



Provided by the author(s) and University College Dublin Library in accordance with publisher policies. Please cite the published version when available.

Title	Joint Modelling of Multiple Network Views
Authors(s)	Gollini, Isabella; Murphy, Thomas Brendan
Publication date	2014-11-17
Publication information	Journal of Computational and Graphical Statistics, :
Publisher	Taylor and Francis
Link to online version	https://www.insight-centre.org/UCD%20Repository
Item record/more information	http://hdl.handle.net/10197/6562
Publisher's statement	This is an electronic version of an article published in Journal of Computational and Graphical Statistics, vol. #, issue #, pg. numbers: #. Journal of Computational and Graphical Statistics is available online at: www.tandfonline.com/doi/abs/[doi]

Downloaded 2022-08-24T19:16:40Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



Joint Modelling of Multiple Network Views

Isabella Gollini* Thomas Brendan Murphy†

September 22, 2014

Abstract

Latent space models (LSM) for network data were introduced by Hoff et al. (2002) under the basic assumption that each node of the network has an unknown position in a D -dimensional Euclidean latent space: generally the smaller the distance between two nodes in the latent space, the greater their probability of being connected. In this paper we propose a variational inference approach to estimate the intractable posterior of the LSM. In many cases, different network views on the same set of nodes are available. It can therefore be useful to build a model able to jointly summarise the information given by all the network views. For this purpose, we introduce the latent space joint model (LSJM) that merges the information given by multiple network views assuming that the probability of a node being connected with other nodes in each network view is explained by a unique latent variable. This model is demonstrated on the analysis of two datasets: an excerpt of 50 girls from ‘Teenage Friends and Lifestyle Study’ data at three time points and the *Saccharomyces cerevisiae* genetic and physical protein-protein interactions.

Keywords: latent space model, latent variable, multiplex networks, social network analysis, variational methods

*Department of Civil Engineering, University of Bristol, England.

†School of Mathematical Sciences, Complex & Adaptive Systems Laboratory and Insight Research Centre, University College Dublin, Ireland

1 Introduction

Network data consists of a set of nodes and a list of edges between the nodes. Recently there has been a growing interest in the modelling of network data. A number of models have been proposed for network data including exponential random graph models (ERGMs) (Holland and Leinhardt, 1981), stochastic blockmodels (Holland et al., 1983; Airoldi et al., 2008) and latent space models (Hoff et al., 2002; Handcock et al., 2007). Recent reviews of various network modeling approaches include Goldenberg et al. (2010) and Salter-Townshend et al. (2012).

Latent space models (LSM) are a well known family of latent variable models for network data introduced by Hoff et al. (2002) under the basic assumption that each node has an unknown position in a D -dimensional Euclidean latent space: generally the smaller the distance between two nodes in the latent space, the greater the probability of them being connected. Unfortunately, the posterior distribution of the LSM cannot be computed analytically. For this reason we propose a variational inferential approach which proves to be less computationally intensive than the MCMC procedure proposed in Hoff et al. (2002) and can therefore easily handle large networks.

In many cases, multiple network link relations on the same set of nodes are available. Multiple network views, also known as multiplex networks (Mucha et al., 2010), can be intended either as multiple link relations among the nodes of the network or a single link relation observed over different conditions, such as one network evolving over time (longitudinal networks). In order to deal with multiplex networks we present a latent space joint model (LSJM) that merges the information given by the multiple network views by assuming that the probability of a node being connected with other nodes in each view is explained by a unique latent variable. To estimate this model we propose an EM algorithm: the parameter estimates obtained from fitting a LSM for each network view independently are used to approximate the joint posterior distribution of the LSJM; then these results are used to update the parameter estimates of every LSM. This process is iterated until convergence. This model has a wide range of applications. For example in computer science

it is of interest to summarize the different relations (e.g. friend, fan, follower or like) that we observe in social media sites like Facebook, Twitter and YouTube (Tang et al., 2011). Another important application is in systems biology where the joint modeling of physical and genetic protein-protein interactions is of wide interest (Bandyopadhyay et al., 2008). Other contexts in which this model can be useful include social sciences and business marketing (Ansari et al., 2011). The LSJM is demonstrated on the analysis of an excerpt of 50 girls from ‘Teenage Friends and Lifestyle Study’ data at three time points (Pearson and Michell, 2000; Pearson and West, 2003), and two *Saccharomyces cerevisiae* networks (Stark et al., 2006). All the methods of this paper are implemented in the `lvm4net` package for R (R Core Team, 2014).

The paper is organized as follows. Section 2 provides an introduction to latent space models for network data with a particular focus on the variational inference approach to fit the latent space model. In Section 3 we introduce the latent space joint model for multiple network view data. In Section 4 we show how missing link data can be managed using the LSJM. In Section 5 we illustrate the capabilities of the LSJM and we analyze its performance in the presence of missing edges by using cross-validation; the model is illustrated on the two example datasets (Section 5.2-5.3). We conclude, in Section 6 with a discussion of the model.

2 Latent Space Model

Latent space models for network data have been introduced by Hoff et al. (2002) under the basic assumption that each node i has an unknown position \mathbf{z}_i in a D -dimensional Euclidean latent space. The *distance model* is an easy-to-interpret LSM which is based on the distance between the nodes in the latent space. Generally the smaller the distance between two nodes in the latent space, the greater the probability that they connect. This model supposes the network to be intrinsically symmetric since the distance between nodes in the latent space is symmetric and thus it has the feature of being reciprocal: if $y_{ij} = 1$ then the probability

of $y_{ji} = 1$ is large, where y_{ij} is the observed variable that is 1 if we observe a link from node i to node j , and 0 otherwise. For this reason, the distance model is particularly suitable for undirected networks or directed networks that exhibit strong reciprocity.

Let N is the number of observed nodes and let \mathbf{Y} be the $N \times N$ adjacency matrix containing the network information, with entries y_{ij} (where $y_{ij} = 0$ or 1), and null diagonal. Let \mathbf{Z} is a $N \times D$ matrix of latent positions where each row is composed by $\mathbf{z}_i = (z_{i1}, \dots, z_{iD})$ the D -dimensional vector indicating the position of observation i in the D -dimensional latent space. The latent space model can be written as

$$p(\mathbf{Y}|\mathbf{Z}, \alpha) = \prod_{i \neq j}^N p(y_{ij}|\mathbf{z}_i, \mathbf{z}_j, \alpha) = \prod_{i \neq j}^N \frac{\exp(\alpha - |\mathbf{z}_i - \mathbf{z}_j|^2)^{y_{ij}}}{1 + \exp(\alpha - |\mathbf{z}_i - \mathbf{z}_j|^2)}.$$

where for ease of notation $\prod_{i \neq j}^N$ is equivalent to $\prod_{i=1}^N \prod_{j=1, j \neq i}^N$.

We assume the following distributions for the model unknowns, where $p(\alpha) = \mathcal{N}(\xi, \psi^2)$, $p(\mathbf{z}_i) \stackrel{iid}{=} \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_D)$ and σ^2, ξ, ψ^2 are fixed parameters, and the squared Euclidean distance between observations i and j is $|\mathbf{z}_i - \mathbf{z}_j|^2 = \sum_{d=1}^D (z_{id} - z_{jd})^2$. The squared Euclidean distance measure is employed instead of the Euclidean distance used in Hoff et al. (2002). This choice has been made for two main reasons: firstly, it allows one to visualize the data more clearly, giving an higher probability of a link between two close nodes in the latent space and lower probabilities to two nodes lying far away from each other (see Figure 1); secondly it requires fewer approximation steps to be made in the estimation procedure. An empirical comparison of the results obtained fitting a LSM using the Euclidean and the squared Euclidean distance is given in the supplementary material.

The posterior probability is of the unknown (\mathbf{Z}, α) is of the form

$$p(\mathbf{Z}, \alpha|\mathbf{Y}) = C p(\mathbf{Y}|\mathbf{Z}, \alpha) p(\alpha) \prod_{i=1}^N p(\mathbf{z}_i),$$

where C is the unknown normalizing constant.




Figure 1: Probability of $y_{ij} = 1$ as a function of the distance metrics. For $d_{ij} < 1$ the solid lines representing the probability of a link based on the squared Euclidean metric are higher than the dotted lines representing the probability of a link based on the Euclidean distance. For $d_{ij} > 1$ the solid lines decrease more rapidly than the dotted lines.

2.1 Variational Inference Approach

Since the posterior distribution cannot be calculated analytically we make use of a variational inference approach to estimate the model. To do this we aim at maximizing a lower bound of the likelihood function. This approach has been proposed for several latent variable models (Attias, 1999; Jordan et al., 1999) and we refer to Beal (2003) for an extensive introduction to the variational methods. In the statistical network models context, Airoldi et al. (2008) proposed the use of the variational method to fit mixed-membership stochastic blockmodels and Salter-Townshend and Murphy (2013) applied variational methods to fit the Latent Position Cluster Model (Handcock et al., 2007); the Latent Position Cluster Model is an extension of the original LSM in which the latent positions are assumed to come from a Gaussian mixture model.

We define the variational posterior $q(\mathbf{Z}, \alpha | \mathbf{Y})$ introducing the variational parameters $\Theta = (\tilde{\xi}, \tilde{\psi}^2), \tilde{\mathbf{z}}_i$ and $\tilde{\Sigma}$:

$$q(\mathbf{Z}, \alpha | \mathbf{Y}) = q(\alpha) \prod_{i=1}^N q(\mathbf{z}_i),$$

where $q(\alpha) = \mathcal{N}(\tilde{\xi}, \tilde{\psi}^2)$ and $q(\mathbf{z}_i) = \mathcal{N}(\tilde{\mathbf{z}}_i, \tilde{\Sigma})$.

The basic idea behind the variational approach is to find a lower bound of the log marginal likelihood $\log p(\mathbf{Y})$ by introducing the variational posterior distribution $q(\mathbf{Z}, \alpha | \mathbf{Y})$. This approach leads to minimize the Kulback-Leibler divergence between the variational posterior $q(\mathbf{Z}, \alpha | \mathbf{Y})$ and the true posterior $p(\mathbf{Z}, \alpha | \mathbf{Y})$:

$$\begin{aligned}
\text{KL}[q(\mathbf{Z}, \alpha | \mathbf{Y}) || p(\mathbf{Z}, \alpha | \mathbf{Y})] &= - \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log \frac{p(\mathbf{Z}, \alpha | \mathbf{Y})}{q(\mathbf{Z}, \alpha | \mathbf{Y})} d(\mathbf{Z}, \alpha) \\
&= \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log \frac{p(\mathbf{Y}, \mathbf{Z}, \alpha)}{p(\mathbf{Y})q(\mathbf{Z}, \alpha | \mathbf{Y})} d(\mathbf{Z}, \alpha) \\
&= \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log \frac{p(\mathbf{Y}, \mathbf{Z}, \alpha)}{q(\mathbf{Z}, \alpha | \mathbf{Y})} d(\mathbf{Z}, \alpha) - \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log p(\mathbf{Y}) d(\mathbf{Z}, \alpha) \\
&= \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log \frac{p(\mathbf{Y}, \mathbf{Z}, \alpha)}{q(\mathbf{Z}, \alpha | \mathbf{Y})} d(\mathbf{Z}, \alpha) - \log p(\mathbf{Y})
\end{aligned}$$

The last line follows as $\log p(\mathbf{Y})$ is neither a function of \mathbf{Z} and α . From this equation it is evident that minimizing $\text{KL}[q(\mathbf{Z}, \alpha | \mathbf{Y}) || p(\mathbf{Z}, \alpha | \mathbf{Y})]$ corresponds to maximizing the following lower bound:

$$\log p(\mathbf{Y}) \geq \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log \frac{p(\mathbf{Y}, \mathbf{Z}, \alpha)}{q(\mathbf{Z}, \alpha | \mathbf{Y})} d(\mathbf{Z}, \alpha)$$

The Kulback-Leibler divergence between the variational posterior and the true posterior for the LSM can be written as:

$$\begin{aligned}
\text{KL}[q(\mathbf{Z}, \alpha | \mathbf{Y}) || p(\mathbf{Z}, \alpha | \mathbf{Y})] &= \text{KL}[q(\alpha) || p(\alpha)] + \sum_{i=1}^N \text{KL}[q(\mathbf{z}_i) || p(\mathbf{z}_i)] - \mathbb{E}_{q(\mathbf{Z}, \alpha | \mathbf{Y})} [\log(p(\mathbf{Y} | \mathbf{Z}, \alpha))] \\
&= \frac{1}{2} \left(\frac{\tilde{\psi}^2}{\psi^2} - \log \frac{\tilde{\psi}^2}{\psi^2} + \frac{(\tilde{\xi} - \xi)^2}{\psi^2} + ND \log(\sigma^2) - N \log(\det(\tilde{\Sigma})) \right) \\
&\quad + \frac{N}{2\sigma^2} \text{tr}(\tilde{\Sigma}) + \frac{\sum_{i=1}^N \tilde{\mathbf{z}}_i^T \tilde{\mathbf{z}}_i}{2\sigma^2} - \mathbb{E}_{q(\mathbf{Z}, \alpha | \mathbf{Y})} [\log(p(\mathbf{Y} | \mathbf{Z}, \alpha))] - \frac{1 + ND}{2},
\end{aligned}$$

where the expected log-likelihood $\mathbb{E}_{q(\mathbf{Z}, \alpha | \mathbf{Y})} [\log(p(\mathbf{Y} | \mathbf{Z}, \alpha))]$ is approximated using the Jensen's

inequality:

$$\begin{aligned}
\mathbb{E}_{q(\mathbf{z}, \alpha | \mathbf{Y})}[\log(p(\mathbf{Y} | \mathbf{Z}, \alpha))] &= \sum_{i \neq j}^N y_{ij} \mathbb{E}_{q(\mathbf{z}, \alpha | \mathbf{Y})}[\alpha - |\mathbf{z}_i - \mathbf{z}_j|^2] - \mathbb{E}_{q(\mathbf{z}, \alpha | \mathbf{Y})}[\log(1 + \exp(\alpha - |\mathbf{z}_i - \mathbf{z}_j|^2))] \\
&\leq \sum_{i \neq j}^N y_{ij} (\mathbb{E}_{q(\mathbf{z}, \alpha | \mathbf{Y})}[\alpha - |\mathbf{z}_i - \mathbf{z}_j|^2]) - \log(1 + \mathbb{E}_{q(\mathbf{z}, \alpha | \mathbf{Y})}[\exp(\alpha - |\mathbf{z}_i - \mathbf{z}_j|^2)]) \\
&= \sum_{i \neq j}^N y_{ij} (\tilde{\xi} - 2\text{tr}(\tilde{\Sigma}) - |\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j|^2) \\
&\quad - \log \left(1 + \frac{\exp\left(\tilde{\xi} + \frac{1}{2}\tilde{\psi}^2\right)}{\det(\mathbf{I} + 4\tilde{\Sigma})^{\frac{1}{2}}} \exp\left(-(\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j)^T (\mathbf{I} + 4\tilde{\Sigma})^{-1} (\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j)\right) \right).
\end{aligned} \tag{1}$$

From this equation it is possible to understand the computational advantage of the squared Euclidean distance model with respect to the Euclidean distance model. In fact, the expected log-likelihood has been approximated using the Jensen’s inequality whereas Salter-Townshend and Murphy (2013) need to use three first-order Taylor-expansions to fit the model with the Euclidean distance. An alternative approach to approximate the expected log-likelihood is given by the Jaakkola & Jordan bound (Jaakkola and Jordan, 2000), but it would require further approximations, and it would be more difficult to compute.

To estimate the model an EM algorithm (Dempster et al., 1977) can be applied. The EM algorithm consists of two main steps: the first step, called the E-step, aims to estimate the parameters $\tilde{\mathbf{z}}_i, \tilde{\Sigma}$ of the posterior distribution of the latent space positions by maximizing the complete data log-likelihood given all the other parameters Θ . The second step is the M-step where Θ is updated maximizing the complete data log-likelihood given $\tilde{\mathbf{z}}_i$ and $\tilde{\Sigma}$. As observed above, in this context maximizing the log-likelihood corresponds to minimize the Kullback-Leibler divergence between the variational posterior and the true posterior. Therefore the EM algorithm can be written as a function of the Kullback-Leibler divergence, this approach is commonly known as Variational EM algorithm (Jordan et al., 1999). This method scales as $\mathcal{O}(N^2)$, but it converges in just a few iterations, and the calculations performed in the estimation procedure are pretty simple (see the supplementary material for a comparison of

CPU times with other methods and models).

The analytical form of the parameter estimates will be found introducing the first and second order Taylor series expansion approximation of the following function:

$$f(\tilde{\mathbf{z}}_i, \tilde{\Sigma}, \tilde{\xi}, \tilde{\psi}^2) = \log \left(1 + \frac{\exp\left(\tilde{\xi} + \frac{1}{2}\tilde{\psi}^2\right)}{\det(\mathbf{I} + 4\tilde{\Sigma})^{\frac{1}{2}}} \exp\left(-(\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j)^T (\mathbf{I} + 4\tilde{\Sigma})^{-1} (\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j)\right) \right) \quad (2)$$

calculated around the estimates calculated at the previous step of the algorithm (see supplementary material).

Here we outline the Variational EM algorithm on the $(i + 1)th$ iteration:

E-Step Estimate the parameters of the latent posterior distributions $\tilde{\mathbf{z}}_i^{(i+1)}$ and $\tilde{\Sigma}^{(i+1)}$ evaluating:

$$\begin{aligned} \Omega(\Theta; \Theta^{(i)}) &= -\text{KL}(q(\mathbf{Z}, \alpha | \mathbf{Y}) || p(\mathbf{Z}, \alpha | \mathbf{Y})) \\ &= \int q(\mathbf{Z}, \alpha | \mathbf{Y}) \log \frac{p(\mathbf{Z}, \alpha | \mathbf{Y})}{q(\mathbf{Z}, \alpha | \mathbf{Y})} d(\mathbf{Z}, \alpha), \end{aligned}$$

where $\Theta = (\tilde{\xi}, \tilde{\psi}^2)$. This gives

$$\tilde{\Sigma}^{(i+1)} = \frac{N}{2} \left[\left(\frac{N}{2\sigma^2} + 2 \sum_{i=1}^N \sum_{j \neq i} y_{ij} \right) \mathbf{I} + J(\tilde{\Sigma}^{(i)}) \right]^{-1},$$

where J is the Jacobian matrix of $f(\tilde{\mathbf{z}}_i^{(i)}, \tilde{\Sigma}^{(i)}, \tilde{\xi}^{(i)}, \tilde{\psi}^{2(i)})$ (Equation 2) evaluated at $\tilde{\Sigma} = \tilde{\Sigma}^{(i)}$.

And

$$\tilde{\mathbf{z}}_i^{(i+1)} = \left[\left(\frac{1}{2\sigma^2} + \sum_{j \neq i} (y_{ji} + y_{ij}) \right) \mathbf{I} + H(\tilde{\mathbf{z}}_i^{(i)}) \right]^{-1} \left[\sum_{j \neq i} (y_{ji} + y_{ij}) \tilde{\mathbf{z}}_j - G(\tilde{\mathbf{z}}_i^{(i)}) + H(\tilde{\mathbf{z}}_i^{(i)}) \tilde{\mathbf{z}}_i^{(i)} \right], \quad (3)$$

where G is the gradient and H is the Hessian matrix of $f(\tilde{\mathbf{z}}_i^{(i)}, \tilde{\Sigma}^{(i+1)}, \tilde{\xi}^{(i)}, \tilde{\psi}^{2(i)})$ evaluated at $\tilde{\mathbf{z}}_i = \tilde{\mathbf{z}}_i^{(i)}$.

M-Step Estimate the parameters of the posterior distribution of α evaluating:

$$\Theta^{(i+1)} = \text{argmax} \Omega(\Theta; \Theta^{(i)}).$$

This gives

$$\tilde{\xi}^{(i+1)} = \frac{\xi + \psi^2 (\sum_{i=1}^N \sum_{j \neq i} y_{ij} - f'(\tilde{\xi}^{(i)}) + \tilde{\xi}^{(i)} f''(\tilde{\xi}^{(i)}))}{1 + \psi^2 f''(\tilde{\xi}^{(i)})},$$

where f' and f'' are the first and the second derivatives of $f(\tilde{\mathbf{z}}_i^{(i+1)}, \tilde{\boldsymbol{\Sigma}}^{(i+1)}, \tilde{\xi}^{(i)}, \tilde{\psi}^{2(i)})$ evaluated at $\tilde{\xi} = \tilde{\xi}^{(i)}$, and

$$\tilde{\psi}^{2(i+1)} = \left(\frac{1}{\psi^2} + 2f'(\tilde{\psi}^{2(i)}) \right)^{-1},$$

where f' is the first derivative of $f(\tilde{\mathbf{z}}_i^{(i+1)}, \tilde{\boldsymbol{\Sigma}}^{(i+1)}, \tilde{\xi}^{(i+1)}, \tilde{\psi}^{2(i)})$ evaluated at $\tilde{\psi}^2 = \tilde{\psi}^{(i)2}$.

3 Latent Space Joint Model

Let us suppose to have K network views on the same set of N nodes. We introduce a model which assumes that a continuous latent variable $\mathbf{z}_i \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_D)$ is able to summarize the information given by all the network views $\mathbf{Y}_1, \dots, \mathbf{Y}_K$ identifying the position of node i in a D -dimensional latent space. In this case the network data assumes conditional independence given the latent variable. Our purpose is to model each network view by using a LSM (see Figure 2).




Figure 2: Latent space joint model for two network views.

This yields to the following joint model:

$$\begin{aligned} p(\mathbf{Y}_1, \dots, \mathbf{Y}_K | \mathbf{Z}, \alpha_1, \dots, \alpha_K) &= \prod_{k=1}^K p(\mathbf{Y}_k | \mathbf{Z}; \alpha_k) \\ &= \prod_{k=1}^K \prod_{i \neq j}^N \frac{\exp(\alpha_k - |\mathbf{z}_i - \mathbf{z}_j|^2)^{y_{ijk}}}{1 + \exp(\alpha_k - |\mathbf{z}_i - \mathbf{z}_j|^2)} \end{aligned}$$

where $p(\alpha_k) = \mathcal{N}(\xi_k, \psi_k^2)$, $p(\mathbf{z}_i) \stackrel{iid}{=} \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_D)$ with $\sigma^2, \xi_k, \psi_k^2$ set to be fixed parameters, and the dyad y_{ijk} takes value 1 if there is a link between node i and node j in network k , and 0 otherwise.

The following identity allows one to find the model parameters $\alpha_1, \dots, \alpha_K$ and the posterior distribution of the latent variable \mathbf{z}_i given the K models:

$$\begin{aligned}
p(\mathbf{z}_i | \mathbf{Y}_1, \dots, \mathbf{Y}_K; \alpha_1, \dots, \alpha_K) &\propto p(\mathbf{z}_i) \prod_{k=1}^K p(\mathbf{Y}_k | \mathbf{z}_i; \alpha_k) \\
&\propto p(\mathbf{z}_i) \prod_{k=1}^K \frac{p(\mathbf{z}_i | \mathbf{Y}_k; \alpha_k)}{p(\mathbf{z}_i)} \\
&= \frac{\prod_{k=1}^K p(\mathbf{z}_i | \mathbf{Y}_k; \alpha_k)}{p(\mathbf{z}_i)^{K-1}}.
\end{aligned} \tag{4}$$

Applying the variational inference approach presented in Section 2.1 introducing the variational parameters $\Theta_k = (\tilde{\xi}_k, \tilde{\psi}_k^2)$, $\tilde{\mathbf{z}}_{ik}$ and $\tilde{\Sigma}_k$, it is possible to approximate $p(\mathbf{z}_i | \mathbf{Y}_k; \alpha_k)$ with $q(\mathbf{z}_i) \sim \mathcal{N}(\tilde{\mathbf{z}}_{ik}, \tilde{\Sigma}_k)$. Recalling that $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_D)$ and Equation 3, we obtain that the posterior distribution of the latent variables given all the network views can be written as:

$$p(\mathbf{z}_i | \mathbf{Y}_1, \dots, \mathbf{Y}_K; \Theta_1, \dots, \Theta_K) \propto \mathcal{N}(\bar{\mathbf{z}}_i, \bar{\Sigma}),$$

where the parameters are

$$\bar{\Sigma} = \left[\sum_{k=1}^K \tilde{\Sigma}_k^{-1} - \frac{K-1}{\sigma^2} \mathbf{I}_D \right]^{-1} \quad \text{and} \quad \bar{\mathbf{z}}_i = \bar{\Sigma} \left[\sum_{k=1}^K \tilde{\Sigma}_k^{-1} \tilde{\mathbf{z}}_{ik} \right]. \tag{5}$$

By fitting LSJM we get information on $p(\mathbf{z}_i | \mathbf{Y}_1, \dots, \mathbf{Y}_K; \Theta_1, \dots, \Theta_K)$ and $p(\mathbf{z}_i | \mathbf{Y}_k; \Theta_k)$ for $k = 1, \dots, K$. This way it is possible to have estimates for both the overall positions $\bar{\mathbf{z}}_i$ and the position given one particular network view $\tilde{\mathbf{z}}_{ik}$.

The estimates of $\tilde{\mathbf{z}}_{ik}$ and $\tilde{\Sigma}_k$ are updated from $\bar{\mathbf{z}}_i$ and $\bar{\Sigma}$, so we can locate the unconnected nodes or subgraphs in the latent space depending on their position conditional on the other network views, avoiding the usual tendency of pushing away the unconnected nodes to maximize the likelihood when fitting the classical LSM. This approach also allows one to plot the positions given each network view in the same latent space, and to look at how the nodes in each network change the positions.

Let $(\Theta_1^{(i)}, \dots, \Theta_K^{(i)})$ be the current estimates of $(\Theta_1, \dots, \Theta_K)$ and initialize $(\Theta_1^{(0)}, \dots, \Theta_K^{(0)})$.

The Variational EM algorithm at the iteration $(i + 1)$ can be summarized as follows:

E-Step Estimate the parameters $\bar{\Sigma}^{(i+1)}$ and $\bar{\mathbf{z}}_i^{(i+1)}$ of the posterior distribution of the latent variables given all the network views evaluating:

$$\begin{aligned}
\mathcal{Q}(\Theta_1, \dots, \Theta_K; \Theta_1^{(i)}, \dots, \Theta_K^{(i)}) &= \\
&= \mathbb{E}_{p(\mathbf{Z}|\mathbf{Y}_1, \dots, \mathbf{Y}_K; \Theta_1^{(i)}, \dots, \Theta_K^{(i)})} [\log(p(\mathbf{Y}_1, \dots, \mathbf{Y}_K, \mathbf{Z}|\Theta_1, \dots, \Theta_K))] \\
&= \sum_{k=1}^K \mathbb{E}_{p(\mathbf{Z}|\mathbf{Y}_1, \dots, \mathbf{Y}_K; \Theta_1^{(i)}, \dots, \Theta_K^{(i)})} [\log(p(\mathbf{Y}_k, \mathbf{Z}|\Theta_k))] \\
&\quad - (K-1) \mathbb{E}_{p(\mathbf{Z}|\mathbf{Y}_1, \dots, \mathbf{Y}_K; \Theta_1^{(i)}, \dots, \Theta_K^{(i)})} [\log(p(\mathbf{Z}))].
\end{aligned} \tag{6}$$

Thus we can estimate the parameters of the posterior distribution $p(\mathbf{z}_i|\mathbf{Y}_k; \Theta_k)$ given each network separately:

$$\tilde{\Sigma}_k^{(i+1)} = \frac{N}{2} \left[\left(\frac{N}{2\sigma^2} + 2 \sum_{i=1}^N \sum_{j \neq i} y_{ijk} \right) \mathbf{I} + J_k(\bar{\Sigma}^{(i)}) \right]^{-1},$$

where J_k is the Jacobian matrix of $f(\bar{\mathbf{z}}_i^{(i)}, \bar{\Sigma}^{(i)}, \tilde{\xi}_k^{(i)}, \tilde{\psi}_k^{2(i)})$ (Equation 2) evaluated at $\tilde{\Sigma} = \bar{\Sigma}^{(i)}$, and,

$$\tilde{\mathbf{z}}_{ik}^{(i+1)} = \left[\left(\frac{1}{2\sigma^2} + \sum_{j \neq i} (y_{jik} + y_{ijk}) \right) \mathbf{I} + H_k(\bar{\mathbf{z}}_i^{(i)}) \right]^{-1} \left[\sum_{j \neq i} (y_{jik} + y_{ijk}) \bar{\mathbf{z}}_j^{(i)} - G_k(\bar{\mathbf{z}}_i^{(i)}) + H_k(\bar{\mathbf{z}}_i^{(i)}) \bar{\mathbf{z}}_i^{(i)} \right], \tag{7}$$

where G_k and H_k are respectively the gradient and the Hessian matrices of $f(\bar{\mathbf{z}}_i^{(i)}, \bar{\Sigma}^{(i+1)}, \tilde{\xi}_k^{(i)}, \tilde{\psi}_k^{2(i)})$ evaluated at $\tilde{\mathbf{z}}_i = \bar{\mathbf{z}}_i^{(i)}$.

The posterior distribution of the latent positions given all the network views is estimated merging the estimates of the single models:

$$p(\mathbf{z}_i|\mathbf{Y}_1, \dots, \mathbf{Y}_K; \Theta_1, \dots, \Theta_K) \propto \mathcal{N}(\bar{\mathbf{z}}_i, \bar{\Sigma}),$$

where

$$\bar{\Sigma}^{(i+1)} = \left[\sum_{k=1}^K [\tilde{\Sigma}_k^{(i+1)}]^{-1} - \frac{K-1}{\sigma^2} \mathbf{I}_D \right]^{-1}, \quad \text{and} \quad \bar{\mathbf{z}}_i^{(i+1)} = \bar{\Sigma}^{(i+1)} \left[\sum_{k=1}^K [\tilde{\Sigma}_k^{(i+1)}]^{-1} \tilde{\mathbf{z}}_{ik}^{(i+1)} \right].$$

M-Step Update the model parameters evaluating

$$(\Theta_1^{(i+1)}, \dots, \Theta_K^{(i+1)}) = \operatorname{argmax} \mathcal{Q}(\Theta_1, \dots, \Theta_K; \Theta_1^{(i)}, \dots, \Theta_K^{(i)}).$$

This gives

$$\tilde{\xi}_k^{(i+1)} = \frac{\xi_k + \psi_k^2 (\sum_{i=1}^N \sum_{j \neq i} y_{ijk} - f'_k(\tilde{\xi}_k^{(i)}) + \tilde{\xi}_k^{(i)} f''_k(\tilde{\xi}_k^{(i)}))}{1 + \psi_k^2 f''_k(\tilde{\xi}_k^{(i)})},$$

where f'_k and f''_k are the first and the second derivatives of $f(\bar{\mathbf{z}}_i^{(i+1)}, \bar{\mathbf{\Sigma}}^{(i+1)}, \tilde{\xi}_k^{(i)}, \tilde{\psi}_k^{2(i)})$ evaluated at $\tilde{\xi} = \tilde{\xi}_k^{(i)}$, and

$$\tilde{\psi}_k^{2(i+1)} = \left(\frac{1}{\psi_k^2} + 2f'_k(\tilde{\psi}_k^{(i)2}) \right)^{-1},$$

where f'_k is the first derivative of $f(\bar{\mathbf{z}}_i^{(i+1)}, \bar{\mathbf{\Sigma}}^{(i+1)}, \tilde{\xi}_k^{(i+1)}, \tilde{\psi}_k^{2(i)})$ evaluated at $\tilde{\psi}^2 = \tilde{\psi}_k^{(i)2}$.

4 Missing Link Data

Recent Bayesian approaches to predict missing links in network data have been introduced by Hoff (2009) who proposed to use multiplicative latent factor models, and Koskinen et al. (2010, 2013) in the context of Bayesian exponential random graph models. However there is lack of methods able to take into account the information given by the presence of multiple network views.

Missing (unobserved) links can be easily managed by the LSJM using the information given by all the network views. To estimate the probability of the presence ($y_{ijk} = 1$) or absence ($y_{ijk} = 0$) of an edge we employ the posterior mean of the α_k and of the latent positions so that we get the following equation:

$$y_{ijk}^* = p(y_{ijk} = 1 | \bar{\mathbf{z}}_i, \bar{\mathbf{z}}_j, \tilde{\xi}_k) = \frac{\exp(\tilde{\xi}_k - |\bar{\mathbf{z}}_i - \bar{\mathbf{z}}_j|^2)}{1 + \exp(\tilde{\xi}_k - |\bar{\mathbf{z}}_i - \bar{\mathbf{z}}_j|^2)}.$$

If we want to infer whether to assign $y_{ijk} = 1$ or $y_{ijk} = 0$, we need to introduce a threshold τ_k , and let $y_{ijk} = 1$ if $p(y_{ijk} = 1 | \tilde{\mathbf{z}}_{ik}, \tilde{\mathbf{z}}_{jk}, \tilde{\xi}_k) > \tau_k$, and $y_{ijk} = 0$ otherwise. We set the threshold to be equal to the median probability of a link for the subset of the actual observed links in network k .

To evaluate link prediction in presence of missing links, we use a 10-fold cross validation procedure consisting of randomly splitting the set of all the possible dyads in each network view into 10 subsets. We can predict the links of each subset given the others fitting a LSM

to each network independently and then fitting the LSJM. Finally we can compare the link prediction performance given by these two methods.

The LSJM allows one to locate in the latent space a missing node (no information about links sent and received by the node) in one network by employing the information provided by the other network views. We evaluate the link prediction approach for the missing nodes applying a 10-fold cross validation on the nodes, randomly dividing the set of nodes in each network into 10 subsets, and then predicting the links using the LSJM. We do not use a single LSM to locate missing nodes in the latent space and to estimate their probabilities of links since the only information that the model would use estimating of the latent positions would be the prior distribution of the nodes $p(\mathbf{z}_i)$.

To facilitate the interpretation of the results we matched the rotation of the latent positions in the single LSM with the ones obtained from the LSJM.

5 Applications

5.1 Computational Aspects

The LSM and the LSJM have been fitted assuming that $p(\alpha) = \mathcal{N}(0, 2)$ and $p(\mathbf{z}_i) \stackrel{iid}{=} \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$, initializing the variational parameters $\tilde{\xi}_k = 0$ and $\tilde{\psi}_k^2 = 2$, and latent positions $\bar{\mathbf{z}}_i$ by random generated numbers from $\mathcal{N}(\mathbf{0}, \mathbf{I}_2)$ and setting $\bar{\Sigma} = \mathbf{I}_2$.

We have set the latent space to be bi-dimensional in order to be able to visualize and easily interpret the results. Ten random starts of the algorithm were used and the solution returning the maximum expected likelihood value was selected. The latent positions are identifiable up to a rotation of the latent space. For this reason to speed up the convergence of the algorithm we matched the estimates of $\tilde{\mathbf{z}}_{ik} \forall k > 1$ with $\tilde{\mathbf{z}}_{i1}$ via singular value decomposition in the first 10 iterations of the EM algorithm.

The EM algorithm was stopped after at least 10 iterations when

$$\|\mathbb{E}_q[\log(p(\mathbf{Y}|\mathbf{Z}, \alpha))]^{(i+1)} - \mathbb{E}_q[\log(p(\mathbf{Y}|\mathbf{Z}, \alpha))]^{(i)}\| < tol,$$

where $\mathbb{E}_q[\log(p(\mathbf{Y}|\mathbf{Z}, \alpha))]$ is given by Equation 1 if we fit the LSM or by Equation 6 if we fit the LSJM, i indicates the iteration, tol is a desired tolerance value (which, in this case, was set $tol = 10^{-2}$).

To assess the fit of the model we evaluated the in-sample predictions producing the ROC curve of the estimated link probabilities and calculating the area under the curve (AUC) and the boxplots of the estimated link probabilities for both the true positive and the true negative links. We estimated the probability of a link under each network view by calculating:

$$p(y_{ijk} = 1|\tilde{\mathbf{z}}_{ik}, \tilde{\mathbf{z}}_{jk}, \tilde{\xi}_k) = \frac{\exp(\tilde{\xi}_k - |\tilde{\mathbf{z}}_{ik} - \tilde{\mathbf{z}}_{jk}|^2)}{1 + \exp(\tilde{\xi}_k - |\tilde{\mathbf{z}}_{ik} - \tilde{\mathbf{z}}_{jk}|^2)}.$$

All the calculations have been done using the R package `lvm4net`.

5.2 Excerpt of 50 girls from ‘Teenage Friends and Lifestyle Study’

Pearson and Michell (2000) and Pearson and West (2003) collected data for a ‘Teenage Friends and Lifestyle Study’. The dataset contains three directed networks about friendship relations between students in a school in Glasgow, Scotland. Each student was asked to name up to six best friends in the cohort. The data comes from three yearly waves, from 1995 to 1997. An extended description of all the data in the study can be found in Pearson and Michell (2000). In this paper we will focus on an excerpt of 50 girls that were present at all three measurement points. This dataset is available at the SIENA software website¹ (Ripley et al. (2011)). The networks are formed of 113, 116 and 122 links respectively, and have density of 0.046, 0.047, 0.049. Their degree distributions are shown in Figure 3.

5.2.1 LSM

In this section we applied the LSM to the excerpt of 50 girls from the ‘Teenage Friends and Lifestyle Study’ data.

We fitted a LSM to each network separately assuming that $p(\alpha) = \mathcal{N}(0, 2)$ and $p(\mathbf{z}_i) \stackrel{iid}{=} \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$. The estimated latent posterior positions $\tilde{\mathbf{z}}_i$ (defined in Equation 3) are shown in

¹<http://www.stats.ox.ac.uk/~snijders/siena/>

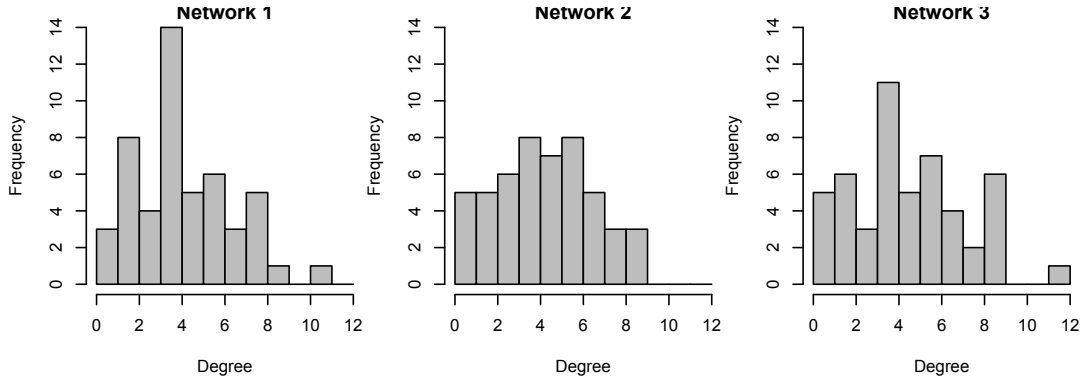


Figure 3: Degree Distribution for the Girls Networks.

Figure 4. A matching rotation of the final estimates was applied in order to facilitate the interpretation of the results. The posterior distributions for the α parameters are quite similar over time $q(\alpha_1) = \mathcal{N}(-0.63, 0.01)$, $q(\alpha_2) = \mathcal{N}(-0.66, 0.01)$ and $q(\alpha_3) = \mathcal{N}(-0.48, 0.01)$; this means that there are not big changes in terms of network density over time. From the




Figure 4: Latent positions $\tilde{\mathbf{z}}_i$ for the Girls Networks fitting the LSM. The grey ellipses represent the 95% approximate credible intervals. Overlapping approximate credible intervals make darker shades of grey.

analysis of the ROC curve (Figure 5 left) and AUC of the in-sample estimated link probabilities, it seems clear that the proposed LSM fits quite well the data separately. The boxplots in Figure 5 (right) show that the estimated probabilities distinguish quite well the true negatives with low probability of forming a link from the true positives with high probability of forming a link.

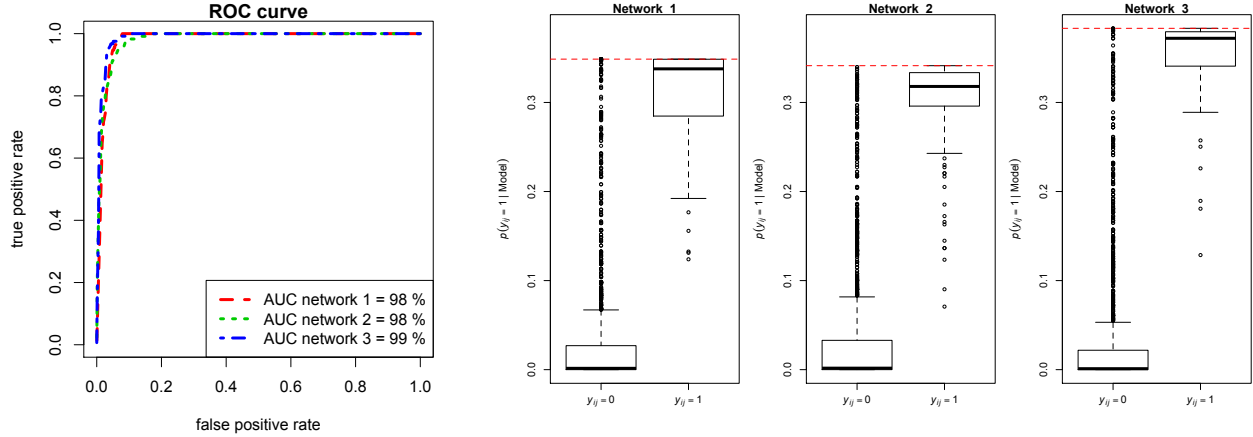


Figure 5: ROC curves (left) and boxplots (right) of the estimated probabilities of a link for the true negatives and true positives obtained by fitting the LSM to the three Girls Networks

5.2.2 LSJM

In this section we fitted the LSJM to the Girls dataset. The overall positions $\bar{\mathbf{z}}_i$ (defined in Equation 5) in the latent space are displayed in Figure 6. In Figure 6 (right) the dots represent the overall positions $\bar{\mathbf{z}}_i$ and the arrows connect the estimated positions $\tilde{\mathbf{z}}_{ik}$ under each model k (defined in Equation 7) from network $k = 1$ to $k = 3$, so that it is possible to see their variation over time.

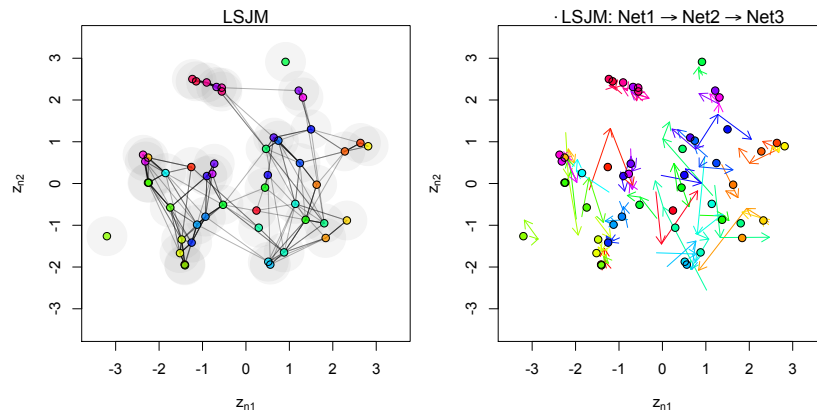


Figure 6: On the left are shown the latent positions $\bar{\mathbf{z}}_i$ for the Girls Networks fitting the LSJM. In the plot on the right the dots represent the overall positions $\bar{\mathbf{z}}_i$ and the arrows connect the estimated position under each model $\tilde{\mathbf{z}}_{i1}, \tilde{\mathbf{z}}_{i2}, \tilde{\mathbf{z}}_{i3}$.

Figure 7 shows the estimated positions $\tilde{\mathbf{z}}_{ik}$ under each model in the latent space. These plots allow a direct comparison between the positions in the latent space $\tilde{\mathbf{z}}_{ik}$ given by the LSJM and the ones obtained by fitting a single LSM for each network view (Figure 4). It is interesting to observe that in the LSJM context the proximity between the disconnected components in network 3 depends on their previous relations. The posterior probabilities of the model parameters $q(\alpha_1) = \mathcal{N}(-0.42, 0.01)$, $q(\alpha_2) = \mathcal{N}(-0.39, 0.01)$ and $q(\alpha_3) = \mathcal{N}(-0.32, 0.01)$ are lower than the single LSM approaches (Section 5.2), implying that a given distance d between two nodes in the latent space corresponds to a higher probability of a link. From the ROC curves, AUC and the boxplots (Figure 8) of the estimated




Figure 7: Estimated latent positions under each model $\tilde{\mathbf{z}}_{i1}, \tilde{\mathbf{z}}_{i2}, \tilde{\mathbf{z}}_{i3}$ for the Girls Networks fitting the LSJM. The grey ellipses represent the 95% approximate credible intervals.

probabilities of a link, it appears clear that the LSJM fits the data quite well. To evaluate the link prediction we applied a 10-fold cross validation setting 10% of the links to be missing at each time point. The area under the ROC curves are 97, 96, 99% fitting the LSJM and 89, 97, 98% fitting the LSM, it shows that the estimates of the links are quite good, especially when applying the LSJM.

Setting $p(y_{ijk} = 1 | \tilde{\mathbf{z}}_{ik}, \tilde{\mathbf{z}}_{jk}, \tilde{\xi}_k) > \tau_k$ where τ_k is equal to the median probability of a link for the subgroup of the actual observed links in network k and applying the LSJM we obtained a misclassification rate of 4% for every network, whereas applying three single LSM we obtain a misclassification rate of 4% for network 1, and 5% for network 2 and 3.

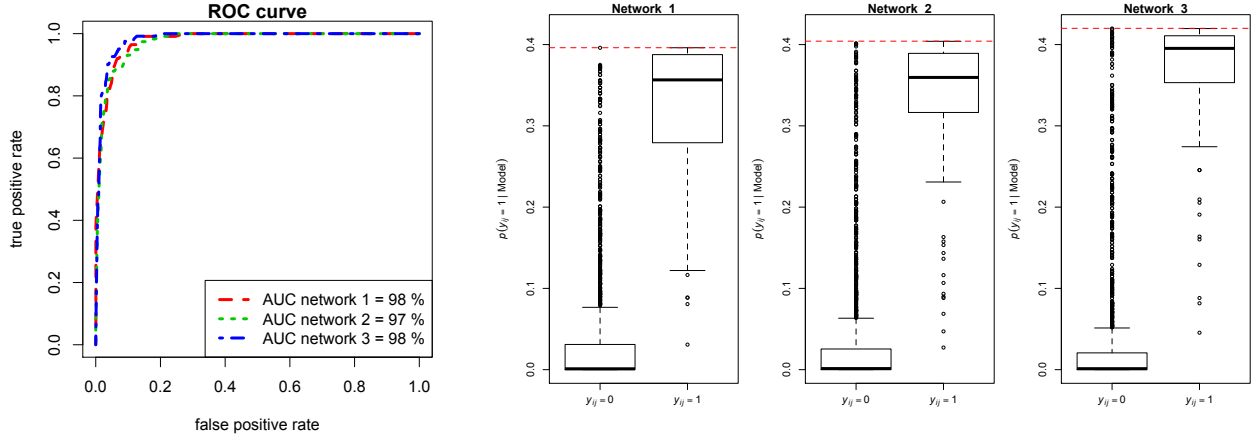


Figure 8: ROC curves and boxplots of the estimated probabilities of a link for the true negatives and true positives obtained by fitting the LSJM to the Girls Networks.

The estimated networks are formed of 109, 115 and 118 links respectively, and have density of 0.044, 0.047, 0.048. Their degree distributions are shown in Figure 9. Comparing these results with the true network statistics (Section 5.2, and Figure 3) it's possible to observe that the distributions are quite similar.




Figure 9: Estimated Degree Distribution for the Girls Networks fitting the LSJM.

The LSJM allows one to also manage missing nodes, to do this we applied a 10-fold cross validation setting the 10% of the nodes in each network to be missing. We obtained a misclassification rate of 9% for all the three networks. As mentioned above in this case the LSM approach would be useless since it would locate the nodes only relying on the prior information.

5.3 Saccharomyces Cerevisiae Protein-Protein Interactions

We analyse a dataset containing two undirected networks formed by genetic and physical protein-protein interactions (PPI) between 67 *Saccharomyces cerevisiae* proteins. The genetic interactions network is formed of 294 links, and its density is 0.066, while the physical interactions network is formed of 190 links, and its density is 0.043. Their degree distributions are shown in Figure 10. The complex relational structure of this dataset has led

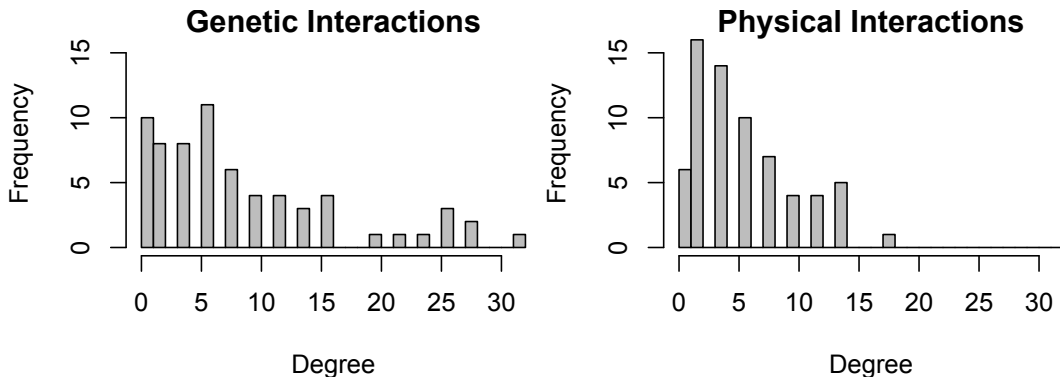


Figure 10: Degree Distribution for the PPI dataset.

to implementation of models aiming at describing the functional relationships between the observations (Bandyopadhyay et al., 2008; Troyanskaya et al., 2003). A list of proteins included in this dataset is displayed in the supplementary material. The dataset is available in the `lvm4net` package, and was downloaded from the Biological General Repository for Interaction Datasets (BioGRID) database² (Stark et al. (2006)). We refer to Stark et al. (2006, 2011) for a description of BioGRID, and for details regarding how the data were collected.

5.3.1 LSM

We fit the LSM to the PPI data working with the genetic and physical interaction networks separately. We assume that $p(\alpha) = \mathcal{N}(0, 2)$ and $p(\mathbf{z}_i) \stackrel{iid}{=} \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$. The estimated latent positions $\tilde{\mathbf{z}}_i$ (defined in Equation 3) are shown in Figure 11. The posterior distributions of the α 's $q(\alpha_1) = \mathcal{N}(-0.332, 0.003)$ and $q(\alpha_2) = \mathcal{N}(-1.001, 0.005)$ indicate that network 1

²<http://thebiogrid.org/>




Figure 11: Latent positions $\tilde{\mathbf{z}}_i$ for the Genetic (left) and Physical (right) interaction data fitting the LSM.

regarding the genetic interactions is much more dense than network 2 which refers to the physical interactions. The ROC curves, AUC and the boxplots (Figure 12) show that the proposed LSM fit the data quite well.

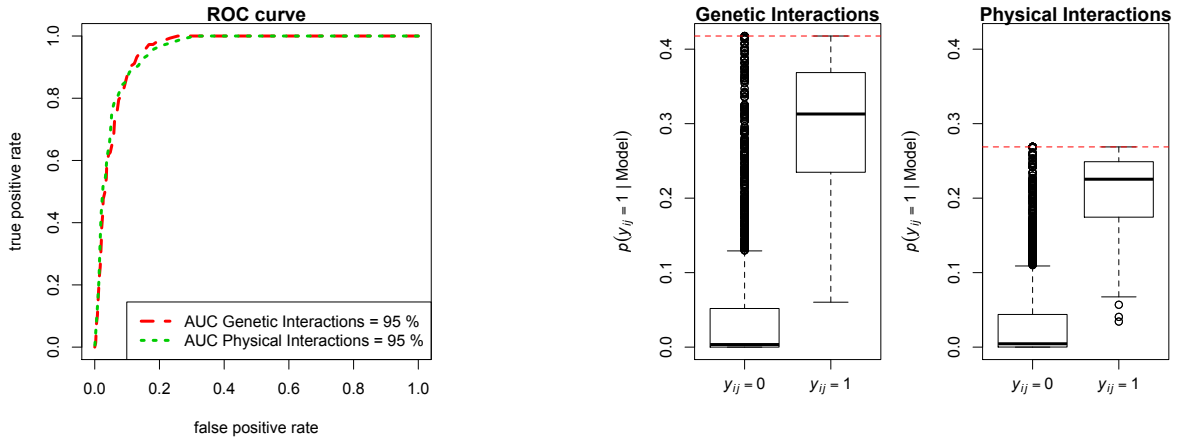


Figure 12: ROC curves and boxplots of the estimated probabilities of a link for the true negatives and true positives obtained by fitting the LSM to the PPI data.

5.3.2 LSJM

We apply the LSJM to the PPI dataset. Figure 13 shows the estimated overall latent positions $\bar{\mathbf{z}}_i$ (defined in Equation 5) and in the plot on the right each arrow starts from the positions $\tilde{\mathbf{z}}_{i1}$ for the genetic interaction dataset and points to the latent positions $\tilde{\mathbf{z}}_{i2}$

for the physical interaction dataset (defined in Equation 7). Figure 14 shows the estimated




Figure 13: On the left are shown the latent positions $\bar{\mathbf{z}}_i$ for the PPI networks fitting the LSJM. In the plot on the right the dots represent the overall positions $\bar{\mathbf{z}}_i$ and the arrows connect the estimated position under each model $\tilde{\mathbf{z}}_{i1}, \tilde{\mathbf{z}}_{i2}$.

position under each model $\tilde{\mathbf{z}}_{ik}$ in the latent space. It is possible to compare these results with Figure 11 in which we have fitted two single LSM to the data. The ROC curve, AUC




Figure 14: Estimated latent positions under each model $\tilde{\mathbf{z}}_{i1}, \tilde{\mathbf{z}}_{i2}$ for the PPI networks fitting the LSJM

and the boxplots in Figure 15 show that the LSJM fits the data quite well. The posterior distributions for the α 's are $q(\alpha_1) = \mathcal{N}(-0.410, 0.003)$ $q(\alpha_2) = \mathcal{N}(-0.940, 0.005)$.

We applied a 10-fold cross validation to evaluate the prediction of missing links. The area under the ROC curve is 94% for both the genetic and the physical interaction networks




Figure 15: ROC curves of the estimated probabilities of a link for the true negatives and true positives obtained by fitting the LSJM to the PPI Networks

fitting the LSJM, fitting the LSM the AUC values are 66% for the genetic interaction network and 97% for the physical interaction network. The results of the LSJM show a much better fit in terms of estimates of missing links for the genetic interaction network compared to the ones obtained by the single LSM approach. Setting the threshold to be equal to the median probability of a link for the subgroup of the actual observed links we obtained similar misclassification rates for the missing links in both LSJM and single LSM approach: 9% for the genetic interaction network, and 6% for the physical interaction network.

The estimated genetic interactions network is formed of 297 links, and its density is 0.067, while the estimated physical interactions network is formed of 185 links, and its density is 0.042. Their degree distributions are shown in Figure 16. Comparing these results with the true network statistics (Section 5.3) it's possible to observe that the simulated degree distributions (Figure 16) are smoother than the original one (Figure 10), but overall they are quite similar.

Applying a 10-fold cross validation for missing nodes using the LSJM we obtained a misclassification rate of 24% for the genetic interactions dataset and 20% for the physical interaction network.




Figure 16: Estimated Degree Distribution for the PPI networks fitting the LSJM.

6 Conclusions

A lot of network data require the introduction of novel models able to describe their complex connectivity structure. On the other hand new inferential methods are needed to carry out estimation efficiently. In this paper, we proposed a latent variable model (LSJM) for multiple network views that extends the latent space model proposed by Hoff et al. (2002), allowing the information given by different relations on the same nodes to be summarized in the same latent space. The use of the variational approach to compute the model parameters allows us to apply the latent space joint model to larger networks (of the order of thousands of nodes). A comparison between the variational method and MCMC for the single view network latent space model is given in Salter-Townshend and Murphy (2013). An alternative variational algorithm that can be used in this context could be derived from the methods outlined in Opper and Archambeau (2009). This model allows the position of each observation in a latent space to be found based on all the available information in the datasets. Further information like the latent positions in each network view separately are obtained updating the single-network estimates given the overall positions. This information allows effective visualization and prediction of the data. The examples presented show how the LSJM facilitates the interpretation of the different positions of the network nodes in the latent space according to longitudinal measurements in the excerpt of 50 girls from ‘Teenage Friends and Lifestyle Study’ example and multiple relations in the *Saccharomyces Cerevisiae*

genetic and physical protein-protein interactions dataset. All the methods presented in this paper are included in the `lvm4net` package for R (Gollini, 2014).

Future work may lead to an extension of the model allowing cluster formation by assuming that the latent positions come from a Gaussian mixture model fitting each network using a Latent Position Cluster Model (Handcock et al., 2007). Or a novel modelling approach which takes explicitly into account the sequential feature as in the dynamic network analysis (Sarkar and Moore, 2005; Hoff, 2011; Westveld and Hoff, 2011).

Acknowledgements

The authors wish to acknowledge the anonymous reviewers for helpful comments. This work was supported by Science Foundation funded Clique Research Cluster (08/SRC/I1407) and Insight Research Centre (SFI/12/RC/2289).

References

- Airoldi, E. M., Blei, D. M., Fienberg, S. E., and Xing, E. P. (2008), “Mixed-membership Stochastic Blockmodels,” *Journal of Machine Learning Research*, 9, 1981–2014.
- Ansari, A., Koenigsberg, O., and Stahl, F. (2011), “Modeling multiple relationships in social networks,” *Journal of Marketing Research*, 48, 713 –728.
- Attias, H. (1999), “Inferring Parameters and Structure of Latent Variable Models by Variational Bayes,” in *Proceedings of the Proceedings of the Fifteenth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, San Francisco, CA: Morgan Kaufmann, pp. 21–30.
- Bandyopadhyay, S., Kelley, R., Krogan, N. J., and Ideker, T. (2008), “Functional Maps of Protein Complexes from Quantitative Genetic Interaction Data,” *PLoS Computational Biology*, 4.

- Beal, M. J. (2003), “Variational algorithms for approximate Bayesian inference,” Ph.D. thesis, Gatsby Computational Neuroscience Unit, University College London.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977), “Maximum likelihood for incomplete data via the EM algorithm (with discussion),” *Journal of the Royal Statistical Society, Series B*, 39, 1–38.
- Goldenberg, A., Zheng, A. X., Fienberg, S. E., and Airoldi, E. M. (2010), “A Survey of Statistical Network Models,” *Foundations and Trends in Machine Learning*, 2, 129–233.
- Gollini, I. (2014), *lvm4net: Latent Variable Models for Networks*, R package Version 0.1.
- Handcock, M., Raftery, A., and Tantrum, J. (2007), “Model-Based Clustering for Social Networks (with discussion),” *Journal of the Royal Statistical Society: Series A*, 170, 1–22.
- Hoff, P. (2009), “Multiplicative latent factor models for description and prediction of social networks,” *Computational and Mathematical Organization Theory*, 15, 261–272.
- (2011), “Hierarchical multilinear models for multiway data,” *Computational Statistics & Data Analysis*, 55, 530–543.
- Hoff, P., Raftery, A., and Handcock, M. (2002), “Latent Space Approaches to Social Network Analysis,” *Journal of the American Statistical Association*, 97, 1090–1098.
- Holland, P. W., Laskey, K., and Leinhardt, S. (1983), “Stochastic Blockmodels: First Steps,” *Social Networks*, 5, 109 – 137.
- Holland, P. W. and Leinhardt, S. (1981), “An Exponential Family of Probability Distributions for Directed Graphs,” *Journal of the American Statistical Association*, 76, pp. 33–50.
- Jaakkola, T. S. and Jordan, M. I. (2000), “Bayesian Parameter Estimation Via Variational Methods,” *Statistics and Computing*, 10, 25–37.

- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999), “An Introduction to Variational Methods for Graphical Models,” *Machine Learning*, 37, 183–233.
- Koskinen, J. H., Robins, G. L., and Pattison, P. E. (2010), “Analysing exponential random graph (p-star) models with missing data using Bayesian data augmentation,” *Statistical Methodology*, 7, 366 – 384.
- Koskinen, J. H., Robins, G. L., Wang, P., and Pattison, P. E. (2013), “Bayesian analysis for partially observed network data, missing ties, attributes and actors,” *Social Networks*, 35, 514 – 527.
- Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., and Onnela, J.-P. (2010), “Community Structure in Time-Dependent, Multiscale, and Multiplex Networks,” *Science*, 328, 876–878.
- Opper, M. and Archambeau, C. (2009), “The Variational Gaussian Approximation Revisited,” *Neural Computation*, 21, 786–792.
- Pearson, M. and Michell, L. (2000), “Smoke Rings: social network analysis of friendship groups, smoking and drug-taking,” *Drugs: education, prevention and policy*, 7, 21–37.
- Pearson, M. and West, P. (2003), “Drifting smoke rings: social network analysis and Markov processes in a longitudinal study of friendship groups and risk-taking,” *Connections*, 25, 59–76.
- R Core Team (2014), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- Ripley, R. M., Snijders, T. A., and Preciado, P. (2011), *Manual for SIENA version 4.0*, Oxford, University of Oxford, Department of Statistics; Nuffield College.
- Salter-Townshend, M. and Murphy, T. B. (2013), “Variational Bayesian inference for the Latent Position Cluster Model for network data,” *Computational Statistics & Data Analysis*, 57, 661–671.

- Salter-Townshend, M., White, A., Gollini, I., and Murphy, T. B. (2012), “Review of statistical network analysis: models, algorithms, and software,” *Statistical Analysis and Data Mining*, 5, 243–264.
- Sarkar, P. and Moore, A. W. (2005), “Dynamic social network analysis using latent space models,” *SIGKDD Explorations*, 7, 31–40.
- Stark, C., Breitkreutz, B. J., Chatr-Aryamontri, A., Boucher, L., Oughtred, R., Livstone, M. S., Nixon, J., Van Auken, K., Wang, X., Shi, X., Reguly, T., Rust, J. M., Winter, A., Dolinski, K., and Tyers, M. (2011), “The BioGRID Interaction Database: 2011 update.” *Nucleic Acids Research*, 39, D698–D704.
- Stark, C., Breitkreutz, B. J., Reguly, T., Boucher, L., Breitkreutz, A., and Tyers, M. (2006), “BioGRID: a general repository for interaction datasets.” *Nucleic Acids Research*, 34, D535–D539.
- Tang, L., Wang, X., and Liu, H. (2011), “Community detection via heterogeneous interaction analysis,” *Data Mining and Knowledge Discovery*, 1–33.
- Troyanskaya, O. G., Dolinski, K., Owen, A. B., Altman, R. B., and Botstein, D. (2003), “A Bayesian framework for combining heterogeneous data sources for gene function prediction (in *Saccharomyces cerevisiae*).” *Proceedings of the National Academy of Sciences of the United States of America*, 100, 8348–8353.
- Westveld, A. and Hoff, P. (2011), “A Mixed Effects Model for Longitudinal Relational and Network Data, with Applications to International Trade and Conflict,” *Annals of Applied Statistics*, 5, 843–872.