrum*, denoted by \mathcal{L}' , and whose construction comes in the form of a two-step aggregation of the values of \mathcal{L} ; see Figure 3 for an overview of this procedure.

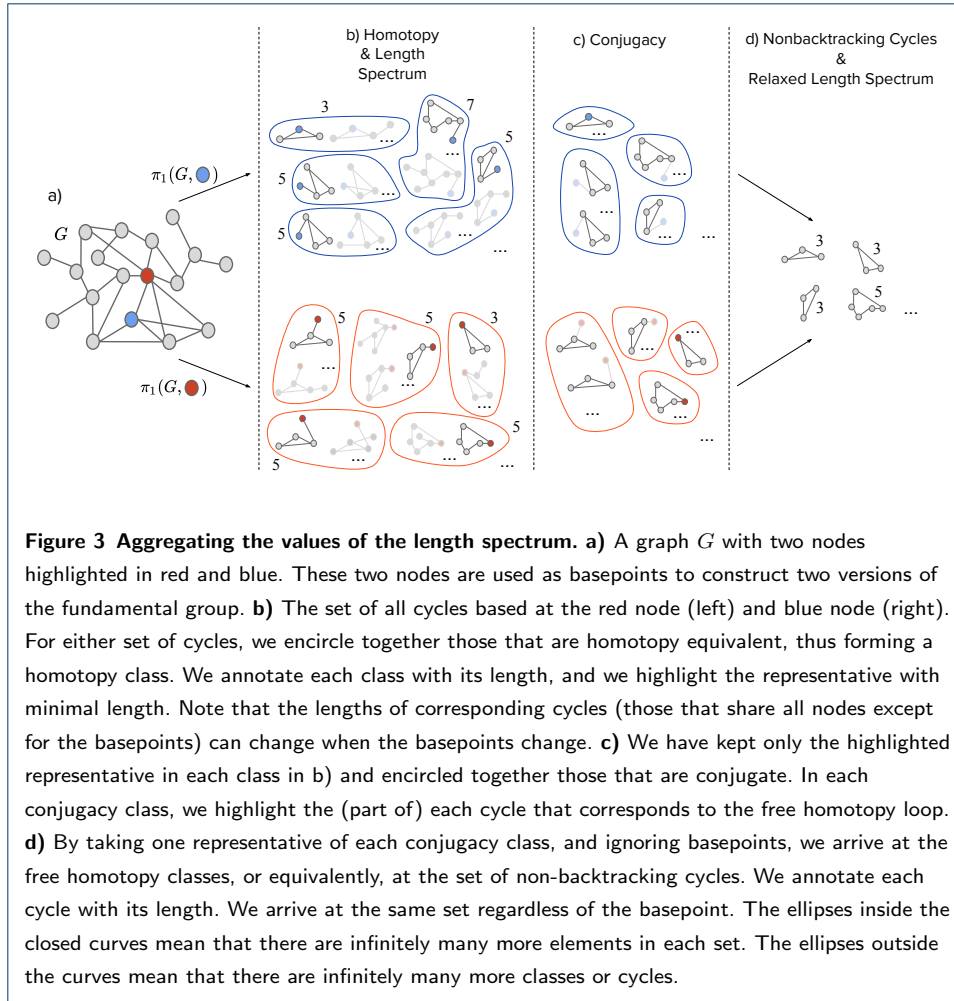
4.1 Relaxed Length Spectrum

The first step of this procedure is to focus on the image of the length spectrum rather than the domain (i.e., focus on the collection of lengths of cycles). The second step is to aggregate these values by considering the size of the level sets of either length spectrum.

Concretely, when comparing two graphs G and H , instead of comparing \mathcal{L}_G and \mathcal{L}_H directly, we compare the number of cycles in G of length 3 with the number of cycles in H of the same length, as well as the number of cycles of length 4, of length 5, etc, thereby essentially comparing the histogram of values that each \mathcal{L} takes. Theoretically, focusing on this histogram provides a common ground to compare the two functions. In practice, this aggregation allows us to reduce the amount of memory needed to store the length spectra because we no longer keep track of the exact composition of each of the infinitely many (free) homotopy classes. Instead, we only keep track of the frequency of their lengths. According to this aggregation, we define the *relaxed* version of the length spectrum as the set of points $\mathcal{L}' = \{(k, n(k)) : k = 1, 2, \dots\}$, where $n(k)$ is the number of conjugacy classes of π_1 (i.e., free homotopy classes) of length k .

The major downside of removing focus from the underlying group structure and shifting it towards (the histogram of values in) the image is that we lose information

^[4]In [13], the authors need an isomorphism between the fundamental group of the spaces that are being compared, which is also computationally prohibitive.



about the combinatorial composition of each cycle. Indeed, $\pi_1(G)$ holds information about the number of cycles of a certain length k in G ; this information is also stored in \mathcal{L}' . However, the group structure of $\pi_1(G)$ also allows us to know how many of those cycles of length k are formed by the concatenation of two (three, four, etc.) cycles of different lengths. This information is lost when considering only the sizes of level sets of the image, i.e., when considering \mathcal{L}' .

The next step makes use of the non-backtracking cycles (NBCs). We rely on NBCs because the set of conjugacy classes of $\pi_1(G)$ is in bijection with the set of NBCs of G (see e.g., [20, 22]). In other words, to compute the set \mathcal{L}' we need only account for the lengths of all NBCs. Indeed, consider the non-backtracking matrix B of G and recall that $\text{tr}(B^k)$ equals the number of NBCs of length k in the graph. This gives us precisely the set $\mathcal{L}' = \{(k, \text{tr}(B^k))\}_{k=1}^{\infty}$. Recall that $\text{tr}(B^k) = \sum_i \lambda_i^k$, where each λ_i is a non-backtracking eigenvalue. Therefore, the eigenvalues of B contain all

the information necessary to compute and compare \mathcal{L}' . In this way, we can study the (eigenvalue) spectrum of B , as a proxy for the (length) spectrum of π_1 .

We have deviated from the original definition of the length spectrum in important ways. Fortunately, our experiments indicate that \mathcal{L}' contains enough discriminatory information to distinguish between real and synthetic graphs effectively; see Sections 5.3.2 and 5.3.5. We discuss this limitation further in our concluding remarks in Section 8. Beyond experimental results, one may ask if there are theoretical guarantees that the relaxed version of the length spectrum will keep some of the discriminatory power of the original. Indeed, even though our inspiration for this work is partially the main rigidity result of [13], we can still trust the eigenvalue spectrum of B to be useful when comparing graphs. On the one hand, the spectrum of B has been found to yield fewer isospectral graph pairs (i.e., non-isomorphic graphs with the same eigenvalues) than the adjacency and Laplacian matrices in the case of small graphs [36]. On the other hand, B is tightly related to the theory of graph zeta functions [22], in particular the Ihara Zeta function, which is known to determine several graph properties such as girth, number of spanning trees, whether the graph is bipartite, regular, or a forest, etc [37]. Thus, both as a relaxed version of the original length spectrum, but also as an object of interest in itself, we find the eigenvalue spectrum of the non-backtracking matrix B to be an effective means of determining the dissimilarity between two graphs. For the rest of this work we focus on the eigenvalues of B and on using them to compare complex networks.

4.2 Computing B

As mentioned previously, the eigenvalues of B different than ± 1 are also the eigenvalues of the $2n \times 2n$ block matrix B' defined in Equation 2. Computing its eigenvalues can then be done with standard computational techniques, and it will not incur a cost greater than computing the eigenvalues of the adjacency or Laplacian matrices, as B' is precisely four times the size of A and it has only $2n$ more non-zero entries (see Equation 2). This is what we will do in Section 5 in order to compare two relaxed length spectra. However, in Section 6 we will need the eigenvectors of B , which cannot be computed from B' . For this purpose, we now present an efficient algorithm for computing B . Computing its eigenvectors can then be done with standard techniques.

Given a graph with n nodes and m undirected edges in edge list format, define the $n \times 2m$ incidence matrices $P_{x,u \rightarrow v} = \delta_{xu}$ and $Q_{x,u \rightarrow v} = \delta_{xv}$, and write $C = P^T Q$. Observe that $C_{k \rightarrow l, u \rightarrow v} = \delta_{vk}$ and therefore,

$$B_{k \rightarrow l, u \rightarrow v} = C_{k \rightarrow l, u \rightarrow v} (1 - C_{u \rightarrow v, k \rightarrow l}) \quad (3)$$

Note that an entry of B may be positive only when the corresponding entry of C is positive. Therefore, we can compute B in a single iteration over the nonzero entries of C .

Now, C has a positive entry for each pair of incident edges in the graph, from which we have the following result.

Lemma 2 *Set $\text{nnz}(C)$ to be the number of non-zero entries in C , and $\langle k^2 \rangle$ the second moment of the degree distribution. Then, $\text{nnz}(C) = n\langle k^2 \rangle$.*

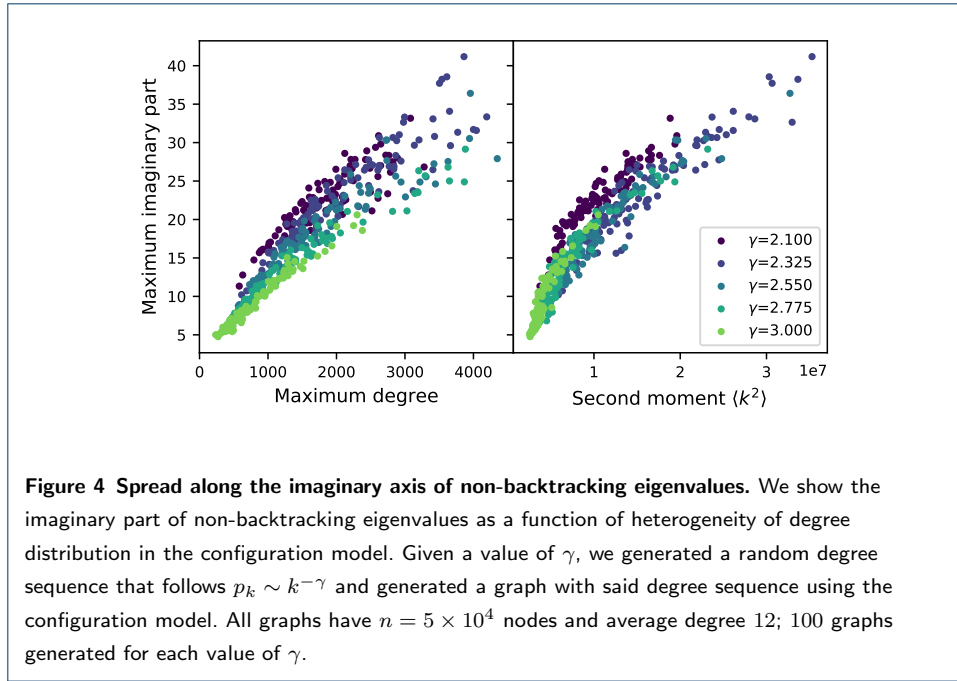
Proof This follows by direct computation: $\sum_{k \rightarrow l, i \rightarrow j} \delta_{kj} = \sum_k (\sum_l a_{lk}) (\sum_i a_{ik}) = n\langle k^2 \rangle$ \square

Computing P and Q takes $O(m)$ time. Since computing the product of sparse matrices takes time proportional to the number of positive entries in the result, computing C takes $O(n\langle k^2 \rangle)$. Thus we can compute B in time $O(m + n\langle k^2 \rangle)$. In the case of a power-law degree distribution with exponent $2 \leq \gamma \leq 3$, the runtime of our algorithm falls between $O(m + n)$ and $O(m + n^2)$. Note that if a graph is given in adjacency list format, one can build B directly from the adjacency list in time $\Theta(n\langle k^2 \rangle - n\langle k \rangle)$ by generating a sparse matrix with the appropriate entries set to 1 in a single iteration over the adjacency list.

4.3 Spectral properties of B

In our experiments we have found that the eigenvalues of B behave nicely with respect to certain fundamental properties of complex networks such as degree distribution and triangles. If the theory of the length spectrum justifies the use of the non-backtracking eigenvalues to compare graphs in general, the properties we present next justify their use for complex networks in particular.

Lemma 3 *We have $\text{nnz}(B) = n(\langle k^2 \rangle - \langle k \rangle)$.*



Proof We directly compute $\sum_{k \rightarrow l, i \rightarrow j} \delta_{kj}(1 - \delta_{il}) = \sum_k (\deg(k) - 1) \deg(k) = n(\langle k^2 \rangle - \langle k \rangle)$. \square

Contrast this to $nnz(A) = n\langle k \rangle$ for the adjacency matrix A . As a consequence of this, we have found experimentally that the larger $\langle k^2 \rangle$, the larger the range of the non-backtracking eigenvalues along the imaginary axis; see Figure 4.

Next, we turn to B 's eigenvalues and their relation to the number of triangles. Write $\lambda_k = a_k + ib_k \in \mathbb{C}$ for the non-backtracking eigenvalues, where i is the imaginary unit and $k = 1, \dots, 2m$. The number of triangles in a network is proportional to $tr(B^3) = \sum_k Re(\lambda_k^3)$, which, by a direct application of the binomial theorem, equals

$$tr(B^3) = \sum_{k=1}^{2m} a_k(a_k^2 - 3b_k^2). \quad (4)$$

In general, the eigenvalues of B with small absolute value tend to fall on a circle in the complex plane [26, 38, 39, 27]. However, if $\sum_k a_k^2$ is large and $\sum_k b_k^2$ is small (implying a large number of triangles), the eigenvalues cannot all fall too close to a circle, since they will need to have different absolute values. Hence, the more triangles in the graph, the less marked the circular shape of the eigenvalues.

Another way of saying the same thing is that the more triangles exist in the graph, the eigenvalues have larger and positive real parts.

4.4 Eigenvalues: visualization

We close this section with a visualization of the non-backtracking eigenvalues. We choose six different random graph models: Erdős-Rényi (ER) [40, 41], Barabási-Albert (BA) [42], Stochastic Kronecker Graphs (KR) [43, 44], Configuration Model with power law degree distribution (CM; $p_k \sim k^{-\gamma}$ with $\gamma = 2.3$) [45], Watts-Strogatz (WS) [46], Hyperbolic Graphs (HG; $\gamma = 2.3$) [47, 48]. From each model, we generate 50 graphs with $n = 10,000$ nodes and approximately $\langle k \rangle = 15$; see Table 2 for more information about these graphs. For each graph, we plot the largest $r = 200$ eigenvalues on the complex plane in Figure 1.

Each model generates eigenvalue distributions presenting different geometric patterns. As expected from the previous section, HG has a less marked circular shape around the origin since it is the model that generates graphs with the most triangles. Furthermore, HG also has the largest maximum degree, and therefore its eigenvalues have a greater spread along the imaginary axis. We also show in Figure 1 a low-dimensional projection of the non-backtracking eigenvalues using UMAP [49]. Notice how, even in two dimensions, the non-backtracking eigenvalues are clustered according to the random graph model they come from. This provides experimental evidence that the non-backtracking eigenvalues shall be useful when comparing complex networks. For a more in-depth explanation of the parameters used for UMAP, see Appendix 1.

5 NBD: Non-backtracking distance

Based on the previous discussion, we propose a method to compute the distance between two complex networks based on the non-backtracking eigenvalues. In this Section, we use d to refer to an arbitrary metric defined on subsets of \mathbb{R}^2 . That is, the distance between two graphs G, H is given by $d(\{\lambda_k\}, \{\mu_k\})$, where λ_k, μ_k are the eigenvalues of G, H , respectively, which we identify with points in \mathbb{R}^2 by using their real and imaginary parts as coordinates. With these notations we are finally ready to propose the Non-Backtracking Spectral Distance, or NBD for short.

Table 2 Data sets used in this work. n : number of nodes in the largest connected component; $\langle k \rangle$: average degree; \bar{c} : average clustering coefficient. *Left*: random graph data sets. We generate 50 graphs of each data set, for a total of 900 random graphs. *Right*: real networks, grouped in four categories.

data set	n	$\langle k \rangle$	\bar{c}
BA-1k	1k	13.90	0.05
CM-1k	1k	10.20	0.17
ER-1k	1k	14.99	0.15
HG-1k	997 ± 2	14.72	0.82
KR-1k	1021 ± 1	14.54	0.03
WS-1k	1k	14.00	0.67
BA-5k	5k	13.98	0.02
CM-5k	5k	11.40	0.13
ER-5k	5k	15.00	0.00
HG-5k	4982 ± 8	14.67	0.82
KR-5k	4089 ± 3	18.27	0.01
WS-5k	5k	14.00	0.67
BA-10k	10k	13.99	0.01
CM-10k	10k	11.97	0.12
ER-10k	10k	14.99	0.00
HG-10k	9958 ± 15	14.96	0.83
KR-10k	8170 ± 11	18.32	0.01
WS-10k	10k	14.00	0.67

data set	network	n	$\langle k \rangle$	\bar{c}
Online social (social)	facebook	4k	43.69	0.61
	sdot8	77k	14.13	0.06
	sdot9	82k	14.18	0.06
	epinions	76k	10.69	0.14
	twitter	81k	33.02	0.57
	wiki	7k	28.32	0.14
Coauthorship (CA)	AstroPh	19k	21.11	0.63
	CondMat	23k	8.08	0.63
	GrQc	5k	5.53	0.53
	HepPh	12k	19.74	0.61
	HepTh	10k	5.26	0.47
Peer-to-peer file-sharing (P2P)	P2P04	11k	7.35	0.01
	P2P05	9k	7.20	0.01
	P2P06	9k	7.23	0.01
	P2P08	6k	6.59	0.01
	P2P09	8k	6.41	0.01
	P2P24	27k	4.93	0.01
	P2P25	23k	4.82	0.01
	P2P30	36k	4.82	0.01
	P2P31	63k	4.73	0.01
Autonomous systems (AS)	AS331	11k	4.12	0.30
	AS407	11k	4.10	0.29
	AS414	11k	4.16	0.30
	AS421	11k	4.19	0.30
	AS428	11k	4.13	0.29
	AS505	11k	4.13	0.29
	AS512	11k	4.12	0.29
	AS519	11k	4.11	0.29
	AS526	11k	4.19	0.30

Definition 1 Consider two graphs G, H , and let $\{\lambda_k\}_{k=1}^{k=r}, \{\mu_k\}_{k=1}^{k=r}$, be the r non-backtracking eigenvalues of largest magnitude of G and H , respectively. We define the NBD between G and H as $NBD(G, H) = d(\{\lambda_k\}, \{\mu_k\})$.

Note that in Definition 1 we leave open two important degrees of freedom: the choice of d and the choice of r . We will discuss the optimal choices for these parameters in the next Section. Regardless of these choices, however, we have the following results.

Proposition 1 If d is a metric over subsets of \mathbb{R}^2 , then NBD is a pseudo-metric over the set of graphs.

Proof NBD inherits several desirable properties from d : non-negativity, symmetry, and, importantly, the triangle inequality. However, the distance between two distinct graphs may be zero when they share all of their r largest eigenvalues. Thus, NBD is not a metric over the space of graphs but a pseudo-metric. \square

5.1 Computing NBD

Algorithm 1 presents our method to compute the NBD between two graphs. It makes use of the following fact, which simplifies the computation. It is known (see e.g. [36]) that the multiplicity of 0 as an eigenvalue of B equals the number of edges outside of the 2-core of the graph. For example, a tree, whose 2-core is empty, has all its eigenvalues equal to 0. On the one hand, we could use this valuable information as part of our method to compare two graphs. On the other hand, the existence of zero eigenvalues does not change the value of $\text{tr}(B^k)$, $k \geq 0$, and thus leaves the relaxed length spectrum $\mathcal{L}' = \{(k, \text{tr}(B^k))\}_k$ intact. Moreover, iteratively removing the nodes of degree one (a procedure called “shaving”) reduces the size of B (or the sparsity of B' ; see Equation 2), which makes the computation of non-zero eigenvalues faster.

Algorithm 1 Non-Backtracking Spectral Distance

Input: Two graphs G, H , distance d , number of eigenvalues r

Output: real number d , distance between G, H

- 1: $\tilde{G}, \tilde{H} \leftarrow \text{shave}(G), \text{shave}(H)$
 - 2: $\{\lambda_k\}_k, \{\mu_k\}_k \leftarrow$ the r largest eigenvalues of B' corresponding to \tilde{G}, \tilde{H}
 - 3: **return** $d(\{\lambda_i\}, \{\mu_i\})$
-

Given two graphs G, H , we first compute their 2-cores by shaving them, that is, iteratively removing the nodes of degree 1 until there are none; call the new graphs \tilde{G} and \tilde{H} . Then for each graph, we compute the B' matrix using Equation 2 to then compute the largest r eigenvalues. Lastly, we compute the distance d between the resulting sets of eigenvalues. We proceed to analyze the runtime complexity of each line of Algorithm 1. Line 1: the shaving step has a worst-case scenario of $O(n^2)$, where n is the number of nodes. Indeed, the shaving algorithm must iterate two steps until completion: first, identify all nodes of degree 1, and second, remove those nodes from the graph. For the first step, querying the degree of all nodes takes $O(n)$ time, while the deletion step takes an amount of time that is linear in the amount of nodes that are removed at that step. In the worst-case scenario,

consider a path graph. At each iteration, the degrees of all nodes must be queried, and exactly two nodes are removed. The algorithm terminates after $n/2$ steps, at each of which it must query the degree of all remaining nodes, thus it takes $O(n^2)$. In several classes of complex networks, however, the number of nodes removed by the algorithm (i.e. those outside the 2-core) will be a small fraction of n . Line 2: the eigenvalue computation can be done in $O(n^3)$ using, for example, power iteration methods, though its actual runtime will greatly depend on the number of eigenvalues and the structure of the matrix. Line 3: finally, computing the distance between the computed eigenvalues depends on which distance is being used; see next Section.

5.2 Choice of distance and dependence on r

Previous works provide evidence for the fact that the majority of the non-backtracking eigenvalues have magnitude less than $\rho = \sqrt{\lambda_1}$, where λ_1 is the largest eigenvalue (which is guaranteed to be real due to the Perron-Frobenius theorem since all of B 's entries are real and non-negative). Indeed, in [50] the authors show that, in the limit of large network size, non-backtracking eigenvalues with magnitude larger than ρ occur with probability 0. Therefore, to compare two graphs, we propose to compute all the eigenvalues that have magnitude larger than ρ . Informally, this means we are comparing the outliers of the eigenvalue distribution, since they occur with less probability the larger the network becomes. This provides a heuristic to choose the value of r .

To test this heuristic, as well as choose the distance d , we perform the following experiment. Given two graphs G_1 and G_2 , we compute the largest 1000 non-backtracking eigenvalues of both graphs. Then, we compute three distance measures between $\{\lambda_i\}_1^r$ and $\{\mu_i\}_1^r$ for increasing values of r . The distances are defined as follows:

- The **Euclidean** metric computes the Euclidean distance between the vectors whose elements are the eigenvalues sorted in order of magnitude: $\text{Euclidean}(\{\lambda_i\}, \{\mu_i\}) = \sqrt{\sum_i |\lambda_i - \mu_i|^2}$, with $|\lambda_i| \geq |\lambda_{i+1}|$ and $|\mu_i| \geq |\mu_{i+1}|$ for all i . If two eigenvalues have the same magnitude, we use their real parts to break ties. Computing this distance takes $O(r)$ computations, where r is the number of eigenvalues.

- The **EMD** metric computes the Earth Mover Distance (a.k.a. Wasserstein distance) [51] which is best understood as follows. Write $\lambda_k = a_k + ib_k$ and $\mu_k = c_k + id_k$ for all k . Consider the sets $\{(a_k, b_k)\}_{k=1}^r$ and $\{(c_k, d_k)\}_{k=1}^r$ as subsets of \mathbb{R}^2 , and imagine a point mass with weight $1/r$ at each of these points. Intuitively, **EMD** measures the minimum amount of work needed to move all point masses at the points (a_k, b_k) to the points (c_k, d_k) . This distance requires the computation of the (Euclidean) distance between every pair of eigenvalues and thus it takes $O(r^2)$ runtime.
- **Hausdorff** metric is a standard distance between compact sets in a metric space [17]; it is defined as

$$\text{Hausdorff}(\{\lambda_i\}, \{\mu_i\}) = \max \left(\max_i \min_j |\lambda_i - \mu_j|, \max_j \min_i |\lambda_i - \mu_j| \right). \quad (5)$$

Similarly as above, this computation takes $O(r^2)$ runtime.

From each of the following data sets in turn we take two graphs and run this experiment: **HG**, **KR**, **P2P**, **AS**, **CA** (see Table 2 and Section 5.3.1 for full description of data sets). We report the results in Figure 5, where we plot the average of the three distances for each data set across several repetitions of the experiment (10 repetitions for random graphs and **CA**, 36 for **P2P** and **AS**), as well as the median value of r after which the eigenvalues have magnitude less than ρ ; we call this value r_0 .

In Figure 5 we see that **Hausdorff** has high variance and a somewhat erratic behavior on random graphs. **Euclidean** has the least variance and most predictable behavior, continuing to increase as the value of r grows. In the **KR** data sets, **Hausdorff** presents an “elbow” near r_0 , after which it levels off without much change for increasing values of r . Interestingly, the plots suggest that, in the case of **EMD**, the steepest decrease usually occurs before r_0 . Also of note is the fact that on **HG**, all distances have a slight increase for large values of r .

Since **Hausdorff** does not appear to have a predictable behavior, and **Euclidean** continues to increase with no discernible inflection point even when using the largest $r = 1000$ eigenvalues, we favor **EMD** when comparing non-backtracking eigenvalues of two graphs. Furthermore, **EMD** has a predictable “elbow” behavior that happens near r_0 . Thus we conclude that this is an appropriate choice for the number of

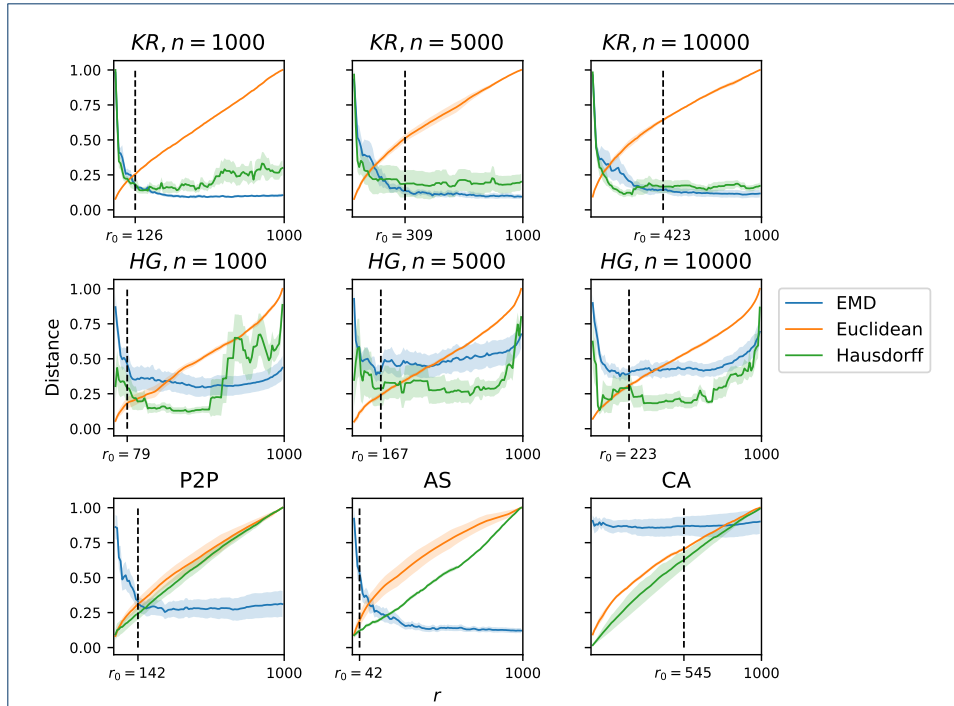


Figure 5 Different ways of comparing non-backtracking eigenvalues. We compare the non-backtracking eigenvalues of two different graphs of the same data set using three different methods: EMD, Euclidean, Hausdorff (see Section 5.2), using the r eigenvalues largest in magnitude. We also show the median value r_0 after which the eigenvalues have magnitude less than $\sqrt{\lambda_1}$. Shaded areas show standard deviation. First two rows show results for random graphs of models KR and HG for different number of nodes n . Bottom row shows results for real data sets. See Table 2 for description of data sets. Values have been normalized by the largest observed distance with each different distance method for ease of comparison.

eigenvalues to use. In the following we use EMD as the distance d and r_0 as the value of r for all our experiments.^[5]

5.3 NBD: Experiments

5.3.1 Data sets and base lines

A description of the data sets we use throughout this article can be found in Table 2. We use random graphs obtained from six different models: Erdős-Rényi (ER) [40, 41], Barabási-Albert (BA) [42], Stochastic Kronecker Graphs (KR) [43, 44], Configuration Model with power law degree distribution (CM) [45], Watts-Strogatz (WS)

^[5]Here we choose both EMD and $r = r_0$ based on the experimental evidence of Figure 5.

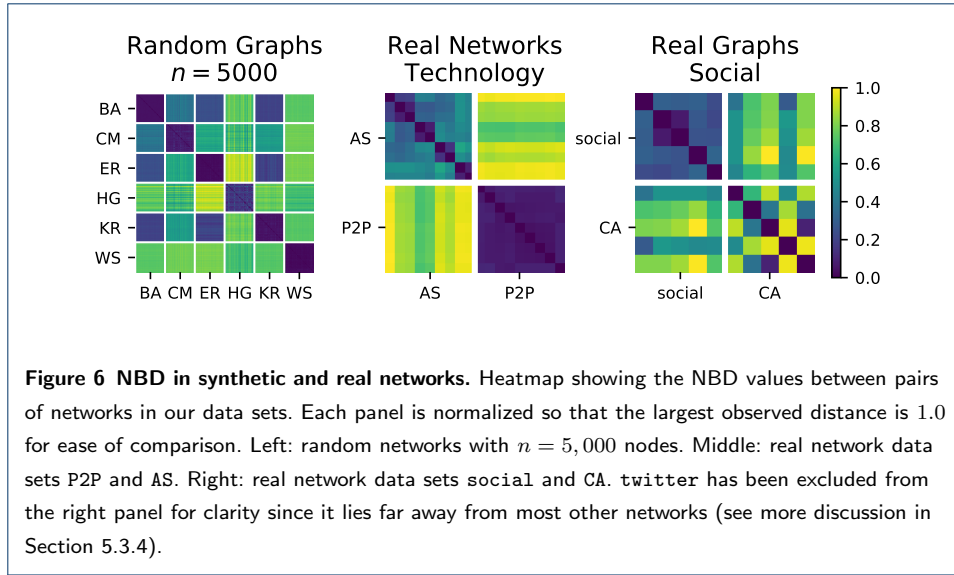
An early version of NBD used Euclidean and a value of r set to a constant [52].

Following that version, other authors have independently proposed using EMD [53].

[46], and Hyperbolic Graphs (HG) [47, 48]. We generated CM and HG with degree distribution $p_k \sim k^{-\gamma}$ with exponent $\gamma = 2.3$, and WS with rewiring probability $p = 0.01$. All graphs have approximate average degree $\langle k \rangle = 15$. From each model, we generate three batches, each with 50 graphs: the first with $n = 1,000$ nodes, the second with $n = 5,000$, and the third with $n = 10,000$. This yields 6×3 data sets that we refer to by model and graph size (see Table 2). We also use real networks, divided in four different data sets: **social**, **CA**, **AS**, and **P2P**. **social** contains six networks obtained from human online social networks (**facebook**, **twitter**, **sdo8**, **sdo9**, **epinions**, and **wiki**) [54, 55, 56, 57], **CA** contains five co-authorship networks (**AstroPh**, **CondMat**, **GrQc**, **HepPh**, **HepTh**) obtained from the arXiv pre-print server [58], **P2P** contains nine snapshots of the peer-to-peer connections of a file-sharing network [58], and **AS** contains nine snapshots of the Internet autonomous systems infrastructure [59]. Note that **P2P** and **AS** are temporal snapshots of the same network and thus we may assume that they have been generated using the same growth process. The same is not true for **social** and **CA**: even though they contain similar kinds of networks, their generating mechanisms may be different. The real networks were obtained through SNAP [60] and ICON [1]. In all, we have 929 different networks (900 random, 29 real), comprising 22 different data sets (6×3 random, 4 real).

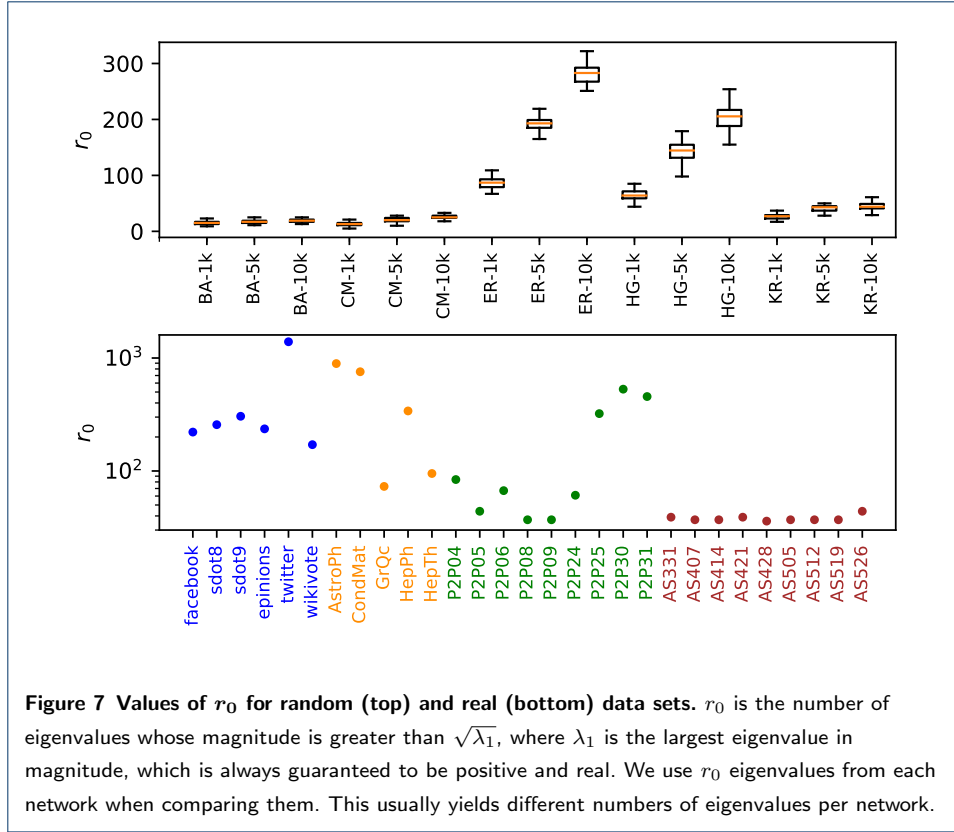
Figure 6 shows the result of computing NBD on different synthetic and real data sets. Observe that NBD is capable of distinguishing between random graphs generated from different models, as well as between the technological data sets **P2P** and **AS**. NBD detects less of a difference between data sets **social** and **CA**, which points toward the homogeneity of connectivity patterns in human-generated networks.

Among all distance methods available, we choose the following three to compare against NBD. The first method is **ESCAPE** [61] which given a graph G computes a 98-dimensional vector v of motif counts, i.e., v_i is the number of times that the i th motif appears in G for $i = 1, \dots, 98$. To compare two graphs we then compute the cosine distance between their **ESCAPE** feature vectors. The second method is Graphlet Correlation Distance [12], or **GCD** for short, which characterizes a graph with a so-called *graphlet correlation matrix*, a 11×11 matrix whose entries are the correlation between the number of per-node occurrences of 11 small motifs. The distance between two graphs is then the Frobenius norm of the difference



of their graphlet correlation matrices. The third method is taken from Schieber et al [8], and will be referred to here as *S*. Distance measure *S* depends on the distribution of shortest path length distances between each pair of nodes in a graph. In this way, we are comparing NBD, which depends on the number and length of non-backtracking cycles, to other distance methods that also depend on the occurrence of a particular type of subgraphs: *ESCAPE* and *GCD* depend on motif counts, while distance *S* depends on lengths of shortest paths. As a baseline, we use the 1-dimensional EMD between the $r = 300$ largest Laplacian eigenvalues of each graph; we call this baseline *Lap*. We report the performance of two versions of NBD: NBD-300 uses the largest $r = 300$ non-backtracking eigenvalues of each graph, while NBD- ρ uses $r = \min\{300, r_0\}$ where r_0 is the number of eigenvalues whose magnitude is greater than $\rho = \sqrt{\lambda_1}$, as explained before. We use a maximum of 300 because most of our networks have $r_0 < 300$; see Figure 7 for values of r_0 for all our data sets. This means that NBD- ρ sometimes uses vastly different numbers of eigenvalues to compare two networks. For example, when comparing a BA-10k graph to *twitter*, we are comparing approximately 15 eigenvalues versus $\min\{300, r_0 = 1393\} = 300$. Note that EMD handles this difference naturally.

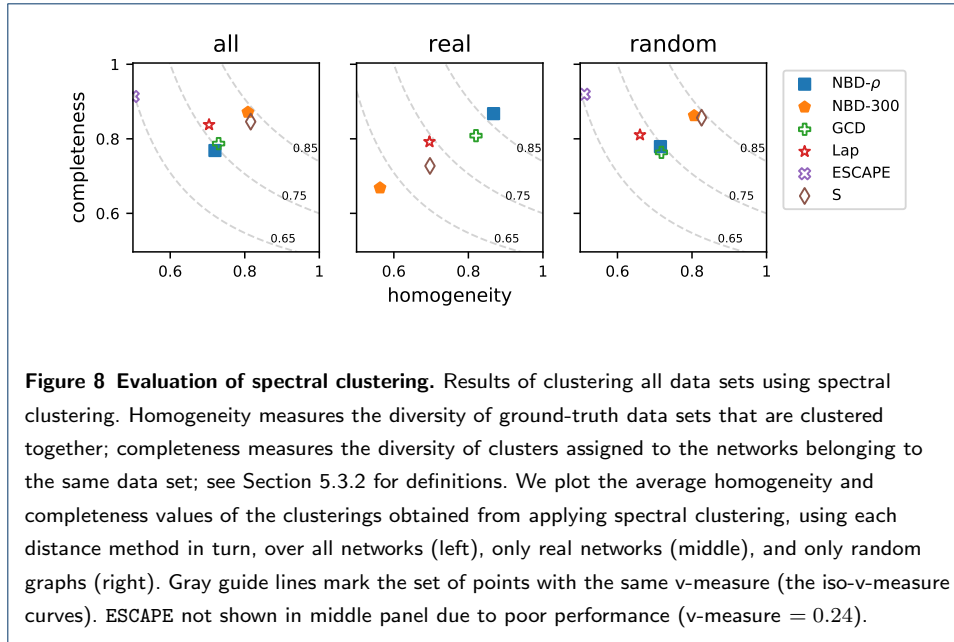
Using each distance method in turn, we obtain a 929×929 distance matrix containing the distance between each pair of networks in our data sets. Using these distance matrices we perform the following experiments.



5.3.2 Clustering networks

Our main experiment is to evaluate the distance methods in a clustering context [62]. We use each distance matrix as the input to a spectral clustering algorithm. To form a proximity graph from the distance matrix, we use the RBF kernel. That is, if the distance between two graphs under the i -th distance method is d , their proximity is given by $\exp(-d^2/2\sigma_i^2)$. We choose σ_i as the mean within-data set average distance, that is, we take the average distance among all pairs of **ER** graphs, the average distance among all pairs of **BA** graphs, and so on for each data set of random graphs, and average all of them together [63]. Note this yields a different σ for each distance method, and that we do not use graphs from real data sets to tune this parameter. We program the algorithm to find the exact number of ground-truth clusters. We report the results of this experiment using all graphs (22 clusters), only real graphs (4 clusters), and only random graphs (18 clusters).

Once we have obtained the clusters, we use as performance metrics *homogeneity*, *completeness*, and *v-measure*. Homogeneity measures the diversity of data sets within a cluster, and is maximized when all data sets in each cluster come from



the same data set (higher is better). Completeness is the converse; it measures the diversity of clusters assigned to the graphs in one data set (higher is better). The v-measure is the harmonic mean of homogeneity and completeness [64]. Concretely, let C be the ground-truth class of a network chosen uniformly at random, and let K be the cluster assigned to a network chosen uniformly at random. Homogeneity is defined as $1 - H(C|K)/H(C)$ where $H(C)$ is the entropy of the class distribution, and $H(C|K)$ is the conditional entropy of the class distribution given the cluster assignment. If the members of each cluster have the same class distribution as the total class distribution, then $H(C|K) = H(C)$, and therefore homogeneity is 0. If each cluster contains elements of the same class, i.e., if the conditional distribution of classes given the clusters is constant, then $H(C|K) = 0$ and homogeneity is 1. If $H(C) = 0$, then homogeneity is defined to be 1 by convention. Similarly, completeness is defined as $1 - H(K|C)/H(K)$, and conventionally equal to 1 when $H(K) = 0$. Note that homogeneity is analogous to *precision*, while completeness is analogous to *recall*, which are used in binary classification. Correspondingly, v-measure is analogous to the F1-score.

In Figure 8 we present the results of this experiment. When clustering all graphs, NBD-300 performs best (v-measure 0.84), followed by S (v-measure 0.83). Distances S and GCD perform well across data sets, however, numerical methods for computing

eigenvalues are orders of magnitude faster than the state-of-the-art methods for computing all shortest path distances (needed to compute **S**) or per-node motif counts (needed for **GCD**); see Section 5.3.3. **ESCAPE** performs well in completeness but not in homogeneity in random graphs, and extremely poorly in real graphs (v-measure 0.27).

Interestingly, when clustering only the real networks, **NBD-300** performs more poorly than when using random graphs, while **NBD- ρ** obtains the highest v-measure score across all our experiments (0.87). This indicates that the choice of $r = r_0$ as the number of eigenvalues larger than $\sqrt{\lambda_1}$ can be both more efficient and more informative than computing a set constant value of, say, $r = 300$.

We conclude from this experiment that (i) counting motifs locally (number of per-node occurrences) is better than counting them globally (number of occurrences in the whole graph), because **GCD** performs better than **ESCAPE** across data sets, (ii) that comparing length distributions of paths or cycles seems to yield better performance than motif counting, because **NBD** and **S** perform better than others in general, and (iii) that non-backtracking eigenvalues provide slightly better performance at distinguishing between real and synthetic graphs than other combinatorial methods, as well as other spectral methods such as **Lap**. We highlight that **NBD** is in general also more efficient than the combinatorial methods, while being only slightly more time-consuming than the spectral method **Lap**; see next Section.

5.3.3 Runtime comparison

As part of Section 5.1 we included a complexity analysis of the runtime of **NBD** (Algorithm 1). Here, we include a direct comparison of runtime of **NBD** and all other baseline algorithms. Observe that all five graph distance algorithms used have something in common: they all run in two steps. The first step is to compute a certain kind of summary statistics of each of the two graphs to be compared, and the second step is to compare those precomputed statistics. For **NBD** and **Lap**, the first step computes eigenvalues of the corresponding matrices, while **GCD** and **ESCAPE** compute motif counts, and **S** computes the shortest path distance between each pair of nodes in the graph. The bulk of the computational time of all methods is spent in the computation of these statistics, rather than in the comparison. Therefore, in this section we compare the runtime of this first step only. Another reason to focus on

the runtime of this first step only is that it depends on one single graph (in all cases, the computation of the corresponding statistics is done on each graph in parallel), and thus we can focus on investigating how the structure of a single graph affects computation time. All experiments were conducted on a computer with 16 cores and 100GB of RAM with Intel Skylake chip architecture, with no other jobs running in the background. Note that all distance algorithms are implemented in Python, except for **ESCAPE** which is implemented in C++. Thus we focus on comparing the four algorithms that are implemented in Python. For big-O complexity analyses we refer the interested reader to each of the corresponding references for each algorithm.

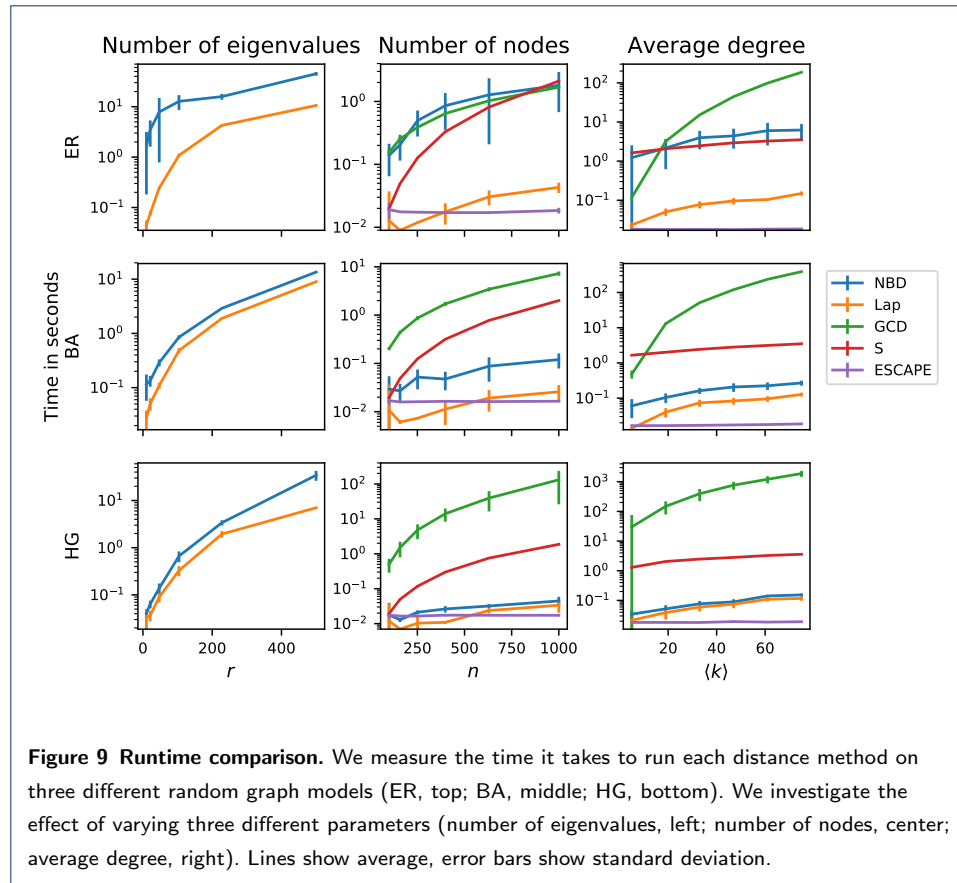


Figure 9 shows results of this experiment. In all panels, the vertical axis shows the running time of the algorithm in seconds, and the horizontal axis presents the value of a structural parameter. Each row corresponds to one of the three following random graph models: ER (top row), BA (middle row), HG (bottom row). Each column corresponds to the variation of a single structural parameter, while keeping the other two fixed. Concretely, the left column shows how the number of eigenvalues

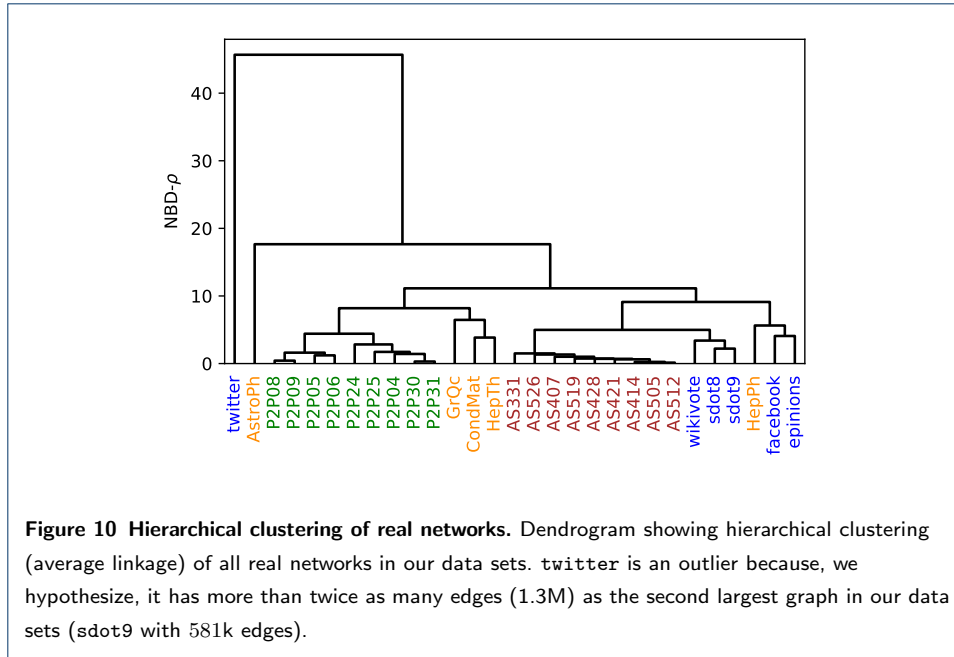
r used for NBD and **Lap** affects their runtime on graphs with $n = 1000$ and $\langle k \rangle = 15$. The center column shows how varying the number of nodes n affects the runtime of all five distances while holding $r = 10$ and $\langle k \rangle = 15$ fixed. The right column shows how varying $\langle k \rangle$ affects the runtime while keeping $r = 10$ and $n = 1000$ fixed. Note that all panels are presented on a log-linear scale.

The left column of Figure 9 shows evidence for the fact that the runtime scaling of NBD is very close to that of **Lap**, as claimed in Section 4.2 because the former uses the non-backtracking matrix, which is always four times the size of the Laplacian matrix. However, their runtimes do not differ by only a constant factor in the cases of ER and HG. We hypothesize that this is because the non-backtracking matrix has a more complicated structure than the Laplacian depending on the structure of the underlying graph (see Section 4.3 and especially Lemma 3). In all other panels we see that NBD and **Lap** present similar scaling behavior, but NBD is usually slower. The middle column presents the effect of varying the number of nodes in the graph on all five distance methods. NBD is consistently the third fastest method for HG and BA graphs, after **ESCAPE** and **Lap**, while **GCD** is consistently the slowest by two or three orders of magnitude in some cases. Note that the scaling behavior of all methods is fairly similar between BA and HG graphs, while for ER it is markedly different. When varying the number of nodes of ER graphs, for example, NBD and **GCD** seem to scale at the same rate, and take the about the same runtime, while **S** scales more poorly than all others. In all other panels, **GCD** seems to scale more poorly than other methods.

From the results presented in Figure 9 we conclude that NBD is a viable alternative in the case of complex networks with heterogeneous degree distributions (modeled by BA and HG) when the application at hand requires comparison of counts of motifs or subgraphs such as non-backtracking paths of arbitrary length.

5.3.4 Nearest neighbors

In Figure 10 we present a dendrogram plot computed on all real networks using a hierarchical clustering (average linkage) of $\text{NBD-}\rho$. We can see that all networks of data sets **P2P** and **AS** are perfectly clustered with each other in one cluster, while five networks of the **social** data set (all except **twitter**) are clustered with each other, although they are broken up into two clusters. Data set **AS** is also broken



up into three different clusters: one with `GrQc`, `CondMat`, and `HepTh`, while `HepPh` and `AstroPh` lie in their own cluster. To test the robustness of these clusters, we ran the same experiment (i) leaving one data set out each time, and (ii) randomly choosing the ordering of the input data sets. These variations are made to address the shortcomings of hierarchical clustering as some implementations may break ties in arbitrary ways. Neither of these variations affected the clusters reported in Figure 10.

To better understand these clusters, we present in Table 3 the $k = 10$ networks that are closest to each of these data sets: `facebook`, `AstroPh`, and `twitter`. We show the nearest networks among all networks, including random graphs, as well as the nearest real networks only. In the case of `facebook`, we see that, even when including random networks, `epinions` is the network closest to it. This is perhaps surprising since these two networks have substantially different values for the number of nodes, average degree, and average clustering coefficient (see Table 2). The networks closest to `facebook`, after `epinions`, all come from the HG random graph model, which is expected because this model has been shown to produce heavy-tailed degree distributions with high clustering, which both `facebook` and `epinions` present as well. When considering only real networks, all but one of the social networks are within the 10 nearest networks to `facebook`. In the case of

AstroPh, the closest network to it is another coauthorship network, **CondMat**. Like in the previous case, it is close to many random graphs of model **HG**. Among the real networks, all other coauthorship graphs are in the nearest 10. Finally, **twitter** is the data set that lies farthest from all others. This may be due to the fact that it has substantially more edges than all others: it has 1.3M edges, while the second largest, **epinions**, has only 748k. We see that the closest random networks to **twitter** all come from the random model **WS**. This may be due to the fact that both **twitter** and **WS** have many large eigenvalues with zero imaginary part, so the earth mover distance between them will be small. We also observe that, even though **twitter** is far from all the networks in the dendrogram, all of the **social** graphs are within the 10 nearest networks to it.

Table 3 For networks **facebook**, **AstroPh**, **twitter** we show the 10 networks closest to each, as determined by $\text{NBD-}\rho$, and report the distance score in parentheses. Note that all social networks except for **twitter** are within the top 5 nearest networks to **facebook**, in spite of their imperfect clustering in Figure 10. Likewise, all coauthorship networks are within the top 7 networks closest to **AstroPh**. Interestingly, **twitter** is also closest to other social networks, as well as some coauthorship networks, even though it is an outlier and the distances are far greater than others.

rank	facebook	AstroPh	twitter
1	epinions (4.08)	CondMat (11.33)	AstroPh (31.38)
2	HepPh (5.26)	epinions (13.72)	CondMat (40.11)
3	sdot9 (6.05)	facebook (14.18)	epinions (41.49)
4	sdot8 (6.24)	GrQc (14.71)	facebook (42.45)
5	wiki (6.56)	HepTh (14.84)	HepTh (42.60)
6	AS414 (8.98)	sdot9 (16.96)	HepPh (44.30)
7	AS505 (9.03)	HepPh (16.97)	GrQc (44.54)
8	AS512 (9.05)	sdot8 (17.37)	sdot9 (44.66)
9	AS421 (9.06)	AS407 (17.62)	wiki (45.64)
10	AS407 (9.16)	AS414 (17.68)	sdot8 (45.88)

From these observations we conclude that NBD is able to identify networks that are generated by similar mechanisms: all technological graphs – i.e., autonomous systems of the Internet (AS) and automatically generated peer-to-peer file sharing networks (P2P) – are perfectly clustered together, and all online social graphs are ranked near each other, even when some of these have vastly different elementary statistics such as number of nodes, number of edges, and average clustering coefficient.

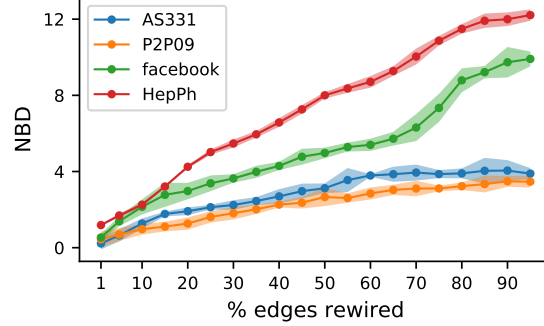


Figure 11 NBD in the presence of noise. For each graph AS331, P2P09, facebook, HepPh we compute 10 rewired versions using degree-preserving randomization, for increasing percentage of edges rewired. We show the average NBD between the original graph and the rewired versions; shaded regions indicate standard deviation. As expected, NBD is an increasing function the percentage of rewired edges. This implies that NBD is detecting graph structural features that are not determined by degree distribution.

5.3.5 Rewiring edges

We are interested in the behavior of NBD in the presence of noise. For this purpose, given a graph G , we compute G' by rewiring some of the edges of G in such a way that the degree distribution is kept intact (degree preserving randomization). We then compute the NBD between G and G' . For this experiment we use four networks, one from each of our real data sets: AS331, P2P09, facebook, and HepPh. We show our results in Figure 11.

As expected, the NBD between a graph and its rewired versions grows steadily as the percentage of rewired edges increases. Note that AS331 and P2P09 seem to plateau when the percentage of rewired edges is 65% or larger, while facebook and HepPh have a slight jump close to 70%. Observe too that the variance of the NBD is kept fairly constant, which indicates that the distribution of non-backtracking cycle lengths in all rewired versions of each graph is very similar to each other. This may be due to the fact that edge rewiring destroys short cycles more quickly than long cycles. Furthermore, observe that for high percentages of rewired edges, the rewired graph is essentially a configuration model graph with the same degree distribution as the original. Therefore, by comparing the NBD values from Figures 10 and 11, we can conclude that the NBD between these four networks and a random graph with the same degree distribution is larger than the average distance to other networks in the same data set (cf. Figure 10 and discussion in previous Section). We conclude

that NBD is detecting graph structural properties – stored in the distribution of non-backtracking cycle lengths and encoded in the non-backtracking eigenvalues –, that are not solely determined by the degree-distribution of the graph.

6 NBED: Non-backtracking embedding dimensions

In this Section we discuss the potential of a second application of the eigendecomposition of the non-backtracking matrix in the form of a new embedding technique. In this case, instead of the eigenvalues, we use the eigenvectors. We propose the Non-Backtracking Embedding Dimensions (or NBED for short) as an edge embedding technique that assigns to each undirected edge of the graph two points in \mathbb{R}^d – one for each orientation –, where d is the embedding dimension. We are able to find many interesting patterns that encode the structure of the network, which we then put to use for anomaly detection. While the visualizations found with NBED are certainly peculiar and rich in information, future work will prove essential in determining the full extent of possible applications of NBED.

Concretely, given a graph G with edge set E and directed edge set \vec{E} , consider a unit eigenvector v corresponding to the non-backtracking eigenvalue λ . Whenever λ is real, the entries of v are guaranteed to be real. Given the d eigenvectors $\{v^i\}_{i=1}^{i=d}$ corresponding to the d eigenvalues with largest magnitude, the embedding of an edge $(k, l) \in E$ is a point in \mathbb{R}^{2d} whose entries are given by

$$(f(v_{k \rightarrow l}^1), f(v_{k \rightarrow l}^2), \dots, f(v_{k \rightarrow l}^d), f(v_{l \rightarrow k}^1), f(v_{l \rightarrow k}^2), \dots, f(v_{l \rightarrow k}^d)), \quad (6)$$

where f is defined as $f(v_{k \rightarrow l}^i) = \text{Re}(\lambda_i)\text{Re}(v_{k \rightarrow l}^i) - \text{Im}(\lambda_i)\text{Im}(v_{k \rightarrow l}^i)$ and Re , Im are the real and imaginary parts of a complex number, respectively. As mentioned previously, λ_1 , the largest eigenvalue, is always guaranteed to be real and positive, and the entries of v^1 are also real and positive. Thus, $f(v_{k \rightarrow l}^1) = v_{k \rightarrow l}^1$ for every $k \rightarrow l$. Whenever any $\lambda_i, i = 2, 3, \dots$, has non-zero imaginary part, the entries of v^i may also be complex numbers in general, and $f(v_{k \rightarrow l}^i)$ is simply a linear combination of the real and imaginary parts.

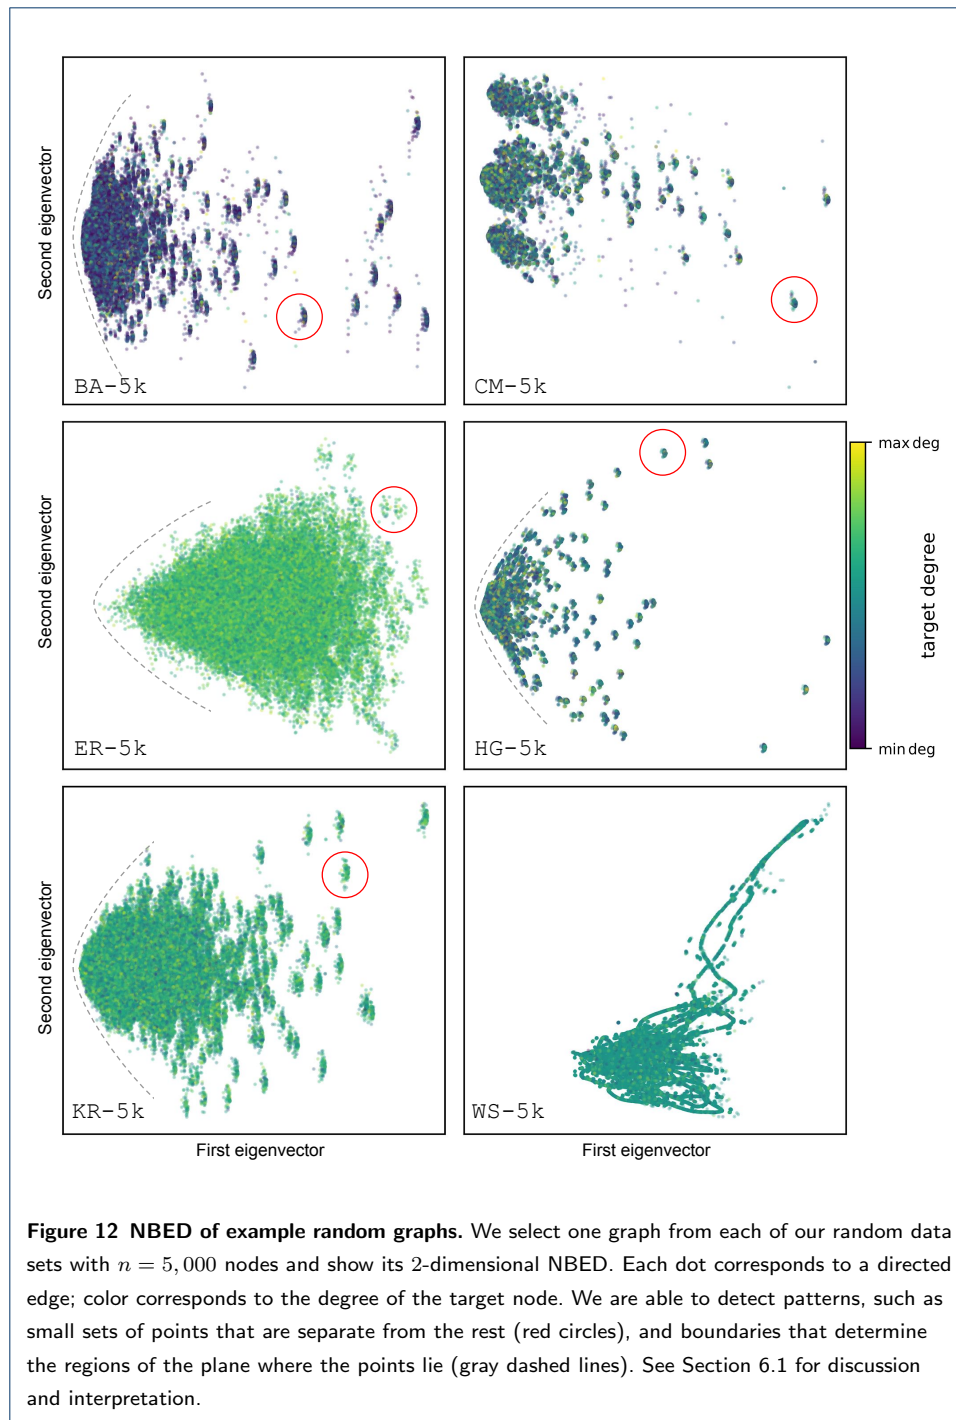
In the following we use the 2-dimensional NBED of real and random graphs. The advantages of this 2-dimensional edge embedding are as follows. First, the first and second eigenvectors of the non-backtracking matrix have been studied before and they can be interpreted in terms of edge centrality [24, 25] and community

detection [26, 27, 28] respectively. Second, NBED provides a deterministic graph embedding, as opposed to other methods that are stochastic due to their dependence on random sampling of subgraphs [65, 66] or due to stochastic optimization of some model [67, 68]. Third, NBED makes a distinction between the two possible orientations of an edge thus providing more information about it. We conjecture that these properties make NBED more robust to the presence of noise in the graph when compared to embeddings generated by popular models such as deep networks, which are in general not interpretable, are stochastic, and usually work on nodes rather than directed edges. Further, the 2-dimensional version of NBED is particularly important because our experimental evidence shows that the second non-backtracking eigenvalue of complex networks tends to be real and positive, which means that we can visualize these embeddings on the real plane without losing any information due to the linear combination f . We proceed to analyze the visual patterns we can identify in NBED visualizations; future research will be necessary for analytic characterizations of them.

6.1 Visualization


In Figure 12 we present the 2-dimensional NBED of one network from each of our random data sets with $n = 5,000$ nodes (see Section 5.3.1 for description). Each edge (k, l) corresponds to two dots in this plot: $(v_{k \rightarrow l}^1, v_{k \rightarrow l}^2)$ and $(v_{l \rightarrow k}^1, v_{l \rightarrow k}^2)$. All graphs used in Figure 12 have their two largest eigenvalues real and positive. We proceed to describe the patterns that we observe in these plots and how they relate to structural properties of the graph.

- 1 In all of the plots except for **WS-5k** we see small sets of dots (red circles) that are markedly separated from the bulk of the dots. Every one of these sets is made up of the embeddings of the directed edges $k \rightarrow l$ for a fixed l , and the horizontal position of this set corresponds to the degree of l . That is, the larger the degree of a fixed node l , the more likely all the embeddings of edges of the form $k \rightarrow l$ are to cluster together separate from the rest, and the further right this cluster will be.
- 2 Two of these sets corresponding to the same graph seem to have similar structures. That is, the relative positions of dots forming a cluster encircled in red is repeated for all such clusters inside each graph. For example, the sets of



BA-5k and KR-5k are much more vertically spread than those of CM-5k or HG-5k. We conclude that each of these sets have a particular internal structure that correlates with global graph structure, as it is repeated across sets. In ER-5K, the internal structure of each set is less well marked, owing to the fact that there are no global structural patterns in the ER model.

- 3 In BA-5k, ER-5k, HG-5k, and KR-5k, we find that the dots inhabit a certain region of the plane that is bounded by a roughly parabolic boundary (gray dashed lines). Furthermore, the cusp of this boundary always lies precisely at the origin. That is, there are never dots to the left of the origin, and the vertical spread of the dots grows at roughly the same rate both in the upper half-plane and the lower half-plane. Since the vertical axis corresponds to the entries of the second eigenvector v^2 which has been shown to describe community structure, we conjecture that the significance of this boundary is related to the existence of communities in the network. (See Section 7 for more discussion on this point.)
- 4 In Figure 12, the color of the embedding of edge $k \rightarrow l$ corresponds to the degree of the source node k . In this way, we can use these plots to gather information about the degree distribution of the graph. For example, the plot for ER-5k looks like a noisy point cloud because the degree distribution is Poisson centered around $\langle k \rangle = 15$, and the plot for WS-5k is very homogeneous in color because all nodes start with the same degree $k = 14$ and only a few of them are rewired. Furthermore, if one compares the color distribution of the plots for BA-5k, CM-5k, and HG-5k, one can differentiate that even though all have heavy-tailed degree distributions, their exponents must be different. Indeed, for BA-5k we have $\gamma = 3$ while the other two were generated with $\gamma = 2.3$.
- 5 From all previous considerations we can analyze the plot of KR-5k as a case study. First, it presents small sets of dots that are markedly separate from the bulk, and whose structure is similar to those of BA-5k. Second, its color distribution is closer to that of ER-5k than to any of the others. Third, it exhibits the boundary effect that we see in both BA-5k and ER-5k. From these observations we conclude that the stochastic Kronecker graph model is a “mix” between the Barabasi-Albert and Erdos-Renyi models, in the sense its plot has elements from both of the plots for BA and ER. Consider this in light of Figure 1 where the graphs of these three data sets are clustered together, even when using the largest $r = 200$ non-backtracking eigenvalues of each graph. This means that the eigendecomposition of the non-backtracking matrix is detecting the similarities and differences between these three random

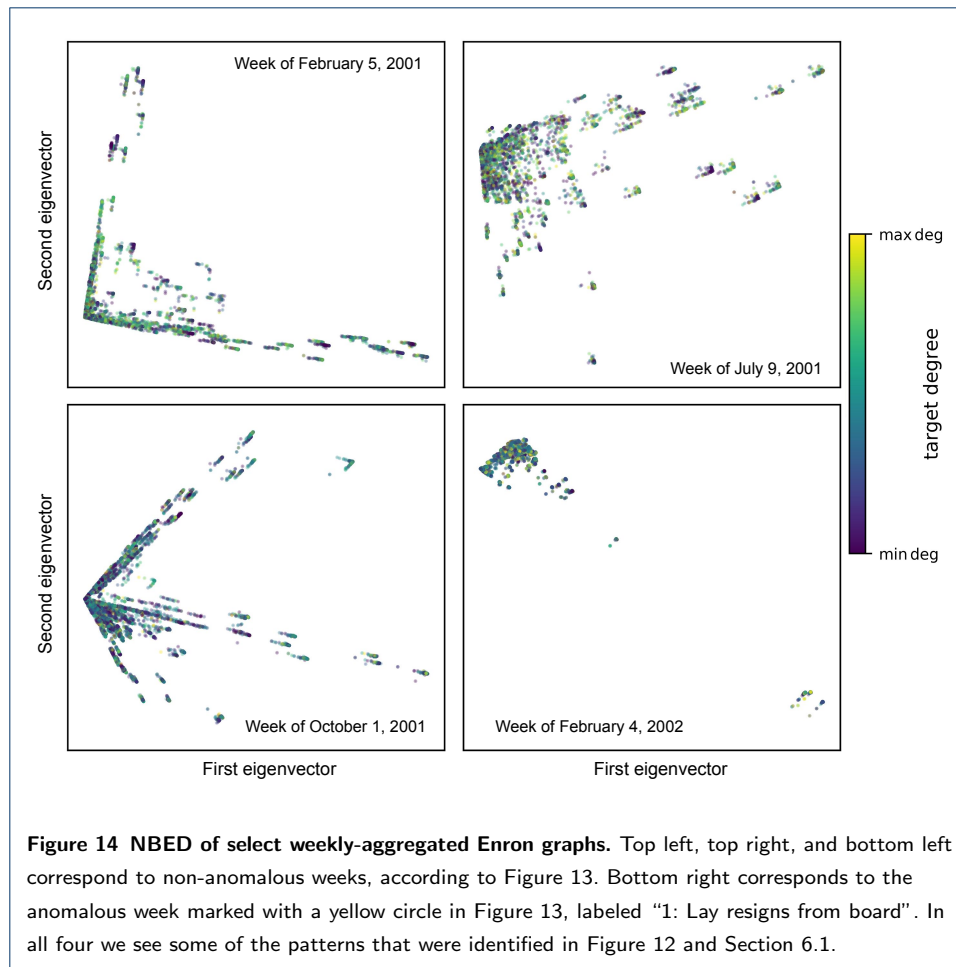


graph models, which we can identify through a visual analysis of Figures 1 and 12.

- 6 Lastly, we focus on the plot for WS-5k in Figure 12. This plot shows none of the characteristics of the others, and in fact some of its points clearly fall on a path determined by a continuous curve. We hypothesize that this is due to the strong symmetry of the ring lattice from which it is generated. Studying the properties of the NBED of highly symmetric graphs and their randomized versions is a task we leave for a future work.

7 Case Study: Enron emails

In this Section we use both NBD and NBED to perform an analysis of the well-known Enron email corpus [15]. This data set is comprised of email communications of Enron corporation employees in the years 2000-2002. From it, we form two different data sets of networks. In all networks each node represents an email address, and an edge between two email addresses represents that at least one email was sent from one account to the other. We aggregate these networks on a weekly basis, as well as on a daily basis, and we proceed to apply our methods NBD and NBED on them.



In the top panel of Figure 13 we show the distance between a daily graph and the graph corresponding to a Sunday. In this plot we clearly see that NBD is able to detect the periodicity of email communications. In the bottom plot we show the distance between a weekly graph and the week starting on Monday January 1st 2001. In this case, we see that peaks in NBD values (or in r_0 values) correlate to several major events that occurred during the Enron bankruptcy scandal [69, 70, 71]. Not all important events are detected this way, however, as the event “1: Skilling becomes CEO” is not identified through high peaks in either NBD or r_0 . Further research is necessary to clarify precisely what kind of events are being detected in this way.

In Figure 14 we show the NBED of four weekly-aggregated graphs. The first three (top right, top left, bottom left) correspond to non-anomalous weeks, while the one on the bottom right, corresponding to the week of February 4th, 2002,

is considered anomalous (cf. Figure 13 where it is marked with a yellow circle and labeled "7: Lay resigns from board"). Most other weekly graphs have NBED plots that generally resemble the first three. In these plots we are able to see once again some of the features identified in Section 12. First, they exhibit the sets of points clustered together separate from the rest with a similar internal structure to each other; like before, these correspond to the incoming edges of high degree nodes. Second, the roughly parabolic boundaries of Figure 12 have become sharply pointed with a cusp at the origin, from which two seemingly straight line segments are born in different directions. Third, the angle that these straight line segments form at the origin is different from graph to graph, and the lines are not symmetric with respect to the horizontal axis. As mentioned previously, the second eigenvector (e.g., the vertical axis in Figures 12 and 14) is related to community structure. We hypothesize that the asymmetry seen in Enron graphs but not in random graphs is due to the presence of unevenly-sized communities; see Appendix 2. Lastly, the anomalous week of February 4th, 2002 (bottom right) looks decidedly different than the others, due to the presence of a few points very far apart from the bulk of the plot.

We have presented a visual analysis of the Enron corpus using NBD and NBED. Using these techniques, we are able to recover properties not only of the underlying network, but of the underlying data set as well, such as periodicity and temporal anomalies. Even though the analytical details of NBED require further consideration, we are still able to interpret the visualizations of Figure 14 to mine important information about the underlying data set.

8 Discussion and conclusions

We have focused on the problem of deriving a notion of graph distance for complex networks based on the length spectrum function. We add to the repertoire of distance methods the Non-Backtracking Spectral Distance (NBD): a principled, interpretable, computationally efficient, and effective technique that takes advantage of the fact that one can interpret the non-backtracking cycles of a graph as its free homotopy classes. NBD is principled because it is backed by the theory of the length spectrum, which characterizes the 2-core of a graph up to isomorphism. It is interpretable because we can study its behavior in the presence of structural fea-

tures such as hubs and triangles, and we can use the resulting geometric features of the eigenvalue distribution to our advantage. It is efficient relative to other similar methods that depend on the combinatorial enumeration of different kinds of sub-graphs. Lastly, we have presented extensive experimental evidence to show that it is effective at discriminating between complex networks in various contexts, including visualization, clustering, and anomaly detection. Performance of NBD is better or comparable to other distance methods such as **ESCAPE**, **S**, **GCD**, and **Lap**; see Section 5.3. We chose to compare against these methods because the first three depend on motif counts in one way or another, as NBD depends on non-backtracking cycle counts, and **Lap** depends on the spectral decomposition of a matrix representation of a graph, as NBD depends on the non-backtracking eigenvalues.

Motivated by the usefulness of NBD due to the connections with the homotopy of graphs and the spectrum of the non-backtracking matrix, we also presented a new embedding technique, Non-Backtracking Embedding Dimensions (or NBED for short) which provides a rich visualization full of interpretable patterns that describe the structural properties of a network. We have provided examples of these patterns as well as their application to anomaly detection. Further research will reveal the full potential of applications of NBED.

An implementation of NBD, NBED and our algorithm for computing the non-backtracking matrix, is available at [16].

Limitations NBD relies on the assumption that the non-backtracking cycles contain enough information about the network. Accordingly, the usefulness of the NBD will decay as the 2-core of the graph gets smaller. For example, trees have an empty 2-core, and all of its non-backtracking eigenvalues are equal to zero. In order to compare trees, and more generally, those nodes outside the 2-core of the graph, the authors of [36] propose several different strategies, for example adding a “cone node” that connects to every other node in the graph. However, many real-world networks are not trees and we extensively showcased the utility of NBD on this class of networks.

The greatest limitation of NBED is that we are not able to provide rigorous derivations for the patterns we identify in Section 6.1. Without a formal theory of the relationship between the eigenvectors of the non-backtracking matrix and

the structural properties of graphs it is difficult to design algorithms that make use of these patterns automatically. Regardless, we have presented evidence for the usefulness of these patterns even with a visual analysis.

Future work There are many other avenues to explore in relation to how to exploit the information stored in the length spectrum and the fundamental group of a graph. As mentioned in Sec. 4.1, the major downside of the relaxed length spectrum \mathcal{L}' is the fact that we lose information stored in the combinatorics of the fundamental group. That is, \mathcal{L}' stores information about the frequency of lengths of free homotopy classes, but no information on their concatenation, i.e., the group operation in $\pi_1(G)$. One way to encapsulate this information is by taking into account not only the frequency of each possible length of non-backtracking cycles, but also the number of non-backtracking cycles of fixed lengths ℓ_1 and ℓ_2 that can be concatenated to form a non-backtracking cycle of length ℓ_3 . It remains an open question whether this information can be computed using the non-backtracking matrix for all values of the parameters ℓ_1, ℓ_2, ℓ_3 , and if so, how to do it efficiently. One alternative is to rely upon efficient motif counting [61, 72].

A different research direction is to focus on the non-backtracking eigenvalues themselves, independently of the length spectrum theory. One standing question is to characterize the behavior of the eigenvalues after the network has been rewired. Section 5.3.5 only scratches the surface of what can be said in this regard. However, spectral analysis of the non-backtracking matrix is exceedingly difficult due to the fact that it is asymmetric and non-normal and therefore most of the usual tools for spectral analysis are not applicable.

In this work, we have focused on introducing and exploiting novel theoretical concepts such as the length spectrum and the fundamental group to the study of complex networks. We are confident this work will pave the road for more research in topological and geometric data analysis in network science.

Appendix 1: UMAP parameter settings

In order to understand the visualizations of the non-backtracking eigenvalues in Figure 1 we will now explain some of the features of the UMAP algorithm. For full detail we refer the interested reader to [49].

UMAP stands for Uniform Manifold Approximation and Projection, authored by Leland McInnes, John Healy, and James Melville. Data are first represented in a high dimensional euclidean space using Laplacian Eigenmaps [73], then they are approximated uniformly using fuzzy simplicial sets and patched together to form a manifold, and finally this manifold is projected to \mathbb{R}^2 . This process reveals topological features of the data and provides flexibility for geometrical specifications to distinguish task-specific quantities of interest.

In this respect, UMAP provides a very flexible algorithm, which has various parameters. In the interest of reproducibility, we list our parameter choices of the subset of parameters that yield results of interest for network science, as they allow us to separate different random graph models from each other. A full explanation of these and other parameters is included in [49] in Section 4.3. In Figure 1 of this reference authors show a comparison of how visualizations change with the choice of parameters.

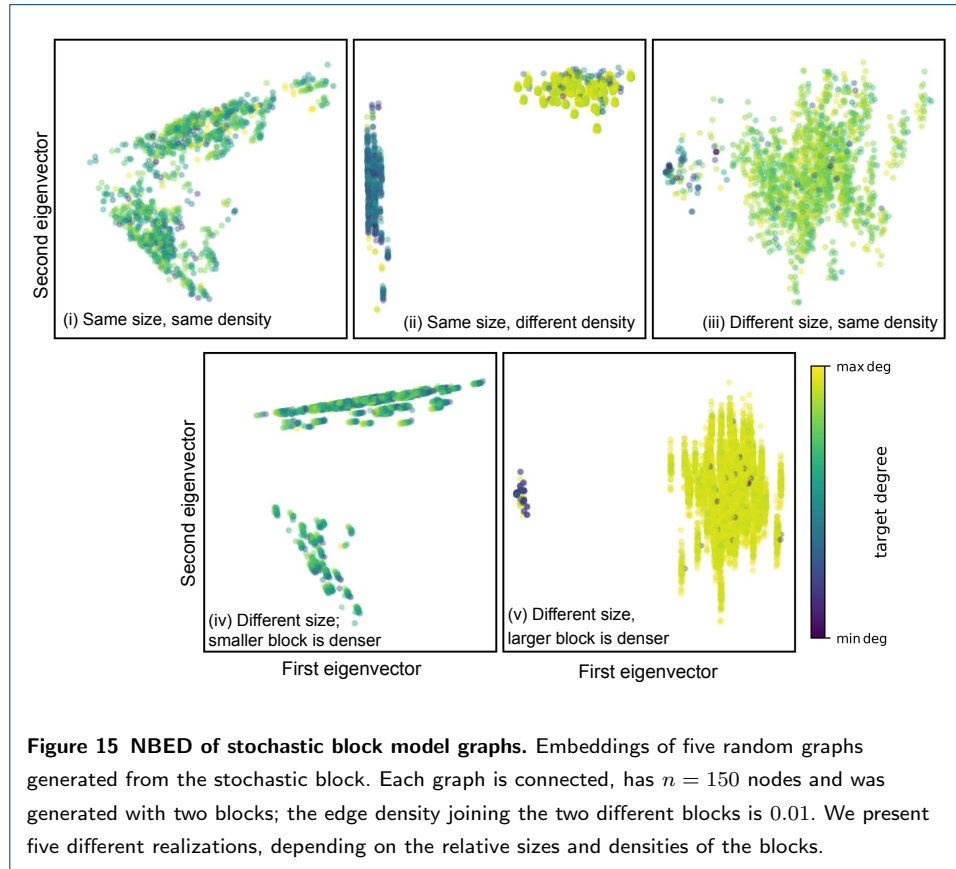
- **n_neighbors:** A weighted k -nearest neighbor graph is constructed from the initial data. The number of neighbors is set with this parameter. We used a value of 75.
- **metric:** Different metrics in the high dimensional space where the data are embedded can be specified. We found good overall results with the Canberra metric (shown in Figure 1. The Chebyshev and Euclidean metrics cluster ER graphs together, for some specific values of the other parameters. However they make HG and CM clusters overlap. A clearly separated projection for the HG and CM graphs can be found using Correlation metric.
- **n_epochs:** Training optimization epochs. More epochs can provide better results, with the usual computational drawbacks. We use 1000 epochs.
- **min_dist:** Alters the minimum distance between embedded points. Smaller values increase clustering, larger values present a more uniformly distributed visualiaztion. We use a distance of 0.01.
- **repulsion_strength:** Values above 1 increase the weight of negative samples, i.e., the importance of the distances between far-apart data points. We use a value of 10.

- **negative_sample_rate**: Sets the number of negative samples to be selected per each positive sample in the optimization process. We use 50 negative samples per positive sample.

We were not able to find a specific set of parameters that improved upon the visualization shown in Figure 1, though it is perhaps the interplay between **min_dist** and **repulsion_strength** what causes ER graphs to split into two different clusters.

Appendix 2: Asymmetric NBED embeddings

As mentioned in Section 7, the NBED of Enron graphs present a feature that we do not see in any of the random graphs we handled for this work. Namely, the NBED of Enron graphs are asymmetric with respect to the horizontal axis (Figure 14), whereas the NBED of all other random graphs (Figure 12) are generally symmetric. In Section 7 we hypothesized that this is due to the fact that Enron graphs contain communities. Here, we present some evidence for this hypothesis.



In Figure 15 we show the NBED of five random graphs generated from the stochastic block model as follows. Each graph has $n = 150$ nodes and was generated with

two blocks. In all cases, the generated graphs were connected and the edge density joining the two different blocks was 0.01. We vary the size and internal density of each block to yield five different realizations: (i) two blocks of the same size and same internal density equal to 0.1, (ii) blocks of same size, one with density 0.5 and the other with density 0.1, (iii) blocks of different size, with the same density 0.1, (iv) blocks of different size where the larger block has density 0.1 and the smaller has density 0.5, and (v) blocks of different size where the larger block has density 0.5 and the smaller has density 0.1. In cases (ii) and (iv), we see NBED embeddings that are markedly asymmetrical with respect to the horizontal axis. Furthermore, case (iv) in particular is visually reminiscent of the Enron embeddings. Further work will elucidate more systematically the role of differently-sized communities with different densities on the eigenvalues and eigenvectors of the non-backtracking matrix.

Availability of data and material

With this work we also release publicly available source code that implements NBD, NBED and the related matrix computations [16].

Competing interests

The authors declare that they have no competing interests.

Funding

Torres and Eliassi-Rad were funded by the National Science Foundation grants NSF CNS-1314603, NSF IIS-1741197. Suárez-Serrato was supported by University of California Institute for Mexico and the United States CN-16-43 and Universidad Nacional Autónoma de México DGAPA-UNAM PAPIIT IN102716, DGAPA-UNAM PASPA.

Author's contributions

This manuscript would not exist without LT; he was the lead contributor in all aspects of the manuscript. PSS introduced the mathematical concept of length spectra and topological data analysis to LT and TER. LT and PSS were the major contributors to the topological graph analysis and graph distance in the manuscript. LT and TER were the major contributors to the graph embedding and experimental design in the manuscript. LT developed the code and conducted all the experiments. LT was the lead contributor in writing the manuscript. All authors read and approved the final manuscript.

Acknowledgements

We thank Evimaria Terzi for her contributions to an earlier version of this work. Torres and Eliassi-Rad were supported by NSF CNS-1314603 and NSF IIS-1741197. Suárez-Serrato was supported by UC-MEXUS (University of California Institute for Mexico and the United States) CN-16-43, DGAPA-UNAM PAPIIT IN102716, and DGAPA-UNAM PASPA program. Part of this research was performed while Suárez Serrato was visiting the Institute for Pure and Applied Mathematics (IPAM), which is supported by NSF DMS-1440415. Eliassi-Rad and Suárez Serrato met during the IPAM Long Program on Culture Analytics in 2016,^[6] of which Eliassi-Rad was a co-organizer.

^[6]<https://www.ipam.ucla.edu/programs/long-programs/culture-analytics/>

Author details

¹ Network Science Institute, Northeastern University, 360 Huntington Ave, 02115-5000, Boston, MA, USA. ² Department of Mathematics, University of California, Santa Barbara, 552 University Rd, 93117 Isla Vista, CA, USA. ³ Instituto de Matemáticas, Universidad Nacional Autónoma de México, Área de la Investigación Científica, Circuito Exterior, Ciudad Universitaria, 04510, CDMX Ciudad de México, México. ⁴ Khoury College of Computer Sciences, 360 Huntington Ave, 02115-5000, Boston, MA, USA.

References

1. Clauset, A., Tucker, E., Sainz, M.: The Colorado Index of Complex Networks. <https://icon.colorado.edu/> Accessed 19 June 2018
2. Kunegis, J.: KONECT: The Koblenz network collection. In: Proc. of the 22nd Int'l World Wide Web Conference (WWW), pp. 1343–1350 (2013)
3. Soundarajan, S., Eliassi-Rad, T., Gallagher, B.: A guide to selecting a network similarity method. In: Proc. of the 2014 SIAM Int'l Conf. on Data Mining (SDM), pp. 1037–1045 (2014)
4. Koutra, D., Shah, N., Vogelstein, J.T., Gallagher, B., Faloutsos, C.: DeltaCon: Principled massive-graph similarity function with attribution. TKDD **10**(3), 28–12843 (2016)
5. Bagrow, J.P., Boltt, E.M.: An information-theoretic, all-scales approach to comparing networks. arXiv:1804.03665 [cs.SI] (2018)
6. Bento, J., Ioannidis, S.: A family of tractable graph distances. In: Proc. of the 2018 SIAM Int'l Conf. on Data Mining (SDM), pp. 333–341 (2018)
7. Onnela, J.-P., Fenn, D.J., Reid, S., Porter, M.A., Mucha, P.J., Fricker, M.D., Jones, N.S.: Taxonomies of networks from community structure. Phys. Rev. E **86**, 036104 (2012)
8. Schieber, T.A., Carpi, L., Díaz-Guilera, A., Pardalos, P.M., Masoller, C., Ravetti, M.G.: Quantification of network structural dissimilarities. Nature communications **8**, 13928 (2017)
9. Chowdhury, S., Mémoli, F.: Distances and isomorphism between networks and the stability of network invariants. arXiv:1708.04727 [cs.DM] (2017)
10. Chowdhury, S., Mémoli, F.: The metric space of networks. arXiv:1804.02820 [cs.DM] (2018)
11. Berlingerio, M., Koutra, D., Eliassi-Rad, T., Faloutsos, C.: Network similarity via multiple social theories. In: Advances in Social Networks Analysis and Mining (ASONAM), pp. 1439–1440 (2013)
12. Yaveroğlu, Ö.N., Malod-Dognin, N., Davis, D., Levnajic, Z., Janjic, V., Karapandza, R., Stojmirovic, A., Pržulj, N.: Revealing the hidden language of complex networks. Scientific reports **4**, 4547 (2014)
13. Constantine, D., Lafont, J.-F.: Marked length rigidity for one-dimensional spaces. J. Topol. Anal. (2018)
14. Leininger, C.J., McReynolds, D.B., Neumann, W.D., Reid, A.W.: Length and eigenvalue equivalence. Int. Math. Res. Not. IMRN **2007**(24), 135 (2007)
15. Klimt, B., Yang, Y.: The enron corpus: A new dataset for email classification research. In: Boulicaut, J.-F., Esposito, F., Giannotti, F., Pedreschi, D. (eds.) Proc. of the 2004 European Conference on Machine Learning (ECML), pp. 217–226 (2004)
16. Torres, L.: SuNBEM: Spectral Non-Backtracking Embedding And pseudo-Metric. GitHub (2018). [url{https://github.com/leotrs/sunbeam}](https://github.com/leotrs/sunbeam)
17. Munkres, J.R.: Topology, 2nd edn. Prentice Hall, Englewood Cliffs, N.J. (2000)
18. Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge; New York (2017)
19. Batagelj, V., Zaversnik, M.: Fast algorithms for determining (generalized) core groups in social networks. Adv. Data. Anal. Classi. **5**(2), 129–145 (2011)
20. Terras, A.: Zeta Functions of Graphs: A Stroll Through the Garden. Cambridge Studies in Advanced Mathematics, p. 239. Cambridge University Press, Cambridge, Cambridge; New York (2011)
21. Lang, S.: Linear Algebra, 3rd edn. Springer, New York (2004)
22. Hashimoto, K.: Zeta functions of finite graphs and representations of p -adic groups. In: Automorphic Forms and Geometry of Arithmetic Varieties. Adv. Stud. Pure Math., vol. 15, pp. 211–280 (1989)
23. Bass, H.: The Ihara-Selberg zeta function of a tree lattice. Internat. J. Math. **3**(6), 717–797 (1992)
24. Martin, T., Zhang, X., Newman, M.E.J.: Localization and centrality in networks. Phys. Rev. E **90**, 052808 (2014)
25. Grindrod, P., Higham, D.J., Noferini, V.: The deformed graph Laplacian and its applications to network

- centrality analysis. *SIAM J. Matrix Anal. Appl.* **39**(1), 310–341 (2018)
26. Krzakala, F., Moore, C., Mossel, E., Neeman, J., Sly, A., Zdeborová, L., Zhang, P.: Spectral redemption in clustering sparse networks. *Proc. Natl. Acad. Sci. USA* **110**(52), 20935–20940 (2013)
 27. Bordenave, C., Lelarge, M., Massoulié, L.: Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs. In: *Proc. of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 1347–1357 (2015)
 28. Kawamoto, T.: Localized eigenvectors of the non-backtracking matrix. *J. Stat. Mech. Theory Exp.* (2), 023404 (2016)
 29. Ren, P., Wilson, R.C., Hancock, E.R.: Graph characterization via Ihara coefficients. *IEEE T. Neural Networ.* **22**(2), 233–245 (2011)
 30. Jiang, F., He, L., Zheng, Y., Zhu, E., Xu, J., Yu, P.S.: On spectral graph embedding: A non-backtracking perspective and graph approximation. In: *Proc. of the 2018 SIAM Int'l Conf. on Data Mining (SDM)*, pp. 324–332 (2018)
 31. Gulikers, L., Lelarge, M., Massoulié, L.: Non-backtracking spectrum of degree-corrected stochastic block models. In: *Proc. of the 8th Innovations in Theoretical Computer Science (ITCS)*, pp. 44–14427 (2017)
 32. Estrada, E.: Spectral moments of the edge adjacency matrix in molecular graphs, 1. definition and applications to the prediction of physical properties of alkanes. *J. Chem. Inf. Comp. Sci.* **36**(4), 844–849 (1996)
 33. Preciado, V.M., Jadbabaie, A., Verghese, G.C.: Structural analysis of Laplacian spectral properties of large-scale networks. *IEEE Trans. Automat. Contr.* **58**(9), 2338–2343 (2013)
 34. Goyal, P., Ferrara, E.: Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems* **151**, 78–94 (2018)
 35. Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.* **40**(3), 52–74 (2017)
 36. Durfee, C., Martin, K.: Distinguishing graphs with zeta functions and generalized spectra. *Linear Algebra Appl.* **481**, 54–82 (2015)
 37. Cooper, Y.: Properties determined by the Ihara zeta function of a graph. *Electron. J. Combin.* **16**(1), 14–84 (2009)
 38. Angel, O., Friedman, J., Hoory, S.: The non-backtracking spectrum of the universal cover of a graph. *Trans. Amer. Math. Soc.* **367**(6), 4287–4318 (2015)
 39. Wood, P.M., Wang, K.: Limiting empirical spectral distribution for the non-backtracking matrix of an Erdős-Rényi random graph. *arXiv:1710.11015 [math.PR]* (2017)
 40. Erdős, P., Rényi, A.: On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.* **5**, 17 (1960)
 41. Bollobás, B.: *Random Graphs*, 2nd edn. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge ; New York (2001)
 42. Barabási, A.-L., Albert, R.: Emergence of scaling in random networks. *Science* **286**(5439), 509–512 (1999)
 43. Leskovec, J., Chakrabarti, D., Kleinberg, J.M., Faloutsos, C., Ghahramani, Z.: Kronecker graphs: An approach to modeling networks. *J. Mach. Learn. Res.* **11**, 985–1042 (2010)
 44. Seshadhri, C., Pinar, A., Kolda, T.G.: An in-depth analysis of stochastic Kronecker graphs. *J. ACM* **60**(2), 13–11332 (2013)
 45. Newman, M.E.J.: The structure and function of complex networks. *SIAM Rev.* **45**(2), 167–256 (2003)
 46. Watts, D.J., Strogatz, S.H.: Collective dynamics of small-world networks. *Nature* **393**(6684), 440 (1998)
 47. Krioukov, D., Papadopoulos, F., Kitsak, M., Vahdat, A., Boguñá, M.: Hyperbolic geometry of complex networks. *Phys. Rev. E* **82**, 036106 (2010)
 48. Aldecoa, R., Orsini, C., Krioukov, D.: Hyperbolic graph generator. *Comput. Phys. Commun.* **196**, 492–496 (2015)
 49. McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv:1802.03426* (2018)
 50. Saade, A., Krzakala, F., Zdeborová, L.: Spectral density of the non-backtracking operator on random graphs. *EPL (Europhys. Lett.)* **107**(5), 50005 (2014)
 51. Rubner, Y., Tomasi, C., Guibas, L.J.: A metric for distributions with applications to image databases. In: *Proc. of the 1998 Int'l Conf. on Computer Vision (ICCV)*, pp. 59–66 (1998)

52. Torres, L., Suarez-Serrato, P., Eliassi-Rad, T.: Graph distance from the topological view of non-backtracking cycles. arXiv:1807.09592 (July, 2018)
53. Mellor, A., Grusovin, A.: Graph comparison via the non-backtracking spectrum. arXiv:1812.05457 (Dec, 2018)
54. McAuley, J.J., Leskovec, J.: Learning to discover social circles in ego networks. In: Bartlett, P.L., Pereira, F.C.N., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems* 25, pp. 548–556 (2012)
55. Leskovec, J., Huttenlocher, D.P., Kleinberg, J.M.: Signed networks in social media. In: Mynatt, E.D., Schoner, D., Fitzpatrick, G., Hudson, S.E., Edwards, W.K., Rodden, T. (eds.) *Proc. of the 28th Int'l Conf. on Human Factors in Computing Systems, CHI 2010, Atlanta, Georgia, USA, April 10-15, 2010*, pp. 1361–1370 (2010)
56. Leskovec, J., Lang, K.J., Dasgupta, A., Mahoney, M.W.: Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics* **6**(1), 29–123 (2009)
57. Richardson, M., Agrawal, R., Domingos, P.M.: Trust management for the semantic web. In: Fensel, D., Sycara, K.P., Mylopoulos, J. (eds.) *The Semantic Web - Proc. of the 2nd Int'l Semantic Web Conference (ISWC)*. *Lecture Notes in Computer Science*, vol. 2870, pp. 351–368 (2003)
58. Leskovec, J., Kleinberg, J.M., Faloutsos, C.: Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)* **1**(1), 2 (2007)
59. Leskovec, J., Kleinberg, J.M., Faloutsos, C.: Graphs over time: densification laws, shrinking diameters and possible explanations. In: Grossman, R., Bayardo, R.J., Bennett, K.P. (eds.) *Proc. of the 11th ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 177–187 (2005)
60. Leskovec, J., Krevl, A.: SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data> (2014)
61. Pinar, A., Seshadhri, C., Vishal, V.: ESCAPE: efficiently counting all 5-vertex subgraphs. In: *Proc. of the 26th Int'l Conf. on World Wide Web (WWW)*, pp. 1431–1440 (2017)
62. Yaveroğlu, Ö.N., Milenković, T., Pržulj, N.: Proper evaluation of alignment-free network comparison methods. *Bioinformatics* **31**(16), 2697–2704 (2015)
63. von Luxburg, U.: A tutorial on spectral clustering. *Statistics and Computing* **17**(4), 395–416 (2007)
64. Rosenberg, A., Hirschberg, J.: V-measure: A conditional entropy-based external cluster evaluation measure. In: Eisner, J. (ed.) *Proc. of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pp. 410–420 (2007)
65. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: Krishnapuram, B., Shah, M., Smola, A.J., Aggarwal, C.C., Shen, D., Rastogi, R. (eds.) *Proc. of the 22nd ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 855–864 (2016)
66. Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: online learning of social representations. In: Macskassy, S.A., Perlich, C., Leskovec, J., Wang, W., Ghani, R. (eds.) *The 20th ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining, (KDD)*, pp. 701–710 (2014)
67. Wang, D., Cui, P., Zhu, W.: Structural deep network embedding. In: Krishnapuram, B., Shah, M., Smola, A.J., Aggarwal, C.C., Shen, D., Rastogi, R. (eds.) *Proc. of the 22nd ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining (KDD)*, pp. 1225–1234 (2016)
68. Cao, S., Lu, W., Xu, Q.: Deep neural networks for learning graph representations. In: Schuurmans, D., Wellman, M.P. (eds.) *Proc. of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1145–1152
69. Marks, R.: Enron Timeline (2008). <http://www.agsm.edu.au/bobm/teaching/BE/Enron/timeline.html> Accessed 2018-06-06
70. The Guardian: Timeline: Enron (2006). <https://www.theguardian.com/business/2006/jan/30/corporatefraud.enron> Accessed 2018-06-06
71. The New York Times: Timeline: A chronology of Enron Corp. (2006). <https://www.nytimes.com/2006/01/18/business/worldbusiness/timeline-a-chronology-of-enron-corp.html> Accessed 2018-06-06
72. Kolda, T.G., Pinar, A., Seshadhri, C.: Triadic measures on graphs: The power of wedge sampling. In: *Proc. of the 13th SIAM Int'l Conf. on Data Mining (SDM)*, pp. 10–18 (2013)
73. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* **15**(6), 1373–1396 (2003)