Tie-Decay Networks in Continuous Time and Eigenvector-Based Centralities

Walid Ahmad, Mason A. Porter, and Mariano Beguerisse-Díaz

Abstract—**Network theory is a useful framework for studying interconnected systems of interacting entities. Many networked systems evolve continuously in time, but most existing methods for the analysis of time-dependent networks rely on discrete or discretized time. In this paper, we propose an approach for studying networks that evolve in continuous time by distinguishing between *interactions*, which we model as discrete contacts, and *ties*, which encode the strengths of relationships over time. To illustrate our tie-decay network formalism, we adapt the well-known PageRank centrality score to our tie-decay framework in a mathematically tractable and computationally efficient way. We apply this framework to a synthetic example and then use it to study a network of retweets during the 2012 National Health Service controversy in the United Kingdom. Our work also provides guidance for similar generalizations of other tools from network theory to continuous-time networks with tie decay, including for applications to streaming data.**

Index Terms—**complex networks, network theory (graphs).**

I. Introduction

NETWORKS provide a versatile framework to model and analyze complex systems of interacting entities [59]. In many complex systems, interaction patterns change in time and the entities can also leave or enter the system at different times. To accurately model and understand such systems, it is essential to incorporate temporal information about their interactions into network representations [5], [9], [13], [17], [38], [73], [81]. See Refs. [32]–[34], [52] for overviews of the study of time-dependent networks, which are often also called *temporal networks* or *dynamic networks*.

A major challenge in the analysis of temporal networks is that one often has to discretize time by aggregating connections into time windows. Given a discrete or discretized set of interactions, one can then analyze communities, important nodes, and other facets of temporal networks by examining a multilayer-network representation of these interactions [1], [32], [40], [75]. An important challenge that arises with aggregation is that there may not be any obvious or even any 'correct' size of a time window (even when such aggregation employs non-uniform time windows [12], [68], [69], [74]). A window that is too small risks missing important network structures (e.g., by construing a signal as noise), but using an overly large window may obscure important temporal features. (See [18] for one discussion.) Moreover, in many social systems, interactions are bursty [4], [33], [41], which is a crucial consideration when aggregating interactions [31] and can be a major source of concern when using homogeneous time windows [69]. Bursty interactions pose a challenge not only when choosing the width of time windows, but also when choosing where to place the boundaries of such windows. Shifting time windows forward or backward may significantly alter the statistics of a data set, even when one does not change the width of the windows [41].

From a modeling perspective, aggregating interactions often may not be an appropriate approach for studying systems with asynchronous activity or systems that evolve continuously in time. See [79] for an investigation of biological contagions, [84] for a study of influential accounts in social networks, [85] and [86] for a generalization of the formalism of 'activity-driven networks' to continuous time, [57] for a study of rankings in competitive sports, and [19] for a general continuous-time framework for temporal networks. In many cases, contacts in a temporal network can have a non-instantaneous duration, and it can be important to take such information into account [61], [71]. For example, the phone-call data that were studied in [28] require contacts to exist for the duration of a phone call. In other cases, interactions can be instantaneous (e.g., a mention in a tweet, a text message, and so on), and their importance decreases over time [11], [46]. In many temporal networks (e.g., ones that involve feeds on social media), there is also a decay in attention span in reactions to content [30], [49], [50].

In the present paper, we introduce a framework for modeling temporal networks in which the strength of a connection (i.e., a tie) can evolve continuously in time. For example, perhaps the strength of a tie decays exponentially after the most recent interaction. (One can also use point-process models such as Hawkes processes [46] to examine similar ideas from a node-centric perspective.) Our mathematical formalism of such 'tie-decay networks' allows us to examine them using analytical calculations and to implement them efficiently in real-world applications with streaming data. We showcase our tie-decay formalism by computing continuous-time PageRank

Manuscript received September 19, 2019; revised December 31, 2020; accepted April 3, 2021. Date of publication April 6, 2021; date of current version July 7, 2021. Recommended for acceptance by Prof. Xianbin Cao. *(Corresponding author: Mason A. Porter.)*

Walid Ahmad was with Mathematical Institute, University of Oxford, Oxford OX2 6GG, U.K. He is now with Reverie Labs, Cambridge, MA 02139 USA (e-mail: wna2103@columbia.edu).

Mason A. Porter is with Department of Mathematics, University of California, Los Angeles, CA 90095 USA (e-mail: mason@math.ucla.edu).

Mariano Beguerisse-Díaz is with Mathematical Institute, University of Oxford, Oxford OX2 6GG, U.K. and also with Spotify Ltd., London WC2N 6AT, U.K. (e-mail: marianob@spotify.com).

Digital Object Identifier 10.1109/TNSE.2021.3071429

centrality scores for both a synthetic temporal network and a temporal network that we construct from a large collection of Twitter interactions over the course of several months.

Our paper proceeds as follows. In Section II, we formalize our discussion of ties, interactions, and temporal networks. We also introduce the notion of *tie-decay networks*, which is the focus of our study. In Section III, we adapt PageRank centrality to tie-decay networks. In Section IV, we construct a synthetic network with known properties to illustrate some of the pitfalls of binning interactions and how tie-decay networks can help avoid them. In Section V, we discuss and compute tie-decay PageRank centralities to examine important Twitter accounts in a National Health Service (NHS) retweet network. In Section VI, we conclude and discuss the implications of our work. We give proofs of our main theoretical results in Appendices A and B. We discuss additional numerical computations in Appendices C, D, and E.

II. Ties, Interactions, and Temporal Networks

Our objective is to construct continuous-time temporal networks that capture the evolution of relationships between entities in a system. To do this, we make an important distinction between 'interactions' and 'ties'. An *interaction* between two entities is an event that takes place during a specific time interval or at a specific point in time. Examples of interactions include face-to-face meetings, text messages, and phone calls. A *tie* between two entities is a relationship between them. A tie can have a weight to represent its strength (such as the strength of a friendship or a collaboration). Ties between entities strengthen with repeated interactions, but they can also deteriorate in their absence [11], [55], [58]. There are many empirically plausible, domain-specific deterioration (i.e., 'decay') functions that one can use; examples include linear decay, power-law decay, and exponential decay [11], [55], [58], [83]. In the present paper, we use exponential decay, which is a common choice in many modeling frameworks (e.g., for the intensity decay function in a Hawkes process [46]). We restrict ourselves to modeling instantaneous interactions, but it is possible to generalize our tie-decay formalism to incorporate noninstantaneous interactions.

Consider a set of n interacting entities (i.e., nodes), and let $B(t)$ be the $n \times n$ time-dependent, real, nonnegative matrix whose entries $b_{ij}(t)$ encode the connection strengths between entities i and j at time t. To construct a continuous-time temporal network of these ties, we make two modeling assumptions about how ties evolve and how interactions strengthen them:

(1) In the absence of interactions, we assume that ties decay exponentially, as proposed by Jin *et al.* [37]. In mathematical terms, $b'_{ij} = -\alpha b_{ij}$ (where the prime represents differentiation with respect to time), so $b_{ij}(t) = b_{ij}(0)e^{-\alpha t}$ for some $\alpha > 0$ and an initial condition $b_{ij}(0)$.

(2) If two entities interact at time $t = \tau$, the strength of the tie between them grows instantaneously by 1, and it then decays as normal. This choice differs from [37], who reset the strength to 1 after each interaction.

Taken together, these assumptions imply that the temporal evolution of a tie satisfies the ordinary differential equation (ODE)

$$b'_{ij} = -\alpha b_{ij} + \delta(t - \tau)e^{-\alpha(t-\tau)} . \tag{1}$$

In equation (1), we represent an instantaneous interaction at $t = \tau$ as an impulse using the Dirac δ-function. If the tie has the resting initial condition $b_{ij}(0) = 0$, the solution to equation (1) is $b_{ij}(t) = H(t - \tau)e^{-\alpha(t-\tau)}$, where $H(t)$ is the Heaviside step function. This formulation is related to the one in Flores and Romance [19], who integrated over functions that represent the temporal evolution of interactions. A related notion of tie decay appears in the work of Sharan and Neville [72], although they considered homogeneous decay of all edges of a network in discrete time, instead of examining the decay of individual edges in continuous time. When there are multiple interactions between entities, we represent them as streams of impulses in an $n \times n$ matrix $\widetilde{A}(t)$ with entries $\tilde{a}_{ij}(t)$. If entity i interacts with entity j at times $\tau_{ij}^{(1)}, \tau_{ij}^{(2)}, \ldots$, then $\tilde{a}_{ij}(t) = \sum_k \delta(t - \tau_{ij}^{(k)})e^{-\alpha(t-\tau_{ij}^{(k)})}$. We rewrite equation (1) as

$$b'_{ij} = -\alpha b_{ij} + \tilde{a}_{ij} , \tag{2}$$

which has the solution $b_{ij}(t) = \sum_k H(t - \tau_{ij}^{(k)})e^{-\alpha(t-\tau_{ij}^{(k)})}$ from a resting[1] initial condition.

In practice—and, specifically, in data-driven applications—one can readily construct $B(t)$ by discretizing time so that there is at most one interaction during each time step of length Δt (e.g., as in a Poisson process). This type of time discretization is common in simulations of stochastic dynamical systems, such as when using Gillespie algorithms [16], [67], [80]. In our case, we let $A(t)$ be the $n \times n$ matrix in which an entry $a_{ij}(t) = 1$ if entity i interacts with entity j at time t and $a_{ij}(t) = 0$ otherwise. At each time step, $A(t)$ has at most one nonzero entry in a directed network (and at most two of them in an undirected network). Therefore,

$$B(t + \Delta t) = e^{-\alpha \Delta t}B(t) + A(t + \Delta t) . \tag{3}$$

Equivalently, if interactions between pairs of entities occur at times $\tau^{(\ell)}$ (it can be a different pair at different times) such that $0 \leq \tau^{(0)} < \tau^{(1)} < \cdots < \tau^{(T)}$, then at $t \geq \tau^{(T)}$, we have

$$B(t) = \sum_{k=0}^{T} e^{-\alpha(t-\tau^{(k)})} A(\tau^{(k)}) . \tag{4}$$

If there are no interactions at time t, then every entry of the matrix $A(t)$ is 0.

Our continuous-time approach avoids having to impose a hard partition of the interactions into bins (i.e., windows). However, one still needs to choose a value of the decay parameter α. Another benefit of our approach is that it eliminates the placement of the time windows as a potential source of bias [41]. When choosing a value of α, it is perhaps intuitive to think about the *half-life* $\eta_{1/2}$ of a tie, as it gives the amount of time that it takes for a tie to lose half of its strength in the absence of new interactions. Given $\alpha > 0$, the half-life of a tie is $\eta_{1/2} = \alpha^{-1} \ln 2$. Our choice of using α to downweight

[1] It is not unlike a Norwegian blue parrot. (Norwegian blues stun easily.)

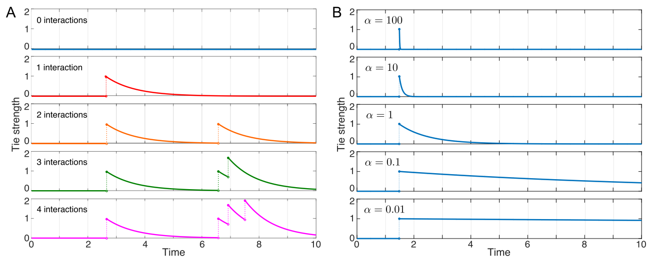

Fig. 1. **A:** Evolution of the strength of a tie with exponential decay [see equation (2)]. **B:** The decay rate α determines how fast the tie decays.

old activity is consistent with the choice of Grindrod and Higham [27], [28], who used an exponential decay factor to reduce the importance of old interactions in the context of dynamic communicability.

In Fig. 1, we illustrate the evolution of the strength of a tie in our tie-decay formalism. If entities i and j have never interacted up to and including time t_0, then $b_{ij}(t_0) = 0$. Suppose that they first interact at time $\tau^{(1)} > t_0$ (such that $a_{ij}(\tau^{(1)}) = 1$). Their tie strength then increases by 1, so $b_{ij}(\tau^{(1)}) = 1$. It subsequently decays exponentially until they interact again, so $b_{ij}(t > \tau^{(1)}) = e^{-\alpha(t-\tau^{(1)})}$ before their next interaction. If entities i and j next interact at time $\tau^{(2)} > \tau^{(1)}$, such that $a_{ij}(\tau^{(2)}) = 1$, their tie strength becomes $b_{ij}(\tau^{(2)}) = e^{-\alpha(\tau^{(2)}-\tau^{(1)})} + 1$. The tie strength between i and j continues to evolve in this way as more interactions occur.

III. Eigenvector-Based Centrality Scores in Tie-Decay Networks

One common question that arises frequently when analyzing networks in scientific and industrial applications is the following: What are the most important nodes? [29] To examine this question, researchers have developed numerous centrality scores to quantify the importance of nodes according to different criteria [59].

An important family of centrality scores arises from spectral properties of the adjacency matrix (or other matrices) of a network [8], [53], [66], [75]. One attractive feature of computing centrality scores using a spectral method is that one can exploit the full structure of a network. By contrast, degree centrality is a simple centrality score that relies only on a network's local structure. Eigenvector-based centrality scores have been insightful in numerous applications, and one can use efficient numerical algorithms to compute eigenvectors and singular vectors of matrices [25], [77]. Popular spectral centrality scores for directed networks include PageRank [23], [63] (which exploits the properties of a random walk on a network) and hub and authority scores [42] (which exploit both random-walk properties and the asymmetry of connections in directed networks).

In temporal networks, centrality scores must incorporate not only which nodes and edges are present in a network, but also when they are present [39], [64]. This makes it challenging to develop and analyze centrality measures in temporal networks. Some approaches have exploited numerical methods for dynamical systems to compute specific scores, such as a Katz centrality for temporal networks [6], [28], and others have employed aggregated or multilayer representations of temporal networks to calculate spectral centrality scores [3], [75]. However, these approaches have either been limited to a specific kind of centrality, or they have relied on the judicious aggregation of interactions into time bins. For example, an early paper [9] on centralities in temporal networks used a time bin of one day. Choosing an appropriate size for such a time bin is far from straightforward and requires deep knowledge of the system under study. Overly coarse bins obfuscate temporal features, whereas bins that are too small may obscure network structures, yielding scores that may result more from noise than from signals.

Our tie-decay network formalism in equation (2) allows us to employ efficient numerical techniques to compute a variety of spectral centrality scores in our temporal networks. One can tune the decay parameter α (which one can also generalize to be node-specific, tie-specific, or time-dependent) to consider different time scales. A key benefit of our approach is that we can easily incorporate both new interactions and new nodes as a network evolves. In the present paper, we showcase an application using PageRank centrality, but it is also worthwhile to study other spectral centrality scores using our tie-decay formalism.

A. Tie-Decay PageRank Centrality

PageRank centrality is a widely used (and historically important) eigenvector-based centrality score for time-independent networks [59], [63]. The PageRank score of a node in a network corresponds to its stationary distribution in a random walk with teleportation [23], [53]. In this type of random walk, a walker departs from a node by following an outgoing edge with probability $\lambda \in (0, 1)$ (where, typically, one chooses the edge with a probability that is proportional to its weight); it 'teleports' to some other node in the network with probability $1 - \lambda$. It is common to choose the destination node uniformly at random, but many other choices are possible [23], [44]. In the present paper, we employ uniform teleportation with $\lambda = 0.85$ (which is a common choice). Let B be the adjacency matrix of a weighted network with n nodes, so b_{ij} encodes the weight of a directed tie from node i to node j. The $n \times 1$ vector $\boldsymbol{\pi}$ of PageRank scores, with $\boldsymbol{\pi} > 0$ and $\|\boldsymbol{\pi}\|_1 = 1$, is the leading-eigenvector solution of the eigenvalue problem

$$G^T \boldsymbol{\pi} = \boldsymbol{\pi}\,, \tag{5}$$

where G is the $n \times n$ transition-rate matrix of a teleporting random walk:

$$\begin{aligned} G &= \lambda\big(D^\dagger B + \mathbf{c}\mathbf{v}^T\big) + (1-\lambda)\mathbb{1}\mathbf{v}^T \\ &= \lambda P + (1-\lambda)\mathbb{1}\mathbf{v}^T\,, \end{aligned} \tag{6}$$

where $P = D^\dagger B + \mathbf{c}\mathbf{v}^T$; the matrix D is the diagonal matrix of weighted out-degrees, so $d_{ii} = \sum_k b_{ik}$ and $d_{ij} = 0$ when $i \neq j$; and $D^\dagger$ is its Moore–Penrose pseudo-inverse. The $n \times 1$ vector $\mathbf{c}$ is an indicator of 'dangling nodes' (i.e., nodes with 0 out-degree): $c_i = 1 - d^{\dagger}_{ii} \sum_k b_{ik}$, so $c_i = 1$ if the out-degree of i is 0 and $c_i = 0$ otherwise. Additionally, $\mathbb{1}$ is the $n \times 1$ vector of 1s, and the $n \times 1$ distribution vector $\mathbf{v}$ encodes the probabilities of each node to receive a teleported walker. In the present paper, we use $v_i = 1/n$ for all i.

The perturbations to $D^{\dagger}B$ from $\mathbf{v}$ and $\mathbf{c}$ ensure the ergodicity of the teleporting random walk, so the Perron–Frobenius theorem guarantees that G^T has a unique right leading eigenvector $\boldsymbol{\pi}$ whose entries are all strictly positive. To calculate $\boldsymbol{\pi}$, one can perform a power iteration on G^T [77], but in practice we do not need to explicitly construct G^T. The iteration

$$\boldsymbol{\pi}^{(k+1)} = \lambda P^T \boldsymbol{\pi}^{(k)} + (1-\lambda)\mathbf{v}\,, \tag{7}$$

with $\boldsymbol{\pi}^{(0)} = \mathbf{0}$ or $\boldsymbol{\pi}^{(0)} = \mathbf{v}$, converges to $\boldsymbol{\pi}$ and preserves the sparsity of P. Equation (7), which ensures that computations are efficient, is equivalent to a power iteration [23].

To compute time-dependent PageRank scores from the tie-strength matrix $B(t)$, we define the temporal transition matrix

$$P(t) = D^{\dagger}(t)B(t) + \mathbf{c}(t)\mathbf{v}^T\,, \tag{8}$$

where $D(t)$ is the diagonal matrix of weighted out-degrees (i.e., the row sums of $B(t)$) at time t. The rank-1 correction $\mathbf{c}(t)\mathbf{v}^T$ depends on time, because the set of dangling nodes can change in time (though $\mathbf{v}$ remains fixed).[2] The iteration to obtain the time-dependent vector of PageRank scores $\boldsymbol{\pi}(t)$ is now given by

$$\boldsymbol{\pi}^{(k+1)}(t) = \lambda P^T(t) \boldsymbol{\pi}^{(k)}(t) + (1-\lambda)\mathbf{v}\,. \tag{9}$$

To understand the temporal evolution of $\boldsymbol{\pi}(t)$, we begin by establishing some properties of the temporal transition matrix $P(t)$ in the following lemma.

Lemma 1: When there are no new interactions between times t and $t+\Delta t$, the entries of $A(t+\Delta t)$ are all 0 and $P(t+\Delta t) = P(t)$. If there is a single new interaction from node i to node j, such that $a_{ij}(t+\Delta t) = 1$, then $P(t+\Delta t) = P(t) + \Delta P$, where

$$\begin{aligned}\Delta P = &\frac{1}{1+e^{-\alpha\Delta t}d_{ii}(t)}\mathbf{e}_i\mathbf{e}_j^T \\ &- \frac{1}{d_{ii}(t)(1+e^{-\alpha\Delta t}d_{ii}(t))}\mathbf{e}_i\mathbf{e}_i^T B(t) - c_i(t)v_i\mathbf{e}_i\mathbb{1}^T\end{aligned} \tag{10}$$

and $\mathbf{e}_i$ and $\mathbf{e}_j$, respectively, are the i-th and j-th canonical vectors.

The first term in the right-hand side of equation (10) is a matrix whose only nonzero entry is the (i,j)-th term, the second term is a rescaling of the i-th row of $B(t)$, and the third term is the perturbation due to teleportation. An important implication of Lemma 1 is that the PageRank scores do not change when there are no new interactions, so $\boldsymbol{\pi}(t+\Delta t) = \boldsymbol{\pi}(t)$. If nodes or ties have different decay rates (so that now we index α as α_i or α_{ij}), then this no longer has to be the case.

When there are new interactions, the following result sets an upper bound on how much the PageRank scores can change.

Theorem 1: Suppose that there is a single interaction between times t and $t+\Delta t$ from node i to node j, such that the change ΔP in the transition matrix satisfies equation (10). It follows that

$$\begin{aligned}&\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \\ &\quad\le \frac{2\lambda}{1-\lambda}\min\left\{\pi_i(t),\, \frac{1}{1+e^{-\alpha\Delta t}d_{ii}(t)} - \frac{c_i(t)}{2}\right\}.\end{aligned} \tag{11}$$

We present two corollaries of Theorem 1.

Corollary 1: If i is a dangling node at time t, then

$$\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \le \frac{2\lambda}{1-\lambda}\min\left\{\pi_i(t), \frac{1}{2}\right\}. \tag{12}$$

Corollary 2: If node i has one or more outgoing edges at time t, then

$$\begin{aligned}&\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \\ &\quad\le \frac{2\lambda}{1-\lambda}\min\left\{\pi_i(t),\, \frac{1}{1+e^{-\alpha\Delta t}d_{ii}(t)}\right\}.\end{aligned} \tag{13}$$

We prove Lemma 1 in Appendix A, and we prove Theorem 1 and Corollaries 1 and 2 in Appendix B.

B Temporal Iteration of PageRank Scores

To calculate the PageRank scores at time $t+\Delta t$, we use the iteration in equation (9) to update the PageRank vector using $\boldsymbol{\pi}(t)$ as the initial value. That is,

$$\boldsymbol{\pi}^{(0)}(t+\Delta t) = \boldsymbol{\pi}(t)\,. \tag{14}$$

The relative error of the computed PageRank vector at iteration k is

$$\left\|e^{(k)}_{\mathrm{rel}}\right\|_1 = \left\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}^{(k)}(t+\Delta t)\right\|_1. \tag{15}$$

A result from [7] (see their Theorem 6.1) implies that $\|e^{(k)}_{\mathrm{rel}}\|_1 \le \lambda^k \|e^{(0)}_{\mathrm{rel}}\|_1$. The relation $\boldsymbol{\pi}^{(0)}(t+\Delta t) = \boldsymbol{\pi}(t)$ and Theorem 1 then imply that

$$\begin{aligned}\left\|e^{(k)}_{\mathrm{rel}}\right\|_1 &\le \lambda^k\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \\ &\le \frac{2\lambda^{k+1}}{1-\lambda}\min\left\{\pi_i(t),\, \frac{1}{1+e^{-\alpha\Delta t}d_{ii}(t)} - \frac{c_i(t)}{2}\right\}.\end{aligned} \tag{16}$$

Therefore, we can select an error tolerance ϵ such that $\|e^{(k^*)}_{\mathrm{rel}}\|_1 \le \epsilon$ for some number k^* of iterations. The value k^* is the maximum number of iterations that we need for the relative error to be at most ϵ. We compute k^* by calculating

[2] Strictly speaking, once a node leaves a dangling-node set, it never returns. Although the tie strength decays exponentially, it never quite reaches 0. In practice, one can opt to remove ties with $b_{ij} \ll 1$ to preserve the sparsity of $B(t)$. When one does this, nodes can return to the dangling-node set.

$$k^* = \frac{\ln(\epsilon) - \ln(2) + \ln(1-\lambda) - \ln\left(\min\left\{\pi_i(t)\,,\, \frac{1}{1+e^{-\alpha\Delta t} d_{ii}(t)} - \frac{c_i(t)}{2}\right\}\right)}{\ln(\lambda)} - 1\,. \tag{17}$$

In practice, we can instead track the residual after k iterations [23]:

$$\begin{aligned}\mathbf{r}^{(k)}(t+\Delta t) &= \boldsymbol{\pi}^{(k+1)}(t+\Delta t) - \boldsymbol{\pi}^{(k)}(t+\Delta t)\\ &= (1-\lambda)\mathbf{v} - \left(I_n - \lambda P^T(t+\Delta t)\right)\boldsymbol{\pi}^{(k)}(t+\Delta t)\,.\end{aligned} \tag{18}$$

We use this residual to bound the relative error,

$$\begin{aligned}\left\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}^{(k)}(t+\Delta t)\right\|_1 &= \left\|\left(I_n - \lambda P^T(t+\Delta t)\right)^{-1}\mathbf{r}^{(k)}(t+\Delta t)\right\|_1\\ &\leq \frac{1}{1-\lambda}\left\|\mathbf{r}^{(k)}(t+\Delta t)\right\|_1\,,\end{aligned} \tag{19}$$

where I_n denotes the $n \times n$ identity matrix. We thereby monitor the convergence of the iterative process.

In our experiments, we always obtain $\|\boldsymbol{\pi}^{(k+1)}(t+\Delta t) - \boldsymbol{\pi}^{(k)}(t+\Delta t)\|_{l_1} < 10^{-6}$ in two iterations or fewer (see Section V-C). See [14] for an alternative approach for approximating PageRank after adding a single edge to a network.

IV. A Synthetic Example

To illustrate some of the features of our tie-decay formalism, we construct an example that illustrates the challenges of binning interactions and how our approach can help overcome them.

Consider a network with node set $V = \{1, 2, 3, 4, 5\}$. These nodes interact with each other cyclically in discrete time, with only a single edge present during each time step. At the initial time t_0, we select a 'target node' and a 'source node' and create a directed edge from the source to the target. At the next time (i.e., $t = t_1$), there is an edge from a different source node to the same target. This occurs with a third source node and the same target at time t_2 and with the final remaining source node and this target at time t_3. We have now exhausted all possible source nodes for this target, so at time t_4, we select a new target node and repeat the above process. Specifically, there is exactly one directed edge at each discrete time, and we select each of the four possible source nodes exactly once between times t_4 and t_7. We repeat this cycle thrice more, and we finish with the final source–target pair at time t_{19}. The process begins again at time t_{20} with the first target node, and it continues periodically. In Fig. 2, we show the interactions in this synthetic temporal network for the first 24 time steps. We expect the cyclical behavior of this network to be reflected in the temporal centralities of its nodes.

One way to examine temporal centralities of the nodes in our synthetic network is by aggregating the interactions into sliding windows of duration w [65], [70]. To do this, we construct time-independent networks with interactions in the time window $(t_k - w, t_k]$, and we then compute the PageRank scores of the nodes for each of these networks. In Fig. 3 A, we

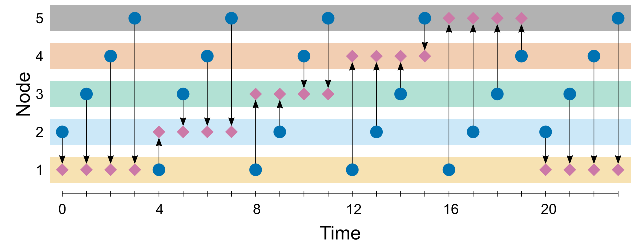

Fig. 2. Schematic illustration of the interactions in our five-node synthetic network. (See the description in the main text.) We show target nodes as pink diamonds and source nodes as blue disks. In this example, there is one target at a time, and the different nodes take turns being the target. After all source nodes interact with the current target, a new node becomes the target, and all source nodes then interact with it. This cycle keeps repeating in this temporal network, which is periodic with a period of 20 time steps.

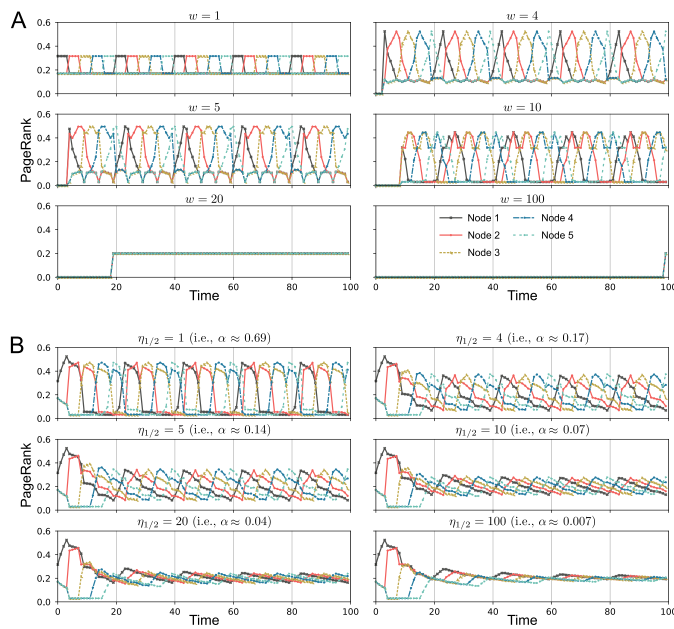

Fig. 3. PageRank scores over time for each node in our five-node synthetic network. (See the description in the main text.) **A:** Calculations of time-independent PageRank scores using sliding windows $(t_k - w, t_k]$ of different duration to aggregate interactions. **B:** Calculations of tie-decay PageRank scores for different values of the tie half-life.

show the PageRank scores for these networks using windows of different lengths. To use this sliding-window approach, we need at least w time steps before we can construct the first window; this is a potential issue for some studies. Additionally, the PageRank scores are sensitive to the value of w. This is another potential difficulty, especially because one does not know appropriate window lengths in advance in many situations. An alternative to aggregating interactions into sliding windows is to use adjacent windows of length w [10] With this choice, we aggregate interactions exactly once every w time steps. We show the resulting time series of PageRank scores in Fig. 10 in Appendix C.

When a window includes individual interactions only ($w = 1$) or matches the length of one cycle ($w = 4$), the PageRank centrality scores capture the sequence of interactions between the nodes. We show these two cases in the top row of

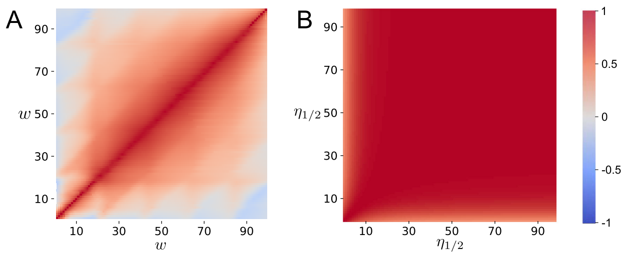

Fig. 4. **A:** Pearson correlation matrix of the PageRank time series of node 1 in our synthetic network using sliding windows $(t_k - w, t_k]$ for $w \in \{1, \ldots, 100\}$. **B:** Pearson correlation matrix of the tie-decay PageRank time series of node 1 for integer half-lives $\eta_{1/2} \in \{1, \ldots, 100\}$.

Fig. 3 A. However, these choices of w are not without issues. Using $w = 1$ neglects important temporal features, and selecting $w = 4$ (the cycle length) and placing it correctly requires prior knowledge of the temporal network, which we possess only because we invented the rules to create this network. The extreme sensitivity of PageRank scores to the value of w also causes other problems. For example, when w is a multiple of the number of nodes and the period (e.g., $w = 20$), this approach fails to give any useful information at all, because all differences in the interaction patterns are 'masked' entirely by the window.

We now use our tie-decay formalism to obtain PageRank time series for several tie half-lives (see Fig. 3 B). The cyclical dynamics of the interactions manifest for all of these values of the half-life $\eta_{1/2}$. Different values of $\eta_{1/2}$ entail different oscillation magnitudes, but the time series have oscillations for all of our choices of $\eta_{1/2}$. Every four time steps, we observe an increase in the PageRank score of the newly-selected target node. Our tie-decay framework also does not have the masking problem that we illustrated above.

We also examine the sensitivity with respect to the choice of window length and half-life by computing the Pearson correlation of the PageRank time series for different values of w and τ (for ordinary PageRank) and $\eta_{1/2}$ (for tie-decay PageRank). In Fig. 4 A, we show the Pearson correlation between the time series of the (ordinary) PageRank centrality of node 1 for each $w \in \{1, 2, \ldots, 100\}$. As the figure shows, binning interactions over time produces a time series that is very sensitive to the choice of w. The mean correlation between time series is 0.414, and the standard deviation is 0.301. Some time series are even correlated negatively with each other. The correlations also depend on the periodicity of our temporal network. For example, choices of w for the two time series that differ by exactly 20 have larger positive correlations with each other than with other window lengths. By comparison, the Pearson correlations between the PageRank time series in the tie-decay networks for a variety of values of the half-life $\eta_{1/2}$ (see Fig. 4 B) paint a very different picture. We observe large positive correlations for many pairs of values of τ. The mean correlation between the time series is 0.945, which is much larger than what we obtain for our calculation with sliding windows; its standard deviation is 0.109, which is much smaller than that for sliding windows.

Our synthetic example of a temporal network highlights some of the advantages of our continuous-time tie-decay network formalism for temporal networks. Results from frameworks that require binning interactions can be extremely sensitive both to the choice of window length and to the placement of windows. By contrast, our tie-decay framework is more robust to parameter choices. Varying $\eta_{1/2}$ (or, equivalently, varying α) adjusts the longevity of ties while maintaining similar PageRank centrality trajectories, lending confidence to investigations even when (as is usually the case) one does not possess precise knowledge of the time scales of the interactions in a system.

V. The National Health Service (NHS) Retweet Network

We now compute tie-decay PageRank scores to track the evolution of node importances over time in a large data set of time-annotated interactions on Twitter.

Twitter is a social-media platform that has become a prominent channel for organizations, individuals, 'bots', 'sockpuppets', and other types of accounts to broadcast events, share ideas, report events, and socialize by posting messages (i.e., 'tweets') of at most 140 characters in length [43].[3] Data from Twitter has allowed researchers to study patterns and trends in a plethora of large-scale political and social phenomena, such as protests and civil unrest, discussions of public health, and information propagation [2], [3], [15], [21], [26], [56], [62], [76].

Twitter accounts (which can represent an individual, an organization, a bot, and so on) can interact in several ways, and there are various ways to encode such interactions in the form of a network. For example, accounts can subscribe to receive other accounts' tweets (a 'follow' connection), can mention each other in a tweet (a 'mention' connection), can spread a tweet that was posted originally by someone else (a 'retweet' connection) to their followers, and so on. These interactions represent an explicit declaration of interest from a source account about a target, and one can thus encode them using directed networks [2]. Because these interactions are time-resolved, it is sensible to analyze Twitter networks as time-dependent networks [3].

We study a retweet network, which we construct from a data set of tweets about the United Kingdom's National Health Service (NHS) that were posted after the controversial Health and Social Care Act of 2012 [48]. Our data set covers over five months of time and consists of tweets in English that include the term 'NHS'. Specifically, we consider retweets[4] by the 10 000 most-active Twitter accounts (according to the number of tweets in our data set) from 5 March 2012 through 21 August 2012 (see Fig. 5). See Appendix D for a discussion of some basic statistical properties of the NHS retweet network. All data were collected by Sinnia,[5] a data-analytics company, using

[3] Twitter subsequently expanded the maximum tweet length to 280 characters, but the maximum was 140 characters at the time that our data set was collected.

[4] We consider only retweets, so if MAP tweets something and MBD retweets it, then only the retweet can be part of our data set.

[5] See http://www.sinnia.com/.

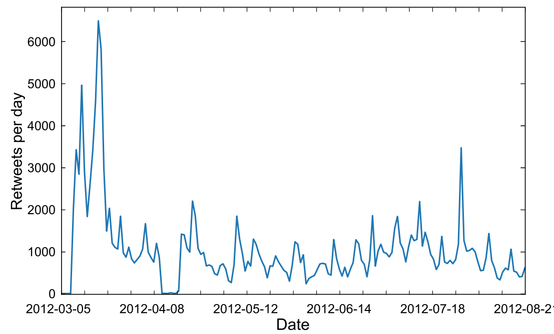

Fig. 5. Number of daily retweets among the 10 000 most-active Twitter accounts in the NHS data set.

Twitter Gnip PowerTrack API.[6] From these data, we construct a tie-decay temporal network in which the interactions are retweets.[7] Our code for constructing tie-decay networks is available at https://bitbucket.org/walid0925/tiedecay/.

A. Tie-Decay PageRank Centrality in the NHS Retweet Network

The temporal network of retweets from the NHS data has the tie-strength matrix $\mathbf{B}(t)$ (see equation (3)) and starts from the initial condition $\mathbf{B}(0) = 0$. We construct three tie-decay networks by using values of α that correspond to tie half-lives of 1 hour, 1 day, and 1 week. We compute tie-decay PageRank scores of all Twitter accounts for each of these networks.

In Fig. 6, we illustrate the effect of the value of α on PageRank scores. We compute tie-decay PageRank scores for networks in which the tie half-life is 1 hour, 1 day, and 1 week. In each panel of Fig. 6, we plot the Twitter account with the largest PageRank score at every time point; the transitions between white and gray shading indicate when there is a change in which account has the top score. When the half-life is short (i.e., α is large), interactions produce feeble ties that fade quickly unless there are frequent and sustained interactions between the accounts. Consequently, the tie-decay PageRank scores of the Twitter accounts change wildly in time, and (as illustrated in the top panel of Fig. 6) such a short half-life implies that the Twitter account with the top tie-decay PageRank score changes frequently. When the half-life is longer (e.g., 1 day), ties are better able to gain traction by strengthening from interactions that otherwise would be too far apart in time. Consequently, there are fewer changes in which Twitter account has the top PageRank score. The middle panel of Fig. 6 illustrates that we indeed observe fewer changes in the top-ranked account when the half-life is 1 day. Finally, when the half-life of a tie is 1 week, two specific accounts (@DrEoinClarke and @marcuschown) dominate the ranking; they alternate between the top and second spots.

In this case study, we reveal two accounts as dominant ones as we tune the half-life of ties to larger values. The first of these, Eoin Clarke (@DrEoinClarke), is a Labour-party activist and

[6] See https://gnip.com/realtime/powertrack/.

[7] In our computations, we set $b_{ij} < 10^{-7}$ to 0 to preserve the sparsity of $B(t)$. Consequently, nodes may rejoin the dangling-node set.

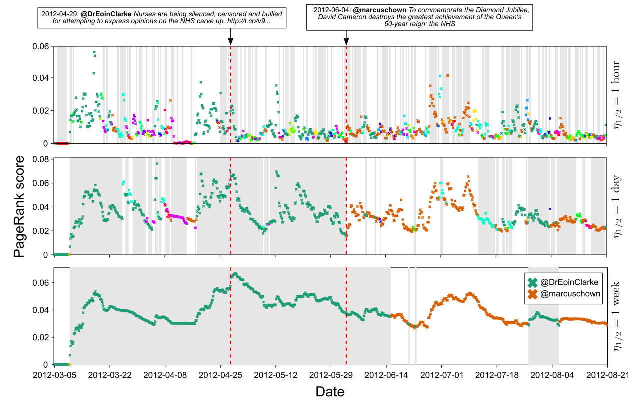

Fig. 6. Top Twitter accounts, according to tie-decay PageRank, in the temporal NHS retweet network with three values of tie half-life. Each color is associated with a unique Twitter account, and the alternating gray and white background color highlights intervals in which the same account has the top tie-decay PageRank score. Transitions in color (from white to gray, and vice versa) indicate when there is a change in which account has the top score. The first dashed red vertical line indicates the time that @DrEoinClarke posted the tweet in the associated box; the second such line corresponds to a tweet by @marcuschown.

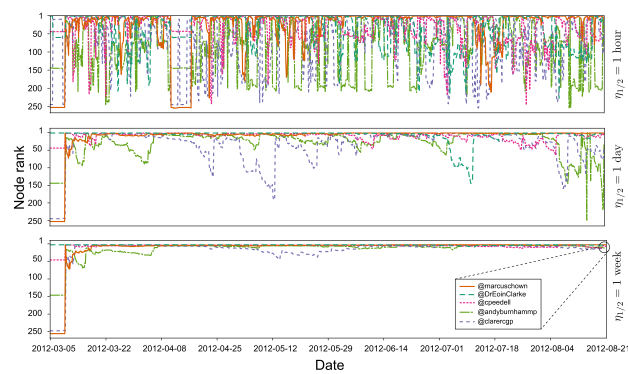

Fig. 7. The effect of the half-life value on the time-resolved PageRank rankings of five prominent Twitter accounts in the NHS retweet network. We show time series for three different values of the half-life. More-important accounts are higher on the vertical axis.

was an outspoken critic of the U.K. coalition government's stance in 2012 on the NHS. Marcus Chown (@marcuschown) is a science writer, journalist, and broadcaster who was also an outspoken critic of the U.K. government's NHS policy in 2012. The dominance of these two accounts becomes apparent as we increase the half-life of the ties. On 29 April, @DrEoinClarke posted a tweet that gathered significant attention. This tweet yields a short-lived boost in his PageRank score when $\eta_{1/2}$ is 1 hour and 1 day, and it yields a more sustained increase when $\eta_{1/2}$ is 1 week. On 4 June, @marcuschown posted a tweet that boosts his PageRank score. When $\eta_{1/2}$ is 1 day, the number of retweets of this tweet are enough to carry him to the top spot. However, when $\eta_{1/2}$ is 1 week, the retweets are not enough to overtake @DrEoinClarke, whose ties remain strong.

In Fig. 7, we show a complementary illustration of the effect of half-life value on the time-resolved PageRank rankings of Twitter accounts. We construct a time-independent network in which we aggregate all of the interactions in our data set—specifically, we consider $B(t)$ for $\alpha = 0$ with the time t set to be 21 August 2012—and we determine the top-5

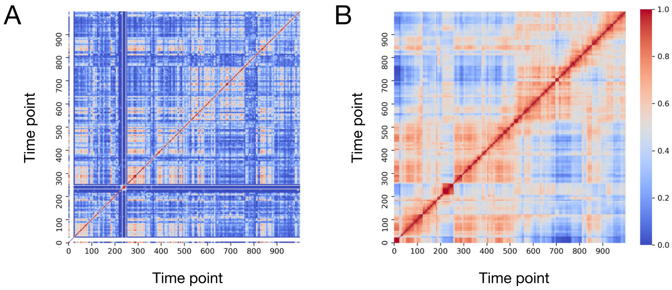

Fig. 8. **A:** Pearson correlations between PageRank vectors in networks that we construct using sliding windows of $w = 1$ day. **B:** Pearson correlations between PageRank vectors in tie-decay networks with $\eta_{1/2} = 1$ day.

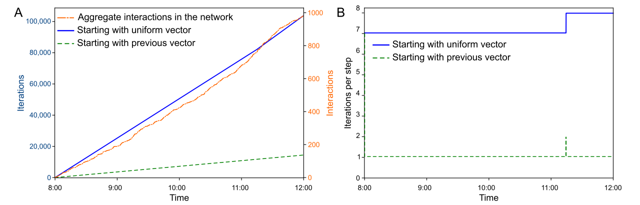

Fig. 9. **A:** Cumulative number of iterations to convergence (where we define 'convergence' as $\|\boldsymbol{\pi}^{(k+1)}(t+\Delta t) - \boldsymbol{\pi}^{(k)}(t+\Delta t)\|_{l_1} < 10^{-6}$) for calculating tie-decay PageRank on the NHS retweet network when the starting vector is the uniform vector (solid blue curve) and the previous PageRank vector (dashed green curve). For context, we also include the aggregate number of interactions in the network (dash–dotted orange curve) and label it on the right vertical axis. **B:** When analyzing the NHS retweet network, we observe that tie-decay PageRank requires at most 2 iterations to converge when we start from the previous time step's PageRank vector, whereas using the uniform vector necessitates 7 or more iterations. In this example, the half-life of a tie is 1 day.

accounts by calculating the standard time-independent version of PageRank (with $\lambda = 0.85$) on this network. We then track the time-dependent PageRank ranks (where rank 1 is the Twitter account with the largest tie-decay PageRank score, and so on) of these five accounts for different values of α (or, equivalently, of $\eta_{1/2}$). When $\eta_{1/2}$ is 1 hour, these accounts often overtake each other in the rankings and the changes in rankings can be rather drastic, as some Twitter accounts drop or rise by almost 250 spots. As we consider progressively longer half-lives, we observe less volatility in the rankings.

The experiments in Figs. 6 and 7 demonstrate how one can use α as a tuning parameter to encode the longevity of relationship values in a temporal network. They also demonstrate the value of our tie-decay formalism for illustrating fluctuations in network structures. When analyzing networks in discrete time, there is a risk that aggregating interactions may conceal important dynamics and nuances of network structure [18]. By contrast, our continuous-time network formalism avoids arbitrary cutoff choices (and potential ensuing biases [41]) when choosing the borders of time windows. It also allows a smoother exploration of network structure at a level of temporal granularity that is encoded in the value of α.

B. *Aggregating Interactions Versus Examining Tie-Decay Networks*

Using tie-decay networks instead of aggregating interactions via sliding windows can have a large impact on the qualitative results of an investigation. To illustrate this using the NHS data, we compare the Pearson correlations between PageRank vectors in the tie-decay networks to the Pearson correlations between PageRank vectors in networks that we construct using sliding windows. We sample K equally-spaced time points between the first and last interactions in the data.[8] At each time point t_k, we construct a time-independent network with interactions in the time window $(t_k - w, t_k]$ for a given window length w, and we compute the PageRank vector for this network. In Fig. 8 A, we show the Pearson correlation matrix between PageRank vectors using $K = 1000$ and $w = 1$ day. The sampled time points are approximately four hours apart, so consecutive time points have overlapping windows. However, the Pearson correlations between the PageRank vectors from the different time-independent networks are relatively small. This is usually the case even for the PageRank vectors from time points that are near each other. Such small temporal correlations are not surprising; previous research [9] has reported that the connections in networks that one constructs by aggregating interactions on different days can vary drastically. When this occurs, the most central nodes in networks from different days can also differ drastically.

We also calculate the tie-decay PageRank vectors using the same time points. In this case, we observe a pronounced block structure in the correlation matrix (see Fig. 8 B). That is, PageRank vectors from time points that are close to each other have large correlations with each other because tie strengths of past interactions persist in time. We can adjust the extent of such persistence by varying the half-life. In Fig. 12 in Appendix E, we show the distributions of pairwise correlations between PageRank vectors using both sliding windows and tie-decay networks.

C. *Computational Efficiency*

In applications (e.g., for streaming data), it is often desirable to update the values of time-dependent centrality measures, such as tie-decay PageRank scores, when there is a new interaction. We know from Theorem 1 that there is a bound on the magnitude of the difference between the PageRank vectors at times t and $t + \Delta t$ when there is a new interaction. Consequently, we expect the iterative procedure for computing PageRank at $t + \Delta t$ to converge faster when we use the PageRank vector from t as our initial vector than if we start the computation from scratch. To demonstrate this, we select a period of time—between 8:00 am and 12:00 pm on 18 March 2012—with high activity in our NHS data and calculate the tie-decay PageRank vector at time $t + \Delta t$ with two different starting vectors: the uniform vector $\boldsymbol{\pi}^{(0)}(t+\Delta t) = \frac{1}{n}\mathbb{1}$ and the previous PageRank vector $\boldsymbol{\pi}^{(0)}(t+\Delta t) = \boldsymbol{\pi}(t)$. In time-independent networks, it has been observed that the uniform vector has better convergence properties than any other starting vector in the absence of prior knowledge about the final

[8] The first interaction occurs at 20:41:46 on 5 March 2012, and the last interaction occurs at 09:09:25 on 21 August 2012.

PageRank vector [23]. In our tie-decay network formalism, given the bound between the magnitudes of the PageRank vectors at times t and $t+\Delta t$, it is intuitive that using the vector from the previous time has computational advantages over other choices. We demonstrate this fact in Fig. 9.

VI. Conclusions and Discussion

We have introduced a continuous-time framework that incorporates tie decay for studying temporal networks, and we used our formalism—which we call 'tie-decay networks'—to generalize PageRank centrality to continuous time.

In our proposed tie-decay formalism, a tie between two nodes strengthens through repeated interactions and it decays in their absence. Such tie-decay networks allow one to tractably analyze time-dependent interactions without having to aggregate interactions into time windows (i.e., bins), as is typically done in existing frameworks for studying temporal networks [32], [34]. We purposely avoided aggregating interactions using time bins, whose sizes and placement are difficult to determine, by modeling the weakening of ties in time as exponential decay with rate α (or, equivalently, with a half-life of $\eta_{1/2}$). In addition to representing the decay of human relationships [11], as we have done in this paper, it is also possible to use our formalism more generally to model the decreasing value of old information, decay in other types of interactions, and so on.

We showcased our tie-decay formalism on both a synthetic temporal network and a network of retweets on Twitter. Our computations illustrated that adjusting the value of the half-life $\eta_{1/2}$ allows one to examine the temporal dynamics of rankings at different time scales of interest. Our study of a synthetic network illustrated that using tie-decay networks can help mitigate serious issues, such as sensitivity to time scales and the masking of interaction patterns, that can arise from binning interactions. To provide a case study for the study of tie-decay PageRank centralities in an empirical temporal network, we investigated the temporal evolution of the ranks of important accounts in a large collection of retweets about the United Kingdom's National Health Service. We also developed a numerical scheme and bounds on the change of tie-decay PageRank scores upon the arrival of each new interaction. Such bounds are important for studying data streams, in which new data arrives at a potentially alarming rate. By tuning the decay rate of interactions, we illustrated that tie-decay PageRank scores can change much more drastically when the half-life is short than when the half-life is long.

Our tie-decay formalism for continuous-time changes in network architecture provides an important step for the study of streaming network data and the development of tools to analyze temporal networks in real time. Streaming data is ubiquitous—it arises in social-media data, sensor streams, communication networks, and more [20], [45]—and analyzing tie-decay networks offers a promising approach for studying it. In the present paper, we illustrated how to perform an update of tie-decay PageRank from one time step to another, and it will be important to develop such ideas further for other types of computations (such as community detection and other types of clustering). In the short term, it will be useful to implement efficient schemes for numerical computations of tie-decay generalizations of other centralities scores (such as hubs and authorities) that are obtained from eigenvectors. For instance, our framework permits a tie-decay generalization of personalized PageRank [23], [36] (by making a different choice of $\mathbf{v}$ in equation (9)), which one can use in turn to develop new, principled methods for studying local community structure in networks that evolve in continuous time. Our tie-decay formalism is also very well-suited to incorporating time-dependent strategies for teleportation in PageRank [24]. An extension of our tie-decay formalism to noninstantaneous interactions (i.e., taking durations into account) is possible by replacing the term with the Dirac δ-function in equation (1) with a function that is nonzero only when a tie exists. For example, if an interaction lasts from τ_{begin} until τ_{end}, then one can use the function $H(t-\tau_{\text{begin}})-H(t-\tau_{\text{end}})$, where H is the Heaviside step function. One can also formulate interactions with time durations using window functions [82] or test functions [35].

A wealth of other avenues are also worth pursuing. For example, it is desirable to systematically investigate heterogeneous decay rates (e.g., individual rates for nodes or ties), fit the decay parameter to data, use decay functions other than exponential ones [22], [84],[9] develop clustering methods for tie-decay networks, analyze localization phenomena (and their impact on centralities and clustering) [51], [75], develop and study random-network null models for tie-decay networks, incorporate noninstantaneous interactions, investigate change-point detection, and examine continuous-time networks with multiplex interactions. It is also desirable to study a variety of dynamical processes—such as contagions and opinion dynamics—on tie-decay networks.

Many empirical networks and data sets from which one can construct networks are time-dependent, and it is important to be able to model such systems in continuous time. Tie-decay networks offer a promising approach for further development of continuous-time temporal networks.

Appendix

A. Proof of Lemma 1

Proof: When there is a new interaction from node i to node j, the rows of $P(t+\Delta t)$ that correspond to nodes $k \neq i$ (i.e., the nodes from which the new connection does not originate) are unchanged: $p_{kh}(t+\Delta t) = p_{kh}(t)$ for all $h \in \{1, \ldots, n\}$.

To determine the change in the i-th row of $P(t+\Delta t)$, we first examine the i-th row of $B(t+\Delta t)$ by calculating

$$b_{ih}(t+\Delta t) = \begin{cases} e^{-\alpha \Delta t} b_{ih}(t)\,, & h \neq j \\ e^{-\alpha \Delta t} b_{ij}(t) + 1\,, & h = j\,. \end{cases} \quad (20)$$

[9] It may be particularly interesting to explore the effects of heavy tails in decay rates. As explained in [54], one can express a power law as a mixture of exponentials; this facilitates the incorporation of heavy tails into our tie-decay formalism.

We then consider the change to the rank-1 correction $\mathbf{c}(t)\mathbf{v}^T$ of equation (8). If i is not a dangling node at time t, then $\mathbf{c}(t+\Delta t) = \mathbf{c}(t)$. However, if i is a dangling node at time t, then $c_i(t) = 1$ and $c_i(t+\Delta t) = 0$. Therefore,

$$c_i(t+\Delta t) - c_i(t) = \begin{cases} 0\,, & c_i(t) = 0 \\ -1\,, & c_i(t) = 1\,. \end{cases} \tag{21}$$

Observe that $c_i(t+\Delta t)$ necessarily equals 0 and that $c_i(t) \in \{0,1\}$. Therefore,

$$c_i(t+\Delta t) - c_i(t) = -c_i(t)\,, \tag{22}$$

so the change to the correction term is

$$c_i(t+\Delta t)v_i - c_i(t)v_i = -c_i(t)v_i\,. \tag{23}$$

The i-th row of $P(t+\Delta t)$ is

$$p_{ih}(t+\Delta t) = \begin{cases} \frac{e^{-\alpha\Delta t} b_{ih}(t)}{1+e^{-\alpha\Delta t}\sum_k b_{ik}(t)} - c_i(t)v_i\,, & h \neq j \\ \frac{e^{-\alpha\Delta t} b_{ij}(t)+1}{1+e^{-\alpha\Delta t}\sum_k b_{ik}(t)} - c_i(t)v_i\,, & h = j\,. \end{cases} \tag{24}$$

For $h \neq j$, the difference between $p_{ih}(t+\Delta t)$ and $p_{ih}(t)$ is

$$\begin{aligned} p_{ih}(t+\Delta t) - p_{ih}(t) &= \frac{e^{-\alpha\Delta t} b_{ih}(t)}{1+e^{-\alpha\Delta t}\sum_k b_{ik}(t)} - \frac{b_{ih}(t)}{\sum_k b_{ik}(t)} - c_i(t)v_i \\ &= \frac{-b_{ih}(t)}{\sum_k b_{ik}(t)\left[1+e^{-\alpha\Delta t}\sum_k b_{ik}(t)\right]} - c_i(t)v_i \\ &= \frac{-b_{ih}(t)}{d_{ii}(t)[1+e^{-\alpha\Delta t} d_{ii}(t)]} - c_i(t)v_i\,. \end{aligned} \tag{25}$$

When $h = j$, we have

$$\begin{aligned} p_{ij}(t+\Delta t) - p_{ij}(t) &= \frac{1+e^{-\alpha\Delta t} b_{ij}(t)}{1+e^{-\alpha\Delta t}\sum_k b_{ik}(t)} - \frac{b_{ij}(t)}{\sum_k b_{ik}(t)} - c_i(t)v_i \\ &= \frac{\sum_k b_{ik}(t) - b_{ij}(t)}{\sum_k b_{ik}(t)\left[1+e^{-\alpha\Delta t}\sum_k b_{ik}(t)\right]} - c_i(t)v_i \\ &= \frac{d_{ii}(t) - b_{ij}(t)}{d_{ii}(t)[1+e^{-\alpha\Delta t} d_{ii}(t)]} - c_i(t)v_i \\ &= \frac{1}{1+e^{-\alpha\Delta t} d_{ii}(t)} - \frac{b_{ij}(t)}{d_{ii}(t)[1+e^{-\alpha\Delta t} d_{ii}(t)]} - c_i(t)v_i\,. \end{aligned} \tag{26}$$

In matrix terms, the change from $P(t)$ to $P(t+\Delta t)$ is thus

$$\begin{aligned} \Delta P = &\frac{1}{1+e^{-\alpha\Delta t} d_{ii}(t)}\mathbf{e}_i\mathbf{e}_j^T \\ &- \frac{1}{d_{ii}(t)(1+e^{-\alpha\Delta t} d_{ii}(t))}\mathbf{e}_i\mathbf{e}_i^T B(t) - c_i(t)v_i\mathbf{e}_i\mathbb{1}^T\,, \end{aligned} \tag{27}$$

which concludes the proof. ∎

B. Proof of Theorem 1

Proof: The change in PageRank scores with one new interaction is

$$\begin{aligned} \boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t) &= \left[\lambda(P(t)^T + \Delta P^T) + (1-\lambda)\mathbf{v}\mathbb{1}^T\right]\boldsymbol{\pi}(t+\Delta t) \\ &\quad - \left[\lambda P(t)^T + (1-\lambda)\mathbf{v}\mathbb{1}^T\right]\boldsymbol{\pi}(t) \\ &= \lambda P(t)^T(\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)) + \lambda\Delta P^T\boldsymbol{\pi}(t+\Delta t)\,. \end{aligned} \tag{28}$$

Rearranging terms gives

$$\left(I_n - \lambda P(t)^T\right)(\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)) = \lambda\Delta P^T\boldsymbol{\pi}(t+\Delta t)\,, \tag{29}$$

which implies that

$$\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t) = \lambda\left(I_n - \lambda P(t)^T\right)^{-1}\Delta P^T\boldsymbol{\pi}(t+\Delta t)\,. \tag{30}$$

From a Neumann-series expansion [78], we see that $\|(I_n - \lambda P(t)^T)^{-1}\|_1$ is bounded above by $1/(1-\lambda)$.

Taking norms on both sides of (30) yields

$$\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \leq \frac{\lambda}{1-\lambda}\left\|\Delta P^T\right\|_1\,. \tag{31}$$

Noting that $\|\Delta P^T\|_1 = \|\Delta P\|_\infty$, we use the definition of ΔP from equation (10) to obtain

$$\begin{aligned} &\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \\ &\quad \leq \frac{\lambda}{(1-\lambda)(1+e^{-\alpha\Delta t} d_{ii}(t))}\left\|\mathbf{e}_i\mathbf{e}_j^T - \frac{1}{d_{ii}(t)}\mathbf{e}_i\mathbf{e}_i^T B(t)\right\|_\infty \\ &\quad - \frac{\lambda c_i(t)v_i}{1-\lambda}\left\|\mathbf{e}_i\mathbb{1}^T\right\|_\infty\,. \end{aligned} \tag{32}$$

Recall that $B(t)$ is the tie-strength matrix and that $\mathbf{e}_i$ and $\mathbf{e}_j$, respectively, are the i-th and j-th canonical vectors. Let $Q = \mathbf{e}_i\mathbf{e}_j^T - \frac{1}{d_{ii}(t)}\mathbf{e}_i\mathbf{e}_i^T B(t)$ be the matrix with entries

$$q_{hk} = \begin{cases} 1 - b_{hk}(t)/d_{ii}(t)\,, & h = i\,,\ k = j \\ -b_{hk}(t)/d_{ii}(t)\,, & h = i\,,\ k \neq j \\ 0\,, & \text{otherwise}\,, \end{cases} \tag{33}$$

so Q has nonzero entries only in row i. Noting that $d_{ii}(t) = \sum_k b_{ik}(t)$ and using

$$\|Q\|_\infty = \max_{1\leq h\leq n}\sum_{k=1}^n |q_{hk}| = \sum_{k=1}^n |q_{ik}|\,, \tag{34}$$

we see that

$$\|Q\|_\infty \leq 2\,. \tag{35}$$

We also observe that $\mathbf{e}_i\mathbb{1}^T$ is the $n\times n$ matrix whose entries are equal to 1 in row i and are equal to 0 elsewhere. Therefore,

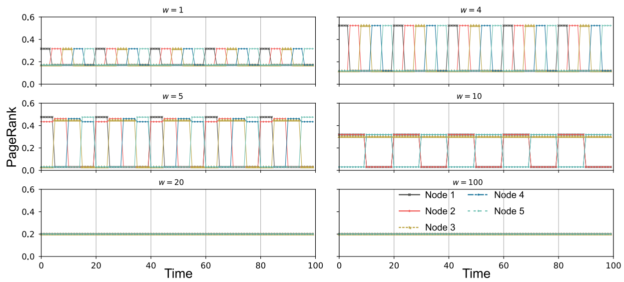

Fig. 10. PageRank scores over time for each node in our five-node synthetic network using adjacent (and hence nonoverlapping) time windows.

$$\left\|\mathbf{e}_i \mathbb{1}^T\right\|_\infty = n\,. \tag{36}$$

With $nv_i = 1$ (i.e., uniform teleportation), it follows from equations (32), (35), and (36) that

$$\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \leq \frac{2\lambda}{(1-\lambda)(1+e^{-\alpha \Delta t} d_{ii}(t))} - \frac{\lambda c_i(t)}{1-\lambda}\,. \tag{37}$$

The change in the PageRank vector is also subject to the bound [47], [60]

$$\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \leq \frac{2\lambda}{1-\lambda} \sum_{s \in \mathcal{S}(t+\Delta t)} \pi_s(t) = \frac{2\lambda}{1-\lambda} \pi_i(t)\,, \tag{38}$$

where $\mathcal{S}(t+\Delta t)$ is the set of nodes (in this case, just node i) that experience a change in transition probabilities (i.e., a change in outgoing edges). Combining the bounds in (37) and (38) yields

$$\begin{aligned}&\|\boldsymbol{\pi}(t+\Delta t) - \boldsymbol{\pi}(t)\|_1 \\ &\leq \frac{2\lambda}{1-\lambda} \min\left\{\pi_i(t),\, \frac{1}{1+e^{-\alpha\Delta t} d_{ii}(t)} - \frac{c_i(t)}{2}\right\},\end{aligned} \tag{39}$$

which completes our proof. ■

Note that $d_{ii}(t) > 0 \Leftrightarrow c_i(t) = 0$ and $d_{ii}(t) = 0 \Leftrightarrow c_i(t) = 1$, which guarantees that the quantity on the right-hand side of (39) is always positive. This gives the results in Corollaries 1 and 2.

C. Synthetic Network: Using Adjacent Time Windows

In our synthetic network from Section IV, suppose that we aggregate the interactions into adjacent (and hence nonoverlapping) windows of length w. We then calculate PageRank time series for the resulting sequence of time-independent networks and show our results in Fig. 10.

D. NHS Retweet Network: Network Statistics

The NHS network includes retweets between the 10 000 most-active accounts, where we quantify their activity as their number of tweets in the data set. There are 181123 retweets between the 10 000 most-active accounts between 5 March 2012 and 21 August 2012. We interpret each of these retweets as one interaction. In total, 6866 of the 10 000 accounts interact with each other via retweets at least once during this time period. There are 6013 accounts that retweet others and 4957 accounts whose tweets are retweeted. In Fig. 11, we show the the distribution of retweets among these accounts.

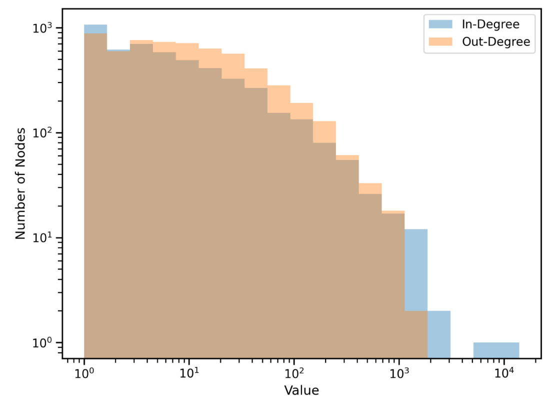

Fig. 11. Distributions of retweets received (i.e., in-degree) and sent (i.e., out-degree) in the NHS retweet network.

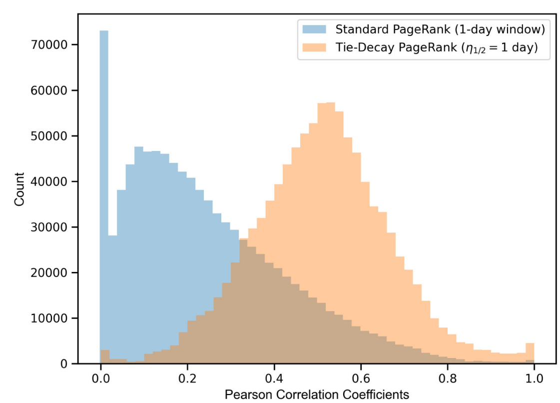

Fig. 12. Distributions of Pearson correlation coefficients for the correlation matrices in Fig. 8. We show the distribution for the standard PageRank score with time windows of 1 day in blue and the distribution for tie-decay PageRank with a half-life of $\eta_{1/2} = 1$ day in orange.

E. NHS Retweet Network: Aggregating Interactions Versus Using Tie-Decay Networks

In Fig. 12, we show the distributions of the Pearson correlations between PageRank scores using both tie-decay networks and networks with sliding time windows.

Acknowledgements

We thank Mauricio Barahona, Alain Goriely, Peter Grindrod, Mikko Kivelä, Renaud Lambiotte, Naoki Masuda, Fabian Ying, and Xinzhe Zuo for useful discussions. We thank Guillermo Garduño Hernández and Sinnia for their help in collecting the data. MBD acknowledges support from the Oxford–Emirates Data Science Laboratory.

References

[1] M. Bazzi, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, "Community detection in temporal multilayer networks, with an application to correlation networks," *Multiscale Model. Simul.*, vol. 14, pp. 1–41, 2016.

[2] M. Beguerisse-Díaz, G. Garduño-Hernández, B. Vangelov, S. N. Yaliraki, and M. Barahona, "Interest communities and flow roles in directed networks: The Twitter network of the UK riots," *J. Roy. Soc.: Interface*, vol. 11, 2014, Art no. 20140940.

[3] M. Beguerisse-Díaz, A. K. McLennan, G. Garduño-Hernández, M. Barahona, and S. J. Ulijaszek, "The 'who' and 'what' of #diabetes on Twitter," *Digit. Health*, vol. 3, 2017, Art. no. 2055207616688841.

[4] M. Beguerisse-Díaz, M. A. Porter, and J.-P. Onnela, "Competition for popularity in bipartite networks," *Chaos*, vol. 20, 2010, Art. no. 043101.

[5] I. V. Belykh, V. N. Belykh, and M. Hasler, "Blinking model and synchronization in small-world networks with a time-varying coupling," *Phys. D.*, vol. 195, pp. 188–206, 2004.

[6] F. Béres, R. Pálovics, A. Oláh, and A. A. Benczúr, "Temporal walk based centrality metric for graph streams," *Appl. Netw. Sci.*, vol. 3, 2018, Art. no. 32.

[7] M. Bianchini, M. Gori, and F. Scarselli, "Inside PageRank," *ACM Trans. Internet Technol.*, vol. 5, pp. 92–128, 2005.

[8] P. Bonacich and P. Lloyd, "Eigenvector-like measures of centrality for asymmetric relations," *Social Netw.*, vol. 23, pp. 191–201, 2001.

[9] D. Braha and Y. Bar-Yam, "From centrality to temporary fame: Dynamic centrality in complex networks," *Complexity*, vol. 12, pp. 59–63, 2006.

[10] D. Braha and Y. Bar-Yam, "Time-dependent complex networks: Dynamic centrality, dynamic motifs, and cycles of social interactions," in *Adaptive Networks: Theory, Models and Applications*, T. Gross and H. Sayama, Eds. Heidelberg, Germany: Springer-Verlag, 2009, pp. 39–50.

[11] R. S. Burt, "Decay functions," *Social Netw.*, vol. 22, pp. 1–28, 2000.

[12] R. S. Caceres, T. Berger-Wolf, and R. Grossman, "Temporal scale of processes in dynamic networks," in *2011 IEEE 11th Int. Conf. Data Mining Workshops*, New York, NY, USA: Institute Electrical Electronics Engineers, 2011, pp. 925–932.

[13] National Research Council, *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, R. Breiger, K. Carley, and P. Pattison, Eds. Nat. Acad. Press, Washington, DC, USA, 2003.

[14] S. Chien, C. Dwork, R. Kumar, D. R. Simon, and D. Sivakumar," *Link Evol.: Anal. Algorithms*, vol. 1, pp. 277–304, 2004.

[15] P. Cihon and T. Yasseri, "A biased review of biases in Twitter studies on political collective action," *Front. Phys.*, vol. 4, 2016, Art. no. 34.

[16] R. Erban, S. J. Chapman, and P. K. Maini, "A. practical guide to stochastic simulations of reaction–diffusion processes," 2007, *arXiv:0704.1908*.

[17] J. D. Farmer, S. A. Kauffman, N. H. Packard, and A. S. Perelson, "Adaptive dynamic networks as models for the immune system and autocatalytic sets," *Ann. New York Acad. Sci.*, vol. 504, pp. 118–131, 1987.

[18] D. J. Fenn *et al.*, "Dynamical clustering of exchange rates," *Quantitative Finance*, vol. 12, pp. 1493–1520, 2012.

[19] J. Flores and M. Romance, "On eigenvector-like centralities for temporal networks: Discrete vs. continuous time scales," *J. Comput. Appl. Math.*, vol. 330, pp. 1041–1051, 2018.

[20] M. M. Gaber, A. Zaslavsky, and S. Krishnaswamy, "Mining data streams: A review," *ACM SIGMOD Rec.*, vol. 34, pp. 18–26, 2005.

[21] J. Giles, "Making the links," *Nature*, 488, no. 7412, p. 448–450, 2012.

[22] J. P. Gleeson, D. Cellai, J.-P. Onnela, M. A. Porter, and F. Reed-Tsochas, "A simple generative model of collective online behavior," *Proc. Nat. Acad. Sci. USA*, vol. 111, pp. 10411–10415, 2014.

[23] D. F. Gleich, "PageRank beyond the Web," *SIAM Rev.*, vol. 57, pp. 321–363, 2015.

[24] D. F. Gleich and R. A. Rossi, A. dynamical system for PageRank with time-dependent teleportation," *Internet Math.*, vol. 10, pp. 188–217, 2014.

[25] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed. Baltimore, MD, USA: Johns Hopkins Univ. Press, 2012.

[26] S. González-Bailón and N. Wang, "Networked discontent: The anatomy of protest campaigns in social media," *Social Netw.*, vol. 44, pp. 95–104, 2016.

[27] P. Grindrod and D. J. Higham, "A matrix iteration for dynamic network summaries," *SIAM Rev.*, vol. 55, pp. 118–128, 2013.

[28] P. Grindrod and D. J. Higham, "A dynamical systems view of network centrality," *Proc. Roy. Soc. A: Math., Phys. Eng. Sci.*, vol. 470, 2014, Art. no. 20130835.

[29] P. Grindrod, D. J. Higham, and P. Laflin, "The graph whisperers," In *UK Success Stories in Industrial Mathematics*, P. Aston, A. Mulholland, and K. Tant, Eds. Cham, Switzerland: Springer International Publishing, pp. 271–279, 2016.

[30] N. O. Hodas and K. Lerman, "How Visibility and Divided Attention Constrain Social Contagion," in *Proc. 2012 ASE/IEEE Int. Conf. Social Comput. and 2012 ASE/IEEE Int. Conf. Privacy, Secur., Risk and Trust*, 2012, pp. 249–257.

[31] T. Hoffmann, M. A. Porter, and R. Lambiotte, "Generalized master equations for non-Poisson dynamics on networks," *Phys. Rev. E*, vol. 86, 2012, Art. no. 046102.

[32] P. Holme, "Modern temporal network theory: A colloquium," *Eur. Phys. J. B.*, vol. 88, 2015, Art. no. 234.

[33] P. Holme and J. Saramäki, "Temporal networks," *Phys. Rep.*, vol. 519, pp. 97–125, 2012.

[34] P. Holme and J. Saramäki, Eds. *Temporal Network Theory*. Cham, Switzerland: Springer International Publishing, 2019.

[35] S. Howison, *Practical Applied Mathematics: Modelling, Analysis, Approximation*. Cambridge Texts in Applied Mathematics, Cambridge, U.K.: Cambridge Univ. Press, 2003.

[36] L. G. S. Jeub, P. Balachandran, M. A. Porter, P. J. Mucha, and M. W. Mahoney, "Think locally, act locally: Detection of small, medium-sized, and large communities in large networks," *Phys. Rev. E.*, vol. 91, 2015, Art. no. 012821.

[37] E. M. Jin, M. Girvan, and M. E. J. Newman, "Structure of growing social networks," *Phys. Rev. E.*, vol. 64, 2001, Art. no. 046132.

[38] L. Katz and C. H. Proctor, "The concept of configuration of interpersonal relations in a group as a time-dependent stochastic process," *Psychometrika*, vol. 24, pp. 317–327, 1959.

[39] H. Kim and R. Anderson, "Temporal node centrality in complex networks," *Phys. Rev. E.*, vol. 852012, Art. no. 026107.

[40] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks," *J. Complex Netw.*, vol. 2, pp. 203–271, 2014.

[41] M. Kivelä and M. A. Porter, "Estimating interevent time distributions from finite observation periods in communication networks," *Phys. Rev. E.*, vol. 92, 2015, Art. no. 052813.

[42] J. M. Kleinberg, "Authoritative sources in a hyperlinked environment," *J. ACM*, vol. 46, pp. 604–632, 1999.

[43] H. Kwak, C. Lee, H. Park, and S. Moon, "What is Twitter, a social network or a news media?," in *Proc. 19th Int. Conf. World Wide Web*, 2010, pp. 591–600.

[44] R. Lambiotte and M. Rosvall, "Ranking and clustering of nodes in networks with smart teleportation," *Phys. Rev. E.*, vol. 85, 2012, Art no. 056107.

[45] M. Latapy, T. Viard, and C. Magnien, "Stream graphs and link streams for the modeling of interactions over time," *Social Netw. Anal. Mining*, vol. 8, 2018, Art. no. 61.

[46] P. J. Laub, T. Taimre, and P. K. Pollett, "Hawkes processes," 2015, *arXiv:1507.02822*.

[47] H. C. Lee and A. Borodin, "Perturbation of the hyper-linked environment," in *Proc. Int. Comput. Combinatorics Conf.*, Heidelberg, Germany: Springer-Verlag, 2003, pp. 272–283.

[48] legislation.gov.uk, "Health and Social Care Act 2012," 2012. Accessed: Jun. 15, 2021. [Online]. Available: https://www.legislation.gov.uk/ukpga/2012/7/contents/enacted/data.htm

[49] K. Lerman, "Information is not a virus, and other consequences of human cognitive limits," *Future Internet*, vol. 8, 2016, Art. no. 21.

[50] K. Lerman, R. Ghosh, and T. Surachawala, "Social contagion: An empirical study of information spread on Digg and Twitter follower graphs," 2012. *arXiv:1202.3162*.

[51] T. Martin, X. Zhang, and M. E. J. Newman, "Localization and centrality in networks," *Phys. Rev. E.*, vol. 90, 2014, Art. no. 052808.

[52] N. Masuda and R. Lambiotte, *A Guide to Temporal Networks*, 2nd ed. Singapore: World Scientific Publishing, 2020.

[53] N. Masuda, M. A. Porter, and R. Lambiotte, "Random walks and diffusion on networks," *Phys. Rep.*, vol. 716–717, pp. 1–58, 2017.

[54] N. Masuda and L. Rocha, A, "Gillespie algorithm for non-Markovian stochastic processes," *SIAM Rev.*, vol. 60, pp. 95–115, 2018.

[55] R. Michalski, B. K. Szymański, P. Kazienko, C. Lebiere, O. Lizardo, and M. Kulisiewicz, "Social networks through the prism of cognition," *Complexity*, vol. 2021, 2021, Art. no. 4963903.

[56] A. J. Morales, J. C. Losada, and R. M. Benito, "Users structure and behavior on an online social network during a political protest," *Phys. A.*, vol. 391, pp. 5244–5253, 2012.

[57] S. Motegi and N. Masuda, "A. network-based dynamical ranking system for competitive sports," *Sci. Rep.*, vol. 2, 2012, Art. no. 904.

[58] H. Navarro, G. Miritello, A. Canales, and E. Moro, "Temporal patterns behind the strength of persistent ties," *Eur. Phys. J. — Data Sci.*, vol. 6, 2017, Art. no. 31.

[59] M. E. J. Newman, *Networks*. 2nd ed. Oxford, U.K.: Oxford Univ. Press, 2018.

[60] A. Y. Ng, A. X. Zheng, and M. I. Jordan, “Link analysis, eigenvectors and stability,” in *Proc. 17th Int. Joint Conf. Artif. Intell.—Vol. 2*, 2001, pp. 903–910.

[61] J.-P. Onnela *et al.*, “Analysis of a large-scale weighted network of one-to-one human communication,” *New J. Phys.*, vol. 9, 2007, Art. no. 179.

[62] D. J. P. O’Sullivan, G. Garduño-Hernández, J. P. Gleeson, and M. Beguerisse-Díaz, “Integrating sentiment and social structure to determine preference alignments: The Irish Marriage Referendum,” *Roy. Soc. Open Sci.*, vol. 4, 2017, Art. no. 170154.

[63] L. Page, S. Brin, R. Motwani, and T. Winograd, The PageRank citation ranking: Bringing order to the Web,” in *Proc. 7th Int. World Wide Web Conf.*, 1998, pp. 161–172.

[64] R. K. Pan and J. Saramäki, “Path lengths, correlations, and centrality in temporal networks,” *Phys. Rev. E.*, vol. 84, 2011, Art. no. 016105.

[65] L. Peel and A. Clauset, “Detecting change points in the large-scale structure of evolving networks,” in *Proc. 29th AAAI Conf. Artif. Intell.*, 2015, pp. 2914–2920.

[66] N. Perra and S. Fortunato, “Spectral centrality measures in complex networks,” *Phys. Rev. E.*, vol. 78, 2008, Art. no. 036107.

[67] M. A. Porter and J. P. Gleeson, *Dynamical Systems on Networks: A Tutorial*. Frontiers in Applied Dynamical Systems: Reviews and Tutorials, vol. 4. Cham, Switzerland: Springer International Publishing, 2016.

[68] I. Psorakis, “Probabilistic inference in ecological networks: Graph discovery, community detection and modelling dynamic sociality, D.Phil. Thesis, Univ. Oxford, Oxford, U.K., 2013.

[69] I. Psorakis, S. J. Roberts, I. Rezek, and B. C. Sheldon, “Inferring social network structure in ecological systems from spatio-temporal data streams,” *J. Roy. Soc. Interface*, vol. 9, pp. 3055–3066, 2012.

[70] S. D. Ridder, B. Vandermarliere, and J. Ryckebusch, “Detection and localization of change points in temporal networks with the aid of stochastic block models,” *J. Statist. Mech: Theory Experiment*, vol. 2016, 2016, Art. no. 113302.

[71] J. Saramäki and E. Moro, “From seconds to months: An overview of multi-scale dynamics of mobile telephone calls,” *Eur. Phys. J. B.*, vol. 88, 2015, Art. no. 164.

[72] U. Sharan and J. Neville, “Exploiting time-varying relationships in statistical relational models,” in *Proc. 9th WebKDD and 1st SNA–KDD Workshop Web Mining Social Netw. Anal.*, 2007, pp. 9–15.

[73] T. A. B. Snijders, “The statistical evaluation of social network dynamics,” *Sociol. Methodol.*, 31, pp. 361–395, 2001.

[74] R. Sulo, T. Berger-Wolf, and R. Grossman, “Meaningful selection of temporal resolution for dynamic networks,” in *Proc. 8th Workshop Mining Learn. Graphs*, 2010, pp. 127–136.

[75] D. Taylor, S. A. Myers, A. Clauset, M. A. Porter, and P. J. Mucha, “Eigenvector-based centrality measures for temporal networks,” *Multiscale Model. Simul.*, vol. 15, pp. 537–574, 2017.

[76] E. Tonkin, H. D. Pfeiffer, and G. Tourte, “Twitter, information sharing and the London riots?,” *Bull. Amer. Soc. Inf. Sci. Technol.*, vol. 38, pp. 49–57, 2012.

[77] L. N. Trefethen and D. Bau, *Numerical Linear Algebra*. Philadelphia, PA, USA: Soc. Ind. Appl. Math., 1997.

[78] E. E. Tyrtyshnikov, *A Brief Introduction to Numerical Analysis*. Basel, Switzerland: Birkhäuser, 1997.

[79] E. Valdano, M. R. Fiorentin, C. Poletto, and V. Colizza, “Epidemic threshold in continuous-time evolving networks,” *Phys. Rev. Lett.*, vol. 120, 2018, Art. no. 068302.

[80] C. L. Vestergaard and M. Génois, “Temporal Gillespie algorithm: Fast simulation of contagion processes on time-varying networks,” *PLoS Comput. Biol.*, vol. 11, 2015, Art. no. e1004579.

[81] S. Wasserman and D. Iacobucci, “Sequential social network data,” *Psychometrika*, vol. 53, pp. 261–282, 1988.

[82] E. W. Weisstein, *CRC Concise Encyclopedia of Mathematics*, 2nd ed. Boca Raton, FL, USA: CRC Press, 2002.

[83] J. T. Wixted and E. B. Ebbesen, “On the form of forgetting,” *Psychol. Sci.*, vol. 2, pp. 409–415, 1991.

[84] X. Yang and J. Fan, “Influential user subscription on time-decaying social streams,” 2018, *arXiv:1802.05305*.

[85] L. Zino, A. Rizzo, and M. Porfiri, “Continuous-time discrete-distribution theory for activity-driven networks,” *Phys. Rev. Lett.*, vol. 117, 2016, Art. no. 228302.

[86] L. Zino, A. Rizzo, and M. Porfiri, “An analytical framework for the study of epidemic models on activity driven networks,” *J. Complex Netw.*, vol. 5, pp. 924–952, 2017.

Walid Ahmad received his B.Sc. degree in Applied Mathematics from Columbia University, New York City, NY, USA and his M.Sc. degree in Mathematical Modelling and Scientific Computing from the University of Oxford. He is currently a Machine-Learning Research Engineer with Reverie Labs, Cambridge, MA, USA, where he works on machine-learning methods for computational drug discovery.

Mason A. Porter is currently a Professor in the Department of Mathematics at the University of California, Los Angeles. His research interests include networks, complex systems, and nonlinear systems. He is a Fellow of the American Mathematical Society, the American Physical Society, and the Society for Industrial and Applied Mathematics.

Mariano Beguerisse-Díaz is currently a Senior Research Scientist with Spotify and a Visiting Fellow at the University of Oxford. He has held Research Fellowships with the University of Oxford and with Imperial College London. His research interests include network science, data science, mathematical biology, recommendation systems, computational social science, and mathematical modeling.