



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Linear Algebra and its Applications

journal homepage: www.elsevier.com/locate/laa

Ranking hubs and authorities using matrix functions

Michele Benzi^{a,*}, Ernesto Estrada^{b,2}, Christine Klymko^{c,3}

^a Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA

^b Department of Mathematics and Statistics and Institute of Complex Systems, University of Strathclyde, Glasgow G1 1XQ, UK

^c Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA

ARTICLE INFO

Article history:

Received 16 July 2012

Accepted 1 October 2012

Available online 6 December 2012

Submitted by R.A. Brualdi

AMS classification:

05C50

15A16

65F60

90B10

Keywords:

Hubs

Authorities

Centrality

Communicability

Matrix exponential

Directed networks

Digraphs

Bipartite graphs

HITS

Katz

PageRank

Gauss quadrature

ABSTRACT

The notions of subgraph centrality and communicability, based on the exponential of the adjacency matrix of the underlying graph, have been effectively used in the analysis of undirected networks. In this paper we propose an extension of these measures to directed networks, and we apply them to the problem of ranking hubs and authorities. The extension is achieved by *bipartization*, i.e., the directed network is mapped onto a bipartite undirected network with twice as many nodes in order to obtain a network with a symmetric adjacency matrix. We explicitly determine the exponential of this adjacency matrix in terms of the adjacency matrix of the original, directed network, and we give an interpretation of centrality and communicability in this new context, leading to a technique for ranking hubs and authorities. The matrix exponential method for computing hubs and authorities is compared to the well known HITS algorithm, both on small artificial examples and on more realistic real-world networks. A few other ranking algorithms are also discussed and compared with our technique. The use of Gaussian quadrature rules for calculating hub and authority scores is discussed.

© 2012 Elsevier Inc. All rights reserved.

* Corresponding author.

E-mail addresses: benzi@mathcs.emory.edu (M. Benzi), ernesto.estrada@strath.ac.uk (E. Estrada), cklymko@emory.edu (C. Klymko).

¹ The work of this author was supported by National Science Foundation Grant DMS 1115692.

² The work of this author was supported in part by the New Professor's Fund, University of Strathclyde, and by the project "Mathematics of Large Technological Evolving Networks (MOLTEN)" funded by the Engineering and Physical Sciences Research Council and by the Research Councils UK Digital Economy programme, Grant ref. EP/I016058/1.

³ The work of this author was supported in part by the Laney Graduate School of Arts and Sciences at Emory University and by the National Science Foundation.

1. Introduction

In recent years, the study of networks has become central to many disciplines [5,8,9,15,16,36–38]. Networks can be used to describe and analyze many different types of interactions, from those between people (social networks), to the flow of goods across an area (transportation networks), to links between websites (the WWW graph), and so forth. In general, a network is a set of objects (nodes) and the connections between them (edges). Often, research is focused on determining and describing important structural characteristics of a network or the interactions among its components.

One common question in network analysis is to determine the most “important” nodes (or edges) in the network, also called *node* or *vertex (edge) centrality*. The interpretation of what is meant by “important” can change from application to application. Due to this, many different measures of centrality have been developed. For an overview, see [8]. A closely related notion is that of *rank* of a node in a network. There exist a number of definitions and algorithms for computing rankings; see, e.g., [23,31,30,32,33,41] for up-to-date overviews.

The main notion of node centrality considered in this paper, *subgraph centrality*, was introduced by Estrada and Rodríguez-Velázquez in [20]. We refer readers to [20] for the motivation behind this notion and for its name; see also the review article [19], and the discussion in Section 4. The interpretation of centrality described in [19] applies mostly to undirected networks. However, many important real-world networks (the World Wide Web, the Internet, citation networks, food webs, certain social networks, etc.) are directed. One goal of this paper is to extend the notions of centrality and communicability described in [17,19] to directed networks, with an eye towards developing new ranking algorithms for, e.g., document collections, web pages, and so forth. We further compare our approach with some standard algorithms, such as HITS (see [29]) and a few others. Methods of quickly determining hub and authority rankings using Gauss-type quadrature rules are also discussed.

2. Basic notions

Here we briefly review some basic graph-theoretic notions; we refer to [13] for a comprehensive treatment. A *graph* $G = (V, E)$ is formed by a set of nodes (vertices) V and edges E formed by unordered pairs of vertices. Every network is naturally associated with a graph $G = (V, E)$ where $|V|$ is the number of nodes in the network and E is the collection of edges between objects, $E = \{(i, j) \mid \text{there is an edge between node } i \text{ and node } j\}$. The *degree* d_i of a vertex i is the number of edges incident to i .

A directed graph, or *digraph* $G = (V, E)$ is formed by a set of vertices V and edges E formed by ordered pairs of vertices. That is, $(i, j) \in E \not\Rightarrow (j, i) \in E$. In the case of digraphs, which model directed networks, there are two types of degree. The *in-degree* of node i is given by the number of edges which point to i . The *out-degree* is given by the number of edges pointing out from i .

A *walk* is a sequence of vertices v_1, v_2, \dots, v_k such that for $1 \leq i < k$, there is an edge between v_i and v_{i+1} (or a directed edge from v_i to v_{i+1} in the case of a digraph). Vertices and edges may be repeated. A walk is *closed* if $v_1 = v_k$. A *path* is a walk consisting only of distinct vertices.

A graph G is *connected* if every pair of vertices is linked by a path in G . A digraph is *strongly connected* if for any pair of vertices v_i and v_k there is a walk starting at v_i and ending at v_k . A digraph is *weakly connected* if the graph obtained by disregarding the orientation of its edges is connected. Unless otherwise specified, every digraph in this paper is simple (unweighted with no multiple edges or loops and connected). Note, however, that most of the techniques and results in the paper can be extended without difficulty to more general digraphs, in particular weighted ones.

The *adjacency matrix* of a graph is a matrix $A \in \mathbb{R}^{|V| \times |V|}$ defined in the following way:

$$A = (a_{ij}); \quad a_{ij} = \begin{cases} 1, & \text{if } (i, j) \text{ is an edge in } G, \\ 0, & \text{else.} \end{cases}$$

Under the conditions imposed on G , A has zeros on the diagonal. If G is an undirected graph, A will be a symmetric matrix and the eigenvalues will be real. In the case of digraphs, A is not symmetric and may have complex (non-real) eigenvalues.

3. Kleinberg’s HITS algorithm

Here we briefly recall the classical *Hypertext Induced Topics Search* (HITS) algorithm, first introduced by J. Kleinberg in [29]. This algorithm provides the motivation for the extension of subgraph centrality to directed graphs given in Section 5.

3.1. The basic iteration

The HITS algorithm is based on the idea that in the World Wide Web, and indeed in all document collections which can be represented by directed networks, there are two types of important nodes: *hubs* and *authorities*. Hubs are nodes which point to many nodes of the type considered important. Authorities are these important nodes. From this comes a circular definition: good hubs are those which point to many good authorities and good authorities are those pointed to by many good hubs.

Thus, the HITS ranking relies on an iterative method converging to a stationary solution. Each node i in the network is assigned two non-negative weights, an *authority weight* x_i and a *hub weight* y_i . To begin with, each x_i and y_i is given an arbitrary nonzero value. Then, the weights are updated in the following ways:

$$x_i^{(k)} = \sum_{j:(j,i) \in E} y_j^{(k-1)} \quad \text{and} \quad y_i^{(k)} = \sum_{j:(i,j) \in E} x_j^{(k)} \quad \text{for } k = 1, 2, 3, \dots \tag{3.1}$$

The weights are then normalized so that $\sum_j (x_j^{(k)})^2 = 1$ and $\sum_j (y_j^{(k)})^2 = 1$.

The above iterations occur sequentially and it can be shown that, under mild conditions, both sequences of vectors $\{x^{(k)}\}$ and $\{y^{(k)}\}$ converge as $k \rightarrow \infty$. In practice, the iterative process is continued until there is no significant change between consecutive iterates.

This iteration sequence shows the natural dependence relationship between hubs and authorities: if a node i points to many nodes with large x -values, it receives a large y -value and, if it is pointed to by many nodes with large y -values, it receives a large x -value.

In terms of matrices, Eq. (3.1) becomes: $x^{(k)} = A^T y^{(k-1)}$ and $y^{(k)} = Ax^{(k)}$, followed by normalization in the 2-norm. This iterative process can be expressed as

$$x^{(k)} = c_k A^T A x^{(k-1)} \quad \text{and} \quad y^{(k)} = c'_k A A^T y^{(k-1)}, \tag{3.2}$$

where c_k and c'_k are normalization factors. A typical choice for the initialization vectors $x^{(0)}, y^{(0)}$ would be the constant vector

$$x^{(0)} = y^{(0)} = [1/\sqrt{n}, 1/\sqrt{n}, \dots, 1/\sqrt{n}],$$

see [21]. Hence, HITS is just an iterative power method to compute the dominant eigenvector for $A^T A$ and for AA^T . The authority scores are determined by the entries of the dominant eigenvector of the matrix $A^T A$, which is called the *authority matrix* and the hub scores are determined by the entries of the dominant eigenvector of AA^T , called the *hub matrix*. Recall that the eigenvalues of both $A^T A$ and AA^T are the squares of the singular values of A . Also, the eigenvectors of $A^T A$ are the right singular vectors of A , and the eigenvectors of AA^T are the left singular vectors of A .

3.2. HITS reformulation

In a digraph the adjacency matrix A is generally nonsymmetric, however, the two matrices used in the HITS algorithm ($A^T A$ and AA^T) are symmetric. Note that, setting

$$\mathcal{A} = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix},$$

a symmetric matrix is obtained. Now,

$$\mathcal{A}^2 = \begin{pmatrix} AA^T & 0 \\ 0 & A^T A \end{pmatrix}; \quad \mathcal{A}^3 = \begin{pmatrix} 0 & AA^T A \\ A^T AA^T & 0 \end{pmatrix}.$$

In general,

$$\mathcal{A}^{2k} = \begin{pmatrix} (AA^T)^k & 0 \\ 0 & (A^T A)^k \end{pmatrix}; \quad \mathcal{A}^{2k+1} = \begin{pmatrix} 0 & A(A^T A)^k \\ (A^T A)^k A^T & 0 \end{pmatrix}.$$

Applying HITS to this matrix \mathcal{A} , $\mathcal{A}^T = \mathcal{A}$ so $\mathcal{A}^T \mathcal{A} = \mathcal{A} \mathcal{A}^T = \mathcal{A}^2$ and introducing the vector $u^{(k)} = \begin{pmatrix} y^{(k)} \\ x^{(k)} \end{pmatrix}$ for $k = 1, 2, 3, \dots$, Eq. (3.2) becomes

$$u^{(k)} = \mathcal{A}^2 u^{(k-1)} = \begin{pmatrix} AA^T & 0 \\ 0 & A^T A \end{pmatrix} u^{(k-1)}, \tag{3.3}$$

followed by normalization of the two vector components of $u^{(k)}$ so that each has 2-norm equal to 1. Now, if A is an $n \times n$ matrix, \mathcal{A} is $2n \times 2n$ and vector $u^{(k)}$ is in \mathbb{R}^{2n} . The first n entries of $u^{(k)}$ correspond to the hub rankings of the nodes, while the last n entries give the authority rankings. Under suitable assumptions (see the discussion in [32, Chapter 11.3]), as $k \rightarrow \infty$ the sequence $\{u^{(k)}\}$ converges to the dominant nonnegative eigenvector of \mathcal{A} , which yields the desired hub and authority rankings.

Hence, in HITS only information obtained from the dominant eigenvector of \mathcal{A} is used. It is natural to expect that taking into account spectral information corresponding to the remaining eigenvalues and eigenvectors of \mathcal{A} may lead to improved results.

Among the limitations of HITS, we mention the possible dependence of the rankings on the choice of the initial vectors $x^{(0)}, y^{(0)}$, see [21] for examples of this, and the fact that HITS hub/authority rankings tend to be “degree-biased”, i.e., they are strongly correlated with the out-/in-degrees of the corresponding nodes [14]. The latter property is in fact shared by most eigenvector-based rankings; for a discussion of this phenomenon in the case of scale-free graphs, see [35].

4. Subgraph centralities and communicabilities

In [19], the authors review several measures to rank the nodes in an undirected network A based on the use of matrix functions, such as the matrix exponential e^A . The *subgraph centrality* [20] of node i is given by $[e^A]_{ii}$ and the *communicability* [17] between nodes i and j ($i \neq j$) is given by $[e^A]_{ij}$. Nodes i corresponding to higher values of $[e^A]_{ii}$ are considered more important than nodes corresponding to lower values. Large values of $[e^A]_{ij}$ indicate that information flows more easily between nodes i and j than between pairs of nodes corresponding to lower values of the same quantity. The *Estrada index* of the graph is given by $\text{Tr}(e^A) = \sum_{i=1}^n [e^A]_{ii}$. This index, which provides a global characterization of a network, is analogous to the *partition function* in statistical mechanics and plays an important role in the study of networks at the *macroscopic* level: quantities such as the *natural connectivity*, the *total energy*, the *Helmholtz free energy* and the *entropy* of a network can all be expressed in terms of the Estrada index [18].

Consider the power series expansion of e^A ,

$$e^A = I + A + \frac{A^2}{2!} + \frac{A^3}{3!} + \dots + \frac{A^k}{k!} + \dots \tag{4.1}$$

From graph theory, it is well known that if A is the adjacency matrix of an undirected graph, $[A^k]_{ij} = [A^k]_{ji}$ counts the number of walks of length k between nodes i and j . Thus, the subgraph centrality of node i , $[e^A]_{ii}$, counts the total number of closed walks starting at node i , penalizing longer walks by scaling walks of length k by the factor $\frac{1}{k!}$. The communicability between nodes i and j , $[e^A]_{ij}$, counts the number of walks between nodes i and j , again scaling walks of length k by a factor of $\frac{1}{k!}$.

It is worth mentioning that normalization of the diagonal entries of e^A by $\text{Tr}(e^A)$ yields a probability distribution on the nodes of the network, which can be given the following interpretation: the i th diagonal entry of $e^A/\text{Tr}(e^A)$ is the probability of selecting any weighted self-returning (closed) walk that starts and ends at node i among all the weighted self-returning walks that start at any node and return to the same node. The weights used (factorial penalization) ensure that the shortest walks receive more weight than the longer ones: hence, the subgraph centrality of node i is proportional to the probability of finding a random walker walking “nearby” node i .

Although the matrix exponential is certainly well-defined for any matrix, whether symmetric or not, the interpretation of the notion of subgraph centrality for directed networks can be problematic. To see this, consider the directed path graph consisting of n nodes, with edge set $E = \{(1, 2), (2, 3), \dots, (n - 1, n)\}$ and adjacency matrix

$$A = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}. \tag{4.2}$$

The entries of e^A are given by

$$[e^A]_{ij} = \begin{cases} 1/(j - i)!, & \text{if } j \geq i, \\ 0, & \text{else.} \end{cases}$$

In particular, the diagonal entries of e^A are all equal to 1. Therefore, it is impossible to distinguish any of the nodes from the others on the basis of this centrality measure; yet, it is clear that the first and last node are rather special, and certainly more “peripheral” (less “central”) than the other nodes. Also, we note that the probabilistic interpretation given above for undirected graphs is no longer meaningful for this example. Part of the problem, of course, is that the path digraph contains no closed walks. In the next section we show one way to extend the notion of subgraph centrality to digraphs that is immune from such shortcomings, and correctly differentiates between nodes in the example above. (On the other hand, it is interesting to note that the interpretation of the off-diagonal entries of e^A in terms of communicabilities is straightforward for the directed path. All entries of e^A below the main diagonal are zero, reflecting the fact that information can only flow from a node to higher-numbered nodes. Also, the entries of e^A decay rapidly away from the main diagonal, reflecting the fact that the “ease” of communication between a node and a higher numbered one decreases rapidly with the distance.)

Another issue when extending the notions of subgraph centrality and communicability to directed graphs is that computational difficulties may arise. While the computations involved do not pose a problem for small networks, many real-world networks are large enough that directly computing the

exponential of the adjacency matrix is prohibitive. In [2], techniques for bounding and estimating individual entries of the matrix exponential using Gaussian quadrature rules are discussed; see also [6] and Section 9 below. The ability to find upper and lower bounds for the entries requires that the matrix be symmetric, thus these bounds cannot be directly computed using the adjacency matrix of a directed network. Again, these difficulties can be circumvented using the approach discussed in the next section.

5. An extension to digraphs

Although the techniques described in [2] cannot be directly applied to non-symmetric matrices, setting

$$\mathcal{A} = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} \tag{5.1}$$

produces a symmetric matrix \mathcal{A} and, thus, upper and lower bounds of individual entries of $e^{\mathcal{A}}$ can be computed. In Proposition 1 below we relate $e^{\mathcal{A}}$ to the underlying hub and authority structure of the original digraph. By B^\dagger we denote the Moore–Penrose generalized inverse of matrix B .

Proposition 1. *Let \mathcal{A} be as described in Eq. (5.1). Then,*

$$e^{\mathcal{A}} = \begin{pmatrix} \cosh(\sqrt{AA^T}) & A(\sqrt{A^T A})^\dagger \sinh(\sqrt{A^T A}) \\ \sinh(\sqrt{A^T A}) (\sqrt{A^T A})^\dagger A^T & \cosh(\sqrt{A^T A}) \end{pmatrix}.$$

Proof. Let $A = U\Sigma V^T$ be the SVD of the original (non-symmetric) adjacency matrix A . Then, \mathcal{A} can be decomposed as $\mathcal{A} = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix} \begin{pmatrix} U^T & 0 \\ 0 & V^T \end{pmatrix}$. Hence,

$$e^{\mathcal{A}} = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \exp \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix} \begin{pmatrix} U^T & 0 \\ 0 & V^T \end{pmatrix}. \tag{5.2}$$

Now,

$$\begin{aligned} \exp \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix} &= \cosh \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix} + \sinh \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix} \\ &= \begin{pmatrix} \cosh(\Sigma) & 0 \\ 0 & \cosh(\Sigma) \end{pmatrix} + \begin{pmatrix} 0 & \sinh(\Sigma) \\ \sinh(\Sigma) & 0 \end{pmatrix}. \end{aligned}$$

Thus,

$$\exp \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix} = \begin{pmatrix} \cosh(\Sigma) & \sinh(\Sigma) \\ \sinh(\Sigma) & \cosh(\Sigma) \end{pmatrix}. \tag{5.3}$$

Putting together Eqs. (5.2) and (5.3),

$$\begin{aligned}
 e^A &= \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} \cosh(\Sigma) & \sinh(\Sigma) \\ \sinh(\Sigma) & \cosh(\Sigma) \end{pmatrix} \begin{pmatrix} U^T & 0 \\ 0 & V^T \end{pmatrix} \\
 &= \begin{pmatrix} \cosh(\sqrt{AA^T}) & A(\sqrt{A^T A})^\dagger \sinh(\sqrt{A^T A}) \\ \sinh(\sqrt{A^T A})(\sqrt{A^T A})^\dagger A^T & \cosh(\sqrt{A^T A}) \end{pmatrix}.
 \end{aligned}$$

The identities involving the off-diagonal blocks can be easily checked using the SVD of A . \square

5.1. Interpretation of diagonal entries

In the context of undirected networks, the interpretation of the entries of the matrix exponential in terms of subgraph centralities and communicabilities is well-established, see e.g. [19]. In the case of directed networks and e^A , things are not as clear. The network behind A can be thought of as follows: take the vertices from the original network A and make two copies of them, V and V' . Then, undirected edges exist between the two sets based on the following rule: $E' = \{(i, j') \mid \text{there is a directed edge from } i \text{ to } j \text{ in the original network}\}$. This creates a bipartite graph with $2n$ nodes: $1, 2, \dots, n, n+1, n+2, \dots, 2n$. We denote by $V(A)$ this set of nodes. The use of bipartization to treat rectangular and structurally unsymmetric matrices is of course standard in numerical linear algebra.

In the undirected case, each node had only one role to play in the network: any information that came into the node could leave by any edge. In the directed case, there are two roles for each node: that of a hub and that of an authority. It is unlikely that a high ranking hub will also be a high ranking authority, but each node can still be seen as acting in both of these roles. In the network A , the two aspects of each node are separated. Nodes $1, 2, \dots, n$ in $V(A)$ represent the original nodes in their role as hubs and nodes $n + 1, n + 2, \dots, 2n$ in $V(A)$ represent the original nodes in their role as authorities.

Given a directed network, an *alternating walk* of length k , starting with an out-edge, from node v_1 to node v_{k+1} is a list of nodes v_1, v_2, \dots, v_{k+1} such that there exists edge (v_i, v_{i+1}) if i is odd and edge (v_{i+1}, v_i) if i is even:

$$v_1 \rightarrow v_2 \leftarrow v_3 \rightarrow \dots$$

An *alternating walk* of length k , starting with an in-edge, from node v_1 to node v_{k+1} in a directed network is a list of nodes v_1, v_2, \dots, v_{k+1} such that there exists edge (v_{i+1}, v_i) if i is odd and edge (v_i, v_{i+1}) if i is even:

$$v_1 \leftarrow v_2 \rightarrow v_3 \leftarrow \dots$$

From graph theory (see also [11]), it is known that $[AA^T A \dots]_{ij}$ (where there are k matrices being multiplied) counts the number of alternating walks of length k , starting with an out-edge, from node i to node j , whereas $[A^T AA^T \dots]_{ij}$ (where there are k matrices being multiplied) counts the number of alternating walks of length k , starting with an in-edge, from node i to node j . That is, $[(AA^T)^k]_{ij}$ and $[(A^T A)^k]_{ij}$ count the number of alternating walks of length $2k$.

In the original network A , if node i is a good hub, it will point to many good authorities, which will in turn be pointed at by many hubs. These hubs will also point to many authorities, which will again be pointed at by many other hubs. Thus, if i is a good hub, it will show up many times in the sets of hubs described above. That is, there should be many even length alternating walks, starting with an

out-edge, from node i to itself. Giving a walk of length $2k$ a weight of $\frac{1}{(2k)!}$, these walks can be counted using the (i, i) entry of the matrix

$$I + \frac{AA^T}{2!} + \frac{AA^TAA^T}{4!} + \dots + \frac{(AA^T)^k}{(2k)!} + \dots$$

Letting $A = U\Sigma V^T$ be the SVD of A , this becomes:

$$U \left(I + \frac{\Sigma^2}{2!} + \frac{\Sigma^4}{4!} + \dots + \frac{\Sigma^{2k}}{(2k)!} + \dots \right) U^T$$

$$= U \cosh(\Sigma) U^T = \cosh(\sqrt{AA^T}).$$

The *hub centrality* of node i (in the original network) is thus given by

$$[e^A]_{ii} = [\cosh(\sqrt{AA^T})]_{ii}.$$

This measures how well node i transmits information to the authoritative nodes in the network.

Similarly, if node i is a good authority, there will be many even length alternating walks, starting with an in-edge, from node i to itself. Giving a walk of length $2k$ a weight of $\frac{1}{(2k)!}$, these walks can be counted using the (i, i) entry of $\cosh(\sqrt{A^T A})$.

Hence, the *authority centrality* of node i is given by

$$[e^A]_{n+i, n+i} = [\cosh(\sqrt{A^T A})]_{ii}.$$

It measures how well node i receives information from the hubs in the network.

Note that the traces of the two diagonal blocks in e^A are identical, so each accounts for half of the Estrada index of the bipartite graph. Also, recalling the well-known fact that the eigenvalues of A are $\pm\sigma_i$ where σ_i denotes the singular values of A , we have

$$\text{Tr}(e^A) = \sum_{i=1}^n e^{\sigma_i} + \sum_{i=1}^n e^{-\sigma_i} = 2 \sum_{i=1}^n \cosh(\sigma_i),$$

an identity that can also be obtained directly from the expression for e^A given in Proposition 1.

Returning to the example of the directed path graph with adjacency matrix A given by (4.2), one finds that using the diagonal entries of e^A to rank the nodes gives node 1 as the least authoritative node, and node n as the one with the lowest hub ranking, with all the other nodes being tied. Thus we see that, while e^A fails to differentiate between the nodes of this graph, using e^A yields a very reasonable hub/authority ranking of the nodes.

5.2. Interpretation of off-diagonal entries

Although not used in the remainder of this paper, for the sake of completeness we give here an interpretation of the off-diagonal entries of e^A . As we will see, this interpretation is rather different from the one usually given for the off-diagonal entries of e^A , and provides information of a different nature on the structure of the underlying graph.

In discussing the off-diagonal entries of A , there are three blocks to consider. First, there are the off-diagonal entries of the upper-left block, $\cosh(\sqrt{AA^T})$, then there are the off-diagonal entries of the

lower-right block, $\cosh(\sqrt{A^T A})$. Finally, there is the off-diagonal block, $A(\sqrt{A^T A})^\dagger \sinh(\sqrt{A^T A})$ (the fourth block in e^A being its transpose).

From Section 5.1, $[e^A]_{ij} = [\cosh(\sqrt{AA^T})]_{ij}$, $1 \leq i, j \leq n$, counts the number of even length alternating walks, starting with an out-edge, from node i to node j , weighting walks of length $2k$ by a factor of $\frac{1}{(2k)!}$. When $i \neq j$, these entries measure how similar nodes i and j are as hubs. That is, if nodes i and j point to many of the same nodes, there will be many short even length alternating walks between them.

The *hub communicability* between nodes i and j , $1 \leq i, j \leq n$, is given by

$$[e^A]_{ij} = [\cosh(\sqrt{AA^T})]_{ij}$$

This measures how similar nodes i and j are in their roles as hubs. That is, a larger value of hub communicability between nodes i and j indicates that they point to many of the same authorities. In other words, they point to nodes which are authorities on the same subjects.

Similarly, $[e^A]_{n+i, n+j} = [\cosh(\sqrt{A^T A})]_{ij}$, $1 \leq i, j \leq n$, counts the number of even length alternating walks, starting with an in-edge, from node i to node j , also weighing walks of length $2k$ by a factor of $\frac{1}{(2k)!}$. When $i \neq j$, these entries measure how similar the two nodes are as authorities. If i and j are pointed at by many of the same hubs, there will be many short even length alternating walks between them.

The *authority communicability* between nodes i and j , $1 \leq i, j \leq n$, is given by

$$[e^A]_{i+n, j+n} = [\cosh(\sqrt{A^T A})]_{ij}$$

This measures how similar nodes i and j are in their roles as authorities. That is, a larger value of authority communicability between nodes i and j means that they are pointed to by many of the same hubs and, as such, are likely to contain information on the same subjects.

Let us now consider the off-diagonal blocks of \mathcal{A} . Here, $[\sinh(\sqrt{A^T A})]_{ij}$ counts the number of odd length alternating walks, starting with an out-edge, from node i to node j , weighing walks of length $2k + 1$ by $\frac{1}{(2k+1)!}$. This measures the communicability between node i as a hub and node j as an authority.

The *hub-authority communicability* between nodes i and j (that is, the communicability between node i as a hub and node j as an authority) is given by:

$$\begin{aligned} [e^A]_{i, n+j} &= [A(\sqrt{A^T A})^\dagger \sinh(\sqrt{A^T A})]_{ij} \\ &= [\sinh(\sqrt{A^T A})(\sqrt{A^T A})^\dagger A^T]_{ji} = [e^A]_{n+j, i}. \end{aligned}$$

A large hub-authority communicability between nodes i and j means that they are likely in the same “part” of the directed network: node i tends to point to nodes that contain information similar to that on which node j is an authority.

5.3. Relationship with HITS

As described in 3.2, the HITS ranking of nodes as hubs and authorities uses only information from the dominant eigenvector of \mathcal{A} . Here we show that when using the diagonal of e^A , we exploit information contained in all the eigenvectors of \mathcal{A} ; moreover, the HITS rankings can be regarded as an approximation of those given by the diagonal entries of e^A .

Assume the eigenvalues of \mathcal{A} can be ordered as $\lambda_1 > \lambda_2 \geq \lambda_3 \geq \dots \geq \lambda_{2n}$. Then, \mathcal{A} can be written as $\mathcal{A} = \sum_{i=1}^{2n} \lambda_i u_i u_i^T$ where u_1, u_2, \dots, u_{2n} are the normalized eigenvectors of \mathcal{A} . Taking the exponential of \mathcal{A} , we get:

$$e^A = \sum_{i=1}^{2n} e^{\lambda_i} u_i u_i^T = e^{\lambda_1} u_1 u_1^T + \sum_{i=2}^{2n} e^{\lambda_i} u_i u_i^T.$$

Now, the hub and authority rankings come from the diagonal entries of e^A :

$$\text{diag}(e^A) = e^{\lambda_1} \text{diag}(u_1 u_1^T) + \sum_{i=2}^{2n} e^{\lambda_i} \text{diag}(u_i u_i^T).$$

Rescaling the hub and authority scores by e^{λ_1} does not alter the rankings; hence, we can instead consider

$$\text{diag}(e^{-\lambda_1} e^A) = \text{diag}(e^{A-\lambda_1 I}) = \text{diag}(u_1 u_1^T) + \sum_{i=2}^{2n} e^{\lambda_i - \lambda_1} \text{diag}(u_i u_i^T).$$

Now, the diagonal entries of the rank-one matrix $u_1 u_1^T$ are just the squares of the (nonnegative) entries of the dominant eigenvector of A ; hence, the rankings provided by the first term in the expansion of e^A in the eigenbasis of A are precisely those given by HITS.

It is also clear that if $\lambda_1 \gg \lambda_2$, then the rankings provided by the diagonal entries of e^A are unlikely to differ much from those of HITS, since the weights $e^{\lambda_i - \lambda_1}$ will be tiny, for all $i = 2, \dots, 2n$. Conversely, if the gap between λ_1 and the rest of the spectrum is small ($\lambda_1 \approx \lambda_2$), then the contribution from the remaining eigenvectors, $\sum_{i=2}^{2n} e^{\lambda_i - \lambda_1} \text{diag}(u_i u_i^T)$, may be non-negligible relative to the first term and therefore the resulting rankings could differ significantly from those obtained using HITS. In Section 8 we will see examples of real networks illustrating both scenarios.

Summarizing, use of the matrix exponential for ranking hubs and authorities amounts to using the (squared) entries of *all* the eigenvectors of A , weighted by the exponential of the corresponding eigenvalues. Of course, in place of the exponential, a number of other functions could be used; see the discussion in the next section. Although using an exponential weighting scheme may at first sight appear to be arbitrary, its use can be rigorously justified; see the discussion in the next section, and [18] for a thorough treatment in the context of undirected graphs. As shown above, the HITS ranking scheme uses the leading term only, corresponding to the approximation $e^A \approx e^{\lambda_1} u_1 u_1^T$. Between these two extremes one could also use approximations of the form

$$e^A \approx \sum_{i=1}^k e^{\lambda_i} u_i u_i^T, \tag{5.4}$$

where $1 < k < n$; indeed, in most cases of practical interest a modest value of $k (\ll n)$ will be sufficient for a very good approximation, since the eigenvalues of A are often observed to decay rapidly from a certain index k onward. We return on this topic in Section 9.

6. Other ranking schemes

In this section we discuss a few other schemes that have been proposed in the literature, and compare them with the hub and authority centrality measures based on the exponential of A .

6.1. Resolvent-based measures

Besides the matrix exponential, another function that has been successfully used to define centrality and communicability measures for an undirected network is the matrix resolvent, which can be defined as

$$R(A; c) = (I - cA)^{-1} = I + cA + c^2 A^2 + \dots + c^k A^k + \dots,$$

with $0 < c < 1/\lambda_{\max}(A)$. This approach was pioneered early on by Katz [28], and variants thereof have since been used by numerous authors; see, e.g., [6,8,18,19,23,41]. Here A is the symmetric adjacency matrix of the undirected network. The condition on the parameter c ensures that $R(A; c)$ is well defined (i.e., that $I - cA$ is invertible and the geometric series converges to its inverse) and nonnegative; indeed, $I - cA$ will be a nonsingular M -matrix. It is hardly necessary to mention the close relationship existing between the resolvent and the exponential function, which can be expressed via the Laplace transform. For the adjacency matrix A of a bipartite graph given by (5.1), the resolvent is easily determined to be

$$R(A; c) = \begin{pmatrix} (I - c^2AA^T)^{-1} & cA(I - c^2A^T A)^{-1} \\ c(I - c^2A^T A)^{-1}A^T & (I - c^2A^T A)^{-1} \end{pmatrix}. \tag{6.1}$$

The condition on c can be expressed as $0 < c < 1/\sigma_1$, where $\sigma_1 = \|A\|_2$ denotes the largest singular value of A , the adjacency matrix of the undirected network. This ensures that the matrix in (6.1) is well-defined and nonnegative, with positive diagonal entries. The diagonal entries of $(I - c^2AA^T)^{-1}$ provide the hub scores, those of $(I - c^2A^T A)^{-1}$ the authority scores. A drawback of this approach is the need to select the parameter c , and the fact that different values of c may lead to different rankings. We have performed numerical experiments with this approach and we found that for certain values of c , particularly those close to the upper limit $1/\sigma_1$, the hub and authority rankings obtained with the resolvent function are not too different from those obtained with the matrix exponential. However, not surprisingly, as the value of c is reduced, one obtains hub and authority rankings that are strongly correlated with the out- and in-degree of the nodes, respectively.⁴ Overall, because the resolvent tends to weigh short walks more heavily than the exponential, and since longer walks contribute relatively little to the centrality scores, it is fair to say that the exponential is less “degree biased” than the resolvent function. Also, since the exponential rankings do not depend on a tuneable parameter, they provide unambiguous rankings.

We note that “Katz” authority and hub scores may also be obtained by considering the column and row sums of the (nonsymmetric) matrix resolvent $(I - cA)^{-1}$, where A is the adjacency matrix of the original digraph and $c > 0$ is again assumed to be small enough for the corresponding Neumann series to converge. Indeed, the row sums of $(I - cA)^{-1}$ count the number of (weighted) walks out of each node, while the column sums count the number of (weighted) walks into each node. Denoting by $\mathbf{1}$ the vector of all ones, hub and authority rankings can be obtained by solving the two linear systems

$$(I - cA)y = \mathbf{1} \quad \text{and} \quad (I - cA^T)x = \mathbf{1}, \tag{6.2}$$

respectively. Here the parameter c must satisfy $0 < c < 1/\rho(A)$, where $\rho(A)$ denotes the spectral radius of A . The results of numerical experiments comparing the Katz scores with those based on the exponential of A are given in Section 8. Here we observe that these Katz scores are also dependent on the choice of the parameter c , and similar considerations to those made for $(I - cA)^{-1}$ apply.

A natural analogue to this approach is the use of row and column sums of the exponential e^A to rank hubs and authorities. Some results obtained with this approach are discussed in Section 8. We note that this method is different from the *Exponentiated Inputs HITS Method* of [21]. The latter method is a modification to HITS which was developed in order to correct the issue of non unique results in certain networks. If the dominant eigenvalue of $A^T A$ (and, consequently, of AA^T) is not simple, then the corresponding eigenspace is multidimensional. This means that the choice of the initial vector affects the convergence of the HITS algorithm and different hub and authority vectors can be produced using different initial vectors. This can occur only when $A^T A$ is reducible, that is, when the original network is not strongly connected. In [21], Farahat *et al.* propose a modification to the HITS algorithm which amounts to replacing A and A^T with $e^A - I$ and $(e^A - I)^T$ in the HITS iteration. They note that, as long as the original network is weakly connected, the dominant eigenvalue of $(e^A - I)^T(e^A - I)$ is simple. Thus, HITS with this exponentiated input produces unique hub and authority rankings. However, a result of this replacement is that nodes with zero in-degree (or a low in-degree) are less important

⁴ Note that if c is taken too small, then the resolvent approaches the identity matrix and it becomes impossible to have meaningful rankings of the nodes.

in the calculation of authority scores than nodes with a high in-degree. When there are many nodes with zero in-degree or whose edges point to only a few other nodes, dropping these edges can greatly affect the HITS rankings. An obvious disadvantage of this algorithm is its cost, since it requires iterating with a matrix exponential and its transpose. It can be implemented using only matrix-vector products involving A and A^T by means of techniques, like Krylov subspace methods, for evaluating the action of a matrix function on a given vector; see, e.g., [26, Chapter 13]. This approach leads to a nested iteration scheme, with HITS as the outer iteration and the Krylov method as the inner one. Generally speaking, we have found HITS with exponentiated inputs to be less reliable and more expensive than the other methods considered in this paper. We refer to [3] for additional discussion and some examples.

6.2. PageRank and reverse PageRank

As is well known, the (now) classical PageRank algorithm provides a means of finding the authoritative nodes in a digraph. In PageRank, the importance of a node v is determined by the importance of the nodes pointing to v . In the most basic formulation, the rank of v is given by

$$r(v) = \sum_{u \in B_v} \frac{r(u)}{|u|} \tag{6.3}$$

where $B_v = \{u : \text{there is a directed link from } u \text{ to } v\}$ and $|u|$ is the out-degree of u . The ranks of the nodes are computed by initially setting, say, $r(v) = \frac{1}{n}$ (where n is the size of the network) and iteratively computing the rankings until convergence. This can also be written as

$$\pi_k^T = \pi_{k-1}^T P, \quad k = 1, 2, \dots \tag{6.4}$$

where π_k is the vector of node ranks at iteration k and P is the matrix given by

$$p_{ij} = \begin{cases} 1/|v_i|, & \text{if there is a directed edge from } v_i \text{ to } v_j, \\ 0, & \text{else.} \end{cases}$$

Here, P can be viewed as a probability transition matrix, where p_{ij} is the probability of traveling from node v_i to node v_j along an edge and the iterations can be understood as the evolution of a Markov chain modeling a random walk on the network.

However, for an arbitrary network, there is no guarantee that the PageRank algorithm will converge. If there are nodes with zero out-degree, P will not be stochastic. To correct this, the matrix \bar{P} is used, where each zero row of P is replaced with $\mathbf{1}^T/n$. Although this guarantees that the algorithm will converge, it does not guarantee the existence of a unique solution. Even with the augmentation, \bar{P} might still be a reducible matrix, corresponding to a reducible Markov chain. When this happens, there are rank sinks, i.e., nodes in which the random walk will become trapped and, subsequently, these nodes will receive a disproportionately high rank. However, if P were irreducible, there would be no rank sinks and the Perron–Frobenius theorem would guarantee that the Markov chain had a unique, positive stationary distribution.

The standard way to form a stochastic, irreducible PageRank matrix $\bar{\bar{P}}$ is to introduce the rank-1 matrix $E = \mathbf{1}\mathbf{1}^T/n$ and to consider instead of \bar{P} the convex combination

$$\bar{\bar{P}} = \alpha \bar{P} + (1 - \alpha)E, \tag{6.5}$$

where α is a constant with $0 < \alpha < 1$. The coefficient $1 - \alpha$ is a measure of the tendency of a person surfing the web to jump from one page to another without following links. In practice, a frequently recommended value is $\alpha = 0.85$. For a more comprehensive overview of the PageRank algorithm, see [23,27,31,32].

It was pointed out in [22] that applying PageRank to the digraph obtained by reversing the direction of the edges provides a natural way to rank the hubs; this is usually referred to as *Reverse PageRank*.

In other words, authority rankings are obtained by applying PageRank to the “Google” matrix derived from A , and hub rankings are obtained by the same process applied to A^T . Like HITS, PageRank and Reverse PageRank are eigenvector-based ranking algorithms that do not take into account information about the network contained in the non-dominant eigenvectors. As already mentioned, it has been argued [35] that eigenvector-based algorithms tend to be degree-biased. Furthermore, like the Katz-type algorithms, the rankings obtained depend on the choice of a tuneable damping parameter. While the success of PageRank in finding authoritative nodes is well known and very well documented, much less is known about the effectiveness of Reverse PageRank in identifying hubs; some references are [1, 10, 42, 43]. We present the results of a few numerical experiments with PageRank and Reverse PageRank in Section 8.

7. Examples

In this section and the next we illustrate the proposed method on some simple networks of small size, as well as on some larger data sets corresponding to real networks. We also compare our approach with HITS and other rankings schemes, including Katz, PageRank and Reverse PageRank.

7.1. Small digraphs

In this section we compare out and in-degree counts, HITS, and our proposed method to obtain hub and authority rankings in a few small digraphs. The purpose of this section is mostly pedagogical.

7.1.1. Example 1

Consider the small directed network in Fig. 7.1 (left panel). The adjacency matrix is given by

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

The corresponding bipartite graph is shown in Fig. 7.1 (right panel). If hubs and authorities are determined simply using in-degree and out-degree counts, the result is as follows:




Fig. 7.1. The original directed network from Example 1, with adjacency matrix A (left) and the bipartite network with adjacency matrix A (right).

Node	Out-degree	In-degree
1	2	1
2	2	3
3	2	2
4	1	1

Under this ranking, the hub ranking of the nodes is: {1, 2, 3 (tie); 4}. The authority ranking of the nodes is: {2; 3; 1, 4 (tie)}. We obtain somewhat different results using the HITS algorithm. The eigenvectors of AA^T and $A^T A$ corresponding to the largest eigenvalue $\lambda_{\max} \approx 3.9563$, which is simple, yield the following rankings for hubs and authorities:

Node	Hub rank	Authority rank
1	.3383	.0965
2	.1729	.4618
3	.2798	.2854
4	.2091	.1562

Here, the ranking for hubs is: {1; 3; 4; 2}. The ranking for authorities is: {2; 3; 4; 1}. Note that node 2, which was given a top hub score by looking just at the out-degrees, is judged by HITS as the node with the lowest hub score.

Using e^A as described above, the rankings for hub centralities and authority centralities are:

Node	Hub centrality = $[e^A]_{ii}$	Authority centrality = $[e^A]_{4+i,4+i}$
1	2.3319	1.5906
2	2.2289	3.0209
3	2.2812	2.2796
4	1.6414	1.5922

With this method, the hub ranking of the nodes is: {1; 3; 2; 4}. The authority ranking is: {2; 3; 4; 1}. On this example, our method produces the same authority ranking as HITS. The hub ranking, however, is slightly different: both methods identify node 1 as the one with the highest hub score, followed by node 3; however, our method assigns the lowest hub score to node 4 rather than node 2. This is arguably a more meaningful ranking.

7.1.2. Example 2

Consider the small directed network in Fig. 7.2 (left panel). The adjacency matrix is given by

$$A = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

The corresponding bipartite graph is shown in Fig. 7.2 (right panel). If hubs and authorities are determined only using in-degrees and out-degrees, the result is as follows:




Fig. 7.2. The original directed network from Example 2, with adjacency matrix A (left) and the bipartite network with adjacency matrix \mathcal{A} (right).

Node	Out-degree	In-degree
1	1	1
2	2	2
3	1	1
4	1	1

Under this criterion, the hub and authority rankings are both $\{2; 1, 3, 4 \text{ (tie)}\}$. While it is intuitive that node 2 should be given a high score (both as an authority and as a hub), just looking at the degrees does not allow one to distinguish the remaining nodes.

Consider now the use of HITS. The largest eigenvalue of AA^T (and $A^T A$) is $\lambda_{\max} = 2$ and it has multiplicity two. Thus, different starting vectors for the HITS algorithm may produce different rankings, as discussed in [21]. Starting from a constant authority vector $x^{(0)}$, as suggested in [29], produces the following scores:

Node	Hub rank	Authority rank
1	.0000	.3333
2	.5000	.3333
3	.2500	.0000
4	.2500	.3333

The ranking for hubs is: $\{2; 3, 4 \text{ (tie)}; 1\}$. The ranking for authorities is the following: $\{1, 2, 4 \text{ (tie)}; 3\}$. If the ranking is determined using $e^{\mathcal{A}}$ as described above, the resulting scores are:

Node	Hub centrality = $[e^{\mathcal{A}}]_{ii}$	Authority centrality = $[e^{\mathcal{A}}]_{4+i,4+i}$
1	1.5431	1.5891
2	2.1782	2.1782
3	1.5891	1.5431
4	1.5891	1.5891

With this method, the hub ranking of the nodes is the same as in HITS: $\{2; 3, 4 \text{ (tie)}; 1\}$. However, in the authority ranking, node 2 is the clear winner rather than being part of a three-way tie for first place: $\{2; 1, 4 \text{ (tie)}; 3\}$. In this example, the method based on the matrix exponential is able to identify a top authority node by making use of additional spectral information.




Fig. 7.3. The original directed network from Example 3, with adjacency matrix A (left) and the bipartite network with adjacency matrix A (right).

7.2. Example 3

Let G be the small directed network in Fig. 7.3. The adjacency matrix is given by

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}.$$

If hubs and authorities are determined using only in-degrees and out-degrees, the result is:

Node	Out-degree	In-degree
1	0	4
2	1	1
3	1	1
4	1	1
5	1	1
6	4	0

The hub ranking of the nodes using degrees is: {6; 2,3,4,5 (tie); 1}. The authority ranking is {1; 2,3,4,5 (tie); 6}.

If the HITS algorithm is used, the resulting rankings are similar, but not exactly the same. Starting with a constant authority vector $x^{(0)}$, the results are:

Node	Hub rank	Authority rank
1	.000	.200
2	.125	.200
3	.125	.200
4	.125	.200
5	.125	.200
6	.500	.000

The hub ranking of the nodes is: {6; 2, 3, 4, 5 (tie); 1}. The authority ranking is: {1,2,3,4,5 (tie); 6}. Here, HITS does not differentiate between node 1 and nodes 2, 3, 4, and 5 in terms of the authority score, even though node 1 has by far the highest in-degree. This appears as a failure of HITS, since it is intuitive that node 1 should be regarded as very authoritative.

When e^A is used to calculate the hub and authority scores, node 1 does get a higher authority ranking than all the other nodes:

Node	Hub centrality = $[e^A]_{ii}$	Authority centrality = $[e^A]_{6+i,6+i}$
1	1.0000	3.7622
2	1.6905	1.6905
3	1.6905	1.6905
4	1.6905	1.6905
5	1.6905	1.6905
6	3.7622	1.0000

Note that, if desired, the value 1 can be subtracted from these scores since it does not affect the relative ranking of the nodes. The hub ranking is {6; 2,3,4,5 (tie); 1}, and the authority ranking is: {1; 2,3,4,5 (tie); 6}.

8. Application to web graphs

Similarly to HITS, and in analogy to subgraph centrality for undirected networks, the rankings produced by the values on the diagonal of e^A can be used to rank websites as hubs and authorities in web searches (many other applications are of course also possible). Three of the data sets considered here are small web graphs consisting of web sites on various topics and can be found at [40] along with the website associated with each node; see also [7]. The experiments for this paper were run on the “Expanded” version of the data sets. Each data set is named after the corresponding topic.⁵ In addition, we include results for the *wb-cs-stanford* data set from the University of Florida Sparse Matrix Collection [12]. This digraph represents a subset of the Stanford University web. In this section, the hub and authority rankings obtained from e^A are compared with those from HITS, Katz (using (6.2) with $c = 1/(\rho(A) + 0.1)$), the row and column sums of the exponential e^A of the nonsymmetric matrix A , and PageRank/Reverse PageRank. For the latter we use the standard value $\alpha = 0.85$ for the damping parameter. All experiments are performed using Matlab Version 7.9.0 (R2009b) on a MacBook Pro running OS X Version 10.6.8, a 2.4 GHZ Intel Core i5 processor and 4 GB of RAM. For the purpose of these tests we use the built-in Matlab function `expm` to compute the matrix exponentials, and `backslash` to compute the Katz scores. Other approximations of e^A are discussed in Section 9.

8.1. Abortion data set

The *abortion* data set contains $n = 2293$ nodes and $m = 9644$ directed edges. The expanded matrix $\mathcal{A} = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$ has order $N = 2n = 4586$ and contains $2m = 19288$ nonzeros. The maximum eigenvalue of \mathcal{A} is $\lambda_N \approx 31.91$ and the second largest eigenvalue is $\lambda_{N-1} \approx 26.04$. In this matrix, the largest eigenvalue is fairly well-separated from the second largest so that one would expect the HITS rankings (which only use information from the dominant eigenpair of \mathcal{A}) to be reasonably close to the rankings from e^A (which use information from all of the eigenvalues and corresponding eigenvectors).

⁵ It should be noted, however, that in the node list for the adjacency matrix, the node labeling begins with 1 and in the list of websites associated with the nodes found at [40], node labeling begins at 0. Thus, node i in the adjacency matrix is associated with website $i - 1$.




Fig. 8.1. Plot of the eigenvalues of the expanded abortion matrix A .

Table 8.1

Top 10 hubs of the abortion web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A row sums and Reverse PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A rs	RPR
48	48	80	80	125
1021	1006	1431	1431	2184
1007	1007	1432	1432	79
1006	1021	1387	1426	81
1053	1053	1388	1425	48
1020	1020	1389	1415	1424
987	960	1397	1388	1447
990	968	1425	1389	78
985	969	1426	1397	134
989	970	1415	1387	1445

A plot of the eigenvalues of the expanded abortion data set matrix can be found in Fig. 8.1. Note the high multiplicity of the zero eigenvalue in this matrix, as well as in the adjacency matrices of the computational complexity and death penalty data sets. Also, quite a few of the nonzero eigenvalues are rather small. Due to this, the numerical rank of the matrix is very low, a property that can be exploited when estimating the entries of e^A using Lanczos-based methods; see Section 9 for further discussion on this.

The top 10 hubs and authorities of the abortion data set, as determined using the diagonal entries of e^A , HITS with constant initial vector, the row/column sums of $(I - cA)^{-1}$ (“Katz”), the row/column sums of e^A and Reverse PageRank/PageRank are shown in Tables 8.1 and 8.2. We observe that there is a good deal of agreement between the e^A rankings and the HITS ones: indeed, both methods identify the websites labeled 48, 1021, 1007, 1006, 1053, 1020 as the top 6 hubs, and both pick web site 48 as the top one. Also, there are 7 web sites identified by both methods as being among the top 10 authorities. The top authority identified by HITS is ranked third by e^A , and conversely the top authority identified by e^A is third in the HITS ranking. The Katz rankings and those based on e^A show considerable agreement with one another, but are very different from the HITS ones and from those based on e^A . Node 48, which is the top-ranked hub according to HITS and e^A , is now not even among the top 100. Conversely, node 80, which is ranked the top hub by Katz and e^A , is not in the top 100 nodes according to HITS or to e^A .

Table 8.2

Top 10 authorities of the abortion web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A column sums and PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A cs	PR
967	939	1430	1430	1609
958	958	1387	1387	1941
939	967	1425	1425	1948
962	961	1426	1426	1608
963	962	1429	1417	587
964	963	1396	1409	1610
961	964	1405	1429	2045
965	965	1406	1406	317
966	966	1409	1396	2191
587	1582	1417	1405	753




Fig. 8.2. Plot of the eigenvalues of the expanded computational complexity matrix \mathcal{A} .

This is not too surprising, since the metrics based on A and those based on \mathcal{A} are obtained by counting rather different types of graph walks. Finally, for this network Reverse PageRank and PageRank return rankings with almost no overlap with any of the other methods.

8.2. Computational complexity data set

The computational complexity data set contains $n = 884$ nodes and $m = 1616$ directed edges. The expanded matrix \mathcal{A} has order $N = 2n = 1768$ and contains $2m = 2232$ nonzeros. The maximum eigenvalue of \mathcal{A} is $\lambda_N \approx 10.93$ and the second largest eigenvalue is $\lambda_{N-1} \approx 9.86$. Here, the (relative) spectral gap between the first and the second eigenvalue is smaller than in the previous example; consequently, we expect the rankings produced using e^A and HITS to be less similar than for the abortion data set. A plot of the eigenvalues of the expanded computational complexity data set matrix can be found in Fig. 8.2.

The top 10 hubs and authorities of the computational complexity data set, determined by the various ranking methods, can be found in Tables 8.3 and 8.4. As expected, we see less agreement between HITS and the diagonals of e^A . Concerning the hubs, both methods agree that the web site labelled 57 is by far the most important hub on the topic of computational complexity. However, the method based on

Table 8.3

Top 10 hubs of the computational complexity web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A row sums and Reverse PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A rs	RPR
57	57	56	57	57
17	634	709	56	56
644	644	57	17	17
643	721	697	51	51
634	643	705	634	21
106	544	690	21	11
119	632	714	255	255
529	801	708	173	12
86	640	712	709	13
162	639	715	45	45

Table 8.4

Top 10 authorities of the computational complexity web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A column sums and PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A cs	PR
1	719	688	673	673
315	717	685	1	664
673	727	673	664	534
148	723	690	534	45
719	808	56	45	2
717	735	686	473	1
2	737	664	315	376
45	1	1	376	341
727	722	45	688	50
534	770	534	599	51

e^A identifies as the second most important hub the web site corresponding to node 17, which is ranked only 39th by HITS. The two methods agree on the next three hubs, but after that they return completely different results. The difference is even more pronounced for the authority rankings. The method based on e^A clearly identifies web site 1 as the most authoritative one, whereas HITS relegates this node to 8th place. The top authority according to HITS, web site 719, places 5th in the ranking obtained by e^A . The two methods agree on only two other web sites as being in the top 10 authorities (717 and 727). The Katz rankings and those based on e^A show little overlap for this data set, although node 57 is clearly considered an important hub by all measures. A natural question is how much these results are affected by the choice of the parameter c used to compute the Katz scores. We found experimentally that, in contrast to the situation for the other data sets, small changes in the value of c can significantly affect the Katz ranking for this particular data set. Changing the value of c to $c = 1/(\rho(A) + 0.3)$ results in hub and authority rankings that are much closer to those given by the column/row sums of e^A . The potential sensitivity to c is a clear drawback of the Katz-based approach compared to the methods based on the matrix exponential. Coming to (Reverse) PageRank, it is interesting to note that for this data set it provides rankings that are at least in partial agreement with some of the other measures, especially those based on e^A . Looking at the authority scores, we also notice a good degree of overlap among all methods, except HITS. Due to the small spectral gap, HITS is probably the least reliable of all ranking methods on this particular data set.

8.3. Death penalty data set

The *death penalty* data set contains $n = 1850$ and $m = 7363$ directed edges. The expanded matrix \mathcal{A} has order $N = 2n = 3700$ and contains $m = 14726$ nonzeros. The maximum eigenvalue of \mathcal{A} is $\lambda_N \approx 28.02$ and the second largest eigenvalue $\lambda_{N-1} \approx 17.68$. In this case, the largest and second largest eigenvalues are quite far apart, and the relative gap is larger than in the previous examples. A plot of the eigenvalues of the expanded death penalty matrix can be found in Fig. 8.3.

Due to the presence of a large spectral gap, much of the information used in forming the rankings of e^A is also used in the HITS ranking, and we expect the two methods to produce similar results; see




Fig. 8.3. Plot of the eigenvalues of the expanded death penalty matrix \mathcal{A} .

Table 8.5

Top 10 hubs of the death penalty web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A row sums and Reverse PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A rs	RPR
210	210	1632	1632	210
637	637	133	133	1632
413	413	1671	1671	70
1586	1586	552	552	95
552	552	1651	1651	135
462	462	1673	210	133
930	930	1328	1673	55
542	542	1653	1653	958
618	618	210	1328	1077
1275	1275	1709	1709	315

Section 5.3. Indeed, as shown in Table 8.5 (hubs) and Table 8.6 (authorities), in this case the top 10 rankings produced by the two methods are actually identical.

Looking at the Katz scores and those based on e^A , we see in this case a great deal of overlap between these two, but almost completely different rankings compared to HITS and e^A (although node 210 is clearly an important hub by any standard). Note that node 1632 is both the top hub and the top authority according to Katz and to e^A . PageRank and Reverse PageRank show a limited amount of overlap with the other measures; nevertheless, nodes 210 and 1632 are also found to be important hubs and nodes 1632, 1 and 4 are found to be authoritative, in agreement with some of the other measures.

8.4. Stanford web graph

The *wb-cs-stanford* data set from the University of Florida sparse matrix collection contains $n = 9914$ nodes and $m = 36854$ directed edges. The expanded matrix \mathcal{A} has order $N = 2n = 19828$ and contains $m = 73708$ nonzeros. The maximum eigenvalue of \mathcal{A} is $\lambda_N \approx 38.38$ and the second largest is $\lambda_{N-1} \approx 32.12$, hence there is a sizeable gap. Tables 8.7–8.8 report the results obtained with the various ranking schemes.

Table 8.6

Top 10 authorities of the death penalty web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A column sums and PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A cs	PR
4	4	1632	1632	993
1	1	1662	1662	667
6	6	1697	1697	3
7	7	1689	1689	736
10	10	1653	1653	735
16	16	1671	1671	1632
2	2	1675	1675	42
3	3	1684	1684	1
44	44	798	789	4
27	27	1652	1654	1212

Table 8.7

Top 10 hubs of the wb-cs-stanford web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A row sums and Reverse PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A rs	RPR
6562	6562	6562	6562	251
6838	6838	6837	6837	252
6840	6837	6838	6838	253
6837	6839	6839	6839	254
6839	6840	6840	6840	271
6616	6616	6669	6669	2240
6765	6615	6668	6668	2241
6615	6765	6670	6670	2242
6669	6669	6616	6616	2243
6731	6731	6615	6615	348

Table 8.8

Top 10 authorities of the wb-cs-stanford web graph, ranked using $[e^A]_{ii}$, HITS, Katz, e^A column sums and PageRank with $\alpha = 0.85$.

$[e^A]_{ii}$	HITS	Katz	e^A cs	PR
6837	6837	6837	6837	2264
6840	6839	6839	6839	8226
6839	6840	6840	6840	8059
6838	6838	6838	6838	8057
6617	6617	6573	6573	4485
6615	6615	6574	6575	5707
6766	6614	6575	6576	8225
6764	6616	6576	6577	6837
6616	6764	6577	6578	6839
6614	6766	6578	6579	6840

The first thing to observe is the remarkable agreement between the HITS, e^A , Katz, and e^A rankings of both hubs and authorities. This is in stark contrast with the results for the previous three data sets. Moreover, many of the nodes that are ranked highly as hubs are also ranked highly as authorities. A plausible explanation of these observations is that the adjacency matrix A for this digraph is much closer to being symmetric than in the other cases. Indeed, the percentage of “bidirectional” edges in the wb-cs-stanford graph is 47.63%; the corresponding percentages for the abortion, computational complexity and death penalty graphs are just 2.72%, 2.97% and 4.02%, respectively.

Interestingly, the (Reverse) PageRank results are now drastically different from the ones provided by all the other measures in nearly all cases. The only (partial) exception is that PageRank finds nodes 6837, 6839 and 6840 to be among the top 10 authorities; these three nodes are identified as the three most authoritative ones by the remaining methods.

9. Approximating the matrix exponential

Several approaches are available for computing the matrix exponential [26]. A commonly used scheme is the one based on Padé approximation combined with the scaling and squaring method [25,26], implemented in Matlab by the `expm` function. For an $n \times n$ matrix, this method requires $O(n^2)$ storage and $O(n^3)$ arithmetic operations; any sparsity in A , if present, is not exploited in currently available implementations. Evaluation of the matrix exponential based on diagonalization also requires $O(n^2)$ storage and $O(n^3)$ operations. Furthermore, these methods cannot be easily adapted to the case where only selected entries (e.g., the diagonal ones) of the matrix exponential are of interest.

For the purpose of ranking hubs and authorities in a directed network, only the main diagonal of e^A is required. This can be done without having to compute *all* the entries in e^A . If some of the off-diagonal entries (communicabilities) are desired, for example those between the highest ranked hubs and/or authorities, it is also possible to compute them without having to compute the whole matrix e^A , which would be prohibitive even for a moderately large network. We further emphasize that in most applications one is not so much interested in computing an exact ranking of *all* the nodes in a digraph, but only in identifying the top k ranked nodes, where the integer k is small compared to n (for example, $k = 10$ or $k = 20$). It is highly desirable to develop methods that are capable of quickly identifying the top k hubs/authorities without having to compute accurate hub/authority scores for each node.

Efficient, accurate methods for estimating (or, in some cases, bounding) arbitrary entries in a matrix function $f(A)$ have been developed by Golub, Meurant and collaborators (see [24] and references therein) and first applied to problems of network analysis by Benzi and Boito in [2]; see also [6]. Here we limit ourselves to a brief description of these methods, referring the reader to [2] and [24] for further details. Let A be a real, symmetric, $n \times n$ matrix and let f be a function defined on the spectrum of A . Consider the eigendecompositions $A = Q\Lambda Q^T$ and $f(A) = Qf(\Lambda)Q^T$, where $Q = [\phi_1, \dots, \phi_n]$ and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$; here we assume that the eigenvalues of A are ordered as $\lambda_1 \leq \dots \leq \lambda_n$. For given vectors u and v we have

$$u^T f(A)v = u^T Qf(\Lambda)Q^T v = w^T f(\Lambda)z = \sum_{k=1}^n f(\lambda_k)w_k z_k, \tag{9.1}$$

where $w = Q^T u = (w_k)$ and $z = Q^T v = (z_k)$. In particular, for $f(A) = e^A$ we obtain

$$u^T e^A v = \sum_{k=1}^n e^{\lambda_k} w_k z_k. \tag{9.2}$$

Choosing $u = v = e_i$ (the vector with the i th entry equal to 1 and all the remaining ones equal to 0) we obtain an expression for the subgraph centrality of node i :

$$SC(i) := \sum_{k=1}^n e^{\lambda_k} \phi_{k,i}^2,$$

where $\phi_{k,i}$ denotes the i th component of vector ϕ_k . Likewise, choosing $u = e_i$ and $v = e_j$ we obtain the following expression for the communicability between node i and node j :

$$C(i, j) := \sum_{k=1}^n e^{\lambda_k} \phi_{k,i} \phi_{k,j}.$$

Analogous expressions hold for other matrix functions, such as the resolvent.

Hence, the problem is reduced to evaluating bilinear expressions of the form $u^T f(A)v$. Such bilinear forms can be thought of as Riemann–Stieltjes integrals with respect to a (signed) spectral measure:

$$u^T f(A)v = \int_a^b f(\lambda) d\mu(\lambda), \quad \mu(\lambda) = \begin{cases} 0, & \text{if } \lambda < a = \lambda_1, \\ \sum_{k=1}^i w_k z_k, & \text{if } \lambda_i \leq \lambda < \lambda_{i+1}, \\ \sum_{k=1}^n w_k z_k, & \text{if } b = \lambda_n \leq \lambda. \end{cases}$$

This integral can be approximated by means of a Gauss-type quadrature rule:

$$\int_a^b f(\lambda) d\mu(\lambda) = \sum_{j=1}^p c_j f(t_j) + \sum_{k=1}^q v_k f(\tau_k) + R[f], \tag{9.3}$$

where $R[f]$ denotes the error. Here the nodes $\{t_j\}_{j=1}^p$ and the weights $\{c_j\}_{j=1}^p$ are unknown, whereas the nodes $\{\tau_k\}_{k=1}^q$ are prescribed. We have

- $q = 0$ for the Gauss rule,
- $q = 1$, $\tau_1 = a$ or $\tau_1 = b$ for the Gauss–Radau rule,
- $q = 2$, $\tau_1 = a$ and $\tau_2 = b$ for the Gauss–Lobatto rule.

For certain matrix functions, including the exponential and the resolvent, these quadrature rules can be used to obtain lower and upper bounds on the quantities of interest; prescribing additional quadrature nodes leads to tighter and tighter bounds, which (in exact arithmetic) converge monotonically to the true values [24]. The evaluation of these quadrature rules is mathematically equivalent to the computation of orthogonal polynomials via a three-term recurrence, or, equivalently, to the computation of entries and spectral information of a certain tridiagonal matrix via the Lanczos algorithm. Here we briefly recall how this can be done for the case of the Gauss quadrature rule, when we wish to estimate the i th diagonal entry of $f(A)$. It follows from (9.3) that the quantity of interest has the form $\sum_{j=1}^p c_j f(t_j)$. This can be computed from the relation (Theorem 3.4 in [24]):

$$\sum_{j=1}^p c_j f(t_j) = e_1^T f(J_p) e_1,$$

where

$$J_p = \begin{pmatrix} \omega_1 & \gamma_1 & & & \\ \gamma_1 & \omega_2 & \gamma_2 & & \\ & \ddots & \ddots & \ddots & \\ & & \gamma_{p-2} & \omega_{p-1} & \gamma_{p-1} \\ & & & \gamma_{p-1} & \omega_p \end{pmatrix}$$

is a tridiagonal matrix whose eigenvalues are the Gauss nodes, whereas the Gauss weights are given by the squares of the first entries of the normalized eigenvectors of J_p . The entries of J_p are computed using the Lanczos algorithm with starting vectors $x_{-1} = 0$ and $x_0 = e_i$. Note that it is not required to compute all the components of the eigenvectors of J_p if one uses the Golub–Welsch QR algorithm; see [24].

For small p (i.e., for a small number of Lanczos steps), computing the $(1, 1)$ entry of $f(J_p)$ is inexpensive. The main cost in estimating one entry of $f(A)$ with this approach is associated with the sparse matrix-vector multiplies in the Lanczos algorithm applied to the adjacency matrix A . If only a small, fixed number of iterations are performed for each diagonal element of $f(A)$, as is usually the case, the computational cost (per node) is at most $O(n)$ for a sparse graph, resulting in a total cost of $O(n^2)$ for computing the subgraph centrality of every node in the network. If only $k < n$ subgraph centralities

Table 9.1

The number of iterations necessary for the top 10 hubs or authorities to be determined (not necessarily in the correct order).

Dataset	Hub (lower bound)	Hub (upper bound)
Abortion	>40	>40
Comp. complex.3	3	3
Death penalty	5	3
Stanford	8	8
	Authority (lower bound)	Authority (upper bound)
Abortion	2	2
Comp. complex.4	2	5
Death penalty	4	2
Stanford	7	8

are wanted, with k independent of n , then the overall cost of the computation will be $O(n)$ provided that sparsity is carefully exploited in the Lanczos algorithm and that only a small number p of iterations (independent of n) is carried out. Note, however, that depending on the connectivity characteristics of the network under consideration, the prefactor in the $O(n)$ estimate could be large. The algorithm can be implemented so that the storage requirements are $O(n)$ for a sparse network—that is, a network in which the total number of links grows linearly in the number n of nodes.

When applying the approach based on Gauss quadrature rules to the $2n \times 2n$ matrix \mathcal{A} , only matrix-vector products with A and its transpose are required, just like in the HITS algorithm. If only the hub scores are wanted, it is also possible to apply the described techniques to the symmetric matrix AA^T using the function $f(\lambda) = \cosh(\sqrt{\lambda})$; the same applies if only the authority scores are wanted, working this time with the matrix $A^T A$. The problem with this approach is that only estimates (rather than increasingly accurate lower and upper bound) can be obtained, due to the fact that the function $f(\lambda) = \cosh(\sqrt{\lambda})$ is not strictly completely monotonic on the positive real axis. We refer to [4] for details. In our experiments we always work with the matrix \mathcal{A} , since we are interested in computing both hub and authority scores.

9.1. Test results

Accurate evaluation of *all* the diagonal entries of e^{-A} using quadrature rules may be too expensive for truly large-scale graphs. In most applications, fortunately, it is not necessary to rank all the nodes in the network: only the top few hubs and authorities are likely to be of interest. When using quadrature rules, the number of quadrature nodes (Lanczos iterations) required to correctly rank the nodes as hubs or authorities varies and depends on both the eigenvalues of e^{-A} and how close the diagonal entries are in value. If the rankings of the nodes are very close, it can take many iterations for the ordering to be exactly determined. However, since estimates for diagonal entries are calculated individually, once the top 10 (say) nodes have been identified, additional iterations can be performed only on these nodes in order to determine their exact ranking.

Our approach exploits the monotonicity of the Gauss-Radau bounds: as soon as the lower bound for node i is above the upper bounds for other nodes, we know that node i will be ranked higher than those other nodes. This observation leads to a simple algorithm for identifying the top- k nodes. The number of Lanczos iterations per node necessary to identify the top $k = 10$ hubs and authorities, using Gauss-Radau lower and upper bounds, for the four data sets from Section 8 is given in Table 9.1. Our implementation is based on Meurant’s Matlab code [34]. From the table it can be seen that, in most cases, only 2-5 iterations per node are needed. An exception is the determination of the top 10 hubs of the abortion data set, for which the number of iterations is large (> 40). This is due to a cluster of nodes (nodes 960 and 968-990) that have nearly identical hub rankings. These nodes’ scores agree to 15 significant digits. However, for most applications, if a subset of nodes are so closely ranked, their exact ordering may not be so important. Table 9.2 reports the number of Lanczos iterations needed for the top $k = 10$ hubs and authorities to be ranked at least in the top 30. Here, the number of iterations per node needed is never more than 7. The total cost is thus $O(n)$ Lanczos iterations, again leading to an $O(n^2)$ overall complexity. Various enhancements can be used to reduce costs, including the use of

Table 9.2

The number of iterations necessary for the top 10 hubs or authorities to be ranked in the top 30.

Dataset	Hub (lower bound)	Hub (upper bound)
Abortion	5	4
Comp. complex.2		2
Death penalty	2	2
Stanford	7	4
	Authority (lower bound)	Authority (upper bound)
Abortion	2	2
Comp. complex.4		2
Death penalty	2	2
Stanford	2	4

sparse-sparse mat-vecs in the Lanczos iteration, and the exclusion of nodes with zero out-degree (for hub computations) and zero in-degree (for authority computations) from the top- k calculations. It is also safe to assume that in most cases of interest, one can also exclude nodes with in- and out-degree 1 from the computations, leading to further savings.

10. Conclusions and outlook

In this paper we have presented a new approach to ranking hubs and authorities in directed networks using functions of matrices. Bipartization is used to transform the original directed network into an undirected one with twice the number of nodes. The adjacency matrix of the bipartite graph is symmetric, and this allows the use of subgraph centrality (and communicability) measures for undirected networks. We showed that the diagonal entries of the matrix exponential provide hub and authority rankings, and we gave an interpretation for the off-diagonal entries (communicabilities). Unlike HITS, the results are independent of any starting vectors; and unlike the Katz-based ranking schemes, there is no dependency on an arbitrary parameter.

Several examples, both synthetic and corresponding to real data sets, have been used to demonstrate the effectiveness of the proposed ranking algorithms relative to HITS and to other ranking schemes based on the matrix resolvent and on the exponential of the adjacency matrix of the original digraph. Our experiments indicate that our method results in rankings that are frequently different from those computed by HITS, at least in the absence of large gaps between the dominant singular value of the adjacency matrix A and the remaining ones. This is to be expected, since our method uses information from all the singular spectrum of the network, not just the dominant left and right singular pairs.

As usual in this field, there is no simple way to compare different ranking schemes, and therefore it is impossible to state with certainty that a ranking scheme will give “better” results than a different scheme in practice. It is, however, certainly the case that the method based on the exponential of A takes into account more spectral information than HITS does; moreover, the rankings so obtained are unambiguous, in that they do not depend on an the choice of an initial guess or on a tuneable parameter. As we saw, the latter is a weak spot of Katz-like methods, and a similar case can be made for PageRank and Reverse PageRank.

Compared to HITS, the new technique has a higher computational cost. We showed how Gaussian quadrature rules can be used to quickly identify the top ranked hubs and authorities for networks involving thousands of nodes. We note that such schemes require a symmetric input matrix and are not readily applicable to nonsymmetric matrices, since in this case one can only hope for estimates instead of lower and upper bounds.

Future work should include a more efficient implementation and tests on larger networks. It is likely that the proposed approach based on Gaussian quadrature will prove to be too expensive for truly large-scale networks with millions of nodes. We hope to explore techniques similar to those presented in [6, 39] in order to extend our methodology to truly large-scale networks. Another relevant question is the study of the rate of convergence of the Lanczos algorithm for estimating bilinear forms associated with adjacency matrices of graphs of different types.

Acknowledgments

We are indebted to Tammy Kolda (Sandia National Laboratories), Stratis Gallopoulos (University of Patras) and David Gleich (Purdue University) for valuable suggestions.

References

- [1] Z. Bar-Yossef, L.-T. Mashiach, Local approximation of PageRank and Reverse PageRank, in: Proceedings CKIM'08, October 26–30, 2008, Napa Valley, California.
- [2] M. Benzi, P. Boito, Quadrature rule-based bounds for functions of adjacency matrices, *Linear Algebra Appl.* 433 (2010) 637–652.
- [3] M. Benzi, E. Estrada, C. Klymko, Ranking Hubs and Authorities Using Matrix Functions, in: Math/CS Technical Report TR-2012-003, Emory University. Available at <http://www.mathcs.emory.edu/publications.php#reports>.
- [4] M. Benzi, G.H. Golub, Bounds on the entries of matrix functions with applications to preconditioning, *BIT* 39 (1999) 417–438.
- [5] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, D.-U. Hwang, Complex networks: structure and dynamics, *Phys. Rep.* 424 (2006) 175–308.
- [6] F. Bonchi, P. Esfandiari, D.F. Gleich, C. Greif, L.V.S. Lakshmanan, Fast matrix computations for pair-wise and column-wise compute times and Katz scores, *Internet Math.* 8 (2012) 73–112.
- [7] A. Borodin, G.O. Roberts, J.S. Rosenthal, P. Tsaparas, Finding authorities and hubs from link structures on the World Wide Web, in: Proceedings WWW10 (2001), ACM 1-58113-348-0/01/0005.
- [8] U. Brandes, T. Erlebach (Eds.), *Network Analysis: Methodological Foundations*, Lecture Notes in Computer Science, vol. 3418, Springer, New York, 2005.
- [9] G. Caldarelli, *Scale-free Networks*, Oxford University Press, Oxford, UK, 2007.
- [10] M.C. Crofoot, D.I. Rubenstein, A.S. Maiya, T.Y. Berger-Wolf, Aggression, grooming and group-level cooperation in white-faced capuchins (*Cebus capucinus*): Insights from social networks, *Amer. J. Primatol.* 73 (2011) 821–833.
- [11] J.J. Crofts, E. Estrada, D.J. Higham, A. Taylor, Mapping directed networks, *Electron. Trans. Numer. Anal.* 37 (2010) 337–350.
- [12] T.A. Davis, Y. Hu, The University of Florida Sparse Matrix Collection, *ACM Trans. Math. Soft.* 38 (1) (2011), Article 1.
- [13] R. Diestel, *Graph Theory*, Springer-Verlag, Berlin, 2000.
- [14] C.H.Q. Ding, H. Zha, X. He, P. Husbands, H.D. Simon, Link analysis: hubs and authorities on the World Wide Web, *SIAM Rev.* 46 (2004) 256–268.
- [15] E. Estrada, *The Structure of Complex Networks*, Oxford University Press, Oxford, UK, 2011.
- [16] E. Estrada, M. Fox, D. Higham, G.-L. Oppo (Eds.), *Network Science, Complexity in Nature and Technology*, Springer, New York, 2010.
- [17] E. Estrada, N. Hatano, Communicability in complex networks, *Phys. Rev. E* 77 (2008), article 036111.
- [18] E. Estrada, N. Hatano, M. Benzi, The physics of communicability in complex networks, *Phys. Rep.* 514 (2012) 89–119.
- [19] E. Estrada, D.J. Higham, Network properties revealed through matrix functions, *SIAM Rev.* 52 (2010) 671–696.
- [20] E. Estrada, J.A. Rodríguez-Velázquez, Subgraph centrality in complex networks, *Phys. Rev. E* 71 (2005) 056103.
- [21] A. Farahat, T. Lofaro, J.C. Miller, G. Rae, L.A. Ward, Authority rankings from HITS, PageRank, and SALSA: existence, uniqueness, and effect of initialization, *SIAM J. Sci. Comput.* 27 (2006) 1181–1201.
- [22] D. Fogaras, Where to start browsing the Web?, in: T. Böhme, G. Heyer, H. Unger (Eds.), Proceedings IICS 2003, Lecture Notes in Computer Science, 2877 (2003) 65–79.
- [23] M. Franceschet, PageRank: Standing on the shoulders of giants, *Commun. ACM* 54 (2011) 92–101.
- [24] G.H. Golub, G. Meurant, *Matrices, Moments and Quadrature with Applications*, Princeton University Press, Princeton, NJ, 2010.
- [25] N.J. Higham, The scaling and squaring method for the matrix exponential revisited, *SIAM J. Matrix Anal. Appl.* 26 (2006) 1179–1193.
- [26] N.J. Higham, *Functions of Matrices. Theory and Computation*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2008.
- [27] S. Kamvar, *Numerical Algorithms for Personalized Search in Self-Organizing Information Networks*, Princeton University Press, Princeton and Oxford, 2010.
- [28] L. Katz, A new status index derived from sociometric data analysis, *Psychometrika* 18 (1953) 39–43.
- [29] J. Kleinberg, Authoritative sources in a hyperlinked environment, *J. ACM* 46 (1999) 604–632.
- [30] G. Kollias, E. Gallopoulos, Functional Rankings with Multidamping: Generalizing PageRank with Inhomogeneous Matrix Products, Tech. Rep. TR HPCLAB-SCG 01/09-11, University of Patras, Greece, 2011.
- [31] A.N. Langville, C.D. Meyer, A survey of eigenvector methods for Web information retrieval, *SIAM Rev.* 47 (2005) 135–161.
- [32] A.N. Langville, C.D. Meyer, *Google's PageRank and Beyond: The Science of Search Engine Rankings*, Princeton University Press, Princeton, NJ, 2006.
- [33] A.N. Langville, C.D. Meyer, *Who's #1? The Science of Rating and Ranking*, Princeton University Press, Princeton, NJ, 2012.
- [34] G. Meurant, MMQ toolbox for MATLAB, <http://pagesperso-orange.fr/gerard.meurant/>.
- [35] M. Mihail, C. Papadimitriou, On the eigenvalue power law, in: J.D.P. Rolim, S. Vadhan (Eds.), Proceedings of RANDOM 2002, Lectures Notes in Computer Science, vol. 2483, 2002, pp. 254–262.
- [36] M.E.J. Newman, The structure and function of complex networks, *SIAM Rev.* 45 (2003) 167–256.
- [37] M.E.J. Newman, *Networks: An Introduction*, Cambridge University Press, Cambridge, UK, 2010.
- [38] M.E.J. Newman, A.-L. Barabasi, D.J. Watts, *The Structure and Dynamics of Networks*, Princeton University Press, Princeton, NJ, 2003.
- [39] B. Savas, I. Dhillon, Clustered low rank approximation of graphs in information science applications, in: Proceedings of the 2011 SIAM Conference on Data Mining, April 2011.
- [40] P. Tsaparas, Datasets for Experiments on Link Analysis Ranking Algorithms. Available at <http://www.cs.toronto.edu/~tsap/experiments/datasets/index.html> (last accessed June 2012).

- [41] S. Vigna, Spectral Ranking, arXiv:0912.0238v9 [cs.IR], 2 March 2011.
- [42] G. Wu, J. Li, SWRank: an approach for ranking semantic web reversely and consistently, Third International Conference on Semantics, Knowledge and Grid, IEEE, 2007, doi: 10.1109/SKG.2007.81.
- [43] G. Wu, J. Li, L. Feng, K. Wang, Identifying potentially important concepts and relations in an ontology, in: A. Sheth et al. (Eds.), Proceedings of ISWC 2008, Lecture Notes in Computer Science, vol. 5318, 2008, pp. 33–49.