Published in final edited form as:

J Am Stat Assoc. 2015 ; 110(511): 1037–1046. doi:10.1080/01621459.2014.965777.

Testing for nodal dependence in relational data matrices

Alexander Volfovsky¹ and Peter D. Hoff²

¹Department of Statistics, University of Washington

²Departments of Statistics and Biostatistics, University of Washington

Abstract

Relational data are often represented as a square matrix, the entries of which record the relationships between pairs of objects. Many statistical methods for the analysis of such data assume some degree of similarity or dependence between objects in terms of the way they relate to each other. However, formal tests for such dependence have not been developed. We provide a test for such dependence using the framework of the matrix normal model, a type of multivariate normal distribution parameterized in terms of row- and column-specific covariance matrices. We develop a likelihood ratio test (LRT) for row and column dependence based on the observation of a single relational data matrix. We obtain a reference distribution for the LRT statistic, thereby providing an exact test for the presence of row or column correlations in a square relational data matrix. Additionally, we provide extensions of the test to accommodate common features of such data, such as undefined diagonal entries, a non-zero mean, multiple observations, and deviations from normality. Supplementary materials for this article are available online.

Keywords

Networks; Matrix normal; hypothesis testing; Maximum likelihood

1 Introduction

Networks or relational data among m actors, nodes or objects are frequently presented in the form of an $m \times m$ matrix $Y = \{y_{ij} : 1 \leq i, j \leq m\}$, where the entry y_{ij} corresponds to a measure of the directed relationship from object i to object j . Such data are of interest in a variety of scientific disciplines: Sociologists and epidemiologists gather friendship network data to study social development and health outcomes among children (Fletcher et al., 2011; Potter et al., 2012), economists study markets by analyzing networks of business interactions among companies or countries (Westveld and Hoff, 2011a; Lazzarini et al., 2001), and biologists study gene-gene interaction networks to better understand biological pathways (Bergmann et al., 2003; Stuart et al., 2003).

Often of interest in the study of such data is a description of the variation and similarity among the objects in terms of their relations. Similarities among rows and among columns

Supplementary Materials

Proofs An appendix containing technical proofs to Theorems 2 and 3. (PDF)

in empirical networks have long been observed (Sampson, 1968; Leskovec et al., 2008), leading to the development of statistical tools to summarize such patterns. CONCOR (CONvergence of iterated CORrelations) is an early example of a procedure that partitions the rows (or columns) of Y into groups based on a summary of the correlations among the rows (or columns) of Y (White et al., 1976; McQuitty and Clark, 1968). The procedure yields a “blockmodel” of the objects, a representation of the original data matrix Y by a smaller matrix that identifies relationships among groups of objects. While this algorithm is still commonly used (Lincoln and Gerlach, 2004; Lafosse and Ten Berge, 2006), it suffers from a lack of statistical interpretability (Panning, 1982), as it is not tied to any particular statistical model or inferential goal.

Several model-based approaches presume the existence of a grouping of the objects such that objects within a group share a common distribution for their outgoing relationships. This is the notion of stochastic equivalence, and is the primary assumption of stochastic blockmodels, a class of models for which the probability of a relationship between two objects depends only on their individual group memberships (Holland et al., 1983; Wang and Wong, 1987; Nowicki and Snijders, 2001). Airoldi et al. (2008) extend the basic blockmodel by allowing each object to belong to several groups. In this model the probability of a relationship between two nodes depends on all the group memberships of each object. This and other variants of stochastic blockmodels belong to the larger class of latent variable models, in which the probability distribution of the relationship between any two objects i and j depends on unobserved object-specific latent characteristics z_i and z_j (Hoff et al., 2002). Statistical models of this type all presume some form of similarity among the objects in the network. However, while such models are widely used and studied, no formal test for similarities among the objects in terms of their relations has been proposed.

Many statistical methods for valued or continuous relational data are developed in the context of normal statistical models. These include, for example, the widely-used social relations model (Kenny and La Voie, 1984; Li and Loken, 2002) and covariance models for multivariate relational data (Li, 2006; Westveld and Hoff, 2011b; Hoff, 2011). Additionally, statistical models for binary and ordinal relational data can be based on latent normal random variables via probit or other link functions (Hoff, 2005, 2008). In this article we propose a novel approach to testing for similarities between objects in terms of the row and column correlation parameters of the matrix normal model. The matrix normal model consists of the multivariate normal distributions that have a Kronecker-structured covariance matrix (Dawid, 1981). Specifically, we say that an $m \times m$ random matrix Y has the mean-zero matrix normal distribution $N_{m \times m}(0, \Sigma_r, \Sigma_c)$ if $\text{vec}(Y) \sim N_{m^2}(0, \Sigma_c \otimes \Sigma_r)$ where “vec” is the vectorization operator and “ $\otimes$ ” denotes the Kronecker product. Under this distribution, the covariance between two relations y_{ij} and y_{kl} is given by $\text{cov}(y_{ij}, y_{kl}) = \Sigma_{r,ik} \Sigma_{c,jl}$. Furthermore, it is straightforward to show that

$$E[Y Y^t] = \sum_r \text{tr}(\Sigma_r) \text{ and } E[Y^t Y] = \sum_c \text{tr}(\Sigma_c).$$

These identities suggest the interpretation of Σ_r and Σ_c as the covariance of the objects as senders of ties and as receivers of ties, respectively. In this article, we evaluate evidence for similarities between objects by testing for non-zero correlations in this matrix normal model. Specifically, we develop a test of

$$H_0: (\Sigma_r, \Sigma_c) \in \mathcal{D}_+^m \times \mathcal{D}_+^m \text{ versus } H_1: (\Sigma_r, \Sigma_c) \in (\mathcal{S}_+^m \times \mathcal{S}_+^m) \setminus (\mathcal{D}_+^m \times \mathcal{D}_+^m)$$

where $\mathcal{D}_+^m$ is the set of $m \times m$ diagonal matrices with positive entries and $\mathcal{S}_+^m$ is the set of $m \times m$ positive definite symmetric matrices. Model H_0 , which we call the Kronecker variance model, represents heteroscedasticity among the rows and the columns while still maintaining their independence. Model H_1 , which we call the full Kronecker covariance model, allows for correlations between all of the rows and all the of columns. Rejection of the null would support further inference via a model that allowed for similarities among the objects, such as a stochastic blockmodel, some other latent variable model or the matrix normal model. Acceptance of the null would caution against fitting such a model in order to avoid spurious inferences.

This goal of evaluating the evidence for row or column correlation is in contrast to that of the existing testing literature for matrix normal distributions. This literature has focused on an evaluation of the null hypothesis that $\text{cov}(\text{vec}(Y)) = \Sigma_c \otimes \Sigma_r$ (our H_1) against an unstructured alternative (Roy and Khattree, 2005; Mitchell et al., 2006; Lu and Zimmerman, 2005; Srivastava et al., 2008). The tests proposed in this literature are likelihood ratio tests that, in the case of an $m \times m$ square matrix Y , require at least $n > m^2$ replications to estimate the covariance under the fully unstructured model. Such tests are not applicable to most relational datasets, which typically consist of at most a few observed relational matrices.

In the next section we derive a hypothesis test of H_0 versus H_1 in the context of the matrix normal model. We show that given a single observed relational matrix Y , the likelihood is bounded under both H_0 and H_1 and so a likelihood ratio test of H_0 against H_1 can be constructed. We further show how the null distribution of the test statistic can be approximated with an arbitrarily high precision via a Monte Carlo procedure. In Section 2.3 we extend these results to the general class of matrix variate elliptically contoured distributions. The power of the test in several different situations is evaluated in Section 3.

Although the development of our testing procedure is based on the mean-zero matrix normal distribution, it is straightforward to extend the test to several other scenarios commonly encountered in the study of relational data, including missing diagonal entries, non-zero mean structure, multiple heteroscedastic observations and binary networks. These extensions and two data examples are discussed in Section 4. A discussion follows in Section 5.

2 Likelihood ratio test

In this section we propose a likelihood ratio test (LRT) for evaluating the presence of correlations among the rows and correlations among the columns of a square matrix. The data matrix Y is modeled as a draw from a mean zero matrix normal distribution $N_{m \times m}(0,$

Σ_r, Σ_c). The parameter space under the null hypothesis H_0 is $\Theta_0 = \mathcal{D}_+^m \times \mathcal{D}_+^m$, the space of all pairs of diagonal $m \times m$ matrices with positive entries. Under the alternative H_1 , the parameter space is $\Theta_1 = (\mathcal{S}_+^m \times \mathcal{S}_+^m) \setminus (\mathcal{D}_+^m \times \mathcal{D}_+^m)$, the collection of all pairs of positive definite matrices of dimension m for which at least one is not diagonal. To derive the LRT statistic, we first obtain the maximum likelihood estimates (MLEs) under the unrestricted parameter space $\Theta = \Theta_0 \cup \Theta_1$ and under the null parameter space Θ_0 . From these MLEs, we construct several equivalent forms of the LRT statistic. While the null distribution of the test statistic is not available in closed form, the statistic is invariant under diagonal rescalings of the data matrix Y , implying that the distribution of the statistic is constant as a function of $(\Sigma_r, \Sigma_c) \in \Theta_0$. This fact allows us to obtain null distributions and p -values via Monte Carlo simulation.

2.1 Maximum likelihood estimates

The density of a mean zero matrix normal distribution $N_{m \times m}(0, \Sigma_r, \Sigma_c)$ is given by

$$p(Y|\Sigma_r, \Sigma_c) = (2\pi)^{-m^2/2} |\Sigma_c \otimes \Sigma_r|^{-1/2} \exp(-\frac{1}{2} \text{tr}(\Sigma_r^{-1} Y \Sigma_c^{-1} Y^t)),$$

where “tr” is the matrix trace and “ $\otimes$ ” is the Kronecker product. Throughout the article we will write $\ell(\Sigma_r, \Sigma_c; Y)$ as minus two times the log likelihood minus $m^2 \log 2\pi$, hereafter referred to as the scaled log likelihood:

$$\ell(\Sigma_r, \Sigma_c; Y) = -2 \log p(Y|\Sigma_r, \Sigma_c) - m^2 \log 2\pi = \text{tr}[\Sigma_r^{-1} Y \Sigma_c^{-1} Y^t] - \log |\Sigma_c^{-1} \otimes \Sigma_r^{-1}|. \quad (1)$$

We will state all the results in this paper in terms of $\ell(\Sigma_r, \Sigma_c; Y)$. For example, an MLE will be a minimizer of $\ell(\Sigma_r, \Sigma_c; Y)$ in (Σ_r, Σ_c) . The following result implies that if Y is a draw from an absolutely continuous distribution on $\mathbb{R}^{m \times m}$, then the scaled likelihood is bounded from below, and achieves this bound on a set of nonunique MLEs:

Theorem 1: *If Y is full rank then $\ell(\Sigma_r, \Sigma_c; Y) \geq m^2 + m \log |YY^t/m|$ for all $(\Sigma_r, \Sigma_c) \in \Theta$, with equality if $\Sigma_r = Y \Sigma_c^{-1} Y^t / m$, or equivalently, $\Sigma_c = Y^t \Sigma_r^{-1} Y / m$.*

Proof: We first look for MLEs at the critical points of $\ell(\Sigma_r, \Sigma_c; Y)$. Setting derivatives of ℓ to zero indicates that critical points satisfy

$$\hat{\Sigma}_r = Y \hat{\Sigma}_c^{-1} Y^t / m \quad (2)$$

$$\hat{\Sigma}_c = Y^t \hat{\Sigma}_r^{-1} Y / m. \quad (3)$$

Note that these equations are redundant: If $(\hat{\Sigma}_r, \hat{\Sigma}_c)$ satisfy Equation 2, then these values satisfy Equation 3 as well. The value of the scaled log likelihood at such a critical point is

$$\begin{aligned}
l(Y\hat{\Sigma}_c^{-1}Y^t/m, \hat{\Sigma}_c; Y) &= \text{tr} \left[(Y\hat{\Sigma}_c^{-1}Y^t/m)^{-1} Y\hat{\Sigma}_c^{-1}Y^t \right] + \log |\hat{\Sigma}_c \otimes Y\hat{\Sigma}_c^{-1}Y^t/m| \\
&= m \text{tr} \left[Y^{-t} \hat{\Sigma}_c Y^{-1} Y\hat{\Sigma}_c^{-1}Y^t \right] + m \log |\hat{\Sigma}_c| - m \log |Y\hat{\Sigma}_c^{-1}Y^t/m| \\
&= m \text{tr} [I] + m \log |Y\hat{\Sigma}_c^{-1}Y^t/m|.
\end{aligned} \tag{4}$$

In the second and third lines, Y^{-1} exists and $|Y\hat{\Sigma}_c^{-1}Y^t/m| > 0$ since Y is square and full rank.

Now we compare the scaled log likelihood at a critical point to its value at any other point.

$$\begin{aligned}
l(\Sigma_r, \Sigma_c; Y) - l(Y\hat{\Sigma}_c^{-1}Y^t/m, \hat{\Sigma}_c; Y) &= \text{tr} [\Sigma_r^{-1}Y\hat{\Sigma}_c^{-1}Y^t] + m \log (|\Sigma_r| |\hat{\Sigma}_c|) - m^2 - m \log |Y\hat{\Sigma}_c^{-1}Y^t/m| \\
&= m^2 \left[\frac{1}{m} \text{tr} [\Sigma_r^{-1}Y\hat{\Sigma}_c^{-1}Y^t/m] - \frac{1}{m} \log |\Sigma_r^{-1}Y\hat{\Sigma}_c^{-1}Y^t/m| - 1 \right]
\end{aligned} \tag{5}$$

The first equality is a simple combination of equations 1 and 4. The second equality is a rearrangement of terms that combines all the determinant in the log terms. This difference can be written as $m^2(a - \log g - 1)$, where a is the arithmetic mean and g is the geometric

mean of the eigenvalues of $(\Sigma_r^{-1/2}Y\hat{\Sigma}_c^{-1}Y^t\Sigma_r^{-1/2})/m$. To complete the proof we show that $a - \log g - 1 \geq 0$. Consider $f(x) = x - 1 - \log x$ and its first and second derivatives with respect to x : $f'(x) = 1 - \frac{1}{x}$, and $f''(x) = \frac{1}{x^2}$. The second derivative is positive at the critical point $x = 1$, so $f(1) = 0$ is a global minimum of the function. Thus $x - 1 - \log x \geq 0$. Now let $\lambda_1,$

$\dots, \lambda_m$ be the eigenvalues of $(\Sigma_r^{-1/2}Y\hat{\Sigma}_c^{-1}Y^t\Sigma_r^{-1/2})/m$ and so $a = \frac{1}{m} \sum \lambda_i$ and $g = (\prod \lambda_i)^{1/m}$. We then have

$$a \geq \log a + 1 = \log \left(\frac{1}{m} \sum x_i \right) + 1 \geq \log \left(\left(\prod x_i \right)^{1/m} \right) + 1 = \log g + 1$$

as $a \geq g$ since $\lambda_i \geq 0 \forall i$. Since $a - 1 - \log g \geq 0$ we have the desired result.

Note that the MLE is not unique, nor is the MLE of $\Sigma_c \otimes \Sigma_r$. For example, $I \otimes Y\hat{\Sigma}_c^{-1}Y^t/m$ is an MLE of $\Sigma_c \otimes \Sigma_r$, as is $Y^t Y \otimes I/m$. Moreover, there is an MLE for each $\Sigma_r \in \mathcal{S}_+^m$ given by $(\Sigma_r, Y^t \Sigma_r^{-1} Y/m)$, and similarly there is an MLE for each $\Sigma_c \in \mathcal{S}_+^m$ given by $(Y \Sigma_c^{-1} Y^t/m, \Sigma_c)$.

Theorem 1 also implies that the likelihood is bounded under the null. Unlike the unrestricted case, the MLE under the null is unique up to scalar multiplication:

Theorem 2: *If Y is full rank then the MLE $\hat{D}_c \otimes \hat{D}_r$ under H_0 is unique, while $\hat{D}_r$ and $\hat{D}_c$ are unique up to a multiplication and division by the same positive scalar.*

A proof is given in the Supplement. To find the MLE under the null model, we obtain the derivatives of the scaled log likelihood l with respect to $(\Sigma_r, \Sigma_c) \in \Theta_0$. For notational convenience, we will refer to diagonal versions of Σ_r and Σ_c as D_r and D_c respectively. Setting these derivatives equal to zero, we establish that the critical points of l must satisfy

$$D_r = Y D_c^{-1} Y^t \circ I / m \quad (6)$$

$$D_c = Y^t D_r^{-1} Y \circ I / m \quad (7)$$

where “ $\circ$ ” is the Hadamard product. The MLE can be found by iteratively solving equations (6) and (7). This procedure can be seen as a type of block coordinate descent algorithm, decreasing l at each iteration (Tseng, 2001).

2.2 Likelihood ratio test statistic and null distribution

Since the scaled log likelihood is bounded below, we are able to obtain a likelihood ratio statistic that is finite with probability 1 when Y is sampled from an absolutely continuous distribution on $\mathbb{R}^{m \times m}$. As usual, a likelihood ratio test statistic can be obtained from the ratio of the unrestricted maximized likelihood to the likelihood maximized under the null. We take our test statistic to be $T(Y) = \ell(\hat{D}_r, \hat{D}_c; Y) - \ell(\hat{\Sigma}_r, \hat{\Sigma}_c; Y)$, where $(\hat{\Sigma}_r, \hat{\Sigma}_c)$ is any unrestricted MLE and $(\hat{D}_r, \hat{D}_c)$ is the MLE under Θ_0 . Since the scaled log likelihood l is minus two times the likelihood, our statistic is a monotonically increasing function of the likelihood ratio.

In Theorem 1 we showed that $\ell(\hat{\Sigma}_r, \hat{\Sigma}_c; Y) = m^2 + m \log |YY^t/m|$ for any unrestricted MLE $(\hat{\Sigma}_r, \hat{\Sigma}_c)$. Similarly, letting $(\hat{D}_r, \hat{D}_c)$ be an MLE under H_0 , we have

$$\begin{aligned} l(\hat{D}_r, \hat{D}_c; Y) &= \text{tr} \left[Y^t \hat{D}_r^{-1} Y \hat{D}_c^{-1} \right] + \log |\hat{D}_c \otimes \hat{D}_r| \\ &= \text{tr} \left[(Y^t \hat{D}_r^{-1} Y) (Y^t \hat{D}_r^{-1} Y / m \circ I)^{-1} \right] + \log |\hat{D}_c \otimes \hat{D}_r| \\ &= \sum_i (Y^t \hat{D}_r^{-1} Y)_{ii} (Y^t \hat{D}_r^{-1} Y / m \circ I)_{ii}^{-1} + \log |\hat{D}_c \otimes \hat{D}_r| \\ &= m \sum_i (Y^t \hat{D}_r^{-1} Y)_{ii} / (Y^t \hat{D}_r^{-1} Y)_{ii} + \log |\hat{D}_c \otimes \hat{D}_r| = m^2 + \log |\hat{D}_c \otimes \hat{D}_r|, \end{aligned}$$

where the second equality stems from MLE satisfying $\hat{D}_c = Y^t \hat{D}_r^{-1} Y / m \circ I$ (Equation (7)). The third equality relies on the following identity for traces: for a diagonal matrix A and unstructured matrix B of the same dimension, $\text{tr} [AB] = \sum_i A_{ii} B_{ii}$. The final line is due to the following identity for Hadamard products: if I is the identity matrix and B is an unstructured matrix, then $(B \circ I)_{ii}^{-1} = 1/B_{ii}$.

The maximized likelihoods under the null and alternative give

$$T(Y) = \log |\hat{D}_c \otimes \hat{D}_r| - m \log |YY^t/m| = m \left(\log |\hat{D}_c| + \log |\hat{D}_r| - \log |YY^t/m| \right). \quad (8)$$

Since no closed form solution exists for $\hat{D}_r$ it is not clear how to obtain the null distributions of T in closed form. However, it is possible to simulate from the null distribution of $T(Y)$, as the distribution of the test statistic is the same for all elements of the null hypothesis. To see this, we show that the test statistic itself is invariant under left and right transformations of

the data by positive diagonal matrices. Let $\tilde{Y} = D_1 Y D_2$ for positive diagonal matrices D_1 and D_2 . Since the MLE $\hat{D}_c \otimes \hat{D}_r$ is unique it is an equivariant function of any matrix Y with respect to left and right multiplication by diagonal matrices (see Eaton (1983) Prop 7.11). In particular, writing $\hat{\theta}(\tilde{Y})$ for the estimate of $\hat{D}_c \otimes \hat{D}_r$ based on a data matrix Y , we have $\hat{\theta}(\tilde{Y}) = (D_2^{1/2} \otimes D_1^{1/2}) \hat{\theta}(Y) (D_2^{1/2} \otimes D_1^{1/2})$. Since the determinant is a multiplicative map, we can write the determinant of the above as $|\hat{\theta}(\tilde{Y})| = |(D_2 \otimes D_1)| |\hat{\theta}(Y)|$. Using the above, the $T(\tilde{Y})$ can be written as in Equation (8):

$$\begin{aligned} T(\tilde{Y}) &= m \log |\hat{\theta}(\tilde{Y})| - m \log |\tilde{Y} \tilde{Y}^t / m| \\ &= m \log |\hat{\theta}(Y)| + m \log |D_2 \otimes D_1| - m \log |Y Y^t / m| - m \log |D_2 \otimes D_1| = T(Y). \end{aligned}$$

Since for a matrix normal random variable $Y \sim N_{m \times m}(0, \Sigma_1, \Sigma_2)$ we have

$Y \stackrel{d}{=} \sum_1^{1/2} Y_0 \sum_2^{1/2}$ for $Y_0 \sim N_{m \times m}(0, I, I)$, the above argument implies that $T(Y) \stackrel{d}{=} T(Y_0)$ under the null. Therefore, the null distribution of T can be approximated via Monte Carlo simulation of Y from any distribution in H_0 . For example, a Monte Carlo approximation to the q^{th} quantile, T_q can be obtained from the following algorithm:

1. Simulate $Y_0^1, \dots, Y_0^S \sim \text{i. i. d. } N_{m \times m}(0, I, I)$;
2. Let $\hat{T}_q = \min\{T(Y_0^q) : \sum_{s=1}^S 1[T(Y_0^q) \geq T(Y_0^s)] / S \geq q\}$.

2.3 Matrix variate elliptically contoured distributions

The results of the previous subsection are immediately extendable to the general class of matrix variate elliptically contoured distributions. In this section we show that under minor regularity conditions on the distributions, the likelihood for a matrix variate elliptically contoured distribution is bounded when the matrix normal distribution is bounded. We provide the form of an MLE for the general class of mean zero square matrix variate elliptically contoured distributions and demonstrate that the likelihood ratio test between H_0 and H_1 has the same form as in Equation (8).

We use the notation of Gupta and Varga (1994) for the matrix variate elliptically contoured distribution. We say that Y has a mean zero square matrix variate elliptically countoured distribution and write $Y \sim E_{m \times m}(0, \Sigma_r, \Sigma_c, h)$ if its density has the form

$$f_Y(Y) = \frac{1}{|\Sigma_r|^{\frac{m}{2}} |\Sigma_c|^{\frac{m}{2}}} h(\text{tr}[Y^t \Sigma_r^{-1} Y \Sigma_c^{-1}]) = \frac{1}{|\Sigma_r|^{\frac{m}{2}} |\Sigma_c|^{\frac{m}{2}}} h(\text{vec}(Y)^t (\Sigma_c^{-1} \otimes \Sigma_r^{-1}) \text{vec}(Y)). \quad (9)$$

For $h(w) = (2\pi)^{m^2/2} \exp(-\frac{1}{2}w)$, Y is a mean zero matrix variate normal variable. Gupta and Varga (1994) show that if $Y \sim E(0, \Sigma_r, \Sigma_c, h)$ then $AYB \sim E(0, A\Sigma_r A^t, B\Sigma_c B^t, h)$ and if second moments exist, then $\text{cov}(\text{vec}(Y)) = c_h(\Sigma_c \otimes \Sigma_r)$. In Gupta and Varga (1995), the authors showed that when Σ_r and h are known, the MLE of Σ_c is proportional to the MLE of Σ_c under normality, but they do not provide results for the boundedness of the likelihood or

existence of MLEs for the case where only h is known. To find the form of the MLE in this case we state a simplified version of Theorem 1 of Anderson et al. (1986):

Theorem: Let Ω be a set in the space of $\mathcal{S}_+^{m^2}$, such that if $V \in \Omega$ then $cV \in \Omega \forall c > 0$ (that is Ω is a cone). Suppose h is such that $h(y^t y)$ is a density in $\mathbb{R}^{m^2}$ and $x^{m^2/2} h(x)$ has a finite positive maximum x_h . Suppose that on the basis of an observation y from $|V|^{-1/2} h(y^t V^{-1} y)$ an MLE under normality $\tilde{V} \in \Omega$ exists and $\tilde{V} > 0$ with probability 1. Then an MLE for h is $\hat{V} = (m^2/x_h) \tilde{V}$ and the maximum of the likelihood is $|V|^{-1/2} h(x_h)$.

In the previous subsection we proved that for the mean zero square matrix normal distribution, the likelihood is bounded for a single observation. A direct application of the above theorem with $\Omega = \mathcal{S}_+^m \times \mathcal{S}_+^m \subset \mathcal{S}_+^{m^2}$ and $y = \text{vec}(Y)$ proves that for the likelihood of a generic matrix variate elliptically contoured distribution with h defined as in 9, the likelihood is bounded and the MLE of $\text{cov}(\text{vec}(Y))$ is proportional to the MLE under normality. Clearly the theorem holds for a smaller space $\omega = \mathcal{D}_+^m \times \mathcal{D}_+^m \subset \mathcal{S}_+^{m^2}$ as well, and thus the likelihood ratio statistic can be constructed as follows

$$\begin{aligned} T(Y) &= 2\log \left[|\hat{V}_\Omega|^{-1/2} h(x_h) \right] - 2\log \left[|\hat{V}_\omega|^{-1/2} h(x_h) \right] \\ &= \log |\hat{V}_\omega| - \log |\hat{V}_\Omega| = \log |\tilde{V}_\omega| - \log |\tilde{V}_\Omega| \end{aligned} \quad (10)$$

where $\tilde{V}_\omega$ and $\tilde{V}_\Omega$ are the MLEs under normality that were previously derived. Equation (10) is identical to the original form of the test (e.g. Equation (8)). As such, to conduct the test for any elliptically contoured distribution, we can construct a reference distribution for the null based on a mean zero matrix variate distribution.

3 Power calculations

In this section we present power calculations for three different types of covariance models. The three covariance models we consider are: (1) Exchangeable row covariance and exchangeable column covariance; (2) Maximally sparse Kronecker structured covariance; and (3) the covariance induced by a nonseparable stochastic blockmodel with two row groups and two column groups. For each covariance model, we consider the power as a function of parameters that control the total correlation within a covariance matrix as well as in terms of m , the dimension of the matrix. In Table 1 we present the 95% quantiles based on the null distributions required for performing level $\alpha = 0.05$ tests.

3.1 Exchangeable row and column covariance structure

We first consider a submodel of the matrix normal model in which Σ_r and Σ_c have exchangeable covariance structure. In this structure, the correlation between any two rows is a constant ρ_r and the correlation between any two columns is a constant ρ_c . Specifically, $\text{cov}(\text{vec}(Y)) = \Sigma_c \otimes \Sigma_r$ where $\Sigma_r = (1 - \rho_r) I + \rho_r \mathbf{1}\mathbf{1}^t$ and $\Sigma_c = (1 - \rho_c) I + \rho_c \mathbf{1}\mathbf{1}^t$, and $\mathbf{1}$ is a vector of ones of length m . We first consider a network with $m = 10$ nodes and present the power as a function of ρ_r and ρ_c ranging from $-1/9$ to 1, where the lower bound guarantees

that the covariance matrices are positive definite. We calculate the power on a 25×25 grid in $[-1/9, 1]^2$ and use a bivariate interpolation to construct the heatmap in the top left panel of Figure 1. From the plot it is evident that the power is an increasing function in $|\rho_r|$ and $|\rho_c|$. In particular, keeping ρ_r constant, the power is an increasing function of $|\rho_c|$ and vice versa.

In the top left panel of Figure 1 we observed that while keeping ρ_c constant, the power is an increasing function of $|\rho_r|$. To study the power for higher dimensional matrices, we set $\rho_c = 0$ and vary m and ρ_r . The dashed line in the top left panel of Figure 1 traces the power function for $m = 10$ and $\rho_c = 0$. This corresponds to the same style dashed line in the top right hand panel of Figure 1. The other four lines in the right hand panel represent the calculated power as a function of ρ_r for different dimensions m , holding $\rho_c = 0$. As is expected, for each m , the power is an increasing function of $|\rho_r|$. Similarly, for each fixed ρ_r value, the power is an increasing function of m . This latter phenomenon is due to the increase in the amount of data information with the increase in the dimension of the sociomatrix.

3.2 Maximally sparse Kronecker covariance structured correlation

While the previous example demonstrates the power of the test in the presence of many nonzero off-diagonal entries in the correlation matrices, it is of interest to see if the test has any power against alternatives that do not exhibit a large amount of correlation. For this purpose we consider a maximally sparse Kronecker covariance structure. We set the columns to be independent and only the first two rows to be correlated. This can be written compactly as $\Sigma_c = I$ and $\Sigma_r = I + \rho E_{12} + \rho E_{21}$, where E_{ij} is the 0 matrix with a 1 in the $(i, j)^{\text{th}}$ entry. For each of the matrix sizes $m \in \{5, 10, 25, 50, 100\}$ we computed power functions for values of ρ ranging between -1 and 1 . Monte Carlo approximations to the corresponding power functions are presented in the bottom left panel of Figure 1. We plot the results for $m = 5$ and $m = 100$ and see that the power increases monotonically as a function of $|\rho|$ for both dimensions. While the power of the test for a fixed ρ appears to decrease as the size of the network m increases, the two curves are nearly identical. We explain this as follows: while one expects that as the dimension m increases there is an increase in data information for identifying the correlation ρ , the power curve is influenced more heavily by the fact that the difference between Σ_r and the identity matrix becomes less pronounced. Additional power curves for a range of m values between 5 and 100 were approximated. All the curves were between the $m = 5$ and $m = 100$ curves that are presented in the plot. For all the power calculations, the lowest calculated power for values of ρ close to 0 was always within two Monte Carlo standard errors of 0.05.

3.3 Misspecified covariance structure

In the Introduction we discussed a popular model for relational data with an underlying assumption of stochastically equivalent nodes called the stochastic blockmodel. Straightforward calculations show that in general the covariance induced by a stochastic blockmodel is nonseparable, but still induces correlations among the rows and among the columns. As such, we are interested in evaluating the power of our test against such nonseparable alternatives.

A stochastic blockmodel can be represented in terms of multiplicative latent variables. Specifically, we can write the relationship $y_{ij} = u_i^t W v_j + \varepsilon_{ij}$ where u_i and v_j are latent vectors representing the row group membership of node i and the column group membership of node j . W is a matrix of means for the different group memberships and ε_{ij} is iid random noise. For the purposes of this power calculation we consider a simple setup where each node belongs to one of two row groups and one of two column groups with equal

probability. We let $W = \begin{pmatrix} 0 & -\mu \\ \mu & 0 \end{pmatrix}$ depend on a single parameter $\mu > 0$. Under this choice of W , we have $E[Y] = 0$ and $E[\mathbf{1}^t Y \mathbf{1}] = 0$. Since there are only two groups, the latent group membership vectors can be written as $u_i = (u_{i1}, 1 - u_{i1})$ and $v_j = (v_{j1}, 1 - v_{j1})$ where u_{i1} and v_{j1} are independent Bernoulli(1/2) random variables.

The bottom right panel of Figure 1 presents the power calculations for dimensions $m \in \{5, 10, 25, 50, 100\}$ and $|\mu| \in [0, 4]$. The power of the level $\alpha = 0.05$ test is increasing in $|\mu|$ which is a desirable property for this blockmodel since as $|\mu|$ grows the difference in the means for the groups becomes greater. The power of the test also increases with the dimension m .

4 Extensions and Applications

In this section we develop several extensions of the proposed test, and illustrate their use in the context of two data analysis examples. In the first example, we show how the test can be extended to accommodate a missing diagonal, an unknown non-zero mean, and heteroscedastic replications. The second example illustrates the use of the test for binary network data, a common type of relational data.

4.1 Extensions and continuous data example

International trade data on the value of exports from country to country is collected by the UN on a yearly basis and disseminated through the UN Comtrade website: <http://comtrade.un.org>. In this section we consider measures of total exports between twenty-six mostly large, well developed countries with high gross domestic product collected from 1996 to 2009 (measured in 2009 dollars). Specifically, we are interested in evaluating evidence for correlations among exporters and among importers. As trade between countries is relatively stable across years, we analyze the yearly change in log trade values, resulting in thirteen measurements (for the fourteen years of data) for every country-country pair. The data takes the form of a three way array $Y = \{Y_{ijk} : i, j \in \{1, \dots, 26\}, k \in \{1, \dots, 13\}\}$ where i and j index the exporting and importing countries respectively and k indexes the year. Exports from a country to itself are not defined, and so entries Y_{iik} are “missing”.

We consider a model for trade of the form, $Y_{ijk} = \beta_1 x_{ik} + \beta_2 x_{jk} + \varepsilon_{ijk}$, where x_{ik} is the difference in log gross domestic product of country i between years k and $k - 1$ (theoretical development of this model is available in the economics literature, see Tinbergen et al. (1962), and Bergstrand (1985)). We use GDP data collected by the World Bank through <http://data.worldbank.org/> to obtain OLS estimates of β_1 and β_2 . To investigate the correlations among importers and among exporters we collect the residuals $e_{ijk} = Y_{ijk} - \hat{Y}_{ijk}$

into thirteen matrices, $E_{..k}$ for $k = 1, \dots, 13$. Figure 2 plots the first two eigenvectors of moment estimates of pairwise row and column correlation matrices based on $E_{..1}, \dots, E_{..(13)}$. We observe systematic geographic patterns in both panels of the figure suggesting evidence that the ε_{ijk} are not independent. To evaluate this evidence formally by testing for dependence of the ε_{ijk} we extend the conditions under which the test developed in this article is applicable. Specifically, we must accommodate the following features of this data: a missing diagonal (Y_{iik} is not defined for all i and k), multiple observations (13 data points), and a nonzero mean structure.

Missing diagonal—In relational datasets, the relationship of an actor to himself is typically undefined, meaning that the relational matrix Y has an undefined diagonal. It is common to treat the entries of an undefined diagonal as missing at random and to use a data augmentation procedure to recover a complete data matrix, applying the analysis to the complete data. In the context of this article, this approach would allow us to treat the whole data matrix as a draw from a matrix normal distribution and perform our test exactly as outlined in Section 2. In this section we describe an augmentation procedure that does not require distributional assumptions for the diagonal elements. The procedure produces a data matrix $\tilde{Y}$ that we use to calculate the test statistic $T(\tilde{Y})$. Specifically, $\tilde{Y}$ replaces the undefined diagonal of Y with zeros, keeping the rest of the data matrix the same. We show that $T(\tilde{Y})$ is invariant under diagonal transformations and thus we can approximate the null distribution of the test statistic based on for data drawn from a matrix normal distribution where the diagonal entries are replaced with zeros.

Consider square matrices Y and $\tilde{Y}$ where $Y \sim N_{m \times m}(0, D_r, D_c)$ while $\tilde{Y}$ is distributed identically to Y except the diagonal entries are replaced with zeros. Similarly define square matrices $Y_0 \sim N_{m \times m}(0, I, I)$ and $\tilde{Y}_0$. We showed in Section 2 that the distribution of the likelihood ratio test statistic is invariant under transformations by diagonal matrices on the left and right, that is $T(Y) \stackrel{d}{=} T(Y_0)$. We now show that $T(\tilde{Y}) \stackrel{d}{=} T(\tilde{Y}_0)$. It is immediate that since $Y \stackrel{d}{=} D_r^{1/2} Y_0 D_c^{1/2}$ we have $\tilde{Y} \stackrel{d}{=} D_r^{1/2} \tilde{Y}_0 D_c^{1/2}$, as zeros on the diagonal are preserved by left and right diagonal transformations and the off diagonal entries (i, j) are normally distributed with variance $D_{r,i} D_{c,j}$. As in Section 2.2, we appeal to the equivariance of a unique MLE to show that $T(\tilde{Y}) = T(\tilde{Y}_0)$. The argument is identical to the one appearing in the paragraph following Equation (8) on page 9 and so we do not reproduce it here. Since $T(\tilde{Y}) \stackrel{d}{=} T(\tilde{Y}_0)$, we can approximate the null distribution and calculate the relevant quantiles for the test statistic with a simple update to the algorithm at the end of Section 2.2:

1. Simulate $\tilde{Y}_0^1, \dots, \tilde{Y}_0^S \stackrel{\text{iid}}{\sim} \mathcal{L}(\tilde{Y}_0)$, where $\mathcal{L}(\tilde{Y}_0)$ denotes the distribution of $\tilde{Y}_0$;
2. Let $\hat{T}_q = \min\{T(\tilde{Y}_0^q) : \sum_{s=1}^S 1[T(\tilde{Y}_0^s) \geq T(\tilde{Y}_0^q)]/S \geq q\}$.

Repeated observations—The test we discussed in this article is designed for a single observation. However, the test conveniently generalizes to the situation in which multiple observations are available. We will consider two types of additional observations: independent homoscedastic observations and independent heteroscedastic observations.

First, if there are p independent identically distributed observations, we note that likelihood equations of Section 2 can be rewritten as

$$\begin{aligned} mp\hat{D}_r &= \sum Y_i \hat{D}_c^{-1} Y_i^t \circ I & mp\hat{D}_c &= \sum Y_i^t \hat{D}_r^{-1} Y_i \circ I \\ mp\hat{\Sigma}_r &= \sum Y_i \hat{\Sigma}_c^{-1} Y_i^t & mp\hat{\Sigma}_c &= \sum Y_i^t \hat{\Sigma}_r^{-1} Y_i. \end{aligned}$$

The likelihood remains bounded and the form of the test statistic is identical to Equation (8). When the observations are heteroscedastic the likelihood equations (included in the proof of Theorem 3 in the Supplement) are more complicated because of the need to estimate the variability along the replications (we refer the reader to Hoff (2011) for an exposition on the general class of array normal distributions and estimation procedures).

Theorem 3: Let $Y_1, \dots, Y_p$ be independent random matrices distributed as $Y_i \sim N_{m \times m}(0, d_i \Sigma_r, \Sigma_c)$. Then $\forall p \geq 1$ the likelihood is bounded as a function of the covariance matrices Σ_r and Σ_c and the variance parameters $d_1, \dots, d_p$.

A proof is in the Supplement. Theorem 3 extends the literature on maximum likelihood estimation for proportional covariance models from natural exponential families to the matrix normal family which is a curved exponential family (Eriksen, 1987; Flury, 1986; Jensen and Johansen, 1987; Jensen and Madsen, 2004). Due to Theorem 3, we can modify the test statistic to test $H_0: D_{\text{obs}} \in \mathcal{D}_+^p, D_c, D_r \in \mathcal{D}_+^m$ vs $H_1: D_{\text{obs}} \in \mathcal{D}_+^p, \Sigma_c, \Sigma_r \in \mathcal{S}_+^m$:

$$\begin{aligned} T(Y_1, \dots, Y_p) &= l(\hat{D}_{\text{obs}}^{\text{null}}, \hat{D}_r, \hat{D}_c; Y) - l(\hat{D}_{\text{obs}}^{\text{alt}}, \hat{\Sigma}_r, \hat{\Sigma}_c; Y) \\ &= \log |\hat{D}_{\text{obs}}^{\text{null}} \otimes \hat{D}_c \otimes \hat{D}_r| - \log |\hat{D}_{\text{obs}}^{\text{alt}} \otimes \hat{\Sigma}_c \otimes \hat{\Sigma}_r|. \end{aligned}$$

We can again approximate the null distribution of the test statistic due to the invariance of the test statistic $T(Y_1, \dots, Y_p)$ under diagonal transformations of the data along all three modes. The results on missing diagonal elements (above) and a non-zero mean (below) are also immediately applicable to $T(Y_1, \dots, Y_p)$ above.

Relaxing the mean zero assumption—The reference distributions for the test statistics developed in Section 2 are based on the assumption that $E[Y] = 0$, a strong assumption that is unlikely to be true for any observed dataset. While treating the mean as a nuisance parameter is tempting, the likelihood function under the alternative model is unbounded when estimating a mean matrix and two covariance matrices simultaneously. We propose to first fit a regression based mean to the data assuming the entries in the data matrix are independently distributed with the same variance parameter and to then perform the test based on the demeaned data. We consider the following regression framework for the mean: $y_{ij} = \beta^T x_{ij} + \varepsilon_{ij}$. The regressor x_{ij} is a p -dimensional vector that can include features of node i , features of node j and dyadic features for nodes i and j . The ε_{ij} are assumed to be independent and identically distributed errors. Writing this in vector notation as $\text{vec}(Y) = X\beta + \text{vec}(\varepsilon)$ where $X = (x_{12}^t \cdots x_{(m-1)m}^t)^t$ and ε is an $m \times m$ matrix, the OLS estimate of β

is $\hat{\beta} = (X'X)^{-1} X'\text{vec}(Y)$. Under mild regularity conditions on the distribution of the explanatory variables X and the row and column variances for new nodes, the OLS estimate $\hat{\beta}$ is a consistent estimate of β . This motivates us to base the test statistic on $\text{vec}(\hat{\varepsilon}) = \text{vec}(Y - X\hat{\beta})$, the residuals of the regression, as we expect the distribution of $\text{vec}(\hat{\varepsilon})$ to be close to the distribution of $\text{vec}(\varepsilon)$ for large m . The null distribution of the test statistic T based on the residuals from the regression is not identical to the one derived in Section 2 and no explicit computation of the new null distribution is readily available. However, we have observed via simulation that the level of the test based on the estimated residuals $\hat{\varepsilon}$ appears to be asymptotically correct and that the test statistic based on ε and $\hat{\varepsilon}$ appear to have the same limiting distributions.

Application to international trade data—Figure 2 suggested that there is evidence of residual dependence among exporters and among importers based on additional information about the data (the relative geographic positions of the countries). Above we developed the tools to test for independence of the ε_{ijk} in the mean model using the likelihood ratio test proposed in Section 2. Formally, we are testing the null hypothesis $H_0: D_{\text{time}} \in \mathcal{D}_+^{13}, D_{\text{imp}}, D_{\text{exp}} \in \mathcal{D}_+^{26}$ versus the alternative hypothesis $H_1: D_{\text{time}} \in \mathcal{D}_+^{13}, \Sigma_{\text{imp}}, \Sigma_{\text{exp}} \in \mathcal{S}_+^{26}$. The approximate 95% quantile of the distribution of the test statistic when the data are missing diagonal entries under the null is 729.8. Setting $E_{ijk} = 0$ for all i and k , the test statistic for the data is $T(E_{..1}, \dots, E_{..13}) = 3354$. This value is much greater than the 95% quantile confirming that we should reject the independence of the ε_{ijk} . It is thus inappropriate to assume that the exporters and importers are independent.

4.2 Application to binary protein-protein interaction network

So far we have developed a testing procedure for the presence of row and column correlations in relational matrices within the framework of matrix normal and general matrix variate elliptically contoured distributions. In this section we propose a methodology that allows us to evaluate the presence of row and column correlations for binary relational data, where the observed network is represented by a sociomatrix matrix A where a_{ij} is a binary indicator of a relationship from node i to node j . In this example we use the protein-protein interaction data of Butland et al. (2005), which consists of a record of interactions between $m = 270$ essential proteins of *E. coli*. The data are organized into a 270×270 binary matrix A , where $a_{ij} = 1$ if protein j binds to protein i and $a_{ij} = 0$ otherwise. The network has one large connected component with 234 nodes as shown in the left hand side of Figure 3. In this case diagonal elements of the matrix A are meaningful since proteins may bind to themselves.

A popular class of models for such data are latent factor models (Hoff, 2005, 2008), which represent the probability that $a_{ij} = 1$ as a function of the vector product of low-dimensional sender and receiver factors u_i and v_j via a probit link, so that $\text{Pr}(a_{i,j}=1) = \Phi(\gamma + u_i'v_j)$, where Φ is the standard normal CDF and γ is an intercept term. This model has a representation as a latent variable model as follows:

$$a_{ij}=1(y_{ij}>-\gamma), \quad y_{ij}=u_i^t v_j + \varepsilon_{ij}, \quad Y=UV^t+E,$$

where $\{\varepsilon_{ij}\} \stackrel{\text{iid}}{\sim} N(0, 1)$ and $u_i, v_j \in \mathbb{R}^R$ for some choice of $R \ll 270$, so that the mean matrix UV^t for Y has rank R . It follows that $E[Y Y^t] = U(V^t V)U^t + I$ and $E[Y^t Y] = V(U^t U)V^t + I$ and so U and V describe the row heterogeneity and column heterogeneity of Y , respectively.

We extend the test developed in this article to evaluate how well a model of a given rank R captures the row and column heterogeneity of the data. If the model is correct for given values of R , U and V , the residual matrix $E = Y - UV^t$ should not exhibit any row or column covariance, and our test statistic $T(E)$ should be small. Conversely, if the data exhibit substantial row and column dependence beyond what is represented by U and V , then the test statistic should be large.

Unfortunately, the value of $E = Y - UV^t$ is not directly observed as Y , U and V are unknown. However, a Bayesian analysis can provide a posterior distribution of $E = Y - UV^t$, from which the posterior distribution of the test statistic $t(E)$ and p -value can be obtained. The resulting conditional distribution of the p -value is called a “fuzzy p -value distribution” (Thompson and Geyer, 2007), and represents the uncertainty in what the p -value would be if Y , U and V were observed. We use the Markov chain Monte Carlo procedure outlined in Hoff (2008) to obtain simulations of $\{Y, U, V\}$ from the posterior distribution given the data A . For each simulated triple $\{Y, U, V\}$, the value of $E = Y - UV^t$ can be computed and a p -value obtained using our testing procedure. If there is strong evidence in the data for row and column dependence beyond what is represented with the rank- R factor model, then the fuzzy p -value distribution should be skewed towards zero (see Thompson and Geyer (2007) for more on the interpretation of fuzzy p -values).

As this is a very sparse network (the interaction rate is $\rho = 0.03$), we expect a low rank approximation to be appropriate. In fact, analysis using cross validation of a symmetrized version of this data identified $R = 3$ to be an appropriate rank in Hoff (2008). In the right hand side of Figure 3 we present the distributions of the fuzzy p -values for $R \in \{0, 1, 2, 3\}$. A visual inspection of the the fuzzy p -values in the four panels of the figure give guidance on the dimension of the latent factors. For example, the fuzzy p -value distribution for the $R = 0$ model is concentrated at low values, suggesting that the likely values of E under this model would reject the hypotheses of no row or column covariance. For $R \in \{1, 2\}$ the fuzzy p -values are no longer concentrated lower than 0.05, but the distribution is skewed to the right, which we take as evidence that there is correlation in Y that is not captured by the rank 1 and rank 2 models. For $R = 3$ the posterior distribution of $E = Y - UV^t$ is not concentrated around values that would give low p -values, indicating a lack of evidence for row and column dependence in Y beyond that provided by the rank-3 mean model UV^t .

5 Discussion

In this article we presented a likelihood ratio test for relational datasets. Unlike the previous testing literature for matrix normal models that required multiple observations, and concentrated on testing a null of separable covariances versus an unstructured alternative, we

proposed testing a null of no row or column correlations versus an alternative of full row and column correlations using a single observation of a network. While the form of the null distribution of the test statistic is intractable, we are able to simulate a reference distribution for the test statistic under the null due to its invariance to left and right diagonal transformations of the data. In the power simulations of Section 3 we demonstrated the power of the test against maximally sparse and nonseparable alternatives.

This test can be applied to a wide variety of relational data. While the test was developed using the matrix-normal model, we have shown that this distributional assumptions can be greatly relaxed. Specifically, if we consider a data matrix Y with an arbitrary matrix variate elliptically contoured distribution that is centered at the zero matrix, the test statistic for testing for correlation among the rows and among the columns of Y is identical to that of the matrix normal case. We have also demonstrated that the test can accommodate frequently observed features of relational data such as non-zero mean, missing diagonal and multiple observations. In Section 4.2 we demonstrated an application of the theory developed in this paper to binary network data where the matrix Y is an adjacency matrix. The method we describe for binary data can be extended to ordinal and discrete data that can be modeled via a latent matrix variate elliptically contoured distribution.

Once we reject the null hypothesis of independence among the rows and among the columns of a relational matrix, we are faced with the challenge of modeling the dependence in the data. As shown in Section 2.1, for a mean zero matrix normal distribution, the MLE is not

unique. Specifically, there is an MLE for each $\Sigma_r \in \mathcal{S}_+^m$ given by $(\Sigma_r, Y^t \Sigma_r^{-1} Y / m)$. We have observed in separate work that it is possible to distinguish between MLEs by considering their risk. However, other than in very specialized cases (such as equal eigenvalues of Σ_r and Σ_c), obtaining analytic results for identifying risk optimal MLEs is difficult. The presence of a non-zero mean leads to an unbounded likelihood and further complicates the problem of estimation. Several authors have recently considered Bayesian and penalized likelihood approaches to this estimation problem. Yu et al. (2007) studied hierarchical Gaussian Process priors in the context of a classification problem. In our context, this approach results in a matrix normal prior for the mean parameters and inverse Wishart priors for the row and column covariance matrices. A second approach based on a mixture of independent L_1 and L_2 penalties on the row and column precision matrices was proposed by Allen and Tibshirani (2010).

Computer code and data for Sections 3 and 4 are available at the authors' websites.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was partially supported by NICHD grant 1R01HD067509-01A1. The authors would like to thank Bailey Fosdick for helpful discussions.

References

- Airoldi E, Blei D, Fienberg S, Xing E. Mixed membership stochastic blockmodels. *The Journal of Machine Learning Research*. 2008; 9:1981–2014. [PubMed: 21701698]
- Allen GI, Tibshirani R. Transposable regularized covariance models with an application to missing data imputation. *The Annals of Applied Statistics*. 2010; 4(2):764–790. [PubMed: 26877823]
- Anderson T, Hsu H, Fang KT. Maximum-likelihood estimates and likelihood-ratio criteria for multivariate elliptically contoured distributions. *Canadian Journal of Statistics*. 1986; 14(1):55–59.
- Bergmann S, Ihmels J, Barkai N. Similarities and differences in genome-wide expression data of six organisms. *PLoS Biology*. 2003; 2(1):e9. [PubMed: 14737187]
- Bergstrand J. The gravity equation in international trade: some microeconomic foundations and empirical evidence. *The review of economics and statistics*. 1985; 67(3):474–481.
- Butland G, Peregrín-Alvarez JM, Li J, Yang W, Yang X, Canadien V, Starostine A, Richards D, Beattie B, Krogan N, et al. Interaction network containing conserved and essential protein complexes in *Escherichia coli*. *Nature*. 2005; 433(7025):531–537. [PubMed: 15690043]
- Dawid A. Some matrix-variate distribution theory: notational considerations and a bayesian application. *Biometrika*. 1981; 68(1):265–274.
- Eaton, M. *Multivariate statistics: a vector space approach*. Wiley; New York: 1983.
- Eriksen PS. Proportionality of covariance matrices. *Ann Statist*. 1987; 15(2):732–748.
- Fletcher A, Bonell C, Sorhaindo A. You are what your friends eat: systematic review of social network analyses of young people's eating behaviours and bodyweight. *Journal of epidemiology and community health*. 2011; 65(6):548–555. [PubMed: 21427455]
- Flury BK. Proportionality of k covariance matrices. *Statistics & probability letters*. 1986; 4(1):29–33.
- Gupta A, Varga T. A new class of matrix variate elliptically contoured distributions. *Statistical Methods & Applications*. 1994; 3(2):255–270.
- Gupta A, Varga T. Some inference problems for matrix variate elliptically contoured distributions. *Statistics: A Journal of Theoretical and Applied Statistics*. 1995; 26(3):219–229.
- Hoff P. Bilinear mixed-effects models for dyadic data. *J Amer Statist Assoc*. 2005; 100(469):286–295.
- Hoff, P. Modeling homophily and stochastic equivalence in symmetric relational data. In: Platt, J.; Koller, D.; Singer, Y.; Roweis, S., editors. *Advances in Neural Information Processing Systems 20*. MIT Press; Cambridge, MA: 2008. p. 657–664.
- Hoff P. Separable covariance arrays via the tucker product, with applications to multivariate relational data. *Bayesian Analysis*. 2011; 6(2):179–196.
- Hoff P, Raftery A, Handcock M. Latent space approaches to social network analysis. *Journal of the American Statistical Association*. 2002; 97(460):1090–1098.
- Holland P, Laskey K, Leinhardt S. Stochastic blockmodels: first steps. *Social networks*. 1983; 5(2):109–137.
- Jensen ST, Johansen S. Estimation of proportional covariances. *Statistics & probability letters*. 1987; 6(2):83–85.
- Jensen ST, Madsen J. Estimation of proportional covariances in the presence of certain linear restrictions. *The Annals of Statistics*. 2004; 32(1):219–232.
- Kenny D, La Voie L. The social relations model. *Advances in experimental social psychology*. 1984; 18:142–182.
- Lafosse R, Ten Berge J. A simultaneous concordance algorithm for the analysis of two partitioned matrices. *Computational statistics & data analysis*. 2006; 50(10):2529–2535.
- Lazzarini S, Chaddad F, Cook M. Integrating supply chain and network analyses: the study of networks. *Journal on chain and network science*. 2001; 1(1):7–22.
- Leskovec, J.; Lang, K.; Dasgupta, A.; Mahoney, M. Statistical properties of community structure in large social and information networks. *Proceedings of the 17th international conference on World Wide Web; ACM*; 2008. p. 695–704.
- Li H. The covariance structure and likelihood function for multivariate dyadic data. *J Multivariate Anal*. 2006; 97(6):1263–1271.

- Li H, Loken E. A unified theory of statistical analysis and inference for variance component models for dyadic data. *Statist Sinica*. 2002; 12(2):519–535.
- Lincoln, J.; Gerlach, M. Japan's network economy: structure, persistence, and change. Cambridge University Press; 2004.
- Lu N, Zimmerman D. The likelihood ratio test for a separable covariance matrix. *Statistics & probability letters*. 2005; 73(4):449–457.
- McQuitty L, Clark J. Clusters from iterative, intercolumnar correlational analysis. *Educational and psychological measurement*. 1968
- Mitchell M, Genton M, Gumpertz M. A likelihood ratio test for separability of covariances. *Journal of Multivariate Analysis*. 2006; 97(5):1025–1043.
- Nowicki K, Snijders T. Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association*. 2001; 96(455):1077–1087.
- Panning W. Fitting blockmodels to data. *Social Networks*. 1982; 4(1):81–101.
- Potter G, Handcock M, Longini I, Halloran M. Estimating within-school contact networks to understand influenza transmission. *Ann Appl Stat*. 2012; 6(1):1–26. [PubMed: 22639701]
- Roy A, Khattree R. On implementation of a test for kronecker product covariance structure for multivariate repeated measures data. *Statistical Methodology*. 2005; 2(4):297–306.
- Sampson, S. PhD thesis. Cornell University; 1968 Sep. A novitiate in a period of change: An experimental and case study of social relationships.
- Srivastava M, von Rosen T, Von Rosen D. Models with a kronecker product covariance structure: estimation and testing. *Mathematical Methods of Statistics*. 2008; 17(4):357–370.
- Stuart J, Segal E, Koller D, Kim S. A gene-coexpression network for global discovery of conserved genetic modules. *Science*. 2003; 302(5643):249–255. [PubMed: 12934013]
- Thompson E, Geyer C. Fuzzy p-values in latent variable problems. *Biometrika*. 2007; 94(1):49–60.
- Tinbergen, J., et al. Shaping the world economy: Suggestions for an international economic policy. Twentieth Century Fund; New York: 1962.
- Tseng P. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*. 2001; 109(3):475–494.
- Wang Y, Wong G. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*. 1987; 82(397):8–19.
- Westveld A, Hoff P. A mixed effects model for longitudinal relational and network data, with applications to international trade and conflict. *Ann Appl Stat*. 2011a; 5(2A):843–872.
- Westveld AH, Hoff PD. A mixed effects model for longitudinal relational and network data, with applications to international trade and conflict. *Ann Appl Stat*. 2011b; 5(2A):843–872.
- White H, Boorman S, Breiger R. Social structure from multiple networks. i. blockmodels of roles and positions. *American journal of sociology*. 1976:730–780.
- Yu K, Chu W, Yu S, Tresp V, Xu Z. Stochastic relational models for discriminative link prediction. *Advances in neural information processing systems*. 2007; 19:1553.

Figure 1.

The top row of panels displays the power of the test under the exchangeable covariance model of Section 3.1. The bottom left panel displays the power of the test for the maximally sparse covariance model of Section 3.2 and the bottom right panel displays the power of the test for the nonseparable stochastic blockmodel of Section 3.3.

Figure 2. Plots of the first two eigenvectors of the estimates $\hat{R}_{row}$ and $\hat{R}_{col}$ of the row and column correlation matrices. Proximity of countries in the eigenspace indicates a positive correlation.

Figure 3. Protein-protein interaction data and histograms of fuzzy p -values for models of ranks $R \in \{0, 1, 2, 3\}$. The bins have a width of 0.05.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1

95% quantile of the null distribution of the test statistic for testing H_0 versus H_1 . Approximation from 100,000 simulated values.

Dimension m	5	10	15	20	25	30	50	100
95% quantile	43.3	144.3	297.4	502.8	760.0	1064.6	2802.1	10668.4