

Accurate directional inference in Gaussian graphical models

Claudia Di Caterina

`claudia.dicaterina@univr.it`

Department of Economics, University of Verona
37129 Verona, Italy

Nancy Reid

`reid@ustat.toronto.edu`

Department of Statistical Sciences, University of Toronto
Toronto, Canada M5S 3G3

Nicola Sartori

`sartori@stat.unipd.it`

Department of Statistical Sciences, University of Padova
35121 Padova, Italy

June 22, 2022

Abstract

Directional tests to compare incomplete undirected graphs are developed in the general context of covariance selection for Gaussian graphical models. The exactness of the underlying saddlepoint approximation leads to exceptional accuracy of the proposed approach. This is verified by simulation experiments with high-dimensional parameters of interest, where inference via standard asymptotic approximations to the likelihood ratio test and some of its higher-order modifications fails. The directional p -value is used to illustrate the assessment of Markovian dependencies in a dataset from a veterinary trial on cattle. A second example with microarray data shows how to select the graph structure related to genetic anomalies due to acute lymphocytic leukemia.

Keywords: Covariance Selection; Exponential Family Model; Higher-order Asymptotics; Likelihood Ratio Test; Saddlepoint Approximation; Undirected Graph.

1 Introduction

Undirected graphical models have gained considerable success in a variety of fields, including medicine, social sciences and physics, due to their flexibility and easy interpretation. Typically, these probabilistic graphs describe complex multivariate distributions of variables (nodes) through the product of simpler sub-models, each referred to a low-dimensional subset of the graph (clique). Book-length expositions on the topic can be found in Lauritzen (1996), Borgelt and Kruse (2002), and Whittaker (2009).

Today, applications of graphical models are challenged by the growth in size and sophistication of modern data. An important question is inferring the structure of large graphs, i.e. the underlying connections (edges) between the variables under examination. This task is well

known in the literature by the name of covariance selection. For reasons of convenience, a graphical model is often expressed by means of the exponential family form. The Gaussian distribution is particularly suitable for continuous responses, as conditional independence in the graph can be easily characterized in terms of assumptions on model parameters (see Section 3.1).

Directional inference on a vector-valued parameter of interest was introduced by Fraser and Massam (1985) in nonnormal linear regression models and then generalized in Skovgaard (1988). Substantial progress from both a methodological and computational perspective was made by Davison et al. (2014), where the computation of the directional p -value by one-dimensional numerical integration proved especially accurate in several settings. The procedure was extended from linear exponential families to nonlinear parameters of interest in general continuous models by Fraser et al. (2016). Besides its accuracy, the directional approach was found to coincide with exact results in several classical situations (McCormack et al., 2019).

In this paper we show how to employ directional tests for covariance selection in Gaussian graphical models. Davison et al. (2014, Section 5.3) derived the directional p -value in the special case of testing an incomplete graph versus the saturated model with connections between all nodes. The corresponding test was later found to be exact by Huang et al. (2021) and is computationally much simpler than the extension proposed here. Specifically, we allow the use of directional p -values in common situations where some edges in the graph are known a priori to be missing. Simulation studies confirm the exceptional performance of the directional approach, even in high-dimensional scenarios where the number of nodes is of the same order of magnitude as the sample size; the likelihood ratio statistic and its higher-order modifications (Skovgaard, 2001) break down in these settings.

Section 2 reviews directional inference in exponential family models. Section 3 sets the notation to deal with Gaussian graphs and derives the steps to carry out covariance selection based on directional p -values. Simulation results comparing the accuracy of the various methods are shown in Section 5, while Section 6 reports applications to data from a veterinary trial and from a microarray study of altered gene expression in acute lymphocytic leukemia. Comments and final remarks can be found in Section 7.

2 Background

2.1 Likelihood ratio tests

Assume that y follows a parametric distribution $f(y; \theta)$, with $\theta \in \mathbb{R}^p$. The log-likelihood function $\ell(\theta) = \ell(\theta; y) = \log f(y; \theta)$ is maximized by the maximum likelihood (ML) estimator $\hat{\theta} = \hat{\theta}(y)$. Whenever appropriate, the notation $\ell^0(\theta) = \ell(\theta; y^0)$ and $\hat{\theta}^0 = \hat{\theta}(y^0)$ will be adopted to stress the dependence of those quantities on the observed data point y^0 . Possibly after a reparameterization, the model parameter can be typically expressed as $\theta = (\psi, \lambda)$, where $\psi(\theta)$ is the d -dimensional component of interest involved in the hypothesis $H_\psi: \psi(\theta) = \psi$. We write $\hat{\theta}_\psi = (\psi, \hat{\lambda}_\psi)$ to denote the constrained ML estimator of θ under H_ψ .

Under usual regularity conditions (see, e.g., Cox and Hinkley, 1974, Section 9.3), the first-order approximation to the distribution of $\hat{\theta}$ is normal with mean θ and estimated covariance matrix $j(\hat{\theta})^{-1}$, with $j(\theta) = -\partial^2 \ell(\theta) / \partial \theta \partial \theta^\top$ the observed Fisher information matrix. The hypothesis H_ψ can be tested via the likelihood ratio statistic

$$w(\psi) = 2\{\ell(\hat{\theta}) - \ell(\hat{\theta}_\psi)\}, \quad (1)$$

which enjoys the property of invariance to parameterizations and has an approximate χ_d^2 distribution under the null hypothesis H_ψ , d being the dimension of the parameter of interest ψ .

Skovgaard (2001) introduced two modifications to (1), namely

$$w^*(\psi) = w(\psi) \left\{ 1 - \frac{\log \gamma(\psi)}{w(\psi)} \right\}^2 \quad \text{and} \quad w^{**}(\psi) = w(\psi) - 2 \log \gamma(\psi), \quad (2)$$

and showed that the limiting distribution of both test statistics based on the correction factor $\gamma(\psi)$ is also χ_d^2 . These modifications were obtained by analogy with the derivations for scalar parameters of interest of modifications to the square root of $w(\psi)$, the so-called r^* approximation of Barndorff-Nielsen (1986), further discussed in Fraser et al. (1999). Skovgaard (2001) emphasized not only the simplicity of computation of the adjustment, especially when compared to Bartlett (1937) correction using moments, but also its excellent large-deviation results.

Tests based on $w(\psi)$, including $w^*(\psi)$, $w^{**}(\psi)$ and the Bartlett-corrected $w(\psi)$, provide omnibus measures of departure of the data from H_ψ : the resulting p -value averages the deviations from the null hypothesis in all potential directions of the parameter space. In the next section, the approach of Davison et al. (2014, Section 3) for measuring the departure from H_ψ only in the direction indicated by the observed data will be reviewed. For a more complete exposition of the difference between omnibus and directional tests, see Fraser and Reid (2006).

2.2 Directional tests in linear exponential families

Following Fraser et al. (2016, Section 3), three steps are in general necessary to obtain the directional p -value when inference is made about a vector parameter of interest in the presence of a finite dimensional nuisance component. Focusing on hypotheses that are linear in the canonical parameter $\varphi = \varphi(\theta)$ of an exponential family model, we shall summarize here the procedure detailed in Davison et al. (2014, Section 3).

Denoting by $u = u(y)$ the sufficient statistic for the p -dimensional vector parameter φ , the first decrease in dimension from the sample size n to p is delivered by considering the marginal density of u and the corresponding log-likelihood function $\ell(\theta; u) = \varphi(\theta)^\top u - K\{\varphi(\theta)\}$, taking the standard exponential family form. Specifically, given $u^0 = u(y^0)$ and the centered statistic $s = u - u^0$, the tilted log-likelihood function is defined by

$$\ell(\theta; s) = \varphi(\theta)^\top s + \ell^0(\theta), \quad (3)$$

where $\ell^0(\theta) = \ell(\theta; u = u^0)$ and the observed value of s equals $s^0 = 0$.

When linearity in φ applies to both the interest and nuisance parameters, meaning $\varphi = \theta = (\psi, \lambda)$, expression (3) can be written as

$$\ell(\varphi; s) = \psi^\top s_1 + \lambda^\top s_2 + \ell^0(\psi, \lambda), \quad (4)$$

where ψ and s_1 have dimension d . Hence, the second dimensionality reduction from p to d follows directly from conditioning on the component of the statistic sufficient for λ . Indeed, the conditional distribution of s_1 given s_2 depends on ψ only and still belongs to the exponential family (cf. equation (5) in Davison et al., 2014, Section 3.1). Such a conditioning translates into fixing $\hat{\varphi}_\psi = (\psi, \hat{\lambda}_\psi)$ at the observed value $\hat{\varphi}_\psi^0 = (\psi, \hat{\lambda}_\psi^0)$.

The saddlepoint approximation to this conditional distribution (Barndorff-Nielsen and Cox, 1979) is typically very accurate, and produces the following density for the reduced model in \mathbb{R}^d :

$$h(s; \psi) = c \exp[\ell(\hat{\varphi}_\psi^0; s) - \ell\{\hat{\varphi}(s); s\}] |J_{\varphi\varphi}\{\hat{\varphi}(s); s\}|^{-1/2}, \quad s \in \mathcal{L}^0, \quad (5)$$

where c is a normalization constant and \mathcal{L}^0 is the d -dimensional plane described by setting $s_2 = 0$, or equivalently $\hat{\varphi}_\psi = \hat{\varphi}_\psi^0$. In addition, $\hat{\varphi}(s)$ solves in φ the score equation from the log-likelihood (4), $s = -\ell_\varphi^0(\varphi) = -\partial \ell^0(\varphi) / \partial \varphi$, and $J_{\varphi\varphi}(\varphi; s) = -\partial^2 \ell(\varphi; s) / \partial \varphi \partial \varphi^\top$.

The third step consists of constructing, within this simpler model, a one-dimensional conditional distribution for s along the direction indicated by the data. With this aim, denote by s_ψ the expectation of s under model (5) if H_ψ holds, that is the value of s for which $\varphi = \hat{\varphi}_\psi^0$ is the constrained ML estimate:

$$s_\psi = -\ell_\varphi^0(\hat{\varphi}_\psi^0) = \begin{bmatrix} -\ell_\psi^0(\hat{\varphi}_\psi^0) \\ 0 \end{bmatrix}, \quad (6)$$

depending on the observed data point y^0 . The line \mathcal{L}^* , in \mathcal{L}^0 , which joins the observed value $s^0 = 0$ and the expected value s_ψ can be parameterized by a scalar $t \in \mathbb{R}$:

$$s(t) = s_\psi + t(s^0 - s_\psi) = (1 - t)s_\psi,$$

and consequently the maximum likelihood estimate $\hat{\varphi}(s)$ in (5) can vary with $s(t)$. The approximation (5) constrained to \mathcal{L}^* is used to compute the p -value, the probability that $s(t)$ is as far or farther from s_ψ than is the observed value $s^0 = 0$. The directional p -value which measures the deviation from H_ψ along the line \mathcal{L}^* is thus

$$p(\psi) = \frac{\int_1^{t_{\sup}} t^{d-1} h\{s(t); \psi\} dt}{\int_0^{t_{\sup}} t^{d-1} h\{s(t); \psi\} dt}, \quad (7)$$

where $t = 0$ and $t = 1$ correspond respectively to $s = s_\psi$ and to the observed value $s^0 = 0$. The factor t^{d-1} is due to the Jacobian of the transformation from the variable s to $(\|s\|, s/\|s\|)$ (Davison et al., 2014, Section 3.2). The upper limit of the integrals in (7) is the largest value of t for which the ML estimator corresponding to $s(t)$ exists, and in some special situations can be determined analytically.

Being based on the saddlepoint approximation (5), the relative error of formula (7) is at most $O(n^{-3/2})$, where n is the number of independent observations from a continuous model. There are however many situations in which the normalized saddlepoint approximation is in fact the exact conditional density of the minimal sufficient statistic (McCormack et al., 2019), and such a normalization is implicit in (7). The theoretical properties of directional p -values in linear exponential families are detailed in Davison et al. (2014, Section 3.2).

Using the notation established in this section, we can also derive the form of the term $\gamma(\psi)$ appearing in (2) under exponential family models. Specifically, equation (13) in Skovgaard (2001) is

$$\gamma(\psi) = \frac{\{(s - s_\psi)^\top J_{\varphi\varphi}^{-1}(\hat{\varphi}_\psi)(s - s_\psi)\}^{d/2}}{w^{d/2-1}(\hat{\varphi} - \hat{\varphi}_\psi)^\top (s - s_\psi)} \left\{ \frac{|J_{\varphi\varphi}(\hat{\varphi}_\psi)|}{|J_{\varphi\varphi}(\hat{\varphi})|} \right\}^{1/2}, \quad (8)$$

to be evaluated at $s = 0$ when computing the corresponding observed p -value.

3 Directional tests for Gaussian graphical models

3.1 Notation and setup

Gaussian graphical models are very useful for describing normal multivariate distributions using the nodes and edges of a related graph. The nodes correspond to variables and the lack of an edge between two nodes models the conditional independence of the two variables, given the remaining ones. This corresponds to a zero entry in the concentration (inverse covariance) matrix (Lauritzen, 1996), and covariance selection involves identifying these conditional independencies.

Let y_1, \dots, y_n be a random sample from the q -variate normal distribution $N_q(\mu, \Omega^{-1})$, where the mean is $\mu \in \mathbb{R}^q$ and the $q \times q$ concentration matrix Ω is positive definite. The log-likelihood for $\theta = (\mu, \Omega)$ is

$$\ell(\theta; y) = \frac{n}{2} \log |\Omega| - \frac{1}{2} \text{tr}(\Omega y^\top y) + 1_n^\top y \Omega \mu - \frac{n}{2} \mu^\top \Omega \mu, \quad (9)$$

where y denotes the $n \times q$ matrix with l th row vector y_l^\top and 1_n is a $n \times 1$ vector of ones. The components of the ML estimate $\hat{\theta}$ are

$$\hat{\mu} = y^\top 1_n / n, \quad \hat{\Omega} = (y^\top y / n - y^\top 1_n 1_n^\top y / n^2)^{-1}.$$

For covariance selection the mean parameter is not of direct interest, so we focus instead on the marginal distribution of the ML estimator for the covariance matrix $\hat{\Omega}^{-1} \sim W_q(n-1, \Omega^{-1}/n)$, where W_q denotes the Wishart random variable of order q . The marginal log-likelihood function for Ω

$$\ell(\Omega; y) = \frac{n-1}{2} \log |\Omega| - \frac{n}{2} \text{tr}(\Omega \hat{\Omega}^{-1}),$$

sometimes referred to as restricted log-likelihood or REML, can then be used to carry out inference just on the concentration matrix. The directional p -value for testing constraints on Ω in Section 3.3 is equal to that obtained from the full log-likelihood function (9). It is also convenient to exploit the symmetry of the concentration matrix and express the restricted log-likelihood as

$$\ell(\omega; u) = \frac{n-1}{2} \log |\Omega| - \frac{n-1}{2} \omega^\top J u, \quad (10)$$

where $\omega = \text{vech } \Omega$, $u = n/(n-1) \text{vech } \hat{\Omega}^{-1}$ and the matrix $J = G^\top G$ is diagonal with elements equal to either 1 or 2. If A is a $q \times q$ symmetric matrix, $\text{vec } A$ is the $q^2 \times 1$ vector which stacks the columns of A on top of one another, while $\text{vech } A$ retains only the $q^* = q(q+1)/2$ entries in the lower triangle of A . The two vectors are linked by the relationship $\text{vec } A = G \text{vech } A$, which also gives the $q^2 \times q^*$ duplication matrix G (see, e.g., Abadir and Magnus, 2005, Section 11.3).

In the saturated case addressed by Davison et al. (2014), i.e. the case of a complete graph where Ω has no particular a priori structure, the condition $n > q$ is required for the existence of $\hat{\Omega}$ (Lauritzen, 1996, Theorem 5.1). On the other hand, if the graph is incomplete with some zero off-diagonal entries in Ω , the ML estimate exists if n is larger than the maximal clique size of the hypothesized graph or its decomposable version (Buhl, 1993; Lauritzen, 1996, Section 5.3.2). In what follows, we focus on comparing nested unsaturated models corresponding to nested incomplete graphs. Therefore we allow the sample size n to be smaller than the number of nodes q , but large enough for the ML estimate of the concentration matrix to exist under the alternative model under study (cf. Section 3.2).

3.2 Likelihood quantities for unsaturated models

Suppose some off-diagonal elements $\Omega_{ij}, 1 \leq i < j \leq q$, in the concentration matrix are known to be zero, meaning that the underlying graph is known to be incomplete. As in Roverato and Whittaker (1996), we can rearrange the elements of ω, u and the leading diagonal of J to simplify the calculations. Specifically, defining the edge sets

$$\begin{aligned} k &= \{(i, j) : \Omega_{ij} \neq 0, i \leq j\}, \\ h &= \{(i, j) : \Omega_{ij} = 0, i < j\}, \end{aligned} \quad (11)$$

and giving any ordering to k and h such that

$$\begin{aligned} k &= \{k_1, k_2, \dots, k_p\}, \\ h &= \{h_1, h_2, \dots, h_w\}, \end{aligned}$$

it is possible to define

$$\omega = \begin{pmatrix} \omega_k \\ \omega_h \end{pmatrix}, \quad u = \begin{pmatrix} u_k \\ u_h \end{pmatrix}, \quad J = \begin{pmatrix} J_{kk} & 0 \\ 0 & J_{hh} \end{pmatrix}.$$

Since in unsaturated models $\omega_h = 0$, we can write $\Omega = \Omega_k = \Omega(\omega_k)$ so that the log-likelihood (10) becomes

$$\ell(\omega_k; u_k) = \frac{n-1}{2} \log |\Omega_k| - \frac{n-1}{2} \omega_k^\top J_{kk} u_k, \quad (12)$$

which is a function of the p -dimensional canonical parameter $\varphi = \omega_k$ only, with $p > q$. Differentiation of (12) with respect to ω_k leads to the score function

$$\ell_{\omega_k}(\omega_k) = \frac{n-1}{2} J_{kk}(\sigma_k - u_k),$$

where σ_k is the partition of $\sigma = \text{vech } \Omega_k^{-1}$ obtained according to (11). Solving the score equation leads to $\hat{\sigma}_k = u_k$ and to the corresponding ML estimate $\hat{\omega}_k$, usually derived numerically (see Section 5.3 of Davison et al., 2014).

As the observed and expected information matrices are equal in canonical exponential families, from the results in Roverato and Whittaker (1996, Section 3) follows that

$$J_{\omega_k \omega_k}(\omega_k) = \frac{n-1}{4} J_{kk} \text{Iss}(\Omega_k^{-1})_{kk} J_{kk}, \quad (13)$$

where $\text{Iss}(\Omega_k^{-1})_{kk}$ is a $p \times p$ partition of the Isserlis matrix of the covariance matrix $\Sigma = \Omega_k^{-1}$ (Isserlis, 1918). The entries of $\text{Iss}(\Sigma)_{kk}$ are

$$\text{Cov}(u_{ij}, u_{rs}) = \Sigma_{ir} \Sigma_{js} + \Sigma_{is} \Sigma_{jr},$$

with $(i, j), (r, s) \in k$.

3.3 Comparison of nested unsaturated models

Consider now the partition $\omega_k = (\psi, \lambda)$ of the canonical parameter, where ψ is the component of interest having dimension $d \leq p - q$. The null hypothesis $H_0 : \psi = 0$ tests whether d additional off-diagonal entries $\Omega_{ij}, i < j$, are zero. Hence, the reduced null model is nested in the alternative unsaturated model of Section 3.2. Starting from (12), the log-likelihood ratio statistic for testing H_0 is

$$w(\psi_0) = -(n-1) \log |\hat{\Omega}_k^{-1} \hat{\Omega}_0|, \quad (14)$$

where $\hat{\Omega}_k = \Omega(\hat{\omega}_k)$ is the ML estimate of Ω obtained from (12), and $\hat{\Omega}_0 = \Omega(\hat{\omega}_{k0})$ is its constrained ML estimate under H_0 , with $\hat{\omega}_{k0} = (0, \hat{\lambda}_0)$. The null asymptotic distribution of $w(\psi_0)$ is χ_d^2 .

For the directional p -value that discriminates between two nested Gaussian graphical models, as specified in (6) we first find the expected value of s under H_0

$$s_\psi = -\ell_{\omega_k}(\hat{\omega}_{k0}) = \frac{n-1}{2} J_{kk}(u_k - \hat{\sigma}_{k0}),$$

where $\hat{\sigma}_{k0} = \text{vech } \hat{\Omega}_0^{-1}$. Then, the log-likelihood function (3) along the line $s(t) = (1-t)s_\psi$ follows from (12):

$$\ell\{\omega_k; s(t)\} = \frac{n-1}{2} \log |\Omega_k| - \frac{n-1}{2} \omega_k^\top J_{kk} \{\hat{\sigma}_{k0} + t(u_k - \hat{\sigma}_{k0})\}. \quad (15)$$

The maximization of (15) entails that $\hat{\sigma}_k\{s(t)\} = \hat{\sigma}_k(t) = \hat{\sigma}_{k0} + t(u_k - \hat{\sigma}_{k0})$ or, equivalently,

$$\hat{\Omega}_k^{-1}\{s(t)\} = \hat{\Omega}_k^{-1}(t) = t\hat{\Omega}_k^{-1} + (1-t)\hat{\Omega}_0^{-1}. \quad (16)$$

Given that $\hat{\Omega}_k(t) = \Omega\{\hat{\omega}_k(t)\}$, by taking the inverse of the matrix resulting in the left-hand side of (16) the value of $\hat{\omega}_k(t)$ is obtained accordingly. The replacement of ω_k in (15) with $\hat{\omega}_k(t)$ and $\hat{\omega}_{k0}$, respectively, delivers the result

$$\begin{aligned} \exp[\ell\{\hat{\omega}_{k0}; s(t)\} - \ell\{\hat{\omega}_k(t); s(t)\}] &\propto |\hat{\Omega}_k(t)|^{-\frac{n-1}{2}} \exp\left[\frac{n-1}{2}\{\hat{\omega}_k(t) - \hat{\omega}_{k0}\}^\top J_{kk}\hat{\sigma}_k(t)\right] \\ &\propto |\hat{\Omega}_k(t)|^{-\frac{n-1}{2}}, \end{aligned}$$

since the function $\{\hat{\omega}_k(t) - \hat{\omega}_{k0}\}^\top J_{kk}\hat{\sigma}_k(t)$ is constant with respect to t (see proof in the Appendix). By (13), we obtain $|J_{\omega_k\omega_k}(\omega_k)| \propto |\text{Iss}(\Omega_k^{-1})_{kk}|$ and consequently

$$|J_{\omega_k\omega_k}\{\hat{\omega}_k(t); s(t)\}|^{-1/2} \propto |\text{Iss}\{\hat{\Omega}_k^{-1}(t)\}_{kk}|^{-1/2}.$$

Thus, following expression (5), the directional test is based on $p(\psi_0)$ in (7) with

$$h\{s(t); \psi_0\} \propto |\hat{\Omega}_k^{-1}(t)|^{\frac{n-1}{2}} |\text{Iss}\{\hat{\Omega}_k^{-1}(t)\}_{kk}|^{-1/2}, \quad (17)$$

and the analytical value of t_{sup} calculated as detailed in Section 4.2. If the alternative model were saturated, with q^* -vector $\omega_k = \omega$, then

$$|\text{Iss}\{\hat{\Omega}_k^{-1}(t)\}_{kk}| = |\text{Iss}\{\hat{\Omega}_k^{-1}(t)\}| = 2^q |\hat{\Omega}_k^{-1}(t)|^{q+1},$$

according to the general expression for computing the determinant of the Isserlis matrix (Roverato and Whittaker, 1998, Section 2). Consequently, (17) would reduce to

$$h\{s(t); \psi_0\} \propto |\hat{\Omega}_k^{-1}(t)|^{\frac{n-1}{2}} |\hat{\Omega}_k^{-1}(t)|^{-\frac{q+1}{2}} = |\hat{\Omega}_k^{-1}(t)|^{(n-q-2)/2},$$

which agrees with the simpler result obtained by Davison et al. (2014, Section 5.3) for testing the absence of some connections in the complete graph.

The saddlepoint approximation for the minimal sufficient statistic of a random sample from a multivariate normal variable is exact when $n > q$ (Barndorff-Nielsen and Cox, 1979, p. 287). Since $s(t)$ is a component of this sufficient statistic, expression (17) gives its unnormalized exact marginal distribution in \mathcal{L}^* . The normalizing constant simplifies in the ratio of integrals in (7), thus the approximation error involved in the calculation of the directional p -value stems only from the one-dimensional numerical integration. Simulations in Section 5 confirm this result, and show that the directional p -value is extremely accurate even in settings with $n < q$.

Finally, the term $\gamma(\psi)$ in Skovgaard's (2001) modified likelihood ratio statistic in (8) is:

$$\gamma(\psi_0) = \frac{2\{(\hat{\sigma}_{k0} - \hat{\sigma}_k)^\top \text{Iss}(\hat{\Omega}_0^{-1})_{kk}^{-1}(\hat{\sigma}_{k0} - \hat{\sigma}_k)\}^{d/2}}{\{-\log |\hat{\Omega}_k^{-1}\hat{\Omega}_0|\}^{d/2-1}(\hat{\omega}_k - \hat{\omega}_{k0})^\top J_{kk}(\hat{\sigma}_{k0} - \hat{\sigma}_k)} \left\{ \frac{|\text{Iss}(\hat{\Omega}_0^{-1})_{kk}|}{|\text{Iss}(\hat{\Omega}_k^{-1})_{kk}|} \right\}^{1/2}. \quad (18)$$

4 Computational aspects

4.1 Calculation of the determinant of the Isserlis matrix

In situations where the dimension p of the canonical parameter ω_k under the alternative model is smaller than q^* but still considerably large, the calculation of the determinant of the matrix $\text{Iss}\{\hat{\Omega}_k^{-1}(t)\}_{kk}$ in (17) can be particularly involving. It is then advisable to exploit some useful results on the Isserlis matrix in order to speed up the computing time for the directional p -value.

Let A be a $q \times q$ symmetric invertible matrix. Roverato and Whittaker (1998, (15)), for any partition (k', k'') of the edge set k in (11) such that $k' \cup k'' = k$ and $k' \cap k'' = \bar{k}$, show that

$$|\text{Iss}(A)_{kk}| = \frac{|\text{Iss}(A)_{k'k'}| |\text{Iss}(A)_{k''k''}|}{|\text{Iss}(A)_{\bar{k}\bar{k}}|},$$

which gives a convenient way to reduce the dimensions of the matrices. If, moreover, the graph induced by k is chordal with vertex set decomposable into cliques C_1, \dots, C_K and separators S_2, \dots, S_K according to definitions in Lauritzen (1996, Section 2.1), this can be further simplified to

$$|\text{Iss}(A)_{kk}| = 2^q \frac{\prod_{i=1}^K |A_{C_i C_i}|^{n_{C_i}+1}}{\prod_{i=2}^K |A_{S_i S_i}|^{n_{S_i}+1}}, \quad (19)$$

where n_{C_i} and n_{S_i} denote the number of nodes in the i th clique and i th separator, respectively, while $A_{C_i C_i}$ and $A_{S_i S_i}$ are submatrices of A with rows and columns corresponding to the relative nodes (Roverato and Whittaker, 1998, (17)).

4.2 Determination of t_{sup} and numerical integration

In this framework, the upper bound t_{sup} of the integrals in (7) is the largest value of t such that the ML estimate $\hat{\Omega}_k(t)$ under the alternative hypothesis is positive definite. Specifically, as shown by Huang et al. (2021, Section 4.1, Lemma 1) for the case of saturated alternative model, the closed-form expression equals $t_{\text{sup}} = 1/(1 - \nu_{(1)})$, where $\nu_{(1)}$ here indicates the smallest among the eigenvalues ν_i , $i = 1, \dots, q$, of the matrix $\hat{\Omega}_0 \hat{\Omega}_k^{-1}$.

Let us express the integrand function appearing in (7) as $\exp\{\bar{g}(t; \psi)\}$, where $\bar{g}(t; \psi) = (d-1) \log t + \log h\{s(t; \psi)\}$. Following Huang et al. (2021, Section 4.2), in order to improve numerical stability it is convenient to rescale $\bar{g}(t; \psi)$ and calculate the directional p -value via the equivalent formula

$$p(\psi) = \frac{\int_1^{t_{\text{sup}}} \exp\{\bar{g}(t; \psi) - \bar{g}(\hat{t}; \psi)\} dt}{\int_0^{t_{\text{sup}}} \exp\{\bar{g}(t; \psi) - \bar{g}(\hat{t}; \psi)\} dt}, \quad \text{where } \hat{t} = \arg \sup_{t \in [0, t_{\text{sup}}]} \bar{g}(t; \psi).$$

In addition, the Gauss–Hermite quadrature (Liu and Pierce, 1994) can be used to perform the numerical integration only in the relevant region $[t_{\min}, t_{\max}] \subseteq [0, t_{\text{sup}}]$ where the integrand function concentrates. An accurate approximation to the directional p -value (7) which can further reduce the computing time is then

$$p(\psi) \doteq \frac{\int_1^{t_{\max}} \exp\{\bar{g}(t; \psi) - \bar{g}(\hat{t}; \psi)\} dt}{\int_{t_{\min}}^{t_{\max}} \exp\{\bar{g}(t; \psi) - \bar{g}(\hat{t}; \psi)\} dt}. \quad (20)$$

Define $\bar{g}_1(t; \psi) = (d-1) \log t - \{(n-1)/2\} \log |\hat{\Omega}_k(t)|$ and $\bar{g}_2(t; \psi) = -(1/2) \log |\text{Iss}\{\hat{\Omega}_k^{-1}(t)\}_{kk}|$, such that $\bar{g}(t; \psi) = \bar{g}_1(t; \psi) + \bar{g}_2(t; \psi)$. The extremes of integration are typically computed as $t_{\min} = \max\{0, \hat{t} - c/q(\hat{t}; \psi)\}$ and $t_{\max} = \min\{\hat{t} + c/q(\hat{t}; \psi), t_{\text{sup}}\}$, where $c = 5$ and $q(t; \psi) = -\partial^2 \bar{g}(t; \psi) / \partial t^2$ (cf. Huang et al., 2021, Appendix B.3). Unfortunately here the second derivative of $\bar{g}_2(t; \psi)$ cannot be derived explicitly, although empirical results show it is positive and increasing with the ratio q/n . Since its numerical approximation seems quite unstable, we set

$$q(t; \psi) = -\frac{\partial^2 \bar{g}_1(t; \psi)}{\partial t^2} = \frac{d-1}{t^2} + \frac{n-1}{2} \sum_{i=1}^q \frac{(1-\nu_i)^2}{(1-t+\nu_i)^2},$$

and $c > 5$ according to the relative magnitude of n and q . Specifically, in our simulations of next section the value of c was chosen in each setting in order to obtain the same results as if

the integration were carried out on the entire interval $[0, t_{\text{sup}}]$. The empirical findings indicate that the integrand function is more concentrated when q is not too large with respect to n . The Gauss–Hermite quadrature is indeed useful in the least challenging settings and seems to give no practical advantage when $n < q$.

Finally, note that in case the function $\bar{g}(t; \psi)$ is monotonic in $[0, t_{\text{sup}}]$ the above strategy cannot be applied, so the directional p -value has to be calculated directly via formula (7). However, we observed this issue only 21 times in the Monte Carlo experiments below.

5 Simulation studies

The performance of the directional approach in terms of covariance selection for Gaussian graphical models is examined here through simulation-based experiments. The focus of the first scenario is put on a small graph with $q = 4$ nodes, as in Eriksen (1996, Section 4). Figure 1 displays the two models under comparison which differ only by $d = 2$ edges. Monte Carlo simulations are based on 100 000 samples of size $n = 7$ generated under the null hypothesis. The empirical p -value distribution of the tests based on $w(\psi_0)$, $w^*(\psi_0)$, $w^{**}(\psi_0)$ and the directional procedure is shown in the left plot of Figure 2 with respect to the reference uniform distribution, zooming on the interval $(0, 0.1)$. The right plot compares the relative errors of the three most accurate methods. Despite the simplicity of the example, the likelihood ratio statistic leads to too many rejections of the null hypothesis, as n is relatively small. The higher-order modifications remedy this, yet the superiority of the directional approach is more evident in terms of relative error.




Figure 1: Graphs for the first simulation scenario where the dimension of the parameter of interest equals $d = 2$. The alternative model on the left is compared against the null model on the right.

The inferential benefits of our approach over the omnibus likelihood-based competitors are particularly appreciated with higher magnitudes of q and d . In the second main scenario we consider the data of Kenward (1987, Table 1) from a study on intestinal parasites of 60 calves, where the weight in kg of each bovine was recorded on 11 occasions during the grazing season. To enable comparison with Davison et al. (2014, Section 5.3), who could only test the saturated model, we draw 100 000 samples of size $n = 60$ from a q -variate Gaussian random variable under the hypothesis of first-order Markovian dependence MD(1) with tridiagonal concentration matrix. For each $q \in \{11, 30, 50\}$, this null model is compared against four different alternative unsaturated structures, using also $w(\psi_0)$, $w^*(\psi_0)$ and $w^{**}(\psi_0)$. These Markovian models MD(m) with $1 < m < q - 1$ correspond to so-called band concentration matrices, whose nonzero entries are confined to m diagonals on either side of the main one. The orders m are chosen to check the




Figure 2: Results based on 100 000 samples simulated under the null model represented by the right graph of Figure 1 with $n = 7$ and $q = 4$. On the left, the empirical p -values lower than 0.1 are compared with the uniform distribution given by the gray diagonal for w (dot-dashed orange line), w^* (dashed green line), w^{**} (long-dashed dark green line) and the directional test (solid blue line). On the right, the corresponding relative errors are plotted in a similar fashion only for w^* , w^{**} and the directional method.

behavior of the various methods for a wide range of dimensions d of the parameter of interest, and consequently of the nuisance component. The simplification (19) is particularly useful for computing the directional p -values with such a high-dimensional parameter of interest.

Table 1 reports experimental results obtained when $q = 11$ as in the original dataset, whereas Tables 2 and 3 refers to cases with data simulated using a larger covariance matrix, $q = 30$ and $q = 50$ respectively. In line with the findings in Davison et al. (2014, Table 6), the empirical distribution of the directional p -values is extremely accurate in all settings, and almost unaffected by the size of q and d . The usual likelihood ratio statistic $w(\psi_0)$ is very sensitive to the dimension of both ψ and λ ; its adjustments $w^*(\psi_0)$ and, particularly, $w^{**}(\psi_0)$ seem to suffer from the increasing dimension d of the parameter of interest. Tables 2 and 3 clearly indicate that, as d grows, the test based on $w(\psi_0)$ becomes too liberal and those based on $w^*(\psi_0)$ and $w^{**}(\psi_0)$ too conservative. For the intermediate case $q = 30$, Figure 3 contrasts the null distribution of the directional p -values with those from $w(\psi_0)$, $w^*(\psi_0)$ and $w^{**}(\psi_0)$.

Before proceeding, let us focus on the implementation of formula (19) to obtain the determinant of the Isserlis matrix of Ω estimated under the alternative hypothesis. When multiplying the determinants of many square matrices of moderate order, some propagation of numerical errors can occur. In our experiments this is visible, to a certain extent, in the intermediate sections of Tables 2 and 3, when the performance of directional tests seems slightly less excellent than in the remaining sections. Indeed, when the null is tested against more extreme Markovian models, the matrices involved in (19) are either many but small (top section) or large but few (bottom section), thus the final product of their determinants is not overly affected by numerical error. That being said, it is important to point out that in all settings the directional approach remains remarkably accurate and brings a great improvement over the competing testing procedures.

The third simulation scenario considers a block diagonal configuration of the concentration matrix under the null hypothesis. Here, 100 000 samples of size $n \in \{40, 60, 90, 120\}$ were drawn from a normal distribution with $q = 50$ components and covariance matrix $\Sigma_0 = \text{diag}\{\Sigma_{01}, \Sigma_{01}\}$, with Σ_{01} sub-matrix 25×25 having diagonal entries equal to 1 and off-diagonal entries equal

Table 1: Empirical p -value distributions (%) based on 100 000 replications. The Markovian model MD(1) is tested against different models MD(m) of orders $m \in \{2, 3, 6, 9\}$, when $n = 60$ observations of a graph with $q = 11$ nodes are available.

Nominal (%)	1.0	2.5	5.0	10.0	25.0	50.0	75.0	90.0	95.0	97.5	99.0
vs MD(2), $d = 9$											
Likelihood ratio, (14)	1.4	3.3	6.3	12.0	28.4	53.7	77.5	91.2	95.6	97.8	99.1
Skovgaard's w^* , (18)	1.0	2.5	5.1	10.0	25.1	50.2	75.1	89.9	94.9	97.4	99.0
Skovgaard's w^{**} , (18)	1.0	2.5	5.1	10.0	25.1	50.2	75.1	89.9	94.9	97.4	99.0
Directional, (20)	1.0	2.5	5.1	10.0	25.2	50.3	75.2	90.1	95.0	97.5	99.0
vs MD(3), $d = 17$											
Likelihood ratio, (14)	1.8	3.9	7.2	13.5	30.4	56.1	79.3	92.0	96.0	98.1	99.2
Skovgaard's w^* , (18)	1.1	2.6	5.0	10.0	24.6	49.6	74.6	89.6	94.7	97.3	98.9
Skovgaard's w^{**} , (18)	1.0	2.5	5.0	9.9	24.5	49.5	74.5	89.5	94.7	97.2	98.9
Directional, (20)	1.0	2.6	5.1	10.1	25.0	50.3	75.4	90.2	95.0	97.5	99.0
vs MD(6), $d = 35$											
Likelihood ratio, (14)	2.5	5.5	9.8	17.4	36.2	62.2	83.3	94.0	97.2	98.6	99.5
Skovgaard's w^* , (18)	0.8	2.1	4.3	8.8	22.4	46.4	71.7	87.8	93.6	96.6	98.5
Skovgaard's w^{**} , (18)	0.8	2.1	4.2	8.6	22.0	45.9	71.2	87.5	93.4	96.4	98.5
Directional, (20)	1.0	2.5	4.9	10.0	25.0	50.3	75.3	90.2	95.1	97.5	99.0
vs MD(9), $d = 44$											
Likelihood ratio, (14)	3.3	6.9	12.0	20.6	40.8	66.2	85.9	95.2	97.8	99.0	99.6
Skovgaard's w^* , (18)	0.7	1.8	3.7	7.8	20.7	43.7	69.1	86.3	92.6	96.1	98.2
Skovgaard's w^{**} , (18)	0.7	1.8	3.6	7.5	20.1	42.8	68.2	85.7	92.2	95.8	98.1
Directional, (20)	1.0	2.4	4.9	9.9	25.2	50.0	75.0	90.1	95.1	97.5	99.0
Standard error	0.0	0.0	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.0	0.0

to 0.5. Such condition clearly implies that $\Omega_0 = \Sigma_0^{-1}$ is also block diagonal, so that the first 25 nodes are conditionally (as well as unconditionally) independent of the last 25 nodes in the graph. On the other hand, our alternative model admits the existence of some conditional dependence between the two subsets of nodes. Specifically, besides the nonzero elements defined in Ω_0 , we also suppose $\Omega_{ij} = \Omega_{ji} \neq 0$ for $i = 16, \dots, 25$ and $j = 26, \dots, 50$. It follows that the dimension of the parameter of interest is $d = 250$ and (19) can be used to speed up calculations of the Isserlis matrix associated with the chordal alternative unsaturated graph.

Simulation results in this framework are presented in Table 4. Given the notable size of d , the relative performance of the approximations under comparison, in terms of the empirical p -value distribution, is analogous to that in the previous experiment, with the only exception that here the version $w^{**}(\psi_0)$ appears generally more reliable than $w^*(\psi_0)$. Although the increase in sample size generates some accuracy improvements for all the competitors as expected, the directional procedure guarantees an almost perfect agreement with the nominal uniform distribution for all values of n considered. The extreme liberality of the standard likelihood ratio test persists, Skovgaard's $w^*(\psi_0)$ does not correct it enough and the version $w^{**}(\psi_0)$ overcorrects it. As in the first simulation scenario, Figure 4 displays the p -values obtained via the likelihood ratio statistic, its modified versions and the directional procedure.

6 Applications

First, we examine the dataset already introduced in the second simulation scenario of Section 5 from the experiment about the control of intestinal parasites in cattle. However, here we focus on the two treatment groups with equal size $n = 30$ separately, in order to investigate differences

Table 2: Empirical p -value distributions (%) based on 100 000 replications. The Markovian model MD(1) is tested against different models MD(m) of orders $m \in \{2, 9, 18, 28\}$, when $n = 60$ observations of a graph with $q = 30$ nodes are available.

Nominal (%)	1.0	2.5	5.0	10.0	25.0	50.0	75.0	90.0	95.0	97.5	99.0
vs MD(2), $d = 28$											
Likelihood ratio, (14)	1.6	3.8	7.2	13.4	30.5	56.4	79.4	92.2	96.2	98.1	99.3
Skovgaard's w^* , (18)	1.0	2.5	5.0	10.0	24.9	50.0	75.1	90.1	95.0	97.5	99.0
Skovgaard's w^{**} , (18)	1.0	2.5	5.0	10.0	24.9	50.0	75.0	90.0	95.0	97.5	99.0
Directional, (20)	1.0	2.4	4.9	10.0	24.9	50.1	75.2	90.2	95.1	97.5	99.0
vs MD(9), $d = 196$											
Likelihood ratio, (14)	11.1	19.1	28.4	41.5	64.6	84.6	95.3	98.7	99.5	99.8	99.9
Skovgaard's w^* , (18)	0.3	0.9	2.0	4.4	13.3	32.3	57.9	78.5	87.1	92.5	96.4
Skovgaard's w^{**} , (18)	0.3	0.8	1.7	3.9	12.1	30.2	55.4	76.5	85.7	91.4	95.7
Directional, (20)	0.9	2.3	4.8	9.7	24.7	50.3	75.8	90.5	95.4	97.7	99.1
vs MD(18), $d = 340$											
Likelihood ratio, (14)	53.8	66.9	76.9	86.0	95.0	98.8	99.8	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	0.0	0.1	0.3	0.7	3.0	10.7	27.4	48.8	62.1	72.8	82.9
Skovgaard's w^{**} , (18)	0.0	0.0	0.1	0.4	1.7	6.9	19.5	38.2	51.2	62.5	74.3
Directional, (20)	0.8	2.2	4.6	9.5	24.7	50.2	76.0	90.8	95.6	97.8	99.2
vs MD(28), $d = 405$											
Likelihood ratio, (14)	86.2	92.3	95.6	97.9	99.5	99.9	100.0	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	0.0	0.0	0.0	0.2	0.9	4.3	13.8	30.0	42.5	53.9	67.0
Skovgaard's w^{**} , (18)	0.0	0.0	0.0	0.0	0.2	1.4	5.9	15.5	24.5	33.9	46.4
Directional, (20)	1.0	2.4	5.1	10.1	25.2	50.1	75.1	90.1	95.1	97.5	99.0
Standard error	0.0	0.0	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.0	0.0

in the underlying temporal dynamics of growth. Recalling that each animal was weighed $q = 11$ consecutive times, we start by assuming a Markovian dependence of order $m = 3$, the simplest model accepted in a test against the saturated one by all the procedures under analysis and in both groups. This model is then compared against the null hypothesis of first-order dependence, implying $d = 17$. For the calves randomly assigned to the first treatment, the likelihood ratio statistic is $w(\psi_0) = 28.384$ with p -value= 0.041, Skovgaard's modifications are $w^*(\psi_0) = 22.977$ with p -value= 0.150 and $w^{**}(\psi_0) = 22.691$ with p -value= 0.160 and the directional p -value is 0.111. For the second group we get instead $w(\psi_0) = 31.895$ with p -value= 0.016, $w^*(\psi_0) = 30.055$ with p -value= 0.026, $w^{**}(\psi_0) = 30.028$ with p -value= 0.026 and directional p -value= 0.029. The standard likelihood ratio test is the only one to reject the MD(1) model at a 5% significance level for both treatments. Conversely, the other statistics recognize a different time pattern and indicate a more complex dependence of the weights in the second group.

We now consider some microarray data from the biostatistical literature (see, e.g., Massa et al., 2010), which characterize gene expression signatures in acute lymphocytic leukemia cells associated with genotypic abnormalities in adult patients. The normalized version of such data, available in the package topologyGSA (Massa and Sales, 2016) of the R software (R Core Team, 2020), is especially useful for analyzing the B-cell receptor (BCR) signaling pathway composed by $q = 35$ gene products. The observed samples are classified according to the presence of molecular rearrangements in their genetic profile.

The conversion of biological pathways into graphical models has become standard practice in biostatistics to separate and compare specific portions of the genetic process under examination. Based on findings in Massa et al. (2010), it seems of interest to investigate whether the graph resulting from the well-known BCR signaling pathway in Figure 5 can be further simplified. In

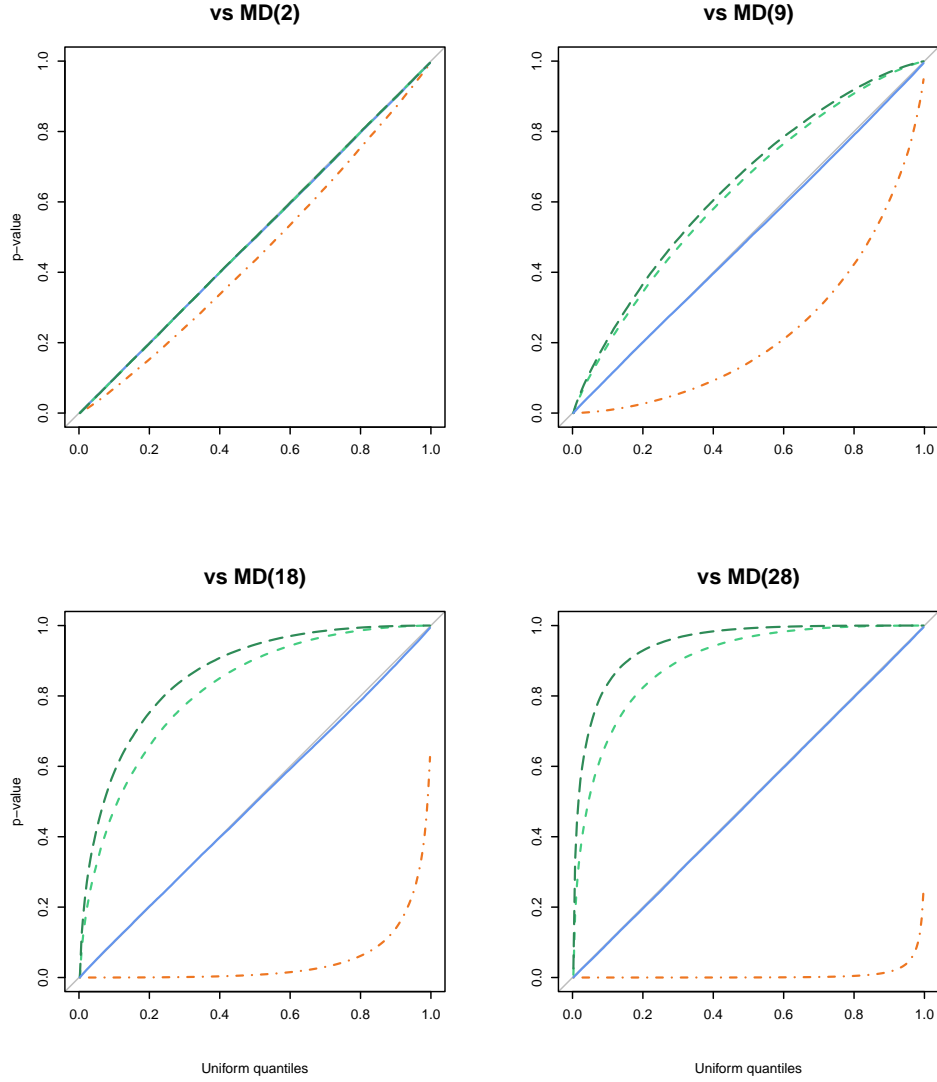


Figure 3: Results based on 100 000 samples simulated under the null model MD(1) with $n = 60$ and $q = 30$. The empirical p -values obtained via w (dot-dashed orange line), w^* (dashed green line), w^{**} (long-dashed dark green line) and the directional test (solid blue line) are compared with the uniform distribution given by the gray diagonal for several alternative hypotheses and related dimensions of ψ : MD(2) and $d = 28$ (top left), MD(9) and $d = 196$ (top right), MD(18) and $d = 340$ (bottom left), MD(28) and $d = 405$ (bottom right).

more detail, the restricted graphical model under the null hypothesis in our analysis corresponds to the identified path starting from nodes CD22 and CD72 and ending at AP1, going through RasGRP3, Ras, Raf, MEK1/2 and ERK enzymes. Such a comparison implies testing the lack of $d = 12$ edges and can be carried out on the subset of patients not suffering from so-called BCR/ABL rearrangements. With $n = 41$, we obtain $w(\psi_0) = 33.520$ with $p\text{-value} = 8.028 \times 10^{-4}$, $w^*(\psi_0) = 32.172$ with $p\text{-value} = 13.018 \times 10^{-4}$, $w^{**}(\psi_0) = 32.158$ with $p\text{-value} = 13.083 \times 10^{-4}$ and directional $p\text{-value} = 13.941 \times 10^{-4}$. Given the modest size of the parameter of interest, the results are consistent with the simulation outputs. Although the shorter biological path is not preferred over the conventional signaling pathway by any method, the evidence against

Table 3: Empirical p -value distributions (%) based on 100 000 replications. The Markovian model MD(1) is tested against different models MD(m) of orders $m \in \{2, 16, 32, 48\}$, when $n = 60$ observations of a graph with $q = 50$ nodes are available.

Nominal (%)	1.0	2.5	5.0	10.0	25.0	50.0	75.0	90.0	95.0	97.5	99.0
vs MD(2), $d = 48$											
Likelihood ratio, (14)	1.8	4.2	7.8	14.5	32.4	58.2	80.9	93.0	96.7	98.4	99.4
Skovgaard's w^* , (18)	1.0	2.5	5.0	9.9	25.1	50.1	74.9	90.0	95.1	97.5	99.0
Skovgaard's w^{**} , (18)	1.0	2.5	5.0	9.9	25.0	50.0	74.9	89.9	95.1	97.5	99.0
Directional, (20)	1.0	2.5	4.9	9.9	25.0	50.1	75.1	90.1	95.2	97.6	99.1
vs MD(16), $d = 615$											
Likelihood ratio, (14)	77.9	86.7	92.1	96.0	99.0	99.8	100.0	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	0.0	0.0	0.1	0.2	1.1	5.1	15.9	33.3	46.2	57.9	70.6
Skovgaard's w^{**} , (18)	0.0	0.0	0.0	0.1	0.5	2.6	9.3	22.2	33.2	44.2	57.4
Directional, (20)	0.8	2.0	4.3	9.1	24.4	50.4	76.3	91.4	96.1	98.1	99.3
vs MD(32), $d = 1023$											
Likelihood ratio, (14)	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.7	1.7	3.4	6.7
Skovgaard's w^{**} , (18)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2
Directional, (20)	0.5	1.4	3.4	8.0	23.5	51.7	78.6	92.8	96.9	98.7	99.5
vs MD(48), $d = 1175$											
Likelihood ratio, (14)	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.4	1.0	2.0	4.0
Skovgaard's w^{**} , (18)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Directional, (20)	0.8	2.2	4.7	9.8	25.4	51.1	76.2	90.9	95.5	97.8	99.2
Standard error	0.0	0.0	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.0	0.0

the former based on the likelihood ratio test appears too striking. A more reliable conclusion should then be drawn according to the directional p -value, which here agrees very closely with Skovgaard's ones, as expected when $d = 12$.

7 Discussion

A directional likelihood-based method for covariance selection in incomplete undirected graphs has been presented in this paper. The directional test for comparing two nested Gaussian graphical models is based on (17), which is the saddlepoint approximation to the kernel of the conditional univariate density of the sufficient statistic. The exactness of the saddlepoint approximation in this general context, at least when $n > q$, is supported by the empirical exceptional accuracy of the proposed method, even with a high-dimensional parameter of interest and when $n < q$.

Our proposal differs from the omnibus approach adopted when using the likelihood ratio tests in one important respect: the conditioning on the direction of departure of the observed data from the hypothesis under examination. In fact, such conditioning improves the accuracy of the directional procedure by simultaneously adjusting for the estimation of nuisance parameters and transforming the component of interest into a scalar (Fraser et al., 2016). Nonetheless, as attested by the simulation results in Section 5 and previous work, the directional p -value is also very accurate marginally.

The proposed directional approach has been also compared to another higher-order solution, which consists of a modification to the likelihood ratio test. Although this provides improvement over the first order approximation to the likelihood ratio statistic, in extreme settings its

Table 4: Empirical p -value distributions (%) based on 100 000 replications. The two-block diagonal structure of the concentration matrix for a graph with $q = 50$ nodes is tested against a more complex structure including $d = 250$ additional edges.

Nominal (%)	1.0	2.5	5.0	10.0	25.0	50.0	75.0	90.0	95.0	97.5	99.0
$n = 40$											
Likelihood ratio, (14)	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	27.1	39.1	50.6	63.5	81.7	93.7	98.5	99.7	99.9	100.0	100.0
Skovgaard's w^{**} , (18)	0.7	1.7	3.4	7.2	19.3	42.0	68.0	85.9	92.5	96.1	98.3
Directional, (20)	1.0	2.5	5.0	10.1	25.2	50.2	75.2	90.0	94.9	97.4	98.9
$n = 60$											
Likelihood ratio, (14)	98.4	99.3	99.7	99.9	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	2.4	5.3	9.6	17.3	36.6	62.9	84.2	94.6	97.6	98.9	99.6
Skovgaard's w^{**} , (18)	0.6	1.7	3.5	7.5	20.4	43.9	70.3	87.4	93.5	96.7	98.6
Directional, (20)	1.0	2.5	5.0	10.0	25.1	50.1	75.2	90.2	95.1	97.6	99.0
$n = 90$											
Likelihood ratio, (14)	65.9	77.1	85.0	91.5	97.3	99.4	99.9	100.0	100.0	100.0	100.0
Skovgaard's w^* , (18)	1.3	3.2	6.1	12.0	28.5	54.2	78.2	91.7	96.0	98.1	99.2
Skovgaard's w^{**} , (18)	0.8	2.1	4.3	8.9	23.0	47.6	73.2	89.0	94.5	97.2	98.8
Directional, (20)	0.9	2.5	5.0	10.1	25.0	50.1	75.1	90.1	95.1	97.6	99.0
$n = 120$											
Likelihood ratio, (14)	36.6	50.0	61.6	73.6	88.6	96.7	99.3	99.9	100.0	100.0	100.0
Skovgaard's w^* , (18)	1.1	2.9	5.6	11.0	26.8	52.2	76.5	90.9	95.5	97.8	99.1
Skovgaard's w^{**} , (18)	0.9	2.3	4.6	9.4	24.0	48.7	73.9	89.4	94.6	97.3	98.9
Directional, (20)	1.0	2.5	5.0	10.1	25.1	50.1	75.0	90.0	95.0	97.5	99.0
Standard error	0.0	0.0	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.0	0.0

accuracy deteriorates and therefore inference based on such a test cannot be trusted in general. A possible alternative could be the parametric bootstrap, but the computational effort required for this might be prohibitive. The saddlepoint approximation to the conditional density was derived explicitly, and the computation of the directional p -value via one-dimensional numerical integration is made especially fast by the expedients described in Section 4.

The directional approach detailed here could be extended to graphical models for discrete data, like those dealt with in Roverato (2017). However, as discreteness prevents the saddlepoint approximation from being exact even upon normalization, one might reasonably not expect the same accuracy of directional p -values observed in this work, at least in the most challenging testing problems.

The present methodology only applies to situations where the number of observations is such that the ML estimate exists with probability one under the alternative hypothesis. In particular, n must be larger than the maximal clique size of the hypothesized graph or its decomposable version (Buhl, 1993). The development of reliable likelihood-based testing procedures, omnibus or directional, in circumstances with p much larger than n is still an open problem to be addressed in future research.

Acknowledgments

The authors are grateful to Alberto Roverato for useful discussions and suggestions on the R code to compute the Isserlis matrix. They also thank Davide Risso for his help with the genetic application and Caizhu Huang for suggesting improvements in the computational part.

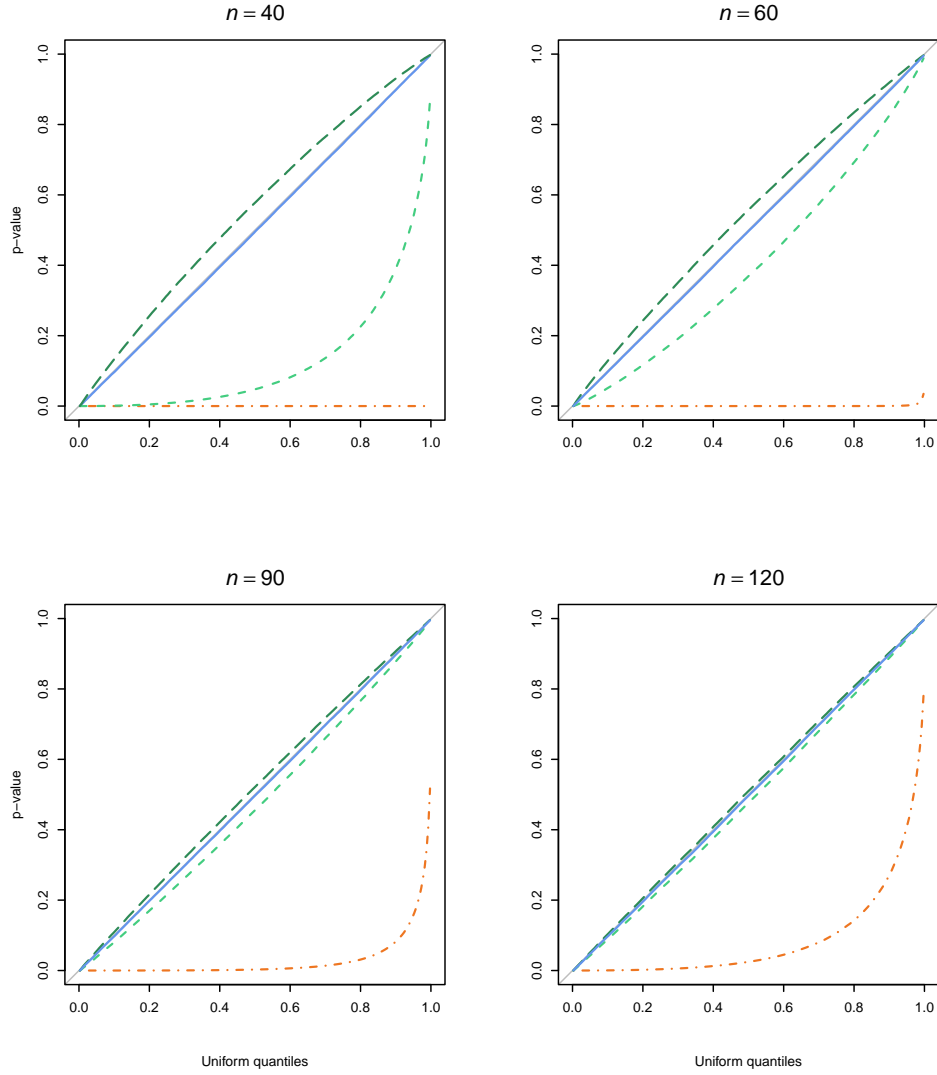


Figure 4: Results based on 100 000 samples simulated under the null model with block diagonal concentration matrix and $q = 50$. The empirical p -values obtained via w (dot-dashed orange line), w^* (dashed green line), w^* (long-dashed dark green line) and the directional test (solid blue line) are compared with the uniform distribution given by the gray diagonal for the same alternative hypothesis implying $d = 250$ and several sample sizes: $n = 40$ (top left), $n = 60$ (top right), $n = 90$ (bottom left), $n = 120$ (bottom right).

Supplementary material

The Supplementary material (available at <https://github.com/cdicaterina/DirInfGGM>) provides the data and the R code to reproduce all numerical results in the paper.

References

Abadir, K. M. and J. R. Magnus (2005). *Matrix Algebra*. Cambridge University Press.

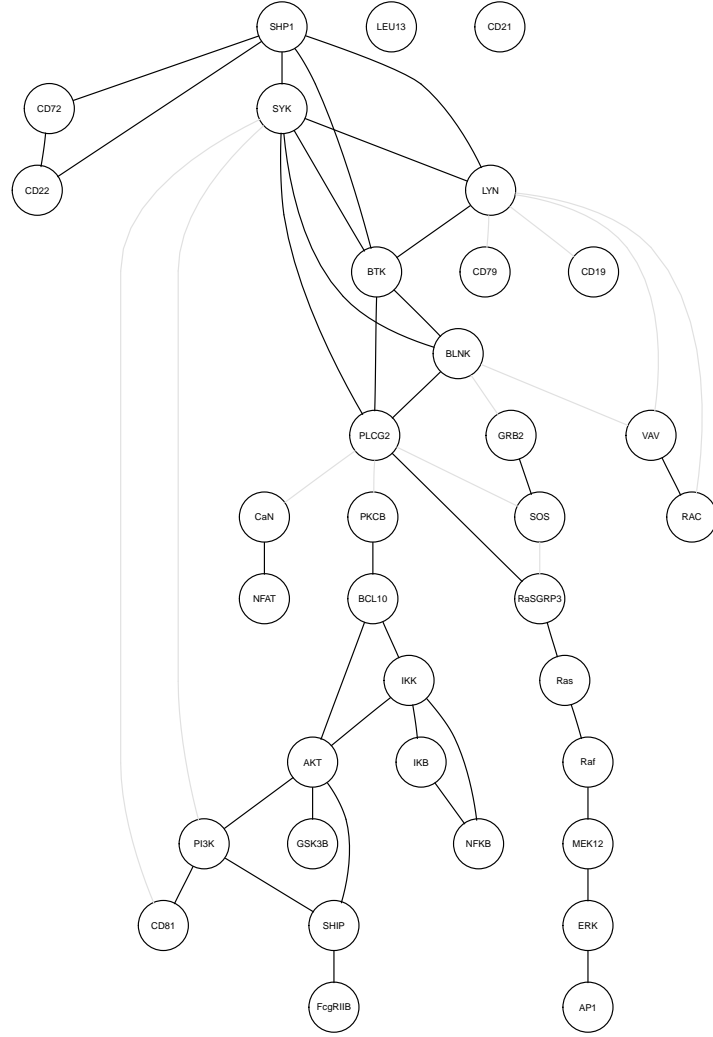


Figure 5: BCR signaling pathway involving $q = 35$ gene products. The interest is on testing whether a simpler path without the $d = 12$ gray edges can be identified.

Barndorff-Nielsen, O. E. (1986). Inference on full or partial parameters based on the standardized signed log likelihood ratio. *Biometrika* 73(2), 307–322.

Barndorff-Nielsen, O. E. and D. R. Cox (1979). Edgeworth and saddle-point approximations with statistical applications (with Discussion). *Journal of the Royal Statistical Society: Series B* 41, 279–312.

Bartlett, M. S. (1937). Properties of sufficiency and statistical tests. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences* 160, 268–282.

Borgelt, C. and R. Kruse (2002). *Graphical Models: Methods for Data Analysis and Mining*. John Wiley & Sons.

Buhl, S. L. (1993). On the existence of maximum likelihood estimators for graphical Gaussian models. *Scandinavian Journal of Statistics* 20, 263–270.

- Cox, D. R. and D. V. Hinkley (1974). *Theoretical Statistics*. London: Chapman & Hall.
- Davison, A. C., D. A. S. Fraser, N. Reid, and N. Sartori (2014). Accurate directional inference for vector parameters in linear exponential families. *Journal of the American Statistical Association* 109, 302–314.
- Eriksen, P. S. (1996). Tests in covariance selection models. *Scandinavian Journal of Statistics* 23, 275–284.
- Fraser, D. A. S. and H. Massam (1985). Conical tests: Observed levels of significance and confidence regions. *Statistische Hefte* 26, 1–17.
- Fraser, D. A. S. and N. Reid (2006). Assessing a vector parameter. *Student* 5, 247–256.
- Fraser, D. A. S., N. Reid, and N. Sartori (2016). Accurate directional inference for vector parameters. *Biometrika* 103, 625–639.
- Fraser, D. A. S., N. Reid, and J. Wu (1999). A simple general formula for tail probabilities for frequentist and Bayesian inference. *Biometrika* 86, 249–264.
- Huang, C., C. Di Caterina, and N. Sartori (2021). Directional testing for high-dimensional multivariate normal distributions. arXiv:2107.09418.
- Isserlis, L. (1918). On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables. *Biometrika* 12, 134–139.
- Kenward, M. G. (1987). A method for comparing profiles of repeated measurements. *Journal of the Royal Statistical Society: Series C* 36, 296–308.
- Lauritzen, S. L. (1996). *Graphical Models*. Oxford: Oxford University Press.
- Liu, Q. and D. A. Pierce (1994). A note on gauss–hermite quadrature. *Biometrika* 81, 624–629.
- Massa, M. S., M. Chiogna, and C. Romualdi (2010). Gene set analysis exploiting the topology of a pathway. *BMC Systems Biology* 4, 121.
- Massa, S. and G. Sales (2016). *topologyGSA: Gene Set Analysis Exploiting Pathway Topology*. R package version 1.4.6.
- McCormack, A., N. Reid, N. Sartori, and S. A. Theivendran (2019). A directional look at F-tests. *The Canadian Journal of Statistics* 47, 619–627.
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Roverato, A. (2017). *Graphical Models for Categorical Data*. Cambridge University Press.
- Roverato, A. and J. Whittaker (1996). Standard errors for the parameters of graphical Gaussian models. *Statistics and Computing* 6, 297–302.
- Roverato, A. and J. Whittaker (1998). The Isserlis matrix and its application to non-decomposable graphical Gaussian models. *Biometrika* 85, 711–725.
- Skovgaard, I. M. (1988). Saddlepoint expansions for directional test probabilities. *Journal of the Royal Statistical Society: Series B* 50, 269–280.
- Skovgaard, I. M. (2001). Likelihood asymptotics. *Scandinavian Journal of Statistics* 28, 3–32.
- Whittaker, J. (2009). *Graphical Models in Applied Multivariate Statistics*. John Wiley & Sons.

Appendix

We show that the scalar function

$$f(t) = \{\hat{\omega}_k(t) - \hat{\omega}_{k0}\}^\top J_{kk} \hat{\sigma}_k(t)$$

does not depend on t . Since $f(t) = \text{tr}\{f(t)\}$ and the two models under comparison are nested, it is equivalent to prove that $\text{tr}\{\{\hat{\omega}(t) - \hat{\omega}_0\}^\top J \hat{\sigma}(t)\}$ is constant in t , where

$$\hat{\omega}(t) = \begin{pmatrix} \hat{\omega}_k(t) \\ 0 \end{pmatrix}, \quad \hat{\omega}_0 = \begin{pmatrix} \hat{\omega}_{k0} \\ 0 \end{pmatrix}, \quad \hat{\sigma}(t) = \begin{pmatrix} \hat{\sigma}_k(t) \\ \hat{\sigma}_h(t) \end{pmatrix}$$

are all vectors of dimension q^* . Letting $\hat{\Omega}_k^{-1}(t) = \Sigma\{\hat{\sigma}(t)\}$, we have:

$$\begin{aligned} \text{tr}\{\{\hat{\omega}(t) - \hat{\omega}_0\}^\top J \hat{\sigma}(t)\} &= \text{tr}[\text{vech}\{\hat{\Omega}_k(t) - \hat{\Omega}_0\}^\top G^\top G \text{vech}\hat{\Omega}_k^{-1}(t)] \\ &= \text{tr}[\{\hat{\Omega}_k(t) - \hat{\Omega}_0\}^\top \hat{\Omega}_k^{-1}(t)] \\ &= \text{tr}[I_q - \hat{\Omega}_0\{t\hat{\Omega}_k^{-1} + (1-t)\hat{\Omega}_0^{-1}\}] \\ &= \text{tr}(I_q) - t\text{tr}(\hat{\Omega}_0\hat{\Omega}_k^{-1}) - (1-t)\text{tr}(I_q) \\ &= q - tq - (1-t)q = 0. \end{aligned}$$

This uses basic matrix algebra (see, for instance, Abadir and Magnus, 2005), and the fact that $\text{tr}(\hat{\Omega}_0\hat{\Omega}_k^{-1}) = \text{tr}(\hat{\Omega}_0\hat{\Omega}_0^{-1}) = \text{tr}(I_q) = q$ because the matrices $\hat{\Omega}_k^{-1}$ and $\hat{\Omega}_0^{-1}$ differ only for the elements that correspond to zeros in the constrained concentration matrix $\hat{\Omega}_0$, as can be checked from the solution of the ML equations.