

Joint latent space models for network data with high-dimensional node variables

BY XUEFEI ZHANG, GONGJUN XU AND JI ZHU

*Department of Statistics, University of Michigan,
1085 South University Avenue, Ann Arbor, Michigan 48109, U.S.A.*

xfzhang@umich.edu gongjun@umich.edu jizhu@umich.edu

SUMMARY

Network latent space models assume that each node is associated with an unobserved latent position in a Euclidean space, and such latent variables determine the probability of two nodes connecting with each other. In many applications, nodes in the network are often observed along with high-dimensional node variables, and these node variables provide important information for understanding the network structure. However, classical network latent space models have several limitations in incorporating node variables. In this paper, we propose a joint latent space model where we assume that the latent variables not only explain the network structure, but are also informative for the multivariate node variables. We develop a projected gradient descent algorithm that estimates the latent positions using a criterion incorporating both network structure and node variables. We establish theoretical properties of the estimators and provide insights into how incorporating high-dimensional node variables could improve the estimation accuracy of the latent positions. We demonstrate the improvement in latent variable estimation and the improvements in associated downstream tasks, such as missing value imputation for node variables, by simulation studies and an application to a Facebook data example.

Some key words: High-dimensional data; Latent space model; Network analysis.

1. INTRODUCTION

Network data that describe the relations or interactions among individuals have been prevalent in many scientific and engineering fields, including but not limited to social media, the world wide web and neuroscience (Newman, 2010; Kolaczyk & Csárdi, 2014). A network is usually represented by nodes and edges, where each node represents an individual and an edge represents the connectivity between two nodes. In recent years, many statistical models have been proposed to analyse network data appearing in various domains; see, for example, Goldenberg et al. (2010) for a review. Many of the existing models are based on the assumption that the formation of network links is driven by nodal latent variables, including stochastic block models (Holland et al., 1983), latent space models (Hoff et al., 2002) and random dot product graph models (Young & Scheinerman, 2007; Athreya et al., 2017). It is critical to estimate the node latent variables accurately, because the estimated latent representations of nodes not only provide insights into the structure of the network, but can also be further used as node features for subsequent tasks, such as node clustering, prediction for node response variables and network link prediction.

In practice, the network link information is often collected along with additional high-dimensional node variables. For example, in an online social network, where nodes represent

users and links represent friendship relationships, we also observe users' multiple personal information such as age, gender and educational institution (Leskovec & McAuley, 2012); and in a citation network where nodes represent papers and links represent citation relationships, word frequencies over a large number of words for each paper are recorded as well (McCallum et al., 2000). The dimension of node variables in these applications can be large in the sense that it is comparable to the number of nodes. Existing studies have shown that node variables provide complementary information to network links and often play important roles for estimating the latent structure of the network (Newman & Clauset, 2016; Zhang et al., 2016; Binkiewicz et al., 2017). Thus, it is important to model the network and node variables jointly such that the node variable information can be utilized for improved understanding of the node latent variables.

This paper proposes a joint latent space model to model network links and high-dimensional node variables simultaneously using shared latent variables. On the one hand, as mentioned above, many commonly used network models assume that the network links are determined through node latent variables. Among these models, latent space models are probably the most popular (Hoff et al., 2002) and have been shown to be powerful for capturing many commonly observed features of real-world networks (Ward & Hoff, 2007; Ward et al., 2007, 2011; Friel et al., 2016), such as node degree heterogeneity, homophily and community structures (Ma et al., 2020). Taking advantage of these properties of the network latent space model, we also assume that each node can be represented by a latent vector in a (low-)dimensional Euclidean space, and that the connecting probability between two nodes depends on the corresponding pair of nodes' positions in the unobserved space. On the other hand, it is also commonly observed that high- or moderate-dimensional variables that are correlated can often be explained in terms of a few unobserved latent variables (Bai & Li, 2012; Wang et al., 2017; Hair et al., 2018). Furthermore, for network data with node variables, the latent variables that explain the observed high-dimensional node variables could be correlated with the latent position variables that explain the network links. This motivates us to consider the joint latent space modelling framework, which uses the shared latent variables to model both parts of the observed information, with the goal of utilizing node variables effectively to help estimate the node latent positions.

Various latent-variable-based network models have been proposed for modelling network data with node variables, such as the latent space model (Hoff et al., 2002) and its variants (Hoff, 2003, 2005, 2008, 2009; Handcock et al., 2007; Krivitsky et al., 2009; Sewell & Chen, 2015, 2016; Ma et al., 2020). When node covariates are present, existing latent space models usually incorporate such information by including pairwise node variable similarities to model link probabilities (Hoff et al., 2002; Ma et al., 2020). Such similarity-based approaches have several limitations when node variables are of high dimension. First, the majority of existing latent space models adopt Bayesian estimation approaches. The high-dimensional similarity vector would introduce a large number of parameters and additional Markov chain Monte Carlo sampling, therefore making the estimation much more computationally challenging. Second, the performance of the similarity-based method would be sensitive to the specific choice of similarity measure; and it does not model the relationship between the observed node variables and the latent variables. In practice, node variables are often correlated with latent variables, and this relationship can be utilized for better understanding of the network structure (Kim & Leskovec, 2012; Xu et al., 2012; Yang et al., 2013). Furthermore, from the theoretical perspective, the existing literature has rarely studied the effects of high-dimensional node variables on estimating network latent representations, which is necessary in modern network data analysis with node variables.

The main contributions of this paper are summarized from the following perspectives. From the modelling perspective, the proposed framework has several advantages in comparison to the existing work. First, we model the relationship between node latent variables and node covariates

by a set of shared latent variables, which provides a natural way of borrowing information from node variables to improve the estimation of the latent variables. Second, the proposed model adopts the framework of generalized linear factor models to model the distribution of node variables and, therefore, can handle multiple types of node variables arising in practice, such as continuous, binary and count variables. Furthermore, we develop an efficient projected gradient descent algorithm to estimate the model parameters and latent representations, by treating the latent representations as fixed effects. Such an estimation method is computationally more efficient than the Bayesian approaches in the literature.

Moreover, from the theoretical perspective, we show that the proposed estimators of the fixed effect joint latent space model are error rate optimal under mild conditions. We also establish the corresponding non-asymptotic upper and lower error bounds. In addition, we provide new findings on how the information from both the network and node variables would balance with each other to affect the estimation of latent variables. In particular, we provide a theoretical guarantee that, when the dimension of signal node variables is large enough, borrowing information from node variables would always achieve improvement in estimating node latent positions, in comparison to the results that use network link information only. We also investigate how the sparsity level of the network would affect the necessity of including node variables for joint estimation. Our theoretical findings are further supported by extensive simulation studies.

2. PROPOSED METHOD

2.1. Joint latent space model

We start by introducing the notation. Assume that we have n data points connected by a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, 2, \dots, n\}$ denotes the node set and \mathcal{E} denotes the edge set. We consider undirected networks and write $(i, i') \in \mathcal{E}$ if there is a link between node i and node i' . The network can be represented by an adjacency matrix $A \in \{0, 1\}^{n \times n}$, where $A_{ii'} = A_{i'i} = 1$ if $(i, i') \in \mathcal{E}$ and $A_{ii'} = A_{i'i} = 0$ otherwise. Furthermore, we assume for each node i that we also observe a vector of covariate variables, denoted by $Y_i \in \mathbb{R}^q$. The matrix of node variables is denoted by $Y = [Y_1, Y_2, \dots, Y_n]^T \in \mathbb{R}^{n \times q}$. For a general matrix $M \in \mathbb{R}^{m \times n}$, we denote its i th row by M_i and its j th column by M_j .

We consider a joint latent space model for modelling network data with high-dimensional node variables. Specifically, we assume that each node $i \in \mathcal{V}$ can be represented by a low-dimensional vector $Z_i \in \mathbb{R}^k$ in an unobserved latent space. The latent variables of all the nodes are denoted by $Z = [Z_1, Z_2, \dots, Z_n] \in \mathbb{R}^{n \times k}$. As commonly assumed in previous network latent space models, two nodes that are close in the latent space are more likely to be connected. Meanwhile, it is also proper to assume that, when the two nodes are close in the latent space, they may display similarity regarding the observed traits. For instance, for two individuals who have close latent representations in a social network, they may choose similar jobs or hold similar political perspectives. This naturally leads to the consideration that the latent variables not only model the network connectivity, but are also informative for the node variables. In particular, we make the assumption that the distribution of network links and that of node variables are driven by the shared latent variables; see Fig. 1.

For network A , we assume that, for each pair of nodes (i, i') , given their latent positions Z_i and $Z_{i'}$, the presence or absence of an edge between them is determined by the corresponding pair of latent variables and is independent of any other edges. Specifically, for $i < i'$, we assume that $A_{ii'} = A_{i'i} \stackrel{\text{ind}}{\sim} \text{Ber}(P_{ii'})$ with $P_{ii'} = f(Z_i, Z_{i'})$ for some function f . Multiple choices for the function f are available; see Hoff et al. (2002), Hoff (2003, 2008) and Ma et al. (2020) for

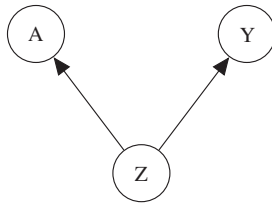


Fig. 1. Graphical representation of the joint latent space model.

examples. In this paper, we consider using the inner-product latent space model (Hoff, 2003; Ma et al., 2020), i.e.,

$$\text{logit } P_{i'j'} = \Theta_{i'j'}^A = \alpha_i + \alpha_{j'} + Z_i^T Z_{j'} \quad (1)$$

Model specification (1) can also be expressed in matrix form as

$$\text{logit } P = \Theta^A = \alpha \mathbf{1}_n^T + \mathbf{1}_n \alpha^T + ZZ^T,$$

where $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)^T$. We choose the inner-product latent space model to model the network part because of its flexibility to capture commonly observed network characteristics. For example, it allows node degree heterogeneity through the parameter α_i , and in general the larger α_i , the more likely that node i connects with other nodes. It also allows for transitivity, i.e., nodes with common neighbours are more likely to connect since their latent positions are more likely to have a larger inner product.

Furthermore, we assume that the same latent variables Z are used to describe the multivariate node variables $Y \in \mathbb{R}^{n \times q}$. Specifically, we assume that, given Z , the entries in Y are independent and Z models Y through generalized linear models (Dunn & Smyth, 2018), with

$$g(\mathbb{E}Y) = \Theta^Y = \mathbf{1}_n \gamma^T + ZB, \quad (2)$$

where $\gamma \in \mathbb{R}^q$ and $B \in \mathbb{R}^{k \times q}$ are the regression coefficients. For example, when the entries in Y are continuous and $g(x) = x$ is the identity mapping, we assume that

$$Y_{ij} \stackrel{\text{ind}}{\sim} \mathcal{N}\{(1_n \gamma^T + ZB)_{ij}, \sigma^2\} \quad (3)$$

for some σ^2 . When the entries in Y are all binary and $g(x) = \log\{x/(1-x)\}$, we have

$$Y_{ij} \stackrel{\text{ind}}{\sim} \text{Ber}\left\{ \frac{\exp(1_n \gamma^T + ZB)_{ij}}{1 + \exp(1_n \gamma^T + ZB)_{ij}} \right\}. \quad (4)$$

More generally, when there is more than one type of variable in Y , we could divide Y into R blocks, i.e., $Y = [Y_1 | \dots | Y_R]$, with each sub-block containing the same type of variables and having its own link function. For the following algorithms and theoretical results, we mainly focus on the case when there is only one type of variable in Y . The corresponding results for Y with multiple variable types can be naturally obtained.

Certain modelling approaches for the community detection problem can be considered as having a similar flavour to jointly modelling the distributions of A and Y via shared latent variables, where the discrete latent variable $Z_i \in \{0, 1\}^k$ represents the unobserved community membership

of node i . For instance, Xu et al. (2012), Yang et al. (2013) and Kim & Leskovec (2012) assumed that, for nodes from the same community or cluster, their edge connections and node variables should follow the common distribution specific to that cluster. In other words, the node latent communities determine the distributions of both A and Y , and therefore information from both A and Y could be used for community detection. Our joint latent space model considers a more general setting, where the latent variables could take continuous values in the unobserved latent space.

We treat Z as fixed effects rather than random for two reasons. First, our method does not require specific assumptions on the distribution of Z , and is therefore more flexible and general; meanwhile, treating Z as random effects usually requires making certain parametric assumptions on the distribution of Z . Second, by treating Z as fixed parameters, gradient descent methods can be adopted for parameter estimation, and the computation is usually efficient and scalable. On the other hand, when Z is viewed as random effects, the popularly used Bayesian estimation approaches may be computationally expensive.

Remark 1. Although we assume that models (1) and (2) share the same set of latent variables Z , this assumption can be easily relaxed. For instance, consider that we have latent variables $Z_A \in \mathbb{R}^{n \times k}$ modelling A through (1). Meanwhile, there exists another set of latent variables $Z_Y \in \mathbb{R}^{n \times k'}$ that are specifically informative for Y :

$$g(\mathbb{E}Y) = 1_n \gamma^T + Z_Y B. \tag{5}$$

If there exists a matrix $W \in \mathbb{R}^{k \times k'}$ such that $Z_Y \approx Z_A W$, then model (5) could be rewritten as $g(\mathbb{E}Y) \approx 1_n \gamma^T + Z_A W B = 1_n \gamma^T + Z_A B'$ with $B' = W B$, which gives a good approximation to the model in (2). Therefore, even when the two sets of latent variables explaining network and node variables are not exactly the same, if there is an approximate linear transformation relationship between them, our proposed joint latent space model is still valid, with Z_A being the shared latent variables that explains both parts of the observed information. Additionally, we could further consider a more general case where Z_A and Z_Y are different, but have some overlapped latent variables. These two sets of parameters, Z_A and Z_Y , can still be jointly estimated, but it is beyond the scope of the current paper and so we leave the investigation for future work.

To ensure that the joint model is identifiable, we need to put additional structural constraints on the latent variables Z . First, we could add a constant term to both Z_i and Z_j , and subtract the corresponding terms from α_i and α_j to keep the distribution of A invariant, so we require that the latent variables are centred, i.e., $JZ = Z$, where $J = I_n - 1_n 1_n^T/n$. This constraint makes Z identifiable up to an orthogonal transformation of its rows. Correspondingly, B is identifiable up to an orthogonal transformation of its columns. Therefore, we further require that the sample covariance of Z , i.e., $Z^T Z/n$, is a diagonal but nonidentity matrix. Then the parameters α , Z , B and γ can be uniquely determined.

2.2. Estimation

The parameters that need to be estimated are Z , α , B and γ . For the network data A , we consider the loss function as its conditional negative loglikelihood:

$$L_A = -\log P(A | Z, \alpha) = - \sum_{1 \leq i, i' \leq n} \{A_{ii'} \Theta_{ii'}^A - f_A(\Theta_{ii'}^A)\} \tag{6}$$

with $f_A(x) = \log\{1 + \exp(x)\}$. For the node variables Y , we have the negative conditional loglikelihood

$$L_Y = -\log P(Y | Z, B, \gamma) = - \sum_{1 \leq i \leq n, 1 \leq j \leq q} \{Y_{ij} \Theta_{ij}^Y - f_Y(\Theta_{ij}^Y)\},$$

where the terms that are irrelevant to Z and γ are omitted. The form of $f_Y(\cdot)$ depends on how the distribution of Y is specified. For example, when Y is continuous as in model (3), $f_Y(x) = x^2/2$; when Y is binary as in model (4), then $f_Y(\cdot)$ takes the same form as $f_A(\cdot)$.

We define the objective function as

$$L(Z, \alpha, B, \gamma) = L_A + \lambda L_Y, \quad (7)$$

where λ is a weight parameter that controls the information contributed from each part. Our goal is to find the estimators \hat{Z} , $\hat{\alpha}$, \hat{B} and $\hat{\gamma}$ that are the solutions to the following optimization problem:

$$\min_{Z \in \mathbb{R}^{n \times k}, JZ=Z, \alpha \in \mathbb{R}^n, B \in \mathbb{R}^{k \times q}, \gamma \in \mathbb{R}^q} L(Z, \alpha, B, \gamma). \quad (8)$$

To optimize (8), we consider using a projected gradient descent algorithm, which is commonly used for solving constrained optimization problems. The term projected means that we project the parameter estimate into the space that satisfies the constraint. Specifically, at each iteration, given all the other parameter estimates, the \hat{B} and $\hat{\gamma}$ that minimize the objective function (7) can be solved directly, and they are therefore updated with those values. As for \hat{Z} and $\hat{\alpha}$, they are updated along the direction of their negative gradients. After the updates, we include an additional projection step such that the parameter estimates satisfy the constraint in (8). The algorithm is summarized in Algorithm 1 below. As pointed out by Ma et al. (2020), where the convergence property of such an algorithm has been studied, this algorithm is not guaranteed to converge to any global optimizer when the objective function is not convex, as in this work. In §4, we use simulation studies to demonstrate that the algorithm generally converges well and provides good estimation results.

Algorithm 1. Projected gradient descent algorithm for parameter estimation.

Inputs: network adjacency matrix $A \in \mathbb{R}^{n \times n}$; node variables $Y \in \mathbb{R}^{n \times q}$; latent space dimension $k \geq 1$; initial estimates $Z^0, \alpha^0, B^0, \gamma^0$; step sizes η_z, η_α ; number of iterations T

Parameters: Z, α, B, γ

For $t = 0, 1, \dots, T - 1$

$$Z^{t+1} = Z^t - \eta_z \nabla_Z L = Z^t + 2\eta_z \{A - f'_A(\Theta^t)\} (Z^t)^\top + \lambda \eta_z \{Y - f'_Y(Z^t B^t)\} (B^t)^\top$$

$$\alpha^{t+1} = \alpha^t - \eta_\alpha \nabla_\alpha L = \alpha^t + 2\eta_\alpha \{A - f'_A(\Theta^t)\} \mathbf{1}_n$$

$$(\gamma_j^{t+1}, B_j^{t+1}) = \arg \max_{b \in \mathbb{R}^{k+1}} \sum_i \{Y_{ij} [1, Z_i^t] b - f_Y([1, Z_i^t] b)\} \quad (j = 1, \dots, q)$$

$$Z^{t+1} = JZ^{t+1}$$

Outputs: $\hat{Z} = Z^T, \hat{\alpha} = \alpha^T, \hat{B} = B^T, \hat{\gamma} = \gamma^T$

Algorithm 1 requires initial inputs of several hyperparameters. For the initialization values $Z^0, \alpha^0, B^0, \gamma^0$ of Z, α, B, γ and choices of step sizes η_z and η_α , we adapt the initialization method and step size choice proposed by Ma et al. (2020), who optimized a regularized version of the negative loglikelihood of the network data as in (6) to obtain initial estimates Z^0 and α^0 . Then we

regress Y on Z^0 to get an initialization of B^0 and γ^0 . For the step size choice, we let $\eta_Z = \eta/\|Z^0\|_F^2$ and $\eta_\alpha = \eta/(2n)$ for a small and fixed constant η . As for the latent space dimension k , it can be selected through cross-validation. As we focus on estimating the latent variables that explain the network links, we consider performing cross-validation on the adjacency matrix. In particular, we randomly remove entries from the adjacency matrix, then fit the joint latent space model and predict those missing links based on the fitted values. The process can be repeated several times and the k value that gives the best link prediction performance is chosen from a set of candidate values.

3. THEORETICAL RESULTS

In this section we state our main theoretical results on the estimation of Z under the joint modelling framework. We first show the error bound of the estimators obtained from (8). Then we discuss how the joint modelling framework could improve the estimation of Z .

To study the theoretical properties, we make the following assumptions on the parameters.

Assumption 1. There exists $M_1 > 0$ such that $-M_1 < \Theta_{i'i'}^A < M_1$ for $1 \leq i, i' \leq n$.

Assumption 2. There exists $M_2 > 0$ such that $-M_2 < \Theta_{ij}^Y < M_2$ for $1 \leq i \leq n, 1 \leq j \leq q$.

Moreover, we introduce a feasible parameter space as

$$\begin{aligned} \mathcal{F}(n, k, M_1, M_2) = & \left\{ \Theta \in \mathbb{R}^{n \times (n+q)} \mid \Theta = [\Theta^A, \Theta^Y], \right. \\ & \Theta^A = \alpha 1_n^T + 1_n \alpha^T + ZZ^T, \Theta^Y = 1_n \gamma^T + ZB, \\ & \left. \max_{1 \leq i, i' \leq n} |\Theta_{i'i'}^A| < M_1, \max_{1 \leq i \leq n, 1 \leq j \leq q} |\Theta_{ij}^Y| < M_2, JZ = Z \right\}. \end{aligned} \quad (9)$$

In the rest of the paper, we use \mathcal{F} as an abbreviation of $\mathcal{F}(n, k, M_1, M_2)$ to denote the parameter space. We denote the ground truth parameter by

$$\Theta^* = [\Theta^{*A}, \Theta^{*Y}] = [\alpha^* 1_n^T + 1_n (\alpha^*)^T + Z^* (Z^*)^T, 1_n (\gamma^*)^T + Z^* B^*] \in \mathcal{F}.$$

We let

$$\hat{\Theta} = [\hat{\Theta}^A, \hat{\Theta}^Y] = [\hat{\alpha} 1_n^T + 1_n^T \hat{\alpha} + \hat{Z} \hat{Z}^T, 1_n \hat{\gamma}^T + \hat{Z} \hat{B}],$$

where $\hat{Z}, \hat{\alpha}, \hat{B}$ and $\hat{\gamma}$ are the solutions obtained from the optimization problem $\min_{\Theta \in \mathcal{F}} L(Z, \alpha, B, \gamma)$. We constrain the true parameters and estimators in the feasible parameter space, mainly for the purpose of theoretical analysis. In practice, when implementing Algorithm 1, we do not include additional constraints regarding $\max_{1 \leq i, i' \leq n} |\hat{\Theta}_{i'i'}^A|$ and $\max_{1 \leq i \leq n, 1 \leq j \leq q} |\hat{\Theta}_{ij}^Y|$, and simulation studies indicate that this does not affect the results.

The following theorem provides upper and lower bounds on the estimation error of $\hat{\Theta}$.

THEOREM 1. *Under Assumptions 1 and 2, we have*

$$\frac{1}{\{n(n+q)\}^{1/2}} \mathbb{E} \|\hat{\Theta} - \Theta^*\|_F \leq \frac{\kappa \max(\lambda, 1)(2k+3)^{1/2}}{\min\{\min_{|v| < M_1} f_A''(v), \lambda \min_{|v| < M_2} f_Y''(v)\}} \frac{1}{n^{1/2}}, \quad (10)$$

where κ is an absolute constant and M_1, M_2 and k are allowed to change with n .

Moreover, let $\bar{\Theta} \in \mathbb{R}^{n \times (n+q)}$ be an arbitrary estimator. When $q = \mathcal{O}(n)$, there exist $\Theta^0 \in \mathcal{F}$, $\epsilon_1 > 0$ and $n_0, q_0 > 0$ such that, for $n > n_0$ and $q > q_0$,

$$P \left[\frac{1}{\{n(n+q)\}^{1/2}} \|\bar{\Theta} - \Theta^0\|_F \geq \frac{\epsilon_1}{n^{1/2}} \right] \geq \frac{1}{2}. \tag{11}$$

The proof is given in the [Supplementary Material](#). When M_1, M_2 and k are fixed constants, the result in (10) implies that $\|\hat{\Theta} - \Theta^*\|_F / \{n(n+q)\}^{1/2} = \mathcal{O}_p(1/n^{1/2})$. Together with the lower bound in (11), we can see that the rate of estimation error obtained in Theorem 1 is optimal. Moreover, the results also indicate that, using the network itself, i.e., $q = 0$, we have $\|\hat{\Theta}^A - \Theta^{*A}\|_F/n = \mathcal{O}_p\{e^{M_1}(k/n)^{1/2}\}$. Therefore, we achieve the same order of estimation error under the joint modelling framework, compared to that obtained with the network information only. In particular, the upper bound of $\|\hat{\Theta}^A - \Theta^{*A}\|_F/n$ is consistent with the results of [Ma et al. \(2020\)](#), who considered the problem of estimating latent variables using the network only, by minimizing L_A as defined in (6).

While Theorem 1 indicates that the error bounds of estimators obtained under the joint modelling framework have the same order as those obtained from the network data, without Y , we are also interested in how the additional node variables Y can help the estimation of latent variables. We evaluate the effect of Y in terms of the one-step update analysis. In particular, assume that we have an estimated \tilde{Z} through the network, for example, from the algorithm proposed by [Ma et al. \(2020\)](#). Suppose that with node variables Y we update \tilde{Z} for one more step based on Algorithm 1,

$$\hat{Z} = \tilde{Z} + (1 - \tilde{\lambda})\eta_z\{A - f'_A(\tilde{\Theta}^A)\}\tilde{Z} + \tilde{\lambda}\eta_z\{Y - f'_Y(\tilde{\Theta}^Y)\}(\tilde{B})^T, \tag{12}$$

where $\tilde{\lambda} = \lambda/(\lambda + 2)$ and \tilde{B} is a reasonable estimate of B . To investigate the properties of \hat{Z} , we focus on (3) where Y is continuous and make the following additional assumptions.

Assumption 3. The dimension of node variables q satisfies the condition that $q = \mathcal{O}(n)$.

Assumption 4. Let $\text{cov}(Z) = Z^T Z/n = \text{diag}(\lambda_1, \dots, \lambda_k) \oplus I_k$ be diagonal, and let the diagonal elements be of constant order $\Theta(1)$.

Assumption 5. Denote the eigendecomposition of BB^T/q as $U\Lambda U^T$, where $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_k)$. The eigenvalues are of constant order $\Theta(1)$.

Assumption 3 allows q to grow at a rate slower than, or of the same order, as n . Assumptions 4 and 5 are standard since $Z \in \mathbb{R}^{n \times k}$ and $B \in \mathbb{R}^{k \times q}$. Assumption 5 implies that node covariates could not be pure noises, but are signals. The following theorem shows that we can achieve more accurate estimation of Z , as long as the dimension of signal node variables is high enough.

THEOREM 2. *Suppose that \tilde{Z} and $\tilde{\alpha}$ are estimated from Algorithm 1 of [Ma et al. \(2020\)](#), and that we have a fixed \tilde{B} satisfying $\|\tilde{B} - B\|_F^2 = \mathcal{O}(1)$. Then, under Assumptions 1–5, there exist positive constants C and $\bar{\lambda}$ such that, when $q > Cn$, we have $\mathbb{E}\|\hat{Z} - Z\|_F^2 < \mathbb{E}\|\tilde{Z} - Z\|_F^2$ for any $\tilde{\lambda} \in (0, \bar{\lambda})$, where \hat{Z} is obtained from (12). Moreover, we can show that, under a proper choice of $\tilde{\lambda}$, whose form is specified in the [Supplementary Material](#), the improvement in estimating Z , i.e., $\mathbb{E}\|\hat{Z} - Z\|_F^2 - \mathbb{E}\|\tilde{Z} - Z\|_F^2$, is at least*

$$\frac{\{\tilde{\sigma}_k \mathbb{E}\|\tilde{Z} - Z\|_F^2 - \|\tilde{Z} - Z\|_F(nq^{-1}\tilde{\sigma}_1\lambda_1)^{1/2}\|B - \tilde{B}\|_F\}^2}{\rho_2[\{\tilde{\sigma}_k \mathbb{E}\|\tilde{Z} - Z\|_F - (nq^{-1}\tilde{\sigma}_1\lambda_1)^{1/2}\|B - \tilde{B}\|_F\}^2 + nq^{-1}\sigma^2\tilde{\sigma}_1]}, \tag{13}$$

where $\rho_2 = \eta_z q$, and $\tilde{\sigma}_1$ and $\tilde{\sigma}_k$ are the maximum and minimum eigenvalues of $\tilde{B}\tilde{B}^\top/q$, respectively.

From Theorem 2, we see that with additional high-dimensional node variables we can achieve more accurate estimation of latent variables, in terms of $\mathbb{E}\|\hat{Z} - Z\|_F^2$. The improvement in (13) depends on the signal and noise contained in the node covariates. Specifically, $\tilde{\sigma}_k$, the minimum eigenvalue of $\tilde{B}\tilde{B}^\top/q$, can be approximately viewed as the signal in node covariates, and the quantity in (13) would increase as $\tilde{\sigma}_k$ increases. Moreover, the quantity monotonically decreases depending on σ^2 , the noise in the node covariates. In practice, a \tilde{B} that satisfies the requirement $\|\tilde{B} - B\|_F^2 = \mathcal{O}(1)$ can be obtained by regressing Y on \tilde{Z} .

Theorem 2 relies on specific \tilde{Z} and $\tilde{\alpha}$. More generally, we consider the scenario where we are given initial estimates of Z, α, B, γ , denoted by $\tilde{Z}, \tilde{\alpha}, \tilde{B}, \tilde{\gamma}$, respectively, satisfying the conditions that $\|\tilde{Z} - Z\|_F^2 = \mathcal{O}(1)$, $\|\tilde{\alpha}1_n^\top - \alpha 1_n^\top\|_F^2 = \mathcal{O}(n)$, $\|\tilde{B} - B\|_F^2 = \mathcal{O}(1)$ and $\|\tilde{\gamma} - \gamma\|_2^2 = \mathcal{O}(1)$. The following proposition provides implications for which scenarios the joint modelling framework can help better estimate the latent variables Z .

PROPOSITION 1. *Given $\tilde{Z}, \tilde{\alpha}, \tilde{B}, \tilde{\gamma}$ that satisfy the above required conditions, we consider updating \tilde{Z} one step further by (12) and obtaining a \hat{Z} . Under Assumptions 1–5, there exists an optimal $\tilde{\lambda}_{\text{opt}}$ such that $\mathbb{E}\|\hat{Z} - Z\|_F^2$ is minimized; $\tilde{\lambda}_{\text{opt}}$ is given in (S8) in the [Supplementary Material](#). When $\tilde{\lambda}_{\text{opt}}$ is strictly positive, the joint modelling framework achieves a mean square error of \hat{Z} that is better than the results obtained when using information from A or Y only.*

The proofs of Theorem 2 and Proposition 1 are given in the [Supplementary Material](#). For Proposition 1, the explicit expression of $\tilde{\lambda}_{\text{opt}}$ is given in (S8) in the [Supplementary Material](#). It depends on the signal and noise contained in both networks and node covariates, indicating the relative importance of the information from both parts. First, a strictly positive $\tilde{\lambda}_{\text{opt}}$ suggests that incorporating information from node variables is preferred. Based on the calculation in the [Supplementary Material](#), we can see that an overall sparser network would more likely lead to a positive $\tilde{\lambda}_{\text{opt}}$. Therefore, when the information from the network part is relatively limited, borrowing information from node variables would be preferred and helpful. Second, when controlling all other parameters and increasing q , we would also obtain a larger $\tilde{\lambda}_{\text{opt}}$. This suggests that, when node variables are of higher dimension or contain richer information, more weight should be put on the node variable part to improve the estimation of Z . Finally, suppose that we do not use the optimal $\tilde{\lambda}$, but fix it in the one-step estimation. In the [Supplementary Material](#) we calculate the difference between taking a nonzero $\tilde{\lambda}$ and $\tilde{\lambda} = 0$, and the result also suggests that, when the dimension of the node variables increases or the network gets sparser, the benefit of incorporating node variables becomes more significant in terms of estimating Z . In the [Supplementary Material](#), we use simulation studies to demonstrate how network information and node variable information affect the estimation of Z . Specifically, we examine the influence of the dimension of the node variables and the network density in terms of estimating Z as well as the relationship between the dimension of the node variables and the optimal weight.

Remark 2. Our work is closely related to the problem of low-rank matrix estimation and completion (Candès & Tao, 2010; Bhaskar & Javanmard, 2015; Chatterjee, 2015), due to the assumption on the mean parameter $\Theta \in \mathbb{R}^{n \times (n+q)}$ defined in (9). Specifically, if we treat node covariates and the network equally, we could view the problem as recovering the low-rank matrix Θ with an observed matrix $[A, Y] \in \mathbb{R}^{n \times (n+q)}$. In fact, the upper bound in (10) depends on k and n through the rate $(k/n)^{1/2}$, which is consistent with that in the low-rank matrix estimation literature. However, we could not directly leverage the existing matrix estimation and completion

method for estimating Z , since the matrix to be estimated is not an arbitrary low-rank matrix, but has a specific structure due to the shared latent factors Z . This is the distinct part of our model assumption in comparison to the general low-rank matrix estimation problem.

4. SIMULATION STUDIES

4.1. Effect of the dimension of node variables

To study how information borrowed from Y can affect estimating latent variables, we compare the estimation of Z using the network latent space model and the joint latent space model. For the network latent space model, we consider the version without covariates to demonstrate how node variables can be useful in improving the estimation of Z .

We first study the effect of the node variable dimension. We set $n = 200$ and $k = 2$ or 4 . We vary q from 2 to 100 to study how the dimension of the node variables affects the estimation of Z . The model parameters are specified as follows.

Step 1. Generate the degree heterogeneity parameters $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$, where $\alpha_1, \alpha_2, \dots, \alpha_n \stackrel{\text{iid}}{\sim} U[-0.25, -0.75]$.

Step 2. Generate k latent vector centres $\mu_1, \dots, \mu_k \in \mathbb{R}^k$ with coordinates independent and identically distributed from $U[-1, 1]$.

Step 3. Generate latent variables $Z \in \mathbb{R}^{n \times k}$. First generate a matrix $Z_0 \in \mathbb{R}^{n \times k}$ such that each entry is independent and identically distributed $\mathcal{N}(0, 1)$. Then we divide n data points equally into k subsets, and for points in each subset, respectively add μ_1, \dots, μ_k to them. Finally, we transform Z by (i) setting $Z = JZ$, (ii) normalizing Z such that $\|ZZ^T\|_F = n$ and (iii) rotating $Z = ZR$ for some rotation matrix R such that the covariance of Z is a diagonal matrix.

Step 4. Generate the coefficients $B \in \mathbb{R}^{k \times q}$, with each entry independent and identically distributed from $\mathcal{N}(0, 1)$.

After setting the parameters, we generate A based on model (1) and generate Y based on either model (3) or model (4). Under the considered parameter settings, the average density of the networks is about 0.30. Each setting is repeated 30 times. For each replication, we fit both the network inner-product latent space model and the joint latent space model to obtain estimations of Z , denoted by Z_{net} and Z_{joint} , respectively. We evaluate the performance of each method using the criterion $\Delta_Z = \|\hat{Z}\hat{Z}^T - ZZ^T\|_F^2 / \|ZZ^T\|_F^2$. Figure 2 shows the average results of the 30 replications; we can see that, as the dimension of Y increases, the estimation of Z improves, and the joint latent space model starts to outperform the network latent space model even when q is relatively small, indicating that the constant C in Theorem 2 is of small value. Figure 2 also demonstrates that overall the improvement of the joint latent space model over the network latent space model is robust to the choice of λ , though the specific value of λ may affect how much improvement we could obtain by incorporating node variables.

4.2. Effect of the network density

Section 4.1 shows that the more information provided from Y , the more improvement we could obtain in the estimation of Z . As a counterpart, in this subsection we demonstrate the effect of the network density on estimating Z when the information from Y is fixed. We fix $q = 100$ and change the parameter settings in the network. The density of the network is controlled by varying the range of the node degree heterogeneity parameters α from $\alpha_1, \dots, \alpha_n \sim U[-0.375, -0.125]$

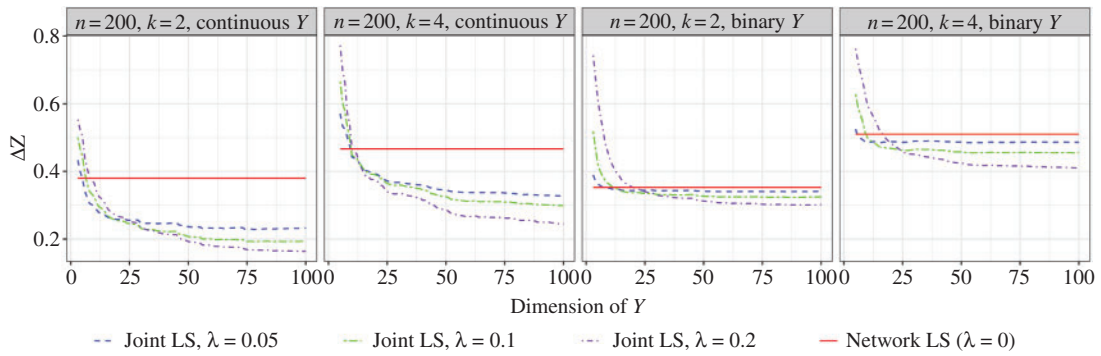


Fig. 2. Estimation of Z versus the dimension of Y .

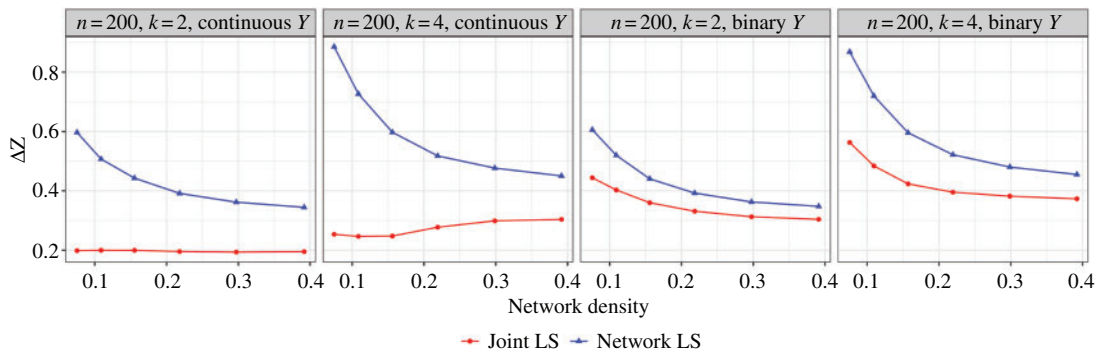


Fig. 3. Estimation of Z versus network density.

to $\alpha_1, \dots, \alpha_n \sim U[-2.25, -0.75]$. We set Z and B in the same manner as in § 4.1. Now, under different settings for α , the network density ranges from 0.08 to 0.39. We again repeat the simulation 30 times under each setting.

Figure 3 shows the estimation of Z based on Z_{net} and Z_{joint} under different network density levels. As the network gets sparser, the result of the Z estimation using A only gets worse, while the performance of Z_{joint} is relatively stable. This especially indicates the benefit of incorporating node variables in estimating Z , when the network is relatively sparse and may not provide enough information.

Additional simulation studies that support our theoretical findings are provided in the [Supplementary Material](#).

5. REAL DATA EXAMPLE

We now demonstrate the proposed model on Facebook social circle data for the task of missing value imputation for node variables. The dataset consists of 10 different networks, each representing an ego network of a selected user, where the ego network is defined as the network between all the user’s friends. See Fig. 4 for an example of an ego network. For each ego network, the data include the users’ anonymized variables. For example, the original dataset contains a variable ‘political = Democratic Party’, and the curated dataset would transform this into ‘political = anonymized feature 1’. The variables associated with each node in this data example are all binary. We analysed eight out of these 10 networks, as two of the networks have relatively few nodes and several variable columns associated with them contain only one or a few 1s. For each

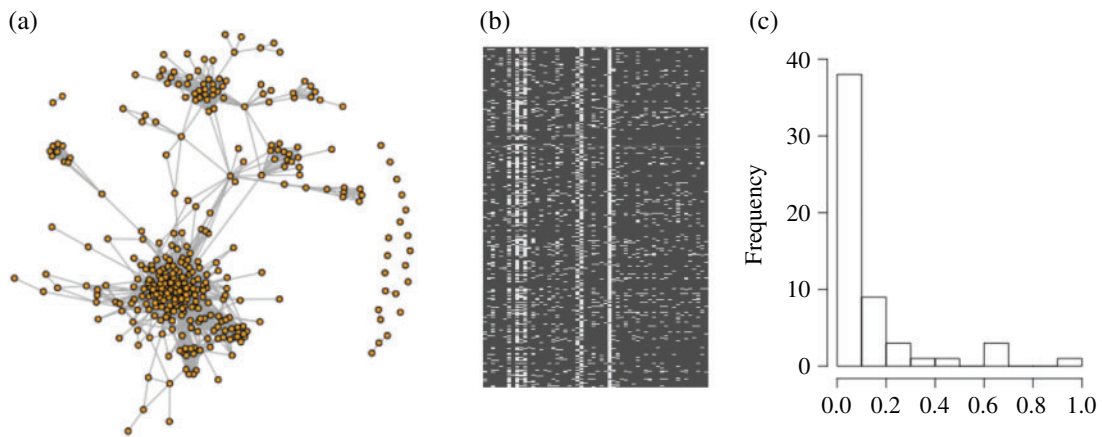


Fig. 4. Example of an ego network: (a) circle-1 network; (b) node variable matrix; (c) number of variables versus the mean of the variable.

Table 1. Node variable missing value imputation results for Facebook data

	n	q	Density	Area under the curve		
				Network	Joint	MICE
Circle 1	347	56	0.042	0.840	0.884	0.707
Circle 2	755	66	0.105	0.865	0.873	0.739
Circle 3	792	100	0.045	0.856	0.901	0.780
Circle 4	1045	153	0.049	0.867	0.871	0.693
Circle 5	547	58	0.03	0.852	0.905	0.738
Circle 6	227	87	0.124	0.832	0.850	0.586
Circle 7	159	55	0.134	0.809	0.836	0.621
Circle 8	170	37	0.115	0.812	0.854	0.698

network, we also remove the variables that have too many or too few 1s in that variable. If a variable of a user is missing, it is of interest to impute the missing value based on the user's own profile as well as his/her social connections.

We randomly sample 5% of the entries in the node variable matrix Y and set them as missing. Then we fit the network latent space model and the joint latent space model to respectively obtain estimated \hat{Z}_{net} and \hat{Z}_{joint} . After obtaining the estimated \hat{Z} for the j th column of Y , $1 \leq j \leq q$, we regress those observed Y_j on \hat{Z}_{obs} to obtain an estimated \hat{B}_j , then we predict those missing entries in Y_j as $\hat{Z}_{\text{miss}}\hat{B}_j$. Moreover, we also apply multiple imputation by chained equations (MICE; Azur et al., 2011), a commonly used method for missing value imputation that utilizes the node variable part only, as a benchmark.

In Table 1 we list the average area under the receiver operating curves over 30 replications. The results indicate that both network-based methods outperform MICE, and that the joint latent space model consistently achieves better performance than the network latent space model across all eight networks. The area under the receiver operating curve obtained by fitting the network latent space model is already promising, which suggests that the network itself contains substantial information. However, we can still achieve noticeable improvement after incorporating the node variable information. This implies that the latent position variables are associated with the node variables, and that the joint estimation can help achieve more accurate estimation and imputation results.

6. DISCUSSION

In our simulation and real data examples, we use all n data points with observed information for model fitting, and we obtain estimated model parameters $\hat{Z} \in \mathbb{R}^{n \times k}$, $\hat{\alpha} \in \mathbb{R}^n$, $\hat{B} \in \mathbb{R}^{q \times k}$ and $\hat{\gamma} \in \mathbb{R}^q$. The predictions for the missing node variables are also performed on these n data points. One possible extension is to consider an inductive setting, where we make a prediction on a new node, which is not present during the training stage, with partially observed information. For example, consider the cold-start scenario where only the new node's variables Y_{n+1} are observed, but its link information to all the other n nodes is unknown. This is a case commonly seen in social networks, where newly registered users only provide their personal information, but have no connection with any other users. To predict links between the $(n+1)$ th node and the previous n nodes, we can first estimate \hat{Z}_{n+1} by regressing Y_{n+1} on the fitted \hat{B} , then compare $\hat{\alpha}_i + \hat{Z}_i^T \hat{Z}_{n+1}$ for $i = 1, \dots, n$ to rank which nodes have higher probabilities of connecting with the new node. Correspondingly, we can also make predictions on the $(n+1)$ th node variables, when we only know its link information with other nodes.

ACKNOWLEDGEMENT

The authors are grateful to the editor, an associate editor and three referees for their valuable comments and suggestions. This research was partially supported by the U.S. National Science Foundation.

SUPPLEMENTARY MATERIAL

[Supplementary Material](#) includes proofs of Theorems 1 and 2, and additional simulation studies.

REFERENCES

- ATHREYA, A., FISHKIND, D. E., TANG, M., PRIEBE, C. E., PARK, Y., VOGELSTEIN, J. T., LEVIN, K., LYZINSKI, V. & QIN, Y. (2017). Statistical inference on random dot product graphs: a survey. *J. Mach. Learn. Res.* **18**, 8393–484.
- AZUR, M. J., STUART, E. A., FRANGAKIS, C. & LEAF, P. J. (2011). Multiple imputation by chained equations: what is it and how does it work? *Int. J. Meth. Psych. Res.* **20**, 40–9.
- BAI, J. & LI, K. (2012). Statistical analysis of factor models of high dimension. *Ann. Statist.* **40**, 436–65.
- BHASKAR, S. A. & JAVANMARD, A. (2015). 1-bit matrix completion under exact low-rank constraint. In *2015 49th Ann. Conf. Info. Sci. Syst. (CISS)*, Baltimore, MD, March 2015, pp. 1–6. Piscataway, NJ: IEEE Press.
- BINKIEWICZ, N., VOGELSTEIN, J. T. & ROHE, K. (2017). Covariate-assisted spectral clustering. *Biometrika* **104**, 361–77.
- CANDÈS, E. J. & TAO, T. (2010). The power of convex relaxation: near-optimal matrix completion. *IEEE Trans. Info. Theory* **56**, 2053–80.
- CHATTERJEE, S. (2015). Matrix estimation by universal singular value thresholding. *Ann. Statist.* **43**, 177–214.
- DUNN, P. K. & SMYTH, G. K. (2018). *Generalized Linear Models with Examples in R*. New York: Springer.
- FRIEL, N., RASTELLI, R., WYSE, J. & RAFTERY, A. E. (2016). Interlocking directorates in Irish companies using a latent space model for bipartite networks. *Proc. Nat. Acad. Sci.* **113**, 6629–34.
- GOLDENBERG, A., ZHENG, A. X., FIENBERG, S. E. & AIROLDI, E. M. (2010). A survey of statistical network models. *Foundat. Trends Mach. Learn.* **2**, 129–233.
- HAIR, J., BLACK, W., BABIN, B. & ANDERSON, R. (2018). *Multivariate Data Analysis*. Andover, UK: Cengage Learning EMEA.
- HANDCOCK, M. S., RAFTERY, A. E. & TANTRUM, J. M. (2007). Model-based clustering for social networks. *J. R. Statist. Soc. A* **170**, 301–54.
- HOFF, P. D. (2003). Random effects models for network data. In *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, R. Breiger, K. Carley & P. Pattison, eds. Washington, DC: National Academies Press, pp. 303–12.
- HOFF, P. D. (2005). Bilinear mixed-effects models for dyadic data. *J. Am. Statist. Assoc.* **100**, 286–95.

- HOFF, P. (2008). Modeling homophily and stochastic equivalence in symmetric relational data. In *Advances in Neural Information Processing Systems*, vol. 20, J. Platt, D. Koller, Y. Singer & S. Roweis, eds. Red Hook; NY: Curran Associates, pp. 1–8.
- HOFF, P. D. (2009). Multiplicative latent factor models for description and prediction of social networks. *Comp. Math. Organ. Theory* **15**, 261.
- HOFF, P. D., RAFTERY, A. E. & HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Am. Statist. Assoc.* **97**, 1090–8.
- HOLLAND, P. W., LASKEY, K. B. & LEINHARDT, S. (1983). Stochastic blockmodels: first steps. *Social Networks* **5**, 109–37.
- KIM, M. & LESKOVEC, J. (2012). Latent multi-group membership graph model. In *Proc. 29th Int. Conference Int. Conf. Mach. Learn.* Madison, WI: Omnipress, pp. 1–8.
- KOLACZYK, E. D. & CSÁRDI, G. (2014). *Statistical Analysis of Network Data with R*, R. Gentleman, K. Hornik & G. Parmigiani, eds. New York: Springer.
- KRIVITSKY, P. N., HANDCOCK, M. S., RAFTERY, A. E. & HOFF, P. D. (2009). Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social Networks* **31**, 204–13.
- LESKOVEC, J. & MCAULEY, J. J. (2012). Learning to discover social circles in ego networks. In *Advances in Neural Information Processing Systems*, vol. 25, F. Pereira, C. J. C. Burges, L. Bottou & K. Q. Weinberger, eds. Red Hook; NY: Curran Associates.
- MA, Z., MA, Z. & YUAN, H. (2020). Universal latent space model fitting for large networks with edge covariates. *J. Mach. Learn. Res.* **21**, 1–67.
- MCCALLUM, A. K., NIGAM, K., RENNIE, J. & SEYMORE, K. (2000). Automating the construction of internet portals with machine learning. *Info. Retr.* **3**, 127–63.
- NEWMAN, M. (2010). *Networks: An Introduction*. Oxford: OUP Oxford.
- NEWMAN, M. E. & CLAUSET, A. (2016). Structure and inference in annotated networks. *Nature Commun.* **7**, doi: 10.1038/ncomms11863.
- SEWELL, D. K. & CHEN, Y. (2015). Latent space models for dynamic networks. *J. Am. Statist. Assoc.* **110**, 1646–57.
- SEWELL, D. K. & CHEN, Y. (2016). Latent space models for dynamic networks with weighted edges. *Social Networks* **44**, 105–16.
- WANG, J., ZHAO, Q., HASTIE, T. & OWEN, A. B. (2017). Confounder adjustment in multiple hypothesis testing. *Ann. Statist.* **45**, 1863–94.
- WARD, M. D. & HOFF, P. D. (2007). Persistent patterns of international commerce. *J. Peace Res.* **44**, 157–75.
- WARD, M. D., SIVERSON, R. M. & CAO, X. (2007). Disputes, democracies, and dependencies: a reexamination of the Kantian peace. *Am. J. Polit. Sci.* **51**, 583–601.
- WARD, M. D., STOVEL, K. & SACKS, A. (2011). Network analysis and political science. *Ann. Rev. Polit. Sci.* **14**, 245–64.
- XU, Z., KE, Y., WANG, Y., CHENG, H. & CHENG, J. (2012). A model-based approach to attributed graph clustering. In *Proc. 2012 ACM SIGMOD Int. Conf. Manag. Data*. New York: Association for Computing Machinery, pp. 505–16.
- YANG, J., MCAULEY, J. & LESKOVEC, J. (2013). Community detection in networks with node attributes. In *2013 IEEE 13th Int. Conf. Data Mining*. Piscataway, NJ: IEEE Press, pp. 1151–6.
- YOUNG, S. J. & SCHEINERMAN, E. R. (2007). Random dot product graph models for social networks. In *Algorithms and Models for the Web-Graph*, A. Bonato & F. R. K. Chung, eds. Berlin: Springer, pp. 138–49.
- ZHANG, Y., LEVINA, E. & ZHU, J. (2016). Community detection in networks with node features. *Electron. J. Statist.* **10**, 3153–78.

[Received on 6 December 2019. Editorial decision on 5 August 2021]