



Link prediction in dynamic networks using random dot product graphs

Francesco Sanna Passino¹ · Anna S. Bertiger² · Joshua C. Neil^{2,3} · Nicholas A. Heard¹

Received: 1 October 2020 / Accepted: 21 July 2021 / Published online: 5 August 2021
© The Author(s) 2021

Abstract

The problem of predicting links in large networks is an important task in a variety of practical applications, including social sciences, biology and computer security. In this paper, statistical techniques for link prediction based on the popular random dot product graph model are carefully presented, analysed and extended to dynamic settings. Motivated by a practical application in cyber-security, this paper demonstrates that random dot product graphs not only represent a powerful tool for inferring differences between multiple networks, but are also efficient for prediction purposes and for understanding the temporal evolution of the network. The probabilities of links are obtained by fusing information at two stages: spectral methods provide estimates of latent positions for each node, and time series models are used to capture temporal dynamics. In this way, traditional link prediction methods, usually based on decompositions of the entire network adjacency matrix, are extended using temporal information. The methods presented in this article are applied to a number of simulated and real-world graphs, showing promising results.

Keywords Adjacency spectral embedding · Dynamic networks · Link prediction · Random dot product graph

Responsible editor: Jingrui He.

✉ Francesco Sanna Passino
f.sannapassino@imperial.ac.uk

¹ Department of Mathematics, Imperial College London, London, United Kingdom

² Microsoft 365 Defender, Microsoft Corporation, Redmond, WA, USA

³ Present Address: Securonix Threat Labs, Securonix Inc., Addison, TX, USA

1 Introduction

Link prediction is defined as the task of predicting the presence of an edge between two nodes in a network, based on latent characteristics of the graph (Liben-Nowell and Kleinberg 2007). The problem has been widely studied in the literature (Lü and Zhou 2011; Menon and Elkan 2011), and has relevant applications in a variety of different fields. In this paper, the discussion about link prediction is motivated by applications in cyber-security and computer network monitoring (Jeske et al. 2018). The ability to correctly predict and associate anomaly scores with the connections in a network is valuable for the cyber-defence of enterprises. In cyber settings, adversaries may introduce changes in the structure of an enterprise network in the course of their attack. Therefore, predicting links in order to identify significant deviations in expected behaviour could lead to the detection of an otherwise extremely damaging network breach. In particular, it is necessary to correctly score *new links* (Metelli and Heard 2019), representing previously unobserved connections. The task is particularly important since it is common to observe malicious activity associated with new links (Neil et al. 2013), and it is therefore crucial to understand the normal process of link formation in order to detect a cyber-attack.

In this article, it is assumed that snapshots of a dynamic network are observed at discrete time points $t = 1, \dots, T$, obtaining a sequence of graphs $\mathbb{G}_t = (V, E_t)$. The set V represents the set of nodes, which is invariant over time, whereas the set E_t is a time dependent edge set, where $(i, j) \in E_t$ for $i, j \in V$, if i is connected to j at least once during the time period $(t - 1, t]$. Each snapshot of the graph can be characterised by the adjacency matrix $\mathbf{A}_t \in \{0, 1\}^{n \times n}$, where $n = |V|$ and for $1 \leq i, j \leq n$, $A_{ijt} = \mathbb{1}_{E_t}\{(i, j)\}$, such that $A_{ijt} = 1$ if a link between the nodes i and j exists in $(t - 1, t]$, and $A_{ijt} = 0$ otherwise. The graph is said to be undirected if $(i, j) \in E_t \iff (j, i) \in E_t$, implying that \mathbf{A}_t is symmetric; otherwise, the graph is said to be directed. It will be assumed that the graph has no self-edges, implying \mathbf{A}_t is a hollow matrix. Similarly, bipartite graphs $\mathbb{G}_t = (V_1, V_2, E_t)$ can be represented using two node sets V_1 and V_2 , and rectangular adjacency matrices $\mathbf{A}_t \in \{0, 1\}^{n_1 \times n_2}$, $n_1 = |V_1|$, $n_2 = |V_2|$, where $A_{ijt} = 1$ if $i \in V_1$ connects to $j \in V_2$ in $(t - 1, t]$.

This paper discusses methods for *temporal link prediction* (Dunlavy et al. 2011): given a sequence of adjacency matrices $\mathbf{A}_1, \dots, \mathbf{A}_T$ observed over time, the main objective is to reliably predict \mathbf{A}_{T+1} . In this article, temporal link prediction techniques based on *random dot product graphs* (RDPG, Young and Scheinerman 2007) are discussed and compared. RDPGs are a class of *latent position models* (Hoff et al. 2002), and have been extensively studied because of their analytical tractability (Athreya et al. 2018). Each node i is given a latent position \mathbf{x}_i in a d -dimensional latent space \mathbb{X} such that $\mathbf{x}^\top \mathbf{x}' \in [0, 1] \forall \mathbf{x}, \mathbf{x}' \in \mathbb{X}$. The edges between pairs of nodes are generated independently, with probability of a link between nodes i and j obtained through the inner product $\mathbb{P}(A_{ij} = 1) = \mathbf{x}_i^\top \mathbf{x}_j$. In matrix notation, the latent position can be arranged in a $n \times d$ matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{X}^n$, and the expected value of a single realised adjacency matrix \mathbf{A} is expressed as $\mathbb{E}(\mathbf{A}) = \mathbf{X}\mathbf{X}^\top$.

Random dot product graph models and spectral embedding methods are often the first step in the analysis of a graph, because of their simplicity and ease of implemen-

tation, since intensive hyperparameter tuning is not required. RDPGs are extensively applied in neuroscience (see, for example, Priebe et al. 2017). Furthermore, they have appealing theoretical statistical properties in terms of consistency of the estimated latent positions. Therefore, it is of interest to understand their performance for link prediction purposes.

RDPGs models for multiple heterogeneous graphs on the same node set have recently been proposed in the literature, but these models have not been formally extended to a dynamic setting for link prediction purposes. Early examples discuss methods for clustering and community detection with multiple graphs (Tang et al. 2009; Shiga and Mamitsuka 2012; Dong et al. 2014). More recently, the focus has been on testing for differences in brain connectivity networks (Arroyo-Reli3n et al. 2019; Ginestet et al. 2017; Durante and Dunson 2018; Kim and Levina 2019). Levin et al. (2017) propose an *omnibus* embedding in which the different graphs are jointly embedded into a common latent space, providing distinct representations for each graph and for each node. Wang et al. (2021) propose the *multiple random eigen graph* (MREG) model, where a common set of d -dimensional latent features \mathbf{X} is shared between the graphs, and the inner product between the latent positions is weighted differently across the networks, obtaining $\mathbb{E}(\mathbf{A}_t) = \mathbf{X}\mathbf{R}_t\mathbf{X}^\top$, where \mathbf{R}_t is a $d \times d$ diagonal matrix. Nielsen and Witten (2018) propose the *multiple random dot product graph* (multi-RDPG), which more naturally extends the RDPG to the multi-graph setting. Their formulation is similar to the MREG of Wang et al. (2021), but \mathbf{X} is modelled as an orthogonal matrix, and \mathbf{R}_t is constrained to be positive semi-definite. The model is further extended in *common subspace independent edge* (COSIE) graphs (Arroyo-Reli3n et al. 2020), in which \mathbf{R}_t does not need to be a diagonal matrix. More recently, Jones and Rubin-Delanchy (2021) proposed the Unfolded Adjacency Spectral Embedding (UASE) for the multilayer random dot product graph (MRDPG), which is also applied to a link prediction example within a cyber-security context. In this work, existing methods for RDPG-based inference, for example omnibus embeddings (Levin et al. 2017) and COSIE graphs (Arroyo-Reli3n et al. 2020), will be analysed for link prediction purposes, and compared to standard spectral embedding techniques.

The main contribution of this work is to adapt the existing methods for multiple RDPG graph inference for temporal link prediction. Furthermore, this article proposes methods to combine the information obtained via spectral methods with time series models, to capture the temporal dynamics of the observed graphs. The proposed methodologies will be extensively compared on real world and simulated networks. It will be shown that this approach significantly improves the predictive performance of multiple RDPG models, especially when the network presents a seasonal or temporal evolution. Overall, this article provides insights into the predictive capability of random dot product graphs, and gives guidelines for practitioners on the optimal choice of the embedding for temporal link prediction.

Importantly, the strategies for combination of individual embeddings, and their time series extensions, can in principle be applied to *any* embedding method for static graphs, despite the main focus on RDPGs of this article. This article primarily focuses on RDPGs because of the wide variety of embedding techniques which have been suggested in the literature under this model, but that so far have not been compared for link prediction purposes.

The article is organised as follows: Sect. 2 discusses related literature around link prediction, and Sect. 3 introduces background on the generalised random dot product graph and adjacency spectral embeddings, the main statistical tools used in this paper. Methods for link prediction based on random dot product graphs are discussed in Sect. 4. Section 5 presents techniques to improve the predictive performance of the RDPG models, based on time series models. Results and applications on simulated and real world networks are finally discussed in Sect. 6.

2 Related literature

Many other models other than RDPGs have been proposed in the literature for link prediction. Traditionally, the temporal link prediction task is tackled using tensor decompositions (Dunlavy et al. 2011). Dynamic models have also been proposed in the literature of Poisson matrix factorisation and recommender systems (Charlin et al. 2015; Hosseini et al. 2020), and extended to Bayesian tensor decompositions (Schein et al. 2015). In general, including time has been shown to significantly improve the predictive performance in a variety of model settings, for example stochastic block-models (Ishiguro et al. 2010; Xu and Hero III 2014; Xing et al. 2010). More generic latent feature models for dynamic networks have also been extensively discussed in the literature (Sarkar and Moore 2006; Krivitsky and Handcock 2014; Sewell and Chen 2015).

Latent features are usually obtained via matrix factorisation, considering constant and time-varying components within the decomposition (Deng et al. 2016; Yu et al. 2017a, b). Usually, a Markov assumption is placed on the evolution of the latent positions (Zhu et al. 2016; Chen and Li 2018). Gao et al. (2011) propose to combine node features into a temporal link prediction framework based on matrix factorisation. Non-parametric (Sarkar et al. 2014; Durante and Dunson 2014) and deep learning (Li et al. 2014) approaches have also been considered.

More recently, advancements have been made in the application of deep learning to graph-valued data. In particular, deep learning methods on graphs are classified by Zhang et al. (2020) into five categories: graph recurrent neural networks, graph convolutional networks, graph autoencoders, graph reinforcement learning, graph adversarial methods. A comprehensive survey of existing static network embedding methods, including deep learning techniques, is provided in Cai et al. (2018). Commonly used static embedding methods in machine learning are DeepWalk (Perozzi et al. 2014), SDNE (Wang et al. 2016), node2vec (Grover and Leskovec 2016), GraphSAGE (Hamilton et al. 2017), graph convolutional networks (GCN, Kipf and Welling 2017), and Watch Your Step with Graph Attention (WYS-GA, Abu-El-Haija et al. 2018). Many of such methods could be unified under a matrix factorisation framework (Qiu et al. 2018). A systematic comparison of some of the aforementioned methodologies for unsupervised network embedding is provided in Khosla et al. (2021). Methodologies have also been recently proposed in the dynamic network setting, within the context of representation learning (for example, Nguyen et al. 2018; Kumar et al. 2019; Liu et al. 2019; Qu et al. 2020), and deep generative models (for example,

Zhou et al. 2020). The interested reader is referred to the survey of Kazemi et al. (2020) and references therein.

Again, it is emphasised that the objective of this paper is *not* to claim that the RDPG is superior to competing models, but to provide guidelines for practitioners using RDPGs in their application domains, offering insights on the performance of these models for link prediction purposes.

3 Random dot product graphs and adjacency spectral embedding

In this section, the generalised random dot product graph (Rubin-Delanchy et al. 2017) and methods for estimation of the latent positions are formally introduced. Suppose $\mathbf{A} \in \{0, 1\}^{n \times n}$ is a symmetric adjacency matrix of an undirected graph with n nodes.

Definition 1 (*Generalised random dot product graph – GRDPG*) Let d_+ and d_- be non-negative integers such that $d = d_+ + d_-$. Let $\mathbb{X} \subseteq \mathbb{R}^d$ such that $\forall \mathbf{x}, \mathbf{x}' \in \mathbb{X}, 0 \leq \mathbf{x}^\top \mathbf{I}(d_+, d_-) \mathbf{x}' \leq 1$, where

$$\mathbf{I}(p, q) = \text{diag}(1, \dots, 1, -1, \dots, -1) \tag{1}$$

with p ones and q minus ones. Let \mathcal{F} be a probability measure on \mathbb{X} , $\mathbf{A} \in \{0, 1\}^{n \times n}$ be a symmetric matrix and $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{X}^n$. Then $(\mathbf{A}, \mathbf{X}) \sim \text{GRDPG}_{d_+, d_-}(\mathcal{F})$ if $\mathbf{x}_1, \dots, \mathbf{x}_n \stackrel{iid}{\sim} \mathcal{F}$ and for $i < j, \mathbb{P}(A_{ij} = 1) = \mathbf{x}_i^\top \mathbf{I}(d_+, d_-) \mathbf{x}_j$ independently.

The adjacency spectral embedding (ASE) provides consistent estimates of the latent positions in GRDPGs (Rubin-Delanchy et al. 2017), up to indefinite orthogonal transformations.

Definition 2 (*Adjacency spectral embedding – ASE*) For $d \in \{1, \dots, n\}$, consider the spectral decomposition $\mathbf{A} = \hat{\mathbf{\Gamma}} \hat{\mathbf{\Lambda}} \hat{\mathbf{\Gamma}}^\top + \hat{\mathbf{\Gamma}}_\perp \hat{\mathbf{\Lambda}}_\perp \hat{\mathbf{\Gamma}}_\perp^\top$, where $\hat{\mathbf{\Lambda}}$ is a $d \times d$ diagonal matrix containing the top d eigenvalues in magnitude, in decreasing order, $\hat{\mathbf{\Gamma}}$ is a $n \times d$ matrix containing the corresponding orthonormal eigenvectors, and the matrices $\hat{\mathbf{\Lambda}}_\perp$ and $\hat{\mathbf{\Gamma}}_\perp$ contain the remaining $n - d$ eigenvalues and eigenvectors. The adjacency spectral embedding $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n]^\top$ of \mathbf{A} in \mathbb{R}^d is

$$\hat{\mathbf{X}} = \hat{\mathbf{\Gamma}} |\hat{\mathbf{\Lambda}}|^{1/2} \in \mathbb{R}^{n \times d}, \tag{2}$$

where the operator $|\cdot|$ applied to a matrix returns the absolute value of its entries.

If the graph is directed, and the adjacency matrix is not symmetric, it could be implicitly assumed that the generating model is $\mathbb{P}(A_{ij} = 1) = \mathbf{x}_i^\top \mathbf{y}_j, \mathbf{x}_i, \mathbf{y}_j \in \mathbb{X}$, where each node is given two different latent positions, representing the behaviour of the node as source or destination of the link. In this case, the embeddings can be estimated using the singular value decomposition (SVD).

Definition 3 (*Adjacency embedding of the directed graph – DASE*) Given a directed graph with adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$, and a positive integer $d, 1 \leq d \leq n$, consider the singular value decomposition

$$\mathbf{A} = [\hat{\mathbf{U}} \ \hat{\mathbf{U}}_{\perp}] \begin{bmatrix} \hat{\mathbf{D}} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{D}}_{\perp} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{V}}^{\top} \\ \hat{\mathbf{V}}_{\perp}^{\top} \end{bmatrix} = \hat{\mathbf{U}}\hat{\mathbf{D}}\hat{\mathbf{V}}^{\top} + \hat{\mathbf{U}}_{\perp}\hat{\mathbf{D}}_{\perp}\hat{\mathbf{V}}_{\perp}^{\top}, \tag{3}$$

where $\hat{\mathbf{D}} \in \mathbb{R}_+^{d \times d}$ is diagonal matrix containing the top d singular values in decreasing order, $\hat{\mathbf{U}} \in \mathbb{R}^{n \times d}$ and $\hat{\mathbf{V}} \in \mathbb{R}^{n \times d}$ contain the corresponding left and right singular vectors, and the matrices $\hat{\mathbf{D}}_{\perp}, \hat{\mathbf{U}}_{\perp}$, and $\hat{\mathbf{V}}_{\perp}$ contain the remaining $n - d$ singular values and vectors. The d -dimensional directed adjacency embedding of \mathbf{A} in \mathbb{R}^d , is defined as the pair

$$\hat{\mathbf{X}} = \hat{\mathbf{U}}\hat{\mathbf{D}}^{1/2}, \quad \hat{\mathbf{Y}} = \hat{\mathbf{V}}\hat{\mathbf{D}}^{1/2}. \tag{4}$$

The DASE can be also naturally extended to bipartite graphs.

4 Dynamic link prediction in random dot product graphs

Given a time series of network adjacency matrices $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_T$, the objective is to correctly predict \mathbf{A}_{T+1} . The most common approach in the literature (Sharan and Neville 2008; Scheinerman and Tucker 2010; Dunlavy et al. 2011) is to analyse a collapsed version $\tilde{\mathbf{A}}$ of the adjacency matrices:

$$\tilde{\mathbf{A}} = \sum_{t=1}^T \psi_{T-t+1} \mathbf{A}_t, \tag{5}$$

where $\psi_1, \dots, \psi_T \in \mathbb{R}$ is a sequence of weights. Scheinerman and Tucker (2010) propose to consider an average adjacency matrix, setting $\psi_t = 1/T \ \forall t = 1, \dots, T$, which corresponds to the maximum likelihood estimate of $\mathbb{E}(\mathbf{A}_t)$ if $\mathbf{A}_1, \dots, \mathbf{A}_T$ are sampled independently from the same Bernoulli($\mathbf{X}\mathbf{X}^{\top}$) distribution. The main limitation of such a model is that it is assumed that the graphs do not display any temporal evolution. Furthermore, if (5) is used, it is assumed that all the possible edges of the adjacency matrix follow the same dynamics. Obtaining the ASE $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n]$ of $\tilde{\mathbf{A}}$ leads to an estimate the scores:

$$\mathbf{S} = \hat{\mathbf{X}}\hat{\mathbf{X}}^{\top}. \tag{6}$$

For simplicity, the inner product (6) is not weighted by the matrix $\mathbf{I}(d_+, d_-)$, implicitly assuming $d_+ = d$ and $d_- = 0$. The estimation approaches in (5) and (6) will be used as baselines for comparison with alternative methods for temporal link prediction techniques using RDPGs, which will be proposed and discussed in the remainder of this section. The proposed methods will be classified in the following two categories:

- averages of inner products of embeddings (AIP),
- inner products of an average of embeddings (IPA).

First, it is possible to consider the individual ASE for each adjacency matrix \mathbf{A}_t , and calculate an AIP score:

$$\mathbf{S}_{\text{AIP}} = \frac{1}{T} \sum_{t=1}^T \hat{\mathbf{X}}_t \hat{\mathbf{X}}_t^\top. \tag{7}$$

A second option is to obtain an averaged embedding $\bar{\mathbf{X}}$ from $\hat{\mathbf{X}}_1, \hat{\mathbf{X}}_2, \dots, \hat{\mathbf{X}}_T$, and use this for predicting the link probabilities. This procedure is slightly more complex than (7), since the embeddings are invariant to orthogonal transformations: given an orthogonal matrix $\mathbf{\Omega}_t \in \mathbb{R}^{d \times d}$, $\mathbb{E}(\mathbf{A}_t) = \mathbf{X}_t \mathbf{X}_t^\top = (\mathbf{X}_t \mathbf{\Omega}_t)(\mathbf{X}_t \mathbf{\Omega}_t)^\top$. Therefore, the embeddings $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_T$ only provide estimates of the corresponding latent positions up to an orthogonal transformation, which could vary for different values of t . Consequently, the embeddings must be suitably *aligned* or *registered*, before a direct comparison can be carried out. Discussion of a technique to align the embeddings is deferred to Appendix A. Assuming that an averaged embedding $\bar{\mathbf{X}}$ is obtained, the matrix of IPA scores for prediction of \mathbf{A}_{T+1} is:

$$\mathbf{S}_{\text{IPA}} = \bar{\mathbf{X}} \bar{\mathbf{X}}^\top. \tag{8}$$

Similar scoring mechanisms can be derived for the techniques of multiple graph inference described in Sect. 1, such as the omnibus embedding (Levin et al. 2017), based on the the omnibus matrix:

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_1 & \frac{\mathbf{A}_1 + \mathbf{A}_2}{2} & \dots & \frac{\mathbf{A}_1 + \mathbf{A}_T}{2} \\ \frac{\mathbf{A}_2 + \mathbf{A}_1}{2} & \mathbf{A}_2 & \dots & \frac{\mathbf{A}_2 + \mathbf{A}_T}{2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\mathbf{A}_T + \mathbf{A}_1}{2} & \frac{\mathbf{A}_T + \mathbf{A}_2}{2} & \dots & \mathbf{A}_T \end{bmatrix}. \tag{9}$$

The ASE $\hat{\mathbf{X}}$ of $\tilde{\mathbf{A}}$ gives T latent positions for each node. The individual estimates $\hat{\mathbf{X}}_t = [\hat{\mathbf{x}}_{1t}, \dots, \hat{\mathbf{x}}_{nt}]$ of the latent positions for the t -th adjacency matrix are represented by the submatrix formed by the estimates between the $((t-1)n+1)$ -th and tn -th row of $\hat{\mathbf{X}}$. Then, from the time series $\hat{\mathbf{X}}_1, \hat{\mathbf{X}}_2, \dots, \hat{\mathbf{X}}_T$ of omnibus embeddings, a matrix of scores can be obtained using either AIP (7) or IPA (8). In this case, the individual embeddings are directly comparable and an alignment step is not required. On the other hand, the omnibus embedding cannot easily be updated when new graphs $\mathbf{A}_{T+1}, \mathbf{A}_{T+2}, \dots$ become available, since $\tilde{\mathbf{A}}$ and the embedding must be recomputed for each new snapshot. The idea of an omnibus embedding can be also easily extended to directed and bipartite graphs, constructing the matrix $\tilde{\mathbf{A}}$ analogously and then calculating the DASE.

Embeddings generated using the more parsimonious COSIE model (Arroyo-Reli3n et al. 2020) are also considered. In COSIE networks, the latent positions are assumed to be common across the T snapshots of the graph, but the link probabilities are scaled by a time-varying matrix $\mathbf{R}_t \in \mathbb{R}^{d \times d}$: $\mathbb{E}(\mathbf{A}_t) = \mathbf{X}\mathbf{R}_t\mathbf{X}^\top$. The common latent positions \mathbf{X} and the time series of weighting matrices $\mathbf{R}_1, \dots, \mathbf{R}_T$ can be estimated via multiple adjacency spectral embedding (MASE, Arroyo-Reli3n et al. 2020).

Definition 4 (*Multiple adjacency spectral embedding – MASE*) Given a sequence of network adjacency matrices $\mathbf{A}_1, \dots, \mathbf{A}_T$, and an integer $d \in \{1, \dots, n\}$, obtain the individual ASEs $\hat{\mathbf{X}}_t = \hat{\mathbf{\Gamma}}_t|\hat{\mathbf{\Lambda}}_t|^{1/2} \in \mathbb{R}^{n \times d}$. Then, construct the $n \times Td$ matrix $\tilde{\mathbf{\Gamma}} = [\hat{\mathbf{\Gamma}}_1, \dots, \hat{\mathbf{\Gamma}}_T] \in \mathbb{R}^{n \times Td}$, and consider its singular value decomposition $\tilde{\mathbf{\Gamma}} = \hat{\mathbf{U}}\hat{\mathbf{D}}\hat{\mathbf{V}}^\top + \hat{\mathbf{U}}_\perp\hat{\mathbf{D}}_\perp\hat{\mathbf{V}}_\perp^\top$, where $\hat{\mathbf{D}} \in \mathbb{R}_+^{d \times d}$ is a diagonal matrix containing the top d singular values in decreasing order, $\hat{\mathbf{U}} \in \mathbb{R}^{n \times d}$ and $\hat{\mathbf{V}} \in \mathbb{R}^{Td \times d}$ contain the corresponding left and right singular vectors, and the matrices $\hat{\mathbf{D}}_\perp$, $\hat{\mathbf{U}}_\perp$, and $\hat{\mathbf{V}}_\perp$ contain the remaining singular values and vectors. The d -dimensional multiple adjacency embedding of $\mathbf{A}_1, \dots, \mathbf{A}_T$ in \mathbb{R}^d is given by $\hat{\mathbf{X}} = \hat{\mathbf{U}}$, which provides an estimate of \mathbf{X} , and the sequence $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_T$, where

$$\hat{\mathbf{R}}_t = \hat{\mathbf{X}}^\top \mathbf{A}_t \hat{\mathbf{X}}. \tag{10}$$

For prediction, the matrix of AIP scores could be obtained from the time series of estimated link probabilities $\hat{\mathbf{X}}\hat{\mathbf{R}}_1\hat{\mathbf{X}}^\top, \dots, \hat{\mathbf{X}}\hat{\mathbf{R}}_T\hat{\mathbf{X}}^\top$:

$$\mathbf{S}_{\text{AIP}} = \frac{1}{T} \sum_{t=1}^T \hat{\mathbf{X}}\hat{\mathbf{R}}_t\hat{\mathbf{X}}^\top. \tag{11}$$

Alternatively, an averaged $\bar{\mathbf{R}}$ could be equivalently obtained from the time series of estimates $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_T$. Combining $\bar{\mathbf{R}}$ with the estimate of the latent positions $\hat{\mathbf{X}}$ yields the IPA scores:

$$\mathbf{S}_{\text{IPA}} = \hat{\mathbf{X}}\bar{\mathbf{R}}\hat{\mathbf{X}}^\top. \tag{12}$$

The COSIE model can also be extended to directed and bipartite graphs assuming $\mathbb{E}(\mathbf{A}_t) = \mathbf{X}\mathbf{R}_t\mathbf{Y}^\top$. This construction leads to estimates $\hat{\mathbf{R}}_t = \hat{\mathbf{U}}^\top \mathbf{A}_t \hat{\mathbf{V}}$, where $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ are estimates of \mathbf{X} and \mathbf{Y} obtained from MASE on the DASE embeddings $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_T$ and $\hat{\mathbf{Y}}_1, \dots, \hat{\mathbf{Y}}_T$, based on two matrices $\tilde{\mathbf{\Gamma}}$ constructed from the left and right singular vectors (*cf.* Definition 4).

In summary, several link prediction schemes based on random dot product graph models have been proposed, corresponding to three different types of spectral embedding:

- scores based on individual embeddings, *cf.* (7) and (8),
- omnibus scores, *cf.* (7) and (8), based on the matrix representation in (9),
- COSIE scores, *cf.* (11) and (12).

Two scores, denoted AIP and IPA, are calculated for each embedding type. All methods will be compared to the popular collapsed adjacency matrix method in (6).

The methods described in this section are all based on truncated eigendecompositions of some form of the adjacency matrix. The full eigendecomposition of a $n \times n$ dense matrix requires a cubic computational cost $\mathcal{O}(n^3)$, but only d eigenvectors and eigenvalues, where d is in general $\mathcal{O}(1)$, are required in the algorithms presented in this section. This reduces the computational effort to $\mathcal{O}(n^2)$ (Yang et al. 2021). Also, graph adjacency matrices are binary and normally highly sparse. In this setting, efficient algorithms based on the power method calculate the required decomposition of a matrix \mathbf{A} in $\mathcal{O}\{\text{nnz}(\mathbf{A})\mathcal{P}(1/\varepsilon)\}$ (Ghashami et al. 2016), where $\text{nnz}(\cdot)$ denotes the number of non-zero entries of the matrix, $\mathcal{P}(\cdot)$ is a polynomial function and $\varepsilon \in (0, 1)$ is an approximation error parameter. This allows the methodologies in this section to be scalable to large networks with the support of modern computer libraries. For example, calculating the individual embeddings $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_T$ only has complexity $\mathcal{O}\{\mathcal{P}(1/\varepsilon) \sum_{t=1}^T \text{nnz}(\mathbf{A}_t)\}$. COSIE adds a further SVD decomposition in the MASE algorithm, which requires further $\mathcal{O}(nd^2)$ operations. On the other hand, calculating the omnibus embedding for large graphs might quickly become cumbersome, especially if T is large, since up to $\mathcal{O}(n^2T^2)$ operations are required.

5 Improving prediction using time series models

The collapsed matrix used in (5) assumes that the underlying dynamics of each link are the same across the entire graph. This assumption is particularly limiting in real world applications, where different behaviours might be associated with different nodes or links. Instead, edge specific matrix parameters $\Psi_1, \dots, \Psi_T \in \mathbb{R}^{n \times n}$ might be able to more reliably capture the behaviour of each edge. A modification of the collapsed matrix $\tilde{\mathbf{A}}$ in (5) is therefore proposed:

$$\tilde{\mathbf{A}} = \sum_{t=1}^T (\Psi_{T-t+1} \odot \mathbf{A}_t), \quad (13)$$

where $\Psi_1, \dots, \Psi_T \in \mathbb{R}^{n \times n}$ is a sequence of weighting matrices, and \odot is the Hadamard element-wise product. The matrix in (13) is denoted *predicted adjacency matrix*. Note that in (13), the weights can only be estimated for those entries such that $A_{ijt} = 1$ for at least one $t \in \{1, \dots, T\}$, but the ASE of $\tilde{\mathbf{A}}$ still allows to estimate non-zero link probabilities even for those edges such that $A_{ijt} = 0 \forall t$.

The idea could be easily extended to all the other prediction settings proposed in Sect. 4, replacing the average link probability or average embedding with an autoregressive combination. For example, from the sequence of standard embeddings $\hat{\mathbf{X}}_1, \hat{\mathbf{X}}_2, \dots, \hat{\mathbf{X}}_T$, it could be possible to obtain the scores as:

$$\mathbf{S}_{\text{PIP}} = \sum_{t=1}^T (\Psi_{T-t+1} \odot \hat{\mathbf{X}}_t \hat{\mathbf{X}}_t^\top). \quad (14)$$

Alternatively, it could be possible to use a similar technique to obtain an estimate $\tilde{\mathbf{X}}_{T+1} = \sum_{t=1}^T (\Psi_{T-t+1} \odot \hat{\mathbf{X}}_t)$ of the embedding \mathbf{X}_{T+1} , where in this case $\Psi_t \in \mathbb{R}^{n \times d}$.

The scores are then obtained as:

$$\mathbf{S}_{IPP} = \tilde{\mathbf{X}}_{T+1} \tilde{\mathbf{X}}_{T+1}^\top. \tag{15}$$

Similarly, for COSIE, the scores could be based on a linear combination of $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_T$ to estimate \mathbf{R}_{T+1} .

The equations (14) and (15) define two different methodologies to extend multiple RDPG inference models to a dynamic setting. Following the same nomenclature used in Sect. 4, the scores based on time series models have been respectively denoted as PIP for the predicted inner product (14), and IPP for the inner product of the prediction (15). The PIP scores have a particularly interesting property: the node-specific embeddings are combined with time series models at the edge levels, fusing information at two different network resolutions.

Note that, for COSIE scores, $\mathbf{S}_{AIP} = \mathbf{S}_{IPA}$ from (11) and (12), but $\mathbf{S}_{PIP} \neq \mathbf{S}_{IPP}$. This is because for \mathbf{S}_{PIP} , the scores are obtained directly from a prediction based on the estimated link probabilities, whereas for \mathbf{S}_{IPA} , the prediction is based on an estimate of \mathbf{R}_{T+1} from $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_T$.

For estimation of the weighting matrices Ψ_1, \dots, Ψ_T , the time series of link probabilities or node embeddings will be modelled independently. Seasonal autoregressive integrated (SARI) processes represent a flexible modelling assumption. A univariate time series $Z_1, \dots, Z_T \in \mathbb{R}$, is a $\text{SARI}(p, b)(P, B)_s$ with period s if the series is a causal autoregressive process defined by the equation

$$\phi(L)\Phi(L^s)(1-L)^b(1-L^s)^B Z_t = \varepsilon_t, \quad \varepsilon_t \stackrel{iid}{\sim} \mathbb{N}(0, \sigma^2), \tag{16}$$

where $\phi(v) = 1 - \phi_1 v - \dots - \phi_p v^p$, $\Phi(v) = 1 - \Phi_1 v - \dots - \Phi_P v^P$, and L is the lag operator $L^k Z_t = Z_{t-k}$ (Brockwell and Davis 1987). For example, consider a process $\text{SARI}(1, 0)(0, 1)_s$. The model equation (16) becomes $\tilde{Z}_t = \phi_1 \tilde{Z}_{t-1} + \varepsilon_t$ with $\tilde{Z}_t = Z_t - Z_{t-s}$. The process is causal if and only if $\phi(v) \neq 0$ and $\Phi(v) \neq 0$ for $|v| \leq 1$. The parameters of the process p, b, P, B and s are all *integer-valued*, and can be interpreted as follows: s is the seasonal periodicity; b corresponds to the differencing required to make the process stationary in mean, variance, and autocorrelations, and B refers to the seasonal differencing; p is to the number of autoregressive terms appearing in the equation, and similarly P refers to the number of seasonal autoregressive terms. More details about SARI models and their interpretation are given in Brockwell and Davis (1987).

The value of s usually depends on the application domain. For example, in computer networks with daily network snapshots, it is reasonable to assume $s = 7$, which represents a periodicity of one week. The remaining parameters, p, b, P and B , could be estimated using information criteria. For small values of T , the corrected Akaike information criterion (AICc) is preferred. The corresponding coefficients of the polynomials $\phi(v)$ and $\Phi(v)$, and the variance σ^2 , can be estimated via maximum likelihood (Brockwell and Davis 1987). For a discussion on automatic selection of the parameters in SARI models, see Hyndman and Khandakar (2008).

For prediction of future values Z_{t+1} conditional on Z_1, \dots, Z_t , the general forecasting equation is obtained from (16), setting ε_{t+1} to its expected value $\mathbb{E}(\varepsilon_{t+1}) = 0$, and obtaining an estimate \hat{Z}_{t+1} solving from the known terms of the equation. Analogously, k -steps ahead forecasts for Z_{t+k} can also be obtained.

In this article, the univariate time series Z_1, \dots, Z_T modelled using SARI are of four different types:

- Time series of estimated link probabilities obtained from *any* embedding method: for example $\hat{\mathbf{x}}_{i1}^\top \hat{\mathbf{x}}_{j1}, \dots, \hat{\mathbf{x}}_{iT}^\top \hat{\mathbf{x}}_{jT}$ for the standard ASE, representing the sequence of estimated scores for the edge (i, j) . This type of time series is used to obtain PIP scores, see (14);
- Time series of node embeddings on a given embedding dimension: for example $\hat{x}_{ir1}, \dots, \hat{x}_{irT}$, obtained considering only the i -th node embedding on the r -th dimension from $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_T$. Such values are used for the IPP scores, see (15);
- Time series of entries of the COSIE matrix, used to obtain IPP scores, see (15). For example: $\hat{R}_{kh1}, \dots, \hat{R}_{khT}$, corresponding to the (k, h) -th entry in $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_T$;
- Times series of binary entries of the network adjacency matrices: for example A_{ij1}, \dots, A_{ijT} for the edge (i, j) , used for (13).

The coefficients for each entry of the weighting matrices Ψ_1, \dots, Ψ_T are obtained by matching (13), (14) and (15) with the model equation (16).

In this work, the binary time series $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_T$ is also modelled using independent SARI models for estimation of (13). Such a modelling approach might not be entirely technically suited for binary-valued time series, but this choice has relevant practical advantages: most programming languages have packages for automatic estimation of the parameters in SARI models, whereas the choice of initial values and estimation of the parameters in most generalised process for binary time series (MacDonald and Zucchini 1997; Kauppi and Saikkonen 2008; Benjamin et al. 2003) is notoriously difficult, which is not desirable when the estimation task should be performed automatically and in parallel over a large set of time series.

The time series modelling extensions presented in this section are computationally expensive, since up to n^2 or nd time series model are fitted, each with complexity $\mathcal{O}(Tk)$, where k is the number of models compared for the purpose of model selection using AICc. This results in a computational cost of $\mathcal{O}(n^2Tk)$ for PIP scores or $\mathcal{O}(ndTk)$ for IPP scores, difficult to manage for large networks. Therefore, with the exception of the IPP COSIE scores, the methodologies proposed in this section do not scale well to large networks, unlike the techniques proposed in Sect. 4. In the case of the IPP COSIE score (12), the cost for prediction of future values of \mathbf{R}_t is only $\mathcal{O}(d^2Tk)$, independent of n .

The methodologies for constructing link prediction scores discussed in this work are summarised in Fig. 1 in a flowchart. In summary, three choices must be made:

- embedding method (collapsed adjacency matrix, standard ASE, omnibus ASE, COSIE embedding);
- combination method (average, *cf.* Sect. 4, or prediction, *cf.* Sect. 5);
- inner product (combination of inner products or inner product of a combination).

Clearly, the same methodology could be applied to *any* embedding method, not necessarily based on the RDPG. Some examples will be given in Sect. 6.2.2.




Fig. 1 Flowchart summarising the construction of a link prediction method using the techniques in Sects. 4 and 5

6 Results

The proposed methods were tested on synthetic data and on real world dynamic networks from different application domains: transportation systems, cyber-security, and co-authorship networks. For all examples, the parameter d was selected using the profile likelihood elbow criterion of Zhu and Ghodsi (2006), unless otherwise specified.

6.1 Simulated data

6.1.1 Seasonal stochastic blockmodel

The performance of the rival link prediction techniques discussed in this article is initially compared on simulated data from stochastic blockmodels. The stochastic blockmodel (Holland et al. 1983) can be interpreted as a special case of a GRDPG (Rubin-Delanchy et al. 2017): each node i is assigned a latent community $z_i \in \{1, \dots, K\}$, with corresponding latent position $\mu_{z_i} \in \mathbb{R}^d$; the probability of a link (i, j) only depends on the community allocation of the two nodes: $\mathbb{P}(A_{ij} = 1) = \mu_{z_i}^\top \mathbf{I}(d_+, d_-) \mu_{z_j}$. To simulate a stochastic blockmodel, a within-community probability matrix $\mathbf{B} = \{B_{ij}\} \in [0, 1]^{K \times K}$ was generated, where $B_{ij} \sim \text{Beta}(1.2, 1.2)$ is the probability of a link between two nodes in communities i and j , and K is the number of communities. The matrix has full rank with probability 1, hence $K = d$. In the simulation, $T = 100$ graph snapshots with $n = 100$ and $K = 5$ were generated. The community allocations were chosen to be time dependent, assuming a seasonality of one week. For each node, community allocations $z_{i,u}$, $u = 0, \dots, s - 1$, with $s = 7$, were sampled uniformly from $\{1, \dots, K\}$. Then, the adjacency matrices were obtained as:

$$\mathbb{P}(A_{ij_t} = 1) = B_{z_{i,t \bmod s}, z_{j,t \bmod s}}, \quad t = 1, \dots, T. \tag{17}$$

Therefore, the link probabilities change over time, with a periodicity of 7 days. The models presented in Sect. 4 were fitted using the first $T' = 80$ snapshots of the graph, with the objective of predicting the remaining $T - T'$ adjacency matrices. The methods that are initially compared are:

- ASE of the *averaged adjacency matrix*, cf. (5) and (6),
- AIP and IPA scores calculated from the *standard ASE*, cf. (7) and (8),
- AIP and IPA scores calculated from the *omnibus embedding*, cf. (9),
- AIP and IPA scores calculated from *COSIE*, cf. (11) and (12).




Fig. 2 Results of the link prediction procedure on the simulated seasonal SBM

The link prediction problem can be framed as a binary classification task. Hence, the performances of the methods presented in this article are evaluated using the area under the receiver operating characteristic (ROC) curve, commonly known as AUC scores. The results are plotted in Fig. 2. The best performance is achieved by the AIP score based on the standard ASE, which outperforms the commonly used method of the collapsed adjacency matrix (5).

It is anticipated that the predictions should be improved upon by the PIP and IPP extensions presented in Sect. 5, since the simulated network has clear dynamics which are not explicitly taken into account using the techniques from Sect. 4. In particular, four methods are discussed:

- ASE of the *predicted adjacency matrix*, cf. (13), with weights obtained from independent seasonal $\text{SARI}(p, b)(P, B)_T$ processes fitted on each binary sequence $A_{ij1}, A_{ij2}, \dots, A_{ijT'}$ such that at least one $A_{ijt} = 1$;
- *PIP scores* calculated from the *standard ASE*, cf. (14), based on the edge time series $\hat{\mathbf{x}}_{i1}^\top \hat{\mathbf{x}}_{j1}, \hat{\mathbf{x}}_{i2}^\top \hat{\mathbf{x}}_{j2}, \dots, \hat{\mathbf{x}}_{iT'}^\top \hat{\mathbf{x}}_{jT'}$ obtained from the individual ASEs on each $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_{T'}$;
- *IPP scores* calculated from the *standard ASE*, cf. (15), based on prediction of the subsequent embeddings $\tilde{\mathbf{X}}_{T'+1}, \tilde{\mathbf{X}}_{T'+2}, \dots$ from the time series of *aligned*¹ individual ASEs $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_{T'}$;
- *IPP scores* calculated from *COSIE*, based on prediction of correction matrices $\tilde{\mathbf{R}}_{T'+1}, \tilde{\mathbf{R}}_{T'+2}, \dots$ from the time series $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_{T'}$, where independent models are fitted to the $d \times d$ time series corresponding to each entry.

¹ The indefinite Procrustes alignment step, described in Appendix A, has been implemented in *python* using *mpy2* and the R codebase developed by Joshua Agterberg, available online at https://github.com/jagterberg/indefinite_procrustes and https://jagterberg.github.io/assets/procrustes_simulation.html.

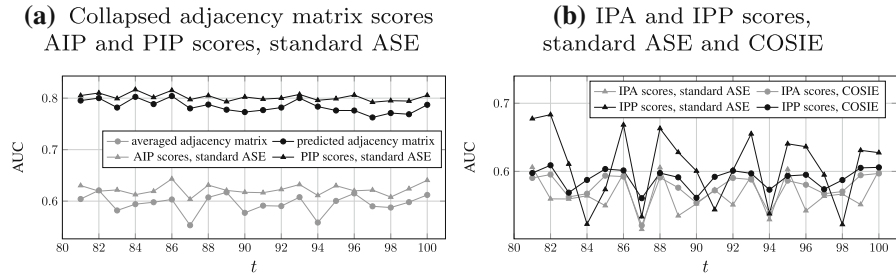


Fig. 3 Comparison between four of the link prediction models in Fig. 2, and their extensions using the methods in Sect. 5, on the synthetic SBM data

The time series models were fit using the function *auto_arima* in the statistical *python* library *pmdarima*, using the corrected AIC criterion to estimate the number of parameters. The results are presented in Fig. 3.

The AIP method (7) which had the best performance in Fig. 2, is significantly improved by the PIP score (14) using time series modelling, and it is overall the only method that reaches values of the AUC well above 0.8. Remarkably, the performance of the predicted adjacency matrix method in (13) outperforms the results based on most of the other methods, despite the issues related to the modelling of binary time series pointed out in Sect. 5. On the other hand, the improvements obtained using the COSIE-based scores and the IPP score (15) seem to be less significant compared to the two other methods. This aspect will also be confirmed on real data examples in the next section. In general, it is clear from the plots in Fig. 3 that adding temporal dynamics to the network via time series modelling is beneficial for link prediction purposes. In particular, including edge specific information from the time series of estimated link probabilities, or from the binary time series of links, has significantly improved link prediction.

6.1.2 Logistic dynamic network model

Next, the effect of the dynamic component on the prediction is evaluated. A directed dynamic graph with $n = 100$ and $T = 100$ is simulated, assuming $A_{ijt} \sim \text{Bernoulli}(v_{ijt})$, $t = 1, \dots, T$, where

$$\text{logit}(v_{ijt}) = b_{ij} + c_{ij}\theta(t - 1), \tag{18}$$

where b_{ij} is a baseline, such that $b_{ij} \sim \text{Uniform}(-6.9, 0)$ independently for all pairs (i, j) , implying $v_{ij1} \in (0.001, 0.5)$. Furthermore, $\theta \in \mathbb{R}$ is a trend parameter common to all the possible edges, and $c_{ij} \in \{-1, 1\}$ is the sign of the trend on each edge, such that $\mathbb{P}(c_{ij} = 1) = \mathbb{P}(c_{ij} = -1) = 1/2$. Note that if $\theta = 0$, the graph does not have any dynamics, whereas if $|\theta|$ increases, the graph dynamics also increases. Note that, asymptotically for $t \rightarrow \infty$, $v_{ijt} \rightarrow 0$ if $c_{ij}\theta \rightarrow -\infty$, and $v_{ijt} \rightarrow 1$ if $c_{ij}\theta \rightarrow \infty$. Dynamic graphs are simulated for $\theta \in \{0, 0.025, 0.05, 0.075\}$, and the *AIP and PIP scores* obtained from *standard ASE* trained on the first T' adjacency matrices are

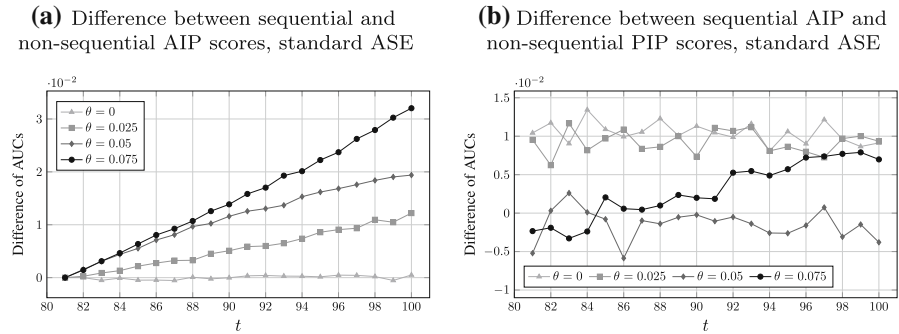


Fig. 4 Difference between sequential AIP scores and (a) non-sequential AIP scores, (b) non-sequential PIP scores, calculated from standard ASE, for $\theta \in \{0, 0.025, 0.05, 0.075\}$ in the logistic dynamic network model (18)

calculated for prediction of the last $T - T'$ network snapshots. The results are then compared with the *AIP scores* with *standard ASE*, sequentially updated when new network snapshots become available, used for a 1-step ahead prediction of the next snapshot. If the network has a relevant dynamic component, the difference between the AUC obtained from the sequential and non-sequential AIP scores increases over time, because the network structure changes and the non-sequential scores cannot capture such evolution. On the other hand, the difference between the sequential AIP scores and the PIP scores should not show an increasing trend over time, since the time dynamics is taken into account via time series modelling.

Figure 4 plots the time series of differences between the sequential AIP scores and the corresponding non-sequential scores (Fig. 4a), and the difference between the sequential AIP scores and the PIP scores (Fig. 4b), for four different values of θ . The plot demonstrates that in presence of a strong dynamic component, the sequential scores outperform non-sequential scores over time, as expected. On the other hand, the PIP scores perform similarly to the sequential scores, and the trend in the differences seem to disappear, up to fluctuations: this is overall remarkable, since the PIP scores are used for up to $(T - T')$ -steps ahead predictions of matrices of scores based only on the initial T' adjacency matrices, whereas the sequential scores also use the last $T - T'$ adjacency matrices sequentially for 1-step ahead predictions.

6.2 Santander bikes

Santander Cycles is a self-service cycle hire scheme operated in central London. Data about usage of the bikes are periodically released by Transport for London.² In this example, data from 7 March, 2018 to 19 March, 2019 were used, for a total of $T = 378$ days. Each bike sharing station is considered as a node, and an undirected edge (i, j, t) is drawn if at least one journey between the stations i and j is completed on day t . The total number of docking stations in London is $n = 840$. The daily graphs are fairly

² The data are freely available at <https://cycling.data.tfl.gov.uk/>, powered by TfL Open Data.




Fig. 5 Results of the link prediction procedure on the Santander Cycles network for different RDPG-based link prediction methods

dense, with an average edge density of approximately 10% across the T networks. The first $T' = 250$ graphs are used as training set.

6.2.1 Averaged scores

Initially, the methods compared are four of the techniques used to produce Fig. 2:

- ASE of the *averaged adjacency matrix*,
- AIP and IPA scores calculated from the *standard ASE*,
- IPA scores calculated from *COSIE*.

For the Santander Cycles network, the results are reported in Fig. 5 for $d = 10$. In Fig. 5, the performance of the classification procedure drops around day 294. This corresponds to Christmas day, which has a different behaviour compared to non-festive days. It is also interesting to point out that COSIE methods tends to perform better on weekdays than weekends, whereas the other methods more accurately predict the links on weekends compared to weekdays. The results of the link prediction procedure in Fig. 5 suggest that the data might not have a long term trend, but only a seasonal component, since the performance does not significantly decrease over time, and the parameters obtained using a training set of size $T' = 250$ seem to reliably predict the structure of the adjacency matrix even at $T = 378$. Overall, this example seems to confirm that the method of AIP scores (7) based on standard ASE has the best performance for link prediction purposes when time dynamics are not included.

6.2.2 Comparison with alternative methods

To provide further comparison, the methods proposed in Sect. 4 are also compared in Fig. 6 to other methods used for link prediction in the literature. In order to demonstrate that the proposed methodology could be readily extended to any embedding technique, the embeddings were calculated from each of the adjacency matrices $\mathbf{A}_1, \dots, \mathbf{A}_T$, and

prediction scores were obtained using the AIP methodology, akin to (7). The alternative methods considered in this section are:

- AIP scores calculated from the *Adamic-Adar* (AA) and *Jaccard* coefficients (used, for example, in Güneş et al. 2016),
- AIP scores calculated from *non-negative matrix factorisation* (NMF; see, for example, Chen et al. 2017; Yu et al. 2017a) with $d = 10$,
- AIP scores calculated from *unsupervised GraphSAGE* (Hamilton et al. 2017), *GCN* (Kipf and Welling 2017) with *Deep Graph Infomax* (DGI, Velickovic et al. 2019), and *Watch Your Step* with *Graph Attention* (WYS-GA, Abu-El-Haija et al. 2018).³ Unsupervised network embeddings $\hat{\mathbf{x}}_{it} \in \mathbb{R}^d$ are obtained independently for each of the T' graphs in the training set, using one of the aforementioned methods with one-hot indicator vectors as node features (when required). For the unsupervised GraphSAGE, a two-layer model with sizes 50 and $d = 10$ is fitted, with 10 walks of length 10 per node, batch size 512 and 10 iterations in each encoder; Adam (Kingma and Ba 2015) is used for learning the embeddings, with learning rate 10^{-2} . Note that the embedding dimension has been chosen to match the dimension of the RDGP-based embeddings. For the GCN, a one-layer network is trained, with layer size $d = 10$ and ReLU activation, optimised by Adam with learning rate 10^{-2} . For WYS-GS, 100 walks per node are used, with $\beta = 0.1$ and $C = 10$ (for the definition of such parameters, see Abu-El-Haija et al. 2018), with embedding dimension $d = 10$; the model is then trained with Adam with learning rate 10^{-3} . For each of the methods, edge features are obtained from the Hadamard product $\hat{\mathbf{x}}_{it} \odot \hat{\mathbf{x}}_{jt}$ between the estimated node embeddings (see, for example, Grover and Leskovec 2016). The link probabilities for each time window are then estimated from T' independent logistic regression models with response A_{ijt} and d predictors of the form $\hat{\mathbf{x}}_{it} \odot \hat{\mathbf{x}}_{jt}$. The link probabilities are then combined using the AIP method (7), and used to predict connections in the last $T - T'$ observed graphs.
- Predictive scores calculated from three methods specifically developed for representation learning of dynamic graphs: the *Deep Embedding Auto-Encoder Method for Dynamic Graphs* (DynGEM, Goyal et al. 2017), the *dynamic graph2vec autoencoder recurrent neural network* model (DyG2V-AERNN, Goyal et al. 2020),⁴ and the *Dynamic Self-Attention Network* (DySAT, Sankar et al. 2020).⁵ The methods were run using the default implementation of the software packages, setting the output embedding dimension to $d = 10$ and the batch size to 100. Link probabilities for the $T - T'$ graphs in the test set were calculated using the same procedure described for GraphSAGE, GCN-DGI and WYS-GA.

The best performance among the alternative methods is achieved by the NMF scores, which achieve an almost equivalent performance to the AIP scores obtained from standard ASE. Slightly inferior performance is achieved by DyG2V-AERNN and DySAT, despite consistently exceeding 0.85 in AUC. GCN, Jaccard, WYS-GA

³ The models were fitted using the implementation in the *python* library *stellargraph* (CSIRO's Data61 2018).

⁴ The models were fitted using the implementation in the *python* library *DynamicGEM* (Goyal et al. 2018).

⁵ The model was fitted using the code in the *GitHub* repository <https://github.com/aravindsankar28/DySAT>.




Fig. 6 Results of the link prediction procedure on the Santander Cycles network for different link prediction methods

and GraphSAGE have a slightly worse performance, but still consistently achieve AUC scores exceeding 0.8. All of the proposed methodologies perform better than the Adamic-Adar and DynGEM, which seem to be largely outperformed by spectral embedding methods. Note that representation learning methods shine in particular when applied to large graphs with rich node features (for example, Hamilton et al. 2017), which is clearly not the case for the Santander cycles network. Furthermore, hyperparameter tuning is necessary to obtain a good link prediction performance, whereas RDGP-based methods only require the embedding dimension as input, and no further tuning is required.

As discussed in Sect. 1, it should be noted that the methodologies of AIP and IPA scores proposed in Sect. 4, and the corresponding PIP and IPP extensions in Sect. 5, could be applied to *any* sequence of individual embeddings, not necessarily obtained using ASE, but also with other embedding methods for static networks, for example NMF, as demonstrated in this section. Since one of the main objectives of this paper is comparing different embedding methods based on RDPGs, the focus in subsequent examples will be *only* on RDPG-based techniques.

6.2.3 Sequential scores

So far, the embeddings learned using the initial T' snapshots have been used to predict *all* the remaining $T - T'$ adjacency matrices. In practical applications, it would be necessary to sequentially update the scores when new snapshots become available over time, improving the predictive performance. This is demonstrated in Fig. 7, where the *AIP scores* for *standard ASE* and the *average adjacency matrix scores* from Fig. 5 are compared with their sequential counterparts, obtained when the model is updated using each new observation A_t in the test set. Clearly, updating the scores sequentially is beneficial, especially towards the end of the test set, whereas the difference between the




Fig. 7 Results for the AIP scores (7) and averaged adjacency matrix scores for the Santander Cycles network, with and without sequential updates

methodologies is negligible in the initial snapshots of the test set. Both methods seem to reliably predict even k -steps ahead network snapshots, since the non-sequential curves in Fig. 7 are fairly close to their sequential counterparts. The method of AIP scores based on standard ASE outperforms the scores based on the averaged adjacency matrix, commonly used in the literature, including sequential settings. Furthermore, the *non-sequential* AIP scores (7) also outperform the *sequential* averaged adjacency matrix. This result is quite remarkable, since the former only use the initial T' snapshots of the network for training, whereas the latter is sequentially updated with the snapshots in the test set.

6.2.4 Predicted scores

The performance of the classifiers can be improved using some of the time series model in Sect. 5. Figure 8a show the results obtained from the prediction of subsequent COSIE correction matrices. The predictive performance is slightly improved by the extended time series models. Again, it is empirically confirmed that adding temporal dynamics is beneficial for the performance of random dot product graph based classifiers.

On the other hand, predicting the subsequent adjacency spectral embeddings from the time series of aligned embeddings $\hat{\mathbf{X}}_1, \hat{\mathbf{X}}_2, \dots, \hat{\mathbf{X}}_{T'}$ does not seem to improve the predictive performance. The results are presented in Figure 8b, and confirms the findings in Figure 3b, where the improvements on the simulated network were less significant compared to other methods. In this case, the time series models are not able to capture the dynamics of the aligned embeddings, and the predictive performance does not improve in AUC.

The limited improvements in the results seem to suggest that the network does not have a strong dynamic component. The tradeoff between performance and the computational effort required to fit multiple independent time series simultaneously, would suggest use of the AIP scores (7) based on standard ASE in practical applications.




Fig. 8 Comparison between two of the link prediction models in Figure 5, and their extensions using the methods in Sect. 5, on the Santander Cycles network

6.3 Los Alamos National Laboratory dataset

The unified host and network dataset (Turcotte et al. 2018) released by the Los Alamos National Laboratory (LANL) consists of a collection of network flow and host event logs generated from their machines running Microsoft Windows. From the host event logs, 90 daily user-authentication bipartite graphs have been constructed, writing $A_{ijt} = 1$ if the user i initiates a connection authenticating to computer j , on day t . This graph is known as the *user – destination IP* graph. A total of $n_1 = 12, 222$ users, $n_2 = 5047$ hosts, and 85, 020 pairs (i, j) are observed, corresponding to approximately 0.137% of all possible links.

6.3.1 Averaged scores and subsampling

The first $T' = 56$ matrices are used as training set. Note that it is computationally difficult to calculate $n_1 \times n_2$ scores for each adjacency matrix, and storing such large dense matrices in memory is also not efficient. Therefore, an estimate of the AUC can be obtained by subsampling the negative class at random from the zeroes in the test set adjacency matrices. Two subsampling techniques are used to construct the negative class for prediction of \mathbf{A}_t :

- (1) the negative class is constructed by sampling pairs (i, j) such that $A_{ijt} = 0$,
- (2) the negative class contains randomly selected pairs (i, j) such that $A_{ijt} = 0$, and all pairs (i, j) such that $A_{ijt} = 0$ and $A_{ijj't'} = 1$ for at least 1 value of $t' \in \{1, \dots, T\}$.

For simplicity, the two techniques are denoted with the numbers (1) and (2) in Fig. 9, which reports the results for the 6 methods considered in Fig. 2 in Sect. 6.1. The former subsampling technique provides an estimate of the ROC curve for the entire matrix, since the scores are sampled at random from the distribution of all scores. On the other hand, the latter method includes in the negative class more elements that tend to have associated high scores, represented by the pairs (i, j, t) such that $A_{ijt} = 0$ and $A_{ijj't'} = 1$ for at least 1 value of $t' \in \{1, \dots, T\}$, giving an unbalanced sampling procedure and therefore a biased estimate of the ROC curve. Clearly, higher AUC scores are obtained using the first subsampling procedure.

Interestingly, in Fig. 9a, the COSIE-based AIP scores seem to have the best predictive performance across the different methods. In particular, COSIE scores tend

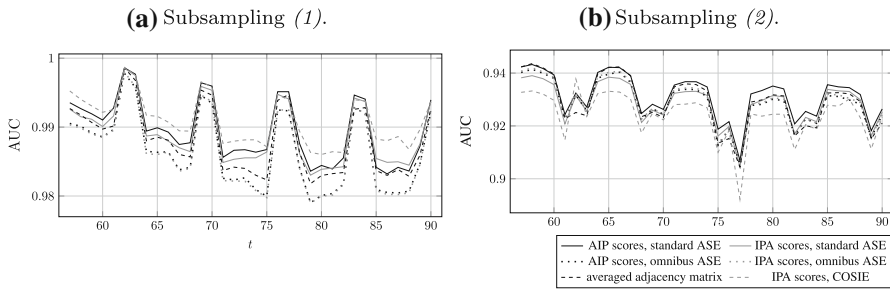


Fig. 9 Results of the link prediction procedure on the LANL network. AUCs are calculated from $\approx 150,000$ links per graph

to largely outperform the other methods during weekdays, whereas the performance during weekends seems almost equivalent, and sometimes inferior, to the AIP scores (7) based on the standard ASE. On the other hand, in Fig. 9b, COSIE scores have the worst performance among the methods, except a spike on day 62. In Fig. 9b, AIP scores (7) based on the standard ASE once again give the best predictive performance. The results of Fig. 9b are of particular interest since these allow for a comparison of the classifiers on a more challenging negative class compared to Fig. 9a. Therefore, it is reasonable to conclude that the AIP scores (7) emerge again as the most suitable method for link prediction based on random dot product graphs.

6.3.2 Significance of the difference between methods

It could be argued that the differences between the methodologies do not appear to be significant, and might be due to the subsampling scheme. In order to assess significance of the differences, the subsampling procedure was repeated for $M = 100$ times, and the corresponding AUCs were calculated. Figure 10 plots the estimated 95% confidence intervals for the difference between the AUCs obtained using different methods. Figure 10a uses the IPA scores calculated with COSIE as reference, and the AUCs are obtained with subsampling (1). Similarly, Fig. 10b uses subsampling (2), and the AIP scores calculated with standard ASE are used as reference. Note that the width of the confidence intervals in Fig. 10 is barely visible, since the standard deviations of the AUC scores are $< 10^{-4}$. This is because the number of subsamples is large enough to obtain extremely precise estimates of the AUC. Since the confidence intervals do not contain zero at any of the time points, the pairwise differences between the performance of different methodologies appear to be statistically significant. Similar confidence intervals, and corresponding p -values $< 10^{-6}$ were observed for all the comparisons in the current and next sections, suggesting that the subsampling procedure for estimation of the AUC is robust.

6.3.3 Streaming link prediction and sequential scores

For this example, it is also demonstrated how the proposed methods can be extended for streaming applications. For example, the method of combining inner products

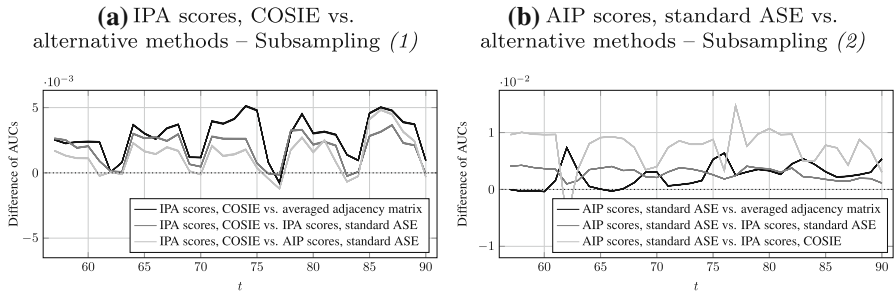


Fig. 10 95% confidence intervals of the difference between AUCs obtained alternative methods, obtained using subsampling (1) in (a) and (2) in (b). Reference method: (a) IPA scores, COSIE, (b) AIP scores, standard ASE

of individual embeddings is particularly suitable for implementation in a streaming fashion, since the inner products are easily updated on the run when new snapshots of the graph are observed. A demonstration using the method of *fixed forgetting factors* (Gama et al. 2013) is given here. The matrix of scores S_t for prediction of A_{t+1} is updated in streaming as follows:

$$w_t = \lambda w_{t-1} + 1, \quad S_t = \left(1 - \frac{1}{w_t}\right) S_{t-1} + \frac{1}{w_t} \hat{X}_t \hat{X}_t^T, \quad (19)$$

starting from $w_1 = 1$ and $S_1 = \hat{X}_1 \hat{X}_1^T$. The forgetting factor $\lambda \in [0, 1]$ is usually chosen to be close to 1 (Gama et al. 2013). Note that $\lambda = 1$ corresponds to the sequential AIP scores in (7), calculated as in Fig. 7, whereas smaller values of λ give more weight to recent observations. In particular, $\lambda = 0$ only gives weight to the most recent snapshot. Schemes similar to (19) could be also implemented to update the collapsed adjacency matrix (5), obtaining the scores from its ASE at each point in time, or to update the matrix R_t in (4) for the COSIE embeddings, or the rotated embeddings \hat{X}_t . The forgetting factor approach might be interpreted as a simplification of the time series scheme proposed in Sect. 5, where the same weight is given to each edge.

The results for the entire observation period for a range of different values of λ are plotted in Fig. 11. The best performance is achieved with the forgetting factor approach with $\lambda \in [0.4, 0.8]$. The performance for $\lambda = 0$ clearly drops around the weekends, since the network has a seasonal component which is not accounted for in the prediction. The difference between the curves for $\lambda = 1$ and $\lambda < 1$ suggests that the graph has temporal dynamics which is captured by the forgetting factor approach, which down-weights past observations in favour of more recent snapshots of the graph. This impression is confirmed by Fig. 12, which shows the sequential and non-sequential AIP scores based on the standard ASE, akin to Fig. 7. The predictive performance slightly deteriorates for snapshots that are further away in time, whereas 1-step ahead predictions based on the sequential scores consistently give better results.




Fig. 11 Results for the AIP scores (7) on the LANL network with streaming updates. AUCs are calculated from $\approx 150,000$ links, using subsampling (2)




Fig. 12 Results for the AIP scores (7) for the LANL network, with and without sequential updates. AUCs are calculated from $\approx 150,000$ links per graph, using subsampling (1)

6.3.4 Predicted scores

The performance can be again improved using the time series models in Sect. 5. Figure 13 shows the results obtained using the two different subsampling schemes for three methods:

- ASE of the *averaged* and *predicted adjacency matrix*,
- *AIP* and *PIP* scores calculated from *standard ASE*,
- *IPA* and *IPP* scores calculated from *COSIE*.

From Fig. 13a, 13c and 13e, it is evident that the extensions do not improve the performance of the classifier when the subsampling scheme (1) is used. On the other hand, Fig. 13b, 13d and 13f show that relevant improvements (especially on day 62) are obtained when the subsampling method (2) is used, which represents a more difficult classification task. Again, the results confirm that the performance of RDPGs for link prediction can be enhanced by time series models.

6.4 Imperial College network flow data

The methodologies proposed in Sect. 4 are also tested on a large computer network, demonstrating that the spectral embedding techniques are scalable to graphs of large

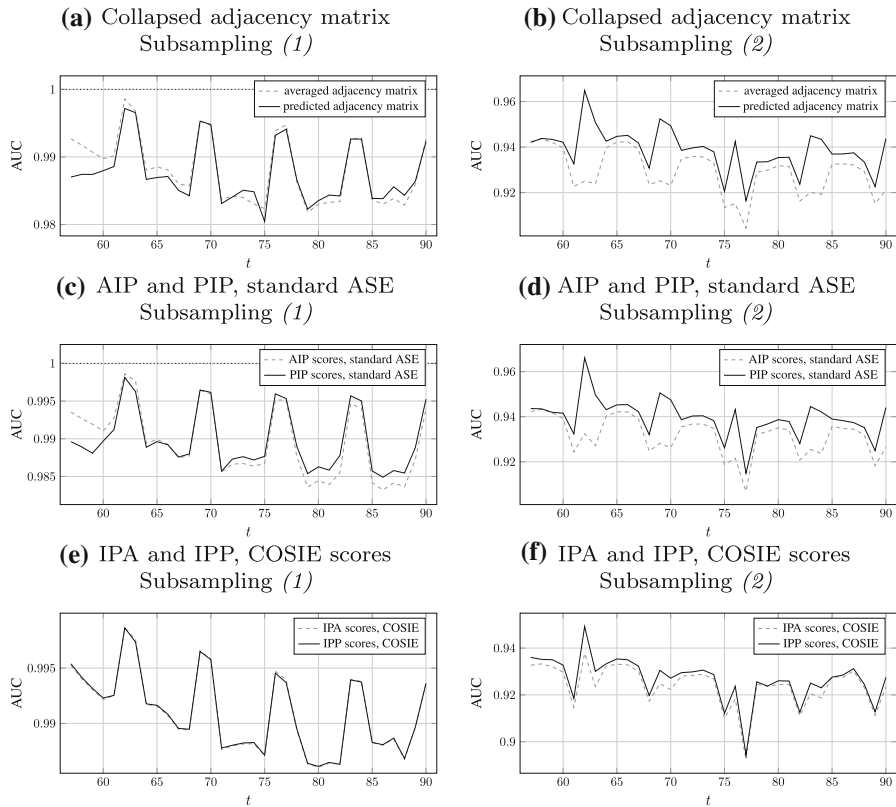


Fig. 13 Comparison between three of the link prediction models in Fig. 9, and their extensions using the methods in Sect. 5, on the LANL network. AUCs are calculated from $\approx 150,000$ links per graph

size. A directed graph with $n = 95,220$ has been constructed using network flow data collected at Imperial College London, limiting the connections to internal hosts only. The data have been collected over $T = 47$ days, between 1st January and 16th February 2020. An edge between a client i and a server j is drawn on day t if the two hosts have connected at least once during the corresponding day. The number of edges ranges from a minimum of 344,565, observed on 1st January, to a maximum of 912,984 on 6th February.

The results are plotted in Fig. 14. In this case, the best performance is achieved by the COSIE scores, similarly to the LANL network in Fig. 9a. The traditional methodology of the averaged adjacency matrix is also outperformed by the AIP scores, which supports the results obtained in all the previous real data examples. As pointed out in Sect. 5, fitting the time series models on the sequence $\hat{\mathbf{R}}_1, \dots, \hat{\mathbf{R}}_{T'}$ of COSIE weighting matrices is inexpensive, and it is therefore possible to scale the time series extension of the IPA scores to a fairly large network. The results of this procedure are plotted in Fig. 15, which shows again that the time series extension is beneficial for link prediction purposes.




Fig. 14 Results of the link prediction procedure on the Imperial College NetFlow network. AUCs are calculated from 6.7 million links per graph




Fig. 15 IPA and IPP scores using COSIE embeddings on the Imperial College NetFlow network. AUCs are calculated from 6.7 million links per graph

6.5 DBLP dataset

Finally, the methodology is evaluated on a version of the DBLP co-authorship dataset⁶, extensively used in the computer science literature. The DBLP network is undirected, with $n = 1, 258, 753$ nodes, corresponding to authors of papers in computer science, for a total of 7, 654, 336 undirected co-authorship edges over $T = 15$ years, starting in 2000. An edge between two authors i and j is drawn if they co-authored a paper in year t . The network is extremely sparse, with edge densities ranging from $1.90 \cdot 10^{-7}$ in 2000 to $7.41 \cdot 10^{-7}$ to 2015. The number of co-authorships consistently increases over the years, corresponding to a steady increase in the number of publications in computer science journals and conferences.

The initial $T' = 10$ graphs, from 2000 to 2009, are used as training set, whereas the last $T - T'$ graphs, from 2010 to 2014, correspond to the test set. The results are presented in Table 1, for the same models examined in Sect. 6.4 on the Imperial College network, for the two different subsampling schemes. Unsurprisingly, the AIP scores with standard ASE achieve the best performance. Limited improvement is obtained considering the time series extension on the IPA scores with COSIE.

The performance of the RDPG on the DBLP network is significantly worse than the other examples considered in this section, suggesting that RDPG-based methods might not be the most appropriate technique for this graph. This is because the graph is

⁶ The data are freely available at <https://projects.csail.mit.edu/dnd/DBLP/>.

Table 1 Results of the link prediction procedure on the DBLP co-authorship network. AUCs are calculated from approximately 17 million links per graph

Methodology	Predicted graph				
	2010	2011	2012	2013	2014
(a) Subsampling (1)					
averaged adjacency matrix	0.6912	0.6431	0.6095	0.5847	0.5777
AIP scores, standard ASE	0.7063	0.6685	0.6413	0.6235	0.6184
IPA scores, standard ASE	0.6637	0.6291	0.6066	0.5906	0.5856
IPA scores, COSIE	0.6744	0.6378	0.6123	0.5955	0.5903
IPP scores, COSIE	0.6893	0.6511	0.6252	0.6072	0.6020
(b) Subsampling (2).					
averaged adjacency matrix	0.6690	0.6202	0.5857	0.5598	0.5508
AIP scores, standard ASE	0.6765	0.6369	0.6083	0.5892	0.5807
IPA scores, standard ASE	0.6416	0.6058	0.5822	0.5653	0.5577
IPA scores, COSIE	0.6502	0.6124	0.5862	0.5687	0.5612
IPP scores, COSIE	0.6629	0.6234	0.5964	0.5773	0.5690

extremely sparse, and a large number of new nodes appears in the network over time. Such authors, before their first published paper, are represented by rows and columns filled with zeros in the adjacency matrix, acting as disconnected components in the graph. Therefore, learning the embedding for such nodes is particularly difficult for the RDPG, since a limited (or null) history of co-authorships is available.

6.6 Discussion and summary of results

The main methodological contribution of this article is an adaptation for temporal link prediction of existing RDPG-based methods for joint inference on multiple graphs. Also, this article proposes techniques to capture the temporal dynamics of the observed graphs, combining the link probabilities and embeddings obtained via spectral methods with time series models. As demonstrated in Sect. 6.2.2, the proposed methods for combination of individual embeddings, and the extension based on time series modes, can be applied on any embedding method for static graphs, not necessarily restricted to RDPG models. The proposed methods have been applied on different synthetic and real-world datasets of different complexity, which highlighted the strengths and weaknesses of the proposed methodologies. In general, in most simulated and real-world networks analysed in this work, the AIP scores (7) based on standard ASE appear to consistently achieve a very good link prediction performance in terms of AUC scores.

Both simulations and applications on real data demonstrated that adding temporal dynamics to the estimated link probabilities via time series modelling is beneficial for link prediction purposes (*cf.* Sects. 6.1, 6.2.4 and 6.3.4). On the other hand, the time series extension is computationally expensive for most RDPG-based models, except the IPP scores based on COSIE scores. The extensions provide significant benefits

only if the network presents a strong dynamic component, as demonstrated in the simulation on the logistic dynamic network model in Sect. 6.1.2.

The application to the Santander bikes data (*cf.* Sect. 6.2) highlighted that the random dot product graphs are an excellent option for link prediction for fairly small graphs without node or edge features. In such cases, the simplicity of RDPG-based models appears to overcome the advantages of deep learning methods (*cf.* Sect. 6.2.2), which instead shine when applied to large graphs with rich node and edge features (for example, Hamilton et al. 2017). Furthermore, it must be remarked that RDPG-based methods require minimal tuning (essentially, only the choice of the latent dimension d), whereas alternative state-of-the-art models require computationally expensive hyperparameter tuning procedures.

The applications on the LANL, ICL and DBLP networks (*cf.* Sects. 6.3, 6.4 and 6.5) demonstrated that the methodologies of AIP and IPA scores are scalable to fairly large graphs, and a good performance is achieved when the set of nodes is stable over time. It has also been shown that the proposed methodologies, in particular the AIP scores, are well suited to streaming applications (*cf.* Sect. 6.3.3). The main limitation of the proposed link prediction framework is its inability to easily deal with new nodes appearing in the network, as exemplified by the application on the DBLP network in Sect. 6.5.

7 Conclusion

In this paper, link prediction techniques based on random dot product graphs have been presented, discussed and compared. In particular, link prediction methods based on sequences of embeddings, COSIE models, and omnibus embeddings have been considered. Applications on simulated and real world data have shown that one of the most common approaches used in the literature, the decomposition of a collapsed adjacency matrix $\tilde{\mathbf{A}}$, is usually outperformed by other methods for multiple graph inference in random dot product graphs.

Estimating link probabilities with the average of the inner products from sequences of individual embeddings of different snapshots of the graph has given the best performance in terms of AUC scores across multiple datasets. This result is particularly appealing for practical applications: calculating the individual ASEs is computationally inexpensive using algorithms for large sparse matrices, and the method seems particularly suitable for implementation in a streaming fashion, since the average inner product could be easily updated on the run when new snapshots of the graph are observed, as demonstrated in Sect. 6.3.

The methods discussed in the article have then been further extended to include temporal dynamics, using time series models. The extensions have shown improvements over standard random dot product graph based link prediction techniques, especially when the graph exhibits a strong dynamic component. The techniques presented in this article could also be readily extended to *any* embedding method for static graphs, following the framework for calculating AIP and IPA scores presented in Section 4, and the PIP and IPP extensions in Section 5. Therefore, our methodology is not only applicable to RDPG embeddings, particularly attractive for their theoretical proper-

ties and ease of implementation, but also to modern static embedding methods for graph adjacency matrices, for example graph neural network techniques like GCN or GraphSAGE. Overall, this article provides valuable guidelines for practitioners for using random dot product graphs as tools for link prediction in networks, providing insights into the predictive capability of such statistical network models.

Acknowledgements FSP acknowledges funding from the EPSRC.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Generalised Procrustes alignment of individual embeddings

As discussed in Sect. 4, for prediction of the future latent positions based on individual embeddings $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_T$, it is first necessary to align the embeddings. This section discusses a popular method for aligning two matrices: Procrustes analysis (Dryden and Mardia 2016) and its generalisation to T matrices (Gower 1975). The alignment step is required because the latent positions \mathbf{X}_t of a single graph are not identifiable up to orthogonal transformations, which leave the inner product unchanged: for an orthogonal matrix $\mathbf{\Omega}_t$, $(\mathbf{X}_t \mathbf{\Omega}_t)(\mathbf{X}_t \mathbf{\Omega}_t)^\top = \mathbf{X}_t \mathbf{X}_t^\top$. Therefore $\hat{\mathbf{X}}_1, \dots, \hat{\mathbf{X}}_T$ only represent estimates of a rotation of the embeddings. Given two embeddings $\hat{\mathbf{X}}_1, \hat{\mathbf{X}}_2 \in \mathbb{R}^{n \times d}$, Procrustes analysis aims to find the optimal rotation of $\hat{\mathbf{X}}_2$ on $\hat{\mathbf{X}}_1$. The following minimisation criterion is utilised:

$$\min_{\mathbf{\Omega}} \left\| \hat{\mathbf{X}}_1 - \hat{\mathbf{X}}_2 \mathbf{\Omega} \right\|_F, \tag{20}$$

where $\mathbf{\Omega}$ is an orthogonal matrix, and $\|\cdot\|_F$ denotes the Frobenius norm. The solution of the minimisation problem has been derived in Schönemann (1966), and is based on the SVD decomposition $\hat{\mathbf{X}}_2^\top \hat{\mathbf{X}}_1 = \tilde{\mathbf{U}} \tilde{\mathbf{D}} \tilde{\mathbf{V}}^\top$. The solution is $\mathbf{\Omega}^* = \tilde{\mathbf{U}} \tilde{\mathbf{V}}^\top$, and it follows that the optimal rotation of $\hat{\mathbf{X}}_2$ onto $\hat{\mathbf{X}}_1$ is $\hat{\mathbf{X}}_2 \tilde{\mathbf{U}} \tilde{\mathbf{V}}^\top$.

Similarly, a set of T embeddings $\hat{\mathbf{X}}_t \in \mathbb{R}^{n \times r}$, $t = 1, \dots, T$ can be superimposed using *generalised Procrustes analysis* (GPA, Gower 1975), which uses the minimisation criterion:

$$\min_{\mathbf{\Omega}_j} \sum_{j=1}^T \left\| \hat{\mathbf{X}}_j \mathbf{\Omega}_j - \tilde{\mathbf{X}} \right\|_F^2 \text{ s.t. } \sum_{j=1}^T \mathbf{S}^2(\hat{\mathbf{X}}_j) = \sum_{j=1}^T \mathbf{S}^2(\hat{\mathbf{X}}_j \mathbf{\Omega}_j), \tag{21}$$

where, similarly to (20), $\mathbf{\Omega}_j$ are orthogonal matrices. Additionally, $\tilde{\mathbf{X}} \in \mathbb{R}^{n \times r}$ is a reference matrix, shared across the T matrices, and $\mathbf{S}(\cdot)$ is the *centroid size* $\mathbf{S}(\mathbf{M}) = \|(\mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top) \mathbf{M}\|_F$. The GPA algorithm solves (21) by iterating standard Procrustes

analysis (Dryden and Mardia 2016), after a suitable initialisation of the reference matrix:

1. *update* the embeddings $\hat{\mathbf{X}}_t$, performing a standard Procrustes superimposition of each $\hat{\mathbf{X}}_j$ on $\tilde{\mathbf{X}}$:

$$\hat{\mathbf{X}}_j \leftarrow \hat{\mathbf{X}}_j \tilde{\mathbf{U}}_j \tilde{\mathbf{V}}_j^\top, \quad (22)$$

where $\hat{\mathbf{X}}_j^\top \tilde{\mathbf{X}} = \tilde{\mathbf{U}}_j \tilde{\mathbf{D}}_j \tilde{\mathbf{V}}_j^\top$;

2. *update* the reference embedding: $\tilde{\mathbf{X}} = \sum_{t=1}^T \hat{\mathbf{X}}_t / T$;
3. *repeat* steps 1 and 2 until the difference between two consecutive values of (21) is within a tolerance η .

The final value of the reference embedding $\tilde{\mathbf{X}}$ can be interpreted as the average of rotations of the initial embeddings. The alignment step increases the computational cost of the operations described in Sect. 4 by a factor of $\mathcal{O}(nd^2)$, caused by the repeated SVD decompositions and matrix multiplications in the GPA algorithm.

When $d_- \neq 0$ in the GRDPG setting, the problem is known as *indefinite Procrustes problem*, and does not have closed form solution (Kintzel 2005). In this setting, the criterion (20) must be optimised numerically for $\Omega \in \mathbb{O}(d_+, d_-)$, the indefinite orthogonal group with signature (d_+, d_-) . The optimisation routine could be applied iteratively in the GPA algorithm to obtain an *indefinite GPA*.

On the other hand, for directed and bipartite graphs, the criteria (20) and (21) must be optimised jointly for the two embeddings obtained using DASE in Definition 3. For two embeddings $(\hat{\mathbf{X}}_1, \hat{\mathbf{Y}}_1)$ and $(\hat{\mathbf{X}}_2, \hat{\mathbf{Y}}_2)$ obtained from a directed graph, the solution is simply given by aligning the stacked matrices $[\hat{\mathbf{X}}_1^\top, \hat{\mathbf{Y}}_1^\top]^\top \in \mathbb{R}^{2n \times r}$ and $[\hat{\mathbf{X}}_2^\top, \hat{\mathbf{Y}}_2^\top]^\top \in \mathbb{R}^{2n \times r}$ using (20). The procedure could be also iterated for more than two embeddings to obtain a *joint generalised Procrustes algorithm*.

References

- Abu-El-Haija S, Perozzi B, Al-Rfou R, Alemi AA (2018) Watch your step: Learning node embeddings via graph attention. In: Advances in Neural Information Processing Systems. vol. 31. Curran Associates, Inc
- Arroyo-Reli3n JD, Athreya A, Cape J, Chen G, Priebe CE, Vogelstein JT (2020) Inference for multiple heterogeneous networks with a common invariant subspace. *Journal of Machine Learning Research* (to appear)
- Arroyo-Reli3n JD, Kessler D, Levina E, Taylor SF (2019) Network classification with applications to brain connectomics. *Ann Appl Stat* 13(3):1648–1677
- Athreya A, Fishkind DE, Tang M, Priebe CE, Park Y, Vogelstein JT, Levin K, Lyzinski V, Qin Y, Sussman DL (2018) Statistical inference on random dot product graphs: a survey. *J Mach Learn Res* 18(226):1–92
- Benjamin MA, Rigby RA, Stasinopoulos DM (2003) Generalized autoregressive moving average models. *J Am Stat Assoc* 98(461):214–223
- Brockwell PJ, Davis RA (1987) Springer series in statistics. Time series: theory and methods. Springer, New York
- Cai H, Zheng VW, Chang K (2018) A comprehensive survey of graph embedding: problems, techniques, and applications. *IEEE Trans Knowl Data Eng* 30(9):1616–1637
- Charlin L, Ranganath R, McInerney J, Blei DM (2015) Dynamic poisson factorization. In: Proceedings of the 9th ACM conference on recommender systems. pp. 155–162

- Chen B, Li F, Chen S, Hu R, Chen L (2017) Link prediction based on non-negative matrix factorization. *PLOS ONE* 12(8):1–18
- Chen H, Li J (2018) Exploiting structural and temporal evolution in dynamic link prediction. In: Proceedings of the 27th ACM International conference on information and knowledge management. pp. 427–436
- CSIRO's Data61: Stellargraph machine learning library. <https://github.com/stellargraph/stellargraph> (2018)
- Deng D, Shahabi C, Demiryurek U, Zhu L, Yu R, Liu Y (2016) Latent space model for road networks to predict time-varying traffic. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 1525–1534
- Dong X, Frossard P, Vandergheynst P, Nefedov N (2014) Clustering on multi-layer graphs via subspace analysis on Grassmann manifolds. *IEEE Trans Signal Process* 62(4):905–918
- Dryden IL, Mardia KV (2016) Statistical shape analysis, with applications in R. John Wiley and Sons, Hoboken
- Dunlavy DM, Kolda TG, Acar E (2011) Temporal link prediction using matrix and tensor factorizations. *ACM Trans Knowl Discov Data* 5(2):1–27
- Durante D, Dunson DB (2014) Nonparametric Bayes dynamic modelling of relational data. *Biometrika* 101(4):883–898
- Durante D, Dunson DB (2018) Bayesian inference and testing of group differences in brain networks. *Bayesian Anal* 13(1):29–58
- Gama J, Sebastião R, Rodrigues PP (2013) On evaluating stream learning algorithms. *Mach Learn* 90(3):317–346
- Gao S, Denoyer L, Gallinari P (2011) Temporal link prediction by integrating content and structure information. In: Proceedings of the 20th ACM International conference on information and knowledge management. pp. 1169–1174
- Ghashami M, Liberty E, Phillips JM (2016) Efficient frequent directions algorithm for sparse matrices. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 845–854
- Ginestet CE, Li J, Balachandran P, Rosenberg S, Kolaczyk ED (2017) Hypothesis testing for network data in functional neuroimaging. *Ann Appl Stat* 11(2):725–750
- Gower JC (1975) Generalized Procrustes analysis. *Psychometrika* 40(1):33–51
- Goyal P, Kamra N, He X, Liu Y (2017) DynGEM: Deep embedding method for dynamic graphs. In: IJCAI International Workshop on Representation Learning for Graphs,
- Goyal P, Rokka Chhetri S, Canedo A (2020) dyngraph2vec: capturing network dynamics using dynamic graph representation learning. *Knowl-Based Syst* 187:104816
- Goyal P, Rokka Chhetri S, Mehrabi N, Ferrara E, Canedo A (2018) DynamicGEM: a library for dynamic graph embedding methods. arXiv e-prints
- Grover A, Leskovec J (2016) node2vec: Scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 855–864
- Güneş İ, Gündüz-Öğüdücü Ş, Çataltepe Z (2016) Link prediction using time series of neighborhood-based node similarity scores. *Data Min Knowl Discov* 30(1):147–180
- Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. In: Proceedings of the 31st International conference on neural information processing systems. pp. 1025–1035
- Hoff PD, Raftery AE, Handcock MS (2002) Latent space approaches to social network analysis. *J Am Stat Assoc* 97(460):1090–1098
- Holland PW, Laskey KB, Leinhardt S (1983) Stochastic blockmodels: first steps. *Soc Netw* 5(2):109–137
- Hosseini SA, Khodadadi A, Alizadeh K, Arabzadeh A, Farajtabar M, Zha H, Rabiee HR (2020) Recurrent Poisson factorization for temporal recommendation. *IEEE Trans Knowl Data Eng* 32(1):121–134
- Hyndman R, Khandakar Y (2008) Automatic time series forecasting: the forecast package for R. *J Stat Softw* 27(3):1–22
- Ishiguro K, Iwata T, Ueda N, Tenenbaum JB (2010) Dynamic infinite relational model for time-varying relational data analysis. *Adv Neural Inf Process Syst* 23:919–927
- Jeske DR, Stevens NT, Tartakovsky AG, Wilson JD (2018) Statistical methods for network surveillance. *Appl Stoch Models Bus Ind* 34(4):425–445
- Jones A, Rubin-Delanchy P (2021) The multilayer random dot product graph
- Kauppi H, Saikkonen P (2008) Predicting U.S. recessions with dynamic binary response models. *Rev Econ Stat* 90(4):777–791
- Kazemi SM, Goel R, Jain K, Kobyzev I, Sethi A, Forsyth P, Poupart P (2020) Representation learning for dynamic graphs: A survey. *J Mach Learn Res* 21(70):1–73

- Khosla M, Setty V, Anand A (2021) A comparative study for unsupervised network representation learning. *IEEE Trans Knowl Data Eng* 33(5):1807–1818
- Kim Y, Levina E (2019) Graph-aware modeling of brain connectivity networks. arXiv e-prints [arXiv:1903.02129](https://arxiv.org/abs/1903.02129)
- Kingma DP, Ba J (2015) Adam: a method for stochastic optimization. In: Bengio Y, LeCun Y (eds) 3rd International conference on learning representations. ICLR. San Diego, CA, USA
- Kintzel U (2005) Procrustes problems in finite dimensional indefinite scalar product spaces. *Linear Algebra Appl* 402:1–28
- Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: 5th International conference on learning representations, ICLR 2017, Conference Track Proceedings
- Krivitsky PN, Handcock MS (2014) A separable model for dynamic networks. *J Royal Stat Soc: Series B (Statistical Methodology)* 76(1):29–46
- Kumar S, Zhang X, Leskovec J (2019) Predicting dynamic embedding trajectory in temporal interaction networks. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery and data mining. pp. 1269–1278. KDD '19
- Levin K, Athreya A, Tang M, Lyzinski V, Park Y, Priebe CE (2017) A central limit theorem for an omnibus embedding of multiple random graphs and implications for multiscale network inference. arXiv e-prints [arXiv:1705.09355](https://arxiv.org/abs/1705.09355)
- Li X, Du N, Li H, Li K, Gao J, Zhang A (2014) A deep learning approach to link prediction in dynamic networks. In: Proceedings of the 2014 SIAM International conference on data mining. pp. 289–297
- Liben-Nowell D, Kleinberg J (2007) The link-prediction problem for social networks. *J Am Soc Inf Sci Technol* 58(7):1019–1031
- Liu Z, Zhou D, He J (2019) Towards explainable representation of time-evolving graphs via spatial-temporal graph attention networks. In: Proceedings of the 28th ACM international conference on information and knowledge management. pp. 2137–2140. CIKM '19
- Lü L, Zhou T (2011) Link prediction in complex networks: a survey. *Phys A: Stat Mech Appl* 390(6):1150–1170
- MacDonald IL, Zucchini W (1997) Hidden Markov and other models for discrete-valued time series. Taylor & Francis, Milton Park
- Menon AK, Elkan C (2011) Link prediction via matrix factorization. *Joint Eur Conf Mach Learn Knowl Discov Datab*. Springer, Berlin, pp 437–452
- Metelli S, Heard NA (2019) On Bayesian new edge prediction and anomaly detection in computer networks. *Ann Appl Stat* 13(4):2586–2610
- Neil J, Hash C, Brugh A, Fisk M, Storlie CB (2013) Scan statistics for the online detection of locally anomalous subgraphs. *Technometrics* 55(4):403–414
- Nguyen GH, Lee JB, Rossi RA, Ahmed NK, Koh E, Kim S (2018) Continuous-time dynamic network embeddings. In: Companion proceedings of the the web conference 2018. pp. 969–976. WWW '18
- Nielsen AM, Witten D (2018) The multiple random dot product graph model. arXiv e-prints [arXiv:1811.12172](https://arxiv.org/abs/1811.12172)
- Perozzi B, Al-Rfou R, Skiena S (2014) Deepwalk: Online learning of social representations. In: Proceedings of the 20th ACM SIGKDD International conference on knowledge discovery and data mining. pp. 701–710. KDD '14
- Priebe CE, Park Y, Tang M, Athreya A, Lyzinski V, Vogelstein JT, Qin Y, Cocanougher B, Eichler K, Zlatic M, Cardona A (2017) Semiparametric spectral modeling of the drosophila connectome
- Qiu J, Dong Y, Ma H, Li J, Wang K, Tang J (2018) Network embedding as matrix factorization: Unifying DeepWalk, LINE, PTE, and node2vec. In: Proceedings of the eleventh ACM International conference on web search and data mining. pp. 459–467. WSDM '18, Association for Computing Machinery
- Qu L, Zhu H, Duan Q, Shi Y (2020) Continuous-time link prediction via temporal dependent graph neural network. In: Proceedings of the web conference 2020. pp. 3026–3032. WWW '20
- Rubin-Delanchy P, Priebe CE, Tang M, Cape J (2017) A statistical interpretation of spectral embedding: the generalised random dot product graph. arXiv e-prints
- Sankar A, Wu Y, Gou L, Zhang W, Yang H (2020) DySAT: deep neural representation learning on dynamic graphs via self-attention networks. In: Proceedings of the 13th International conference on web search and data mining. pp. 519–527
- Sarkar P, Chakrabarti D, Jordan M (2014) Nonparametric link prediction in large scale dynamic networks. *Electr J Stat* 8(2):2022–2065

- Sarkar P, Moore AW (2006) Dynamic social network analysis using latent space models. *Adv Neural Inf Process Syst* 18:1145–1152
- Schein A, Paisley J, Blei DM, Wallach H (2015) Bayesian Poisson tensor factorization for inferring multilateral relations from sparse dyadic event counts. In: *Proceedings of the 21th ACM SIGKDD International conference on knowledge discovery and data mining*. pp. 1045–1054
- Scheinerman ER, Tucker K (2010) Modeling graphs using dot product representations. *Comput Stat* 25(1):1–16
- Schönemann PH (1966) A generalized solution of the orthogonal Procrustes problem. *Psychometrika* 31(1):1–10
- Sewell DK, Chen Y (2015) Latent space models for dynamic networks. *J Am Stat Assoc* 110(512):1646–1657
- Sharan U, Neville J (2008) Temporal-relational classifiers for prediction in evolving domains. In: *Proceedings of the 2008 Eighth IEEE International conference on data mining*. pp. 540–549
- Shiga M, Mamitsuka H (2012) A variational Bayesian framework for clustering with multiple graphs. *IEEE Trans Knowl Data Eng* 24(4):577–590
- Tang W, Lu Z, Dhillon IS (2009) Clustering with multiple graphs. In: *Proceedings of the 2009 Ninth IEEE International conference on data mining*. pp. 1016–1021. ICDM '09, IEEE Computer Society, Washington, DC, USA
- Turcotte MJM, Kent AD, Hash C (2018) Unified host and network data set, chap. 1, pp. 1–22. World Scientific
- Velickovic P, Fedus W, Hamilton WL, Liò P, Bengio Y, Hjelm RD (2019) Deep graph infomax. In: 7th International conference on learning representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019
- Wang D, Cui P, Zhu W (2016) Structural deep network embedding. In: *Proceedings of the 22nd ACM SIGKDD International conference on knowledge discovery and data mining*. pp. 1225–1234. KDD '16
- Wang S, Arroyo J, Vogelstein JT, Priebe CE (2021) Joint embedding of graphs. *IEEE Trans Pattern Anal Mach Intell* 43(4):1324–1336
- Xing EP, Fu W, Song L (2010) A state-space mixed membership blockmodel for dynamic network tomography. *Ann Appl Stat* 4(2):535–566
- Xu KS, Hero AO III (2014) Dynamic stochastic blockmodels for time-evolving social networks. *IEEE J Select Topics Signal Process* 8(4):552–562
- Yang C, Priebe CE, Park Y, Marchette DJ (2021) Simultaneous dimensionality and complexity model selection for spectral graph clustering. *J Comput Graph Stat*. <https://doi.org/10.1080/10618600.2020.1824870>
- Young SJ, Scheinerman ER (2007) Random dot product graph models for social networks. *Algorithms and models for the web-graph*. Springer, Berlin, pp 138–149
- Yu W, Aggarwal CC, Wang W (2017a) Temporally factorized network modeling for evolutionary network analysis. In: *Proceedings of the Tenth ACM International conference on web search and data mining*. pp. 455–464
- Yu W, Cheng W, Aggarwal CC, Chen H, Wang W (2017b) Link prediction with spatial and temporal consistency in dynamic networks. In: *Proceedings of the twenty-sixth international joint conference on artificial intelligence*. pp. 3343–3349
- Zhang Z, Cui P, Zhu W (2020) Deep learning on graphs: a survey. *IEEE Trans Knowl Data Eng*. <https://doi.org/10.1109/TKDE.2020.2981333>
- Zhou D, Zheng L, Han J, He J (2020) A data-driven graph generative model for temporal interaction networks. In: *Proceedings of the 26th ACM SIGKDD International conference on knowledge discovery and data mining*. pp. 401–411. KDD '20
- Zhu L, Guo D, Yin J, Steeg GV, Galstyan A (2016) Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE Trans Knowl Data Eng* 28(10):2765–2777
- Zhu M, Ghodsi A (2006) Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Comput Stat Data Anal* 51(2):918–930