

NONLOCAL PAGERANK

STEFANO CIPOLLA¹, FABIO DURASTANTE^{2,*} AND FRANCESCO TUDISCO³

Abstract. In this work we introduce and study a nonlocal version of the PageRank. In our approach, the random walker explores the graph using longer excursions than just moving between neighboring nodes. As a result, the corresponding ranking of the nodes, which takes into account a *long-range interaction* between them, does not exhibit concentration phenomena typical of spectral rankings which take into account just local interactions. We show that the predictive value of the rankings obtained using our proposals is considerably improved on different real world problems.

Mathematics Subject Classification. 05C82, 68R10, 94C15, 60J20.

Received January 29, 2020. Accepted October 5, 2020.

1. INTRODUCTION

Identifying and quantifying important components in a dataset or a complex system modeled by a network, using only the topological structure of nodes and edges, is a very important issue in exploratory data analysis.

Various specific tasks have been designed around this quite general problem, at a global scale – with the aim of providing insightful summary statistics such as clustering coefficients, robustness or total communicability [7, 18, 33] – at an intermediate (or meso) scale – by identifying structures such as communities, anti-communities or core-periphery [30, 33, 53, 60] – and at a local level – where we aim at quantifying various node or edge properties such as triadic closure or edge communicability [18, 25]. Here we focus on the so called *centrality* problem, where we aim at assigning an importance score to each node in the network in order to discover the most relevant nodes. While this task aims at unveiling network features that take place at a very local scale (nodes), it is nowadays apparent that complex networks feature an intrinsic higher-order organization [5] and that centrality scores should exploit the structure of the network as a whole in order to unveil insightful network properties that are otherwise overlooked. To this end, in this work we propose a simple generalization of the renowned PageRank centrality which forces the global network structure into this local scale centrality model.

PageRank had its fortune due to its employment in early versions of the Google search engine [48]. The main idea of this centrality model is a mutually reinforcing definition of importance: the importance of a node is influenced by the importances of the nodes it connects to. Equivalently, PageRank centrality can be interpreted as the average amount of time that a random walker spends on each node as the length of the

Keywords and phrases. Complex network, nonlocal dynamics, Markov chain, Perron–Frobenius.

¹ Department of Mathematics “Tullio Levi-Civita”, University of Padua, Padova 35121, Italy.

² Consiglio Nazionale delle Ricerche – Istituto per le Applicazioni del Calcolo “M. Picone”, Naples, Italy.

³ GSSI Gran Sasso Science Institute, L’Aquila 67100, Italy.

*Corresponding author: f.durastante@na.iac.cnr.it

walks tend to infinity. While this recursive definition clearly involves the global structure of connections in the network, at each time step the random walker moves from one node to another taking into account only the direct neighbors of that node. Thus, while this definition implies that each node score is influenced by all the other node importances, the classical PageRank centrality often results into a localized measure, as it typically happens for typical eigenvector-based centrality scores [43].

Different strategies have been considered in recent literature to overcome this issue. On the one end, higher-order adjacency tensors have been employed to model higher-order neighborhoods made by hyperedges containing three or more nodes. This approach is typically characterized by the use of hypergraphs or simplicial complexes [3, 4, 11, 17]. On the other hand, non Markovian stochastic processes with memory have been used to model random walks that take into account longer paths of connections [2, 6, 16, 31].

Following this second line of research, in this work we propose a nonlocal version of the classical PageRank model based on the usage of the Lévy random walk, *i.e.* the usage of an *anomalous* nonlocal diffusion that employs a one-parameter family of decaying transition probabilities. In the case of undirected networks, this type of random walk was considered for example in [51]. The main idea of our approach is to move from the original exploration strategy of the network exploited in the PageRank, where the random walker moves between neighboring nodes with uniform probability, to a strategy that permits longer excursions between the nodes of the network, *i.e.* it allows us to move from a node i to any other node that is connected to i through a path of any length. These types of longer length interactions enhance the navigability of the network and thus allow for a faster exploration, as observed in other related contexts, including fractal small-worlds networks [52], lattices [39] and general multi-hopper models on digraphs [26, 63]. In particular, very related to our approach is the concept of path Laplacians introduced by Estrada *et al.* [24] and further analyzed in [27, 28]; another related approach is the generalized PageRank [40], which is based a matrix function of the transition probability matrix. The proposed nonlocal PageRank is also based on a function of the transition matrix, however, unlike generalized PageRank, it uses a entrywise function.

The reminder of the paper is structured as follows: We start by fixing the notation and by recalling the standard PageRank model in Section 2. Then, in Section 3 we introduce the proposed nonlocal PageRank model, which is based on the choice of a distance function between the nodes and a one-parameter family of Lévy-type random walks on the graph. In Section 4 we perform an asymptotic analysis on the selection of the parameter that defines the transition probabilities and we show that the classical PageRank follows as a special case of the new model for large values of the parameter and when the chosen distance is the standard shortest-path distance. However, different values of the parameter allow us to obtain models that are more stable and less localized, as discussed in Section 4.1. These properties help improving a number of network mining tasks: In Section 4.2 we show how different values of the parameter affect the behavior of the model in the context of link prediction. Whereas in Section 5 we discuss how different choices of the distance function can affect the model. In particular, choosing suitable – and possibly problem related – distance functions gives us an additional flexibility which allows us to improve the quality of the resulting centrality assignment. We highlight this by considering the London underground train test problem, where we design a “metro distance” that takes into account the multilayer structure of the network (given by the several underground train lines) and that compares favorably with other PageRank-like centralities.

Data and software

All data and software used in this work are available online. For convenience, we list below all the network data used in the different sections of the paper. We detail additional information on the various datasets when appropriate in the text.

- **USAir97**, this is a directed network of air traffic in the U.S.A. available from Pajek repository [54].
- **EUair**, this is a thirty-seven layer network, each one corresponding to a different airline operating in Europe [10]. We use the aggregate graph, consisting of the union of all the layers.
- **Barcelona**, this is the directed transportation network for the city of Barcelona (Spain) from the Research Core Team collection [57].

- `adjnoun`, this undirected network represents common occurrences of adjectives and nouns in the “David Copperfield” novel by C. Dickens [46].
- `zachary`, this is a (small) social network of a university karate club [46].
- `gre_115`, this directed network is available from the Harwell-Boeing collection [23].
- `cake9`, this is a network obtained from a DNA electrophoresis model [61] in which a polymer is modeled as a chain of “monomers” connected by bonds.
- `delaunay_n10/delaunay_n12`, are networks obtained from the adjacency matrices relative to two Delaunay triangulations of random points in the unit square [36].
- `3elt`, is an example graph from the AG-Monien Graph Collection by Ralf Diekmann and Robert Preis (see [19]).
- `tube` is the undirected multilayer network of London’s underground trains connections.

The `tube` network was created by us starting from the dataset developed in [20] and it is made of 13 layers: one layer for Docklands Light Railway (DLR) trains, one layer for overground trains and eleven underground train layers, one for each line. The dataset includes also geographical coordinates of the nodes and passengers usage statistics across several years (2008–2017), obtained from ORR London Datastore [47]. Both this dataset and the software we developed for the experiments shown in the paper are available at <https://github.com/Cirdans-Home/NonLocalPageRank>.

2. THE PAGERANK ALGORITHM

A *digraph*, or directed graph, $\Gamma = (V, E)$ is defined by a set of n nodes $V = \{v_1, \dots, v_n\} \equiv \{1, \dots, n\}$, and a set of ordered edges $E = \{(i, j) : i, j \in V\} \subseteq V \times V$ representing the connections between the nodes. A *walk* of length k in Γ is a list of nodes i_1, \dots, i_k, i_{k+1} such that $(i_j, i_{j+1}) \in E, \forall j = 1, \dots, k$. If the first and the last edge coincides then the walk is called a closed walk. If no repeated nodes appear in the sequence then the walk is called a *path*, while a path in which only the first and last node coincide is called a *cycle*. We consider here only *loop-less* graphs, *i.e.* edges of the form $(i, i) \in E$ are not allowed. We say that a digraph is *strongly connected* if there exists a path between every pair of nodes.

Every graph Γ can be represented as a binary adjacency matrix $A = (a_{i,j})$ with

$$a_{i,j} = \begin{cases} 1 & (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

As we allow directed graphs, such matrix will not be symmetric, in general.

Let $\mathbf{1}$ be the vector of all ones of size n , and let A be the adjacency matrix of the digraph $\Gamma = (V, E)$ with $|V| = n$. We define the diagonal matrix D of the out-degrees of Γ as

$$D_{\text{out}} = \text{diag}(A\mathbf{1}) = \text{diag}(d_1^{\text{out}}, \dots, d_n^{\text{out}}), \quad (2.1)$$

in which each diagonal entry d_i^{out} represents the number of outgoing edges from the node i .

We can now build a random walk on Γ by considering the following transition matrix $P = (p_{i,j})$

$$p_{i,j} = \begin{cases} 1/d_i^{\text{out}} & (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (2.2)$$

Note that P is tightly related to the adjacency matrix and the diagonal matrix of the out degrees. In fact, a compact form for (2.2) reads $P = D_{\text{out}}^{-1}A$, where the matrix D_{out}^{-1} is defined by setting the inverse of zero diagonal entries to zero by convention.

A random walker that obeys the transition matrix P has equal chance of moving from a node to any of its out-neighbors. Note that by following this transition we could end in a *cul-de-sac*, represented by a node i with




FIGURE 1. Pictorial representation of the PageRank vector \mathbf{s} for the Zachary's karate club digraph [33], computed with $c = 0.85$.

$d_i^{\text{out}} = 0$. To avoid this circumstance we modify P in (2.2) to permit the walker to teleport to any other location in the graph with some probability, *i.e.* we define the *PageRank* transition matrix

$$G = c\tilde{P} + \frac{1-c}{n}\mathbf{1}\mathbf{1}^T, \quad c \in (0, 1], \quad (2.3)$$

where \tilde{P} is the matrix P in which each zero row has been replaced by the uniform vector $\mathbf{1}^T/n$. The matrix G is a positive, row stochastic matrix. Thus, by the Perron–Frobenius Theorem, it admits a unique positive and dominant left eigenvector \mathbf{s} corresponding to the eigenvalue 1 (see *e.g.* [8]). In other words, the Markov chain with transition matrix G has a unique and positive stationary distribution

$$\mathbf{s}^T = \mathbf{s}^T G, \quad \mathbf{s}^T \mathbf{1} = 1, \quad s_i > 0, \quad \forall i = 1, \dots, n.$$

Such \mathbf{s} is called the *PageRank* vector of Γ , and its i th entry provides a measure of the importance of node i in the digraph Γ . Figure 1 illustrates an example PageRank vector on the Zachary's karate club network.

From the modeling point of view, the teleportation factor c included in (2.3) can be interpreted in terms of a “surfer” navigating the digraph Γ of the web hyperlinks. When moving from a page to another, the surfer may decide to follow one of the hyperlinks that are listed in the web page, choosing among them with uniform probability. The surfer makes this choice with probability c . Whilst, with probability $1 - c$, they may decide to “teleport” onto a new web page, in principle not connected with the current one, choosing again uniformly at random. This is an example of a nonlocal behavior in which it is possible to end up being in nodes that are far away from the original starting node. However, we have only partial control on this longer jumps. In fact, we are only allowed to either tune the parameter c or to introduce a “personalized version of the PageRank” where the surfer chooses to teleport onto a new page j with a probability v_j that depends on the destination page j , rather than choosing among all the possible web pages uniformly at random. From a mathematical viewpoint, this second choice means we replace the PageRank matrix (2.3) with $c\tilde{P} + (1-c)\mathbf{1}\mathbf{v}^T$, being $\mathbf{v} = (v_1, \dots, v_n)$.

In the next Section 3 we introduce a modification of this model that allows us to include one additional level of nonlocality to the original PageRank centrality.

3. THE NONLOCAL PAGERANK MODEL

Let us suppose that we are at a node i in the digraph Γ . Following from the discussion of Section 2 we know that we have now two possibilities: going to a node that is connected to the present one, or to teleport away in the graph, possibly giving preference to certain nodes over others. However, this preference does not depend on the current node i and, thus, does not take into account the fact that one is typically more inclined to teleport to some page that is somewhat related to i . To overcome this issue, we define here a process where the probability to move to a node j , that lies on a walk that contains i , is large when j is close to i and decreases the further away we move from i .

To this end, we let $\delta : V \times V \rightarrow \mathbb{R}_+$ be a distance function on Γ and define the transition probability matrix $P_\alpha = (p_{i,j})$ as

$$p_{i,j} = \begin{cases} \frac{f_\alpha(\delta(i,j))}{\sum_{k \neq i} f_\alpha(\delta(i,k))} & \delta(i,j) < \infty \\ 0 & \text{otherwise,} \end{cases} \quad (3.1)$$

where $f_\alpha(x) : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a family of nonnegative and nonincreasing functions parametrized by $\alpha \in \mathbb{R}_+$. The family f_α will be used to regulate the trade-off between the importance of the nodes at a short-range and the ones that are further away. Note that the distance δ does not need to define a metric on Γ and, for example, does not need to be symmetric, *i.e.* we allow $\delta(i,j) \neq \delta(j,i)$. In particular, observe that this is the situation if we choose δ to be the shortest path distance, as we will discuss in more details in the next section. Also note that, if Δ denotes the distance matrix such that $\Delta_{ij} = \delta(i,j)$, then it holds

$$P_\alpha = \text{diag}(f_\alpha(\Delta)\mathbf{1})^{-1}f_\alpha(\Delta)$$

where f_α is applied entrywise and the inverse of $+\infty$ is set to zero by convention.

As for the standard PageRank case, a random walker obeying the transition rule (3.1) could end up in a node i such that $\delta(i,j) = +\infty$ for all other $j \in V$ (*i.e.* the i -th row of P_α is all zero) and get stuck there. Thus we modify the state transitions to include a teleportation, *i.e.* we let

$$G_\alpha = c\tilde{P}_\alpha + \frac{1-c}{n}\mathbf{1}\mathbf{1}^T, \quad c \in (0, 1], \quad (3.2)$$

where \tilde{P}_α is the matrix P_α in which each zero row has been replaced by $\mathbf{1}^T/n$. The matrix G_α is again a positive row stochastic matrix and, by the Perron–Frobenius Theorem, there exists a unique \mathbf{s}_α with positive entries, such that $\mathbf{s}_\alpha^T = \mathbf{s}_\alpha^T G_\alpha$, $\mathbf{s}_\alpha^T \mathbf{1} = 1$. We call \mathbf{s}_α the *nonlocal PageRank* vector.

Clearly, the nonlocal PageRank depends on the choice of the distance δ and on the choice of the family of decaying functions f_α . In the next section we show that, when δ is the shortest path distance and f_α is suitably defined, the nonlocal PageRank interpolates the standard PageRank and the parameter α can be used to tune the “amount of nonlocality” we want to take into account.

4. SHORTEST PATH DISTANCE AND CONVERGENCE TO THE PAGERANK

As we will see in Section 5, different choices of the distance δ can be used, depending on the application. In fact, this flexibility is one of the key advantages of our proposed framework, which allows us to model nonstandard node-node interactions and thus to capture node properties that are overlooked otherwise. On the other hand, the arguably most natural choice for distance δ is the *shortest path* distance, whose definition we recall below

Definition 4.1. Given a digraph $\Gamma = (V, E)$ the *shortest-path* distance $\delta_\Gamma(i, j)$ between any two nodes $i, j \in V$ is the smallest length of any path from i to j . If there exists no such a path, we let $\delta_\Gamma(i, j) = +\infty$.

This choice of the distance function allows us to retrieve the classical PageRank as a special case of our nonlocal PageRank model, provided the family f_α satisfies the following nonrestrictive decaying condition.

Definition 4.2. We say that f_α is a smoothing family of functions for the distance δ if $f_\alpha(x) \rightarrow 0$ and $f_\alpha(x)/f_\alpha(1) \rightarrow 0$ as $\alpha \rightarrow \infty$, for all x in the set

$$\Omega_\delta = \{x : x \geq \delta(i, j), \text{ for all } (i, j) \notin E\} \subseteq \mathbb{R}_+.$$

When $\delta = \delta_\Gamma$ we have $\Omega_{\delta_\Gamma} = [2, \infty]$, and we can easily produce examples of smoothing families by considering any nonincreasing function f such that $f(2) < 1$ and then defining $f_\alpha(x) = f(x)^\alpha$. For example, we can choose, as in [51], the functions

$$f_\alpha(x) = \frac{1}{x^\alpha}.$$

In this case, the resulting transition matrix P_α describes a Lévy random walk on Γ . Another example is the exponential function

$$f_\alpha(x) = e^{-\alpha x}.$$

For any smoothing family of functions we have

Lemma 4.3. Let f_α be a smoothing family of nondecreasing functions for the distance $\delta = \delta_\Gamma$. Then $P_\alpha \rightarrow P$ entrywise as $\alpha \rightarrow \infty$, where P is the PageRank transition matrix (2.3).

Proof. Note that $\sum_k f_\alpha(\delta_\Gamma(i, k)) = d_i^{\text{out}} f_\alpha(1) + \sum_{k: \delta_\Gamma(i, k) \geq 2} f_\alpha(\delta_\Gamma(i, k))$. Thus, simple algebraic manipulations show that $P_\alpha = P + Y_\alpha$, with

$$(Y_\alpha)_{ij} = \begin{cases} \frac{1}{d_i^{\text{out}} d_i^{\text{out}} f_\alpha(1) + \sum_{k: \delta_\Gamma(i, k) \geq 2} f_\alpha(\delta_\Gamma(i, k))} \sum_{k: \delta_\Gamma(i, k) \geq 2} f_\alpha(\delta_\Gamma(i, k)) & (i, j) \in E \\ \frac{f_\alpha(\delta_\Gamma(i, j))}{\sum_{k \neq i} f_\alpha(\delta_\Gamma(i, k))} & 2 \leq \delta_\Gamma(i, j) < \infty \\ 0 & \delta_\Gamma(i, j) = \infty. \end{cases}$$

As $f_\alpha(x) \rightarrow 0$ for $\alpha \rightarrow \infty$ and $x \geq 2$, we have $(Y_\alpha)_{ij} \rightarrow 0$, and the proof is complete. \square

Note that, when $f_\alpha(x) = f(x)^\alpha$ is defined *via* a power-law (and $f(x)$ is bounded) we obviously have $f_\alpha(x) \rightarrow 1$ as $\alpha \rightarrow 0$, implying that P_α converges to the uniform transition matrix where i transitions to every node j at finite distance from i with equal probability. This observation, combined with the lemma above, shows that choosing values of α reasonably far from 0 and ∞ allows us to define a model that interpolates between a standard Markov chain on the graph and a purely uniformly random model. This is also shown by Figure 2, where we plot the Kendall's τ correlation coefficient between the nonlocal PageRank centrality score – obtained with different values of α – and the standard PageRank centrality on the datasets `netscience`, `astro-ph` and `hep-th`. In particular, the figure highlights the result in Lemma 4.3 showing that as $\alpha \rightarrow +\infty$ the correlation between the nonlocal PageRank and the PageRank converges to 1, thus showing the convergence to the standard PageRank algorithm.

Clearly, while the two extreme choices $\alpha \rightarrow 0$ and $\alpha \rightarrow \infty$ do not require to compute δ_Γ , one of the drawbacks of using the nonlocal PageRank transition matrix P_α for other choices of α is the computation of the shortest-path distance matrix from Definition 4.1. The algorithm of choice for computing all the distances between the nodes is either the Floyd–Warshall algorithm [32, 62] or the Johnson's algorithm [37]. The first one works on weighted digraphs with positive or negative edge weights, and no negative cycles, and it has a running time of $O(n^3)$. The second one allows only for few negative edge weights and has a running time of $O(n^2 \log(n) + nm)$ where $m = |E|$. As it is clear from the running times estimates, the Floyd–Warshall algorithm is best suited for dense networks ($m \approx n$), while the Johnson's algorithm should be preferred in the case of sparse networks ($m \ll n$). The storage of $O(n^2)$ machine numbers has to be expected.

In the next section we show how different values of α can improve the stability and reduce the localization phenomenon of the standard PageRank.




FIGURE 2. The figure shows Kendall's τ correlation coefficient of PageRank *vs.* Nonlocal PageRank for different $f_\alpha(x)$ and different values of α for the shortest path distance. Problems sizes (left to right) $n = 1589$, $n = 16\,706$ and $n = 8361$.

4.1. Stability and nonlocality

The classical PageRank centrality measure shares the same principal flaw of all the eigenvector centralities. The leading eigenvector of the associated transition matrix can suffer the localization phenomenon [35, 43, 56], *i.e.* most of the measure weight tends to concentrate around few most important nodes while giving to all the other nodes a small and numerically identical value, and thus ranking. This phenomenon has also the secondary effect of producing high variations in the value of the entries of the PageRank vector when the network topology faces a small perturbation, for example few edges are added or removed from the graph [42, 49, 50].

In a linear algebra terminology, this type of graph modification is equivalent to a structured perturbation of the transition matrix and is related to the concept of *condition number of a Markov chain* [45].

For a given α , a given distance δ and a given smoothing function f_α , we define a condition number of the nonlocal PageRank as a coefficient κ_α that quantifies the relative change in the nonlocal PageRank vector when changes occur in the graph. More precisely, suppose that $\bar{\Gamma} = (V, \bar{E})$ is an edge-perturbation of $\Gamma = (V, E)$ and let \bar{G}_α be the nonlocal transition matrix (3.2) of $\bar{\Gamma}$. Then κ_α is a condition number of the nonlocal PageRank if

$$\frac{\|\mathbf{s}_\alpha - \bar{\mathbf{s}}_\alpha\|}{\|\mathbf{s}_\alpha\|} \leq \kappa_\alpha \frac{\|G_\alpha - \bar{G}_\alpha\|}{\|G_\alpha\|} \quad (4.1)$$

and this bound holds for all possible perturbations of Γ . In (4.1) \mathbf{s}_α and $\bar{\mathbf{s}}_\alpha$ are the nonlocal PageRank of G_α and \bar{G}_α respectively. Note that property (4.1) does not define a unique condition number κ_α and in fact several comparisons between different possible choices of κ_α have been studied [14, 38, 45]. Here we use a result of Seneta [55] to characterize the stability of the nonlocal PageRank in terms of its norm-1 condition number.

To this end, we use the *ergodicity coefficient* of the matrix sequence $\{G_\alpha\}_\alpha$. For a fixed α , this coefficient is defined as

$$\tau_1(G_\alpha) = \sup_{\substack{\|\delta\|_1=1 \\ \delta^T \mathbf{1}=0}} \|\delta^T G_\alpha\|_1 = \frac{1}{2} \max_j \|G_\alpha^T (I - \mathbf{e}_j \mathbf{1}^T)\|_1 = \frac{1}{2} \max_{i,j} \sum_k |(G_\alpha)_{ik} - (G_\alpha)_{jk}| \quad (4.2)$$

and the following quantity

$$\text{cond}_1(\mathbf{s}_\alpha) = \frac{1}{1 - \tau_1(G_\alpha)}$$

is a condition number for the nonlocal PageRank, in the sense that inequality (4.1) holds for $\kappa_\alpha = \text{cond}_1(\mathbf{s}_\alpha)$ and $\|\cdot\| = \|\cdot\|_1$, see *e.g.* [55]. Therefore, if we define the remainder $R = G_\alpha - \bar{G}_\alpha$ and we express G_α as the

structured perturbation $G_\alpha = \overline{G}_\alpha + R$, with R such that $R\mathbf{1} = 0$, we can recover the following perturbation bound on the associated PageRank vectors

$$\|\mathbf{s}_\alpha - \overline{\mathbf{s}}_\alpha\|_1 \leq \|R\|_1(1 - \tau_1(G_\alpha))^{-1}. \quad (4.3)$$

In particular, as $\|\mathbf{s}_\alpha^T\|_1 = \|G_\alpha\|_1 = 1$ and $|\lambda(G_\alpha)| \leq \tau_1(G_\alpha)$, for any eigenvalue $\lambda(G_\alpha)$ of G_α such that $\lambda(G_\alpha) \neq 1$ (see *e.g.* [59]), the relation (4.3) gives us a bound on the condition number of \mathbf{s}_α for relative changes in G_α in the norm sense, *i.e.*

$$\text{cond}_1(\mathbf{s}_\alpha) \leq \frac{1}{1 - |\lambda(G_\alpha)|}, \quad \forall \lambda(G_\alpha) \neq 1. \quad (4.4)$$

It is then interesting to evaluate the behavior of such quantity with respect to the parameter α , and indeed it is possible to prove the following result.

Theorem 4.4. *Let f_α be a family of decreasing smoothing functions for the distance $\delta = \delta_\Gamma$. Then, there exists α_0 such that $\tau_1(G_\alpha) \leq \tau_1(G)$ for all $\alpha \geq \alpha_0$.*

Proof. Since G_α and G are rank-one modifications of \tilde{P}_α and \tilde{P} , using the second part of (4.2), the thesis holds if and only if $\tau_1(\tilde{P}_\alpha) \leq \tau_1(\tilde{P})$. Moreover, since P and P_α have the same zero rows and \tilde{P} and \tilde{P}_α coincide on those rows, using the third part in (4.2), we see that it is enough to show that $\tau_1(P_\alpha) \leq \tau_1(P)$, for all $\alpha \geq \alpha_0$.

We will obtain the thesis modifying one row at time of the matrix P to obtain the matrix P_α . Recall, moreover, that by Lemma 4.3, we have $\lim_{\alpha \rightarrow \infty} P_\alpha = P$.

If $\delta_\Gamma(i, j) = 1$ for all i, j then $P_\alpha = P$ for all α and thus the thesis is straightforward.

Assume that there exist i and j such that $\delta_\Gamma(i, j) > 1$. Define the matrix \hat{P}_α obtained modifying the i -th row of P into the i -th row of P_α , *i.e.* let

$$(\hat{P}_\alpha)_{i,:} := \left[\frac{f_\alpha(\delta_\Gamma(i, 1))}{\sum_{\ell \neq i} f_\alpha(\delta_\Gamma(i, \ell))} \cdots 0 \cdots \frac{f_\alpha(\delta_\Gamma(i, n))}{\sum_{\ell \neq i} f_\alpha(\delta_\Gamma(i, \ell))} \right]$$

where the zero entry is the diagonal $(\hat{P}_\alpha)_{i,i}$, and let $(\hat{P}_\alpha)_{j,:} := P_{j,:}$ for all $j \neq i$. Observe that if $k \in E_i := \{k \text{ s.t. } (i, k) \in E\}$, it holds

$$(\hat{P}_\alpha)_{i,k} = \frac{1}{d_i + \sum_{\ell: \delta_\Gamma(i, \ell) \geq 2} f_\alpha(\delta_\Gamma(i, \ell)) / f_\alpha(1)} < \frac{1}{d_i} = P_{i,k}, \quad (4.5)$$

where the strict inequality holds from our assumption. More precisely,

$$(\hat{P}_\alpha)_{i,k} = P_{i,k} + D_i^\alpha \quad (4.6)$$

where $D_i^\alpha := 1 / (d_i + \sum_{\ell: \delta_\Gamma(i, \ell) \geq 2} f_\alpha(\delta_\Gamma(i, \ell)) / f_\alpha(1)) - 1 < 0$ for all $\alpha > 0$. Let us define $i_0(\alpha), j_0(\alpha)$ such that

$$\tau(\hat{P}_\alpha) = \frac{1}{2} \sum_k |(\hat{P}_\alpha)_{i_0(\alpha), k} - (\hat{P}_\alpha)_{j_0(\alpha), k}|.$$

Since \hat{P}_α converges, there exists $\bar{\alpha}_0$ s.t. $i_0(\alpha) = i_0$ and $j_0(\alpha) = j_0$ for all $\alpha \geq \bar{\alpha}_0$. Moreover, since \hat{P}_α differs from P only in the i -th row, to prove $\tau_1(\hat{P}_\alpha) \leq \tau_1(P)$, it is enough to consider the case $i_0 = i, j_0 \neq i$.

In the setting we are considering we have $(\hat{P}_\alpha)_{j_0, k} = P_{j_0, k}$ for all k . So, for $i = i_0$, let us define the following sets

$$\begin{aligned} \mathcal{I} &:= \{k : (\hat{P}_\alpha)_{i,k} \geq P_{j_0, k}\}, & V \setminus \mathcal{I} &:= \{k : (\hat{P}_\alpha)_{i,k} < P_{j_0, k}\}, \\ \mathcal{J} &:= \{k : P_{i,k} \geq P_{j_0, k}\}, & V \setminus \mathcal{J} &:= \{k : P_{i,k} < P_{j_0, k}\}. \end{aligned}$$

We then observe that

- since both P and \widehat{P}_α are stochastic, we have $\sum_{k \in V} \{(\widehat{P}_\alpha)_{i,k} - P_{j_0,k}\} = 0$, and thus the row sum on the two complementary sets \mathcal{I} and $V \setminus \mathcal{I}$ yields $\sum_{k \in \mathcal{I}} \{(\widehat{P}_\alpha)_{i,k} - P_{j_0,k}\} = \sum_{k \in V \setminus \mathcal{I}} \{P_{j_0,k} - (\widehat{P}_\alpha)_{i,k}\}$;
- since by (4.5) when $k \in E_i$ we have $P_{i,k} > (\widehat{P}_\alpha)_{i,k}$, then $\mathcal{I} \cap E_i \subseteq \mathcal{J} \cap E_i$.

From the above observations, we find

$$\begin{aligned}
\tau_1(\widehat{P}_\alpha) &= \sum_{k \in \mathcal{I}} ((\widehat{P}_\alpha)_{i,k} - P_{j_0,k}) \\
&= \sum_{k \in \mathcal{I} \cap E_i} ((\widehat{P}_\alpha)_{i,k} - P_{j_0,k}) + \sum_{k \in \mathcal{I} \cap E_i^c} ((\widehat{P}_\alpha)_{i,k} - P_{j_0,k}) \\
&= \sum_{k \in \mathcal{I} \cap E_i} (P_{i,k} - P_{j_0,k}) + \sum_{k \in \mathcal{I} \cap E_i} D_i^\alpha + \sum_{k \in \mathcal{I} \cap E_i^c} ((\widehat{P}_\alpha)_{i,k} - P_{j_0,k}) \\
&\leq \sum_{k \in \mathcal{I} \cap E_i} (P_{i,k} - P_{j_0,k}) + \sum_{k \in \mathcal{I} \cap E_i} D_i^\alpha + \sum_{k \in \mathcal{I} \cap E_i^c} \frac{f_\alpha(\delta_\Gamma(i, k))/f_\alpha(1)}{d_i + \sum_{\ell: \delta_\Gamma(i, \ell) \geq 2} f_\alpha(\delta_\Gamma(i, \ell))/f_\alpha(1)} \\
&\leq \sum_{k \in \mathcal{J} \cap E_i} (P_{i,k} - P_{j_0,k}) + \sum_{k \in \mathcal{I} \cap E_i} D_i^\alpha + \sum_{k \in \mathcal{I} \cap E_i^c} \frac{f_\alpha(\delta_\Gamma(i, k))/f_\alpha(1)}{d_i + \sum_{\ell: \delta_\Gamma(i, \ell) \geq 2} f_\alpha(\delta_\Gamma(i, \ell))/f_\alpha(1)},
\end{aligned} \tag{4.7}$$

where in the first inequality we used (4.6) and the definition of $(\widehat{P}_\alpha)_{i,k}$ and, in the second one, the fact that $\mathcal{I} \cap E_i \subseteq \mathcal{J} \cap E_i$.

Since $\sum_{k \in \mathcal{I} \cap E_i} D_i^\alpha < 0$ for all $\alpha > 0$ and $\sum_{k \in \mathcal{I} \cap E_i^c} \left(\frac{f_\alpha(\delta_\Gamma(i, k))/f_\alpha(1)}{d_i + \sum_{\ell: \delta_\Gamma(i, \ell) \geq 2} f_\alpha(\delta_\Gamma(i, \ell))/f_\alpha(1)} \right) \rightarrow 0$ as $\alpha \rightarrow \infty$ (f_α is a smoothing family), there exists $\alpha_0^{(1)} \geq \bar{\alpha}_0$ s.t.

$$\sum_{k \in \mathcal{I} \cap E_i} D_i^\alpha + \sum_{k \in \mathcal{I} \cap E_i^c} \frac{f_\alpha(\delta_\Gamma(i, k))/f_\alpha(1)}{d_i + \sum_{\ell: \delta_\Gamma(i, \ell) \geq 2} f_\alpha(\delta_\Gamma(i, \ell))/f_\alpha(1)} < 0 \text{ for all } \alpha \geq \alpha_0^{(1)}.$$

Finally, observing that $\sum_{k \in \mathcal{J} \cap E_i} (P_{i,k} - P_{j_0,k}) \leq \frac{1}{2} \sum_k |P_{i,k} - P_{j_0,k}| \leq \frac{1}{2} \max_{s,m} \sum_k |P_{s,k} - P_{m,k}| = \tau_1(P)$, we obtain from (4.7), that $\tau_1(\widehat{P}_\alpha) \leq \tau_1(P)$ for all $\alpha \geq \alpha_0^{(1)}$.

Let us now define $\widehat{P}_\alpha^{(1)} = \widehat{P}_\alpha$. Observe that the above reasoning is still valid using $\widehat{P}_\alpha^{(1)}$ in place of P and considering as \widehat{P}_α the matrix obtained modifying one more row of the new $\widehat{P}_\alpha^{(1)}$ into the same row of P_α (overall modifying exactly two rows of the original P into the corresponding rows of P_α). We call the matrix obtained in this way $\widehat{P}_\alpha^{(2)}$. In this case we have that there exists $\alpha_0^{(2)}$ such that, for all $\alpha \geq \alpha_0^{(2)}$, $\tau_1(\widehat{P}_\alpha^{(2)}) \leq \tau_1(\widehat{P}_\alpha^{(1)})$.

Applying this argument n times we obtain that there exists $\alpha_0^{(n)} =: \alpha_0$ such that, for all $\alpha \geq \alpha_0$, the following chain of inequalities holds

$$\tau_1(P_\alpha) = \tau_1(\widehat{P}_\alpha^{(n)}) \leq \dots \leq \tau_1(\widehat{P}_\alpha^{(1)}) \leq \tau_1(P),$$

concluding the proof. \square

The result above shows that when the nonlocality parameter α is large enough, the corresponding PageRank vector is more stable with respect to structured perturbations of Γ or, equivalently, with respect to perturbations of the adjacency matrix. In fact, if \mathbf{s} and \mathbf{s}_α are the PageRank and the nonlocal PageRank vectors, respectively, the inequality $\tau_1(G_\alpha) \leq \tau_1(G)$ implies $\text{cond}_1(\mathbf{s}_\alpha) \leq \text{cond}_1(\mathbf{s})$. On the other hand, extensive numerical experiments show that this inequality actually holds for all values of α . We illustrate this result in Figure 3 where we compare the ergodicity coefficient $\tau_1(G_\alpha)$ of the nonlocal PageRank (for $\delta = \delta_\Gamma$) with the coefficient $\tau_1(G)$ for

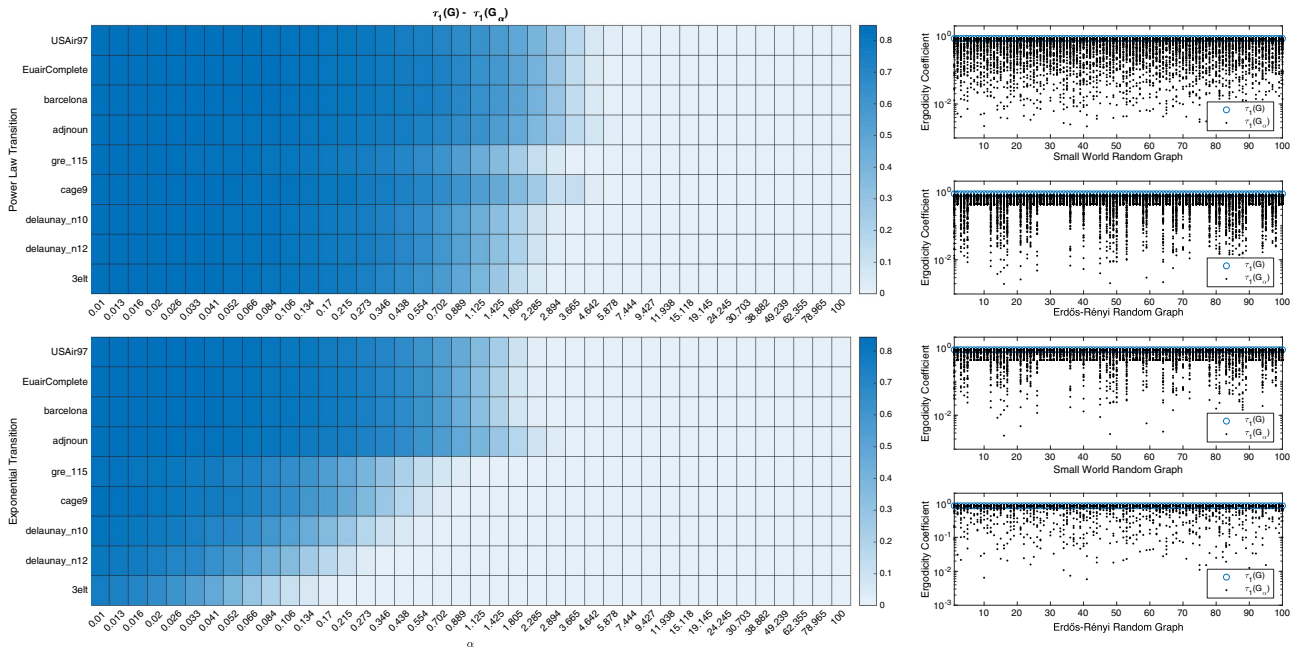


FIGURE 3. Behavior of the ergodicity coefficient τ_1 for both the power law $f_\alpha(x) = x^{-\alpha}$ and the exponential $f_\alpha(x) = \exp(-\alpha x)$ transitions. The heat-maps on the left-hand side represent the difference $\tau_1(G) - \tau_1(G_\alpha)$ for growing values of α on a large set of test networks. Whereas, the panels on the right-hand side report the behavior of the ergodic coefficients for one hundred random graphs of Small World and Erdős–Rényi type. For every graph the ergodic coefficient for the nonlocal PageRank is computed for one hundred different random alphas and, in every case, we observe that $\tau_1(G_\alpha) \leq \tau_1(G)$.

the standard PageRank, for a broad range of values of the parameter α and for both the example smoothing families of functions $f_\alpha(x) = x^{-\alpha}$ and $f_\alpha(x) = e^{-\alpha x}$. This figure shows that the nonlocal PageRank is actually more stable than its local counterpart for all values of α .

Clearly, the norm-1 bound we have obtained provides an indication on the entry-wise variation of the PageRank vector as well. To showcase this effect we consider the case of the undirected cycle graph \mathcal{C}_n , *i.e.* the graph on n nodes containing a single cycle through all of them. It is easy to observe that for such graph the non-normalized PageRank vector is $\mathbf{s}_\alpha = \mathbf{1}$, $\forall \alpha \geq 0$. Let us now add a directed edge between the ℓ th and the 1st node of the graph, and evaluate how the PageRank vector for the relative value of α is modified. It is known that for $\alpha = \infty$ this modification produces a strong localization effect on the standard PageRank vector \mathbf{s}_∞ ([49], Thm. 8.1). The bound obtained in Theorem 4.4 suggests that the relative change on the vector \mathbf{s}_α decreases with smaller values of α , *i.e.* the localization effect is milder on the nonlocal PageRank vector. This is shown in Figure 4 where we plot the PageRank vector of the cycle \mathcal{C}_{100} perturbed with one additional edge that connects nodes 1 and 40, for different values of α and for $\delta = \delta_\Gamma$. Consistently with the analysis carried out in [49], we observe two localized peaks that appear in the PageRank vector \mathbf{s}_∞ entries corresponding to the nodes 1 and 40. At the same time, nonlocal versions of PageRank smooth out these peaks as soon as we let the value of α decrease.

In Figure 5 we show the localization behavior of nonlocal PageRank vectors on the `adjnoun` real-world network. As expected, the standard PageRank algorithm concentrates the measure around few nodes (hubs)




FIGURE 4. Locality of the PageRank vector for the undirected cycle graph with $n = 100$ nodes when a single directed edge between nodes 40 and 1 is added to the graph (*right* panel). The PageRank vector for the undirected cycle is the vector $\mathbf{1}$. The smaller the value of α is, the more stable the measure is with respect to the addition of this new edge since the curves for the smaller values of α tends to the PageRank of the unmodified graph. On the other hand, for growing values of α the the curve representing the value of the nonlocal PageRank converges to curve of the standard PageRank (represented as the dot-dashed black line on the *left* panel).

and assigns small and almost indiscernible values to the nodes that occupy the lower part of the ranking, whereas tuning the parameter α allows us spread the measure more evenly.

The nonlocality properties of \mathbf{s}_α are useful in a number of situations. As an example of the improvements one can obtain by tuning the parameter α , while fixing the choice of the distance to the shortest-path distance δ_Γ , in the next section we consider the link prediction task with rooted PageRank similarity.

4.2. Link prediction

Link prediction is an important task in network analysis, which consists of the problem of predicting the existence of one or more missing (unobserved) edges in a given instance of a network Γ [1, 41]. More precisely, we suppose having a snapshot $\Gamma_0 = (V, E_0)$ at time t_0 of a graph, and we want to guess what edges will be added at a subsequent time step t_1 , in which the graph becomes $\Gamma_1 = (V, E_1)$, with $E_1 = E_0 \cup E_{\text{add}}$, and $E_{\text{add}} \subset V \times V \setminus E_0$. Two typical scenarios where this problem applies are the case of an evolving network and the case of network data affected by noise, where it is suspected that a certain number of edges are missing.

A successful approach for link prediction works by first assigning every edge $e_{i \rightarrow j}$ in $V \times V \setminus E_0$ a score, $\text{score}(i, j)$, based on the graph Γ_0 . In this way a ranked list of edges is produced, in decreasing order of $\text{score}(i, j)$, and the new edges defining E_1 are taken as the edges with higher score.

Rooted (or seeded) PageRank is an established method for assigning such scores, based on the PageRank transition matrix. By extending that method, we consider a nonlocal rooted PageRank similarity, where the whole matrix of the scores of all the edges of type $e_{i \rightarrow j}$ for the parameter α is defined by

$$(S_\alpha)_{i,j} = \text{score}_\alpha(i, j) = (X_\alpha + X_\alpha^T)_{i,j}, \quad X_\alpha = (1 - c)(I - cP_\alpha^T)^{-1}, \quad (4.8)$$

being P_α the nonlocal transition probability matrix in (3.1). Note that, due to Lemma 4.3, we have that score_α coincides with the standard rooted PageRank similarity score when $\alpha \rightarrow \infty$.

In what follows we compare the link prediction performance of the nonlocal PageRank with the one obtained by using the standard PageRank. The test networks are both real-world examples: `adjnoun`, `USAir97`, `ca99`, `EuAirComplete`; and synthetic graphs `deLaunay_n10`, `deLaunay_n12`, and `3elt`. The network `EuAirComplete` is

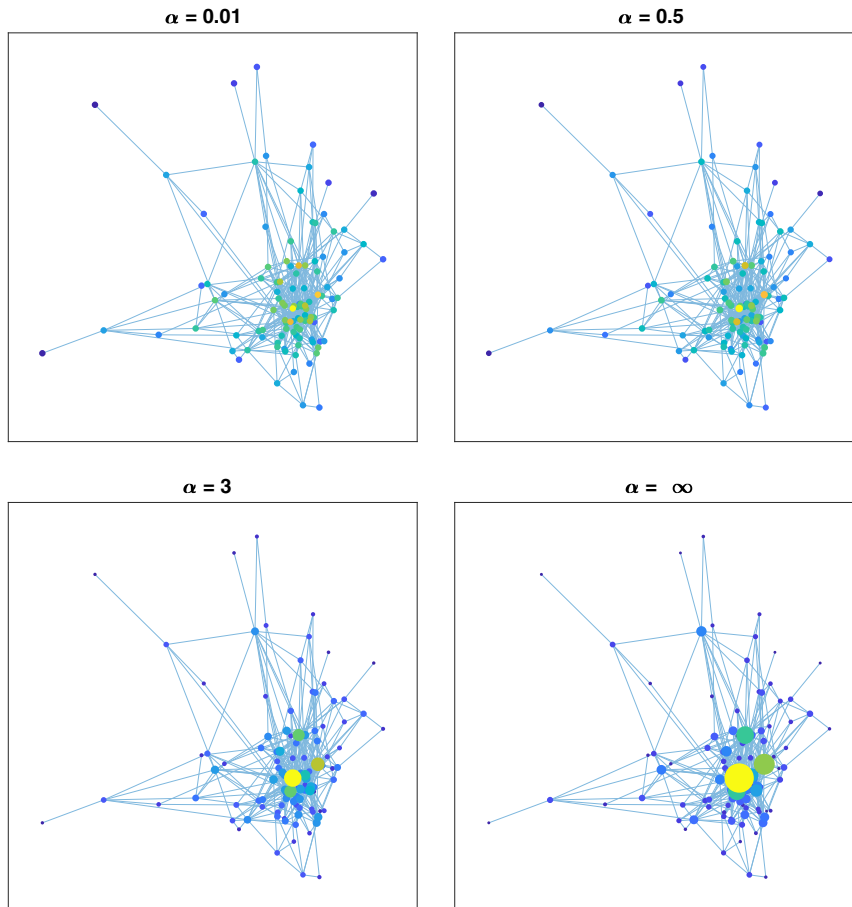


FIGURE 5. Localization of the PageRank values for the graph Newman/adjnoun contains the network of common adjective and noun adjacencies for the novel “David Copperfield” by C. Dickens, as described in [46]. We have computed the PageRank vector for different values of α and plotted here the same graph with the nodes of the size representing the value assigned them by the algorithm. All values are consistently scaled. We observe that as α decreases the localization of the measure is sensibly reduced.

the flattened version of the multilayer network EuAir from [10], representing the connection of airports in the European Union by the different airlines operating there. We subtract from each of the graphs 10% of the edges chosen uniformly at random and we try to guess back all of them. To determine the parameters needed for the predictions, namely the α and c parameters for the nonlocal PageRank, and the c parameter for the rooted PageRank, we apply a 10-fold cross-validation procedure on the smaller graph. The parameters are chosen on the grids $c \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.85, 0.9, 0.99\}$ and $\alpha \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 2, 3, 5\}$. The whole procedure is then repeated fifteen times and the results obtained are then reported in Figure 6 as boxplots showing median and quartiles of the obtained prediction accuracy on the fifteen trials. What we observe is that harvesting information from far-away nodes, *i.e.* exploiting the slower decay of the similarity matrix S_α in (4.8), yields better results. When such behavior is not observed the cross-validation pushes then the solution to be the one obtained with the rooted PageRank. The selected α and c parameters from the cross-validation are available in the repository together with the code. For what concerns the cost, here it is dominated by




FIGURE 6. Boxplot showing median and quartiles of prediction accuracy over 15 trials on seven synthetic benchmarks and real-world datasets, for standard rooted PageRank versus nonlocal PageRank with smoothing function $f_\alpha(x) = x^{-\alpha}$. For each trial we remove 10% of edges chosen uniformly at random and we try to guess back all of them, choosing the parameters α and c for nonlocal PageRank and the parameter c for rooted PageRank *via* 10-fold cross validation on the smaller graph (with the removed edges) using the grids $c \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.85, 0.9, 0.99\}$ and $\alpha \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 2, 3, 5\}$.

the $O(n^3)$ matrix inversion in (4.8), thus the overall cost of both the standard and the nonlocal PageRank are comparable.

5. CHOOSING THE APPROPRIATE DISTANCE

The shortest path distance in Definition 4.1 is not the only possible choice for generating the transition matrix in (3.1). Being able to choose the distance function δ is an additional degree of flexibility of the proposed nonlocal PageRank which allows us to adapt the model more tightly to the problem. In principle any graph distance, respectively metric, can be adopted to define the transition probability matrix and the choice is a matter of modeling reasons.

For illustration purposes, in the next section we first describe an example of graph metric obtained by using the logarithmic distance. Then, in Section 5.2 we consider the London underground train multilayer network and define a problem-dependent *metro distance*. This distance takes into account for the multiple layers and allows us to improve the centrality assignment of the train stations, when compared to independent station usage data we collected from [47].

5.1. An example of digraph metric: Logarithmic distance

We illustrate here the behaviour of the nonlocal PageRank obtained with a distance that generates a metric on the graph, *i.e.* that is both symmetric, and satisfies the triangle inequality. To introduce such metric, called Logarithmic distance [12, 13], we need to start from a particular *proximity* (*similarity*) measure, $S = (s_{i,j}) = (s(i, j))$. Such measure is defined in terms of the following Laplacian matrix L of the digraph Γ

$$L = D_{\text{out}} - A, \quad (5.1)$$

as

$$S = (I + L)^{-1}, \quad (5.2)$$

where D_{out} is the diagonal matrix of the out degrees in (2.1) and A is the adjacency matrix of Γ .




FIGURE 7. HB/gre_115 digraph. The figure shows the adjacency and distance matrices for the loop-less version of the HB/gre_115 digraph.

This measure still accounts for the lack of symmetry of the underlying digraph, *i.e.* of its adjacency matrix, and satisfies both the *transition inequality* and *graph bottleneck identity*, *i.e.* it is such that $s_{i,j}s_{j,k} \leq s_{i,k}s_{j,j}$, and $s_{i,j}s_{j,k} = s_{i,k}s_{j,j}$ if and only if every path from i to k contains j . In the case of undirected graphs (5.2) is usually called the regularized Laplacian kernel. Such definition is indeed well posed, *i.e.* we are ensured that the matrix $I + L$ can be inverted, because $I + L$ is an example of a nonsingular M -matrix (*cf.* [8]). Note that this is sufficient to guarantee also that all the elements of S are nonnegative; see [8] for further details.

To obtain the Logarithmic distance based on the similarity S we then build the matrix $H = (h_{i,j})$ and the vector \mathbf{h} defined as

$$h_{i,j} = \log(s_{i,j}) \leq 0, \quad \mathbf{h} = (h_{1,1}, h_{2,2}, \dots, h_{n,n})^T \quad (5.3)$$

from which we define the *logarithmic distance* as

$$\delta_{\log}(i, j) = \left(\frac{1}{2}(U + U^T) \right)_{i,j}, \quad U = \mathbf{h}\mathbf{1}^T - H. \quad (5.4)$$

We stress that δ_{\log} , differently from the shortest-path distance in Definition 4.1, generates a metric on the digraph, see *e.g.* Theorem 1 of [13]. This can also be observed by the example network in Figure 7 in which we show the comparison of the entries of the distance matrix for the shortest path distance (central panel) and for the logarithmic distance (right panel) as compared to the pattern of the adjacency matrix of the real world network gre_115 (left panel).

Instead, in order to better grasp the differences between the centralities obtained, in Figure 8 we scatter plot the behavior of the nonlocal PageRank on the dataset USAir97, for the two different smoothing functions $f(x)_\alpha = 1/x^\alpha$ and $f_\alpha(x) = e^{-\alpha x}$ and the two choices of graph distances $\delta = \delta_\Gamma$ and $\delta = \delta_{\log}$.

5.2. Metro distance and the London's train multilayer network

In this final example scenario of applicability for the proposed nonlocal PageRank model we consider the ranking in term of importance of urban railways. Tools from network science have proved to be valuable in the study of urban transport [22, 58] and here we consider the use of nonlocal PageRank network centrality in the case of the London train network. Nodes in the network represent stations, and we seek a distance δ so that the resulting centrality measure correlates well with passenger usage. Such a measure, which requires only the topological connectivity structure, offers helpful information at the design stage. More importantly, it can be used in what-if-scenario testing, in order to predict the effect of changes, including unplanned network disruptions.

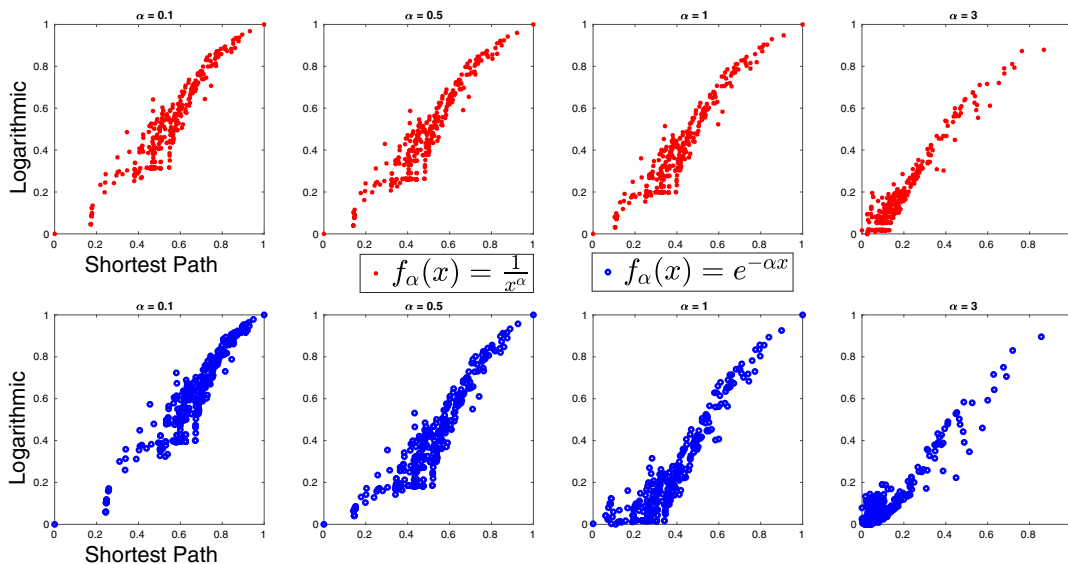


FIGURE 8. USAir97 digraph. The figure shows scatter plots of Nonlocal PageRank obtained using shortest path distance *vs.* Nonlocal PageRank obtained using logarithmic distance for different $f_\alpha(x)$ and different values of α .

The London train network is an undirected transportation network representing connections between tube train stations of the city of London. The dataset `tube` we consider here is a multilayer version of the train network, where each underground line corresponds to one layer. The dataset has been generated using the data from [20] as baseline. The aggregate network consists of one connected component with 271 nodes and 315 edges with nonzero weights. Each edge encodes the information about the membership of a given node to one – or more than one in the case of intersections – of the $k = 11$ different underground lines. We collect additional passenger data from [47]: for each train station of the above network, we collect the number of passengers entering or exiting that station per year. We collect data for ten years: from 2008 to 2017. Our data is publicly available at <https://github.com/Cirdans-Home/NonLocalPageRank>.

Motivated by the question of whether we can identify highly populated stations by exploiting only the topology of connections between stations, in this section we study the behavior of the nonlocal PageRank centrality with different choices of the distance function δ and of the decaying parameter α . In particular, in the next subsection we will design a distance that is specifically conceived for this issue and we will show that this indeed helps boosting the performance of the centrality model in this context.

The metro distance

Every day experience using underground trains suggests that not all the paths between two connected nodes i and j of the network are equally attractive: we posit that users tend to prefer paths that avoid line changes (or minimize them) even if that means choosing a path that is longer in terms of the number of stations that the path involves.

This argument suggests that the shortest path distance $\delta_\Gamma(i, j)$ is not appropriate to faithfully model the graph exploration of a rail traveler. Hence, we consider here a natural modification of the shortest path distance δ_Γ , which we will call *metro distance* and which we define as follows

$$\delta_M(i, j) = \begin{cases} \delta_\Gamma(i, j) & \text{if } i, j \text{ are on the same metro line,} \\ \delta_\Gamma(i, j) + C(i, j) & \text{if } i, j \text{ are on different metro lines,} \end{cases} \quad (5.5)$$




FIGURE 9. Non-local PageRank ranking performance. Year 2017.

where $C(i, j)$ is the number of times a traveler needs to change train line when traveling from node i to node j . In other words, the metro distance penalizes paths using as penalization parameter the number of times the passenger changes line.

Our aim is now to test the extent to which the nonlocal PageRank centrality with δ_M can identify nodes that perform better, in terms of total passenger usage, than other PageRanks. In order to measure this type of performance for the different ranking strategies, we consider the *Intersection Similarity* (ISIM), whose definition we briefly recall. The intersection similarity is a measure that compares two ranked lists that may not contain the same elements. It is defined as follows [29]: let \mathbf{p}, \mathbf{q} be two ranked lists of n elements we want to compare. The intersection similarity of \mathbf{p} and \mathbf{q} is the vector $\text{ISIM}(\mathbf{p}, \mathbf{q})$ with entries

$$\text{ISIM}(\mathbf{p}, \mathbf{q})_k = \frac{1}{k} \sum_{j=1}^k \frac{|\Delta((p_1, \dots, p_j), (q_1, \dots, q_j))|}{2j}, \quad k = 1, \dots, n$$

where, for sets S, T , $|S|$ denotes the cardinality and Δ is the symmetric difference operator $\Delta(S, T) := (S \setminus T) \cup (T \setminus S)$.

When the first k entries in \mathbf{p} and \mathbf{q} are completely different $\text{ISIM}(\mathbf{p}, \mathbf{q})_k$ is equal to 1, whereas $\text{ISIM}(\mathbf{p}, \mathbf{q})_k = 0$ if and only if the first k entries in \mathbf{p} and \mathbf{q} coincide exactly. More in general, lower values in $\text{ISIM}(\mathbf{p}, \mathbf{q})$ imply a better matching between \mathbf{p} and \mathbf{q} .

In Figure 9 we aim at comparing the top fifteen stations identified by the following three models

- the nonlocal PageRank with δ_Γ (“*SP distance*”);
- the nonlocal PageRank with δ_M (“*Metro distance*”);
- the standard PageRank.

with the “ground truth”, *i.e.* the actual fifteen stations with the highest number of passengers, whose corresponding ranked list we denote by ρ . In particular, we compute the intersection similarity between the ranking of any of the three above PageRanks and ρ , which we denote respectively by $\text{ISIM}(\delta_\Gamma)$, $\text{ISIM}(\delta_M)$ and $\text{ISIM}(pr)$.

The two curves in the left panel of Figure 9 show the ratios

$$\frac{\text{ISIM}(\delta_\Gamma)_{15}}{\text{ISIM}(pr)_{15}} \quad \text{and} \quad \frac{\text{ISIM}(\delta_M)_{15}}{\text{ISIM}(pr)_{15}}$$

with blue and red lines, respectively, for different values of the parameter α , computed with the power-law decaying function $f_\alpha(x) = 1/x^\alpha$.

Whereas, the central panel compares the cumulative sum of passenger usage values for the 15 top ranked stations for the three rankings. This comparison is done by showing both the ratio between the cumulative sum

TABLE 1. Millions of passengers per year using the top $k = 5, 15, 45$ stations ranked according to the standard “local PageRank” the nonlocal PageRank with δ_Γ ($\alpha = 1.7$) and the nonlocal PageRank with δ_M ($\alpha = 1.7$). We consider the range 2017–2008.

	k	2017	2016	2015	2014	2013	2012	2011	2010	2009	2008
Ground truth	5	421.6895	433.7675	437.0499	443.6953	413.9785	399.7917	386.0158	366.9199	353.7508	359.7531
	15	891.2556	918.167	920.4402	918.3901	867.6314	835.8205	805.7167	755.9158	731.6231	739.0077
	45	1614.8047	1665.0602	1615.3159	1615.757	1542.1181	1490.7425	1430.1159	1374.3307	1329.411	1348.2863
“Local” PageRank	5	286.787	294.6265	288.0692	286.9402	273.3404	264.2259	255.6953	245.6262	229.4468	230.0058
	15	580.5896	592.7364	588.1613	589.1657	545.3233	522.9533	504.7835	480.4868	462.1433	473.7567
	45	1260.4716	1294.76	1259.7501	1259.2124	1204.2262	1159.854	1116.8103	1070.5588	1036.5845	1050.4783
Nonlocal PageRank δ_Γ	5	341.4023	349.6481	349.839	353.9939	323.0698	311.4126	297.0677	279.9141	268.7073	271.4097
	15	746.1333	768.6693	762.5181	751.9844	723.7442	699.2322	675.9891	648.6213	624.0347	629.1298
	45	1302.5842	1320.4892	1320.0771	1324.7913	1269.6994	1222.4786	1181.0192	1150.9596	1123.1939	1133.568
Nonlocal PageRank δ_M	5	341.4023	349.6481	349.839	353.9939	323.0698	311.4126	297.0677	279.9141	268.7073	271.4097
	15	758.6355	785.5156	774.1294	761.5849	733.9973	709.7501	686.8022	659.9881	634.5561	639.799
	45	1352.7352	1373.8639	1362.7264	1366.1506	1300.7856	1251.9737	1203.3251	1154.5498	1124.5568	1139.3027

of passenger usage values corresponding to δ_Γ and the one corresponding to the standard PageRank (blue line) and the ratio between the one corresponding to δ_M and the one corresponding to the standard PageRank (red line), for different values of α . In both panels we observe the metro distance δ_M performing generally better than the standard shortest path distance, with relative optimal performance obtained for $\alpha = 1.7$. This is further highlighted by the panel on the right of Figure 9, where we compare the first 25 entries of the three different intersection similarity vectors $\text{ISIM}(\delta_\Gamma)$, $\text{ISIM}(\delta_M)$ and $\text{ISIM}(pr)$, for $\alpha = 1.7$.

As shown in Figure 9, a proper choice of the parameter α , allows for a more accurate ranking than the local PageRank. In particular, choosing $\alpha = 1.7$ and the metro distance δ_M , we are able to match the ground truth ranking more closely if compared with the local PageRank and the nonlocal PageRank employing the distance δ_Γ . Finally, as the right panel of Figure 9 shows, even though the ISIM performance with respect to the ground truth ranking coincides on the top ranked nodes for the standard PageRank and the nonlocal PageRank with *Metro distance*, the nonlocal model allows us to obtain sensibly better performance in terms of ISIM when a greater number of ranked nodes is taken into account. This issue is further corroborated from the results presented in Table 1 where a similar behavior is observed across the ten-year span usage data; here we compare the number of passengers using the top $k = 5, 15, 45$ stations identified by local/nonlocal Pagerank as before. Finally, in Table 2 and Figure 10 we show the name and the geographic collocations of the top 10 ranked stations according to the considered different ranking algorithms. The overall emerging experimental evidence from Figures 9, 10 and Tables 1, 2, highlights how the flexibility of the proposed model allows us to design centrality that better adapt to the specific problem. While the nonlocal PageRank model with the metro distance does not yield a perfect matching with the ground truth, it outperforms other PageRank models and obtains remarkable performance which we find particularly interesting given that the model exploits only the topological structure of nodes and edges.

Computation of the metro distance

To obtain the penalized version of the shortest path distance $\delta_M(i, j)$, see (5.5), we synthesized from a *multilayer* interpretation for the graph Γ . This is a particularly useful approach and, in this case, enabled us to encode the information coming from the connectedness of two nodes i and j *via* the line k . The multilayer structure of the graph can be naturally represented using a tensor $\mathcal{T} \in \mathbb{R}^{n \times n \times k}$ such that

$$\mathcal{T}_{ij\ell} = \begin{cases} 1 & \text{if } i \sim j \text{ on the line } \ell \\ 0 & \text{otherwise.} \end{cases} \quad (5.6)$$

In the following we will use the following Matlab notation: $\mathcal{T}(:, :, \ell) := (\mathcal{T}(i, j, \ell))_{i, j \in \{1, \dots, n\}}$. Observe that for every $\ell = 1, \dots, k$, since we are considering undirected connections among nodes, we have $\mathcal{T}(:, :, \ell) = \mathcal{T}(:, :, \ell)^T$;

TABLE 2. Ten London train stations with highest ranking value, according to the standard “*local PageRank*”, the nonlocal PageRank with δ_Γ and the nonlocal PageRank with δ_M ($\alpha = 1.7$ in both cases). Year 2017.

Ground truth		“Local” PageRank		Nonlocal PageRank δ_Γ		Nonlocal PageRank δ_M	
King’s Cross	97.9183	King’s Cross	97.9183	Green Park	39.3382	King’s Cross	97.9183
Waterloo	91.2706	Baker Str.	28.7846	Baker Str.	28.7846	Baker Str.	28.7846
Oxford Circus	84.0906	Paddington	48.8225	Oxford Circus	84.0906	Green Park	39.3382
Victoria	79.3593	Earl’s Court	19.991	King’s Cross	97.9183	Oxford Circus	84.0906
London Bridge	69.0507	Waterloo	91.2706	Waterloo	91.2706	Waterloo	91.2706
Liverpool Str.	67.7402	Turnham Green	6.1552	Bond Str.	38.8027	Paddington	48.8225
Stratford	61.9904	Green Park	39.3382	Bank	30.8981	Bank	30.8981
Canary Wharf	50.9136	Oxford Circus	84.0906	Westminster	25.5954	Bond Str.	38.8027
Paddington	48.8225	Stockwell	11.6971	Paddington	48.8225	Earl’s Court	19.991
Euston	43.0737	Liverpool Str.	67.7402	Liverpool Str.	67.7402	Euston	43.0737

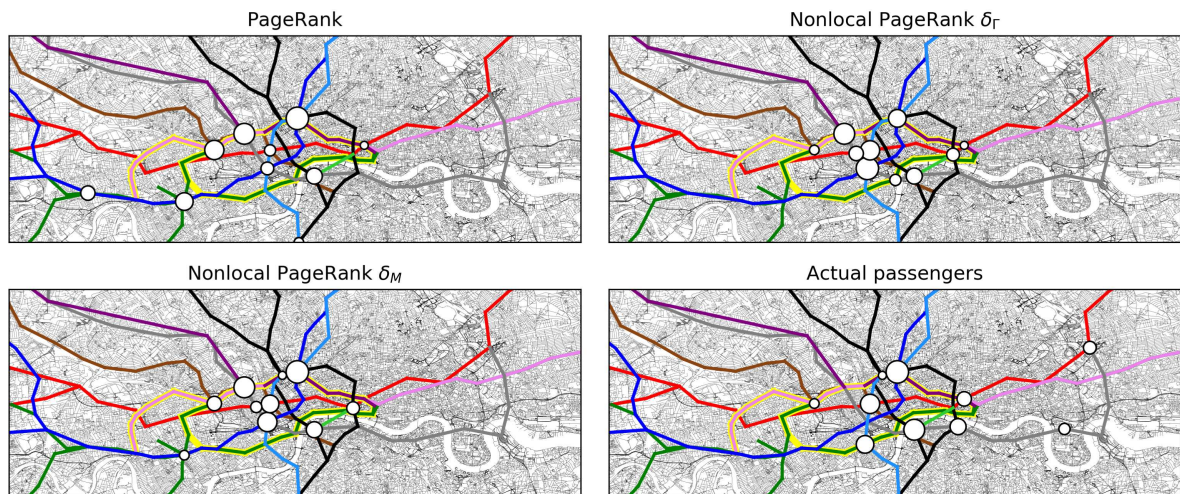


FIGURE 10. This figure shows the name and the geographic collocations of the top 10 ranked stations according to the considered different ranking algorithms together with the top 10 ranked stations with respect to the actual number of passengers.

moreover, $\text{sign}(\sum_{\ell=1}^k T(:, :, \ell))$ returns exactly the adjacency matrix of the full graph Γ . By mapping each node $i \in \{1, \dots, n\}$ of the graph in i_1, \dots, i_k , it is possible to form the block matrix

$$T = \begin{bmatrix} T_{1,1} & \dots & T_{1,k} \\ \vdots & \vdots & \vdots \\ T_{k,1} & \dots & T_{k,k} \end{bmatrix} \in \mathbb{R}^{nk \times nk} \quad (5.7)$$

such that $T_{\ell,\ell} = \mathcal{T}(:, :, \ell)$ and $(T_{\ell_1,\ell_2})_{i,i} = (T_{\ell_2,\ell_1})_{i,i} = 1$ if the node $i \in \{1, \dots, n\}$ is at the intersection of the metro lines ℓ_1 and ℓ_2 ; the remaining elements of T_{ℓ_1,ℓ_2} and T_{ℓ_2,ℓ_1} are set to zero. Now, in the *expanded* graph $\bar{\Gamma} = \text{adj}(T)$ with nodes $\{1_1, \dots, n_1, \dots, 1_k, \dots, n_k\}$, let us analyze the path $j_{\ell_1} \rightarrow i_{\ell_2} \rightarrow u_{\ell_3} \rightarrow e_{\ell_4}$ with $\{j, i, u, e\} \in \{1, \dots, n\}$; it is easy to recognize that the case $i = u$ corresponds to the case where a traveler changes metro line ℓ_2 into metro line ℓ_3 at the node i . The metro distance $\delta_M(i, j)$ is then obtained considering $\delta_M(i, j) = \min_{\ell_1=1, \dots, k, \ell_2=1, \dots, k} \delta_{\Gamma}(i_{\ell_1}, j_{\ell_2})$, being δ_{Γ} the shortest path distance of the nodes i_{ℓ_1}, j_{ℓ_2} computed in $\bar{\Gamma}$. As in most of the other examples, the main computational cost is represented by the computation of the shortest path distance on $\bar{\Gamma}$.

6. CONCLUSIONS AND FUTURE WORK

In this work we have introduced a nonlocal version of the classic PageRank model. The new version encompasses a nonlocal navigation strategy of the underlying network, permitting the usage of any suitable graph distance. Generalizing the Lévy and exponential transition models, we have introduced a general definition for a class of functions which can be used to modulate the range of the interactions.

With the approach presented here it is possible to mitigate several typical phenomena occurring in eigenvector centralities, such as the phenomenon of localization of the measure, and the issue of the assignation of numerically indistinguishable values to nodes that are in the lower part of the ranking. The mitigation of these behaviors increases the predictive power of the nonlocal PageRank when compared with the local PageRank, as it has been confirmed by the real-world applications presented in this work.

Even if the model we present encodes a completely different dynamics for the network interactions, if compared with the local approach, it is still possible to use a standard series of numerical tools for the efficient computations of PageRank vectors, see for example [9, 15, 21, 34]. The main differences are represented by the setup phase of the algorithm, *i.e.* by the need of computing the distance matrix for the underlying network, and by the fact that nonlocal transition matrices are in general not sparse. On the other hand, suitable choices of the smoothing functions f_{α} may lead to structured transition matrices (as in the case of fractional derivatives *e.g.* [44]) and exploring this line of research may lead to efficient methods for using nonlocal PageRank on large scale problems, an issue that will be object of future investigations.

Acknowledgements. The work of F.D. has been supported by the GNCS-INdAM project “Nonlocal models for the analysis of complex networks”.

REFERENCES

- [1] M. Al Hasan and M.J. Zaki, A survey of link prediction in social networks. In: Social Network Data Analytics. Springer, Boston (2011) 243–275.
- [2] F. Arrigo, P. Grindrod, D.J. Higham and V. Noferini, Non-backtracking walk centrality for directed networks. *J. Complex Networks* **6** (2017) 54–78.
- [3] F. Arrigo, D.J. Higham and F. Tudisco, A framework for second order eigenvector centralities and clustering coefficients. *Proc. R. Soc. A* **476** (2020) 20190724.
- [4] A.R. Benson, Three hypergraph eigenvector centralities. *SIAM J. Math. Data Sci.* **1** (2019) 293–312.
- [5] A.R. Benson, D.F. Gleich and J. Leskovec, Higher-order organization of complex networks. *Science* **353** (2016) 163–166.
- [6] A.R. Benson, D.F. Gleich and L.-H. Lim, The spacey random walk: a stochastic process for higher-order data. *SIAM Rev.* **59** (2017) 321–345.
- [7] M. Benzi and C. Klymko, Total communicability as a centrality measure. *J. Complex Networks* **1** (2013) 124–149.

- [8] A. Berman and R.J. Plemmons, Nonnegative Matrices in the Mathematical Sciences. In: Vol. 9 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA (1994).
- [9] C. Brezinski and M. Redivo-Zaglia, The PageRank vector: properties, computation, approximation, and acceleration. *SIAM J. Matrix Anal. Appl.* **28** (2006) 551–575.
- [10] A. Cardillo, J. Gómez-Gardeñes, M. Zanin, M. Romance, D. Papo, F. Del Pozo and S. Boccaletti, Emergence of network features from multiplexity. *Sci. Rep.* **3** (2013) 1344.
- [11] T.-H.H. Chan, A. Louis, Z.G. Tang and C. Zhang, Spectral properties of hypergraph Laplacian and approximation algorithms. *J. ACM* **65** (2018) 1–48.
- [12] P. Chebotarev, A class of graph-geodesic distances generalizing the shortest-path and the resistance distances. *Disc. Appl. Math.* **159** (2011) 295–302.
- [13] P. Chebotarev, The graph bottleneck identity. *Adv. Appl. Math.* **47** (2011) 403–413.
- [14] G.E. Cho and C.D. Meyer, Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra App.* **335** (2001) 137–150.
- [15] S. Cipolla, C. Di Fiore and F. Tudisco, Euler-Richardson method preconditioned by weakly stochastic matrix algebras: a potential contribution to Pagerank computation. *Electron. J. Linear Algebra* **32** (2017) 254–272.
- [16] S. Cipolla, M. Redivo-Zaglia and F. Tudisco, Extrapolation methods for fixed-point multilinear PageRank computations. *Numer. Linear Algebra App.* **27** (2020) e2280.
- [17] S. Cipolla, M. Redivo-Zaglia and F. Tudisco, Shifted and extrapolated power methods for tensor ℓ^p -eigenpairs. *Electron. Trans. Numer. Anal.* **53** (2020) 1–27.
- [18] R. Cohen and S. Havlin, Complex Networks: Structure, Robustness and Function. Cambridge University Press (2010).
- [19] T.A. Davis and Y. Hu, The University of Florida sparse matrix collection. *ACM Trans. Math. Softw.* **38** (2011) 1–25.
- [20] M. De Domenico, A. Solé-Ribalta, S. Gómez and A. Arenas, Navigability of interconnected networks under random failures. *Proc. Natl. Acad. Sci.* **111** (2014) 8351–8356.
- [21] G.M. Del Corso, A. Gulli and F. Romani, Fast PageRank computation via a sparse linear system. *Internet Math.* **2** (2005) 251–273.
- [22] S. Derrible, Network centrality of metro systems. *PLoS one* **7** (2012) e40575.
- [23] I.S. Duff, R.G. Grimes and J.G. Lewis, Sparse matrix test problems. *ACM Trans. Math. Softw.* **15** (1989) 1–14.
- [24] E. Estrada, Path Laplacian matrices: introduction and application to the analysis of consensus in networks. *Linear Algebra App.* **436** (2012) 3373–3391.
- [25] E. Estrada, The Structure of Complex Networks: Theory and Applications. Oxford University Press (2012).
- [26] E. Estrada, J.-C. Delvenne, N. Hatano, J.L. Mateos, R. Metzler, A.P. Riascos and M.T. Schaub, Random multi-hopper model: super-fast random walks on graphs. *J. Complex Networks* **6** (2017) 382–403.
- [27] E. Estrada, E. Hameed, N. Hatano and M. Langer, Path Laplacian operators and superdiffusive processes on graphs. I. One-dimensional case. *Linear Algebra App.* **523** (2017) 307–334.
- [28] E. Estrada, E. Hameed, M. Langer and A. Puchalska, Path Laplacian operators and superdiffusive processes on graphs. II. Two-dimensional lattice. *Linear Algebra App.* **555** (2018) 373–397.
- [29] R. Fagin, R. Kumar and D. Sivakumar, Comparing top k lists. *SIAM J. Disc. Math.* **17** (2003) 134–160.
- [30] D. Fasino and F. Tudisco, A modularity based spectral method for simultaneous community and anti-community detection. *Linear Algebra App.* **542** (2018) 605–623.
- [31] D. Fasino and F. Tudisco, Ergodicity coefficients for higher-order stochastic processes. *SIAM J. Math. Data Sci.* **2** (2020) 740–769.
- [32] R.W. Floyd, Algorithm 97: shortest path. *Commun. ACM* **5** (1962) 345.
- [33] M. Girvan and M.E. Newman, Community structure in social and biological networks. *Proc. Natl. Acad. Sci.* **99** (2002) 7821–7826.
- [34] G.H. Golub and C. Greif, An Arnoldi-type algorithm for computing PageRank. *BIT Numer. Math.* **46** (2006) 759–771.
- [35] D.J. Higham, A matrix perturbation view of the small world phenomenon. *SIAM Rev.* **49** (2007) 91–108.
- [36] M. Holtgrewe, P. Sanders and C. Schulz, Engineering a scalable high quality graph partitioner. In: *2010 IEEE International Symposium on Parallel Distributed Processing (IPDPS)* (2010) 1–12.
- [37] D.B. Johnson, Efficient algorithms for shortest paths in sparse networks. *J. ACM* **24** (1977) 1–13.
- [38] S. Kirkland, On a question concerning condition numbers for Markov chains. *SIAM J. Matrix Anal. App.* **23** (2002) 1109–1119.
- [39] G. Li, S.D.S. Reis, A.A. Moreira, S. Havlin, H.E. Stanley and J.S. Andrade, Towards design principles for optimal transport networks. *Phys. Rev. Lett.* **104** (2010) 018701.
- [40] P. Li, I. Chien and O. Milenkovic, Optimizing generalized pagerank methods for seed-expansion community detection, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox and R. Garnett. In: Vol. 32 of *Advances in Neural Information Processing Systems*. Curran Associates, Inc. (2019) 11710–11721.
- [41] D. Liben-Nowell and J. Kleinberg, The link-prediction problem for social networks. *J. Am. Soc. Inf. Sci. Technol.* **58** (2007) 1019–1031.
- [42] X. Liu, G. Strang and S. Ott, Localized eigenvectors from widely spaced matrix modifications. *SIAM J. Disc. Math.* **16** (2003) 479–498.
- [43] T. Martin, X. Zhang and M.E.J. Newman, Localization and centrality in networks. *Phys. Rev. E* **90** (2014) 052808.

- [44] S. Massei, M. Mazza and L. Robol, Fast solvers for two-dimensional fractional diffusion equations using rank structured matrices. *SIAM J. Sci. Comput.* **41** (2019) A2627–A2656.
- [45] M. Neumann and J. Xu, Improved bounds for a condition number for Markov chains. *Linear Algebra App.* **386** (2004) 225–241.
- [46] M.E.J. Newman, Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* **74** (2006) 036104.
- [47] Office of Railand Road. Estimate of London’s tube station usage. <https://dataportal.orr.gov.uk/statistics/usage/estimates-of-station-usage/>.
- [48] L. Page, S. Brin, R. Motwani and T. Winograd, The PageRank citation ranking: bringing order to the web. Technical report, Stanford InfoLab (1999).
- [49] M. Paton, K. Akartunali and D.J. Higham, Centrality analysis for modified lattices. *SIAM J. Matrix Anal. App.* **38** (2017) 1055–1073.
- [50] S. Pozza and F. Tudisco, On the stability of network indices defined by means of matrix functions. *SIAM J. Matrix Anal. App.* **39** (2018) 1521–1546.
- [51] A.P. Riascos and J.L. Mateos, Long-range navigation on complex networks using Lévy random walks. *Phys. Rev. E* **86** (2012) 056110.
- [52] M.R. Roberson and D. ben Avraham, Kleinberg navigation in fractal small-world networks. *Phys. Rev. E* **74** (2006) 017101.
- [53] M.P. Rombach, M.A. Porter, J.H. Fowler and P.J. Mucha, Core-periphery structure in networks. *SIAM J. Appl. Math.* **74** (2014) 167–190.
- [54] R.A. Rossi and N.K. Ahmed, The network data repository with interactive graph analytics and visualization. In: *AAAI* (2015).
- [55] E. Seneta, Perturbation of the stationary distribution measured by ergodicity coefficients. *Adv. Appl. Probab.* **20** (1988) 228–230.
- [56] K.J. Sharkey, Localization of eigenvector centrality in networks with a cut vertex. *Phys. Rev. E* **99** (2019) 012315.
- [57] B. Stabler, H. Bar-Gera and E. Sall, Transportation networks for research. <https://github.com/bstabler/TransportationNetworks>.
- [58] W.M. To, Centrality of an urban rail system. *Urban Rail Transit* **1** (2015) 249–256.
- [59] F. Tudisco, A note on certain ergodicity coefficients. *Spec. Matrices* **3** (2015).
- [60] F. Tudisco and D.J. Higham, A nonlinear spectral method for core–periphery detection in networks. *SIAM J. Math. Data Sci.* **1** (2019) 269–292.
- [61] A. van Heukelum, G. Barkema and R. Bisseling, DNA electrophoresis studied with the cage model. *J. Comput. Phys.* **180** (2002) 313–326.
- [62] S. Warshall, A theorem on boolean matrices. *J. ACM* **9** (1962) 11–12.
- [63] T. Weng, M. Small, J. Zhang and P. Hui, Lévy walk navigation in complex networks: a distinct relation between optimal transport exponent and network dimension. *Sci. Rep.* **5** (2015) 1–9.