Latent Space Models for Dynamic Networks with Weighted Edges[◇]

Daniel K. Sewell and Yuguo Chen*

*Department of Statistics, University of Illinois at Urbana-Champaign, 725 S. Wright
Street, Champaign, IL 61820 USA*

Abstract

Longitudinal binary relational data can be better understood by implementing a latent space model for dynamic networks. This approach can be broadly extended to many types of weighted edges by using a link function to model the mean of the dyads, or by employing a similar strategy via data augmentation. To demonstrate this, we propose models for count dyads and for non-negative real dyads, analyzing simulated data and also both mobile phone data and world export/import data. The model parameters and latent actors' trajectories, estimated by Markov chain Monte Carlo algorithms, provide insight into the network dynamics.

Keywords: Embedding; Markov chain Monte Carlo; Network dynamics; Visualization; Weighted network; Valued dyad

[◇]This research was partly supported by the National Science Foundation grants DMS-1106796 and DMS-1406455.

*Telephone: +1 217 244 3932.

Fax: +1 217 244 7190.

Email address: yuguo@illinois.edu.

1 Introduction

Representing relational data by networks is extremely useful and widely implemented. The dyadic relations which compose these networks are viewed as a set of actors and a set of edges between the actors. The edges can vary in many ways, such as being directed or undirected, static or temporal, binary or weighted. Binary networks, where between each actor an edge either does or does not exist, are encountered more often in the literature, although many such networks are by nature weighted. Weighted networks, also referred to as valued networks, consist of actors connected by edges which can take more than two values. By accounting for the weight, or strength, of the edges, the richness of the data can be better exploited. Examples of analyses of real world weighted networks include food webs (Krause et al. 2003), gene expression data (Zhang and Horvath 2005), airline networks (Barrat et al. 2005), mobile phone networks (Onnela et al. 2007), and many more.

Often in binary networks it is of interest to compute various network measures, and recently there has been increasing work in extending these measures to weighted networks. Opsahl et al. (2010) derived for weighted networks measures for degree, closeness, and betweenness. Yang and Knoke (2001) derived a method for computing path length in the case of weighted edges. Opsahl and Panzarasa (2009) developed a method for analyzing the clustering that exists within a network with weighted edges. Other interesting works include Kunegis et al. (2009), which analyzed the case where edges took values in $\{-1, 0, 1\}$, and Newman (2004), which showed how to model networks whose edges are counts by representing them as multigraphs. To fully model the network, Krivitsky (2012) extended the commonly used exponential random graph model (ERGM) to account for networks whose dyads are counts; Krivitsky and Butts (2012) extended the ERGM to account for networks whose dyads are rankings.

Network data are most often inherently dynamic, even though it is frequently the case that the data are simply aggregated over time into one static network. Many popular static networks have been extended to longitudinal network data. Examples of this include the temporal exponential random graph model developed by Hanneke et al. (2010) and the separable temporal exponential graph model by Krivitsky and Handcock (2014), the mixed

membership stochastic blockmodel for dynamic networks by Xing et al. (2010), and the latent space model for dynamic networks by several authors including Sarkar and Moore (2005), Sewell and Chen (2015b), Morgan (2014) and Durante and Dunson (2014).

This paper is focused on network data that is dynamic, weighted, and possibly directed. There are few resources available to the researcher investigating such data. Most approaches in existence focus on latent space models for dynamic undirected networks. Latent space models assume the dependence of the network is induced by a set of latent variables. Such approaches are typically intuitive and have the advantage of producing meaningful visualizations, allowing the researcher to better understand the network structure as well as the behavior of individual actors.

Sarkar et al. (2007) extended the CODE model of Globerson et al. (2004) for dynamic undirected networks. This method is an approximate filtering algorithm which models the longitudinal count networks, embedding the actors in a latent space. This method is not easily generalizable to other sorts of co-occurrence data besides counts, however, and cannot handle directed edges. Hoff (2011) described a multilinear model for undirected longitudinal networks. In this work, Hoff showed how to model undirected edges or ranked edges, where each dyad is an element from a finite ordered set, though it should be feasible to extend his approach to other types of dyads. Sewell and Chen (2015a) developed a latent space model for directed ranked dynamic networks, where each actor ranks each other actor, although it is not obvious how to extend this approach beyond this specific context.

The remainder of the paper is organized as follows. Section 2 extends the latent space model for dynamic networks with valued edges. Section 3 gives a method of estimation. Section 4 describes an approximation to reduce computational cost for large networks. Section 5 gives simulation results. Section 6 gives the results for analyzing a mobile phone network and world trade data. Section 7 provides a brief discussion.

2 Models

We assume here that each actor exists within some latent space which can be interpreted as a characteristic space, or a social space. When actors are closer together in this latent space,

the probability of a stronger edge is increased (where a “stronger edge” means a stronger relationship, though the actual form of this is context specific).

We first introduce some general notation to be used throughout. Assume we have a set of actors $\mathcal{N}$ and a set of edges $\mathcal{E}$. Let $n = |\mathcal{N}|$ be the number of actors, and let Y_t be the $n \times n$ adjacency matrix of the observed network at time t whose entries y_{ijt} correspond to the weight of the edge from actor i to actor j for $t \in \{1, 2, \dots, T\}$. Let $\mathbf{X}_{it} \in \mathbb{R}^p$ be the position vector of the i^{th} actor at time t within the p dimensional latent space. Let $\mathcal{X}_t$ be the matrix whose i^{th} row is $\mathbf{X}_{it}$. Finally, let Ψ be the vector of unknown parameters (which will vary depending on dyadic type).

As in Sarkar and Moore (2005) and Sewell and Chen (2015b), we assume the latent actor positions transition according to a Markov process, where the initial distribution is

$$\pi(\mathcal{X}_1 | \Psi) = \prod_{i=1}^n N(\mathbf{X}_{i1} | \mathbf{0}, \tau^2 I_p), \quad (1)$$

and the transition equation is

$$\pi(\mathcal{X}_t | \mathcal{X}_{t-1}, \Psi) = \prod_{i=1}^n N(\mathbf{X}_{it} | \mathbf{X}_{i(t-1)}, \sigma^2 I_p), \quad (2)$$

for $t = 2, 3, \dots, T$, where I_p is the $p \times p$ identity matrix, and $N(\mathbf{x} | \boldsymbol{\mu}, \Sigma)$ denotes the multivariate normal probability density function with mean $\boldsymbol{\mu}$ and covariance matrix Σ evaluated at $\mathbf{x}$. While this is the latent dependence structure used throughout the remainder of the paper, other dependence structures could be defined, such as the latent path model given by Morgan (2014).

In most dynamic network models it is assumed that the dependence structure of the network is fully induced by the latent positions of the actors. This assumption, along with the Markovian properties of the latent positions, leads to the state space temporal dependence structure given in Figure 1, as well as the conditional independence of each dyad within a time period. The ranked networks of the form analyzed by Krivitsky and Butts (2012) and Sewell and Chen (2015a) are a counter example of where there is an extra dependency constraint in the data, but we will not discuss further these rare data types. What remains then is to derive an appropriate conditional likelihood function, $\pi(Y_1, \dots, Y_T | \mathcal{X}_1, \dots, \mathcal{X}_T, \Psi) = \prod_{t=1}^T \prod_{i \neq j} \pi(y_{ijt} | \mathcal{X}_t, \Psi)$.

Figure 1: Illustration of the dependence structure for the latent space model. Y_t is the observed graph, $\mathcal{X}_t$ is the unobserved latent actor positions, and Ψ is the vector of model parameters.

Most latent space approaches have the conditional likelihood constructed by writing the logit of the edge probability as a linear form of covariates and a function of the latent variables, i.e., $\text{logit}(\pi(y_{ijt}|\cdot)) = \alpha' \mathbf{w}_{ijt} + f_{\Psi}(\mathbf{X}_{it}, \mathbf{X}_{jt})$, where α is a vector of unknown parameters, $\mathbf{w}_{ijt}$ is a vector of dyad specific covariates, and $f_{\Psi} : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is a function taking as its arguments two actors' latent variables. Our generalization of this has the basic form

$$g(\mathbb{E}(y_{ijt})) = \alpha' \mathbf{w}_{ijt} + f_{\Psi}(\mathbf{X}_{it}, \mathbf{X}_{jt}), \quad (3)$$

for some link function g . We can utilize the same types of link functions found in generalized linear mixed models. For example if our dyads are in the form of continuous data, we may set g to be the identity; this may arise in, for instance, proximity networks (see, e.g., Olgun et al. 2009), where the distance between individuals is recorded on a regular basis. The common case of modeling binary dyads through the logit link function is yet another example. In Section 2.1 we will go into detail for the context of count data, using a log link function.

In some cases, however, the dyads cannot be modeled directly through a link function as in (3). Instead we can introduce additional latent variables, and then adopt a similar strategy. For example, we may consider a zero inflated model. The zero inflated model is a two component mixture model, where one could introduce additional latent indicator variables which determine whether the observation is coming from the component which is a point mass at zero or the component that has some other density function π^* (e.g., π^* is the Poisson density). We could then model $g(\mathbb{E}_{\pi^*}(y_{ijt}))$ as in (3). This situation may arise in large sparse weighted network data, such as company wide email count networks. Zero-inflated models are certainly not the only possibility of this type of data augmentation,

as we will see in Section 2.2.

For the remainder of the paper we will focus on count data and non-negative continuous edges. We will furthermore utilize the conditional likelihood given by Sewell and Chen (2015b), determined by

$$f_{\Psi}(\mathbf{X}_{it}, \mathbf{X}_{jt}) = \beta_{IN} \left(1 - \frac{d_{ijt}}{r_j}\right) + \beta_{OUT} \left(1 - \frac{d_{ijt}}{r_i}\right), \quad (4)$$

where $d_{ijt} = \|\mathbf{X}_{it} - \mathbf{X}_{jt}\|$ is the distance between actors i and j at time t within the latent space, and $\mathbf{r} = (r_1, r_2, \dots, r_n)$ is a vector of positive actor specific parameters constrained such that $\sum_{i=1}^n r_i = 1$ for model identifiability. For the remainder of the paper we will also, for simplicity, ignore the covariate information $\boldsymbol{\alpha}'\mathbf{w}_{ijt}$. It is straightforward to reincorporate such information into the work that follows.

Each r_i can be thought of as the i^{th} actor’s social reach. That is, a larger value of r_i implies that it is more likely for an edge, either $y_{i,t}$ or y_{it} , to take a larger value. These r_i ’s also hold a geometric interpretation within the latent space, specifically a radius. For example, in the context of binary networks, this radius can be understood to imply that actors inside of each other’s radii have a greater than 1/2 probability of an edge, and actors are outside of each other’s radii have a smaller than 1/2 probability of an edge. The coefficients β_{IN} and β_{OUT} can help in understanding the global structure of the network, insofar as telling us whether activity (tendency to send stronger edges) or popularity (tendency to receive stronger edges) is more important in forming high strength edges. Specifically, $\beta_{IN} > \beta_{OUT}$ implies popularity is more important than activity in the edge formation process, and $\beta_{OUT} > \beta_{IN}$ implies the opposite. If the edges are undirected, then setting $\mathbb{P}(y_{ijt}|\cdot) = \mathbb{P}(y_{jit}|\cdot)$ is equivalent to constraining $\beta_{IN} = \beta_{OUT}$. See Sewell and Chen (2015b) for more details on parameter interpretation.

2.1 Counts

A commonly encountered dyadic type which can be modeled by (3) is where y_{ijt} is a count. This context may exist in the form of counting the number of phone calls, the number of emails, the number of cosponsored legislative bills, the number of passengers or of flights in airline data, etc. We can use the canonical link for a Poisson random variable to determine

the likelihood function in the following way:

$$\mathbb{P}(y_{ijt}|\mathcal{X}_t, \Psi) = \frac{\lambda_{ijt}^{y_{ijt}} \exp(-\lambda_{ijt})}{y_{ijt}!}, \quad y_{ijt} = 0, 1, 2, \dots \quad (5)$$

where

$$\log(\lambda_{ijt}) = \beta_{IN} \left(1 - \frac{d_{ijt}}{r_j}\right) + \beta_{OUT} \left(1 - \frac{d_{ijt}}{r_i}\right). \quad (6)$$

Here $\Psi = (\beta_{IN}, \beta_{OUT}, \mathbf{r}, \tau^2, \sigma^2)$ is the vector of parameters. Thus the likelihood is

$$\mathbb{P}(Y_1, Y_2, \dots, Y_T | \mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T, \Psi) = \prod_{t=1}^T \prod_{i \neq j} \frac{\lambda_{ijt}^{y_{ijt}} \exp(-\lambda_{ijt})}{y_{ijt}!}. \quad (7)$$

2.2 Non-Negative Continuous Edges

Here we consider the case of non-negative continuous real valued edges. These types of networks can occur in many biological contexts, in economic contexts, in length of phone calls, etc. The latent space framework provides a natural way to think about such a weighted network, in that we can consider two actors with large weighted edges between them as very close in the latent space, and two actors with smaller weighted edges as more separated within the latent space.

By embedding the network into a latent space we can better differentiate between zero valued edges. Consider as an example a longitudinal sequence of social networks, where the dyadic variable measured is the amount of time two individuals spent speaking with each other: Suppose at a particular time, person i has a weighted edge of zero with two others, persons j and k . Now persons i and j are potential friends though they have not currently met; meanwhile, persons i and k know each other already and strongly dislike each other. In both cases the measured edges between i and j and between i and k will be the same (zero), but we can differentiate them in two ways. First and foremost we can compare edge probabilities (e.g., $\mathbb{P}(y_{ijt} = 0) \ll \mathbb{P}(y_{ikt} = 0)$). Second, viewing the latent variables as unobserved actor attributes, we can determine the dissimilarity between each pair (e.g., $d_{ijt} \ll d_{ikt}$). The key point here is that we are using all the data, not just the data from the pairs (i, j) and (i, k) , to learn more about such observed zeros; i.e., we are letting all dyads help inform us as to the position of each actor within the latent space. This can be better understood by considering if i and j have many links to the same actors, then the geometric

constraints within the latent space imply that i and j will be close together, whereas the same would not be true if, say, i and k do not have many links to the same actors.

Network data with non-negative continuous edges is a context where there is not an obvious link function g to be applied to the mean of y_{ijt} , but by introducing an additional latent variable we can adopt a similar strategy. In particular, we apply a Tobit model when formulating the likelihood function, letting $y_{ijt} = y_{ijt}^* 1_{\{y_{ijt}^* > 0\}}$, where $1_{\{\cdot\}}$ is the indicator function and y_{ijt}^* is a continuous normal random variable. This type of approach may be most appropriate when the weighted dyads we observe are really proxies for some underlying relationship between the two actors, but we can only observe the effects from positive relationships (e.g., length of phone calls can only serve as a proxy for a relationship between friends, and not between enemies). We then apply (3) to the latent variables y_{ijt}^* , letting g be the identity function, obtaining

$$y_{ijt}^* = \beta_{IN} \left(1 - \frac{d_{ijt}}{r_j}\right) + \beta_{OUT} \left(1 - \frac{d_{ijt}}{r_i}\right) + \epsilon_{ijt}, \quad (8)$$

$$\epsilon_{ijt} | (\mathcal{X}_t, \Psi) \stackrel{iid}{\sim} N(0, \gamma^2). \quad (9)$$

With this we have

$$\pi(y_{ijt} | \mathcal{X}_t, \Psi) = [N(y_{ijt} | \mathbb{E}(y_{ijt}^* | \mathcal{X}_t, \Psi), \gamma^2)]^{1_{\{y_{ijt} > 0\}}} \left[1 - \Phi\left(\frac{\mathbb{E}(y_{ijt}^* | \mathcal{X}_t, \Psi)}{\gamma}\right)\right]^{1_{\{y_{ijt} = 0\}}}, \quad (10)$$

where Φ is the standard normal cumulative distribution function, and $\mathbb{E}(y_{ijt}^* | \mathcal{X}_t, \Psi) = \beta_{IN} (1 - d_{ijt}/r_j) + \beta_{OUT} (1 - d_{ijt}/r_i)$ is the conditional expectation of y_{ijt}^* . The vector of parameters is now supplemented by γ^2 such that $\Psi = (\beta_{IN}, \beta_{OUT}, \gamma^2, \mathbf{r}, \tau^2, \sigma^2)$. Since the ϵ_{ijt} 's are conditionally i.i.d., we have that the observation equation is

$$\mathbb{P}(Y_1, Y_2, \dots, Y_T | \mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T, \Psi) = \prod_{t=1}^T \prod_{i \neq j} \pi(y_{ijt} | \mathcal{X}_t, \Psi). \quad (11)$$

3 Estimation

To obtain estimates of the latent space positions and of the unknown parameters, we sample via a Markov chain Monte Carlo (MCMC) algorithm from the posterior

$$\pi(\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T, \Psi | Y_1, Y_2, \dots, Y_T). \quad (12)$$

The general strategy is to find reasonable estimates of the latent positions and of the model parameters to initialize the chain, and then use a Metropolis-Hastings (MH) within Gibbs sampling to obtain the posterior samples.

The prior for $\mathbf{r}$ was a Dirichlet distribution. The priors for τ^2 , σ^2 and, in the case of continuous data, γ^2 were chosen to be inverse gamma (IG), as these distributions are conjugate for τ^2 and σ^2 . The shape and scale parameters for τ^2 were set to be equal to $2 + \delta$ and $(1 + \delta)\tau_0^2$ respectively for some small δ (we used 0.05) and some positive constant τ_0^2 , and were similarly set for σ^2 and γ^2 . With this parameterization, the prior variances of τ^2 , σ^2 and γ^2 are kept large. The prior set on β_{IN} was a normal distribution with mean ν_{IN} and (large) variance ξ_{IN} , and similarly for β_{OUT} .

In some cases, by taking a preliminary look at the data we can set these hyperparameters to reasonable values; specifically, we may match the prior means to the initialized values given in the next section. This “first glance” allows us to form our prior beliefs about the scale of the parameters, as well as about the individual actor effects. This idea follows the same underlying concept as empirical Bayes methodology, in that we are (albeit to a small degree in comparison to standard empirical Bayes methods) using the data to construct the hyperparameters. This idea was used in Sewell and Chen (2015b) to good effect, and also yielded good results in our analyses. For some other parameters there is not an obvious way in which we can set the hyperparameters in this fashion. As will be described in the simulation study, however, the results were not sensitive to the selection of these hyperparameter values.

3.1 Initialization

We initialized the radii as

$$r_i = \frac{\sum_{t=1}^T \sum_{j \neq i} (y_{ijt} + y_{jit})/2}{\sum_{t'=1}^T \sum_{i'=1}^n \sum_{j' \neq i'} y_{i'j't'}}. \quad (13)$$

The Dirichlet hyperparameters $\mathbf{a} = (a_1, a_2, \dots, a_n)$ for $\mathbf{r}$ were set to be equal to these initial estimates. Doing so sets the prior expected value of r_i to be the initial estimate of r_i , which reflects our prior intuition; additionally, as each a_j would be small (averaging $1/n$), the prior variance for each r_i will be large (leading to a “flat” prior).

To find initial latent positions, we implemented the generalized multidimensional scaling

algorithm (GMDS), as described in Sarkar and Moore (2005). GMDS starts by taking a distance matrix at time 1 and performing classical multidimensional scaling. Then, for each subsequent time period t , $t = 2, 3, \dots, T$, GMDS balances the position matrix from the previous time point with the classical multidimensional scaling result obtained from the distance matrix at time t .

The original distance matrices can be found in a number of ways, but we offer the following suggestion. We treated the data as binary, where

$$y_{ijt}^{(binary)} = \begin{cases} 1 & \text{if } y_{ijt} > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

We then computed each distance d_{ijt} according to

$$d_{ijt} = \begin{cases} \frac{1}{2} \min\{r_i, r_j\} & \text{if } y_{ijt}^{(binary)} = y_{jit}^{(binary)} = 1 \\ \frac{1}{2}(r_i + r_j) & \text{if } y_{ijt}^{(binary)} + y_{jit}^{(binary)} = 1 \\ \frac{3}{2}(r_i + r_j) & \text{if } y_{ijt}^{(binary)} = y_{jit}^{(binary)} = 0. \end{cases} \quad (15)$$

The general idea here is that positive edges indicate a closeness between the actors; using the radii as measures of closeness accounts for the individual effects as well as keeps the distances on the same scale as the radii, as would seem reasonable based on (4). With the T distance matrices computed, we can then implement GMDS to obtain initial latent positions.

The initial estimate for τ^2 was computed (using the initial estimates of $\mathcal{X}_1$) as

$$\frac{1}{np} \sum_{i=1}^n \|\mathbf{X}_{i1}\|^2. \quad (16)$$

In our analyses we also used this value to determine the hyperparameter τ_0^2 , the prior mean of τ^2 .

We found that the initial value of σ^2 did not make a noticeable difference in the performance, nor did the value of σ_0^2 . Similarly, the initial estimates for γ^2 , β_{IN} and β_{OUT} and the values of their hyperparameters did not significantly affect the number of iterations required to reach convergence.

3.2 Posterior Sampling

We implement a MH within Gibbs sampling scheme. The algorithm is

0. Set the initial values of the latent positions and parameters as given in Section 3.1.
1. For $t = 1, 2, \dots, T$ and for $i = 1, 2, \dots, n$, draw $\mathbf{X}_{it}$ via MH.
2. Draw τ^2 from $\pi(\tau^2|\mathcal{X}_1)$.
3. Draw σ^2 from $\pi(\sigma^2|\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T)$.
4. Draw $\mathbf{r}$ via MH.
5. Draw β_{IN} via MH.
6. Draw β_{OUT} via MH.

If data is non-negative continuous

7. Draw γ^2 via MH.

Repeat steps 1-7.

The full conditional distributions needed for steps 2-7 are given in the Appendix. Regarding the proposal distributions, $\mathbf{X}_{it}$, β_{IN} , and β_{OUT} can come from a symmetric proposal (e.g., normal random walk). For γ^2 , however, some asymmetric proposal such as a log-normal (what we used in our analyses) or an inverse gamma distribution ought to be used to ensure positive valued proposals; this asymmetric proposal will then need to be accounted for in the acceptance probability. Because of the constraint on $\mathbf{r}$, a Dirichlet proposal is suggested for the radii, which also will be an asymmetric proposal. Suggested parameters for this Dirichlet proposal are $\kappa \mathbf{r}^{curr}$, where $\mathbf{r}^{curr}$ are the current values for $\mathbf{r}$ and κ is some large value.

One final note is that, as is the case for any such latent space model, the posterior is invariant under rotations, reflections and translations of the latent positions $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T$. In order to make the MCMC iterations comparable, after each iteration of steps 1-7 we perform a Procrustes transformation on the $nT \times p$ matrix $(\mathcal{X}'_1, \dots, \mathcal{X}'_T)'$. The Procrustean transformation finds a set of rotations, reflections and translations to minimize the difference between a given matrix and some target matrix. In our analyses, we constructed the target matrix from the initialized latent position trajectories.

4 Scalability

The MCMC algorithm of Section 3.2 can handle many data sets, including the two that are fully described in Section 6. However, in cases where the network is very large, the MCMC algorithm may prove to be too slow to be viable. For static binary latent space network

models, Raftery et al. (2012) described a method for approximating the log likelihood using case-control principles. Sewell and Chen (2015b) also used this method for binary dynamic latent space network models, adapting it slightly to allow for missing data. Here we extend this for models whose dyads can be described by an exponential family of distributions.

For the MCMC algorithm, the MH steps required in updating the latent positions, $\mathbf{r}$, β_{IN} , β_{OUT} and other likelihood related parameters (e.g., γ^2 in the case of non-negative real dyads) all require $O(Tn^2)$ terms to be summed. In this discussion we will assume here that a non-relationship between two actors implies that $y_{ijt} = 0$, otherwise y_{ijt} is some positive value; the principles discussed next ought to hold even if this is not the case. Generalizing the approximation method first proposed by Raftery et al. (2012), we can reduce this computational cost to $O(Tn)$.

Suppose that, conditional on the latent positions, the y_{ijt} 's are independent with

$$\pi(y_{ijt}|\cdot) = h(y_{ijt}) \exp(\boldsymbol{\eta}'_{ijt} \mathbf{T}(y_{ijt}) + A(\boldsymbol{\eta}_{ijt})), \quad (17)$$

where $\boldsymbol{\eta}_{ijt}$ is a vector valued function of $(\mathcal{X}_t, \boldsymbol{\Psi})$ and $\mathbf{T}(y_{ijt})$ is a vector of sufficient statistics. Then we can rewrite the loglikelihood of $(Y_1, \dots, Y_T)$ as

$$\begin{aligned} \ell(\mathcal{X}_1, \dots, \mathcal{X}_T, \boldsymbol{\Psi}) = & \sum_{t=1}^T \sum_{i=1}^n \left[\sum_{j: y_{ijt} > 0} (\boldsymbol{\eta}'_{ijt} \mathbf{T}(y_{ijt}) + A(\boldsymbol{\eta}_{ijt})) \right. \\ & \left. + \sum_{j: y_{ijt} = 0} (\boldsymbol{\eta}'_{ijt} \mathbf{T}(y_{ijt}) + A(\boldsymbol{\eta}_{ijt})) \right] + \text{constant}. \end{aligned} \quad (18)$$

It is reasonable to assume that as the network gets larger and larger, the number of edges of each node does not grow at the same rate (i.e., the network gets sparser). Hence we make the assumption that either the maximum degree is fixed or is of $o(n)$. If this is the case, then we can, for each i and t , take a subsample $\{j_k\}_{k=1}^{N_{i,t,0}}$ from the set $\{j : y_{ijt} = 0\}$ and use a simple Monte Carlo estimate of the final summation of (18) to reduce the computational cost to linear with respect to n . Then the approximation we use of the log likelihood is

$$\begin{aligned} \ell(\mathcal{X}_1, \dots, \mathcal{X}_T, \boldsymbol{\Psi}) \approx & \sum_{t=1}^T \sum_{i=1}^n \left[\sum_{j: y_{ijt} > 0} (\boldsymbol{\eta}'_{ijt} \mathbf{T}(y_{ijt}) + A(\boldsymbol{\eta}_{ijt})) \right. \\ & \left. + \frac{n_{i,t,0}}{N_{i,t,0}} \sum_{k=1}^{N_{i,t,0}} (\boldsymbol{\eta}'_{ij_k t} \mathbf{T}(y_{ij_k t}) + A(\boldsymbol{\eta}_{ij_k t})) \right] + \text{constant}, \end{aligned} \quad (19)$$

where $n_{i,t,0} = |\{j : y_{ijt} = 0\}|$. In most cases, $\mathbf{T}(y_{ijt}) = 0$ if $y_{ijt} = 0$ and hence the above can be simplified such that the second summation is only $\frac{n_{i,t,0}}{N_{i,t,0}} \sum_{k=1}^{N_{i,t,0}} A(\boldsymbol{\eta}_{ijk t})$. Also, there could potentially be multiple methods of selecting the subsequences $\{j_k\}_{k=1}^{N_{i,t,0}}$; see Raftery et al. (2012) for more details.

For $T = 1$ and $y_{ijt} \in \{0, 1\}$, this leads to Raftery et al.'s approximation. For the context presented in Section 2.1, we can approximate the log likelihood as

$$\ell(\mathcal{X}_1, \dots, \mathcal{X}_T, \boldsymbol{\Psi}) \approx \sum_{t=1}^T \sum_{i=1}^n \left[\sum_{j: y_{ijt} > 0} \left(y_{ijt} \log(\lambda_{ijt}) + \lambda_{ijt} \right) + \frac{n_{i,t,0}}{N_{i,t,0}} \sum_{k=1}^{N_{i,t,0}} \lambda_{ijk t} \right] + \text{constant}. \quad (20)$$

For the context presented in Section 2.2, we can approximate the log likelihood as

$$\begin{aligned} \ell(\mathcal{X}_1, \dots, \mathcal{X}_T, \boldsymbol{\Psi}) \approx & \sum_{t=1}^T \sum_{i=1}^n \left\{ \sum_{j: y_{ijt} > 0} \left[-\frac{1}{2} \log(\gamma^2) - \frac{1}{2\gamma^2} (y_{ijt} - \mathbb{E}(y_{ijt}^* | \mathcal{X}_t, \boldsymbol{\Psi}))^2 \right] \right. \\ & \left. + \frac{n_{i,t,0}}{N_{i,t,0}} \sum_{k=1}^{N_{i,t,0}} \log \left(1 - \Phi(\mathbb{E}(y_{ijk t}^* | \mathcal{X}_t, \boldsymbol{\Psi}) / \gamma) \right) \right\} + \text{constant}. \end{aligned} \quad (21)$$

An interesting point is that if we assume that the network becomes more sparse as n grows, then it may be more appropriate to utilize a zero-inflated model, such as was mentioned in Section 2. Suppose we can augment the data by component indicator variables $z_{ijt} \in \{1, 2\}$, such that $\pi(y_{ijt} | z_{ijt} = 1, \cdot) = \delta(y_{ijt})$, and $\pi(y_{ijt} | z_{ijt} = 2, \cdot)$ can be constructed according to (17), where δ is the Dirac delta function. Letting $\pi(z_{ijt} = 1) = \alpha$, we can then write the approximated complete log likelihood (i.e., $\pi(y_{ijt}, z_{ijt} : 1 \leq i \neq j \leq n, 1 \leq t \leq T | \cdot)$) as

$$\begin{aligned} & \ell(\mathcal{X}_1, \dots, \mathcal{X}_T, \boldsymbol{\Psi}) \\ = & \sum_{t=1}^T \sum_{i=1}^n \left\{ \sum_{j: y_{ijt} > 0} [\log(1 - \alpha) + \boldsymbol{\eta}'_{ijt} \mathbf{T}(y_{ijt}) + A(\boldsymbol{\eta}_{ijt})] \right. \\ & \left. + \frac{n_{i,t,0}}{N_{i,t,0}} \sum_{k=1}^{N_{i,t,0}} \left[1_{\{z_{ijk t}=1\}} \log(\alpha) + 1_{\{z_{ijk t}=2\}} \left(\log(1 - \alpha) + \boldsymbol{\eta}'_{ijk t} \mathbf{T}(y_{ijk t}) + A(\boldsymbol{\eta}_{ijk t}) \right) \right] \right\} + \text{constant}. \end{aligned} \quad (22)$$

Since the z_{ijt} 's are nuisance parameters, we need not sample all of them in the Gibbs sampler, but rather only the z_{ijt} 's corresponding to each of the n subsequences $\{j_k\}_{k=1}^{N_{i,t,0}}$, thus

maintaining the computational cost of $O(Tn)$.

5 Simulations

We analyzed simulated data for both count data and non-negative continuous data. In each case we simulated twenty data sets where the number of actors was 100 and the number of time points was 10. The values used in these simulations, given in the next two sections, were chosen to create data that were similar to the real data we analyzed.

5.1 Simulated Count Data

For each of the twenty simulations, the parameter values were set at $\beta_{IN} = 3$ and $\beta_{OUT} = 1$. The transition variance was set to be $\sigma^2 = 1 \times 10^{-6}$. The latent positions at time 1 were drawn from a mixture of 10 normals with equal mixture component weights, where the cluster means were drawn randomly from a multivariate normal distribution with mean zero and covariance $((9/10) \cdot 2 \times 10^{-5})I_p$, and where the cluster covariances were $((1/10) \cdot 2 \times 10^{-5})I_p$, and $p = 2$ is the dimension of the latent space. After the initial latent positions $\mathcal{X}_1$ were drawn, the radii $\mathbf{r}$ were drawn from a Dirichlet distribution whose i^{th} parameter was equal to $n(1/\|\mathbf{X}_{i1}\|)/\max_k\{1/\|\mathbf{X}_{k1}\|\}$, thus giving those centrally located actors a large individual effect, which reflects a reality. Subsequent latent positions $\mathcal{X}_t$, $t \geq 2$, were drawn according to (2). The adjacency matrices Y_1 to Y_T were then generated according to (5) and (6). The mean proportion of edges that were positive over the simulations was 0.56, ranging from 0.34 to 0.69.

For each simulation we drew σ_0^2 from $U(1 \times 10^{-4}, 1 \times 10^{-2})$, where $U(a, b)$ is the uniform distribution over the interval (a, b) . Both hyperparameters ν_{IN} and ν_{OUT} were for each simulation drawn from $U(1, 15)$; ξ_{IN} and ξ_{OUT} were set to be 1,000.

To evaluate the simulation results, we compared the estimates of the coefficients β_{IN} and β_{OUT} with the truth, evaluated the pseudo R^2 , and evaluated the pairwise ratios of estimated distances to true distances corresponding to the latent positions. The pseudo R^2 value is the deviance based pseudo R^2 for count data found, and recommended, in Cameron and

Windmeijer (1996). This is calculated as

$$R^2 = \frac{\sum_{t=1}^T \sum_{i \neq j} y_{ijt} \log(\hat{\lambda}_{ijt}/\bar{y}) - (\hat{\lambda}_{ijt} - \bar{y})}{\sum_{t'=1}^T \sum_{i' \neq j'} y_{i'j't'} \log(y_{i'j't'}/\bar{y})}, \quad (23)$$

where $\bar{y} = \sum_{t=1}^T \sum_{i \neq j} y_{ijt}$ and $\hat{\lambda}_{ijt}$ is found by plugging in the posterior mean estimates in (6). To clarify what is meant by the distance ratios, note that for each simulation there are $Tn(n-1)/2$ distances within the latent space. We calculate all these pairwise distances using the posterior mean latent positions as well as using the true latent positions. So for each simulation we can plot a curve corresponding to the distribution of these ratios. We would hope for this curve to be narrow and centered at 1.

The posterior mean estimate, averaged over the twenty simulations, for β_{IN} (β_{OUT}) whose true value was 3 (1), was 2.95 (1.01), ranging from 2.84 to 3.01 (0.969 to 1.06). The pseudo R^2 values' average was 0.930, ranging from 0.908 to 0.944, implying that the posterior means fit the data well. The distributions of the ratios of pairwise distances are given in Figure 2(a), where each curve corresponds to a simulation. From this figure we see that the picture we obtain from the estimated latent space is close to the true latent space, up to a rotation/translation, for all but perhaps one simulation; this outlying simulation still yields a narrow distribution, implying that the picture of the latent space is close to the truth up to a scaling factor. In each simulation we are satisfied with the results; considering that σ_0^2 , ν_{IN} , and ν_{OUT} were randomized in each case, we can conclude that the results are not sensitive to these hyperparameters.

5.2 Simulated Continuous Data

For each of the twenty simulations, the parameter values were set at $\beta_{IN} = 3$, $\beta_{OUT} = 1$, $\gamma^2 = 4$, and $\sigma^2 = 1 \times 10^{-6}$. The latent positions at time 1 were drawn from a mixture of 10 normals with equal mixture component weights, where the cluster means were drawn randomly from a multivariate normal distribution with mean zero and covariance $((9/10) \cdot 2 \times 10^{-5})I_p$, and where the cluster covariances were $((1/10) \cdot 2 \times 10^{-5})I_p$, and $p = 2$ is the dimension of the latent space. After the initial latent positions $\mathcal{X}_1$ were drawn, the radii $\mathbf{r}$ were drawn from a Dirichlet distribution whose i^{th} parameter was equal to

Figure 2: Distributions of ratios of pairwise distances for simulated data with count and non-negative real edges. Each curve corresponds to a simulation.

$n(1/\|\mathbf{X}_{i1}\|)/\max_k\{1/\|\mathbf{X}_{k1}\|\}$. Subsequent latent positions $\mathcal{X}_t$, $t \geq 2$, were drawn according to (2). The adjacency matrices Y_1 to Y_T were then constructed by generating y_{ijt}^* according to (8) and letting $y_{ijt} = y_{ijt}^* 1_{\{y_{ijt}^* > 0\}}$. The mean proportion of edges that were positive over the simulations was 0.480, ranging from 0.293 to 0.590.

For each simulation we drew σ_0^2 from $U(1 \times 10^{-4}, 1 \times 10^{-2})$, γ_0^2 from $U(1, 5)$, and both ν_{IN} and ν_{OUT} from $U(1, 15)$; both ξ_{IN} and ξ_{OUT} were set to be 1,000.

To evaluate the simulation results, we compared the estimates of the coefficients β_{IN} and β_{OUT} with the truth, evaluated the pseudo R^2 , and evaluated the pairwise ratios of estimated distance to true distance. In this context of continuous non-negative data, we used the pseudo R^2 value recommended in Veall and Zimmermann (1994), originally derived by McKelvey and Zavoina (1975). This is calculated as

$$R^2 = \frac{\sum_{t=1}^T \sum_{i \neq j} (\hat{y}_{ijt}^* - \hat{y}^*)^2}{\sum_{t'=1}^T \sum_{i' \neq j'} (\hat{y}_{i'j't'}^* - \hat{y}^*)^2 + Tn(n-1)\hat{\gamma}^2}, \quad (24)$$

where $\hat{y}_{ijt}^* = \hat{\beta}_{IN}(1 - \hat{d}_{ijt}/\hat{r}_j) + \hat{\beta}_{OUT}(1 - \hat{d}_{ijt}/\hat{r}_i)$ and $\hat{y}^* = 1/(Tn(n-1)) \sum_{t=1}^T \sum_{i \neq j} \hat{y}_{ijt}^*$. The $\hat{\cdot}$ symbol over the model parameters implies the posterior mean estimate.

The posterior mean estimate, averaged over the twenty simulations, for β_{IN} (β_{OUT}) whose true value was 3 (1), was 2.96 (1.01), ranging from 2.63 to 3.12 (0.938 to 1.11). The pseudo R^2 values' average was 0.854, ranging from 0.660 to 0.982, implying that the posterior means fit the data well. The distributions of the ratios of pairwise distances are given in Figure 2(b),

where each curve corresponds to a simulation. Nearly all of these are narrow and centered near one, and all seem narrow, implying that the picture we obtain from the estimated latent space is close to the true latent space up to a rotation/translation and sometimes a small scalar. In each simulation we are satisfied with the results; considering that σ_0^2 , γ_0^2 , ν_{IN} , and ν_{OUT} were randomized in each case, we can conclude that the results are not sensitive to these hyperparameters.

6 Data Analysis

6.1 Friends and Family Data

We consider the Friends and Family data collected by the MIT Human Dynamics Lab (Aharony et al. 2011). We looked at the mobile phone log, counting the number of calls between each (directed) pair of individuals from October, 2010, to May, 2011. The context of the study is a community of couples, around half of who have children, where one member of each couple is associated with a nearby major research university in North America. Of the approximately 200 applicants, 130 actors of the network were selected in such a way as to represent the full community and sub-communities. The entire community consists of 400 residents of a young family living community. This study captured many aspects of the subjects beyond just the phone log, and among these we will look at religion and race. For more details on the data and the collection process see Aharony et al. (2011).

The edges y_{ijt} of the adjacency matrices Y_t represent the number of phone calls from subject i to subject j . These counts were binned by month, and hence we had $T = 8$ time points. We eliminated any subjects who averaged less than one phone call, incoming or outgoing, per month. This left 119 subjects. Using counts rather than simply considering whether subject i did or did not call subject j during the t^{th} month gives more insight into how gregarious each subject is, as well as better insight into how close each actor is to each other actor with whom they conversed via phone.

Initialization was performed according to Section 3, setting $\sigma_0^2 = 1 \times 10^{-3}$, $\nu_{IN} = \nu_{OUT} = 10$, and $\xi_{IN} = \xi_{OUT} = 1000$. We ran the MCMC algorithm until we obtained 500,000

samples, using a burn-in of 300,000. Figure 3 gives the trace plots for β_{IN} , β_{OUT} , τ^2 and σ^2 ; from this we can visually confirm that the MCMC chain has reached convergence.

Figure 3: MCMC trace plots for the model parameters corresponding to the Friends and Family data. Horizontal axis is in iterations $\times 10^4$.

The pseudo R^2 value was 0.819, implying a very good fit of the data. The posterior means of the coefficients were $\beta_{IN} = 4.17$ and $\beta_{OUT} = 1.29$, implying that, in the friendship network structure, popularity is more important than social activity.

Figure 4 gives a plot of the posterior mean latent positions at times 1, 4, and 8. The actors' shapes indicate the race (Asian, black, Hispanic, middle eastern, or white), and the boxes or circles around the actors indicate their level of religion (either not at all religious or very religious). From these figures we can see that there is some association between race and social position, as well as between religion and social position. To verify this visual inspection, we performed a Mantel test between the posterior means of the latent positions and these two exogenous variables at each time point. More specifically, we compared the distance matrix whose entries are given by $\|\hat{\mathbf{X}}_{it} - \hat{\mathbf{X}}_{jt}\|$, where $\hat{\mathbf{X}}_{it}$ is the posterior mean of $\mathbf{X}_{it}$, to the distance matrix whose entries are 1 if actors i and j are not of the same race and

Figure 4: Posterior means of latent positions for the Friends and Family data at times (a)-(b) 1, (c)-(d) 4, and (e)-(f) 8. Each figure on the right is the zoomed in figure of the dotted box in the figures to their left. Asians are indicated by +, blacks by asterisks, hispanics by solid squares, middle eastern by solid circles, and whites by solid triangles. Actors that are boxed indicated that they are not at all religious, and actors that are circled indicated they are very religious.

Figure 5: Testing association between social position and (a) Race, (b) Religion. The top row is the test statistic, and the bottom row is the p -value. The dotted line is 0.1.

0 otherwise, as well as to the distance matrix whose entries are 1 if actors i and j are not of the same religious dedication and 0 otherwise. The test statistic as well as the p -values are given in Figure 5. Thus from this analysis we have empirical evidence that one's social position is associated with race and religious dedication.

6.2 World Trade Data

World trade data, measuring annual exports/imports between countries in the years 1991-2000, was analyzed. The data were obtained through the Economic Web Institute at <http://www.economicswbinstitute.org/worldtrade.htm>, originally obtained through the IMF Direction of Trade (DOT) Yearbook. Through this site, annual import/export data is available from 1948 to 2000. We selected the most modern subset of this data which provided a reasonable number of countries that were present through all time points (e.g., not considering states that become independent in the midst of the selected time period) as a pedagogical example. The bilateral trade was measured in current millions of U.S. dollars; we analyzed the log of the trade amount, as is common in this context (see, e.g., Ward et al. 2013).

To account for global inflation/deflation and any other non-relational economic effects, the data was rescaled such that the total quantity of annual world trade is constant. What we end up with then is a network consisting of 107 countries who were all involved in world trade through the 10 time years, 1991 to 2000, whose edges are non-negative reals. For more information on the data see Gleditsch (2002).

Ward et al. (2013) developed a complex model for world trade data, combining a common economic model for world trade called the gravity model with aspects of the latent space model developed by Hoff (2005). Their approach used one set of latent variables to model the incidence of trade and another set of latent variables to model the volume of trade. However, if we view the amount of trade between two countries as a positive-valued proxy indicating the strength of the relationship between the two countries' economies, our approach may be more appropriate. Regardless, as the primary purpose of analyzing the world trade data described above is to serve as a pedagogical example of our more general methodology, we have maintained the more simple model framework of Sewell and Chen (2015b) with our extension for weighted network data, demonstrating the data augmentation scheme of Section 2.2.

The hyperparameters for the priors of σ^2 , τ^2 , γ^2 , and $\mathbf{r}$ were formulated according to the description in Section 3. We set $\sigma_0^2 = 1 \times 10^{-3}$, $\gamma_0^2 = 25$, $\nu_{IN} = \nu_{OUT} = 10$, and $\xi_{IN} = \xi_{OUT} = 1000$. Figure 6 gives the trace plots for β_{IN} , β_{OUT} , σ^2 , τ^2 , and γ^2 . A burn-in of 125,000 was used, leaving a chain of length 75,000; from this we can visually confirm that the MCMC chain has reached convergence.

The pseudo R^2 value was 0.970, indicating a very good fit of the data. The estimates for β_{IN} and β_{OUT} were 2.33 and 2.10 respectively, implying that the amount of trade is determined more by the importing country than the exporting country, but only slightly so. If β_{IN} had been much larger than β_{OUT} then this would have suggested that the importer was in larger control of the trade relationship, and if β_{OUT} was much larger than β_{IN} we would say the same about the exporter. However, in our case we see that the two coefficients are close to each other, suggesting that the trade relationship is closely balanced.

Figure 7 gives plots of the posterior mean latent positions, broken up into three time periods: from 1991 to 1993, from 1994 to 1996, and from 1997 to 2000. Temporal direction

Figure 6: MCMC trace plots for the model parameters corresponding to the world trade data. Horizontal axis is in iterations $\times 10^4$.

is shown via arrows. The size of the actor corresponds to its r_i value. Each color represents a geographical region, where green is Africa, yellow is Asia, dark red is Eurasia, blue is Europe, red is North America and the Caribbean, sea green is Oceania, and brown is South America. It is apparent that the actors move within the latent space much less during each of these three periods than during the transition from 1993 to 1994 and from 1996 to 1997. These two major shrinkage events occurring both coincide with major events in world trade. In 1993, the General Agreement on Tariffs and Trade was updated, which would later lead to the creation of the World Trade Organization (WTO) (see <http://www.wto.org>). Looking at Figure 7, we can see that there is already some shrinkage happening during the year 1993 which then continues going into 1994. Specifically we see that certain continents (Africa, Asia, and Europe) come together during this time. Europe is arguably the clearest case, and it turns out there is a good reason for this: the European Economic Area was established on January 1, 1994. Regarding the second shrinkage event in 1997, a publication from the WTO states that “the volume of world merchandise exports grew by 9.5 per cent in 1997.” This is seen visually in Figure 7. However, since the original data had been scaled to account

Figure 7: Latent positions of each nation from the world trade import/export relational data. Each figure on the right is the zoomed in figure of the dotted box in the figures to their left. See main text for key to colors. 23

for such growth, we conclude that the reason for this growth in world exports is not due to existing relationships getting stronger, but rather to the formation of many more trading relationships.

7 Discussion

Using the weights associated with edges makes better use of data than only incorporating the existence or non-existence of an edge. The weighted data is more informative and should lead to more accurate inference. It also eliminates the need to make an arbitrary user-defined cutoff for determining whether an edge should be one or zero.

We have described a general strategy for applying the latent space model for dynamic networks to data consisting of weighted edges. This can be applied either directly or indirectly through additional latent variables. We have demonstrated the flexibility of the latent space models for dynamic networks by modeling cosponsorship count data and non-negative continuous world trade data.

Our latent space models can handle directed edges of many edge types, model both local and global structures, inherently account for transitivity, and yield a rich visualization of the data. An additional note is that using the MH within Gibbs sampling allows edges missing at random or missing completely at random to be incorporated into the model and estimated (see Sewell and Chen 2015b, for details).

Appendix: Full Conditional Distributions

The full conditional distributions for τ^2 and σ^2 are respectively

$$\pi(\tau^2|\mathcal{X}_1) \sim \text{IG}(2 + \delta + np/2, (1 + \delta)\tau_0^2 + \frac{1}{2} \sum_{i=1}^n \|\mathbf{X}_{i1}\|^2), \quad (25)$$

$$\pi(\sigma^2|\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_T) \sim \text{IG}(2 + \delta + \frac{np(T-1)}{2}, (1 + \delta)\sigma_0^2 + \frac{1}{2} \sum_{t=2}^T \sum_{i=1}^n \|\mathbf{X}_{it} - \mathbf{X}_{i(t-1)}\|^2), \quad (26)$$

for $\delta, \tau_0^2, \sigma_0^2 > 0$.

We let $\pi_{y_{ijt}} \triangleq \pi(y_{ijt}|\mathcal{X}_t, \Psi)$ as given in (5) and (6) if we have count edges, or as given in (10) if we have non-negative real edges. Then the full conditional distribution for $\mathbf{X}_{it}$ in these two cases is

$$\pi(\mathbf{X}_{it}|Y_{1:T}, \Psi) \propto \begin{cases} \left(\prod_{j \neq i} \pi_{y_{ijt}} \pi_{y_{jit}} \right) \cdot N(\mathbf{X}_{it}|\mathbf{0}, \tau^2 I_p) \cdot N(\mathbf{X}_{i(t+1)}|\mathbf{X}_{it}, \sigma^2 I_p), & \text{if } t = 1 \\ \left(\prod_{j \neq i} \pi_{y_{ijt}} \pi_{y_{jit}} \right) \cdot N(\mathbf{X}_{i(t+1)}|\mathbf{X}_{it}, \sigma^2 I_p) \cdot N(\mathbf{X}_{it}|\mathbf{X}_{i(t-1)}, \sigma^2 I_p), & \text{if } 1 < t < T \\ \left(\prod_{j \neq i} \pi_{y_{ijt}} \pi_{y_{jit}} \right) \cdot N(\mathbf{X}_{it}|\mathbf{X}_{i(t-1)}, \sigma^2 I_p), & \text{if } t = T. \end{cases} \quad (27)$$

The full conditional distribution for each of the model parameters follows the form

$$\pi(\psi|Y_{1:T}, \mathcal{X}_{1:T}, \Psi \setminus \{\psi\}) \propto \left[\prod_{t=1}^T \pi(Y_t|\mathcal{X}_t, \Psi) \right] \cdot \pi(\psi) \quad (28)$$

where $\Psi \setminus \{\psi\}$ is the set of parameters excluding ψ , and for count data $\psi \in \{\beta_{OUT}, \beta_{IN}, \mathbf{r}\}$ and for non-negative real data $\psi \in \{\beta_{OUT}, \beta_{IN}, \gamma^2, \mathbf{r}\}$.

References

- Aharony, N., Pan, W., Ip, C., Khayal, I., and Pentland, A. (2011), “Social fMRI: investigating and shaping social mechanisms in the real world,” *Pervasive and Mobile Computing*, 7, 643–659.
- Barrat, A., Barthélemy, M., and Vespignani, A. (2005), “The effects of spatial constraints on the evolution of weighted complex networks,” *Journal of Statistical Mechanics: Theory and Experiment*, 2005, P05003.
- Cameron, A. C. and Windmeijer, F. A. (1996), “R-squared measures for count data regression models with applications to health-care utilization,” *Journal of Business & Economic Statistics*, 14, 209–220.
- Durante, D. and Dunson, D. B. (2014), “Nonparametric Bayes dynamic modelling of relational data,” *Biometrika*, 101, 125–138.

- Gleditsch, K. S. (2002), “Expanded trade and GDP data,” *Journal of Conflict Resolution*, 46, 712–724.
- Globerson, A., Chechik, G., Pereira, F., and Tishby, N. (2004), “Euclidean embedding of co-occurrence Data,” *Advances in Neural Information Processing Systems*, 17, 497–504.
- Hanneke, S., Fu, W., and Xing, E. P. (2010), “Discrete temporal models of social networks,” *Electronic Journal of Statistics*, 4, 585–605.
- Hoff, P. D. (2005), “Bilinear mixed-effects models for dyadic data,” *Journal of the American Statistical Association*, 100, 286–295.
- (2011), “Hierarchical multilinear models for multiway data,” *Computational Statistics & Data Analysis*, 55, 530–543.
- Krause, A. E., Frank, K. A., Mason, D. M., Ulanowicz, R. E., and Taylor, W. W. (2003), “Compartments revealed in food-web structure,” *Nature*, 426, 282–285.
- Krivitsky, P. N. (2012), “Exponential-family random graph models for valued networks,” *Electronic Journal of Statistics*, 6, 1100–1128.
- Krivitsky, P. N. and Butts, C. T. (2012), “Exponential-family random graph models for rank-order relational data,” *arXiv:1210.0493*.
- Krivitsky, P. N. and Handcock, M. S. (2014), “A separable model for dynamic networks,” *Journal of the Royal Statistical Society: Series B*, 76, 29–46.
- Kunegis, J., Lommatzsch, A., and Bauckhage, C. (2009), “The slashdot zoo: mining a social network with negative edges,” in *Proceedings of the 18th International Conference on World Wide Web*, ACM, pp. 741–750.
- McKelvey, R. D. and Zavoina, W. (1975), “A statistical model for the analysis of ordinal level dependent variables,” *Journal of Mathematical Sociology*, 4, 103–120.
- Morgan, J. W. (2014), “The latent path model for dynamic social networks with an application to party switching in Poland,” in *New Faces in Political Methodology VII Conference Papers*.

- Newman, M. E. J. (2004), “Analysis of weighted networks,” *Physical Review E*, 70, 056131.
- Olgun, D. O., Gloor, P. A., and Pentland, A. S. (2009), “Capturing individual and group behavior with wearable sensors,” in *Proceedings of the 2009 AAAI Spring Symposium on Human Behavior Modeling, SSS*, vol. 9, pp. 68–74.
- Onnela, J.-P., Saramäki, J., Hyvönen, J., Szabó, G., De Menezes, M. A., Kaski, K., Barabási, A.-L., and Kertész, J. (2007), “Analysis of a large-scale weighted network of one-to-one human communication,” *New Journal of Physics*, 9, 179.
- Opsahl, T., Agneessens, F., and Skvoretz, J. (2010), “Node centrality in weighted networks: Generalizing degree and shortest paths,” *Social Networks*, 32, 245–251.
- Opsahl, T. and Panzarasa, P. (2009), “Clustering in weighted networks,” *Social Networks*, 31, 155–163.
- Raftery, A. E., Niu, X., Hoff, P. D., and Yeung, K. Y. (2012), “Fast inference for the latent space network model using a case-control approximate likelihood,” *Journal of Computational and Graphical Statistics*, 21, 901–919.
- Sarkar, P. and Moore, A. (2005), “Dynamic social network analysis using latent space models,” *ACM SIGKDD Explorations Newsletter*, 7, 31–40.
- Sarkar, P., Siddiqi, S. M., and Gordon, G. J. (2007), “A latent space approach to dynamic embedding of co-occurrence data,” in *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics AISTATS*, vol. 7.
- Sewell, D. K. and Chen, Y. (2015a), “Analysis of the formation of the structure of social networks using latent space models for ranked dynamic networks,” *Journal of the Royal Statistical Society: Series C*, in press, DOI: 10.1111/rssc.12093.
- (2015b), “Latent space models for dynamic networks,” *Journal of the American Statistical Association*, in press, DOI: 10.1080/01621459.2014.988214.
- Veall, M. R. and Zimmermann, K. F. (1994), “Goodness of fit measures in the Tobit model,” *Oxford Bulletin of Economics and Statistics*, 56, 485–499.

- Ward, M. D., Ahlquist, J. S., and Rozenas, A. (2013), “Gravity’s rainbow: A dynamic latent space model for the world trade network,” *Network Science*, 1, 95–118.
- Xing, E. P., Fu, W., and Song, L. (2010), “A state-space mixed membership blockmodel for dynamic network tomography,” *The Annals of Applied Statistics*, 4, 535–566.
- Yang, S. and Knoke, D. (2001), “Optimal connections: strength and distance in valued graphs,” *Social Networks*, 23, 285–295.
- Zhang, B. and Horvath, S. (2005), “A general framework for weighted gene co-expression network analysis,” *Statistical Applications in Genetics and Molecular Biology*, 4, Article 17.