



An Expectation Conditional Maximization Approach for Gaussian Graphical Models

Zehang Richard Li & Tyler H. McCormick

To cite this article: Zehang Richard Li & Tyler H. McCormick (2019) An Expectation Conditional Maximization Approach for Gaussian Graphical Models, *Journal of Computational and Graphical Statistics*, 28:4, 767-777, DOI: [10.1080/10618600.2019.1609976](https://doi.org/10.1080/10618600.2019.1609976)

To link to this article: <https://doi.org/10.1080/10618600.2019.1609976>

 View supplementary material 

 Published online: 19 Jun 2019.

 Submit your article to this journal 

 Article views: 430

 View related articles 

 View Crossmark data 
CrossMark

 Citing articles: 1 View citing articles 

An Expectation Conditional Maximization Approach for Gaussian Graphical Models

Zehang Richard Li^a and Tyler H. McCormick^b

^aDepartment of Biostatistics, Yale School of Public Health, New Haven, CT; ^bDepartments of Statistics & Sociology, University of Washington, Seattle, WA

ABSTRACT

Bayesian graphical models are a useful tool for understanding dependence relationships among many variables, particularly in situations with external prior information. In high-dimensional settings, the space of possible graphs becomes enormous, rendering even state-of-the-art Bayesian stochastic search computationally infeasible. We propose a deterministic alternative to estimate Gaussian and Gaussian copula graphical models using an expectation conditional maximization (ECM) algorithm, extending the EM approach from Bayesian variable selection to graphical model estimation. We show that the ECM approach enables fast posterior exploration under a sequence of mixture priors, and can incorporate multiple sources of information. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received September 2017
Revised April 2019

KEYWORDS

Copula graphical model;
Sparse precision matrix;
Spike-and-slab prior

1. Introduction

For high-dimensional data, graphical models (Lauritzen 1996) provide a convenient characterization of the conditional independence structure among variables. In settings where the rows in the data matrix $X \in \mathbb{R}^{n \times p}$ follow an iid multivariate Gaussian distribution, $\text{Normal}(\mathbf{0}, \Sigma)$, the zeros in off-diagonal elements of the precision matrix $\Omega = \Sigma^{-1}$ correspond to pairs of variables that are conditionally independent. Standard maximum likelihood estimators of the sparse precision matrix behave poorly and do not exist when $n < p$, leading to extensive work on algorithms (and their properties) for estimating Ω (e.g., Meinshausen and Bühlmann 2006; Yuan and Lin 2007; Friedman, Hastie, and Tibshirani 2008; Rothman et al. 2008; Friedman, Hastie, and Tibshirani 2010; Cai, Zhang, and Zhou 2010; Witten, Friedman, and Simon 2011; Mazumder and Hastie 2012, etc.).

In the Bayesian literature, structure learning in high-dimensional Gaussian graphical models has also gained popularity in the past decade. Broadly speaking, two main classes of priors have been studied for inference of the precision matrix in Gaussian graphical models, namely the *G*-Wishart prior and shrinkage priors. The *G*-Wishart prior (Roverato 2002) extends the Wishart distribution by restricting its support to the space of positive definite matrices with zeros specified by a graph. It is attractive in Bayesian modeling due to its conjugacy with the Gaussian likelihood. Posterior inference under the *G*-Wishart distribution, though computationally challenging, can be carried out via various algorithms, including shotgun stochastic search (Jones et al. 2005), reversible jump MCMC (Lenkoski and Dobra 2011; Dobra, Lenkoski, and Rodriguez 2011; Wang and Li 2012), birth-death MCMC (Mohammadi et al. 2017), etc. More recently, shrinkage priors for precision matrices have gained much popularity, as they provide Bayesian interpretations to some of the widely used penalized likelihood estimators. As a Bayesian analogy to graphical lasso (Yin and

Li 2011; Witten, Friedman, and Simon 2011; Mazumder and Hastie 2012), Bayesian graphical lasso was proposed in Wang (2012) and Peterson et al. (2013). Wang (2015) later drew the connection between the Bayesian variable selection (George and McCulloch 1993) and Bayesian graphical model estimation, and proposed a new class of spike-and-slab prior for precision and covariance matrices. This class of priors was also later explored in Peterson, Stingo, and Vannucci (2015) to estimate the dependence structures among regression coefficients, and in Lukemire et al. (2017) to estimate multiple networks. This type of spike-and-slab prior enables a fast block Gibbs sampler that significantly improves the scalability of the model, but such flexibility is at the cost of prior interpretability since the implied marginal distribution of each elements in the precision matrix is intractable due to the positive definiteness constraint. Wang (2015) provided some heuristics and discussions on prior choices, but it is still not clear how to choose the hyperparameters for practical problems or how those choices affect parameter estimation.

In this article, we introduce a new algorithm to estimate sparse precision matrices with spike-and-slab priors (Wang 2015) using a deterministic approach, EM graph selection (EMGS), based on the expectation conditional maximization (ECM) algorithm (Meng and Rubin 1993). We also show that a stochastic variation of the EMGS approach can be extended to copula graphical model estimation. Our work extends the EM approach to variable selection (EMVS) (Ročková and George 2014) to general graphical model estimation.

The proposed ECM algorithm is closely connected to frequentist penalized likelihood methods. Similar to the algorithms with concave penalized regularization, such as SCAD (Fan, Feng, and Wu 2009), the spike-and-slab prior used in our method yields sparse inverse covariance matrix where large values are estimated with less bias (see Figure 1).

CONTACT Zehang Richard Li  zehang.li@yale.edu  Department of Biostatistics, Yale School of Public Health, 47 College St, Suite 18, New Haven, CT 06520-8034.
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/JCGS.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JCGS.
This article was submitted prior to Tyler McCormick being named Editor of JCGS.

© 2019 American Statistical Association, Institute of Mathematical Statistics, and Interface Foundation of North America

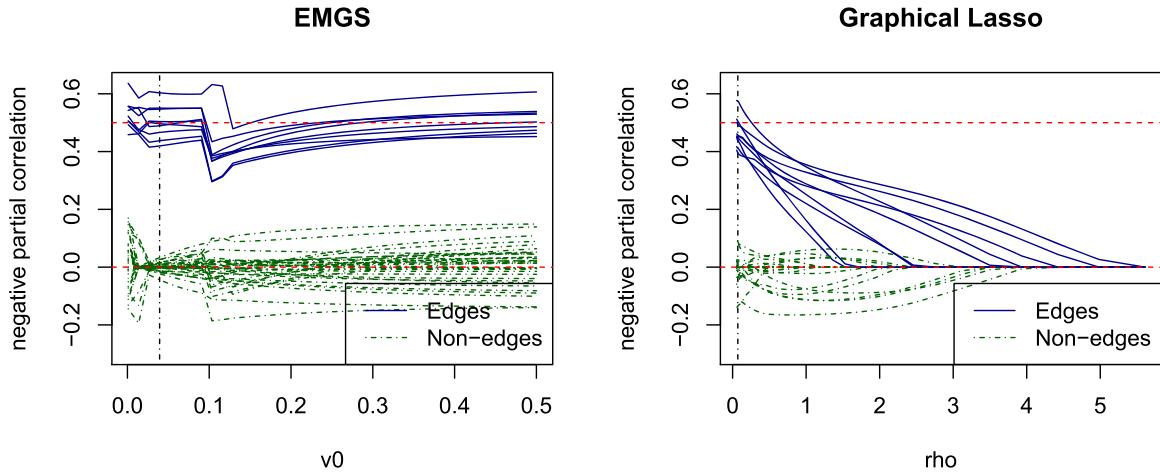


Figure 1. Comparing partial correlation path using EMGS and graphical lasso, on a 10-node graph. The red dashed line at 0.5 is the true value for the nonzero negative partial correlations. The nonzero off-diagonal elements are plotted with blue solid lines. The vertical line indicates the tuning parameter selected with cross-validation.

Similar work has been concurrently developed by Deshpande, Rockova, and George (2017) using spike-and-slab lasso prior in the multivariate linear regression models. The proposed approach in this article differs from Deshpande, Rockova, and George (2017) in two ways: First, we use a mixture of Gaussian distributions instead of the Laplace distributions as the prior on the off-diagonal elements of the precision matrix, which allows us to construct a closed-form conditional maximization step using coordinate descent, rather than relying on additional algorithms solving a graphical lasso problem at each iteration. Second, and more importantly, our work also differs in scope, as we extended the algorithm to non-Gaussian outcomes, the scenarios where informative priors exist, and to incorporate the imputation of missing values.

The rest of the article is organized as follows: In Section 2, we describe the spike-and-slab prior we use for the precision matrix. Section 3 presents the main ECM framework and algorithms for Gaussian graphical model estimation, and Section 4 proposes the extension to the copula graphical model and the modified stochastic ECM algorithm. Then in Section 5 we explore the incorporation of informative prior knowledge into the model. We discuss briefly about single model selection in Section 6. Section 7 examines the performance of our method through numerical simulations. Sections 8 and 9 further illustrate our model using two examples from scientific settings. Section 8 compares our method and alternatives in terms of structure learning and prediction of missing values in a dataset of hourly bike/pedestrian traffic volumes along a busy trail in Seattle. Section 9 discusses our method in the context of learning latent structures among binary symptoms from a dataset of verbal autopsy (VA) surveys, which are used to estimate a likely cause of death in places where most deaths occur outside of medical facilities. Finally, in Section 10 we discuss the limitations of the approach and provide some future directions for improvements.

2. Spike-and-Slab Prior for Gaussian Graphical Model

First, we review the *stochastic search structure learning* (SSSL) prior proposed in Wang (2015) for sparse precision matrices.

Consider the standard Gaussian graphical model setting, with observed data $\mathbf{X} \in \mathbb{R}^{n \times p}$. Each observation follows a multivariate Gaussian distribution, that is, $\mathbf{x}_i \sim \text{Normal}(\mathbf{0}, \boldsymbol{\Omega}^{-1})$, where \mathbf{x}_i is the i th row of the \mathbf{X} , and $\boldsymbol{\Omega}$ is the precision matrix. Given hyperparameter v_0 , v_1 , and π_δ , the prior on $\boldsymbol{\Omega}$ is defined as

$$p(\boldsymbol{\Omega}|\boldsymbol{\delta}) = C_\delta^{-1} \prod_{j < k} \text{Normal}(\omega_{jk}|0, v_{\delta_{jk}}^2) \times \prod_j \text{Exp}(\omega_{jj}|\lambda/2) \mathbf{1}_{\Omega \in M^+}, \quad (1)$$

$$p(\boldsymbol{\delta}|\pi_\delta) \propto C_\delta \prod_{j < k} \pi_\delta^{\delta_{jk}} (1 - \pi_\delta)^{1 - \delta_{jk}}, \quad (2)$$

where δ_{jk} are latent indicator variables, and π_δ is the prior sparsity parameter. The C_δ term is the normalizing constant that ensures the integration of $p(\boldsymbol{\Omega}|\boldsymbol{\delta})$ on M^+ is one. This formulation places a Gaussian mixture prior on the off-diagonal elements of $\boldsymbol{\Omega}$, similar to the spike-and-slab prior used in the Bayesian variable selection literature. By setting $v_1 \gg v_0$, the mixture prior imposes a different strength of shrinkage for elements drawn from the “slab” (v_1) and “spike” (v_0), respectively. This representation allows us to shrink elements in $\boldsymbol{\Omega}$ to 0 if they are small in scale, while not biasing the large elements significantly.

The spike-and-slab formulation of $\boldsymbol{\Omega}$ provides an efficient computation strategy via block Gibbs sampling. However, a main limitation is that parameter estimation can be sensitive to the choice of prior parameters. Unlike the variable selection problem in regression, information on the scale of the elements in the precision matrix typically cannot be easily solicited from domain knowledge. As shown in Wang (2015), there is no analytical relationship between the prior sparsity parameter π_δ and the induced sparsity from the joint distribution. This complexity results from the positive definiteness constraint on the precision matrix. Thus even if the sparsity of the precision matrix is known before fitting the model, additional heuristics and explorations are required to properly select the prior π_δ . Similarly, the induced marginal distribution of the elements in $\boldsymbol{\Omega}$ is intractable as well. The supplementary materials contain an simple illustration of such differences. Thus although

fully Gibbs sampler is attractive for high-dimensional problems, in practice researchers will usually need to evaluate the model fit under multiple prior choices, adding substantially to the computational burden.

3. Fast Deterministic Algorithm for Graph Selection

Consider spike-and-slab priors on Ω as described in the previous section and let the hyperprior on the sparsity parameter to be $\pi_\delta \sim \text{Beta}(a, b)$, the complete-data posterior distribution can be expressed as

$$p(\Omega, \delta, \pi_\delta | X) = p(X | \Omega) p(\Omega | \delta, \nu_0, \nu_1, \lambda) p(\delta | \pi_\delta) p(\pi_\delta | a, b).$$

To perform posterior sampling in the fully Bayesian fashion, the block Gibbs algorithm in Wang (2015) reduces the problem to iteratively sampling from $(p - 1)$ -dimensional multivariate Gaussian distributions for each column of Ω , which can still be computationally expansive for large p or if the sampling needs to be repeated for multiple prior setups. Inspired by the EM approach for variable selection proposed in Ročková and George (2014), we propose a EMGS algorithm to identify the posterior mode of $p(\Omega, \pi_\delta | X)$ directly without the full stochastic search. We iteratively maximize the following objective function

$$\begin{aligned} Q(\Omega, \pi_\delta | \Omega^{(l)}, \pi_\delta^{(l)}) &= E_{\delta | \Omega^{(l)}, \pi_\delta^{(l)}, X} (\log p(\Omega, \delta, \pi_\delta | X) | \Omega^{(l)}, \pi_\delta^{(l)}, X) \\ &= \text{constant} + \frac{n}{2} \log |\Omega| - \frac{1}{2} \text{tr}(X^T X \Omega) \\ &\quad - \frac{1}{2} \sum_{j < k} \omega_{jk}^2 E_{\cdot | \cdot} \left[\frac{1}{\nu_0^2(1 - \delta_{jk}) + \nu_1^2 \delta_{jk}} \right] - \frac{\lambda}{2} \sum_j \omega_{jj} \\ &\quad + \sum_{j < k} \log \left(\frac{\pi_\delta}{1 - \pi_\delta} E_{\cdot | \cdot} [\delta_{jk}] \right) + \frac{p(p-1)}{2} \log(1 - \pi_\delta) \\ &\quad + (a-1) \log(\pi_\delta) + (b-1) \log(1 - \pi_\delta), \end{aligned}$$

where $E_{\cdot | \cdot} [\cdot]$ denotes $E_{\delta | \Omega^{(l)}, \pi_\delta^{(l)}, X} [\cdot]$. This objective function can be easily estimated using ECM algorithm, and the algorithm can naturally handle missing values in the E-step. We present the details of the proposed algorithm in the next subsection and then compare the algorithm with the coordinate ascent algorithm for solving graphical lasso problem in Section 3.2.

3.1. The ECM Algorithm

3.1.1. The E-Step

We start by computing the conditional expectations $E_{\delta | \Omega^{(l)}, \pi_\delta^{(l)}, X} [\delta_{jk}]$ and $E_{\delta | \Omega^{(l)}, \pi_\delta^{(l)}, X} \left[\frac{1}{\nu_0^2(1 - \delta_{jk}) + \nu_1^2 \delta_{jk}} \right]$. This proceeds in the similar fashion as the standard EMVS,

$$E_{\delta_{jk} | \Omega^{(l)}, \pi_\delta^{(l)}, X} [\delta_{jk}] = p_{jk}^* \equiv \frac{a_{jk}}{a_{jk} + b_{jk}}, \quad (3)$$

where $a_{jk} = p(\omega_{jk} | \delta_{jk} = 1) \pi_\delta^{(l)}$ and $b_{jk} = p(\omega_{jk} | \delta_{jk} = 0) (1 - \pi_\delta^{(l)})$, and

$$E_{\delta | \Omega^{(l)}, \pi_\delta^{(l)}, X} \left[\frac{1}{\nu_0^2(1 - \delta_{jk}) + \nu_1^2 \delta_{jk}} \right] = \frac{1 - p_{jk}^*}{\nu_0^2} + \frac{p_{jk}^*}{\nu_1^2} \equiv d_{jk}^*. \quad (4)$$

3.1.2. Modified E-Step With Missing Data

When missing data exist in the data matrix X , the E-step can be easily extended to find the expectation of the missing values as well. In that case, the conditional expectations of δ remains unaffected, and we only need to additionally obtain the expectation for the $X^T X \Omega$ term as

$$\begin{aligned} E_{\delta, X | \Omega} (X^T X \Omega) &= E_{\delta, X | \Omega} \left(\left(\sum_i^n \mathbf{x}_i \mathbf{x}_i^T \right) \Omega \right) \\ &= \left(\sum_i^n E_{\mathbf{x}_{i,m} | \mathbf{x}_{i,o}, \Omega} (\mathbf{x}_i \mathbf{x}_i^T) \right) \Omega, \end{aligned}$$

where $\mathbf{x}_{i,o}$ and $\mathbf{x}_{i,m}$ denote the observed and missing cells in \mathbf{x}_i respectively. Without loss of generality, if we let $\mathbf{x}_i^T = [\mathbf{x}_{i,o}^T, \mathbf{x}_{i,m}^T]$, we know

$$E_{\mathbf{x}_{i,m} | \mathbf{x}_{i,o}, \Omega} (\mathbf{x}_{i,m}) = -\Omega_{oo}^{-1} \Omega_{mo} \mathbf{x}_{i,o}$$

$$E_{\mathbf{x}_{i,m} | \mathbf{x}_{i,o}, \Omega} (\mathbf{x}_i \mathbf{x}_i^T) = E_{\cdot | \cdot} (\mathbf{x}_i) E_{\cdot | \cdot} (\mathbf{x}_i)^T + \begin{pmatrix} \mathbf{0}_{oo} & \mathbf{0}_{om} \\ \mathbf{0}_{mo} & \Omega_{mm}^{-1} \end{pmatrix},$$

where Ω_{oo} , Ω_{mo} , and Ω_{mm} are the corresponding submatrices of Ω .

3.1.3. The CM-Step

After the E-step is performed, the CM-step performs the maximization of (Ω, π_δ) in a coordinate ascent fashion. First, the maximization of π_δ has the close-form solution

$$\pi_\delta^{(l+1)} = (a + \sum_{j < k} \delta_{jk} - 1) / (a + b + p(p-1)/2 - 2). \quad (5)$$

The joint maximization of Ω has no closed-form solution, but if we denote

$$\Omega = \begin{pmatrix} \Omega_{11} & \omega_{12} \\ \omega_{12}^T & \omega_{22} \end{pmatrix} \quad X^T X = \begin{pmatrix} S_{11} & s_{12} \\ s_{12}^T & s_{22} \end{pmatrix},$$

Wang (2015) showed that the conditional distribution of the last column satisfies

$$\omega_{12} \sim \text{Normal}(-C s_{12}, C), \quad C = ((s_{22} + \lambda) \Omega^{-1} + \text{diag}(v_{\delta_{12}}))^{-1},$$

where $v_{\delta_{12}}$ are the inclusion indicators for ω_{12} and

$$\omega_{22} - \omega_{12}^T \Omega_{11}^{-1} \omega_{12} \sim \text{Gamma} \left(1 + \frac{n}{2}, \frac{\lambda + s_{22}}{2} \right).$$

This enables us to perform conditional maximization (Meng and Rubin 1993) for the last column holding the rest of Ω fixed. That is, starting with $\Omega^{(l+1)} = \Omega^{(l)}$, we iteratively permute each column to the last and update it with

$$\omega_{12}^{(l+1)} = ((s_{22} + \lambda) (\Omega_{11}^{(l+1)})^{-1} + \text{diag}(d_{jk}^*))^{-1} s_{12} \quad (6)$$

and

$$\omega_{22}^{(l+1)} = (\omega_{12}^{(l+1)})^T (\Omega_{11}^{(l+1)})^{-1} \omega_{12}^{(l+1)} + \frac{n}{\lambda + s_{22}}. \quad (7)$$

Finally, be iterating between the E-step and the CM-steps until convergence, we obtain our estimator of the posterior mode $\hat{\Omega}$ and $\hat{\pi}_\delta$.

3.2. Connection to the Graphical Lasso

This column-wise update resembles the penalized likelihood approach in frequentist settings. In the graphical lasso algorithm (Mazumder and Hastie 2012), for example, the goal is to minimize the l_1 -penalized negative log-likelihood

$$f(\boldsymbol{\Omega}) = -\log |\boldsymbol{\Omega}| + \text{tr}(\mathbf{S}\boldsymbol{\Omega}) + \|\boldsymbol{\Omega}\|_1,$$

which can be solved via a block coordinate descent that iteratively solves the lasso problem

$$\boldsymbol{\omega}_{12} = \underset{\boldsymbol{\alpha} \in \mathbb{R}^{m-1}}{\text{argmin}} \boldsymbol{\alpha}^T \boldsymbol{\Omega}_{11}^{-1} \boldsymbol{\alpha} + \boldsymbol{\alpha}^T \mathbf{s}_{12} + \lambda \|\boldsymbol{\alpha}\|_1.$$

The updates at each iteration in the EMGS framework solve the optimization problem for $\boldsymbol{\omega}_{12}$ under an adaptive ridge penalty

$$\boldsymbol{\omega}_{12} = \underset{\boldsymbol{\alpha} \in \mathbb{R}^{m-1}}{\text{argmin}} \boldsymbol{\alpha}^T \boldsymbol{\Omega}_{11}^{-1} \boldsymbol{\alpha} + \boldsymbol{\alpha}^T \mathbf{s}_{12} + \sum_{j=1}^{m-1} d_j^* \alpha_j^2.$$

The penalty parameters d_j^* are the corresponding d_{jk}^* estimated from the E-step and are informed by data. That is, instead of choosing a fixed penalty parameter for all precision matrix elements, the EMGS approach learns the element-wise penalization parameter at each iteration based on the magnitude of the current estimated $\boldsymbol{\Omega}$ and the hyperpriors placed on θ . Thus, as long as the signal from data is not too weak, the EMGS procedure can estimate large elements in the precision matrix with much lower bias than graphical lasso, as the adaptive penalties associated with large ω_{jk} are small. To illustrate the diminished bias, we fit the EMGS algorithm to a simple simulated example, where $n = 100$, $p = 10$ and $\boldsymbol{\Omega}$ is constructed by $\omega_{jj} = 1$, and $\omega_{jk} = 0.5$ if $|j - k| = 1$. We fix $v_1 = 100$ and compare the regularization path with various v_0 values with graphical lasso, as shown in Figure 1. This simple example illustrates two main advantages of EMGS. First, it identifies the set of nonzero elements quickly and estimates the partial correlations correctly around 0.5 under all values of v_0 . The clear separation of the truly nonzero edges regardless of v_0 also makes it straightforward to threshold $|\omega_{jk}|$ to recover the true graph structures. Graphical lasso, on the other hand, shrinks the nonzero partial correlations significantly under large penalties, and thus lead to worse graph selection if the tuning parameter is not properly chosen. Second, to select and compare a single model, we also identified the optimal tuning parameter using 5-fold cross-validation for both methods, and it can be seen that the graphical lasso estimator suffers from the weak penalty and contains more noise than using EMGS.

4. ECM Algorithm for Copula Graphical Models

In this section, we extend the framework to non-Gaussian data with Gaussian copulas (Nelsen 1999). Denote the observed data $\mathbf{X} \in \mathbb{R}^{n \times p}$, and each of the p variables could be either continuous, ordinal, or binary. We model each observation as following a Gaussian copula model, that is, there exists a set of monotonically increasing transformations $f = \{f_1, \dots, f_p\}$ such that $\mathbf{Z} = f(\mathbf{X}) \sim \text{Normal}(\mathbf{0}, \mathbf{R})$, where \mathbf{R} is a correlation matrix. Following the same setup as before, we let \mathbf{R} be the induced

correlation matrix from $\boldsymbol{\Omega}$ with the spike-and-slab prior defined as before, that is,

$$\mathbf{R}_{[j,k]} = \boldsymbol{\Omega}_{[j,k]}^{-1} / \sqrt{\boldsymbol{\Omega}_{[j,j]}^{-1} \boldsymbol{\Omega}_{[k,k]}^{-1}}.$$

The explicit form of f is typically unknown, thus we impose no restrictions on the class of marginal transformations. Instead, we follow the extended rank likelihood method proposed in Hoff (2007), decomposing the complete data likelihood into

$$p(\mathbf{X}|\mathbf{R}, f) = \text{Pr}(\mathbf{Z} \in \mathbf{S}|\mathbf{R})p(\mathbf{X}|\mathbf{Z} \in \mathbf{S}, \mathbf{R}, f), \quad (8)$$

where \mathbf{S} is the support of \mathbf{Z} induced by the ranking of \mathbf{X} defined by

$$S_{ij} = [\max\{z_{i'j'} : x_{i'j'} < x_{ij}\}, \min\{z_{i'j'} : x_{i'j'} > x_{ij}\}].$$

Since our goal is to recover the structure in $\boldsymbol{\Omega}$, we can estimate the parameters using only the first part of (8) without estimating the nuisance parameter f . Moreover, since the latent Gaussian variable \mathbf{Z} is constructed to be centered at $\mathbf{0}$, the rank likelihood remains unchanged when multiplying columns of \mathbf{X} by any constant. Thus, inference could be performed without restricting \mathbf{R} to be a correlation matrix (Hoff 2007). In this way, the target function to maximize is the extended rank likelihood function

$$p(\boldsymbol{\Omega}, \boldsymbol{\delta}, \pi_\delta, \mathbf{Z}|\mathbf{X}) = p(\mathbf{Z} \in \mathbf{S}|\boldsymbol{\Omega}, \mathbf{S})p(\boldsymbol{\Omega}|\boldsymbol{\delta})p(\boldsymbol{\delta}|\pi_\delta).$$

This is immediately analogous to the EMGS framework with latent Gaussian variable \mathbf{Z} as additional missing data. That is, we maximize the objective function defined as

$$\begin{aligned} & Q(\boldsymbol{\Omega}, \pi_\delta | \boldsymbol{\Omega}^{(l)}, \pi_\delta^{(l)}) \\ &= E_{\boldsymbol{\delta}, \mathbf{Z} | \boldsymbol{\Omega}^{(l)}, \pi_\delta^{(l)}, \mathbf{X}} (\log p(\boldsymbol{\Omega}, \boldsymbol{\delta}, \pi_\delta, \mathbf{Z}|\mathbf{X}) | \boldsymbol{\Omega}^{(l)}, \pi_\delta^{(l)}, \mathbf{X}) \\ &= \text{constant} + Q_1 - \frac{1}{2} \sum_{j < k} \omega_{jk}^2 E_{\cdot \cdot} \left[\frac{1}{v_0^2(1 - \delta_{jk}) + v_1^2 \delta_{jk}} \right] \\ &\quad - \frac{\lambda}{2} \sum_i \omega_{ii} \\ &\quad + \sum_{j < k} \log \left(\frac{\pi_\delta}{1 - \pi_\delta} E_{\cdot \cdot} [\delta_{jk}] \right) + \frac{p(p-1)}{2} \log(1 - \pi_\delta) \\ &\quad + (a-1) \log(\pi_\delta) + (b-1) \log(1 - \pi_\delta), \end{aligned}$$

where $E_{\cdot \cdot} [\cdot]$ denotes $E_{\boldsymbol{\delta}, \mathbf{Z} | \boldsymbol{\Omega}^{(l)}, \pi_\delta^{(l)}, \mathbf{X}} [\cdot]$, and the only term different from the standard EMGS objective function is

$$\begin{aligned} Q_1 &= E_{\mathbf{Z} | \boldsymbol{\Omega}^{(l)}, \pi_\delta^{(l)}, \mathbf{X}} (\log p(\mathbf{Z}|\boldsymbol{\Omega}, \mathbf{S})) \\ &= \text{constant} + \frac{n}{2} \log |\boldsymbol{\Omega}| - \frac{1}{2} E_{\mathbf{Z} | \boldsymbol{\Omega}^{(l)}, \mathbf{X}} [\text{tr}(\mathbf{Z}^T \mathbf{Z} \boldsymbol{\Omega})]. \end{aligned}$$

Exact computation for this expectation is intractable as $\mathbf{Z}|\mathbf{X}$ is a Gaussian random matrix where each row is conditionally Gaussian and the within column ranks are fixed by \mathbf{S} . Alternatively, posterior samples of \mathbf{Z} are easy to obtain from the conditional truncated Gaussian distribution (Hoff 2007), so we can adopt stochastic variants of the EM algorithm (Wei and Tanner 1990; Delyon, Lavielle, and Moulines 1999; Nielsen 2000; Levine and Casella 2001). We present one such algorithm in the subsequent subsection.

4.1. The SAE-Step for Non-Gaussian Variables

Among the many variations of the EM with stochastic approximation, we discuss estimation steps using stochastic approximation EM (SAEM) algorithm (Delyon, Lavielle, and Moulines 1999). SAEM calculates the E-step at each iteration as a weighted average of the current objective function and new stochastic samples using a decreasing sequence of weights for the stochastic averages, in a similar fashion as simulated annealing. In the stochastic E-step, we compute an additional term $Q(\boldsymbol{\Omega}^{(l)}) = E_{\mathbf{Z}|\boldsymbol{\Omega}^{(l)}, \mathbf{X}}[\mathbf{Z}^T \mathbf{Z}]$ as

$$Q(\boldsymbol{\Omega}^{(l)}) = (1 - t_l)Q(\boldsymbol{\Omega}^{(l)}) + \frac{t_l}{B_l} \sum_{b=1}^{B_l} \mathbf{Z}_{(b)}^T \mathbf{Z}_{(b)},$$

where t_l is an decreasing step-size sequence such that $\sum t_l = \infty$, $\sum t_l^2 < \infty$, and B_l is the number of stochastic samples drawn at each iteration. The rank constrained Gaussian variables can be drawn using the same procedure described in Hoff (2007).

The CM-step then proceeds as before, except that the empirical cross-product matrix \mathbf{S} is replaced by its expectation $Q(\boldsymbol{\Omega}^{(l)})$. For the numerical examples in this article, we set fixed B_l and $t_l = 1/l$. Other weighting schemes could also be explored and may yield different rates of convergence.

5. Incorporating Edge-Wise Informative Priors

The exchangeable beta-binomial prior discussed so far assumes no prior structure on $\boldsymbol{\Omega}$ and prior sparsity controlled by a single parameter for all off-diagonal elements. For many problems in practice, informative priors may exist for pairwise interactions of the variables. For example, Peterson et al. (2013) infer cellular metabolic networks based on prior information in the form of reference network structures. Bu and Lederer (2017) improve estimation of brain connectivity network by incorporating the distance between regions of the brain. In problems with small sample sizes, such prior information can help algorithms identify the high probability edges more quickly and provide more interpretable model. More generally, we can consider a situation where certain groupings exist among variables. For example, when the variables represent log sales of p products on the market, one might expect that the products within the same brand are more likely to be more strongly correlated. If we define a fixed index function $g_j \in \{1, \dots, G\}, j \in \{1, \dots, p\}$, where G denotes the total number of groups, we can modify the prior into

$$\begin{aligned} p(\boldsymbol{\Omega}|\boldsymbol{\delta}) &= C_{\boldsymbol{\delta}}^{-1} \prod_{j < k} \text{Normal} \left(\omega_{jk} | 0, \frac{\nu_{\delta_{jk}}^2}{\tau_{g_j g_k}} \right) \\ &\quad \times \prod_j \text{Exp}(\omega_{jj} | \lambda/2) \mathbf{1}_{\Omega \in M^+} \quad (9) \\ p(\boldsymbol{\delta}|\pi_{\boldsymbol{\delta}}) &\propto C_{\boldsymbol{\delta}} \prod_{j < k} \pi_{\boldsymbol{\delta}}^{\delta_{jk}} (1 - \pi_{\boldsymbol{\delta}})^{1 - \delta_{jk}} \\ p(\boldsymbol{\tau}) &= \prod_{g < g'} \text{Gamma}(a_{\tau}, b_{\tau}). \end{aligned}$$

The block-wise rescaling parameter $\tau_{g_j g_k}$ of the variance parameter allows us to model within- and between-block elements of

$\boldsymbol{\Omega}$ adaptively with different scales. This is particularly useful in applications where block dependence structures have different strengths. Take the example of sales of products for example. Products within the same brand or category are more likely to be conditional dependent, yet the within group sparsity and the scale of the off-diagonal elements may differ for different brands. In the special case where the full edge-level prior probabilities of connection are known, as considered by Peterson et al. (2013) and Bu and Lederer (2017), we can also equivalently let $G = P$ and parameterize $p(\boldsymbol{\tau})$ with the edge-specific priors.

The ECM algorithm discussed above only requires minor modifications to include the additional scale parameter so that the penalties for each block are allowed to vary (e.g., Ishwaran and Rao 2003; Wakefield, De Vocht, and Hung 2010). The new objective function could be similarly estimated with ECM algorithm by including this additional update in the CM-step

$$\tau_{gg'}^{(l+1)} = \frac{a_{\tau} - 1 + \frac{1}{2} \sum_{j < k} \mathbf{1}_{j,k,g,g'} \mathbf{1}_{j,k,g,g'}}{b_{\tau} + \frac{1}{2} \sum_{j < k} \omega_{jk}^2 d_{jk}^* \mathbf{1}_{j,k,g,g'}}, \quad (10)$$

where $\mathbf{1}_{j,k,g,g'} = 1$ if $g_j = g, g_k = g'$, or $g_j = g', g_k = g$. To illustrate the behavior of this block rescaled prior, we simulate data with $n = 200, p = 60$, with the precision matrix to be block diagonal with three equal-sized blocks. We simulate the three block sub-matrices of $\boldsymbol{\Omega}$ to correspond to random graphs with sparsity 0.4, as described in Section 7. Figure 2 shows the effect of the structured prior. It can be seen that the estimated $1/\hat{\tau}_{gg'}$ are much larger where $g = g'$, which leads to weaker shrinkage effects for within cluster cells. Accordingly the resulting graph using the structured prior shows fewer false positives for the off-diagonal blocks, and better discovery of the true positives within blocks.

6. Posterior Summary of the ECM Output

One of the main computational advantage of the ECM approach over stochastic search is that the posterior mode is fast to obtain. Thus it provides a more efficient alternative to experimenting multiple choices of priors with full MCMC, as discussed before. In practice, we fix ν_1 to be a large constant and vary the choice of ν_0 to reflect different levels of shrinkage on the off-diagonal elements of $\boldsymbol{\Omega}$ that are close to 0. Intuitively, a larger ν_0 increases the probability of small parameters being drawn from the spike distribution and thus leads to sparse models. By fitting a sequence of ν_0 , we can create regularization plots, for example, Figure 1, similar to that used in penalized regression literature to visually examine the influence of the prior choices. Choosing a single tuning parameter ν_0 is possible with standard model selection criterion, such as AIC (Akaike 1998), BIC (Schwarz 1978), RIC (Lysen 2009), StARS (Liu, Roeder, and Wasserman 2010), etc., or K -fold cross-validation using the average log-likelihood of the validation sets. In the rest of the article, we select a single tuning parameter ν_0 using 5-fold cross-validation. In the case of non-Gaussian data or data with missing values, the likelihood on test data can be evaluated by the average of the expected covariance $\frac{1}{m} \sum_{v_0} E_{\mathbf{X}_{\text{test}}|\mathbf{X}_{\text{train}}, v_0} (\mathbf{X}_{\text{test}}^T \mathbf{X}_{\text{test}})$ under the sequence of m tuning parameters. This term can be easily calculated by plugging in the test data in the E-step of the algorithm. It is worth noting that since the mixture of Gaussian

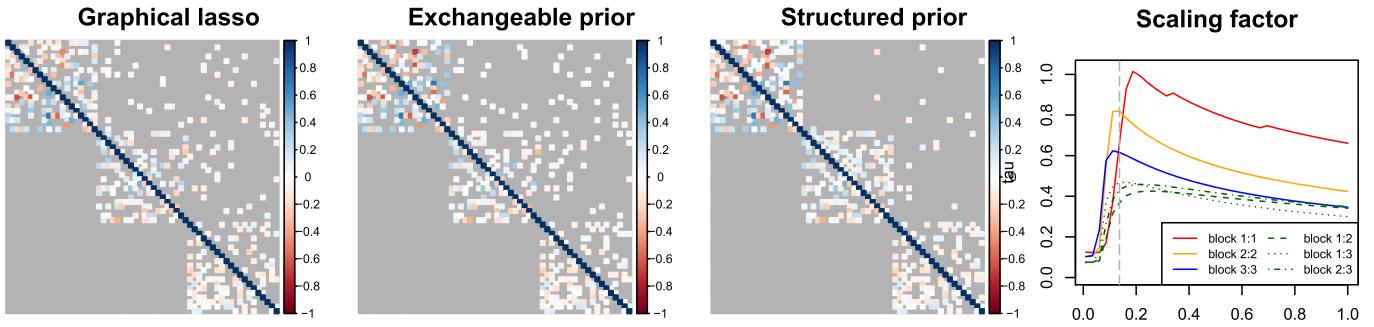


Figure 2. Comparing the estimated and true precision matrix using graphical lasso, EMGS with exchangeable prior, and with structured prior for block-wise rescaling. In each plot of the precision matrix comparison, the upper triangle shows the estimated matrix and the lower triangle shows the true precision matrix. All the tuning parameters are selected by cross-validation. The presented edges are further thresholded to have the same number of edges compared to the true graph. The forth plot shows the change of $1/\hat{\tau}_{gg'}$ over different choices of v_0 . The blocks are labeled 1 to 3 from top left to bottom right.

prior does not lead to exact sparsity, in scenarios where graph structure is of direct interest, we further determine the graph structure by thresholding the off-diagonal elements, $|\omega_{jk}|$, as the posterior inclusion probability p_{jk}^* conditional on ω_{jk} is a monotone function of ω_{jk} .

7. Simulation

We follow a similar simulation setup to Mohammadi et al. (2017) with different graph structures. We compare the performance of our method with graphical lasso for Gaussian data and graphical lasso with nonparanormal transformation (Liu, Lafferty, and Wasserman 2009), and the rank-based extension proposed in Xue and Zou (2012) for non-Gaussian data. We consider the following sparsity patterns in our simulation:

- AR(1): A graph with $\sigma_{jk} = 0.7^{|j-k|}$.
- AR(2): A graph with $\omega_{jj} = 1$, $\omega_{j,j-1} = \omega_{j-1,j} = 0.5$, and $\omega_{j,j-2} = \omega_{j-2,j} = 0.25$, and $\omega_{jk} = 0$ otherwise.
- Random: A graph in which the edge set E is randomly generated from independent Bernoulli distributions with probability 0.2 and the corresponding precision matrix is generated from $\Omega \sim W_G(3, I_p)$.
- Cluster: A graph in which the number of clusters is $\max\{2, [p/20]\}$. Each cluster has the same structure as a random graph. The corresponding precision matrix is generated from $\Omega \sim W_G(3, I_p)$.

We simulate data with sample size $n \in \{100, 200, 500\}$, and of dimension $p \in \{50, 100, 200\}$, using the each types of precision matrices above that are rescaled to have unit variances. We generate both Gaussian and non-Gaussian data for each configuration. For the non-Gaussian case, we perform the marginal transformation of the latent Gaussian variables so that the variables follow a marginal distribution of $\text{Poisson}(\theta)$, with $\theta = 10$ or 2.

For each generated graph, we fit our ECM algorithm with a sequence of 40 increasing v_0 's, and fix $v_1 = 100$, $\lambda = 1$, and $a = b = 1$. We select the final v_0 using 5-fold cross-validation. We also select the tuning parameter for graphical lasso using cross-validation (GL-CV). We then evaluate the bias of EMGS and graphical lasso estimator of precision matrices compared to the truth in terms of the matrix Frobenius norm, $\|\hat{\Omega} - \Omega\|_F$.

$\sqrt{\sum_j \sum_k |\hat{\omega}_{jk} - \omega_{jk}|^2}$. Because of the excess biased induced by a single penalty parameter, cross-validation tend to choose small penalties for graphical lasso, leading to massive false positives in edge discovery. Thus to allow a fair comparison, we compare the area under the ROC curve (AUC) by increasingly thresholding elements in $\hat{\Omega}$ obtained by cross-validation for both EMGS and graphical lasso. Besides selecting tuning parameter by cross-validation for graphical lasso and the nonparanormal transformed estimator, we also consider $\hat{\Omega}$ selected using two popular model selection criterion: rotation information criterion (GL-RIC) (Lysen 2009), and stability approach (GL-STARS) (Liu, Roeder, and Wasserman 2010). For the copula graphical model, we also compare the rank-based extension proposed in Xue and Zou (2012) of graphical lasso (GL-rank) with the tuning parameter selected with cross-validation.

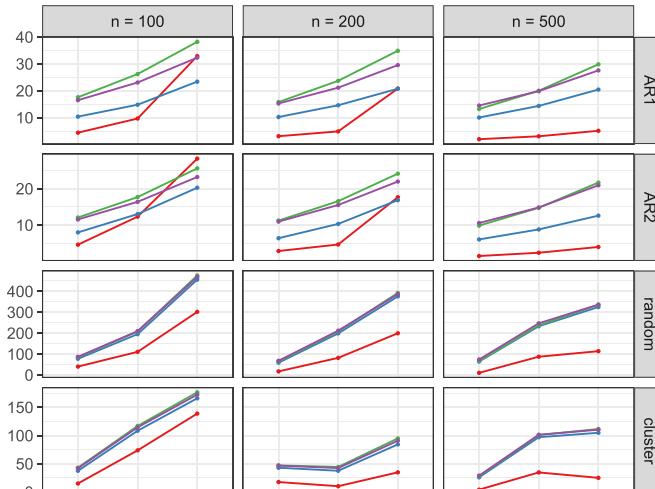
The simulation results are summarized in Figures 3 and 4. Less bias in parameter estimation are indicated by smaller F -norm values and better graph learning is indicated by larger AUC values. In almost all cases of our simulation study, we observe significantly reduced biases in the estimator from EMGS estimators, as well as better graph selection performance in most cases. We also include additional comparisons in the supplementary materials that examine the bias in matrix spectral norms, the F_1 -score for graphical lasso estimators at the selected penalty levels, as well as the F_1 -score when all estimators are thresholded to have the correct number of edges.

All computation are conducted in the R statistical programming environment (R Core Team 2018). The graphs are simulated using the R packages BDgraph (Mohammadi and Wit 2015) and tmvtnorm (Wilhelm and Manjunath 2015). EMGS is implemented with the Rcpp package (Eddelbuettel and François 2011). The graphical lasso estimation is implemented with the R package huge (Zhao et al. 2012) and glasso (Friedman, Hastie, and Tibshirani 2018). Visualizations are created with ggplot2 (Wickham 2016) and corplot (Wei and Simko 2017). The AUC values are calculated with the ROCR package (Sing et al. 2005).

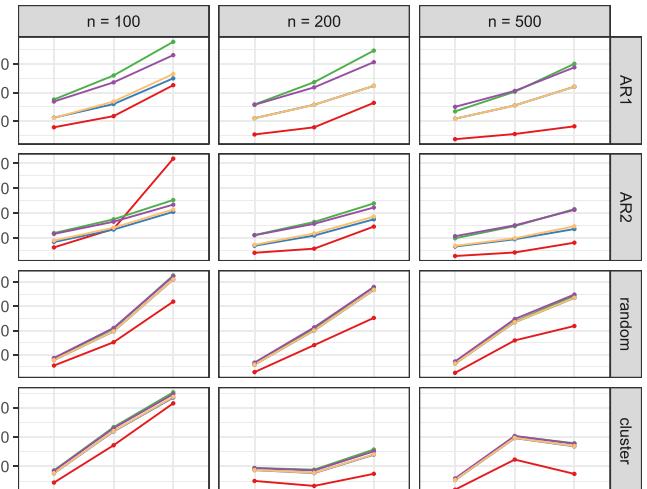
8. Traffic on the Burke Gilman Trail

In this section, we consider graph estimation and prediction for the hourly traffic on the Burke Gilman Trail in Seattle. We use

Average bias in F-norm: the Gaussian case



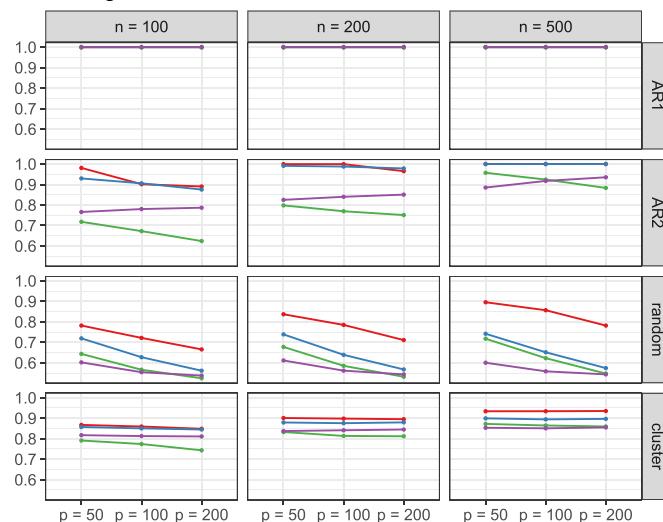
Average bias in F-norm: the copula case



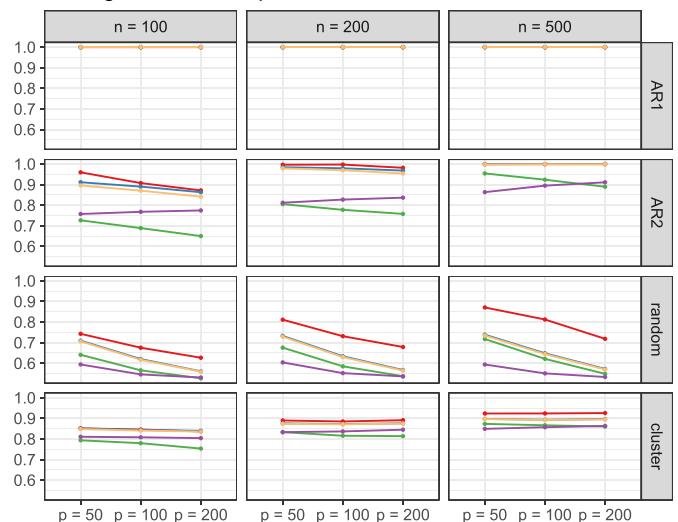
Estimator — EMGS — GL-CV — GL-RIC — GL-StARS — GL-rank

Figure 3. Comparing estimation of the precision matrix for both the Gaussian and Gaussian copula case under different simulation setups. Five estimators are considered: the proposed method (EMGS), Gaussian and nonparanormal graphical lasso with penalty selected by cross-validation (GL-CV), RIC (GL-RIC), stability approach (GL-StARS), and rank-based extension of graphical lasso proposed in Xue (2012) selected by cross-validation for the copula case (GL-rank). EMGS shows lower bias in almost all cases.

Average AUC: the Gaussian case



Average AUC: the copula case



Estimator — EMGS — GL-CV — GL-RIC — GL-StARS — GL-rank

Figure 4. Comparing estimation of the graph structure for both the Gaussian and Gaussian copula case under different simulation setups. EMGS shows higher AUC in almost all cases.

the hourly counts of bikes and pedestrians traveling on the trail through north of NE 70th Street using data from the Seattle Open Data program.¹

The data are captured by sensors that detect both bikes and pedestrians, and their directions of travel. At each hour, the sensors record four counts of travelers: by bike or foot, and toward north or south. We used all the data from 2014 that contain $n = 365$ observations of $24 \times 4 = 96$ measurements. We first performed a log transformation on the raw counts, and subtracted the hourly average from the log counts. The data

are reformatted from the original format with the R package `reshape2` (Wickham 2007).

We estimated the joint distribution of the 96 measurements using EMGS with both the beta-binomial prior and the group-wise structured priors, with 4 groups defined by the mode of travel/direction pairs. Figure 5 shows the estimated graphs and the induced covariance matrices. Graphical lasso estimates many edges with small ω_{jk} , while EMGS allows us to pick out large ω_{jk} , especially those that correspond to the edges between the number of pedestrians traveling within the same hour in opposite directions, and the number of bikes traveling in adjacent hours in the same direction during morning and

¹<http://www.seattle.gov/tech/initiatives/open-data/>

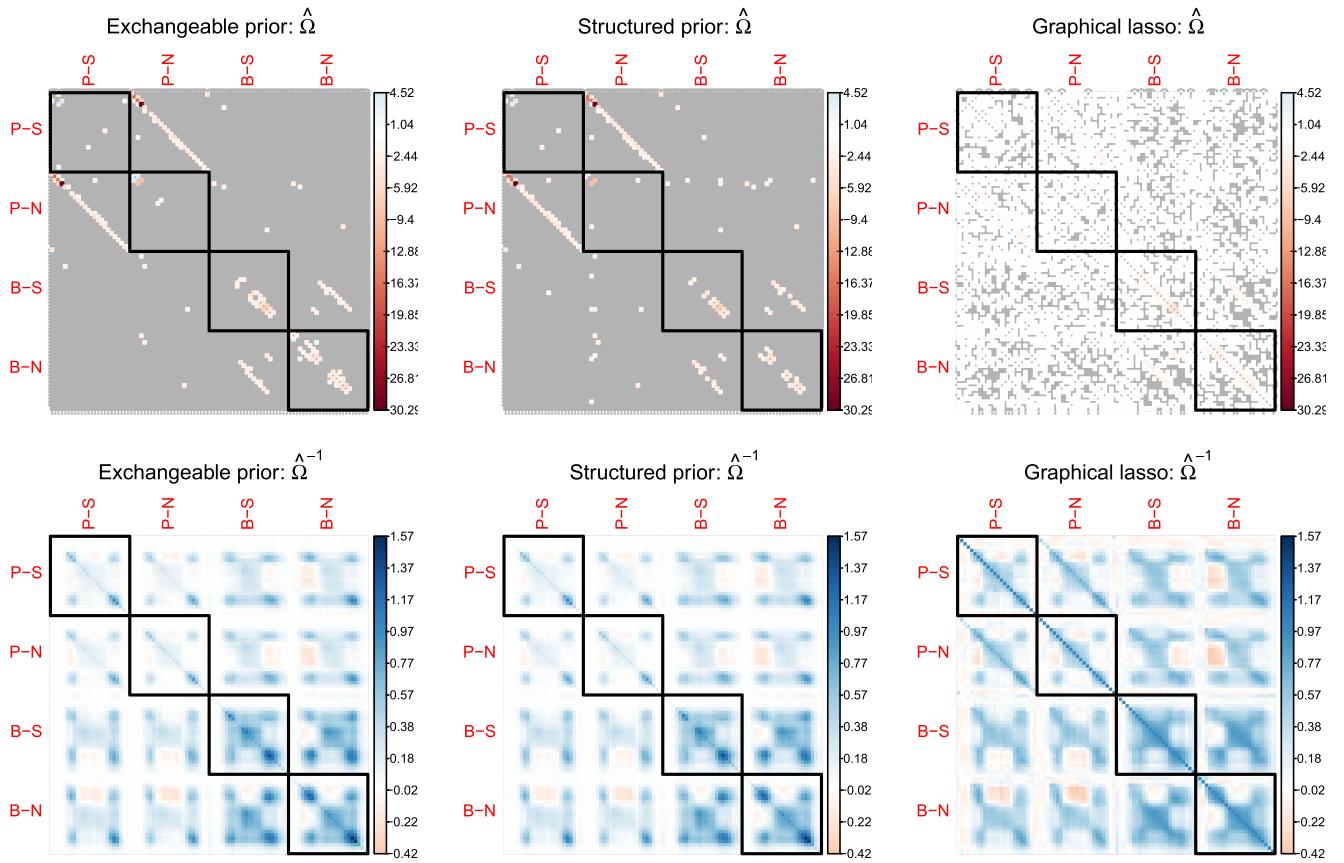


Figure 5. Comparing the estimated precision matrices from cross-validation. The blocks correspond to travel mode and direction pairs. From upper left to lower right: southbound pedestrians, northbound pedestrians, southbound bikes, and northbound bikes. Within each block, the entries correspond to 24 hourly intervals starting from midnight. Top row: estimated covariance matrix. Edges with less than 0.5 probability of being from the slab distributions in EMGS output, and exact zeros in graphical lasso output are marked with gray color. Bottom row: estimated precision matrix with highlighted graph selection.

Table 1. Average and standard deviation of the mean squared errors from 100 cross-validation experiments, with the best performance highlighted in bold.

	EMGS			
	Exchangeable	Structured	GLasso	Empirical
Average MSE	0.2828	0.2809	0.4262	0.4602
Standard deviation of the MSEs	0.0052	0.0050	0.0064	0.0096

afternoon commute hours. In this analysis, the structured priors lead to a slightly more concentrated set of entries, but both priors lead to similar graph estimation for EMGS. We also compare the performance of predicting missing values using $\hat{\Omega}$, by randomly removing half of the measurements on half of the days. The missing observations can be imputed by the EMGS algorithm described in Section 3, and similarly we can estimate $\hat{\Omega}$ by either the empirical covariance matrix or from graphical lasso using only the observed variables. We compare the predictive performance by the mean squared error defined as $\frac{1}{n_{\text{miss}}} \sum_{i,j} (X_{ij} - \hat{X}_{ij})^2$.

Intuitively, predictions based on penalized estimators that are over shrunk toward zero is likely to increase bias, while with little penalization, the estimated covariance matrix is more likely to be noisy, as shown in Figure 5. Table 1 shows the average MSE and their standard deviations using different estimators over 100 replications, and it confirms the improved

prediction performance from EMGS compared to graphical lasso.

9. Symptom Structure in Verbal Autopsies

In this section, we use EMGS to learn the latent dependence structure among symptoms reported on VA surveys. VA surveys collect information about a deceased person's health history through an interview with caregivers or family members of the decedent. VAs are widely used in countries without full-coverage civil registration and vital statistics systems. About 2/3 of deaths worldwide occur in such settings (Horton 2007). VA data consist primarily of binary indicators of symptoms and conditions leading to the death (e.g., Did the decedent have a fever? Was there pain in the lower belly?).

Several algorithms have been proposed to assign causes of death using such binary input (Byass et al. 2012; Serina et al. 2015; McCormick et al. 2016), but these algorithms typically assume that the binary indicators are independent. We use data from the Physicians Health Metrics Research Consortium (Murray et al. 2011). We created 107 variables from the binary questions in the dataset of 7841 adults using the R package openVA (Li, McCormick, and Clark 2019), and removed the variables with more than 50% of values missing, leaving us with 90 indicators. There are many missing

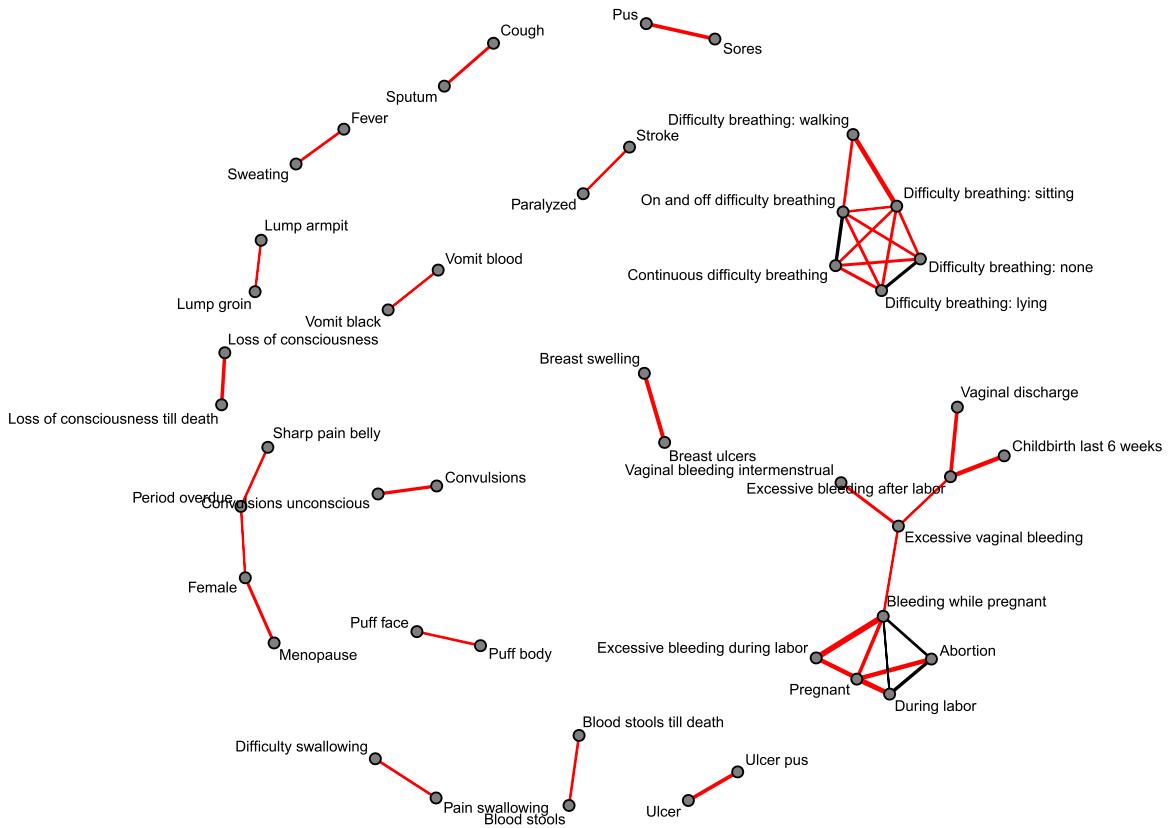


Figure 6. Estimated edges between the indicators in the VA dataset. The width of the edges is proportional to the value of $|\omega_{jk}|$. Red edges correspond to negative values of ω_{jk} , or positive partial correlations. Black edges correspond to positive values of ω_{jk} , or negative partial correlations.

values even after reducing the number of indicators, so there is only one observation with answers for all 90 indicators. This high proportion of missing data makes it difficult to directly apply different types of rank-based estimators for the latent precision matrix that only uses complete observations. Instead, we focus on exploration of the joint distribution of the binary variables under the latent Gaussian framework described in Section 4. We first rescale the dataset by the marginal means of the indicators to remove the different levels of prevalence among the symptoms. We then apply the EMGS algorithm to the rescaled dataset with the same hyperpriors used in Section 7, and select the final v_0 using cross-validation. The resulting conditional dependence graph with 46 indicators and 42 edges is shown in Figure 6, where several main symptom pairs (e.g., fever and sweating, stroke and paralysis, etc.) and symptom groups (e.g., indicators related to pregnancy) are discovered, indicating the existence of some symptom clusters that are strongly dependent in the dataset. Further incorporation of the ECM framework into a classification framework could improve accuracy over existing methods for automatic cause-of-death assignment. The visualization of the symptom network is made with R package network (Butts 2008).

10. Discussion

We propose a deterministic approach for graphical model estimation that builds upon the recently proposed class of

spike-and-slab prior for precision matrices. By drawing the connection between the conditional maximization updates under the spike-and-slab prior and the graphical lasso algorithm, we illustrate that EM type algorithm can be used to efficiently obtain posterior modes of the precision matrix under adaptive penalization. It also allows us to build richer class of models that incorporate prior information and extend to copula graphical models. The computational speed of the EGMS algorithm allows us to explore multiple prior choices without fitting many time-consuming MCMC chains. However, it also comes at the price of two potential limitations. First, characterization of posterior uncertainty is nontrivial due to the deterministic nature of the algorithm. As in Ročková and George (2014), one may choose to fit a Bayesian model “locally” from the posterior mode obtained by the ECM procedure, though this may still be challenging in high-dimensional problems. Another limitation is that like the EM algorithm, ECM algorithm also converges only to local modes, thus the precision matrix initialization is critical. In this article, we used the same initialization as the P-Glasso algorithm described in Mazumder and Hastie (2012). Other heuristics for initialization and warm start may also be explored. Finally, multimodal posteriors are common with spike-and-slab priors. The proposed method could be extended to introduce perturbations in the algorithm, possibly drawing from the variable selection literature (see, e.g., Ročková and George 2014; Ročková 2018).

Replication code for the numerical examples in this article is available at <https://github.com/richardli/EMGS>.

Supplementary Materials

The supplementary materials provide additional details on data and simulation results.

Acknowledgments

We would like to thank Jon Wakefield, Sam Clark, Johannes Lederer, Adrian Dobra, Daniela Witten, and Matt Taddy for helpful discussions and feedback.

Funding

The authors gratefully acknowledge grants SES-1559778 and DMS-1737673 from the National Science foundation, and grants from the National Institute of Child Health and Human Development (NICHD), NIH funding: K01 HD078452 and R01 HD086227.

References

Akaike, H. (1998), "Information Theory and an Extension of the Maximum Likelihood Principle," in *Selected Papers of Hirotugu Akaike*, New York: Springer, pp. 199–213. [\[771\]](#)

Bu, Y., and Lederer, J. (2017), "Integrating Additional Knowledge Into Estimation of Graphical Models," arXiv no. 1704.02739. [\[771\]](#)

Butts, C. T. (2008), "network: A Package for Managing Relational Data in R," *Journal of Statistical Software*, 24, 1–36. [\[775\]](#)

Byass, P., Chandramohan, D., Clark, S. J., D'Ambruoso, L., Fottrell, E., Graham, W. J., Herbst, A. J., Hodgson, A., Hounton, S., Kahn, K., and Krishnan, A. (2012), "Strengthening Standardised Interpretation of Verbal Autopsy Data: The New InterVA-4 Tool," *Global Health Action*, 5, 19281. [\[774\]](#)

Cai, T. T., Zhang, C. H., and Zhou, H. H. (2010), "Optimal Rates of Convergence for Covariance Matrix Estimation," *Annals of Statistics*, 38, 2118–2144. [\[767\]](#)

Delyon, B., Lavielle, M., and Moulines, E. (1999), "Convergence of a Stochastic Approximation Version of the EM Algorithm," *Annals of Statistics*, 27, 94–128. [\[770,771\]](#)

Deshpande, S. K., Rockova, V., and George, E. I. (2017), "Simultaneous Variable and Covariance Selection With the Multivariate Spike-and-Slab Lasso," arXiv no. 1708.08911. [\[768\]](#)

Dobra, A., Lenkoski, A., and Rodriguez, A. (2011), "Bayesian Inference for General Gaussian Graphical Models With Application to Multivariate Lattice Data," *Journal of the American Statistical Association*, 106, 1418–1433. [\[767\]](#)

Eddelbuettel, D., and François, R. (2011), "Rcpp: Seamless R and C++ Integration," *Journal of Statistical Software*, 40, 1–18. [\[772\]](#)

Fan, J., Feng, Y., and Wu, Y. (2009), "Network Exploration via the Adaptive Lasso and Scad Penalties," *The Annals of Applied Statistics*, 3, 521. [\[767\]](#)

Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation With the Graphical Lasso," *Biostatistics*, 9, 432–441. [\[767\]](#)

——— (2010), "Applications of the Lasso and Grouped Lasso to the Estimation of Sparse Graphical Models," Technical Report, pp. 1–22. [\[767\]](#)

——— (2018), "glasso: Graphical Lasso: Estimation of Gaussian Graphical Models," R Package Version 1.10. [\[772\]](#)

George, E. I., and McCulloch, R. E. (1993), "Variable Selection via Gibbs Sampling," *Journal of the American Statistical Association*, 88, 881–889. [\[767\]](#)

Hoff, P. D. (2007), "Extending the Rank Likelihood for Semiparametric Copula Estimation," *The Annals of Applied Statistics*, 1, 265–283. [\[770,771\]](#)

Horton, R. (2007), "Counting for Health," *Lancet*, 370, 1526. [\[774\]](#)

Ishwaran, H., and Rao, J. S. (2003), "Detecting Differentially Expressed Genes in Microarrays Using Bayesian Model Selection," *Journal of the American Statistical Association*, 98, 438–455. [\[771\]](#)

Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C., and West, M. (2005), "Experiments in Stochastic Computation for High-Dimensional Graphical Models," *Statistical Science*, 20, 388–400. [\[767\]](#)

Lauritzen, S. L. (1996), *Graphical Models* (Vol. 17), New York: Clarendon Press. [\[767\]](#)

Lenkoski, A., and Dobra, A. (2011), "Computational Aspects Related to Inference in Gaussian Graphical Models With the G-Wishart Prior," *Journal of Computational and Graphical Statistics*, 20, 140–157. [\[767\]](#)

Levine, R. A., and Casella, G. (2001), "Implementations of the Monte Carlo EM Algorithm," *Journal of Computational and Graphical Statistics*, 10, 422–439. [\[770\]](#)

Li, Z. R., McCormick, T., and Clark, S. (2019), "openVA: Automated Method for Verbal Autopsy," R Package Version 1.0.8. [\[774\]](#)

Liu, H., Lafferty, J., and Wasserman, L. (2009), "The Nonparanormal: Semiparametric Estimation of High Dimensional Undirected Graphs," *Journal of Machine Learning Research*, 10, 2295–2328. [\[772\]](#)

Liu, H., Roeder, K., and Wasserman, L. (2010), "Stability Approach to Regularization Selection (StARS) for High Dimensional Graphical Models," in *Advances in Neural Information Processing Systems*, pp. 1432–1440. [\[771,772\]](#)

Lukemire, J., Kundu, S., Pagnoni, G., and Guo, Y. (2017), "Bayesian Joint Modeling of Multiple Brain Functional Networks," arXiv no. 1708.02123. [\[767\]](#)

Lysen, S. (2009), "Permuted Inclusion Criterion: A Variable Selection Technique," Publicly Accessible Penn Dissertations, p. 28. [\[771,772\]](#)

Mazumder, R., and Hastie, T. (2012), "The Graphical Lasso: New Insights and Alternatives," *Electronic Journal of Statistics*, 6, 2125. [\[767,770,775\]](#)

McCormick, T. H., Li, Z. R., Calvert, C., Crampin, A. C., Kahn, K., and Clark, S. J. (2016), "Probabilistic Cause-of-Death Assignment Using Verbal Autopsies," *Journal of the American Statistical Association*, 111, 1036–1049. [\[774\]](#)

Meinshausen, N., and Bühlmann, P. (2006), "High-Dimensional Graphs and Variable Selection With the Lasso," *The Annals of Statistics*, 34, 1436–1462. [\[767\]](#)

Meng, X.-L., and Rubin, D. B. (1993), "Maximum Likelihood Estimation via the ECM Algorithm: A General Framework," *Biometrika*, 80, 267–278. [\[767,769\]](#)

Mohammadi, A., Abegaz, F., van den Heuvel, E., and Wit, E. C. (2017), "Bayesian Modelling of Dupuytren Disease by Using Gaussian Copula Graphical Models," *Journal of the Royal Statistical Society, Series C*, 66, 629–645. [\[767,772\]](#)

Mohammadi, A., and Wit, E. C. (2015), "BDgraph: An R Package for Bayesian Structure Learning in Graphical Models," arXiv no. 1501.05108. [\[772\]](#)

Murray, C. J., Lopez, A. D., Black, R., Ahuja, R., Ali, S. M., Baqui, A., Dandona, L., Dantzer, E., Das, V., Dhingra, U., and Dutta, A. (2011), "Population Health Metrics Research Consortium Gold Standard Verbal Autopsy Validation Study: Design, Implementation, and Development of Analysis Datasets," *Population Health Metrics*, 9, 27. [\[774\]](#)

Nelsen, R. B. (1999), *An Introduction to Copulas*, Lecture Notes in Statistics (Vol. 139), New York: Springer-Verlag. [\[770\]](#)

Nielsen, S. F. (2000), "The Stochastic EM Algorithm: Estimation and Asymptotic Results," *Bernoulli*, 6, 457–489. [\[770\]](#)

Peterson, C. B., Stingo, F. C., and Vannucci, M. (2015), "Joint Bayesian Variable and Graph Selection for Regression Models With Network-Structured Predictors," *Statistics in Medicine*, 35, 1017–1031. [\[767\]](#)

Peterson, C., Vannucci, M., Karakas, C., Choi, W., Ma, L., and Meletić-Savatić, M. (2013), "Inferring Metabolic Networks Using the Bayesian Adaptive Graphical Lasso With Informative Priors," *Statistics and Its Interface*, 6, 547. [\[767,771\]](#)

R Core Team (2018), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing. [\[772\]](#)

Rocková, V. (2018), "Particle EM for Variable Selection," *Journal of the American Statistical Association*, 113, 1684–1697. [\[775\]](#)

Ročková, V., and George, E. I. (2014), "EMVS: The EM Approach to Bayesian Variable Selection," *Journal of the American Statistical Association*, 109, 828–846. [\[767,769,775\]](#)

Rothman, A. J., Bickel, P. J., Levina, E., and Zhu, J. (2008), "Sparse Permutation Invariant Covariance Estimation," *Electronic Journal of Statistics*, 2, 494–515. [\[767\]](#)

Roverato, A. (2002), "Hyper Inverse Wishart Distribution for Non-decomposable Graphs and Its Application to Bayesian Inference for Gaussian Graphical Models," *Scandinavian Journal of Statistics*, 29, 391–411. [\[767\]](#)

Schwarz, G. (1978), "Estimating the Dimension of a Model," *The Annals of Statistics*, 6, 461–464. [\[771\]](#)

Serina, P., Riley, I., Stewart, A., Flaxman, A. D., Lozano, R., Mooney, M. D., Luning, R., Hernandez, B., Black, R., Ahuja, R., and Alam, N. (2015), "A Shortened Verbal Autopsy Instrument for Use in Routine Mortality Surveillance Systems," *BMC Medicine*, 13, 1. [\[774\]](#)

Sing, T., Sander, O., Beerenwinkel, N., and Lengauer, T. (2005), "ROCR: Visualizing Classifier Performance in R," *Bioinformatics*, 21, 7881. [\[772\]](#)

Wakefield, J., De Vocht, F., and Hung, R. J. (2010), "Bayesian Mixture Modeling of Gene-Environment and Gene-Gene Interactions," *Genetic Epidemiology*, 34, 16–25. [\[771\]](#)

Wang, H. (2012), "Bayesian Graphical Lasso Models and Efficient Posterior Computation," *Bayesian Analysis*, 7, 867–886. [\[767\]](#)

——— (2015), "Scaling It Up: Stochastic Search Structure Learning in Graphical Models," *Bayesian Analysis*, 10, 351–377. [\[767,768,769\]](#)

Wang, H., and Li, S. Z. (2012), "Efficient Gaussian Graphical Model Determination Under G-Wishart Prior Distributions," *Electronic Journal of Statistics*, 6, 168–198. [\[767\]](#)

Wei, G. C., and Tanner, M. A. (1990), "A Monte Carlo Implementation of the EM Algorithm and the Poor Man's Data Augmentation Algorithms," *Journal of the American Statistical Association*, 85, 699–704. [\[770\]](#)

Wei, T., and Simko, V. (2017), "R Package 'corrplot': Visualization of a Correlation Matrix," Version 0.84. [\[772\]](#)

Wickham, H. (2007), "Reshaping Data With the Reshape Package," *Journal of Statistical Software*, 21, 1–20. [\[773\]](#)

——— (2016), *ggplot2: Elegant Graphics for Data Analysis*, New York: Springer-Verlag. [\[772\]](#)

Wilhelm, S., and Manjunath, B. G. (2015), "tmvtnorm: Truncated Multivariate Normal and Student t Distribution," R Package Version 1.4-10. [\[772\]](#)

Witten, D. M., Friedman, J. H., and Simon, N. (2011), "New Insights and Faster Computations for the Graphical Lasso," *Journal of Computational and Graphical Statistics*, 20, 892–900. [\[767\]](#)

Xue, L. (2012), "Regularized Learning of High-Dimensional Sparse Graphical Models," PhD Thesis, University of Minnesota. [\[773\]](#)

Xue, L., and Zou, H. (2012), "Regularized Rank-Based Estimation of High-Dimensional Nonparanormal Graphical Models," *The Annals of Statistics*, 40, 2541–2571. [\[772\]](#)

Yin, J., and Li, H. (2011), "A Sparse Conditional Gaussian Graphical Model for Analysis of Genetical Genomics Data," *The Annals of Applied Statistics*, 5, 2630. [\[767\]](#)

Yuan, M., and Lin, Y. (2007), "Model Selection and Estimation in the Gaussian Graphical Model," *Biometrika*, 94, 19–35. [\[767\]](#)

Zhao, T., Liu, H., Roeder, K., Lafferty, J., and Wasserman, L. (2012), "The Huge Package for High-Dimensional Undirected Graph Estimation in R," *Journal of Machine Learning Research*, 13, 1059–1062. [\[772\]](#)