



## Bayesian Structure Learning in Undirected Gaussian Graphical Models: Literature Review with Empirical Comparison

Lucas Vogels, Reza Mohammadi, Marit Schoonhoven & Ş. İlker Birbil

To cite this article: Lucas Vogels, Reza Mohammadi, Marit Schoonhoven & Ş. İlker Birbil (2024) Bayesian Structure Learning in Undirected Gaussian Graphical Models: Literature Review with Empirical Comparison, Journal of the American Statistical Association, 119:548, 3164-3182, DOI: [10.1080/01621459.2024.2395504](https://doi.org/10.1080/01621459.2024.2395504)

To link to this article: <https://doi.org/10.1080/01621459.2024.2395504>



© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 10 Oct 2024.



[Submit your article to this journal](#)



Article views: 2091



[View related articles](#)



[View Crossmark data](#)



Citing articles: 3 [View citing articles](#)

# Bayesian Structure Learning in Undirected Gaussian Graphical Models: Literature Review with Empirical Comparison

Lucas Vogels , Reza Mohammadi , Marit Schoonhoven , and Ş. İlker Birbil 

Faculty of Economics and Business, University of Amsterdam, Amsterdam, Netherlands

## ABSTRACT

Gaussian graphical models provide a powerful framework to reveal the conditional dependency structure between multivariate variables. The process of uncovering the conditional dependency network is known as structure learning. Bayesian methods can measure the uncertainty of conditional relationships and include prior information. However, frequentist methods are often preferred due to the computational burden of the Bayesian approach. Over the last decade, Bayesian methods have seen substantial improvements, with some now capable of generating accurate estimates of graphs up to a thousand variables in mere minutes. Despite these advancements, a comprehensive review or empirical comparison of all recent methods has not been conducted. This article delves into a wide spectrum of Bayesian approaches used for structure learning and evaluates their efficacy through a comprehensive simulation study. We also demonstrate how to apply Bayesian structure learning to a real-world dataset and provide directions for future research. This study gives an exhaustive overview of this dynamic field for newcomers, practitioners, and experts.

## ARTICLE HISTORY

Received July 2023  
Accepted August 2024

## KEYWORDS

Bayesian model selection; Covariance selection; Link prediction; Markov chain Monte Carlo; Markov random fields

## 1. Introduction



One can depict conditional dependencies between a large number of variables using undirected graphical models. Using data to recover the structure of a graphical model is called *structure learning*. There are many applications of undirected graphical models. In biology, they are used to recover gene networks (Bhadra and Mallick 2013; Li, Craig, and Bhadra 2019a; Chandra, Mueller, and Sarkar 2022); in neuroscience, graphical models illustrate the connectivity of the brain (Hinne et al. 2014; Dyrba et al. 2020); in economics, they map relationships between purchases of customers (Giudici and Castelo 2003); in finance, they discover how risks of financial institutions are related (Wang 2015); and, in psychology, they map relationships between psychological variables (Epskamp et al. 2018; Waldorp and Marsman 2022).


Within structure learning, Bayesian methods offer several advantages over frequentist approaches. First, Bayesian methods integrate prior beliefs, which can enhance the precision of learning the structure. For example, one can have an idea about the risk factors of a disease before seeing any patient data. Second, Bayesian methods estimate the entire probability distribution of parameters, whereas frequentist methods only provide point estimates. This allows Bayesian methods to incorporate model uncertainty. Despite these advantages, frequentist methods in structure learning are more popular than Bayesian methods due to their computational efficiency and simplicity. However, in the last decade, new Bayesian methods were proposed that offer computational efficiency and accuracy, even for large-scale

graphs. Meanwhile, the development of several software packages makes Bayesian structure learning accessible and practical; see, for example, the R packages BDgraph (Mohammadi, Wit, and Dobra 2022), ssgraph (Mohammadi 2022), BGGM (Williams and Mulder 2019) and baygel (Smith, Arashi, and Bekker 2023a).

Despite the recent developments in Bayesian structure learning, there is no comprehensive review or empirical comparison of all methods. For newcomers to the field, this review provides a clear introduction. For practitioners, we discuss typical challenges when applying Bayesian methods to real data and show the benefits of Bayesian structure learning in practice. For experts, we provide an in-depth overview of Bayesian methods coupled with an extensive simulation study.

In this article, we consider general undirected graphs. Moreover, we assume Gaussian data, that is, we work with Gaussian graphical models (GGMs). GGMs are the most widely researched undirected graphical models due to their simplicity and variety of applications. The focus of this review is on Bayesian methods in GGMs. Besides, we do include a short overview of relevant frequentist methods as they form the starting point for some of their Bayesian counterparts. Moreover, the simulation study contains a frequentist approach to show how Bayesian methods perform relative to the better-known frequentist alternative. We focus on the Bayesian methods of the last two decades since these contributed most to the progress in Bayesian structure learning. GGMs form the foundation of several sub-fields that we briefly touch upon too.

**CONTACT** Lucas Vogels  [l.f.vogels@uva.nl](mailto:l.f.vogels@uva.nl)  Faculty of Economics and Business, University of Amsterdam, Amsterdam, Netherlands.

 Supplementary materials for this article are available online. Please go to [www.tandfonline.com/r/JASA](http://www.tandfonline.com/r/JASA).

© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Section 2 provides a general introduction to frequentist and Bayesian structure learning in GGMs. Section 3 provides a review of the recent literature, and Section 4 contains a simulation study comparing the performance of seven state-of-the-art Bayesian algorithms and one frequentist algorithm. In Section 5, we apply Bayesian methods to a gene dataset and discuss typical challenges that arise in real-world applications of Bayesian structure learning. We end with a short overview of related sub-fields and give future perspectives and recommendations in Section 6.

## 2. Structure Learning

This section introduces structure learning. We start with the problem and notation in Section 2.1, discuss frequentist approaches in Section 2.2, and end with an introduction to Bayesian structure learning in Section 2.3.

### 2.1. Problem and Notation

Let  $(X_1, \dots, X_p)$  be a vector of random variables, and  $X_{-ij}$  denote this vector without the variables  $X_i$  and  $X_j$  with  $i, j = 1, \dots, p$  and  $i \neq j$ . We say that the variables  $X_i$  and  $X_j$  are conditionally independent, when

$$P(X_i|X_j, X_{-ij}) = P(X_i|X_{-ij}).$$

In other words, when the values of the variables  $X_{-ij}$  are given, knowing the value of  $X_j$  does not change the probability distribution of  $X_i$ . When two variables  $X_i$  and  $X_j$  are conditionally (in)dependent, they are also referred to as partially (un)correlated (Lauritzen 1996).

One can depict the conditional dependencies between any pair of variables in an undirected graph  $G = (V, E)$ . Here, every node in the node set  $V = \{1, \dots, p\}$  corresponds to a random variable, and the edge set  $E$  is defined by  $E = \{(i, j) \in V \times V : i < j \text{ and } X_i \text{ and } X_j \text{ are conditionally dependent}\}$ . We slightly abuse notation by saying  $(i, j) \in G$  when we mean  $(i, j) \in E$ . We denote the true graph with  $G^* = (V, E^*)$ . In general, we do not know the true graph. We only observe a sample of  $n$  observations of the variables  $X_1, \dots, X_p$ . Let  $x_{li}$  denote the value of the random variable  $X_i$  for observation  $l$  with  $i = 1, \dots, p$  and  $l = 1, \dots, n$ . We denote all observations of the  $i$ th variable as  $\mathbf{X}_i = (x_{1i}, \dots, x_{ni})^T$ . Lastly,  $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_p)$  denotes the  $n \times p$  matrix containing all observations of all variables. We assume that all observations are drawn from a multivariate normal distribution with mean  $\mathbf{0}$  and an unknown  $p \times p$  covariance matrix  $\Sigma^*$ . That is,  $(x_{l1}, x_{l2}, \dots, x_{lp}) \sim \mathcal{N}(\mathbf{0}, \Sigma^*)$  for all  $l = 1, \dots, n$ .

The problem of finding the graph  $G^*$  was introduced by Dempster (1972) and coined as covariance selection. It is also referred to as structure learning or graphical model determination. The goal of structure learning is to use the sample  $\mathbf{X}$  to infer the properties of  $G^*$ . Let  $\mathbf{K}^* = \Sigma^{*-1}$  with elements  $k_{ij}^*$ ,  $i, j = 1, \dots, p$ , be the true and unknown precision matrix or concentration matrix. The precision matrix is a helpful tool in unraveling the structure of the true graph, since

$$\rho_{ij} = -\frac{k_{ij}^*}{\sqrt{k_{ii}^* k_{jj}^*}}, \tag{1}$$

where  $\rho_{ij}$  is the partial correlation between variables  $X_i$  and  $X_j$ . For a proof, see Lauritzen (1996). Due to (1), we obtain the following relationship between  $\mathbf{K}^*$  and  $G^*$ :

$$k_{ij}^* = 0 \iff (i, j) \notin G^*. \tag{2}$$

### 2.2. Frequentist Structure Learning

This section gives a short overview of frequentist methods and their connections to their Bayesian counterparts. Most frequentist methods in structure learning for GGMs obtain an estimate  $\hat{G} = (V, \hat{E})$  of the true graph  $G^*$  by either solving multiple regressions or employing the penalized likelihood (Liang and Jia 2024, chap. 2.2). The regression-based approaches start with the observation that we can express each variable  $X_i$  in terms of the other variables  $X_{-i}$  as follows:

$$X_i = \sum_{j \neq i} \beta_{ij} X_j + \epsilon_i, \quad i = 1, \dots, p, \tag{3}$$

where  $\epsilon_i$  is uncorrelated with  $X_{-i}$  and  $\beta_{ij} = \frac{k_{ij}^*}{k_{ii}^*}$  (Lauritzen 1996). Finding the regression coefficients  $\beta_{ij}$  in (3) is thus enough for structure learning, since  $\beta_{ij} = 0 \iff k_{ij}^* = 0 \iff (i, j) \notin G^*$ . In their neighborhood selection method Meinshausen and Bühlmann (2006) use the Lasso regression by Tibshirani (1996) to obtain estimates for all regression coefficients by solving a sequence of optimization problems given by

$$\min_{\beta_{ij}, j \neq i} \|\mathbf{X}_i - \sum_{j \neq i} \beta_{ij} \mathbf{X}_j\|^2 + \lambda \sum_{j \neq i} |\beta_{ij}|, \quad i = 1, \dots, p, \tag{4}$$

where the first term of the objective function is the residual term and the second term (with the regularization parameter  $\lambda \geq 0$ ) is the  $\ell_1$ -norm regularization term enabling sparse solutions. An edge  $(i, j)$  is then included, if  $\beta_{ij} \neq 0$ . Zhao and Yu (2006) show that this approach is consistent, that is, it recovers the true graph with probability going to 1 as  $n \rightarrow \infty$ . Similar approaches have been proposed by Yuan (2010), Sun and Zhang (2012), and Liu and Wang (2017), each using different methods to find the coefficients  $\beta_{ij}$  in (3). The so-called pseudo-likelihood can be derived from (3) and is used by several Bayesian methods discussed in Sections 3.1.1 and 3.1.2. Moreover, the Bayesian regression approaches discussed in Section 3.3 also use (3), but use Bayesian methods to estimate the regression coefficients.

The penalized likelihood approach starts with the observation that the log-likelihood can be written as

$$\log P(\mathbf{X}|\mathbf{K}) = C + \frac{n}{2} \log |\mathbf{K}| - \frac{n}{2} \text{tr}(\mathbf{S}\mathbf{K}),$$

where  $C$  is a constant,  $\mathbf{S} = \frac{1}{n} \mathbf{X}^T \mathbf{X}$  is the sample covariance matrix, and  $\text{tr}(\mathbf{A})$  and  $|\mathbf{A}|$  are the trace and determinant of a square matrix  $\mathbf{A}$ , respectively. Penalized likelihood methods aim to find the maximum likelihood estimator (MLE), that is, the positive definite precision matrix  $\mathbf{K}$  that maximizes the log-likelihood. However, this MLE is generally not sparse. Penalized likelihood methods therefore add a penalty function  $p(\mathbf{K})$  as a regularization term and minimize the negative penalized log-likelihood:

$$\min_{\mathbf{K} \in \mathcal{K}} -\frac{n}{2} \log |\mathbf{K}| + \frac{n}{2} \text{tr}(\mathbf{S}\mathbf{K}) + \lambda p(\mathbf{K}), \tag{5}$$

where  $\mathcal{K}$  is the set of all  $p \times p$  positive definite matrices. In the case of  $\ell_1$ -norm regularization, the penalty term is given by  $p(\mathbf{K}) = \sum_{i,j,i \neq j} |k_{ij}|$  and the optimization problem (5) becomes a convex program that Yuan and Lin (2007) solve using the interior point algorithm (Vandenberghe, Boyd, and Wu 1998). More efficient is the block coordinate descent approach introduced by Banerjee, El Ghaoui, and d'Aspremont (2008) and improved by Friedman, Hastie, and Tibshirani (2008), who coined it the graphical Lasso (glasso) algorithm. To this day, the glasso algorithm remains the most popular in structure learning and is therefore included in the simulation study in Section 4. Performance bounds of the graphical lasso estimate  $\hat{\mathbf{K}}$  have been proposed in the Frobenius norm (Rothman et al. 2008, Theorem 1), in the element-wise  $\ell_\infty$  norm (Ravikumar et al. 2011, Corollary 1), spectral norm (Ravikumar et al. 2011, eq. (43b)) and element-wise  $\ell_1$  norm Maathuis et al. (2019, Theorem 14.1.2). Moreover, Ravikumar et al. (2011, Theorem 2) prove that the graph  $\hat{G}$  inferred by  $\hat{\mathbf{K}}$  converges in probability to the true graph  $G^*$ . Several other penalty functions  $p(\mathbf{K})$  have been proposed, see, for example, Fan, Feng, and Wu (2009). The tuning parameter  $\lambda$  directly impacts the sparsity of the estimate  $\hat{\mathbf{K}}$ . Several strategies have been proposed for the choice of  $\lambda$ , including the extended Bayesian information criterion (Foygel and Drton 2010), the rotation information criterion (Zhao et al. 2012), and  $k$ -fold cross validation (Bien and Tibshirani 2011). The Bayesian methods discussed in Section 3.1.3 have a direct connection to these frequentist penalized likelihood approaches: their Bayesian estimate of the precision matrix is equal to the one obtained by solving the minimization problem given by (5).

### 2.3. Bayesian Structure Learning

In contrast to frequentist methods that aim to provide an estimate of the true graph, Bayesian methods strive to uncover the probability distribution of the true graph given the sample  $\mathbf{X}$ . This is called the posterior distribution and is denoted with  $P(G|\mathbf{X})$ . This is the probability that, given the data  $\mathbf{X}$ , the true graph is equal to some graph  $G$ . It is given by Bayes' rule:

$$P(G|\mathbf{X}) = \frac{P(\mathbf{X}|G)P(G)}{P(\mathbf{X})}. \quad (6)$$

Here,  $P(\mathbf{X}|G)$  is the likelihood and denotes the probability of observing the data  $\mathbf{X}$  when the true graph is equal to  $G$ . For each  $G$ , the prior  $P(G)$  describes the given belief that  $G$  is the true graph. Lastly,  $P(\mathbf{X})$  is the normalizing constant given by  $P(\mathbf{X}) = \sum_{G \in \mathcal{G}} P(\mathbf{X}|G)P(G)$ , where  $\mathcal{G}$  denotes the set of all graphs with  $p$  nodes. To evaluate the likelihood  $P(\mathbf{X}|G)$  in (6), one could use

$$P(\mathbf{X}|G) = \int_{\mathbf{K} \in \mathcal{K}} P(\mathbf{X}|G, \mathbf{K})P(\mathbf{K}|G)d\mathbf{K}. \quad (7)$$

$P(\mathbf{K}|G)$  is the prior distribution on the precision matrix. For each  $G$  and  $\mathbf{K}$ , it describes the belief that when the true graph is equal to  $G$ , the precision matrix is equal to  $\mathbf{K}$ . Often the integral in (7) is hard to evaluate. Other Bayesian methods therefore focus on the joint posterior distribution of the true graph and precision matrix given by

$$P(G, \mathbf{K}|\mathbf{X}) = \frac{P(\mathbf{X}|\mathbf{K}, G)P(\mathbf{K}|G)P(G)}{P(\mathbf{X})}. \quad (8)$$

Here,  $P(\mathbf{X}|\mathbf{K}, G)$  is the joint likelihood and denotes the probability density of  $n$  samples from a multivariate normal distribution with mean  $\mathbf{0}$  and covariance matrix  $\mathbf{K}^{-1}$ . It is given by

$$P(\mathbf{X}|\mathbf{K}, G) = (2\pi)^{-pn/2} |\mathbf{K}|^{n/2} \exp \left\{ -\frac{n}{2} \text{tr}(\mathbf{K}\mathbf{S}) \right\}. \quad (9)$$

### 3. Review of Methodology

The history of Bayesian structure learning is one of progress. From inefficient approaches restricted to small and decomposable graphs in the 1990s and 2000s to the efficient, accurate, and generally applicable methods of today. This chapter presents the story of this progress. We briefly cover the most relevant approaches before 2010 before we give a detailed outline of all recent methods.

Bayesian approaches in structure learning aim to uncover the posterior distribution  $P(G|\mathbf{X})$  of the graph. For that, one could either evaluate the posterior (6) or the joint posterior (8). For both, the prior distribution  $P(\mathbf{K}|G)$  is required. Early Bayesian structure learning focuses on proposing this prior. Ideally, this prior is conjugate, that is, the prior and posterior follow the same distribution family. Dawid and Lauritzen (1993) define such prior and call it the Hyper Inverse-Wishart (HIW) prior. Using the HIW prior, Giudici and Green (1999) and Carvalho, Massam, and West (2007) design algorithms for Bayesian structure learning that output the posterior for small graphs. The HIW prior, however, is only defined on decomposable graphs, which form a small subset of the whole graph space. Roverato (2002) solves this by extending the HIW to general graphs, a distribution that Atay-Kayis and Massam (2005) later coin the  $G$ -Wishart distribution. Its density is given by

$$P(\mathbf{K}|G) = \frac{1}{I_G(b, \mathbf{D})} |\mathbf{K}|^{\frac{b}{2}-1} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{K}\mathbf{D}) \right\} I(\mathbf{K} \in \mathcal{P}_G), \quad (10)$$

where  $\mathcal{P}_G \subset \mathcal{K}$  denotes the set of symmetric positive definite matrices  $\mathbf{K}$  that have  $k_{ij} = 0$  when  $(i, j) \notin G$ . Here,  $I(\mathbf{K} \in \mathcal{P}_G)$  is an indicator function that is equal to one, if  $\mathbf{K} \in \mathcal{P}_G$ , and zero otherwise. The symmetric positive definite matrix  $\mathbf{D}$  and the scalar  $b > 2$  are called the scale and shape parameters of the  $G$ -Wishart distribution. They are mostly set to  $b = 3$  and  $\mathbf{D} = \mathbf{I}$ , where  $\mathbf{I}$  is the  $p$ -dimensional identity matrix. The normalizing constant

$$I_G(b, \mathbf{D}) = \int_{\mathbf{K} \in \mathcal{P}_G} |\mathbf{K}|^{\frac{b-2}{2}} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{K}\mathbf{D}) \right\} d\mathbf{K} \quad (11)$$

ensures that  $\int_{\mathbf{K} \in \mathcal{P}_G} P(\mathbf{K}|G)d\mathbf{K} = 1$ . Wong, Moffa, and Kuipers (2024) present techniques to evaluate the normalizing constant for certain types of graphs. For general graphs, however, the normalizing constant is troublesome to evaluate. Roverato (2002), Atay-Kayis and Massam (2005), and Lenkoski and Dobra (2011) propose approximations. Using the  $G$ -Wishart prior, the marginal likelihood (7) becomes

$$P(\mathbf{X}|G) = (2\pi)^{-np/2} \frac{I_G(b+n, \mathbf{D} + \mathbf{X}^T \mathbf{X})}{I_G(b, \mathbf{D})}. \quad (12)$$

Using the approximations for the normalizing constants in this ratio, one can approximate  $P(\mathbf{X}|G)$ . Jones et al. (2005) and

Lenkoski and Dobra (2011) both design stochastic search algorithms that in each iteration favor a move to a graph with a high posterior probability  $P(G|\mathbf{X}) \propto P(\mathbf{X}|G)P(G)$ . However, the approximation of the normalizing constants in (12) is unstable and computationally expensive, limiting structure learning to small graphs with less than 20 nodes.

That is why others focus on inference on the joint posterior (8). For this and for inference on the precision matrix in general, it is required to sample from the  $G$ -Wishart distribution (10), which is challenging due to the normalizing constant (11). This led to a string of articles on efficient sampling from the  $G$ -Wishart distribution, see Wang and Li (2012) for a review of those. However, sampling from the  $G$ -Wishart distribution remained computationally expensive.

The period after 2010 saw several discoveries accelerating Bayesian structure learning. They include solutions to the challenges discussed before: computationally efficient approximation of the normalizing constant (11) and computationally efficient  $G$ -Wishart sampling. They also contain entirely new perspectives and approaches, enabling the creation of new methods that push the boundaries of accuracy and computational efficiency. These approaches form the topic of this chapter. We group them into four different categories: Markov chain Monte Carlo (MCMC), Expectation Conditional Maximization (ECM), regression, and hypothesis approaches. The most common Bayesian methods in each category are listed in Table 1. Next, we provide an overview of the workings and theoretical properties of these methods.

### 3.1. MCMC-based Approaches

This section gives a short introduction to Markov chain Monte Carlo (MCMC) sampling before describing the main MCMC-based approaches in Bayesian structure learning.

Most methods in Bayesian structure learning use MCMC sampling, because of the variety of inference it provides. Specifically, when using MCMC sampling, one can approximate the expected value of  $h(G^*)$  for any information function  $h : \mathcal{G} \rightarrow \mathbb{R}$ . For example, one can choose  $h(G)$  equal to 1 when an edge  $e$  is in  $G$ , and 0 otherwise. This will give the posterior edge inclusion probabilities

$$p_e := P(e \in G^* | \mathbf{X}) = E((h(G^*) | \mathbf{X})). \tag{13}$$

Similarly, with other information functions  $h$ , one can evaluate different posterior probabilities. For example, the probability that the true graph will have more than 10 edges, the probability that the true graph is connected, and so on. However, this flexibility comes at a cost: MCMC algorithms are computationally demanding because they iteratively explore large parameter spaces.

Let  $P(\mathbf{X}|\boldsymbol{\theta})$  be the likelihood conditional on some parameter vector  $\boldsymbol{\theta} \in \mathcal{T}$ . Let  $\boldsymbol{\theta}^*$  be the true and unknown parameter vector. We want to perform inference on the posterior  $P(\boldsymbol{\theta}|\mathbf{X})$ . We are interested in obtaining  $E(h(\boldsymbol{\theta}^*)|\mathbf{X}) = \int_{\boldsymbol{\theta}} h(\boldsymbol{\theta})P(\boldsymbol{\theta}|\mathbf{X})d\boldsymbol{\theta}$  for some information function  $h : \mathcal{T} \rightarrow \mathbb{R}$ . MCMC algorithms start at an initial realization of the parameter vector  $\boldsymbol{\theta}^{(0)}$  and then move to  $\boldsymbol{\theta}^{(1)}$ ,  $\boldsymbol{\theta}^{(2)}$ , all the way to  $\boldsymbol{\theta}^{(S)}$  with  $S \in \mathbb{N}$ . We call

**Table 1.** Overview of all the state-of-the-art methods for Bayesian structure learning in Gaussian graphical models.

Category	Method name	Reference	Included in simulation
MCMC on $\mathcal{G} \times \mathcal{K}$	RJ - O	Dobra, Lenkoski, and Rodriguez (2011)	
	RJ - WL	Wang and Li (2012)	
	RJ - CL	Cheng and Lenkoski (2012)	
	RJ - D	Lenkoski (2013)	
	RJ - DCBF	Hinne et al. (2014)	
	RJ - A	Mohammadi, Wit, and Dobra (2022)	
	RJ - WWA	Van den Boom, Beskos, and De Iorio (2022a)	✓
	RJ - UB	Van den Boom et al. (2022b)	
	BD - O	Mohammadi and Wit (2015)	
	BD - A	Mohammadi, Massam, and Letac (2023a)	✓
	SS - O	Wang (2015)	✓
	SS - AT	Atchadé (2019)	
	SS - BC	Jalali, Khare, and Michailidis (2020)*	
MCMC on $\mathcal{G}$	G - Stingo	Stingo and Marchetti (2014)	
	G - MPLRJ	Mohammadi et al. (2023b)*	
	G - MPLBD	Mohammadi et al. (2023b)*	✓
MCMC on $\mathcal{K}$	K - BGLasso1	Wang (2012)	
	K - BGLasso2	Khondker et al. (2013)	
	K - RIW	Kundu, Mallick, and Baladandayuthapani (2018)	
	K - Horseshoe	Li, Craig, and Bhadra (2019a)	✓
	K - LRD	Chandra, Mueller, and Sarkar (2022)*	
	K - BAGR	Smith, Arashi, and Bekker (2023b)*	
	K - Horseshoe-like	Sagar et al. (2024)	
ECM	ECM - EMGS	Li and McCormick (2019)	
	ECM - BAGUS	Gan, Narisetty, and Liang (2019)	✓
	ECM - HL	Sagar et al. (2024)	
Regression	R - BLNRE	Li and Zhang (2017)	
	R - P	Williams et al. (2018)	
Hypothesis	H - LR	Leday and Richardson (2018)	
	H - BGGM	Williams and Mulder (2020)	✓

NOTE: They are listed per category and chronologically. The last column indicates whether the method is included in the simulation study in Section 4. Papers marked with \* were not yet published by September 2024.

the resulting sequence  $(\theta^{(0)}, \dots, \theta^{(S)})$  a Markov chain. For every  $s = 1, \dots, S$ ,  $\theta^{(s)}$  is called a state of the Markov chain. All states are elements of the state space  $\mathcal{T}$ . The transition kernel  $P(x, A)$  denotes the probability that the next state will be part of the set  $A \subset \mathcal{T}$ , given that the current state is  $x \in \mathcal{T}$ . A suitable transition kernel ensures that the chain converges to the so-called stationary distribution. There are three conditions sufficient for this convergence (Tierney 1994, Theorem 1), the most crucial being the balance condition, which is often replaced by a stricter condition called the detailed balance condition, see eq. (1) in Green (1995). Now, if the balance condition holds and if the stationary distribution is equal to the posterior distribution, the distribution of our Markov Chain will get arbitrarily close to the posterior as we increase the chain length  $S$ . This holds no matter the starting position  $\theta^{(0)}$  of our Markov Chain. We can then get arbitrarily close to our desired expected value using

$$E(h(\theta^*)|\mathbf{X}) = \lim_{S \rightarrow \infty} \frac{1}{S} \sum_{s=1}^S h(\theta^{(s)}) \approx \frac{1}{S} \sum_{s=1}^S h(\theta^{(s)}). \quad (14)$$

### 3.1.1. MCMC-based Approaches on the Joint Space

Most of the MCMC-based approaches in Bayesian structure learning create a Markov Chain  $((G^{(0)}, \mathbf{K}^{(0)}), \dots, (G^{(S)}, \mathbf{K}^{(S)}))$  over the joint space of graphs and precision matrices that converges to the joint posterior (8). Using (14) these methods can perform inference on almost all characteristics of both the graph and the precision matrix. We call these methods joint MCMC methods. They come in three types: reversible jump (RJ), birth-death (BD), and spike-and-slab (SS) methods. All joint MCMC methods are listed in Table 2.

RJ methods are the first MCMC-based methods in Bayesian structure learning. They use the  $G$ -Wishart prior given by (10). RJ methods are a variant of Metropolis-Hastings methods that were introduced by Green (1995). Algorithm 1 represents the pseudo-code for this class of methods.

One can define the acceptance probabilities  $\alpha(G', G)$  such that the balance condition holds and the resulting MCMC chain converges to the joint posterior. The first three RJ methods to use this strategy are the RJ-Original (RJ-O) method by Dobra, Lenkoski, and Rodriguez (2011), the RJ-Wang Li (RJ-WL) method by Wang and Li (2012), and the RJ-Cheng Lenkoski (RJ-CL) method by Cheng and Lenkoski (2012). These methods form the foundation for all later reversible jump

---

#### Algorithm 1: Reversible Jump (RJ) algorithm

---

**Input:** Data  $\mathbf{X}$  and an initial graph  $G^{(0)}$ .

**for** state  $s = 1, \dots, S$  **do**

    Set  $G = G^{(s-1)}$ ;

    Propose a new graph  $G'$  by adding or deleting one edge from  $G$ ;

    Accept the proposal with acceptance probability  $\alpha(G', G)$ ;

    Set  $G^{(s)} = G'$  if accepted, set  $G^{(s)} = G$  if rejected;

    Sample  $\mathbf{K}^{(s)}$  from the  $G^{(s)}$ -Wishart distribution;

**Output:** Samples  $(G^{(0)}, \mathbf{K}^{(0)}), \dots, (G^{(S)}, \mathbf{K}^{(S)})$  from the joint posterior (8)

---

methods. Their approach to sample from the  $G$ -Wishart distribution, however, is not scalable to large-scale graphs and lacks a theoretical guarantee of convergence. To address these issues, Lenkoski (2013) introduces an efficient exact sampler from the  $G$ -Wishart distribution. The resulting RJ-Double (RJ-D) method thereby becomes the first joint method to achieve detailed balance, and hence, theoretical convergence. Hinne et al. (2014) combine several existing techniques to improve the overall efficiency in the RJ-Double Continuous Bayes Factor (RJ-DCBF) method. Despite these improvements, RJ methods are still computationally demanding for graphs with more than 50 variables. One of the reasons is that the acceptance probabilities contain a ratio of normalizing constants that is time-consuming to evaluate. Mohammadi, Massam, and Letac (2023a) tackle this by developing a closed-form approximation of the ratio used in the RJ - Approximation (RJ-A) method by Mohammadi and Wit (2019). Due to the approximation, the RJ-A method loses theoretical convergence but gains computational efficiency. The RJ-Weighted Proposal (RJ-WWA) method by Van den Boom, Beskos, and De Iorio (2022a) combines the best of both worlds: theoretical convergence and computational efficiency. It achieves this with two main improvements. First, the RJ-WWA algorithm is more likely to propose graphs with higher acceptance probabilities thereby improving the low acceptance probabilities of earlier algorithms. Second, they introduce a pre-acceptance step, in which the acceptance probabilities are approximated using the computationally

**Table 2.** Overview of joint MCMC methods for Bayesian structure learning in GGMs with their corresponding prior on  $\mathbf{K}$ , type of Markov chain (discrete or continuous), name, and reference. The last column indicates whether an algorithm theoretically converges to the posterior distribution.

Prior on $\mathbf{K}$	Discr/Cont.	Method name	Reference	Conv.	
G-Wishart	Discrete	RJ - O	Dobra, Lenkoski, and Rodriguez (2011)		
		RJ - WL	Wang and Li (2012)		
		RJ - CL	Cheng and Lenkoski (2012)		
		RJ - D	Lenkoski (2013)	✓	
		RJ - DCBF	Hinne et al. (2014)	✓	
		RJ - A	Mohammadi, Wit, and Dobra (2022)		
		RJ - WWA	Van den Boom, Beskos, and De Iorio (2022a)	✓	
		RJ - UB	Van den Boom et al. (2022b)	✓	
		Continuous	BD - O	Mohammadi and Wit (2015)	✓
			BD - A	Mohammadi, Massam, and Letac (2023a)	
Spike-and-slab	Discrete	SS - O	Wang (2015)		
		SS - AT	Atchadé (2019)		
		SS - BC	Jalali, Khare, and Michailidis (2020)*		

NOTE: Papers marked with \* were not yet published by September 2024.

efficient approximation of Mohammadi, Massam, and Letac (2023a). Only when pre-accepted, the exact acceptance probabilities are calculated. The most recent reversible jump method is the RJ-Unbiased (RJ-UB) method by Van den Boom et al. (2022b). The RJ-UB method introduces unbiased estimation to Bayesian structure learning by applying a technique of Glynn and Rhee (2014). It requires the construction of two Markov chains  $(G^{(0)}, \dots, G^{(S)})$  and  $(\bar{G}^{(0)}, \dots, \bar{G}^{(S)})$  that meet at some time  $\tau < \infty$ , that is,  $G^{(s)} = \bar{G}^{(s)}$  for all  $s \geq \tau$ . Now, for any  $q = 1, 2, \dots$ , the quantity

$$\hat{h}_q = h(G^{(q)}) + \sum_{s=q+1}^{\tau} [h(G^{(s)}) - h(\bar{G}^{(s)})] \quad (15)$$

is an unbiased estimator of  $E(h(G^*)|\mathbf{X})$ . To create the two Markov chains, Van den Boom et al. (2022b) use a combination of a sequential Monte Carlo (SMC) method and a particle-independent Metropolis-Hastings method. Van den Boom et al. (2022b) are not the first to use SMC techniques in Bayesian structure learning. Tan et al. (2017) also use an SMC sampler. Their method, however, is computationally cumbersome and comes with no theoretical guarantees.

Mohammadi and Wit (2015) develop a different strategy to tackle the issue of low acceptance probabilities and create a new class of joint MCMC methods: Birth-Death (BD) methods. Their BD-Original (BD-O) method abolishes acceptance probabilities altogether while maintaining theoretical convergence. The BD-O method designs a continuous time Markov chain (Cappé, Robert, and Rydén 2003). This chain spends a continuous time  $W(s)$  in each state  $(G^{(s)}, K^{(s)})$  before moving to the next state. The next state is determined by adding or removing an edge based on so-called birth-death rates that are recalculated at every iteration. Its pseudo-code is shown in Algorithm 2.

---

**Algorithm 2:** Birth-death (BD) algorithm

---

**Input:** Data  $\mathbf{X}$  and an initial graph  $G^{(0)}$ .  
**for** state  $s = 1, \dots, S$  **do**  
    Calculate the birth-death rate  $r_e$  for every edge  $e = (i, j)$ ;  
    Calculate the waiting time  $W(s - 1) = \frac{1}{\sum_e r_e}$ ;  
    Select one edge  $e$  with a probability proportional to its birth-death rate;  
    Obtain  $G^{(s)}$  by flipping edge  $e$  in  $G^{(s-1)}$ ;  
    Sample  $\mathbf{K}^{(s)}$  from the  $G^{(s)}$ -Wishart distribution;

**Output:** Samples  $(G^{(0)}, \mathbf{K}^{(0)}), \dots, (G^{(S)}, \mathbf{K}^{(S)})$  and waiting times  $(W^{(0)}, \dots, W^{(S)})$

---

The BD-O method meets the balance condition and therefore converges to the joint posterior. The desired expected value can now be obtained using

$$E(h(G^*, \mathbf{K}^*)|\mathbf{X}) \approx \frac{1}{W} \sum_{s=1}^S W(s)h(G^{(s)}, \mathbf{K}^{(s)}), \quad (16)$$

where  $W$  is the sum of the waiting times of all states. Per iteration, the BD-O algorithm is computationally costly because it calculates the birth and death rates for all edges. However,

its MCMC chain needs less iterations to converge because it moves to a new graph every iteration. Moreover, the birth and death rates can be calculated in parallel. Just like the acceptance probabilities of the RJ methods, the birth and death rates of the BD-O method contain a ratio of normalizing constants. Using the approximation of this ratio by Mohammadi, Massam, and Letac (2023a), the BD-Approximation (BD-A) method achieves a significant reduction in computational cost, but the balance condition no longer holds, and there is no theoretical convergence.

Despite all the mentioned improvements, RJ and BD methods still suffer from a high computational burden. This has two reasons: first, at every iteration, a new precision matrix needs to be sampled from the  $G$ -Wishart distribution, which remains, despite the direct sampler of Lenkoski (2013), computationally costly. Second, at every MCMC iteration, the graph changes by at most one edge. This leads to poor mixing and convergence of the MCMC chain, especially for large-scale graphs. Both problems are overcome by a new class of joint MCMC methods called Spike-and-Slab (SS) methods. The key element of SS methods is their prior on the precision matrix, called the SS prior (Tadesse and Vannucci 2021). An SS prior samples each element  $k_{ij}$  of the precision matrix individually, such that  $k_{ij}$  is likely to be close to zero for  $(i, j) \notin G$  (the spike) and likely to be away from zero for  $(i, j) \in G$  (the slab). The elements of the precision matrix are still dependent, since the resulting matrix  $\mathbf{K}$  needs to be positive definite. Moreover, like the  $G$ -Wishart prior, SS priors generally contain an intractable normalizing constant. SS priors allow for the design of efficient algorithms that sample an entirely new graph at each iteration. The Markov chain's mixing is therefore greatly improved compared to the RJ and BD methods. Algorithm 3 presents the pseudo-code for the class of SS algorithms. In the SS-Original (SS-O) method, Wang (2015) introduces SS

---

**Algorithm 3:** Spike-and-slab (SS) algorithm

---

**Input:** Data  $\mathbf{X}$  and an initial graph  $G^{(0)}$ .  
**for** state  $s = 1, \dots, S$  **do**  
    Sample a new precision matrix  $\mathbf{K}^{(s)}$ ;  
    For every edge  $e = (i, j)$ , calculate the probability  $l_e$ ;  
    Obtain  $G^{(s)}$  by including every edge  $e$  with probability  $l_e$ ;

**Output:** Samples  $(G^{(0)}, \mathbf{K}^{(0)}), \dots, (G^{(S)}, \mathbf{K}^{(S)})$  from the joint posterior (8)

---

priors to Bayesian structure learning. This prior samples off-diagonal elements  $k_{ij}$  from a normal distribution with variance  $v_0$  when  $(i, j) \notin G$  and with variance  $v_1 > v_0$  when  $(i, j) \in G$ . Diagonal elements are sampled from an exponential distribution with parameter  $\eta$ . The SS-Atchadé (SS-AT) method by Atchadé (2019) and the SS-BCONCORD (SS-BC) by Jalali, Khare, and Michailidis (2020) set off-diagonal elements  $k_{ij}$  to zero when  $(i, j) \notin G$ . These so-called *discrete* SS priors were introduced by Talluri, Baladandayuthapani, and Mallick (2014). The SS-AT and SS-BC methods do not use the exact likelihood (9), but instead consider an approximation. These approximations are variants of the pseudo-likelihood approximation (Besag 1975)

given by

$$\prod_{i=1}^p P(\mathbf{X}_i | \mathbf{X}_{-i}, G, \mathbf{K}) \approx P(\mathbf{X} | \mathbf{K}, G), \quad (17)$$

where  $\mathbf{X}_{-i}$  denotes the data matrix  $\mathbf{X}$  with the column  $\mathbf{X}_i$  removed. Both the SS-AT and SS-BC methods motivate the use of this approximation by proving a theoretical property called posterior contraction. This means that the posterior probability of precision matrices  $\mathbf{K}$  that are “far away” from the true precision matrix  $\mathbf{K}^*$  goes to zero, as the number of observations goes to infinity. Or, put differently, the posterior distribution  $P(\mathbf{K} | \mathbf{X})$  is concentrated around the true value  $\mathbf{K}^*$ . The SS-AT and SS-BC methods combine the discrete SS prior with the approximated likelihood to construct computationally efficient MCMC algorithms. On top of that, Jalali, Khare, and Michailidis (2020) prove that their SS-BC method provides a consistent estimate of the graph, that is, as the number of observations goes to infinity, the estimated edge inclusion probabilities  $p_e$  (13) converge (in probability) to 1 for edges in the true graph and to 0 for edges not in the true graph.

### 3.1.2. MCMC-based Approaches on the Graph Space

A benefit of joint MCMC methods is that with the samples  $(G^{(0)}, \mathbf{K}^{(0)}), \dots, (G^{(S)}, \mathbf{K}^{(S)})$ , one can perform inference on the precision matrix. However, one is often just interested in retrieving the structure of the true graph. In that case, obtaining a new precision matrix at every iteration is a computational burden. The class of methods in this section overcomes this burden by creating a Markov chain  $(G^{(0)}, \dots, G^{(S)})$  only over the graph space. We also refer to these types of methods as G-methods.

Stingo and Marchetti (2014) introduce this strategy in their G-Stingo method. If a graph  $G_d$  is decomposable, the pseudo-likelihood approximation (17) becomes exact. Stingo and Marchetti (2014) use this fact to design an MCMC algorithm over the space of decomposable graphs. They then design a framework to generalize this chain to the space of all graphs  $\mathcal{G}$ .

The pseudo-likelihood approximation (17) can also be used to approximate the marginal likelihood  $P(\mathbf{X} | G)$ :

$$\begin{aligned} P(\mathbf{X} | G) &= \int_{\mathbf{K}} P(\mathbf{X} | G, \mathbf{K}) P(\mathbf{K} | G) d\mathbf{K} \\ &\approx \int_{\mathbf{K}} \prod_{i=1}^p P(\mathbf{X}_i | \mathbf{X}_{-i}, \mathbf{K}, G) P(\mathbf{K} | G) d\mathbf{K} \quad (18) \\ &:= \hat{P}(\mathbf{X} | G), \end{aligned}$$

where  $\hat{P}(\mathbf{X} | G)$  is called the marginal pseudo-likelihood (MPL). Leppä-aho et al. (2017) introduce the concept of marginal pseudo-likelihood to GGMs. They use the Wishart prior with scale parameter  $b = p$  and  $D = \mathbf{S}$ . The Wishart prior is equivalent to the G-Wishart prior (10), where  $G$  is the complete graph. Note that this prior uses the data  $\mathbf{X}$ . Such priors are called fractional priors. Leppä-aho et al. (2017) choose this prior because it leads to a closed-form expression of  $\hat{P}(\mathbf{X} | G)$ . They design a hill-climbing algorithm that, at each iteration, moves to a graph with a higher marginal pseudo-likelihood.

The G-MPL-Reversible Jump (G-MPLRJ) method by Mohammadi et al. (2023b) combines the closed-form expression

of  $\hat{P}(\mathbf{X} | G)$  with an MCMC algorithm to sample from the so-called pseudo-posterior  $\hat{P}(G | \mathbf{X}) := \hat{P}(\mathbf{X} | G) P(G) / P(\mathbf{X})$ . Their discrete time MCMC algorithm creates the chain  $(G^{(0)}, \dots, G^{(S)})$ . The pseudo-code resembles Algorithm 1 with two differences. First, no precision matrix is sampled at every iteration. Second, the acceptance probabilities  $\alpha(G', G)$  are calculated using the marginal pseudo-likelihood, greatly improving the computational efficiency. The same work also presents the G-MPL-Birth Death (G-MPLBD) method. Again using the approximation in (18), this method designs a continuous time MCMC similar to Algorithm 2. Both the G-MPLRJ and the G-MPLBD algorithm meet the balance conditions and therefore converge to the pseudo-posterior. Similar to the SS - BC by Jalali, Khare, and Michailidis (2020) (see Section 3.1.1), Mohammadi et al. (2023b) prove two theoretical results: the pseudo-posterior concentrates around the true graph  $G^*$  and the edge inclusion probabilities are consistent.

### 3.1.3. MCMC-based Approaches on the Precision Matrix Space

Methods on the space of precision matrices remove the necessity of sampling a graph at every iteration. They create an MCMC chain  $(\mathbf{K}^{(1)}, \dots, \mathbf{K}^{(S)})$  solely over the space of positive definite  $p \times p$  matrices. The methods are computationally efficient and have appealing theoretical properties. We also refer to these types of methods as K-methods. K-methods do not use a prior on the graph and are, therefore, not able to incorporate a prior belief in the graphical structure. Moreover, since K-methods do not provide MCMC samples in the graph space, they still require some operation to select a graph  $G$  from the posterior samples of the precision matrix. Two examples of such operations are thresholding and credible intervals. The thresholding method (Wang 2012) first computes the edge inclusion probability  $p_e$  of an edge  $e$  and then includes it in the estimate of the graph if  $p_e > \gamma$  where  $\gamma \in (0, 1)$  is some threshold. The credible intervals method (Li, Craig, and Bhadra 2019a) first constructs intervals, for every edge  $(i, j)$ , that contains  $\gamma\%$  of the values  $\{k_{ij}^{(1)}, k_{ij}^{(2)}, \dots, k_{ij}^{(S)}\}$  for some user-defined  $\gamma$ . An edge is then excluded if and only if the interval includes zero.

For K-methods, the prior on the precision matrix  $P(\mathbf{K} | G)$  no longer depends on  $G$  and can therefore be written as  $P(\mathbf{K})$ . We call these priors K-only priors. Most K-only priors are defined by two probability distributions. One for off-diagonal elements and one for diagonal elements. The elements of K-only priors are not independent due to the positive definite constraint and these priors also contain an intractable normalizing constant. Specific choices of K-only priors enable algorithms to be computationally efficient and/or have attractive theoretical properties.

All K-methods share one or more of three useful theoretical properties. First, all K-methods meet the balance condition and therefore converge to the posterior distribution of the precision matrix (Tierney 1994, Theorem 1). Second, most K-methods have a connection with the frequentist graphical lasso algorithm. Specifically, priors of K-methods are often chosen such that the precision matrix that maximizes the posterior  $P(\mathbf{K} | \mathbf{X})$  is equal to the graphical lasso estimate obtained by solving the optimization problem (5) for some regularization parameter  $\lambda$  and penalty

**Table 3.** K-methods for Bayesian structure learning with their corresponding reference, name, K-only prior, and theoretical properties.

Reference	Name	$P(k_{ij})$		Theor. Prop.		
		$i \neq j$	$i = j$	i	ii	iii
Wang (2012)	K - BGLasso1	Double exp.	Exp.	✓	✓	
Khondker et al. (2013)	K - BGLasso2	Double exp.	Exp.	✓	✓	
Kundu, Mallick, and Baladandayuthapani (2018)	K - RIW		Wishart	✓		
Li, Craig, and Bhadra (2019a)	K - Horseshoe	Horseshoe	$\propto 1$	✓		
Chandra, Mueller, and Sarkar (2022)*	K - LRD	NA	NA	✓		✓
Smith, Arashi, and Bekker (2023b)*	K - BAGR	Norm.	Trunc. norm.	✓	✓	
Sagar et al. (2024)	K - Horseshoe-like	Horseshoe-like	$\propto 1$	✓	✓	✓

NOTE: All marginal K-only priors are not equal to the distributions mentioned in the table due to the positive definite constraint. The horseshoe and horseshoe-like distribution are defined in the corresponding articles. The theoretical properties are (i) the Markov chain converges to the posterior distribution, (ii) the maximum a posteriori estimate equals the maximum likelihood estimate (5), and (iii) the posterior concentrates around the true value. Papers marked with \* were not yet published by September 2024.

**Table 4.** ECM methods for Bayesian structure learning with their corresponding reference, name, prior on K, theoretical properties, and output.

Reference	Name	$P(k_{ij} G)$			Theory		Output	
		$(i, j) \in G$	$(i, j) \notin G$	$i = j$	i	ii	$K_{MAP}$	$p_e$
Gan, Narisetty, and Liang (2019)	ECM - BAGUS	DE(0, $v_1$ )	DE(0, $v_0$ )	Exp( $\eta$ )	✓	✓	✓	✓
Li and McCormick (2019)	ECM - EMGS	N(0, $v_1$ )	N(0, $v_0$ )	Exp( $\eta$ )			✓	
Sagar et al. (2024)	ECM - HL		Horseshoe-like	$\propto 1$	✓		✓	

NOTE: DE( $a, b$ ) denotes the double exponential distribution with mean  $a$  and scale parameter  $b$ . N( $a, b$ ) denotes the normal distribution with mean  $a$  and standard deviation  $b$ . Exp( $\eta$ ) is the exponential distribution with rate  $\eta$ . Note that all marginal priors are not equal to the distributions mentioned in the table due to the positive definite constraint. The theoretical properties are i) the MAP estimate is consistent, and ii) the estimate of the graph is consistent. The output is the MAP estimate of the precision matrix ( $K_{MAP}$ ) and/or the edge inclusion probabilities  $p_e$  of (13).

function  $p(\mathbf{K})$ . The third and last theoretical property is the concentration of the precision matrix around the true value as in the SS-BC algorithm by Jalali, Khare, and Michailidis (2020) (see also Section 3.1.1). Table 3 contains all K-methods with their corresponding K-only priors and theoretical properties.

K-methods are not scale-invariant, that is, scaling the observations of one or more variables affects the posterior distribution. K-methods typically deal with this by standardizing the data to have unit variance. Carter, Rossell, and Smith (2023) observe, however, that this standardization can adversely affect inference. Instead, Carter, Rossell, and Smith (2023) propose to perform inference directly on the partial correlations and define a class of priors for which the resulting posterior is scale-invariant.

### 3.2. Expectation Conditional Maximization Approaches

MCMC-based approaches are popular due to the wide variety of inference they can provide using (14). They are, however, computationally demanding because of the large parameter space they explore. Expectation Conditional Maximization (ECM) methods are a more efficient alternative. They allow for the quick computation of the maximum a posteriori (MAP) estimate of the precision matrix but do not provide inference on the rest of the posterior. The three ECM methods in Bayesian structure learning are listed in Table 4.

At every iteration  $s = 1, \dots, S$ , an ECM algorithm renders a new precision matrix  $\mathbf{K}^{(s)}$ . The sequence  $(\mathbf{K}^{(1)}, \mathbf{K}^{(2)}, \dots)$  converges to the MAP estimator. In each iteration two steps are applied: the Expectation step (E-step) and the Conditional-Maximization step (CM-step). In iteration  $s$ , the E-step evaluates  $Q(\mathbf{K}|\mathbf{K}^{(s-1)})$ , defined as the expected value of the log posterior with respect to the current conditional distribution of  $G$  given  $\mathbf{K}$  and  $\mathbf{X}$ . In the CM-step the new precision matrix  $\mathbf{K}^{(s)}$  is set to the

matrix that optimizes  $Q(\mathbf{K}|\mathbf{K}^{(s-1)})$ . The pseudo-code is shown in Algorithm 4.

#### Algorithm 4: Expectation Conditional Maximization (ECM) algorithm

**Input:** Data  $\mathbf{X}$  and an initial precision matrix  $\mathbf{K}^{(0)}$ .

**for** iteration  $s = 1, \dots, S$  **do**

    E-step: evaluate  $Q(\mathbf{K}|\mathbf{K}^{(s-1)})$ ;

    CM-step: set  $\mathbf{K}^{(s)} = \arg \max_{\mathbf{K} \in \mathcal{K}} Q(\mathbf{K}|\mathbf{K}^{(s-1)})$ ;

**Output:** A sequence  $\mathbf{K}^{(0)}, \mathbf{K}^{(1)}, \dots, \mathbf{K}^{(S)}$  that converges to the MAP estimate.

ECM methods in Bayesian structure learning differ in three aspects from each other: their prior, their theoretical properties, and their output. See Table 4 for an overview. First, each ECM method uses a different SS prior (see Section 3.1.1). Second, some methods provide consistent results on the graph and/or the precision matrix. For example, with their ECM-BAGUS method Gan, Narisetty, and Liang (2019) show that the estimate of the precision matrix  $\hat{\mathbf{K}}$  gets arbitrarily close to the true precision matrix  $\mathbf{K}^*$  as the number of observations goes to infinity. To be specific, under certain assumptions and for a constant  $C$ , they show that

$$\|\hat{\mathbf{K}} - \mathbf{K}^*\|_\infty \leq C \sqrt{\frac{\log p}{n}}, \tag{19}$$

where  $\|\cdot\|_\infty$  denotes the  $l_\infty$  norm. The ECM-Horseshoelike (ECM-HL) method by Sagar et al. (2024) proves a similar result. The ECM - BAGUS method is the only ECM method that also provides a consistent result for the graph. Third, ECM methods differ in the output they provide. Like the K-methods of Section 3.1.3, ECM methods require some post-hoc operation to provide inference on the graph. The ECM-BAGUS method is

thereby the only ECM method that outputs the edge inclusion probabilities (13).

A limitation of ECM methods is that they only work for specifically designed priors on the precision matrix. This is overcome by Sagar et al. (2023), who present a method that outputs a MAP estimate of the precision matrix for a large class of priors.

### 3.3. Regression Approaches

In Section 2.2, we saw that structure learning is equivalent to finding the coefficients of  $p$  different regressions. That is, estimating the coefficients  $\beta_{ij}$  in (3) allows for direct estimation of the graphical model, since  $\beta_{ij} = \frac{k_{ij}}{k_{ii}}$  and hence  $\beta_{ij} = 0 \iff (i, j) \notin G$ . Regression methods use a Bayesian approach to find the coefficients  $\beta_{ij}$ , thereby offering a simple approach to Bayesian structure learning. The resulting estimates of the graph and the precision matrix, however, come with limited theoretical guarantees.

There are two regression methods: The R - Bayesian Lasso Neighborhood Regression Estimate (R - BLNRE) method by Li and Zhang (2017) and the R - Projection (R-P) method by Williams et al. (2018). Both put priors on each coefficient  $\beta_{ij}$  and create Markov chains  $(\beta_{ij}^{(1)}, \dots, \beta_{ij}^{(S)})$  for each  $\beta_{ij}$ . These MCMC chains provide approximate samples from the posterior distribution  $P(\beta_{ij}|\mathbf{X})$ . Like MCMC-based approaches on the space of precision matrices (Section 3.1.3), regression methods now use a post-MCMC operation to obtain an estimate of the graph and the precision matrix.

The R-BLNRE method finds the  $\beta_{ij}$  in (3) using the Bayesian lasso, a popular algorithm by Park and Casella (2008). In the Bayesian lasso the prior on each coefficient  $\beta_{ij}$  is set to the double exponential distribution. The mode of the resulting posterior distribution  $P(\beta_{ij}|\mathbf{X})$  equals the frequentist estimate obtained by solving the optimization problems in (4).

The R-P method by Williams et al. (2018) uses the horseshoe prior for the regression coefficients. The same prior is used for the elements of the precision matrix in the K- Horseshoe method by Li, Craig, and Bhadra (2019a). Using so-called projection predictive selection, the samples from the resulting posterior  $P(\beta_{ij}|\mathbf{X})$  are used to output an estimate of the graph and precision matrix.

### 3.4. Hypothesis Approaches

Hypothesis methods circumvent the time-consuming exploration of the model space necessary for MCMC-based methods. Instead, they perform inference on the graph directly. To do so, they first formulate, for every  $(i, j)$ , a null hypothesis  $H_0 : k_{ij} = 0$  and an alternative hypothesis  $H_1 : k_{ij} \neq 0$ . Then, they calculate the Bayes factor in favor of  $H_1$ , given by

$$BF_{ij} = \frac{P(\mathbf{X}|H_1)}{P(\mathbf{X}|H_0)}.$$

This strategy was introduced by Giudici (1995), who derived a closed-form expression of  $BF_{ij}$  when the number of observations is larger than the number of variables, that is,  $n > p$ . The Hypothesis Leday and Richardson (H-LR) method by Leday and Richardson (2018) builds upon this result by deriving a

closed-form expression for  $BF_{ij}$  that also holds for  $p > n$  and is consistent, that is,  $\lim_{n \rightarrow \infty} BF_{ij} = 0$  if  $H_0$  is true, and  $\lim_{n \rightarrow \infty} BF_{ij} = \infty$  if  $H_1$  is true. The H-LR method scales  $BF_{ij}$  to obtain a scale-invariant Bayes factor  $sBF_{ij}$ . By computing  $sBF_{ij}$  for all edges and selecting a threshold  $\gamma$ , one can obtain a graph estimate  $\hat{G}$  by including an edge  $(i, j)$  in  $\hat{G}$  if  $sBF_{ij} > \gamma$ . For a prior on the precision matrix, they use the Wishart distribution, which is equivalent to the G-Wishart distribution  $P(\mathbf{K}|G)$  in (10) when  $G$  is a complete graph. The H-LR method does not use a prior on the graph and is therefore, like the K-methods, not able to incorporate a prior belief on the graphical structure.

The Hypothesis-BGGM (H-BGGM) method by Williams and Mulder (2020) formulates the hypothesis using the partial correlations  $\rho_{ij}$  in (1). The resulting hypotheses are  $H_0 : \rho_{ij} = 0$ ,  $H_1 : \rho_{ij} > 0$ , and  $H_2 : \rho_{ij} < 0$ . They also formulate an unrestricted hypothesis  $H_u : \rho_{ij} \in (-1, 1)$ . These hypotheses allow for the computation of a scale-invariant Bayes factor. The edge exclusion probability  $P((i, j) \notin G)$  can now be calculated as follows:

$$\begin{aligned} P((i, j) \notin G) &= P(H_0|\mathbf{X}) \\ &= \frac{P(\mathbf{X}|H_0)P(H_0)}{P(\mathbf{X}|H_0)P(H_0) + P(\mathbf{X}|H_1)P(H_1) + P(\mathbf{X}|H_2)P(H_2)} \\ &= \frac{BF_{0u}P(H_0)}{BF_{0u}P(H_0) + BF_{1u}P(H_1) + BF_{2u}P(H_2)}. \end{aligned}$$

Here,  $BF_{ku} = \frac{P(\mathbf{X}|H_k)}{P(\mathbf{X}|H_u)} = \frac{P(H_k|\mathbf{X})}{P(H_k)}$  for  $k = 0, 1, 2$ .  $P(H_k)$  is the prior distribution for accepting hypothesis  $H_k$ . The Bayes factor  $BF_{ku}$  is therefore just a ratio of the posterior and prior of the hypotheses  $H_k$ . Using the Wishart prior for the precision matrix, they show that the prior and posterior partial correlations approximately follow normal distributions. The Bayes factors  $BF_{ku}$  can then be computed analytically or by sampling repeatedly from the Wishart distribution and using (1). The Wishart distribution, however, is not able to incorporate prior beliefs on the graph. Therefore, they propose two other priors: the Corrected Wishart prior and the matrix-F prior. The H-BGGM method can also be used for confirmatory hypothesis tests. These kinds of hypothesis tests involve more than one partial correlation; for example,  $H_0 : \rho_{12} > \rho_{34}$  and  $H_1 : \rho_{12} \leq \rho_{34}$ .

## 4. Empirical Comparison

The simulation study in this section contributes to the literature in three ways. First, it includes large-scale instances, that is, instances with a thousand variables. Until now, Bayesian algorithms are only compared on 250 variables or less, an exception being the K-LRD algorithm by Chandra, Mueller, and Sarkar (2022). Second, we report information on how long it takes algorithms to achieve good results. This information is missing in Bayesian literature. Third, and most importantly, never before has such a variety of Bayesian algorithms been compared in a single simulation study. All our results can be reproduced by the specified scripts on our GitHub page.<sup>1</sup> We discuss the simulation setup in Section 4.1 and the results in Section 4.2.

<sup>1</sup> <https://github.com/lucasvogels33/Review-paper-Bayesian-Structure-Learning-in-GGMs>

**Table 5.** Existing and missing empirical comparisons among different classes of Bayesian methods.

	Frequentist	RJ	BD	SS	G-methods	K-methods	ECM	Regression	Hypothesis
Frequentist	████████	missing	2015	2020	missing	2024	2024	2018	2020
RJ		████████	2022a	missing	2014	missing	missing	missing	2018
BD			████████	missing	missing	missing	missing	2018	2020
SS				████████	missing	missing	missing	missing	missing
G-methods					████████	missing	missing	missing	missing
K-methods						████████	2024	2018	missing
ECM							████████	missing	missing
Regression								████████	missing
Hypothesis									████████

**4.1. Simulation Setup**

**Selected algorithms:** In Bayesian structure learning literature, new algorithms tend to be compared to frequentist algorithms, or Bayesian algorithms of similar nature, for example, K-algorithms are compared to K-algorithms, ECM algorithms to ECM algorithms, and so on. We group Bayesian method into seven categories and show in Table 5 that comparisons between different categories are largely missing in the literature.

In this section, we complete most of this table by comparing seven Bayesian algorithms (RJ-WWA, BD-A, SS-O, G-MPLBD, K-Horseshoe, ECM-BAGUS, and H-BGGM) and one frequentist algorithm (glasso). For this simulation study, Bayesian algorithms are chosen based on three criteria: (i) the algorithm demonstrates outstanding performance in its category, (ii) there is publicly accessible working code for the algorithm, and (iii) the algorithm is documented in a published article. No regression algorithm is selected because there is no code available. As a prior for the graph we sample the edges from independent and identically distributed Bernoulli distributions. The resulting prior becomes

$$P(G) = \delta^{|E|} (1 - \delta)^{p(p-1)/2 - |E|}, \tag{20}$$

where  $\delta$  is the prior sparsity of the graph. In literature this prior sparsity is often set to either the uninformative value  $\delta = 0.5$  (Gan, Narisetty, and Liang 2019) or to a function that decreases with  $p$ , for example  $\delta = 2/(p - 1)$  (Van den Boom, Beskos, and De Iorio 2022a). In this simulation study, we opt for a middle ground and set  $\delta = 0.2$ . Supplementary materials S1 and S2 contain the references and all other settings of the algorithms.

**Performance metrics:** A metric should measure the distance between an estimate and the truth. In the case of Bayesian structure learning, however, the true posterior distribution  $P(G|X)$  is unknown. In the absence of a true posterior distribution, it is therefore common in the literature to opt for the second best choice: compare the estimate of the posterior against the true graph  $G^*$ . For lack of a better alternative, we adopt the same strategy too.

We compare the selected algorithms based on the accuracy of the edge inclusion probabilities  $p_e$  defined in (13). Let  $P$  be the matrix with elements  $p_e$ . We refer to  $P$  as the edge inclusion matrix. We use three metrics for the accuracy of the edge inclusion probabilities: AUC,  $Pr^+$ , and  $Pr^-$ . The AUC represents the area under the Receiver Operating Characteristic (ROC) curve. It reflects to what degree the edge inclusion probabilities of links in the true graph are higher than those of links not in the true graph. It ranges from 0 (worst) to 1 (best). An

AUC of 0.5 means the edge inclusion matrix is as good as a random guess. The glasso and K-Horseshoe algorithms do not output edge inclusion probabilities. To calculate the AUC of the glasso algorithm, we solve the minimization problem in (5) for different values of  $\lambda$ . The area under the resulting ROC curve gives us the AUC. For the K-Horseshoe algorithm we use the same strategy as Li, Craig, and Bhadra (2019a), that is, the ROC curve is generated by varying the length of posterior credible intervals from 1% to 99%.

The AUC is a measure for the ranking of the edge inclusion probabilities but does not convey their magnitude. We therefore add two metrics for the magnitude: one measuring the algorithm’s ability to predict the presence of an edge ( $Pr^+$ ), the other the ability to predict the absence of an edge ( $Pr^-$ ). Formally, we have

$$Pr^+ = \frac{1}{|E^*|} \sum_{(i,j) \in G^*} p_e \tag{21}$$

and

$$Pr^- = \frac{1}{|E^{*-}|} \sum_{(i,j) \notin G^*} p_e, \tag{22}$$

where  $|A|$  denotes the number of elements in the set  $A$  and  $E^{*-}$  denotes the set of links that are not in the true graph.  $Pr^+$  ranges from 0 (worst) to 1 (best), whereas  $Pr^-$  ranges from 1 (worst) to 0 (best). As before, these metrics need to be redefined for the glasso and K-Horseshoe algorithms, since these algorithms do not output edge inclusion probabilities, but a point estimate of the true graph. By setting  $p_e = 1$  if  $e \in \hat{G}$ , and zero if  $e \notin \hat{G}$ , we can still calculate  $Pr^+$  and  $Pr^-$ . Note that  $Pr^+$  ( $Pr^-$ ) then becomes the true (false) positive rate of the estimate  $\hat{G}$ . Concerns might be raised about the fairness when comparing the  $Pr^+$  and  $Pr^-$  metrics from a point estimate  $\hat{G}$  with those from the posterior edge inclusion probabilities  $p_e$ . However, in both cases, the  $Pr^+$  and  $Pr^-$  metrics do provide valuable information and are thus included. We do recommend the reader to proceed with caution when directly comparing the  $Pr^+$  and  $Pr^-$  metrics from the glasso and K-horseshoe method with those from other algorithms.

We also measure the computational cost of each algorithm. We first run each Bayesian algorithm until visual inspection shows convergence and denote with  $P$  the final edge inclusion matrix. The computational cost  $T$  of an algorithm is then defined as the time it takes the algorithm to produce an AUC within 0.01 of its final AUC. That is,

$$T = \min(t : |AUC(P_t) - AUC(P)| < \epsilon), \tag{23}$$

**Table 6.** Computational cost  $T$  (23) in seconds of the algorithms for different instances.

$p$	Graph	Density	$n$	glasso	RJ-WWA	BD-A	SS-O	ECM-BAGUS	G-MPLBD	K-Horseshoe	H-BGGM
10	Random	10 %	20	2	1	0	0	1	0	0	0
		10 %	350	2	0	0	0	1	0	0	0
	Cluster	10 %	20	2	1	0	0	1	0	0	0
		10 %	350	2	0	0	0	1	0	0	0
100	Random	1 %	40	3	3623	5313	5	84	4	49	–
		1 %	700	3	147	1835	19	76	2	3	11
		10 %	40	5	2794	9018	10	252	6	32	–
		10 %	700	5	746	6093	12	482	8	18	10
	Cluster	1 %	40	3	620	4496	13	83	4	11	–
		1 %	700	3	159	3013	19	70	1	5	11
		10 %	40	4	3489	10,368	13	186	5	17	–
		10 %	700	3	1569	6465	10	349	6	10	11
1000	Random	0.1 %	60	606	–	–	9063	15,057	624	24,064	–
		0.1 %	1050	354	–	–	4007	8907	492	12,300	–
	Cluster	0.1 %	60	672	–	–	8510	15,710	623	7451	–
		0.1 %	1050	366	–	–	3823	8224	481	12,389	–

NOTE: The values are averages over 16 replications. The “–” entry indicates that an algorithm did not produce results on that instance.

**Table 7.** AUC scores of the algorithms for different instances.

$p$	Graph	Density	$n$	glasso	RJ-WWA	BD-A	SS-O	ECM-BAGUS	G-MPLBD	K-Horseshoe	H-BGGM	
10	Random	10 %	20	0.7	0.72	0.73	0.73	0.7	0.72	0.7	0.67	
		10 %	350	0.94	0.95	0.95	0.93	0.91	0.95	0.94	0.93	
	Cluster	10 %	20	0.83	0.82	0.82	0.81	0.79	0.82	0.81	0.75	
		10 %	350	0.93	0.93	0.93	0.91	0.92	0.94	0.92	0.92	
100	Random	1 %	40	0.83	0.83	0.81	0.84	0.76	0.81	0.84	–	
		1 %	700	0.96	0.96	0.96	0.94	0.94	0.95	0.96	0.95	
		10 %	40	0.67	0.7	0.68	0.71	0.59	0.67	0.71	–	
		10 %	700	0.78	0.92	0.91	0.89	0.8	0.9	0.92	0.87	
	Cluster	1 %	40	0.83	0.85	0.83	0.83	0.78	0.78	0.81	0.85	–
		1 %	700	0.95	0.96	0.96	0.94	0.94	0.94	0.94	0.95	0.94
		10 %	40	0.71	0.77	0.72	0.75	0.61	0.7	0.73	–	
		10 %	700	0.84	0.94	0.93	0.91	0.84	0.92	0.93	0.87	
1000	Random	0.1 %	60	0.86	–	–	0.84	0.84	0.72	0.84	–	
		0.1 %	1050	0.97	–	–	0.95	0.96	0.95	0.96	–	
	Cluster	0.1 %	60	0.84	–	–	0.82	0.84	0.7	0.84	–	
		0.1 %	1050	0.96	–	–	0.95	0.95	0.95	0.96	–	

NOTE: The AUC reaches its best score at 1 and its worst at 0. The values are averages over 16 replications. The “–” entry indicates that an algorithm did not produce results on that instance.

where we set  $\epsilon = 0.01$ . Here,  $P_t$  denotes the edge inclusion matrix after  $t$  seconds of running the algorithm and  $AUC(P)$  denotes the AUC corresponding to an edge inclusion matrix  $P$ . In real-world applications, the true graph is unknown and the computational cost  $T$  is not defined. For convergence assessments without knowing the true graph, we refer to Section 5. The computational cost of the H-BGGM and glasso algorithms can not be defined using (23). Instead, we define their computation cost  $T$  as the total time it takes these algorithms to provide their edge inclusion probabilities. We note that both the glasso and ECM-BAGUS algorithms require parameter tuning. We include this time in the computational cost.

**Simulation instances:** We include small ( $p = 10$ ), medium ( $p = 100$ ) and large ( $p = 1000$ ) scale graphs with a low ( $n = 20 \log p$ ) and high ( $n = 350 \log p$ ) number of observations. The use of the logarithm here is motivated by the performance bound given by (19). We consider the following two graph types:

1. Random: a graph in which every edge  $(i, j)$  is drawn from an independent Bernoulli distribution with probability  $\delta$ .
2. Cluster: a graph with two clusters for  $p \in \{10, 100\}$  and eight clusters for  $p = 1000$ . Each cluster has the same structure as the Random graph.

The value of  $\delta$  is chosen so that the resulting edge density in the graphs is 10% for  $p = 10$ , 1% and 10% for  $p = 100$ , and 0.1% for  $p = 1000$ . For every generated graph  $G^*$ , the corresponding precision matrix  $\mathbf{K}^*$  is then sampled from the  $G$ -Wishart distribution. That is,  $\mathbf{K}^* \sim W_{G^*}(3, \mathbf{I})$ . For each generated  $G^*$  and  $\mathbf{K}^*$ , we then sample  $n$  observations from the  $p$ -variate normal distribution with covariance matrix  $\Sigma^* = \mathbf{K}^{*-1}$  and mean zero. For  $p \in \{10, 100, 1000\}$ , each  $n$ , each graph type, and each density, we repeat the process 16 times which is generally enough to reach low ( $< 0.025$ ) standard errors on AUC,  $Pr^+$ , and  $Pr^-$ .

## 4.2. Results

The results on the computational costs  $T$ , AUC,  $Pr^-$ , and  $Pr^+$  are presented, respectively, in Tables 6–9. If there is one conclusion to be drawn from these results, then it is this: the idea that Bayesian algorithms lack the computational advantage of their frequentist counterparts is outdated. Whether with 10, 100, or 1000 variables, one or more Bayesian algorithms can estimate the full posterior distribution or any function thereof at the same time that the glasso algorithm can provide a point

**Table 8.**  $Pr^-$  scores (22) of the algorithms for different instances.

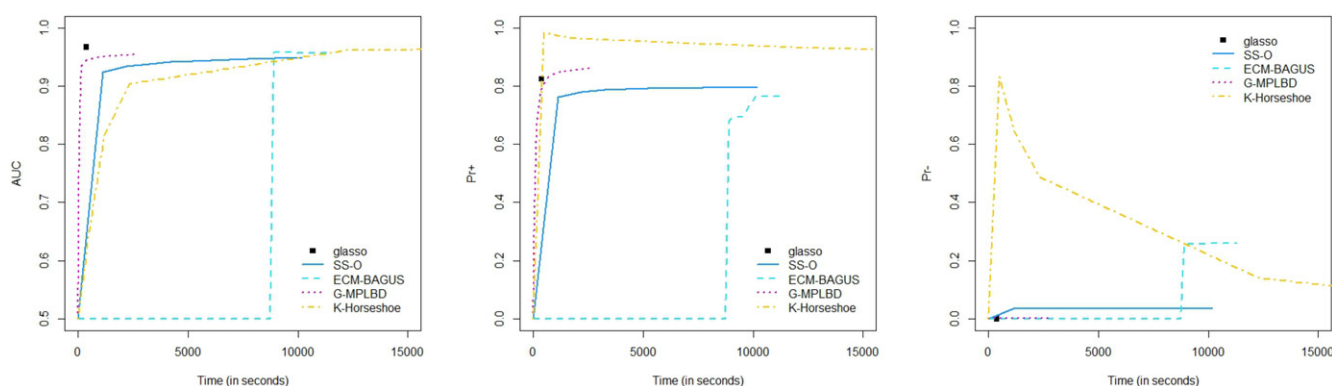
$p$	Graph	Density	$n$	glasso	RJ-WWA	BD-A	SS-O	ECM-BAGUS	G-MPLBD	K-Horseshoe	H-BGGM
10	Random	10%	20	0.05	0.09	0.1	0.08	0.18	0.08	0.06	0.88
		10%	350	0.05	0.02	0.03	0.03	0.15	0.01	0.09	0.33
	Cluster	10%	20	0.04	0.08	0.09	0.08	0.2	0.07	0.02	0.88
		10%	350	0.06	0.02	0.03	0.03	0.12	0.01	0.07	0.31
100	Random	1%	40	0	0.05	0.07	0.09	0.16	0.03	0	-
		1%	700	0.01	0.01	0.02	0.03	0.11	0.01	0	0.23
		10%	40	0.13	0.1	0.12	0.12	0.12	0.04	0.01	-
		10%	700	0.17	0.03	0.04	0.05	0.15	0	0.04	0.23
	Cluster	1%	40	0	0.05	0.07	0.09	0.18	0.03	0	-
		1%	700	0	0.01	0.02	0.03	0.14	0.01	0	0.23
		10%	40	0.06	0.08	0.11	0.11	0.12	0.04	0.01	-
		10%	700	0.11	0.02	0.03	0.05	0.12	0	0.03	0.23
1000	Random	0.1%	60	0	-	-	0.18	0.29	0.1	0.05	-
		0.1%	1050	0	-	-	0.04	0.26	0	0.05	-
	Cluster	0.1%	60	0	-	-	0.18	0.32	0.1	0.05	-
		0.1%	1050	0	-	-	0.04	0.23	0	0.05	-

NOTE: The  $Pr^-$  reaches its best score at 0 and its worst at 1. The values are averages over 16 replications. The “-” entry indicates that an algorithm did not produce results on that instance.

**Table 9.**  $Pr^+$  scores (21) of the algorithms for different instances.

$p$	Graph	Density	$n$	glasso	RJ-WWA	BD-A	SS-O	ECM-BAGUS	G-MPLBD	K-Horseshoe	H-BG-GM
10	Random	10%	20	0.38	0.38	0.4	0.35	0.31	0.4	0.39	0.9
		10%	350	0.87	0.77	0.78	0.56	0.69	0.81	0.85	0.9
	Cluster	10%	20	0.45	0.45	0.47	0.4	0.35	0.5	0.45	0.91
		10%	350	0.87	0.79	0.8	0.58	0.73	0.82	0.88	0.89
100	Random	1%	40	0.3	0.46	0.5	0.5	0.46	0.47	0.41	-
		1%	700	0.83	0.84	0.84	0.63	0.82	0.84	0.87	0.9
		10%	40	0.3	0.31	0.32	0.29	0.22	0.22	0.24	-
		10%	700	0.59	0.73	0.72	0.57	0.58	0.68	0.79	0.76
	Cluster	1%	40	0.26	0.56	0.51	0.5	0.46	0.47	0.43	-
		1%	700	0.81	0.83	0.84	0.61	0.83	0.84	0.81	0.9
		10%	40	0.27	0.33	0.35	0.32	0.24	0.26	0.26	-
		10%	700	0.59	0.73	0.73	0.59	0.63	0.68	0.8	0.77
1000	Random	0.1%	60	0.24	-	-	0.61	0.52	0.38	0.62	-
		0.1%	1050	0.83	-	-	0.8	0.76	0.86	0.91	-
	Cluster	0.1%	60	0.2	-	-	0.59	0.5	0.36	0.62	-
		0.1%	1050	0.8	-	-	0.78	0.81	0.85	0.91	-

NOTE: The  $Pr^+$  reaches its best score at 1 and its worst at 0. The values are averages over 16 replications. The “-” entry indicates that an algorithm did not produce results on that instance.



**Figure 1.** AUC (left),  $Pr^+$  (middle) and  $Pr^-$  (right) over time. The figures show averages over 16 replications on an instance with  $p = 1000$ ,  $n = 1050$  on the Random graph with density 0.1%.

estimate. Moreover, there is no evidence of the glasso algorithm outperforming Bayesian algorithms on AUC,  $Pr^+$ , or  $Pr^-$  either. It underlines the staggering progress Bayesian algorithms have made over the last decade. This progress is illustrated too in Figure 1, which shows the convergence of performance metrics for each algorithm on a large scale instance ( $p = 1000$ ): the

AUC,  $Pr^+$ , and  $Pr^-$  values of the G-MPLBD, SS-O, and K-Horseshoe algorithms converge in approximately the same time as the glasso algorithm.

Now, let us look into the results in more detail, starting with the computational cost in Table 6. Unsurprisingly, the glasso algorithm is fast. It provides a point estimate of the graph within

seconds for small and medium-sized problems and within 10 min for large-scale problems. The majority of these 10 min is spent on tuning the regularization parameter  $\lambda$  in (5). More surprising is the outstanding performance of the G-MPLBD algorithm that allows inference on the entire posterior within the same time as the glasso provides a point estimate. The most time-consuming algorithms are the RJ-WWA and BD-A algorithms, which are slow to run on large-scale instances. This is due to the computationally expensive sampling from the G-Wishart distribution at every iteration and the slow edge-by-edge exploration of the state space. Moreover, both algorithms benefit from parallel computing but are run on one core in this simulation study. The H-BGGM algorithm gives errors for the high dimensional case  $n < p$  and for the case  $p = 1000$ . The fields corresponding to these cases are therefore left empty.

In terms of the AUC (Table 7), most algorithms perform similarly. Noteworthy is the overall underperformance of the ECM-BAGUS algorithm and the underperformance of the glasso algorithm in medium-sized problems ( $p = 100$ ) with a high edge density (10%). The  $Pr^-$  values (Table 8) are comparable across most algorithms too with estimates close to the desired value of zero. Only the ECM-BAGUS and H-BGGM algorithms perform significantly worse, the latter even assigning in some cases an average edge inclusion probability of 0.88 to edges that are not in the true graph. Looking at the  $Pr^+$  values in Table 9, we see more variability between algorithms. For instances with a high number of observations, the K-Horseshoe algorithm performs best. For large-scale graphs, the K-horseshoe algorithm outperforms the others too. The ECM-BAGUS algorithm performs slightly worse than others in most instances. The  $Pr^+$  values of the H-BGGM algorithm seem close to optimal but are only marginally more than their  $Pr^-$  values.

The results also reveal that it is not so much the structure of the graph (Random or Cluster) as the density that influences the performance of algorithms. As expected, more observations lead to better results. Generally,  $n = 350 \log p$  observations are sufficient to yield good results.

## 5. Case Study: Human Gene Expression

In this section we demonstrate the capabilities of Bayesian structure learning on a real-world dataset. We discuss the challenges that arise when applying Bayesian approaches in practice and present ways to tackle these challenges. Consequently, this section serves as a guideline for practitioners who are new to the field.

We use the SS-O algorithm and the BD-A algorithm to perform inference on a gene network between  $p = 100$  genes using genetic data of  $n = 60$  unrelated individuals. Several Bayesian structure learning methods are analyzed on this dataset; see, for example, (Mohammadi and Wit 2015; Li, Craig, and Bhadra 2019a; Van den Boom, Beskos, and De Iorio 2022a; Mohammadi et al. 2023b). This dataset can therefore be seen as a benchmark for Bayesian structure learning. We refer to Stranger et al. (2007) and Bhadra and Mallick (2013) for more information on the collection of the data.

Genes are specific sequences of DNA. Gene expression is the process by which a gene produces a protein that influences

the functioning of an organism. Some of these proteins directly influence the organism by, for example, initializing a process to break down food. Other proteins, however, only serve to “activate” other genes, which in turn make proteins that activate other genes, and so on. These activation relationships can be shown in a gene network, where every node corresponds to a gene and every edge indicates an activation relationship between two genes. Discovering these gene networks plays a key role in understanding disease susceptibility and therefore, ultimately, in treatment and public health (Stranger et al. 2007). We obtain a sample  $\mathbf{X}$  of our variables  $X_i$ ,  $i = 1, \dots, p$ , where  $\mathbf{X}$  is an  $n \times p$  matrix with elements  $x_{li}$  denoting the level of protein corresponding to gene  $i = 1, \dots, p$  measured in individual  $l = 1, \dots, n$ .

Structure learning depends on the assumption that the data comes from a multivariate normal distribution with a mean zero. However, in the gene dataset, the univariate histograms of each gene are clearly not normally distributed. Neither are their means zero, see Figure 2. In fact, the normality and zero mean assumptions are rarely met in practice. Fortunately, one can transform our non-Gaussian continuous variables  $X_1, \dots, X_p$  to Gaussian variables  $Z := (Z_1, \dots, Z_p)$  with mean zero, that is,  $Z \sim \mathcal{N}(\mathbf{0}, \Sigma_Z)$ , such that the sparsity of  $\mathbf{K}_Z = \Sigma_Z^{-1}$  still encodes the conditional (in)dependence between the variables  $X_1, \dots, X_p$  as in (2); see, for example Liu, Lafferty, and Wasserman (2009) and Liu et al. (2012). The *huge.npn()* function in the *Huge* R package and the *bdgraph.npn()* function in the *BDgraph* R package perform this transformation efficiently. One can now apply Bayesian structure learning algorithms on the transformed dataset  $\mathbf{Z}$ . These transformations are only possible for continuous data. For discrete or binary data, Gaussian copula graphical models can be used to perform structure learning; see Section 6.1 for an overview.

We choose to work with MCMC algorithms to perform inference on the posterior of the gene network. These algorithms offer a wide variety of output and are computationally feasible on this 100 node graph. Specifically, we use the SS-O algorithm and the BD-A algorithm, but the same steps apply to other MCMC algorithms.

First, we select the tuning parameters. For the SS-O algorithm, we follow Wang (2015) and select  $\nu_0 = 0.02$ ,  $\nu_1 = 2$  and  $\eta = 1$ . One could also opt for the more time-consuming approach and run the algorithm several times for different parameters, and select the parameters that perform best according to some criterion (Gan, Narisetty, and Liang 2019). The BD-A algorithm has no tuning parameters.

Second, we need to select a prior for the graph. Generally, the prior is selected using prior knowledge about the sparsity of the true graph. Here, we select the prior given by (20), where we select the prior sparsity equal to  $\delta = 0.2$ , in line with previous studies on the gene expression dataset; see, for example, Van den Boom, Beskos, and De Iorio (2022a) and Li, Craig, and Bhadra (2019a).

Now, the correct MCMC settings need to be selected. That is, the number of burn-in iterations, the number of MCMC iterations, and the initial graph  $G^{(0)}$  of the Markov chain. The initial graph should not be important, since the Markov chain should converge to the posterior distribution regardless of the

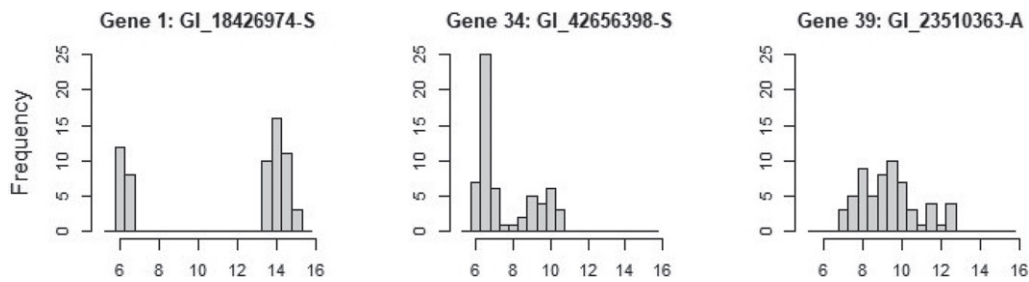


Figure 2. Univariate histograms of three randomly selected genes in the human gene expression dataset.

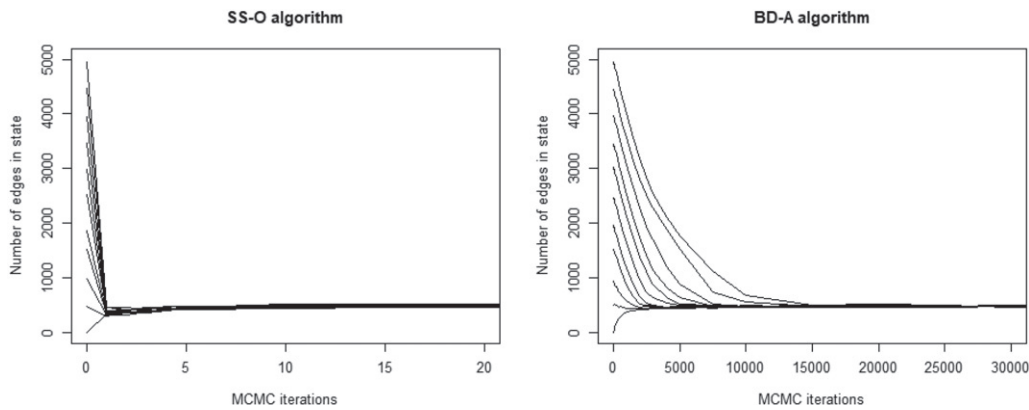


Figure 3. Number of edges in the visited graph per MCMC iteration for 10 replications. The density of the initial graph differs per replication, ranging from 0% (empty graph) to 100% (full graph).

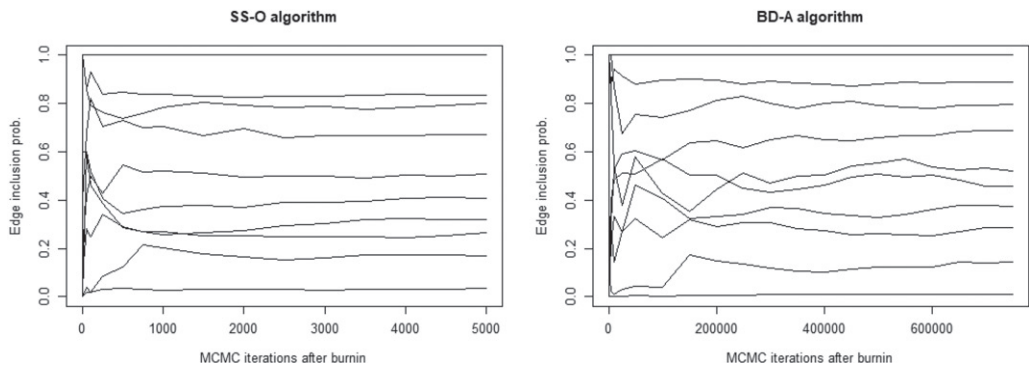


Figure 4. The edge inclusion probabilities of 10 randomly selected edges for every state in the Markov chain.

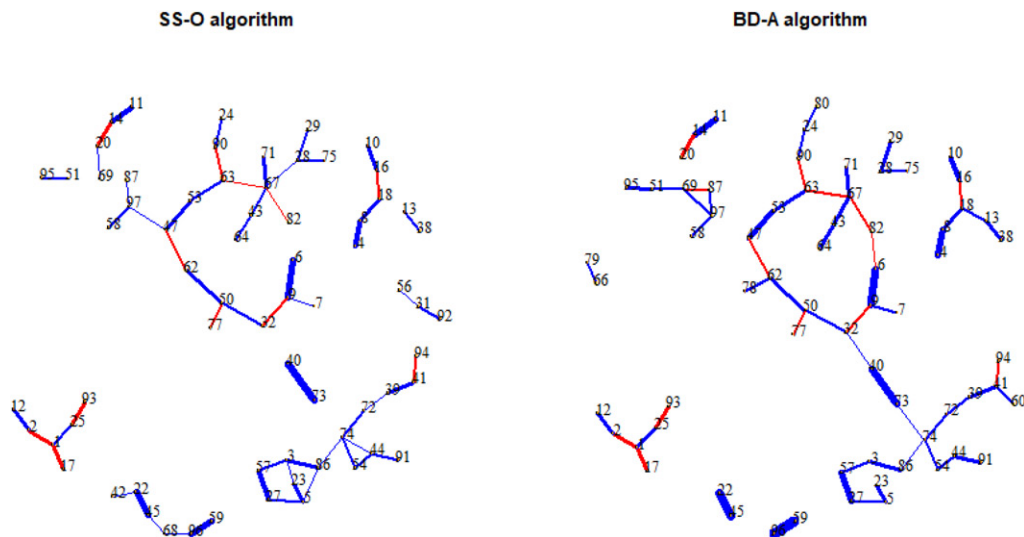
starting graph. To confirm that this indeed happens, we run the algorithms several times, each time from a different initial graph, and report the number of edges in each Markov chain state. The resulting plots (Figure 3) indeed show that, regardless of the initial graph, the Markov chain converges to the same neighborhood. Moreover, we can deduce that approximately the first 10 (SS-O algorithm) and the first 15,000 iterations (BD-A algorithm) can be discarded as burn-in iterations. To be on the safe side, we select burn-in periods of 100 (SS-O) and 25,000 (BD-A) iterations.

To select the number of MCMC iterations (after burn-in), we plot the edge inclusion probabilities of ten randomly selected links at every state. The resulting plots (Figure 4) show that convergence happens in less iterations for the SS-O algorithm (approximately 5000) than for the BD-A algorithm (approximately 750,000). This is due to the good mixing of the SS-O algorithm versus the slow edge-by-edge exploration of the BD-A algorithm. We set the amount of MCMC iterations accordingly:

5000 for the SS-O algorithm and 750,000 for the BD-A algorithm.

Figure 5 shows the gene networks inferred by the BD-A algorithm and the SS-O algorithm. The networks show similarities in both the sign (positive or negative) and the size of the partial correlations. Moreover, the networks show similarities with the existing literature. The five most connected genes of the K-Horseshoe algorithm (Li, Craig, and Bhadra 2019a) all appear among the most connected genes of our analysis. Moreover, most tree structures of Figure 5 appear exactly in the recovered network given by Bhadra and Mallick (2013, Figure 4).

The impact of the prior density on the posterior remains relatively unexplored in Bayesian structure learning literature. In order to test the sensitivity of the algorithms to the priors, we ran the algorithms with a sparse (1% density), average (20% density) and dense (50% density) graph prior. We compare the boxplots of the resulting edge inclusion probabilities in Figure S1 (SS-O algorithm) and in Figure S2 (BD-A algorithm) in Section



**Figure 5.** Gene networks inferred by the SS-O algorithm (left) and the BD-A algorithm (right). Only the 60 edges with the highest edge inclusion probabilities are shown. The width of the edge indicates the strength of the partial correlation. The color represents the sign of the partial correlation: red (negative) and blue (positive). See supplementary material S5 for the gene names corresponding to the node numbering.

S4 in the supplementary materials. As expected, a higher prior density results in higher edge inclusion probabilities with the SS-O algorithm being slightly more sensitive to a higher prior density. For all prior densities, however, both algorithms are able to distinguish between edges with a low and high edge inclusion probability.

## 6. Conclusion and Future Research

In structure learning, Bayesian methods constitute a powerful alternative to frequentist ones by providing inference on the entire posterior. The argument that the speed and simplicity of frequentist methods are superior is waning as Bayesian methods have improved significantly in both aspects. Bayesian methods that provide accurate solutions to thousand variable problems within mere minutes are now a reality and easily accessible in software packages. This section starts with a short overview of structure learning literature beyond the basic Gaussian case and ends with suggestions for further research.

### 6.1. Extension Beyond Basic GGMs

While the majority of research on graphical modeling has concentrated on Gaussian graphical models, significant efforts have been directed toward extending beyond basic GGMs over the past decade. We provide a brief overview of these extensions here. For comprehensive reviews, works such as Liang and Jia (2024) and Maathuis et al. (2019) offer detailed insights. Additionally, these developments merit a dedicated review paper, highlighting their importance in the field.

**Beyond Gaussianity assumption:** Graphical modeling for non-Gaussian datasets, including types like single nucleotide polymorphisms, RNA sequencing, and household income data, is increasingly relevant in data science. Over the past decade, adapting models to handle such non-Gaussian data has been a major research area. For continuous non-Gaussian datasets,

like the gene interaction data discussed in Section 5, a two-step method involving a nonparametric data-Gaussianized transformation (Liu, Lafferty, and Wasserman 2009) followed by GGM application is effective. This method also extends to multivariate discrete data, as shown in Jia et al. (2017) and Liang and Jia (2024, chap. 5). For multivariate discrete or count data, Gaussian copula graphical models offer an alternative, as detailed in Vinciotti, Behrouzi, and Mohammadi (2022), Dobra and Mohammadi (2017) and Dobra and Lenkoski (2011). For mixed data types, Bayesian structure learning methods like those in Dobra and Lenkoski (2011); Mohammadi et al. (2017), that use a latent GGM with a copula approach Hoff (2007), are suitable.

**Multiple Gaussian graphical models:** Multiple GGMs, as an extension to standard GGMs, are useful for jointly analyzing data from multiple sources (also known as heterogeneous data), for example, neurological data measured at multiple timescales, or joint neurological, genetic, and phenotypic data. For these types of data, Peterson, Stingo, and Vannucci (2015) use Markov random field prior (Li and Zhang 2010) to model a super-graph linking different graphical models. Tan et al. (2017) uses a logistic regression model to link the connectivity of nodes to covariates specific to each graph. More recently, Li, McCormick, and Clark (2019b) developed doubly spike-and-slab mixture priors as a new class of priors for joint estimation of multiple graphical models.

**Colored Gaussian graphical models:** These are a class of GGMs with additional symmetry restrictions in the form of equality constraints on the parameters (Højsgaard and Lauritzen 2008). Most of the methods mentioned in Section 3 can be applied to the colored GGMs. For example, Li, Gao, and Massam (2020) implemented an RJ method by using the colored  $G$ -Wishart prior as the Diaconis-Ylvisaker conjugate prior for GGMs. Roverato and Nguyen (2022) applied a stepwise model search procedure to estimate a brain network from fMRI data.

**Time-varying GGMs:** While extensive literature revolves around learning static graphical time-invariant models, the

change of interdependencies with a covariate (e.g., time or space) is often the rule rather than the exception for real-world data, such as friendships between individuals in a social community, communication between genes in a cell, equity trading between companies, and computer network traffic. Bayesian approaches in time-varying GGMs are computationally burdensome with time complexity. Thus, Yu, Wu, and Dauwels (2022) developed a low-complexity Bayesian approach by imposing temporally dependent spike-and-slab priors on the graphs such that they are sparse and vary smoothly across time.

**Multivariate regressions:** GGMs, and sparse graphical models more broadly, are applicable for statistical inference in high-dimensional regression. This application is straightforward when the response variables and explanatory variables are collectively viewed as a set of  $p$ -dimensional variables within our graphical models, as discussed in sec. 6 of Lenkoski and Dobra (2011), for example. In such instances, the inference process can address both the selection of relevant variables and the quantification of uncertainty for the regression coefficients, as outlined in the works of Liang, Xue, and Jia (2022) and Liang, Song, and Yu (2013).

## 6.2. Future Research

In the coming decade, Bayesian methods have the potential to become even faster, expand their theoretical guarantees, grow beyond the Gaussian case, and make more impact with their applications. This section speculates how.

First, new MCMC methods can increase the efficiency and feasible dimension of Bayesian structure learning. Mohammadi et al. (2023b) show the potential of MCMC methods that move only over the graph space, and not over the space of precision matrices. These methods, however, still only allow graphs to change at most one edge per MCMC iteration. It is an open question whether the balance conditions still hold when allowing changes of multiple edges. If true, this could lead to a significant reduction in computation time. Similarly, Van den Boom, Beskos, and De Iorio (2022a) show the benefit of the informed proposal and delayed acceptance techniques for the reversible jump approach on the joint space of graphs and precision matrices. It remains to be shown whether methods on the graph space can benefit from the same techniques. Wang (2015) introduces the spike-and-slab prior in Bayesian structure learning. He circumvents the calculation of the normalizing constant by putting its inverse in the prior of the graph. This leads to a computationally efficient method but creates a challenge when incorporating any knowledge of the graph structure in the prior of the graph. This raises two questions: (i) Could the normalizing constant of the spike-and-slab prior be approximated? (ii) Could  $G$ -Wishart methods on the joint space benefit from a similar trick?

Second, moving away from the standard MCMC approach offers a promising and unexplored perspective to Bayesian structure learning. Dai et al. (2022) argue that Sequential Monte Carlo (SMC) methods remain under-used in statistics, despite several advantages. Van den Boom et al. (2022b) and Tan et al. (2017) show that the SMC approach works for Bayesian structure learning. However, their SMC methods cannot yet compete in terms of computational efficiency with the MCMC methods.

Likewise, Nemeth and Fearnhead (2021) outline the benefits of the stochastic gradient MCMC (SGMCMC) method, and Tan and Friel (2020) design an SGMCMC method for exponential random graph models. As of now, however, no SGMCMC structure learning method exists.

Third, Bayesian structure learning needs more theoretical results. Although the majority of MCMC methods meet the balance conditions, not one of them provides rates of convergence as in Tierney (1994, sec. 3.2). Moreover, among the methods in Table 1, only one (Van den Boom et al. 2022b) provides an unbiased estimate. In the future, more theoretical results are essential to the credibility of Bayesian structure learning.

Fourth, recent improvements in structure learning in GGMs could enhance methods beyond the general Gaussian case. The most apparent is the non-Gaussian case, in which the state-of-the-art methods discussed in this article can be directly applied. This avenue, however, is barely investigated. Similarly, other related fields can benefit from the recent strides made in the general Gaussian case. They include multiple Gaussian graphical models (Tan et al. 2017), colored Gaussian graphical models (Li, Gao, and Massam 2020), and graphical models with external network data (Jewson et al. 2022).

Lastly, the increase in the feasible dimension of Bayesian methods will enhance their applications. In Section 5, we reduce the dataset containing several thousands of genes to a mere 100 genes to make the data feasible for computation, potentially losing valuable information. This kind of dimension reduction will be decreasingly necessary as Bayesian methods improve. Likewise, in neuroscience, the models of the brain no longer have to be simplified to just 100 areas (Dyrba et al. 2020). This could potentially improve the understanding of cognitive diseases like Alzheimer's. The enhancements in Bayesian structure learning also open up new applications. Especially exciting but yet unexplored examples are graph neural networks and large language models, which make use of the dependency networks among a large number of variables.

## Supplementary Materials

**GitHub repository:** All scripts used to perform the simulation study of Section 4 and the case study of Section 5 can be found on <https://github.com/lucasvogels33/Review-paper-Bayesian-Structure-Learning-in-GGMs>.

**Supplementary materials:** The supplementary materials provide (S1) the references to the code of all algorithms compared in Section 4, (S2) all algorithm-specific settings for the simulation study of Section 4, (S3) the data dictionary for the data used in Section 5, (S4) boxplots showing prior sensitivity of the BD-A and SS-O algorithms on the human gene dataset in Section 5, (S5) a table to convert the gene numbers of Figure 5 to the names of the genetic transcripts.

## Acknowledgments

The authors thank the Editor, Prof. Lee, the Associate Editor, and the two referees for their insightful comments, which have significantly improved this article.

## Disclosure Statement

The authors report there are no competing interests to declare.

## ORCID

Lucas Vogels  <http://orcid.org/0009-0005-3715-3937>

Reza Mohammadi  <http://orcid.org/0000-0001-9538-0648>

Marit Schoonhoven  <http://orcid.org/0000-0002-2268-7486>

Ş. İlker Birbil  <http://orcid.org/0000-0001-7472-7032>

## References

- Atay-Kayis, A., and Massam, H. (2005), "A Monte Carlo Method for Computing the Marginal Likelihood in Non-decomposable Gaussian Graphical Models," *Biometrika*, 92, 317–335. [3166]
- Atchadé, Y. F. (2019), "Quasi-Bayesian Estimation of Large Gaussian Graphical Models," *Journal of Multivariate Analysis*, 173, 656–671. [3167,3168,3169]
- Banerjee, O., El Ghaoui, L., and d'Aspremont, A. (2008), "Model Selection through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data," *Journal of Machine Learning Research*, 9, 485–516. [3166]
- Besag, J. (1975), "Statistical Analysis of Non-lattice Data," *Journal of the Royal Statistical Society, Series D*, 24, 179–195. [3169]
- Bhadra, A., and Mallick, B. (2013), "Joint High-Dimensional Bayesian Variable and Covariance Selection with an Application to eQTL Analysis," *Biometrics*, 69, 447–457. [3164,3176,3177]
- Bien, J., and Tibshirani, R. J. (2011), "Sparse Estimation of A Covariance Matrix," *Biometrika*, 98, 807–820. [3166]
- Cappé, O., Robert, C. P., and Rydén, T. (2003), "Reversible Jump, Birth-and-Death and More General Continuous Time Markov Chain Monte Carlo Samplers," *Journal of the Royal Statistical Society, Series B*, 65, 679–700. [3169]
- Carter, J. S., Rossell, D., and Smith, J. Q. (2023), "Partial Correlation Graphical LASSO," *Scandinavian Journal of Statistics*, 51, 32–63. [3171]
- Carvalho, C. M., Massam, H., and West, M. (2007), "Simulation of Hyper-Inverse Wishart Distributions in Graphical Models," *Biometrika*, 94, 647–659. [3166]
- Chandra, N. K., Mueller, P., and Sarkar, A. (2022), "Bayesian Scalable Precision Factor Analysis for Massive Sparse Gaussian Graphical Models," Unpublished manuscript, arXiv: 2107.11316. [3164,3167,3171,3172]
- Cheng, Y., and Lenkoski, A. (2012), "Hierarchical Gaussian Graphical Models: Beyond Reversible Jump," *Electronic Journal of Statistics*, 6, 2309–2331. [3167,3168]
- Dai, C., Heng, J., Jacob, P. E., and Whiteley, N. (2022), "An Invitation to Sequential Monte Carlo Samplers," *Journal of the American Statistical Association*, 117, 1587–1600. [3179]
- Dawid, A. P., and Lauritzen, S. L. (1993), "Hyper Markov Laws in the Statistical Analysis of Decomposable Graphical Models," *The Annals of Statistics*, 21, 1272–1317. [3166]
- Dempster, A. P. (1972), "Covariance Selection," *Biometrics*, 28, 157–175. [3165]
- Dobra, A., and Lenkoski, A. (2011), "Copula Gaussian Graphical Models and their Application to Modeling Functional Disability Data," *The Annals of Applied Statistics*, 5, 969–993. [3178]
- 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. [3167,3168]
- Dobra, A., and Mohammadi, R. (2017), "Loglinear Model Selection and Human Mobility," *The Annals of Applied Statistics*, 12, 815–845. [3178]
- Dyrba, M., Mohammadi, R., Grothe, M. J., Kirste, T., and Teipel, S. J. (2020), "Gaussian Graphical Models Reveal Inter-Modal and Inter-Regional Conditional Dependencies of Brain Alterations in Alzheimer's Disease," *Frontiers in Aging Neuroscience*, 12, 99. [3164,3179]
- Epskamp, S., Waldorp, L. J., Möttus, R., and Borsboom, D. (2018), "The Gaussian Graphical Model in Cross-Sectional and Time-Series Data," *Multivariate Behavioral Research*, 53, 453–480. [3164]
- Fan, J., Feng, Y., and Wu, Y. (2009), "Network Exploration via the Adaptive LASSO and SCAD Penalties," *The Annals of Applied Statistics*, 3, 521 – 541. [3166]
- Foygel, R., and Drton, M. (2010), "Extended Bayesian Information Criteria for Gaussian Graphical Models," in *Advances in Neural Information Processing Systems* (Vol. 23). [3166]
- Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation with the Graphical Lasso," *Biostatistics*, 9, 432–441. [3166]
- Gan, L., Narisetty, N. N., and Liang, F. (2019), "Bayesian Regularization for Graphical Models with Unequal Shrinkage," *Journal of the American Statistical Association*, 114, 1218–1231. [3167,3171,3173,3176]
- Giudici, P. (1995), "Bayes Factors for Zero Partial Covariances," *Journal of Statistical Planning and Inference*, 46, 161–174. [3172]
- Giudici, P., and Castelo, R. (2003), "Improving Markov Chain Monte Carlo Model Search for Data Mining," *Machine Learning*, 50, 127–158. [3164]
- Giudici, P., and Green, P. J. (1999), "Decomposable Graphical Gaussian Model Determination," *Biometrika*, 86, 785–801. [3166]
- Glynn, P. W., and Rhee, C.-H. (2014), "Exact Estimation for Markov Chain Equilibrium Expectations," *Journal of Applied Probability*, 51A, 377–389. [3169]
- Green, P. (1995), "Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination," *Biometrika*, 82, 711–732. [3168]
- Hinne, M., Lenkoski, A., Heskes, T. M., and van Gerven, M. (2014), "Efficient Sampling of Gaussian Graphical Models Using Conditional Bayes Factors," *Stat*, 3, 326–336. [3164,3167,3168]
- Hoff, P. D. (2007), "Extending the Rank Likelihood for Semiparametric Copula Estimation," *The Annals of Applied Statistics*, 1, 265–283. [3178]
- Højsgaard, S., and Lauritzen, S. L. (2008), "Graphical Gaussian Models with Edge and Vertex Symmetries," *Journal of the Royal Statistical Society, Series B*, 70, 1005–1027. [3178]
- Jalali, P., Khare, K., and Michailidis, G. (2020), "B-concord - A Scalable Bayesian High-Dimensional Precision Matrix Estimation Procedure," unpublished manuscript, arXiv:2005.09017. [3167,3168,3169,3170,3171,3173]
- Jewson, J., Li, L., Battaglia, L., Hansen, S., Rossell, D., and Zwiernik, P. (2022), "Graphical Model Inference with External Network Data," unpublished manuscript, arXiv: 2210.11107. [3179]
- Jia, B., Xu, S., Xiao, G., Lamba, V., and Liang, F. (2017), "Learning Gene Regulatory Networks from Next Generation Sequencing Data," *Biometrics*, 73, 1221–1230. [3178]
- 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. [3166]
- Khondker, Z., Zhu, H., Chu, H., Lin, W., and Ibrahim, J. (2013), "The Bayesian Covariance Lasso," *Statistics and Its Interface*, 6, 243–259. [3167,3171]
- Kundu, S., Mallick, B., and Baladandayuthapani, V. (2018), "Efficient Bayesian Regularization for Graphical Model Selection," *Bayesian Analysis*, 14, 449–476. [3167,3171]
- Lauritzen, S. L. (1996), *Graphical Models*, Oxford, UK: Oxford University Press. [3165]
- Leday, G., and Richardson, S. (2018), "Fast Bayesian Inference in Large Gaussian Graphical Models," *Biometrics*, 75, 1288–1298. [3167,3172,3173]
- Lenkoski, A. (2013), "A Direct Sampler for G-Wishart Variates," *Stat*, 2, 119–128. [3167,3168,3169]
- 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. [3166,3167,3179]
- Leppä-aho, J., Pensar, J., Roos, T., and Corander, J. (2017), "Learning Gaussian Graphical Models with Fractional Marginal Pseudo-likelihood," *International Journal of Approximate Reasoning*, 83, 21–42. [3170]
- Li, F., and Zhang, N. R. (2010), "Bayesian Variable Selection in Structured High-Dimensional Covariate Spaces with Applications in Genomics," *Journal of the American Statistical Association*, 105, 1202–1214. [3178]
- Li, F. and Zhang, X. (2017), "Bayesian Lasso with Neighborhood Regression Method for Gaussian Graphical Model," *Acta Mathematicae Applicatae Sinica, English Series*, 33, 485–496. [3167,3172]

- Li, Q., Gao, X., and Massam, H. (2020), “Bayesian Model Selection Approach for Coloured Graphical Gaussian Models,” *Journal of Statistical Computation and Simulation*, 90, 2631–2654. [3178,3179]
- Li, Y., Craig, B. A., and Bhadra, A. (2019a), “The Graphical Horseshoe Estimator for Inverse Covariance Matrices,” *Journal of Computational and Graphical Statistics*, 28, 747–757. [3164,3167,3170,3172,3173,3176,3177]
- Li, Z., McCormick, T., and Clark, S. (2019b), “Bayesian Joint Spike-and-Slab Graphical Lasso,” in *International Conference on Machine Learning*, PMLR, pp. 3877–3885. [3178]
- Li, Z. R., and McCormick, T. H. (2019), “An Expectation Conditional Maximization Approach for Gaussian Graphical Models,” *Journal of Computational and Graphical Statistics*, 28, 767–777. [3167,3171]
- Liang, F., and Jia, B. (2024), *Sparse Graphical Modeling for High Dimensional Data: A Paradigm of Conditional Independence Tests*, Boca Raton, FL: CRC Press. [3165,3178]
- Liang, F., Song, Q., and Yu, K. (2013), “Bayesian Subset Modeling for High-Dimensional Generalized Linear Models,” *Journal of the American Statistical Association*, 108, 589–606. [3179]
- Liang, F., Xue, J., and Jia, B. (2022), “Markov Neighborhood Regression for High-Dimensional Inference,” *Journal of the American Statistical Association*, 117, 1200–1214. [3179]
- Liu, H., Han, F., Yuan, M., Lafferty, J., and Wasserman, L. (2012), “The Nonparanormal Skeptic,” *Proceedings of the 29th International Conference on Machine Learning, ICML 2012 (Vol. 2)*, pp. 1415–1422. [3176]
- 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. [3176,3178]
- Liu, H. and Wang, L. (2017), “TIGER: A Tuning-Insensitive Approach for Optimally Estimating Gaussian Graphical Models,” *Electronic Journal of Statistics*, 11, 241–294. [3165]
- Maathuis, M., Drton, M., Lauritzen, S., and Wainwright, M. (2019), *Handbook of Graphical Models*, Boca Raton, FL: CRC Press. [3166,3178]
- Meinshausen, N., and Bühlmann, P. (2006), “High-Dimensional Graphs and Variable Selection with the Lasso,” *The Annals of Statistics*, 34, 1436–1462. [3165]
- Mohammadi, A., Abegaz, F., van den Heuvel, E., and Wit, E. (2017), “Bayesian Modelling of Dupuytren Disease by Using Gaussian Copula Graphical Models,” *Journal of the Royal Statistical Society, Series C*, 66, 629–645. [3178]
- Mohammadi, A., and Wit, E. C. (2015), “Bayesian Structure Learning in Sparse Gaussian Graphical Models,” *Bayesian Analysis*, 10, 109–138. [3167,3168,3169,3173,3176]
- Mohammadi, R. (2022), *ssgraph: Bayesian Graph Structure Learning Using Spike-and-Slab Priors*, R package version 1.15. [3164]
- Mohammadi, R., Massam, H., and Letac, G. (2023a), “Accelerating Bayesian Structure Learning in Sparse Gaussian Graphical Models,” *Journal of the American Statistical Association*, 118, 1345–1358. [3167,3168,3169]
- Mohammadi, R., Schoonhoven, M., Vogels, L., and Birbil, I. S. (2023b), “High-Dimensional Bayesian Structure Learning in Gaussian Graphical Models Using Marginal Pseudo-Likelihood,” Unpublished manuscript, arXiv:2307.00127. [3167,3170,3176,3179]
- Mohammadi, R., and Wit, E. C. (2019), “Bdgraph: An R Package for Bayesian Structure Learning in Graphical Models,” *Journal of Statistical Software*, 89, 1–30. [3168]
- Mohammadi, R., Wit, E., and Dobra, A. (2022), *BDgraph: Bayesian Structure Learning in Graphical Models Using Birth-Death MCMC*, R package version 2.72. [3164,3167,3168]
- Nemeth, C., and Fearnhead, P. (2021), “Stochastic Gradient Markov Chain Monte Carlo,” *Journal of the American Statistical Association*, 116, 433–450. [3179]
- Park, T., and Casella, G. (2008), “The Bayesian Lasso,” *Journal of the American Statistical Association*, 103, 681–686. [3172]
- Peterson, C., Stingo, F. C., and Vannucci, M. (2015), “Bayesian Inference of Multiple Gaussian Graphical Models,” *Journal of the American Statistical Association*, 110, 159–174. [3178]
- Ravikumar, P., Wainwright, M., Raskutti, G., and Yu, B. (2011), “High-Dimensional Covariance Estimation by Minimizing L1-Penalized Log-Determinant Divergence,” *Electronic Journal of Statistics*, 5, 935–980. [3166]
- Rothman, A., Bickel, P., Levina, E., and Zhu, J. (2008), “Sparse Permutation Invariant Covariance,” *Electronic Journal of Statistics*, 2, 494–515. [3166]
- 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. [3166]
- Roverato, A., and Nguyen, D. N. (2022), “Model Inclusion Lattice of Coloured Gaussian Graphical Models for Paired Data,” in *International Conference on Probabilistic Graphical Models*, PMLR, pp. 133–144. [3178]
- Sagar, K., Banerjee, S., Datta, J., and Bhadra, A. (2024), “Precision Matrix Estimation under the Horseshoe-Like Prior-Penalty Dual,” *Electronic Journal of Statistics*, 18, 1–46. [3167,3171,3173]
- Sagar, K., Datta, J., Banerjee, S., and Bhadra, A. (2023), “Maximum a Posteriori Estimation in Graphical Models Using Local Linear Approximation,” unpublished manuscript, arXiv: 2104.10750. [3172]
- Smith, J., Arashi, M., and Bekker, A. (2023a), *baygel: Bayesian Shrinkage Estimators for Precision Matrices in Gaussian Graphical Models*, R package version 0.3.0. [3164]
- Smith, J., Arashi, M., and Bekker, A. (2023b), “A Data Driven Bayesian Graphical Ridge Estimator,” unpublished manuscript, arXiv:2210.16290. [3167,3171]
- Stingo, F., and Marchetti, G. (2014), “Efficient Local Updates for Undirected Graphical Models,” *Statistics and Computing*, 25, 159–171. [3167,3170,3173]
- Stranger, B. E., Nica, A. C., Forrest, M. S., Dimas, A., Bird, C. P., Beazley, C., Ingle, C. E., Dunning, M., Flicek, P., Koller, D., Montgomery, S., Tavaré, S., Deloukas, P., and Dermitzakis, E. T. (2007), “Population Genomics of Human Gene Expression,” *Nature Genetics*, 39, 1217–1224. [3176]
- Sun, T., and Zhang, C.-H. (2012), “Sparse Matrix Inversion with Scaled Lasso,” *Journal of Machine Learning Research*, 14, 3385–3418. [3165]
- Tadesse, M., and Vannucci, M. (2021), *Handbook of Bayesian Variable Selection*, Boca Raton, FL: CRC Press. [3169]
- Talluri, R., Baladandayuthapani, V., and Mallick, B. (2014), “Bayesian Sparse Graphical Models and their Mixtures,” *Stat*, 3, 109–125. [3169]
- Tan, L. S. L., and Friel, N. (2020), “Bayesian Variational Inference for Exponential Random Graph Models,” *Journal of Computational and Graphical Statistics*, 29, 910–928. [3179]
- Tan, L. S. L., Jasra, A., Iorio, M. D., and Ebbels, T. M. D. (2017), “Bayesian Inference for Multiple Gaussian Graphical Models with Application to Metabolic Association Networks,” *The Annals of Applied Statistics*, 11, 2222–2251. [3169,3178,3179]
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [3165]
- Tierney, L. (1994), “Markov Chains for Exploring Posterior Distributions,” *The Annals of Statistics*, 22, 1701–1728. [3168,3170,3179]
- Van den Boom, W., Beskos, A., and De Iorio, M. (2022a), “The G-Wishart Weighted Proposal Algorithm: Efficient Posterior Computation for Gaussian Graphical Models,” *Journal of Computational and Graphical Statistics*, 31, 1215–1224. [3167,3168,3173,3176,3179]
- Van den Boom, W., Jasra, A., De Iorio, M., Beskos, A., and Eriksson, J. (2022b), “Unbiased Approximation of Posteriors via Coupled Particle Markov Chain Monte Carlo,” *Statistics and Computing*, 32, 1–19. [3167,3168,3169,3179]
- Vandenbergh, L., Boyd, S., and Wu, S.-P. (1998), “Determinant Maximization with Linear Matrix Inequality Constraints,” *SIAM Journal on Matrix Analysis and Applications*, 19, 499–533. [3166]
- Vinciotti, V., Behrouzi, P., and Mohammadi, R. (2022), “Bayesian Inference of Microbiota Systems From Count Metagenomic Data,” unpublished manuscript, arXiv:2307.00127. [3178]
- Waldorp, L., and Marsman, M. (2022), “Relations between Networks, Regression, Partial Correlation, and the Latent Variable Model,” *Multivariate Behavioral Research*, 57, 994–1006. [3164]
- Wang, H. (2012), “Bayesian Graphical Lasso Models and Efficient Posterior Computation,” *Bayesian Analysis*, 7, 867–886. [3167,3170,3171]

- (2015), “Scaling It Up: Stochastic Search Structure Learning in Graphical Models,” *Bayesian Analysis*, 10, 351–377. [3164,3167,3168,3169,3176,3179]
- Wang, H. and Li, S. Z. (2012), “Efficient Gaussian Graphical Model Determination under G-Wishart Prior Distributions,” *Electronic Journal of Statistics*, 6, 168–198. [3167,3168]
- Williams, D., Piironen, J., Vehtari, A., and Rast, P. (2018), “Bayesian Estimation of Gaussian Graphical Models with Projection Predictive Selection,” Unpublished manuscript, arXiv:1801.05725. [3167,3172,3173]
- Williams, D. R., and Mulder, J. (2019), *BGGM: An R Package for Bayesian Gaussian Graphical Models*, R package version 2.0.0. [3164]
- (2020), “Bayesian Hypothesis Testing for Gaussian Graphical Models: Conditional Independence and Order Constraints,” *Journal of Mathematical Psychology*, 99, 102441. [3167,3172,3173]
- Wong, C., Moffa, G., and Kuipers, J. (2024), “A New Way to Evaluate G-Wishart Normalising Constants via Fourier Analysis,” unpublished manuscript, arXiv:2404.06803. [3166]
- Yu, H., Wu, S., and Dauwels, J. (2022), “Efficient Variational Bayes Learning of Graphical Models with Smooth Structural Changes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45, 475–488. [3179]
- Yuan, M. (2010), “High Dimensional Inverse Covariance Matrix Estimation via Linear Programming,” *Journal of Machine Learning Research*, 11, 2261–2286. [3165]
- Yuan, M., and Lin, Y. (2007), “Model Selection and Estimation in the Gaussian Graphical Model,” *Biometrika*, 94, 19–35. [3166]
- Zhao, P., and Yu, B. (2006), “On Model Selection Consistency of Lasso,” *Journal of Machine Learning Research*, 7, 2541–2563. [3165]
- 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. [3166]