# Consistent model selection of discrete Bayesian networks from incomplete data 

Nikolay Balov<br>Department of Biostatistics and Computational Biology University of Rochester Medical Center<br>601 Elmwood Ave<br>Rochester, NY-14642<br>e-mail: nikolay_balov@urmc.rochester.edu


#### Abstract

A maximum likelihood based model selection of discrete Bayesian networks is considered. The structure learning is performed by employing a scoring function $S$, which, for a given network $G$ and $n$-sample $D_{n}$, is defined as the maximum marginal log-likelihood $l$ minus a penalization term $\lambda_{n} h$ proportional to network complexity $h(G)$,


$$
S\left(G \mid D_{n}\right)=l\left(G \mid D_{n}\right)-\lambda_{n} h(G)
$$

An available case analysis is developed with the standard log-likelihood replaced by the sum of sample average node log-likelihoods. The approach utilizes partially missing data records and allows for comparison of models fitted to different samples.

In missing completely at random settings the estimation is shown to be consistent if and only if the sequence $\lambda_{n}$ converges to zero at a slower than $n^{-1 / 2}$ rate. In particular, the BIC model selection $\left(\lambda_{n}=0.5 \log (n) / n\right)$ applied to the node-average log-likelihood is shown to be inconsistent in general. This is in contrast to the complete data case when BIC is known to be consistent. The conclusions are confirmed by numerical experiments.

AMS 2000 subject classifications: Primary 62F12; secondary 62H12.
Keywords and phrases: Bayesian networks, categorical data, model selection, penalized maximum likelihood, missing completely at random.

Received December 2012.

## Contents

1 Introduction . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 1048
2 Problem formulation and motivation . . . . . . . . . . . . . . . . . . 1050
2.1 Basic definitions . . . . . . . . . . . . . . . . . . . . . . . . . . 1050
2.2 Node-average log-likelihood . . . . . . . . . . . . . . . . . . . . 1053
2.3 Relation between NAL maximization and EM algorithm . . . . 1054

3 MLE and model selection . . . . . . . . . . . . . . . . . . . . . . . . 1055
3.1 Identifiability . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 1057
3.2 NAL-based scoring functions . . . . . . . . . . . . . . . . . . . . . 1058

4 Estimation consistency . . . . . . . . . . . . . . . . . . . . . . . . . . 1060
5 Numerical experiments . . . . . . . . . . . . . . . . . . . . . . . . . . 1063
5.1 Simulated 2-node network . . . . . . . . . . . . . . . . . . . . . 1063
5.2 Consistent estimation of the ALARM network . . . . . . . . . . 1064

6 Conclusion . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 1067
7 Proofs . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 1067
Acknowledgements . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 1076
References . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 1076

## 1. Introduction

The continuing interest in developing sparse statistical models, with the notable presence of Bayesian networks among them, is well motivated by a number of pressing practical problems coming from gene/protein expression analysis and medical imaging, to mention a few. Although graphical probability models based on directed connections between random variables provide efficient joint distribution description, the application of such models is often limited by the ambiguity of their observed behavior which makes the learning rather difficult.

One of the prevailing approaches to graphical model selection is through optimization of some scoring functions. In the context of Bayesian networks, the usual choice is the log of posterior. Let $(G, \theta)$ be a Bayesian network with graph structure $G$ and probability model parameter $\theta \in \Theta$. Following the Bayesian paradigm (see for example [5] and [11]), one specifies prior probability distributions $\pi$ for $G$ and $\theta$. Then, for a sample $D_{n}$ of size $n$, one considers the Bayesian scoring function

$$
S\left(G \mid D_{n}\right)=\log \pi(G)+\log \mathcal{L}\left(G \mid D_{n}\right)
$$

where

$$
\mathcal{L}\left(G \mid D_{n}\right)=\int_{\theta \in \Theta} \mathcal{L}\left(G, \theta \mid D_{n}\right) \pi(\theta) d \theta
$$

is the so-called marginal likelihood of $G$, while $\mathcal{L}\left(G, \theta \mid D_{n}\right)$ is the usual likelihood of $(G, \theta)$. The Bayesian scoring function measures the posterior certainty under the chosen prior system and the model with maximum score is thus a natural estimator.

The main virtue of the Bayesian approach is in counter-balancing the tendency of the maximum likelihood estimation to choose the most complex model fitting the data. As first noticed by [10], when the probability parameter space $\Theta$ constitutes an exponential family in an Euclidean space, the marginal loglikelihood of a model $M$ admits the approximation

$$
\log \mathcal{L}\left(M \mid D_{n}\right)=\operatorname{BIC}\left(M \mid D_{n}\right)+O_{p}(1)
$$

based on the so-called Bayesian Information Criterion (BIC),

$$
\operatorname{BIC}\left(M \mid D_{n}\right) \equiv \log \mathcal{L}\left(M, \hat{\theta}_{M} \mid D_{n}\right)-0.5 \log (n) \operatorname{dim}(M)
$$

where $\hat{\theta}_{M}$ is the value of $\theta$ that maximizes the log-likelihood for given $M$ and $D_{n}$, and $\operatorname{dim}(M)$ is the dimension of $M$. The immediate application of this result to discrete and conditional Gaussian Bayesian networks was postponed because of the non-Euclidean structure of the parameter space for these models.

This obstacle was later overcome by [7], who showed the validity of the BIC approximation for a much large family of curved exponential distributions.

In a later work, [6] applied this result to several families of Bayesian network models including the discrete ones, thus showing the asymptotic consistency of BIC. In its generality, the parameter space $\Theta_{G}$ of a discrete Bayesian network $G$ comprises a collection of multinomial distributions and the total number of parameters needed to specify them is what is understood as dimension of $\Theta_{G}$. The BIC approximation is then expressed as

$$
\begin{equation*}
\log \mathcal{L}\left(G \mid D_{n}\right)=\log \mathcal{L}\left(G, \hat{\theta}_{G} \mid D_{n}\right)-0.5 \log (n) \operatorname{dim}\left(\Theta_{G}\right)+O_{p}(1) \tag{1.1}
\end{equation*}
$$

Equation (1.1) suggests a more direct estimating procedure - selecting a model $G$ in $\mathcal{G}$ with maximal BIC score. There are two typical arguments in favor of this route versus the Bayesian one. The first one is methodological prior based inference is not universally accepted. The other one is computational - calculating marginal likelihoods can be prohibitive, especially so in the framework of large dimensional graphical models.

These observations have motivated us to pursue the latter, non-Bayesian approach - maximum likelihood estimation followed by model selection according to some scoring criteria. To generalize it, we reformulate the right-hand side of (1.1) and consider the following estimation problem

$$
\begin{equation*}
\hat{G}=\arg \max _{G \in \mathcal{G}}\left\{n^{-1} \log \mathcal{L}\left(G, \hat{\theta}_{G} \mid D_{n}\right)-\lambda_{n} h(G)\right\} \tag{1.2}
\end{equation*}
$$

where $\lambda_{n}$ is some positive sequence and $h$ is a function measuring the complexity of $G$. The class of problems (1.2) is known as extended (or penalized) likelihood approach [3]. Typical penalization parameters are $\lambda_{n}=0.5 n^{-1} \log (n)$ (BIC) and $\lambda_{n}=n^{-1}$ (AIC), while $\operatorname{dim}\left(\Theta_{G}\right)$ is a usual choice for $h$. We briefly remark that, in order to be useful in practice, the estimation problem (1.2) relies on two assumptions: (1) for a fixed $G$, the MLE $\hat{\theta}_{G}$ can be easily found, and (2), the set of networks $\mathcal{G}$ is not prohibitively large, which usually requires imposing some network structure restrictions. In this paper however, we are mainly concerned with the theoretical aspects of (1.2) - to our knowledge, the consistency properties of $\hat{G}$ are not investigated in presence of missing values - and present results which are relevant to all estimation algorithms involving penalized log-likelihood of this form.

The paper contributes in three main directions. First, in order to more efficiently handle data with incomplete records, we modify the scoring based model selection (1.2) by replacing the log-likelihood function with what we tentatively call node-average log-likelihood (NAL) - a sum of sample average node loglikelihoods relative to the node parents. The NAL statistics utilizes partially incomplete sample records instead of discarding them and provides means for comparing models fitted to different samples. We argue that when the number of nodes is large in comparison to the parent sizes, the NAL-based estimation achieves efficiency close to that of the computationally more demanding Expectation Maximization (EM) procedure [8]. Second, we focus on missing
completely at random data models for they essentially guarantee network identifiability. More general missing at random mechanisms, in most cases, obscure the underlying network structure and render the network unidentifiable. Third, we generalize the scoring criteria by allowing the complexity measure $h$ to be any positive function (as long as it is increasing for $G$ as defined later) and a continuum of penalization parameters $\lambda_{n}=O\left(n^{-\alpha}\right)$ by specifying a range of possible values $\alpha$ for which the estimation is consistent.

In Section 2 we introduce the notion of node-average log-likelihood and describe the model selection problem in the context of Bayesian networks. Then, Section 3, we consider the question of network identifiability and formulate consistency in terms of scoring criteria. For the latter we follow [7] and [4]. We show in Section 3.1 that if the data is missing completely at random, the identifiability arises under some natural conditions. Section 4 presents the main result in this paper, Theorem 4.1, claiming that the estimation is asymptotically consistent provided that $\lambda_{n}$ goes to zero at slower rate than $n^{-1}$, in the complete data case, and $n^{-1 / 2}$, in presence of missing data. We also show the necessity of the later in missing completely at random settings. Thus, the inconsistency of AIC is (re)confirmed along with somewhat unexpected conclusion regarding the BIC criteria - in the context of NAL optimization, BIC is consistent when applied to complete data but inconsistent otherwise. In Section 5 we present some numerical results in confirmation of the theory which are carried out with the catnet package for $\mathbf{R}$. We conclude with a short discussion on possible extensions of the presented approach beyond the class of discrete Bayesian networks.

## 2. Problem formulation and motivation

### 2.1. Basic definitions

Let $\mathbf{X}=\left(X_{i}\right)_{i=1}^{N}$ be a $N$-vector of discrete random variables. Any directed acyclic graph (DAG) $G$ with nodes $\mathbf{X}$ is a collection of directed edges from parent to child nodes such that there are no cycles. We denote with $P a_{i}$ the parents of node $X_{i}$ in $G$; then $G$ is completely described by the parent sets $\left\{P a_{i}\right\}_{i=1}^{N}$. The set of all DAGs with nodes $\mathbf{X}$ admits partial ordering. We say that $G_{1}$ is included in $G_{2}$ and write $G_{1} \subseteq G_{2}$ if all directed edges of $G_{1}$ are present in $G_{2}$ as well. An element $G$ of a set of DAGs $\mathcal{G}$ is called minimal if there is no $\tilde{G} \in \mathcal{G}$ such that $\tilde{G} \subset G$; similarly defined are maximal DAGs. In a set of nested DAGs, the minimum and maximum DAG are always uniquely defined.

Discrete Bayesian network (DBN) on $\mathbf{X}$ is any pair $(G, P)$ consisting of DAG $G$ and probability distribution $P$ on $\mathbf{X}$ subject to two conditions:
(1) the joint distribution of $\mathbf{X}$ given by $P$ satisfies the so-called local Markov property (LMP) with respect to $G$ - any node-variable is independent of its non-descendants given its parents,
(2) $G$ is a minimal DAG compatible with $P$, that is, there is no $\tilde{G} \subset G$ such that $P$ satisfies LMP with respect to $\tilde{G}$.

For any DAG $G$, there is an order of its nodes, called causality order, such that the parents of each one appear earlier in that order. We say that $G$ is compatible with an order $\Omega$ if $P a_{\Omega(1)}=\emptyset$ and for all $i=2, \ldots, N, P a_{\Omega(i)} \subset$ $\left\{X_{\Omega(1)}, \ldots, X_{\Omega(i-1)}\right\}$. For $i<j$, we denote with $X_{\Omega(i)} \prec X_{\Omega(j)}$ the fact that $X_{\Omega(i)}$ appears before $X_{\Omega(j)}$ in the order $\Omega$.

In its generality, the discreteness of our model implies that for each state $x_{P a_{i}}$ of the parents of $X_{i}$, the probability distribution of $X_{i}$ conditional on $P a_{i}=x_{P a_{i}}$ is multinomial. Moreover, the conditional probability tables $\left\{P\left(X_{i} \mid P a_{i}\right)\right\}_{i=1}^{N}$ fully specify the joint distribution of $\mathbf{X}$. Indeed, let $G$ be compatible with an order $\Omega$, that is, $X_{\Omega(1)} \prec X_{\Omega(2)} \prec \cdots \prec X_{\Omega(N)}$. Then, taking into account the LMP, it is evident that with respect to $G$, the joint probability distribution permits the factorization
$P(\mathbf{X})=\prod_{i=1}^{N} P\left(X_{\Omega(i)} \mid X_{\Omega(1)}, \ldots, X_{\Omega(i-1)}\right)=\prod_{i=1}^{N} P\left(X_{\Omega(i)} \mid P a_{\Omega(i)}\right)=\prod_{i=1}^{N} P\left(X_{i} \mid P a_{i}\right)$.
Depending on the context, in a pair $(G, P)$, we shall refer to $P$ either as a joint distribution, such as in the left-hand side of the above display, or as a set of conditional probability tables, as in the right-hand side above.

For any DAG $G$, there is a maximal set $\mathcal{I}(G)$ of (structural) conditional independence relations of the form $(A \Perp B \mid C)$, for $A, B, C \subset \mathbf{X}$ and $A, B \neq \emptyset$, determined by LMP [9]. On the other hand, $P$ also defines a set of (distributional) independence constraints on $\mathbf{X}$. Condition (2) in the definition of DBN is needed for assuring that the sets of structural and distributional independence statements in fact coincide. We say that two DAGs $G_{1}$ and $G_{2}$ are equivalent and write $G_{1} \cong G_{2}$ if $\mathcal{I}\left(G_{1}\right)=\mathcal{I}\left(G_{2}\right)$. With $[G]$ we shall denote the class of DAGs equivalent to $G$. Necessary and sufficient conditions for DAG equivalence can be found in $[13,4]$. We call two DBNs $\left(G_{1}, P_{1}\right)$ and $\left(G_{2}, P_{2}\right)$ equivalent if their joint distributions are equal, $P_{1}=P_{2}$, which implies equivalence between their graph structures, $G_{1} \cong G_{2}$. The essential problem of BN learning is the recovery of the equivalence class $[G]$ from data.

Another useful notion is that of network complexity. The complexity of a DBN $G$ is typically measured by the number of parameters $d f(G)$ needed to specify the conditional probability table of $G$. Let $q\left(X_{i}\right)$ be the number of states, or discrete levels, of $X_{i}$ and $q\left(P a_{i}\right)=\prod_{X \in P a_{i}} q(X)$ be the number of states of the parent set $P a_{i}$. Since for every state of $P a_{i}, q\left(X_{i}\right)-1$ parameters are needed to define the corresponding multinomial distribution for $X_{i}$, we have $d f(G)=\sum_{i=1}^{N} q\left(P a_{i}\right)\left(q\left(X_{i}\right)-1\right)$.

Next we formulate the maximum likelihood estimation (MLE) in the context of DBNs. Let $D_{n}=\left\{x^{s}\right\}_{s=1}^{n}$ be a sample of $n$ independent observations on the vector $\mathbf{X}$. Then, the log-likelihood of a $\mathrm{DBN}(G, P)$ with respect to $D_{n}$ is

$$
\begin{equation*}
\log \mathcal{L}\left(G, P \mid D_{n}\right)=\sum_{i=1}^{N} \sum_{s=1}^{n} \log P\left(X_{i}=x_{i}^{s} \mid P a_{i}=x_{P a_{i}}^{s}\right), \tag{2.1}
\end{equation*}
$$

where $x_{i}^{s}$ and $x_{P a_{i}}^{s}$ are the states of $X_{i}$ and its parent set $P a_{i}$ in the $s$-th record $x^{s}$. According to the ML principle, a DBN estimator can be obtained by
maximizing (2.1). Before proceeding with the inference in presence of missing values we need to introduce some useful statistics and convenient notations.

We write $k \in X_{i}$ to index the states of $X_{i}$ and adopt a multi-index notation, $j \in P a_{i}$, for the parent configurations of $X_{i}$. Let $1_{i, k j}$ be the indicator function of the event ( $X_{i}=k, P a_{i}=j$ ). For a given sample $D_{n}$ let us define the counts $n_{i, k j} \equiv \sum_{s=1}^{n} 1_{i, k j}\left(x^{s}\right), n_{i, j} \equiv \sum_{k \in X_{i}} n_{i, k j}$ and $n_{i} \equiv \sum_{j \in P a_{i}} n_{i, j}$. A record $x^{s}$ in $D_{n}$ we shall call incomplete if some of the values $x_{i}^{s}$ are missing. By convention, if the value of $X_{i}$ in $x^{s}$ is missing, then $1_{i, k j}\left(x^{s}\right)=0$, while if some of the parents in $P a_{i}$ are missing, then both $1_{i, k j}\left(x^{s}\right)=0$ and $1_{i, j}\left(x^{s}\right)=0$. It is always the case then that $n_{i} \leq n$. We shall consider an inference framework using the counts $n_{i, j}$ and $n_{i, k j}$ as statistics summarizing the information in the sample $D_{n}$.

Let $\mathbf{Z}=\left(Z_{i}\right)_{i=1}^{N}$ be a binary random vector such that $Z_{i}=1$ if $X_{i}$ is observed and $Z_{i}=0$ if it is missing. For an index set $A$ we define $Z_{A}=\prod_{i \in A} Z_{i}$. The joint distribution of $(\mathbf{X}, \mathbf{Z})$ describes all incomplete samples $D_{n}$ of observations on $\mathbf{X}$.

Let us introduce the probabilities $\theta_{i}, \theta_{i, j}$ and $\theta_{i, k j}$ as

$$
\begin{gather*}
\theta_{i} \equiv P\left(Z_{i}=1, Z_{P a_{i}}=1\right) \\
\theta_{i, j} \equiv P\left(P a_{i}=j \mid Z_{i}=1, Z_{P a_{i}}=1\right) \\
\theta_{i, k j} \equiv P\left(X_{i}=k \mid P a_{i}=j, Z_{i}=1, Z_{P a_{i}}=1\right) \tag{2.2}
\end{gather*}
$$

With $\theta$ we shall denote the set $\left\{\theta_{i}, \theta_{i, j}, \theta_{i, k j}\right\}_{i, k, j}$ and call it observed conditional probability table of $G$.

For a sample of fixed size $n$, the random variables $n_{i}$ and the random vectors $\left\{n_{i, j}\right\}_{j \in P a_{i}}$ and $\left\{n_{i, k j}\right\}_{k \in X_{i}}$ then satisfy

$$
\begin{align*}
n_{i} \mid n & \sim \operatorname{Binom}\left(\theta_{i}, n\right) \\
\left\{n_{i, j}\right\}_{j} \mid n_{i} & \sim \operatorname{Multinom}\left(\left\{\theta_{i, j}\right\}_{j}, n_{i}\right) \\
\left\{n_{i, k j}\right\}_{k} \mid n_{i, j} & \sim \operatorname{Multinom}\left(\left\{\theta_{i, k j}\right\}_{k}, n_{i, j}\right) \tag{2.3}
\end{align*}
$$

Therefore, as long as $n_{i}, n_{i, j}$ and $n_{i, k j}$ are of interest, the table $\theta$ is all we need to know about the DBN and the mechanism of missingness.

The usual point estimators of $\theta_{i}, \theta_{i, k j}$ and $\theta_{i, j}$ are

$$
\hat{\theta}_{i}=\frac{n_{i}}{n}, \hat{\theta}_{i, j}=\frac{n_{i, j}}{n_{i}}, \hat{\theta}_{i, k j}=\frac{n_{i, k j}}{n_{i, j}} .
$$

We shall denote the conditional table defined by $\hat{\theta}$ 's with $\hat{\theta}\left(G \mid D_{n}\right)$ to emphasize that it is estimated for the DAG $G$ from the sample $D_{n}$. The statistics $\hat{\theta}_{i, j}$ and $\hat{\theta}_{i, k j}$ are unbiased estimators of $\theta_{i, j}$ and $\theta_{i, k j}$, respectively

$$
\begin{equation*}
E \hat{\theta}_{i, j}=E_{n_{i}} E\left(\left.\frac{n_{i, j}}{n_{i}} \right\rvert\, n_{i}\right)=\theta_{i, j}, E \hat{\theta}_{i, k j}=E_{n_{i, j}} E\left(\left.\frac{n_{i, k j}}{n_{i, j}} \right\rvert\, n_{i, j}\right)=\theta_{i, k j} \tag{2.4}
\end{equation*}
$$

The missing data distribution usually belongs to one of the following categories:
(i) The data is missing completely at random (MCAR) when the missing probabilities are unrelated to either the observed or the unobserved values. In this case $\mathbf{Z}$ is independent of $\mathbf{X}$ and we have $\theta_{i, j}=P\left(P a_{i}=j\right)$ and $\theta_{i, k j}=$ $P\left(X_{i}=k \mid P a_{i}=j\right)$.
(ii) The data is missing at random (MAR) when the missing probabilities depend on the observed values but not on the unobserved ones. Let us consider a special case of MAR when for each $i$, there is $C_{i} \subset \mathbf{X}$ such that $X_{i} \notin C_{i}$, $C_{i} \cap P a_{i}=\emptyset$ and $\left(Z_{i}, Z_{P a_{i}}\right)$ is independent of $\left(X_{i}, P a_{i}\right)$ given $C_{i}$. If furthermore $C_{i}$ has no descendants of $X_{i}$, then, by application of LMP, $\theta_{i, k j}=P\left(X_{i}=\right.$ $\left.k \mid P a_{i}=j\right)$ holds. For a general MAR however the latter may not be true.
(iii) If the missing probabilities depend on the unobserved values we have not missing at random (NMAR) case and then neither $\theta_{i, j}=P\left(P a_{i}=j\right)$ nor $\theta_{i, k j}=P\left(X_{i}=k \mid P a_{i}=j\right)$ hold anymore.

As we discuss in Section 3.1, the missing data distribution is implicated in network identifiability. In this regard, the MCAR model is the most transparent one for it does not interfere with the network topology.

### 2.2. Node-average log-likelihood

We consider two objective functions for estimating DBNs based on the loglikelihood (2.1). The first one is the sample average log-likelihood

$$
\begin{align*}
\tilde{l}\left(G \mid D_{n}\right) & =\frac{1}{n} \sum_{i=1}^{N} \sum_{j \in P a_{i}} \sum_{k \in X_{i}} n_{i, k j} \log \hat{\theta}_{i, k j} \\
& =\sum_{i=1}^{N} \sum_{j \in P a_{i}} \frac{n_{i, j}}{n} \sum_{k \in X_{i}} \hat{\theta}_{i, k j} \log \hat{\theta}_{i, k j} \tag{2.5}
\end{align*}
$$

When the data has no missing values we have $n \tilde{l}\left(G \mid D_{n}\right)=\max _{\theta} \log \mathcal{L}\left(G, \theta \mid D_{n}\right)$.
The second objective function is the sum of sample average node log-likelihoods

$$
\begin{align*}
l\left(G \mid D_{n}\right) & =\sum_{i=1}^{N} \frac{1}{n_{i}} \sum_{j \in P a_{i}} \sum_{k \in X_{i}} n_{i, k j} \log \hat{\theta}_{i, k j} \\
& =\sum_{i=1}^{N} \sum_{j \in P a_{i}} \hat{\theta}_{i, j} \sum_{k \in X_{i}} \hat{\theta}_{i, k j} \log \hat{\theta}_{i, k j}=\sum_{i=1}^{N} l\left(X_{i} \mid P a_{i}, D_{n}\right) \tag{2.6}
\end{align*}
$$

where $l\left(X_{i} \mid P a_{i}, D_{n}\right) \equiv \sum_{j \in P a_{i}} \hat{\theta}_{i, j} \sum_{k \in X_{i}} \hat{\theta}_{i, k j} \log \hat{\theta}_{i, k j}$ is known as negative conditional entropy of node $X_{i}$. Hereafter, we drop the qualifier 'sample average' from (2.5) and (2.6) and call (2.6) node-average log-likelihood (NAL).

If $D_{n}$ is a complete sample, then for every $i, n_{i}=\sum_{j \in P a_{i}} n_{i, j}=n$. Hence $\hat{\theta}_{i, j} \hat{\theta}_{i, k j}=n_{i, k j} / n$ and consequently $\tilde{l}\left(G \mid D_{n}\right)=l\left(G \mid D_{n}\right)$. If the data is incomplete however, we may have $n_{i}<n$ and then (2.5) and (2.6) will be different. In the latter case, the log-likelihood (2.5) may have imbalanced representation of
the potential parent sets. For example, if for two different parent sets $P a_{i}$ and $P a_{i}^{\prime}$ of the i-th node $n_{i}\left(P a_{i}^{\prime}\right)<n_{i}\left(P a_{i}\right)$, then $P a_{i}^{\prime}$ might be preferably selected due to the smaller size of the subsample that represents it in (2.5) even when $P a_{i}^{\prime}$ has worse fit than $P a_{i}$, i.e. $l\left(X_{i} \mid P a_{i}^{\prime}, D_{n}\right)<l\left(X_{i} \mid P a_{i}, D_{n}\right)$. The simplest solution to this problem - discarding all incomplete records in the sample - may drastically reduce the effective sample size. On the other hand, (2.6) can utilize all $n_{i}$ sample records for estimation of $\theta_{i, k j}$ 's. Essentially, NAL exploits the decomposable nature of the log-likelihood (2.5) and, by adjusting for the sample size, allows comparison of models fitted to different samples. We mention that, similarly, NAL can be adopted in other decomposable log-likelihood based models.

It can be easily demonstrated that the maximum likelihood principle alone is inefficient for estimating DBNs. Let us assume for simplicity that $\mathcal{G}$ comprises all DBNs with node order compatible with the index order, $X_{i} \prec X_{2} \prec \cdots \prec X_{N}$. The maximum NAL equation, $\hat{G}=\arg \max l\left(G \mid D_{n}\right)$, will then result in the following estimates for the parents set $P a_{i}$

$$
\hat{P} a_{i}=\arg \max _{P a_{i} \subseteq\{1, \ldots, i-1\}} \hat{\theta}_{i, j} \sum_{k \in X_{i}} \hat{\theta}_{i, k j} \log \hat{\theta}_{i, k j} .
$$

From the increasing property of the conditional log-likelihood (see Lemma 7.1 below) it follows that the solution of the above equation is $P a_{i}=\{1, \ldots, i-1\}$, for every $i>1$. Thus, the MLE solution will be the most complex DBN in $\mathcal{G}$ and will overestimate the true $G$. In the remainder of this paper we shall investigate more closely the properties of NAL-based estimation in a model selection context and shall provide criteria for asymptotically consistent estimation.

### 2.3. Relation between NAL maximization and EM algorithm

In missing data settings, the standard way to utilize all of the available data is to apply an EM algorithm - see [8] for application of EM to Bayesian networks. For a sample $D_{n}$ let $D_{n}^{o b s}$ be the observed part of the data. The EM algorithm involves the following conditional expectation

$$
\begin{aligned}
Q\left(G, P \mid G^{\prime}, P^{\prime}\right) & \equiv E\left(\log P\left(D_{n} \mid G, P\right) \mid D_{n}^{o b s}, G^{\prime}, P^{\prime}\right) \\
& =\sum_{i=1}^{N} \sum_{j \in P a_{i}, k \in X_{i}} E\left(\sum_{s=1}^{n} 1_{\left\{x_{i}^{s}=k, x_{P a_{i}}^{s}=j\right\}} \mid D_{n}^{o b s}, G^{\prime}, P^{\prime}\right) \log P_{i, k j}
\end{aligned}
$$

where $P_{i, k j}=P\left(X_{i}=k \mid P a_{i}=j\right)$. Finding $Q$ implements the E-step of the algorithm. The M-step maximizes $Q\left(G, P \mid G^{\prime}, P^{\prime}\right)$ for $G$ and $P$. Solutions of the EM algorithm are all $(\hat{G}, \hat{P})$ such that $Q(\hat{G}, \hat{P} \mid \hat{G}, \hat{P})=\max _{G, P} Q(G, P \mid \hat{G}, \hat{P})$.

It can be shown that NAL maximization is equivalent to solving a sub-optimal EM algorithm with $\sum_{s=1}^{n} 1_{\left\{x_{i}^{s}=k, x_{P a_{i}}^{s}=j\right\}}$ replaced by the sum $n_{i, k j}+n_{i, k j}^{m i s}$, where $n_{i, k j}$ and $n_{i, k j}^{m i s}$ are the number of records in $D_{n}$ for which the event $\left(X_{i}=k, P a_{i}=j\right)$ is observed and missing, respectively. For each $i$, this is
equivalent to replacing $D_{n}^{o b s}$ by a sub-sample $D_{n, i}^{o b s}$ with all $x^{s}$ from $D_{n}^{o b s}$ for which $\left(X_{i}, P a_{i}\right)$ is not fully observed being removed. Let $n_{i}^{m i s}=\sum_{k, j} n_{i, k j}^{m i s}=$ $n-n_{i}$. Given $D_{n}^{o b s}, n_{i, k j}, n_{i}$ and $n_{i}^{m i s}$ are fixed but $n_{i, k j}^{m i s}$ is random. In fact, $n_{i, k j}^{m i s}$ conditional on $\left(n_{i}=n_{i}^{m i s}, G^{\prime}, P^{\prime}\right)$ follows a Binomial distribution. Since $E\left(\sum_{s=1}^{n} 1_{\left\{x_{i}^{s}=k, x_{P a_{i}}^{s}=j\right\}} \mid D_{n, i}^{o b s}, G^{\prime}, P^{\prime}\right)=n_{i, k j}+E\left(n_{i, k j}^{m i s} \mid n_{i}=n_{i}^{m i s}, G^{\prime}, P^{\prime}\right)$, we define

$$
Q^{*}\left(G, P \mid G^{\prime}, P^{\prime}\right) \equiv \sum_{i=1}^{N} \sum_{j \in P a_{i}, k \in X_{i}}\left(n_{i, k j}+E\left(n_{i, k j}^{m i s} \mid n_{i}=n_{i}^{m i s}, G^{\prime}, P^{\prime}\right)\right) \log P_{i, k j}
$$

Under the MAR assumption $P\left(X_{i}, P a_{i} \mid Z_{i}, Z_{P a_{i}}, G^{\prime}, P^{\prime}\right)=P\left(X_{i}, P a_{i} \mid G^{\prime}, P^{\prime}\right)$, we have $E\left(n_{i, k j}^{m i s}\right)=E_{n_{i, j}^{m i s}}\left(n_{i, j}^{m i s} P_{i, k j}^{\prime}\right)=n_{i}^{m i s} P_{i, j}^{\prime} P_{i, k j}^{\prime}$. Therefore

$$
\begin{equation*}
Q^{*}\left(G, P \mid G^{\prime}, P^{\prime}\right)=\sum_{i=1}^{N} \sum_{j \in P a_{i}, k \in X_{i}}\left(n_{i, k j}+\left(n-n_{i}\right) P_{i, j}^{\prime} P_{i, k j}^{\prime}\right) \log P_{i, k j} \tag{2.7}
\end{equation*}
$$

We then observe that $P \mapsto Q^{*}(G, P \mid G, P)$ is maximized for $\hat{P}_{i, j}=n_{i, j} / n_{i}=\hat{\theta}_{i, j}$ and $\hat{P}_{i, k j}=n_{i, k j} / n_{i, j}=\hat{\theta}_{i, k j}$, and consequently

$$
Q^{*}(\hat{G}, \hat{P} \mid \hat{G}, \hat{P})=n \sum_{i=1}^{N} \sum_{j \in P a_{i}} \hat{P}_{i, j} \sum_{k \in X_{i}} \hat{P}_{i, k j} \log \hat{P}_{i, k j}=n l\left(\hat{G} \mid D_{n}\right)
$$

We hence conclude that the EM algorithm based on $Q^{*}$ essentially maximizes the NAL function (2.6). Of course, $Q$ utilizes all of the available data, while $Q^{*}$ does not - when even one component of $\left(X_{i}, P a_{i}\right)$ is missing, $Q^{*}$ treats the entire record as missing, while $Q$ tries to use the available information by calculating (often costly) conditional expectations. Nevertheless, the NAL-based inference is much more efficient than the naive approach that ignores all records for which at least one component of $\mathbf{X}$ is missing; even more so in cases when the dimensionality $N$ is much higher that the maximum size of $\left|P a_{i}\right|$ 's (the so-called in-degree). In such cases the difference between $Q$ and $Q^{*}$ is less pronounced (if $n_{i} \gg n-n_{i}$ then $\left.\left|Q-Q^{*}\right| \ll\left|Q^{*}\right|\right)$ and so is the difference between NAL maximization and EM algorithm. Moreover, the sub-optimality of NAL maximization is counterbalanced by its computational simplicity. The EM algorithm is usually intractable for data with number of nodes in the thousands while NAL optimization may still be a possibility. In conclusion, the NAL-based learning seems to be an effective and computationally more affordable alternative of EM for estimating high dimensional, low in-degree Bayesian networks.

## 3. MLE and model selection

Let $\left(G_{0}, P_{0}\right)$ be a DBN with nodes $\mathbf{X}$, parent sets $P a_{i}^{0}$, and observed conditional probability table $\theta_{0}$. For an arbitrary DAG $G$ with nodes $\mathbf{X}$ and parents $P a_{i}$, we
consider probability distribution $P_{G \mid G_{0}}$ on $\mathbf{X}$ induced by $G_{0}$ which, for a state $x$ of $\mathbf{X}$, is given by

$$
\begin{equation*}
P_{G \mid G_{0}}(x) \equiv \prod_{i=1}^{N} P_{0}\left(X_{i}=x_{i} \mid P a_{i}=x_{P a_{i}}\right) \tag{3.1}
\end{equation*}
$$

and compare it to

$$
P_{0}(x)=\prod_{i=1}^{N} P_{0}\left(X_{i}=x_{i} \mid P a_{i}^{0}=x_{P a_{i}^{0}}\right)
$$

In general, $P_{G \mid G_{0}}$ is different from $P_{0}$ and $\left(G, P_{G \mid G_{0}}\right)$ may not be well defined DBN, because $G$ is not necessarily a minimal DAG compatible with $P_{G \mid G_{0}}$ (see condition (2) from the definition of DBN). However, if $G$ is a minimal DAG such that $P_{G \mid G_{0}}=P_{0}$, then $G \cong G_{0}$.

We also consider the following observation probabilities of $G$ induced by $G_{0}$

$$
\begin{gathered}
\theta_{i}\left(G \mid G_{0}\right) \equiv P\left(Z_{i}=1, Z_{P a_{i}}=1\right) \\
\theta_{i, j}\left(G \mid G_{0}\right) \equiv P\left(P a_{i}=j \mid Z_{i}=1, Z_{P a_{i}}=1\right) \\
\theta_{i, k j}\left(G \mid G_{0}\right) \equiv P\left(X_{i}=k \mid P a_{i}=j, Z_{i}=1, Z_{P a_{i}}=1\right)
\end{gathered}
$$

where the probabilities are with respect to the joint distribution of $\mathbf{Z}$ and $\mathbf{X} \mid\left(G_{0}, P_{0}\right)$. Recall that according to $(2.2)$ the entries of $\theta_{0}$ are

$$
\begin{gathered}
\theta_{i}^{0} \equiv P\left(Z_{i}=1, Z_{P a_{i}^{0}}=1\right) \\
\theta_{i, j}^{0}=P\left(P a_{i}^{0}=j \mid Z_{i}=1, Z_{P a_{i}^{0}}=1\right) \\
\theta_{i, k j}^{0}=P\left(X_{i}=k \mid P a_{i}^{0}=j, Z_{i}=1, Z_{P a_{i}^{0}}=1\right)
\end{gathered}
$$

Let $\theta\left(G \mid G_{0}\right)$ denote the corresponding conditional probability table with entries $\theta_{i}\left(G \mid G_{0}\right), \theta_{i, j}\left(G \mid G_{0}\right)$ and $\theta_{i, j k}\left(G \mid G_{0}\right)$. Clearly, we can write $\theta\left(G_{0} \mid G_{0}\right)=\theta_{0}$. Moreover, in the important case when $Z$ is MCAR we have

$$
\begin{gathered}
\theta_{i, j}\left(G \mid G_{0}\right) \stackrel{\text { mcar }}{=} P_{0}\left(P a_{i}=j\right) \\
\theta_{i, k j}\left(G \mid G_{0}\right) \stackrel{\text { mcar }}{=} P_{0}\left(X_{i}=k \mid P a_{i}=j\right)
\end{gathered}
$$

and $\theta\left(G \mid G_{0}\right)$ is the conditional probability table corresponding to $P_{G \mid G_{0}}$.
Next, we define the NAL of $G$ with respect to $G_{0}$ given by

$$
\begin{align*}
l\left(G \mid G_{0}\right) & \equiv \sum_{i=1}^{N} l\left(X_{i} \mid P a_{i}, G_{0}\right) \\
l\left(X_{i} \mid P a_{i}, G_{0}\right) & \equiv \sum_{j \in P a_{i}} \theta_{i, j} \sum_{k \in X_{i}} \theta_{i, k j} \log \theta_{i, k j} \tag{3.2}
\end{align*}
$$

where $\theta_{i, j}=\theta_{i, j}\left(G \mid G_{0}\right)$ and $\theta_{i, k j}=\theta_{i, k j}\left(G \mid G_{0}\right)$. Essentially, $l\left(X_{i} \mid P a_{i}, G_{0}\right)$ is the observed population negative entropy of $X_{i}$ conditional on $P a_{i}$ and $l\left(G \mid G_{0}\right)$ is the population version of (2.6). For brevity, we shall write $l\left(G_{0}\right)$ instead of $l\left(G_{0} \mid G_{0}\right)$.

### 3.1. Identifiability

Let $G_{0}$ belong to a collection $\mathcal{G}$ of DAGs with nodes $\mathbf{X}$. If $D_{n}$ is an independent sample from a DBN $\left(G_{0}, P_{0}\right)$, by the strong law of large numbers, for any fixed $G \in \mathcal{G}, \hat{\theta}_{i, k j}\left(G \mid D_{n}\right) \rightarrow \theta_{i, k j}\left(G \mid G_{0}\right)$, a.s., and hence, $l\left(G \mid D_{n}\right) \rightarrow l\left(G \mid G_{0}\right)$, a.s. as $n \rightarrow \infty$. A necessary condition for MLE consistency is the identifiability of $G_{0}$, which in its usual sense requires $l\left(G \mid G_{0}\right)<l\left(G_{0}\right)$ for all $G \in \mathcal{G}$ such that $G \neq G_{0}$. The latter is a strong requirement however, for thus defined the identifiability will never hold unless $G_{0}$ is a maximal DAG in $\mathcal{G}$ that contains $G_{0}$ - as we show later (Lemma 7.1) $l\left(X_{i} \mid P a_{i}, G_{0}\right)$ is a non-decreasing function of $P a_{i}$. In the light of this observation we shall adopt a more appropriate definition of identifiability, one that assumes smaller likelihoods only for the DAGs not containing the true one. To simplify the notation, hereafter we shall refer to the DBN $\left(G_{0}, P_{0}\right)$ simply as $G_{0}$.

Definition 3.1. We say that $G_{0}$ is identifiable in $\mathcal{G}$, if for any $G \in \mathcal{G}$ we have $l\left(G \mid G_{0}\right) \leq l\left(G_{0}\right)$ when $G_{0} \subseteq G$ and $l\left(G \mid G_{0}\right)<l\left(G_{0}\right)$ when $G_{0} \nsubseteq G$.

Note that the identifiability of $G_{0}$ depends on the joint distribution of $\mathbf{X}$ and $\mathbf{Z}$. The utility of this definition is due to the following observation. If $G_{0}$ is identifiable in $\mathcal{G}$, then

$$
\begin{equation*}
G^{*} \equiv \min \left\{\tilde{G} \in \mathcal{G} \mid l\left(\tilde{G} \mid G_{0}\right)=\max _{G \in \mathcal{G}} l\left(G \mid G_{0}\right)\right\}=G_{0} \tag{3.3}
\end{equation*}
$$

implicitly assuming the existence of unique such minimum $G^{*}$ (in general we may have multiple minimal $\tilde{G}$ maximizing the NAL). Moreover, it is easy to check that (3.3) is a necessary and sufficient condition for identifiability. In 'learning from data' settings, we can replace $l\left(G \mid G_{0}\right)$ in (3.3) with $l\left(G \mid D_{n}\right)$ and find an estimator $\hat{G}^{*}$ of the minimal DAG $G^{*}$, exhaustively in $\mathcal{G}$ or by some more efficient algorithm. Then $\hat{G}^{*}$ would be an estimator of $G_{0}$ as well. In this way, the identifiability assures the principal possibility of recovering $G_{0}$.

It is intuitively clear that in order to recover the graph structure $G_{0}$ from incomplete samples, the missing data mechanism should not interfere with the associations between $X_{i}$ 's determined by $G_{0}$. This condition is satisfied for any MCAR model. In more general MAR settings, the identifiability of $G_{0}$ depends on the interaction between $\mathbf{X}$ and $\mathbf{Z}$ and can not be judged without actually knowing $G_{0}$. We thus regard the MAR assumption as not significant generalization over MCAR due to the practical impossibility to check it prior to learning.

The next result shows that in MCAR settings the population NAL does not increase when the true DBN is nested in a larger one, and moreover, that its maximum is achieved only for DAGs equivalent to the true one.

Proposition 3.1. If $\mathbf{Z}$ is $M C A R$, we have the following:
(i) if $G_{0} \subseteq G$ then $l\left(G \mid G_{0}\right)=l\left(G_{0}\right)$;
(ii) $\max _{G} l\left(G \mid G_{0}\right)=l\left(G_{0}\right)$, where the maximum is over all $D A G s$ on $\mathbf{X}$;
(iii) if $l\left(G \mid G_{0}\right)=l\left(G_{0}\right)$, then $P_{G \mid G_{0}}=P_{0}$.

From these properties of the NAL of $G$ with respect to $G_{0}$ we can draw two immediate conclusions as stated in the next two corollaries.

Corollary 3.1. If $\mathbf{Z}$ is $M C A R$ then $G_{0}$ is identifiable in any set of DAGs compatible with its order.

Therefore, provided a true node order is known (that is an order with which $G_{0}$ is compatible; there might be many such orders), $G_{0}$ can be recovered from the set of all DAGs compatible with that order.

We can further extend Definition 3.1 to account for classes of equivalent DBNs. Recall that, ultimately, it is the independence relation set $\mathcal{I}\left(G_{0}\right)$, shared among all equivalent to $G_{0}$ DBNs, that is of main interest. In the view of condition (3.3), we say that $\left[G_{0}\right]$ is identifiable in $\mathcal{G}$ if

$$
\begin{equation*}
\min \left\{\tilde{G} \in \mathcal{G} \mid l\left(\tilde{G} \mid G_{0}\right)=\max _{G \in \mathcal{G}} l\left(G \mid G_{0}\right)\right\} \cong G_{0} \tag{3.4}
\end{equation*}
$$

in the sense that any minimal $\tilde{G}$ that maximizes the NAL $l\left(G \mid G_{0}\right)$ is equivalent to $G_{0}$ (we also assume that the set on the left is not empty). Proposition 3.1, cases (ii) and (iii), implies that the maximum NAL is $l\left(G_{0}\right)$ and any $G$ that attains this maximum satisfies $P_{G \mid G_{0}}=P_{0}$. If in addition $G$ is minimal, then $\left(G, P_{G \mid G_{0}}\right)$ is a well defined DBN which is equivalent to $\left(G_{0}, P_{0}\right)$ and hence (3.4) is satisfied. We have thus obtained the following.

Corollary 3.2. If $\mathbf{Z}$ is $M C A R$, then $\left[G_{0}\right]$ is identifiable in any $\mathcal{G}$ that contains at least one element of $\left[G_{0}\right]$. In particular, $\left[G_{0}\right]$ is (globally) identifiable in the set of all DAGs on $\mathbf{X}$.

As defined, the identifiability of the equivalence class [ $G_{0}$ ] depends implicitly on the choice of log-likelihood proxy function. Note that $\left[G_{0}\right]$ is not guaranteed to be identifiable, even in MCAR settings, if in (3.4) we replace the NAL $l$ with the standard $\log$-likelihood $\tilde{l}$ from (2.5).

### 3.2. NAL-based scoring functions

As we have observed earlier, the MLE criteria selects the most complex BN in $\mathcal{G}$ containing $G_{0}$ and unless some complexity penalization is imposed, the MLE is prone to overfitting. Methodologically, there are two approaches addressing the model selection problem. The first one is provided by the Bayesian paradigm, where the parameter $(G, \theta)$ is assumed coming from some prior distribution and one looks for the maximum posterior estimator. The second, frequentist, approach is to optimize a scoring function based on the log-likelihood and additional complexity penalization term - a penalized log-likelihood. We consider a general scoring function of the form

$$
\begin{equation*}
S\left(G \mid D_{n}\right)=l\left(G \mid D_{n}\right)-\lambda_{n} h(G) \tag{3.5}
\end{equation*}
$$

where $\lambda_{n}$ are positive numbers indexed by the sample size $n$ and $h(G)$ is a positive function accounting for the complexity of the $G$. When needed, we
shall write $S_{h}$ to specify what $h$ is meant. The role of the sequence $\lambda_{n}$ is to apply a proper amount of penalty that guarantees estimation consistency.

One can employ different measures for network complexity. Any complexity function $h$ is assumed to be increasing in the following sense: for any two DAGs $G_{1}$ and $G_{2}$ such that $G_{1} \subset G_{2}, G_{1} \neq G_{2}$, we have $h\left(G_{1}\right)<h\left(G_{2}\right)$. In regard to DBNs, a typical choice is the total number of parameters $d f(G)$ needed to specify the multinomial conditional distributions of $G$, that is, the number of independent parameters in $\theta$.

We return to (3.5) with some typical examples. Since the NAL $l\left(G \mid D_{n}\right)$, being sum of node sample averages, is normalized by the sample size, the standard model selection criteria AIC and BIC, formulated in terms of the scoring function (3.5) are given by $\lambda_{n}=1 / n$ and $\lambda_{n}=0.5 \log (n) / n$, respectively. The so called minimum description length (MDL) score, representing the information content of a model, is given by $\log (n) d f(G) / n$ and is equivalent to BIC.

Similarly to NAL, often, the chosen overall DBN complexity can also be represented as a sum of node-wise complexities. For example, $d f(G)=\sum_{i} d f\left(X_{i} \mid P a_{i}\right)$, $d f\left(X_{i} \mid P a_{i}\right) \equiv\left(q\left(X_{1}\right)-1\right) q\left(P a_{i}\right)$. In such cases it might be more appropriate to replace $\lambda_{n}$ with node-specific penalization $\lambda_{n_{i}}$ 's

$$
\begin{equation*}
S\left(G \mid D_{n}\right)=\sum_{i=1}^{N}\left\{l\left(X_{i} \mid P a_{i}, D_{n}\right)-\lambda_{n_{i}} h\left(X_{i} \mid P a_{i}\right)\right\} \tag{3.6}
\end{equation*}
$$

We shall refer to these as decomposable scores. Typically, one uses one and the same function of $n$ to express $\lambda_{n_{i}}$ 's, such as $\lambda_{n}=\lambda_{0} n^{-\alpha}, \alpha \in(0,0.5)$. The decomposable BIC criteria then is

$$
\begin{array}{r}
S_{B I C}\left(X_{i} \mid P a_{i}, D_{n, i}\right)=l\left(X_{i} \mid P a_{i}, D_{n}\right)-0.5 \frac{\log \left(n_{i}\right)}{n_{i}} d f\left(X_{i} \mid P a_{i}\right) \\
S_{B I C}\left(G \mid D_{n}\right)=\sum_{i=1}^{N} S_{B I C}\left(X_{i} \mid P a_{i}, D_{n}\right)=\sum_{i=1}^{N} \frac{1}{n_{i}} B I C\left(X_{i} \mid P a_{i}, D_{n, i}\right) \tag{3.7}
\end{array}
$$

where $D_{n, i}$ is the sub-sample of $D_{n}$ of size $n_{i}$ for which $\left(X_{i}, P a_{i}\right)$ is observed and $B I C\left(X_{i} \mid P a_{i}, D_{n, i}\right)$ is the original BIC criteria, (1.1), applied to the regression model $X_{i} \mid P a_{i}$.

As we have stated in the introduction, we consider an MLE based model selection by maximizing $S$ as a function of $G$ given a sample $D_{n}$,

$$
\begin{equation*}
\hat{G}=\arg \max _{G \in \mathcal{G}} S\left(G \mid D_{n}\right) \tag{3.8}
\end{equation*}
$$

Note that we do not maximize $S$ for $G$ and $\theta$ simultaneously. We estimate $\theta$ for each $G$ using the plug-in estimator $\hat{\theta}\left(G \mid D_{n}\right)$ and then the DAG with maximal score is chosen as graph structure estimator. In what follows we show that, by solving (3.8) for proper $\lambda_{n}$, we can obtain consistent estimation of the true model with no further conditions on $h$.

Let $\hat{G}$ be the estimator (3.8) for a sample $D_{n}$ coming from a DBN $G_{0}$. Then the following claim is immediate.

Proposition 3.2 (Consistency Criteria). Provided for any $G_{1} \in \mathcal{G}$ and $G_{2} \in \mathcal{G}$ the following two conditions are satisfied
(C1) if $G_{0} \subseteq G_{1}$ but $G_{0} \nsubseteq G_{2}$, then $P\left(S\left(G_{1} \mid D_{n}\right)>S\left(G_{2} \mid D_{n}\right)\right) \rightarrow 1$, as $n \rightarrow \infty$,
(C2) if $G_{0} \subseteq G_{1}, G_{0} \subset G_{2}$ and $h\left(G_{1}\right)<h\left(G_{2}\right)$, then $P\left(S\left(G_{1} \mid D_{n}\right)>S\left(G_{2} \mid D_{n}\right)\right) \rightarrow$ 1 , as $n \rightarrow \infty$,
$\hat{G}$ is a consistent estimator of $G_{0}$, that is, $P\left(\hat{G} \neq G_{0}\right) \rightarrow 0$, as $n \rightarrow \infty$.
The conditions (C1) and (C2) are relaxed versions of those used in [4]. In fact, the consistent scoring criterion in [4] is a special case of the more abstract formulation of model selection consistency in [7]. We end this section with the following important observation.

Corollary 3.3. If conditions (C1) and (C2) are satisfied for any $D A G$ equivalent to $G_{0}$, then $[\hat{G}]$ is a consistent estimator of $\left[G_{0}\right]$, that is, $P(\mathcal{I}(\hat{G}) \neq$ $\left.\mathcal{I}\left(G_{0}\right)\right) \rightarrow 0$, as $n \rightarrow \infty$.

## 4. Estimation consistency

Let $\left(G_{0}, P_{0}\right)$ be a DBN with conditional table $\theta_{0}$ in a set of DAGs $\mathcal{G}$ and $D_{n}$ be an independent sample drawn from it. In this section we investigate the consistency of the estimators $\hat{G}$ and $[\hat{G}]$ with respect to a scoring function $S$, where $\hat{G}$ is given by (3.8).

As we have observed earlier, if the data has missing values, it is not anymore true that $l\left(G \mid D_{n}\right)=\tilde{l}\left(G \mid D_{n}\right)$, the usual sample average log-likelihood (2.5). Therefore, $(\hat{G}, \hat{\theta})$ is no longer an MLE for $\left(G_{0}, \theta_{0}\right)$ and the standard consistency results from the asymptotic theory are not directly applicable. A proper account for the incompleteness of the data is thus needed.

For a sample of fixed size $n$, the random variables $n_{i}$ and the random vectors $\left\{n_{i, j}\right\}_{j \in P a_{i}}$ and $\left\{n_{i, k j}\right\}_{k \in X_{i}}$ satisfy

$$
\begin{align*}
\left\{n_{i, j}\right\}_{j} \mid n_{i} & \sim \text { Multinom }\left(\left\{\theta_{i, j}\left(G \mid G_{0}\right)\right\}_{j}, n_{i}\right) \\
\left\{n_{i, k j}\right\}_{k} \mid n_{i, j} & \sim \text { Multinom }\left(\left\{\theta_{i, k j}\left(G \mid G_{0}\right)\right\}_{k}, n_{i, j}\right) \tag{4.1}
\end{align*}
$$

and the statistics $\hat{\theta}_{i, j}$ and $\hat{\theta}_{i, k j}$ are unbiased estimators of $\theta_{i, j}\left(G \mid G_{0}\right)$ and $\theta_{i, k j}\left(G \mid G_{0}\right)$, respectively. Moreover, if $G_{0}$ is identifiable in $\mathcal{G}$, then for each $i$, the probability of the event ' $\left(X_{i}, P a_{i}^{0}\right)$ is observed' must be strictly positive, i.e. $\theta_{i}^{0}>0$. Since $\mathcal{G}$ is always finite, the following is well defined

$$
\begin{equation*}
\beta(\mathcal{G}) \equiv \min _{G \in \mathcal{G}} \min _{i=1}^{N}\left\{\theta_{i}\left(G \mid G_{0}\right) \mid \theta_{i}\left(G \mid G_{0}\right)>0\right\} \tag{4.2}
\end{equation*}
$$

and $\beta(\mathcal{G})>0$. The complete data case can be thus represented as $\beta(\mathcal{G})=1$. Note that $\beta$ depends implicitly on the distribution of $\mathbf{Z}$.

The next result establishes the rate of convergence of the empirical NAL to the population one without imposing any restrictions on the distribution of $\mathbf{Z}$ or on $G_{0}$ ( $G_{0}$ need not be identifiable).

Lemma 4.1. Let $D_{n}$ be sample from a $D B N\left(G_{0}, P_{0}\right)$. Then for any $D A G G$

$$
\begin{equation*}
l\left(G \mid D_{n}\right)-l\left(G \mid G_{0}\right)=O_{p}\left(n^{-1 / 2}\right) \tag{4.3}
\end{equation*}
$$

which implies $l\left(G \mid D_{n}\right) \rightarrow_{p} l\left(G \mid G_{0}\right)$.
Providing conditions for scoring function consistency is our next goal. Let us assume that $G_{0}$ is identifiable in $\mathcal{G}$. In the light of Lemma 4.1, if $G$ does not contain $G_{0}$, then there is a positive constant $\delta$ such that $l\left(G_{0} \mid D_{n}\right)-l\left(G \mid D_{n}\right)>\delta$ with probability going to 1 , as $n \rightarrow \infty$. It is evident therefore that if the sequence $\lambda_{n}$ diminishes with $n, \lambda_{n} \rightarrow 0$, then, asymptotically, the scoring function $S_{h}$ will select an estimator that contains the true model $G_{0}$ regardless of the chosen complexity function $h$. In addition however, we want that estimator to get close (in sense of the complexity measured by $h$ ) to $G_{0}$ with the increase of the sample size. Since for any $G$ such that $G_{0} \subset G$ we have $l\left(G \mid D_{n}\right)-l\left(G_{0} \mid D_{n}\right) \rightarrow_{p} 0$, the latter can be assured if we require $\lambda_{n}$ to diminish at a slower rate than that of $l\left(G \mid D_{n}\right)-l\left(G_{0} \mid D_{n}\right)$. We show that this rate is $n^{-1}$ for complete samples and $n^{-1 / 2}$ in case of missing data.

We moreover show that the consistency sufficient conditions, $\lambda_{n}=o(1)$ and $n^{-1 / 2} \lambda_{n}^{-1}=o(1)$, become essentially necessary. More precisely, the necessity is guaranteed if the following condition is satisfied. As usual $P a_{i}$ and $P a_{i}^{0}$ denote the parent sets of $G$ and $G_{0}$, respectively.

Condition 4.1. There are $G \in \mathcal{G}$ with $G_{0} \subset G$ and $i \in\{1, \ldots, N\}$ such that $P a_{j}=P a_{j}^{0}$ for all $j \neq i, P a_{i} \backslash P a_{i}^{0} \neq \emptyset$ and $P\left(Z_{P a_{i} \backslash P a_{i}^{0}}=1 \mid Z_{i}=1, Z_{P a_{i}^{0}}=\right.$ 1) $\in(0,1)$.

In words, the condition refers to the possibility of extending the parent set of a node of $G_{0}$ by one or more new nodes that are, conditionally, neither always observed nor never observed (thus $G_{0}$ must not be a maximal DAG in $\mathcal{G}$ ).

Next, we summarize the above observations in the following theorem.
Theorem 4.1. Let $G_{0}$ be identifiable in $\mathcal{G}$ and $S$ be a scoring function (3.5) with penalization parameter $\lambda_{n}$ such that $\lambda_{n} \rightarrow 0$. The following are satisfied.
(i) If $\beta(\mathcal{G}) \in(0,1)$ and $\sqrt{n} \lambda_{n} \rightarrow \infty$, then $\hat{G}$ is consistent estimator of $G_{0}$.
(ii) If $\beta(\mathcal{G})=1$ and $n \lambda_{n} \rightarrow \infty$, then $\hat{G}$ is consistent estimator of $G_{0}$.
(iii) If $\mathbf{Z}$ is MCAR, Condition 4.1 holds and $\underline{\lim } \sqrt{n} \lambda_{n}<\infty$, then $\hat{G}$ is inconsistent estimator of $G_{0}$.
The complete data case of the theorem, (ii), also follows from a more general result by [7] (Proposition 1.2 and Remark 1.2). There, the consistency result is derived using the properties of MLE for exponential families and central limit theorem. The essential contribution of the above theorem is in the missing data cases $(i)$ and (iii). We emphasize that case $(i)$ holds for a general $\mathcal{G}$ and missing data distribution as long as $G_{0}$ is identifiable in $\mathcal{G}$. In (iii) however, we require for $\mathbf{Z}$ to be MCAR in order to guarantee that the condition $\sqrt{n} \lambda_{n} \rightarrow \infty$ is necessary for consistent estimation. Below we make some further remarks.

The claims of the theorem are established by verifying conditions (C1) and (C2) from Proposition 3.2 for $G_{0}$ and hence, for any DAG equivalent to $G_{0}$.

Therefore, it follows from Corollary 3.3 that the theorem remains true if we replace $G_{0}$ by $\left[G_{0}\right]$ and $\hat{G}$ by $[\hat{G}]$. The theorem thus provides conditions for consistent estimation of the equivalence class of $G_{0}$.

As evident from the proof of the theorem, the requirement $\lambda_{n} \rightarrow 0$ is needed for guaranteeing the first, ( C 1 ), consistency condition in Proposition 3.2, while $\sqrt{n} \lambda_{n} \rightarrow \infty\left(n \lambda_{n} \rightarrow \infty\right)$ is required for the second one (C2). The AIC selection criterion, $\lambda_{n}=1 / n$, is not consistent for it satisfies (C1) but fails to satisfy (C2), regardless of $\beta$. It will thus recover the true structure but will tend to select networks with higher complexities than the true one. Therefore AIC is prone to overfitting and so is any scoring function with $n \lambda_{n}=O(1)$. At the other end of the consistency spectrum of $\alpha, \lim _{n} \sup \lambda_{n}>0$, the estimated networks will tend to have complexities below the true one. Due to the missingness, there is an implication regarding the $\mathrm{BIC}(\mathrm{MDL})$ criterion, $\lambda_{n}=0.5 \log (n) / n$. Because $n \log (n) / n \rightarrow \infty$ but $\sqrt{n} \log (n) / n \rightarrow 0$, BIC is guaranteed to be consistent only in the complete data case and it will be, in general, inconsistent in MCAR settings (see the corollaries that follow). The numerical results presented in Section 5 confirm this conclusion.

Theorem 4.1 requires the observation probability $\beta(\mathcal{G})$ to be fixed. If we allow it to depend on $n$, case (ii) of the theorem arises from (i) if we have $\lim _{n} \beta_{n}(\mathcal{G})=1$. Then $n \lambda_{n} \rightarrow \infty$ is a sufficient consistency condition. There is no contradiction with case (iii), since then it must be that $P\left(Z_{X_{i}}=1, Z_{P a_{i}}=\right.$ $1)=1$ and $P\left(Z_{X_{i}}=1, Z_{P a_{i}^{0}}=1\right)=1$, and hence Condition 4.1 fails. As evident from the proof of the theorem, when $\overline{\lim } \beta_{n}<1,(i)$ and (iii) still hold. We leave undecided the last alternative $0<\underline{\lim } \beta_{n}<\varlimsup \beta_{n}=1$.

Next, we argue that Condition 4.1 arises naturally in MCAR settings. In the probability space of all MCAR distributions for $\mathbf{Z}$ defined by the Borel sets in $\left\{u \in[0,1]^{2^{N}-1}, \sum_{k=1}^{2^{N}-1} u_{k} \leq 1\right\}$ (a distribution $u$ is defined by assigning each of the $2^{N}$ states of $\mathbf{Z}$ a probability value in $[0,1]$ such that their sum is 1 ), the subspace of distributions for which $P\left(Z_{P a_{i} \backslash P a_{i}^{0}}=1 \mid Z_{i}=1, Z_{P a_{i}^{0}}=1\right)=0$ or 1 has Borel measure zero. We thus have the following consequences of Theorem 4.1 which extend Corollary 3.1 and 3.2 , and essentially summarize the practical contribution of this investigation.

Corollary 4.1. Let $\left(G_{0}, P_{0}\right)$ be a non-maximal DBN and $\mathcal{G}$ consist of all DAGs compatible with a node order of $G_{0}$. Then, for almost all MCAR distributions, $\hat{G}$ is consistent estimator of $G_{0}$ if and only if $\lambda_{n} \rightarrow 0$ and $\sqrt{n} \lambda_{n} \rightarrow \infty$.

In the last statement we assume that $\mathcal{G}$ comprises all DAGs on $\mathbf{X}$ and use the global identifiability of $\left[G_{0}\right]$.

Corollary 4.2. Provided that $\mathcal{I}\left(G_{0}\right)$ is non-empty, for almost all MCAR distributions, $[\hat{G}]$ is consistent estimator of $\left[G_{0}\right]$ if and only if $\lambda_{n} \rightarrow 0$ and $\sqrt{n} \lambda_{n} \rightarrow \infty$.

Note that, the non-emptiness of $\mathcal{I}\left(G_{0}\right)$ is required in order for any DAG equivalent to $G_{0}$ to be non-maximal and hence, for Condition 4.1 to hold.

## 5. Numerical experiments

With the number of possible DAGs being super-exponential to the number of nodes, the task of reconstructing a DBN from data is in general NP-hard. The MLE based problem (3.8) essentially requires exhausting all DAGs in $\mathcal{G}$. For the purpose of numerical illustration in this section we make two simplifying the inference assumptions - that the causal order of the nodes of the original $\operatorname{DBN} G_{0}$ is known, as well as the maximum size of the parent sets of $G_{0}$, its in-degree. We thus assume that the search set $\mathcal{G}$ comprises all DAGs compatible with a true node order. By Corollary 3.1, when the missing data model is MCAR, $G_{0}$ is identifiable in $\mathcal{G}$. In our numerical experiments we use exclusively the complexity function $d f$, which recall is given by $d f(G, \theta)=\operatorname{dim}\left(\Theta_{G}\right)$, and the decomposable scoring function (3.6). Then (3.8) can be solved by an efficient exhaustive search via dynamic programming, an approach that is implemented in the catnet package for $R,[1]$. We are aware that more general learning algorithms are available in the literature that can also accommodate available case analysis based on NAL. For example, one can implement a search based on local optimizations as described in [4] by replacing the usual log-likelihood with NAL. However, our goal here is not to compare different learning strategies but to empirically verify the conclusions of Theorem 4.1, which hold for all NAL-based estimators (3.8).

The standard AIC and BIC model selection criteria are compared to scoring functions with $\lambda_{n}=(1 / N) n^{-\alpha}$ for different choices of $\alpha \in(0,1)$. The factor $1 / N$, to some extent arbitrary, makes the penalization relatively small for not large $n$ (note that NAL is of rate $O(N)$ ). For this choice of $\lambda_{n}$ and small $n$, the estimator $\hat{G}$ therefore may over-fit the data but, provided the scoring criteria is consistent, $d f(\hat{G})$ should approach the true complexity as $n$ increases.

### 5.1. Simulated 2-node network

Here we consider a simplest possible example to verify the consistency of the NAL estimator (3.8). We generate samples from a model $G_{0}$ with 2 independent binary variables $X_{1}$ and $X_{2}$ (that is $P a_{1}=P a_{2}=\emptyset$ ) with marginal probabilities $\theta_{1}=(0.4,0.6)$ and $\theta_{2}=(0.3,0.7)$, respectively. We assume that $X_{1} \prec X_{2}$ and then the only alternative to $G_{0} \mathrm{BN}$ model is $G_{1}$ with $P a_{1}=\emptyset$ and $P a_{2}=\left\{X_{1}\right\}$. We also assume that $X_{2}$ is always observed $\left(P\left(Z_{2}=1\right)=1\right)$ but $Z_{1}$ is MCAR with different missing probabilities $P\left(Z_{1}=0\right) \in\{0,0.01,0.05,0.10,0.25\}$. The sample observation probability is then $\beta=1-P\left(Z_{1}=0\right)$. For each $\beta$ and sample size $n \in\left\{10^{2}, 10^{3}, 10^{4}, 10^{5}\right\}$, we generate 1000 samples $D_{n}$ and count how many times $S\left(G_{1} \mid D_{n}\right)>S\left(G_{0} \mid D_{n}\right)$, that is, $\hat{G}=G_{1}$ and $G_{1}$ is erroneously selected instead of $G_{0}$. Table 1 summarizes results for different choices of the penalization parameter $\alpha$ as well as BIC and AIC. As expected, all considered scoring functions except AIC are consistent in the no missing data case $(\beta=1)$. In presence of missing values however, for scoring functions with $\alpha>0.5$ the percent of false model selections is significant; moreover, it increases when the proportion of missing values increases, suggesting inconsistency. In particular,

Table 1
Consistency results for a simulated 2-node network. Two possible models $G_{0}$ and $G_{1}$ are considered as described in the main text. Shown are the percents of wrong selections (choosing the alternative $G_{1}$ instead of the true model $G_{0}$ )

| $\alpha$ | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | BIC | AIC |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $n=10^{2}$ |  |  |  |  |  |  |  |  |  |
| $\beta=1$ | 0.0 | 0.0 | 0.0 | 0.3 | 0.9 | 3.5 | 10.6 | 2.8 | 16.0 |
| $\beta=0.99$ | 0.0 | 0.0 | 0.0 | 0.5 | 1.7 | 6.6 | 17.0 | 4.5 | 22.9 |
| $\beta=0.95$ | 0.0 | 0.0 | 0.2 | 0.9 | 3.8 | 12.8 | 24.0 | 8.7 | 31.2 |
| $\beta=0.90$ | 0.0 | 0.0 | 0.0 | 0.7 | 6.9 | 16.6 | 31.5 | 12.5 | 37.0 |
| $\beta=0.75$ | 0.0 | 0.0 | 1.1 | 7.0 | 18.5 | 29.9 | 40.2 | 27.3 | 44.4 |
| $n=10^{3}$ |  |  |  |  |  |  |  |  |  |
| $\beta=1$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.2 | 3.6 | 0.7 | 13.9 |
| $\beta=0.99$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 1.6 | 13.3 | 2.9 | 33.5 |
| $\beta=0.95$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.4 | 12.1 | 28.8 | 17.1 | 42.0 |
| $\beta=0.90$ | 0.0 | 0.0 | 0.0 | 0.1 | 3.6 | 19.1 | 34.7 | 23.0 | 43.9 |
| $\beta=0.75$ | 0.0 | 0.0 | 0.0 | 1.9 | 15.8 | 33.2 | 42.4 | 36.2 | 47.2 |
| $n=10^{4}$ |  |  |  |  |  |  |  |  |  |
| $\beta=1$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.8 | 0.2 | 15.0 |
| $\beta=0.99$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 2.7 | 24.7 | 13.9 | 44.1 |
| $\beta=0.95$ | 0.0 | 0.0 | 0.0 | 0.0 | 1.5 | 21.3 | 37.7 | 31.6 | 47.8 |
| $\beta=0.90$ | 0.0 | 0.0 | 0.0 | 0.0 | 7.0 | 28.9 | 41.5 | 36.5 | 47.5 |
| $\beta=0.75$ | 0.0 | 0.0 | 0.0 | 1.8 | 22.3 | 41.2 | 50.5 | 47.3 | 53.8 |
| $n=10^{5}$ |  |  |  |  |  |  |  |  |  |
| $\beta=1$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 17.7 |
| $\beta=0.99$ | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 11.8 | 37.8 | 35.8 | 50.4 |
| $\beta=0.95$ | 0.0 | 0.0 | 0.0 | 0.0 | 3.5 | 31.3 | 44.6 | 43.3 | 49.8 |
| $\beta=0.90$ | 0.0 | 0.0 | 0.0 | 0.0 | 13.1 | 36.1 | 47.2 | 46.0 | 50.5 |
| $\beta=0.75$ | 0.0 | 0.0 | 0.0 | 1.0 | 21.4 | 38.5 | 45.5 | 45.3 | 48.2 |

the inconsistency of BIC is very pronounced for all $\beta<1$. Even when the proportion of missing values is only 1 percent, $\beta=0.99$, the percent of wrong selections start from 4.5 for $n=10^{2}$ and climbs to 35.8 for $n=10^{5}$. The presented results are in strong support of the predictions of Theorem 4.1.

### 5.2. Consistent estimation of the ALARM network

Here we consider a well known in the literature benchmark network. ALARM, a medical diagnostic alarm message system for patient monitoring developed by [2], is a typical example of belief propagation network as those employed in many expert systems. The DAG of ALARM has 37 nodes, 45 directed edges, varying number of categories $(2,3$ and 4$)$ and complexity of 473 . We perform network reconstruction using both complete and MCAR missing data simulated from the network, in order to confirm the effect of missingness on the model selection as predicted by Theorem 4.1. The graph structures of the estimated networks are compared to the original one by the so-called $F$-score, the harmonic mean of precision $(T P /(T P+F P))$ and recall $(T P /(T P+F N))$, where $P$ and $N$ refer to the presence and absence of directed edges. $F$-score of 1 represents perfect reconstruction.

Missing data samples are simulated by deleting 1, 2 and 4 values from each sample record, completely at random (so, there is not even 1 fully complete

Table 2
Model selection results for the ALARM network using complete samples $(\beta=1)$ and samples following MCAR models with $\beta=0.84$ and $\beta=0.70$. Shown are the $F$-scores between the true network and the estimated ones

| n | 5 e 2 | 2.5 e 3 | 5e3 | 2.5 e 4 | 5 e 4 | 1 e 5 | 2.5 e 5 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| no missing values, $\beta=1$ |  |  |  |  |  |  |  |
| $\alpha=0.25$ | 0.85 | 0.95 | 0.97 | 0.97 | 0.97 | 0.97 | 0.98 |
| $\alpha=0.3$ | 0.88 | 0.97 | 0.97 | 0.97 | 0.98 | 0.98 | 0.99 |
| $\alpha=0.35$ | 0.86 | 0.97 | 0.97 | 0.98 | 0.99 | 1.00 | 1.00 |
| $\alpha=0.4$ | 0.81 | 0.97 | 0.98 | 1.00 | 1.00 | 1.00 | 1.00 |
| $\alpha=0.45$ | 0.75 | 0.96 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| $\alpha=0.5$ | 0.73 | 0.92 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 |
| $\alpha=0.75$ | 0.57 | 0.65 | 0.62 | 0.62 | 0.65 | 0.63 | 0.65 |
| BIC | 0.85 | 0.97 | 0.98 | 1.00 | 1.00 | 1.00 | 1.00 |
| AIC | 0.80 | 1.84 | 0.82 | 0.81 | 0.79 | 0.80 | 0.80 |
| MCAR, $\beta=0.84$ |  |  |  |  |  |  |  |
| $\alpha=0.25$ | 0.79 | 0.86 | 0.91 | 0.97 | 0.97 | 0.97 | 0.97 |
| $\alpha=0.3$ | 0.79 | 0.88 | 0.88 | 0.93 | 0.93 | 0.97 | 0.99 |
| $\alpha=0.35$ | 0.79 | 0.82 | 0.85 | 0.90 | 0.92 | 0.95 | 1.00 |
| $\alpha=0.4$ | 0.74 | 0.78 | 0.81 | 0.85 | 0.86 | 0.91 | 0.92 |
| $\alpha=0.45$ | 0.69 | 0.72 | 0.77 | 0.80 | 0.79 | 0.83 | 0.83 |
| $\alpha=0.5$ | 0.65 | 0.70 | 0.70 | 0.72 | 0.71 | 0.78 | 0.74 |
| $\alpha=0.75$ | 0.56 | 0.60 | 0.62 | 0.61 | 0.61 | 0.63 | 0.61 |
| BIC | 0.77 | 0.82 | 0.82 | 0.77 | 0.71 | 0.71 | 0.66 |
| AIC | 0.74 | 0.68 | 0.67 | 0.62 | 0.61 | 0.63 | 0.61 |
| $\mathrm{MCAR}, \beta=0.70$ |  |  |  |  |  |  |  |
| $\alpha=0.25$ | 0.76 | 0.83 | 0.88 | 0.93 | 0.97 | 0.97 | 0.97 |
| $\alpha=0.3$ | 0.75 | 0.80 | 0.86 | 0.89 | 0.91 | 0.93 | 1.00 |
| $\alpha=0.35$ | 0.69 | 0.73 | 0.84 | 0.87 | 0.89 | 0.93 | 0.95 |
| $\alpha=0.4$ | 0.68 | 0.73 | 0.78 | 0.82 | 0.82 | 0.87 | 0.89 |
| $\alpha=0.45$ | 0.66 | 0.65 | 0.74 | 0.76 | 0.79 | 0.78 | 0.79 |
| $\alpha=0.5$ | 0.58 | 0.61 | 0.69 | 0.70 | 0.72 | 0.70 | 0.73 |
| $\alpha=0.75$ | 0.56 | 0.57 | 0.61 | 0.61 | 0.62 | 0.61 | 0.63 |
| BIC | 0.74 | 0.75 | 0.79 | 0.74 | 0.74 | 0.69 | 0.67 |
| AIC | 0.70 | 0.61 | 0.66 | 0.62 | 0.63 | 0.61 | 0.63 |

record in the samples). Since the maximum parent size is 3 , in the first case, the probability to have no missing 3 -node subset $\left(X_{i}, P a_{i}\right)$ is $\binom{36}{3} /\binom{37}{3}$ and hence, the effective observation probability $\beta$ from (4.2) is about 0.92 . When 2 values per record are deleted, $\beta$ drops to $\binom{35}{3} /\binom{37}{3} \approx 0.84$; when 4 values are deleted, $\beta$ is about 0.70 . The target set of models $\mathcal{G}$ includes all DAGs with 37 nodes, maximum of 3 parents per node, compatible with the true node order. Under these constraints, the number of DAGs in $\mathcal{G}$ is 1133 . For each possible complexity $t$, the optimal network $\hat{G}(t)$ is found and a final selection is made according to their scores.

Table 2 shows comparison results for 9 scoring criteria ( $\alpha=0.25,0.3,0.35,0.4$, $0.45,0.5,0.75, \mathrm{BIC}$ and AIC ) and 7 samples sizes, from 5 e 2 to 2.5 e 5 . In the complete data case, the scoring functions with $\alpha \in[0.4,0.5]$ and BIC reconstruct the true network for all samples with $n \geq 2.5 e 4$. As predicted, in the missing data cases the score function for $\alpha=0.5$ and BIC become inconsistent due to overfitting. This effect is more clearly demonstrated in Figures 1 and 2 that show the complexity profile functions for different experimental cases. According


FIG 1. Estimating the ALARM network from complete and MCAR samples of size $2.5 e 5$. For $\beta=1,0.92,0.84,0.70$, the so called complexity profile - the complexity of the estimated network on $y$-axis as a function of the penalization parameter $\alpha$ on $x$-axis - is shown in the range $[0.25,0.5]$. In presence of missing values $(\beta<1)$, the BIC selection (dash, horizontal) tends to move up and away from the true complexity of 473 (solid, horizontal), demonstrating the inconsistency of BIC.


Fig 2. Estimating $A L A R M$ from $M C A R$ samples with fixed observation probability $\beta=0.84$. Shown are the complexity profiles of 6 samples of increasing size $n$. For $n>1 e 6$, the complexity of the BIC estimates are off charts (>900).
to Theorem 4.1, the complexity profiles in the $(0,0.5)$ range should converge to the horizontal line of true complexity. In Figure 1 the sample size is kept fixed and we see that with the increase of the proportion of missing values ( $\beta$ decreasing), the profiles depart from the line of true complexity. On the other hand, Figure 2 shows profiles of samples with fixed proportion of missing values but of increasing size. We observe that, although slow, the profiles get closer to the line of true complexity as $n$ increases. It is also evident that the BIC selected complexity drifts up and away from the true one with the increase of the sample size, an indication for its inconsistency.

## 6. Conclusion

We have addressed the problem of discrete Bayesian network estimation from incomplete data by maximizing a penalized log-likelihood scoring function. The essential step in our approach is replacing the usual log-likelihood with a sum of node-average log-likelihoods, the so-called NAL. We have motivated our decision with a more efficient utilization of the available data and have shown the connection between NAL optimization and EM algorithm. Although our setup allows the missing data distribution to be arbitrary as long as the true DAG structure remains identifiable, the latter rarely holds for general MAR models. As we have demonstrated however, in MCAR settings, the identifiability of the set of independence relations, which characterizes all networks equivalent to the true one, is always guaranteed. We have shown, Theorems 4.1, that in presence of missing values the NAL-based estimator (3.8) requires more stringent conditions on the penalization parameter $\lambda_{n}$ to achieve consistency than in the complete data case. The discrepancy is due to the fact that in NAL each node may utilize different data subset for estimation thus reducing the overall convergence rate. Although the theorem guarantees consistency for penalties in a continuous range, choosing an optimal penalization parameter that performs well in finite sample settings is an open problem deserving further investigation.

The scope of this article has been limited to discrete BNs for which selfcontained proofs of the results have been derived. It is straightforward however to apply NAL-based estimation to other classes of parametric BNs, such as linear Gaussian networks. Then, as long as for any $G, G_{0} \subset G, l\left(G \mid D_{n}\right)-$ $l\left(G_{0} \mid D_{n}\right)$ is $O_{p}\left(n^{-1 / 2}\right)$, in the missing, and $O_{p}\left(n^{-1}\right)$, in the complete data case, Theorem 4.1, with some technical modification of the proofs, seems to remain valid. Formulating identifiability and consistency for available case analysis in more general graphical model settings is thus a subject of continuing interest.

## 7. Proofs

The next lemma is instrumental in the proof of Proposition 3.1. It shows that the (population) node log-likelihood $l(X \mid A), X \in\left\{X_{i}\right\}_{i=1}^{N}, A \subset\left\{X_{i}\right\}_{i=1}^{N}$, is an increasing function of $A$ with respect to the set inclusion operation. In complete
data settings, this result is better know as non-negativity of the Kullback-Leibler divergence.
Lemma 7.1. For any $A, B \subset \mathbf{X}$ and $X \in \mathbf{X}$ such that $Z_{B}$ is independent of $(X, A, B)$ given $\left(Z=1, Z_{A}=1\right)$, we have $l(X \mid A) \leq l(X \mid A, B)$. The inequality is strict if $P\left(X \mid A, Z=1, Z_{A}=1\right) \neq P\left(X \mid A, B, Z=1, Z_{A}=1\right)$.
Proof. Let

$$
\theta_{k a} \equiv P\left(X=k \mid A=a, Z=1, Z_{A}=1\right)
$$

By assumption

$$
\begin{aligned}
\theta_{k a b} & \equiv P\left(X=k \mid A=a, B=b, Z=1, Z_{A}=1, Z_{B}=1\right) \\
& =P\left(X=k \mid A=a, B=b, Z=1, Z_{A}=1\right)
\end{aligned}
$$

We can therefore write the expression $\theta_{k a}=\sum_{b \in B} P\left(B=b \mid A=a, Z=1, Z_{A}=\right.$ 1) $\theta_{\text {kab }}$. By the convexity of the function $t \mapsto t \log (t)$ we have

$$
\theta_{k a} \log \left(\theta_{k a}\right) \leq \sum_{b \in B} P\left(B=b \mid A=a, Z=1, Z_{A}=1\right) \theta_{k a b} \log \left(\theta_{k a b}\right)
$$

and the claim follows from

$$
\begin{aligned}
l(X \mid A)= & \sum_{a \in A} P\left(A=a \mid Z=1, Z_{A}=1\right) \sum_{k \in X} \theta_{k a} \log \left(\theta_{k a}\right) \\
\leq & \sum_{a \in A, b \in B} P\left(B=b \mid A=a, Z=1, Z_{A}=1\right) \\
& \times P\left(A=a \mid Z=1, Z_{A}=1\right) \sum_{k \in X} \theta_{k a b} \log \left(\theta_{k a b}\right) \\
= & \sum_{a \in A, b \in B} P\left(A=a, B=b \mid Z=1, Z_{A}=1, Z_{B}=1\right) \sum_{k \in X} \theta_{k a b} \log \left(\theta_{k a b}\right) \\
= & l(X \mid A, B)
\end{aligned}
$$

The last inequality is strict if $P\left(X \mid A, Z=1, Z_{A}=1\right) \neq P(X \mid A, B, Z=1$, $Z_{A}=1$ ).
Proof of Proposition 3.1. Part (i)
Let $G_{0} \subset G$. The MCAR condition on $\mathbf{Z}$ and LMP imply that for every $i$ and $Y \subset \mathbf{X}$ such that $Y \prec_{G_{0}} X_{i}$ and $Y \cap P a_{i}^{0}=\emptyset$, the following two conditions hold
(i) $\left(Y, Z_{Y}\right)$ is independent of $X_{i}$ given $\left(P a_{i}^{0}, Z_{i}=1, Z_{P a_{i}^{0}}=1\right)$.
(ii) $Z_{Y}$ is independent of $P a_{i}^{0}$ given $\left(Z_{i}=1, Z_{P a_{i}^{0}}=1\right)$.

For each $i$, since $P a_{i}^{0} \subset P a_{i}, P a_{i} \backslash P a_{i}^{0} \prec X_{i}$, by (i) we have that $\left(P a_{i} \backslash P a_{i}^{0}\right.$, $Z_{P a_{i} \backslash P a_{i}^{0}}$ ) and $X_{i}$ are independent conditionally on $P a_{i}^{0}$, and therefore for each $j \in P a_{i}^{0}$ and $j^{\prime} \in P a_{i} \backslash P a_{i}^{0}$,

$$
\begin{aligned}
\theta_{i, k\left(j j^{\prime}\right)}\left(G \mid G_{0}\right) & =P\left(X_{i}=k \mid P a_{i}=\left(j j^{\prime}\right), Z_{i}=1, Z_{P a_{i}}=1\right) \\
& =P\left(X_{i}=k \mid P a_{i}^{0}=j, Z_{i}=1, Z_{P a_{i}^{0}}=1\right)=\theta_{i, k j}^{0}
\end{aligned}
$$

Moreover, by (ii) applied to $Z_{P a_{i} \backslash P a_{i}^{0}}$ and $P a_{i}^{0}$

$$
\begin{aligned}
\sum_{j^{\prime} \in P a_{i} \backslash P a_{i}^{0}} \theta_{i, j j^{\prime}} & =P\left(P a_{i}^{0}=j \mid Z_{i}=1, Z_{P a_{i}}=1\right) \\
& =P\left(P a_{i}^{0}=j \mid Z_{i}=1, Z_{P a_{i}^{0}}=1\right)=\theta_{i, j}^{0}
\end{aligned}
$$

which implies

$$
\begin{aligned}
l\left(X_{i} \mid P a_{i}\right) & =\sum_{j \in P a_{i}^{0}} \sum_{j^{\prime} \in P a_{i} \backslash P a_{i}^{0}} \theta_{i, j j^{\prime}} \sum_{k \in X_{i}} \theta_{i, k\left(j j^{\prime}\right)} \log \theta_{i, k\left(j j^{\prime}\right)} \\
& =\sum_{j \in P a_{i}^{0}} \sum_{k \in X_{i}} \theta_{i, k(j)} \log \theta_{i, k(j)}=l\left(X_{i} \mid P a_{i}^{0}\right)
\end{aligned}
$$

We thus have $l\left(G \mid G_{0}\right)=l\left(G_{0}\right)$.

## Part (ii) and Part (iii)

By the definition of $l\left(G \mid G_{0}\right)$ in (3.2) and some summation manipulations we obtain

$$
\begin{aligned}
l\left(G \mid G_{0}\right) & =\sum_{i=1}^{N} \sum_{x_{P a_{i}} \in P a_{i}} \sum_{x_{i} \in X_{i}} P_{0}\left(X_{i}=x_{i}, P a_{i}=x_{P a_{i}}\right) \log P_{0}\left(X_{i}=x_{i} \mid P a_{i}=x_{P a_{i}}\right) \\
& =\sum_{i=1}^{N} \sum_{x \in \mathbf{X}} P_{0}(\mathbf{X}=x) \log P_{0}\left(X_{i}=x_{i} \mid P a_{i}=x_{P a_{i}}\right)=\sum_{x \in \mathbf{X}} P_{0}(x) \log P_{G \mid G_{0}}(x),
\end{aligned}
$$

where $x$ indexes the states of $\mathbf{X}$. Since $\sum_{x} P_{0}(x)=1, \sum_{x} P_{G \mid G_{0}}(x)=1$ and the log-function is concave, we have

$$
l\left(G \mid G_{0}\right)=\sum_{x} P_{0}(x) \log P_{G \mid G_{0}}(x) \leq \sum_{x} P_{0}(x) \log P_{0}(x)=l\left(G_{0}\right)
$$

with equality that is achieved only when $P_{G \mid G_{0}}=P_{0}$.
Proof of Corollary 3.1. Let $G_{0} \in \mathcal{G}$ be DBN with a node order $\Omega$ and $\mathcal{G}$ be a set of DAGs compatible with $\Omega$. We need to show that $l\left(G \mid G_{0}\right)<l\left(G_{0}\right)$ for all $G \in \mathcal{G}$ for which $G_{0} \nsubseteq G$.

Note that for any $G \in \mathcal{G}, G \cup G_{0}$ is also a DAG compatible with $\Omega$ and by Proposition 3.1,

$$
l\left(X_{i} \mid P a_{i}\right) \leq l\left(X_{i} \mid P a_{i} \cup P a_{i}^{0}\right)=l\left(X_{i} \mid P a_{i}^{0}\right)
$$

for all $i$. Moreover, because $G_{0} \nsubseteq G$, there is an $i$ such that $P a_{i}=\left(P a_{i}^{0} \backslash Y\right)$ for $Y, \emptyset \neq Y \subset P a_{i}^{0}$. Since $P\left(X_{i} \mid P a_{i}, Z_{i}=1, Z_{P a_{i}}=1\right) \neq P\left(X_{i} \mid P a_{i}^{0}, Z_{i}=\right.$ $1, Z_{P a_{i}^{0}}=1$ ), because by definition $G_{0}$ is a minimal DAG compatible with $P_{0}$, by Lemma 7.1, $l\left(X_{i} \mid P a_{i}^{0} \backslash Y\right)<l\left(X_{i} \mid P a_{i}^{0}\right)$, implying $l\left(G \mid G_{0}\right)<l\left(G_{0}\right)$.

The next result is used in the proof of Lemma 4.1.

Lemma 7.2. Let $\hat{\phi}_{n}=\phi+O_{p}\left(n^{-1 / 2}\right)$ and $\hat{\theta}_{n}=\theta+O_{p}\left(n^{-1 / 2}\right)$ for $\theta>0$. Then

$$
\begin{equation*}
\hat{\phi}_{n} \log \left(\hat{\theta}_{n}\right)-\phi \log (\theta)=O_{p}\left(n^{-1 / 2}\right) \tag{7.1}
\end{equation*}
$$

Proof. By applying Taylor expansion to the logarithm function, we can write

$$
\left.\hat{\phi}_{n} \log \left(\hat{\theta}_{n}\right)-\phi \log (\theta)=\left(\hat{\phi}_{n}-\phi\right) \log (\theta)+\frac{\phi_{n}}{\eta_{n}}\left(\hat{\theta}_{n}-\theta\right)\right)
$$

for some $\eta_{n}$ between $\hat{\theta}_{n}$ and $\theta$. Since $\phi_{n} / \eta_{n} \rightarrow_{p} \phi / \theta<\infty$, the claims follows from the assumptions.

Proof of Lemma 4.1. Let $G \in \mathcal{G}$ has parent sets $P a_{i}$. For all $i=1, \ldots, N$, $j \in P a_{i}$ and $k \in X_{i}$, we define

$$
\xi_{i, k j} \equiv \hat{\theta}_{i, j} \hat{\theta}_{i, k j} \log \hat{\theta}_{i, k j}-\theta_{i, j} \theta_{i, k j} \log \theta_{i, k j}
$$

where $\theta_{i, j}=\theta_{i, j}\left(G \mid G_{0}\right)$ and $\theta_{i, k j}=\theta_{i, k j}\left(G \mid G_{0}\right)$, and $\hat{\theta}_{i, j}$ and $\hat{\theta}_{i, k j}$ are the corresponding estimates. In this notation we have

$$
l\left(X_{i} \mid P a_{i}, D_{n}\right)-l\left(X_{i} \mid P a_{i}, G_{0}\right)=\sum_{i=1}^{N} \sum_{j \in P a_{i}} \sum_{k \in X_{i}} \xi_{i, k j} .
$$

Note that when either $\theta_{i}=0, \theta_{i, j}=0$ or $\theta_{i, k j}=0$ holds, then the state $(i, k j)$ will be unobservable and $\xi_{i, k j}=0$. We thus may assume without loss of generality that $\theta_{i}>0, \theta_{i, j}>0$ and $\theta_{i, k j}>0$ for all $i, j$ and $k$. Moreover, by Hoeffding's inequality, for all $k, j, \epsilon>0, P\left(\left|\hat{\theta}_{i, k j}-\theta_{i, k j}\right| \geq \epsilon\right) \leq 2 \exp \left(-2 n_{i, j} \epsilon^{2}\right)$ and hence $\hat{\theta}_{i, j}=\theta_{i, j}+O_{p}\left(n^{-1 / 2}\right)$ and $\hat{\theta}_{i, k j}=\theta_{i, k j}+O_{p}\left(n^{-1 / 2}\right)$. We can therefore apply Lemma 7.2 to $\hat{\theta}_{i, j} \hat{\theta}_{i, k j}$ and $\hat{\theta}_{i, k j}$ to infer that $\xi_{i, k j}=O_{p}\left(n^{-1 / 2}\right)$, from which the claim follows.

The following two lemmas are essential for the proof of Theorem 4.1. The first one extends Lemma 7.2.

Lemma 7.3. Let for $m=1, \ldots, k, \hat{\theta}_{m}=\theta_{0}+O_{p}\left(n^{-1 / 2}\right), \theta_{0}>0$ and $\hat{\theta}=$ $\sum_{m} \gamma_{m} \hat{\theta}_{m}$, for $\gamma_{m} \geq 0$ such that $\sum_{m} \gamma_{m}=1$. Then

$$
\begin{equation*}
\Delta \equiv \sum_{m} \gamma_{m} \hat{\theta}_{m} \log \left(\hat{\theta}_{m}\right)-\hat{\theta} \log (\hat{\theta})=O_{p}\left(n^{-1}\right) \tag{7.2}
\end{equation*}
$$

Proof. Note that $\hat{\theta}^{m}, \hat{\theta}$ and $\gamma_{m}$ are all considered to be random variables. By applying Taylor expansion to the logarithm function, we can write

$$
\hat{\theta}_{m} \log \left(\hat{\theta}_{m}\right)=\hat{\theta}_{m}\left[\log (\hat{\theta})+\frac{1}{\hat{\theta}}\left(\hat{\theta}_{m}-\hat{\theta}\right)-\frac{1}{2 \eta_{m}^{2}}\left(\hat{\theta}_{m}-\hat{\theta}\right)^{2}\right],
$$

for some $\eta_{m}$ between $\hat{\theta}_{m}$ and $\hat{\theta}$. Since $\eta_{m} \rightarrow_{p} \theta_{0}>0$, by assumption, we have $\left(\hat{\theta}_{m}-\hat{\theta}\right)^{2} / \eta_{m}^{2}=O_{p}\left(n^{-1}\right)$. After some algebra we obtain

$$
\begin{aligned}
\Delta & =\sum_{m} \gamma_{m}\left[\hat{\theta}_{m} \log \left(\hat{\theta}_{m}\right)-\hat{\theta}_{m} \log (\hat{\theta})\right]=\frac{1}{\hat{\theta}}\left[\sum_{m} \gamma_{m} \hat{\theta}_{m}^{2}-\hat{\theta}^{2}\right]+O_{p}\left(n^{-1}\right) \\
& =\frac{1}{\hat{\theta}}\left[\sum_{m} \gamma_{m}\left(\hat{\theta}_{m}-\theta_{0}\right)^{2}-\left(\hat{\theta}-\theta_{0}\right)^{2}\right]+O_{p}\left(n^{-1}\right)
\end{aligned}
$$

Finally, since $1 / \hat{\theta} \rightarrow_{p} 1 / \theta_{0}<\infty$ and $\left(\hat{\theta}-\theta_{0}\right)^{2}=O_{p}\left(n^{-1}\right)$, (7.2) follows.
The next lemma presents a central limit result for difference between sample and sub-sample averages. It consequently establishes a variability rate of $n^{-1 / 2}$ for such differences.
Lemma 7.4. Let for each $n>0, x_{n}^{1}, \ldots, x_{n}^{n}$ be i.i.d. random variables with mean $\mu_{n}$ and variance $\sigma_{n}^{2}$ such that $\sigma_{n}^{2} \rightarrow \sigma^{2}<\infty$. Let for some fixed $\beta \in$ $(0,1), k_{n} \sim \operatorname{Binom}(\beta, n)$ and $\alpha_{n}$ be a random draw without replacements of $k_{n}$ elements from the set $\{1, \ldots, n\}$. Then for

$$
S_{n}=\frac{1}{k_{n}} \sum_{i \in \alpha_{n}} x_{i}-\frac{1}{n} \sum_{i=1}^{n} x_{i}
$$

for almost every sequence $\left\{\alpha_{n}\right\}_{n \geq 1}$, we have

$$
\begin{equation*}
\sqrt{n} S_{n} \rightarrow_{d} \mathcal{N}\left(0, \frac{1-\beta}{\beta} \sigma^{2}\right) \tag{7.3}
\end{equation*}
$$

Proof. First, we rearrange the elements of $S_{n}$

$$
S_{n}=\frac{n-k_{n}}{n k_{n}} \sum_{i \in \alpha_{n}} x_{i}-\frac{1}{n} \sum_{i \notin \alpha_{n}} x_{i}
$$

and define $\left\{y_{n}^{j}\right\}_{j=1}^{k_{n}}$ to be $\frac{n-k_{n}}{n k_{n}} x_{i}$ for $i \in \alpha_{n}$, and $\left\{z_{n}^{j}\right\}_{j=1}^{n-k_{n}}$ to be $-\frac{1}{n} x_{i}$ for $i \notin \alpha_{n}$. Hence $S_{n}=y_{n}^{1}+\cdots+y_{n}^{k_{n}}+z_{n}^{1}+\cdots+z_{n}^{n-k_{n}}$.

We are going to apply the Lindeberg-Feller CLT (see for example Th. 2.27 in [12]) to the triangular sequence

$$
\sqrt{n} y_{n}^{1}, \ldots, \sqrt{n} y_{n}^{k_{n}}, \sqrt{n} z_{n}^{1}, \ldots, \sqrt{n} z_{n}^{n-k_{n}}
$$

A key observation is that for each $n$, since $\alpha_{n}$ is a random draw without replacements, conditionally on $\alpha_{n}, y_{n}$ 's and $z_{n}$ 's are independent. Moreover, it is easy to verify that

$$
E\left(S_{n} \mid k_{n}\right)=\frac{1}{n}\left(\frac{n-k_{n}}{k_{n}} k_{n} \mu-\left(n-k_{n}\right) \mu\right)=0
$$

and
$\operatorname{Var}\left(\sqrt{n} S_{n} \mid k_{n}\right)=\sum_{i=1}^{k_{n}} n \operatorname{Var}\left(y_{n}^{i}\right)+\sum_{i=1}^{n-k_{n}} n \operatorname{Var}\left(z_{n}^{i}\right)=\frac{1}{n}\left(\frac{\left(n-k_{n}\right)^{2}}{k_{n}}+\left(n-k_{n}\right)\right) \sigma_{n}^{2}$.

By the law of large numbers, $k_{n} / n \rightarrow \beta$ almost surely, and therefore

$$
\operatorname{Var}\left(\sqrt{n} S_{n} \mid k_{n}\right) \rightarrow_{a . s .} \frac{1-\beta}{\beta} \sigma^{2}
$$

It is left to verify that for any $\epsilon>0$

$$
\sum_{i=1}^{k_{n}} n E\left(y_{n}^{i}\right)^{2} 1_{\left\{\sqrt{n}\left|y_{n}^{i}\right|>\epsilon\right\}}+\sum_{i=1}^{n-k_{n}} n E\left(z_{n}^{i}\right)^{2} 1_{\left\{\sqrt{n}\left|z_{n}^{i}\right|>\epsilon\right\}} \rightarrow 0
$$

almost surely for $\left\{\alpha_{n}\right\}_{n \geq 1}$. Indeed, the left-hand side sum equals

$$
\frac{\left(n-k_{n}\right)^{2}}{n k_{n}} E\left[\left(x_{n}\right)^{2} 1_{\left\{\frac{n-k_{n}}{k_{n}}\left|x_{n}\right|>\sqrt{n} \epsilon\right\}}\right]+\frac{n-k_{n}}{n} E\left[\left(x_{n}\right)^{2} 1_{\left\{\left|x_{n}\right|>\sqrt{n} \epsilon\right\}}\right]
$$

and, by the assumptions, the above expectations converge to 0 for almost every sequence $\left\{k_{n}\right\}_{n \geq 1}$, and hence, for almost every $\left\{\alpha_{n}\right\}_{n \geq 1}$. Therefore the Lindeberg-Feller CLT is applicable and (7.3) holds.

Proof of Theorem 4.1. We shall first outline some key steps in the proof. The following notion will be useful: for a given sample $D_{n}$ and two DAGs $G_{1}=\left\{P a_{i}^{1}\right\}$ and $G_{2}=\left\{P a_{i}^{2}\right\}$, we say that the NALs $l\left(G_{1} \mid D_{n}\right)$ and $l\left(G_{2} \mid D_{n}\right)$ are estimated upon one and the same sample, if for every $i, l\left(X_{i} \mid P a_{i}^{1}\right)$ and $l\left(X_{i} \mid P a_{i}^{2}\right)$ are estimated from one and the same subsample of $D_{n}$, that is, for every $x^{s} \in D_{n}$, $\left(X_{i}, P a_{i}^{1}\right)$ and $\left(X_{i}, P a_{i}^{2}\right)$ are either both observed or both unobserved (missing) in $x^{s}$. When $\left(X_{i}, P a_{i}\right)$ is observed in all $x^{s} \in D_{n}$, we say that $D_{n}$ is complete with respect to $\left(X_{i}, P a_{i}\right)$.

The essential problem of achieving consistent estimation is to decide between the true model $G_{0}$ and a more complex model $G_{1}$ in which $G_{0}$ is nested. According to Lemma 4.1, the NAL scores $l\left(G_{1} \mid D_{n}\right)$ and $l\left(G_{0} \mid D_{n}\right)$ both converge to $l\left(G_{0}\right)$ at a rate of $n^{-1 / 2}$ and so does their difference - this is essentially Lemma 7.2. If therefore the scoring function penalty $\lambda_{n}$ converges to 0 at a slower than $n^{-1 / 2}$ rate, in the limit, $G_{0}$ would be preferred to $G_{1}$ as a less complex model. The latter condition is sufficient for both complete and missing data (claim (i) of the theorem). However, it turns out that when $l\left(G_{1} \mid D_{n}\right)$ and $l\left(G_{0} \mid D_{n}\right)$ are estimated upon one and the same sample, as in the complete data case, their difference converges at a faster rate of $n^{-1}$ - this is essentially due to the result of Lemma 7.3. Then we can relax the necessary convergence rate of $\lambda_{n}$ and still achieve consistency (claim (ii)). The application of Lemma 7.3 crucially depends on the condition: for every node $X_{i}$, if a sample is complete with respect to $\left(X_{i}, P a_{i}^{0}\right)$ so it is with respect to $\left(X_{i}, P a_{i}^{1}\right)$. In MCAR settings the latter does not hold and the difference $l\left(G_{1} \mid D_{n}\right)-l\left(G_{0} \mid D_{n}\right)$ has a persistent variability of order $n^{-1 / 2}$, due to a central limit result, Lemma 7.4. Consequently the condition on $\lambda_{n}$ to diminish at a rate slower than $n^{-1 / 2}$ becomes both sufficient and necessary (claim (iii)).

The more formal proof follows. We shall prove consistency by verifying conditions (C1) and (C2) in Proposition 3.2. We first assume that $G_{1}, G_{2} \in \mathcal{G}$ are such
that $G_{0} \subseteq G_{1}$ and $G_{0} \nsubseteq G_{2}$. Then, by the identifiability of $G_{0}$, Definition 3.1, we have $l\left(G_{2} \mid G_{0}\right)<l\left(G_{1} \mid G_{0}\right)$. Moreover, by Lemma 4.1, regardless of the observation probability $\beta(\mathcal{G})>0, l\left(G_{1} \mid D_{n}\right) \rightarrow_{p} l\left(G_{1} \mid G_{0}\right)$ and $l\left(G_{2} \mid D_{n}\right) \rightarrow_{p} l\left(G_{2} \mid G_{0}\right)$, and hence, for $\delta=\left(l\left(G_{1} \mid G_{0}\right)-l\left(G_{2} \mid G_{0}\right)\right) / 2, P\left(l\left(G_{1} \mid D_{n}\right)>l\left(G_{2} \mid D_{n}\right)+\delta\right) \rightarrow 1$, as $n \rightarrow \infty$. The consistency condition (C1) therefore holds because the sequence $\lambda_{n}$ diminishes with $n$.

As in the proof of Lemma 4.1, without loss of generality we may assume that for all $G \in \mathcal{G}, i=1, \ldots, N, j \in P a_{i}(G)$ and $k \in X_{i}, \theta_{i, j}\left(G \mid G_{0}\right)>0$ and $\theta_{i, k j}\left(G \mid G_{0}\right)>0$.

## Part (i)

Let now assume $G_{0} \subseteq G_{1}, G_{0} \subseteq G_{2}$ and $h\left(G_{1}\right)<h\left(G_{2}\right)$. To verify the consistency condition (C2) we need to find the rate of convergence of the random variable $l\left(G_{1} \mid D_{n}\right)-l\left(G_{2} \mid D_{n}\right)$. This rate depends on whether the data is complete or not.

By Lemma 4.1, $l\left(G \mid D_{n}\right)=l\left(G \mid G_{0}\right)+O_{p}\left(n^{-1 / 2}\right)$ and since $l\left(G_{1} \mid G_{0}\right)=l\left(G_{2} \mid G_{0}\right)$, we have

$$
l\left(G_{1} \mid D_{n}\right)-l\left(G_{2} \mid D_{n}\right)=O_{p}\left(n^{-1 / 2}\right)
$$

The latter holds regardless of the observation probability $\beta(\mathcal{G})>0$. Therefore, the condition $\sqrt{n} \lambda_{n} \rightarrow \infty$ implies that the positive sequence $\lambda_{n}\left(h\left(G_{2}\right)-h\left(G_{1}\right)\right)$ will overcome the likelihood difference $l\left(G_{2} \mid D_{n}\right)-l\left(G_{1} \mid D_{n}\right)$ with probability approaching 1 , that is, $P\left(S\left(G_{1} \mid D_{n}\right)>S\left(G_{2} \mid D_{n}\right)\right) \rightarrow 1$, as $n \rightarrow \infty$. This proves the first part of the theorem.

## Part (ii)

In case of complete data, $\beta(\mathcal{G})=1$, we shall obtain a faster convergence rate of $n^{-1}$ for the difference $l\left(G \mid D_{n}\right)-l\left(G_{0} \mid D_{n}\right), G_{0} \subseteq G$, which will prove the second part (ii) of the claim.

We consider the sample average log-likelihood of the node $X_{i}$. Since $G_{0} \subseteq G$, we have that $P a_{i}=P a_{i}^{0} \cup Y$ for some $Y \subset\left\{X_{i}\right\}_{i=1}^{N}, Y \cap P a_{i}^{0}=\emptyset$. Observe that in $G_{0}, X_{i}$ cannot have descendants in $Y$. Indeed, if there is a directed path $X_{i}$ to $X_{s} \in Y$ in $G_{0}$, this path cannot be in $G$ also, for otherwise one would have the loop $X_{i}$ to $X_{s}$ to $X_{i}$ in $G$ and $G$ would not be a DAG. But if there is a path in $G_{0}$ that is not in $G$, then $G_{0} \nsubseteq G$, a contradiction. Therefore, by LMP, $X_{i}$ and $Y$ are independent given $P a_{i}^{0}$.

In the usual notation, for $j \in P a_{i}^{0}, m \in Y$ and $k \in X_{i}, \hat{\theta}_{i, k j m}$ denotes the estimator of $P\left(X_{i}=k \mid P a_{i}^{0}=j, Y=m\right)$ and $\hat{\theta}_{i, j m}$ is the estimator of $P\left(P a_{i}^{0}=j, Y=m\right)$. We start with the expression

$$
\begin{align*}
& l\left(X_{i} \mid P a_{i}, D_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, D_{n}\right) \\
& \quad=\sum_{m \in Y} \sum_{j \in P a_{i}^{0}} \sum_{k \in X_{i}} \hat{\theta}_{i, j m} \hat{\theta}_{i, k j m} \log \left(\hat{\theta}_{i, k j m}\right)-\sum_{j \in P a_{i}^{0}} \sum_{k \in X_{i}} \hat{\theta}_{i, j} \hat{\theta}_{i, k j} \log \left(\hat{\theta}_{i, k j}\right) \\
& \quad=\sum_{j \in P a_{i}^{0}} \sum_{k \in X_{i}} \hat{\theta}_{i, j} \Delta_{i, k j} \tag{7.4}
\end{align*}
$$

where

$$
\Delta_{i, k j} \equiv \sum_{m \in Y} \frac{\hat{\theta}_{i, j m}}{\hat{\theta}_{i, j}} \hat{\theta}_{i, k j m} \log \left(\hat{\theta}_{i, k j m}\right)-\hat{\theta}_{i, k j} \log \left(\hat{\theta}_{i, k j}\right)
$$

By definition

$$
\hat{\theta}_{i, j m}=\frac{n_{i, j m}}{\sum_{j^{\prime}, m^{\prime}} n_{i, j^{\prime} m^{\prime}}}, \hat{\theta}_{i, j}=\frac{n_{i, j}}{\sum_{j^{\prime}} n_{i, j^{\prime}}} \text { and } \hat{\theta}_{i, k j m}=\frac{n_{i, k j m}}{n_{i, j m}}
$$

By the sample completeness, we have $\sum_{m} n_{i, j m}=n_{i, j}$ and $\sum_{m} n_{i, k j m}=n_{i, k j}$, implying

$$
\begin{equation*}
\sum_{m \in Y} \frac{\hat{\theta}_{i, j m}}{\hat{\theta}_{i, j}} \hat{\theta}_{i, k j m}=\frac{\sum_{j^{\prime}} n_{i, j^{\prime}}}{\sum_{j^{\prime}, m} n_{i, j^{\prime} m}} \frac{\sum_{m} n_{i, k j m}}{n_{i, j}}=\frac{n_{i, k j}}{n_{i, j}}=\hat{\theta}_{i, k j} \tag{7.5}
\end{equation*}
$$

If we set $\gamma_{m}=\hat{\theta}_{i, j m} / \hat{\theta}_{i, j}$, then $\sum_{m} \gamma_{m}=1$ and $\sum_{m} \gamma_{m} \hat{\theta}_{i, k j m}=\hat{\theta}_{i, k j}$. However, the latter are not guaranteed in incomplete settings because then we may have $\sum_{m} n_{i, j m}<n_{i, j}$ and(or) $\sum_{m} n_{i, k j m}<n_{i, k j}$.

We can now apply Lemma 7.3 with $\gamma_{m}, \hat{\theta}_{i, k j m}$ and $\hat{\theta}_{i, k j}$ to infer that $\Delta_{i, k j}=$ $O_{p}\left(n^{-1}\right)$ and

$$
\begin{equation*}
l\left(X_{i} \mid P a_{i}, D_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, D_{n}\right)=O_{p}\left(n^{-1}\right) \tag{7.6}
\end{equation*}
$$

Therefore $l\left(G \mid D_{n}\right)-l\left(G_{0} \mid D_{n}\right)=O_{p}\left(n^{-1}\right)$ holds for all $G$ such that $G_{0} \subseteq G$.
If both $G_{1}$ and $G_{2}$ contain $G_{0}$, it follows that $l\left(G_{2} \mid D_{n}\right)-l\left(G_{1} \mid D_{n}\right)=O_{p}\left(n^{-1}\right)$. Therefore, the condition $n \lambda_{n} \rightarrow \infty$ implies that the positive sequence $\lambda_{n}\left(h\left(G_{2}\right)-\right.$ $\left.h\left(G_{1}\right)\right)$ will overcome the likelihood difference $l\left(G_{2} \mid D_{n}\right)-l\left(G_{1} \mid D_{n}\right)$ with probability approaching 1 , which concludes the second part (ii) of the theorem.

Remark 7.1. (7.6) holds even for incomplete samples $D_{n}$ if they satisfy the property: $\left(X_{i}, P a_{i}\right)$ is complete in $D_{n}$ whenever $\left(X_{i}, P a_{i}^{0}\right)$ is complete, or equivalently, $l\left(X_{i} \mid P a_{i}, D_{n}\right)$ and $l\left(X_{i} \mid P a_{i}^{0}, D_{n}\right)$ are calculated upon one and the same subsample of $D_{n}$. Lemma 7.3 is applicable in such cases because we still have $\sum_{m} n_{i, j m}=n_{i, j}$ and $\sum_{m} n_{i, k j m}=n_{i, k j}$, and consequently, $\sum_{m} \gamma_{m}=1$ and $\sum_{m} \gamma_{m} \hat{\theta}_{i, k j m}=\hat{\theta}_{i, k j}$. Interestingly, if $Z$ is MCAR with $P\left(Z_{Y} \mid Z_{i}=1, Z_{P a_{i}^{0}}=\right.$ $1)=a \in(0,1)$, then $\sum_{m} n_{i, j m} / n_{i, j} \rightarrow_{p} a$ and $\sum_{m} n_{i, k j m} / n_{i, k j} \rightarrow_{p} a$, and consequently $\sum_{m} \gamma_{m} \rightarrow_{p} 1$ and $\sum_{m} \gamma_{m} \hat{\theta}_{i, k j m}-\hat{\theta}_{i, k j} \rightarrow_{p} 0$; this however is not enough for the claim in Lemma 7.3 to hold.

## Part (iii)

The last part of the theorem claims the necessity of condition (i) in case of incomplete sample that also satisfies Condition 4.1. Let $i$ be an unique node index of $G \in \mathcal{G}$ for which the condition holds, that is, for $j \neq i, P a_{j}=P a_{j}^{0}$, but $P a_{i} \backslash P a_{i}^{0} \neq \emptyset$. Without loss of generality we may assume that $h(G)=h\left(G_{0}\right)+1$ and that $D_{n}$ is complete with respect to $\left(X_{i}, P a_{i}^{0}\right)$. We shall show that for $\lambda_{n}$ such that $\underline{\lim } \sqrt{n} \lambda_{n}<\infty$

$$
\begin{equation*}
\varlimsup_{n \rightarrow \infty} P\left(l\left(X_{i} \mid P a_{i}, D_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, D_{n}\right)-\lambda_{n}>0\right)>0 \tag{7.7}
\end{equation*}
$$

which is equivalent to $S$ to be inconsistent.

Let $\tilde{D}_{n}=\left\{\tilde{x}^{t}\right\}_{t=1}^{\tilde{n}}$ be the $\tilde{n}$-subsample of $D_{n}$ for which $\left(X_{i}, P a_{i}\right)$ is observed. Then we have $l\left(X_{i} \mid P a_{i}, D_{n}\right)=l\left(X_{i} \mid P a_{i}, \tilde{D}_{n}\right)$. Note that $\tilde{n}$ is random and $\tilde{n} \sim$ $\operatorname{Binom}(a, n)$, for $a=P\left(Z_{P a_{i}}=1 \mid Z_{i}=1, Z_{P a_{i}^{0}}=1\right) \in(0,1)$, by Condition 4.1.

For every probability table $\theta$, we denote

$$
l\left(X_{i} \mid P a_{i}^{0}, \theta, \tilde{D}_{n}\right) \equiv \frac{1}{\tilde{n}} \sum_{t=1}^{\tilde{n}} l\left(\tilde{x}^{t} \mid P a_{i}^{0}, \theta\right)
$$

and

$$
l\left(X_{i} \mid P a_{i}^{0}, \theta, D_{n}\right) \equiv \frac{1}{n} \sum_{s=1}^{n} l\left(x^{s} \mid P a_{i}^{0}, \theta\right)
$$

where

$$
l\left(x \mid P a_{i}^{0}, \theta\right) \equiv \sum_{j \in P a_{i}^{0}} \sum_{k \in X_{i}} 1_{x_{i}=k, p a_{i}^{0}=j} \log \left(\theta_{i, k j}\right) .
$$

We have

$$
l\left(X_{i} \mid P a_{i}^{0}, D_{n}\right)=l\left(X_{i} \mid P a_{i}^{0}, \hat{\theta}_{n}, D_{n}\right), \text { for } \hat{\theta}_{n}=\arg \max _{\theta} l\left(X_{i} \mid P a_{i}^{0}, \theta, D_{n}\right)
$$

and

$$
l\left(X_{i} \mid P a_{i}^{0}, \tilde{D}_{n}\right)=l\left(X_{i} \mid P a_{i}^{0}, \tilde{\theta}_{n}, \tilde{D}_{n}\right), \text { for } \tilde{\theta}_{n}=\arg \max _{\theta} l\left(X_{i} \mid P a_{i}^{0}, \theta, \tilde{D}_{n}\right)
$$

Next, we show that the convergence rate of the difference $l\left(X_{i} \mid P a_{i}^{0}, \hat{\theta}_{n}, D_{n}\right)-$ $l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, D_{n}\right)$ is $n^{-1}$. The function $f(\theta) \equiv l\left(X_{i} \mid P a_{i}^{0}, \theta, D_{n}\right)$ has continuous first and second derivatives in a neighborhood of $\theta_{0}$. Since $\left.\frac{\partial f}{\partial \theta}\right|_{\hat{\theta}_{n}}=0(f$ has a maximum at $\hat{\theta}_{n}$, the Taylor's expansion of $f$ at $\theta=\hat{\theta}_{n}$ is

$$
f\left(\theta_{0}\right)=f\left(\hat{\theta}_{n}\right)+\left.0.5\left(\theta_{0}-\hat{\theta}_{n}\right)^{T} \frac{\partial^{2} f}{\partial \theta \partial \theta^{T}}\right|_{\theta_{n}^{*}}\left(\theta_{0}-\hat{\theta}_{n}\right), \text { for } \theta_{n}^{*} \in\left[\theta_{0}, \hat{\theta}_{n}\right]
$$

Moreover, the Hessian at $\theta_{0}$ is bounded because $\theta_{0}$ is bounded away from 0 and $\theta_{n}^{*} \rightarrow_{p} \theta_{0}$. Hence

$$
\left.\frac{\partial^{2} f}{\partial \theta \partial \theta^{T}}\right|_{\theta_{n}^{*}}=\left.\frac{\partial^{2} f}{\partial \theta \partial \theta^{T}}\right|_{\theta_{0}}+O_{p}(1)
$$

Because $\left\|\theta_{0}-\hat{\theta}_{n}\right\|^{2}=O_{p}\left(n^{-1}\right)$, we infer

$$
\sqrt{n}\left(l\left(X_{i} \mid P a_{i}^{0}, \hat{\theta}_{n}, D_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, D_{n}\right)\right)=O_{p}\left(n^{-1 / 2}\right)
$$

Similarly we have

$$
\begin{equation*}
\sqrt{n}\left(l\left(X_{i} \mid P a_{i}^{0}, \tilde{\theta}_{n}, \tilde{D}_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, \tilde{D}_{n}\right)\right)=O_{p}\left(n^{-1 / 2}\right) \tag{7.8}
\end{equation*}
$$

Due to the MCAR assumption, $Z_{P a_{i} \backslash P a_{i}^{0}}$ is independent of $\left(X_{i}, P a_{i}^{0}\right)$. This and Condition 4.1 imply that $\tilde{D}_{n}$ is obtained from $D_{n}$ by random draws without
replacements. Therefore, we can apply Lemma 7.4 to the set $\left\{l\left(x^{s} \mid P a_{i}^{0}, \theta_{0}\right)\right\}_{s=1}^{n}$ of i.i.d. random variables and its subset $\left\{l\left(\tilde{x}^{t} \mid P a_{i}^{0}, \theta_{0}\right)\right\}_{t=1}^{\tilde{n}}$. We thus infer

$$
\begin{equation*}
\sqrt{n}\left(l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, \tilde{D}_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, D_{n}\right)\right) \rightarrow_{d} \mathcal{N}\left(0, \gamma \operatorname{Var}\left(l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}\right)\right)\right) \tag{7.9}
\end{equation*}
$$

where $\gamma=(1-a) / a>0$. Also note that $\operatorname{Var}\left(l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}\right)\right)>0$ by the identifiability of $G_{0}$.

Moreover, by Remark 7.1, (7.6) applied to $\tilde{D}_{n}$ yields

$$
\begin{equation*}
\sqrt{n}\left(l\left(X_{i} \mid P a_{i}, \tilde{D}_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, \tilde{D}_{n}\right)\right)=O_{p}\left(n^{-1 / 2}\right) \tag{7.10}
\end{equation*}
$$

Finally, we consider the difference implicated in (7.7)

$$
\sqrt{n}\left(l\left(X_{i} \mid P a_{i}, \tilde{D}_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, D_{n}\right)\right)-\sqrt{n} \lambda_{n}
$$

apply (7.10)

$$
=\sqrt{n}\left(l\left(X_{i} \mid P a_{i}^{0}, \tilde{\theta}_{n}, \tilde{D}_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, \hat{\theta}_{n}, D_{n}\right)\right)-\sqrt{n} \lambda_{n}+O_{p}\left(n^{-1 / 2}\right)
$$

then use (7.8)

$$
\begin{aligned}
& =\sqrt{n}\left(l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, \tilde{D}_{n}\right)-l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}, D_{n}\right)\right)-\sqrt{n} \lambda_{n}+O_{p}\left(n^{-1 / 2}\right) \\
& =: T_{n}+O_{p}\left(n^{-1 / 2}\right)
\end{aligned}
$$

Since $\underline{\lim } \sqrt{n} \lambda_{n}<\infty$, there is a subsequence $n^{\prime}$ such that $\lim \sqrt{n^{\prime}} \lambda_{n^{\prime}}=\lambda_{0}<\infty$ and for which, taking into account the convergence (7.9), we have

$$
T_{n^{\prime}} \rightarrow_{d} \mathcal{N}\left(-\lambda_{0}, \gamma \operatorname{Var}\left(l\left(X_{i} \mid P a_{i}^{0}, \theta_{0}\right)\right)\right) .
$$

Therefore $\lim P\left(T_{n^{\prime}}>0\right) \rightarrow 1-\Phi\left(\lambda_{0} / \sqrt{\gamma \operatorname{Var}\left(l\left(X \mid P a_{i}^{0}, \theta_{0}\right)\right)}\right)>0$, where $\Phi$ is the c.d.f. of the standard normal distribution. Hence (7.7) is verified and with this the proof of the theorem.

## Acknowledgements

This work was supported by NIH grant K99LM009477 from the National Library Of Medicine. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Library Of Medicine or the National Institutes of Health. The author thanks an anonymous reviewer whose comments greatly improved the paper and who suggested the 2-node simulated example in Section 5, and also Peter Salzman for many stimulating discussions.

## References

[1] Balov, N., Salzman, P. (2013). catnet: Categorical Bayesian Network Inference. R package version 1.13.8.
[2] Beinlich, I., Suermondth, G., Chavez, R., Cooper, G. (1989). The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks. In 2-nd European Conference on AI and Medicine.
[3] Buntine, W. (1996). A guide to the literature on learning graphical models. IEEE Transactions on Knowledge and Data Engineering, 8:195-210
[4] Chickering, D. M. (2002). Optimal structure identification with greedy search. Journal of Machine Learning Research, 3:507-554. MR1991085
[5] Cooper, G., Herskovits, E. (1992). A Bayesian method for the induction of probabilistic networks from data. Machine Learning, 9(4):309-347.
[6] Geiger, D., Heckerman, D., King, H. and Meek, C. (2001). Stratified exponential families: Graphical models and model selection. The Annals of Statistics, 29(2):505-529. MR1863967
[7] Haughton, Dominique M. A. (1988). On the choice of a model to fit data from an exponential family. The Annals of Statistics, 16(1):342-355. MR0924875
[8] Lauritzen, S.,L. (1995). The EM algorithm for graphical association models with missing data. Computational Statistics and Data Analysis, 19(2):191-201.
[9] Pearl, J. (1988). Probabilistic reasoning in intelligent systems: networks of plausible reasoning. Morgan Kaufmann, San Mateo, CA. MR0965765
[10] Schwartz, G. (1978). Estimating the dimension of a model. The Annals of Statistics, 6(2):461-464. MR0468014
[11] Spiegelhalter, D., Dawid, A., Lauritzen, S., Cowell, R. (1993). Bayesian analysis in expert systems. Statistical Science, 8(3):219-247. MR1243594
[12] van der Vaart, A.W. (2007). Asymptotic Statistics. Cambridge University Press. MR1652247
[13] Verma, T. and Pearl, J. (1990). Equivalence and synthesis of causal models. In Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence, 255-268.

