



HAL
open science

Enhancing the Power of Gaussian Graphical Model Inference by Modeling the Graph Structure

Valentin Kilian, Tabea Rebafka, Fanny Villers

► **To cite this version:**

Valentin Kilian, Tabea Rebafka, Fanny Villers. Enhancing the Power of Gaussian Graphical Model Inference by Modeling the Graph Structure. 2024. hal-04460593

HAL Id: hal-04460593

<https://hal.science/hal-04460593v1>

Preprint submitted on 28 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Enhancing the Power of Gaussian Graphical Model Inference by Modeling the Graph Structure

Valentin Kilian, Tabea Rebafka and Fanny Villers*

LPSM, Sorbonne Université, Paris, France

Abstract

For the problem of inferring a Gaussian graphical model (GGM), this work explores the application of a recent approach from the multiple testing literature for graph inference. The main idea of the method by Rebafka et al. (2022) is to model the data by a latent variable model, the so-called noisy stochastic block model (NSBM), and then use the associated ℓ -values to infer the graph. The inferred graph controls the false discovery rate, that means that the proportion of falsely declared edges does not exceed a user-defined nominal level. Here it is shown that any test statistic from the GGM literature can be used as input for the NSBM approach to perform GGM inference. To make the approach feasible in practice, a new, computationally efficient inference algorithm for the NSBM is developed relying on a greedy approach to maximize the integrated complete-data likelihood. Then an extensive numerical study illustrates that the NSBM approach outperforms the state of the art for any of the here considered GGM-test statistics. In particular in sparse settings and on real datasets a significant gain in power is observed.

Keywords: Gaussian graphical models; noisy stochastic block model; structured ℓ -values; power enhancement; integrated complete-data likelihood; greedy algorithm

*The authors gratefully acknowledge *please list all relevant funding sources in the unblinded version*

1 Introduction

Graphical modeling offers a convenient way to describe complex conditional dependence structures between variables using a graph. This is of great interest in a large variety of domains of application. One of the main question is to attempt to discover the links between variables from a set of repeating observations of these variables. Gaussian Graphical Models (GGM) provide a suitable framework in this context. The goal of our work is to propose an approach for GGM inference that is both powerful to detect significant links and allows a control of the error rate. More precisely, we aim at providing a control of the false discovery rate (FDR), that is we want to guarantee that the proportion of falsely declared dependencies between variables does not exceed a user-defined nominal level. The approach that we propose here is based on a method for graph inference developed in a general statistical framework by Rebafka et al. (2022). We make this approach usable in practice for GGM inference by developing a new computationally efficient estimation algorithm for the so-called noisy stochastic block model and show that the approach detects more dependencies among variables than state-of-the-art methods.

1.1 Inference in Gaussian graphical models

We consider the setting, where observations are independent realizations of a multivariate Gaussian distribution. Its precision matrix, that is the inverse of the covariance matrix, provides information on the conditional correlations of the variables. Notably, two variables are conditionally independent given all other variables if and only if the associated entry of the precision matrix is zero. Now, GGM inference amounts to recover the non zero entries of the precision matrix from the realizations of the Gaussian distribution. The problem can also be viewed as a graph inference problem, where the edges of a binary graph indicate the non zero entries of the precision matrix and this graph is to be inferred from the Gaussian observations.

In the literature there are mainly three different approaches for GGM inference. First, some methods aim at estimating the precision matrix, as done in graphical Lasso (Banerjee et al., 2008; Friedman et al., 2007) with a penalized maximum likelihood approach, or in (Cai et al., 2011; Li et al., 2022) with a constrained ℓ_1 minimization method. Second, there are methods based on neighborhood estimation, as the Lasso procedure by Meinshausen and Bühlmann (2006), that performs a conditional regression for every variable against all other variables. Third, GGM inference can be done using multiple-testing procedures based on appropriate estimators of the precision matrix. When the sample size is large and the sample covariance matrix invertible, there are natural test statistics for the entries of the precision matrix (Drton and Perlman, 2007). However, in the high-dimensional setting, the sample covariance matrix is typically nonsingular, and there have been several propositions to circumvent this issue. The use of bagging or shrinkage provides more reliable estimates of the covariance matrix or its inverse (Schäfer and Strimmer, 2005; Schäfer and Strimmer, 2005). Alternatively, one can test the first-order partial correlation, that is, the conditional correlation of two variables given only another third one (Wille and Bühlmann, 2006). More recently, new estimators of the precision matrix were obtained by improving the original Lasso-regularized estimators (Liu, 2013; Ren et al., 2015; Jankova and van de Geer, 2018). As these estimates are asymptotically normal when the precision matrix is sparse, they can serve to infer the graph while controlling the FDR (Liu, 2013).

1.2 Graph inference by a multiple testing procedure

Rebafka et al. (2022) propose a multiple testing procedure for graph inference in the general statistical framework, where a real-valued graph is observed resulting from a perturbation of some underlying binary latent graph, which is to be recovered. We will show that this approach can be used for GGM inference, by using test statistics from the GGM literature as input to this method.

The procedure by Rebafka et al. (2022) consists in modeling the data by a suitable latent variable model and then using the associated structured ℓ -values, that is, the posterior probabilities of the presence of an edge, to infer a graph such that the FDR is controlled. This method is more powerful than other existing multiple testing procedures, since the ℓ -value of an entry (i, j) of the graph depends not only on the observed value for (i, j) , but also on the learned statistical graph model.

Data are modeled by the so-called noisy stochastic block model (NSBM), where the underlying binary graph is a classical stochastic block model, to which the following blurring mechanism is applied: absent edges are replaced with pure random noise, and in place of present edges, some effect or signal is observed. For the testing procedure the NSBM has to be fitted to the data, but inference is involved as the model is quite complex. In Rebafka et al. (2022) a variational EM-algorithm is proposed, but the implemented algorithm is rather slow and not scalable to large graphs. Therefore, it is not convenient for GGM inference, where graphs can be huge. In this work, a new inference algorithm is developed according to the approach proposed by Côme and Latouche (2015). While the variational EM-algorithm provides an approximation of the maximum likelihood estimate, the aforementioned method is a Bayesian approach and relies on a new objective function, the so-called exact integrated complete-data likelihood (ICL). Its maximization is a discrete optimization problem, that can be solved by a greedy algorithm. An important advantage of the algorithm is that it automatically selects the best model, that is, the best number of blocks in the NSBM.

1.3 Contributions and organization of the paper

In Section 2, the definition of the NSBM and the associated multiple testing procedure for graph inference are recalled. In Section 3 a fast greedy algorithm is developed to maximize the ICL criterion in the NSBM. Then, Section 4 shows how to use the NSBM for

GGM inference with control of the FDR. Namely, any known test statistic from the GGM literature can be used as an input for the NSBM approach. Finally, Section 5 provides a numerical study to assess the performance of our NSBM approach for GGM inference and a comparison to the state of the art on both simulated data and a real dataset.

2 Model and graph inference procedure

In this section we present notations, the model and the graph inference procedure.

2.1 Model definition

In the noisy stochastic block model (NSBM) the observation matrix $X = (X_{i,j})_{i,j} \in \mathbb{R}^{p \times p}$ is a perturbed version of some underlying binary latent graph represented by its adjacency matrix $A = (A_{i,j})_{i,j} \in \{0, 1\}^{p \times p}$, which is modeled by a stochastic block model (SBM). Here only undirected graphs without loops are considered. Denote $\mathcal{A} = \{(i, j) : 1 \leq i < j \leq p\}$ the set of all possible undirected edges among the p nodes of the graph.

The nodes are partitioned into Q blocks by a vector $Z = (Z_1, \dots, Z_p)$ which is composed of independent and identically distributed discrete latent variables Z_i with probabilities $\pi_q = \mathbb{P}(Z_1 = q)$ for $q \in \llbracket Q \rrbracket$ for some parameter $\pi = (\pi_1, \dots, \pi_Q) \in [0, 1]^Q$ such that $\sum_{q \in \llbracket Q \rrbracket} \pi_q = 1$. The entries $A_{i,j}$ of the latent graph A are independent conditionally on Z with Bernoulli distribution, that is,

$$(A_{i,j})_{(i,j) \in \mathcal{A}} \mid Z \sim \bigotimes_{(i,j) \in \mathcal{A}} \mathcal{B}(w_{Z_i, Z_j}),$$

for some symmetric parameter matrix $w = (w_{q,l})_{q,l} \in [0, 1]^{Q \times Q}$. The blurring mechanism is the following. If there is no edge between nodes i and j , i.e. $A_{i,j} = 0$, then $X_{i,j}$ is standard Gaussian noise, otherwise (when $A_{i,j} = 1$) some signal is observed and the signal strength

is modulated by the latent variables Z_i and Z_j . That is,

$$(X_{i,j})_{(i,j) \in \mathcal{A}} | Z, A \sim \bigotimes_{(i,j) \in \mathcal{A}} (1 - A_{i,j})\mathcal{N}(0, 1) + A_{i,j}\mathcal{N}(\mu_{Z_i, Z_j}, \sigma^2),$$

for an unknown symmetric parameter matrix $\mu = (\mu_{q,l})_{q,l} \in \mathbb{R}^{Q \times Q}$ and known variance $\sigma^2 > 0$. The unknown model parameter in the NSBM is denoted by $\theta = (\pi, w, \mu)$ and the distribution and density of (X, A, Z) are denoted by \mathbb{P}_θ and p_θ , respectively. In the Supplementary Material an extension of the NSBM is presented, where the variance of the Gaussian under the alternative is unknown and assumed to depend on the latent variables.

2.2 Graph inference via multiple testing

Graph inference in the NSBM can be viewed as the multiple testing problem, where for all pairs $(i, j) \in \mathcal{A}$ we consider the hypotheses of the absence or presence of an edge between nodes i and j in the latent graph A . That is, we test $H_{0,i,j} : A_{i,j} = 0$ against $H_{1,i,j} : A_{i,j} \neq 0$.

Rebafka et al. (2022) propose a powerful testing procedure based on the posterior probabilities that the null hypothesis is true, also called ℓ -values. They are given by

$$\ell_{i,j}(X, Z, \theta) = \mathbb{P}_\theta(A_{i,j} = 0 | X, Z) = \frac{(1 - w_{Z_i, Z_j})f_{\mathcal{N}(0,1)}(X_{i,j})}{(1 - w_{Z_i, Z_j})f_{\mathcal{N}(0,1)}(X_{i,j}) + w_{Z_i, Z_j}f_{\mathcal{N}(\mu_{Z_i, Z_j}, \sigma^2)}(X_{i,j})},$$

where $f_{\mathcal{N}(m, s^2)}$ denotes the Gaussian density with mean m and variance s^2 . The procedure consists in rejecting $H_{0,i,j}$ if the ℓ -value is sufficiently small, that is $\ell_{i,j}(X, Z; \theta) \leq t$ for some threshold t , which can be chosen such that the FDR is controlled. For more details we refer the reader to Rebafka et al. (2022).

The point of the ℓ -value $\ell_{i,j}$ for the node pair (i, j) is that its value does not only depend on the observation $X_{i,j}$, but on the entire NSBM through the model parameter θ and the node clustering Z . This results in an important increase of the power compared to the classical Benjamini-Hochberg procedure. Rebafka et al. (2022) also provide results on the optimality of this ℓ -values procedure in terms of FDR and the power. The approach

is typically used with estimated model parameters $\hat{\theta}$ and estimated latent variables \hat{Z} . Therefore, an efficient inference algorithm for the NSBM is required.

3 New inference algorithm in the NSBM

While Rebafka et al. (2022) propose a variational EM-algorithm limited to networks of very moderate size, in this section a much more efficient and more convenient inference algorithm for the NSBM is developed. The approach relies on an integrated likelihood criterion and comes with a fast greedy algorithm that computes the best clustering Z , estimates the model parameter θ and automatically performs model selection, that is the choice of the best number of blocks in the SBM.

3.1 ICL criterion in the NSBM

The integrated complete-data log-likelihood (ICL) is defined in a Bayesian framework, where π_θ denotes a prior on the parameter θ . In the NSBM the ICL is given by

$$\text{ICL}^{\text{NSBM}}(X, A, Z) = \log p(X, A, Z) = \log \left(\int_{\Theta} p_\theta(X, A, Z) \pi_\theta(\theta) d\theta \right).$$

Notably, the criterion only depends on the observations and the latent variables. The value of Z that maximize the ICL can be used to derive an estimator of θ , which in turn will be used to compute the ℓ -values for the final test procedure.

To compute the maximum it is convenient to have an analytic expression of ICL^{NSBM} . This is obtained by considering a factorized prior of the form $\pi_\theta(\pi, w, \mu) = \pi_\pi(\pi) \pi_w(w) \pi_\mu(\mu)$ and by using conjugate priors for the different parameters. Here we choose the Dirichlet distribution for the cluster proportions π , Beta distributions for the edge probabilities w and Gaussian distributions for the means μ . More precisely, for convenient hyperparameters

$n_0, \eta_0, \zeta_0, \rho_0, \tau_0$, we consider

$$\begin{aligned}\pi_\pi(\pi) &= \text{Dir}(n_0, \dots, n_0), & \pi_w(w) &= \bigotimes_{q \leq l} \pi_{w_{q,l}}(w_{q,l}) = \bigotimes_{q \leq l} \text{Beta}(\eta_0, \zeta_0), \\ \pi_\mu(\mu) &= \bigotimes_{q \leq l} \pi_{\mu_{q,l}}(\mu_{q,l}) = \bigotimes_{q \leq l} \mathcal{N}(\rho_0, \tau_0^2).\end{aligned}\tag{1}$$

To state the formulae of ICL^{NSBM} , we introduce some further notations and useful count statistics. Let $N = p(p-1)/2$ be the maximal number of possible edges, $M^A = \sum_{(i,j) \in \mathcal{A}} A_{i,j}$ the total number of edges and $N_Q = Q(Q+1)/2$ the number of pairs of blocks. Let $Z_{i,q} = \mathbb{1}\{Z_i = q\}$ indicate whether node i is assigned to block q or not, and $\mathcal{Z}_{q,l}^{i,j} = \mathbb{1}\{Z_{i,q}Z_{j,l} + Z_{i,l}Z_{j,q} > 0\}$ whether nodes i and j belong to blocks q and l . Then, $n_q = \sum_{i \in \llbracket p \rrbracket} Z_{i,q}$ is the number of nodes associated with block q , and $n_{q,l} = \sum_{(i,j) \in \mathcal{A}} \mathcal{Z}_{q,l}^{i,j} A_{i,j}$ (resp. $\bar{n}_{q,l} = \sum_{(i,j) \in \mathcal{A}} \mathcal{Z}_{q,l}^{i,j} (1 - A_{i,j})$) is the number of present (resp. absent) edges between nodes in q and l . Finally, $I_{q,l} = \{(i,j) \in \mathcal{A}, A_{i,j} = 1, \mathcal{Z}_{q,l}^{i,j} = 1\}$ is the set of present edges between nodes belonging to blocks q and l (such that $|I_{q,l}| = n_{q,l}$), and $S_{q,l}$ denotes the empirical variance associated with $\{X_{i,j}, (i,j) \in I_{q,l}\}$.

Proposition 1 *In the NSBM with the priors defined in (1), the ICL is given by*

$$\text{ICL}^{\text{NSBM}}(X, A, Z) = \text{ICL}^{\text{SBM}}(A, Z) + \text{ICL}^{\text{noise}}(X, A, Z),$$

where ICL^{SBM} is the ICL in the SBM given by

$$\begin{aligned}\text{ICL}^{\text{SBM}}(A, Z) &= \log p(Z) + \log p(A|Z) \\ &= \log \frac{C(n_0 + n_1, \dots, n_0 + n_Q)}{C(n_0, \dots, n_0)} + \sum_{q \leq l} \log \frac{B(\eta_0 + n_{q,l}, \xi_0 + \bar{n}_{q,l})}{B(\eta_0, \xi_0)},\end{aligned}\tag{2}$$

where $C(x_1, \dots, x_Q) = \frac{\prod_q \Gamma(x_q)}{\Gamma(\sum_q x_q)}$, $\Gamma(\cdot)$ is the Gamma function and $B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$ the Beta function. Furthermore, $\text{ICL}^{\text{noise}}(X, A, Z) = \log p(X|A, Z)$ is the part of the ICL criterion

associated with the noise layer in the NSBM given by

$$\begin{aligned} \text{ICL}^{\text{noise}}(X, A, Z) = & -\frac{N}{2} \log(2\pi) - \frac{1}{2} \sum_{(i,j) \in \mathcal{A}} (1 - A_{i,j}) X_{i,j}^2 - (M^A - N_Q) \log \sigma \quad (3) \\ & - \frac{1}{2} \sum_{q \leq l} \left\{ \log(\sigma^2 + \tau_0^2 n_{q,l}) + \frac{\tau_0^2 n_{q,l}^2}{\tau_0^2 \sigma^2 n_{q,l} + \sigma^4} S_{q,l} + \frac{1}{\sigma^2 + \tau_0^2 n_{q,l}} \sum_{(i,j) \in I_{q,l}} (X_{i,j} - \rho_0)^2 \right\}. \end{aligned}$$

3.2 Greedy NSBM inference algorithm

The maximization of the ICL is a discrete optimization problem and we propose a greedy algorithm for this task, which is in line with the algorithm developed by Côme and Latouche (2015). The algorithm successively picks a node and searches its best block assignment in terms of the ICL criterion while keeping all other node assignments fixed. This amounts to evaluate the variation of the ICL denoted by $\Delta_{g \rightarrow h}$ due to swapping the selected node i^* with current block membership g to block h , where $\Delta_{g \rightarrow h}$ is defined by

$$\Delta_{g \rightarrow h} = \text{ICL}^{\text{NSBM}}(X, Z^{\text{swap}}, A^{\text{swap}}) - \text{ICL}^{\text{NSBM}}(X, Z, A),$$

where $Z = (Z_1, \dots, Z_p)$ denotes the current partition of the nodes with $Z_{i^*} = g$ and Z^{swap} the block assignments after the swap, that is, $Z_{i^*}^{\text{swap}} = h$ and $Z_i^{\text{swap}} = Z_i$ for all $i \neq i^*$. Likewise, A denotes the current adjacency matrix and A^{swap} the one after the swap obtained by thresholding the posterior probabilities of edges involving node i^* , that is,

$$A_{i^*,j}^{\text{swap}} = A_{j,i^*}^{\text{swap}} = \mathbb{1}\{\rho_{i^*,j}^{\text{swap}} > 0.5\}, \quad \forall j \in \llbracket p \rrbracket, \quad A_{i,j}^{\text{swap}} = A_{i,j}, \quad \forall i \neq i^*, j \neq i^*. \quad (4)$$

where $\rho_{i,j}^{\text{swap}} = \mathbb{P}_{\hat{\theta}}(A_{i,j} = 1 \mid X, Z^{\text{swap}})$ and current parameter estimate $\hat{\theta}$ composed of

$$\hat{w}_{q,l} = \frac{\sum_{i < j} \mathcal{Z}_{q,l}^{i,j} \rho_{i,j}}{\sum_{i < j} \mathcal{Z}_{q,l}^{i,j}}, \quad \hat{\mu}_{q,l} = \frac{\sum_{i < j} \mathcal{Z}_{q,l}^{i,j} \rho_{i,j} X_{i,j}}{\sum_{i < j} \mathcal{Z}_{q,l}^{i,j} \rho_{i,j}}, \quad \hat{\sigma}_{q,l}^2 = \frac{\sum_{i < j} \mathcal{Z}_{q,l}^{i,j} \rho_{i,j} (X_{i,j} - \hat{\mu}_{q,l})^2}{\sum_{i < j} \mathcal{Z}_{q,l}^{i,j} \rho_{i,j}}, \quad (5)$$

where the posterior probabilities $\rho_{i,j} = \mathbb{P}_{\hat{\theta}}(A_{i,j} = 1 \mid X, Z)$ are used as weights to define estimates that have the form of weighted means.

In the greedy algorithm, as the variation $\Delta_{g \rightarrow h}$ is evaluated many times, its computation represents a huge part of the computational burden of the algorithm. Thus, it is important

that the numerical evaluation of $\Delta_{g \rightarrow h}$ is fast. Details on the efficient computation of $\Delta_{g \rightarrow h}$ are provided in the Supplementary Material.

It occurs that the algorithm successively moves all nodes from one block to other blocks resulting in an empty block. Thus, starting from a large number Q_{init} of blocks, the approach automatically performs model selection by emptying blocks and providing a final clustering with an appropriate number of clusters. The ICL criterion can be considered as a penalized criterion encouraging the use of a parsimonious model.

The greedy ICL algorithm is summarized in Algorithm 1. The estimates of the model parameter θ and the clustering Z are then used to compute the ℓ -values (obtained in fact directly from the posterior edges probabilities ρ) to infer a graph such that the FDR is controlled. The code will soon be available in the R package `noisysbmGGM`.

A numerical comparison of the method with the variational EM algorithm implemented in the package `noisySBM` showed that both methods provide essentially the same estimates when applied to the same dataset, while the greedy ICL algorithm largely outperforms EM in terms of computing time. Furthermore, graphs of larger size can be treated with this new algorithm, which makes it suitable for GGM inference.

4 GGM inference using the NSBM approach

This section describes how to use the multiple testing approach using the NSBM to infer a GGM. In particular, different test statistics from the GGM literature are presented, which are convenient for building the observation matrix X provided to the NSBM algorithm.

4.1 GGM test statistics for the NSBM

Let $\mathbf{Y} = (Y_1, \dots, Y_p)^t$ be a centered random vector with multivariate Gaussian distribution with nonsingular covariance matrix Σ . The inverse of the latter, denoted by $\Omega = \Sigma^{-1} =$

Algorithm 1: Greedy ICL algorithm for the NSBM

Input : Observation X , initial number of latent blocks Q_{init}

Output: Clustering Z into Q clusters, estimator θ , edge probabilities ρ

Initialize Z , θ , $\rho = (\rho_{ij})_{i,j \in \mathcal{A}}$ and A ; Set $Q = Q_{\text{init}}$;

repeat

 Select a node $i^* \in \llbracket p \rrbracket$; Set $g \leftarrow Z_{i^*}$;

for $h \neq g$ **do**

 Calculate $\Delta_{g \rightarrow h}$;

end

 Find h such that $\Delta_{g \rightarrow h}$ is maximum;

if $\Delta_{g \rightarrow h} > 0$ **then**

 Update clustering Z by setting $Z_{i^*} = h$;

if *block g is empty* **then**

 Shrink the number of blocks: $Q = Q - 1$;

end

 Update θ , $\rho = (\rho_{ij})_{i,j \in \mathcal{A}}$ and A based on the updated clustering Z ;

end

until *convergence*;

$(\omega_{i,j})_{i,j}$, is called the precision matrix and represents the partial correlations. More precisely,

$$\text{Corr}(Y_i, Y_j | \mathbf{Y}_{-(i,j)}) = -\frac{\omega_{i,j}}{\sqrt{\omega_{i,i}\omega_{j,j}}}, \quad \text{where} \quad \mathbf{Y}_{-(i,j)} = \{Y_k, k \notin \{i, j\}\}.$$

The interpretation is that $\omega_{i,j} = 0$ if and only if Y_i and Y_j are independent given all other variables $\mathbf{Y}_{-(i,j)}$. In other words, if $\omega_{i,j} \neq 0$, then there is a direct dependency or link between Y_i and Y_j (Lauritzen, 1996). Now, the aim of GGM inference is to detect the pairs of random variables with nonzero partial correlation on the basis of a sample of \mathbf{Y} . This amounts to infer the binary graph with adjacency matrix A with $A_{i,j} = \mathbf{1}\{\omega_{i,j} \neq 0\}$.

The literature provides various test statistics computed on a sample of size n of \mathbf{Y} for

testing whether partial correlations are zero or not. They can be used to directly build a multiple testing procedure, but in this work we use them to build the observation matrix X fed to the greedy NSBM algorithm. The following test statistics are considered:

- When the sample size n exceeds the number of variables p , the inverse of the sample covariance matrix is an estimate of the precision matrix. Applying Fisher’s z -transformations to the associated partial correlations provides asymptotically normal test statistics (Anderson, 2003).
- Ren et al. (2015) use the residuals of bivariate regressions that model the outcome of two variables Y_i and Y_j as a function of all other variables to estimate the precision matrix. We refer to this test statistic as **Ren**. Asymptotic normality is obtained in the sparse setting, where the maximum degree s of the graph satisfies $s = o(\sqrt{n}/\log p)$.
- Jankova and van de Geer (2018) propose test statistics inspired by the debiasing approach in high-dimensional linear regression. Two variants are available, the first one uses graphical Lasso, here referred to as **JankovaGL**, the second is a nodewise Lasso approach, denoted by **JankovaNW**. In the same sparse regime as above, these test statistics are asymptotically normal.
- Liu (2013) uses a bias correction of the sample covariance of residuals when considering the conditional regression for the variables Y_i and Y_j . Two variants are available depending on the method for the estimation of the regression parameters: Lasso regression and a scaled Lasso regression, denoted by **LiuSL** resp. **LiuL**. Contrary to the other methods, the asymptotic variance of these test statistics is unknown. Thus, we will apply an extension of Algorithm 1 that estimates the limit variance (see Supplementary Material).

4.2 Multiple testing procedures for GGM inference

Liu (2013) proposes a classical multiple testing approach for GGM inference based on the test statistics he proposed and an adjustment of the significance level of the tests. The approach achieves the asymptotic control of the FDR at the desired level α in the sparse setting. It is straightforward to use Liu’s procedure with all other test statistics described above and this is done in our numerical study in the next section, where Liu’s multiple testing approach serves as a baseline to which our NSBM approach will be compared. The main difference of Liu’s approach and ours is that Liu only applies a correction of the significance level, but does not model neither the joint distribution of the test statistics nor the graph structure.

In detail, in the numerical study we consider the challenging setting where $p > n$. We denote by `Ren-classical`, `JankovaNW-classical`, `JankovaGL-classical`, `LiuSL-classical`, `LiuL-classical` the classical multiple testing procedure by Liu (2013) applied to the different test statistics. Computations are performed with the R package `SILGGM` (Zhang et al., 2018). Then, we apply the analogous procedures with our NSBM approach, that is Algorithm 1 applied to the same test statistics, followed by the multiple testing procedure based on ℓ -values. We denote them by `Ren-NSBM`, `JankovaNW-NSBM`, `JankovaGL-NSBM`, `LiuSL-NSBM`, `LiuL-NSBM`. For `LiuSL-NSBM` and `LiuL-NSBM` the version of the NSBM with unknown variances under the alternative (see Section 2 of the Supplementary Material) is used. This is done with our R package `noisysbmGGM`. Furthermore, the following popular procedures that do not control any error rate are applied: the procedure by Meinshausen and Bühlmann (2006) (`MB`) and graphical Lasso (`glasso`). Finally, the CLIME procedure (`clime`) (Cai et al., 2011; Li et al., 2022) is applied to complete the picture.




Figure 1: Illustration of the graph structures in the four settings.

5 Numerical study

In this section various numerical experiments are conducted to assess the performance of our NSBM approach for GGM inference, to compare to the state of the art, and to evaluate robustness to model misspecification.

5.1 Comparison of test statistics on synthetic data

Data are generated from a Gaussian distribution $\mathcal{N}_p(0, \Sigma)$ with different choices for the precision matrix $\Omega = (\omega_{i,j})_{i,j} = \Sigma^{-1}$. More precisely, we consider different underlying graphs $A = (\mathbb{1}\{\omega_{i,j}\})_{i,j}$ illustrated in Figure 1. In the SBM setting (a), the graph A is the realization of a SBM with five communities of equal size and varying probabilities of connection among the blocks. To evaluate the robustness of our procedure to model misspecification, we also consider three settings where A is not a SBM. In the hub setting (b), the set of nodes is partitioned into groups of size 10. In each group one node is a hub and as such connected to all other nodes in the group. In the band setting (c), every node i is connected to all nodes j at distance 3, that is $A_{i,j} = \mathbb{1}\{1 \leq |i - j| \leq 3\}$. This yields a matrix A with bands of ones around the diagonal. Finally, in the scale-free setting (d), the graph is generated by preferential attachment (Barabasi and Albert, 1999). That is, the graph is built iteratively starting from a small initial graph, here consisting of only two connected nodes. Then new nodes are added one by one and each of them is connected to

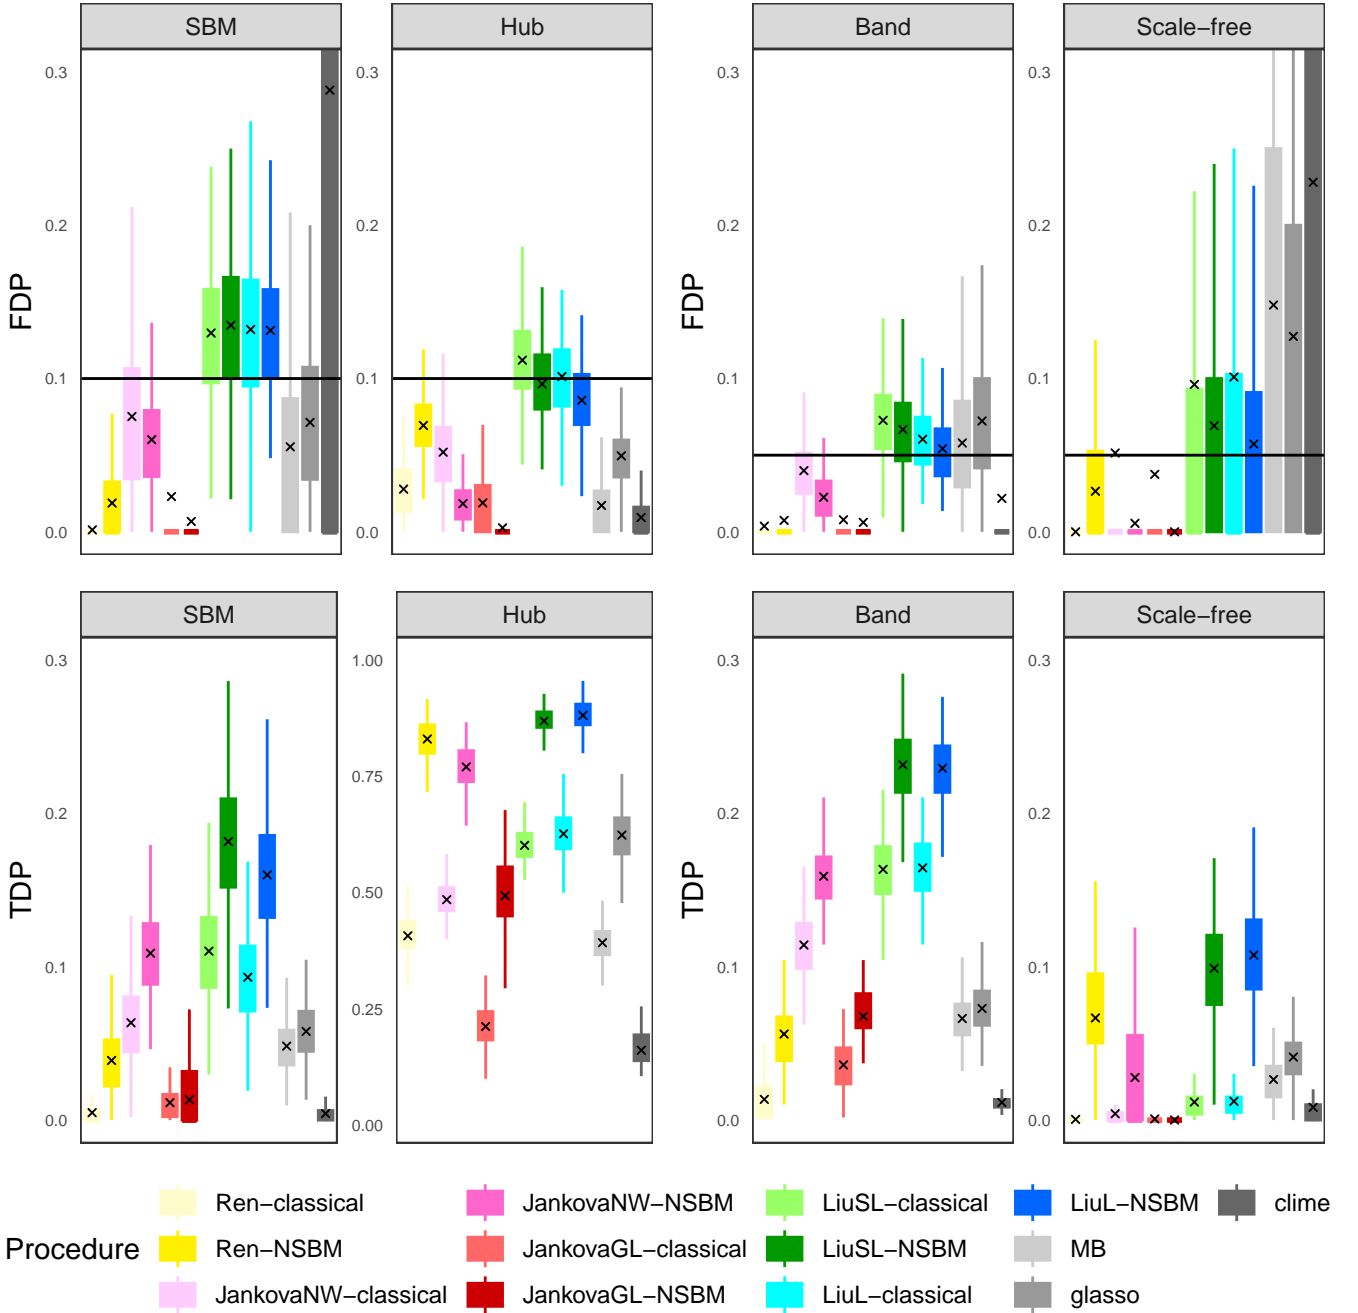


Figure 2: Boxplots of the FDP and TDP for all procedures described in Section 4.2 for the settings SBM, hub, band and scale-free for 200 simulated datasets, where $p = 200$ and $n = 100$. Horizontal lines represent the nominal level α , with $\alpha = 0.1$ in the SBM and the hub setting and $\alpha = 0.05$ in the band and scale-free cases. The crosses in the boxplots correspond to the sample FDR (resp. TDR) defined as the mean of the FDP (resp. TDP).

a single node in the graph with probabilities proportional to the current node degrees.

Now, to generate a precision matrix Ω and a covariance matrix Σ associated with a given adjacency matrix A , we proceed as in the R package `huge`. That is, first a positive definite matrix Ω is obtained by the transformation

$$\Omega = \gamma A + \text{diag} \left(\left| \min_{\lambda \in \text{Sp}(\gamma A)} (\lambda) \right| + \beta \right),$$

where $\gamma = 0.3$, $\beta = 0.2$ and $\text{Sp}(\gamma A)$ denotes the spectrum of γA . In other words, the ones in A are replaced by γ and the diagonal is such that Ω is invertible. Hence, $\Sigma = \Omega^{-1}$ is well defined and realizations of the desired Gaussian distribution can be simulated.

In the simulations there are $p = 200$ variables and the sample size is $n = 100$. In every setting all procedures described in Section 4.2 are evaluated on 200 datasets with nominal level $\alpha = 0.1$ in the SBM and the hub setting and $\alpha = 0.05$ in the band and scale-free cases. The performance of a procedure is evaluated by comparing the inferred graph \hat{A} to the true adjacency matrix A . In particular, the false discovery proportion (FDP) is the proportion of errors among the edges declared as significant, whereas the proportion of edges that are correctly declared as significant among the true edges in A is referred to as the true discovery proportion (TDP). They are defined as

$$\text{FDP} = \frac{|\{(i, j) : \hat{A}_{i,j} = 1, A_{i,j} = 0\}|}{|\{(i, j) : \hat{A}_{i,j} = 1\}| \vee 1}, \quad \text{TDP} = \frac{|\{(i, j) : \hat{A}_{i,j} = 1, A_{i,j} = 1\}|}{|\{(i, j) : A_{i,j} = 1\}|}.$$

For procedures that are supposed to control the FDR, the FDP should in average be lower or equal than α , while the TDP representing the power of the procedure is desired to be as large as possible.

Figure 2 displays the results. First, we observe that concerning the FDR, for all Ren- and Jankova-procedures the FDR is always controlled at the nominal level α . The FDR of all Liu-procedures slightly exceeds α , but the Liu-NSBM procedures mostly do better than their classical inference counterparts. The last three procedures in the figure are not conceived for controlling the FDR, so the comparison to the nominal level does not really make sense since the output of these procedures is the same for any α .

Second, in terms of power, in every setting the largest TDRs are achieved by NSBM based procedures, in particular our approach with both of Liu’s test statistics appears to be the most powerful procedures. Notably, also the power of the methods that do not require the control of the FDR, namely the one by Meinshausen & Bühlmann, graphical Lasso and CLIME, is largely below those of most other procedures. Moreover, in all settings the NSBM approach is more powerful than the classical inference procedure using the same test statistics. That is, whatever test statistic is used, there is a gain in power when using our NSBM approach. This is in line with what has been observed for other multiple testing procedures, that is that using a latent structure and learning the model is beneficial (Sun and Cai, 2009; Cai and Sun, 2009; Liu et al., 2016; Rebafka et al., 2022). Furthermore, these results confirm the observation done in Rebafka et al. (2022) that the SBM is an appropriate choice for graph inference, as the SBM is a highly flexible random graph model accommodating a wide spectrum of graph topologies. Here we see that the SBM is also appropriate in the framework of GGM inference. In other words, our approach is robust to model misspecification as good results are also obtained in the hub, band and scale-free settings in our study.

5.2 Sparse setting

As in practice, most graphs are rather sparse, we now investigate the behaviour of our procedure in the sparse setting described in terms of the maximum degree in the graph. Indeed, most of the test statistics that we consider are known to be asymptotically normal in the sparse setting when $s = o(\sqrt{n}/\log(p))$, where s is the maximum node degree in the graph. For this, we simulate data with given maximum degree. More precisely, given a maximum degree s , we first generate a sequence of degrees according to a power law with exponent 2 and then try to simulate a graph A with this degree sequence. Then, we proceed as previously to generate Gaussian vectors.




Figure 3: TDR (Average TDP) evaluated for different procedures over 200 simulations versus the maximum degree s of the graphs ($p = 100, n = 100, \alpha = 0.1$). In blue and filled red are respectively represented the classical inference procedures and our NSBM procedures with the different test statistics.

We carry out simulations with graphs composed of $p = 100$ nodes and $n = 100$ observations, for a maximum degree s in the range from 3 to 20. The average FDR and average TDR are estimated over 200 simulations and results are presented in Figure 3. As expected, the larger the maximum degree s , the lower the power. This is true for all methods, but the decrease in power of the NSBM procedures is less important than for the classical ones. Thus, there is a clear benefit of using our procedure for GGM inference in sparse settings.

5.3 Validation on a human T-cell dataset

Now, we investigate the performance of our procedure on a real dataset. We consider the multivariate flow cytometry data produced by Sachs et al. (2005) concerning a human T-cell signalling pathway, whose deregulation may lead to carcinogenesis. The data are available from <https://www.science.org/doi/10.1126/science.1105809> and were extensively studied in the literature (see e.g. Sachs et al., 2005; Henao and Winther, 2011; Desgranges, 2015; Eaton and Murphy, 2007; Ramsey and Andrews, 2018). The amount of 11 specific proteins in single cells was simultaneously measured under different perturbation conditions of the cell. We focus on a general perturbation (anti-CD3/CD28 + ICAM-2) that overall stimulates the cellular signalling network. In this condition measurements for 902 cells are available.

As the dataset is rather huge, it can be used to first compute a benchmark on the full dataset and then assess the performance of the inference procedures on small subsets of the data. More precisely, on the full dataset various methods, namely the classical procedure with Liu’s test statistics LiuL and nominal level $\alpha = 0.05$, yield the graph displayed in Figure 4. The inferred graph contains ten edges, nine of which are interactions well known from the literature. The additional edge p38-JNK was also inferred by Sachs et al. (2005) but with low confidence, and also was found by Henao and Winther (2011); Desgranges (2015); Eaton and Murphy (2007).

Now to evaluate the ability to recover these edges from a smaller dataset, we sample subsets of $n = 20$ cells and apply both the classical and the NSBM procedures with the test statistic LiuL with $\alpha = 0.05$. Table 1 reports the number of times the edges of the benchmark network from Figure 4 are detected with each method on 200 sampled datasets. Obviously, the NSBM procedure detects all edges more often than the classical procedure. That is, there is a real gain in power also on real datasets. This confirms that using our procedure helps to improve GGM inference.

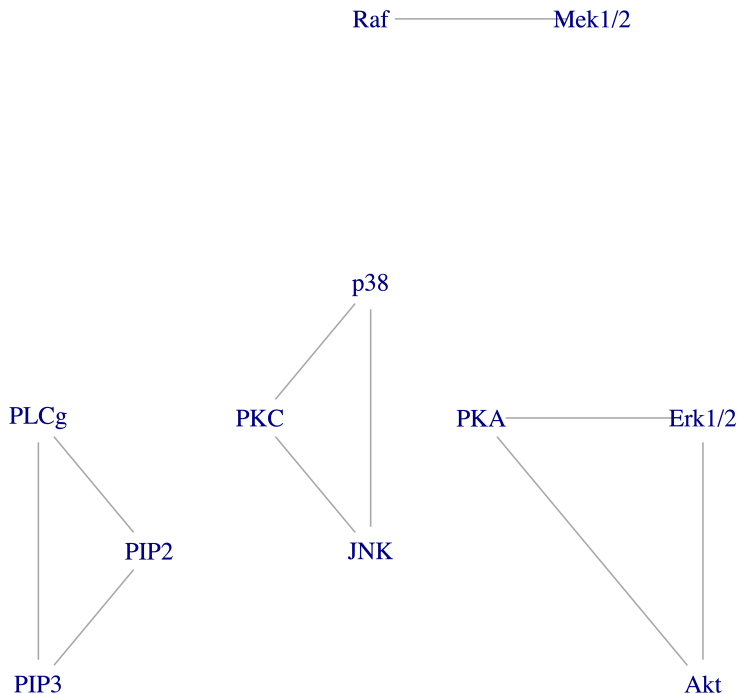


Figure 4: Benchmark obtained on the full dataset.

6 Conclusion

In this work a new approach for GGM inference is proposed that relies on the modeling of the graph structure by a latent variable model. It allows to control the false discovery rate and outperforms state-of-the-art methods in terms of power, that is, more “true” edges are detected than with classical methods. The performance is assessed in various numerical experiments and the procedure is particularly suited for real datasets, since it is robust to model misspecification and works on sparse networks as well as on small datasets.

It is noteworthy that this is a generic approach that can be used in combination with any test statistic for GGM inference. It can be viewed as a layer on top of the traditional problem of constructing a test statistic. So the user may consider his or her favorite test statistic and provide them to our NSBM method for GGM inference. We have shown that this consistently yields more powerful results than with the classical inference approach.

Moreover, the new greedy ICL algorithm for fitting the NSBM to a dataset is much more

edge	LiuL-classical	LiuL-NSBM
Raf - Mek1/2	200	200
PLCg - PIP2	28	44
PLCg - PIP3	121	143
PIP2 - PIP3	167	172
Erk1/2 - Akt	199	199
Erk1/2 -PKA	42	54
Akt - PKA	92	121
PKC - p38	146	158
PKC - JNK	107	121
p38 - JNK	102	115

Table 1: The number of times the 10 edges of Figure 4 are detected over 200 subsampled datasets with sample size $n = 20$ for both the classical and the NSBM procedure with Liu’s test statistics LiuL with $\alpha = 0.05$.

attractive in terms of computing time and scalability to large graphs than the variational EM algorithm. It is noteworthy that the ICL algorithm can be used beyond the context of NSBM-based GGM inference.

References

- Anderson, T. (2003). *An Introduction to Multivariate Statistical Analysis*. Wiley Series in Probability and Statistics. Wiley.
- Banerjee, O., L. El Ghaoui, and A. d’Aspremont (2008). Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal Machine Learning Research* 9, 485–516.

- Barabasi, A. and R. Albert (1999). Emergence of scaling in random networks. *Science* 286(5439), 509–512.
- Cai, T., W. Liu, and X. Luo (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association* 106(494), 594–607.
- Cai, T. T. and W. Sun (2009). Simultaneous testing of grouped hypotheses: Finding needles in multiple haystacks. *Journal of the American Statistical Association* 104(488), 1467–1481.
- Côme, E. and P. Latouche (2015). Model selection and clustering in stochastic block models based on the exact integrated complete data likelihood. *Statistical Modelling* 15(6), 564–589.
- Desgranges, N. (2015). *Generalization of the PC algorithm for non-linear and non-Gaussian data and its application to biological data*. Ph. D. thesis, ING - Scuola di Ingegneria Industriale e dell’Informazione.
- Drton, M. and M. D. Perlman (2007). Multiple testing and error control in Gaussian graphical model selection. *Statistical Science* 22(3), 430 – 449.
- Eaton, D. and K. Murphy (2007). Exact Bayesian structure learning from uncertain interventions. In M. Meila and X. Shen (Eds.), *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, Volume 2 of *Proceedings of Machine Learning Research*, San Juan, Puerto Rico, pp. 107–114. PMLR.
- Friedman, J., T. Hastie, and R. Tibshirani (2007). Sparse inverse covariance estimation with the graphical Lasso. *Biostatistics* 9(3), 432–441.
- Henao, R. and O. Winther (2011). Sparse linear identifiable multivariate modeling. *Journal of Machine Learning Research* 12(24), 863–905.

- Jankova, J. and S. van de Geer (2018). Inference in high-dimensional graphical models. In M. Maathuis, M. Drton, and M. Lauritzen, S. Wainwright (Eds.), *Handbook of Graphical Models*. CRC Press.
- Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press.
- Li, S., T. T. Cai, and H. Li (2022). Transfer learning in large-scale Gaussian graphical models with false discovery rate control. *Journal of the American Statistical Association* *0*(0), 1–13.
- Liu, J., C. Zhang, and D. Page (2016). Multiple testing under dependence via graphical models. *The Annals of Applied Statistics* *10*(3), 1699–1724.
- Liu, W. (2013). Gaussian graphical model estimation with false discovery rate control. *The Annals of Statistics* *41*(6), 2948 – 2978.
- Meinshausen, N. and P. Bühlmann (2006). High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics* *34*(3), 1436 – 1462.
- Ramsey, J. and B. Andrews (2018). FASK with interventional knowledge recovers edges from the Sachs model. ArXiv:1805.03108.
- Rebafka, T., É. Roquain, and F. Villers (2022). Powerful multiple testing of paired null hypotheses using a latent graph model. *Electronic Journal of Statistics* *16*(1), 2796 – 2858.
- Ren, Z., T. Sun, C.-H. Zhang, and H. H. Zhou (2015). Asymptotic normality and optimalities in estimation of large Gaussian graphical models. *The Annals of Statistics* *43*(3), 991–1026.
- Sachs, K., O. Perez, D. Pe’er, D. Lauffenburger, and G. Nolan (2005). Causal protein-signaling networks derived from multiparameter single-cell data. *Science (New York, N. Y.)* *308*, 523–9.

- Schäfer, J. and K. Strimmer (2005). An empirical Bayes approach to inferring large-scale gene association networks. *Bioinformatics* 21(6), 754–764.
- Schäfer, J. and K. Strimmer (2005). A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Statistical Applications in Genetics and Molecular Biology* 4(1).
- Sun, W. and T. T. Cai (2009). Large-scale multiple testing under dependence. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 71(2), 393–424.
- Wille, A. and P. Bühlmann (2006). Low-order conditional independence graphs for inferring genetic networks. *Statistical Applications in Genetics and Molecular Biology* 5, 1–34.
- Zhang, R., Z. Ren, and W. Chen (2018). SILGGM: An extensive R package for efficient statistical inference in large-scale gene networks. *PLOS Computational Biology* 14(8), 1–14.