

Bayesian inference for general Gaussian graphical models with application to multivariate lattice data

Adrian Dobra*, Alex Lenkoski† and Abel Rodriguez‡

Abstract

We introduce efficient Markov chain Monte Carlo methods for inference and model determination in multivariate and matrix-variate Gaussian graphical models. Our framework is based on the G-Wishart prior for the precision matrix associated with graphs that can be decomposable or non-decomposable. We extend our sampling algorithms to a novel class of conditionally autoregressive models for sparse estimation in multivariate lattice data, with a special emphasis on the analysis of spatial data. These models embed a great deal of flexibility in estimating both the correlation structure across outcomes and the spatial correlation structure, thereby allowing for adaptive smoothing and spatial autocorrelation parameters. Our methods are illustrated using simulated and real-world examples, including an application to cancer mortality surveillance.

Keywords: CAR model; Gaussian graphical model; G-Wishart distribution; Lattice data; Markov chain Monte Carlo (MCMC) simulation; Spatial statistics.

1 Introduction

Graphical models (Lauritzen, 1996), which encode the conditional independence among variables using an undirected graph, have become a popular tool for sparse estimation in both the statistics and machine learning literatures. Implementing model selection approaches in the context of graphical models typically allows a dramatic reduction in the number of parameters under consideration, preventing overfitting and improving predictive capability. In particular, Bayesian approaches to inference in graphical models generate regularized estimators that incorporate uncertainty of the model structure.

The focus of this paper is Bayesian inference in Gaussian graphical models (Dempster, 1972) based on the G-Wishart prior (Roverato, 2002; Atay-Kayis & Massam, 2005; Letac & Massam,

*Departments of Statistics, Biobehavioral Nursing and Health Systems and the Center for Statistics and the Social Sciences, Box 354322, University of Washington, Seattle, WA 98195, U.S.A. (email: adobra@uw.edu).

†Institut für Angewandte Mathematik, Universität Heidelberg, Im Neuenheimer Feld 294, 69115 Heidelberg, Germany (email: alex.lenkoski@uni-heidelberg.de).

‡Department of Applied Mathematics and Statistics, University of California, Mailstop SOE2, Santa Cruz, California 95064, U.S.A. (email: abel@soe.ucsc.edu).

2007). This class of distributions is extremely attractive since it represents the conjugate family for the precision matrix whose elements associated with edges in the underlying graph are constrained to be equal to zero. In addition, the G-Wishart naturally accommodates decomposable and non-decomposable graphs in a coherent manner.

Many recent papers have described various stochastic search methods in Gaussian graphical models (GGMs) with the G-Wishart prior based on marginal likelihoods which, in this case, are given by the ratio of the normalizing constants of the posterior and prior G-Wishart distributions – see Atay-Kayis & Massam (2005); Jones et al. (2005); Carvalho & Scott (2009); Lenkoski & Dobra (2010) and the references therein. Although the computation of marginal likelihoods for decomposable graphs is straightforward, similar computations for non-decomposable graphs raise significant numerical challenges that are a consequence of having to approximate the normalizing constants of the corresponding G-Wishart posterior distributions. In that regard, Lenkoski & Dobra (2010) point out that the Monte Carlo method of Atay-Kayis & Massam (2005) is fast and accurate when employed for the computations of the normalizing constant of G-Wishart priors, although it is slow to converge for computations that involve G-Wishart posteriors. This remark leads to the idea of devising a Bayesian model determination method that avoids the computation of posterior normalizing constants.

The contributions of this paper are twofold. First, we devise a new reversible jump sampler (Green, 1995) for model determination and estimation in multivariate and matrix-variate GGMs that avoids the calculation of posterior marginal likelihoods. The algorithm relies on a novel Metropolis-Hastings algorithm for sampling from the G-Wishart distribution associated with an arbitrary graph. Related methods are the direct sampling algorithm of Carvalho et al. (2007) and the block Gibbs sampler algorithm of Piccioni (2000) and Ascì & Piccioni (2007). However, known applications of the former method involve exclusively decomposable graphs, while the latter method relies on the identification of all the cliques of the graph, a problem that is extremely difficult in the case of graphs with many vertices or edges (Tomita et al., 2006). Our method does not require the determination of the cliques and works well for any graph as it is based on the Cholesky decompositions of the precision matrix developed in Roverato (2000, 2002); Atay-Kayis & Massam (2005). Second, we devise a new flexible class of conditionally autoregressive models (CAR) (Besag, 1974) for lattice data that rely on our novel sampling algorithms. The link between conditionally autoregressive models and GGMs was originally pointed out by Besag & Kooperberg (1995). However, since typical neighborhood graphs are non-decomposable, fully exploiting this connection within a Bayesian framework requires that we are able to estimate Gaussian graphical based on general graphs. Our main focus is on applications to multivariate lattice data, where our approach based on matrix-variate GGMs provides a natural approach to create sparse multivariate CAR models.

The organization of the paper is as follows. Section 2 formally introduces GGMs and the G-Wishart distribution, along with a description of our novel sampling algorithm in the case of multivariate GGMs. This algorithm represents a generalization of the work of Giudici & Green (1999) and is applicable beyond decomposable GGMs. Section 3 discusses inference and model determination in matrix-variate GGMs. Unlike the related framework developed in Wang & West (2009) (which requires computation of marginal likelihoods and involves exclusively decompos-

able graphs), our sampler operates on the joint distribution of the row and column precision matrices, the row and column conditional independence graphs and the auxiliary variable that needs to be introduced to solve the underlying non-identifiability problem associated with matrix-variate normal distributions. Section 4 reviews conditional autoregressive priors for lattice data and discuss their connection to GGMs. This section discusses both univariate and multivariate models for continuous and discrete data based on generalized linear models. Section 5 presents three illustrations of our methodology: a simulation study, a spatial linear regression for modeling statewide SAT scores, and a cancer mortality mapping model. Finally, Section 6 concludes the paper by discussing some of the limitations of the approach and future directions for our work.

2 Gaussian graphical models and the G-Wishart distribution

We let $\mathbf{X} = \mathbf{X}_{V_p}$, $V_p = \{1, 2, \dots, p\}$, be a random vector with a p -dimensional multivariate normal distribution $N_p(\mathbf{0}, \mathbf{K}^{-1})$. We consider a graph $G = (V_p, E)$, where each vertex $i \in V$ corresponds with a random variable X_i and $E \subset V \times V$ are undirected edges. Here “undirected” means that $(i, j) \in E$ if and only if $(j, i) \in E$. We denote by \mathcal{G}_p the set of all $p(p-1)/2$ undirected graphs with p vertices. A Gaussian graphical model with conditional independence graph G is constructed by constraining to zero the off-diagonal elements of \mathbf{K} that do not correspond with edges in G (Dempster, 1972). If $(i, j) \notin E$, X_i and X_j are conditional independent given the remaining variables. The precision matrix $\mathbf{K} = (K_{ij})_{1 \leq i, j \leq p}$ is constrained to the cone \mathbf{P}_G of symmetric positive definite matrices with off-diagonal entries $K_{ij} = 0$ for all $(i, j) \notin E$.

We consider the G -Wishart distribution $\text{Wis}_G(\delta, \mathbf{D})$ with density

$$p(\mathbf{K}|G, \delta, \mathbf{D}) = \frac{1}{I_G(\delta, \mathbf{D})} (\det \mathbf{K})^{(\delta-2)/2} \exp \left\{ -\frac{1}{2} \langle \mathbf{K}, \mathbf{D} \rangle \right\}, \quad (1)$$

with respect to the Lebesgue measure on \mathbf{P}_G (Roverato, 2002; Atay-Kayis & Massam, 2005; Letac & Massam, 2007). Here $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}^T \mathbf{B})$ denotes the trace inner product. The normalizing constant $I_G(\delta, \mathbf{D})$ is finite provided $\delta > 2$ and \mathbf{D} positive definite (Diaconnis & Ylvisaker, 1979).

We write $\mathbf{K} \in \mathbf{P}_G$ as

$$\mathbf{K} = \mathbf{Q}^T (\boldsymbol{\Psi}^T \boldsymbol{\Psi}) \mathbf{Q}, \quad (2)$$

where $\mathbf{Q} = (Q_{ij})_{1 \leq i \leq j \leq p}$ and $\boldsymbol{\Psi} = (\Psi_{ij})_{1 \leq i \leq j \leq p}$ are upper triangular, while $\mathbf{D}^{-1} = \mathbf{Q}^T \mathbf{Q}$ is the Cholesky decomposition of \mathbf{D}^{-1} . We see that $\mathbf{K} = \boldsymbol{\Phi}^T \boldsymbol{\Phi}$ with $\boldsymbol{\Phi} = \boldsymbol{\Psi} \mathbf{Q}$ is the Cholesky decomposition of \mathbf{K} . The zero constraints on the off-diagonal elements of \mathbf{K} associated with G induce well-defined sets of free elements $\boldsymbol{\Phi}^{\nu(G)} = \{\Phi_{ij} : (i, j) \in \nu(G)\}$ and $\boldsymbol{\Psi}^{\nu(G)} = \{\Psi_{ij} : (i, j) \in \nu(G)\}$ of the matrices $\boldsymbol{\Phi}$ and $\boldsymbol{\Psi}$ (Roverato, 2002; Atay-Kayis & Massam, 2005), where

$$\nu(G) = \{(i, i) : i \in V_p\} \cup \{(i, j) : i < j \text{ and } (i, j) \in E\}.$$

We denote by $M^{\nu(G)}$ the set of incomplete triangular matrices whose elements are indexed by $\nu(G)$ and whose diagonal elements are strictly positive. Both $\boldsymbol{\Phi}^{\nu(G)}$ and $\boldsymbol{\Psi}^{\nu(G)}$ must belong to $M^{\nu(G)}$. The

non-free elements of Ψ are determined through the completion operation from Lemma 2 of Atay-Kayis & Massam (2005) as a function of the free elements $\Psi^{v(G)}$. Each element Ψ_{ij} with $i < j$ and $(i, j) \notin E$ is a function of the other elements $\Psi_{i'j'}$ that precede it in lexicographical order.

Roverato (2002) proves that the Jacobean of the transformation that maps $\mathbf{K} \in \mathbf{P}_G$ to $\Phi^{v(G)} \in M^{v(G)}$ is

$$J(\mathbf{K} \rightarrow \Phi^{v(G)}) = 2^p \prod_{i=1}^p \Phi_{ii}^{v_i^G+1},$$

where v_i^G is the number of elements in the set $\{j : j > i \text{ and } (i, j) \in E\}$. Furthermore, Atay-Kayis & Massam (2005) show that the Jacobean of the transformation that maps $\Phi^{v(G)} \in M^{v(G)}$ to $\Psi^{v(G)} \in M^{v(G)}$ is given by

$$J(\Phi^{v(G)} \rightarrow \Psi^{v(G)}) = \prod_{i=1}^p Q_{ii}^{d_i^G+1},$$

where d_i^G is the number of elements in the set $\{j : j < i \text{ and } (i, j) \in E\}$. We have $\det \mathbf{K} = \prod_{i=1}^p \Phi_{ii}^2$ and $\Phi_{ii} = \Psi_{ii} Q_{ii}$. It follows that the density of $\Psi^{v(G)}$ with respect to the Lebesgue measure on $M^{v(G)}$ is

$$\begin{aligned} p(\Psi^{v(G)} | \delta, \mathbf{D}) &= p(\mathbf{K} | G, \delta, \mathbf{D}) J(\mathbf{K} \rightarrow \Phi^{v(G)}) J(\Phi^{v(G)} \rightarrow \Psi^{v(G)}), \\ &= \frac{2^p}{I_G(\delta, \mathbf{D})} \prod_{i=1}^p Q_{ii}^{v_i^G+d_i^G+\delta} \prod_{i=1}^p \Psi_{ii}^{v_i^G+\delta-1} \exp\left(-\frac{1}{2} \sum_{i,j=1}^p \Psi_{ij}^2\right). \end{aligned}$$

We note that $v_i^G + d_i^G$ represents the number of neighbors of vertex i in the graph G .

2.1 Sampling from the G-Wishart distribution

We introduce a Metropolis-Hastings algorithm for sampling from the G-Wishart distribution (1). We consider a strictly positive precision parameter σ_m . We denote by $\mathbf{K}^{[s]} = \mathbf{Q}^T (\Psi^{[s]})^T \Psi^{[s]} \mathbf{Q}$ the current state of the chain with $(\Psi^{[s]})^{v(G)} \in M^{v(G)}$. The next state $\mathbf{K}^{[s+1]} = \mathbf{Q}^T (\Psi^{[s+1]})^T \Psi^{[s+1]} \mathbf{Q}$ is obtained by sequentially perturbing the free elements $(\Psi^{[s]})^{v(G)}$. A diagonal element $\Psi_{i_0 i_0}^{[s]} > 0$ is updated by sampling a value γ from a $\mathcal{N}(\Psi_{i_0 i_0}^{[s]}, \sigma_m^2)$ distribution truncated below at zero. We define the upper triangular matrix Ψ' such that $\Psi'_{ij} = \Psi_{ij}^{[s]}$ for $(i, j) \in v(G) \setminus \{(i_0, i_0)\}$ and $\Psi'_{i_0 i_0} = \gamma$. The non-free elements of Ψ' are obtained through the completion operation from $(\Psi')^{v(G)}$. The Markov chain moves to $\mathbf{K}' = \mathbf{Q}^T (\Psi')^T \Psi' \mathbf{Q}$ with probability $\min\{R_m, 1\}$, where

$$R_m = \frac{p((\Psi')^{v(G)} | \delta, \mathbf{D}) p(\Psi_{i_0 i_0}^{[s]} | \Psi'_{i_0 i_0})}{p((\Psi^{[s]})^{v(G)} | \delta, \mathbf{D}) p(\Psi'_{i_0 i_0} | \Psi_{i_0 i_0}^{[s]})} = \frac{\phi(\Psi_{i_0 i_0}^{[s]} / \sigma_m)}{\phi(\Psi'_{i_0 i_0} / \sigma_m)} \left(\frac{\Psi'_{i_0 i_0}}{\Psi_{i_0 i_0}^{[s]}}\right)^{v_{i_0}^G + \delta - 1} R'_m. \quad (3)$$

Here $\phi(\cdot)$ represents the CDF of the standard normal distribution and

$$R'_m = \exp\left\{-\frac{1}{2} \sum_{i,j=1}^p \left[(\Psi'_{ij})^2 - (\Psi_{ij}^{[s]})^2\right]\right\}. \quad (4)$$

A free off-diagonal element $\Psi_{i_0 j_0}^{[s]}$ is updated by sampling a value $\gamma' \sim \mathcal{N}(\Psi_{i_0 j_0}^{[s]}, \sigma_m^2)$. We define the upper triangular matrix Ψ' such that $\Psi'_{ij} = \Psi_{ij}^{[s]}$ for $(i, j) \in \nu(G) \setminus \{(i_0, j_0)\}$ and $\Psi'_{i_0 j_0} = \gamma'$. The remaining elements of Ψ' are determined by completion from $(\Psi')^{\nu(G)}$. The proposal distribution is symmetric $p(\Psi'_{i_0 j_0} | \Psi_{i_0 j_0}^{[s]}) = p(\Psi_{i_0 j_0}^{[s]} | \Psi'_{i_0 j_0})$, thus we accept the transition of the chain from $\mathbf{K}^{[s]}$ to $\mathbf{K}' = \mathbf{Q}^T (\Psi')^T \Psi' \mathbf{Q}$ with probability $\min\{R'_m, 1\}$, where R'_m is given in equation (4). Since $(\Psi^{[s]})^{\nu(G)} \in M^{\nu(G)}$, we have $(\Psi')^{\nu(G)} \in M^{\nu(G)}$ which implies $\mathbf{K}' \in \mathcal{P}_G$.

We denote by $\mathbf{K}^{[s+1]}$ the precision matrix obtained after completing all the Metropolis-Hastings updates associated with the free elements indexed by $\nu(G)$. The acceptance probabilities depend only on the free elements of the upper diagonal matrices $\Psi^{[s]}$ and Ψ' . An efficient implementation of our sampling procedure calculates only the final matrix $\mathbf{K}^{[s+1]}$ and avoids determining the intermediate candidate precision matrices \mathbf{K}' . There is another key computational aspect that relates to the dependence of the Cholesky decomposition on a particular ordering of the variables. Consider two free elements $\Psi_{i_0 j_0}^{[s]}$ and $\Psi_{i'_0 j'_0}^{[s]}$ such that $(i_0, j_0) < (i'_0, j'_0)$ in lexicographical order. The completion operation from Lemma 2 of Atay-Kayis & Massam (2005) shows that perturbing the value of $\Psi_{i_0 j_0}^{[s]}$ leads to a change in the value of all non-free elements $\{\Psi_{ij}^{[s]} : (i, j) \notin E \text{ and } (i_0, j_0) < (i, j)\}$. On the other hand, perturbing the value of $\Psi_{i'_0 j'_0}^{[s]}$ implies a change in the smaller set of elements $\{\Psi_{ij}^{[s]} : (i, j) \notin E \text{ and } (i'_0, j'_0) < (i, j)\}$. This means that the chain is less likely to move when performing the update associated with $\Psi_{i_0 j_0}^{[s]}$ than when updating $\Psi_{i'_0 j'_0}^{[s]}$. As such, the ordering of the variables is modified at each iteration of the MCMC sampler. More specifically, a permutation ν is uniformly drawn from the set of all possible permutations Υ_p of V . The row and columns of \mathbf{D} are reordered according to ν and a new Cholesky decomposition of \mathbf{D}^{-1} is determined. The set $\nu(G)$ and $\{d_i^G : i \in V\}$ are recalculated given the ordering of the vertices induced by ν , i.e. $i <_\nu j$ if and only if $\nu(i) < \nu(j)$. Changing the ordering guarantees that across iterations it is equally likely that (i_0, j_0) precedes or succeeds (i'_0, j'_0) .

Our later developments from Section 3 involve sampling from the G-Wishart distribution (1) subject to the constraint $K_{11} = 1$. We have

$$K_{11}^{[s]} = 1 \Leftrightarrow (\Psi^{[s]} \mathbf{Q})_{11} = 1 \Leftrightarrow \Psi_{11}^{[s]} = 1/Q_{11}.$$

We subsequently obtain the next state $\mathbf{K}^{[s+1]}$ of the Markov chain by perturbing the free elements $(\Psi^{[s]})^{\nu(G) \setminus \{(1,1)\}}$. When defining the triangular matrix Ψ' we set $\Psi'_{11} = 1/Q_{11}$ which implies that the corresponding candidate matrix \mathbf{K}' has $K'_{11} = 1$. Thus $\mathbf{K}^{[s+1]}$ also obeys the constraint $K_{11}^{[s+1]} = 1$. The random orderings of the variables need to be drawn from the set $\Upsilon_p^{(1,1)}$ of permutations $\nu \in \Upsilon_p$ such that $\nu(1) = 1$. This way the (1, 1) element of \mathbf{K} always occupies the same position.

2.2 Bayesian inference in GGMs

We let $\mathcal{D} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$ be the observed data of n independent samples from $\mathcal{N}_p(\mathbf{0}, \mathbf{K}^{-1})$. The likelihood function is

$$p(\mathcal{D} | \mathbf{K}) = (2\pi)^{-np/2} (\det \mathbf{K})^{n/2} \exp \left\{ -\frac{1}{2} \langle \mathbf{K}, \mathbf{U} \rangle \right\}, \quad (5)$$

where $\mathbf{U} = \sum_{j=1}^n \mathbf{x}^{(j)}(\mathbf{x}^{(j)})^T$ is the observed sum-of-products matrix. Given a graph G , we assume a G-Wishart prior $\text{Wis}_G(\delta_0, \mathbf{D}_0)$ for the precision matrix \mathbf{K} . We take $\delta_0 = 3 > 2$ and $\mathbf{D}_0 = \mathbf{I}_p$, the p -dimensional identity matrix. This choice implies that prior for \mathbf{K} is equivalent with one observed sample, while the observed variables are assumed to be apriori independent of each other. Since the G-Wishart prior for \mathbf{K} is conjugate for the likelihood (5), the posterior of \mathbf{K} given G is $\text{Wis}_G(n + \delta_0, \mathbf{U} + \mathbf{D}_0)$. We also assume a uniform prior $p(G) = 2/p(p-1)$ on \mathcal{G}_p .

We develop a MCMC algorithm for sampling from the joint distribution

$$p(\mathcal{D}, \mathbf{K}, G | \delta_0, \mathbf{D}_0) \propto p(\mathcal{D} | \mathbf{K}) p(\mathbf{K} | G, \delta_0, \mathbf{D}_0) p(G),$$

that is well-defined if and only if $\mathbf{K} \in \mathbf{P}_G$. We denote the current state of the chain by $(\mathbf{K}^{[s]}, G^{[s]})$, $\mathbf{K}^{[s]} \in \mathbf{P}_{G^{[s]}}$. Its next state $(\mathbf{K}^{[s+1]}, G^{[s+1]})$, $\mathbf{K}^{[s+1]} \in \mathbf{P}_{G^{[s+1]}}$, is generated by sequentially performing the following two steps. We make use of two strictly positive precision parameters σ_m and σ_g that remain fixed at some suitable small values. We assume that the ordering of the variables has been changed according to a permutation ν selected at random from the uniform distribution on Υ_p . We denote by $(\mathbf{U} + \mathbf{D}_0)^{-1} = (\mathbf{Q}^*)^T \mathbf{Q}^*$ the Cholesky decomposition of $(\mathbf{U} + \mathbf{D}_0)^{-1}$, where the rows and columns of this matrix have been permuted according to ν .

We denote by $\text{nbd}_p^+(G)$ the graphs that can be obtained by adding an edge to a graph $G \in \mathcal{G}_p$ and by $\text{nbd}_p^-(G)$ the graphs that are obtained by deleting an edge from G . We call the one-edge-way set of graphs $\text{nbd}_p(G) = \text{nbd}_p^+(G) \cup \text{nbd}_p^-(G)$ the neighborhood of G in \mathcal{G}_p . These neighborhoods connect any two graphs in \mathcal{G}_p through a sequence of graphs such that two consecutive graphs in this sequence are each others' neighbors.

Step 1: Resample the graph. We sample a candidate graph $G' \in \text{nbd}_p(G^{[s]})$ from the proposal

$$q(G' | G^{[s]}) = \frac{1}{2} \text{Uni}(\text{nbd}_p^+(G^{[s]})) + \frac{1}{2} \text{Uni}(\text{nbd}_p^-(G^{[s]})), \quad (6)$$

where $\text{Uni}(A)$ represents the uniform distribution on the discrete set A . The distribution (6) gives an equal probability of proposing to delete an edge from the current graph and of proposing to add an edge to the current graph. We favor (6) over the more usual proposal distribution $\text{Uni}(\text{nbd}_p(G^{[s]}))$ that is employed, for example, by Madigan & York (1995). If $G^{[s]}$ contains a very large or a very small number of edges, the probability of proposing a move that adds or, respectively, deletes an edge from $G^{[s]}$ is extremely small when sampling from $\text{Uni}(\text{nbd}_p(G^{[s]}))$, which could lead to poor mixing in the resulting Markov chain.

We assume that the candidate G' sampled from (6) is obtained by adding the edge (i_0, j_0) , $i_0 < j_0$, to $G^{[s]}$. Since $G' \in \text{nbd}_p^+(G^{[s]})$ we have $G^{[s]} \in \text{nbd}_p^-(G')$. We consider the decomposition of the current precision matrix

$$\mathbf{K}^{[s]} = (\mathbf{Q}^*)^T \left((\boldsymbol{\Psi}^{[s]})^T \boldsymbol{\Psi}^{[s]} \right) \mathbf{Q}^*,$$

with $(\boldsymbol{\Psi}^{[s]})^{\nu(G^{[s]})} \in M^{\nu(G^{[s]})}$. Since the vertex i_0 has one additional neighbor in G' , we have $d_{i_0}^{G'} = d_{i_0}^{G^{[s]}}$, $d_{j_0}^{G'} = d_{j_0}^{G^{[s]}} + 1$, $v_{i_0}^{G'} = v_{i_0}^{G^{[s]}} + 1$, $v_{j_0}^{G'} = v_{j_0}^{G^{[s]}}$ and $\nu(G') = \nu(G^{[s]}) \cup \{(i_0, j_0)\}$. We define an upper

triangular matrix Ψ' such that $\Psi'_{ij} = \Psi_{ij}^{[s]}$ for $(i, j) \in \nu(G^{[s]})$. We sample $\gamma \sim N(\Psi_{i_0 j_0}^{[s]}, \sigma_g^2)$ and set $\Psi'_{i_0 j_0} = \gamma$. The rest of the elements of Ψ' are determined from $(\Psi')^{\nu(G')}$ through the completion operation. The value of the free element $\Psi'_{i_0 j_0}$ was set by perturbing the non-free element $\Psi_{i_0 j_0}^{[s]}$. The other free elements of Ψ' and $\Psi^{[s]}$ coincide.

We take $\mathbf{K}' = (\mathbf{Q}^*)^T ((\Psi')^T \Psi') \mathbf{Q}^*$. Since $(\Psi')^{\nu(G')} \in M^{\nu(G')}$, we have $\mathbf{K}' \in \mathcal{P}_{G'}$. The dimensionality of the parameter space increases by one as we move from $(\mathbf{K}^{[s]}, G^{[s]})$ to (\mathbf{K}', G') , thus we must make use of the reversible jump algorithm of Green (1995). The Markov chain moves to (\mathbf{K}', G') with probability $\min\{R_g^+, 1\}$ where R_g^+ is given by

$$R_g^+ = \frac{p(\mathcal{D}|\mathbf{K}')}{p(\mathcal{D}|\mathbf{K}^{[s]})} \frac{p(\mathbf{K}'|G', \delta_0, \mathbf{D}_0)}{p(\mathbf{K}^{[s]}|G^{[s]}, \delta_0, \mathbf{D}_0)} \frac{|\text{nb}d_p^+(G^{[s]})|}{|\text{nb}d_p^-(G')|} \times \quad (7)$$

$$\times \frac{J(\mathbf{K}' \rightarrow (\Psi')^{\nu(G')})}{J(\mathbf{K}^{[s]} \rightarrow (\Psi^{[s]})^{\nu(G^{[s]})})} \frac{J((\Psi^{[s]})^{\nu(G^{[s]})}, \gamma) \rightarrow (\Psi')^{\nu(G')})}{\frac{1}{\sigma_g \sqrt{2\pi}} \exp\left(-\frac{(\Psi'_{i_0 j_0} - \Psi_{i_0 j_0}^{[s]})^2}{2\sigma_g^2}\right)},$$

where $|A|$ denotes the number of elements of the set A . Otherwise the chain stays at $(\mathbf{K}^{[s]}, G^{[s]})$. Since $(\Psi')^{\nu(G')} = (\Psi^{[s]})^{\nu(G^{[s]})}$, the Jacobean of the transformation from $((\Psi^{[s]})^{\nu(G^{[s]})}, \gamma)$ to $(\Psi')^{\nu(G')}$ is equal to 1. Moreover, $\Psi^{[s]}$ and Ψ' have the same diagonal elements, hence $\det \mathbf{K}^{[s]} = \prod_{i=1}^p (Q_{ii}^* \Psi_{ii}^{[s]})^2 = \det \mathbf{K}'$. It follows that equation (7) becomes

$$R_g^+ = \sigma_g \sqrt{2\pi} Q_{i_0 i_0}^* Q_{j_0 j_0}^* \Psi_{i_0 i_0}^{[s]} \frac{I_{G^{[s]}}(\delta_0, \mathbf{D}_0)}{I_{G'}(\delta_0, \mathbf{D}_0)} \frac{|\text{nb}d_p^+(G^{[s]})|}{|\text{nb}d_p^-(G')|} \times \quad (8)$$

$$\times \exp\left\{-\frac{1}{2} \left[\langle \mathbf{K}' - \mathbf{K}^{[s]}, \mathbf{U} + \mathbf{D}_0 \rangle - \frac{(\Psi'_{i_0 j_0} - \Psi_{i_0 j_0}^{[s]})^2}{\sigma_g^2} \right]\right\}.$$

Next we assume that the candidate G' is obtained by deleting the edge (i_0, j_0) from $G^{[s]}$. We have $d_{i_0}^{G'} = d_{i_0}^{G^{[s]}}$, $d_{j_0}^{G'} = d_{j_0}^{G^{[s]}} - 1$, $\nu_{i_0}^{G'} = \nu_{i_0}^{G^{[s]}} - 1$, $\nu_{j_0}^{G'} = \nu_{j_0}^{G^{[s]}}$ and $\nu(G') = \nu(G^{[s]}) \setminus \{(i_0, j_0)\}$. We define an upper triangular matrix Ψ' such that $\Psi'_{ij} = \Psi_{ij}^{[s]}$ for $(i, j) \in \nu(G')$. The rest of the elements of Ψ' are determined through completion. The free element $\Psi_{i_0 j_0}^{[s]}$ becomes non-free in Ψ' , hence the parameter space decreases by 1 as we move from $(\Psi^{[s]})^{\nu(G^{[s]})}$ to $(\Psi')^{\nu(G')} \in M^{\nu(G')}$. As before, we take $\mathbf{K}' = (\mathbf{Q}^*)^T ((\Psi')^T \Psi') \mathbf{Q}^*$. The acceptance probability of the transition from $(\mathbf{K}^{[s]}, G^{[s]})$ to (\mathbf{K}', G') is $\min\{R_g^-, 1\}$ where

$$R_g^- = \left(\sigma_g \sqrt{2\pi} Q_{i_0 i_0}^* Q_{j_0 j_0}^* \Psi_{i_0 i_0}^{[s]}\right)^{-1} \frac{I_{G^{[s]}}(\delta_0, \mathbf{D}_0)}{I_{G'}(\delta_0, \mathbf{D}_0)} \frac{|\text{nb}d_p^-(G^{[s]})|}{|\text{nb}d_p^+(G')|} \times \quad (9)$$

$$\times \exp\left\{-\frac{1}{2} \left[\langle \mathbf{K}' - \mathbf{K}^{[s]}, \mathbf{U} + \mathbf{D}_0 \rangle + \frac{(\Psi'_{i_0 j_0} - \Psi_{i_0 j_0}^{[s]})^2}{\sigma_g^2} \right]\right\}.$$

We denote by $(\mathbf{K}^{\lceil s+1/2 \rceil}, G^{\lceil s+1 \rceil})$, $\mathbf{K}^{\lceil s+1/2 \rceil} \in G^{\lceil s+1 \rceil}$, the state of the chain at the end of this step.

Step 2: Resample the precision matrix. Given the updated graph $G^{\lceil s+1 \rceil}$, we update the precision matrix $\mathbf{K}^{\lceil s+1/2 \rceil} = (\mathbf{Q}^*)^T (\boldsymbol{\Psi}^{s+1/2})^T \boldsymbol{\Psi}^{s+1/2} \mathbf{Q}^*$ by sequentially perturbing the free elements $(\boldsymbol{\Psi}^{\lceil s+1/2 \rceil})^{v(G^{\lceil s+1 \rceil})}$. For each such element, we perform the corresponding Metropolis-Hastings step from Section 2.2 with $\delta = n + \delta_0$, $\mathbf{D} = \mathbf{U} + \mathbf{D}_0$ and $\mathbf{Q} = \mathbf{Q}^*$. The standard deviation of the normal proposals is σ_m . We denote by $\mathbf{K}^{\lceil s+1 \rceil} \in \mathcal{P}_{G^{\lceil s+1 \rceil}}$ the precision matrix obtained after all the updates have been performed.

3 Matrix-variate Gaussian graphical models

We extend our framework to the case when the observed data $\mathcal{D} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$ are associated with a $p_R \times p_C$ random matrix $\mathbf{X} = (X_{ij})$ that follows a matrix-variate normal distribution

$$\text{vec}(\mathbf{X}^T) | \mathbf{K}_R, \mathbf{K}_C \sim \mathcal{N}_{p_R p_C}(\mathbf{0}, [\mathbf{K}_R \otimes \mathbf{K}_C]^{-1})$$

with p.d.f. (Gupta & Nagar, 2000):

$$p(\mathbf{X} | \mathbf{K}_R, \mathbf{K}_C) = (2\pi)^{-p_R p_C / 2} (\det \mathbf{K}_R)^{p_C / 2} (\det \mathbf{K}_C)^{p_R / 2} \exp \left\{ -\frac{1}{2} \text{tr} [\mathbf{K}_R \mathbf{X} \mathbf{K}_C \mathbf{X}^T] \right\}. \quad (10)$$

Here \mathbf{K}_R is a $p_R \times p_R$ row precision matrix and \mathbf{K}_C is a $p_C \times p_C$ column precision matrix. Furthermore, we assume that $\mathbf{K}_R \in \mathcal{P}_{G_R}$ and $\mathbf{K}_C \in \mathcal{P}_{G_C}$ where $G_R = (V_{p_R}, E_R)$ and $G_C = (V_{p_C}, E_C)$ are two graphs with p_R and p_C vertices, respectively. We consider the rows $\mathbf{X}_{1*}, \dots, \mathbf{X}_{p_R*}$ and the columns $\mathbf{X}_{*1}, \dots, \mathbf{X}_{*p_C}$ of the random matrix \mathbf{X} . From Theorem 2.3.12 of Gupta & Nagar (2000) we have $\mathbf{X}_{i*}^T \sim \mathcal{N}_{p_C}(\mathbf{0}, (\mathbf{K}_R^{-1})_{ii} \mathbf{K}_C^{-1})$ and $\mathbf{X}_{*j} \sim \mathcal{N}_{p_R}(\mathbf{0}, (\mathbf{K}_C^{-1})_{jj} \mathbf{K}_R^{-1})$. The graphs G_R and G_C define GGMs for the rows and columns of \mathbf{X} (Wang & West, 2009):

$$\begin{aligned} \mathbf{X}_{i_1*} \perp\!\!\!\perp \mathbf{X}_{i_2*} \mid \mathbf{X}_{(V_{p_R} \setminus \{i_1, i_2\})^*} &\Leftrightarrow (\mathbf{K}_R)_{i_1 i_2} = (\mathbf{K}_R)_{i_2 i_1} = 0 \Leftrightarrow (i_1, i_2) \notin E_R, \text{ and} \\ \mathbf{X}_{*j_1} \perp\!\!\!\perp \mathbf{X}_{*j_2} \mid \mathbf{X}_{(V_{p_C} \setminus \{j_1, j_2\})^*} &\Leftrightarrow (\mathbf{K}_C)_{j_1 j_2} = (\mathbf{K}_C)_{j_2 j_1} = 0 \Leftrightarrow (j_1, j_2) \notin E_C. \end{aligned} \quad (11)$$

Any prior specification for \mathbf{K}_R and \mathbf{K}_C must take into account the fact that the two precision matrices are not uniquely identified from their Kronecker product which means that, for any $z > 0$, $(z^{-1} \mathbf{K}_R) \otimes (z \mathbf{K}_C) = \mathbf{K}_R \otimes \mathbf{K}_C$ represents the same precision matrix for $\text{vec}(\mathbf{X}^T)$ – see equation (10). We follow the basic idea laid out in Wang & West (2009) and impose the constraint $(\mathbf{K}_C)_{11} = 1$. Furthermore, we define a prior for \mathbf{K}_C through parameter expansion by assuming a G-Wishart prior $\text{Wis}_{G_C}(\delta_C, \mathbf{D}_C)$ for the matrix $z \mathbf{K}_C$ with $z > 0$, $\delta_C > 2$ and $\mathbf{D}_C \in \mathcal{P}_{G_C}$. It is immediate to see that the Jacobean of the transformation from $z \mathbf{K}_C$ to (z, \mathbf{K}_C) is

$$J((z \mathbf{K}_C) \rightarrow (z, \mathbf{K}_C)) = z^{|\nu(G_C)|-1}.$$

It follows that our joint prior for (z, \mathbf{K}_C) is given by

$$p(z, \mathbf{K}_C | G_C, \delta_C, \mathbf{D}_C) = \frac{1}{I_{G_C}(\delta_C, \mathbf{D}_C)} (\det \mathbf{K}_C)^{\frac{\delta_C - 2}{2}} \exp \left\{ -\frac{1}{2} \langle \mathbf{K}_C, z \mathbf{D}_C \rangle \right\} z^{\frac{p_C(\delta_C - 2)}{2} + |\nu(G_C)| - 1}.$$

The elements of $\mathbf{K}_R \in \mathbf{P}_{G_R}$ are not subject to any additional constraints, hence we assume a G-Wishart prior $\text{Wis}_{G_R}(\delta_R, \mathbf{D}_R)$ for \mathbf{K}_R . We take $\delta_C = \delta_R = 3$, $\mathbf{D}_C = \mathbf{I}_{p_C}$ and $\mathbf{D}_R = \mathbf{I}_{p_R}$. We complete our prior specification by assuming uniform priors $p(G_C) = 2/p_C(p_C-1)$ and $p(G_R) = 2/p_R(p_R-1)$ for the row and column graphs, where $G_R \in \mathcal{G}_{p_R}$ and $G_C \in \mathcal{G}_{p_C}$.

We perform Bayesian inference for matrix-variate GGMs by developing a MCMC algorithm for sampling from the joint distribution of the data, row and column precision matrices and row and column graphs

$$p(\mathcal{D}, \mathbf{K}_R, z, \mathbf{K}_C, G_R, G_C | \delta_C, \mathbf{D}_C, \delta_R, \mathbf{D}_R) \propto p(\mathcal{D} | \mathbf{K}_R, \mathbf{K}_C) \times \quad (12)$$

$$\times p(\mathbf{K}_R | G_R, \delta_R, \mathbf{D}_R) p(z, \mathbf{K}_C | G_C, \delta_C, \mathbf{D}_C) p(G_C) p(G_R),$$

that is well-defined for $\mathbf{K}_R \in \mathbf{P}_{G_R}$, $\mathbf{K}_C \in \mathbf{P}_{G_C}$ with $(\mathbf{K}_C)_{11} = 1$ and $z > 0$. Equation (12) is written as:

$$p(\mathcal{D}, \mathbf{K}_R, z, \mathbf{K}_C, G_R, G_C | \delta_C, \mathbf{D}_C, \delta_R, \mathbf{D}_R) \propto z^{\frac{p_C(\delta_C-2)}{2} + |\mathcal{Y}(G_C)|-1} (\det \mathbf{K}_R)^{\frac{np_C + \delta_R - 2}{2}} (\det \mathbf{K}_C)^{\frac{np_R + \delta_C - 2}{2}} \times \quad (13)$$

$$\times \exp \left\{ -\frac{1}{2} \text{tr} \left[\sum_{j=1}^n \mathbf{K}_R \mathbf{x}^{(j)} \mathbf{K}_C (\mathbf{x}^{(j)})^T + \mathbf{K}_R \mathbf{D}_R + \mathbf{K}_C (z \mathbf{D}_C) \right] \right\}.$$

Our sampling scheme is comprised of the following five steps that explain the transition of the Markov chain from its current state $(\mathbf{K}_R^{[s]}, \mathbf{K}_C^{[s]}, G_R^{[s]}, z^{[s]})$ to its next state $(\mathbf{K}_R^{[s+1]}, \mathbf{K}_C^{[s+1]}, G_R^{[s+1]}, z^{[s+1]})$. We use four strictly positive precision parameters $\sigma_{m,R}$, $\sigma_{m,C}$, $\sigma_{g,R}$ and $\sigma_{g,C}$.

Step 1: Resample the row graph. We denote $n_R^* = np_C$ and $\mathbf{U}_R^* = \sum_{j=1}^n \mathbf{x}^{(j)} \mathbf{K}_C^{[s]} (\mathbf{x}^{(j)})^T$. We generate a random permutation $\nu_R \in \Upsilon_{p_R}$ of the row indices V_{p_R} and reorder the row and columns of the matrix $\mathbf{U}_R^* + \mathbf{D}_R$ according to ν_R . We determine the Cholesky decomposition $(\mathbf{U}_R^* + \mathbf{D}_R)^{-1} = (\mathbf{Q}_R^*)^T \mathbf{Q}_R^*$. We proceed as described in Step 1 of Section 2.2. Given the notations we used in that section, we take $p = p_R$, $n = n_R^*$, $\mathbf{U} = \mathbf{U}_R^*$, $\delta_0 = \delta_R$, $\mathbf{D}_0 = \mathbf{D}_R$ and $\sigma_g = \sigma_{g,R}$. We denote the updated row precision matrix and graph by $(\mathbf{K}_R^{[s+1/2]}, G_R^{[s+1]})$, $\mathbf{K}_R^{[s+1/2]} \in \mathbf{P}_{G_R^{[s+1]}}$.

Step 2: Resample the row precision matrix. We denote $n_R^* = np_C$ and $\mathbf{U}_R^* = \sum_{j=1}^n \mathbf{x}^{(j)} \mathbf{K}_C^{[s]} (\mathbf{x}^{(j)})^T$.

We determine the Cholesky decomposition $(\mathbf{U}_R^* + \mathbf{D}_R)^{-1} = (\mathbf{Q}_R^*)^T \mathbf{Q}_R^*$ after permuting the row and columns of $\mathbf{U}_R^* + \mathbf{D}_R$ according to a random ordering in Υ_{p_R} . The conditional distribution of $\mathbf{K}_R \in \mathbf{P}_{G_R^{[s+1]}}$ is G-Wishart $\text{Wis}_{G_R^{[s+1]}}(n_R^* + \delta_R, \mathbf{U}_R^* + \mathbf{D}_R)$. We make the transition from $\mathbf{K}_R^{[s+1/2]} \in \mathbf{P}_{G_R^{[s+1]}}$ to $\mathbf{K}_R^{[s+1]} \in \mathbf{P}_{G_R^{[s+1]}}$ using Metropolis-Hastings updates described in Section 2.1. Given the notations we used in that section, we take $p = p_R$, $\delta = n_R^* + \delta_R$, $\mathbf{D} = \mathbf{U}_R^* + \mathbf{D}_R$, $\mathbf{Q} = \mathbf{Q}_R^*$ and $\sigma_m = \sigma_{m,R}$.

Step 3: Resample the column graph. We denote $n_C^* = np_R$ and $\mathbf{U}_C^* = \sum_{j=1}^n (\mathbf{x}^{(j)})^T \mathbf{K}_R^{[s+1]} \mathbf{x}^{(j)}$. The

relevant conditional distribution is (see equation (13)):

$$p\left(\mathbf{K}_C^{[s]}, G_C^{[s]}, z^{[s]} | \mathcal{D}, \mathbf{K}_R^{[s+1]}, \delta_C, \mathbf{D}_C\right) \propto \frac{1}{I_{G_C^{[s]}}(\delta_C, \mathbf{D}_C)} \left(z^{[s]}\right)^{|\nu(G_C)|} \left(\det \mathbf{K}_C^{[s]}\right)^{\frac{n_C^* + \delta_C - 2}{2}} \times \quad (14)$$

$$\times \exp\left\{-\frac{1}{2} \langle \mathbf{K}_C^{[s]}, \mathbf{U}_C^* + z^{[s]} \mathbf{D}_C \rangle\right\}.$$

We sample a candidate column graph $G'_C \in \text{nb}_{p_C}(G_C^{[s]})$ from the proposal

$$q\left(G'_C | G_C^{[s]}, z^{[s]}\right) = \frac{1}{2} \frac{\left(z^{[s]}\right)^{|\nu(G'_C)|}}{\sum_{G''_C \in \text{nb}_{p_C}^+(G_C^{[s]})} \left(z^{[s]}\right)^{|\nu(G''_C)|}} \mathbf{1}_{\{G'_C \in \text{nb}_{p_C}^+(G_C^{[s]})\}} + \frac{1}{2} \frac{\left(z^{[s]}\right)^{|\nu(G'_C)|}}{\sum_{G''_C \in \text{nb}_{p_C}^-(G_C^{[s]})} \left(z^{[s]}\right)^{|\nu(G''_C)|}} \mathbf{1}_{\{G'_C \in \text{nb}_{p_C}^-(G_C^{[s]})\}}, \quad (15)$$

where $\mathbf{1}_A$ is equal to 1 if A is true and is 0 otherwise. The proposal (15) gives an equal probability that the candidate graph is obtained by adding or deleting an edge from the current graph – see also equation (6).

We assume that G'_C is obtained by adding the edge (i_0, j_0) to $G_C^{[s]}$. We generate a random permutation $\nu_C \in \Upsilon_{p_C}^{(1,1)}$ of the row indices V_{p_C} and reorder the row and columns of the matrix $\mathbf{U}_C^* + z^{[s]} \mathbf{D}_C$ according to ν_C . The permutation ν_C is such that $\nu_C(1) = 1$, hence the $(1, 1)$ element of $\mathbf{K}_C^{[s]}$ remains in the same position. We determine the Cholesky decomposition $(\mathbf{U}_C^* + z^{[s]} \mathbf{D}_C)^{-1} = (\mathbf{Q}_C^*)^T \mathbf{Q}_C^*$ of $(\mathbf{U}_C^* + z^{[s]} \mathbf{D}_C)^{-1}$. We consider the decomposition of the column precision matrix

$$\mathbf{K}_C^{[s]} = (\mathbf{Q}_C^*)^T (\boldsymbol{\Psi}_C^{[s]})^T \boldsymbol{\Psi}_C^{[s]} \mathbf{Q}_C^*,$$

with $(\boldsymbol{\Psi}_C^{[s]})^{\nu(G_C^{[s]})} \in M^{\nu(G_C^{[s]})}$. We define an upper triangular matrix $\boldsymbol{\Psi}'_C$ such that $(\boldsymbol{\Psi}'_C)_{ij} = (\boldsymbol{\Psi}_C^{[s]})_{ij}$ for $(i, j) \in \nu(G_C^{[s]})$. We sample $\gamma \sim \mathcal{N}\left((\boldsymbol{\Psi}_C^{[s]})_{i_0 j_0}, \sigma_{g, C}^2\right)$ and set $(\boldsymbol{\Psi}'_C)_{i_0 j_0} = \gamma$. The rest of the elements of $\boldsymbol{\Psi}'_C$ are determined from $(\boldsymbol{\Psi}'_C)^{\nu(G'_C)}$ through the completion operation. We consider the candidate column precision matrix

$$\mathbf{K}'_C = (\mathbf{Q}_C^*)^T (\boldsymbol{\Psi}'_C)^T \boldsymbol{\Psi}'_C \mathbf{Q}_C^*. \quad (16)$$

We know that $\mathbf{K}_C^{[s]} \in \mathbf{P}_{G_C^{[s]}}$ must satisfy $(\mathbf{K}_C^{[s]})_{11} = 1$. The last equality implies $(\boldsymbol{\Psi}_C^{[s]})_{11} = 1 / (\mathbf{Q}_C^*)_{11}$, hence $(\boldsymbol{\Psi}'_C)_{11} = 1 / (\mathbf{Q}_C^*)_{11}$. Therefore we have $\mathbf{K}'_C \in \mathbf{P}_{G'_C}$ and $(\mathbf{K}'_C)_{11} = 1$.

We make the transition from $(\mathbf{K}_C^{[s]}, G_C^{[s]})$ to (\mathbf{K}'_C, G'_C) with probability $\min\{R_C^+, 1\}$ where

$$R_C^+ = \sigma_{g,C} \sqrt{2\pi} (\mathbf{Q}_C^*)_{i_0 i_0} (\mathbf{Q}_C^*)_{j_0 j_0} (\Psi_C^{[s]})_{i_0 i_0} \frac{I_{G_C^{[s]}}(\delta_C, \mathbf{D}_C) \sum_{G'_C \in \text{enbd}_{p_C}^+(G_C^{[s]})} (z^{[s]})^{|\nu(G'_C)|}}{I_{G'_C}(\delta_C, \mathbf{D}_C) \sum_{G'_C \in \text{enbd}_{p_C}^-(G'_C)} (z^{[s]})^{|\nu(G'_C)|}} \times \quad (17)$$

$$\times \exp \left\{ -\frac{1}{2} \left[\langle \mathbf{K}'_C - \mathbf{K}_C^{[s]}, \mathbf{U}_C^* + z^{[s]} \mathbf{D}_C \rangle - \frac{\left((\Psi'_C)_{i_0 j_0} - (\Psi_C^{[s]})_{i_0 j_0} \right)^2}{\sigma_{g,C}^2} \right] \right\}.$$

Next we assume that G'_C is obtained by deleting the edge (i_0, j_0) from $G_C^{[s]}$. We define an upper triangular matrix Ψ'_C such that $(\Psi'_C)_{ij} = (\Psi_C^{[s]})_{ij}$ for $(i, j) \in \nu(G'_C)$. The candidate \mathbf{K}'_C is obtained from Ψ'_C as in equation (16). We make the transition from $(\mathbf{K}_C^{[s]}, G_C^{[s]})$ to (\mathbf{K}'_C, G'_C) with probability $\min\{R_C^-, 1\}$ where

$$R_C^- = \left\{ \sigma_{g,C} \sqrt{2\pi} (\mathbf{Q}_C^*)_{i_0 i_0} (\mathbf{Q}_C^*)_{j_0 j_0} (\Psi_C^{[s]})_{i_0 i_0} \right\}^{-1} \frac{I_{G_C^{[s]}}(\delta_C, \mathbf{D}_C) \sum_{G'_C \in \text{enbd}_{p_C}^-(G_C^{[s]})} (z^{[s]})^{|\nu(G'_C)|}}{I_{G'_C}(\delta_C, \mathbf{D}_C) \sum_{G'_C \in \text{enbd}_{p_C}^+(G'_C)} (z^{[s]})^{|\nu(G'_C)|}} \times \quad (18)$$

$$\times \exp \left\{ -\frac{1}{2} \left[\langle \mathbf{K}'_C - \mathbf{K}_C^{[s]}, \mathbf{U}_C^* + z^{[s]} \mathbf{D}_C \rangle + \frac{\left((\Psi'_C)_{i_0 j_0} - (\Psi_C^{[s]})_{i_0 j_0} \right)^2}{\sigma_{g,C}^2} \right] \right\}.$$

We denote the updated column precision matrix and graph by $(\mathbf{K}_C^{[s+1/2]}, G_C^{[s+1]})$.

Step 4: Resample the column precision matrix. We denote $n_C^* = np_R$ and $\mathbf{U}_C^* = \sum_{j=1}^n (\mathbf{x}^{(j)})^T \mathbf{K}_R^{[s+1]} \mathbf{x}^{(j)}$.

We determine the Cholesky decomposition $(\mathbf{U}_C^* + z^{[s]} \mathbf{D}_C)^{-1} = (\mathbf{Q}_C^*)^T \mathbf{Q}_C^*$ after permuting the row and columns of $\mathbf{U}_C^* + z^{[s]} \mathbf{D}_C$ according to a random ordering in $\Upsilon_{p_C}^{(1,1)}$. The conditional distribution of $\mathbf{K}_C \in \mathcal{P}_{G_C^{[s+1]}}$ with $(\mathbf{K}_C)_{11} = 1$ is G-Wishart $\text{Wis}_{G_C^{[s+1]}}(n_C^* + \delta_C, \mathbf{U}_C^* + z^{[s]} \mathbf{D}_C)$. We make the transition from $\mathbf{K}_C^{[s+1/2]}$ to $\mathbf{K}_C^{[s+1]} \in \mathcal{P}_{G_C^{[s+1]}}$ using Metropolis-Hastings updates from Section 2.1. Given the notations we used in that section, we take $p = p_C$, $\delta = n_C^* + \delta_C$, $\mathbf{D} = \mathbf{U}_C^* + z^{[s]} \mathbf{D}_C$, $\mathbf{Q} = \mathbf{Q}_C^*$ and $\sigma_m = \sigma_{m,C}$. The constraint $(\mathbf{K}_C)_{11} = 1$ is accommodated as described at the end of Section 2.1.

Step 5: Resample the auxiliary variable. The conditional distribution of $z > 0$ is

$$\text{Gamma} \left(\frac{p_C(\delta_C - 2)}{2} + |\nu(G_C^{[s+1]})|, \frac{1}{2} \text{tr}(\mathbf{K}_C^{[s+1]} \mathbf{D}_C) \right). \quad (19)$$

Here $\text{Gamma}(\alpha, \beta)$ has density $f(x|\alpha, \beta) \propto \beta^\alpha x^{\alpha-1} \exp(-\beta x)$. We sample from (19) to obtain $z^{[s+1]}$.

4 Sparse models for lattice data

Conditional autoregressive (CAR) models (Besag, 1974; Mardia, 1988) are routinely used in spatial statistics to model lattice data. In the case of a single observed outcome in each region, the data is associated with a vector $\mathbf{X}^T = (X_1, \dots, X_{p_R})^T$ where X_i corresponds to region i . The zero-centered CAR model is implicitly defined by the set of full conditional distributions

$$X_i | \{X_{i'} : i' \neq i\} \sim \mathcal{N} \left(\sum_{i' \neq i} b_{ii'} X_{i'}, \lambda_i^2 \right), \quad i = 1, \dots, p_R.$$

Therefore, CAR models are just two-dimensional Gaussian Markov random fields. According to Brook's (1964) theorem, this set of full-conditional distributions implies that the joint distribution for \mathbf{X} satisfies

$$p(\mathbf{X} | \boldsymbol{\lambda}) \propto \exp \left\{ -\frac{1}{2} \mathbf{X}^T \boldsymbol{\Lambda}^{-1} (\mathbf{I} - \mathbf{B}) \mathbf{X} \right\},$$

where $\boldsymbol{\Lambda} = \text{diag}\{\lambda_1^2, \dots, \lambda_{p_R}^2\}$ and \mathbf{B} is a $p_R \times p_R$ matrix such that $\mathbf{B} = (b_{ii'})$ and $b_{ii} = 0$. In order for $\boldsymbol{\Lambda}^{-1}(\mathbf{I} - \mathbf{B})$ to be a symmetric matrix we require that $b_{ii'} \lambda_{i'}^2 = b_{i'i} \lambda_i^2$ for $i \neq i'$; therefore the matrix \mathbf{B} and vector $\boldsymbol{\lambda}$ must be carefully chosen. A popular approach is to begin by constructing a symmetric proximity matrix $\mathbf{W} = (w_{ii'})$, and then set $b_{ii'} = w_{ii'}/w_{i+}$ and $\lambda_i^2 = \tau^2/w_{i+}$ where $w_{i+} = \sum_{i'} w_{ii'}$ and $\tau^2 > 0$. In that case, $\boldsymbol{\Lambda}^{-1}(\mathbf{I} - \mathbf{B}) = \tau^{-2}(\mathbf{E}_W - \mathbf{W})$, where $\mathbf{E}_W = \text{diag}\{w_{1+}, \dots, w_{p_R+}\}$. The proximity matrix \mathbf{W} is often constructed by first specifying a neighborhood structure for the geographical areas under study; for example, when modeling state or county level data it is often assumed that two geographical units are neighbors if they share a common border. Given that neighborhood structure, the proximity matrix is specified as

$$w_{ii'} = \begin{cases} 1, & \text{if } i' \in \partial i, \\ 0, & \text{otherwise,} \end{cases} \quad (20)$$

where ∂i denotes the set of neighbors of i .

Specifying the joint precision matrix for \mathbf{X} using the proximity matrix derived from the neighborhood structure is very natural; essentially, it implies that observations collected on regions that are not neighbors are conditionally independent from each other given the rest. However, note that the specification in (20) implies that $(\mathbf{E}_W - \mathbf{W})\mathbf{1}_{p_R} = \mathbf{0}$ and therefore the joint distribution on \mathbf{X} is improper. Proper CAR models Cressie (1973); Sun et al. (2000); Gelfand & Vounatsou (2003) can be obtained by including a spatial autocorrelation parameter ρ , so that

$$X_i | \{X_{i'} : i' \neq i\} \sim \mathcal{N} \left(\rho \sum_{i' \neq i} \frac{w_{ii'}}{w_{i+}} X_{i'}, \frac{\tau^2}{w_{i+}} \right), \quad (21)$$

The joint distribution on \mathbf{X} is then multivariate normal $\mathcal{N}_{p_R}(\mathbf{0}, \mathbf{D}_W^{-1})$ where $\mathbf{D}_W = \tau^2(\mathbf{E}_W - \rho\mathbf{W})^{-1}$. This distribution is proper as long as ρ is between the reciprocals of the minimum and maximum eigenvalues for \mathbf{W} . In particular, note that taking $\rho = 0$ leads to independent random effects.

In the spirit of Besag & Kooperberg (1995), an alternative but related approach to the construction of models for lattice data is to let $\mathbf{X} \sim \mathcal{N}_{p_R}(\mathbf{0}, \mathbf{K}^{-1})$ and assign \mathbf{K} a G-Wishart prior $\text{Wis}_{G_W}(\delta, (\delta - 2)\mathbf{D}_W)$ where the graph G_W is implied by the neighborhood matrix \mathbf{W} defined in (20). Following Lenkoski & Dobra (2010), the mode of $\text{Wis}_{G_W}(\delta, (\delta - 2)\mathbf{D}_W)$ is the unique positive definite matrix \mathbf{K} that satisfies the relations

$$\mathbf{K}_{i' i'}^{-1} = (\mathbf{D}_W)_{i' i'}, \text{ if } i' \in \partial i, \text{ and } \mathbf{K}_{i' i'} = 0, \text{ if } i' \notin \partial i. \quad (22)$$

The matrix \mathbf{D}_W^{-1} verifies the system (22), hence it is the mode of $\text{Wis}_{G_W}(\delta, (\delta - 2)\mathbf{D}_W)$. As such, the mode of the prior for \mathbf{K} induces the same prior specification for \mathbf{X} as (21).

It is easy to see that, conditional on $\mathbf{K} \in \mathcal{P}_{G_W}$, we have

$$X_i | \{X_{i'} : i' \neq i\} \sim \mathcal{N}\left(-\sum_{i' \in \partial i} \frac{K_{i' i}}{K_{ii}} X_{i'}, \frac{1}{K_{ii}}\right). \quad (23)$$

Hence, by modeling \mathbf{X} using a Gaussian graphical model and restricting the precision matrix \mathbf{K} to belong to the cone \mathcal{P}_{G_W} , we are inducing a mixture of CAR priors on \mathbf{X} where the priors on

$$b_{i' i} = \begin{cases} -K_{i' i}/K_{ii}, & \text{if } i' \in \partial i, \\ 0, & \text{otherwise,} \end{cases}$$

and $\lambda_i^2 = 1/K_{ii}$ are induced by the G-Wishart prior $\text{Wis}_{G_W}(\delta, (\delta - 2)\mathbf{D}_W)$.

The specification of CAR models through G-Wishart priors solves the impropriety problem and preserves the computational advantages derived from standard CAR specifications while providing greater flexibility. Indeed, the prior is trivially proper because the matrix $\mathbf{K} \in \mathcal{P}_{G_W}$ is invertible by construction. The computational advantages are preserved because the full conditional distributions for each X_i can be easily computed for any matrix \mathbf{K} without the need to perform matrix inversion, and they depend only on a small subset of neighbors $\{X_j : j \in \partial i\}$. Additional flexibility is provided because the weights $b_{i' i}$ for $i' \in \partial i$ and smoothing parameters λ_i^2 are being estimated from the data rather than being assumed fixed, allowing for adaptive spatial smoothing. Indeed, our approach provides what can be considered as a nonparametric alternative to the parametric estimates of the proximity matrix proposed by Cressie & Chan (1989).

A similar approach can be used to construct proper multivariate conditional autoregressive (MCAR) models (Mardia, 1988; Gelfand & Vounatsou, 2003). In this case, we are interested in modeling a $p_R \times p_C$ matrix $\mathbf{X} = (X_{ij})$ where X_{ij} denotes the value of the j -th outcome in region i . We let \mathbf{X} follow a matrix-variate normal distribution with row precision matrix \mathbf{K}_R capturing the spatial structure in the data (which, as in univariate CAR models, is restricted to the cone \mathcal{P}_{G_W} defined by the neighborhood matrix \mathbf{W}), and column precision matrix \mathbf{K}_C , which controls the multivariate dependencies across outcomes. It can be easily shown that the row vector \mathbf{X}_{i*} of \mathbf{X} depends only on the row vectors associated with those regions that are neighbors with i — see also equation (11):

$$\mathbf{X}_{i*}^T | \{\mathbf{X}_{i'*}^T : i' \neq i\} \sim \mathcal{N}_{p_C}\left(-\sum_{i' \in \partial i} \frac{(\mathbf{K}_R)_{i' i}}{(\mathbf{K}_R)_{ii}} \mathbf{X}_{i'*}^T, \frac{1}{(\mathbf{K}_R)_{ii}} \mathbf{K}_C^{-1}\right).$$

This formulation for spatial models can also be used as part of more complex hierarchical models. Indeed, CAR and MCAR models are most often used as a prior for the random effects of a generalized linear model (GLM) to account for residual spatial structure not accounted for by covariates. When no covariates are available, the model can be interpreted as a spatial smoother where the spatial covariance matrix controls the level of spatial smoothing in the underlying (latent) surface. Similarly, MCAR models can be used to construct sparse multivariate spatial GLMs.

As an example, consider the $p_R \times p_C$ matrix $\mathbf{Y} = (Y_{ij})$ of discrete or continuous outcomes, and let $Y_{ij} \sim h_j(\cdot|\eta_{ij})$ where h_j is a probability mass or probability density function that belongs to the exponential family with location parameter η_{ij} . The spatial GLM is then defined through the linear predictor

$$g^{-1}(\eta_{ij}) = \mu_j + X_{ij} + \mathbf{Z}_{ij}\boldsymbol{\beta}_j,$$

where g is the link function, μ_j is an outcome-specific intercept, X_{ij} is a zero-centered spatial random effect associated with location i , \mathbf{Z}_{ij} is a matrix of observed covariates for outcome j at location i , and $\boldsymbol{\beta}_j$ is the vector of fixed effects associated with outcome j . We could further take $y_{ij} \sim \text{Poi}(\eta_{ij})$, $g^{-1}(\cdot) = \log(\cdot)$ and assign a sparse matrix-variate normal distribution for \mathbf{X} . This choice leads to a sparse multivariate spatial log-linear model for count data, which is often used for disease mapping in epidemiology (see Section 5.3).

Due to the structure associated with our centering matrix \mathbf{D}_R , and the presence of the expanded scaling parameter z , prior elicitation for this class of spatial models is relatively straightforward. Indeed, elicitation of the matrix \mathbf{D}_R requires only the elicitation of the neighborhood matrix \mathbf{W} , along with reasonable values for the spatial autocorrelation parameter ρ and scale τ^2 . In particular, in the applications we discuss below we assume that, *a priori*, there is a strong degree of positive spatial association, and set $\rho = 0.99$. Also, and as is customary in the literature on spatial models for lattice data, τ^2 is selected to roughly reflect the scale of data being modeled. In the case of MCAR models, it is common to assume that the prior value for the conditional covariance between variables is zero, which leads to choosing a diagonal \mathbf{D}_C . When the scale of all the outcome variables is the similar (as in the applications below), it is then sensible to pick $\mathbf{D}_C = \mathbf{I}_{p_C}$, as the actual scale is irrelevant.

5 Illustrations

5.1 Simulation study

We consider a simulation study involving matrix-variate normal graphical models where $p_R = 5$ and $p_C = 10$. We sample 100 observations of matrices with row graph G_R defined by the edges $\{(1, i) : 1 < i \leq 5\} \cup \{(2, 3), (3, 4), (4, 5), (2, 5)\}$ and column graph C_{10} . Here C_{10} denotes the 10-dimensional cycle graph with edges $\{(i, i + 1) : 1 \leq i \leq 9\} \cup \{(1, 10)\}$. We note that both the row and column graphs are non-decomposable, with the row graph being relatively dense, while column graph is relatively sparse. In this case $\mathbf{K}_R \in \mathbf{P}_{G_R}$ is set so that $(\mathbf{K}_R)_{ii} = 1$, $(\mathbf{K}_R)_{i,j} = 0.4$ for $(i, j) \in G_R$. $\mathbf{K}_C \in \mathbf{P}_{C_{10}}$ is defined in a similar manner. To generate the 100 samples from this matrix-variate distribution, we sample $\text{vec}(\mathbf{X}^T) \sim \mathbf{N}_{p_R p_C}(\mathbf{0}, [\mathbf{K}_R \otimes \mathbf{K}_C]^{-1})$ and rearrange this vector to obtain the appropriate matrices.

Table 1: Results of simulation study in estimates for the row graph and \mathbf{K}_R . In the matrix below, the values along the diagonal and in the upper triangle correspond to the average of the estimates of entries in \mathbf{K} across the 100 repetitions, while the values below the diagonal correspond to the average of the edge probabilities across the 100 iterations.

0	1	2	3	4	5
1	1.107	0.442	0.438	0.436	0.44
2	1	1.105	0.439	0	0.441
3	1	1	1.1	0.437	0
4	1	0.026	1	1.086	0.438
5	1	1	0.04	1	1.096

We ran the matrix-variate GGM search described in Section 3 for 50,000 iterations after a burn in of 5,000 iterations. We set $\sigma_{m,R}$, $\sigma_{m,C}$, $\sigma_{g,R}$ and $\sigma_{g,C}$ all to the value .5, which achieved an average rejection rate on updates of elements for both \mathbf{K}_C and \mathbf{K}_R of about 30 percent. We repeated the study 100 times to reduce sampling variability. Tables 1 and 2 show the resulting estimates for the row and column GGMs respectively. We see that the matrix-variate GGM search outlined above performs extremely well at recovering the structure of both the column and row graphs as well as the values of both \mathbf{K}_C and \mathbf{K}_R . Since both row and column GGMs are defined by non-decomposable graphs, we have illustrated that our methodology is capable of searching the entire graph space of matrix-variate GGMs.

5.2 Modeling statewide SAT scores

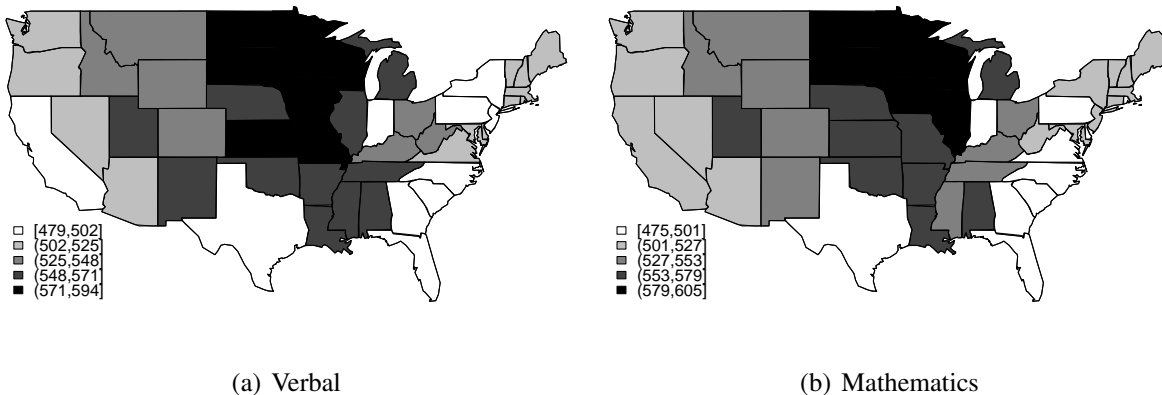
This section develops a model for the scores on the SAT college entrance exam for the year 1999. The data consists of statewide average verbal and mathematics SAT scores, along with the percentage of eligible students taking the test, for each of the 48 contiguous states plus the District of Columbia. The verbal scores were originally analyzed by Wall (2004), who used them to illustrate the differences among different models for lattice data. Figure 1 shows choropleth maps of the two scores; in both cases there is strong evidence of spatial association (note the clear clusters of midwestern, central, western and southeastern states).

In this section we develop a bivariate spatial regression where both verbal and mathematics scores are treated as response variables, while the percentage of eligible students are treated as the explanatory variable. Figure 2 presents pairwise scatterplots of all three variables. They demonstrate that both verbal and mathematics scores are positively associated with each other, and negatively associated with the percentage of eligible students who take the test. This association between SAT scores and the percentage of eligible students could be explained by selection bias: in states where fewer students take the exam, it is typically only the best qualified who do, which in turn results in better averages than those from states who provide a wider access to testing. Hence,

Table 2: Results of simulation study for estimates of the column graph and \mathbf{K}_C . In the matrix below, the values along the diagonal and above it to the average of the estimates of entries in \mathbf{K}_C across the 100 repetitions, while the values below the diagonal correspond to the average of the edge probabilities across the 100 iterations.

	1	2	3	4	5	6	7	8	9	10
1	1	0.395	-0.001	0.001	0	-0.001	0	0.001	-0.001	0.396
2	1	0.941	0.364	0	0.001	0	0	0.001	0.002	0.001
3	0.024	1	0.916	0.365	0	0.001	-0.001	0	0	0
4	0.046	0.032	1	0.932	0.37	-0.001	0	0	0.001	-0.001
5	0.022	0.046	0.033	1	0.936	0.371	-0.001	0	-0.001	0.001
6	0.04	0.038	0.036	0.034	1	0.927	0.362	-0.001	0	-0.001
7	0.037	0.03	0.052	0.025	0.042	1	0.928	0.366	0	0
8	0.047	0.046	0.024	0.053	0.032	0.049	1	0.935	0.377	0
9	0.072	0.082	0.039	0.029	0.036	0.068	0.039	1	0.93	0.364
10	1	0.057	0.037	0.03	0.047	0.044	0.035	0.066	1	0.946

Figure 1: SAT scores in the 48 contiguous states. Left panel shows verbal scores, while the right panel shows mathematics scores.



our model attempts to jointly understand the behavior of the scores adjusting for this selection bias.

More specifically, we let Y_{ij} be the score on test j ($j = 1$ for verbal and $j = 2$ for math) and state $i = 1, \dots, 49$. We use the notation from Section 3 and take $p_R = 49$, $p_C = 2$. The observed scores form a $p_R \times p_C$ matrix $\mathbf{Y} = (Y_{ij})$. Since Figure 2 suggests a quadratic relationship between scores and the percentage of eligible students, we let

$$Y_{ij} = \beta_{0j} + \beta_{1j}Z_i + \beta_{2j}Z_i^2 + X_{ij},$$

where Z_i represents the percentage of eligible students taking the exam and X_{ij} is an error term. The vectors $\boldsymbol{\beta}_j = (\beta_{0j}, \beta_{1j}, \beta_{2j})^T$, $j = 1, 2$, are given independent normal priors $N_3(\mathbf{b}_j, \boldsymbol{\Omega}_j^{-1})$ with $\mathbf{b}_j = (550, 0, 0)^T$ and $\boldsymbol{\Omega}_j^{-1} = \text{diag}\{225, 10, 10\}$. The residual matrix $\mathbf{X} = (X_{ij})$ is given a matrix variate GGM

$$\begin{aligned} \text{vec}(\mathbf{X}^T) | \mathbf{K}_R, \mathbf{K}_C &\sim N_{p_R p_C}(\mathbf{0}, [\mathbf{K}_R \otimes \mathbf{K}_C]^{-1}), \\ \mathbf{K}_R | \delta_R, \mathbf{D}_R &\sim \text{Wis}_{G_R}(\delta_R, \mathbf{D}_R), \\ (z\mathbf{K}_C) | \delta_C, \mathbf{D}_C &\sim \text{Wis}_{G_C}(\delta_C, \mathbf{D}_C). \end{aligned}$$

where $z > 0$ is an auxiliary variable needed to impose the identifiability constraint $(\mathbf{K}_C)_{11} = 1$. The graph G_R , which captures the spatial dependencies among states, is fixed to the one induced by the neighborhood matrix $\mathbf{W} = (w_{ii'})$, where $w_{ii'} = 1$ if and only if states i and i' share a border, while we set a uniform prior on G_C . The degrees of freedom for the G-Wishart priors are set as $\delta_C = \delta_R = 3$, while the centering matrices are chosen as $\mathbf{D}_C = \mathbf{I}_2$ and, in the spirit of CAR models, $\mathbf{D}_R = (\delta_R - 2)w_{1+}(\mathbf{E}_W - 0.99\mathbf{W})^{-1}$. It follows that

$$\text{vec}(\mathbf{Y}^T) | \mathbf{K}_R, \mathbf{K}_C \sim N_{p_R p_C}(\text{vec}((\mathbf{Z}\boldsymbol{\beta})^T), [\mathbf{K}_R \otimes \mathbf{K}_C]^{-1}), \quad (24)$$

where \mathbf{Z} is a $p_R \times 3$ matrix whose i -th row corresponds to the vector $(1, Z_i, Z_i^2)$ and $\boldsymbol{\beta}$ is a 3×2 matrix whose j -th column is $\boldsymbol{\beta}_j$.

A MCMC algorithm for sampling from the posterior distribution of this model's parameters is described in Appendix A. This algorithm was run for 20000 iterations after a burn in of 1000 iterations. To assess convergence, we ran 10 instances of the algorithm at these settings, each initiated at different starting points and with different initial random seeds. Figure 8 shows the mean value, by iteration of the parameter z for each chain, which display significant agreement.

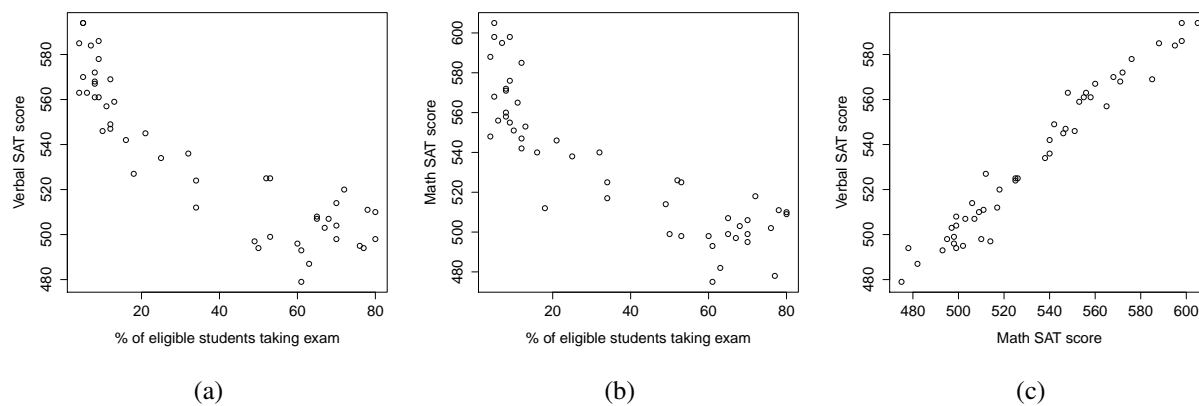
Table 3 shows the posterior estimates for the regression coefficients $\boldsymbol{\beta}_0$, $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$, along with the values reported by Wall (2004) for the verbal scores using two different models for the spatial structure: 1) a conventional CAR model, 2) an isotropic exponential variogram where the centroid of each state is used to calculate distances between states (see Wall, 2004 for details). Note that the posterior means obtained from our model roughly agree with those produced by the more conventional approaches to spatial modeling. However, the Bayesian standard errors obtained from our model are consistently larger than the standard errors reported in Wall (2004). We believe that this can be attributed to the richer spatial structure in our model, with the uncertainty in the spatial structure manifesting itself as additional uncertainty in other model parameters.

As expected, the data suggest a positive association between SAT scores. The posterior probability of an edge between verbal and math scores is 80%, suggesting moderate evidence of association between scores, while the posterior mean for the correlation between scores (conditional on it being different from zero) is 0.35 with a 95% credible interval of (0.09, 0.58).

Figure 3 shows the fitted values of math and verbal scores by state. In comparison to Figure 1, we see considerable shrinkage of the fitted scores, particularly for the mathematics SAT. We note that, even after accounting for the percentage of students taking the exam, substantial spatial structure remains; generally speaking, central states (except Texas) tend to have higher SAT scores than coastal states on either side. This is confirmed by the results in Figure 4, which compares

the spatial interaction structure implied by the prior parameter \mathbf{D}_R^{-1} to the structure implied by the posterior mean of \mathbf{K}_R . Since both matrices share the same underlying conditional independence graph, both matrices have the same pattern of zeroes. If no spatial structure were present in the data we would expect the non-zero entries to have shrunk towards zero. However the posterior for \mathbf{K}_R shows considerably larger variation in the spatial relationship structure than the prior. Indeed, while the prior imposes negative values for the off-diagonal of \mathbf{D}_R^{-1} (implying a positive correlation), the posterior of \mathbf{K}_R has several positive values, the largest of which is the entry between Colorado and New Mexico, implying a negative correlation between deviations in the SAT scores of these two states. More generally, note that the posterior mean for \mathbf{K}_R is very different from its prior mean, which suggest that the model is able to learn from the data about the spatial correlation structure.

Figure 2: Scatterplots for the SAT data set. Panel (a) shows the average verbal SAT score in the state vs. the number of eligible students that take the test, panel (b) shows the average math SAT score against the number of eligible students, and panel (c) shows the average verbal vs. the average math scores.



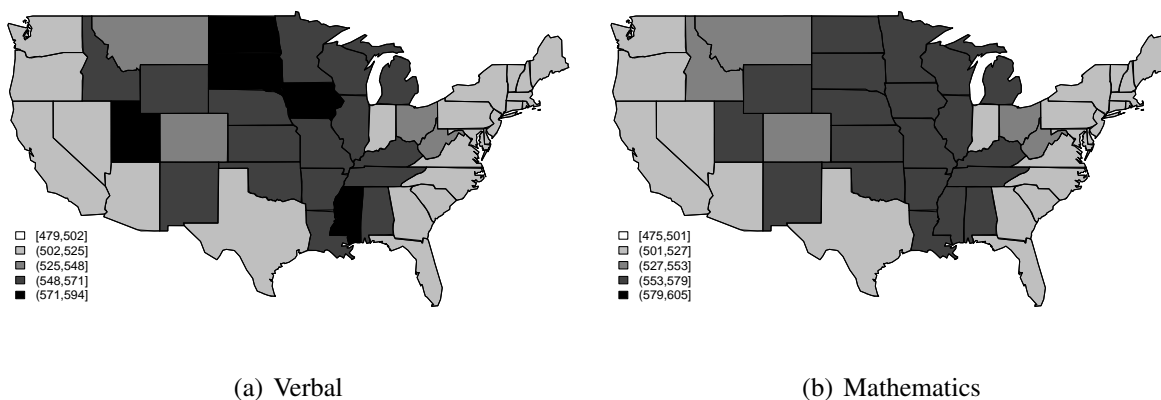
5.3 Mapping cancer mortality in the U.S.

Accurate and timely counts of cancer mortality are very useful in the cancer surveillance community for purposes of efficient resource allocation and planning. Estimation of current and future cancer mortality broken down by geographic area (state) and tumor have been discussed in a number of recent articles, including Tiwari et al. (2004), Ghosh & Tiwari (2007), Ghosh et al. (2007) and Ghosh et al. (2008). This section considers a multivariate spatial model on state-level cancer mortality rates in the United States for 2000. These mortality data are based on death certificates that are filed by certifying physicians and is collected and maintained by the National Center for Health Statistics (<http://www.cdc.gov/nchs>) as part of the National Vital Statistics Sys-

Table 3: Posterior Estimates for the regression parameters in the SAT data. The columns labeled GGM-MCAR correspond to the estimates obtained from the bivariate CAR model described in Section 5.2, while the column labeled Wall partially reproduces the results from Wall (2004) for the verbal SAT scores.

	GGM-MCAR		Wall (Verbal)	
	Math	Verbal	CAR	Variogram
Intercept (β_{0j})	580.12 (17.3)	583.26 (15.91)	584.63 (4.86)	583.64 (6.49)
Percent Eligible (β_{1j})	-2.147 (0.394)	-2.352 (0.498)	-2.48 (0.32)	-2.19 (0.31)
Percent Eligible Squared (β_{2j})	0.0154 (0.0055)	0.0171 (0.0069)	0.0189 (0.004)	0.0146 (0.004)

Figure 3: Fitted SAT scores in the 48 contiguous states. Left panel shows verbal scores, while the right panel shows mathematics scores. The scale is the same as in Figure 1

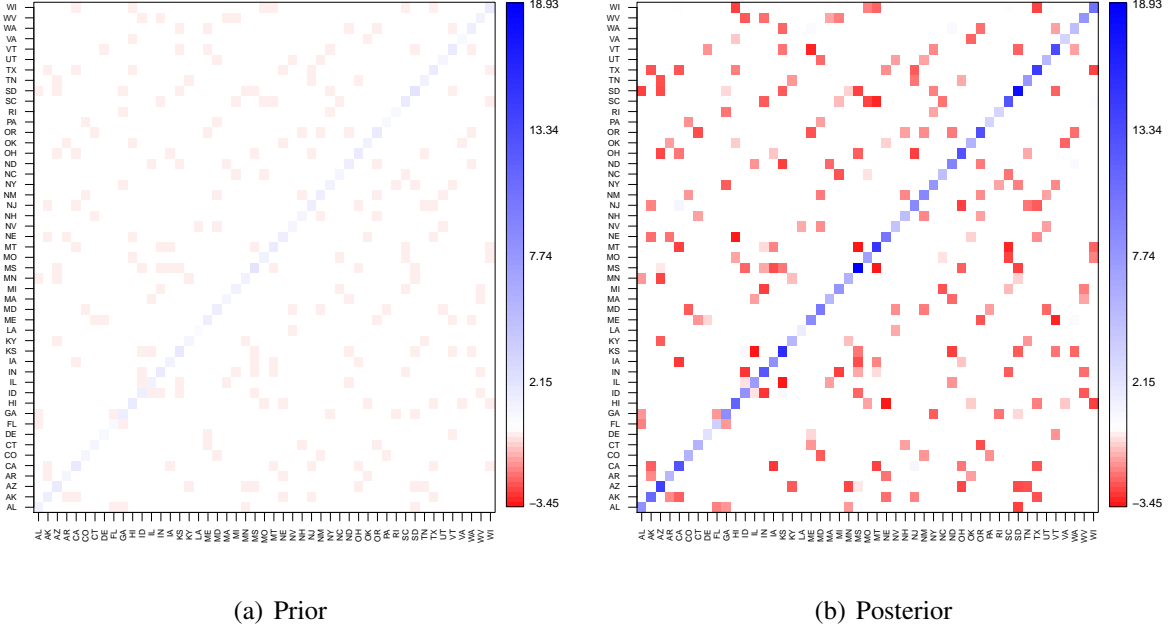


tem. The data is available from the Surveillance, Epidemiology and End Results (SEER) program of the National Cancer Institute (<http://seer.cancer.gov/seerstat>).

The data we analyze consists of mortality counts for 11 types of tumors recorded on the 50 states plus the District of Columbia. Along with mortality counts, we obtained population counts in order to model death risk. Figure 5 shows raw mortality rates for four of the most common types of tumors (colon, lung, breast and prostate). Although the pattern is slightly different for each of these cancers, a number of striking similarities are present; for example, Colorado appears to be a state with low mortality for all of these common cancers, while West Virginia, Pennsylvania and Florida present relatively high mortality rates across the board.

A sparse Poisson multivariate loglinear model with spatial random effects was fitted to this

Figure 4: A comparison of \mathbf{D}_W^{-1} and the posterior mean of \mathbf{K}_R in the SAT example. For comparison, the heat maps have the same scale, showing the considerable added variability in spatial dependence imposed by the model.



data. The model is constructed along the lines described in Section 4. In particular, we let Y_{ij} be the number of deaths in state $i = 1, \dots, p_R = 51$ for tumor type $j = 1, \dots, p_C = 25$. We set

$$\begin{aligned}
 Y_{ij} | \eta_{ij} &\sim \text{Poi}(\eta_{ij}), \\
 \log(\eta_{ij}) &= \log(m_i) + \mu_j + X_{ij},
 \end{aligned} \tag{25}$$

Here m_i is the population of state i , μ_j is the intercept for tumor type j , and X_{ij} is a zero-mean spatial random effect associated with location i and tumor j . We follow the notation from Section 3 and denote $\tilde{X}_{ij} = \mu_j + X_{ij}$. We further model $\tilde{\mathbf{X}} = (\tilde{X}_{ij})$ with a matrix-variate Gaussian graphical model prior:

$$\begin{aligned}
 \text{vec}(\tilde{\mathbf{X}}^T) | \boldsymbol{\mu}, \mathbf{K}_R, \mathbf{K}_C &\sim \text{N}_{p_R p_C} \left(\text{vec} \left\{ (\mathbf{1}_{p_R} \boldsymbol{\mu}^T)^T \right\}, [\mathbf{K}_R \otimes \mathbf{K}_C]^{-1} \right), \\
 \mathbf{K}_R | \delta_R, \mathbf{D}_R &\sim \text{Wis}_{G_R}(\delta_R, \mathbf{D}_R), \\
 (z \mathbf{K}_C) | \delta_C, \mathbf{D}_C &\sim \text{Wis}_{G_C}(\delta_C, \mathbf{D}_C).
 \end{aligned} \tag{26}$$

Here $\mathbf{1}_l$ is the column vector of length l with all elements equal to 1, $\boldsymbol{\mu} = (\mu_1, \dots, \mu_{p_C})^T$ and $z > 0$ is an auxiliary variable needed to impose the identifiability constraint $(\mathbf{K}_C)_{11} = 1$. The row graph G_R is again fixed and derived from the incidence matrix \mathbf{W} derived from neighborhood structure across U.S. states, while the column graph G_C is given a uniform distribution. The degrees of

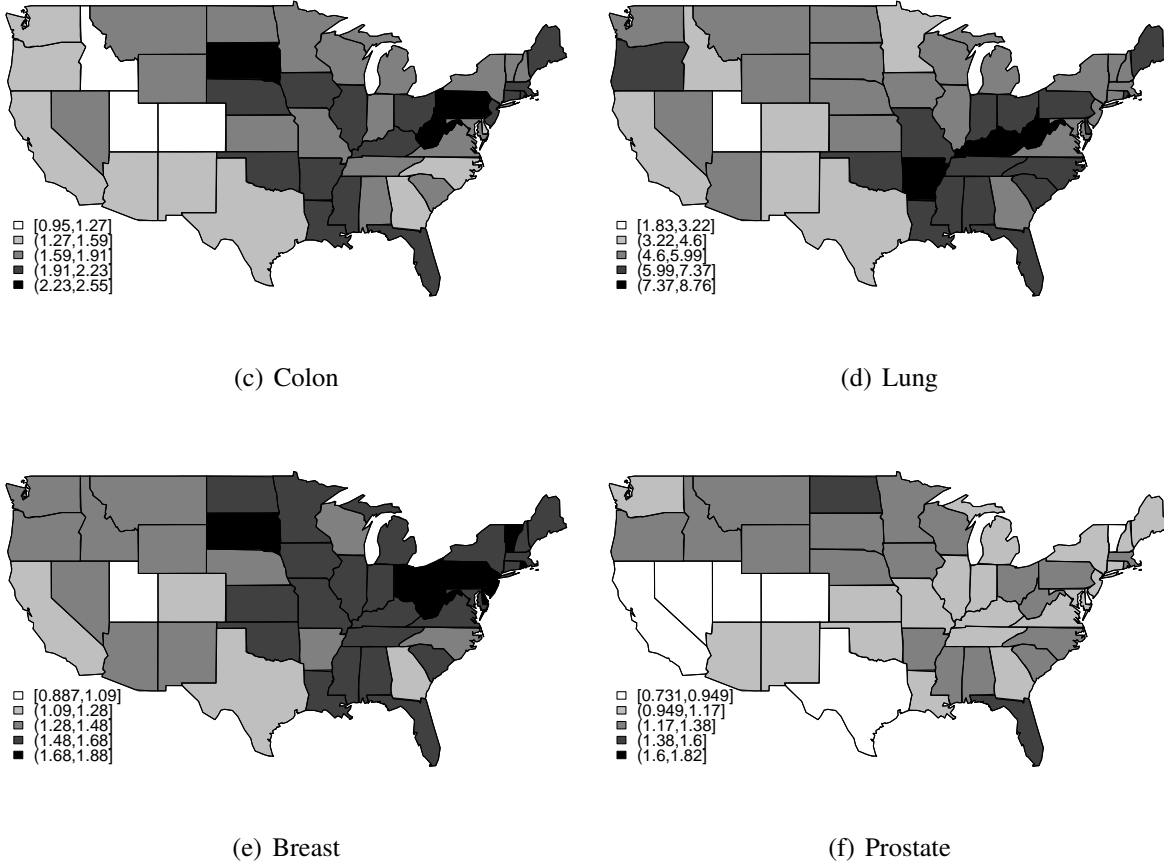


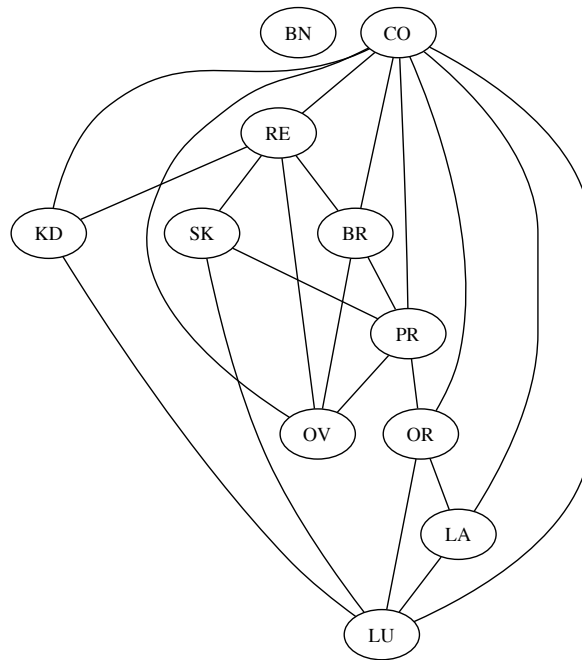
Figure 5: Mortality rates (per 10,000 habitants) in the 48 contiguous states corresponding to four common cancers during 2000.

freedom for the G-Wishart priors are set as $\delta_R = \delta_C = 3$, while the centering matrices are chosen as $\mathbf{D}_C = \mathbf{I}_{p_C}$ and $\mathbf{D}_R = (\delta_R - 2)w_{1+}(\mathbf{E}_W - 0.99\mathbf{W})^{-1}$. We complete the model specification by choosing a uniform prior for the column graph $p(G_C) \propto 1$ and a multivariate normal prior for the mean rates vector $\boldsymbol{\mu} \sim N_{p_C}(\boldsymbol{\mu}_0, \boldsymbol{\Omega}^{-1})$ where $\boldsymbol{\mu}_0 = \mu_0 \mathbf{1}_{p_C}$ and $\boldsymbol{\Omega} = \omega^{-2} \mathbf{I}_{p_C}$. We set μ_0 to the median log incidence rate across all cancers and regions, and ω to be twice the interquartile range in raw log incidence rates.

Posterior inferences on this model are obtained by extending the sampling algorithm described in Section 3; details are presented in Appendix B. Since mortality counts below 25 are often deemed unreliable in the cancer surveillance community, we treated them as missing and sampled them as part of our algorithm. Inferences are based on 50000 samples obtained after a discarding the first 5000 samples. We ran ten separate chains to monitor convergence. Figure 8 shows the mean value of z by log iteration for each of the ten chains. All ten agree by 50000 iterations.

Figure 6 shows the graphical model constructed by adding an edge between any two variables

Figure 6: Graphical model from the cancer example. The eleven cancers are: brain (BN), breast (BR), colon (CO), kidney (KD), larynx (LA), lung (LU), oral (OR), ovarian (OV), prostate (PR), skin (SK), and rectum (RE)




if this edge was present in more than 50% of the total repetitions. This figure shows several interesting patterns. First, we see that lung cancer (LU) and colon cancer (CO) have the highest degree of connectivity. This is a useful result, since lung and colon cancer are the two most prevalent cancers in our dataset. This implies that variability in these cancers, after controlling for spatial effects, is a good predictor of variability in other lower count cancers. We also see a clique formed between lung, larynx (LA) and oral (OR) cancers, which is not surprising, as many cases of these cancers are believed to be related to chronic smoking. Another connexion we expected to see corresponds to the edge between breast (BR) and ovarian (OV) cancer. Indeed, both of these cancers have been related to mutations in the BRAC genes, and it therefore is reasonable that increased levels of one would be correlated to an increase in the other. We find that Brain (BN) cancer rates are not correlated with the rates of any other cancer. Finally, Figure 7 shows the average fitted values for the cancer mortality by state for the four cancers shown in Figure 5. In a result similar to those show in the SAT example above, we see some shrinkage towards the mean, but a preservation of the spatial variability present in the underlying data.


Table 4 shows the distribution of the imputed counts for those cancer/state combinations where the reported count was below 25.

Table 4: Distribution of imputed counts for cancers with fewer than 25 reported counts in a given state


State	Cancer	Median	95% Interval
District of Columbia	Brain	13	(3 , 24)
North Dakota	Brain	19	(9 , 25)
District of Columbia	Rectum	19	(11 , 24)
North Dakota	Rectum	20	(13 , 25)
Wyoming	Rectum	15	(9 , 22)
Delaware	Skin	20	(12 , 25)
District of Columbia	Skin	15	(7 , 23)
North Dakota	Skin	17	(10 , 23)
South Dakota	Skin	20	(13 , 25)
Vermont	Skin	17	(10 , 24)
Wyoming	Skin	13	(8 , 20)
District of Columbia	Kidney	21	(14 , 25)
Wyoming	Kidney	20	(13 , 25)
Delaware	Oral	19	(12 , 25)
North Dakota	Oral	15	(9 , 22)
Rhode Island	Oral	22	(16 , 25)
South Dakota	Oral	18	(11 , 24)
Vermont	Oral	16	(10 , 23)
Wyoming	Oral	11	(6 , 18)
Delaware	Larynx	12	(6 , 19)
District of Columbia	Larynx	11	(5 , 20)
Idaho	Larynx	12	(6 , 18)
Maine	Larynx	20	(11 , 25)
Montana	Larynx	10	(4 , 16)
Nebraska	Larynx	18	(11 , 25)
New Hampshire	Larynx	17	(11 , 24)
New Mexico	Larynx	13	(7 , 21)
North dakota	Larynx	6	(2 , 11)
Rhode Island	Larynx	16	(9 , 23)
South Dakota	Larynx	8	(4 , 14)
Utah	Larynx	8	(3 , 15)
Vermont	Larynx	9	(4 , 15)
Wyoming	Larynx	5	(2 , 9)




(a) Colon



(b) Lung



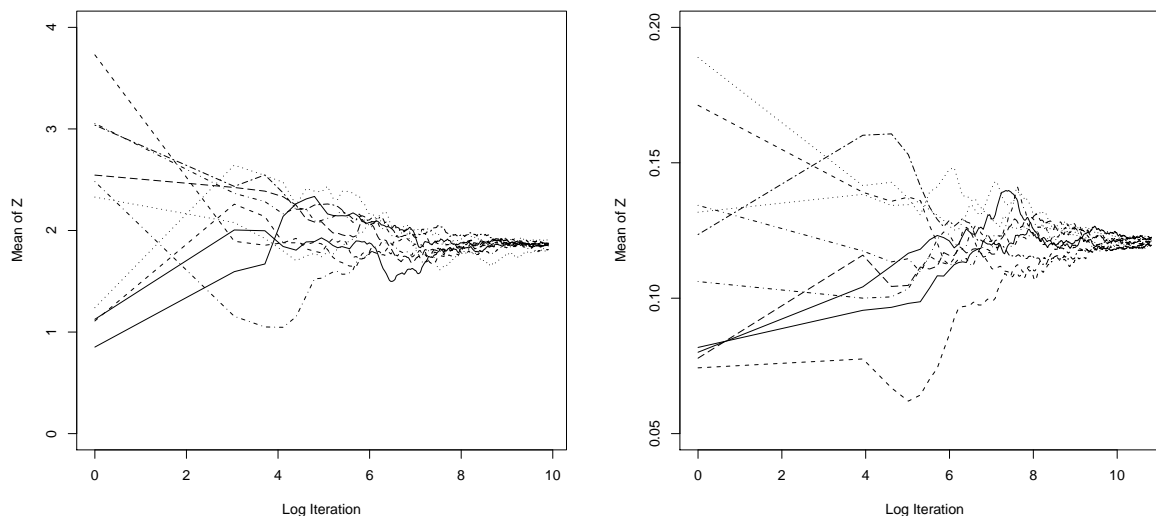
(c) Breast



(d) Prostate

Figure 7: Fitted Mortality rates (per 10,000 habitants) in the 48 contiguous states corresponding to four common cancers during 2000.

Figure 8: Convergence plots for the SAT (left) and cancer examples (right). These plots show the mean of z , by log iteration for each of ten chains.



6 Discussion

This paper has developed and illustrated a surprisingly powerful computational approach for multivariate and matrix-variate GGMs, and we believe that our application to the construction of spatial models for lattice data makes for particularly appealing illustrations. On one hand, the examples we have considered suggest that the additional flexibility in the spatial correlation structure provided by our approach is necessary to accurately model some spatial data sets. Indeed, our approach allows for “non-stationary” CAR models, where the spatial autocorrelation and spatial smoothing parameters vary spatially. On the other hand, to the best of our knowledge, we are unaware of any approach in the literature to construct and estimate sparse MCAR, particularly under a Bayesian approach. In addition to providing insights into the mechanisms underlying the data generation process, it has been repeatedly argued in the literature (see, for example, Carvalho & West, 2007 and Rodriguez et al., 2010) that sparse models tend to provide improved predictions, which is particularly important in such a relevant application as cancer mortality surveillance.

The proposed algorithms work well in the examples discussed here where the number of geographical areas is moderate (that is, around fifty or less). Convergence seems to be achieved quickly (usually within the first 5000 iterations) and running times, although longer than for conventional MCAR models, are still short enough as to make routine implementation feasible (about 1 hour for the simpler SAT dataset, and about 1.5 days for the more complex cancer dataset, both running on a dual-core 2.8 Ghz computer running Linux). However, computation can still be very challenging when the number of units in the lattice is very high. In the future we plan to explore parallel implementations (both on regular clusters and GPU cards), and to exploit the sparse structure of the variance-covariance matrix to reduce computational times.

A Details of the MCMC algorithm for the sparse multivariate spatial Gaussian model for SAT scores

The Markov chain Monte Carlo sampler for the spatial regression model in Section 5.2 involves iterative updating of model parameters from six full conditional distributions.

Step 1: Resample the regression coefficients $\beta_1 = (\beta_{01}, \beta_{11}, \beta_{21})$. We denote by \mathbf{Y}_{*j} the j -th column of the $p_R \times p_C$ matrix \mathbf{Y} . From Theorem 2.3.11 of Gupta & Nagar (2000) applied to equation (24) we have

$$\mathbf{Y}_{*1} | \mathbf{Y}_{*2} \sim N_{p_R} \left(\mathbf{Z} \beta_1 + (\mathbf{Y}_{*2} - \mathbf{Z} \beta_2) \frac{(\mathbf{K}_C^{-1})_{21}}{(\mathbf{K}_C^{-1})_{22}}, [(\mathbf{K}_C)_{11} \mathbf{K}_R]^{-1} \right).$$

Given the prior $\beta_2 \sim N_3(\mathbf{b}_2, \Omega_2^{-1})$ and the constraint $(\mathbf{K}_C)_{11} = 1$, it follows that β_1 is updated by sampling from the multivariate normal $N_3(\mathbf{m}_{\beta_2}, \Omega_{\beta_2}^{-1})$ where

$$\Omega_{\beta_2} = \mathbf{Z}^T \mathbf{K}_R \mathbf{Z} + \Omega_2, \quad \mathbf{m}_{\beta_2} = \Omega_{\beta_2}^{-1} \left\{ \mathbf{Z}^T \mathbf{K}_R \left[\mathbf{Y}_{*1} - (\mathbf{Y}_{*2} - \mathbf{Z} \beta_2) \frac{(\mathbf{K}_C^{-1})_{21}}{(\mathbf{K}_C^{-1})_{22}} \right] + \Omega_2 \mathbf{b}_2 \right\}.$$

Step 2: Resample the regression coefficients $\beta_2 = (\beta_{02}, \beta_{12}, \beta_{22})$. Since apriori $\beta_1 \sim N_3(\mathbf{b}_1, \Omega_1^{-1})$, it follows that β_2 is updated by sampling from the multivariate normal $N_3(\mathbf{m}_{\beta_1}, \Omega_{\beta_1}^{-1})$ where

$$\Omega_{\beta_1} = (\mathbf{K}_C)_{22} \mathbf{Z}^T \mathbf{K}_R \mathbf{Z} + \Omega_1, \quad \mathbf{m}_{\beta_1} = \Omega_{\beta_1}^{-1} \left\{ (\mathbf{K}_C)_{22} \mathbf{Z}^T \mathbf{K}_R \left[\mathbf{Y}_{*2} - (\mathbf{Y}_{*1} - \mathbf{Z} \beta_1) \frac{(\mathbf{K}_C^{-1})_{12}}{(\mathbf{K}_C^{-1})_{11}} \right] + \Omega_1 \mathbf{b}_1 \right\}.$$

Step 3: Resample the row precision matrix \mathbf{K}_R encoding the spatial structure.

Step 4: Resample the column graph G_C .

Step 5: Resample the column precision matrix \mathbf{K}_C .

Step 6: Resample the auxiliary variable z .

After updating the regression coefficients in Steps 1 and 2, we recalculate the residual matrix $\mathbf{X} = (X_{ij})$, where $X_{ij} = Y_{ij} - \beta_{0j} - \beta_{1j} Z_i - \beta_{2j} Z_i^2$. The updated \mathbf{X} represents a matrix variate dataset with a sample size $n = 1$, based on which \mathbf{K}_R , G_C , \mathbf{K}_C and z are updated. Thus Steps 3, 4, 5 and 6 above correspond to Steps 2, 3, 4 and 5 from Section 3, and no new algorithm needs to be described. Note that Step 1 of Section 3 is unnecessary as the row graph G_R is assumed known, and given by the neighborhood structure.

B Details of the MCMC algorithm for the sparse multivariate spatial Poisson count model

The Markov chain Monte Carlo sampler for the sparse multivariate spatial Poisson count model from Section 5.3 involves iterative updating of model parameters in six steps.

Step 1: Resample the mean rates $\boldsymbol{\mu}$. The full conditional for $\boldsymbol{\mu}$ is

$$p(\boldsymbol{\mu} | \widetilde{\mathbf{X}}, \mathbf{K}_R, \mathbf{K}_C) \propto \exp \left\{ -\frac{1}{2} \text{tr} \left[\mathbf{K}_R (\widetilde{\mathbf{X}} - \mathbf{1}_{p_R} \boldsymbol{\mu}^T) \mathbf{K}_C (\widetilde{\mathbf{X}} - \mathbf{1}_{p_R} \boldsymbol{\mu}^T)^T + \mathbf{1}_{p_R} (\boldsymbol{\mu} - \boldsymbol{\mu}_0)^T \boldsymbol{\Omega} (\boldsymbol{\mu} - \boldsymbol{\mu}_0) \mathbf{1}_{p_R}^T \right] \right\}.$$

It follows that $\boldsymbol{\mu}$ is updated by direct sampling from the multivariate normal $N_{p_C}(\mathbf{m}_\mu, \mathbf{K}_\mu^{-1})$ where

$$\mathbf{K}_\mu = (\mathbf{1}_{p_R}^T \mathbf{K}_R \mathbf{1}_{p_R}) \mathbf{K}_C + p_R \boldsymbol{\Omega}, \quad \mathbf{m}_\mu = \mathbf{K}_\mu^{-1} \left[\mathbf{K}_C \widetilde{\mathbf{X}}^T \mathbf{K}_R \mathbf{1}_{p_R} + p_R \boldsymbol{\Omega} \boldsymbol{\mu}_0 \right].$$

Step 2: Resample $\widetilde{\mathbf{X}}$. We update the centered random effects $\widetilde{\mathbf{X}}$ by sequentially updating each row vector $\widetilde{\mathbf{X}}_{i^*}$, $i = 1, \dots, p_R$. The matrix-variate GGM prior from equation (26) implies the following conditional prior for $\widetilde{\mathbf{X}}_{i^*}$ (Gupta & Nagar, 2000):

$$\widetilde{\mathbf{X}}_{i^*}^T | \{ \widetilde{\mathbf{X}}_{i'^*} : i' \neq i \}, \boldsymbol{\mu}, \mathbf{K}_R, \mathbf{K}_C \sim N_{p_C}(\mathbf{M}_i, \boldsymbol{\Omega}_i^{-1}), \quad (27)$$

where

$$\mathbf{M}_i = \boldsymbol{\mu} - \sum_{i' \in \partial i} \frac{(\mathbf{K}_R)_{i'i'}}{(\mathbf{K}_R)_{ii}} (\widetilde{\mathbf{X}}_{i'^*}^T - \boldsymbol{\mu}), \quad \boldsymbol{\Omega}_i = (\mathbf{K}_R)_{ii} \mathbf{K}_C.$$

It follows that the full conditional distribution of $\widetilde{\mathbf{X}}_{i^*}$ is written as

$$p(\widetilde{\mathbf{X}}_{i^*} | \{ \widetilde{\mathbf{X}}_{i'^*} : i' \neq i \}, \mathbf{Y}_{i^*}, \boldsymbol{\mu}, \mathbf{K}_R, \mathbf{K}_C) \propto \prod_{j=1}^{p_C} \exp \{ Y_{ij} (\log(m_i) + \widetilde{X}_{ij}) - m_i \exp(\widetilde{X}_{ij}) \} \exp \left\{ -\frac{1}{2} [\widetilde{\mathbf{X}}_{i^*} - (\mathbf{M}_i)^T] \boldsymbol{\Omega}_i [(\widetilde{\mathbf{X}}_{i^*})^T - \mathbf{M}_i] \right\}. \quad (28)$$

Since the distribution (28) does not correspond to any standard distribution, we design a Metropolis-Hastings algorithm to sample from it. We consider a strictly positive precision parameter $\widetilde{\sigma}$. For each $j = 1, \dots, p_C$, we update the j -th element of $\widetilde{\mathbf{X}}_{i^*}$ by sampling $\gamma \sim N(\widetilde{X}_{ij}, \widetilde{\sigma}^2)$. We define a candidate row vector $\widetilde{\mathbf{X}}_{i^*}^{new}$ by replacing \widetilde{X}_{ij} with γ in $\widetilde{\mathbf{X}}_{i^*}$. We update the current i -th row of $\widetilde{\mathbf{X}}$ with $\widetilde{\mathbf{X}}_{i^*}^{new}$ with probability

$$\min \left\{ 1, \frac{p(\widetilde{\mathbf{X}}_{i^*}^{new} | \{ \widetilde{\mathbf{X}}_{i'^*} : i' \neq i \}, \mathbf{Y}_{i^*}, \boldsymbol{\mu}, \mathbf{K}_R, \mathbf{K}_C)}{p(\widetilde{\mathbf{X}}_{i^*} | \{ \widetilde{\mathbf{X}}_{i'^*} : i' \neq i \}, \mathbf{Y}_{i^*}, \boldsymbol{\mu}, \mathbf{K}_R, \mathbf{K}_C)} \right\}.$$

Otherwise the i -th row of $\widetilde{\mathbf{X}}$ remains unchanged.

Step 3: Resample the row precision matrix \mathbf{K}_R encoding the spatial structure.

Step 4: Resample the column graph G_C .

Step 5: Resample the column precision matrix \mathbf{K}_C .

Step 6: Resample the auxiliary variable z .

After completing Steps 1 and 2, we recalculate the zero-mean spatial random effects $X_{ij} = \widetilde{X}_{ij} - \mu_j$. The updated $\mathbf{X} = (X_{ij})$ represents a matrix variate dataset with a sample size $n = 1$, based on which \mathbf{K}_R , G_C , \mathbf{K}_C and z are updated. Steps 3, 4, 5 and 6 above are Steps 2, 3, 4 and 5 from Section 3.

References

- ASCI, C. & PICCIONI, M. (2007). Functionally compatible local characteristics for the local specification of priors in graphical models. *Scand. J. Statist.* **34**, 829–40.
- ATAY-KAYIS, A. & MASSAM, H. (2005). A Monte Carlo method for computing the marginal likelihood in nondecomposable Gaussian graphical models. *Biometrika* **92**, 317–35.
- BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society, Series B* **36**, 192–236.
- BESAG, J. & KOOPERBERG, C. (1995). On conditional and intrinsic autoregressions. *Biometrika* **82**, 733–746.
- BROOK, D. (1964). On the distinction between the conditional probability and the joint probability approaches in the specification of the nearest neighbour systems. *Biometrika* **51**, 481–489.
- CARVALHO, C. M., MASSAM, H. & WEST, M. (2007). Simulation of hyper-inverse Wishart distributions in graphical models. *Biometrika* **94**, 647–659.
- CARVALHO, C. M. & SCOTT, J. G. (2009). Objective Bayesian model selection in Gaussian graphical models. *Biometrika* **96**.
- CARVALHO, C. M. & WEST, M. (2007). Dynamic matrix-variate graphical models. *Bayesian Analysis* **2**, 69–98.
- CRESSIE, N. A. C. (1973). *Statistics for Spatial Data*. New York: Wiley.
- CRESSIE, N. A. C. & CHAN, N. H. (1989). Spatial modeling of regional variables. *Journal of American Statistical Association* **84**, 393–401.
- DEMPSTER, A. P. (1972). Covariance selection. *Biometrics* **28**, 157–75.
- DIACONNIS, P. & YLVISAKER, D. (1979). Conjugate priors for exponential families. *Ann. Statist.* **7**, 269–81.
- GELFAND, A. E. & VOUNATSOU, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* **4**, 11–25.
- GHOSH, K., GHOSH, P. & TIWARI, R. C. (2008). Comment on “The nested Dirichlet process” by Rodriguez, Dunson and Gelfand. *Journal of the American Statistical Association* **103**, 1147–1149.
- GHOSH, K. & TIWARI, R. C. (2007). Prediction of U.S. cancer mortality counts using semiparametric Bayesian techniques. *Journal of the American Statistical Association* **102**, 7–15.

- GHOSH, K., TIWARI, R. C., FEUER, E. J., CRONIN, K. & JEMAL, A. (2007). Predicting U.S. cancer mortality counts using state space models. In *Computational Methods in Biomedical Research*, Eds. R. Khattree & D. N. Naik, pp. 131–151. Boca Raton, FL: Chapman & Hall / CRC.
- GIUDICI, P. & GREEN, P. J. (1999). Decomposable graphical Gaussian model determination. *Biometrika* **86**, 785–801.
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**.
- GUPTA, A. K. & NAGAR, D. K. (2000). *Matrix Variate Distributions*, volume 104 of *Monographs and Surveys in Pure and Applied Mathematics*. London: Chapman & Hall/CRC.
- JONES, B., CARVALHO, C., DOBRA, A., HANS, C., CARTER, C. & WEST, M. (2005). Experiments in stochastic computation for high-dimensional graphical models. *Statist. Sci.* **20**, 388–400.
- LAURITZEN, S. L. (1996). *Graphical Models*. Oxford University Press.
- LENKOSKI, A. & DOBRA, A. (2010). Computational aspects related to inference in Gaussian graphical models with the G-wishart prior. *Journal of Computational and Graphical Statistics* To appear.
- LETAC, G. & MASSAM, H. (2007). Wishart distributions for decomposable graphs. *Ann. Statist.* **35**, 1278–323.
- MADIGAN, D. & YORK, J. (1995). Bayesian graphical models for discrete data. *International Statistical Review* **63**, 215–232.
- MARDIA, K. V. (1988). Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *Journal of Multivariate Analysis* **24**, 265–284.
- PICCIONI, M. (2000). Independence structure of natural conjugate densities to exponential families and the Gibbs Sampler. *Scand. J. Statist.* **27**, 111–27.
- RODRIGUEZ, A., DOBRA, A. & LENKOSKI, A. (2010). Sparse covariance estimation in heterogeneous samples. Technical report, University of California - Santa Cruz.
- ROVERATO, A. (2000). Cholesky decomposition of a hyper inverse wishart matrix. *Biometrika* **87**, 99–112.
- ROVERATO, A. (2002). Hyper inverse Wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models. *Scand. J. Statist.* **29**, 391–411.
- SUN, D., TSUTAKAWA, R. K., KIM, H. & HE, Z. (2000). Bayesian analysis of mortality rates with disease maps. *Statistics and Medicine* **19**, 2015–2035.

- TIWARI, R. C., GHOSH, K., JEMAL, A., HACHEY, M., WARD, E., THUN, M. J. & FEUER, E. J. (2004). A new method for predicting US and state-level cancer mortality counts for the current calendar year. *CA: A Cancer Journal for Clinicians* **54**, 30–40.
- TOMITA, E., TANAKA, A. & TAKAHASHI, H. (2006). The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical Computer Science* **363**, 28–42.
- WALL, M. M. (2004). A close look at the spatial structure implied by the CAR and SAR models. *Journal of Statistical Planning and Inference* **121**.
- WANG, H. & WEST, M. (2009). Bayesian analysis of matrix normal graphical models. *Biometrika* **96**, 821–834.