

Ranking in evolving complex networks

Hao Liao ^a, Manuel Sebastian Mariani ^{b,a,*}, Matúš Medo ^{c,d,b,*}, Yi-Cheng Zhang ^b,
Ming-Yang Zhou ^a

^a National Engineering Laboratory for Big Data System Computing Technology, Guangdong Province Key Laboratory of Popular High Performance Computers, College of Computer Science and Software Engineering, Shenzhen University, Shenzhen 518060, PR China

^b Department of Physics, University of Fribourg, 1700 Fribourg, Switzerland

^c Institute of Fundamental and Frontier Sciences, University of Electronic Science and Technology of China, Chengdu 610054, PR China

^d Department of Radiation Oncology, Inselspital, Bern University Hospital and University of Bern, 3010 Bern, Switzerland

Keywords:

Complex networks
Ranking
Centrality metrics
Temporal networks
Recommendation
Network science

Complex networks have emerged as a simple yet powerful framework to represent and analyze a wide range of complex systems. The problem of ranking the nodes and the edges in complex networks is critical for a broad range of real-world problems because it affects how we access online information and products, how success and talent are evaluated in human activities, and how scarce resources are allocated by companies and policymakers, among others. This calls for a deep understanding of how existing ranking algorithms perform, and which are their possible biases that may impair their effectiveness. Many popular ranking algorithms (such as Google's PageRank) are static in nature and, as a consequence, they exhibit important shortcomings when applied to real networks that rapidly evolve in time. At the same time, recent advances in the understanding and modeling of evolving networks have enabled the development of a wide and diverse range of ranking algorithms that take the temporal dimension into account. The aim of this review is to survey the existing ranking algorithms, both static and time-aware, and their applications to evolving networks. We emphasize both the impact of network evolution on well-established static algorithms and the benefits from including the temporal dimension for tasks such as prediction of network traffic, prediction of future links, and identification of significant nodes.

Contents

1.	Introduction.....	3
2.	Ranking with static centrality metrics	4
2.1.	Getting started: basic language of complex networks	4
2.2.	Degree and other local centrality metrics	4
2.2.1.	Degree.....	4
2.2.2.	H-index	5
2.2.3.	Other local centrality metrics	5
2.3.	Metrics based on shortest paths.....	5
2.3.1.	Closeness centrality	6
2.3.2.	Betweenness centrality	6
2.4.	Coreness centrality and its relation with degree and H-index	6

* Corresponding authors at: Department of Physics, University of Fribourg, 1700 Fribourg, Switzerland.
E-mail addresses: manuel.mariani@unifr.ch (M.S. Mariani), matus.medo@unifr.ch (M. Medo).

2.5.	Eigenvector-based centrality metrics	7
2.5.1.	Eigenvector centrality	7
2.5.2.	Katz centrality	8
2.5.3.	Win-loss scoring systems for ranking in sport.....	8
2.5.4.	PageRank	9
2.5.5.	PageRank variants.....	10
2.5.6.	HITS algorithm	11
2.6.	A case study: node centrality in the Zachary's karate club network	11
2.7.	Static ranking algorithms in bipartite networks	12
2.7.1.	Co-HITS algorithm	13
2.7.2.	Method of reflections	13
2.7.3.	Fitness-complexity metric	13
2.8.	Rating-based ranking algorithms on bipartite networks	14
3.	The impact of network evolution on static ranking algorithms	15
3.1.	The first-mover advantage in preferential attachment and its suppression.....	15
3.2.	PageRank's temporal bias and its suppression.....	16
3.3.	Illusion of influence in social systems	18
4.	Time-dependent ranking algorithms	19
4.1.	Striving for time-balance: Node-based time-rescaled metrics	19
4.2.	Metrics with explicit penalization for older edges and/or nodes	20
4.2.1.	Time-weighted degree and its use in predicting future trends.....	20
4.2.2.	Penalizing old edges: Effective Contagion Matrix	21
4.2.3.	Penalizing old edges: TimedPageRank	21
4.2.4.	Focusing on a temporal window: T-Rank, SARA	22
4.2.5.	PageRank with time-dependent teleportation: CiteRank.....	22
4.2.6.	Time-dependent reputation algorithms	22
4.3.	Model-based ranking of nodes	23
5.	Ranking nodes in temporal networks.....	25
5.1.	Getting started: how to represent a temporal network?	25
5.2.	Memory effects and time-preserving paths in temporal networks.....	26
5.3.	Ranking based on higher-order network representations	27
5.3.1.	Revisiting the Markovian assumption.....	27
5.3.2.	A second-order Markov model	28
5.3.3.	Second-order PageRank	28
5.3.4.	The effects of memory on diffusion speed	29
5.3.5.	Generalization to higher-order Markov models and prediction of real-world flows	29
5.4.	Ranking based on a multilayer representation of temporal networks	31
5.4.1.	TempoRank	31
5.4.2.	Coupling temporal layers	32
5.4.3.	Dynamic network-based scoring systems for ranking in sport.....	32
5.5.	Extending shortest-path-based metrics to temporal networks.....	33
5.5.1.	Shortest and fastest time-respecting paths	33
5.5.2.	Temporal betweenness centrality	33
5.5.3.	Temporal closeness centrality	33
6.	Temporal aspects of recommender systems.....	34
6.1.	Temporal dynamics in matrix factorization	34
6.1.1.	Matrix factorization method.....	34
6.1.2.	Including time in matrix-factorization methods.....	35
6.2.	Temporal dynamics in network-based recommendation	36
6.3.	The impact of recommendation ranking on network growth.....	37
7.	Applications of ranking to evolving social, economic and information networks	38
7.1.	Quantifying and predicting impact in academic networks	38
7.1.1.	Quantifying the significance of scientific papers.....	38
7.1.2.	Predicting the future success of scientific papers	40
7.1.3.	Predicting the future success of researchers	41
7.1.4.	Ranking scientific journals through non-Markovian PageRank	42
7.2.	Prediction of economic development of countries	42
7.2.1.	Prediction of GDP through linear regression	43
7.2.2.	Prediction of GDP through the method of analogues.....	43
7.3.	Link prediction with temporal information	45
7.3.1.	Time-dependent node similarity	45
7.3.2.	Time-dependent structural perturbation.....	45
8.	Perspectives and conclusions.....	46
8.1.	Which metric to use? Benchmarking the ranking algorithms	46
8.2.	Systemic impact of ranking algorithms: avoiding information bias	47
8.3.	Summary and conclusions	48

1. Introduction

In a world where we need to make many choices and our attention is inherently limited, we often rely on, or at least use the support of, automated scoring and ranking algorithms to orient ourselves in the copious amount of available information. One of the major goals of ranking algorithms is to *filter* [1] the large amounts of available data in order to *retrieve* [2] the most relevant information for our purposes. For example, quantitative metrics for assessing the relevance of websites are motivated by the impossibility, for an individual web user, to browse manually large portions of the Web to find the information that he or she needs.

The resulting rankings influence our choices of items to purchase [3–5], the way we access scientific knowledge [6] and online content [7–11]. At a systemic level, ranking affects how success and excellence in human activity are recognized [12–17], how funding is allocated in scientific research [18], how to characterize and counteract the outbreak of an infectious disease [19–21], and how undecided citizens vote for political elections [22], for example. Given the strong influence of rankings on various aspects of our lives, achieving a deep understanding of how respective ranking methods work, realizing their basic assumptions and limitations, and designing improved methods is of critical importance. Rankings' disadvantages, such as the biases that they can induce, can otherwise outweigh the benefits that they bring. Achieving such understanding becomes more urgent as a huge and ever-growing amount of data is created and stored every day and, as a result, automated methods to extract relevant information are poised to further gain importance in the future [23,24].

An important class of ranking algorithms is based on the *network representation* of the input data. Complex networks have recently emerged as one of the leading frameworks to analyze a wide range of complex social, economic and information systems [25–28]. A network is composed of a set of *nodes* (also referred to as *vertices*) and a set of relationships (referred to as *edges* or *links*) between the nodes. A network representation of a given system greatly reduces the complexity of the original system and, if the representation is well-designed, it allows us to gain fundamental insights into the structure and dynamics of the original system. Fruitful applications of the complex networks approach include prediction of spreading of infectious diseases [19,29], characterization and prediction of economic development of countries [30–32], characterization and detection of early signs of financial crises [33–35], and design of robust infrastructure and technological networks [36,37], among many others.

Network-based ranking algorithms use network representations of the data to compute the inherent value, relevance, or importance of individual nodes in the system [38,39]. The idea of using the network structure to infer node status is long known in social sciences [40,41], where node-level ranking algorithms are typically referred to as *centrality metrics* [42]; in this review, we use the labels “ranking algorithms” and “centrality metrics” interchangeably. Today, research on ranking methods is being carried out by scholars from several fields (such as physicists, mathematicians, social scientists, economists, computer scientists). Network-based ranking algorithms have drawn conspicuous attention from the physics community due to direct applicability of statistical-physics concepts such as random walk and diffusion [43–48], as well as percolation [49–51]. Real-world applications of network-based ranking algorithms cover a vast spectrum of problems, including the design of search engines [38,43,52] and recommender systems [53–55], evaluation of academic research [13,56,57], and identification of the influential spreaders [50,51,58,59]. A paradigmatic example of a network-based ranking algorithm is the Google's popular PageRank algorithm. Although originally devised to rank web pages [43], it has since then found applications in a vast range of real systems [14,57,60–68].

While there are excellent reviews and surveys that address network-based ranking algorithms [38,39,61,62,69] and random walks in complex networks [46], the impact of network evolution on ranking algorithms is discussed there only marginally. This gap deserves to be filled because real networks usually evolve in time [70–76], and temporal effects have been shown to strongly influence properties and effectiveness of ranking algorithms [44,77–83]. The aim of this review is to fill this gap and thoroughly present the temporal aspects of network-based ranking algorithms in social, economic and information systems that evolve in time. The motivation for this work is manifold. First, we want to present in a cohesive fashion information that is scattered across literature from diverse fields, providing scholars from diverse communities with a reference point on the temporal aspects of network-based ranking. Second, we want to highlight the essential role of the temporal dimension, whose omission can produce misleading results when it comes to evaluating the centrality of a node. Third, we want to highlight open challenges and suggest new research directions. Particular attention will be devoted to network-based ranking algorithms that build on classical statistical physics concepts such as random walks and diffusion.

In this review, we will survey both *static* and *time-aware* network-based ranking algorithms. *Static* algorithms take into account only the network's adjacency matrix A (see Section 2.1 where the basic notation used in this review is provided) to compute node centrality, and they constitute the topic of Section 2. The dynamic nature of most real networks can cause static algorithms to fail in a number of real-world systems, as outlined in Section 3. *Time-aware* algorithms take as input the network's adjacency matrix at a given time t and the list of node and/or edge time-stamps. Time-aware algorithms for growing unweighted networks are presented in Section 4. When multiple interactions occur among the network's nodes, algorithms based on *temporal-network representations* of the data are often needed; this class of ranking

algorithms is presented in Section 5 alongside a brief introduction to the temporal network representation of time-stamped datasets. Network-based time-aware recommendation methods and their applications are covered by Section 6; the impact of recommender systems on network evolution is also discussed in this section.

Due to fundamental differences between the datasets where network centrality metrics are applied, no universal and all-encompassing ranking algorithm can exist. The best that we can do is to identify the specific ranking task (or tasks) that a metric is able to address. As it would be impossible to cover all the applications of network centrality metrics studied in the literature, we narrow our focus to applications where the temporal dimension plays an essential role (Section 7). For example, we discuss how time-aware metrics can be used to early identify significant nodes (Section 7.1), how node importance can be used as a predictor of future GDP of countries (Section 7.2), and how time-aware metrics outperform static ones in the classical link-prediction problem (Section 7.3). We provide a general discussion on the problem of ranking in Section 8. In this discussion, we focus on possible ways of validating the existing centrality metrics, and on the importance of detecting and suppressing the biases of the rankings obtained with various metrics.

While this review covers a broad range of interwoven topics, the chapters of this review are as self-contained as possible, and in principle, they can be read independently. We do discuss the underlying mathematical details only occasionally and refer to the corresponding specialized references instead where the interested reader can find more details.

2. Ranking with static centrality metrics

Centrality metrics aim at identifying the most important, or central, nodes in the network. In this section, we review the centrality metrics that only take into account the topology of the network encoded in the network's adjacency matrix A . These metrics neglect the temporal dimension and are thus referred to as *static* metrics here. Static centrality metrics have a long-standing tradition in social network analysis: Katz centrality (Section 2.5) was introduced in the 50s [41], betweenness centrality (Section 2.3) in the 70s [84], and indegree (citation count, Section 2.2) is used to rank scholarly publications since the 70s. The crucial role played by the PageRank algorithm in the outstanding success of the web search engine Google was one of the main motivating factors for a new wave of interest in centrality metrics starting in the late 90s. Today, we are aware that static centrality metrics may exhibit severe shortcomings when applied to dynamically evolving systems – as will be examined in depth in the next sections – and, for this reason, they should never be applied without considering the specific properties of the system in hand. Static metrics are nevertheless important both as basic heuristic tools of network analysis, and as building blocks for time-aware metrics.

2.1. Getting started: basic language of complex networks

Since excellent books [26–28,85] and reviews [25,72,73,86] on complex networks and their applications have already been written, we introduce here only the network-analysis concepts and tools that are indispensable for the presented ranking methods. A network (usually referred to as graph in mathematics literature) composed of N nodes and E edges is denoted as $\mathcal{G}(N, E)$. The network is completely described by its *adjacency matrix* A . In the case of an *undirected* network, matrix element $A_{ij} = 1$ if an edge exists that connects i and j , zero otherwise. The *degree* k_i of node i is then defined as the number of connections of node i . In the case of a *directed* network, $A_{ij} = 1$ if there is a directed edge $j \rightarrow i$ from j to i , zero otherwise. The *outdegree* k_i^{out} (*indegree* k_i^{in}) of node i is defined as the number of outgoing (incoming) edges that connect node i to other nodes.

We introduce now four structural notions that will be useful in what follows: *path*, *path length*, *shortest path*, and *geodesic distance between two nodes*. The definitions provided below apply to undirected networks, but they can be straightforwardly generalized to directed networks (in that case, we speak of directed paths). A network path is defined as a sequence of nodes such that consecutive nodes in the path are connected by an edge. The *length of a path* is defined as the number of edges that compose the path. The *shortest path* between two nodes i and j is the path of the smallest possible length that starts in node i and ends in node j – note that multiple distinct shortest paths between two given nodes may exist. The (*geodesic*) *distance* between two nodes i and j is the length of the shortest paths between i and j . We will use these basic definitions in Section 2.3 to define the closeness and betweenness centrality.

2.2. Degree and other local centrality metrics

We begin with centrality metrics that are *local* in nature. By local metrics, we mean that only the immediate neighborhood of a given node is taken into account for the calculation of the node's centrality score. This is fundamentally different from, for example, eigenvector-based metrics that consider all network's paths (Section 2.5) and metrics that consider all network's shortest paths that pass through a given node to determine its score (Section 2.3).

2.2.1. Degree

The simplest structural centrality metric is degree. For an undirected network, the degree of a node i is defined as

$$k_i = \sum_j A_{ij}. \quad (1)$$

That is, the degree of a node is given by the number of edges attached to it. Using degree as a centrality metric assumes that a node is important if it has many connections. Even though the importance of neighbors of a given node is presumably relevant to determine its centrality [40,60], and this aspect is ignored by degree, one can argue that node degree still captures some part of node importance. In particular, degree is highly correlated with more sophisticated centrality metrics for uncorrelated networks [87,88], and can perform better than more sophisticated network-based metrics in specific tasks such as identifying influential spreaders [89,90].

In the case of a directed network, one can define node degree either by counting the number of incoming edges or the number of outgoing edges – the outcomes of the two different counting procedures are referred to as node indegree and outdegree, respectively. In most cases, incoming edges can be interpreted as recommendation, or endorsement, received from the node from which the edge comes from. While the assumption that edges represent endorsements is not always true,¹ it is reasonable enough for a broad range of real systems and, unless otherwise stated, we assume it throughout this review. For example, in networks of academic publications, a paper’s incoming edges represent its received citations, and their number can be considered as a rough proxy for the paper’s impact.² For this reason, indegree is often used as a centrality metric in directed networks

$$k_i^{in} = \sum_j A_{ij}. \quad (2)$$

Indegree centrality is especially relevant in the context of quantitative evaluation of academic research, where the number of received citations is often used to gauge scientific impact (see [17] for a review of bibliometric indices based on citation count).

2.2.2. *H-index*

A limitation of the degree centrality is that it only considers the node’s number of received edges, regardless of the centrality of the neighbors. By contrast, in many networks it is plausible to assume that a node is important if it is connected to other important nodes in the network. This thesis is enforced by eigenvector-based centrality metrics (Section 2.5) by taking into account the whole network topology to determine a node’s score. A simpler way to implement this idea is to consider nodes’ neighborhood in the node score computation, without including the global network structure. The *H-index* is a centrality metric based on this idea.

The *H-index* for unipartite networks [94] is based on the metric of the same name introduced by Hirsch for bipartite networks [12] to assess researchers’ scientific output. Hirsch’s index of a given researcher is defined as the maximum integer h such that at least h publications authored by that researcher received at least h citations. In a similar way, Lu et al. [94] define the *H-index* of a given node in a monopartite network as the maximum integer h such that there exist at least h neighbors of that node with at least h neighbors each. In this way, a node’s centrality score is determined by the neighbors’ degree values, without the need for inspecting the whole network topology. The *H-index* brings some advantage with respect to degree in identifying the most influential spreaders in real networks [90,94], yet the correlation between the two metrics is high for uncorrelated networks (*i.e.*, networks with negligible degree–degree correlations) [88]. More details on the *H-index* and its relation with other centrality metrics can be found in Section 2.4.

2.2.3. *Other local centrality metrics*

Several other methods have been proposed in the literature to build local centrality metrics. The main motivation to prefer local metrics to “global” metrics (such as eigenvector-based metrics) is that local metrics are much less computationally expensive and thus can be evaluated in relatively short time even in huge graphs. As paradigmatic examples, we mention here the local centrality metric introduced by Chen et al. [95] which considers information up to distance four from a given node to compute the node’s score, and the graph expansion technique introduced by Chen et al. [96] to estimate a website’s PageRank score, which will be defined in Section 2.5. As this review is mostly concerned with time-aware methods, a more detailed description of static local metrics goes beyond its scope. An interested reader can refer to a recent review article by Lu et al. [90] for more information.

2.3. *Metrics based on shortest paths*

This Section presents two important metrics – closeness and betweenness centrality – that are based on the shortest paths in the network that we defined in Section 2.1.

¹ In a signed social network, for example, an edge between two individuals can be either positive (friendship) or negative (animosity). See [91], among others, for generalizations of popular centrality metrics to signed networks.

² This assumption neglects the multi-faceted nature of the scientific citation process [92], and the fact that citations can be both positive and negative – a citation can be considered as “negative” if the citing paper points out flaws of the cited paper, for example [93]. We describe in the following sections some of the shortcomings of citation count and possible ways to counteract them.

2.3.1. Closeness centrality

The basic idea behind closeness centrality is that a node is central if it is “close” (in a network sense) to many other nodes. The first way to implement this idea is to define the closeness centrality [97] score of node i as the reciprocal of its average geodesic distance from the other nodes

$$c_i = \frac{N - 1}{\sum_{j \neq i} d_{ij}}, \quad (3)$$

where d_{ij} denotes the geodesic distance between i and j . A potential problem with Eq. (3) is that if there is a node that cannot be reached from node i , i 's score as determined by Eq. (3) is identically zero. This makes Eq. (3) useless for networks composed of more than one component. A common way to overcome this difficulty [98,99] is to define node i 's centrality as a harmonic mean of its distance from the other nodes

$$c_i = \frac{1}{N - 1} \sum_{j \neq i} \frac{1}{d_{ij}}. \quad (4)$$

With this definition, a given node i receives zero contribution from nodes that are not reachable with a path starting at node i , whereas it receives a large contribution from close nodes – node i 's direct neighbors give contribution one to i 's closeness score, for example.

The main assumption behind closeness centrality deserves some attention. If we are interested in node centrality as an indicator of the expected hitting time for an information (or a disease) flowing through the network, closeness centrality implicitly assumes that information only flows through the shortest paths [42,100]. As pointed out by Borgatti [100], this assumption makes the metric useful only for assessing node importance in situations where flow from any source node has a prior global knowledge of the network and thus knows how to reach the target node through the shortest possible path – this could be the case for a parcel delivery process or a passenger trip, for example. However, this assumption is unrealistic for real-world spreading processes where information flows typically have no prior knowledge of the whole network topology (think of an epidemic disease) and can often reach target nodes through longer paths [20]. Empirical experiments reveal that closeness centrality typically (but not for all datasets) underperforms with respect to local metrics in identifying influential spreaders for SIS and SIR diffusion processes [90,94].

2.3.2. Betweenness centrality

The main assumption of betweenness centrality is that a given node is central if many shortest paths pass through it. To define the betweenness centrality score, we denote by σ_{st} the number of shortest paths between nodes s and t , and by $\sigma_{st}^{(i)}$ the number of these paths that pass through node i . The betweenness centrality score of node i is given by

$$b_i = \frac{1}{(N - 1)(N - 2)} \sum_{s \neq t \neq i} \frac{\sigma_{st}^{(i)}}{\sigma_{st}}. \quad (5)$$

Similarly to closeness centrality, betweenness centrality implicitly assumes that only the shortest paths matter in the process of information transmission (or traffic) between two nodes. While a variant of betweenness centrality which accounts for all network paths has also been considered, it often gives similar results as the original betweenness centrality [101], and no evidence has been found that these metrics accurately reproduce information flows in real networks [26]. In addition, betweenness centrality performs poorly with respect to other centrality metrics in identifying influential spreaders in numerical simulations with the standard SIR and SIS models [90,94]. Another drawback of closeness and betweenness centrality is their high computational complexity mostly due to the calculation of the shortest paths between all pairs of nodes in the network.

2.4. Coreness centrality and its relation with degree and H -index

The coreness centrality (also known as the k -shell centrality [102]) is based on the idea of decomposing the network [103] to distinguish among the nodes that form the core of the network and the peripheral nodes. The decomposition process is usually referred to as the k -shell decomposition [102]. The k -shell decomposition starts by removing all nodes with only one connection (together with their edges), until no more such nodes remain, and assigns them to the 1-shell. For each remaining node, the number of edges connecting to other remaining nodes is called its residual degree. After having assigned the 1-shell, all nodes with residual degree 2 are recursively removed and the 2-shell is created. This procedure continues as the residual degree increases until all nodes in the networks have been assigned to one of the shells. The index c_i of the shell to which a given node i is assigned to is referred to as the node coreness (or k -shell) centrality. Coreness centrality can also be generalized to weighted networks [104], and refined to achieve better performance in identifying influential spreaders [105].

There is an interesting mathematical connection between k -shell centrality, degree and H -index. In Section 2.2, we have introduced the H -index of a node i as the maximum value h such that there exist h neighbors of i , each of them having at least h neighbors. This definition can be formalized by introducing an operator \mathcal{H} which acts on a finite number of real numbers (x_1, x_2, \dots, x_n) by returning the maximum integer $y = \mathcal{H}(x_1, x_2, \dots, x_n) > 0$ such that there exist at least y elements in

$\{x_1, \dots, x_n\}$ each of which is larger or equal than y [94]. Denoting by $(k_{j_1}, k_{j_2}, \dots, k_{j_{k_i}})$ the degrees of the neighbors of a given node i , the H -index of node i can be written in terms of \mathcal{H} as

$$h_i = \mathcal{H}(k_{j_1}, k_{j_2}, \dots, k_{j_{k_i}}). \quad (6)$$

Lu et al. [94] generalize the h index by recursively defining the n th-order H -index of node i as

$$h_i^{(n)} = \mathcal{H}(h_{j_1}^{(n-1)}, h_{j_2}^{(n-1)}, \dots, h_{j_{k_i}}^{(n-1)}), \quad (7)$$

where $h_i^{(0)} = k_i$ and $h_i^{(1)}$ is the H -index of node i . For arbitrary undirected networks, it can be proven [94] that as n grows, $h_i^{(n)}$ converges to coreness centrality c_i ,

$$\lim_{n \rightarrow \infty} h_i^{(n)} = c_i. \quad (8)$$

The family of H -indexes defined by Eq. (7) thus bridges from node degree ($h_i^{(0)} = k_i$) through H -index ($h_i^{(1)} = h_i$) to coreness, that is from a local centrality metric to a global one.

In many real-world networks, the H -index outperforms both degree and coreness in identifying influential spreaders as defined by classical SIR and SIS network spreading models. Recently, Pastor-Satorras and Castellano [88] performed an extensive analytic and numerical investigation of mathematical properties of the H -index, finding analytically that the H -index is expected to be strongly correlated with degree in uncorrelated networks. The same strong correlation is found also in real networks, and Pastor-Satorras and Castellano [88] point out that the H -index is a poor indicator of node spreading influence as compared to the non-backtracking centrality [106]. While the comparison of different metrics with respect to their ability to identify influential spreaders is not one of the main focus of the present review, we point the interested reader to [51,88,90] for recent reports on this important problem.

2.5. Eigenvector-based centrality metrics

Differently from the local metrics presented in Section 2.2 and the metrics based on shortest paths presented in Section 2.3, eigenvector-based metrics take into account all network paths to determine node importance. We refer to the metrics presented in this Section as eigenvector-based metrics because their respective score vectors can be interpreted as the principal eigenvector of a matrix which only depends on the network's adjacency matrix A .

2.5.1. Eigenvector centrality

Often referred to as Bonacich centrality [41], the eigenvector centrality assumes that a node is important if it is connected to other important nodes. To build the eigenvector centrality, we start from a uniform score value $s_i^{(0)} = 1$ for all nodes i . At each subsequent iterative step, we update the scores according to the equation

$$s_i^{(n+1)} = \sum_j A_{ij} s_j^{(n)} = (\mathbf{A}\mathbf{s})_i. \quad (9)$$

The vector of eigenvector centrality scores is defined as the fixed point of this equation. The iterations converge to a vector proportional to the principal eigenvector \mathbf{v}_1 of the adjacency matrix A . To prove this, we first observe that for connected networks, \mathbf{v}_1 is unique³ due to the Perron–Frobenius theorem [107].

We then expand the initial condition $\mathbf{s}^{(0)}$ in terms of the eigenvectors $\{\mathbf{v}_\alpha\}$ of A as

$$\mathbf{s}^{(0)} = \sum_\alpha c_\alpha \mathbf{v}_\alpha. \quad (10)$$

Denoting the eigenvalue associated with the eigenvector \mathbf{v}_α as λ_α and assuming that the eigenvalues are ordered and thus $|\lambda_{\alpha+1}| < |\lambda_\alpha|$, Eq. (9) implies

$$\mathbf{s}^{(n)} = \mathbf{A}^n \mathbf{s}^{(0)} = c_1 \lambda_1^n \mathbf{v}_1 + \sum_{\alpha>1} c_\alpha \lambda_\alpha^n \mathbf{v}_\alpha. \quad (11)$$

When n grows, the second term on the r.h.s. of the previous equation is exponentially small with respect to the r.h.s.'s first term, which results in

$$\mathbf{s}^{(n)} / \lambda_1^n \rightarrow c_1 \mathbf{v}_1. \quad (12)$$

³ For networks with more than one component, there can exist more than one eigenvector associated with the principal eigenvalue of A . However, it can be shown [26] that only one of them can have all components larger than zero, and thus it is the only one that can be obtained by iterating Eq. (9) since we assumed $s_i^{(1)} = 1$ for all i .

By iterating Eq. (9) and continually normalizing the score vector, the scores converge to the adjacency matrix's principal eigenvector \mathbf{v}_1 . Consequently, the vector of eigenvector centrality scores \mathbf{v}_1 satisfies the self-consistent equation

$$s_i = \lambda_1^{-1} \sum_j A_{ij} s_j, \quad (13)$$

which can be rewritten in a matrix form as

$$\mathbf{s} = \lambda_1^{-1} \mathbf{A} \mathbf{s}. \quad (14)$$

The assumptions of the eigenvector centrality seem plausible: high centrality nodes give high contribution to the centrality of the nodes they connect to. However, the metric has an important shortcoming when applied to directed networks. If a node has no incoming links, it has zero score according to Eq. (14) and consequently gives zero contribution to the scores of the nodes it connects to. This makes the metric particularly useless for directed acyclic graphs (such as the network of scientific papers) where it assigns zero score to all papers. The Katz centrality metric, that we introduce below, solves this issue by adding a constant term to the r.h.s. of Eq. (14), and thus achieves that the score of all nodes is positive.

2.5.2. Katz centrality

Katz centrality metric [40] builds on the same premise as eigenvector centrality – a node is important if it is connected to other important nodes – but, differently from eigenvector centrality, it assigns certain minimum score to each node. Different variants of the Katz centrality have been proposed in the literature [40,41,108,109] – we refer the reader interested in the subtleties associated with various possible definitions of the Katz centrality to the original papers, and we present here the variant proposed in [109]. The vector of Katz-centrality (referred to as *alpha-centrality* in [109]) scores is there defined as

$$\mathbf{s} = \alpha \mathbf{A} \mathbf{s} + \beta \mathbf{e} \quad (15)$$

where \mathbf{e} is the N -dimensional vector whose elements are all equal to one. The solution \mathbf{s} of this equation exists only if $\alpha < 1/\lambda_1$ where λ_1 is the principal eigenvalue of \mathbf{A} . For this reason, we restrict the following analysis to this range of α values. The solution reads

$$\mathbf{s} = \beta (1 - \alpha \mathbf{A})^{-1} \mathbf{e}, \quad (16)$$

where $\mathbf{1}$ denotes the $N \times N$ identity matrix. This solution can be written in the form of a geometric series

$$\mathbf{s} = \beta \sum_{k=0}^{\infty} \alpha^k \mathbf{A}^k \mathbf{e}. \quad (17)$$

The score s_i of node i can be thus expressed as

$$s_i = \beta \left(1 + \alpha \sum_j A_{ij} + \alpha^2 \sum_{j,l} A_{ij} A_{jl} + \alpha^3 \sum_{j,l,m} A_{ij} A_{jl} A_{lm} + \mathcal{O}(\alpha^4) \right). \quad (18)$$

Eq. (17) (or, equivalently, (18)) shows that the Katz centrality score of a node is determined by all network paths that pass through node i . Note that $\sum_j A_{ij}$, $\sum_{j,l} A_{ij} A_{jl}$ and $\sum_{j,l,m} A_{ij} A_{jl} A_{lm}$ represent the paths of length one, two and three, respectively, that have node i as the last node. The parameter α determines how contributions of various paths attenuate with path length. For small values of α , long paths are strongly penalized and node score is mostly determined by the shortest possible paths of length one (i.e., by node degree). By contrast, for large values of α (but still $\alpha < \lambda_1^{-1}$, otherwise the series in Eq. (18) does not converge) long paths contribute highly to node score.

Known in social science since the 50s [40], the Katz centrality has been used in a number of applications. For example, a variant of the Katz centrality has been recently shown to substantially outperform other static centrality metrics in predicting neuronal activity [110]. While the metric overcomes the drawback of eigenvector centrality and produces a meaningful ranking also when applied to directed acyclic graphs or disconnected networks, the Katz centrality may still provide unsatisfactory results in networks where the outdegree distribution is heterogeneous. An important node is then able to manipulate the other nodes' importance by simply creating many edges towards a selected group of target nodes, and thus improving their score. Google's PageRank centrality overcomes this limitation by weighting less the edges that are created by nodes with many outgoing connections. Before introducing PageRank centrality, we briefly discuss a simple variant of the Katz centrality that has been applied to networks of sport professional teams and players.

2.5.3. Win-loss scoring systems for ranking in sport

Network-based win-loss scoring systems aim to rank competitors (for simplicity, we refer to players in this Section) in sports based on one-versus-one matches (as in tennis, football, or baseball). In terms of complex networks, each player i is represented by a node and the weight of the directed edge $j \rightarrow i$ represents the number of wins of player j against player i . The main idea of a network-based win-loss scoring scheme is that a player is strong if it is able to defeat other strong

players (i.e., players that have defeated many opponents [111]). This idea led Park and Newman [111] to define the vector of win scores \mathbf{w} and the vector of loss scores \mathbf{l} through a variant of the Katz centrality equation

$$\begin{aligned}\mathbf{w} &= (\mathbf{1} - \alpha \mathbf{A}^\top)^{-1} \mathbf{k}^{out}, \\ \mathbf{l} &= (\mathbf{1} - \alpha \mathbf{A})^{-1} \mathbf{k}^{in},\end{aligned}\tag{19}$$

where $\alpha \in (0, 1)$ is a parameter of the method. The vector \mathbf{s} of player scores is defined as the win–loss differential $\mathbf{s} = \mathbf{w} - \mathbf{l}$. To understand the meaning of the win score, we use again the geometric series of matrices to write

$$w_i = \sum_j A_{ji} + \alpha \sum_{j,k} A_{kj} A_{ji} + \alpha^2 \sum_{j,k,l} A_{kj} A_{jl} A_{li} + \mathcal{O}(\alpha^3).\tag{20}$$

From this equation, we realize that the leading-order contribution to w_i is simply the total number of wins of player i . The successive terms represent “indirect wins” achieved by player i . For example, the term proportional to $\mathcal{O}(\alpha)$ represents the total number of wins of players beaten by player i . An analogous reasoning applies to the loss score \mathbf{l} . Based on similar assumptions as the win–loss score, PageRank centrality metric has been applied to rank sport competitors by Radicchi [14], leading to the Tennis Prestige Score.⁴ A time-dependent generalization of the win–loss score will be presented in Section 5.4.3.

2.5.4. PageRank

PageRank centrality metric has been introduced by Brin and Page [43] with the aim to rank web pages in the World Wide Web graph, and there is general agreement on the essential role played by this metric in the outstanding success of Google’s Web search engine [52,60]. For a directed network, the vector \mathbf{s} of PageRank scores is defined by the equation

$$\mathbf{s} = \alpha \mathbf{P} \mathbf{s} + (1 - \alpha) \mathbf{v}\tag{21}$$

where $P_{ij} = A_{ij}/k_j^{out}$ is the network’s transition matrix,⁵ \mathbf{v} is called the teleportation vector and α is called the teleportation parameter.⁶ A uniform teleportation vector ($v_i = 1/N$ for all nodes i) is the original choice by Brin and Page, and it is arguably the most common choice for \mathbf{v} , although benefits and drawbacks of other teleportation vectors \mathbf{v} have been also explored in the literature [113]. This equation is conceptually similar to the equation that defines Katz centrality, with the important difference that the adjacency matrix \mathbf{A} has been replaced by the transition matrix \mathbf{P} . From Eq. (21), it follows that the vector of PageRank scores can be interpreted as the leading eigenvector of the matrix

$$\mathbf{G} = \alpha \mathbf{P} + (1 - \alpha) \mathbf{v} \mathbf{e}^\top\tag{22}$$

where $\mathbf{e} = (1, \dots, 1)_N$. \mathbf{G} is often referred to as the *Google matrix*. We refer to the review article by Ermann et al. [62] for a detailed review of the mathematical properties of the spectrum of \mathbf{G} .

Beyond its interpretation as an eigenvalue problem, there exists also an illustrative physical interpretation of the PageRank equation. Eq. (21) can be viewed as the stationary-state equation of a stochastic process on the network where a random walker located at a certain node can make two moves: (1) with probability α , jump to another node by following a randomly chosen edge starting at the node where the walker is located; (2) with probability $1 - \alpha$, “teleport” to a randomly chosen node. The master equation of this random walk process reads

$$\mathbf{s}^{(n+1)} = \alpha \mathbf{P} \mathbf{s}^{(n)} + (1 - \alpha) \mathbf{v}.\tag{23}$$

The analogy of PageRank with physical diffusion on the network is probably one of the properties that have most stimulated physicists’ interest in studying the PageRank’s properties.

In most real-world directed networks (such as the Web graph), there are nodes (referred to as *dangling nodes* [114]) without outgoing edges. These nodes with zero out-degree receive score from other nodes, but do not redistribute it to other nodes. In Eq. (21), the existence of dangling nodes makes the transition matrix ill-defined as its elements corresponding to transitions from the dangling nodes are 0/0. There are several strategies how to deal with dangling nodes. For example, one can remove them from the system; this removal is unlikely to significantly affect the ranking of the remaining nodes as by definition they receive no score from the dangling nodes. One can also artificially set $k_j^{out} = 1$ for the dangling nodes which does not affect the score of the other nodes [26]. Another strategy is to replace the ill-defined elements of the transition matrix with uniform entries; in other words, the transition matrix is re-defined as

$$P_{ij} = \begin{cases} A_{ij}/k_j^{out} & \text{if } k_j^{out} > 0, \\ 1/N & \text{if } k_j^{out} = 0. \end{cases}\tag{24}$$

We direct the interested reader to the survey article by Berkhin [114] for a presentation of possible alternative strategies to account for the dangling nodes.

⁴ <http://tennisprestige.soic.indiana.edu/>

⁵ The case $k_j^{out} = 0$ (dangling nodes) is discussed below.

⁶ Some authors [112] refer to α as *damping factor* or *damping parameter*, for a reason that will become clear later in this Section.

Using the Perron–Frobenius theorem, one can show that the existence and uniqueness of the solution of Eq. (21) is guaranteed if $\alpha \in [0, 1)$ [61,114]. We assume here α to be in the range $(0, 1)$. The solution of Eq. (21) reads

$$\mathbf{s} = (1 - \alpha)(1 - \alpha \mathbf{P})^{-1} \mathbf{v}. \quad (25)$$

Similarly as we did for Eq. (16), \mathbf{s} can be expanded by using the geometric series to obtain

$$\mathbf{s} = (1 - \alpha) \sum_{k=0}^{\infty} \alpha^k \mathbf{P}^k \mathbf{v}. \quad (26)$$

As was the case for the Katz centrality, also the PageRank score of a given node is determined by all the paths that pass through that node. The teleportation parameter α controls the exponential damping of longer paths. For small α (but larger than zero), long paths give a negligible contribution and – if outdegree fluctuations are sufficiently small – the ranking by PageRank approximately reduces to the ranking by indegree [56,115]. When α is close to (but smaller than) one, long paths give a substantial contribution to node score. While there is no universal criterion for choosing α , a number of studies (see [112,116,117], for example) have pointed out that large values of α might lead to a ranking that is highly sensible to small perturbations of the network’s structure, which suggests [117] that values around 0.5 should be preferred to the original value $\alpha = 0.85$ used by Brin and Page.

An instructive interpretation of α comes from viewing the PageRank algorithm as a stochastic process on the network (see Eq. (23)). At each iteration step, the random walker has to decide whether to follow a network edge or whether to teleport to a randomly chosen node. On average, the length of the network paths covered by the walker before teleporting to a randomly chosen node is given by [39]

$$\langle l \rangle_{RW} = (1 - \alpha) \sum_{k=0}^{\infty} k \alpha^k = \frac{\alpha}{1 - \alpha}, \quad (27)$$

which is exactly the ratio between the probability of following a link and the probability of random teleportation. For $\alpha = 0.85$, we obtain $\langle l \rangle_{RW} = 5.67$ which corresponds to following six edges before teleporting to a random node. Some researchers ascribe the PageRank’s success to the fact that the algorithm provides a reasonable model of the real behavior of Web users, who surf the Web both by following the hyperlinks that they find in the webpages, and – perhaps when they get bored or stuck – by restarting their search (teleportation). For $\alpha = 0.5$, we obtain $\langle l \rangle_{RW} = 1$. This choice of α may better reflect the behavior of researchers on the academic network of papers (one may feel compelled to check the references in a work of interest but rarely to check the references in one of the referred papers). Chen et al. [56] use $\alpha = 0.5$ for a citation network motivated by the finding that half of the entries in the reference list of a typical publication cite at least one other article in the same reference list, which might further indicate that researchers tend to cover paths up to the length of one when browsing the academic literature.

2.5.5. PageRank variants

PageRank is built on three basic ingredients: the transition matrix \mathbf{P} , the teleportation parameter α , and the teleportation vector \mathbf{v} . With respect to the original algorithm, a majority of PageRank variants are still based on Eq. (21), and modify one or more of these three elements. In this Section, we focus on PageRank variants that do not explicitly depend on edge time stamps. We present four such variants; the reader interested in further variations is referred to the review article by Gleich [61]. Time-dependent variants of PageRank and Katz centrality are the focus of Section 4.

Reverse PageRank. The reverse PageRank [118] score vector for a network described by the adjacency matrix \mathbf{A} is defined as the PageRank score vector for the network described by the transpose \mathbf{A}^T of \mathbf{A} . In other words, reverse-PageRank scores flow in the opposite edge direction with respect to PageRank scores. Reverse PageRank has been also referred to as CheiRank in the literature [62,119]. In general, PageRank and reverse PageRank provide different information on the node role in a network. In particular, large reverse PageRank values mean that the nodes can reach many other nodes in the network. An interested reader can find more details in the review article by Ermann et al. [62].

LeaderRank. In the LeaderRank algorithm [120], a “ground-node” is added to the original network and connected with each of the network’s N nodes by an outgoing and incoming link. A random walk is performed on the resulting network; the LeaderRank score of a given node is given by the fraction of time the random walker spends on that node. With respect to PageRank, LeaderRank keeps the idea of performing a random walk on network’s links, but it does not feature any teleportation mechanism and, as a consequence, it is parameter-free. By adding the ground node, however, the algorithm effectively enforces a degree-dependent teleportation probability.⁷ Lu et al. [120] analyzed online social network data to show that the ranking by LeaderRank is less sensitive than that by PageRank to perturbations of the network structure and to malicious manipulations on the system. Further variants based on degree-dependent weights have been used to detect influential and robust nodes [121]. Furthermore, the “ground-node” notion of LeaderRank has also been generalized to bipartite networks by network-based recommendation systems [122].

⁷ From a node of the original degree k_i , the probability to follow one of the actual links is $k_i/(k_i + 1)$ as opposed to the probability $1/(k_i + 1)$ to move to the ground node and then consequently to a random one of the network’s N nodes.

PageRank with “smart” teleportation. In the original PageRank formulation, the teleportation vector \mathbf{v} represents the amount of score that is assigned to each node by default, independently of network structure. The original PageRank algorithm implements the simplest possible teleportation vector, $\mathbf{v} = \mathbf{e}/N$, where $e_i = 1$ for all components i . This is not the only possible choice, and one might instead assume that higher-degree nodes should be given higher baseline score than obscure nodes. This thesis has been referred to as “smart” teleportation by Lambiotte and Rosvall [113], and it can be implemented by setting $v_i \propto k_i^n$. Smart-teleportation PageRank score can still be expressed in terms of network paths through Eq. (26); differently from PageRank, the zero-order contribution to a given node’s smart-teleportation PageRank score is exactly its indegree. The ranking by smart-teleportation PageRank results to be remarkably more stable than original PageRank with respect to variations in the teleportation parameter α [113]. A detailed comparison of the two rankings as well as a discussion of alternative teleportation strategies can be found in [113].

Pseudo-PageRank. It is worth mentioning that PageRank variants are sometimes built on a different definition of node score, based on a column sub-stochastic matrix⁸ \mathbf{Q} and a non-normalized additive vector \mathbf{f} ($f_i \geq 0$). The corresponding equation that defines the vector \mathbf{s} of nodes’ score is

$$\mathbf{s} = \alpha \mathbf{Q} \mathbf{s} + \mathbf{f}. \quad (28)$$

In agreement with the review by Gleich [61], we refer the problem of solving Eq. (28) as the pseudo-PageRank problem.

One can prove that the pseudo-PageRank problem is equivalent to a PageRank problem (Theorem 2.5 in [61]). In other words, if some authors use the pseudo-PageRank equation (28) to define nodes’ score, the resulting scores are proportional to the scores that solve a PageRank equation that depends on \mathbf{Q} and \mathbf{f} .

More precisely, let \mathbf{y} be the solution of a pseudo-PageRank system with parameter α , column sub-stochastic matrix \mathbf{Q} and additive vector \mathbf{f} . If we define $\mathbf{v} := \mathbf{f}/(\mathbf{e}^T \mathbf{f})$ and $\mathbf{x} := \mathbf{y}/(\mathbf{e}^T \mathbf{y})$, then \mathbf{x} is the solution of a PageRank system with teleportation parameter α , stochastic matrix $\mathbf{P} = \mathbf{Q} + \mathbf{v} \mathbf{c}^T$, and teleportation vector \mathbf{v} , where $\mathbf{c}^T = \mathbf{e}^T - \mathbf{e}^T \mathbf{Q}$ is the correction vector needed to make \mathbf{Q} stochastic. The (unique) solution \mathbf{y} of Eq. (28) thus inherits the properties of the corresponding PageRank solution \mathbf{x} – the two solutions \mathbf{y} and \mathbf{x} only differ by a uniform normalization factor. In the special case where \mathbf{Q} is column stochastic and \mathbf{f} is not normalized to one, we have $\mathbf{P} = \mathbf{Q}$ and $\mathbf{v} := \mathbf{f}/(\mathbf{e}^T \mathbf{f})$ – see [65] and Section 4.2.5 for an example. Since PageRank and pseudo-PageRank problem are equivalent, we use the two descriptions interchangeably in the following.

2.5.6. HITS algorithm

HITS (Hyperlink-Induced Topic Search) algorithm [123] is a popular eigenvector-based ranking method originally aimed at ranking webpages in the WWW. In a unipartite directed network, the HITS algorithm assigns two scores to each node in a self-consistent fashion. Score h – referred to as node hub-centrality score – is large for nodes that point to many authoritative nodes. The other score a – referred to as node authority-centrality score – is large for nodes that are pointed by many hubs. The corresponding equations for node scores a and h read

$$\begin{aligned} \mathbf{a} &= \alpha \mathbf{A} \mathbf{h}, \\ \mathbf{h} &= \beta \mathbf{A}^T \mathbf{a}, \end{aligned} \quad (29)$$

where α and β are parameters of the method. The previous equations can be rewritten as

$$\begin{aligned} \mathbf{A} \mathbf{A}^T \mathbf{a} &= \lambda \mathbf{a}, \\ \mathbf{A}^T \mathbf{A} \mathbf{h} &= \lambda \mathbf{h}, \end{aligned} \quad (30)$$

where $\lambda = (\alpha\beta)^{-1}$. The resulting score vectors \mathbf{a} and \mathbf{h} are thus eigenvectors of the matrices $\mathbf{A} \mathbf{A}^T$ and $\mathbf{A}^T \mathbf{A}$, respectively. The HITS algorithm has been used, for example, to rank publications in citation networks by the literature search engine Citeseer [124]. In the context of citation networks, it is natural to identify topical reviews as hubs, since they contain many references to influential papers in the literature. A variant of the HITS algorithm where a constant additive term is added to both node scores is explored by Ng et al. [125] and results in a more stable ranking of the nodes with respect to small perturbations in network structure. An extension of HITS to bipartite networks [126] is presented in Section 2.7.1.

2.6. A case study: node centrality in the Zachary’s karate club network

In this subsection, we illustrate the differences of the above-discussed node centrality metrics on the example of the small Zachary karate club network [127]. The Zachary Karate Club network captures friendships between the members of a US karate club (see Fig. 1 for its visualization). Interestingly, the club has split in two parts as a result of a conflict between its the instructor (node 1) and the administrator (node 34); the network is thus often used as one of a test cases for community detection methods [128, 129]. The rankings of the network’s nodes produced by various node centrality metrics are shown in Fig. 1. The first thing to note here that nodes 1 and 34 are correctly identified as being central to the network by all metrics. The k -shell metric produces integer values distributed over a limited range (in this network, the smallest and largest k -shell value are 1 and 4, respectively) which leads to many nodes obtaining the same rank (a highly degenerate ranking).

⁸ A matrix \mathbf{Q} is column sub-stochastic if and only if $\sum_i Q_{ij} \leq 1$.

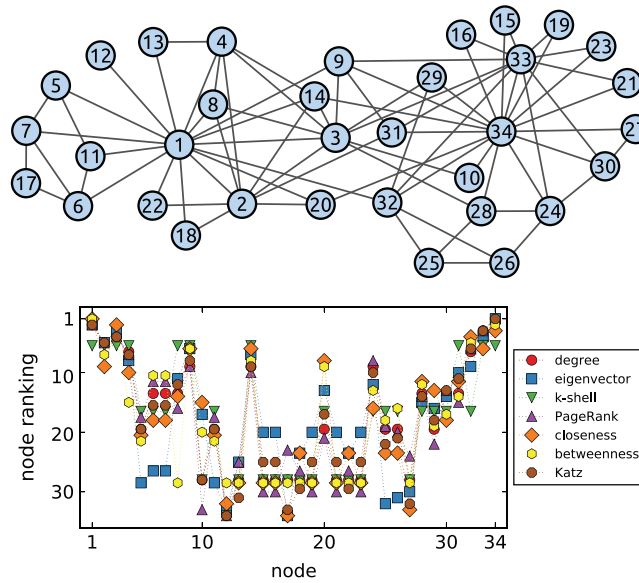


Fig. 1. *Top:* The often-studied Zachary’s karate club network has 34 nodes and 78 links (here visualized with the Gephi software). *Bottom:* Ranking of the nodes in the Zachary karate club network by the centrality metrics described in this section. Node labels on the horizontal axis correspond to the node labels in the top panel.

While most nodes are ranked approximately the same by all metrics, there are also nodes for which the metrics produced widely disparate results. Node 8, for example, is ranked high by k -shell but low by betweenness. An inspection of the node’s direct neighborhood reveals why this is the case: node 8 is connected to high-degree nodes 1, 2, 3, and 4, which directly assures its high k -shell value of 4. At the same time, those neighboring nodes themselves are connected, which creates shortcuts circumventing node 8 and implies that node 8 lies on very few shortest paths. Further, node 10 has low degree and thus also low PageRank score (among the evaluated metrics, these two have the highest Spearman’s rank correlation). By contrast, its closeness is comparatively high as the node is centrally located (e.g., the two central nodes 1 and 34 can be reached from node 10 in two and one steps, respectively). Finally, while PageRank and eigenvector centrality are closely related and their values are quite closely correlated (Pearson’s correlation 0.89), they still produce rather dissimilar rankings (their Spearman’s correlation 0.68 is the lowest among the evaluated metrics).

2.7. Static ranking algorithms in bipartite networks

There are many systems that are naturally represented by *bipartite networks*: users are connected with the content that they consumed online, customers are connected with the products that they purchased, scientists are connected with the papers that they authored, and many others. In a bipartite network, two groups of nodes are present (such as the customer and product nodes, for example) and links exist only between the two groups, not within them. The current state of a bipartite network can be captured by the network’s adjacency matrix B , whose element $B_{i\alpha}$ is one if node i is connected to node α and zero otherwise. Note that to highlight the difference between the two groups of nodes, we use Latin and Greek letters, respectively, to label them. Since a bipartite network’s adjacency matrix is a different mathematical object than a monopartite network’s adjacency matrix, we label them differently as B and A , respectively. While A is by definition a square matrix, B in general is not. Due to this difference, the adjacency matrix of a bipartite network is sometimes referred to as the *incidence matrix* in graph theory [26].

Since bipartite networks are composed of two kinds of nodes, one might be interested in scoring and ranking the two groups of nodes separately. Node-scoring metrics for bipartite networks thus usually give two vectors of scores, one for each kind of nodes, as output. In this Section, we present three ranking algorithms (coHITS, method of reflections, and the fitness-complexity algorithm) specifically aimed at bipartite networks, together with their possible interpretation in real socio-economic networks.

Before presenting the metrics, we stress that neglecting the connections between nodes of the same kind can lead to a considerable loss of information. For example, many bipartite networks are embedded in a social network of the participating users. In websites like [Digg.com](#) and [Last.fm](#), users can consume content and at the same time they can also select other users as their friends and potentially be influenced by their friends’ choices [130]. Similarly, the scholarly network of papers and their authors is naturally influenced by personal acquaintances and professional relationships among the scientists [131]. In addition, more than two network layers are often necessary to properly account for different types of interactions between the nodes [132,133]. We focus here on the simplest settings where a bipartite network representation of the input

data provides sufficiently insightful results, and refer to [134–136] for representative examples of ranking algorithms for multilayer networks.

2.7.1. Co-HITS algorithm

Similarly to the HITS algorithm for monopartite networks, co-HITS algorithm [126] works with two kinds of scores. Unlike HITS, co-HITS assigns one kind of scores to nodes in one group and the other kind of scores to nodes in the other group; each node is thus given one score. Denoting the two groups of nodes as 1 and 2, respectively, and their respective two score vectors as $\mathbf{x} = \{x_i\}$ and $\mathbf{y} = \{y_\alpha\}$, the general co-HITS equations take the form

$$x_i = (1 - \lambda_1)x_i^0 + \lambda_1 \sum_{\alpha} w_{i\alpha}^{21} y_\alpha, \quad y_\alpha = (1 - \lambda_2)y_\alpha^0 + \lambda_2 \sum_i w_{\alpha i}^{12} x_i. \quad (31)$$

Here λ_1 and λ_2 are analogues of the PageRank’s teleportation parameter and x_i^0 and y_α^0 are provide baseline, in general node-specific, scores that are given to all nodes regardless of the network topology. The summation terms represent the redistribution of scores from group 2 to group 1 (through $w_{i\alpha}^{21}$) and vice versa (through $w_{\alpha i}^{12}$). The transition matrices w^{21} and w^{12} as well as the baseline score vectors \mathbf{x}^0 and \mathbf{y}^0 are column-normalized, which implies that the final scores can be obtained by iterating the above-specified equations without further normalization.

Similarly to HITS, equations for x_i and y_α can be combined to obtain a set of equations where only node scores of one node group appear. Connections between the original iterative framework and various regularization schemes, where node scores are defined through optimization problems motivated by the co-HITS equations, are explored by Deng et al. [126]. One of the studied regularization frameworks is eventually found to yield the best performance in a query suggestion task when applied to real data.

2.7.2. Method of reflections

The method of reflections (MR) was originally devised by Hidalgo and Hausmann to quantify the competitiveness of countries and the complexity of products based on the network of international exports [30]. Although the metric can be applied to any bipartite network, in order to present the algorithm, we use the original terminology where countries are connected with the products that they export. The method of reflections is an iterative algorithm and its equations read

$$k_i^{(n)} = \frac{1}{k_i} \sum_{\alpha} B_{i\alpha} k_{\alpha}^{(n-1)}, \quad k_{\alpha}^{(n)} = \frac{1}{k_{\alpha}} \sum_i B_{i\alpha} k_i^{(n-1)} \quad (32)$$

where $k_i^{(n)}$ is the score of country i at iteration step n , and $k_{\alpha}^{(n)}$ is the score of product α at step n . Both scores are initialized with node degree ($k_i^{(0)} = k_i$ and $k_{\alpha}^{(0)} = k_{\alpha}$). In the original method, a threshold value is set, and when the total change of the scores is smaller than this value, the iterations stop. The choice of the threshold is important because the scores converge to a trivial fixed point [137]. The threshold has to be big enough so that rounding errors do not exceed the differences among the scores, as discussed in [137,138]. An eigenvector-based definition was later given by Hausmann et al. [139], which yields the same results as the original formulation and it has the advantage of making the threshold choice unnecessary.

Another drawback of the iterative procedure is that while one might expect that subsequent iterations “refine” the information provided by this metric, Cristelli et al. [138] point out that the iterations of the metric discard information instead of refining it. Furthermore, Mariani et al. [140] noticed that stopping the computation after two iterations maximizes the agreement between the country ranking by the MR score and the country ranking by their importance for the structural stability of the system. While the country scores by the MR have been shown to provide better predictions of the economical growth of countries as compared to traditional economic indicators [139], both the economic interpretation of the metric and its use as a predictive tool have raised some criticism [31,138]. The use of this metric as a predictor of GDP will be discussed in Section 7.2.

2.7.3. Fitness-complexity metric

Similarly to the method of reflections, the Fitness-Complexity (FC) metric aims to simultaneously measure the competitiveness of countries (referred to as country fitness) and the products’ level of sophistication (referred to as product complexity) based on international trade data [31]. The basic thesis of the algorithm is that competitive countries tend to diversify their export basket, whereas sophisticated products tend to be exported only by the most diversified (and thus the most fit) countries. The rationale of this thesis is that the production of a given good requires a certain set of diverse technological capabilities. More complex products require more capabilities to be produced and, for this reason, they are less likely to be produced by developing countries whose productive system has a limited range of capabilities. This idea is described more in detail and supported by a toy model in [138].

The metric’s assumptions can be represented by the following iterative set of equations

$$F_i^{(n)} = \sum_{\alpha} B_{i\alpha} Q_{\alpha}^{(n-1)}, \quad Q_{\alpha}^{(n)} = \frac{1}{\sum_i B_{i\alpha} / F_i^{(n-1)}} \quad (33)$$

where F_i and Q_α represent the fitness of country i and the complexity of product α , respectively. After each iterative step n , both sets of scores $\{F_i^{(n)}\}$ and $\{Q_\alpha^{(n)}\}$ are further normalized by their average value $F^{(n)}$ and $Q^{(n)}$, respectively.

To understand the effects of non-linearity on product score, consider a product α_1 with two exporters i_1 and i_2 whose fitness values are 0.1 and 10, respectively. Before normalizing the scores, product α_1 achieves the score of 0.099 which is largely determined by the score of the less-fit country. By contrast, a product α_2 that is only exported by country i_2 achieves the score of 10 which is much higher than the score of product α_1 . This simple example demonstrates the economic interpretation of the metric. First, if there is a low-score country that can export a given product, the complexity level of this product is likely to be low. By contrast, if only high-score countries are able to export the product, the product is presumably difficult to be produced and it should be thus assigned a high score. By replacing the $1/F_i$ terms in (33) with $1/F_i^\gamma$ ($\gamma > 0$), one can tune the role of the least-fit exporter in determining the product score. As shown in [140], when the exponent γ increases, the ranking of the nodes better reproduces their structural importance, but at the same time it is more sensitive to noisy data.

Variants of the algorithm have been proposed in order to penalize more heavily products that are exported by low-fitness countries [140,141] and to improve the convergence properties of the algorithm [142]. Importantly, the metric has been shown [140,143] to outperform other existing centrality metrics (such as degree, method of reflections, PageRank, and betweenness centrality) in ranking the nodes according to their importance for the structural stability of economic and ecological networks. In particular, in the case of ecological plant–pollinator networks [143], the “fitness score” F of pollinators can be interpreted as their importance, and the “complexity score” Q of plants as their vulnerability. When applied to plant–pollinator networks, the fitness–complexity algorithm (referred to as *MusRank* by Dominguez et al. [143]) reflects the idea that important insects pollinate many plants, whereas vulnerable plants are only pollinated by the most diversified – “generalists” in the language of biology – insects.

2.8. Rating-based ranking algorithms on bipartite networks

In the current computerized society where individuals separated by hundreds or thousands kilometers who have never met each other can easily engage in mutual interactions, reputation systems are crucial in creating and maintaining the level of trust among the participants that is necessary for the system to perform well [144–146]. For example, buyers and sellers in online auction sites are asked to evaluate each others’ behavior and the obtained information is used to obtain their trustworthiness scores. User feedback is typically assumed in the form of ratings in a given rating scale, such as the common 1–5 star system where 1 star and 5 stars represent the worst and best possible rating, respectively (this scale is employed by the important e-commerce websites such as Amazon and Netflix, for example).

Reputation systems have been shown to help the buyers avoid fraud [147] and the sellers with high reputation fare on average better than the sellers with low reputation [148]. Note that the previously discussed PageRank algorithm and other centrality metrics discussed in Section 2.5 also represent particular ways of assessing the reputation of nodes in a network. In that case, however, no explicit evaluations are required because links between nodes of the network are interpreted as implicit endorsements of target nodes by the source nodes.

The most straightforward way to aggregate the ratings collected by a given object is to compute their arithmetic mean (this method is referred to as *mean* below). However, such direct averaging is rather sensitive to noisy information (provided by the users who significantly differ from the prevailing opinion about the object) and manipulation (ratings intended to influence the resulting ranking of objects) [149,150]. To enhance the robustness of results, one usually introduces a reputation system where reputation of users is determined along with the ranking of objects. Ratings by little reputed users are then assumed to have potentially adverse affect on the final results, and they are thus given correspondingly low weights which helps to limit their impact on the system [144,145,151].

A particularly simple reputation system, *iterative refinement* (IR), has been introduced by Laureti et al. [152]. In IR, a user’s reputation score is inversely proportional to the average difference between the ratings given by this user and the estimated quality values of the corresponding objects. The estimates of user reputation and object quality are iterated until a stable solution is found. We denote the rating given by user i to object α as $r_{i\alpha}$ and the number of ratings given by user i as $k_i = \sum_\alpha B_{i\alpha}$ where $B_{i\alpha}$ are elements of the bipartite network’s adjacency matrix. The above-mentioned arithmetic mean can thus be represented as $Q_\alpha = \sum_i r_{i\alpha} B_{i\alpha} / k_\alpha$ where the ratings of all users are assigned the same weight. The iterative refinement algorithm [152] assigns users weights w_i and estimates the quality of object α as

$$Q_\alpha = \frac{\sum_i w_i B_{i\alpha} r_{i\alpha}}{\sum_i w_i B_{i\alpha}} \quad (34)$$

where the weights are computed as

$$w_i = \left(\varepsilon + \frac{1}{k_i} \sum_\alpha B_{i\alpha} (r_{i\alpha} - Q_\alpha)^2 \right)^{-\beta} . \quad (35)$$

Here $\beta \geq 0$ controls how strongly are the users penalized for giving ratings that differ from the estimated object quality and ε is a small parameter that prevents the user weight from diverging in the (unlikely) case when $r_{i\alpha} = Q_\alpha$ for all objects evaluated by user i .

Eqs. (34) and (35) represent an interconnected system of equations that can be solved, similarly to HITS and PageRank, by iterations (the initial user weights are usually assumed to be identical, $w_i = 1$). When $\beta = 0$, IR simplifies to arithmetic mean. As noted by Yu et al. [153], while $\beta = 1/2$ provides better numerical stability of the algorithm as well as translational and scale invariance, $\beta = 1$ is equivalent to a maximum likelihood estimate of object quality assuming that the individual “rating errors” of a user are normally distributed with unknown user-dependent variance. As shown by Medo and Wakeling [154], the IR’s efficacy is importantly limited by the fact that most real settings use a limited range of integer values. Yu et al. [153] generalize the IR by assigning a weight to each individual rating.

There are two further variants of IR. The first variant is based on computing a user’s weight as the correlation between the user’s ratings and the quality estimates of the corresponding objects [155]. The second variant [156] is the same but it suppresses the weight of users who have rated only a few items, and further multiplies the estimated object quality with $\max_{x_i \in \mathcal{L}_i} w_i$ (i.e., the highest weight among the users who have rated a given object). The goal of this modification is to suppress the objects that have been rated by only a few users (if an item is rated well by one or two users, it can appear at the top of the object ranking which is in most circumstances not a desired feature).

3. The impact of network evolution on static ranking algorithms

Many systems grow (or otherwise change) in time and it is natural to expect that this growth leaves an imprint on the metrics that are computed on static snapshots of these systems. This imprint often takes form of a time bias that can significantly influence the results obtained with a metric: nodes may score low or high largely due to the time at which they entered the system. For example, a paper published a few months ago is destined to rank badly by citation count because it has not have the time yet to show its potential by attracting a corresponding number of citations. In this section, we describe a few examples of systems where various forms of time bias exist and, importantly, discuss how the bias can be reduced or even entirely removed.

3.1. The first-mover advantage in preferential attachment and its suppression

Node in-degree, introduced in Section 2.2, is the simplest node centrality metric in a directed network. It relies on the assumption that if a node has been of interest to many nodes that have linked to it, the node itself is important. Node out-degree, by contrast, reflects the activity of a node and thus it cannot be generally interpreted as the node’s importance. Nevertheless, node in-degree too becomes problematic as a measure of node importance when the nodes are not created equal. Perhaps the best example of this is a growing network where nodes appear gradually. Early nodes then enjoy a two-fold advantage over late nodes: (1) they have more time to attract incoming links, (2) in the early phase of the network’s growth, nodes face less competition because there are fewer nodes to choose from. Note that any advantage given to a node is often further magnified by preferential attachment [70,157], also dubbed as cumulative advantage, which is present in many real systems [158,159].

We assume now the usual preferential attachment setting where at every time step, one node is introduced in the system and establishes directed links to a small number m of existing nodes. The probability of choosing node i with in-degree k_i is assumed to be proportional to $k_i + C$, where $C > 0$ is a model parameter. Nodes are labeled with their appearance time; node i has been introduced at time i . The current network size n can be thus used as the time counter. The continuum approximation (see [160] and Section VII.B in [73]) is the simplest approach to study the evolution of the node mean in-degree $\langle k \rangle_i(n)$. The change of $\langle k \rangle_i(n)$ in time step n is equal to the probability of receiving a link at time n , implying

$$\frac{d\langle k \rangle_i(n)}{dn} = m \frac{\langle k \rangle_i(n) + C}{\sum_j k_j(n) + C} \quad (36)$$

where the multiplication with m is due to m links being introduced in each time step. Since each of those links increases the in-degree of one node by one, the denominator can be exactly computed as

$$\sum_j [k_j(n) + C] = mn + Cn. \quad (37)$$

The resulting differential equation for $\langle k \rangle_i(n)$ can be then easily solved to show that the expected in-degree of node i at step n , $\overline{k_i(n)}$, is

$$\overline{k_i(n)} = \overline{k(t_i)} = C[t_i^{-1/(1+C/m)} - 1], \quad (38)$$

where we introduced the “relative birth time” $t_i := i/n$ of node i . Newman [77] used the master equation formalism [72] to find the same result as well as to compute higher moments of $k_i(n)$, and the divergence of $\overline{k_i(n)} := \overline{k(t)}$ for $t \rightarrow 0$ (corresponding to early papers in the thermodynamic limit of the system size) has been dubbed as the “first-mover advantage”. The generally strong dependence of $\overline{k(t)}$ on t for any network size n implies that only the first handful of nodes have the chance to become exceedingly popular (see Fig. 2 for an illustration). Parameters m and C can be in principle determined from the data. In the case of the APS citation data analyzed by Newman [77], for example, the best fit of the overall indegree distribution is obtained with $C = 6.38$ and $m = 22.8$ (note that the value of C is substantially larger than the often-assumed $C = 1$).




Fig. 2. (A) Cumulative node indegree distributions of directed preferential attachment networks grown from a single node to 10^6 nodes at different parameter values. Indicative straight lines have slope $1 + C/m$ predicted for the cumulative degree distribution by theory. (B) The probability that the first node has the largest in-degree in the network (estimated from 10,000 model realizations).

To counter the strong advantage of early nodes in systems under influence of the preferential attachment mechanism, Newman [77] suggests to quantify a node's performance by computing the number of standard deviations by which the node's in-degree exceeds the mean indegree corresponding to the node's appearance time, a quantity of this kind is commonly referred to as the z-score (see Section 4.1 for a general discussion of this approach). The z-score can be thus used to identify the nodes that "beat" preferential attachment and attract more links than one would expect based on their time of arrival. While this is not clear from the published manuscript [77], the preprint version available at <https://arxiv.org/abs/0809.0522> makes it clear that paper z scores are actually obtained by evaluating the mean indegree, $\mu_i(k^{in})$, and the standard deviation of indegree, $\sigma_i(k^{in})$, empirically by considering "a Gaussian-weighted window of width 100 papers around the paper of interest". The z-score of paper i , referred to as the rescaled indegree score $R_i(k^{in})$ in this review, is thus defined as

$$R_i(k^{in}) = \frac{k_i^{in} - \mu_i(k^{in})}{\sigma_i(k^{in})}. \quad (39)$$

Computing μ_i and σ_i analytically as done before in (38) would require knowing accurately the model behind the data which, despite the progress in network modeling [76,161,162], is not feasible. Since a very recent paper can score high on the basis of a single citation, a minimum citation count of five has been imposed in a later analysis in [163]. How best to counter the noise due to young papers and how to choose the evaluation window shape and size is yet to be studied.

The resulting z score was found to perform well in the sense of selecting the nodes with arrival times rather uniformly distributed over the system's lifespan, and the existence of those papers is said to provide "a hopeful sign that we as scientists do pay at least some attention to good papers that come along later" [77]. The selected papers have been revised by Newman [163] five years later and they have been found to outperform randomly drawn control groups with the same prior citation counts. An analogous approach is used in the following Section to correct the bias of PageRank. A general discussion of rescaling techniques for static centrality metrics is presented in Section 4.1.

3.2. PageRank's temporal bias and its suppression

As explained in Section 2.5, the PageRank score of a node can be interpreted as the probability that a random walker is found at the node during its hopping on the directed network. The network's evolution strongly influences the resulting score values. For example, a recent node that has had little time to attract incoming links is likely to achieve a low score. The corresponding bias of PageRank against recent nodes has been documented in the World Wide Web data, for example [164]. Such bias can potentially override the useful information derived by the algorithm and render the resulting scores useless or even misleading. Understanding PageRank's biases and the situations in which they appear is therefore essential for the use of the algorithm in practice.

Bias detection. While PageRank's bias against recent nodes can be understood as a natural consequence of these nodes still lacking the links that they will have the chance to collect in the future, temporal features of the network's growth can further pronounce the bias. The network of citations among scientific papers provides a particularly striking example because a paper can only cite papers that have been published in the past. All links in the citation network thus aim back in time. This particular network feature makes it hard for random walkers on the network to reach recent nodes (the only way this can happen is through the PageRank's teleportation mechanism) and ultimately results in a strong bias towards old nodes that has been reported in the literature [56,165].

Mariani et al. [82] parametrized the space of network features by assuming that the decay of node relevance (which determines the rate at which a node receives new incoming links) and the decay of node activity (which determines the rate at which a node creates new outgoing links) occur at generally different time scales Θ_R and Θ_A , respectively. This

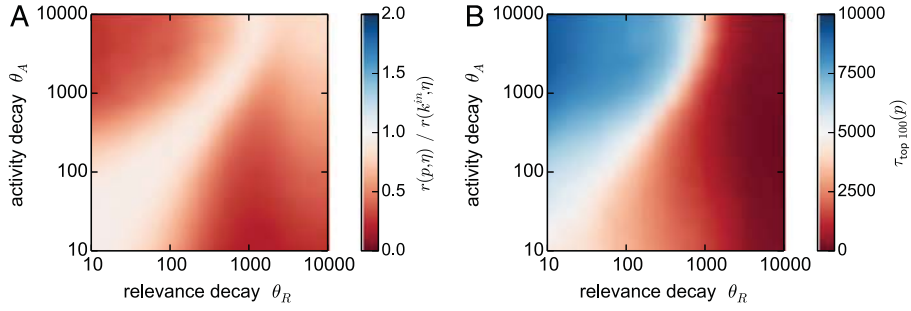


Fig. 3. Performance of PageRank in growing preferential attachment networks with 10,000 nodes where both node relevance and node activity decay. (A) After computing PageRank for the resulting networks, one can evaluate the correlation between node PageRank score and node fitness as well as the correlation between node in-degree and node fitness. This panel shows that when $\theta_A \gg \theta_R$ or $\theta_A \ll \theta_R$, PageRank correlates with node fitness less than in-degree (i.e., PageRank uncovers node properties worse than in-degree). (B) The average appearance time of top 100 nodes in the ranking by PageRank, $\tau_{\text{top } 100}(p)$, shows that PageRank strongly favors recent nodes when $\theta_A \gg \theta_R$ and it favors old nodes when $\theta_A \ll \theta_R$. The parameter regions where PageRank is strongly biased are those where it is outperformed by in-degree. Source: (Adapted from [82]).

parametrization is built on the fitness model with aging that has been originally developed to model citation data [74]. The main finding by Mariani et al. [82] is that when the two time scales have similar magnitude, the average time span of the links that aim forward and backward in time is approximately the same, and the random walk that underlies the PageRank algorithm can thus uncover the nodes' intrinsic fitness (see Fig. 3 for an illustration). When $\theta_R \gg \theta_A$, the network's links tend to point *backward* in time (i.e., from recent to older nodes) and, as a result, strong bias towards old nodes can emerge. As a consequence of this bias, PageRank captures the intrinsic node fitness worse than node indegree. PageRank is similarly outperformed by node in-degree when $\theta_R \ll \theta_A$ for the very opposite reason as before: in this parameter setting, the network's links tend to point *forward* in time (i.e., from old to more recent nodes) and, as a consequence, PageRank favors recent nodes over the old ones.

Bias suppression. Having found that the scores produced by the PageRank algorithm may be biased by node age, the natural question that emerges is whether the bias can be somehow removed. While several modifications of the PageRank algorithm have been proposed to attenuate the advantage of old nodes (they will be the main subject of Section 4), a simple method to effectively suppress the bias is to compare each node's score with the scores of nodes of similar age, in the same spirit as the z-score for citation count introduced in the previous Section. This has been addressed by [166] where the authors study the directed network of citations among scientific papers—a system where the bias towards the old nodes is particularly strong.⁹ Mariani et al. [166] propose to compute the PageRank score of all papers and then rescale the scores by comparing the score of each paper with the scores of other papers published short before and short after. In agreement with the findings presented by Parolo et al. [167], the window of papers included in the rescaling of a paper's score is best defined on the basis of the number of papers in the window (the authors use the window of 1000 papers). Denoting the PageRank score of paper i as p_i and mean and standard deviation of PageRank scores in the window around paper i as $\mu_i(p)$ and $\sigma_i(p)$, respectively, the proposed rescaled score of node i reads

$$R_i(p) = \frac{p_i - \mu_i(p)}{\sigma_i(p)}. \quad (40)$$

Note that unlike the input PageRank scores, the rescaled score can also be negative; this happens when a paper's PageRank score is lower than the average for the papers published in the corresponding window. Rescaled PageRank scores, $R_i(p)$, depend on the choice of window size; this dependence is however rather weak [166].

Section 4.1 provides more details on rescaled PageRank, while its application to the identification of milestone papers in scholarly citation data is presented in Section 7.1.1. In addition to these results and the results presented in [166], we illustrate here (Fig. 4) the metric's ability to actually remove PageRank's time bias in the model data that were used to produce Fig. 3. Fig. 4A shows that the average age $\tau_{100}(R(p))$ of the top 100 nodes as ranked by $R(p)$ is essentially independent of model parameters, especially when compared with the broad range of $\tau_{100}(R(p))$ observed for the ranking by PageRank in Fig. 3A. Fig. 4B compares the correlation between the rescaled PageRank score and node fitness with that between the PageRank score and node fitness. In the region $\theta_R \approx \theta_A$, where PageRank is not biased towards nodes of specific age, there is no advantage to gain by using rescaled PageRank. There are nevertheless extended regions (the top left corner and the right side of the studied parameter range) where the original PageRank is heavily biased and the rescaled PageRank is thus able to better reflect the intrinsic node fitness. One can conclude that the proposed rescaling procedure removes time bias in model data and that this removal improves the ranking of nodes by their significance. The bias of indegree and PageRank

⁹ In the context of the model discussed in the previous Section, $\theta_A = 0$ in the citation network because a paper can establish links to other papers only at the moment of its appearance; it thus holds that $\theta_R \gg \theta_A$ and the bias towards old nodes is indeed expected.




Fig. 4. Performance of rescaled PageRank $R(p)$ in artificial growing networks with 10,000 nodes where both node relevance and node activity decay. (A) The average age of the top 100 nodes as ranked by rescaled PageRank, $\tau_{100}(R(p))$. The small range of values (note the color scale range and compare with Fig. 3B) confirms that the age bias of PageRank is effectively removed by rescaling. (B) Rescaled PageRank outperforms PageRank in uncovering node fitness in most of the investigated parameter space.

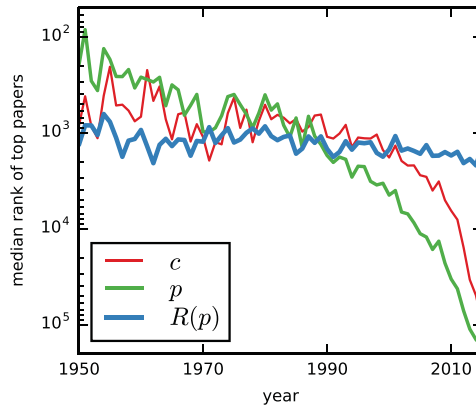


Fig. 5. For the APS citation data from the period 1893–2015 (560,000 papers in total), we compute the ranking of papers according to various metrics—citation count c , PageRank centrality p (with the teleportation parameter $\alpha = 0.5$), and rescaled PageRank $R(p)$. The figure shows the median ranking position of the top 1% of papers from each year. The three curves show three distinct patterns. For c , the median rank is stable until approximately 1995; then it starts to grow because the best young papers have not yet reached sufficiently high citation counts. For p , the median rank grows during the whole displayed time period because PageRank applied on an acyclic time-ordered citation network favors old papers. By contrast, the curve is approximately flat for $R(p)$ during the whole period which confirms that the metric is not biased by paper age and gives equal chances to all papers.

in real data as well as its removal by rescaled PageRank are exemplified in Fig. 5. Further benefits of removing the time bias in real data are presented in Section 7.1.1.

3.3. Illusion of influence in social systems

We close this Section with an example of an edge ranking problem where the task is to rank connections in a social network by the level of social influence that they represent. In the modern information-driven era, the study of how information and opinions spread in the society is more important than ever. The study of social influence is very active with networks playing an important role [168–171]. A common approach to assess the strength of social influence in a social network where every user is connected with their friends is based on measuring the probability that user i collects item α if $f_{i\alpha}$ of i 's friends have already collected it [172–174]. Here $f_{i\alpha}$ is usually called *exposure* as it quantifies how much a user has been exposed to an item (the underlying assumption is that each friend who has collected item α can tell user i about it). Note that the described setting can be effectively represented as a multilayer network [133] where users form a (possibly directed) social network and, at the same time, participate in bipartite user–item network which captures their past content consumption.

However, the described measurement is influenced by preferential attachment that is present in many social and information networks. Even when no social influence takes place and thus the social network has no influence on the items collected by individual users, the probability that a user has collected item α is k_α/U where k_α is the degree of item α and U is the total number of users. Denoting the number of friends of user i as f_i (this is equal to the degree of user i in the social network), there are on average $f_i k_\alpha/U$ friends who have collected item α . At the same time, if preferential attachment is in effect in the system, the probability that user i himself collects item α is also proportional to the item degree k_α . Positive correlation between the likelihood of collecting an item and exposure is therefore bound to exist even when no real social influence takes place.

The normalized exposure

$$n_{i\alpha} = \max_{t < t_{i\alpha}} \left(\frac{f_{i\alpha}(t)}{f_i k_\alpha(t)} k_{\min} \right) \quad (41)$$

was proposed by Vidmer et al. [130] to solve this problem. Here $k_\alpha(t)$ is the degree of item α at time t , $t_{i\alpha}$ is the time when user i has collected item α (if user i has not collected item α at all, the maximum is taken over all t values), and k_{\min} is the minimum required item degree.¹⁰ In Eq. (41), the division by $k_\alpha(t)$ is intended to factor out the impact of preferential attachment on exposure.

Social influence measurements based on exposure and normalized exposure are compared on both real and model data by Vidmer et al. [130]. When standard exposure, $f_{i\alpha}$, is used, real data show social influence even when they are randomized—a clear indication that the bias introduced by the presence of preferential attachment is too strong to be ignored. By contrast, the positive correlation between item collection probability and normalized exposure, $n_{i\alpha}$, disappears when the input data are randomized. In parallel, when no social influence is built in the model data, item collection probability is correctly independent of normalized exposure whereas it grows markedly with standard exposure. By taking time into account and explicitly factoring out the influence of preferential attachment, normalized exposure avoids the problems of the original exposure metric. In the future, the logic could be reversed and normalized exposure could be used to rank items by their likelihood to be actually consumed by a given user.

4. Time-dependent ranking algorithms

In the previous section, we have presented important shortcomings of static metrics when applied to evolving systems, and score-rescaling methods to suppress these biases. However, a direct rescaling of the final scores is not the only way to deal with the shortcomings of static metrics. In fact, a substantial amount of research has been devoted to ranking algorithms that explicitly include time in their defining equation. These time-dependent algorithms mostly aim at suppressing the advantage of old nodes and thus improving the chance of recent nodes to acquire visibility. Assuming that the network is unweighted, time-dependent ranking algorithms take as input the adjacency matrix $A(t)$ of the network at a given time t and information on node age and/or edge time-stamp. They can be classified into three (not all-inclusive) categories: (1) node-based rescaled scores (Section 4.1); (2) metrics that explicitly penalize older nodes or older edges (Section 4.2); (3) metrics based on models of network growth that assume the existence of a latent fitness parameter, which represents a proxy for the node’s success in the system (Section 4.3). Time-dependent generalizations of the static reputation systems introduced in Section 2.8 are also discussed in this section (Section 4.2.6).

4.1. Striving for time-balance: Node-based time-rescaled metrics

As we have discussed in the previous section, network centrality metrics can be heavily biased by node age when applied to evolving networks. In other words, large part of the variation in node score is explained by node age, which may be an undesirable property of centrality score if we are interested in untangling the intrinsic fitness, or quality, of the nodes. In this Section we present metrics that explicitly require that node score is not biased by node age. These metrics simply take the original static centrality scores as input, and rescale them by comparing each node’s score with only the scores of nodes of similar age.

Rescaling indegree. In Section 3, we have seen (Fig. 5) that in real information networks, node indegree can be biased by node age due to the preferential attachment mechanism. This effect can be particularly strong for citation networks, as discussed in [77,163] and in Section 3. In [77], the bias of paper citation count (node in-degree) in the scientific paper citation network is removed by evaluating the mean $\mu_i(c)$ and standard deviation $\sigma_i(c)$ of the citation count for papers published in a similar time as paper i . The systematic dependence of c_i on paper appearance time is then removed by computing the rescaled citation count

$$R_i(c) = \frac{c_i - \mu_i(c)}{\sigma_i(c)}. \quad (42)$$

Note that $R_i(c)$ represents the z-score of paper i within its evaluation window. Values of $R(c)$ larger or smaller than zero indicate whether the paper is out- or under-performing, respectively, with respect to papers of similar age. There exists no principled criterion to choose the papers “published in a similar time” over which $\mu_i(c)$ and $\sigma_i(c)$ are computed. In the case of citation data, one can constrain this computation to a fixed number, Δ , of papers published just before and just after paper i [166], or weight papers score with a Gaussian function centered on the focal paper [77]. A statistical test in [166] shows that in the APS data, differently from the ranking by citation count, the ranking by rescaled citation count $R(c)$ is not biased by paper age.

A different rescaling equation has been used in [175,176], where the citation count of each scientific paper i is divided by the average citation count of papers published in the same year and belonging to the same scientific domain as paper i .

¹⁰ Without this constraint, items with low degree can reach high normalized exposure through normalization with $k_\alpha(t)$; such estimates based on a few events are noisy and thus little useful.

The resulting rescaled score, called c_f in [175], has been found to follow a universal lognormal distribution – independent of paper field and age – in the Web of Science citation data. This finding has been challenged by subsequent studies [177–179], which leaves it open whether an unbiased ranking of academic publications can be achieved through direct rescaling of citation count. Since this review mostly focus on temporal effects, we refer the interested reader to the review article by Waltman [17] for details on a different field-normalization procedure introduced in the bibliometrics literature. As there are various possibilities how to rescale indegree, one could also use a reverse-engineering approach: instead of using a rescaling equation and verifying whether it leads to an unbiased ranking of the nodes, one can directly require that the scores of papers of different age follow the same distribution and then infer the correct rescaling equation [180]. However, this procedure has been only applied to suppress the bias by field of papers published in the same year (see [180] for details), and it remains unexplored for the time-bias suppression problem.

Rescaling PageRank score. Differently from the indegree’s bias towards old nodes, PageRank can develop a temporal bias towards either old or recent nodes (see Section 3.2 and [82]). As shown in [166] and illustrated here in Fig. 4, the algorithm’s temporal bias can be suppressed with the equation

$$R_i(p) = \frac{p_i - \mu_i(p)}{\sigma_i(p)}, \quad (43)$$

which is analogous to Eq. (42) above. Assuming that the nodes are labeled in order of decreasing age, μ_i and σ_i are computed over a “moving” temporal window $[i - \Delta/2, i + \Delta/2]$ centered on paper i . Similarly to what was observed for $R(c)$, a statistical test demonstrates that in the APS data, the ranking by rescaled PageRank R_c , unlike the ranking by PageRank, is not biased by paper age [166]. The performance of rescaled PageRank and rescaled indegree in identifying groundbreaking papers will be discussed in detail and compared with the performance of other ranking methods in Section 7.1.1. Vaccario et al. [179] generalized the rescaling procedure based on the z-score to suppress the bias of scientific papers’ PageRank score by both age and field.

We emphasize that in principle any structural centrality metric can be rescaled with an equation akin to Eq. (43). The only parameter of the rescaled score is the size Δ of the temporal window used for the calculation of μ_i and σ_i ; Mariani et al. [166] point out that too large Δ values results in a ranking biased towards old nodes in a similar way as the ranking by PageRank,¹¹ whereas too small Δ values results in a ranking heavily affected by statistical fluctuations: a node that happens to have only low-degree papers in its small temporal window would score high. One should always avoid both extremes when using a rescaled metric; however, a statistically-grounded guideline to choose Δ is still lacking.

4.2. Metrics with explicit penalization for older edges and/or nodes

We now consider variants of degree (Section 4.2.1) and of eigenvector-based centrality metrics (Sections 4.2.2–4.2.5) where edges and/or nodes are weighted according to their age directly in the equation that defines node score. Hence, differently from rescaled metrics that *a posteriori* normalize the scores by static centrality metrics, the metrics presented in this Section directly include time information in the score calculation.

4.2.1. Time-weighted degree and its use in predicting future trends

Indegree (or degree) is the simplest method to rank nodes in a network. As we have seen in Section 3.1, node degree is biased by age in growing networks that exhibit preferential attachment. As a consequence, degree can fail in individuating the nodes that will become popular in the future [181], since node preferences may shift over time [182] and node attractiveness may decay over time [74]. One can consider a variant of degree where received links are weighted by a function of the edge age $\Delta t_{ij} := t - t_{ij}$, where t and t_{ij} denote the time at which the scores are computed and the time at which the edge from j to i was created, respectively. The time-weighted degree score $s_i(t)$ of node i at time t can be defined as

$$s_i(t) = \sum_j A_{ij}(t) f(t - t_{ij}), \quad (44)$$

where $f(\Delta t)$ is a function of edge age $\Delta t_{ij} = t - t_{ij}$. Zeng et al. [181] set $f(\Delta t) = 1 - \lambda \Theta(\Delta t - T_r)$, where $\Theta(\cdot)$ denotes the Heaviside function – $\Theta(x)$ is equal to one if $x \geq 0$, zero otherwise – and T_r is a parameter of the method. The resulting score $s_i^{\text{rec}}(t; T_r)$ (referred to as Popularity-Based Predictor PBP in [181, 183]) is thus a convex combination of node degree $k_i(t)$ and node recent degree increase $\Delta k_i(t, T_r) = \sum_j A_{ij} \Theta(T_r - (t - t_{ij}))$,

$$s_i^{\text{rec}}(t; T_r) = k_i(t) - \lambda k_i(t - T_r) = (1 - \lambda) k_i(t) + \lambda \Delta k_i(t, T_r). \quad (45)$$

The parameter T_r specifies the length of the “time window” over which the recent degree increase is computed. Note that s_i^{rec} reduces to node degree or recent degree increase for $\lambda = 0$ or $\lambda = 1$, respectively. Closely similar exponential penalization is considered by Zhou et al. [183], leading to

$$s_i^{\text{exp}}(t) = \sum_j A_{ij}(t) e^{-\gamma(t-t_{ij})} \quad (46)$$

¹¹ The ranking by PageRank and the ranking by rescaled PageRank are identical for $\Delta \geq N$.

which is referred to as Temporal-Based Predictor in [181,183]; here $\gamma > 0$ is a parameter of the method which is analogous to T_r before. Results obtained with time-weighted degree scores s^{rec} and s^{exp} on data from Netflix, Movielens and Facebook show that: (1) s^{rec} with λ close to one performs better than degree in identifying the most popular nodes in the future [181]; (2) this performance is further improved by s^{exp} (Fig. 3 in [183]). Metric s^{exp} (albeit with a different parameter calibration) is also used in [184], where it is referred to as the Retained Adjacency Matrix (RAM). s^{exp} is shown in [184] to outperform indegree, PageRank and other static centrality metrics in ranking scientific papers according to their future citation count increase and in identifying the top-cited papers in the future. The performance of s^{exp} is comparable to that of a time-dependent variant of Katz centrality, called Effective Contagion Matrix (ECM), that is discussed below (Section 4.2.2).

Another type of weighted degree is the *long-gap citation count* [16] which measures the number of edges received by a node when it is at least T_{min} years old. Long-gap citation count has been applied to the network of citations between American movies and has been shown to be a good predictor for the presence of movies in the list of milestone movies edited by the National Film Registry (see [16] for details).

4.2.2. Penalizing old edges: Effective Contagion Matrix

The Effective Contagion Matrix centrality is a modification of the classical Katz centrality where the paths that connect a node with another node are not only weighted according to their length, but also according to the ages of their links [184]. For a tree-like network (like citation networks), one defines a modified adjacency matrix $R(t, \gamma)$ (referred to as retained adjacency matrix by Ghosh et al. [184]) whose elements are $R_{ij}(t, \gamma) = A_{ij}(t) \gamma^{\Delta t_{ij}}$, where Δt_{ij} is the age of the edge $j \rightarrow i$ at the time t when the scores are computed. The aggregate score s_i of node i is defined as

$$s_i(t) = \sum_{k=0}^{\infty} \alpha^k \sum_{j=1}^N [R(t, \gamma)^k]_{ij}. \quad (47)$$

According to this definition, a path of length k thus has a weight which is given not only by α^k (as for the Katz centrality), but also by the age of the edges that compose the path. Using the geometric series of matrices, and denoting by \mathbf{e} the vector whose components are all equal to one, one can rewrite Eq. (47) as

$$\mathbf{s}(t) = \sum_{k=0}^{\infty} \alpha^k R(t, \gamma)^k \mathbf{e} = (1 - \alpha R(t, \gamma))^{-1} \mathbf{e}. \quad (48)$$

This can be used to write

$$s_i(t) = \alpha \sum_j [R(t, \gamma)]_{ij} s_j(t) + 1 = \alpha \sum_j A_{ij}(t) \gamma^{\Delta t_{ij}} s_j(t) + 1, \quad (49)$$

which shows that the score of a node is mostly determined by the scores of the nodes that recently pointed to it. ECM score has been found [184] to outperform other metrics (degree, weighted indegree, PageRank, age-based PageRank and CiteRank – see below for CiteRank’s definition) in ranking scientific papers according to their future citation count increase, and in identifying papers that will become popular in the future; the second best-performing metric is weighted indegree as defined in Eq. (46). Ghosh et al. [184] also note that the agreement between the ranking and the future citation count increase can be improved by using individual metrics as features in a support vector regression model. The analysis in [184] does not include TimedPageRank (described in the next Section); in future research, it will be interesting to compare the predictive power of all the existing metrics, and compare it with the predictive power of network growth models [76] (see Section 7.1.2).

4.2.3. Penalizing old edges: TimedPageRank

Similarly to the methods presented in the previous Section, one can introduce time-dependent weights of edges in PageRank’s definition. A natural way to enforce this idea was proposed by Yu et al. [185] in the form

$$s_i(t) = c \sum_j A_{ij}(t) f(\Delta t_{ij}) \frac{s_j(t)}{k_j^{out}} + 1 - c, \quad (50)$$

where Δt_{ij} denotes the age of the link between nodes i and j .¹² Yu et al. [185] used Eq. (50) with $f(\Delta t) = 0.5^{\Delta t}$, and defined node i ’s TimedPageRank score as the product between its score s_i determined by Eq. (50) and a factor $a(t - t_i) \in [0.5, 1]$ which decays with node age $t - t_i$. A similar idea of an explicit aging factor to penalize older nodes was also implemented by Baeza-Yates [164]. Yu et al. [185] used scientific publication citation data to show that the top 30 papers by TimedPageRank are more cited in the future years than the top papers by PageRank.

¹² Note that since $\sum_j A_{ij} f(\Delta t_{ij}) / k_j^{out} \leq 1$ upon this definition, the vector of scores \mathbf{s} is not normalized to one. Eq. (50) thus defines a *pseudo-PageRank problem*, which can be put in a directed correspondence with a PageRank problem as defined in Section 2.5.5 (see Section 2.5.5 and [61] for more details).

4.2.4. Focusing on a temporal window: T-Rank, SARA

T-Rank [186,187] is a variant of PageRank whose original aim is to rank websites. The algorithm assumes that we are mostly interested in the events that happened in a particular temporal window of interest $[t_1, t_2]$. For this reason, the algorithm favors the pages that received many incoming links and were often updated within the temporal window of interest. Time-stamps are thus weighted according to their temporal distance from $[t_1, t_2]$, leading to a “freshness” function $f(t)$ which is equal to or smaller than one if the link’s creation time lies within the range $[t_1, t_2]$ or outside, respectively. If a timestamp t does not belong to the temporal window, its freshness $f(t)$ monotonously decreases with its distance from the temporal window of interest. The freshness of timestamps are then used to compute the freshness of nodes, of links and of page updates, which are the elements needed for the computation of the T-Rank scores. As the temporal window of interest can be chosen at will, the algorithm is flexible and can provide interesting insights into the web users’ reactions after an unexpected event, such as a terror attack [187].

A simpler idea is to use only the links that belong to a certain temporal window to compute the score of a node. This idea has been applied to the citation network of scientific authors to compute an author-level score for a certain time window using a weighted variant of PageRank called Science Author Rank Algorithm (SARA) [13]. The main finding by Radicchi et al. [13] is that the authors that won a Nobel Prize are better identified by SARA than by indices based on indegree (citation count). The temporal evolution of the ranking positions assigned to physics researchers by SARA can be explored at the interactive Web platform <http://www.physauthorsrank.org/authors/show>.

4.2.5. PageRank with time-dependent teleportation: CiteRank

In the previous Sections, we have presented variants of classical centrality metrics (degree, Katz centrality, and PageRank) where edges are penalized according to their age. For the PageRank algorithm, an alternative way to suppress the advantage of older nodes is to introduce an explicit time-dependent penalization of older nodes in the teleportation term. This idea has led to the CiteRank algorithm which was introduced by Walker et al. [65] to rank scientific publications. The vector of CiteRank scores \mathbf{s} can be found as the stationary solution of the following set of recursive linear equations¹³

$$s_i^{(n+1)}(t) = c \sum_j A_{ij}(t) \frac{s_j^{(n)}(t)}{k_j^{out}} + (1 - c) v(t, t_i) \quad (51)$$

where, denoting by t_i the publication date of paper i and t the time at which the scores are computed, we defined

$$v(t, t_i) = \frac{\exp[-(t - t_i)/\tau]}{\sum_{j=1}^N \exp[-(t - t_j)/\tau]}. \quad (52)$$

To choose the values of c and τ , the authors compare papers’ CiteRank score with papers’ future indegree increase Δk^{in} , and find $c = 0.5$, $\tau = 2.6$ years as the optimal parameters. With this choice of parameters, the age distribution of paper CiteRank score is in agreement with the age distribution of the papers’ future indegree increase. In particular, both distributions feature a two-step exponential decay (see Fig. 6). The theoretical analysis based on CiteRank in [65] suggests that the two-step decay of Δk^{in} can be explained by two distinct and co-existing citation mechanisms: researchers cite a paper because they found it either directly or by following the references of a more recent paper. This example shows that well-designed time-dependent metrics are not only useful tools to rank the nodes, but can also shed light into the behavior of the agents in the system.

4.2.6. Time-dependent reputation algorithms

A reputation system ranks individuals by their reputation from the most to the least reputed one based on their past interactions or evaluations. Many reputation systems weight recent evaluations more than old ones and thus produce a time-dependent ranking. For example, the generic reputation formula

$$R(t) = \sum_n W_n \frac{f(t_n, t)}{\sum_m f(t_m, t)} \quad (53)$$

presented by Sabater and Sierra [188] assigns evaluation $W_n \in [-1, 1]$ (here -1 , 0 , and $+1$ represent absolutely negative, neutral, and absolutely positive evaluation, respectively, and n is the evaluation counter) weight proportional to $f(t_n, t)$ where t_i is the evaluation time and t is the time when reputation is computed. The weight function $f(t_n, t)$ should be chosen to favor evaluations that were made short before t ; a simple example of such a function is $f(t_n, t) = t_n/t$. Sabater and Sierra [188] proceed by analyzing an artificial scenario where a user first behaves reliably until reaching a high reputation value and then starts to commit fraud. Thanks to assigning high weight to recent actions, their reputation system is able to quickly reflect the change in the user’s behavior and correspondingly lower the user’s reputation. Instead of weighting the evaluations on the basis of the time when they were made, it is also possible to order evaluations by their time stamps $t(W_1) < t(W_2) < \dots < t(W_N)$ and then assign weight λ^{N-n} to evaluation n . This scheme is referred to as *forgetting* as it corresponds to gradually lowering the weight of evaluations as more recent evaluations arrive [189]. When $\lambda = 1$, no

¹³ This definition is slightly different from that provided in [65] – the scores resulting from the definition adopted here are normalized to one and differ from those obtained with the definition in [65] by a uniform normalization factor.

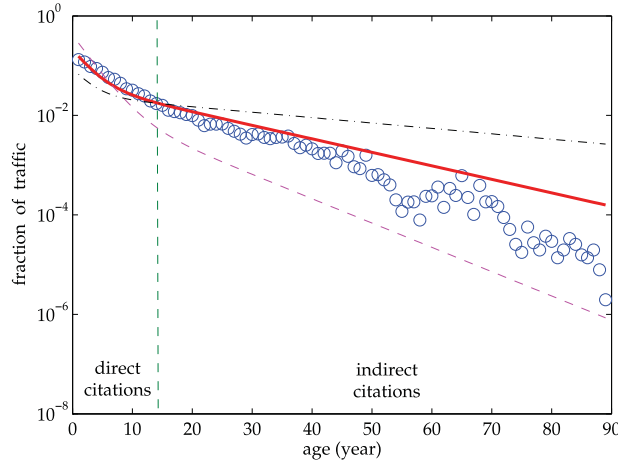


Fig. 6. (Reprinted from [65].) A comparison of the age distribution of future incoming citations (red line) and of CiteRank score (blue dots). Both curves exhibit a two-step exponential decay. According to the analytic results by Walker et al. [65], this behavior can be interpreted in terms of two distinct citation mechanisms (see the text).

forgetting is effectively employed. When $\lambda = 0$, the most recent evaluation alone determines the result. Note that the time information is also extensively used in otherwise time-unaware reputation systems to detect spammers and other anomalous and potentially malicious behavior [190]. Interestingly, there is little research of reputation systems where both network structure and time play important roles.

4.3. Model-based ranking of nodes

The ranking algorithms described up to now make assumptions that are regarded as plausible and thus they are rarely tested on real systems. For instance, indegree assumes that a node is important, or central, if it is able to acquire many incoming links; PageRank assumes that a node is important if it is pointed by other important nodes; rescaled PageRank assumes that a node is important if its PageRank score exceeds PageRank scores of nodes of similar age. To evaluate the degree to which these node centrality definitions are reliable, one can *a posteriori* evaluate the performance of the ranking algorithms in identifying high-quality nodes in the network or in predicting the future evolution of the system (examples of both approaches will be discussed in Section 7).

On the other hand, if one has a mathematical model which well describes how the network evolves, then one can assess node importance by simply fitting the model to the given data. This section will review this approach in detail. We will focus on preferential-attachment models where each node is endowed with a fitness parameter [71], which represents the perceived importance of the node by other nodes in the network.

Ranking based on the fitness model. Bianconi and Barabási [71], extended the original Barabási–Albert model by assigning to each node an intrinsic quality parameter, called *fitness*. Node fitness represents the nodes' inherent ability to acquire new incoming links. In the directed-network version of the model (hereafter *fitness model*), the probability $\Pi_i(t)$ that a new incoming link will be received by node i at time t is given by

$$\Pi_i(t) \sim (k_i^{\text{in}}(t) + C) \eta_i, \quad (54)$$

where η_i represents node i 's fitness and C is an additive constant. By denoting as t_i the appearance time of node i , this attachment rule leads to the following equation for the dependence of the expected node degree on time

$$\langle k_i^{\text{in}}(t) \rangle = m \left(\frac{t}{t_i} \right)^{\beta(\eta_i)}, \quad (55)$$

where if the support of the fitness distribution is bounded, $\beta(\eta)$ is solution of the self-consistent equation of the form $\beta(\eta) = \eta / C[\beta(\eta)]$ and $C[\beta]$ only depends on the chosen fitness distribution (see [71] for the details). Remarkably, the model is flexible and can reproduce several degree distributions by appropriately choosing the functional form of the fitness distribution [71].

Using Eq. (55), we can estimate node fitness in real data. To do this, it is sufficient to follow the degree evolution $k^{\text{in}}(t)$ of the nodes and properly fit $k^{\text{in}}(t)$ to a power-law to obtain the exponent β , which is proportional to node fitness. This procedure has been applied by Kong et al. [191] to data from the WWW, leading to two interesting conclusions: (1) websites' fitness distribution is narrow; (2) fitness distribution is not dependent on time, and can be well approximated by an exponential distribution.




Fig. 7. An illustration of the relevance decay for movies posted in different years (reprinted from [192]). Time is divided into equal-duration time windows (“rating snapshots” in the horizontal axis), and the average relevance $r(t)$ within each snapshot for movies released in a given year is shown on the vertical axis. The left panel shows the average relevance decay for MovieLens movies released from 2000 to 2008; the right panel shows the same for Netflix movies released between 2000 and 2005. The snapshot time interval is one month in Netflix, 102 days in MovieLens respectively.

Ranking based on the relevance model. While the fitness model describes the growth of information networks better than the original Barabási–Albert model [71,191], it is still an incomplete description of degree evolution for many evolving systems. The missing element in the fitness model is node aging [193] which for diverse reasons occurs in most information networks. For example, the ideas contained in scientific papers are incorporated in subsequent papers, reviews and books, which results in a gradual decay of papers’ attractiveness to new citations. A today’s surprising news will in a few days become old and irrelevant.

The *relevance model* introduced by Medo et al. [74] incorporates node aging by assuming

$$\Pi_i(t) \sim (k_i^{in}(t) + C) \eta_i f(t - t_i), \quad (56)$$

where $f(t - t_i)$ is a function of node age which tends to zero [74,76] or to a small constant value [194] for large age. For an illustration of node relevance decay in real data from MovieLens¹⁴ and Netflix,¹⁵ see Fig. 7 and [192]. Assuming that the normalization factor of $\Pi_i(t)$ converges to a finite value Ω^* , one can use Eq. (56) to show that

$$\langle k_i^{in}(t) \rangle \simeq \exp\left(\frac{\eta_i}{\Omega^*} \int_0^t f(t') dt'\right). \quad (57)$$

The value of Ω^* depends on the distribution of node fitness and the aging function (see [74] for details).

This model can be fitted to real citation data in different ways. Medo et al. [74] computed the relevance $r_i(t)$ of node i at time t as the ratio between the fraction of links $l_i(t)$ obtained by node i in a relatively short time window and the fraction of links $l_i^{PA}(t)$ expected from pure preferential attachment. That is, one writes

$$r_i(t) = \frac{l_i(t)}{l_i^{PA}(t)} \quad (58)$$

where $l_i^{PA}(t) = k_i^{in}(t) / \sum_j k_j^{in}(t)$. A fitness estimate for node i is then provided by the node’s total relevance $T_i = \sum_t r_i(t)$; similarly to the fitness distribution found in the WWW [191], the total relevance distribution has an exponential decay in the APS citation data [74]. Alternatively, one can estimate node fitness through standard maximum likelihood estimation (MLE) [194]. Wang et al. [76] use MLE to analyze the citation data from the APS and Web of Science, and find that the time-dependent factor $f(t)$ follows a universal lognormal form

$$f(t) = \frac{1}{\sqrt{2\pi t \sigma_i}} \exp\left[-\frac{(\log(t) - \mu_i)^2}{2 \sigma_i^2}\right] \quad (59)$$

where μ_i and σ_i are parameters specific to the temporal citation pattern of paper i . By fitting the relevance model with Eq. (59) to the data, the authors estimate paper fitness and use this estimate to predict the future impact of papers (we refer to [76] and Section 7.1.2 for more details). Medo [194] uses MLE to compare various network growth models with respect to their ability to fit data from the Econophysics forum by Yi-Cheng Zhang,¹⁶ and the Relevance Model turns out to fit the data better than other existing models.

¹⁴ Data available at <https://grouplens.org/datasets/movielens/10m/>.

¹⁵ Data available at <http://www.netflixprize.com/>.

¹⁶ A website for researchers interested in econophysics and other interdisciplinary branches of physics; see www.unifr.ch/econophysics.

Ranking based on a model of PageRank score evolution. In the previous Section, we discussed various procedures to estimate node fitness based on the same general idea: if we know the relation between node degree and node fitness, we can properly fit the relation to the data to infer the fitness values. In principle, the same idea can find application for network centrality metrics other than degree. For PageRank score s , for example, one can assume a growth model in the form [195]

$$s_i(t) = \exp\left[\int_0^t dt \alpha_i(t)\right], \quad (60)$$

where the behavior of the function $\alpha_i(t)$ has to be inferred from the data. Berberich et al. [195] assume that $\alpha_i(t)$ is independent of time within the system time span $[t_{min}, t_{max}]$ and show that this assumption is reasonable in the academic citation network they analyzed. With this assumption, Eq. (60) becomes

$$s_i(t) = s_i(t_0) \exp[\alpha_i(t - t_{min})], \quad (61)$$

where $\alpha_i(t) = \alpha_i$ if $t \in [t_{min}, t_{max}]$. A fit of Eq. (60) to the data allows us to estimate α_i , which is referred to as the BuzzRank score [195] of node i . Eq. (60) is conceptually similar to Eq. (57) for the RM: the growth of nodes' centrality score depends exponentially on a quantity – node fitness η in the case of the RM, node growth rate α in the case of BuzzRank – that can be interpreted as a proxy for node's success in the system. On the other hand, differently from Eq. (57), Eq. (60) does not explicitly rely on a network growth model.

5. Ranking nodes in temporal networks

In the previous sections, we have illustrated several real-world and model-based examples that show that the temporal ordering of interactions is a key element that cannot be neglected when analyzing an evolving network. For instance, we have seen that some networks exhibit a first-mover advantage (Section 3.1), a diffusion process on a network can be strongly biased by the network's temporal linking patterns (Section 3.2), and neglecting the order of interactions may lead to wrong conclusions when evaluating the social influence between two individuals (Section 3.3). Various approaches to solve the consequent drawbacks of static centrality metrics (such as indegree and PageRank) have been presented in Section 4. In particular, we have introduced metrics (such as the rescaled metrics and CiteRank) that take as input not only the network's adjacency matrix A , but also the time stamps of the edges and/or the age of the nodes. These approaches take into account the temporal dimension of the data. They are useful for their ability to suppress the bias of static metrics (see Section 3) and to highlight recent contents in growing networks (see Section 4).

However, in the previous sections, we have focused on *unweighted* networks, where at most one interaction occurs between two nodes i and j . This assumption is limiting in many real scenarios: for example, individuals in social networks typically interact repeatedly (through emails, mobile phone calls, text messages, etc.), and the number and order of interactions can bring important information on the strength of a social tie. In this section, we focus on systems where repeated node–node interactions can occur and the timing of interactions plays a critical role. To study this kind of systems, we adopt now the temporal networks' framework [196]. A *temporal network* $\mathcal{G}^{(temp)} = (N, E^{(temp)})$ is completely specified by the number N of nodes in the system and the list $E^{(temp)}$ of time-stamped edges (or contacts) (i, j, t) between the nodes.¹⁷ In the literature, temporal networks are also referred to as temporal graphs [197,198], time-varying networks/graphs [199], dynamic networks/graphs [199], scheduled networks [200], among others.

5.1. Getting started: how to represent a temporal network?

Recent studies [44,79] have pointed out that projecting a temporal network $\mathcal{G}^{(temp)}$ to a static weighted (time-aggregated) network discards the essential information and, as a result, often yields misleading results. To properly deal with the temporal nature of temporal networks, there are (at least) two possible strategies. The first is to revisit the way in which the network representation of the data is constructed. While in principle several representation representations of temporal networks can be adopted (see [201] for an overview), a simple and useful representation is the *higher-order network representation* of the temporal network, (also referred to as *memory network representation* in the literature). Differently from the time-aggregated representation which only preserves the number of contacts between pairs of nodes for a given temporal network, a higher-order network representation [44,79,202] transforms the edge list into a set of *time-respecting paths* (Section 5.2) on the network's nodes, and it preserves the statistics of time-respecting paths up to a certain topological length (details are provided below). Standard ranking algorithms (such as PageRank or closeness centrality) can be thus run on higher-order networks (Section 5.3). Higher-order networks are particularly powerful in highlighting the effects of memory on the speed of network diffusion processes (Section 5.3.4) and on the ranking of the nodes (Sections 5.3.5 and 7.1.4).

The second possible strategy is to divide the system's time span into temporal slices (or layers) of equal duration, and consider a set $\{w(t)\}$ of adjacency matrices, one for each layer – we refer to this representation as the *multilayer representation*

¹⁷ For some systems, each edge is also endowed with a temporal duration. For our purposes, by assuming that there exist a basic time unit u in the system, edges endowed with a temporal duration can be still represented through time-stamped edges. We simply represent an edge that spans the temporal interval $[t_1, t_2]$ as a sequence of $(t_2 - t_1)/u$ subsequent time-stamped edges.

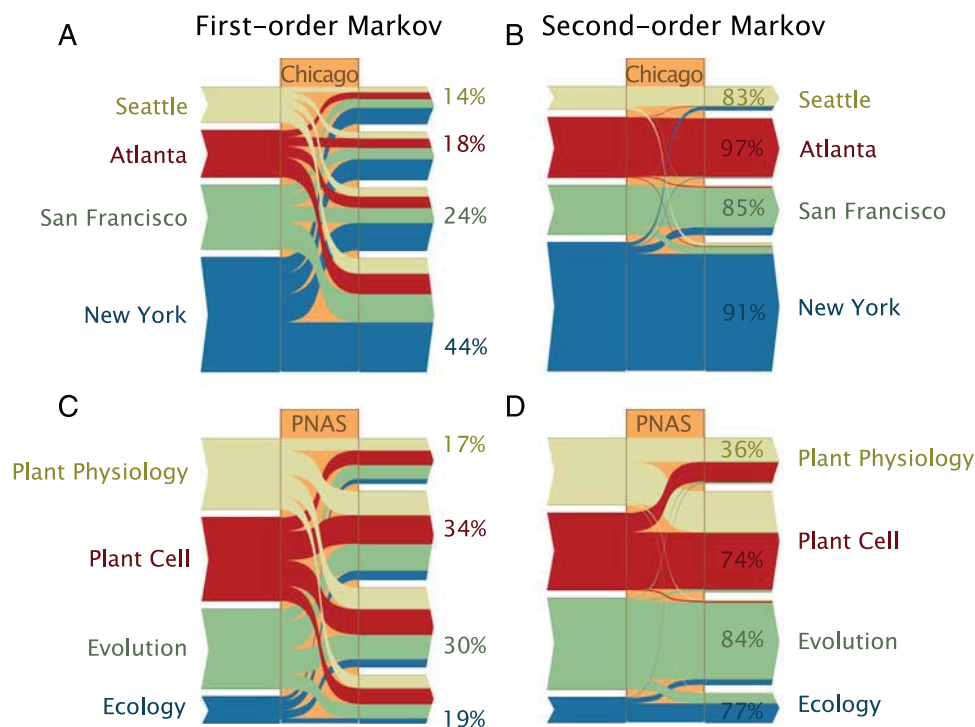


Fig. 8. An illustration of the difference between the first-order Markovian (time-aggregated) and second-order network representation of the same data (reprinted from [44]). Panels A–B represent the destination cities (the right-most column) of flows of passengers from Chicago to other cities, given the previous location (the left-most column). When including memory effects (panel B), the fraction of passengers coming back to the original destination is large, in agreement with our intuition. A similar effect is found for the network of academic journals (panels C–D)– see [44] and Section 7.1.4 for details.

of the original temporal network [132]. The original temporal network is then represented by a number of static networks where usual network analysis methods can be employed. Note that the information encoded in the multilayer representation is exactly the same as that present in the original temporal network only if the temporal span of each layer is equal to the basic time unit of the original dataset. The problem of choosing a suitable time span for the temporal layers is a non-trivial and dataset-dependent problem. In this review, we do not address this issue; the reader can find interesting insights into this problem in [201,203,204], among others. In the following, when presenting methods based on the multilayer representation of temporal networks, we simply assume that each layer’s duration has been chosen appropriately and, as a result, we have a meaningful partition of the original data into temporal layers. Diffusion-based ranking methods based on this construction are described in Section 5.4. We also present generalizations to temporal networks of centrality metrics based on shortest paths (Section 5.5).

5.2. Memory effects and time-preserving paths in temporal networks

Considering the order of events in a temporal network has been shown to be fundamental to properly describe how different kinds of flow processes (e.g., flow of information, flow of traveling passengers) unfold on a temporal network. In this section, we present a paradigmatic example that has been provided by Rosvall et al. [44] and highlights the importance of memory effects in a temporal network. The example is followed by a general definition of time-preserving paths for temporal networks, which constitute a building block for the causality-preserving diffusive centrality metrics that are discussed in Section 5.3.

An example. The example shown in Fig. 8 is taken from [44] and well exemplifies the role of memory in temporal networks. The figure considers two datasets: (1) the network of cities connected by airplane passengers’ itineraries and (2) the JSTOR¹⁸ network of scientific journals connected by citations between journals’ articles. We consider first the network of cities. The Markovian assumption [44] is that a flow’s destination depends only on the current location, and not on its past history. According to this assumption, if we consider the passengers who started from Seattle and passed through Chicago, for example, the fraction of them coming back is only 14% (Fig. 8(A)). One might conjecture that this is a poor description of real passenger traffic as travelers are usually based in a city and are likely to come back to the same city at the end of their

¹⁸ www.jstor.org

itineraries. To capture this effect, one needs to consider a “second-order” description of the flow. This simply means that instead of counting how many passengers move from Chicago to Seattle, we count how many move to Seattle, given that they started in a certain city. We observe that the fraction of passengers who travels from Chicago to Seattle, given that they reached Chicago from Seattle, is 83%, which is much larger than what was expected under the Markovian assumption (14%). Similar considerations can be made for the other cities (Fig. 8A–B). To capture memory effects in the citation network of scientific journals, Rosvall et al. [44] analyze the statistics of triples of journals $j_1 \rightarrow j_2 \rightarrow j_3$ resulting from triples of articles $a_1 \rightarrow a_2 \rightarrow a_3$ such that a_1 cites a_2 and a_2 cites a_3 , and paper a_i is published in journal j_i . Based on this construction, Rosvall et al. [44] show that memory effects heavily impact flows of citations between scientific journals (Fig. 8C–D). We can conclude that including memory is critical for the prediction of flows in real world (see [202] and Section 5.3.5).

Time-respecting paths. The example above shows that when building a network representation from a certain dataset, memory effects need to be taken into account to properly describe the system. To account for memory, in the example of traveling passengers, we are no more interested in how many passengers travel between two locations, but in how many passengers follow a certain (time-ordered) itinerary across different locations. To enforce this paradigm shift, let us introduce a general formalism to describe time-preserving (or time-respecting, or time-ordered) paths [79,197,205]. Given a static network and three nodes a, b, c , the existence of the links (a, b) and (b, c) directly implies that a is connected to c through at least one path $a \rightarrow b \rightarrow c$ (*transitivity* property). The same does not hold in a temporal network [206,207]: for a node a to be connected with node c through a path $a \rightarrow b \rightarrow c$, the link (a, b, t_1) must happen *before* link (b, c, t_2) , i.e., $t_1 \leq t_2$. Time-aggregated networks thus overestimate the number of paths between pairs of nodes – Lentz et al. [206] present an elegant framework to quantify this property. The first general definition of time respecting path of length n is the following: $\{(i_1, i_2, t_1), (i_2, i_3, t_2), \dots, (i_n, i_{n+1}, t_n)\}$ is a time-respecting path if and only if $t_1 < t_2 < \dots < t_n$ [79,197,208,209].

In many practical applications, we are only interested in dynamical processes that occur at timescales much smaller than the total time over which we observe the network. For this reason, the existence of a time-respecting path $\{(i_1, i_2, t_1), (i_2, i_3, t_2)\}$ with a long waiting time between t_1 and t_2 may not be informative at all about the influence of node i_1 on i_3 . A practical solution is to set a threshold δ for the maximum acceptable inter-event time [79,210,211]. Accordingly, one says that $\{(i_1, i_2, t_1), (i_2, i_3, t_2), \dots, (i_n, i_{n+1}, t_n)\}$ is a time-respecting path (with limited waiting time δ) if and only if $\forall n : 0 < t_{n+1} - t_n \leq \delta$. The suitability of a particular choice of δ depends on the system being examined. While too small δ values can result in a too sparse network from which it is hard to extract useful information, too large δ values can produce also paths that are not consistent with the real flows of information. In practice, values as different as six seconds (used for pairwise interactions in an ant colony [79]) and one hour (used for e-mail exchanges between employees in a company [79]) can be used for δ , depending on the context.

5.3. Ranking based on higher-order network representations

In this Section, we use the notion of a time-preserving path presented in the previous Section to introduce centrality metrics built on higher-order representation of temporal networks. Before introducing the higher-order representations (Section 5.3.2), let us briefly revisit the Markovian assumption. Most of the methods presented in this Section are implemented in the Python library `pyTempNets`.¹⁹

5.3.1. Revisiting the Markovian assumption

When studying a diffusion process on a (weighted) time-aggregated network, standard network analysis implicitly assumes the *Markovian property*: the next move of a random walker does not depend on the walker’s previous steps, but only on the node i where he is currently located and on the weights of the links attached to node i . Despite being widely used in network centrality metrics, community detection algorithms and spreading models, this assumption can be unrealistic for a wide range of real systems (see Fig. 8, for example).

Let us revisit the Markovian assumption in general terms. In the most general formulation of a walk on a time-aggregated network, the probability $P_i^{(n+1)}$ that a random walker is located node i at diffusion time-step $n + 1$ depends on the whole past. Using the same notation as [44], if we denote by X_n the random variable which marks the position of the walker at time n , we write

$$P_i^{(n+1)} = P(X_{n+1} = i | X_n = i_n, X_{n-1} = i_{n-1}, \dots, X_1 = i_1, X_0 = i_0), \quad (62)$$

where $X_0 = i_0$ is the initial position of the walker and i_1, \dots, i_n are all his past positions. The Markovian assumption reads

$$P_i^{(n+1)} = P(X_{n+1} = i | X_n = i_n) = P(X_1 = i | X_0 = i_n), \quad (63)$$

which implies that the probability of finding a random walker at node i at step $n + 1$ is given by

$$P_i^{(n+1)} = \sum_j P(j \rightarrow i) P_j^{(n)} = \sum_j P_{ij} P_j^{(n)}, \quad (64)$$

¹⁹ <https://github.com/IngoScholtes/pyTempNets>

where

$$P_{ij} := P(j \rightarrow i) = P(X_{t+1} = i | X_t = j) = \frac{W_{ij}}{\sum_k W_{kj}} \quad (65)$$

are the elements of the (column-stochastic) transition matrix P of the network ($P_{ij} = W_{ij} / \sum_k W_{kj}$, where W_{ij} is the observed number of times a passenger traveled from node j to node i).²⁰

As we have seen in Section 2, PageRank and its variants build on the random walk defined by Eq. (64) and add an additional element, the teleportation term. While the time-dependent PageRank variants defined in Section 4 are based on a different transition matrix than the original PageRank algorithm, they also implicitly assume the Markovian property. We shall now discuss how to relax the Markovian assumption and construct higher-order Markov models where also the previous trajectory of the walker influences where the walker will jump next.

5.3.2. A second-order Markov model

Following [44], we explain in this section how to build a second-order Markov process on a given network. Consider a random walker located at a certain node j . According to the second-order Markov model, the probability that the walker will jump from node j to node i at time $n + 1$ depends also on the node k where the walker was located at the previous step of the dynamics. In other words, the move $\vec{j}i$ from node j to i actually depends on the edge $\vec{k}j$. It follows that the second-order Markov model can be represented as a first-order Markov model on edges in the network. That is to say, the random walker jumps and locates not on nodes, but on edges.

The probability $P(j \rightarrow i, n + 1 | k \rightarrow j, n)$ that the walker performs a jump from node j to node i at time $n + 1$ after having performed a jump from node k to node j at time n can be written as

$$P(j \rightarrow i, n + 1 | k \rightarrow j, n) = P(\vec{k}j \rightarrow \vec{j}i) = \frac{W(\vec{k}j \rightarrow \vec{j}i)}{\sum_{\vec{j}i} W(\vec{k}j \rightarrow \vec{j}i)}, \quad (66)$$

where $W(\vec{k}j \rightarrow \vec{j}i)$ denotes the observed number of occurrences of the directed time-respecting path $\vec{k}j \rightarrow \vec{j}i$ in the data – with the notation introduced in Section 5.2, $\vec{k}j \rightarrow \vec{j}i$ denotes here any length-two path $(k, j, t_1) \rightarrow (j, i, t_2)$ such that $0 < t_2 - t_1 \leq \delta$. Eq. (66) allows us to study a non-Markov diffusion process that preserves the observed frequencies of time-respecting paths of length two. This non-Markov diffusion process can be interpreted as a Markov diffusion process on the *second-order representation* of the temporal network (also referred to as the *memory network* or *second-order network*) whose nodes (also referred to as *memory nodes* or *second-order nodes*) are the edges of the original network. For this reason, one can use standard static network tools to simulate diffusion and compute the corresponding node centrality (see below) on second (and higher) order networks.

The probability of finding the random walker at the memory node $\vec{j}i$ at time $n + 1$ is thus given by

$$P^{(n+1)}(\vec{j}i) = \sum_k P(\vec{k}j \rightarrow \vec{j}i) P^{(n)}(\vec{k}j). \quad (67)$$

The probability $P_i^{(n+1)}$ of finding the random walker at node i at time $n + 1$ is thus

$$P_i^{(n+1)} = \sum_j P^{(n+1)}(\vec{j}i) = \sum_{jk} P(\vec{k}j \rightarrow \vec{j}i) P^{(n)}(\vec{k}j). \quad (68)$$

The random walk based on the second-order model presented in this Section is much more predictable than a random walk based on a network without memory, in the sense that the next step of the walker, given its current position, has much smaller uncertainty as measured by conditional entropy (this holds for all the temporal-network datasets studied by [44]). An attentive reader may have noticed that while constructing the higher-order temporal network has the advantage to allow us to use standard network techniques, this comes at the price of an increased computational complexity. Indeed, in order to describe a non-Markovian process with a Markovian process, we had to increase the dimension of the system: the second-order network comprises $E > N$ memory nodes, where E is the number of edges that have been active at least once during the whole observation period. This also points out that the problem of increased complexity requires more attention for dense networks than for sparse networks.

5.3.3. Second-order PageRank

In this Section, we briefly discuss how the PageRank algorithm can be generalized to higher-order networks. The PageRank score $s(\vec{j}i)$ of a memory node $\vec{j}i$ can be defined as the corresponding component of the stationary state of the following

²⁰ Note that we assume here also that the random walk process is *stationary*, i.e., the transition probabilities do not depend on diffusion step n . This assumption still holds for the random walk process based on higher-order network representations presented below (Sections 5.3.2–5.3.5), whereas it does not hold for random walks based on multilayer representations of the underlying time-stamped data (Section 5.4).

process [44]

$$s(\vec{j}i, n+1) = c \sum_k s(\vec{k}j, n) P(\vec{k}j \rightarrow \vec{j}i) + (1-c) \frac{\sum_l W_{il}}{\sum_{lm} W_{lm}}, \quad (69)$$

where the term multiplied by $1-c$ is the component i of a “smart” teleportation vector (see Section 2.5.5). The second-order PageRank score of node i , s_i , is defined as

$$s_i = \sum_j s(\vec{j}i). \quad (70)$$

The effects of second-order flow constraints of the ranking of scientific journals by PageRank will be presented in Section 7.1.4.

5.3.4. The effects of memory on diffusion speed

From a physicist’s point of view, it is interesting to study how the temporal order of interactions affects the speed of diffusion on the network. Recent literature used standard models for spreading processes [44,212], simulated random walks [213], and graph Laplacian spectral properties [214], to show that considering the timing of interactions can slow down spreading.

However, diffusion can also speed up when memory effects are included. A general analytic framework introduced by Scholtes et al. [79] predicts whether including second-order memory effects will slow-down or speed-up diffusion. The main idea is to exploit the property that the convergence time of a random walk process is related to the second largest eigenvalue λ_2 of its transition matrix. First, one constructs the transition matrix $T^{(2)}$ associated to the second-order Markov model of the network. It can be proven [79], that for any stochastic matrix with eigenvalues $1 = \lambda_1, \lambda_2, \dots, \lambda_m$, and for any $\epsilon > 0$, the number of diffusion steps after which the distance between the vector $\mathbf{p}^{(n)}$ of visitation frequencies and the random-walk stationary vector \mathbf{p} becomes smaller than ϵ is proportional to $1/\log|\lambda_2|$. To establish whether second-order memory effects will slow down or speed up diffusion, one thus constructs a randomized second-order transition matrix $\tilde{T}^{(2)}$ where the expected frequency of physical paths of length two is consistent with the time-aggregated network. The second largest eigenvalue of the resulting matrix is denoted as $\tilde{\lambda}_2$. The analytic prediction for the change of diffusion speed is thus given by

$$S(\tilde{T}^{(2)}) := \frac{\log|\tilde{\lambda}_2|}{\log|\lambda_2|}. \quad (71)$$

Values of $S(\tilde{T}^{(2)})$ larger (smaller) than one predict slow-down (speed-up) of diffusion. The resulting predictions accurately match the results of diffusion simulations on real data (see Fig. 2 in [79]).

5.3.5. Generalization to higher-order Markov models and prediction of real-world flows

In Section 5.3.2, we have considered random walks based on a second order Markov model. It is natural to generalize this model to take into account higher order dependencies. If we are interested, for example, in a third-order Markov model, we could simply build a higher-order network whose memory nodes are triples of physical nodes [79]. The memory node $\vec{i}j\vec{k}$ would then represent a time-respecting path $i \rightarrow j \rightarrow k$. Xu et al. [202] consider a more general higher-order model where, after a maximum order of dependency is fixed, all orders of dependencies can coexist (see Fig. 9C). The authors use global shipping transportation data to show that by considering more than two steps of memory, the random walk is even more predictable, which results in an enhanced accuracy in out-of-sample prediction of real-world flows (Fig. 9D).

Interestingly, Xu et al. [202] also discuss the effects of including memory effects on ranking websites in the World Wide Web. Similarly to Eq. (70), the authors define the higher-order PageRank score of a given node i as the sum of the PageRank scores of higher-order nodes that represent node i . Fig. 10 shows that despite being positively correlated, standard and higher-order PageRank bring substantially different information on node importance. Since higher-order PageRank describes user traffic better than standard PageRank, one could conjecture that it provides a better estimate of website relevance; however, further research is needed to test this hypothesis and investigate possible limitations of the higher-order network approach.

The possibility of including all order of dependencies up to a certain order leads to a natural question: is there an optimal order to reliably describe a temporal-network dataset? Ideally, we would like our model to describe the system accurately enough with the smallest possible memory order; in other words, an ideal model would have good predictive power and, at the same time, low complexity. The problem of determining the optimal order is non-trivial and it is addressed by Scholtes [215] with maximum likelihood estimation techniques. The resulting statistical framework allows us to determine the optimal order k_{opt} to describe a temporal network, thereby distinguishing among systems for which the traditional time-aggregated network representation is satisfactory ($k_{opt} = 1$) and systems for which temporal effects play an essential role ($k_{opt} > 1$).




Fig. 9. An illustration of prediction of real-world ship trajectories between three locations A, B, and C (reprinted from [202]). A part of the data is used to build the higher-order network (training); the remaining part is used to test the predictive power of different network representations (testing, panel A). In the first-order (time-aggregated) network (panel B), the nodes represent the three locations and the directed edge from location A and location B is weighted with the number of occurrences of the displacement $A \rightarrow B$ (see Section 5.3.1). Differently from the higher-order networks used in [44,79], in the higher-order network introduced by Xu et al. [202], nodes of different order co-exist (panel C). Accounting for memory effects dramatically improves prediction accuracy (panel D). More details can be found in [202].

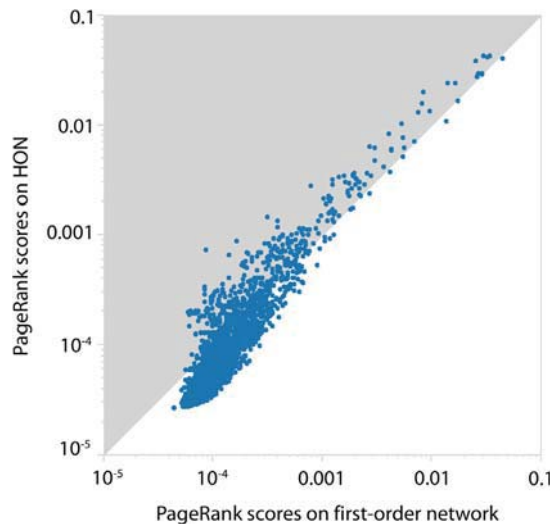


Fig. 10. The relation between higher-order PageRank and first-order PageRank score in WWW data (reprinted from [202]). As reported in [202], when higher-order effects are included, more than 90% of the webpages decrease their PageRank score. This may indicate that standard PageRank score overestimate the traffic on the majority of websites, which may mean that the distribution of the attention received by websites is more uneven than previously thought.

5.4. Ranking based on a multilayer representation of temporal networks

The higher-order Markov models presented above are particularly convenient because: (1) they preserve the temporal order of interactions up to a certain order; (2) centrality scores on the respective networks can be computed by means of standard network techniques. However, one can instead consider the temporal network as a sequence of temporal snapshots, each of which is composed of the set of links that occur within a temporal window of given length [45,136,214]. Within this framework, a temporal network composed of N nodes and of total temporal duration T is represented by L adjacency matrices $\{w(t)\}$ ($t = 1, \dots, L$) with dimension $N \times N$, where L is the number of temporal snapshots the total observation time is divided into. The element $w_{ij}(t)$ of the adjacency matrix for the snapshot t represents the number of interactions between i and j within the time period $[t, t + T/L)$. The set of adjacency matrices is typically referred to as an adjacency tensor [136,201].

5.4.1. TempoRank

Building on the random walk process proposed by Starnini et al. [213], Rocha and Masuda [45] design a random walk-based ranking method for undirected networks, called TempoRank. The method assumes that at each diffusion time t , the random walker only “sees” the network edges active within the temporal layer t ($w_{ij}(t) = 1$ for an edge (i, j) active at layer t), which means that there is an exact correspondence between the diffusion time and the real time – diffusion and network evolution unfold with the same time-scale. This is very different from diffusion-based static ranking algorithms which implicitly assume that diffusion takes place at a much faster timescale than the time scale at which the network changes.

With respect to the random walk processes studied in [213], TempoRank features a “sojourn” (a temporary stay) probability according to which a random walker located at node i has the probability $q^{k_i(t)}$ of remaining located at node i , where $k_i(t) = \sum_j w_{ij}(t)$ is the number of edges of node i in layer t . The (row-stochastic) transition matrix $P_{ij}(t)$ of the random walk at time t is given by

$$P_{ij}(t) = \begin{cases} \delta_{ij} & \text{if } k_i(t) = 0, \\ q^{k_i(t)} & \text{if } k_i(t) > 0, i = j, \\ (1 - q^{k_i(t)}) w_{ij}(t) / k_i(t) & \text{if } k_i(t) > 0, i \neq j. \end{cases} \quad (72)$$

It can be shown that a non-zero sojourn probability ($q > 0$) is sufficient for the convergence of the random walk process if and only if the time-aggregated network is connected; by contrast, if $q = 0$, the convergence of the random walk process to a stationary vector is not guaranteed even if the time-aggregated network is connected.

The need for a non-zero sojourn probability is motivated in Section 2.3 of [45] and well illustrated by the following example. Consider a network of three nodes 1, 2, 3 whose time-aggregated representation is triangular ($w_{12} = w_{23} = w_{13} = 1$). Say that only one edge is active at each temporal layer: $w_{12}(1) = w_{21}(1) = w_{23}(2) = w_{32}(2) = w_{13}(3) = w_{31}(3) = 1$, while all the other $w_{ij}(t)$ are equal to zero. For this simple temporal network, if $q = 0$, a random walker starting at node 1 will always end at node 1 after three steps, which implies that the random walk is not mixing. This issue is especially troublesome for temporal networks with fine temporal resolution for which there exist only a limited number of possible pathways for the random walker, and it is solved by imposing $q > 0$. We remark that by defining the sojourn probability as $q^{k_i(t)}$, we assume that the random decision whether the walker located at node i will move (probability q) or not (probability $1 - q$) is made independently for each edge.

It is important to study the asymptotic properties of the walk defined by Eq. (72). However, since TempoRank’s diffusion time is equal to the real-system time t , it is not a priori clear how the random walk should behave after L diffusion steps, *i.e.*, after the temporal end of the dataset is reached. To deal with this ambiguity, Rocha and Masuda [45] impose periodic boundary conditions on the transition matrix: $P(t + L) = P(t)$. With this choice, when the last temporal layer of the dataset $t = L$ is reached, the random walker continues with the oldest temporal layer $t = 1$. The random walk process that defines the TempoRank score thus runs cyclically multiple times on the dataset. Below, we denote by m the number of cycles performed by the random walk on the dataset; we are interested in the properties of the random walk in the limit of large m . The one-cycle transition matrix $\mathbf{P}^{(temp,1c)}$ is defined as

$$\mathbf{P}^{(temp,1c)} = \prod_{t=1}^L P(t) = P(1) \dots P(L). \quad (73)$$

The TempoRank vector after n time steps is thus given by the leading left eigenvector²¹ $\mathbf{s}(n)$ of matrix $\mathbf{P}^{(temp)}(n)$, where $\mathbf{P}^{(temp)}(n)$ is given by the (periodic) time-ordered product of $n = mL + t$ transition matrices $P(t)$. This eigenvector corresponds to the stationary density of the random walk after m cycles over the dataset and t additional temporal layers. Even when m is large, the vector $\mathbf{s}(n)$ fluctuates within a cycle due to the displacements of the walker along the network temporal links.

²¹ We choose the left-eigenvector here because the transition matrix defined by Eq. (72) is row-stochastic.

The TempoRank vector of scores \mathbf{s} is thus defined as the average within one cycle of the leading eigenvector of the matrix $\mathbf{P}^{(temp)}(n)$ for large m

$$\mathbf{s} = \frac{1}{L} \sum_{t=1}^L \mathbf{s}(mL + t). \quad (74)$$

As shown in [45], the TempoRank score is well approximated by a local time-aware variable, referred to as the “in-strength approximator”, yet how this result depends on network’s structural and temporal patterns remains unexplored. Importantly, TempoRank scores are little correlated with the scores obtained with the time-aggregated network representation. This finding constitutes yet another evidence that node centrality computed with a temporal network representation can substantially differ from node centrality computed with the time-aggregated representation, which is in a qualitative agreement with the findings obtained with higher-order Markov models [44,79].

5.4.2. Coupling temporal layers

Taylor et al. [136] generalize the multilayer approach to node centrality by including the coupling between consecutive temporal network layers. The aim of this coupling is to connect each node with itself between two consecutive layers and to rule the magnitude of variations of node centrality over time. Each node’s centrality for a certain temporal layer is thus dependent not only on the centrality of its neighbors in that layer, but also on its centrality in the precedent and subsequent temporal layer. The method can be in principle applied to any structural centrality metric. A full description of the mathematics of the method involves tensorial calculations and can be found in [136]. The method has been applied to rank mathematics departments and the Supreme Court decisions. In the case of mathematics departments, the regime of strong coupling between layers represents the basic idea the prestige of a mathematics department should not widely fluctuate over time. This method provides us with an analytic framework to identify not only the most central nodes in the network, but also the “top-movers”, *i.e.*, the nodes whose centrality is rising most quickly.

5.4.3. Dynamic network-based scoring systems for ranking in sport

In the context of sport ranking (see Section 2.5.3), Motegi and Masuda [216] have shown that both Park and Newman’s win–loss score [111] and Radicchi’s prestige score [14] can be generalized to include the temporal dimension. There are at least two main reasons to include time in sport ranking: (1) old performances are less indicative of the players’ current strength [217]; (2) it is more rewarding to beat a certain player when he/she is at the peak of his/her performance than doing so when he/she is a novice or short before retirement [216]. To account for the former effect, some time-aware variants of structural centrality metrics simply introduce time-dependent penalization for older interactions [217], in the same spirit as the time-dependent algorithms presented in Section 4.2.

In the dynamic win–loss score introduced by Motegi and Masuda [216], the latter effect is taken into account by assuming that the contribution to player i ’s score coming from a win against player j at time t is proportional to the score of player j at that time t . The equations that define the vector \mathbf{w} of win scores and the vector \mathbf{l} of loss scores are provided in [216]. By assuming that we have a reasonable partition of time into temporal layers $1, \dots, t$, the vector of dynamic win scores $\mathbf{w}(t)$ at time t is defined as²²

$$\mathbf{w}(t) = \mathbf{W}^T(t) \mathbf{e}, \quad (75)$$

where

$$\begin{aligned} \mathbf{W}^T(t) = & \mathbf{A}(t) + e^{-\beta} \sum_{m_t \in \{0,1\}} \alpha^{m_t} \mathbf{A}(t-1) \mathbf{A}(t)^{m_t} \\ & + e^{-2\beta} \sum_{m_t, m_{t-1} \in \{0,1\}} \alpha^{m_t + m_{t-1}} \mathbf{A}(t-2) \mathbf{A}(t-1)^{m_{t-1}} \mathbf{A}(t)^{m_t} \\ & + \dots \\ & + e^{-\beta(t-1)} \sum_{m_2, \dots, m_t \in \{0,1\}} \alpha^{\sum_{j=2}^t m_j} \mathbf{A}(1) \mathbf{A}(2)^{m_2} \dots \mathbf{A}(t)^{m_t}. \end{aligned} \quad (76)$$

The vector of loss scores $\mathbf{l}(t)$ is defined analogously by replacing $\mathbf{A}(t)$ by $\mathbf{A}^T(t)$. As in Section 2.5.3, the dynamic win–loss score $\mathbf{s}(t)$ is defined as the difference $\mathbf{s}(t) := \mathbf{w}(t) - \mathbf{l}(t)$.

We can understand the rationale behind this definition by considering the first terms of Eq. (76). The contributions to i ’s score that come from wins at a past time t' , with $t' \leq t$, are modulated by a factor $\exp[-\beta(t-t')]$, where β is a parameter of the method. Similarly to Park and Newman’s static win–loss score, also indirect wins determine player i ’s score; differently

²² Motegi and Masuda [216] set the temporal duration of each layer to one day. Other choices are possible in principle: for example, one may not want each tennis player’s score to change within a tournament – note that the Association of Tennis Professionals updates international players’ scores only 45 times per year [<http://www.atpworldtour.com/en/rankings/rankings-faq>, accessed 27.02.2017], which approximately corresponds to one ranking update per week.

from the static win-loss score, only indirect wins that happened *after* player's i respective direct win are included. With these assumptions, the contributions to player i 's score at time t that come from the two past time layers $t - 1$ and $t - 2$ are

$$\begin{aligned}
w_i(t) &= \sum_j A_{ij}(t) \\
&+ e^{-\beta} \sum_j A_{ij}(t-1) + \alpha e^{-\beta} \sum_{j,k} A_{ij}(t-1)A_{jk}(t) \\
&+ e^{-2\beta} \sum_j A_{ij}(t-2) + \alpha e^{-2\beta} \sum_{j,k} A_{ij}(t-2)A_{jk}(t) \\
&+ \alpha e^{-2\beta} \sum_{j,k} A_{ij}(t-2)A_{jk}(t-1) \\
&+ \alpha^2 e^{-2\beta} \sum_{j,k,l} A_{ij}(t-2)A_{jk}(t-1)A_{kl}(t) + \mathcal{O}(e^{-3\beta})
\end{aligned} \tag{77}$$

where $\mathcal{O}(e^{-3\beta})$ terms correspond to wins that happened before time $t - 2$. Motegi and Masuda [216] applied different player-level metrics to the network of tennis players, and found that the dynamic win-loss score has a better predictive power than its static counterpart by Park and Newman (see Fig. 2 in [216]).

5.5. Extending shortest-path-based metrics to temporal networks

Up to now, we have only discussed diffusion-based metrics for temporal networks. It is also possible to generalize shortest-path-based centrality metrics to temporal networks, as it will be outlined in this Section.

5.5.1. Shortest and fastest time-respecting paths

It is natural to define the length of a static path as the number of edges traveled through the path. As we have seen in Section 5.2, a natural way to enforce the causality of interactions is to consider time-respecting paths instead of static paths. There are two possible ways to extend the notion of shortest paths to time-respecting paths. In analogy with the definition of length for static paths, one can simply define the length of a time-respecting path as the number of edges traversed along the path. This is called the *topological length* of path $\{(i_1, i_2, t_1), (i_2, i_3, t_2), \dots, (i_n, i_{n+1}, t_n)\}$ which is simply equal to n [209,211]. Alternatively, one can also define the *temporal length* of path $\{(i_1, i_2, t_1), (i_2, i_3, t_2), \dots, (i_n, i_{n+1}, t_n)\}$ which is the time distance between the first and the last path's edge, i.e., $t_n - t_1 + 1$ [210].

Hence, the two alternatives to generalize the notion of the shortest path to temporal networks correspond to choosing the time-respecting paths with the smallest topological length [211] or those with the smallest temporal duration [196,208], respectively. The former and latter choice lead to the definition of *shortest* and *fastest* time-respecting paths, respectively. More details and subtleties associated to the definition of a distance in temporal networks can be found in [201] in Section 4.3.

5.5.2. Temporal betweenness centrality

The static betweenness score of a given node is defined as the fraction of geodesic paths that pass through the node (see Section 2.3). It is natural to extend this definition to temporal networks by considering the fraction of shortest [211] or fastest [208] time-respecting paths that pass through the node. Denoting by $\sigma^{(i)}(j, k)$ the number of shortest (or fastest) time-respecting paths between j and k that pass through node i , we can define the temporal betweenness centrality $B_i^{(temp)}$ of node i as [211]

$$B_i^{(temp)} = \sum_{j \neq i \neq k} \sigma^{(i)}(j, k). \tag{78}$$

Temporal betweenness centrality can be also normalized by the total number of shortest paths that have not i as one of the two extreme nodes [218]. A different definition of temporal betweenness centrality, based on fastest time-respecting paths, has been introduced by Tang et al. [208].

5.5.3. Temporal closeness centrality

In Section 2.3, we defined the static closeness score of a given node as the inverse of the mean geodesic distance of the node from the other nodes in the network. Analogously, one can define the temporal closeness score of a given node as the inverse of the temporal distance (see Section 5.5.1) of the node from the other nodes in the network. Let us denote the temporal distance between two nodes i and j at time t as $d_{ij}^{(temp)}$ – here temporal distance $d_{ij}^{(temp)}$ is a placeholder variable to represent either the topological length of the shortest paths²³ or the duration of the fastest paths between i and j . The

²³ In this case, the use of the adjective “temporal” might feel inappropriate since we are using it to label a topology-related quantity. However, the adjective “temporal” reflects the fact that we are only considering time-respecting paths.

temporal closeness score $C_i^{(temp)}$ of node i thus reads [196]

$$C_i^{(temp)} = \frac{N-1}{\sum_{j \neq i} d_{ij}^{(temp)}}. \quad (79)$$

This definition suffers from the same problem encountered for Eq. (3): if two nodes are not connected by any temporal path within the observation period, their temporal distance is infinite which results in zero closeness score. In analogy with Eq. (4), to overcome this problem, closeness score can be redefined in terms of the harmonic average of the temporal distance between the nodes [211,218]

$$C_i^{(temp)} = \frac{1}{N-1} \sum_{j \neq i} \frac{1}{d_{ij}^{(temp)}}. \quad (80)$$

In general, the correlation between temporal and time-aggregated closeness centrality is larger than the correlation between temporal and time-aggregated betweenness centrality [211]. This is arguably due to the fact that node closeness depends only on paths' length, whereas betweenness centrality takes also into account the paths' actual structure, which is deeply affected when considering the temporal ordering of interactions [211]. A different definition of temporal closeness based on the average of the temporal duration of fastest time-respecting paths is provided by Tang et al. [208].

Computing centrality metrics defined by Eqs. (78) and (80) preserves the causal order of interaction, but requires a detailed analysis of all time-respecting paths in the network. On the other hand, centrality metrics based on higher-order aggregated networks (Section 5.3) can be computed through static methods which are computationally faster. A natural question arises: can we use metrics computed on higher-order time-aggregated networks to approximate temporal centrality metrics? The first investigation has been carried out by Scholtes et al. [211], who show that static betweenness and closeness centrality computed on second-order time-aggregated networks tend to better approximate the corresponding temporal metrics than the respective static metrics computed on the first-order time-aggregated network. These findings suggest that taking into account the first few Markovian orders of memory could be sufficient to approximate reasonably well the actual temporal centrality metrics.

6. Temporal aspects of recommender systems

While most of this review concerns with constructing a general ranking of nodes in a network, there are situations where such a general ranking is of a limited applicability. This is particularly relevant in bipartite user–item networks where users' personal tastes and preferences play an important role. For example, a reader is more likely to be interested in a new book by their favorite author than in a book of a recently announced laureate of the Nobel Prize in literature. For this reason, one seeks not to establish a general ranking of all items in the network but rather multiple *personalized* item rankings—ideally as many of them as there are different users in the system. This is precisely the goal of recommender systems that aim to use data on users' past preferences to identify the items that are likely to be appreciated by a given individual user in the future [219,220].

The conventional approach to recommendation is static in nature: the input data are considered without taking time into account (see [69,221] for reviews of the field where the role of time is not discussed). While this typically produces satisfactory results, there is now increasing evidence that the best results are obtained by considering the time information [222]. Winners of several recent high-profile machine-learning competitions, such as the NetflixPrize [223], have emphasized the importance of taking temporal dynamics of the input data into account [224]. In this section, we present the role played by time in recommendation.

6.1. Temporal dynamics in matrix factorization

Recommendation methods have been traditionally classified as memory and model based [225]. *Memory-based methods* rely on selecting a set of similar users (“neighbors”) for every user and recommend items on the basis of what the neighbors have favored in the past. *Model-based methods* rely on formulating a statistical model for user behavior and then use this model to choose which items to recommend. However, this classification now seems outdated as modern matrix factorization methods combine the two approaches: they are model based (Eqs. (81) and (84) below are examples of how the ratings can be modeled), yet they also have memory-based features as they sum over all memory-stored ratings [220].

Motivated by its good performance in many practical contexts [226–228], we describe here recommendation methods based on matrix factorization. The input data of this method take form of ratings given by users to items; there are U users and I items in total. The method assumes that the rating $r_{i\alpha}$ of user i to item α is an outcome of a matching between the user's preferences and the item's properties.

6.1.1. Matrix factorization method

In matrix factorization methods, both users and items are assumed to be representable in a latent (*i.e.*, hidden, not directly manifested) space of dimension D ; the vectors corresponding to user i and item α are \mathbf{p}_i and \mathbf{q}_α , respectively. Assuming that

the vectors for all users and items are known, a yet unexpressed rating is estimated as

$$\hat{r}_{i\alpha} = \mathbf{p}_i \cdot \mathbf{q}_\alpha. \quad (81)$$

Globally, we factorize the sparse rating matrix R into a product of two matrices—the $U \times D$ matrix P with user taste vectors and the $I \times D$ matrix Q with item feature vectors as $R = PQ^T$. This simple recipe can be further extended to include other factor such as, for example, which items have been actually rated by the users [229]. The increased model complexity is then justified by an improved prediction accuracy. Note that the D dimensions of the latent space have no predefined meaning and it is thus impossible to interpret individual elements of the taste and feature vectors.

Of course, the vectors $\{\mathbf{p}_i\}$ and $\{\mathbf{q}_\alpha\}$ are initially unknown and need to be learned (estimated) from the available data. The learning procedure is formally represented as an optimization problem where the task is to minimize the difference between the estimated and actually expressed ratings. By adopting the mean squared error, the goal is thus to minimize

$$E = \sum_{r_{i\alpha} \in \mathcal{R}} (r_{i\alpha} - \mathbf{p}_i \cdot \mathbf{q}_\alpha)^2 \quad (82)$$

where the sum is over ratings $r_{i\alpha}$ in the set of all ratings \mathcal{R} . Since this is a high-dimensional optimization problem, it is useful to employ some regularization technique to prevent overfitting [230]. This is usually done by adding regularization terms that penalize user and item vectors with large elements, leading to the “regularized” error function

$$\tilde{E} = \sum_{r_{i\alpha} \in \mathcal{R}} (r_{i\alpha} - \mathbf{p}_i \cdot \mathbf{q}_\alpha)^2 + \lambda(\|\mathbf{p}\|^2 + \|\mathbf{q}\|^2) \quad (83)$$

where parameter λ is determined by cross-validation (*i.e.*, by comparing the predictions with a hidden part of data) [231,232]. While Eq. (83) presents a seemingly difficult optimization problem with many parameters, simple optimization by gradient descent [230] works well and quickly converges to a good solution. This optimization method can be implemented efficiently because partial derivatives of Eq. (83) with respect to elements of vectors \mathbf{p}_i and \mathbf{q}_α can be written down analytically (see [226] for a particular example).

When modeling ratings in a bipartite user–item system, it is useful to be more specific than the simple model provided by Eq. (81). A common form [182] of estimated rating is

$$\hat{r}_{i\alpha} = \mu + d_i + d_\alpha + \mathbf{p}_i \cdot \mathbf{q}_\alpha \quad (84)$$

where μ is the overall average rating, and d_i and d_α are the average observed deviations of ratings given by user i and to item α , respectively. The interpretation of this “separation of effects” is straightforward: the rating of a strict user to a generally badly rated movie is likely to be low unless the match between the user’s taste and the movie’s features turns out to be particularly favorable. A more detailed discussion of the baseline estimates can be found in [233]. It is useful to separate the baseline part of the prediction

$$\hat{r}_{i\alpha} = \mu + d_i + d_\alpha \quad (85)$$

for a comparison in recommendation evaluation which can tell us how much the present matrix factorization method can improve over this elementary approach.

6.1.2. Including time in matrix-factorization methods

While generally well-performing, the above-described matrix factorization framework is *static* in nature: all ratings regardless of their age enter the estimation process with the same weight and all parameters are assumed constant. In particular, even when estimated parameter values can change when the input data changes (because of, for example, accumulating more data), a given set of parameter values is assumed to model all available input data—the oldest and the most recent alike. An analysis of the data shows that such a stationarity assumption is generally not satisfied. In the classical NetflixPrize dataset, for example, two non-stationary features become quickly apparent. First, the average item rating has increased suddenly by approximately 0.3 (in the 1–5 scale) in 2004. Second, item rating generally increases with item age. Koren [182] connected the former effect with improving the function of the Netflix’s internal recommendation system and, also, the company’s change of word explanations given to different ratings (from “superb” to “loved it” for rating 5, for example). Similarly, the latter effect is linked with young items being chosen at random at a higher proportion than old items for which the users apparently have more information and only decide to consume and rate them when the chances are good that they might like the item.

To include time effects, the baseline prediction provided by Eq. (85) can be readily generalized to the form

$$\hat{r}_{i\alpha} = \mu + d_i(t) + d_\alpha(t) \quad (86)$$

where the difference terms $d_i(t)$ and $d_\alpha(t)$ are assumed to change with time. How to model the time effects that influence $d_i(t)$ and $d_\alpha(t)$ is an open and context-dependent problem. In any given setting, there are multiple sources of variation with time that can be considered (such as a slowly-changing appreciation of a movie, rapidly-changing mood of a user, and so forth). Which of them have significant impact on the resulting performance is best decided by formulating a complex model

and measuring the prediction precision when components of this model are “turned on” one after another. The evolution of slowly-changing features can be modeled by, for example, dividing the data into bins. If time t falls in bin $B(t)$, we can write $d_\alpha(t) = d_\alpha + d_{\alpha,B(t)}$ where the overall difference d_α is complemented with the difference in a given bin. By contrast, the bins needed to model fast-changing features need to be very short and therefore suffer from poor statistic for each of them; Fast-changing features are thus better modeled with smooth functions. For example, the gradual shift in a user’s bias can be modeled by writing $d_i(t) = d_i + \delta_i \text{sign}(t - t_i) |t - t_i|^\beta$ where t_i is the average time stamp of i ’s ratings. As before, parameter β is to be set by cross validation. For further examples on which temporal effects to model and how, see [182] for an example.

We can now finally come back to the matrix factorization part and introduce time effects there. Out of the two set of computed vectors—user preferences \mathbf{p}_i and item features \mathbf{q}_α , it is reasonable to assume that only the user preferences change. To keep things simple, we follow here the same approach as above to model user differences $d_i(t)$. Element k of i ’s preference vector thus becomes

$$p_{ik} = p_{ik} + \delta_{ik} \text{sign}(t - t_i) |t - t_i|^\beta. \quad (87)$$

The complete prediction thus reads

$$\hat{r}_{i\alpha}(t) = \mu + d_i(t) + d_\alpha(t) + \mathbf{p}_i(t) \cdot \mathbf{q}_\alpha. \quad (88)$$

Note that there are many variants of matrix factorization and also many ways how to incorporate time into them. For the sake of simplicity, we have omitted here the day-specific effect which models the fact that a user’s mood on a given single day can be anything from pessimistic through neutral to optimistic and the predictions should include this. When this effect is included in the model, the accuracy of results can improve considerably (see [182] for details). The matrix model can be further modified to include, for example, the information on which items have been chosen and rated by the user. Koren [229,233] has shown that this modification also improves the prediction accuracy.

To achieve further improvements, several well-performing methods which can be generally very diverse and unrelated can be combined in a unique predictor [234,235]; in the literature, this is referred to as “blending” or “bagging”. Notably, some sort of method combination have been employed by all successful teams in the NetflixPrize [222]. This is a sign that to obtain ultimate achievements in the recommendation and prediction task, there is no single method that outperforms all the others but rather that all methods can capture some aspects of the actual user behavior.

6.2. Temporal dynamics in network-based recommendation

A popular approach to recommendation is based on representing the input data with bipartite networks and studying various diffusive processes on them [54,236]. Provided a set of ratings $\{r_{i\alpha}\}$ given by users to items (the same ratings that served as input for the matrix factorization methods in the previous section), all ratings above a chosen threshold are represented as links between the respective users and items. When computing the recommendation scores of items for a given user i in the original “probabilistic spreading” method [236], one assigns one unit of “resource” to all items connected with the given user and zero resource to the others. The resource is then propagated to the user side and back in a random walk-like process and the resulting amounts of resource on the item side are interpreted as item recommendation scores; items with high score are assumed to be likely to be appreciated by the given user i . The main difference from matrix factorization methods is that network-based methods assume binary user–item data as input²⁴ and use physics-motivated processes on them.

Despite high activity in this field [237], little attention has been devoted to the role that time plays in the recommendation process. While it has been noticed, for example, that the most successful network-based methods tend to recommend popular items [238,239], this has not been explicitly connected with the fact that to gain popularity, items need time and therefore the recommended items are likely to be old. The role of time in network-based recommendation has been finally directly discussed by Vidmer and Medo [83]. The authors point out a deficiency of the usual recommendation evaluation procedure that is based on removing links from the network at random and examining how they are reproduced by a recommendation method. Random removal of links targets preferentially popular items—the evaluation procedure thus shares the bias that is innate to network-based evaluation methods which then leads to an overestimation of the methods’ performance. To make the evaluation more fair and also more relevant from the practical point of view, the latest links need to be removed instead. All the remaining (preceding) links are then used to reconstruct “near future” of the network. Note that this procedure, albeit in line with standards of the field of information filtering, is fundamentally different from the evaluation based on random link removal where both past and future links are used to reconstruct the missing part.

As shown in [83], measured performance of recommendation indeed deteriorates once time-based evaluation is employed. The main reason why network-based methods fail in reproducing the most recent links in a network is that while the methods generally favor popular items, the growth of real networks is typically determined not only by node popularity but also by node age that helps recent nodes to gain popularity despite the presence of old and already popular nodes [74]. In other words, the omnipresent preferential attachment mechanism [158] is modulated with other influences—node age

²⁴ This makes the methods applicable also in systems where the users give no explicit ratings but only attach themselves with items that they have bought, watched, or otherwise interacted with.

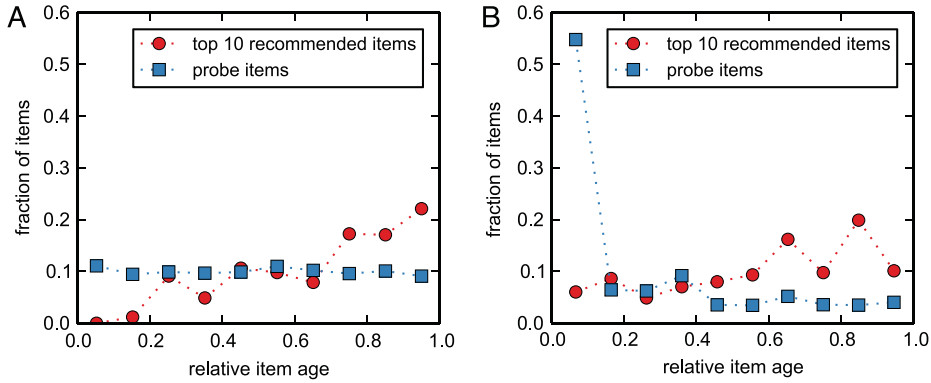


Fig. 11. For two distinct datasets (the popular Netflix dataset [223] in panel A and the Econophysics Forum dataset [240] in panel B), we remove the most recent 1% of links and use the remaining 99% to compute recommendations with the classical probability spreading method [236]. The curves show the age distribution among the items at top 10 places of users' recommendation lists and among the items in the removed 1% of links, respectively (item age is computed relatively with respect to the whole dataset's duration). Recent items are underrepresented in recommendations (and, correspondingly, old items are overrepresented) with both datasets. The disparity is greater for the Econophysics Forum data where aging is faster than it is in the Netflix data.

and node fitness, for example. The difference between the age of items attached with the latest links and the age of items scoring highly in a classical network-based recommendation method is illustrated in Fig. 11.

This suggests that the performance of network-based methods could be improved by enhancing the standing of recent nodes in the recommendation lists. While there are multiple ways how to achieve that, Vidmer and Medo [83] present a straightforward combination of a traditional method's recommendation scores with item degree increase in a recent time period (due to the aging of nodes, high degree increase values tend to be achieved by recent nodes; this is also illustrated in Fig. 11). Denoting the recommendation score of item α for user i at time t as $f_{i\alpha}(t)$ and the degree increase of item α over the last τ time steps as $\Delta k_{\alpha}(t, \tau)$, the best-performing modified score was

$$f_{i\alpha}(t)' = f_{i\alpha}(t) \frac{\Delta k_{\alpha}(t, \tau)}{k_{\alpha}(t)}. \quad (89)$$

Besides multiplying the original score with the recent degree increase, we divide here with the current item degree $k_{\alpha}(t)$ which aims to correct the original method's bias towards popular items (both baseline methods utilized in [83], preferential spreading and similarity-preferential diffusion, exhibit this bias). The above-described modification significantly improves recommendation accuracy (as measured by, for example, recommendation recall) and, at the same time, it tends to recommend less popular items than the original baseline methods and thus answers the call for diversity-favoring recommendation methods [241–243]. The described time-aware recommendation method is very elementary; How best to include time in network-based recommendation methods is still an open question.

6.3. The impact of recommendation ranking on network growth

Once we succeed in finding recommendation methods with satisfactory accuracy, the natural question to ask is what would happen if the users would actually follow the recommendations. In other words, once input data are used to compute recommendations, what is the potential effect of the recommender system on the system's future evolution? This can be addressed by coupling a recommendation method with a simple agent-based model where a user can become active, choose an item based on the received recommendations or at random, and thus establish a new link in the network that affects the future recommendations for this user as well as, indirectly, for all the other users.

Zeng et al. [244] used this idea to show that iterating recommendation results tends to magnify popularity differences among the items with popular items benefiting most from recommendation. To be able to study long-term effects of iterating recommendation, Zeng et al. [5] simulated a closed system where new links are added and, to prevent the bipartite user-item network from eventually becoming complete, the oldest links are removed. The authors show that popularity-favoring recommendation methods generally lead to a pathological stationary state where only a handful of items occupy the attention of all users. This concentration of the users' attention is measured by the Gini index which is a widely-used measure of inequality [245]; its values of zero and one correspond to the perfectly equal setting (where in our case, the degree of all items is the same) and the maximal inequality (where one item monopolizes all the links), respectively. Fig. 12 shows the stationary value of the item degree Gini coefficient that emerges by iterating the above-described network rewiring process.

The authors use the popular ProBS-HeatS hybrid method [55] which has one parameter, θ , that can be used to tune from accurate and popularity-favoring recommendations for $\theta = 0$ to little accurate and diversity-favoring recommendations for $\theta = 1$. As can be seen in the figure, θ values of below approximately 0.6 produces recommendations of high precision but the stationary state is very unequal. By contrast, θ close to one produces recommendations of low precision and a low Gini

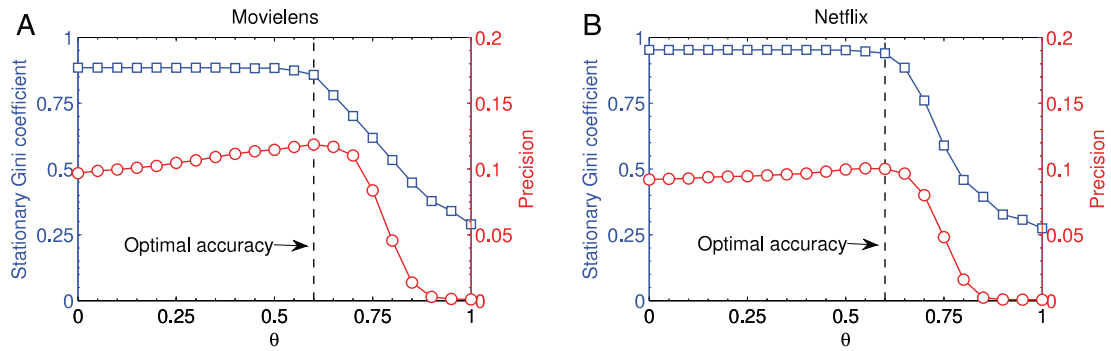


Fig. 12. Results of the rewiring process with the initial setting corresponding to two different real datasets, (A) Movielens and (B) Netflix, as a function of the θ parameter of the HeatS-ProbS recommendation method that is used to compute recommendations that drive the rewiring process. The blue left and red right axes show the stationary Gini coefficient in the long-term limit and the short-term recommendation, respectively. Source: (reprinted from [5]).

coefficient in the stationary state. Fig. 12 shows a hopeful intermediate regime where the sacrifice of some little fraction of the recommendation precision can lead to a stationary state with substantially lower Gini coefficient. This kind of analysis is important as it can highlight potentially harmful long-term effects of information filtering algorithms and it suggests possible ways how to overcome it.

7. Applications of ranking to evolving social, economic and information networks

While complex (temporal) networks constitute a general mathematical framework to describe a wide class of real complex systems, intrinsic differences between the meaning of nodes and their relationships across different datasets make it impossible to devise a unique, all-encompassing metric for node centrality. This has resulted in a plethora of static (Section 2) and time-dependent (Sections 4 and 5) centrality metrics, each of them based on specific assumptions on what important nodes represent. We are left with a fundamental question: given a complex, possibly time-varying, network, which metric shall we use to quantify node importance in the system?

The answer certainly depends on the goals we aim to achieve with a given metric. On the other hand, for a specific goal, the relative performance of two metrics could be very different for different datasets. This makes it necessary to take a close look at the data, establish suitable and well-defined benchmarking criteria, and proceed to performance evaluation. The goal of this section is to present some case studies where this evaluation procedure can be carried out, and to show the benefits from including the temporal dimension into ranking algorithms and/or into the evaluation procedure.

In this section, we focus on three applications of node centrality metrics: ranking of agents (papers, researchers, and journals) in academic networks (Section 7.1), prediction of economic development of countries (Section 7.2), and link prediction in online systems (Section 7.3).

7.1. Quantifying and predicting impact in academic networks

Scientific citation data are a popular object of network study for three main reasons. Firstly, accurate citation records are available from a period that spans over more than hundred years (the often-used APS data start in 1893). Secondly, all researchers have a direct experience with citing and being cited, which makes it easier for them to understand and study this system. Thirdly, researchers are naturally keen on knowing how is their work perceived by the others, for which paper citation count is an obvious but rather simplistic measure. Over time, a whole discipline of *scientometrics* has developed whose goal is to study the ways of quantifying the research impact at the level of individual papers and research journals, as well as individual researchers and research institutions [246,247]. We review here particularly those impact quantification aspects that are strongly influenced by the evolution (in this case largely growth) of the underlying scientific citation data.

7.1.1. Quantifying the significance of scientific papers

The most straightforward way to estimate the impact of a scientific paper is by computing its citation count (*i.e.*, its in-degree in the citation network). One of the problems of citation count is that it is very broadly distributed [246] which makes it difficult to interpret and use it further to, for example, evaluate researchers. For example, if researcher A authored two papers with citation counts 1000 and 10, respectively, and researcher B authored two papers with citation counts 100 and 100, respectively, the average citation count of authored papers points clearly in favor of author A. However, it is long known that the self-reinforcing preferential attachment process strongly contributes to the evolution of paper citation count [158,159,248] which suggests that one should approach high citation counts with some caution. Indeed, a recent study shows that the aggregation of logarithmic citation counts leads to better identification of able authors than the aggregation

of original citation counts [161]. The little reliability of a single highly-cited paper was one of the reasons that motivated the introduction of the popular h -index that aims to estimate the productivity and research impact of authors.

To estimate paper impact beyond the mere citation count is therefore an important challenge. The point that first begs an improvement – already recognized in the 70s by the scientometrics community [249] – is that while the citation count weights all citations equally, the common sense tells us for that a citation from a highly influential paper should weight more than a citation from an obscure paper. The same logic is at the basis of PageRank [43,60] (see Section 2.5.4), that has been therefore applied on the directed network of citations among scientific papers [56,61,165]. While PageRank can indeed identify some papers (referred to as “scientific gems” by Chen et al. [56]) whose PageRank score is high in comparison with their citation count [56], the interpretation of results is made difficult because of the PageRank’s strong bias towards old papers.

The age bias of PageRank in growing networks [82] (see Section 3.2 for details) can be countered by either introducing explicit penalization of old nodes (see the CiteRank algorithm described in Section 4.2.5) or by a posteriori rescaling the resulting PageRank scores as proposed in [166] (see Section 4.1 for a detailed description of the rescaling procedure). Mariani et al. [166] provided a comparison of five different metrics – citation count, PageRank, rescaled citation count (Section 4.1), rescaled PageRank (Section 4.1), CiteRank (Section 4.2.5) – with respect to their ability of identifying an expert-selected group of highly-influential papers. To this end, they used a small group of 87 milestone papers that have been selected in 2016 by editors of the prime American Physical Society (APS) journal, Physical Review Letters, on the occasion of the journal’s 50th anniversary.²⁵ All these papers have been published between 1958 (when Physical Review Letters have been founded) and 2001. The editors choosing the milestone letters thus acted in 2015 with the hindsight of 15 years and more, which adds credibility to their choice.

To compare the rankings of papers by various metric, one can use the input data to produce the rankings and then assess the milestone papers’ positions in each of them. Denoting the milestone papers with $i = 1, \dots, P$ and the ranking of the milestone paper i by metric $m = 1, \dots, M$ as $r_i^{(m)}$, one can for example compute the average ranking of all milestone papers for each metric, $r_{MS}^{(m)} = \frac{1}{P} \sum_i r_i^{(m)}$. A low value of $r_{MS}^{(m)}$ indicates that metric m assigns the milestone papers good (low) positions in the ranking of all papers. The results shown in Fig. 13A indicate that the original PageRank, followed by its rescaled variant and the citation count, are the best metrics in this respect.

However, using the complete input data to compare the metrics inevitably provides a particular view: it only reflects the ranking of milestone papers after a fixed rather long time period of time (recall that many milestone papers are now more than 30 years old). Furthermore, a metric with strong bias towards old papers (such as PageRank, for example), is given undue advantage which makes it difficult to compare its result with another metric which has a different level of bias. To overcome these problems, Mariani et al. [166] provide a “dynamical” evaluation of the metrics where the evolution of each milestone’s ranking position is followed as a function of its age (time since publication). The simplest quantity to study in this way is the *identification rate*: the fraction of milestone papers that appear in the top 1% of the ranking as a function of the time since publication (one can of course choose a different top percentage). This dynamical approach not only provides us with more detailed information but also avoids being influenced by the age distribution of the milestone papers.

The corresponding results are shown in Fig. 13. There are two points to note: time aware quantities, rescaled PageRank $R(p)$ and CiteRank T , rank the milestone papers better than time unaware quantities, citation count c and PageRank p , in the first 15 years after publication. Second, network-based quantities, p and $R(p)$, outperform their counterparts based on simply counting the incoming citations, c and $R(c)$, which indicates that taking the complete network structure of scientific citations is beneficial for our ability to identify significant scientific contributions. As shown recently by Ren et al. [250], it is not given that PageRank performs better than citation count: no performance improvement is found, for example, in the movie citation network constructed and analyzed by Wasserman et al. [16]. Ren et al. [250] argue that the main difference lies in the lack of degree-degree correlations in the movie citation networks, which in turn undermines the basic assumptions of the PageRank algorithm.

The interactive website scienconow.info makes the results obtained with rescaled PageRank and other metrics on regularly updated APS data (as of the time of writing this review, the database includes all papers published until December 2016) publicly available. The very recent paper on the first empirical observation of gravitational waves, for example, reaches a particularly high rescaled PageRank value of 28.6 which is only matched by a handful of papers.²⁶ Despite being very recent, the paper is ranked 17 out of 593,443 papers, as compared to its low rankings by the citation count (rank 1,763) and PageRank (rank 12,277). Among the recent papers that score well by $R(p)$ are recent seminal contributions to the study of graphene (The electronic properties of graphene from 2009) and topological insulators (two papers from 2010 and 2011); both these topics have recently received Nobel Prize in physics. In summary, rescaled PageRank allows one to appraise the significance of published works much earlier than other methods. Thanks to its lack of age bias, the metric can be also used to select both old and recent seminal contributions to a given research field.

The rescaling procedure has also its pitfalls. First, it rewards the papers with immediate impact and, conversely, penalizes the papers that need long time to prove their worthiness, such as the so-called *sleeping beauties* [251,252]. Second, the early evaluation of papers with rescaled quantities naturally results in an increased rate of false positives—papers that due to a

²⁵ See the accompanying website <http://journals.aps.org/prl/50-years/milestones>.

²⁶ Recall that the $R(p)$ value quantifies by how many standard deviations a paper outperforms other papers of similar age.

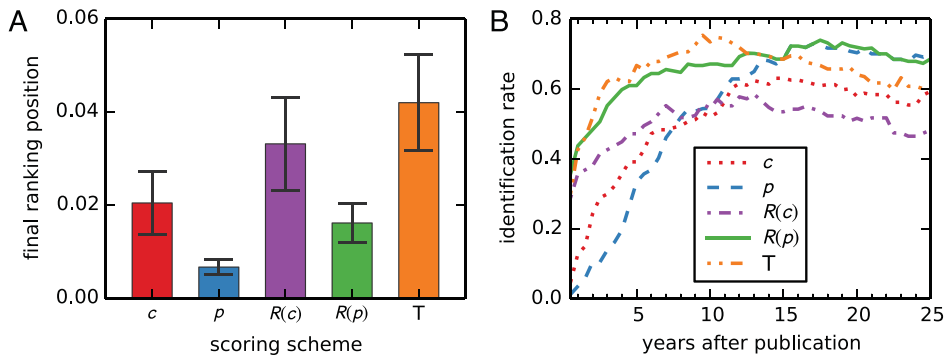


Fig. 13. Ranking performance comparison for a number of paper centrality metrics: citation count, c , PageRank with the teleportation parameter 0.5, p , rescaled citation count, $R(c)$, rescaled PageRank, $R(p)$, and CiteRank, T . (A) The average ranking position (normalized to the number of papers) computed on the complete APS data spanning from 1893 until 2009. (B) The dependence of the milestone identification rate on the age of milestone papers. Source: (Adapted from [166].)

few quickly obtained citations initially reach high values of $R(p)$ and $R(c)$ but these values substantially decrease later.²⁷ If we take into account that those early citations can be also “voices of disagreement”, it is clear that the rescaled scores of very young papers need to be used with caution. A detailed study of these issues remains a future challenge.

7.1.2. Predicting the future success of scientific papers

When instead of paper significance, we are only interested in predicting the eventual number of citations we may follow the approach demonstrated by D. Wang et al. [76] where the network growth model with preferential attachment and aging [74] was calibrated on citation data and the aging was found to decay log-normally with time. While the papers differ in the parameters of their log-normal aging decay, these differences do not influence the predicted long-term citation count which is shown to depend only on paper fitness f_i and overall system parameters. If paper fitness is obtained from, for example, a short initial period of the citation count growth, the model thus can be used to predict the eventual citation count. As the authors show, papers with similar estimated fitness indeed achieve much more similar eventual citation count than the papers that had similar citation counts at the moment of estimation. While J. Wang et al. [254] point to a lack of the method’s prediction power, D. Wang et al. [255] explain that this is a consequence of model over-fitting and discuss how it can be avoided.

A recent prediction method proposed by Cao et al. [256] predicts future citations of papers by building paper popularity profiles and finding the best-matching profile for any given paper; its results are shown to outperform those obtained by the approach by D. Wang et al. [76] (see Fig. 14). Cao’s method has the advantage of not being built on any specific model of paper popularity growth as it directly uses empirical paper popularity profiles. This can be an important advantage, given that the network growth model used in [76, 161] certainly lacks some features of the real citation network. For example, Petersen et al. [257] considered the impact of author reputation on the citation dynamics and found that the author’s reputation dominates the paper citation growth rate of little cited papers. The crossover between the reputation- and preferential attachment-dominated dynamics is found to be approximately at 40 citations (see [257] for details). Golosovsky and Solomon [162] went one step further and proposed a more complex citation growth model built on a copying-redirection-triadic closure mechanism. The use of these models in the prediction of paper success has not been tested yet.

A normalized variant of indegree which accounts for publication time, labeled as rescaled indegree in Section 4.1, has been used by Newman [77] to find the papers that are likely to prove highly successful in the future (see a detailed discussion in Section 3.1). Five years later [163], the previously identified outstanding papers were found to significantly outperform a group of randomly chosen papers as well as a group of randomly chosen papers with the same number of citations. To specifically address the question of predicting the standing of papers in the citation network consisting solely of the citations made in near future, Ghosh et al. [184] use a model of the reference list creation which is based on navigating the citation network by following references in existing papers whilst preferring the references to recent papers (see Section 4.2.2 for details).

An interesting departure from works that build on the citation network among the papers is presented by Sarigol et al. [131] who attempt to predict the paper’s success solely on the basis of the paper’s authors position in the co-authorship network. On a dataset of more than 100,000 computer science papers published between 1996 and 2007, they show that among the papers authored by an author who is in the top 10% by (simultaneously) betweenness centrality, degree centrality, k -core centrality and eigenvector centrality, 36% of these papers are eventually among the top 10% most cited papers five

²⁷ A spectral-clustering method to classify the papers according to their citation trajectories has been recently proposed by Colavizza and Franceschet [253]. According to their terminology, papers with large short-term impact and fast decay of attractiveness for new citations are referred to as *sprinters*, as opposed to *marathoner* papers that exhibit slow decay.

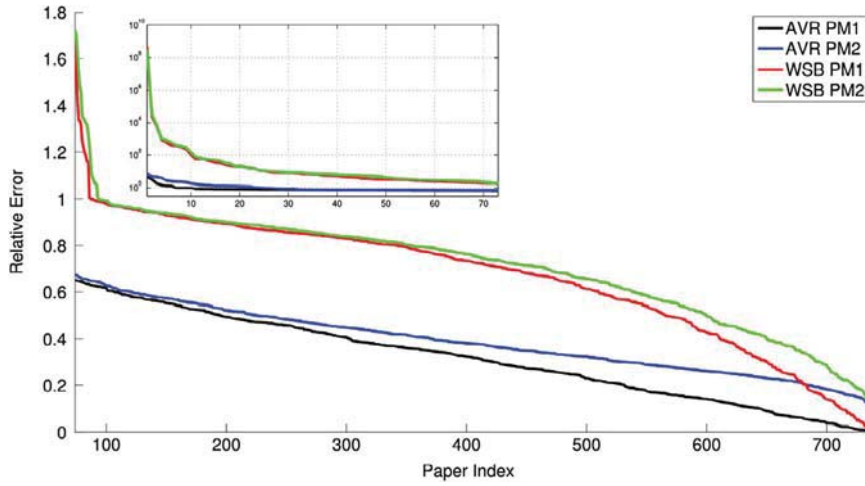


Fig. 14. The relative citation count prediction error (vertical axis) obtained with the method by D. Wang et al. [76], labeled as WSB, as compared to the basic “data analytic” approach by Cao et al. [256], labeled as AVR. Papers published by the Science journal in 2001 were used for the error evaluation. To visualize the error distribution among the papers, the papers are ranked by their relative error for each prediction method separately. Of note is the small group of papers with large relative error shown in the inset. The AVR method suffers from this problem less than the WSB method. Source: (reprinted from [256]).

years after publication. By combining the centrality of paper co-authors with a machine-learning approach, this fraction (which represents the prediction’s precision) can be further increased to 60% whilst the corresponding recall is 18% (*i.e.*, out of the 10% top cited papers, 18% are found with the given predictor).

7.1.3. Predicting the future success of researchers

Compared with the prediction of the future success of a scientific publication, the prediction of the future success of an individual researcher is much more consequential because it is crucial in distributing the limited research resources (research funds, permanent positions, etc.) among a comparatively large number of researchers. Will a researcher with a highly cited paper produce more papers like that or not? Early on, the predictive power of h -index has been studied by Hirsch himself in [258] with the conclusion that this metric predicts a researcher’s success in the future better than other simple metrics (total citation count, citations per paper, and total paper count).

It is probably naive to expect that a single metric, regardless of how well designed, can capture and predict the success in such a complex endeavor as research is. Acuna et al. [259] attempted to predict the future h -index of researchers by combining multiple features using linear regression. For example, the prediction formula for the next year’s h -index of neuroscientists had the form

$$h_{+1} = 0.76 + 0.37\sqrt{n} + 0.97h - 0.07y + 0.02j + 0.03q \quad (90)$$

where n , h , y , j , and q are the number of papers, current h -index, number of years since the first paper, number of distinct journals where the author has published, and the number of articles in top journals, respectively. However, this work was later criticized by Penner et al. [260]. One of the discussed problems is that to predict the h -index is relatively simple as this quantity changes slowly and can only grow. The baseline prediction of unchanged h -index values is thus likely to be rather accurate. According to Penner et al., one should instead aim to predict the change of the author’s h -index and evaluate the predictions accordingly; in this way, no excess self-confidence will result from giving ourselves a simple task. Another problem of the original evaluation is that it considered all researchers and thus masked the substantial level of errors specific to young researchers who, at the same time, are likely candidates to be screened by hiring committees according to quantitative criteria. The young researchers who have published their first first-author paper recently are precisely the target of Zhang et al. [261] who use a machine-learning approach to predict which young researchers will be most cited in the near future, and show that their method outperforms a number of benchmark methods (such as the researcher’s total number of citations and the average citation count of the venues where the researcher has published, for example).

As recently shown by Sinatra et al. [262], the most-cited paper of a researcher is equally likely to appear at the beginning, in the middle, or at the end of the researcher’s career. In other words, researchers seem not to become better, or worse, in producing high-impact papers during their careers. Consequently, the authors develop a model which assigns scientists the ability factor Q that multiplicative improves ($Q > 1$) or lowers ($Q < 1$) the potential impact of their next work. The new Q metric is empirically shown to be stable over time and thus suitable for predicting the authors’ future success, as documented by predicting the h -index. Note that Q is by definition an age-independent model parameter and thus allows to compare scientists of different age without the need for any additional treatment (such as the one presented in Section 3.2 for the PageRank metric).

The literature on this topic is large and we present here only a handful of relevant works. An interested reader is referred to a very recent essay on data-driven predictions of science development [263] which summarizes the achieved progress, future opportunities, and the potential impact.

7.1.4. Ranking scientific journals through non-Markovian PageRank

In Section 5, we have described node centrality metrics for temporal networks. We have shown that higher-order networks describe and predict actual flows in real networks better than memory-less networks (Section 5.2). In this section, we focus on a particular application of memory-aware centrality metrics, the ranking of scientific journals [44].

Weighted variants of the PageRank algorithm [264,265] attempt to infer the importance of a journal based on the assumption that a journal is important if its articles are cited relatively many times by papers published in other important journals. However, such PageRank-based methods are implicitly built on the Markovian assumption, which may provide a poor description of real network traffic (see Section 5.2). As pointed out by [44], this is the case also for multi-disciplinary journals like *Science* or *PNAS* which receive citation flows from journals from various fields. If we neglect memory effects, multi-disciplinary journals redistribute the received citation flow towards journals from diverse fields. This effect causes the flow to repeatedly cross field boundaries, which is arguably unlikely to be done by a real scientific reader and, for this reason, makes the algorithm unreliable as a model for scholars' browsing behavior.

By contrast, in a second-order Markov model, if a multi-disciplinary journal J receives a certain flow from a journal which belongs to field F , then journal J tends to redistribute this flow to journals that belong to the same field F (see panels c–d of Fig. 4 in [44] for an example). This results in a sharp reduction of the flow that crosses disciplinary boundaries (referred to as “flow leak” in [44]) and in a clear increase of the within-community flow. As a result, when including second-order effects, multidisciplinary journals (such as *PNAS* and *Science*) receive smaller net flow, whereas specialized journals (such as *Ecology* and *Econometrica*) gain flow and thus improve their ranking (see [44] for the detailed results).

This property makes also the ranking less subject to manipulation, which is a worrying problem in quantitative research evaluation [266,267]. As we have just seen, if a paper written in a journal specialized in Ecology cites a multi-disciplinary journal, this credit is likely to be redistributed across journals of diverse fields in a first-order Markov model. For a memory-less ranking of journals, a specialized journal would be thus tempted to attempt to boost its score through editorial policies that discourage citations to multi-disciplinary journals. By contrast, the second-order Markov model suppresses large part of flow leaks which makes it less convenient this strategic citation manipulation. The ranking of journals based on second-order Markov model has also been shown to be more robust with respect to the selection of included journals in the analysis [268], which is an additional positive property for the second-order model.

In a scientific landscape where metrics of productivity and impact are proliferating [269], assessing the biases and performance [161] of widely used metrics and devising improved ones is critical. For example, the widely used impact factor suffers from a number of worrying issues [270–272]. The results presented by Rosvall et al. [44] indicate that the inclusion of temporal effects holds promise to improve the current rankings of journals, and hopefully future research will continue exploring this direction.

7.2. Prediction of economic development of countries

Predicting the future economic development of countries is of crucial importance in economics. The GDP per capita is the traditional monetary variable used by economists to assess a country's development [273]. Even though it does not account for critical elements such as income equality, environmental impact or social costs [274], GDP is traditionally regarded as a reasonable proxy for the country's economical wealth, and especially of its industry [275]. Twice a year, in April and in October, the *International Monetary Fund* (IMF) makes projections for the future GDP growth of countries [276]. The web page does not detail the precise procedure, but it indicates that it combines together many factors, such as oil price, food price, and others. Development predictions based on complex networks do not aim to beat the projections made by the IMF. They rather aim to extract valuable information from international trade data with a limited number of network-based metrics, attempting thus to isolate the key factors that drive economic development. The respective research area, known as *Economic Complexity*, has attracted attention in recent years both from the economics [277,278] and from the physics community [31,140].

Two metrics, the Method of Reflections [30], and the Fitness-Complexity metric [31], have been designed to rank countries and products that they export in the international trade network. The definitions of the two metrics (and some of their variants [140,141]) can be found in Section 2.7.2 and Section 2.7.3, respectively. Both algorithms use a complex network approach to represent the international trade data. One builds a bipartite trade network where the countries are connected with the products that they export. The network is built using the concept of Revealed Compared Advantage, RCA (we refer to [30] for details), which distinguishes between “significant” and “non-significant” links in the country-product export network. The resulting network is binary: a country is either considered to be an exporter of a product or it is not.

Once we have defined network-based metrics to rank countries and products in world trade data, a natural question arises: can we use these metric to predict the future development of countries? This question is not univocal: several prediction-related tasks can be designed and the relative performance of the metrics can be different for different tasks. We shall present two possible approaches below.

7.2.1. Prediction of GDP through linear regression

Hidalgo and Hausmann [30] have employed the standard linear-regression approach to the GDP prediction problem: one designs a linear model where the future GDP, $\log[GDP(t + \Delta t)]$, depends linearly on the current GDP, $\log[GDP(t)]$, and on network-based metrics. The linear-regression equations used by [30] read

$$\log[GDP(t + \Delta t)] - \log[GDP(t)] = a + b_1 GDP(t) + b_2 k_i^{(n)} + b_3 k_i^{(n+1)}, \quad (91)$$

where $k_i^{(n)}$ denotes the n th-order country Method-of-Reflection score (hereafter MR-score) as determined with Eq. (32), while a, b_1, b_2, b_3 are coefficients to be determined with a fit to the data. The results by [30,139] show that the Method-of-Reflection score contributes to the variance of economic growth significantly more than measures of governance and institutional quality, education system quality, and standard competitiveness indexes such as the World Economic Forum Global Competitiveness Index. We refer the reader to [139] for the detailed results, and to <http://atlas.media.mit.edu/en/> for a web platform with rankings by the MR and visualizations of the international trade data.

7.2.2. Prediction of GDP through the method of analogues

The linear-regression approach by [30,139] has been criticized by Cristelli et al. [32] who point out that the linear-regression approach assumes that there is a *uniform* trend to uncover. In other words, it assumes that all the countries respond to a variation in their score in the same way, *i.e.*, with the same regression coefficients. Cristelli et al. [32] use the country fitness score defined by Eq. (33) to show that instead, the dynamics of countries in the Fitness-GDP plane is very heterogeneous, and the predictability of countries' future position is only possible in a limited region of the plane. This undermines the very basic assumption of the linear regression technique [32], and calls for a different prediction analysis tool.

Cristelli et al. [32] borrow a well-known method from the literature on dynamical systems [279–282], called *method of analogues*, to study and predict the dynamics of countries in the Fitness-GDP plane. The basic idea of the method is that we can try to predict the evolution of the system by using our knowledge of the history of the system. The main goal of the method is to reveal simple patterns and attempt prediction in systems for which we do not know the laws of evolution. While [32] only considers the Fitness-GDP plane, with Fitness determined through Eq. (33), we consider here the score-GDP plane for three different scores: fitness, degree, and the Method-of-Reflections (MR) score.²⁸

In the following, we analyze the international trade BACI dataset (1998–2014) composed of $N = 261$ countries and $M = 1241$ products.²⁹ From this dataset, we use the procedure introduced by Hidalgo and Hausmann [30] to construct, year by year, the country–product network that connects the countries with the products that they export. The first insights into the heterogeneity of the country dynamics in the score-GDP plane can be gained through the following procedure:

- We divide the score-GDP plane (in a logarithmic scale) in equally-sized boxes.
- We record the box at which each country is located at a given year, as well as the box at which the country is located ten years later.
- We measure the total number of countries N_i that have been located in box i at a certain year as well as the number n_i of different boxes occupied after ten years by countries originating from box i .
- For each box, we compute the dispersion $C_i = (n_i - 1)/(N_i - 1)$.

A box whose all countries will occur in the same box after 10 years has zero dispersion, and a box whose N_i countries will occur in N_i different boxes after ten years has the dispersion of one. Therefore, for a given box, a small value of C_i means that the future evolution of countries that are located at that box is more predictable than the evolution of countries located at boxes with large dispersion values.

In agreement with the results presented in [32], we observe (Fig. 15) an extended “low-dispersion” (*i.e.*, predictable) region in the Fitness-GDP plane, which corresponds essentially to high-fitness countries. This means that for high-fitness countries, fitness is a reliable variable to explain the future evolution of the system. By contrast, for low-fitness countries, fitness fails to capture the future dynamics of the system and additional variables (*e.g.*, institutional or education quality indicators) may be necessary to improve the prediction performance. Cristelli et al. [32] point out that this strongly heterogeneous dynamics goes undetected with standard forecasting tools such as linear regressions. Interestingly, an extended low-dispersion region is also found in the degree-GDP plane, whereas the low-dispersion region is comparatively small in the MR-score-GDP plane.

The complex information provided by the dispersion values in a two-dimensional plane can be reduced by computing a weighted mean of dispersion, which weights the contribution of each box with the number of events that fall into that box. The weighted mean of the box dispersion values in Fig. 15 reads

$$\bar{c} = \frac{\sum_i C_i N_i}{\sum_i N_i}. \quad (92)$$

²⁸ To obtain the MR score, we perform two iterations of Eq. (32). We have checked that the performance of the Method of Reflections does not improve by increasing the number of iterations.

²⁹ In international trade datasets, different classifications of products can be found. In the BACI dataset used here, products are classified according to the Harmonized System 2007 scheme, which classifies hierarchically the products by assigning them a six-digits code. For the analysis presented in this Section, we aggregated the products to the fourth digit.




Fig. 15. Box dispersion in the score-GDP plane for three different country-scoring metrics: degree, Method of Reflections, and Fitness. The results were obtained on the BACI international trade data ranging from 1998 to 2014 [283] (see the main text for details).

The obtained values of \bar{c} are 0.41, 0.71, and 0.35 for the degree, Method of Reflections, and Fitness, respectively. These numbers suggest that while both degree and Fitness significantly outperform the predictive power of the Method of Reflections, Fitness is the best metric by some margin. The results are qualitatively the same when a different number of boxes is used. These findings indicate that fitness and, to a smaller extent, degree can be used to make predictions in the score-GDP plane, whereas the Method of Reflections cannot be used to make reliable predictions with the method of analogues. The simple aggregation method used here to quantify the prediction performance of metrics can be in the future improved to better understand the methods' strong and weak points. An illustrative video of the evolution of countries in the Fitness-GDP plane³⁰ provides a general perspective on the two above-described economic-complexity approaches to GDP prediction.

³⁰ <http://www.nature.com/news/physicists-make-weather-forecasts-for-economies-1.16963>

7.3. Link prediction with temporal information

Link prediction is an important problem in network analysis [284,285], which aims to identify the potential missing links in a network. This is particularly useful in settings where to obtain information about the network is costly, such as in gene interaction experiments in biology [286,287], and it is thus useful to prioritize the experimentation by first studying the node pairs that are the most likely to be actually connected with a link. The problem has its relevance also in the study of social networks [288,289], for example, where it can tell us about their evolution in the future. There is also an interesting convoluted setting where information diffuses over the social network and this diffusion can influence the process of social link creation [290]. A common approach to link prediction is based on choosing a suitable node similarity metric, computing the similarity for all yet unconnected node pairs, and finally assuming that the links are most likely to exist among the nodes that are the most similar by the chosen metric. The simple common neighbor similarity of nodes i and j can be computed as $s_{ij} = \sum_{\alpha} A_{i\alpha} A_{j\alpha}$, where A is the network's adjacency matrix. See [285] for more information on node similarity estimation and link prediction in general.

7.3.1. Time-dependent node similarity

Usual node similarity metrics ignore the time information: the old and new links are assumed to have the same importance, and the same holds for nodes. There are various ways how time can be included in the computation of node similarity. The simple common neighbors metric, for example, can be generalized by assuming that the weight of common neighbors decays exponentially with time, yielding $s_{ij}(t) = \sum_{n \in \Gamma(i,t) \cap \Gamma(j,t)} e^{-\lambda(t-t_{ij})}$ where $\Gamma(i,t)$ is the set of neighbors of node i at time t , t_{ij} is the time when edge between i and j has been established, and λ is a tunable decay parameter. Liu et al. [291] introduce time in the resource allocation dynamics on bipartite networks (first described in [236]). Other originally static similarity metrics can be also modified in a similar way [292]. Note that in constructing time-reflecting similarity metrics, there exist a number of problems: How to characterize the time-decay parameters, whether edges with different timestamps play different roles in the dynamics, whether and how to include users' potential preference for fresher or older objects, among others [293,294].

7.3.2. Time-dependent structural perturbation

We illustrate now how time can be easily introduced in a recently proposed link prediction method based on a perturbed matrix [295]; this time-aware generalization has been recently presented in [296]. The original method assumes an undirected network of N nodes. A randomly chosen fraction p^H of its links constitute a perturbation set with the adjacency matrix ΔA , and the remaining links correspond to the adjacency matrix A^R ; the complete network's adjacency matrix is $A = A^R + \Delta A$. Since A^R is real and symmetric, it has a set of N orthonormal eigenvectors \mathbf{x}_k corresponding to eigenvalues λ_k and can be written as $A^R = \sum_{k=1}^N \lambda_k \mathbf{x}_k \mathbf{x}_k^T$. The complete adjacency matrix, A , has generally different (perturbed) eigenvalues and eigenvectors that can be written as $\lambda_k + \Delta\lambda_k$ and $\mathbf{x}_k + \Delta\mathbf{x}_k$, respectively. Assuming that the perturbation is small, one can find $\Delta\lambda_k = (\mathbf{x}_k^T \Delta A \mathbf{x}_k) / (\mathbf{x}_k^T \mathbf{x}_k)$ and A can be approximated as

$$\tilde{A} = \sum_{k=1}^N (\lambda_k + \Delta\lambda_k) \mathbf{x}_k \mathbf{x}_k^T, \quad (93)$$

where we kept the perturbed eigenvectors unchanged. This matrix can be considered as a linear approximation (extrapolation) of the network's adjacency matrix based on A^R . As discussed in [295], the elements of \tilde{A} can be considered as link score values in the link prediction problem.

However, the original method divides training and probe sets randomly regardless of time, implying an unattainable scenario that future edges are utilized to predict future edges. It is more appropriate to choose the most recent links to constitute the perturbation set (again, the fraction p^H of all links is chosen). To evaluate link prediction in a system where time plays a strong role, one must not choose the probe at random, but instead divide the network by time so that strictly the future links are there to be predicted (see a similar discussion in Section 6.2). To reflect the temporal dynamics of the network's growth, and thus appropriately address the time-based probe division, Wang et al. [296] introduce the "activity" of node i as $s_i = k_i^H / k_i$ where k_i^H and k_i are the degree of node i in the perturbation set³¹ and the complete data, respectively. Node activity s_i aims to capture the trend of i 's popularity growth: whether it is on a rise or it already slows down. In addition, active nodes are more likely to be connected by a link than little active ones.

Here elements of the eigenvector \mathbf{x}_k can be modified by including node activity as

$$x'_{k,i} = x_{k,i} (1 + \alpha s_i) \quad (94)$$

³¹ We stress that the perturbation set corresponds to a temporal window close to the end of the training set and, accordingly, k^H represents the recent degree increase of a node.

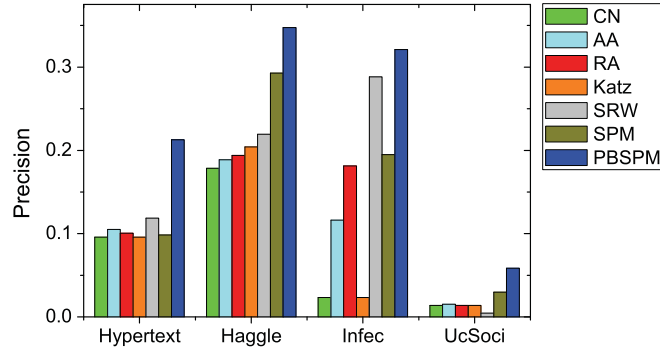


Fig. 16. Precision of traditional methods vs. time-aware link prediction based on structural perturbation (PBSPM) for different networks. In the prediction, 90% historical links of the whole network are utilized to predict 10% future links. The six state-of-the-art methods are Common neighbors (CN) [285], Adamic Adar Index (AA) [301], Resource Allocation Index (RA) [236], Katz Index [40], Superposed Random Walk (SRW) [302] and Structural Perturbation Method (SPM) [295].

where α is a tunable parameter (by setting $\alpha = 0$, node activity is ignored and the original eigenvectors are recovered). The modified eigenvectors can be plugged in (93), yielding the time-aware perturbed matrix

$$\tilde{\mathbf{A}}' = \sum_{k=1}^N (\lambda_k + \Delta\lambda_k) \mathbf{x}'_k \mathbf{x}'_k{}^T. \quad (95)$$

Similarly as before, the elements of $\tilde{\mathbf{A}}'$ can be interpreted as link prediction scores for the respective links.

In [296], a number of different link prediction methods are compared on four distinct datasets: Hypertext data [297], Haggle data [298], Infec data [297] and UcSoci data [299]. Fig. 16 summarizes the results obtained by the authors and shows that taking the time evolution of node popularity into account significantly improves the prediction performance. Apart from the methods introduced above, several further time-aware approaches to link prediction have been proposed recently, such as similarity fusion of multilayer networks [300], real-world information flow feature in link prediction [290], and others.

8. Perspectives and conclusions

After having reviewed the recent progress in the actively studied problem of node ranking in complex networks, the reader may arguably feel overwhelmed by the sheer number and variability of the available static and time-aware metrics. Indeed, a clear understanding of which metric to use in which situation is often lacking. “Proofs” of metric validity are often anecdotal, based on a qualitative comparison between the rankings provided by different metrics [303]. This issue is quite worrying as a wrongly used metric can have immediate real consequences as is the case, for example, for quantitative evaluation of scientific impact, where the used metrics have consequences on funding decisions and academic careers. Waltman [17] points out that since the information that can be extracted from citation networks is inherently limited, researchers should avoid introducing new metrics unless they bring a clear added value to existing metrics. We can only add that this recommendation applies equally to other fields where metrics tend to multiply.

Coming back to the issue of metric validation, we believe that extensive and thorough performance evaluation of (both static and time-aware) network ranking metrics is an important direction future research. Of course, any “performance” evaluation depends on which task we want to achieve with the metric. We outline possible methods to validate centrality metrics in Section 8.1. Whichever validation procedure we choose for the metrics, we cannot neglect the possible bias of the metrics and their impact on the system’s evolution. In this review, we focused on the bias by age of static algorithms and possible ways to counteract it. However, the temporal dimension is not the only possible source of bias for ranking algorithms. We discuss other possible sources of information bias and their relevance in our highly-connected world in Section 8.2.

8.1. Which metric to use? Benchmarking the ranking algorithms

Let us briefly discuss three basic ways that can be used to validate network-based ranking techniques.

Prediction of future edges. While prediction is vital in physics where wrong predictions on real-world behavior can rule out competing theories, it has not yet gained a central role in social and economic sciences [304]. Even though many metrics in the last years have been evaluated according to their predictive power, as outlined in several Sections of Section 7, wide differences among the adopted predictive evaluation procedures make it often impossible to assess the relative usefulness of the metrics. In agreement with [304], we believe that an increased standardization of prediction evaluation would foster a better understanding of the predictive power of the different metrics. With the growing importance of data-driven prediction [263,304,305], benchmarking ranking algorithms through their predictive power may deepen our understanding about which metrics better reproduce the actual nodes’ behavior and eventually lead to useful applications, perhaps in combination with data-mining techniques [306].

Using external information to evaluate the metrics. External information can be beneficial both to compare different rankings and to interpret which properties of the system they actually reflect. For example, network centrality metrics can be evaluated according to their ability to *identify expert-selected significant nodes*. Benchmarks of this kind include (but are not limited to) identification of expert-selected movies [16,250], identification of awarded conference papers [307,308] or of editor-selected milestone papers [166], identification of researchers awarded with international prizes [13,309].

The rationale behind this kind of benchmarking is that if a metric well captures the importance of the nodes, it should be able to rank at the top the nodes whose outstanding importance is undeniable. Results on different datasets [166,250] show that there is no unique metric that outperforms the others in all the studied problems. An extensive thorough investigation of which network properties make different metrics fail or succeed in a given task is a needed step in future research. First steps have been done in [250], which introduces a time-aware null model to reveal the role of network structural correlations in determining the relative performance of local and eigenvector-based metrics in identifying significant nodes in citation networks.

Expert judgment is not the only source of external information that can be used to validate the metrics. For example, Smith et al. [310] and Mao et al. [64] have applied network centrality metrics to mobile-phone communication networks and compared the region-level scores produced by the metrics with the income levels of the region. The effectiveness of centrality metrics can be also validated according to some prior knowledge of the nodes' function in the network. For example, the scores by network centrality metrics have been compared with known neuronal activity in neural network [110].

External information is beneficial not only to evaluate the metrics but also to move towards a better understanding of what the metrics actually represent. For instance, the long-standing assumption that citation counts are proxies for scientific impact has been recently investigated in a social experiment [311]. Radicchi et al. [311] asked 2000 researchers to perform pairwise comparisons between papers, and then aggregated the results over the researchers, quantifying thereby the papers' perceived impact independently of their citation counts. The results of this experiment show that when authors are asked to judge their published papers, perceived impact and citation count correlate well, which indicates that the assumption that citation count represents paper impact is in agreement with experts' judgment.

Using model data to evaluate the metrics. One can also use models of network growth or network diffusion to evaluate the ability of the metrics to rank the nodes according to their importance in the model. This strategy opens different alternative possibilities. One can use models of network growth [74,161] where the nodes are endowed with some fitness parameter, grow synthetic networks of a given size and, a posteriori, evaluate the metrics' ability to uncover the nodes' latent fitness [82,161] (see also Section 3).

Another possibility is to use real data and test the metrics' ability to unearth the intrinsic spreading ability of the nodes according to some model of diffusion on the network. An extensive analysis has been carried out recently in the review article by Lu et al. [90] for static metrics and the classical SIR network diffusion model [26]. In particular, Lu et al. [90] use real data to compare different metrics with respect to their ability to reflect nodes' spreading power as determined by classical network diffusion models. Analogous studies are still missing for temporal networks. For example, it would be interesting to compare different (static and time-aware) metrics with respect to their ability to identify influential spreaders in suitable temporal-network spreading models [44,312,313] similarly to what has been already done for static metrics. The relative performance of the metrics in identifying of influential spreaders in temporal networks can provide us with important insights on which metrics can be consistently more useful than others in real-world applications.

8.2. Systemic impact of ranking algorithms: avoiding information bias

Besides metrics' performance, it is often of critical importance to detect the possible biases of the metrics and, wherever possible, suppress them. For example, in the context of quantitative analysis of academic data, the bias of paper- and researcher-level quantitative metrics by age and scientific field is long known [314,315], and many methods have been proposed in the scientometrics literature to counteract it (see [175,179,316], among others). In this review, we have mostly focused on the bias of ranking by node age (Section 3) and possible methods to reduce or suppress it (Section 4). On the other hand, other types of bias are present in large information systems. In particular, possible biases of ranking that come from social effects [317] are largely unexplored in the literatures.

Several shortcomings of automated ranking techniques can stem from social effects. For example, there is the risk that online users are only exposed to content similar to what they already know and that confirms their points of view, creating a "filter bubble" [318] of information. This is particularly dangerous for our society since online information is often not reliable.³² This issue may be exacerbated by the fact that users more prone to consume unreliable content tend to cluster together, creating groups of highly-polarized clusters referred to as "echo chambers" [319,320]. Designing statistically-grounded procedures to detect these biases and effective methods to suppress them, thereby preserving information diversity and improving information quality, are vital topics for future research.

The impact of centrality metrics and ranking on the behavior of agents in socio-economics systems deserves attention as well. Just to mention a few examples, online users are heavily influenced by recommender systems when choosing which

³² In a recent report, the World Economic Forum has indicated massive digital misinformation as one of the main risks for our interconnected society: <http://reports.weforum.org/global-risks-2013/risk-case-1/digital-wildfires-in-a-hyperconnected-world/>.

items to purchase [69,321], ranking can affect the decision of undecided electors in political elections [322], the impact factor of scientific journals heavily affects scholars' research activity [323]. Careful considerations of possible future effects are thus needed before adopting a metric for some purpose.

To properly assess the potential impact of ranking on systems' evolution, we would need an accurate model of the nodes' behavior. Homogeneous models of network growth assume that all nodes are driven by the same mechanism when choosing which node to connect to. Examples of homogeneous models include models based on the preferential attachment mechanism [74,76,194,324], models where nodes tend to connect with the most central nodes and to delete their connections with the least central nodes [325,326], among others. While such homogeneous models can be used to generate benchmark networks to test the performance of network-based algorithms [82,161], they neglect the multi-faceted nature of node behavior in real systems and they might provide us with an oversimplified description of real systems' evolution. Complex-network models that account for heterogeneous node behavior have already been proposed [327–329], yet we still lack clear guidelines on how best to use these models to predict the impact of a given centrality metric on the properties of a given network.

8.3. Summary and conclusions

Ranking in complex networks plays a crucial role in real-world applications – we have seen many examples of such applications throughout this review (see also the review by Lu et al. [90]). These applications cover problems in a broad range of research areas, including economics [30,31,33], neuroscience [110,330] and social sciences [40,109], among others. From a methodological point of view, we have mostly focused on methods inspired by statistical physics, and we have studied three broad classes of node-level ranking methods: static algorithms (Section 2), time-dependent algorithms based on a time-aggregated network representation of the data (Section 4), and algorithms based on a temporal-network representation of the data (Section 5). We have also discussed examples of edge-level ranking algorithms and their application to the problem of link prediction in online systems (Section 7.3).

In their recent review on community detection algorithms [129], Fortunato and Hric write that “as long as there will be networks, there will be people looking for communities in them”. We might as well say that “as long as there will be networks, there will be people ranking their nodes”. The broad range of applications of network centrality metrics and the increasing availability of high-quality datasets suggest that research on ranking algorithms will not slow down in the forthcoming years. We expect scholars to become more and more sensitive to the problem of understanding and assessing the performance of the metrics for a given task. While we hope that the discussion provided in the previous two Sections will be useful as a guideline, our discussion does not pretend to be complete and several other dimensions can be added to the evaluation of the methods.

We conclude by stressing the essential role of the *temporal dimension* in the design and evaluation of ranking algorithms for evolving networks. The main message of our work is that disregarding the temporal dimension can lead to sub-optimal or even misleading results. The number of methods presented in this review demonstrates that there are many ways to effectively include time in the computation of a ranking. We hope that this work will be useful for scholars from all research fields where networks and ranking are tools of primary importance.

Acknowledgments

We wish to thank Alexander Vidmer for providing us with data and tools used for the analysis presented in Section 7.2, and for his early contribution to the text of that Section. We would also like to thank Claudio J. Tessone for his valuable feedback on the manuscript.

We are also grateful to all those researchers with whom we have exchanged ideas about the topics presented in this review, in particular: Giulio Cimini, Matthieu Cristelli, Alberto Hernando de Castro, Flavio Iannelli, Francois Lafond, Luciano Pietronero, Zhuo-Ming Ren, Andrea Tacchella, Giacomo Vaccario, Zong-Wen Wei, An Zeng.

HL and MYZ acknowledge financial support from the National Natural Science Foundation of China (Grant No. 11547040), Guangdong Province Natural Science Foundation (Grant Nos. 2016A030310051, 2015KQNCX143), Shenzhen Fundamental Research Foundation (Grant Nos. JCYJ20150625101524056, JCYJ20160520162743717, JCYJ20150529164656096, JCYJ20150731160834611), Natural Science Foundation of SZU (Grant No. 2016-24), CCF-Tencent fund (Grant No. AGR20160201). MSM and MM and YCZ acknowledge financial support from the EU FET-Open Grant No. 611272 (project Growthcom) and the Swiss National Science Foundation Grant No. 200020-156188.

References

- [1] U. Hanani, B. Shapira, P. Shoval, Information filtering: Overview of issues, research and systems, *User Model. User-Adapt. Interact.* 11 (3) (2001) 203–259.
- [2] R. Baeza-Yates, B. Ribeiro-Neto, et al., *Modern Information Retrieval*, vol. 463, ACM press, New York, 1999, pp. 1–944.
- [3] P.-Y. Chen, S.-y. Wu, J. Yoon, The impact of online recommendations and consumer feedback on sales, in: *ICIS 2004 Proceedings*, 2004, p. 58.
- [4] D. Fleder, K. Hosanagar, Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity, *Manage. Sci.* 55 (5) (2009) 697–712.
- [5] A. Zeng, C.H. Yeung, M. Medo, Y.-C. Zhang, Modeling mutual feedback between users and recommender systems, *J. Stat. Mech. Theory Exp.* 2015 (7) (2015) P07020.

- [6] D. Feenberg, I. Ganguli, P. Gaule, J. Gruber, It' good to be first: Order bias in reading and citing NBER working papers, *Rev. Econ. Stat.* 99 (2016) 32–39.
- [7] J. Cho, S. Roy, Impact of search engines on page popularity, in: *Proceedings of the 13th International Conference on World Wide Web*, ACM, 2004, pp. 20–29.
- [8] S. Fortunato, A. Flammini, F. Menczer, A. Vespignani, Topical interests and the mitigation of search engine bias, *Proc. Natl. Acad. Sci.* 103 (34) (2006) 12684–12689.
- [9] B. Pan, H. Hembrooke, T. Joachims, L. Lorigo, G. Gay, L. Granka, In google we trust: Users' decisions on rank, position, and relevance, *J. Comput.-Mediat. Commun.* 12 (3) (2007) 801–823.
- [10] J. Murphy, C. Hofacker, R. Mizerski, Primacy and recency effects on clicking behavior, *J. Comput.-Mediat. Commun.* 11 (2) (2006) 522–535.
- [11] R. Zhou, S. Khemmarat, L. Gao, The impact of YouTube recommendation system on video views, in: *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ACM, 2010, pp. 404–410.
- [12] J.E. Hirsch, An index to quantify an individual's scientific research output, *Proc. Natl. Acad. Sci.* 102 (46) (2005) 16569–16572.
- [13] F. Radicchi, S. Fortunato, B. Markines, A. Vespignani, Diffusion of scientific credits and the ranking of scientists, *Phys. Rev. E* 80 (5) (2009) 056103.
- [14] F. Radicchi, Who is the best player ever? a complex network analysis of the history of professional tennis, *PLoS One* 6 (2) (2011) e17249.
- [15] A. Spitz, E.-Á. Horvát, Measuring long-term impact based on network centrality: Unraveling cinematic citations, *PLoS One* 9 (10) (2014) e108857.
- [16] M. Wasserman, X.H.T. Zeng, L.A.N. Amaral, Cross-evaluation of metrics to estimate the significance of creative works, *Proc. Natl. Acad. Sci.* 112 (5) (2015) 1281–1286.
- [17] L. Waltman, A review of the literature on citation impact indicators, *J. Informetrics* 10 (2) (2016) 365–391.
- [18] J. Wilsdon, *The Metric Tide: Independent Review of the Role of Metrics in Research Assessment and Management*, SAGE, 2016, pp. 1–190.
- [19] D. Brockmann, D. Helbing, The hidden geometry of complex, network-driven contagion phenomena, *Science* 342 (6164) (2013) 1337–1342.
- [20] F. Iannelli, A. Koher, D. Brockmann, P. Hoevel, I.M. Sokolov, Effective distances for epidemics spreading on complex networks, 2016. *ArXiv preprint arXiv:1608.06201*.
- [21] Z. Wang, C.T. Bauch, S. Bhattacharyya, A. d'Onofrio, P. Manfredi, M. Perc, N. Perra, M. Salathé, D. Zhao, Statistical physics of vaccination, *Phys. Rep.* 664 (2016) 1–113.
- [22] R. Epstein, R.E. Robertson, The search engine manipulation effect and its possible impact on the outcomes of elections, *Proc. Natl. Acad. Sci.* 112 (33) (2015) E4512–E4521.
- [23] D. Lazer, A.S. Pentland, L. Adamic, S. Aral, A.L. Barabási, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, M. Van Alstyne, *Computational social science*, *Science* 323 (5915) (2009) 721.
- [24] A. McAfee, E. Brynjolfsson, T.H. Davenport, D. Patil, D. Barton, Big data, *Harv. Bus. Rev.* 90 (10) (2012) 61–67.
- [25] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, D.-U. Hwang, *Complex networks: Structure and dynamics*, *Phys. Rep.* 424 (4) (2006) 175–308.
- [26] M. Newman, *Networks: An Introduction*, Oxford University Press, 2010, pp. 1–720.
- [27] M.O. Jackson, *Social and Economic Networks*, Princeton University Press, 2010.
- [28] A.-L. Barabási, *Network Science*, Cambridge University Press, 2016, pp. 1–474.
- [29] D. Balcan, H. Hu, B. Gonçalves, P. Bajardi, C. Poletto, J.J. Ramasco, D. Paolotti, N. Perra, M. Tizzoni, W. Van den Broeck, et al., Seasonal transmission potential and activity peaks of the new influenza A (H1N1): a Monte Carlo likelihood analysis based on human mobility, *BMC Med.* 7 (1) (2009) 45.
- [30] C.A. Hidalgo, R. Hausmann, The building blocks of economic complexity, *Proc. Natl. Acad. Sci.* 106 (26) (2009) 10570–10575.
- [31] A. Tacchella, M. Cristelli, G. Caldarelli, A. Gabrielli, L. Pietronero, A new metrics for countries' fitness and products' complexity, *Sci. Rep.* 2 (2013) 723.
- [32] M. Cristelli, A. Tacchella, L. Pietronero, The heterogeneous dynamics of economic complexity, *PLoS One* 10 (2) (2015) e0117174.
- [33] S. Battiston, M. Puliga, R. Kaushik, P. Tasca, G. Caldarelli, Debtrank: Too central to fail? financial networks, the fed and systemic risk, *Sci. Rep.* 2 (2012) 541.
- [34] D.Y. Kenett, M. Raddant, T. Lux, E. Ben-Jacob, Evolvement of uniformity and volatility in the stressed global financial village, *PLoS One* 7 (2) (2012) e31144.
- [35] S. Battiston, J.D. Farmer, A. Flache, D. Garlaschelli, A.G. Haldane, H. Heesterbeek, C. Hommes, C. Jaeger, R. May, M. Scheffer, Complexity theory and financial regulation, *Science* 351 (6275) (2016) 818–819.
- [36] C.M. Schneider, A.A. Moreira, J.S. Andrade, S. Havlin, H.J. Herrmann, Mitigation of malicious attacks on networks, *Proc. Natl. Acad. Sci. USA* 108 (10) (2011) 3838–3841.
- [37] S.D. Reis, Y. Hu, A. Babino, J.S. Andrade Jr., S. Canals, M. Sigman, H.A. Makse, Avoiding catastrophic failure in correlated networks of networks, *Nat. Phys.* 10 (10) (2014) 762–767.
- [38] N. Duhan, A. Sharma, K.K. Bhatia, Page ranking algorithms: A survey, in: *Advance Computing Conference, 2009. IACC 2009. IEEE International, IEEE, 2009*, pp. 1530–1537.
- [39] M. Medo, Network-based information filtering algorithms: Ranking and recommendation, in: *Dynamics on and of Complex Networks*, vol. 2, Springer, 2013, pp. 315–334.
- [40] L. Katz, A new status index derived from sociometric analysis, *Psychometrika* 18 (1) (1953) 39–43.
- [41] P. Bonacich, Power and centrality: A family of measures, *Amer. J. Sociol.* (1987) 1170–1182.
- [42] S.P. Borgatti, Centrality and AIDS, *Connections* 18 (1) (1995) 112–114.
- [43] S. Brin, L. Page, The anatomy of a large-scale hypertextual Web search engine, *Comput. Netw. ISDN Syst.* 30 (1) (1998) 107–117.
- [44] M. Rosvall, A.V. Esquivel, A. Lancichinetti, J.D. West, R. Lambiotte, Memory in network flows and its effects on spreading dynamics and community detection, *Nature Commun.* 5 (2014) 4630.
- [45] L.E. Rocha, N. Masuda, Random walk centrality for temporal networks, *New J. Phys.* 16 (6) (2014) 063023.
- [46] N. Masuda, M.A. Porter, R. Lambiotte, *Random walks and diffusion on networks*, 2016. *ArXiv preprint arXiv:1612.03281*.
- [47] Z.-K. Zhang, C. Liu, X.-X. Zhan, X. Lu, C.-X. Zhang, Y.-C. Zhang, Dynamics of information diffusion and its applications on complex networks, *Phys. Rep.* 651 (2016) 1–34.
- [48] Z.-M. Ren, Y. Kong, M.-S. Shang, Y.-C. Zhang, A generalized model via random walks for information filtering, *Phys. Lett. A* 380 (34) (2016) 2608–2614.
- [49] M. Piraveenan, M. Prokopenko, L. Hossain, Percolation centrality: Quantifying graph-theoretic impact of nodes during percolation in networks, *PLoS One* 8 (1) (2013) e53095.
- [50] F. Morone, H. Makse, Influence maximization in complex networks through optimal percolation, *Nature* 524 (7563) (2015) 65–68.
- [51] F. Radicchi, C. Castellano, Leveraging percolation theory to single out influential spreaders in networks, 2016. *ArXiv preprint arXiv:1605.07041*.
- [52] A.N. Langville, C.D. Meyer, *Google's PageRank and Beyond: The Science of Search Engine Rankings*, Princeton University Press, 2006, pp. 1–240.
- [53] Z.-K. Zhang, T. Zhou, Y.-C. Zhang, Personalized recommendation via integrated diffusion on user–item–tag tripartite graphs, *Physica A* 389 (1) (2010) 179–186.
- [54] Y.-C. Zhang, M. Blattner, Y.-K. Yu, Heat conduction process on community networks as a recommendation model, *Phys. Rev. Lett.* 99 (15) (2007) 154301.
- [55] T. Zhou, Z. Kuscsik, J.-G. Liu, M. Medo, J.R. Wakeling, Y.-C. Zhang, Solving the apparent diversity-accuracy dilemma of recommender systems, *Proc. Natl. Acad. Sci.* 107 (10) (2010) 4511–4515.
- [56] P. Chen, H. Xie, S. Maslov, S. Redner, Finding scientific gems with Google's PageRank algorithm, *J. Informetrics* 1 (1) (2007) 8–15.

- [57] Y. Ding, E. Yan, A. Frazho, J. Caverlee, PageRank for ranking authors in co-citation networks, *J. Am. Soc. Inf. Sci. Technol.* 60 (11) (2009) 2229–2243.
- [58] Z.-M. Ren, A. Zeng, D.-B. Chen, H. Liao, J.-G. Liu, Iterative resource allocation for ranking spreaders in complex networks, *Europhys. Lett.* 106 (4) (2014) 48005.
- [59] S. Pei, L. Muchnik, J.S. Andrade Jr., Z. Zheng, H.A. Makse, Searching for superspreaders of information in real-world social media, *Sci. Rep.* 4 (2014) 5547.
- [60] M. Franceschet, PageRank: Standing on the shoulders of giants, *Commun. ACM* 54 (6) (2011) 92–101.
- [61] D.F. Gleich, PageRank beyond the Web, *SIAM Rev.* 57 (3) (2015) 321–363.
- [62] L. Ermann, K.M. Frahm, D.L. Shepelyansky, Google matrix analysis of directed networks, *Rev. Modern Phys.* 87 (4) (2015) 1261.
- [63] B. Jiang, S. Zhao, J. Yin, Self-organized natural roads for predicting traffic flow: A sensitivity study, *J. Stat. Mech. Theory Exp.* 2008 (07) (2008) P07008.
- [64] H. Mao, X. Shuai, Y.-Y. Ahn, J. Bollen, Quantifying socio-economic indicators in developing countries from mobile phone communication data: Applications to Côte d'Ivoire, *EPJ Data Sci.* 4 (1) (2015) 1–16.
- [65] D. Walker, H. Xie, K.-K. Yan, S. Maslov, Ranking scientific publications using a model of network traffic, *J. Stat. Mech. Theory Exp.* 2007 (06) (2007) P06010.
- [66] J. Bollen, M.A. Rodriguez, H. Van de Sompel, Journal status, *Scientometrics* 69 (3) (2006) 669–687.
- [67] Y. Jing, S. Baluja, VisualRank: Applying PageRank to large-scale image search, *IEEE Trans. Pattern Anal. Mach. Intell.* 30 (11) (2008) 1877–1890.
- [68] G. Iván, V. Grolmusz, When the Web meets the cell: using personalized PageRank for analyzing protein interaction networks, *Bioinformatics* 27 (3) (2011) 405–407.
- [69] L. Lü, M. Medo, C.H. Yeung, Y.-C. Zhang, Z.-K. Zhang, T. Zhou, Recommender systems, *Phys. Rep.* 519 (1) (2012) 1–49.
- [70] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, *Science* 286 (5439) (1999) 509–512.
- [71] G. Bianconi, A.-L. Barabási, Competition and multiscaling in evolving networks, *Europhys. Lett.* 54 (4) (2001) 436.
- [72] S.N. Dorogovtsev, J.F. Mendes, Evolution of networks, *Adv. Phys.* 51 (4) (2002) 1079–1187.
- [73] R. Albert, A.-L. Barabási, Statistical mechanics of complex networks, *Rev. Modern Phys.* 74 (1) (2002) 47.
- [74] M. Medo, G. Cimini, S. Gualdi, Temporal effects in the growth of networks, *Phys. Rev. Lett.* 107 (23) (2011) 238701.
- [75] F. Papadopoulos, M. Kitsak, M.Á. Serrano, M. Boguná, D. Krioukov, Popularity versus similarity in growing networks, *Nature* 489 (7417) (2012) 537–540.
- [76] D. Wang, C. Song, A.-L. Barabási, Quantifying long-term scientific impact, *Science* 342 (6154) (2013) 127–132.
- [77] M. Newman, The first-mover advantage in scientific publication, *Europhys. Lett.* 86 (6) (2009) 68001.
- [78] N. Perra, A. Baronchelli, D. Mocanu, B. Gonçalves, R. Pastor-Satorras, A. Vespignani, Random walks and search in time-varying networks, *Phys. Rev. Lett.* 109 (23) (2012) 238701.
- [79] I. Scholtes, N. Wider, R. Pfitzner, A. Garas, C.J. Tessone, F. Schweitzer, Causality-driven slow-down and speed-up of diffusion in non-Markovian temporal networks, *Nature Commun.* 5 (2014).
- [80] R. Lambiotte, V. Salnikov, M. Rosvall, Effect of memory on the dynamics of random walks on networks, *J. Complex Netw.* 3 (2) (2015) 177–188.
- [81] J.-C. Delvenne, R. Lambiotte, L.E. Rocha, Diffusion on networked systems is a question of time or structure, *Nature Commun.* 6 (2015).
- [82] M.S. Mariani, M. Medo, Y.-C. Zhang, Ranking nodes in growing networks: When PageRank fails, *Sci. Rep.* 5 (2015) 16181.
- [83] A. Vidmer, M. Medo, The essential role of time in network-based recommendation, *Europhys. Lett.* 116 (2016) 30007.
- [84] L.C. Freeman, A set of measures of centrality based on betweenness, *Sociometry* (1977) 35–41.
- [85] G. Caldarelli, Scale-Free Networks: Complex Webs in Nature and Technology, Oxford University Press, 2007, pp. 1–336.
- [86] M. E.J. Newman, The structure and function of complex networks, *SIAM Rev.* 45 (2003) 167.
- [87] S. Fortunato, M. Boguñá, A. Flammini, F. Menczer, Approximating PageRank from in-degree, in: *International Workshop on Algorithms and Models for the Web-Graph*, Springer, 2006, pp. 59–71.
- [88] R. Pastor-Satorras, C. Castellano, Topological structure of the H -index in complex networks, 2016. ArXiv preprint arXiv:1610.00569.
- [89] K. Klemm, M.Á. Serrano, V.M. Eguiluz, M. San Miguel, A measure of individual role in collective dynamics, *Sci. Rep.* 2 (292) (2012).
- [90] L. Lü, D. Chen, X.-L. Ren, Q.-M. Zhang, Y.-C. Zhang, T. Zhou, Vital nodes identification in complex networks, *Phys. Rep.* 650 (2016) 1–63.
- [91] J. Kunegis, A. Lommatzsch, C. Bauckhage, The slashdot zoo: mining a social network with negative edges, in: *Proceedings of the 18th International Conference on World Wide Web*, ACM, 2009, pp. 741–750.
- [92] L. Bornmann, H.-D. Daniel, What do citation counts measure? A review of studies on citing behavior, *J. Doc.* 64 (1) (2008) 45–80.
- [93] C. Catalini, N. Lacetera, A. Oettl, The incidence and role of negative citations in science, *Proc. Natl. Acad. Sci. USA* 112 (45) (2015) 13823–13826.
- [94] L. Lü, T. Zhou, Q.-M. Zhang, H.E. Stanley, The H -index of a network node and its relation to degree and coreness, *Nature Commun.* 7 (2016) 10168.
- [95] D. Chen, L. Lü, M.-S. Shang, Y.-C. Zhang, T. Zhou, Identifying influential nodes in complex networks, *Physica A* 391 (4) (2012) 1777–1787.
- [96] Y.-Y. Chen, Q. Gan, T. Suel, Local methods for estimating PageRank values, in: *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management*, ACM, 2004, pp. 381–389.
- [97] G. Sabidussi, The centrality index of a graph, *Psychometrika* 31 (4) (1966) 581–603.
- [98] Y. Rochat, Closeness centrality extended to unconnected graphs: The harmonic centrality index, in: *ASNA*, no. EPFL-CONF-200525, 2009.
- [99] P. Boldi, S. Vigna, Axioms for centrality, *Internet Math.* 10 (3–4) (2014) 222–262.
- [100] S.P. Borgatti, Centrality and network flow, *Social Networks* 27 (1) (2005) 55–71.
- [101] M.E. Newman, A measure of betweenness centrality based on random walks, *Social Networks* 27 (1) (2005) 39–54.
- [102] M. Kitsak, L.K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H.E. Stanley, H.A. Makse, Identification of influential spreaders in complex networks, *Nat. Phys.* 6 (11) (2010) 888–893.
- [103] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, E. Shir, A model of internet topology using k -shell decomposition, *Proc. Natl. Acad. Sci.* 104 (27) (2007) 11150–11154.
- [104] A. Garas, F. Schweitzer, S. Havlin, A k -shell decomposition method for weighted networks, *New J. Phys.* 14 (8) (2012) 083030.
- [105] J.-G. Liu, Z.-M. Ren, Q. Guo, Ranking the spreading influence in complex networks, *Physica A* 392 (18) (2013) 4154–4159.
- [106] T. Kawamoto, Localized eigenvectors of the non-backtracking matrix, *J. Stat. Mech. Theory Exp.* 2016 (2) (2016) 023404.
- [107] O. Perron, Zur theorie der matrices, *Math. Ann.* 64 (2) (1907) 248–263.
- [108] C.H. Hubbell, An input–output approach to clique identification, *Sociometry* (1965) 377–399.
- [109] P. Bonacich, P. Lloyd, Eigenvector-like measures of centrality for asymmetric relations, *Social Networks* 23 (3) (2001) 191–201.
- [110] J.M. Fletcher, T. Wennekers, From structure to activity: Using centrality measures to predict neuronal activity, *Int. J. Neural Syst.* 0 (16) (2016) 1750013.
- [111] J. Park, M.E. Newman, A network-based ranking system for US college football, *J. Stat. Mech. Theory Exp.* 2005 (10) (2005) P10014.
- [112] P. Boldi, M. Santini, S. Vigna, PageRank as a function of the damping factor, in: *Proceedings of the 14th International Conference on World Wide Web*, ACM, 2005, pp. 557–566.
- [113] R. Lambiotte, M. Rosvall, Ranking and clustering of nodes in networks with smart teleportation, *Phys. Rev. E* 85 (5) (2012) 056107.
- [114] P. Berkhin, A survey on PageRank computing, *Internet Math.* 2 (1) (2005) 73–120.
- [115] N. Perra, S. Fortunato, Spectral centrality measures in complex networks, *Phys. Rev. E* 78 (3) (2008) 036107.
- [116] M. Bianchini, M. Gori, F. Scarselli, Inside pagerank, *ACM Trans. Internet Technol.* 5 (1) (2005) 92–128.

- [117] K. Avrachenkov, N. Litvak, K.S. Pham, A singular perturbation approach for choosing the PageRank damping factor, *Internet Math.* 5 (1–2) (2008) 47–69.
- [118] D. Fogaras, Where to start browsing the web? in: *International Workshop on Innovative Internet Community Systems*, Springer, 2003, pp. 65–79.
- [119] A. Zhurov, O. Zhurov, D. Shepelyansky, Two-dimensional ranking of Wikipedia articles, *Eur. Phys. J. B* 77 (4) (2010) 523–531.
- [120] L. Lü, Y.-C. Zhang, C.H. Yeung, T. Zhou, Leaders in social networks, the Delicious case, *PLoS One* 6 (6) (2011) e21202.
- [121] Q. Li, T. Zhou, L. Lü, D. Chen, Identifying influential spreaders by weighted LeaderRank, *Physica A* 404 (2014) 47–55.
- [122] Y. Zhou, L. Lü, W. Liu, J. Zhang, The power of ground user in recommender systems, *PLoS One* 8 (8) (2013) e70094.
- [123] J.M. Kleinberg, Authoritative sources in a hyperlinked environment, *J. ACM* 46 (5) (1999) 604–632.
- [124] H. Li, I.G. Councill, L. Bolelli, D. Zhou, Y. Song, W.-C. Lee, A. Sivasubramaniam, C.L. Giles, CiteSeer χ : a scalable autonomous scientific digital library, in: *Proceedings of the 1st International Conference on Scalable Information Systems*, ACM, 2006, p. 18.
- [125] A.Y. Ng, A.X. Zheng, M.I. Jordan, Stable algorithms for link analysis, in: *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM, 2001, pp. 258–266.
- [126] H. Deng, M.R. Lyu, I. King, A generalized Co-HITS algorithm and its application to bipartite graphs, in: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2009, pp. 239–248.
- [127] W.W. Zachary, An information flow model for conflict and fission in small groups, *J. Anthropol. Res.* 33 (4) (1977) 452–473. URL <http://www.jstor.org/stable/3629752>.
- [128] S. Fortunato, Community detection in graphs, *Phys. Rep.* 486 (3) (2010) 75–174.
- [129] S. Fortunato, D. Hric, Community detection in networks: A user guide, *Phys. Rep.* 659 (2016) 1–44.
- [130] A. Vidmer, M. Medo, Y.-C. Zhang, Unbiased metrics of friends' influence in multi-level networks, *EPJ Data Sci.* 4 (1) (2015) 1–13.
- [131] E. Sarigöl, R. Pfitzner, I. Scholtes, A. Garas, F. Schweitzer, Predicting scientific success based on coauthorship networks, *EPJ Data Sci.* 3 (1) (2014) 1.
- [132] M. Kivela, A. Arenas, M. Barthelemy, J.P. Gleeson, Y. Moreno, M.A. Porter, Multilayer networks, *J. Complex Netw.* 2 (3) (2014) 203–271.
- [133] S. Boccaletti, G. Bianconi, R. Criado, C.I. Del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendiña-Nadal, Z. Wang, M. Zanin, The structure and dynamics of multilayer networks, *Phys. Rep.* 544 (1) (2014) 1–122.
- [134] A. Solé-Ribalta, M. De Domenico, S. Gómez, A. Arenas, Centrality rankings in multiplex networks, in: *Proceedings of the 2014 ACM Conference on Web Science*, ACM, 2014, pp. 149–155.
- [135] M. De Domenico, A. Solé-Ribalta, E. Omodei, S. Gómez, A. Arenas, Ranking in interconnected multilayer networks reveals versatile nodes, *Nature Commun.* 6 (2015) 6868.
- [136] D. Taylor, S.A. Myers, A. Clauset, M.A. Porter, P.J. Mucha, Eigenvector-based centrality measures for temporal networks, 2015. ArXiv preprint arXiv:1507.01266.
- [137] G. Caldarelli, M. Cristelli, A. Gabrielli, L. Pietronero, A. Scala, A. Tacchella, A network analysis of countries' export flows: Firm grounds for the building blocks of the economy, *PLoS One* 7 (10) (2012) e47278.
- [138] M. Cristelli, A. Gabrielli, A. Tacchella, G. Caldarelli, L. Pietronero, Measuring the intangibles: A metrics for the economic complexity of countries and products, *PLoS One* 8 (8) (2013) e70726.
- [139] R. Hausmann, C.A. Hidalgo, S. Bustos, M. Coscia, A. Simoes, M.A. Yildirim, *The Alas of Economic Complexity: Mapping Paths to Prosperity*, MIT Press, 2014, pp. 1–368.
- [140] M.S. Mariani, A. Vidmer, M. Medo, Y.-C. Zhang, Measuring economic complexity of countries and products: Which metric to use? *Eur. Phys. J. B* 88 (11) (2015) 1–9.
- [141] R.-J. Wu, G.-Y. Shi, Y.-C. Zhang, M.S. Mariani, The mathematics of non-linear metrics for nested networks, *Physica A* 460 (2016) 254–269.
- [142] V. Stojkoski, Z. Utkovski, L. Kocarev, The impact of services on economic complexity: Service sophistication as route for economic growth, 2016. ArXiv preprint arXiv:1604.06284.
- [143] V. Domínguez-García, M.A. Muñoz, Ranking species in mutualistic networks, *Sci. Rep.* 5 (2015) 8182.
- [144] P. Resnick, K. Kuwabara, R. Zeckhauser, E. Friedman, Reputation systems, *Commun. ACM* 43 (12) (2000) 45–48.
- [145] A. Jøsang, R. Ismail, C. Boyd, A survey of trust and reputation systems for online service provision, *Dec. Support Syst.* 43 (2) (2007) 618–644.
- [146] I. Pinyol, J. Sabater-Mir, Computational trust and reputation models for open multi-agent systems: a review, *Artif. Intell. Rev.* 40 (1) (2013) 1–25.
- [147] D.G. Gregg, J.E. Scott, The role of reputation systems in reducing on-line auction fraud, *Int. J. Electron. Commer.* 10 (3) (2006) 95–120.
- [148] C.G. McDonald, V.C. Slawson, Reputation in an Internet auction market, *Economic Inquiry* 40 (4) (2002) 633–650.
- [149] G. Wang, S. Xie, B. Liu, S.Y. Philip, Review graph based online store review spammer detection, in: *IEEE 11th International Conference on Data Mining, IEEE, 2011*, pp. 1242–1247.
- [150] F. Benevenuto, T. Rodrigues, V. Almeida, J. Almeida, M. Gonçalves, Detecting spammers and content promoters in online video social networks, in: *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM, 2009, pp. 620–627.
- [151] H. Masum, Y.-C. Zhang, Manifesto for the reputation society, *First Monday* 9 (7) (2004).
- [152] P. Laureti, L. Moret, Y.-C. Zhang, Y.-K. Yu, Information filtering via iterative refinement, *Europhys. Lett.* 75 (6) (2006) 1006.
- [153] Y.-K. Yu, Y.-C. Zhang, P. Laureti, L. Moret, Decoding information from noisy, redundant, and intentionally distorted sources, *Physica A* 371 (2) (2006) 732–744.
- [154] M. Medo, J.R. Wakeling, The effect of discrete vs. continuous-valued ratings on reputation and ranking systems, *Europhys. Lett.* 91 (4) (2010) 48004.
- [155] Y.B. Zhou, T. Lei, T. Zhou, A robust ranking algorithm to spamming, *Europhys. Lett.* 94 (4) (2011) 48002.
- [156] H. Liao, A. Zeng, R. Xiao, Z.-M. Ren, D.-B. Chen, Y.-C. Zhang, Ranking reputation and quality in online rating systems, *PLoS One* 9 (5) (2014) e97146.
- [157] D.d.s. Price, A general theory of bibliometric and other cumulative advantage processes, *J. Amer. Soc. Inf. Sci.* 27 (5) (1976) 292–306.
- [158] H. Jeong, Z. Neda, A.-L. Barabási, Measuring preferential attachment in evolving networks, *Europhys. Lett.* 61 (4) (2003) 567.
- [159] S. Redner, Citation statistics from 110 years of physical review, *Phys. Today* 58 (2005) 49.
- [160] P.L. Krapivsky, S. Redner, Organization of growing random networks, *Phys. Rev. E* 63 (6) (2001) 066123.
- [161] M. Medo, G. Cimini, Model-based evaluation of scientific impact indicators, *Phys. Rev. E* 94 (3) (2016) 032312.
- [162] M. Golosovsky, S. Solomon, Growing complex network of citations of scientific papers: Modeling and measurements, *Phys. Rev. E* 95 (2017) 012324.
- [163] M. Newman, Prediction of highly cited papers, *Europhys. Lett.* 105 (2) (2014) 28002.
- [164] R. Baeza-Yates, C. Castillo, F. Saint-Jean, Web dynamics, structure, and page quality, in: *Web Dynamics*, Springer, 2004, pp. 93–109.
- [165] S. Maslov, S. Redner, Promise and pitfalls of extending Google's PageRank algorithm to citation networks, *J. Neurosci.* 28 (44) (2008) 11103–11105.
- [166] M.S. Mariani, M. Medo, Y.-C. Zhang, Identification of milestone papers through time-balanced network centrality, *J. Informetrics* 10 (4) (2016) 1207–1223.
- [167] P.D.B. Parolo, R.K. Pan, R. Ghosh, B.A. Huberman, K. Kaski, S. Fortunato, Attention decay in science, *J. Informetrics* 9 (4) (2015) 734–745.
- [168] P.V. Marsden, N.E. Friedkin, Network studies of social influence, *Sociol. Methods Res.* 22 (1) (1993) 127–151.
- [169] D. Kempe, J. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2003, pp. 137–146.
- [170] J. Tang, J. Sun, C. Wang, Z. Yang, Social influence analysis in large-scale networks, in: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2009, pp. 807–816.

- [171] M. Cha, H. Haddadi, F. Benevenuto, P.K. Gummadi, Measuring user influence in Twitter: the million follower fallacy, *ICWSM 10 (10–17)* (2010) 30.
- [172] J. Leskovec, L.A. Adamic, B.A. Huberman, The dynamics of viral marketing, *ACM Trans. Web 1 (1)* (2007) 5.
- [173] M. Cha, A. Mislove, K.P. Gummadi, A measurement-driven analysis of information propagation in the flickr social network, in: *Proceedings of the 18th International Conference on World Wide Web*, ACM, 2009, pp. 721–730.
- [174] G.V. Steeg, R. Ghosh, K. Lerman, What stops social epidemics? 2011. ArXiv preprint arXiv:1102.1985.
- [175] F. Radicchi, S. Fortunato, C. Castellano, Universality of citation distributions: Toward an objective measure of scientific impact, *Proc. Natl. Acad. Sci. 105 (45)* (2008) 17268–17272.
- [176] F. Radicchi, C. Castellano, Rescaling citations of publications in physics, *Phys. Rev. E 83 (4)* (2011) 046116.
- [177] P. Albarrán, J.A. Crespo, I. Ortuño, J. Ruiz-Castillo, The skewness of science in 219 sub-fields and a number of aggregates, *Scientometrics 88 (2)* (2011) 385–397.
- [178] L. Waltman, N.J. van Eck, A.F. van Raan, Universality of citation distributions revisited, *J. Amer. Soc. Inf. Sci. Technol. 63 (1)* (2012) 72–77.
- [179] G. Vaccario, M. Medo, N. Wider, M.S. Mariani, Quantifying and suppressing ranking bias in a large citation network, 2017. arXiv:1703.08071.
- [180] F. Radicchi, C. Castellano, A reverse engineering approach to the suppression of citation biases reveals universal properties of citation distributions, *PLoS One 7 (3)* (2012) e33833.
- [181] A. Zeng, S. Gualdi, M. Medo, Y.-C. Zhang, Trend prediction in temporal bipartite networks: The case of MovieLens, Netflix, and Digg, *Adv. Complex Syst. 16 (04n05)* (2013) 1350024.
- [182] Y. Koren, Collaborative filtering with temporal dynamics, *Commun. ACM 53 (4)* (2010) 89–97.
- [183] Y. Zhou, A. Zeng, W.-H. Wang, Temporal effects in trend prediction: Identifying the most popular nodes in the future, *PLoS One 10 (3)* (2015) e0120735.
- [184] R. Ghosh, T.-T. Kuo, C.-N. Hsu, S.-D. Lin, K. Lerman, Time-aware ranking in dynamic citation networks, in: *IEEE 11th International Conference on Data Mining Workshops, ICDMW, IEEE, 2011*, pp. 373–380.
- [185] P.S. Yu, X. Li, B. Liu, Adding the temporal dimension to search a case study in publication search, in: *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence*, IEEE Computer Society, 2005, pp. 543–549.
- [186] K. Berberich, M. Vazirgiannis, G. Weikum, T-rank: Time-aware authority ranking, in: *Algorithms and Models for the Web-Graph*, Springer, 2004, pp. 131–142.
- [187] K. Berberich, M. Vazirgiannis, G. Weikum, Time-aware authority ranking, *Internet Math. 2 (3)* (2005) 301–332.
- [188] J. Sabater, C. Sierra, Regret: A reputation model for gregarious societies, in: *Fourth Workshop on Deception Fraud and Trust in Agent Societies*, vol. 70, 2001, pp. 61–69.
- [189] A. Jø sang, R. Ismail, The beta reputation system, in: *Proceedings of the 15th Bled Electronic Commerce Conference*, vol. 5, 2002, pp. 2502–2511.
- [190] Y. Liu, Y. Sun, Anomaly detection in feedback-based reputation systems through temporal and correlation analysis, in: *2010 IEEE Second International Conference on Social Computing*, IEEE, 2010, pp. 65–72.
- [191] J.S. Kong, N. Sarshar, V.P. Roychowdhury, Experience versus talent shapes the structure of the Web, *Proc. Natl. Acad. Sci. 105 (37)* (2008) 13724–13729.
- [192] Z.-M. Ren, Y.-Q. Shi, H. Liao, Characterizing popularity dynamics of online videos, *Physica A 453* (2016) 236–241.
- [193] S.N. Dorogovtsev, J.F.F. Mendes, Evolution of networks with aging of sites, *Phys. Rev. E 62 (2)* (2000) 1842.
- [194] M. Medo, Statistical validation of high-dimensional models of growing networks, *Phys. Rev. E 89 (3)* (2014) 032801.
- [195] K. Berberich, S. Bedathur, M. Vazirgiannis, G. Weikum, BuzzRank... and the trend is your friend, in: *Proceedings of the 15th International Conference on World Wide Web*, ACM, 2006, pp. 937–938.
- [196] P. Holme, J. Saramäki, Temporal networks, *Phys. Rep. 519 (3)* (2012) 97–125.
- [197] D. Kempe, J. Kleinberg, A. Kumar, Connectivity and inference problems for temporal networks, in: *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, ACM, 2000, pp. 504–513.
- [198] V. Kostakos, Temporal graphs, *Physica A 388 (6)* (2009) 1007–1023.
- [199] A. Casteigts, P. Flocchini, W. Quattrociocchi, N. Santoro, Time-varying graphs and dynamic networks, *Int. J. Parallel Emergent Distrib. Syst. 27 (5)* (2012) 387–408.
- [200] K.A. Berman, Vulnerability of scheduled networks and a generalization of Menger's theorem, *Networks 28 (3)* (1996) 125–134.
- [201] P. Holme, Modern temporal network theory: A colloquium, *Eur. Phys. J. B 88 (9)* (2015) 1–30.
- [202] J. Xu, T.L. Wickramaratne, N.V. Chawla, Representing higher-order dependencies in networks, *Sci. Adv. 2 (5)* (2016) e1600028.
- [203] G. Krings, M. Karsai, S. Bernhardsson, V.D. Blondel, J. Saramäki, Effects of time window size and placement on the structure of an aggregated communication network, *EPJ Data Sci. 1 (1)* (2012) 4.
- [204] V. Sekara, A. Stopczynski, S. Lehmann, Fundamental structures of dynamic social networks, *Proc. Natl. Acad. Sci. 113 (36)* (2016) 9977–9982.
- [205] J. Moody, The importance of relationship timing for diffusion, *Soc. Forces 81 (1)* (2002) 25–56.
- [206] H.H. Lentz, T. Selhorst, I.M. Sokolov, Unfolding accessibility provides a macroscopic approach to temporal networks, *Phys. Rev. Lett. 110 (11)* (2013) 118701.
- [207] R. Pfitzner, I. Scholtes, A. Garas, C.J. Tessone, F. Schweitzer, Betweenness preference: Quantifying correlations in the topological dynamics of temporal networks, *Phys. Rev. Lett. 110 (19)* (2013) 198701.
- [208] J. Tang, M. Musolesi, C. Mascolo, V. Latora, V. Nicosia, Analysing information flows and key mediators through temporal centrality metrics, in: *Proceedings of the 3rd Workshop on Social Network Systems*, ACM, 2010, p. 3.
- [209] V. Nicosia, J. Tang, C. Mascolo, M. Musolesi, G. Russo, V. Latora, Graph metrics for temporal networks, in: *Temporal Networks*, Springer, 2013, pp. 15–40.
- [210] R.K. Pan, J. Saramäki, Path lengths, correlations, and centrality in temporal networks, *Phys. Rev. E 84 (1)* (2011) 016105.
- [211] I. Scholtes, N. Wider, A. Garas, Higher-order aggregate networks in the analysis of temporal networks: Path structures and centralities, *Eur. Phys. J. B 89 (3)* (2016) 1–15.
- [212] M. Karsai, M. Kivelä, R.K. Pan, K. Kaski, J. Kertész, A.-L. Barabási, J. Saramäki, Small but slow world: How network topology and burstiness slow down spreading, *Phys. Rev. E 83 (2)* (2011) 025102.
- [213] M. Starnini, A. Baronchelli, A. Barrat, R. Pastor-Satorras, Random walks on temporal networks, *Phys. Rev. E 85 (5)* (2012) 056115.
- [214] N. Masuda, K. Klemm, V.M. Eguíluz, Temporal networks: Slowing down diffusion by long lasting interactions, *Phys. Rev. Lett. 111 (18)* (2013) 188701.
- [215] I. Scholtes, When is a network a network? multi-order graphical model selection in pathways and temporal networks, 2017. ArXiv preprint arXiv:1702.05499.
- [216] S. Motegi, N. Masuda, A network-based dynamical ranking system for competitive sports, *Sci. Rep. 2* (2012) 904.
- [217] P.S.P. Júnior, M.A. Gonçalves, A.H. Laender, T. Salles, D. Figueiredo, Time-aware ranking in sport social networks, *J. Inf. Data Manag. 3 (3)* (2012) 195.
- [218] H. Kim, R. Anderson, Temporal node centrality in complex networks, *Phys. Rev. E 85 (2)* (2012) 026107.
- [219] J.B. Schafer, D. Frankowski, J. Herlocker, S. Sen, Collaborative filtering recommender systems, in: *The Adaptive Web*, Springer, 2007, pp. 291–324.
- [220] Y. Koren, R. Bell, Advances in collaborative filtering, in: *Recommender Systems Handbook*, Springer, 2011, pp. 145–186.
- [221] J. Bobadilla, F. Ortega, A. Hernando, A. Gutiérrez, Recommender systems survey, *Knowl.-Based Syst. 46* (2013) 109–132.
- [222] R.M. Bell, Y. Koren, Lessons from the netflix prize challenge, *ACM SIGKDD Explor. Newsl. 9 (2)* (2007) 75–79.
- [223] J. Bennett, S. Lanning, The netflix prize, in: *Proceedings of KDD Cup and Workshop*, vol. 2007, 2007, p. 35.

- [224] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, A. Bouchachia, A survey on concept drift adaptation, *ACM Comput. Surv.* 46 (4) (2014) 44.
- [225] J.S. Breese, D. Heckerman, C. Kadie, Empirical analysis of predictive algorithms for collaborative filtering, in: *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc., 1998, pp. 43–52.
- [226] G. Takács, I. Pilászy, B. Németh, D. Tikk, Major components of the Gravity recommendation system, *ACM SIGKDD Explor. News.* 9 (2) (2007) 80–83.
- [227] Y. Koren, R. Bell, C. Volinsky, et al., Matrix factorization techniques for recommender systems, *Computer* 42 (8) (2009) 30–37.
- [228] W. Pan, Z. Liu, Z. Ming, H. Zhong, X. Wang, C. Xu, Compressed knowledge transfer via factorization machine for heterogeneous collaborative recommendation, *Knowl.-Based Syst.* 85 (2015) 234–244.
- [229] Y. Koren, Factorization meets the neighborhood: a multifaceted collaborative filtering model, in: *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2008, pp. 426–434.
- [230] Y.S. Abu-Mostafa, M. Magdon-Ismael, H.-T. Lin, *Learning from data*, AMLBook, New York, NY, USA, 2012.
- [231] R.R. Picard, R.D. Cook, Cross-validation of regression models, *J. Amer. Statist. Assoc.* 79 (387) (1984) 575–583.
- [232] S. Arlot, A. Celisse, et al., A survey of cross-validation procedures for model selection, *Statist. Surv.* 4 (2010) 40–79.
- [233] Y. Koren, Factor in the neighbors: Scalable and accurate collaborative filtering, *ACM Trans. Knowl. Discov. Data* 4 (1) (2010) 1.
- [234] L. Breiman, Bagging predictors, *Mach. Learn.* 24 (2) (1996) 123–140.
- [235] M. Jahrer, A. Töschner, R. Legenstein, Combining predictions for accurate recommender systems, in: *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2010, pp. 693–702.
- [236] T. Zhou, J. Ren, M. Medo, Y.-C. Zhang, Bipartite network projection and personal recommendation, *Phys. Rev. E* 76 (4) (2007) 046115.
- [237] F. Yu, A. Zeng, S. Gillard, M. Medo, Network-based recommendation algorithms: A review, *Physica A* 452 (2016) 192.
- [238] J.-G. Liu, T. Zhou, Q. Guo, Information filtering via biased heat conduction, *Phys. Rev. E* 84 (2011) 037101.
- [239] T. Qiu, G. Chen, Z.-K. Zhang, T. Zhou, An item-oriented recommendation algorithm on cold-start problem, *Europhys. Lett.* 95 (2011) 58003.
- [240] H. Liao, R. Xiao, G. Cimini, M. Medo, Network-driven reputation in online scientific communities, *PLoS One* 9 (12) (2014) e112022.
- [241] C.-N. Ziegler, S.M. McNee, J.A. Konstan, G. Lausen, Improving recommendation lists through topic diversification, in: *Proceedings of the 14th International Conference on World Wide Web*, ACM, 2005, pp. 22–32.
- [242] M. Zhang, N. Hurley, Avoiding monotony: Improving the diversity of recommendation lists, in: *Proceedings of the 2008 ACM Conference on Recommender Systems*, ACM, 2008, pp. 123–130.
- [243] G. Adomavicius, Y. Kwon, Improving aggregate recommendation diversity using ranking-based techniques, *IEEE Trans. Knowl. Data Eng.* 24 (2012) 896–911.
- [244] A. Zeng, C.H. Yeung, M.-S. Shang, Y.-C. Zhang, The reinforcing influence of recommendations on global diversification, *Europhys. Lett.* 97 (1) (2012) 18005.
- [245] L. Ceriani, P. Verme, The origins of the Gini index: extracts from *Variabilità e Mutabilità* (1912) by Corrado Gini, *J. Econ. Inequal.* 10 (3) (2012) 421–443.
- [246] J. Bar-Ilan, Informetrics at the beginning of the 21st century – A review, *J. Informetrics* 2 (1) (2008) 1–52.
- [247] J. Mingers, L. Leydesdorff, A review of theory and practice in scientometrics, *European J. Oper. Res.* 246 (1) (2015) 1–19.
- [248] D.J. de Solla Price, Networks of scientific papers, *Science* 149 (3683) (1965) 510–515.
- [249] G. Pinski, F. Narin, Citation influence for journal aggregates of scientific publications: Theory, with application to the literature of physics, *Inform. Process. Lett.* 12 (5) (1976) 297–312.
- [250] Z.-M. Ren, M.S. Mariani, Y.-C. Zhang, M. Medo, A time-respecting null model to explore the properties of growing networks, 2017. arXiv:1703.07656.
- [251] A.F. Van Raan, Sleeping beauties in science, *Scientometrics* 59 (3) (2004) 467–472.
- [252] Q. Ke, E. Ferrara, F. Radicchi, A. Flammini, Defining and identifying sleeping beauties in science, *Proc. Natl. Acad. Sci.* 112 (24) (2015) 7426–7431.
- [253] G. Colavizza, M. Franceschet, Clustering citation histories in the physical review, *J. Informetrics* 10 (4) (2016) 1037–1051.
- [254] J. Wang, Y. Mei, D. Hicks, Comment on “Quantifying long-term scientific impact”, *Science* 345 (6193) (2014) 149.
- [255] D. Wang, C. Song, H.-W. Shen, A.-L. Barabási, Response to Comment on “Quantifying long-term scientific impact”, *Science* 345 (6193) (2014) 149.
- [256] X. Cao, Y. Chen, K.R. Liu, A data analytic approach to quantifying scientific impact, *J. Informetrics* 10 (2) (2016) 471–484.
- [257] A.M. Petersen, S. Fortunato, R.K. Pan, K. Kaski, O. Penner, A. Rungi, M. Riccaboni, H.E. Stanley, F. Pammolli, Reputation and impact in academic careers, *Proc. Natl. Acad. Sci.* 111 (43) (2014) 15316–15321.
- [258] J.E. Hirsch, Does the *h* index have predictive power? *Proc. Natl. Acad. Sci.* 104 (49) (2007) 19193–19198.
- [259] D.E. Acuna, S. Allesina, K.P. Kording, Future impact: Predicting scientific success, *Nature* 489 (7415) (2012) 201–202.
- [260] O. Penner, R.K. Pan, A.M. Petersen, K. Kaski, S. Fortunato, On the predictability of future impact in science, *Sci. Rep.* 3 (2013) 3052.
- [261] C. Zhang, C. Liu, L. Yu, Z.-K. Zhang, T. Zhou, Identifying the academic rising stars, 2016. ArXiv preprint arXiv:1606.05752.
- [262] R. Sinatra, D. Wang, P. Deville, C. Song, A.-L. Barabási, Quantifying the evolution of individual scientific impact, *Science* 354 (6312) (2016) aaf5239.
- [263] A. Clauset, D.B. Larremore, R. Sinatra, Data-driven predictions in the science of science, *Science* 355 (6324) (2017) 477–480.
- [264] C.T. Bergstrom, J.D. West, M.A. Wiseman, The eigenfactor metrics, *J. Neurosci.* 28 (45) (2008) 11433–11434.
- [265] B. González-Pereira, V.P. Guerrero-Bote, F. Moya-Anegón, A new approach to the metric of journals’ scientific prestige: The SJR indicator, *J. Informetrics* 4 (3) (2010) 379–391.
- [266] M.E. Falagas, V.G. Alexiou, The top-ten in journal impact factor manipulation, *Archivum Immunologiae Et Therapiae Experimentalis* 56 (4) (2008) 223.
- [267] C. Wallner, Ban impact factor manipulation, *Science* 323 (5913) (2009) 461.
- [268] L. Bohlin, A. Viamontes Esquivel, A. Lancichinetti, M. Rosvall, Robustness of journal rankings by network flows with different amounts of memory, *J. Assoc. Inf. Sci. Technol.* 67 (10) (2015) 2527–2535.
- [269] R. Van Noorden, Metrics: A profusion of measures, *Nature* 465 (7300) (2010) 864–866.
- [270] M. Amin, M.A. Mabe, Impact factors: use and abuse, *Medicina (Buenos Aires)* 63 (4) (2003) 347–354.
- [271] P.M. Editors, et al., The impact factor game, *PLoS Med.* 3 (6) (2006) e291.
- [272] R. Adler, J. Ewing, P. Taylor, et al., Citation statistics, *Statist. Sci.* 24 (1) (2009) 1.
- [273] S. Kuznets, National income, 1929–1932, in: *National Income, 1929–1932*, NBER, 1934, pp. 1–12.
- [274] R. Costanza, I. Kubiszewski, E. Giovannini, H. Lovins, J. McGlade, K. Pickett, K. Ragnarsdóttir, D. Roberts, R. De Vogli, R. Wilkinson, Development: Time to leave GDP behind, *Nature* 505 (7483) (2014) 283–285.
- [275] D. Coyle, *GDP: A Brief but Affectionate History*, Princeton University Press, 2015, pp. 1–184.
- [276] World economic outlook of the international monetary fund, <http://www.imf.org/external/pubs/ft/weo/2016/01/index.htm>. (Accessed 27 June 2016).
- [277] R. Hausmann, C.A. Hidalgo, The network structure of economic output, *J. Econ. Growth* 16 (4) (2011) 309–342.
- [278] J. Felipe, U. Kumar, A. Abdon, M. Bacate, Product complexity and economic development, *Struct. Change Econ. Dyn.* 23 (1) (2012) 36–68.
- [279] E.N. Lorenz, Atmospheric predictability as revealed by naturally occurring analogues, *J. Atmos. Sci.* 26 (4) (1969) 636–646.
- [280] E.N. Lorenz, Three approaches to atmospheric predictability, *Bull. Amer. Meteor. Soc.* 50 (3454) (1969) 349.
- [281] A. Wolf, J.B. Swift, H.L. Swinney, J.A. Vastano, Determining lyapunov exponents from a time series, *Physica D* 16 (3) (1985) 285–317.
- [282] M. Cencini, F. Cecconi, A. Vulpiani, Chaos: From simple models to complex systems, 17 (2010).

- [283] G. Gaulier, S. Zignago, BACI: International Trade Database at the Product-Level. The 1994–2007 Version. Working Papers 2010-23, CEPII (October 2010). URL <http://www.cepii.fr/CEPII/fr/publications/wp/abstract.asp?NoDoc=2726>.
- [284] G.T. Hart, A.K. Ramani, E.M. Marcotte, How complete are current yeast and human protein–interaction networks? *Genome Biol.* 7 (11) (2006) 1.
- [285] L. Lü, T. Zhou, Link prediction in complex networks: A survey, *Physica A* 390 (6) (2011) 1150–1170.
- [286] B. Barzel, A.-L. Barabási, Network link prediction by global silencing of indirect correlations, *Nature Biotechnol.* 31 (8) (2013) 720–725.
- [287] J. Menche, A. Sharma, M. Kitsak, S.D. Ghiassian, M. Vidal, J. Loscalzo, A.-L. Barabási, Uncovering disease–disease relationships through the incomplete interactome, *Science* 347 (6224) (2015) 1257601.
- [288] D. Liben–Nowell, J. Kleinberg, The link–prediction problem for social networks, *J. Amer. Soc. Inf. Sci. Technol.* 58 (7) (2007) 1019–1031.
- [289] J. Scott, *Social Network Analysis*, Sage, 2012, pp. 1–248.
- [290] D. Li, Y. Zhang, Z. Xu, D. Chu, S. Li, Exploiting information diffusion feature for link prediction in sina weibo, *Sci. Rep.* 6 (2016) 20058.
- [291] J. Liu, G. Deng, Link prediction in a user–object network based on time–weighted resource allocation, *Physica A* 388 (17) (2009) 3643–3650.
- [292] T. Tylenda, R. Angelova, S. Bedathur, Towards time–aware link prediction in evolving social networks, in: *Proceedings of the 3rd Workshop on Social Network Mining and Analysis*, ACM, 2009, p. 9.
- [293] Y. Dhote, N. Mishra, S. Sharma, Survey and analysis of temporal link prediction in online social networks, in: *2013 International Conference on Advances in Computing, Communications and Informatics*, IEEE, 2013, pp. 1178–1183.
- [294] L. Munasinghe, R. Ichise, Time aware index for link prediction in social networks, in: *International Conference on Data Warehousing and Knowledge Discovery*, Springer, 2011, pp. 342–353.
- [295] L. Lü, L. Pan, T. Zhou, Y.-C. Zhang, H.E. Stanley, Toward link predictability of complex networks, *Proc. Natl. Acad. Sci.* 112 (8) (2015) 2325–2330.
- [296] T. Wang, X.-S. He, M.-Y. Zhou, Z.-Q. Fu, Link prediction in evolving networks based on the popularity of nodes, 2016. ArXiv preprint [arXiv:1610.05347](https://arxiv.org/abs/1610.05347).
- [297] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, W. Van den Broeck, What’s in a crowd? Analysis of face-to-face behavioral networks, *J. Theoret. Biol.* 271 (1) (2011) 166–180.
- [298] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott, Impact of human mobility on opportunistic forwarding algorithms, *IEEE Trans. Mob. Comput.* 6 (6) (2007) 606–620.
- [299] T. Opsahl, P. Panzarasa, Clustering in weighted networks, *Social Networks* 31 (2) (2009) 155–163.
- [300] B. Moradabadi, M.R. Meybodi, Link prediction based on temporal similarity metrics using continuous action set learning automata, *Physica A* 460 (2016) 361–373.
- [301] L.A. Adamic, E. Adar, Friends and neighbors on the web, *Social Networks* 25 (3) (2003) 211–230.
- [302] W. Liu, L. Lü, Link prediction based on local random walk, *Europhys. Lett.* 89 (5) (2010) 58007.
- [303] K. Zweig, Good versus optimal: Why network analytic methods need more systematic evaluation, *Open Comput. Sci.* 1 (1) (2011) 137–153.
- [304] J.M. Hofman, A. Sharma, D.J. Watts, Prediction and explanation in social systems, *Science* 355 (6324) (2017) 486–488.
- [305] V. Subrahmanian, S. Kumar, Predicting human behavior: The next frontiers, *Science* 355 (6324) (2017) 489.
- [306] M. Zanin, A. Papo, P.A. Sousa, E. Menasalvas, A. Nicchi, E. Kubik, S. Boccaletti, Combining complex networks and data mining: Why and how, *Phys. Rep.* 635 (2016) 1–44.
- [307] A. Sidiropoulos, Y. Manolopoulos, Generalized comparison of graph–based ranking algorithms for publications and authors, *J. Syst. Softw.* 79 (12) (2006) 1679–1700.
- [308] M. Dunaiski, W. Visser, J. Geldenhuys, Evaluating paper and author ranking algorithms using impact and contribution awards, *J. Informetrics* 10 (2) (2016) 392–407.
- [309] D. Fiala, Time–aware pageRank for bibliographic networks, *J. Informetrics* 6 (3) (2012) 370–388.
- [310] C. Smith–Clarke, A. Mashhadi, L. Capra, Poverty on the cheap: Estimating poverty maps using aggregated mobile communication networks, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2014, pp. 511–520.
- [311] F. Radicchi, A. Weissman, J. Bollen, Quantifying perceived impact of scientific publications, 2016. ArXiv preprint [arXiv:1612.03962](https://arxiv.org/abs/1612.03962).
- [312] P. Van Mieghem, R. Van de Bovenkamp, Non–markovian infection spread dramatically alters the susceptible–infected–susceptible epidemic threshold in networks, *Phys. Rev. Lett.* 110 (10) (2013) 108701.
- [313] A. Koher, H.H. Lentz, P. Hövel, I.M. Sokolov, Infections on temporal networks matrix–based approach, *PONE* 11 (4) (2016) e0151209.
- [314] A. Schubert, T. Braun, Relative indicators and relational charts for comparative assessment of publication output and citation impact, *Scientometrics* 9 (5–6) (1986) 281–291.
- [315] P. Vinkler, Evaluation of some methods for the relative assessment of scientific publications, *Scientometrics* 10 (3–4) (1986) 157–177.
- [316] Z. Zhang, Y. Cheng, N.C. Liu, Comparison of the effect of mean–based method and z–score for field normalization of citations at the level of web of science subject categories, *Scientometrics* 101 (3) (2014) 1679–1693.
- [317] I. Scholtes, R. Pfitzner, F. Schweitzer, The social dimension of information ranking: a discussion of research challenges and approaches, in: *Socioinformatics: The Social Impact of Interactions Between Humans and IT*, Springer, 2014, pp. 45–61.
- [318] E. Pariser, *The filter bubble: what the internet is hiding from you*, Penguin UK, 2011, pp. 1–304.
- [319] M. Del Vicario, G. Vivaldo, A. Bessi, F. Zollo, A. Scala, G. Caldarelli, W. Quattrociocchi, Echo chambers: Emotional contagion and group polarization on Facebook, *Sci. Rep.* 6 (2016) 37825.
- [320] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H.E. Stanley, W. Quattrociocchi, The spreading of misinformation online, *Proc. Natl. Acad. Sci. USA* 113 (3) (2016) 554–559.
- [321] S. Piramuthu, G. Kapoor, W. Zhou, S. Mauw, Input online review data and related bias in recommender systems, *Decis. Support Syst.* 53 (3) (2012) 418–424.
- [322] O. Ruusuvirta, M. Rosema, Do online vote selectors influence electoral participation and the direction of the vote, in: *ECPR General Conference*, September, 2009 13–12.
- [323] B.K. Peoples, S.R. Midway, D. Sackett, A. Lynch, P.B. Cooney, Twitter predicts citation rates of ecological research, *PLoS One* 11 (11) (2016) e0166570.
- [324] S. Fortunato, A. Flammini, F. Menczer, Scale–free network growth by ranking, *Phys. Rev. Lett.* 96 (21) (2006) 218701.
- [325] M.D. König, C.J. Tessone, Network evolution based on centrality, *Phys. Rev. E* 84 (5) (2011) 056108.
- [326] M.D. König, C.J. Tessone, Y. Zenou, Nestedness in networks: A theoretical model and some applications, *Theoretical Economics* 9 (3) (2014) 695–752.
- [327] I. Sendiña–Nadal, M.M. Danziger, Z. Wang, S. Havlin, S. Boccaletti, Assortativity and leadership emerge from anti–preferential attachment in heterogeneous networks, *Sci. Rep.* 6 (2016) 21297.
- [328] M. Medo, M.S. Mariani, A. Zeng, Y.-C. Zhang, Identification and modeling of discoverers in online social systems, *Sci. Rep.* 6 (2016) 34218.
- [329] M. Tomasello, G. Vaccario, F. Schweitzer, Data–driven modeling of collaboration networks: A cross–domain analysis, 2017. [arXiv:1704.01342](https://arxiv.org/abs/1704.01342).
- [330] G. Lohmann, D.S. Margulies, A. Horstmann, B. Pleger, J. Lepsien, D. Goldhahn, H. Schloegl, M. Stumvoll, A. Villringer, R. Turner, Eigenvector centrality mapping for analyzing connectivity patterns in fMRI data of the human brain, *PLoS One* 5 (4) (2010) e10232.