

On the Prior and Posterior Distributions Used in Graphical Modelling

Marco Scutari *

Abstract. Graphical model learning and inference are often performed using Bayesian techniques. In particular, learning is usually performed in two separate steps. First, the graph structure is learned from the data; then the parameters of the model are estimated conditional on that graph structure. While the probability distributions involved in this second step have been studied in depth, the ones used in the first step have not been explored in as much detail.

In this paper, we will study the prior and posterior distributions defined over the space of the graph structures for the purpose of learning the structure of a graphical model. In particular, we will provide a characterisation of the behaviour of those distributions as a function of the possible edges of the graph. We will then use the properties resulting from this characterisation to define measures of structural variability for both Bayesian and Markov networks, and we will point out some of their possible applications.

Keywords: Markov Networks, Bayesian Networks, Random Graphs, Structure Learning, Multivariate Discrete Distributions.

Graphical models (Pearl 1988; Lauritzen 1996) stand out among other classes of statistical models because of their use of graph structures in modelling and performing inference on multivariate, high-dimensional data. The close relationship between their probabilistic properties and the topology of the underlying graphs represents one of their key features, as it allows an intuitive understanding of otherwise complex models.

In a Bayesian setting, this duality leads naturally to splitting model estimation (which is usually called *learning*) in two separate steps (Cowell et al. 2007). In the first step, called *structure learning*, the graph structure \mathcal{G} of the model is estimated from the data. The presence (absence) of a particular edge between two nodes in \mathcal{G} implies the conditional (in)dependence of the variables corresponding to such nodes. In the second step, called *parameter learning*, the parameters Θ of the distribution assumed for the data are estimated conditional to the graph structure obtained in the first step. If we denote a graphical model with \mathcal{M} , so that $\mathcal{M} = (\mathcal{G}, \Theta)$, then we can write graphical model estimation from a data set \mathcal{D} as

$$P(\mathcal{M} | \mathcal{D}) = P(\mathcal{G} | \mathcal{D}) P(\Theta | \mathcal{G}, \mathcal{D}).$$

Furthermore, following Heckerman et al. (1995), we can rewrite structure learning as

$$P(\mathcal{G} | \mathcal{D}) \propto P(\mathcal{G}) P(\mathcal{D} | \mathcal{G}). \quad (1)$$

The prior distribution $P(\mathcal{G})$ and the corresponding posterior distribution $P(\mathcal{G} | \mathcal{D})$ are defined over the space of the possible graph structures, say \mathbf{G} . Since the dimension of \mathbf{G}

*Genetics Institute, University College London, United Kingdom, m.scutari@ucl.ac.uk

grows super-exponentially with the number of nodes in the graph (Harary and Palmer 1973), it is common practice to choose

$$P(\mathcal{G}) = \frac{1}{|\mathbf{G}|} \quad \text{for every } \mathcal{G} \in \mathbf{G} \quad (2)$$

as a non-informative prior, and then to search for the graph structure \mathcal{G} that maximises $P(\mathcal{G} | \mathcal{D})$. Unlike such a *maximum a posteriori* (MAP) approach, a full Bayesian analysis is computationally unfeasible in most real-world settings (Friedman et al. 1999a; Koller and Friedman 2009). Therefore, inference on most aspects of $P(\mathcal{G})$ and $P(\mathcal{G} | \mathcal{D})$ is severely limited by the nature of the graph space.

In this paper, we approach the analysis of those probability distributions from a different angle. We start from the consideration that, in a graphical model, the presence of particular edges and their layout are the most interesting features of the graph structure. Therefore, investigating $P(\mathcal{G})$ and $P(\mathcal{G} | \mathcal{D})$ through the probability distribution they induce over the set \mathcal{E} of their possible edges (identified by the set of unordered pairs of nodes in \mathcal{G}) provides a better basis from which to develop Bayesian inference on \mathcal{G} . This can be achieved by modelling \mathcal{E} as a multivariate discrete distribution encoding the joint state of the edges. Then, as far as inference on \mathcal{G} is concerned, we may rewrite Equation 1 as

$$P(\mathcal{G}(\mathcal{E}) | \mathcal{D}) \propto P(\mathcal{G}(\mathcal{E})) P(\mathcal{D} | \mathcal{G}(\mathcal{E})).$$

As a side effect, this shift in focus reduces the effective dimension of the sample space under consideration from super-exponential (the dimension of \mathbf{G}) to polynomial (the dimension of \mathcal{E}) in the number of nodes. The dimension of the parameter space for many inferential tasks, such as the variability measures studied in this paper, is likewise reduced.

The content of the paper is organised as follows. Basic definitions and notations are introduced in Section 1. The multivariate distributions used to model \mathcal{E} are described in Section 2. Some properties of the prior and posterior distributions on the graph space, $P(\mathcal{G}(\mathcal{E}))$ and $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$, are derived in Section 3. We will focus mainly on those properties related with the first and second order moments of the distribution of \mathcal{E} , and we will use them to characterise several measures of structural variability in Section 4. These measures may be useful for several inferential tasks for both Bayesian and Markov networks; some will be sketched in Section 4. Conclusions are summarised in Section 5, and proofs for the theorems in Sections 2 to 4 are reported in Appendix 5. Appendix 5 lists the exact values for some quantities of interest for $P(\mathcal{G}(\mathcal{E}))$, computed for several graph sizes.

1 Definitions and notations

Graphical models (Lauritzen 1996; Pearl 1988) are a class of statistical models which combine the rigour of a probabilistic approach with the intuitive representation of relationships given by graphs. They are composed by a set $\mathbf{X} = \{X_1, \dots, X_n\}$ of *random variables* describing the data \mathcal{D} and a *graph* $\mathcal{G} = (\mathbf{V}, E)$ in which each *vertex* or *node*

$v \in \mathbf{V}$ is associated with one of the random variables in \mathbf{X} . Nodes and the corresponding variables are usually referred to interchangeably. The *edges* $e \in E$ are used to express the dependence relationships among the variables in \mathbf{X} . Different classes of graphs express these relationships with different semantics, having in common the principle that graphical separation of two vertices implies the conditional independence of the corresponding random variables (Pearl 1988). The two examples most commonly found in literature are *Markov networks* (Whittaker 1990; Edwards 2000), which use undirected graphs (UGs, see Diestel 2005), and *Bayesian networks* (Neapolitan 2003; Korb and Nicholson 2010), which use directed acyclic graphs (DAGs, see Bang-Jensen and Gutin 2009). In the context of Bayesian networks, edges are often called *arcs* and denoted with $a \in A$; we will adopt this notation as well.

The structure of \mathcal{G} (that is, the pattern of the nodes and the edges) determines the probabilistic properties of a graphical model. The most important, and the most used, is the factorisation of the *global distribution* (the joint distribution of \mathbf{X}) into a set of lower-dimensional *local distributions*. In Markov networks, local distributions are associated with *cliques* (maximal subsets of nodes in which each element is adjacent to all the others); in Bayesian networks, each local distribution is associated with one node conditional on its *parents* (nodes linked by an incoming arc). In Markov networks the factorisation is unique; different graph structures correspond to different probability distributions. This is not so in Bayesian networks, where DAGs can be grouped into *equivalence classes* which are statistically indistinguishable. Each such class is uniquely identified by the underlying UG (i.e. in which arc directions are disregarded, also known as the *skeleton*) and by the set of *v-structures* (i.e. converging connections of the form $v_i \rightarrow v_j \leftarrow v_k$, $i \neq j \neq k$, in which v_i and v_k are not connected by an arc) common to all elements of the class. In addition, Bayesian networks can be represented by their *moral graphs* (Pearl 1988), the UGs obtained by “marrying” the parents in the v-structures and disregarding arc directions, to facilitate comparisons with Markov networks.

As for the global and the local distributions, there are many possible choices depending on the nature of the data and the aims of the analysis. However, the literature has focused mostly on two cases: the *discrete case* (Whittaker 1990; Heckerman et al. 1995), in which both the global and the local distributions are multinomial random variables, and the *continuous case* (Whittaker 1990; Geiger and Heckerman 1994), in which the global distribution is multivariate normal and the local distributions are univariate (in Bayesian networks) or multivariate (in Markov networks) normal random variables. In the former, the parameters of interest Θ are the *conditional probabilities* associated with each variable, usually represented as conditional probability tables. In the latter, the parameters of interest Θ are the *partial correlation coefficients* between each variable and its neighbours in \mathcal{G} . Conjugate distributions (Dirichlet and Wishart, respectively) are then used for learning and inference in a Bayesian setting.

2 Multivariate discrete distributions

The choice of an appropriate probability distribution for the set \mathcal{E} of the possible edges is crucial to make the derivation and the interpretation of the properties of \mathcal{E} and $\mathcal{G}(\mathcal{E})$ easier. We will first note that a graph is uniquely identified by its edge set E (or by its arc set A for a DAG), and that each edge e_{ij} or arc a_{ij} is uniquely identified by the nodes v_i and v_j , $i \neq j$ it is incident on. Therefore, if we model \mathcal{E} with a random variable we have that any edge set E (or arc set A) is just an element of its sample space; and since there is a one-to-one correspondence between graphs and edge sets, probabilistic properties and inferential results derived for traditional graph-centric approaches can easily be adapted to this new edge-centric approach and vice versa. In addition, if we denote $\mathcal{E} = \{(v_i, v_j), i \neq j\}$, we can clearly see that $|\mathcal{E}| = \mathcal{O}(|\mathbf{V}|^2)$. On the other hand, $|\mathbf{G}| = \mathcal{O}(2^{|\mathbf{V}|^2})$ for UGs and even larger for DAGs (Robinson 1973; Harary and Palmer 1973) and their equivalence classes (Gillispie and Perlman 2002).

We will also note that an edge or an arc has only few possible states:

- an edge can be either present ($e_{ij} \in E$) or missing from an UG ($e_{ij} \notin E$);
- in a DAG, an arc can be present in one of its two possible directions ($\overleftarrow{a}_{ij} \in A$ or $\overrightarrow{a}_{ij} \in A$) or missing from the graph ($\overleftarrow{a}_{ij} \notin A$ and $\overrightarrow{a}_{ij} \notin A$).

This leads naturally to the choice of a Bernoulli random variable for the former,

$$e_{ij} \sim E_{ij} = \begin{cases} 1 & e_{ij} \in E \text{ with probability } p_{ij} \\ 0 & e_{ij} \notin E \text{ with probability } 1 - p_{ij} \end{cases}, \quad (3)$$

and to the choice of a Trinomial random variable for the latter,

$$a_{ij} \sim A_{ij} = \begin{cases} -1 & \overleftarrow{a}_{ij} \in A \text{ with probability } \overleftarrow{p}_{ij} \\ 0 & \overleftarrow{a}_{ij}, \overrightarrow{a}_{ij} \notin A \text{ with probability } \overrightarrow{p}_{ij} \\ 1 & \overrightarrow{a}_{ij} \in A \text{ with probability } \overrightarrow{p}_{ij} \end{cases}, \quad (4)$$

where \overrightarrow{a}_{ij} is the arc $v_i \rightarrow v_j$ and \overleftarrow{a}_{ij} is the arc $v_j \rightarrow v_i$. Therefore, a graph structure can be modelled through its edge or arc set as follows:

- UGs, such as Markov networks or the skeleton and the moral graph of Bayesian networks, can be modelled by a *multivariate Bernoulli random variable*;
- directed graphs, such as the DAGs used in Bayesian networks, can be modelled by a *multivariate Trinomial random variable*.

In addition to being the natural choice for the respective classes of graphs, these distributions integrate smoothly with and extend other approaches present in the literature. For example, the probabilities associated with each edge or arc correspond to the *confidence coefficients* from Friedman et al. (1999a) and the *arc strengths* from Imoto et al. (2002). In a frequentist setting, they have been estimated using bootstrap resampling (Efron and Tibshirani 1993); in a Bayesian setting, Markov chain Monte Carlo (MCMC) approaches (Friedman and Koller 2003; Melançon and Fabrice 2004) have been used instead.

2.1 Multivariate Bernoulli

Let B_1, \dots, B_k , $k \in \mathbb{N}$ be Bernoulli random variables with marginal probabilities of success p_1, \dots, p_k , that is $B_i \sim Ber(p_i)$, $i = 1, \dots, k$. Then the distribution of the random vector $\mathbf{B} = [B_1, \dots, B_k]^T$ over the joint probability space of B_1, \dots, B_k is a *multivariate Bernoulli random variable* (Krummenauer 1998), denoted as $Ber_k(\mathbf{p})$. Its probability function is uniquely identified by the parameter collection

$$\mathbf{p} = \{p_I : I \subseteq \{1, \dots, k\}, I \neq \emptyset\},$$

which represents the *dependence structure* among the B_i in terms of simultaneous successes for every non-empty subset I of elements of \mathbf{B} . Other characterisations and fundamental properties of the multivariate Bernoulli distribution can be found in Johnson et al. (1997). Kocherlakota and Kocherlakota (1992) focus on the bivariate models specific to $Ber_2(\mathbf{p})$. Additional characterisations and results specific to particular applications can be found in George and McCulloch (1997, variable selection), Farrell and Rogers-Stewart (2008, longitudinal studies), Rubinstein (1999, combinatorial optimisation) and Agresti and Klingenberg (2005, clinical trials), among others.

From the literature we know that the expectation and the covariance matrix of \mathbf{B} are immediate extensions of the corresponding univariate Bernoulli ones;

$$\mathbf{E}(\mathbf{B}) = [p_1, \dots, p_k]^T \quad \text{and} \quad \text{COV}(\mathbf{B}) = [\sigma_{ij}] = p_{ij} - p_i p_j.$$

In particular, the covariance matrix $\Sigma = [\sigma_{ij}]$ has some interesting numerical properties. From basic probability theory, we know its diagonal elements σ_{ii} are bounded in the interval $[0, \frac{1}{4}]$; the maximum is attained for $p_i = \frac{1}{2}$, and the minimum for both $p_i = 0$ and $p_i = 1$. For the Cauchy-Schwarz theorem then $|\sigma_{ij}| \in [0, \frac{1}{4}]$. As a result, we can derive similar bounds for the eigenvalues $\lambda_1, \dots, \lambda_k$ of Σ , as shown in the following theorem.

Lemma 2.1. *Let $\mathbf{B} \sim Ber_k(\mathbf{p})$, and let Σ be its covariance matrix. Let λ_i , $i = 1, \dots, k$ be the eigenvalues of Σ . Then*

$$0 \leq \sum_{i=1}^k \lambda_i \leq \frac{k}{4} \quad \text{and} \quad 0 \leq \lambda_i \leq \frac{k}{4}.$$

Proof. See Appendix 5. □

These bounds define a closed convex set in \mathbb{R}^k , described by the family

$$\mathcal{L} = \left\{ \Delta^{k-1}(c) : c \in \left[0, \frac{k}{4}\right] \right\}$$

where $\Delta^{k-1}(c)$ is the non-standard $k - 1$ simplex

$$\Delta^{k-1}(c) = \left\{ (\lambda_1, \dots, \lambda_k) \in \mathbb{R}^k : \sum_{i=1}^k \lambda_i = c, \lambda_i \geq 0 \right\}. \tag{5}$$

2.2 Multivariate Trinomial

Construction and properties of the multivariate Trinomial random variable are similar to the ones illustrated in the previous section for the multivariate Bernoulli. For this reason, and because it is a particular case of the multivariate multinomial distribution, the multivariate Trinomial distribution is rarely the focus of research efforts in literature. Some of its fundamental properties are covered either in [Johnson et al. \(1997\)](#) or in monographs on contingency tables analysis such as [Bishop et al. \(2007\)](#).

Let T_1, \dots, T_k , $k \in \mathbb{N}$ be Trinomial random variables assuming values $\{-1, 0, 1\}$ and denoted as $T_i \sim \text{Tri}(p_{i(-1)}, p_{i(0)}, p_{i(1)})$ with $p_{i(-1)} + p_{i(0)} + p_{i(1)} = 1$. Then the distribution of the random vector $\mathbf{T} = [T_1, \dots, T_k]^T$ over the joint probability space of T_1, \dots, T_k is a *multivariate Trinomial random variable*, denoted as $\text{Tri}_k(\mathbf{p})$. The parameter collection \mathbf{p} which uniquely identifies the distribution is

$$\mathbf{p} = \left\{ p_{I(T)} : I \subseteq \{1, \dots, k\}, T \in \prod_{i=1}^{|I|} \{-1, 0, 1\}, I \neq \emptyset \right\}$$

where \times denotes the cartesian product. The reduced parameter collection we will need to study its first and second order moments is

$$\tilde{\mathbf{p}} = \{p_{ij(T)} : i, j = 1, \dots, k, T \in \{-1, 0, 1\}^2\}.$$

From the definition, we can easily derive the expected value and the variance of T_i ,

$$\begin{aligned} \mathbb{E}(T_i) &= p_{i(1)} - p_{i(-1)} \\ \text{VAR}(T_i) &= p_{i(1)} + p_{i(-1)} - [p_{i(1)} - p_{i(-1)}]^2 \end{aligned}$$

and the covariance between two variables T_i and T_j ,

$$\begin{aligned} \text{COV}(T_i, T_j) &= [p_{ij(1,1)} - p_{i(1)}p_{j(1)}] + [p_{ij(-1,-1)} - p_{i(-1)}p_{j(-1)}] \\ &\quad - [p_{ij(-1,1)} - p_{i(-1)}p_{j(1)}] - [p_{ij(1,-1)} - p_{i(1)}p_{j(-1)}]. \end{aligned}$$

Again, the diagonal elements of the covariance matrix Σ are bounded. This can be proved either by solving the constrained maximisation problem

$$\max_{p_{i(1)}, p_{i(-1)}} \text{VAR}(T_i) \quad \text{s.t.} \quad p_{i(1)} \geq 0, p_{i(-1)} \geq 0, p_{i(1)} + p_{i(-1)} \leq 1$$

or as an application of the following theorem by [Moors and Muilwijk \(1971\)](#).

Theorem 2.1. *If a discrete random variable X can take values only in the segment $[x_1, x_n]$ of the real axis, the maximum standard deviation of X equals $\frac{1}{2}(x_n - x_1)$. The maximum is reached if X takes the values x_1 and x_n with probabilities $\frac{1}{2}$ each.*

Proof. See [Moors and Muilwijk \(1971\)](#). □

In both cases we obtain that the maximum variance is achieved for $p_{i(1)} = p_{i(-1)} = \frac{1}{2}$ and is equal to 1, so $\sigma_{ii} \in [0, 1]$ and $|\sigma_{ij}| \in [0, 1]$. Furthermore, we can also prove that the eigenvalues of Σ are bounded using the same arguments as in Lemma 2.1.

Lemma 2.2. *Let $\mathbf{T} \sim Tri_k(\mathbf{p})$, and let Σ be its covariance matrix. Let $\lambda_i, i = 1, \dots, k$ be the eigenvalues of Σ . Then*

$$0 \leq \sum_{i=1}^k \lambda_i \leq k \qquad \text{and} \qquad 0 \leq \lambda_i \leq k.$$

Proof. See the proof of Lemma 2.1 in Appendix 5. □

These bounds define again a closed convex set in \mathbb{R}^k , described by the family

$$\mathcal{L} = \{ \Delta^{k-1}(c) : c \in [0, k] \},$$

where $\Delta^{k-1}(c)$ is the non-standard $k - 1$ simplex from Equation 5.

Another useful result, which we will use in Section 3.2 to link inference on UGs and DAGs, is introduced below.

Theorem 2.2. *Let $\mathbf{T} \sim Tri_k(\mathbf{p})$; then $|\mathbf{T}| = \mathbf{B} \sim Ber_k(\mathbf{p}^*)$ and $|T_i| = B_i \sim Ber(p_i^*)$.*

Proof. See Appendix 5. □

It follows that the variance of each T_i can be decomposed in two parts:

$$\text{VAR}(T_i) = \text{VAR}(B_i) + 4p_{i(1)}p_{i(-1)}. \tag{6}$$

The first is a function of the corresponding component $|T_i| = B_i$ of the transformed random vector, while the second depends only on the probabilities associated with -1 and 1 (which correspond to \vec{a}_{ij} and \vec{a}_{ij} in Equation 4).

3 Properties of $P(\mathcal{G}(\mathcal{E}))$ and $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$

The results derived in the previous section provide the foundation for characterising $P(\mathcal{G}(\mathcal{E}))$ and $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$. To this end, it is useful to distinguish three cases corresponding to different configurations of the probability mass among the graph structures $\mathcal{G}(\mathcal{E}) \in \mathbf{G}$:

- *minimum entropy*: the probability mass is concentrated on a single graph structure. This is the best possible configuration for $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$, because only one edge set E (or one arc set A) has a non-zero posterior probability. In other words, the data \mathcal{D} provide enough information to identify a single graph \mathcal{G} with posterior probability 1;
- *intermediate entropy*: several graph structures have non-zero probabilities. This is the case for informative priors $P(\mathcal{G}(\mathcal{E}))$ and for the posteriors $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$ resulting from real-world data sets;
- *maximum entropy*: all graph structures in \mathbf{G} have the same probability. This is the worst possible configuration for $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$, because it corresponds to the non-informative prior from Equation 2. In other words, the data \mathcal{D} do not provide any information useful in identifying a high-posterior graph \mathcal{G} .

Clearly, *minimum* and *maximum entropy* are limiting cases for $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$; the latter is non-informative about $\mathcal{G}(\mathcal{E})$, while the former identifies a single graph in \mathbf{G} . As we will show in Sections 3.1 (for UGs) and 3.2 (for DAGs), they provide useful reference points in determining which edges (or arcs) have significant posterior probabilities and in analysing the variability of the graph structure.

3.1 Undirected graphs

In the *minimum entropy* case, only one configuration of edges E has non-zero probability, which means that

$$p_{ij} = \begin{cases} 1 & \text{if } e_{ij} \in E \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \Sigma = \mathbf{O}$$

where \mathbf{O} is the zero matrix.

The uniform distribution over \mathbf{G} arising from the *maximum entropy* case has been studied extensively in random graph theory (Bollobás 2001); its two most relevant properties are that all edges e_{ij} are independent and have $p_{ij} = \frac{1}{2}$. As a result, $\Sigma = \frac{1}{4}I_k$; all edges display their maximum possible variability, which along with the fact that they are independent makes this distribution non-informative for \mathcal{E} as well as $\mathcal{G}(\mathcal{E})$.

The *intermediate entropy* case displays a middle-ground behaviour between the *minimum* and *maximum entropy* cases. The expected value and the covariance matrix of \mathcal{E} do not have a definite form beyond the bounds derived in Section 2.1. When considering posteriors arising from real-world data, we have in practice that most edges in \mathcal{E}

represent conditional dependence relationships that are completely unsupported by the data. This behaviour has been explained by Pearl (2009) with the tendency of “good” graphical models to represent the causal relationships underlying the data, which are typically sparse. As a result, we have that $\mathbb{E}(e_{ij}) = 0$ and $\text{VAR}(e_{ij}) = 0$ for many e_{ij} , so Σ is almost surely singular unless such edges are excluded from the analysis. Edges that appear with $p_{ij} \simeq \frac{1}{2}$ have about the same marginal probability and variance as in the *maximum entropy* case, so their marginal behaviour is very close to random noise. On the other hand, edges with probabilities near 0 or 1 can be considered to have a good support (against or in favour, respectively). As p_{ij} approaches 0 or 1, e_{ij} approaches its *minimum entropy*.

The closeness of a multivariate Bernoulli distribution to the *minimum* and *maximum entropy* cases can be represented in an intuitive way by considering the eigenvalues $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_k]^T$ of its covariance matrix Σ . Recall that the $\boldsymbol{\lambda}$ can assume values in the convex set \mathcal{L} defined in Equation 5, which corresponds to the region of the first orthant delimited by the non-standard simplex $\Delta^{k-1}(\frac{k}{4})$. In the *minimum entropy* case we have that $\Sigma = \mathbf{O}$, so $\lambda_1 = \dots = \lambda_k = 0$, and in the *maximum entropy case* $\Sigma = \frac{1}{4}I_k$, so $\lambda_1 = \dots = \lambda_k = \frac{1}{4}$; both points lie on the boundary of \mathcal{L} , the first in the origin and the second in the middle of $\Delta^{k-1}(\frac{k}{4})$. The distance between $\boldsymbol{\lambda}$ and these two points provides an intuitive way of measuring the variability of \mathcal{E} and, indirectly, the entropy of the corresponding probability distributions $\mathbb{P}(\mathcal{G}(\mathcal{E}) | \mathcal{D})$ and $\mathbb{P}(\mathcal{G}(\mathcal{E}))$. It is important to note, however, that different distributions over \mathbf{G} may have identical first and second order moments when modelled through \mathcal{E} . Such distributions will have the same $\boldsymbol{\lambda}$ and will therefore map to the same point in \mathcal{L} .

A simple example comprising three different distributions over a set of two edges is illustrated below.

Example 3.1. Consider three multivariate Bernoulli distributions $\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3$ over two edges (denoted with $e_1 \sim E_1$ and $e_2 \sim E_2$ for brevity) with covariance matrices

$$\Sigma_1 = \begin{bmatrix} 0.24 & 0.04 \\ 0.04 & 0.24 \end{bmatrix}, \quad \Sigma_2 = \begin{bmatrix} 0.1056 & -0.0336 \\ -0.0336 & 0.2016 \end{bmatrix}, \quad \Sigma_3 = \begin{bmatrix} 0.1056 & 0.1456 \\ 0.1456 & 0.2016 \end{bmatrix}$$

and eigenvalues

$$\boldsymbol{\lambda}_1 = \begin{bmatrix} 0.28 \\ 0.20 \end{bmatrix}, \quad \boldsymbol{\lambda}_2 = \begin{bmatrix} 0.2121 \\ 0.095 \end{bmatrix}, \quad \boldsymbol{\lambda}_3 = \begin{bmatrix} 0.3069 \\ 0.0003 \end{bmatrix}.$$

Their positions in \mathcal{L} are shown in Figure 1. \mathbf{B}_1 is the closest to $(\frac{1}{4}, \frac{1}{4})$, the point corresponding to the *maximum entropy* case, while \mathbf{B}_2 and \mathbf{B}_3 are farther from $(\frac{1}{4}, \frac{1}{4})$ than \mathbf{B}_1 due to the increasing correlation between e_1 and e_2 (which are independent in the *maximum entropy* case). The correlation coefficients for $\mathbf{B}_1, \mathbf{B}_2$ and \mathbf{B}_3 are $\text{COR}_{\mathbf{B}_1}(E_1, E_2) = 0.1666$, $\text{COR}_{\mathbf{B}_2}(E_1, E_2) = -0.2303$, $\text{COR}_{\mathbf{B}_3}(E_1, E_2) = 0.9978$, and they account for the increasing difference between the eigenvalues of each covariance matrix. In fact, Σ_3 is nearly singular because of the strong linear relationship between e_1 and e_2 , and it is therefore very close to one of the axes delimiting the first quadrant.




Figure 1: The covariance matrices Σ_1 , Σ_2 and Σ_3 from Example 3.1 represented as functions of their eigenvalues in the convex set \mathcal{L} . The points $(0, 0)$ and $(\frac{1}{4}, \frac{1}{4})$ correspond to the *minimum entropy* and *maximum entropy* cases.

If we denote with $E_{00} = \{\emptyset\}$, $E_{01} = \{e_2\}$, $E_{10} = \{e_1\}$, and $E_{11} = \{e_1, e_2\}$ all possible edge sets and with p_{00} , p_{01} , p_{10} and p_{11} the associated probabilities, for \mathbf{B}_1 we have

$$p_{00} = 0.2, \quad p_{01} = 0.2, \quad p_{10} = 0.2 \quad \text{and} \quad p_{11} = 0.4.$$

This is indeed close to a uniform distribution. The probability of both e_1 and e_2 is 0.6 and the variance is 0.24, which are again similar to the reference values for the maximum entropy case. On the other hand, for \mathbf{B}_2 we have

$$p_{00} = 0, \quad p_{01} = 0.12, \quad p_{10} = 0.28 \quad \text{and} \quad p_{11} = 0.6.$$

These probabilities are markedly different from a uniform distribution; the probabilities of e_1 and e_2 are respectively 0.88 and 0.72. Considering also the correlation between e_1 and e_2 , it is intuitively clear why Σ_2 is not as close as Σ_1 to $(\frac{1}{4}, \frac{1}{4})$. This is also true for \mathbf{B}_3 , which has the same marginal distributions as \mathbf{B}_2 but with a much stronger correlation.

3.2 Directed acyclic graphs

The behaviour of the multivariate Trinomial distribution in the *minimum* and *intermediate entropy* cases is similar to the one of the multivariate Bernoulli in many respects, but presents profound differences in the *maximum entropy* case. The reason for these differences is that the structure of a Bayesian network is assumed to be acyclic. Therefore, the state of each arc (i.e. whether it is present in the DAG and its direction) is influenced by the state of all other possible arcs even in the *maximum entropy* case, when otherwise they would be independent. Furthermore, the acyclicity constraint cannot be written in closed form, making the derivation of exact results on the moments of the distribution of \mathcal{E} particularly difficult.

To obtain some simple expressions for the expected value and the covariance matrix, we will first prove a simple theorem on DAGs, which essentially states that if we reverse the direction of every arc the resulting graph is still a DAG.

Theorem 3.1. *Let $G = (\mathbf{V}, A)$ be a DAG, and let $G^* = (\mathbf{V}, A^*)$ be another directed graph such that*

$$\overrightarrow{a_{ij}} \in A^* \iff \overleftarrow{a_{ij}} \in A \quad \text{and} \quad \overleftarrow{a_{ij}} \in A^* \iff \overrightarrow{a_{ij}} \in A$$

for every $a_{ij} \in A$. Then G^* is also acyclic.

Proof. See Appendix 5. □

An immediate consequence of this theorem is that for every DAG including the arc $\overrightarrow{a_{ij}}$ there exists another DAG including the arc $\overleftarrow{a_{ij}}$. Since all DAGs have the same probability in the *maximum entropy case*, this implies that both directions of every arc have the same probability,

$$\overrightarrow{p_{ij}} = \overleftarrow{p_{ij}} \quad \text{for every possible } a_{ij}, i \neq j. \tag{7}$$

Then the expected value of each marginal Trinomial distribution is equal to

$$E(A_{ij}) = \overrightarrow{p_{ij}} - \overleftarrow{p_{ij}} = 0$$

and its variance is equal to

$$\text{VAR}(A_{ij}) = \overrightarrow{p_{ij}} + \overleftarrow{p_{ij}} - (\overrightarrow{p_{ij}} - \overleftarrow{p_{ij}})^2 = 2\overrightarrow{p_{ij}}.$$

The joint probabilities associated with each pair of arcs are also symmetric in the maximum entropy case, again due to Theorem 3.1. Denote with a_{ij}° the event that arc a_{ij} is not present in the DAG. If we consider that both directions of every arc have the same probability and that there is no explicit ordering among the arcs, we have

$$\begin{aligned} P(\overrightarrow{a_{ij}}, \overrightarrow{a_{kl}}) &= P(\overleftarrow{a_{ij}}, \overleftarrow{a_{kl}}), & P(\overrightarrow{a_{ij}}, \overleftarrow{a_{kl}}) &= P(\overleftarrow{a_{ij}}, \overrightarrow{a_{kl}}), \\ P(a_{ij}^\circ, \overrightarrow{a_{kl}}) &= P(\overrightarrow{a_{ij}}, a_{kl}^\circ) = P(a_{ij}^\circ, \overleftarrow{a_{kl}}) = P(\overleftarrow{a_{ij}}, a_{kl}^\circ). \end{aligned} \tag{8}$$

Then the expression for the covariance simplifies to

$$\text{COV}(A_{ij}, A_{kl}) = 2 [P(\overrightarrow{a_{ij}}, \overrightarrow{a_{kl}}) - P(\overrightarrow{a_{ij}}, \overleftarrow{a_{kl}})],$$

which can be interpreted as the difference in probability between a *serial connection* (i.e. $v_i \rightarrow v_j \rightarrow v_l$, if $j = k$) and a *converging connection* (i.e. $v_i \rightarrow v_j \leftarrow v_l$) if the arcs are incident on a common node (Jensen and Nielsen 2007). This is interesting because v-structures are invariant within equivalence classes, while other patterns of arcs are not (Chickering 1995); indeed, equivalence classes are usually represented as *partially directed acyclic graphs* (PDAGs) in which only arcs belonging to v-structures are directed. All other arcs, with the exclusion of those which could introduce additional v-structures (known as *compelled arcs*), are replaced with the corresponding (undirected)




Figure 2: Exact (dashed line) and approximate (solid line) probabilities of an arc being present in a DAG with 3, 4, 5, 6, and 7 nodes.

edges. Therefore, the combination of high values of $|\text{COV}(A_{ij}, A_{kl})|$ and p_{ij}° is indicative of the belief that the corresponding arcs are directed in the PDAG identified by the equivalence class. Along with $\text{VAR}(A_{ij})$ and $\text{VAR}(A_{kl})$, it is also indicative of the stability of the graph structure, both in the arcs and their directions. In an uninformative prior, such as the distribution we are now considering in the *maximum entropy* case, we expect all covariances to be small; we will show this is the case in Theorem 3.4. On the other hand, in an informative distribution such as the ones considered in the *intermediate entropy* case, we expect covariances to be closer to their upper bounds for arcs that are compelled or part of a converging connection, and closer to zero for arcs whose direction is not determined in the equivalence class. Note that the sign of $\text{COV}(A_{ij}, A_{kl})$ depends on the way the two possible directions of each arc are associated with 1 and -1 ; a simple way to obtain a consistent parameterisation is to follow the natural ordering of the variables (i.e. if $i \leq j$ then the arc incident on these nodes is taken to be A_{ij} , $\overrightarrow{a_{ij}}$ is associated with 1 and $\overleftarrow{a_{ij}}$ with -1).

The equalities in Equations 7 and 8 drastically reduce the number of free parameters in the *maximum entropy* case. The marginal distribution of each arc now depends only on $\overrightarrow{p_{ij}}$, whose value can be derived from the following numerical approximation by Melançon et al. (2000).

Theorem 3.2. *The average number of arcs in a DAG with n nodes is approximately $\frac{1}{4}n^2$ in the maximum entropy case.*

Proof. See Melançon et al. (2000). □

Theorem 3.3. *Let $G = (\mathbf{V}, A)$ be a DAG with n nodes. Then for each possible arc*




Figure 3: Estimated (dashed line) and approximate (solid line) probabilities of an arc being present in a DAG with 8 to 50 nodes. The dotted line represents the limiting value in the number of nodes.

$a_{ij}, i \neq j$ we have that in the maximum entropy case

$$\overrightarrow{p}_{ij} = \overleftarrow{p}_{ij} \simeq \frac{1}{4} + \frac{1}{4(n-1)} \quad \text{and} \quad p_{ij}^{\circ} \simeq \frac{1}{2} - \frac{1}{2(n-1)}.$$

Proof. See Appendix 5. □

The quality of this approximation is examined in Figure 2 and Figure 3. In Figure 2, the values provided by Theorem 3.3 for DAGs with 3, 4, 5, 6 and 7 nodes are compared to the corresponding true values. The latter have been computed by enumerating all possible DAGs of that size (i.e. the whole population) and computing the relative frequency of each possible arc. In Figure 3, the values provided by Theorem 3.3 for DAGs with 8 to 50 nodes are compared with the corresponding estimated values computed over a set of 10^9 DAGs of the same size. The latter have been generated with uniform probability using the algorithm from Melançon and Fabrice (2004) as implemented in the bnlearn package (Scutari 2010, 2012) for R (R Development Core Team 2012).

We can clearly see that the approximate values are close to the corresponding true (in Figure 2) or estimated (in Figure 3) values for DAGs with at least 6 nodes. This is not a significant limitation; the true values can be easily computed via exhaustive enumeration for DAGs with 3, 4 and 5 nodes (they are reported in Appendix 5, along with other relevant quantities). Furthermore, it is evident both from Theorem 3.3 and from Figures 2 and 3 that, as the number of nodes diverges,

$$\lim_{n \rightarrow \infty} \overrightarrow{p}_{ij} = \lim_{n \rightarrow \infty} \overleftarrow{p}_{ij} = \frac{1}{4} \quad \text{and} \quad \lim_{n \rightarrow \infty} p_{ij}^{\circ} = \frac{1}{2}. \quad (9)$$

If we take the absolute value of this asymptotic Trinomial distribution, the resulting random variable is $Ber(p_{ij})$ with $p_{ij} = \frac{1}{2}$, which is the marginal distribution of an edge in an UG in the *maximum entropy* case. The absolute value transformation can be interpreted as ignoring the direction of the arc; the events $\overleftarrow{a}_{ij} \in A$ and $\overrightarrow{a}_{ij} \in A$ collapse into $e_{ij} \in E$, while $\overleftarrow{a}_{ij}, \overrightarrow{a}_{ij} \notin A$ maps to $e_{ij} \notin E$. As a result, the marginal distribution of an arc is remarkably similar to the one of the corresponding edge in an undirected graph for sufficiently large DAGs; in both cases, the nodes v_i and v_j are linked with probability $\frac{1}{2}$.

No result similar to Theorem 3.2 has been proved for arbitrary pairs of arcs in a directed acyclic graph; therefore, the structure of the covariance matrix Σ can be derived only in part. Variances can be approximated using the approximate probabilities from Theorem 3.3:

$$\text{VAR}(A_{ij}) = 2\overrightarrow{p}_{ij} \simeq \frac{1}{2} + \frac{1}{2(n-1)} \rightarrow \frac{1}{2} \text{ as } n \rightarrow \infty. \quad (10)$$

Therefore, maximum variance (of each arc) and maximum entropy (of the graph structure) are distinct, as opposed to what happens in UGs. However, we can use the decomposition of the variance introduced in Equation 6 to motivate why the *maximum entropy* case is still a “worst case” outcome for $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$. As we can see from Figure 4, the contributions of the presence of an arc (given by the transformation $|A_{ij}|$) and its direction (given by the $4\overrightarrow{p}_{ij}\overleftarrow{p}_{ij} = 4\overrightarrow{p}_{ij}^2$ term) to the variance are asymptotically equal. This is a consequence of the limits in Equation 9, which imply that an arc (modulo its direction) has the same probability to be present in or absent from the DAG and that its directions also have the same probability. As a result, we are not able to make any decision about either the presence of the arc or its direction. On the contrary, when $\text{VAR}(A_{ij})$ reaches its maximum at 1 we have that $P(\{\overrightarrow{a}_{ij}, \overleftarrow{a}_{ij}\}) = 1$ and $P(a_{ij}^o) = 0$, so we are sure that the arc will be present in the DAG in one of its two possible directions.

As for the covariances, it is possible to obtain tight bounds using *Hoeffding's identity* (Hoeffding 1940; Fisher and Sen 1994),

$$\text{COV}(X, Y) = \iint_{\mathbb{R}^2} F_{X,Y}(x, y) - F_X(x)F_Y(y) dx dy, \quad (11)$$

and the decomposition of the joint distribution of dependent random variables provided by the *Farlie-Morgenstern-Gumbel* (FMG) family of distributions (Mari and Kotz 2001), which has the form

$$F_{X,Y}(x, y) = F_X(x)F_Y(y) [1 + \varepsilon(1 - F_X(x))(1 - F_Y(y))], \quad |\varepsilon| \leq 1. \quad (12)$$

In Equations 11 and 12, $F_{X,Y}$, F_X and F_Y denote the cumulative distribution functions of the joint and marginal distributions of X and Y , respectively.

Theorem 3.4. *Let $G = (V, A)$ be a DAG, and let a_{ij} , $i \neq j$ and a_{kl} , $k \neq l$ be two possible arcs. Then in the maximum entropy case we have that*

$$|\text{COV}(A_{ij}, A_{kl})| \lesssim 4 \left[\frac{3}{4} - \frac{1}{4(n-1)} \right]^2 \left[\frac{1}{4} + \frac{1}{4(n-1)} \right]^2 \quad (13)$$




Figure 4: Decomposition of the asymptotic variance of an arc in the part that depends only on its presence (dashed line) and the part that depends only on its direction (solid line). The dots correspond to the respective values in the maximum entropy case.




Figure 5: Bounds for the absolute value of the covariance and the correlation coefficient of two arcs in a DAG with 6 to 50 nodes. The dotted lines represent the respective limiting values.




Figure 6: Approximate values for $P(\vec{a}_{ij}, \vec{a}_{kl})$ (solid line) and $P(\vec{a}_{ij}, \vec{a}_{kl})$ (dashed line) for DAGs with 8 to 50 nodes. The dotted line represents their asymptotic value.

and

$$|\text{COR}(A_{ij}, A_{kl})| \approx 2 \left[\frac{3}{4} - \frac{1}{4(n-1)} \right]^2 \left[\frac{1}{4} + \frac{1}{4(n-1)} \right]. \quad (14)$$

Proof. See Appendix 5. □

The bounds obtained from this theorem appear to be tight in the light of the true values for the covariance and correlation coefficients (computed again by enumerating all possible DAGs of size 3 to 7). Figure 5 shows that the bounds for DAGs with 6 to 50 nodes; for DAGs with 3, 4 and 5 nodes the approximation of \vec{p}_{ij} the bounds are based on is loose, and the true values of covariance and correlation are known. Non-null covariances range from ± 0.08 (for DAGs with 3 nodes) to ± 0.08410 (for DAGs with 7 nodes), while non-null correlation coefficients vary from ± 0.125 (for DAGs with 3 nodes) to ± 0.1423 (for DAGs with 7 nodes). Both covariance and correlation appear to be strictly increasing in modulus as the number of nodes increases, and converge to the limiting values of the bounds (0.140625 and 0.28125, respectively) from below.

Some other interesting properties are apparent from the true values of the covariance coefficients reported in Appendix 5. They are reported below as conjectures because, while they describe a systematic behaviour that emerges from the DAGs whose sizes we have a complete enumeration for, we were not able to substantiate them with formal proofs.

Conjecture 3.1. *Arcs that are not incident on a common node are uncorrelated.*

This is a consequence of the fact that if we consider A_{ij} and A_{kl} with $i \neq j \neq k \neq l$, we have $P(\vec{a}_{ij}, \vec{a}_{kl}) = P(\vec{a}_{ij}, \vec{a}_{kl})$. Therefore $\text{COV}(A_{ij}, A_{kl}) = 0$. This property seems to generalise to DAGs with more than 7 nodes. Figure 6 shows approximate estimates for

$P(\vec{a}_{ij}, \vec{a}_{kl})$ and $P(\vec{a}_{ij}, \overleftarrow{a}_{kl})$ for DAGs with 8 to 50 nodes, obtained again from 10^9 DAGs generated with uniform probability. The curves for the two probabilities are overlapping and very close to each other for all the considered DAG sizes, thus supporting Conjecture 3.1.

Conjecture 3.2. *The covariance matrix Σ is sparse.*

The proportion of arcs incident on a common node converges to zero as the number of nodes increases; therefore, if we assume Conjecture 3.1 is true, the proportion of elements of Σ that are equal to 0 has limit

$$1 \geq \lim_{n \rightarrow \infty} \frac{\binom{n}{2} \binom{n-2}{2}}{\binom{n}{2} \binom{n}{2} - \binom{n}{2}} \geq \lim_{n \rightarrow \infty} \frac{(n-2)(n-3)}{n(n-1)} = 1. \tag{15}$$

Furthermore, even arcs that are incident on a common node are not strongly correlated.

Conjecture 3.3. *Both covariance and correlation between two arcs incident on a common node are monotonically increasing in modulus.*

Conjecture 3.4. *The covariance between two arcs incident on a common node takes values in the interval $[0.08, 0.140625]$ in modulus, while the correlation takes values in $[0.125, 0.28125]$ in modulus.*

These intervals can be further reduced to $[0.08410, 0.140625]$ and $[0.1423, 0.28125]$ for DAGs larger than 7 nodes due to Conjecture 3.3.

As far as the other two cases are concerned, in the *minimum entropy* case we have that

$$E(A_{ij}) = \begin{cases} -1 & \text{if } \overleftarrow{a}_{ij} \in A \\ 0 & \text{if } \overleftarrow{a}_{ij}, \vec{a}_{ij} \notin A \\ 1 & \text{if } \vec{a}_{ij} \in A \end{cases} \quad \text{and} \quad \Sigma = \mathbf{O}$$

as in the *minimum entropy* case of UGs. The *intermediate entropy* case again ranges from being very close to the *minimum entropy* case (when the graph structure displays little variability) to being very close to the *maximum entropy* case (when the graph structure displays substantial variability). The bounds on the eigenvalues of Σ derived in Lemma 2.2 allow a graphical representation of the variability of the network structure, equivalent to the one illustrated in Example 3.1 for UGs.

4 Measures of variability

Several functions have been proposed in literature as univariate measures of spread of a multivariate distribution, usually under the assumption of multivariate normality; for some examples see Mardia et al. (1979) and Bilodeau and Brenner (1999). Three of them in particular can be used as descriptive statistics for the multivariate Bernoulli and Trinomial distributions: the *generalised variance*,

$$\text{VAR}_G(\Sigma) = \det(\Sigma);$$

the *total variance*,

$$\text{VAR}_T(\Sigma) = \text{tr}(\Sigma);$$

and the squared *Frobenius matrix norm* of the difference between Σ and a target matrix Ψ ,

$$\text{VAR}_F(\Sigma, \Psi) = \|\Sigma - \Psi\|_F^2.$$

Both generalised variance and total variance associate high values of the statistic to unstable network structures, and are bounded due to the properties of the multivariate Bernoulli and Trinomial distributions. For total variance, it is easy to show that either $\text{VAR}_T(\Sigma) \in [0, \frac{k}{4}]$ (for the multivariate Bernoulli) or $\text{VAR}_T(\Sigma) \in [0, k]$ (for the multivariate Trinomial), due to the bounds on the variances σ_{ii} and on the eigenvalues λ_i derived in Sections 2.1 and 2.2. Generalised variance is similarly bounded due to Hadamard's theorem on the determinant of a non-negative definite matrix (Seber 2008): $\text{VAR}_G(\Sigma) \in [0, (\frac{1}{4})^k]$ for the multivariate Bernoulli distribution and $\text{VAR}_G(\Sigma) \in [0, 1]$ for the multivariate Trinomial. They reach the respective maxima in the *maximum entropy* case and are equal to zero only in the *minimum entropy* case. Generalised variance is also strictly convex, but it is equal to zero when Σ is rank deficient. For this reason it may be convenient to reduce Σ to a smaller, full rank matrix (say Σ^*) and consider $\text{VAR}_G(\Sigma^*)$ instead of $\text{VAR}_G(\Sigma)$; using a regularised estimator for Σ such as the one presented in Ledoit and Wolf (2003) is also a viable option.

The behaviour of the squared Frobenius matrix norm, on the other hand, depends on the choice of the target matrix Ψ . For $\Psi = \mathbf{O}$ (the covariance matrix arising from the *minimum entropy* case for both the multivariate Bernoulli and the multivariate Trinomial), $\text{VAR}_F(\Sigma, \Psi)$ associates high values of the statistic to unstable network structures, like $\text{VAR}_T(\Sigma)$ and $\text{VAR}_G(\Sigma)$; however, $\text{VAR}_F(\Sigma, \mathbf{O})$ does not have a unique maximum and none of its maxima correspond to the *maximum entropy* case, making its interpretation unclear. A better choice seems to be a multiple of the covariance matrix arising from the *maximum entropy* case, say $\Psi = k\Sigma_{max}$, associating high values of $\text{VAR}_F(\Sigma, k\Sigma_{max})$ to stable network structures. For the multivariate Bernoulli, if we let $\Psi = \frac{k}{4}I_k$, $\text{VAR}_F(\Sigma, \frac{k}{4}I_k)$ can be rewritten as

$$\text{VAR}_F\left(\Sigma, \frac{k}{4}I_k\right) = \sum_{i=1}^k \left(\lambda_i - \frac{k}{4}\right)^2.$$

It has both a unique global minimum (because it is a convex function),

$$\min_{\mathcal{L}} \text{VAR}_F\left(\Sigma, \frac{k}{4}I_k\right) = \text{VAR}_F\left(\frac{1}{4}I_k\right) = \sum_{i=1}^k \left(\frac{1}{4} - \frac{k}{4}\right)^2 = \frac{k(k-1)^2}{16},$$

and a unique global maximum,

$$\max_{\mathcal{L}} \text{VAR}_F\left(\Sigma, \frac{k}{4}I_k\right) = \text{VAR}_F(\mathbf{O}) = \sum_{i=1}^k \left(\frac{k}{4}\right)^2 = \frac{k^3}{16},$$

which correspond to the *maximum* and *minimum entropy* covariance matrices, respectively. Similar results can be derived for the multivariate Trinomial distribution, using an approximate estimate for Σ_{max} based on the results presented in Section 3.2.

All the descriptive statistics introduced in this section can be normalised as follows:

$$\begin{aligned} \overline{\text{VAR}}_T(\Sigma) &= \frac{\text{VAR}_T(\Sigma)}{\max_{\Sigma} \text{VAR}_T(\Sigma)}, & \overline{\text{VAR}}_G(\Sigma) &= \frac{\text{VAR}_G(\Sigma)}{\max_{\Sigma} \text{VAR}_G(\Sigma)}, \\ \overline{\text{VAR}}_F(\Sigma, k\Sigma_{max}) &= \frac{\max_{\Sigma} \text{VAR}_F(\Sigma, k\Sigma_{max}) - \text{VAR}_F(\Sigma, k\Sigma_{max})}{\max_{\Sigma} \text{VAR}_F(\Sigma, k\Sigma_{max}) - \min_{\Sigma} \text{VAR}_F(\Sigma, k\Sigma_{max})}. \end{aligned} \tag{16}$$

These normalised statistics vary in the $[0, 1]$ interval and associate high values to graphs whose structures display a high variability. Since they vary on a known and bounded scale, they are easy to interpret as absolute quantities (i.e. goodness-of-fit statistics) as well as relative ones (i.e. proportions of total possible variability).

They also have a clear geometric interpretation as distances in \mathcal{L} , as they can all be rewritten as functions of the eigenvalues $\lambda_1, \dots, \lambda_k$. This allows, in turn, to provide an easy interpretation of otherwise complex properties of $P(\mathcal{G}(\mathcal{E}))$ and $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$ and to derive new results. First of all, the measures introduced in Equation 16 can be used to select the best learning algorithm \mathcal{A} in terms of structure stability for a given data set \mathcal{D} . Different algorithms make use of the information present in the data in different ways, under different sets of assumptions and with varying degrees of robustness. Therefore, in practice different algorithms learn different structures from the same data and, in turn, result in different posterior distributions on \mathbf{G} . If we rewrite Equation 1 to make this dependence explicit,

$$P(\mathcal{G}(\mathcal{E}) | \mathcal{D}, \mathcal{A}) \propto P(\mathcal{G}(\mathcal{E})) P(\mathcal{D} | \mathcal{G}(\mathcal{E}), \mathcal{A}),$$

and denote with $\Sigma_{\mathcal{A}}$ the covariance matrix of the distribution of the edges (or the arcs) induced by $P(\mathcal{G}(\mathcal{E}) | \mathcal{D}, \mathcal{A})$, then we can choose the optimal structure learning algorithm \mathcal{A}^* as

$$\mathcal{A}^* = \underset{\mathcal{A}}{\operatorname{argmin}} \overline{\text{VAR}}_T(\Sigma_{\mathcal{A}})$$

or, equivalently, using $\overline{\text{VAR}}_G(\Sigma_{\mathcal{A}})$ or $\overline{\text{VAR}}_F(\Sigma_{\mathcal{A}}, k\Sigma_{max})$ instead of $\overline{\text{VAR}}_T(\Sigma_{\mathcal{A}})$. Such an algorithm has the desirable property of maximising the information gain from the data, as measured by the distance from the non-informative prior $P(\mathcal{G}(\mathcal{E}))$ in \mathcal{L} . In other words, \mathcal{A}^* is the algorithm that uses the data in the most efficient way. Furthermore, an optimal \mathcal{A}^* can be identified even for data sets without a “golden standard” graph structure to use for comparison; this is not possible with the approaches commonly used in the literature, which rely on variations of the Hamming distance (Jungnickel 2008) and knowledge of such a “golden standard” to evaluate learning algorithms (see, for example Tsamardinos et al. 2006).

Similarly, it is possible to study the influence of different values of a tuning parameter for a given structure learning algorithm (and again a given data set). Such parameters include, for example, restrictions on the degrees of the nodes (Friedman et al. 1999b) and regularisation coefficients (Koller and Friedman 2009). If we denote these tuning parameters with τ , we can again choose an optimal τ^* as

$$\tau^* = \underset{\tau}{\operatorname{argmin}} \overline{\text{VAR}}_T(\Sigma_{\mathcal{A}(\tau)}).$$

Another natural application of the variability measures presented in Equation 16 is the study of the consistency of structure learning algorithms. It has been proved in the literature that most structure learning algorithms are increasingly able to identify a single, minimal graph structure as the sample size diverges (see, for example Chickering 2002). Therefore, $P(\mathcal{G}(\mathcal{E})|\mathcal{D})$ converges towards the *minimum entropy* case and all variability measures converge to zero. However, convergence speed has never been analysed and compared across different learning algorithms; any one of $\overline{\text{VAR}}_T(\Sigma_{\mathcal{A}})$, $\overline{\text{VAR}}_G(\Sigma_{\mathcal{A}})$ or $\overline{\text{VAR}}_F(\Sigma_{\mathcal{A}}, k\Sigma_{max})$ provides a coherent way to perform such an analysis.

Lastly, we may use the variability measures from Equation 16 as a basis to investigate different prior distributions for real-world data modelling and to define new ones. Relatively little attention has been paid in the literature to the choice of the prior over \mathbf{G} , and the uniform *maximum entropy* distribution is usually chosen for computational reasons. Its only parameter is the *imaginary sample size*, which expresses the weight assigned to the prior distribution as the size of an imaginary sample size supporting it (Heckerman et al. 1995).

However, choosing a uniform prior also has some drawbacks. Firstly, Steck and Jaakkola (2002) and Steck (2008) have shown that both large and small values of the imaginary sample size have counter-intuitive effects on the sparsity of a Bayesian network even for large sample sizes. For instance, large values of the imaginary sample size may favour the presence of an arc over its absence even when both $P(\mathcal{G}(\mathcal{E}))$ and \mathcal{D} imply the variables the arc is incident on are conditionally independent. Secondly, a uniform prior assigns a non-null probability to all possible models. Therefore, it often results in a very flat posterior which is not able to discriminate between networks that are well supported by the data and networks that are not (Koller and Friedman 2009).

Following Pearl (1988)'s suggestion that "good" graphical models should be sparse, sparsity-inducing priors such as the ones in Buntine (1991) and Friedman and Koller (2003) should be preferred to the *maximum entropy* distribution, as should informative priors (Mukherjee and Speed 2008). For example, the prior proposed in Buntine (1991) introduces a prior probability β to include (independently) each arc in a Bayesian network with a given topological ordering, which means $\overline{p}_{ij} = \beta$ and $\overline{p}_{ij} = 0$ for all $i < j$ in $P(\mathcal{G}(\mathcal{E}))$. Thus, $\text{VAR}(A_{ij}) = \beta - \beta^2$, $\text{VAR}_T(\Sigma) = k(\beta - \beta^2)$ and $\text{VAR}_G(\Sigma) = (\beta - \beta^2)^k$. The prior proposed in Friedman et al. (1999a), on the other hand, controls the number of parents of each node for a given topological ordering. Therefore, it favours low values of $P(\overline{a}_{ij}, \overline{a}_{jk})$ in $P(\mathcal{G}(\mathcal{E}))$ and again $\overline{p}_{ij} = 0$ for all $i < j$. Clearly, the amount of sparsity induced by the hyperparameters of these priors determines the variability of both the prior and the posterior, and can be controlled through the variability measures from Equation 16.

Furthermore, these measures can provide inspiration in devising new priors with the desired form and amount of sparsity. For instance, first order moments can be used to model prior knowledge on specific arcs or edges, and second order moments can be used to control their interplay. In the case of Bayesian networks, it may also be of interest to define a prior distribution assigning the same probability to equivalent models. Clearly, first and second order moments are insufficient for this task unless

all the higher order moments of $P(\mathcal{G}(\mathcal{E}))$ are completely determined by $P(A_{ij})$ and $P(A_{ij}, A_{kl})$. If such an assumption holds, at least approximately, then assigning equal probabilities to equivalent models entails $\overline{p_{ij}^{\rightarrow}} = \overline{p_{ij}^{\leftarrow}}$ for all the arcs that are undirected in the PDAG of the equivalence class. Implementing this condition is problematic for two reasons. First, $P(A_{ij}, A_{kl})$ needs to be specified conditional on the PDAG as well, as it must be consistent with $P(A_{ij})$; and the correct way to do that is not as intuitively clear as in the case of $P(A_{ij})$. Furthermore, identifying which arcs belong to a v-structure has a computational complexity of $\mathcal{O}(|A|^3)$, as it requires checking every triplet of arcs $\{a_{ij}, a_{jk}, a_{ik}\}$. In the case of dense graphs, $\mathcal{O}(|A|^3) \rightarrow \mathcal{O}(|\mathbf{V}|^6)$, and even for sparse graphs $\mathcal{O}(|A|^3) \simeq \mathcal{O}(|\mathbf{V}|^3)$. While better than the super-exponential complexity of dealing with the full parameterisation of $P(\mathcal{G}(\mathcal{E}))$, this is a serious limitation in high dimensional settings.

5 Conclusions

Bayesian inference on the structure of graphical models is challenging in most situations due to the difficulties in defining and analysing prior and posterior distributions over the spaces of undirected or directed acyclic graphs. The dimension of these spaces grows super-exponentially in the number of variables considered in the model, making even MAP analyses problematic.

In this paper, we propose an alternative approach to the analysis of graph structures which focuses on the set of possible edges \mathcal{E} of a graphical model $\mathcal{M} = (\mathcal{G}(\mathcal{E}), \Theta)$ instead of the possible graph structures themselves. The latter are uniquely identified by the respective edge sets; therefore, the proposed approach integrates smoothly with and extends both frequentist and Bayesian results present in the literature. Furthermore, this change in focus provides additional insights on the behaviour of individual edges (which are usually the focus of inference) and reduces the dimension of the sample space from super-exponential to quadratic in the number of variables.

For many inference problems the parameter space is reduced as well, and makes complex inferential tasks feasible. As an example, we characterise several measures of structural variability for both Bayesian and Markov networks using the second order moments of $P(\mathcal{G}(\mathcal{E}))$ and $P(\mathcal{G}(\mathcal{E}) | \mathcal{D})$. These measures have several possible applications and are easy to interpret from both an algebraic and a geometric point of view.

Acknowledgments

The author would like to thank to Adriana Brogini (University of Padova) and David Balding (University College London) for proofreading this article and providing many useful comments and suggestions. Furthermore, the author would also like to thank Giovanni Andreatta and Luigi Salce (University of Padova) for their assistance in the development of the material.

References

- Agresti, A. and Klingenberg, B. (2005). “Multivariate Tests Comparing Binomial Probabilities, with Application to Safety Studies for Drugs.” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(4): 691–706. 509
- Bang-Jensen, J. and Gutin, G. (2009). *Digraphs: Theory, Algorithms and Applications*. Springer-Verlag, 2nd edition. 507
- Bilodeau, M. and Brenner, D. (1999). *Theory of Multivariate Statistics*. Springer-Verlag. 521
- Bishop, Y. M. M., Fienberg, S. E., and Holland, P. W. (2007). *Discrete Multivariate Analysis: Theory and Practice*. Springer. 510
- Bollobás, B. (2001). *Random Graphs*. Cambridge University Press, 2nd edition. 512
- Buntine, W. (1991). “Theory Refinement on Bayesian Networks.” In *Proceedings of the 7th Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-91)*, 52–60. Morgan Kaufmann. 524
- Chickering, D. M. (1995). “A Transformational Characterization of Equivalent Bayesian Network Structures.” In Besnard, P. and Hanks, S. (eds.), *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, 87–98. Morgan Kaufmann. 515
- (2002). “Optimal Structure Identification with Greedy Search.” *Journal of Machine Learning Research*, 3: 507–554. 524
- Cowell, R. G., Dawid, P., Lauritzen, S. L., and Spiegelhalter, D. J. (2007). *Probabilistic Networks and Expert Systems*. Springer. 505
- Diestel, R. (2005). *Graph Theory*. Springer, 3rd edition. 507
- Edwards, D. I. (2000). *Introduction to Graphical Modelling*. Springer, 2nd edition. 507
- Efron, B. and Tibshirani, R. (1993). *An Introduction to the Bootstrap*. Chapman & Hall. 508
- Farrell, P. and Rogers-Stewart, K. (2008). “Methods for Generating Longitudinally Correlated Binary Data.” *International Statistical Review*, 76(1): 28–38. 509
- Fisher, N. I. and Sen, P. K. (1994). *The Collected Works of Wassily Hoeffding*. Springer-Verlag. 518
- Friedman, N., Goldszmidt, M., and Wyner, A. (1999a). “Data Analysis with Bayesian Networks: A Bootstrap Approach.” In Laskey, K. B. and Prade, H. (eds.), *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, 196–205. Morgan Kaufmann. 506, 508, 524
- Friedman, N. and Koller, D. (2003). “Being Bayesian about Bayesian Network Structure: A Bayesian Approach to Structure Discovery in Bayesian Networks.” *Machine Learning*, 50(1–2): 95–126. 508, 524

- Friedman, N., Pe'er, D., and Nachman, I. (1999b). "Learning Bayesian Network Structure from Massive Datasets: The 'Sparse Candidate' Algorithm." In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, 206–215. Morgan Kaufmann. 523
- Geiger, D. and Heckerman, D. (1994). "Learning Gaussian Networks." Technical Report MSR-TR-94-10, Microsoft Research, Redmond, Washington. 507
- George, E. I. and McCulloch, R. E. (1997). "Approaches for Bayesian Variable Selection." *Statistica Sinica*, 7: 339–373. 509
- 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. 508
- Harary, F. and Palmer, E. M. (1973). *Graphical Enumeration*. Academic Press. 506, 508
- Heckerman, D., Geiger, D., and Chickering, D. M. (1995). "Learning Bayesian Networks: The Combination of Knowledge and Statistical Data." *Machine Learning*, 20(3): 197–243. 505, 507, 524
- Hoeffding, W. (1940). "Masstabinvariante Korrelationstheorie." *Schriften des Mathematischen Instituts und des Instituts für Angewandte Mathematik der Universität Berlin*, 5(3): 179–223. 518
- Imoto, S., Kim, S. Y., Shimodaira, H., Aburatani, S., Tashiro, K., Kuhara, S., and Miyano, S. (2002). "Bootstrap Analysis of Gene Networks Based on Bayesian Networks and Nonparametric Regression." *Genome Informatics*, 13: 369–370. 508
- Jensen, F. V. and Nielsen, T. D. (2007). *Bayesian Networks and Decision Graphs*. Springer, 2nd edition. 515
- Johnson, N. L., Kotz, S., and Balakrishnan, N. (1997). *Discrete Multivariate Distributions*. Wiley. 509, 510
- Jungnickel, D. (2008). *Graphs, Networks and Algorithms*. Springer-Verlag, 3rd edition. 523
- Kocherlakota, S. and Kocherlakota, K. (1992). *Bivariate Discrete Distributions*. CRC Press. 509
- Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press. 506, 523, 524
- Korb, K. and Nicholson, A. (2010). *Bayesian Artificial Intelligence*. Chapman & Hall, 2nd edition. 507
- Krummenauer, F. (1998). "Limit Theorems for Multivariate Discrete Distributions." *Metrika*, 47(1): 47–69. 509

- Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press. 505, 506
- Ledoit, O. and Wolf, M. (2003). “Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection.” *Journal of Empirical Finance*, 10: 603–621. 522
- Mardia, K. V., Kent, J. T., and Bibby, J. M. (1979). *Multivariate Analysis*. Academic Press. 521
- Mari, D. D. and Kotz, S. (2001). *Correlation and Dependence*. Imperial College Press. 518
- Melançon, G., Dutour, I., and Bousquet-Mélou, M. (2000). “Random Generation of DAGs for Graph Drawing.” Technical Report INS-R0005, Centre for Mathematics and Computer Sciences, Amsterdam. 516
- Melançon, G. and Fabrice, P. (2004). “Generating Connected Acyclic Digraphs Uniformly at Random.” *Information Processing Letters*, 90(4): 209–213. 508, 517
- Moors, J. J. A. and Muilwijk, J. (1971). “An Inequality for the Variance of a Discrete Random Variable.” *Sankhyā: The Indian Journal of Statistics, Series B*, 33(3/4): 385–388. 510
- Mukherjee, S. and Speed, T. P. (2008). “Network Inference using Informative Priors.” *Proceedings of the National Academy of Sciences (PNAS)*, 105: 14313–14318. 524
- Neapolitan, R. E. (2003). *Learning Bayesian Networks*. Prentice Hall. 507
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann. 505, 506, 507, 524
- (2009). *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition. 513
- R Development Core Team (2012). *R: A Language and Environment for Statistical Computing*.
URL <http://www.R-project.org> 517
- Robinson, R. W. (1973). “Counting Labeled Acyclic Digraphs.” In *New Directions in the Theory of Graphs: Proceedings of the 3rd Ann Arbor Conference on Graph Theory*, 239–273. Academic Press. 508
- Rubinstein, R. Y. (1999). “The Cross-Entropy Method for Combinatorial and Continuous Optimization.” *Methodology and Computing in Applied Probability*, 1: 127–190. 509
- Scutari, M. (2010). “Learning Bayesian Networks with the bnlearn R Package.” *Journal of Statistical Software*, 35(3): 1–22. 517
- (2012). *bnlearn: Bayesian Network Structure Learning*. R package version 3.2.
URL <http://www.bnlearn.com/> 517

- Seber, G. A. F. (2008). *A Matrix Handbook for Statisticians*. Wiley. 522
- Steck, H. (2008). “Learning the Bayesian Network Structure: Dirichlet Prior versus Data.” In *Proceedings of the 24th Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, 511–518. AUAI Press. 524
- Steck, H. and Jaakkola, T. (2002). “On the Dirichlet Prior and Bayesian Regularization.” In *Advances in Neural Information Processing Systems (NIPS)*, 697–704. MIT Press. 524
- Tsamardinos, I., Brown, L. E., and Aliferis, C. F. (2006). “The Max-Min Hill-Climbing Bayesian Network Structure Learning Algorithm.” *Machine Learning*, 65(1): 31–78. 523
- Whittaker, J. (1990). *Graphical Models in Applied Multivariate Statistics*. Wiley. 507

Appendix I: Proofs

Proof of Lemma 2.1. Since Σ is a real, symmetric, non-negative definite matrix, its eigenvalues λ_i are non-negative real numbers; this proves the lower bound in both inequalities.

The upper bound in the first inequality holds because

$$\sum_{i=1}^k \lambda_i = \sum_{i=1}^k \sigma_{ii} \leq \max_{\{\sigma_{ii}\}} \sum_{i=1}^k \sigma_{ii} = \sum_{i=1}^k \max \sigma_{ii} = \frac{k}{4},$$

as the sum of the eigenvalues is equal to the trace of Σ . This in turn implies

$$\lambda_i \leq \sum_{i=1}^k \lambda_i \leq \frac{k}{4},$$

which completes the proof. \square

Proof of Theorem 2.2. It is easy to show that each $|T_i| = B_i$, with $p_{i(1)} + p_{i(-1)} = p_i^*$ and $p_{i(0)} = 1 - p_i^*$. It follows that the parameter collection \mathbf{p} of \mathbf{T} reduces to

$$\begin{aligned} \mathbf{p}^* &= \left\{ p_{I(T)} : I \subseteq \{1, \dots, k\}, T \in \{0, 1\}^{|I|}, I \neq \emptyset \right\} \\ &= \{p_I : I \subseteq \{1, \dots, k\}, I \neq \emptyset\} \end{aligned}$$

after the transformation. Therefore, $|\mathbf{T}| \sim \text{Ber}_k(\mathbf{p}^*)$ is a uniquely identified multivariate Bernoulli random variable according to the definition introduced at the beginning of Section 2.1. \square

Proof of Theorem 3.1. Let's assume by contradiction that G^* is cyclic; this implies that there are one or more nodes $v_i \in \mathbf{V}$ such that

$$v_i \xrightarrow{\overrightarrow{a_{ij}}} v_j \rightarrow \dots \rightarrow v_k \xrightarrow{\overrightarrow{a_{ki}}} v_i$$

for some $v_j, v_k \in \mathbf{V}$. However, this would mean that in G we would have

$$v_i \xrightarrow{\overleftarrow{a_{ki}}} v_k \rightarrow \dots \rightarrow v_j \xrightarrow{\overleftarrow{a_{ij}}} v_i$$

which is not possible since G is assumed to be acyclic. \square

Proof of Theorem 3.3. Each possible arc can appear in the graph in only one direction at a time, so a directed acyclic graph with n nodes can have at most $\binom{n}{2} = \frac{1}{2}n(n-1)$ arcs. Therefore

$$\overrightarrow{p_{ij}} + \overleftarrow{p_{ij}} \simeq \frac{\frac{1}{4}n^2}{\frac{1}{2}n(n-1)} = \frac{1}{2} + \frac{1}{2(n-1)}.$$

But in the *maximum entropy* case we also have that $\overrightarrow{p_{ij}} = \overleftarrow{p_{ij}}$, so

$$\overrightarrow{p_{ij}} = \overleftarrow{p_{ij}} \simeq \frac{1}{4} + \frac{1}{4(n-1)} \quad \text{and} \quad p_{ij}^\circ = 1 - 2\overrightarrow{p_{ij}} \simeq \frac{1}{2} - \frac{1}{2(n-1)},$$

which completes the proof. \square

Proof of Theorem 3.4. In the maximum entropy case, all arcs have the same marginal distribution function,

$$F_A(a_{ij}) \simeq \begin{cases} 0 & \text{in } (-\infty, -1] \\ \frac{1}{4} + \frac{1}{4(n-1)} & \text{in } (-1, 0] \\ \frac{3}{4} - \frac{1}{4(n-1)} & \text{in } (0, 1] \\ 1 & \text{in } (1, +\infty) \end{cases}, \quad (17)$$

so the joint distribution of any pair of arcs a_{ij} and a_{kl} can be written as a member of the Farlie-Morgenstern-Gumbel family of distributions as

$$F_{A_{ij}, A_{kl}}(a_{ij}, a_{kl}) = F_A(a_{ij})F_A(a_{kl})[1 + \varepsilon(1 - F_A(a_{ij}))(1 - F_A(a_{kl}))]. \quad (18)$$

Then if we apply Hoeffding's identity from Equation 11 and replace the joint distribution

function $F_{A_{ij}, A_{kl}}(a_{ij}, a_{kl})$ with the right hand of Equation 18 we have that

$$\begin{aligned}
 & |\text{COV}(A_{ij}, A_{kl})| = \\
 & = \left| \sum_{a_{ij} \in \{-1, 0, 1\}} \sum_{a_{kl} \in \{-1, 0, 1\}} F_{A_{ij}, A_{kl}}(a_{ij}, a_{kl}) - F_A(a_{ij})F_A(a_{kl}) \right| \\
 & \leq \sum_{a_{ij} \in \{-1, 0, 1\}} \sum_{a_{kl} \in \{-1, 0, 1\}} |F_{A_{ij}, A_{kl}}(a_{ij}, a_{kl}) - F_A(a_{ij})F_A(a_{kl})| \\
 & = \sum_{a_{ij} \in \{-1, 0, 1\}} \sum_{a_{kl} \in \{-1, 0, 1\}} |F_A(a_{ij})F_A(a_{kl}) \\
 & \quad \times [1 + \varepsilon(1 - F_A(a_{ij}))(1 - F_A(a_{kl}))] - F_A(a_{ij})F_A(a_{kl})| \\
 & = \sum_{\{-1, 0\}} \sum_{\{-1, 0\}} (1 - F_A(a_{ij}))(1 - F_A(a_{kl})).
 \end{aligned}$$

We can now compute the bounds for $|\text{COV}(a_{ij}, a_{kl})|$ and $|\text{COR}(a_{ij}, a_{kl})|$ using only the marginal distribution function F_A from Equation 17 and the variance from Equation 10, thus obtaining the expressions in Equation 13 and Equation 14. \square

Appendix II: Moments and parameters of the multivariate Trinomial distribution in the maximum entropy case

Below are reported the exact values of the parameters of the marginal Trinomial distributions and of the first and second order moments of the multivariate Trinomial distribution in the maximum entropy case. All these quantities have been computed by a complete enumeration of the directed acyclic graphs of a given size (3, 4, 5, 6 and 7).

5.1 Moments for the 3-dimensional distribution

$$A_{ij} = \begin{cases} -1 & \text{with probability 0.32} \\ 0 & \text{with probability 0.36} \\ 1 & \text{with probability 0.32} \end{cases} \quad \begin{aligned} E(A_{ij}) &= 0 \\ \text{VAR}(A_{ij}) &= 0.64 \\ |\text{COV}(A_{ij}, A_{kl})| &= 0.08 \end{aligned}$$

5.2 Moments for the 4-dimensional distribution

$$A_{ij} = \begin{cases} -1 & \text{with probability 0.309392} \\ 0 & \text{with probability 0.381215} \\ 1 & \text{with probability 0.309392} \end{cases} \quad \begin{aligned} E(A_{ij}) &= 0 \\ \text{VAR}(A_{ij}) &= 0.618784 \end{aligned}$$

$$|\text{COV}(A_{ij}, A_{kl})| = \begin{cases} 0 & \text{if } i \neq j \neq k \neq l \\ 0.081031 & \text{otherwise} \end{cases}$$

5.3 Moments for the 5-dimensional distribution

$$A_{ij} = \begin{cases} -1 & \text{with probability } 0.301082 \\ 0 & \text{with probability } 0.397834 \\ 1 & \text{with probability } 0.301082 \end{cases} \quad \begin{aligned} E(A_{ij}) &= 0 \\ \text{VAR}(A_{ij}) &= 0.602165 \end{aligned}$$

$$|\text{COV}(A_{ij}, A_{kl})| = \begin{cases} 0 & \text{if } i \neq j \neq k \neq l \\ 0.081691 & \text{otherwise} \end{cases}$$

5.4 Moments for the 6-dimensional distribution

$$A_{ij} = \begin{cases} -1 & \text{with probability } 0.294562 \\ 0 & \text{with probability } 0.410875 \\ 1 & \text{with probability } 0.294562 \end{cases} \quad \begin{aligned} E(A_{ij}) &= 0 \\ \text{VAR}(A_{ij}) &= 0.589124 \end{aligned}$$

$$|\text{COV}(A_{ij}, A_{kl})| = \begin{cases} 0 & \text{if } i \neq j \neq k \neq l \\ 0.082121 & \text{otherwise} \end{cases}$$

5.5 Moments for the 7-dimensional distribution

$$A_{ij} = \begin{cases} -1 & \text{with probability } 0.289390 \\ 0 & \text{with probability } 0.421220 \\ 1 & \text{with probability } 0.289390 \end{cases} \quad \begin{aligned} E(A_{ij}) &= 0 \\ \text{VAR}(A_{ij}) &= 0.578780 \end{aligned}$$

$$|\text{COV}(A_{ij}, A_{kl})| = \begin{cases} 0 & \text{if } i \neq j \neq k \neq l \\ 0.82410 & \text{otherwise} \end{cases}$$