

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/304350743>

A review of undirected and acyclic directed Gaussian Markov model selection and estimation

Article · June 2016

CITATIONS

0

READS

154

3 authors, including:



Concha Bielza

Universidad Politécnica de Madrid

351 PUBLICATIONS 5,342 CITATIONS

SEE PROFILE



Pedro Larranaga

Universidad Politécnica de Madrid

481 PUBLICATIONS 19,178 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Thermal Camera [View project](#)



Regularized Model learning in EDAs for continuous and multi-objective optimization [View project](#)

A review of undirected and acyclic directed Gaussian Markov model selection and estimation

Irene Córdoba-Sánchez
Concha Bielza
Pedro Larrañaga

Abstract

Markov models lie at the interface between statistical independence in a probability distribution and graph separation properties. We review model selection and estimation in directed and undirected Markov models with Gaussian parametrization, emphasizing the main similarities and differences. These two model types are foundationally similar but not equivalent, as we highlight. We report existing results from a historical perspective, taking into account literature from both the artificial intelligence and statistics research communities, which first developed these models. Finally, we point out the main active research areas and open problems now existing with regard to these traditional, albeit rich, Markov models.

Keywords. Conditional independence, Gaussian Bayesian network, Gaussian graphical model, Gaussian Markov model, model selection, parameter estimation

1 Introduction

Markov models, or probabilistic graphical models, explicitly establish a correspondence between statistical independence in a probability distribution and certain separation criteria holding in a graph. They originated at the interface between statistics, dominated by Markov random fields (Darroch et al., 1980), and artificial intelligence, with a focus on Bayesian networks (Pearl, 1985, 1986). Markov random fields and Bayesian networks are now both considered traditional models. Nevertheless, they are still widely applied and attract a significant amount of research nowadays (Daly et al., 2011; Uhler, 2012). A feature that they have in common is that they both model conditional independence: Bayesian networks relate conditional independence to acyclic directed graphs, whereas it is associated with undirected graphs in Markov fields. However, the models that they represent are only equivalent under additional assumptions on the respective graphs.

In this paper, we review the existing methods for model selection and estimation in undirected and acyclic directed Markov models with Gaussian parametrization. The multivariate Gaussian distribution is one of the most widely developed and applied statistical families in this context (Werhli et al., 2006; Ibáñez et al., 2016). It provides for an explicit parametric comparison of their similarities and differences. The terminology used in methodological developments and theoretical results varies considerably due to the highly interdisciplinary nature of these Markov models. With a few exceptions (Wermuth, 1980; Pearl, 1988), they are usually studied separately, and most unifying works (Sadeghi and Lauritzen, 2014; Wermuth, 2015) are characterized by a high-level view, where the models are embedded in other more expressive models and the focus is on the properties of those container models. In this paper, on the other hand, we review these models from a low-level perspective. In doing so, we use a unified notation by means of which we

can directly compare the two types of models, in terms of both methodological and theoretical developments.

The paper is structured as follows. A historical introduction to Markov models is presented in Section 2, emphasizing the different research areas that contributed to their birth. Preliminary graph theory concepts are then presented in Section 3. In Section 4, undirected and acyclic directed Markov models are introduced under distribution-free assumptions because many foundational relationships between these models can be established from this general perspective. Next, their parametrization is restricted to multivariate Gaussian distributions, and we explore the main properties derived from this constraint in Section 5. We review maximum likelihood estimation in Section 6. These estimates are used for model selection via hypothesis testing, as reported in Section 7. When maximum likelihood estimators are not guaranteed to exist, a popular technique is to employ regularisation, outlined in Section 8. Finally, the alternative Bayesian approach for model selection and estimation is treated in Section 9. We conclude the paper with a discussion of the current main active lines of research and open problems in Section 10.

2 A historical perspective

We now introduce the main terminology applied to Gaussian Markov models currently in use from a historical perspective. Figure 1 is a timeline illustrating the origins of Markov models, including most of the key works that are cited in this section.

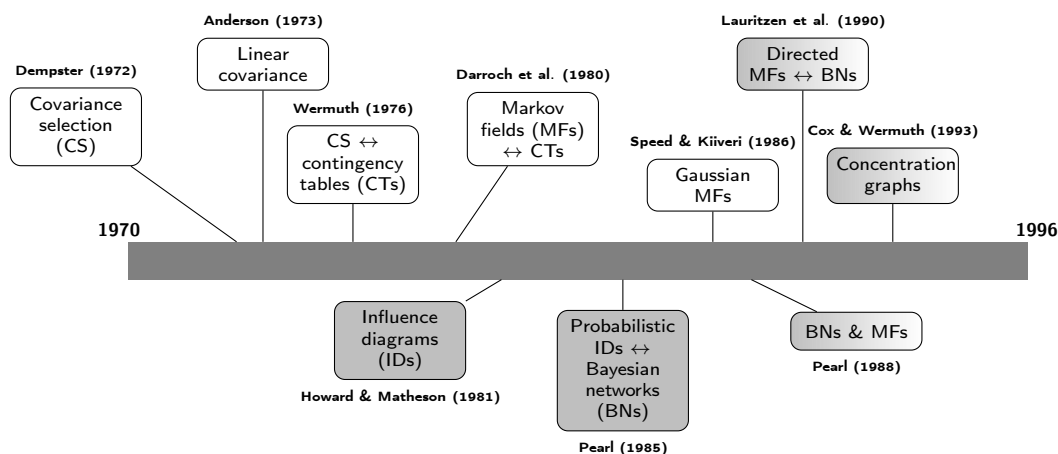


Figure 1: Timeline of the origins of Gaussian Markov models. Papers from the statistical community appear above the line and papers from other research areas are shown below the line. Thematically, grey filled boxes are papers about acyclic directed Markov models, white boxes refer to undirected models, and gradient filled boxes deal with both types of models.

Undirected Markov models for conditional independence are the oldest type of Markov models. They were preceded only by special cases such as the Ising model for ferromagnetic materials (Kindermann and Snell, 1980; Isham, 1981). They are actually a generalization of the Ising model, which is likewise a generalization of Markov chains. Originally, undirected Markov models were called *Markov random fields* (Grimmett, 1973), since they generalized the correspondence between Gibbs measures (Besag, 1974) and Markov properties. The term *graphical*

model was not introduced until Darroch et al. (1980) linked the graphical ideas for contingency tables with Markov properties of discrete Markov fields. They are also termed *Markov networks* (Pearl, 1988) by artificial intelligence researchers by analogy with *Bayesian networks*, used to refer to acyclic directed Markov fields.

Regarding Gaussian parametrization, we find that Anderson (1973) was one of the first researchers to impose some sort of structure on the covariance matrix of a multivariate Gaussian distribution. He aimed to reduce the number of parameters to be estimated and considered the mean vector and covariance matrix to be linear combinations of known linearly independent vectors and matrices, respectively. Along similar lines, Dempster (1972) suggested estimating the inverse of the covariance matrix (concentration matrix) by assuming certain entries to be equal to zero. This was motivated by the representation of the multivariate Gaussian distribution as an exponential family. His work was later referred to as *covariance selection* models. Interestingly, although Dempster did not have any graphical interpretation in mind, such zero entries in the concentration matrix are directly associated with missing edges in undirected Gaussian Markov models, and this correspondence was analysed some years later by Wermuth (1976). On this ground, these Markov models with a Gaussian parameterization are sometimes still called covariance selection models today.

Lastly, in further developments of Markov models, Cox and Wermuth (1993) gave a top-down, unifying view of some graphical models used for representing linear dependence. They equated a special case of graphs with Dempster’s covariance selection models, when assuming a multivariate Gaussian distribution. They referred to these graphs as *concentration graphs*, resembling the zero entries in the concentration matrix.

In contrast, acyclic digraphs were often used as models for multivariate probability distributions after the definition of *influence diagrams*, introduced by Howard and Matheson in 1981 (article reprinted in Howard and Matheson (2005)) and used to model decision-making processes. The probabilistic reduction of influence diagrams coincides with acyclic directed Markov models. This was subsequently studied at length by Pearl (Pearl, 1988), who renamed probabilistic influence diagrams as *Bayesian networks* or *influence networks* (Pearl, 1985). Some researchers working on Markov fields also developed a theory regarding these directed counterparts, which they called *directed Markov fields* (Lauritzen et al., 1990).

There are other works published before the research outlined above employing or referencing acyclic directed Markov models. They were studied implicitly by Wermuth (1980) in the Gaussian case as *linear recursive regression systems*, although the main focus was actually on covariance selection models. The use of directed graphs as graphical models for dependence among random variables can actually be traced back to the work of geneticist Sewall Wright in 1918. Wright developed the method of *path coefficients* (Wright, 1934), nowadays known as *path analysis*. Linearly related variables were represented using a directed acyclic graph, whereas their correlation was represented by bi-directed edges joining the variables.

3 Graph preliminaries

A graph is defined as a pair $\mathcal{G} = (V, E)$ where V is the set of vertices and E is the set of edges. Throughout the whole paper, the graphs will, unless otherwise stated, be labelled and simple. This means that the elements in V are labelled, for example, as $1, \dots, p$; and E is composed of pairs of distinct elements in V . A graph is called *undirected* if these pairs are unordered ($E \subseteq \{\{u, v\} : u, v \in V\}$), and *directed* or *digraph* otherwise ($E \subseteq \{(u, v) : u, v \in V\}$). Edges $\{u, v\}$ in an undirected graph are usually denoted by uv and drawn as a line (see Figure 2a). In a digraph, however, they are called *arcs* or *directed edges* and represented as arrows (Figure 2b

and 2c).




Figure 2: Examples of an undirected graph and two digraphs.

3.1 Undirected graphs

In an undirected graph $\mathcal{G} = (V, E)$, if $uv \in E$, u and v are called *neighbours*. For $v \in V$, the set of its neighbours is denoted by $\text{ne}(v)$, and the *closure* of v is $\text{cl}(v) := \{v\} \cup \text{ne}(v)$. \mathcal{G} is called *complete* if for every $u, v \in V$, $uv \in E$. A maximal $C \subseteq V$ such that \mathcal{G}_C is complete is called a *clique*. Let $\mathcal{H} = (V_{\mathcal{H}}, E_{\mathcal{H}})$ be another undirected graph. \mathcal{H} is a *subgraph* of \mathcal{G} (written as $\mathcal{H} \subseteq \mathcal{G}$) if $V_{\mathcal{H}} \subseteq V$ and $E_{\mathcal{H}} \subseteq E$. If $E_{\mathcal{H}} = \{uv \in E : u, v \in V_{\mathcal{H}}\}$, then \mathcal{H} is called the *induced subgraph* and denoted by $\mathcal{G}_{V_{\mathcal{H}}}$.

A *walk* between u and v is an ordered sequence of vertices $(u =)u_0, u_1, \dots, u_{k-1}, u_k (= v)$, where $u_{i-1}u_i \in E$ for $i \in \{1, \dots, k\}$. The number k is called the *length* of the walk. If $u = v$, the walk is *closed*, and when u_0, \dots, u_{k-1} are distinct, the walk is called a *path*. A closed path of length $k \geq 3$ is called a *cycle* or k -cycle. \mathcal{G} is called *chordal* or *triangulated* if all minimal k -cycles are of length $k = 3$. A *chordal cover* of a graph \mathcal{G} is a graph \mathcal{G}^* such that $\mathcal{G} \subseteq \mathcal{G}^*$ and \mathcal{G}^* is chordal.

$S \subseteq V$ *separates* u and v in $\mathcal{G} = (V, E)$ if there is no path between u and v in the subgraph $\mathcal{G}_{V \setminus S}$. If we consider $A, B, S \subseteq V$, A and B are said to be *separated* by S if u and v are separated by S for all $u \in A, v \in B$. Let V be partitioned into disjoint sets $A, B, S \subseteq V$. (A, B, S) is called a *decomposition* of \mathcal{G} if S separates A and B in \mathcal{G} and \mathcal{G}_S is complete. If $A \neq \emptyset$ and $B \neq \emptyset$, the decomposition is said to be *proper*. An undirected graph is *decomposable* if: (i) it is complete or (ii) it admits a proper decomposition into *decomposable* subgraphs. An undirected graph is decomposable if and only if it is chordal.

3.2 Acyclic digraphs

In a digraph $\mathcal{D} = (V, A)$, the definitions of (induced) subgraph, walk, path, and cycle are analogous to the undirected case. The undirected graph $\mathcal{D}^U := (V, A^U)$ with $A^U := \{uv : (u, v) \in A\}$ is called the *skeleton* of \mathcal{D} , and \mathcal{D} is one of its *orientations*. A digraph \mathcal{D} is said to be *complete* when \mathcal{D}^U is complete.

In the following, assume that \mathcal{D} is acyclic (see Figure 2b and Figure 2c for a cyclic and an acyclic digraph, respectively). The *parent set* of $v \in V$ is $\text{pa}(v) := \{u \in V : (u, v) \in A\}$. Conversely, the *child set* is $\text{ch}(v) := \{u \in V : (v, u) \in A\}$. The *ancestors* of v , $\text{an}(v)$, are those $u \in V$ such that there exists a directed path from u to v ; the *descendants* of v , $\text{de}(v)$, are those $u \in V$ such that there exists a directed path from v to u . The set $\text{nd}(v) := V \setminus (\{v\} \cup \text{de}(v))$ will be the set of *non-descendants* of $v \in V$, and $\text{An}(A) := A \cup (\cup_{a \in A} \text{an}(a))$ the *ancestral set* of $A \subseteq V$. Note that a total order \prec can be defined over the set of vertices V in an acyclic digraph

$\mathcal{D} = (V, A)$, such that if $(u, v) \in A$, then $u \prec v$. This ordering is usually called *ancestral*, and it is a linear extension of the partial order naturally defined as $u \preceq v$ if $u \in \text{an}(v)$. For $v \in V$, the set of *successors* of v with respect to \prec is $\text{su}(v) = \{u \in V : u \succ v\}$; the set of *predecessors* of v is $\text{pr}(v) = \{u \in V : u \prec v\}$.

Finally, let $u, w_1, w_2 \in V$ with $(w_1, u), (w_2, u) \in A$ and $(w_1, w_2), (w_2, w_1) \notin A$ (see vertices 1, 2 and 3 in Figure 2c). Such configurations are usually called *v-structures* and denoted by $w_1 \rightarrow u \leftarrow w_2$. The *moral graph* of \mathcal{D} is defined as the undirected graph $\mathcal{D}^m = (V, A^m)$ with $A^m := A^U \cup \{w_1 w_2 : w_1 \rightarrow u \leftarrow w_2 \text{ for some } u \in V\}$.

4 Undirected and acyclic directed Markov models

The Markov models that we review associate conditional independence in random vectors $\mathbf{X} = (X_1, \dots, X_p)^t$ with undirected graph and acyclic digraph separation properties. This is explicitly specified by the *Markov properties* of the distribution of \mathbf{X} , which are in turn based on what are known as *independence relations*.

In the following, for arbitrary $I \subseteq \{1, \dots, p\}$, we will denote the $|I|$ -dimensional subvector of \mathbf{X} by $\mathbf{X}_I := (X_i)_{i \in I}$. Conditional independence will be expressed as $\mathbf{X}_I \perp\!\!\!\perp \mathbf{X}_J \mid \mathbf{X}_K$, which represents the statement ‘ \mathbf{X}_I is conditionally independent from \mathbf{X}_J given \mathbf{X}_K ’ (see Dawid (1979)).

4.1 Independence relations

An *independence relation* over a set $V = \{1, \dots, p\}$ is a collection \mathcal{I} of triples (A, B, C) , where A, B and C are pairwise disjoint subsets of V . It is called a *semi-graphoid* when the following conditions are met:

- if $(A, B, C) \in \mathcal{I}$ then $(B, A, C) \in \mathcal{I}$,
- if $(A, B \cup C, D) \in \mathcal{I}$ then $(A, C, D) \in \mathcal{I}$ and $(A, B, C \cup D) \in \mathcal{I}$,
- if $(A, B, C \cup D) \in \mathcal{I}$ and $(A, C, D) \in \mathcal{I}$ then $(A, B \cup C, D) \in \mathcal{I}$,

and a *graphoid* when, additionally, if $(A, B, C \cup D) \in \mathcal{I}$ and $(A, C, B \cup D) \in \mathcal{I}$ then $(A, B \cup C, D) \in \mathcal{I}$ (Pearl and Paz, 1987).

Independence relations occur in different contexts that are relevant for Markov models. Specifically, an independence relation \mathcal{I} over $V = \{1, \dots, p\}$ is said to be *induced* by

- an undirected graph $\mathcal{G} = (V, E)$ if $(A, B, S) \in \mathcal{I} \iff A$ and B are separated by S in \mathcal{G} ,
- an acyclic digraph $\mathcal{D} = (V, A)$ if $(A, B, S) \in \mathcal{I} \iff A$ and B are separated by S in $(\mathcal{D}_{\text{An}(A \cup B \cup S)})^m$,
- a p -dimensional random vector \mathbf{X} if $(A, B, S) \in \mathcal{I} \iff \mathbf{X}_A \perp\!\!\!\perp \mathbf{X}_B \mid \mathbf{X}_S$.

Graph-induced independence relations are always graphoids, while probabilistic independence relations are always semi-graphoids. Additional assumptions on the probability spaces involved are required for probabilistic independence relations to be graphoids (Dawid, 1980).

The core of Markov models is the relationship between induced independence relations, which we will denote by $\mathcal{I}(\cdot)$ with the argument being the inducing element. Specifically, if \mathcal{G} is an undirected (acyclic directed) graph, an undirected (directed) *Markov model* is defined as

$$\mathcal{M}(\mathcal{G}) := \{P_{\mathbf{X}} : \mathcal{I}(\mathcal{G}) \subseteq \mathcal{I}(\mathbf{X})\},$$

where the random vectors \mathbf{X} are defined over the same probability space and $P_{\mathbf{X}}$ denotes their distribution. These models are non-empty (Geiger and Pearl, 1990, 1993), that is, for any undirected or acyclic directed graph, there is always a probability distribution whose independence model contains the one generated by the graph.

4.2 Markov properties

When a distribution $P_{\mathbf{X}}$ belongs to $\mathcal{M}(\mathcal{G})$ for an undirected or acyclic directed graph \mathcal{G} , it is said that $P_{\mathbf{X}}$ is *globally \mathcal{G} -Markov* or satisfies the *global Markov property* with respect to \mathcal{G} . Other weaker Markov properties can be defined that are usually capable of simplifying the model. Specifically, if $\mathcal{G} = (V, E)$ is an undirected graph, then the probability distribution $P_{\mathbf{X}}$ of \mathbf{X} is said to be

- *pairwise \mathcal{G} -Markov* if $X_u \perp\!\!\!\perp X_v \mid \mathbf{X}_{V \setminus \{u,v\}}$ for all $uv \notin E$,
- *locally \mathcal{G} -Markov* if $X_v \perp\!\!\!\perp \mathbf{X}_{V \setminus \text{cl}(v)} \mid \mathbf{X}_{\text{ne}(v)}$ for all $v \in V$,

whereas if \mathcal{G} is an acyclic digraph, then $P_{\mathbf{X}}$ is called

- *pairwise \mathcal{G} -Markov* if $X_u \perp\!\!\!\perp X_v \mid \mathbf{X}_{\text{nd}(u) \setminus \{v\}}$ for all $u \in V, v \in \text{nd}(u) \setminus \text{pa}(u)$,
- *locally \mathcal{G} -Markov* if $X_v \perp\!\!\!\perp \mathbf{X}_{\text{nd}(v) \setminus \text{pa}(v)} \mid \mathbf{X}_{\text{pa}(v)}$ for all $v \in V$.

The three Markov properties are equivalent when \mathcal{G} is acyclic directed (Lauritzen et al., 1990), while, if \mathcal{G} is undirected, this equivalence is only guaranteed when $\mathcal{I}(\mathbf{X})$ is a graphoid (Pearl, 1988). Suffice it for $P_{\mathbf{X}}$ to admit a continuous and strictly positive density for this to occur. This sufficient condition was obtained in different forms by several authors, but it is usually attributed to Hammersley and Clifford (1971), who were the first to outline the proof for the discrete case (Speed, 1979). It relies on an additional characterization of a probability distribution with respect to \mathcal{G} : if $\mathfrak{C}(\mathcal{G})$ denotes the class of cliques of \mathcal{G} , the density function f of $P_{\mathbf{X}}$ is said to *factorize according to \mathcal{G}* when there exists a set $\{\psi_C(\mathbf{x}_C) : C \in \mathfrak{C}(\mathcal{G}), \psi_C \geq 0\}$ such that

$$f(\mathbf{x}) = \prod_{C \in \mathfrak{C}(\mathcal{G})} \psi_C(\mathbf{x}_C). \quad (1)$$

When Equation (1) holds, then $P_{\mathbf{X}}$ is globally \mathcal{G} -Markov, while if f is continuous and strictly positive, the pairwise Markov property implies Equation (1). This yields the equivalence of Markov properties.

Finally, recall that the nodes of an acyclic digraph $\mathcal{D} = (V, A)$ can be totally ordered such that if $(u, v) \in A$, then $u \in \text{pr}(v)$. This gives rise to another Markov property, exclusive to acyclic digraphs: $P_{\mathbf{X}}$ is said to be *ordered \mathcal{D} -Markov* if $X_v \perp\!\!\!\perp \mathbf{X}_{\text{pr}(v) \setminus \text{pa}(v)} \mid \mathbf{X}_{\text{pa}(v)}$ for all $v \in V$. This property is also equivalent to the global, local and pairwise Markov properties (Lauritzen et al., 1990).

4.3 Independence and Markov equivalence

When the Markov models defined by two graphs \mathcal{G} and $\tilde{\mathcal{G}}$, with the same vertex set V , coincide, such graphs are said to be *Markov equivalent*. A simpler notion, also implying Markov equivalence, is *independence equivalence*, holding when $\mathcal{I}(\mathcal{G}) = \mathcal{I}(\tilde{\mathcal{G}})$. Independence equivalence is implied by Markov equivalence under fairly general circumstances (Studený, 2005, §6.1), which is why most authors treat them as the same notion. The best suited graph for the Markov model can be chosen based on these equivalences.

We will first characterize equivalence within undirected graphs. For each graphoid \mathcal{I} over V , there exists a unique edge-minimal undirected graph \mathcal{G} such that $\mathcal{I}(\mathcal{G}) \subseteq \mathcal{I}$ (Pearl and Paz, 1987). It follows that $\mathcal{I}(\mathcal{G}) = \mathcal{I}(\tilde{\mathcal{G}})$ (independence equivalence) if and only if \mathcal{G} and $\tilde{\mathcal{G}}$ are identical. Furthermore, if we assume that $\mathcal{I}(\mathbf{X})$ is a graphoid for all $P_{\mathbf{X}} \in \mathcal{M}(\mathcal{G})$, then a unique edge-minimal $\tilde{\mathcal{G}}$ exists, with $\tilde{\mathcal{G}} \subseteq \mathcal{G}$, such that $\mathcal{M}(\mathcal{G}) = \mathcal{M}(\tilde{\mathcal{G}})$ (Markov equivalence), that is, a unique undirected graph can be chosen as representative of each undirected Markov model.


In contrast, acyclic digraphs are not, generally, unique representations of a Markov model, since $\mathcal{I}(\mathcal{D}) = \mathcal{I}(\tilde{\mathcal{D}})$ if and only if \mathcal{D} and $\tilde{\mathcal{D}}$ have the same skeleton and the same v-structures (Verma and Pearl, 1991). However, unique representatives can be constructed: Let \mathfrak{D}_p be the set of acyclic digraphs over $V = \{1, \dots, p\}$ and define an equivalence relation \sim in \mathfrak{D}_p as $\mathcal{D} \sim \tilde{\mathcal{D}} \iff \mathcal{I}(\mathcal{D}) = \mathcal{I}(\tilde{\mathcal{D}})$. The quotient space of \sim is $\mathfrak{D}_p/\sim = \{[\mathcal{D}] : \mathcal{D} \in \mathfrak{D}_p\}$, where $[\mathcal{D}] := \{\tilde{\mathcal{D}} \in \mathfrak{D}_p : \tilde{\mathcal{D}} \sim \mathcal{D}\}$ is the *Markov equivalence class*; thus, $\mathcal{M}(\tilde{\mathcal{D}}) = \mathcal{M}(\mathcal{D})$ for all $\tilde{\mathcal{D}} \in [\mathcal{D}]$, that is, $[\mathcal{D}]$ is the unique representative of the directed Markov model.

Finally, we characterize the equivalence between directed and undirected graphs, first obtained by Wermuth (1980) for multivariate Gaussian distributions and by Wermuth and Lauritzen (1983) for contingency tables, and then generalized by Frydenberg (1990) for graphoid-inducing distributions. When \mathcal{G} is an undirected graph, $\mathcal{M}(\mathcal{G}) = \mathcal{M}(\mathcal{D})$ for some acyclic digraph \mathcal{D} if and only if \mathcal{G} is chordal. Conversely, an acyclic digraph \mathcal{D} is Markov equivalent to its skeleton \mathcal{D}^U if and only if \mathcal{D} contains no v-structures. Furthermore, \mathcal{D} can be related to its moral graph. This requires an analogous formula to Equation (1): a density function f is said to *recursively factorize* according to \mathcal{D} when

$$f(\mathbf{x}) = \prod_{v \in V} f(x_v \mid \mathbf{x}_{\text{pa}(v)}).$$

This characterization is equivalent to the Markov properties, and also implies that f factorizes as in Equation (1) with respect to the moral graph \mathcal{D}^m (Lauritzen et al., 1990). This means that $P_{\mathbf{X}}$ is always globally \mathcal{D}^m -Markov for continuous \mathbf{X} , and thus $\mathcal{M}(\mathcal{D}) \subseteq \mathcal{M}(\mathcal{D}^m)$, with the equality only holding when $\mathcal{D}^m = \mathcal{D}^U$.

Example 1. Figure 3 illustrates the above concepts. The graph in Figure 3a is not chordal, and thus there is no Markov equivalent acyclic digraph. Figure 3b is a chordal cover of Figure 3a, and a Markov equivalent orientation is depicted in Figure 3c. The acyclic digraph in Figure 3d has v-structures, emphasized in dark grey, and thus cannot be Markov equivalent to its skeleton (Figure 3a). The moral graph in Figure 3d is Figure 3e, which in fact is another chordal cover of Figure 3a, and thus none of its orientations will be Markov equivalent to Figure 3c.



(a) Chordless cycle (b) Chordal cover (c) Orientation (d) V-structures (e) Moral graph

Figure 3: Markov equivalence.

5 Gaussian parametrization

When restricting to multivariate Gaussian distributions, we find connections between conditional independence and vanishing parameters. This correspondence can be used to provide a direct interpretation of Markov properties in both the undirected and directed cases, allowing an enhanced manipulation of these Markov models.

In the following, the elements of a real $q \times r$ matrix $\mathbf{M} \in \mathbb{M}_{q \times r}(\mathbb{R})$ will be denoted by m_{ij} , where $i \in \{1, \dots, q\}$ and $j \in \{1, \dots, r\}$. \mathbf{M}_{IJ} will be the $|I| \times |J|$ submatrix of \mathbf{M} , where $I \subseteq \{1, \dots, q\}$ and $J \subseteq \{1, \dots, r\}$, and we will use \mathbf{M}_{IJ}^{-1} as $(\mathbf{M}_{IJ})^{-1}$. $\mathbb{S}^{>0}$ and $\mathbb{S}^{\geq 0}$ will represent the sets of positive and semi-positive definite symmetric matrices, respectively. The p -variate Gaussian distribution is denoted by $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu} \in \mathbb{R}^p$ is the mean vector and $\boldsymbol{\Sigma} \in \mathbb{S}^{>0}$ is the covariance matrix. \mathbf{I}_p will denote the $p \times p$ identity matrix. Dimensional subscripts will often be dropped if the dimension of the respective object is clear from the context.

5.1 Markov properties and the multivariate Gaussian distribution

Let $V = \{1, \dots, p\}$. When a random vector \mathbf{X} is distributed as $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, then for $i, j \in V$, $X_i \perp\!\!\!\perp X_j$ is equivalent to $\sigma_{ij} = 0$. If we consider a partition (I, J) of V , then $\mathbf{X}_I \mid \mathbf{x}_J$ is distributed as $\mathcal{N}_{|I|}(\boldsymbol{\mu}_I, \boldsymbol{\Sigma}_{I \cdot J})$, where $\boldsymbol{\Sigma}_{I \cdot J} = \boldsymbol{\Sigma}_{II} - \boldsymbol{\Sigma}_{IJ} \boldsymbol{\Sigma}_{JJ}^{-1} \boldsymbol{\Sigma}_{JI}$ (Anderson, 2003). Thus, for $i, k \in I$, we have that $X_i \perp\!\!\!\perp X_k \mid \mathbf{x}_J$ is equivalent to $\sigma_{ik \cdot J} = 0$, the (i, k) element in the conditional covariance matrix $\boldsymbol{\Sigma}_{I \cdot J}$.

A correspondence can be established between the zeros in $\boldsymbol{\Sigma}_{I \cdot J}$ and zero patterns in other representative matrices (Wermuth, 1976, 1980) as follows. Let the concentration matrix of \mathbf{X} be $\boldsymbol{\Omega} = \boldsymbol{\Sigma}^{-1}$, with elements ω_{uv} for $u, v \in V$. The matrix $\boldsymbol{\Sigma}_{IJ} \boldsymbol{\Sigma}_{JJ}^{-1}$ is usually denoted by $\mathbf{B}_{I \cdot J}$ and called the matrix of *regression coefficients* of \mathbf{X}_I on \mathbf{X}_J . Let $\boldsymbol{\Omega}_{I \cdot J} := \boldsymbol{\Sigma}_{JJ}^{-1}$, then we have the following matrix identity (Horn and Johnson, 2012).

$$\boldsymbol{\Omega} = \begin{pmatrix} \boldsymbol{\Sigma}_{II} & \boldsymbol{\Sigma}_{IJ} \\ \boldsymbol{\Sigma}_{JI} & \boldsymbol{\Sigma}_{JJ} \end{pmatrix}^{-1} = \begin{pmatrix} \boldsymbol{\Sigma}_{I \cdot J}^{-1} & -\boldsymbol{\Sigma}_{I \cdot J}^{-1} \mathbf{B}_{I \cdot J} \\ -\mathbf{B}_{I \cdot J}^t \boldsymbol{\Sigma}_{I \cdot J}^{-1} & \boldsymbol{\Omega}_{I \cdot J} + \mathbf{B}_{I \cdot J}^t \boldsymbol{\Sigma}_{I \cdot J}^{-1} \mathbf{B}_{I \cdot J} \end{pmatrix}.$$

Hence $\boldsymbol{\Sigma}_{I \cdot J}$ can be related to $\boldsymbol{\Omega}$ and $\mathbf{B}_{I \cdot J}$, as

$$\boldsymbol{\Sigma}_{I \cdot J} = \boldsymbol{\Omega}_{II}^{-1}, \quad (2)$$

$$\mathbf{B}_{I \cdot J} = -\boldsymbol{\Omega}_{II}^{-1} \boldsymbol{\Omega}_{IJ}. \quad (3)$$

This implies that, dually, $\boldsymbol{\Omega}_{II}$ is identically equal to the concentration matrix of $\mathbf{X}_I \mid \mathbf{x}_J$, while $\boldsymbol{\Omega}_{I \cdot J}$ is the concentration matrix of \mathbf{X}_J .

Using the above correspondences, conditional independence can also be related with $\boldsymbol{\Omega}$ and $\mathbf{B}_{I \cdot J}$. From Equation (2) we get, for $i, k \in V$,

$$X_i \perp\!\!\!\perp X_j \mid \mathbf{X}_{V \setminus \{i, k\}} \iff \omega_{ik} = 0, \quad (4)$$

whereas from Equation (3) it follows that, for $J \subseteq V$, $i, k \in V \setminus J$,

$$X_i \perp\!\!\!\perp X_k \mid \mathbf{X}_J \iff \beta_{ik \cdot J \cup \{k\}} = 0, \quad (5)$$

where $\beta_{ik \cdot J \cup \{k\}}$ is the v entry in the vector $\boldsymbol{\beta}_{i \cdot J \cup \{k\}}^t$, that is, the coefficient of X_k on the regression of X_i on $\mathbf{x}_{J \cup \{k\}}$. The original notation for this, introduced by Yule (1907), was $\beta_{ik \cdot J}$, that is, k is implicitly considered as included in the conditioning indexes. However, we have opted for the alternative explicit notation $\beta_{ik \cdot J \cup \{k\}}$, since it simplifies the notation in later sections.

5.2 Gaussian Markov models

In the Gaussian case, undirected Markov models are in correspondence with the concentration matrix, while for acyclic digraphs this correspondence is with the regression coefficients. Both rely on the auxiliary Markov properties that we presented in Section 4.2.

Let $\mathcal{G} = (V, E)$ be an undirected graph and consider \mathbf{X} distributed as $P_{\mathbf{X}} \equiv \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $P_{\mathbf{X}} \in \mathcal{M}(\mathcal{G})$. Since $P_{\mathbf{X}}$ is globally \mathcal{G} -Markov, it is also pairwise \mathcal{G} -Markov, and thus Equation (4) directly yields $\omega_{uv} = 0$ for all $u, v \in V$ such that $uv \in E$. This means that, if we define the set $\mathbb{S}^{>0}(\mathcal{G}) := \{\mathbf{M} \in \mathbb{S}^{>0} : m_{uv} = 0 \text{ for all } uv \notin E\}$, we have $\boldsymbol{\Omega} \in \mathbb{S}^{>0}(\mathcal{G})$ if and only if $P_{\mathbf{X}}$ is pairwise \mathcal{G} -Markov. Furthermore, since the multivariate Gaussian distribution has positive density, $\mathcal{I}(\mathbf{X})$ is a graphoid and thus the three Markov properties are equivalent. In this manner, we can redefine the *Gaussian undirected Markov model* as

$$\mathcal{N}(\mathcal{G}) = \{\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma}) : \boldsymbol{\Sigma}^{-1} \in \mathbb{S}^{>0}(\mathcal{G}), \boldsymbol{\mu} \in \mathbb{R}^p\}. \quad (6)$$

In the directed case, the redefinition is not so direct. Let $\mathcal{D} = (V, A)$ be an acyclic digraph, and assume, for notational simplicity, that the nodes are already ancestrally ordered as $1 \preceq \dots \preceq p$. If \mathbf{X} is distributed as $P_{\mathbf{X}} \equiv \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $P_{\mathbf{X}} \in \mathcal{M}(\mathcal{D})$, it satisfies the ordered Markov property. Thus, whenever $v \in \text{pr}(u) \setminus \text{pa}(u)$, we have $X_u \perp\!\!\!\perp X_v \mid \mathbf{X}_{\text{pa}(u)}$, which is equivalent to $\beta_{uv \cdot \text{pa}(u) \cup \{v\}} = 0$ as in Equation (5). Since we have assumed an ancestral order, $\beta_{uv \cdot \text{pa}(u) \cup \{v\}} = \beta_{uv \cdot \text{pr}(u)}$ for all $u \in V, v \in \text{pr}(u) \setminus \text{pa}(u)$, which leads to $P_{\mathbf{X}}$ being ordered \mathcal{D} -Markov if and only if $\beta_{uv \cdot \text{pr}(u)} = 0$ for all $u \in V, v \in \text{pr}(u) \setminus \text{pa}(u)$. This triangular requirement on the regression coefficients can be expressed with the triangular matrix \mathbf{B} defined as $b_{uv} = 0$ for $v < u, v \notin \text{pa}(u)$, and $b_{uv} = \beta_{uv \cdot \text{pa}(u)} \equiv \beta_{uv \cdot \text{pr}(u)}$ otherwise.

If $v_u := \sigma_{uu \cdot \text{pr}(u)}$, the above characterization leads to a matrix form of the respective linear regressions, as $\mathbf{X} = \boldsymbol{\mu} + \mathbf{B}(\mathbf{X} - \boldsymbol{\mu}) + \mathbf{E}$, where $E_u \sim \mathcal{N}(0, v_u)$. We can rearrange it as

$$\mathbf{X} = \mathbf{U}^{-1}\boldsymbol{\xi} + \mathbf{U}^{-1}\mathbf{E}, \quad (7)$$

where $\boldsymbol{\xi} := \mathbf{U}\boldsymbol{\mu}$ and $\mathbf{U} := \mathbf{I}_p - \mathbf{B}$, since \mathbf{U} is invertible. Let \mathbf{V} be the diagonal matrix of conditional variances v . Sometimes $\boldsymbol{\xi}, \mathbf{B}$ and \mathbf{V} are called the \mathcal{D} -parameters of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ (Andersson and Perlman, 1998). In fact, using \mathbf{U} and \mathbf{V} , $\boldsymbol{\Sigma}$ (and $\boldsymbol{\Sigma}^{-1}$) can be decomposed as $\boldsymbol{\Sigma} = \mathbf{U}^{-1}\mathbf{V}\mathbf{U}^{-t}$. Furthermore, this decomposition uniquely determines $\boldsymbol{\Sigma}$ via \mathbf{U}/\mathbf{B} and \mathbf{V} (Horn and Johnson, 2012). Thus, by analogy with Equation (6), if we define the set $\mathbb{M}(\mathcal{D}) := \{\mathbf{M} \in \mathbb{M}_{p \times p}(\mathbb{R}) : m_{uv} = 0 \text{ for all } (u, v) \notin A\}$ and the set $\boldsymbol{\Delta}_p$ of $p \times p$ diagonal matrices, we can redefine the *Gaussian directed Markov model* as

$$\mathcal{N}(\mathcal{D}) = \{\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma}) : \boldsymbol{\Sigma}^{-1} = (\mathbf{I}_p - \mathbf{B})^t \mathbf{V}^{-1} (\mathbf{I}_p - \mathbf{B}), \mathbf{B} \in \mathbb{M}(\mathcal{D}), \mathbf{V} \in \boldsymbol{\Delta}_p\}. \quad (8)$$

6 Maximum likelihood estimation

Exponential family theory greatly simplifies maximum likelihood estimation (Barndorff-Nielsen, 1978). The multivariate Gaussian distribution is a regular exponential family, and thus both undirected and directed Gaussian Markov models can be expressed as special subfamilies of this family.

6.1 The Gaussian family and maximum likelihood

The canonical parameter in the multivariate Gaussian family is $\boldsymbol{\eta} = (\boldsymbol{\Omega}\boldsymbol{\mu}, -\boldsymbol{\Omega}/2)$ over the space $\mathcal{H} = \{(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2) : \boldsymbol{\eta}_1 \in \mathbb{R}^p, -\boldsymbol{\eta}_2 \in \mathbb{S}^{>0}\}$, and the sufficient statistics are $\mathbf{t}(\mathbf{X}) = (\mathbf{X}, \mathbf{X}\mathbf{X}^t)$.

Let $\{\mathbf{x}^{(n)} : 1 \leq n \leq N\}$ be N independent observations, where $\mathbf{X}^{(n)} \sim \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ for each $n \in \{1, \dots, N\}$, arranged in $\mathbf{x} \in \mathbb{M}_{p \times N}(\mathbb{R})$, the respective random matrix being \mathbf{X} . The random sample is also a regular exponential family with canonical parameter $\boldsymbol{\eta} = (\boldsymbol{\Omega}\boldsymbol{\mu}, -\boldsymbol{\Omega}/2)$ over the space \mathcal{H} . The sufficient statistics in this case are $\mathbf{t}(\mathbf{X}) = (N\bar{\mathbf{X}}, \mathbf{X}\mathbf{X}^t)$ with $N\bar{\mathbf{X}} = \sum_{n=1}^N \mathbf{X}^{(n)}$.

In a regular exponential family $\mathcal{F}_{\mathcal{H}}$, a maximum of the likelihood function, $\mathcal{L}(\boldsymbol{\eta})$, given a random sample $\mathbf{X} = \mathbf{x}$, is reached in \mathcal{H} if and only if $\mathbf{t}(\mathbf{x})$ belongs to the interior of $\mathcal{C}(\mathbf{t})$, the closed convex hull of the support of the distribution of \mathbf{t} , denoted by $\text{int}(\mathcal{C}(\mathbf{t}))$. When this occurs, the maximum is unique and given by $\boldsymbol{\eta} \in \mathcal{H}$ satisfying $\mathbb{E}[\mathbf{t}(\mathbf{X})] = \mathbf{t}(\mathbf{x})$.

For the multivariate Gaussian random sample, we have that $\mathbb{E}[N\bar{\mathbf{X}}] = N\boldsymbol{\mu}$ and $\mathbb{E}[\mathbf{X}\mathbf{X}^t] = N\boldsymbol{\Sigma} + N\boldsymbol{\mu}\boldsymbol{\mu}^t$, thus the convex support of $\mathbf{t}(\mathbf{X}) = (N\bar{\mathbf{X}}, \mathbf{X}\mathbf{X}^t)$ is $\mathcal{C}(\mathbf{t}) = \{(\mathbf{v}, \mathbf{M}) \in \mathbb{R}^p \times \mathbb{S}_p : \mathbf{M} - \mathbf{v}\mathbf{v}^t/N \in \mathbb{S}^{\geq 0}\}$. This indicates that the maximum likelihood estimator for $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ exists if and only if $\mathbf{x}\mathbf{x}^t - N\bar{\mathbf{x}}\bar{\mathbf{x}}^t \in \mathbb{S}^{\geq 0}$, whose probability of occurrence is one whenever $N > p$ and zero otherwise. The solution in this case is $(\bar{\mathbf{x}}, \mathbf{Q}/N)$, where

$$\mathbf{Q} = \sum_{n=1}^N \left(\mathbf{X}^{(n)} - \bar{\mathbf{X}} \right) \left(\mathbf{X}^{(n)} - \bar{\mathbf{X}} \right)^t = \mathbf{X}\mathbf{X}^t - N\bar{\mathbf{X}}\bar{\mathbf{X}}^t.$$

A particular scenario, which is usually assumed, is when $\boldsymbol{\mu} = \mathbf{0}$. In this case, the canonical parameter is $\boldsymbol{\eta} = -\boldsymbol{\Omega}/2$ in the space $\{\boldsymbol{\eta} : -\boldsymbol{\eta} \in \mathbb{S}^{\geq 0}\}$, and the sufficient statistic is $\mathbf{t}(\mathbf{X}) = \mathbf{X}\mathbf{X}^t$. The maximum likelihood estimator exists if and only if $\mathbf{x}\mathbf{x}^t \in \mathbb{S}^{\geq 0}$ and, when it does exist, it is $\mathbf{X}\mathbf{X}^t/N = \mathbf{Q}/N$.

6.2 Gaussian Markov models as exponential families

When \mathcal{G} is an undirected graph, the set $\mathbb{S}^{\geq 0}(\mathcal{G})$ is a convex (linear) cone inside the positive definite cone $\mathbb{S}^{\geq 0}$ (Uhler, 2012), which means that $\mathbb{R}^p \times \mathbb{S}^{\geq 0}(\mathcal{G})$ is an affine subspace of $\mathbb{R}^p \times \mathbb{S}^{\geq 0}$, and thus $\mathcal{N}(\mathcal{G})$ is also a regular exponential family (Barndorff-Nielsen, 1978). Assume that $\boldsymbol{\mu} = \mathbf{0}$ and let $\mathbf{Q}^{\mathcal{G}}$ be the projection of \mathbf{Q} on $E \cup \{uu : u \in V\}$, that is, such that $q_{uv}^{\mathcal{G}} = 0$ for all $uv \notin E$ with $u \neq v$. Since $\mathcal{L}(\boldsymbol{\Omega}) \propto \det(\boldsymbol{\Omega})^{\frac{1}{2}} \exp(-\text{tr}(\boldsymbol{\Omega}\mathbf{Q}))$ and $\boldsymbol{\Omega} \in \mathbb{S}^{\geq 0}(\mathcal{G})$, we have $\text{tr}(\boldsymbol{\Omega}\mathbf{Q}) = \text{tr}(\boldsymbol{\Omega}\mathbf{Q}^{\mathcal{G}})$ and the sufficient statistic for $\mathcal{N}(\mathcal{G})$ is $\mathbf{t}(\mathbf{x}) = \mathbf{Q}^{\mathcal{G}}$ (Lauritzen, 1996). Its convex support is $\mathcal{C}(\mathbf{t}) = \{\mathbf{P}^{\mathcal{G}} : \mathbf{P} \in \mathbb{S}^{\geq 0}\}$, alternatively referred to as the set of projections extensible to full positive definite matrices. Thus, the maximum likelihood estimator for $\boldsymbol{\Sigma}$ exists if and only if $\mathbf{Q}^{\mathcal{G}} \in \text{int}(\mathcal{C}(\mathbf{t}))$. This is equivalent to $\mathbf{Q}^{\mathcal{G}}$ being extensible to a full positive definite matrix. Whenever it exists, it is the only extensible matrix $\hat{\boldsymbol{\Sigma}}$ that also satisfies the model restriction $\hat{\boldsymbol{\Sigma}}^{-1} \in \mathbb{S}^{\geq 0}(\mathcal{G})$. A sufficient condition thus is that $\mathbf{Q} \in \mathbb{S}^{\geq 0}$, which almost surely holds for $N \geq p$.

Now we turn to the case where the random sample \mathbf{X} is assumed to follow multivariate Gaussian distribution constrained by the separation properties in an acyclic digraph. The restriction in Equation (8), however, is not linear in the canonical parameter. In fact, Geiger et al. (2001) show that they are curved exponential families. Multivariate linear regression theory can be applied to obtain the maximum likelihood estimates (Andersson and Perlman, 1998). Recall that if $\mathbf{X} \sim \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \in \mathcal{N}(\mathcal{D})$, then \mathbf{X} can be expressed as Equation (7). Thus,

we can estimate the \mathcal{D} -parameters for $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ as the ordinary least squares estimators:

$$\begin{aligned}\widehat{\boldsymbol{\beta}}_{u,\text{pa}(u)}^t &= \left(\sum_{n=1}^N (x_u^{(n)} - \bar{x}_u) (\mathbf{x}_{\text{pa}(u)}^{(n)} - \bar{\mathbf{x}}_{\text{pa}(u)})^t \right) \left(\sum_{n=1}^N (\mathbf{x}_{\text{pa}(u)}^{(n)} - \bar{\mathbf{x}}_{\text{pa}(u)}) (\mathbf{x}_{\text{pa}(u)}^{(n)} - \bar{\mathbf{x}}_{\text{pa}(u)})^t \right)^{-1}, \\ \widehat{\xi}_u &= \bar{x}_u - \widehat{\boldsymbol{\beta}}_{u,\text{pa}(u)}^t \bar{\mathbf{x}}_{\text{pa}(u)}, \\ N\widehat{v}_{uu} &= N\widehat{\sigma}_{u,\text{pa}(u)} = \sum_{n=1}^N (x_u^{(n)} - \bar{x}_u)^2 - \widehat{\boldsymbol{\beta}}_{u,\text{pa}(u)}^t \left(\sum_{n=1}^N (\mathbf{x}_{\text{pa}(u)}^{(n)} - \bar{\mathbf{x}}_{\text{pa}(u)}) (x_u^{(n)} - \bar{x}_u) \right),\end{aligned}$$

respectively for each $u \in V$. We can then directly obtain the maximum likelihood estimator for $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ from its respective \mathcal{D} -parameter estimators (see Andersson and Perlman, 1998, for an algorithm). As opposed to the undirected case, $(\widehat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\Sigma}})$ exist with probability one if and only if $N \geq p + \max\{|\text{pa}(u)| : u \in V\}$.

7 Model selection via hypothesis testing

The maximum likelihood estimators presented in Section 6 can be used to address the problem of model estimation. A statistical procedure is required to previously select the graph that will define the Markov model. In this section, we will review the main hypothesis testing methods for this purpose.

In the undirected case, we are interested in testing the hypothesis $H_0 : \boldsymbol{\Omega} \in \mathbb{S}^{\succ 0}(\mathcal{G}_0)$ against $H_1 : \boldsymbol{\Omega} \in \mathbb{S}^{\succ 0}(\mathcal{G})$, where $\mathcal{G}_0 = (V, E_{\mathcal{G}_0}) \subseteq \mathcal{G} = (V, E_{\mathcal{G}})$. The result of this test determines whether the edges in $E_{\mathcal{G}} \setminus E_{\mathcal{G}_0}$ should be excluded from the selected model. On this ground, these tests are usually known as *edge exclusion tests*. Note also that this is *backward* model selection, since our null hypothesis consists of a subgraph. Let $\widehat{\boldsymbol{\Sigma}}_0$ and $\widehat{\boldsymbol{\Sigma}}$ be the maximum likelihood estimators for a covariance matrix in the Markov model determined by \mathcal{G}_0 and \mathcal{G} , respectively. The likelihood ratio statistic is

$$T_L = \frac{\det(\widehat{\boldsymbol{\Sigma}})^{\frac{N}{2}}}{\det(\widehat{\boldsymbol{\Sigma}}_0)^{\frac{N}{2}}}.$$

Under H_0 , $-2 \log(T_L)$ is asymptotically distributed as $\chi^2_{|E_{\mathcal{G}}| - |E_{\mathcal{G}_0}|}$ up to terms of order $N^{-1/2}$. However, this is a poor approximation in many cases (Porteous, 1989). More accurate distributional results, which are in fact exact in the chordal case, have been derived (Porteous, 1989). If \mathcal{G} and \mathcal{G}_0 are chordal, and $\mathcal{G}_0 \subset \dots \subset \mathcal{G}_k (= \mathcal{G})$ is a sequence, where, for $1 \leq i \leq k$, $\mathcal{G}_i = (V, E_{\mathcal{G}_i})$ is chordal and $E_{\mathcal{G}_{i-1}} = E_{\mathcal{G}_i} \setminus \{e_i\}$ for some $e_i = u_i v_i \in E_{\mathcal{G}_i}$ (sequence of edge deletions), then, under H_0 , $T_L^{2/N}$ is distributed as the product $\prod_{i=1}^k B_i$ of univariate Beta variables, where, for $1 \leq i \leq k$,

$$B_i \sim \mathcal{B} \left(\frac{1}{2}(N - |\text{ne}(u_i) \cap \text{ne}(v_i)| - 1), \frac{1}{2} \right).$$

When \mathcal{G}_0 and \mathcal{G} have the same non-chordal minimal subgraphs, this result is a better approximation than the χ^2 (Eriksen, 1996). When performing model selection with these tests, multiple testing error rates need to be controlled. To do this, Drton and Perlman (2004) propose an alternative to the previous stepwise methods: a set of simultaneous p -values and confidence intervals is obtained such that the edge set is estimated, for a significance level α , as

$$\widehat{E}^\alpha := \left\{ uv : \sqrt{N-p} |z_{uv \cdot V \setminus \{u,v\}}| > \Phi^{-1} \left(\frac{1}{2} (1 - \alpha)^{\frac{2}{p(p-1)}} + \frac{1}{2} \right) \right\}, \quad (9)$$

where $z_{uv \cdot V \setminus \{u,v\}}$ is the Z -transform of $r_{uv \cdot V \setminus \{u,v\}}$, the sample partial correlation,

$$z_{uv \cdot V \setminus \{u,v\}} = \frac{1}{2} \log \left(\frac{1 + r_{uv \cdot V \setminus \{u,v\}}}{1 - r_{uv \cdot V \setminus \{u,v\}}} \right),$$

and Φ is the cumulative distribution function of a standard Gaussian. Denoting $\widehat{\mathcal{G}}^\alpha = (V, \widehat{E}^\alpha)$, $\liminf_{N \rightarrow \infty} P(\widehat{\mathcal{G}}^\alpha = \mathcal{G}) \geq 1 - \alpha$ holds if the distribution under consideration $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \in \mathcal{N}(\mathcal{G})$ conforms to \mathcal{G} , that is, if $\omega_{uv} = 0 \iff uv \notin E$ (this condition is also known as *faithfulness*). If it does not conform, then the result holds with respect to the smallest graph \mathcal{H} such that $\mathcal{G} \subseteq \mathcal{H}$ and $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ conforms to \mathcal{H} .

In the case of a directed Gaussian Markov model over an acyclic digraph $\mathcal{D} = (V, A)$, most of the results are adaptations from analogues in multivariate linear Gaussian models. The likelihood ratio, whose moments are also characterized in Andersson and Perlman (1998), is

$$T_L = \frac{\det(\widehat{\boldsymbol{\Sigma}})^{\frac{N}{2}}}{\det(\widetilde{\boldsymbol{\Sigma}})^{\frac{N}{2}}} = \frac{\prod_{v \in V} |\widehat{\sigma}_{vv} - \widehat{\boldsymbol{\sigma}}_{v, \text{pa}(v)}^t \widehat{\boldsymbol{\Sigma}}_{\text{pa}(v)}^{-1} \widehat{\boldsymbol{\sigma}}_{v, \text{pa}(v)}|}{\prod_{v \in V} |\widetilde{\sigma}_{vv} - \widetilde{\boldsymbol{\sigma}}_{v, \text{pa}(v)}^t \widetilde{\boldsymbol{\Sigma}}_{\text{pa}(v)}^{-1} \widetilde{\boldsymbol{\sigma}}_{v, \text{pa}(v)}|},$$

where $\widetilde{\boldsymbol{\Sigma}}$ and $\widehat{\boldsymbol{\Sigma}}$ are the respective maximum likelihood estimators for $\widetilde{\mathcal{D}}$ and \mathcal{D} , $\widetilde{\mathcal{D}} \subseteq \mathcal{D}$. The multiple testing procedure in Equation (9) has also been extended in Drton and Perlman (2008), obtaining an estimate of the arc set as

$$\widehat{A}^\alpha := \left\{ (v, u) : v < u \text{ and } \sqrt{N - u - 1} |z_{uv \cdot \text{pr}(u) \setminus \{v\}}| > \Phi^{-1} \left(\frac{1}{2} (1 - \alpha)^{\frac{2}{p(p-1)}} + \frac{1}{2} \right) \right\}, \quad (10)$$

where an ancestral ordering \prec is being assumed in V such that the resulting permutation is the identity, that is, such that $v \prec u \iff v < u$. Consistency is established as in the undirected case; note the symmetry with Equation (9). See Drton and Perlman (2007) for a general discussion on some variations of Equation (9) and Equation (10) and their impact on overall error control.

8 Regularization

Regularization approaches, which simultaneously perform model selection and estimation, have become popular in the context of Markov models. They are usually applied when $N < p$, and thus the existence of the maximum likelihood estimator is not guaranteed. The main consistency results available for both the directed and undirected cases share sparseness and high-dimensionality assumptions, as we will see below. There are two different approaches: methods that penalize likelihood and methods that instead focus on the regression coefficients.

Throughout this section, we will employ the asymptotic notation, specifically symbols $O(\cdot)$ and $\Theta(\cdot)$, which stand for asymptotic inferiority and equivalence, respectively. For $\mathbf{M} \in \mathbb{M}_{q \times r}(\mathbb{R})$, $\text{vec}(\mathbf{M})$ will denote the vectorized function of \mathbf{M} , $(m_{11}, \dots, m_{q1}, \dots, m_{1r}, \dots, m_{qr})^t$. This way, the operator norm of \mathbf{M} will be denoted by $\|\mathbf{M}\|$; whereas $\|\mathbf{M}\|_{q+r}$ will be used to denote $\|\text{vec}(\mathbf{M})\|_{q+r}$, $\|\cdot\|_p$ being the p -norm function. If \mathbf{v} is a p -vector, $\text{diag}(\mathbf{v})$ will denote the matrix \mathbf{M} in $\boldsymbol{\Delta}_p$ with main diagonal \mathbf{v} ; analogously, $\text{diag}(\mathbf{M}) \in \boldsymbol{\Delta}_p$ will have the same diagonal as \mathbf{M} , and \mathbf{M}^- will be used for $\mathbf{M} - \text{diag}(\mathbf{M})$.

$\mathcal{N}(\mathcal{G})$ (Meinshausen and Bühlmann, 2006)	$\mathcal{N}(\mathcal{D})$ (van de Geer and Bühlmann, 2013)
l_1 regularization	l_0 regularization
$p \in O(N^\epsilon)$	$p \in O(N/\log(N))$
Lower bound on $ \rho_{uv \cdot V \setminus \{u,v\}} $	Lower bound on $ \beta_{uv \cdot \text{pr}(u) \setminus \{v\}} $
Upper bound on $ \text{ne}(v) $	Upper bound on $ \text{pa}(v) $
Bounded neighbourhood perturbations	Bounded permutation perturbations

Table 1: Comparison of assumptions for consistency results on penalized estimation in Gaussian Markov models.

8.1 Node-wise regression

Let $\mathcal{G} = (V, E)$ be an undirected graph, with $V = \{1, \dots, p\}$. Meinshausen and Bühlmann (2006) penalize the regression function, as

$$\hat{\mathbf{b}}_u^\lambda := \arg \min \left\{ \frac{1}{N} \|\mathbf{x}_u - \mathbf{b}_u \mathbf{x}\|_2^2 + \lambda f(\mathbf{b}_u) : b_{uu} = 0 \right\}, \quad (11)$$

where $\lambda \geq 0$ and $f(\cdot)$ is the penalty function. Let $\widehat{\text{ne}}(v) := \{u \in V : \hat{b}_{vu} \neq 0\}$ and

$$\begin{aligned} \widehat{E}_\wedge &:= \{uv : u \in \widehat{\text{ne}}(v) \text{ and } v \in \widehat{\text{ne}}(u)\}, \\ \widehat{E}_\vee &:= \{uv : u \in \widehat{\text{ne}}(v) \text{ or } v \in \widehat{\text{ne}}(u)\}. \end{aligned}$$

If E_\wedge, E_\vee denote the respective population versions, then $E_\wedge = E_\vee$, since $u \in \text{ne}(v) \iff v \in \text{ne}(u)$, for all $u, v \in V$, whereas this may not be true with the estimated sets. Let $f(\mathbf{B}) = \|\mathbf{B}\|_1$, commonly known as the *lasso* penalty (Tibshirani, 1996), and assume that

- (a) there exists $\epsilon > 0$ such that $p \in O(N^\epsilon)$ as $N \rightarrow \infty$;
- (b) $|\text{ne}(v)|$ and $\|\beta_{v \cdot \text{ne}(u)}\|_1$ are upper-bounded for all $v \in V$ and $u \in \text{ne}(v)$;
- (c) $|\rho_{uv \cdot V \setminus \{u,v\}}|$ is lower-bounded for all $uv \in E$;
- (d) $\left| \sum_{z \in \text{ne}(v)} \text{sign}(\beta_{vz \cdot \text{ne}(v)}) \beta_{uz \cdot \text{ne}(v)} \right| < 1$, for all $u, v \in V$ with $u \notin \text{ne}(v)$.

Then, both \widehat{E}_\wedge and \widehat{E}_\vee are consistent estimators of E for a choice of λ under assumptions (a)-(d) (Meinshausen and Bühlmann, 2006). This result was also independently discovered by Zhao and Yu (2006), who termed assumption (d) as the *irrepresentable condition*.

As an alternative to the lasso penalty, van de Geer and Bühlmann (2013) use l_0 regularization in the context of directed Gaussian Markov models. Like Meinshausen and Bühlmann (2006), their approach uses regression coefficients, more generally, \mathcal{D} -parameters. As such, the assumptions required for the consistency of both methods share some symmetry, as outlined in Table 1. The estimators in this case are obtained as

$$(\widehat{\mathbf{V}}, \widehat{\mathbf{B}}) = \arg \min \left\{ -N \log |\boldsymbol{\Omega}| + \text{tr}(\boldsymbol{\Omega} \mathbf{S}) + \lambda^2 |\mathcal{A}| : \boldsymbol{\Omega} = (\mathbf{I}_p - \mathbf{B})^t \mathbf{V}^{-1} (\mathbf{I}_p - \mathbf{B}), \mathbf{B} \in \mathbb{M}(\mathcal{D}) \right\},$$

where $\lambda \geq 0$, and they are equal among Markov equivalent models (van de Geer and Bühlmann, 2013).

8.2 Penalized likelihood

In Equation (11) the zero entries are estimated in \mathbf{B} . Alternatively, conditional independences can be obtained from $\mathbf{\Omega}$, for $\lambda \geq 0$, as

$$\widehat{\mathbf{\Omega}}^\lambda = \arg \min \left\{ -N \log \det(\mathbf{\Omega}) + \text{tr}(\mathbf{\Omega}\mathbf{S}) + \lambda f(\mathbf{\Omega}) : \mathbf{\Omega} \in \mathbb{S}^{\succ 0}(\mathcal{G}) \right\},$$

where $\mathbf{S} = \mathbf{Q}/N$. Yuan and Lin (2007) were the first to pursue this approach, and they chose $f(\mathbf{\Omega}) = \|\mathbf{\Omega}^-\|_1$. Later, Banerjee et al. (2008) included the diagonal elements in the regularization function, that is, $f(\mathbf{\Omega}) = \|\mathbf{\Omega}\|_1$. In the former case, the convergence rate is (Rothman et al., 2008)

$$\left\| \widehat{\mathbf{\Omega}}^\lambda - \mathbf{\Omega} \right\|_2 \in O \left(\sqrt{\frac{(|E| + p) \log(p)}{N}} \right) \text{ as } N \rightarrow \infty. \quad (12)$$

A relaxation of Equation (12) can be obtained based on the correlation matrix as follows. Since $\mathbf{\Sigma} = \mathbf{D}\mathbf{P}\mathbf{D}$ with \mathbf{P} the correlation matrix and \mathbf{D} is the diagonal matrix of standard deviations, if we let the corresponding sample estimates be $\widehat{\mathbf{D}}^2 = \text{diag}(\widehat{\mathbf{\Sigma}})$ and $\widehat{\mathbf{P}} = \widehat{\mathbf{D}}^{-1}\widehat{\mathbf{\Sigma}}\widehat{\mathbf{D}}^{-1}$, then we can estimate $\mathbf{K} = \mathbf{P}^{-1}$ as

$$\widehat{\mathbf{K}}^\lambda = \arg \min \left\{ -N \log \det(\mathbf{K}) + \text{tr}(\mathbf{K}\widehat{\mathbf{P}}) + \lambda f(\mathbf{K}) : \mathbf{K} \in \mathbb{S}^{\succ 0}(\mathcal{G}) \right\},$$

for $\lambda \geq 0$. The concentration matrix can then be alternatively estimated as $\widehat{\mathbf{\Omega}}^\lambda = \widehat{\mathbf{D}}^{-1}\widehat{\mathbf{K}}^\lambda\widehat{\mathbf{D}}^{-1}$. Both regularized estimators of $\mathbf{\Omega}$ are consistent when $f(\mathbf{\Omega}) = \|\mathbf{\Omega}^-\|_1$ and $f(\mathbf{K}) = \|\mathbf{K}^-\|_1$ (Rothman et al., 2008): assuming that the eigenvalues of $\mathbf{\Sigma}$ are positive and bounded, then, if $\lambda \in \Theta \left(\sqrt{\log(p)/N} \right)$ as $N \rightarrow \infty$,

$$\left\| \widehat{\mathbf{\Omega}}^\lambda - \mathbf{\Omega} \right\| \in O \left(\sqrt{\frac{(|E| + 1) \log(p)}{N}} \right) \text{ as } N \rightarrow \infty. \quad (13)$$

9 Bayesian model selection and estimation

In Bayesian model selection, the graph \mathcal{G} with highest ‘a posteriori’ probability, $\mathbb{P}(\mathcal{G} \mid \mathbf{X})$, (usually called the *score*) is chosen.

In the following, the *p-variate Wishart distribution* will be denoted by $\mathcal{W}_p(n, \mathbf{\Lambda})$ with $n \in \mathbb{R}$, $n > p - 1$ and $\mathbf{\Lambda} \in \mathbb{M}_{p \times p}(\mathbb{R})$, $\mathbf{\Lambda} \succ 0$; analogously, the *p-variate inverse Wishart distribution* will be $\mathcal{W}_p^{-1}(\nu, \mathbf{\Psi})$ with $\nu \in \mathbb{R}$, $\nu > p - 1$ and $\mathbf{\Psi} \in \mathbb{M}_{p \times p}(\mathbb{R})$, $\mathbf{\Psi} \succ 0$.

9.1 Hyper Markov laws

When \mathcal{G} is undirected, the target probability is $\mathbb{P}(\mathcal{G}, \mathbf{\Omega} \mid \mathbf{X}) \propto \mathbb{P}(\mathbf{X} \mid \mathcal{G}, \mathbf{\Omega}) \mathbb{P}(\mathbf{\Omega} \mid \mathcal{G}) \mathbb{P}(\mathcal{G})$. A uniform probability over the space of undirected graphs is usually chosen for $\mathbb{P}(\mathcal{G})$. However, this choice is biased towards medium-sized graphs, and thus other prior distributions have been proposed (Jones et al., 2005; Mohammadi and Wit, 2015).

For $\mathbb{P}(\mathbf{\Omega} \mid \mathcal{G})$, Dawid and Lauritzen (1993) defined what are known as the *hyper Markov laws* for chordal graphs, which Roverato and Whittaker (1998) generalised to the non-chordal case. Let $\boldsymbol{\theta}$ be a random variable taking values over $\mathcal{N}(\mathcal{G})$. The probability distribution of $\boldsymbol{\theta}$ is said to be (*weakly*) *hyper \mathcal{G} -Markov* if, for any decomposition (A, B, S) of \mathcal{G} , it holds that $\boldsymbol{\theta}_A \perp\!\!\!\perp \boldsymbol{\theta}_B \mid \boldsymbol{\theta}_S$; if it further holds that $\boldsymbol{\theta}_{B|A} \perp\!\!\!\perp \boldsymbol{\theta}_A$, it is called *strongly hyper \mathcal{G} -Markov*. For chordal graphs, if the probability distribution of $\boldsymbol{\theta}$ is strongly hyper \mathcal{G} -Markov with respect to

\mathcal{G} , then the probability distribution of $\boldsymbol{\theta} \mid \mathbf{x}$ is the unique (strong) hyper \mathcal{G} -Markov distribution specified by the clique-marginal distributions $\{\mathbb{P}(\boldsymbol{\theta}_C \mid \mathbf{x}_C^{(1)}, \dots, \mathbf{x}_C^{(N)}) : C \in \mathfrak{C}(\mathcal{G})\}$; and, when densities exist, $f(\boldsymbol{\theta}_C \mid \mathbf{x}) \propto f(\mathbf{x}_C^{(1)}, \dots, \mathbf{x}_C^{(N)} \mid \boldsymbol{\theta}_C) f(\boldsymbol{\theta}_C)$ (Dawid and Lauritzen, 1993). In other words, under these assumptions, it is possible to localize computations over the graph cliques when performing Bayesian inference.

Assuming $\boldsymbol{\mu} = \mathbf{0}$, the inverse Wishart is a conjugate prior for $\boldsymbol{\Omega}$, that is, if $\boldsymbol{\Sigma} \sim \mathcal{W}_p^{-1}(\nu, \boldsymbol{\Psi})$, then $\boldsymbol{\Sigma} \mid \mathbf{Q}/N \sim \mathcal{W}_p^{-1}(N + \nu, \mathbf{Q} + \boldsymbol{\Psi})$ (recall $\mathbf{Q} = \mathbf{x}\mathbf{x}^t$). We can thus construct the *hyper inverse Wishart distribution*, as the unique hyper Markov distribution associated with inverse Wishart clique marginals: $\boldsymbol{\Sigma}_{CC} \sim \mathcal{W}_{|C|}^{-1}(\nu, \boldsymbol{\Psi}^C)$, for each clique $C \in \mathfrak{C}(\mathcal{G})$. This hyper Markov distribution is denoted by $\mathcal{HW}_p^{-1}(\nu, \boldsymbol{\Psi})$, where $\boldsymbol{\Psi} \in \mathbb{S}^{\succ 0}$ such that $\boldsymbol{\Psi}_{CC} = \boldsymbol{\Psi}^C$ for each clique $C \in \mathfrak{C}(\mathcal{G})$. This distribution is strongly hyper \mathcal{G} -Markov. The main advantage of this prior is that it has many properties that mirror those for Markov models, since hyper Markov distributions are also defined in terms of an underlying graph.

An explicit expression for the hyper inverse Wishart density is devised in Giudici (1996). In order to set the parameters of the hyper inverse Wishart distribution, we might adopt a hierarchical Bayesian method such as is reported by Giudici and Green (1999), where δ and $\boldsymbol{\Psi}$ are assumed to have a gamma and Wishart distribution, respectively.

9.2 Conjugate priors

In the directed case, the Bayesian approach has mainly been developed for model selection. Thus, the target probability is $\mathbb{P}(\mathcal{D} \mid \mathbf{X}) \propto \mathbb{P}(\mathbf{X} \mid \mathcal{D}) \mathbb{P}(\mathcal{D})$, where \mathcal{D} denotes an acyclic digraph,

$$\mathbb{P}(\mathbf{X} \mid \mathcal{D}) = \int_{(\boldsymbol{\mu}, \boldsymbol{\Omega}) \in \mathbb{R}^p \times \mathbb{S}^{\succ 0}} \mathbb{P}(\mathbf{X} \mid \boldsymbol{\mu}, \boldsymbol{\Omega}, \mathcal{D}) \mathbb{P}(\boldsymbol{\mu}, \boldsymbol{\Omega} \mid \mathcal{D}) \, d\boldsymbol{\mu} \, d\boldsymbol{\Omega}, \quad (14)$$

and $\mathbb{P}(\mathcal{D})$ is usually set as uniform. Consider directed Markov models over a general parametric family of distributions, \mathcal{F}_θ (as is the case of Gaussian directed Markov models). The following assumptions are used to simplify the computation of Equation (14):

- (a) Let $\mathcal{D}, \tilde{\mathcal{D}}$ be two complete acyclic digraphs. Then $\mathcal{M}(\mathcal{D}) = \mathcal{M}(\tilde{\mathcal{D}})$ and the transformation between their parameters is regular (complete model equivalence and regularity).
- (b) Let $\mathcal{D} = (V, E_{\mathcal{D}})$ and $\tilde{\mathcal{D}} = (V, E_{\tilde{\mathcal{D}}})$ be acyclic digraphs and $v \in V$ such that $\text{pa}_{\mathcal{D}}(v) = \text{pa}_{\tilde{\mathcal{D}}}(v)$. Then, $\mathbb{P}(X_v \mid \mathbf{X}_{\text{pa}(v)}, \theta_v, \mathcal{D}) \equiv \mathbb{P}(X_v \mid \mathbf{X}_{\text{pa}(v)}, \theta_v, \tilde{\mathcal{D}})$ and $\mathbb{P}(\theta_v \mid \mathcal{D}) \equiv \mathbb{P}(\theta_v \mid \tilde{\mathcal{D}})$ (likelihood and prior modularity).
- (c) Let $\mathcal{M}(\mathcal{D})$ be a directed Markov model over \mathcal{F}_θ . Then $\mathbb{P}(\boldsymbol{\theta} \mid \mathcal{D}) = \prod_{v \in V} \mathbb{P}(\theta_v \mid \mathcal{D})$ (global parameter independence).

When they are satisfied, if \mathcal{D}_c is an arbitrary complete digraph, then

$$\mathbb{P}(\mathbf{X} \mid \mathcal{D}) = \mathbb{P}(\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(N)} \mid \mathcal{D}) = \prod_{v \in V} \frac{\mathbb{P}(\mathbf{X}_{\{v\} \cup \text{pa}_{\mathcal{D}}(v)}^{(1)}, \dots, \mathbf{X}_{\{v\} \cup \text{pa}_{\mathcal{D}}(v)}^{(N)} \mid \mathcal{D}_c)}{\mathbb{P}(\mathbf{X}_{\text{pa}_{\mathcal{D}}(v)}^{(1)}, \dots, \mathbf{X}_{\text{pa}_{\mathcal{D}}(v)}^{(N)} \mid \mathcal{D}_c)}, \quad (15)$$

and Equation (15) is equal among independence equivalent directed Markov models (usually called *score equivalence* property) (Geiger and Heckerman, 2002).

When \mathcal{F}_θ is the Gaussian family, the conjugate prior for $(\boldsymbol{\mu}, \boldsymbol{\Omega})$ unknown is the normal-Wishart distribution, where $\boldsymbol{\Omega} \sim \mathcal{W}_p(\alpha_{\boldsymbol{\Omega}}, \mathbf{A})$ and $\boldsymbol{\mu} \mid \boldsymbol{\Omega} \sim \mathcal{N}_p(\boldsymbol{\mu}_0, (\alpha_{\boldsymbol{\mu}} \boldsymbol{\Omega})^{-1})$. Geiger and Heckerman (1994) obtain an explicit expression for each factor in Equation (15): for $U \subseteq V$,

$$f(\mathbf{x}_U^{(1)}, \dots, \mathbf{x}_U^{(N)} \mid \mathcal{D}) = \left(\frac{\alpha_{\boldsymbol{\mu}}}{\alpha_{\boldsymbol{\mu}} + N} \right)^{\frac{|U|}{2}} 2\pi^{-\frac{|U|N}{2}} \frac{\Gamma_{|U|} \left(\frac{N + \alpha_{\boldsymbol{\Omega}} - p + |U|}{2} \right)}{\Gamma_{|U|} \left(\frac{\alpha_{\boldsymbol{\Omega}} - p + |U|}{2} \right)} \frac{|\mathbf{A}_{UU}|^{\frac{\alpha_{\boldsymbol{\Omega}} - p + |U|}{2}}}{|\mathbf{R}_{UU}|^{\frac{\alpha_{\boldsymbol{\Omega}} - p + |U| + N}{2}}},$$

where $\Gamma_p(\cdot)$ is the p -dimensional Gamma function, and

$$\mathbf{R} = \mathbf{\Lambda} + \mathbf{Q} + \frac{N\alpha_{\Omega}}{N + \alpha_{\Omega}}(\boldsymbol{\mu}_0 - \bar{\mathbf{x}})(\boldsymbol{\mu}_0 - \bar{\mathbf{x}})^t.$$

Furthermore, Geiger and Heckerman (2002) characterize the normal-Wishart prior for $(\boldsymbol{\mu}, \boldsymbol{\Omega})$ as the only distribution satisfying the global parameter independence assumption.

Another score commonly used in both the directed and undirected cases is the *Bayesian information criterion* (BIC) (Schwarz, 1978). It was originally defined and proven to be consistent for model selection in linear exponential families. Haughton (1988) extended this consistency to curved exponential families. Thus, since undirected Gaussian Markov models are linear, and directed ones are curved, it is a valid scoring metric for both of them. It is an approximation of $-2\log(\mathbb{P}(\mathbf{X} | \mathcal{D}))$, for choosing the model with highest $\mathbb{P}(\mathcal{D} | \mathbf{X})$, as $-2\log(\mathcal{L}(\hat{\boldsymbol{\theta}})) + k\log(N)$, where $\hat{\boldsymbol{\theta}}$ is the maximum likelihood estimator of $\boldsymbol{\theta}$ and k is the dimension of the model.

10 Open problems

Figure 4 illustrates the main papers containing the methods and theory that we have reviewed in this paper on a timeline following on from the chronology illustrated in Figure 1 at the beginning of the paper.

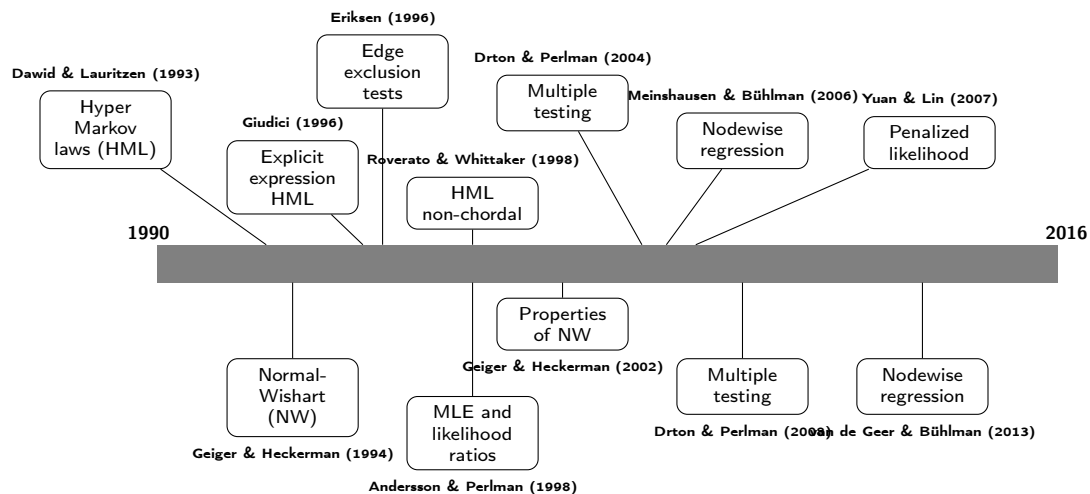


Figure 4: Timeline of Gaussian Markov model selection and estimation since their origins. Papers concerning undirected Markov models appear at the top, while papers dealing with the acyclic directed case appear underneath. MLE stands for maximum likelihood estimation.

Although directed and undirected Markov models are considered as traditional models, several open problems remain.

At the foundational level, independence relations can be generalised to what are known as *separoids* (Dawid, 2001), a relaxation of graphoids. Separoids usually appear whenever a notion of ‘irrelevance’ is being mathematically treated, although they are sometimes too restrictive (Cozman and Walley, 2005). Furthermore, they can be expressed as abstract axiom sets (Córdoba-Sánchez et al., 2016), closely related to the recently defined *independence logic* (Grädel and Väänänen, 2013). Further research on these axiom systems from an abstract point

of view could shed more light on how the apparently different mathematical contexts in which such structures arise are related.

Hammersley and Clifford’s theorem (Hammersley and Clifford, 1971) is a straightforward tool for checking whether an independence model originated from a distribution is a graphoid; however, the condition is not necessary and sufficient. Further relaxations remain to be found, although positivity is generally essential (Moussouris, 1974). Graphoids are also essentially necessary for characterizing Markov equivalence classes, since most of them assume this condition. It will also be interesting to get deeper insights into how these assumptions influence such characterizations. Furthermore, also regarding independence models, recall the space of Markov equivalence classes of acyclic digraphs, \mathfrak{D}_p/\sim . The asymptotic ratio $l = \lim_{p \rightarrow \infty} |\mathfrak{D}_p|/|\mathfrak{D}_p/\sim|$ influences the computational gain obtained by using \mathfrak{D}_p/\sim instead of \mathfrak{D}_p as a search space for model selection. Steinsky (2004) analytically calculates an upper bound of l as 13.65. Exact computations by Gillispie and Perlman (2002), for $p \leq 10$, and approximations by Sonntag et al. (2015), up to $p = 31$, seem to indicate that $l \sim 3.7$. However, its analytical deduction remains an open problem. Note that the computational gain is not only influenced by l , but also by other factors, such as how the element size in \mathfrak{D}_p/\sim is distributed.

The existence of the maximum likelihood estimator in Gaussian undirected Markov models has not yet been completely characterized. Since the problem lies at the interface between statistics and linear algebra, several of the existing results have been independently discovered by researchers in both areas. For chordal graphs, the problem was solved separately by Grone et al. (1984) and Frydenberg and Lauritzen (1989): if $\mathfrak{C}(\mathcal{G}^*)$ is the class of cliques in \mathcal{G} , \mathcal{G}^* is a chordal cover of \mathcal{G} , and $q^* := \max\{|C| : C \in \mathfrak{C}(\mathcal{G}^*)\}$, $q := \max\{|C| : C \in \mathfrak{C}(\mathcal{G})\}$; then there is a probability of one that $\widehat{\Sigma}$ exists if $N \geq q^*$, and does not exist if $N < q$. This result does not account for the case $q \leq N < q^*$ for non-chordal graphs, since otherwise $q = q^*$. Special subtypes of non-chordal graphs are p -cycles, which were addressed by Barrett et al. (1993) (from the viewpoint of linear algebra), and separately by Buhl (1993) (from the perspective of statistics): there is a probability strictly between zero and one of $\widehat{\Sigma}$ existing if $N = 2$. Finally, Uhler (2012) recently detailed, from the algebraic viewpoint, results paralleling findings reported by Buhl (1993) for bipartite graphs.

Regarding both regularization and hypothesis testing, with a few exceptions (Drton and Perlman, 2007), most methods have been asymmetrically developed for either undirected or directed models, as we have shown. A theoretical comparison of how these approaches complement each other would be a key contribution to the field. Specifically, the implications of their assumptions is an essential consideration in regularization when applied to real data, and, as shown in Table 1, there are plenty of similarities that could lead to a unified method.

Finally, Bayesian estimation was actively researched after the seminal papers by Dawid and Lauritzen (1993) (undirected graphs) and Geiger and Heckerman (1994) (directed graphs). Some researchers have extended hyper Markov properties to other Markov models (Rajaratnam et al., 2008). It would be interesting to see how these two widely researched families of distributions are related, for example, to check if hyper Markov distributions satisfy the assumptions stated in Geiger and Heckerman (1994) and vice versa.

Acknowledgements

Research partially supported by the Spanish Ministry of Economy and Competitiveness through the Cajal Blue Brain (C080020-09; the Spanish partner of the Blue Brain initiative from EPFL) and TIN2013-41592-P projects, by the Regional Government of Madrid through the S2013/ICE-2845-CASI-CAM-CM project, and by the European Union’s Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project).

References

- Anderson, T. W. (1973). Asymptotically efficient estimation of covariance matrices with linear structure. *Ann. Statist.*, 1(1):135–141.
- Anderson, T. W. (2003). *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons, 3rd edition.
- Andersson, S. A. and Perlman, M. D. (1998). Normal linear regression models with recursive graphical Markov structure. *Journal of Multivariate Analysis*, 66(2):133 – 187.
- Banerjee, O., El Ghaoui, L., and d’Aspremont, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516.
- Barndorff-Nielsen, O. (1978). *Information and Exponential Families in Statistical Theory*. John Wiley & Sons.
- Barrett, W., Johnson, C. R., and Tarazaga, P. (1993). The real positive definite completion problem for a simple cycle. *Linear Algebra and Its Applications*, 192:3 – 31.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *J. R. Statist. Soc. B*, 36(2):192–236.
- Buhl, S. L. (1993). On the existence of maximum likelihood estimators for graphical Gaussian models. *Scand. J. Statist.*, 20(3):263–270.
- Cox, D. R. and Wermuth, N. (1993). Linear dependencies represented by chain graphs. *Statistical Science*, 8(3):204–218.
- Cozman, F. G. and Walley, P. (2005). Graphoid properties of epistemic irrelevance and independence. *Annals of Mathematics and Artificial Intelligence*, 45(1):173–195.
- Córdoba-Sánchez, I., Bielza, C., and Larrañaga, P. (2016). Graphoids and separoids in model theory. Technical Report TR:UPM-ETSINF/DIA/2016-1, Universidad Politécnica de Madrid.
- Daly, R., Shen, Q., and Aitken, S. (2011). Learning Bayesian networks: Approaches and issues. *The Knowledge Engineering Review*, 26:99–157.
- Darroch, J. N., Lauritzen, S. L., and Speed, T. P. (1980). Markov fields and log-linear interaction models for contingency tables. *Ann. Statist.*, 8(3):522–539.
- Dawid, A. P. (1979). Conditional independence in statistical theory. *J. R. Statist. Soc. B*, 41(1):1–31.
- Dawid, A. P. (1980). Conditional independence for statistical operations. *Ann. Statist.*, 8(3):598–617.
- Dawid, A. P. (2001). Separoids: A mathematical framework for conditional independence and irrelevance. *Annals of Mathematics and Artificial Intelligence*, 32(1):335–372.
- Dawid, A. P. and Lauritzen, S. L. (1993). Hyper Markov laws in the statistical analysis of decomposable graphical models. *Ann. Statist.*, 21(3):1272–1317.
- Dempster, A. P. (1972). Covariance selection. *Biometrics*, 28(1):157–175.

- Drton, M. and Perlman, M. D. (2004). Model selection for Gaussian concentration graphs. *Biometrika*, 91(3):591–602.
- Drton, M. and Perlman, M. D. (2007). Multiple testing and error control in Gaussian graphical model selection. *Statistical Science*, 22(3):430–449.
- Drton, M. and Perlman, M. D. (2008). A SINful approach to Gaussian graphical model selection. *J. Statist. Plannng Inf.*, 138(4):1179–1200.
- Eriksen, P. S. (1996). Tests in covariance selection models. *Scand. J. Statist.*, 23(3):275–284.
- Frydenberg, M. (1990). The chain graph Markov property. *Scand. J. Statist.*, 17(4):333–353.
- Frydenberg, M. and Lauritzen, S. L. (1989). Decomposition of maximum likelihood in mixed graphical interaction models. *Biometrika*, 76(3):539–555.
- Geiger, D. and Heckerman, D. (1994). Learning Gaussian networks. In *Proc. of the Tenth Conference on Uncertainty in Artificial Intelligence*, pages 235–243, San Francisco. Morgan Kaufmann.
- Geiger, D. and Heckerman, D. (2002). Parameter priors for directed acyclic graphical models and the characterization of several probability distributions. *Ann. Statist.*, 30(5):1412–1440.
- Geiger, D., Heckerman, D., King, H., and Meek, C. (2001). Stratified exponential families: Graphical models and model selection. *Ann. Statist.*, 29(2):505–529.
- Geiger, D. and Pearl, J. (1990). On the logic of causal models. In *Proc. of the Fourth Annual Conference on Uncertainty in Artificial Intelligence*, pages 3–14, Corvallis. AUAI Press.
- Geiger, D. and Pearl, J. (1993). Logical and algorithmic properties of conditional independence and graphical models. *Ann. Statist.*, 21(4):2001–2021.
- Gillispie, S. B. and Perlman, M. D. (2002). The size distribution for Markov equivalence classes of acyclic digraph models. *Artificial Intelligence*, 141(1–2):137 – 155.
- Giudici, P. (1996). Learning in graphical Gaussian models. *Bayesian Statistics*, 5:621–628.
- Giudici, P. and Green, P. J. (1999). Decomposable graphical Gaussian model determination. *Biometrika*, 86(4):785–801.
- Grädel, E. and Väänänen, J. (2013). Dependence and independence. *Studia Logica*, 101(2):399–410.
- Grimmett, G. R. (1973). A theorem about random fields. *Bulletin of the London Mathematical Society*, 5(1):81–84.
- Grone, R., Johnson, C. R., Sá, E. M., and Wolkowicz, H. (1984). Positive definite completions of partial Hermitian matrices. *Linear Algebra and Its Applications*, 58:109 – 124.
- Hammersley, J. M. and Clifford, P. (1971). Markov fields on finite graphs and lattices. Unpublished manuscript.
- Haughton, D. M. A. (1988). On the choice of a model to fit data from an exponential family. *Ann. Statist.*, 16(1):342–355.

- Horn, R. A. and Johnson, C. R. (2012). *Matrix Analysis*. Cambridge University Press, 2nd edition.
- Howard, R. A. and Matheson, J. E. (2005). Influence diagrams. *Decision Analysis*, 2(3):127–143.
- Ibáñez, A., Armañanzas, R., Bielza, C., and Larrañaga, P. (2016). Genetic algorithms and Gaussian Bayesian networks to uncover the predictive core set of bibliometric indices. *Journal of the Association for Information Science and Technology*, 67(7):1703–1721.
- Isham, V. (1981). An introduction to spatial point processes and Markov random fields. *International Statistical Review*, 49(1):21–43.
- Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C., and West, M. (2005). Experiments in stochastic computation for high-dimensional graphical models. *Statistical Science*, 20(4):388–400.
- Kindermann, R. and Snell, J. L. (1980). *Markov Random Fields and their Applications*. American Mathematical Society.
- Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press.
- Lauritzen, S. L., Dawid, A. P., Larsen, B. N., and Leimer, H.-G. (1990). Independence properties of directed Markov fields. *Networks*, 20(5):491–505.
- Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.*, 34(3):1436–1462.
- Mohammadi, A. and Wit, E. C. (2015). Bayesian structure learning in sparse Gaussian graphical models. *Bayesian Analysis*, 10(1):109–138.
- Moussouris, J. (1974). Gibbs and Markov random systems with constraints. *Journal of Statistical Physics*, 10(1):11–33.
- Pearl, J. (1985). Bayesian networks: A model of self-activated memory for evidential reasoning. Technical Report R-43, University of California, Los Angeles.
- Pearl, J. (1986). Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3):241 – 288.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Pearl, J. and Paz, A. (1987). Graphoids: A graph-based logic for reasoning about relevance relations. In *Advances in Artificial Intelligence*, volume 2, pages 357–363. Elsevier.
- Porteous, B. T. (1989). Stochastic inequalities relating a class of log-likelihood ratio statistics to their asymptotic χ^2 distribution. *Ann. Statist.*, 17(4):1723–1734.
- Rajaratnam, B., Massam, H., and Carvalho, C. M. (2008). Flexible covariance estimation in graphical Gaussian models. *Ann. Statist.*, 36(6):2818–2849.
- Rothman, A. J., Bickel, P. J., Levina, E., and Zhu, J. (2008). Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515.
- Roverato, A. and Whittaker, J. (1998). The Isserlis matrix and its application to non-decomposable graphical Gaussian models. *Biometrika*, 85(3):711–725.

- Sadeghi, K. and Lauritzen, S. (2014). Markov properties for mixed graphs. *Bernoulli*, 20(2):676–696.
- Schwarz, G. (1978). Estimating the dimension of a model. *Ann. Statist.*, 6(2):461–464.
- Sonntag, D., Peña, J. M., and Gómez-Olmedo, M. (2015). Approximate counting of graphical models via MCMC revisited. *International Journal of Intelligent Systems*, 30(3):384–420.
- Speed, T. P. (1979). A note on nearest-neighbour Gibbs and Markov probabilities. *Sankhya A*, 41(3/4):184–197.
- Steinsky, B. (2004). Asymptotic behaviour of the number of labelled essential acyclic digraphs and labelled chain graphs. *Graphs and Combinatorics*, 20(3):399–411.
- Studený, M. (2005). *On Probabilistic Conditional Independence Structures*. Springer London.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Statist. Soc. B*, 58(1):267–288.
- Uhler, C. (2012). Geometry of maximum likelihood estimation in Gaussian graphical models. *Ann. Statist.*, 40(1):238–261.
- van de Geer, S. and Bühlmann, P. (2013). ℓ_0 -penalized maximum likelihood for sparse directed acyclic graphs. *Ann. Statist.*, 41(2):536–567.
- Verma, T. and Pearl, J. (1991). Equivalence and synthesis of causal models. In *Proc. of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, pages 255–270, Corvallis. AUAI Press.
- Werhli, A. V., Grzegorzcyk, M., and Husmeier, D. (2006). Comparative evaluation of reverse engineering gene regulatory networks with relevance networks, graphical Gaussian models and Bayesian networks. *Bioinformatics*, 22(20):2523–2531.
- Wermuth, N. (1976). Analogies between multiplicative models in contingency tables and covariance selection. *Biometrics*, 32(1):95–108.
- Wermuth, N. (1980). Linear recursive equations, covariance selection, and path analysis. *J. Am. Statist. Ass.*, 75(372):963–972.
- Wermuth, N. (2015). Graphical Markov models, unifying results and their interpretation. *Wiley StatsRef: Statistics Reference Online*, pages 1–29.
- Wermuth, N. and Lauritzen, S. L. (1983). Graphical and recursive models for contingency tables. *Biometrika*, 70(3):537–552.
- Wright, S. (1934). The method of path coefficients. *Ann. Math. Statist.*, 5(3):161–215.
- Yuan, M. and Lin, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35.
- Yule, G. U. (1907). On the theory of correlation for any number of variables, treated by a new system of notation. *Proc. R. Soc. A*, 79(529):182–193.
- Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *Journal of Machine Learning Research*, 7:2541–2563.