


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

A Bayesian Hierarchical score for structure learning from related data sets

Laura Azzimonti*, Giorgio Corani, Marco Scutari

Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), USI/SUPSI, Lugano, Switzerland

ARTICLE INFO

Article history:

Received 15 January 2021

Received in revised form 2 November 2021

Accepted 30 November 2021

Available online xxxx

Keywords:

Bayesian networks

Structure learning

Hierarchical priors

Dirichlet mixtures

Network scores

ABSTRACT

Score functions for learning the structure of Bayesian networks in the literature assume that data are a homogeneous set of observations; whereas it is often the case that they comprise different related, but not homogeneous, data sets collected in different ways. In this paper we propose a new Bayesian Dirichlet score, which we call Bayesian Hierarchical Dirichlet (BHD). The proposed score is based on a hierarchical model that pools information across data sets to learn a single encompassing network structure, while taking into account the differences in their probabilistic structures. We derive a closed-form expression for BHD using a variational approximation of the marginal likelihood, we study the associated computational cost and we evaluate its performance using simulated data. We find that, when data comprise multiple related data sets, BHD outperforms the Bayesian Dirichlet equivalent uniform (BDeu) score in terms of reconstruction accuracy as measured by the Structural Hamming distance, and that it is as accurate as BDeu when data are homogeneous. This improvement is particularly clear when either the number of variables in the network or the number of observations is large. Moreover, the estimated networks are sparser and therefore more interpretable than those obtained with BDeu thanks to a lower number of false positive arcs.

© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Investigating challenging problems at the forefront of science increasingly requires large amounts of data that can only be gathered through collaborations between several institutions. This naturally leads to heterogeneous data sets that are in fact the collation of related, but not identical, subsets of data that will necessarily differ in the details of how they are collected. Examples can be found in multi-centre clinical trials, in which protocols are applied in slightly different ways to different patient populations [1,2]; population genetics, which studies the architecture of phenotypic traits across populations and their evolution [3–6]; ecology and environmental sciences, which produce different patterns of measurement errors and limitations in different environments [7–9]. A common goal in analysing these complex data is to construct a mechanistic model that elucidates the interplay between different elements under investigation, either as a step towards building a causal model or to perform accurate prediction from a purely probabilistic perspective.

The task of efficiently modelling such related data sets is usually tackled by hierarchical models [10], which pool the information common to the different subsets of the data while encoding the information that is specific to each subset. For instance multilevel regression models estimate the conditional distribution of the response variables in these cases.

* Corresponding author.

E-mail address: laura.azzimonti@idsia.ch (L. Azzimonti).

<https://doi.org/10.1016/j.ijar.2021.11.013>

0888-613X/© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Bayesian networks (BNs) [11] provide a rigorous approach for modelling joint distributions, by representing variables as nodes and probabilistic dependencies as arcs in a graph. They can be used for both causal and predictive modelling. To the best of our knowledge, however, no method has been proposed in the literature to combine these two approaches to learn a single BN structure from a set of related data sets and get the best of both worlds.

Available methods focus on learning an ensemble of BNs that have similar structures by penalising differences in their arc sets [12,13]. Parameter learning from related data sets has been investigated in [14] for Gaussian BNs and in [15] for discrete BNs. However, they only consider a naive Bayes structure and they initialise their hyperprior with maximum likelihood point estimates. On the other hand, [16] discusses how to answer causal queries by fusing observations from different domains, possibly characterised by different causal mechanisms. However, they assume the estimated graphs to be already available.

In this paper, we show instead how to learn the structure of a BN from related data sets, containing the same variables, by building on our previous work on parameter learning in [17]. The proposed approach is particularly suited to deal with multiple related data sets characterised by few observations per data set or by an unbalanced number of observations across data sets. In these settings, it is important to share information across data sets to obtain robust estimates of both the parameters and the structure of the BN. First, we briefly introduce BNs and hierarchical models in the context of discrete data as well as prior work on parameter learning from related data sets in Section 2. We propose a score function for related data sets in Section 3 and we study the associated computational complexity in both a theoretical and empirical way in Section 4. Then, we show an example of structure learning by means of BHD in Section 5 and we study its performance on different simulation studies in Section 6. Finally, we discuss our results and possible future research directions in Section 7.

2. Background and notation

Bayesian networks (BNs) are a class of graphical models that use a directed acyclic graph (DAG) \mathcal{G} to model a set of random variables $\mathbf{X} = \{X_1, \dots, X_N\}$: each node is associated with one $X_i \in \mathbf{X}$ and arcs represent direct dependence relationships. Graphical separation of two nodes implies the conditional independence of the corresponding random variables. In principle, there are many possible choices for the joint distribution of \mathbf{X} ; literature has focused mostly on discrete BNs [18], in which both \mathbf{X} and the X_i are categorical (multinomial) random variables. Other possibilities include Gaussian BNs and conditional linear Gaussian BNs [19], which include both discrete and Gaussian BNs as particular cases.

The task of learning a BN from a data set \mathcal{D} of n observations is performed in two steps in an inherently Bayesian fashion:

$$\underbrace{P(\mathcal{G}, \Theta | \mathcal{D})}_{\text{learning}} = \underbrace{P(\mathcal{G} | \mathcal{D})}_{\text{structure learning}} \cdot \underbrace{P(\Theta | \mathcal{G}, \mathcal{D})}_{\text{parameter learning}}, \quad (1)$$

where Θ are the parameters of \mathbf{X} . *Structure learning* consists in finding the DAG \mathcal{G} that encodes the dependence structure of the data. In this paper we will focus on score-based algorithms, which are typically heuristic search algorithms that use a goodness-of-fit score such as BIC [20] or the Bayesian Dirichlet equivalent uniform (BDeu) marginal likelihood [18] to find an optimal \mathcal{G} . *Parameter learning* involves the estimation of the parameters Θ given the DAG \mathcal{G} learned in the first step. Thanks to the Markov property, this step is computationally efficient because if the data are complete the *global distribution* of \mathbf{X} decomposes into

$$P(\mathbf{X} | \mathcal{G}) = \prod_{i=1}^N P(X_i | \Pi_{X_i}) \quad (2)$$

and the *local distribution* associated with each node X_i depends only on the configurations of its parents Π_{X_i} . If we estimate the parameters $\Theta = \{\Theta_{X_1}, \dots, \Theta_{X_N}\}$ in such a way that they are independent across local distributions, parameter learning simplifies into a collection of low-dimensional estimation problems for the Θ_{X_i} associated with each $X_i | \Pi_{X_i}$ given the data available for those variables.

2.1. Classic multinomial-Dirichlet parameterisation

In the case of discrete BNs, we assume that each $X_i | \Pi_{X_i}$ follows a categorical distribution for each configuration of Π_{X_i} . Hence the parameters of $X_i | \Pi_{X_i}$ are the conditional probabilities $\theta_{X_i | \Pi_{X_i}} = \{\theta_{X_i | j}, j = 1, \dots, |\Pi_{X_i}|\}$, whose k th element corresponds to $P(X_i = k | \Pi_{X_i} = j)$, for which we assume a conjugate Dirichlet prior:

$$\begin{aligned} \theta_{X_i | \Pi_{X_i}} | \boldsymbol{\alpha}_i &\sim \text{Dirichlet}(\boldsymbol{\alpha}_i) \\ X_i | \Pi_{X_i}, \theta_{X_i | \Pi_{X_i}} &\sim \text{Categorical}(\theta_{X_i | \Pi_{X_i}}), \end{aligned} \quad (3)$$

where $\boldsymbol{\alpha}_i = \{\alpha_{ijk}, j = 1, \dots, |\Pi_{X_i}|; k = 1, \dots, |X_i|\}$, with $i = 1, \dots, N$, is a hyperparameter vector defined over a simplex with sum $\sum_{j,k} \alpha_{ijk} = s_i > 0$. The posterior estimator of $\theta_{X_i | \Pi_{X_i}}$ is:

$$[\widehat{\theta}_{X_i|j}]_k = \frac{\alpha_{ijk} + n_{ijk}}{\alpha_{ij} + n_{ij}}, \quad \text{where } n_{ij} = \sum_k n_{ijk}, \quad \alpha_{ij} = \sum_k \alpha_{ijk}, \quad (4)$$

and n_{ijk} represents the number of observations for which $X_i = k$ and $\Pi_{X_i} = j$. It is common to set $\alpha_{ijk} = s_i / (|X_i| |\Pi_{X_i}|)$ with the same *imaginary sample size* $s_i = s$ for all X_i .

In the context of structure learning, we have $P(\mathcal{G}|\mathcal{D}) \propto P(\mathcal{D}|\mathcal{G})P(\mathcal{G})$ and we can use $P(\mathcal{D}|\mathcal{G})$ as a score function. (Implicitly, we are assuming a uniform prior for \mathcal{G} $P(\mathcal{G}) \propto 1$ by disregarding it while still searching for the *maximum a posteriori* DAG.) Assuming *positivity* ($\theta_{X_i|\Pi_{X_i}} > 0$), *parameter independence* (columns of $\theta_{X_i|\Pi_{X_i}}$ associated with different parent configurations are independent), *parameter modularity* ($\theta_{X_i|\Pi_{X_i}}$ associated with different nodes are independent) and *complete data*, [18] derived a closed form expression for $P(\mathcal{D}|\mathcal{G})$ known as the *Bayesian Dirichlet* (BD) family of scores:

$$\text{BD}(\mathcal{G}, \mathcal{D}; \boldsymbol{\alpha}) = \prod_{i=1}^N \text{BD}(X_i | \Pi_{X_i}; \boldsymbol{\alpha}_i) = \prod_{i=1}^N \prod_{j=1}^{|\Pi_{X_i}|} \left[\frac{\Gamma(\alpha_{ij})}{\Gamma(\alpha_{ij} + n_{ij})} \prod_{k=1}^{|X_i|} \frac{\Gamma(\alpha_{ijk} + n_{ijk})}{\Gamma(\alpha_{ijk})} \right]. \quad (5)$$

Choosing again $\alpha_{ijk} = s / (|X_i| |\Pi_{X_i}|)$ gives the *Bayesian Dirichlet equivalent uniform* (BDeu) score. A default value of $s = 1$ has been recommended by [21]. Although commonly accepted, these assumptions can have a marked impact on the accuracy of the learned structures [22], especially for sparse data that are likely to lead to violations of the positivity assumption [23]. Furthermore, they are taken to represent lack of prior knowledge, and result in BDeu giving the same score value to BNs in the same equivalence class (*score-equivalence* [24]). Equivalence classes are characterised by the skeleton of \mathcal{G} (its underlying undirected graph) and its v -structures (patterns of arcs of the type $X_j \rightarrow X_i \leftarrow X_k$, with no arc between X_j and X_k), and group DAGs that encode the same global distribution.

2.2. Hierarchical multinomial-Dirichlet parameterisation for related data sets

The classic Multinomial-Dirichlet model in (3) can be extended to handle related data sets by treating it as a particular case of the hierarchical Multinomial-Dirichlet (hierarchical MD) model presented in [17]. For this purpose, we introduce an auxiliary variable F which identifies the $|F|$ related data sets. Assuming that the data sets contain the same variables and that F is always observed, we can learn a BN with a common structure \mathcal{G} but with different parameter estimates for each related data set.

For simplicity, we apply the hierarchical model independently to each local distribution to estimate the joint distribution of (X_i, Π_{X_i}) conditional on F , $\theta_{X_i, \Pi_{X_i} | F} = \{\theta_{X_i, \Pi_{X_i}}^f, f = 1, \dots, |F|\}$, by pooling information between different data sets. The resulting hierarchical model is shown in the top panel of Fig. 1. Specifically, for each node X_i we assume $\boldsymbol{\alpha}_i$ to be a latent random vector and we add a Dirichlet hyperprior to make $\theta_{X_i, \Pi_{X_i}}^f$ a mixture of Dirichlet distributions:

$$\begin{aligned} \boldsymbol{\alpha}_i | s_i, \boldsymbol{\alpha}_{0,i} &\sim s_i \cdot \text{Dirichlet}(\boldsymbol{\alpha}_{0,i}), \\ \theta_{X_i, \Pi_{X_i}}^f | \boldsymbol{\alpha}_i &\sim \text{Dirichlet}(\boldsymbol{\alpha}_i) \quad f = 1, \dots, |F|, \\ X_i, \Pi_{X_i} | F = f, \theta_{X_i, \Pi_{X_i}}^f &\sim \text{Categorical}(\theta_{X_i, \Pi_{X_i}}^f) \quad f = 1, \dots, |F|, \end{aligned} \quad (6)$$

where this time $\boldsymbol{\alpha}_i = \{\alpha_{ijk}\}$ is a latent random vector defined over a simplex of dimension $|X_i| |\Pi_{X_i}| - 1$ with sum s_i . The new hyperparameters of this model are the imaginary sample size s_i and the parameter vector $\boldsymbol{\alpha}_{0,i}$, which in turn is defined over a simplex with sum $s_{0,i}$. The two parameters $s_{0,i}$ and s_i control respectively the concentration of the $\boldsymbol{\alpha}_i$ random vectors around the discrete distribution $\boldsymbol{\alpha}_{0,i}/s_{0,i}$ and the variance of $\theta_{X_i, \Pi_{X_i}}^f$, with $f = 1, \dots, |F|$, around the normalised random vector $\boldsymbol{\alpha}_i/s_i$. Larger values of s_i yield $\theta_{X_i, \Pi_{X_i}}^f$ that are more similar to each other, while larger values of $s_{0,i}$ provide $\theta_{X_i, \Pi_{X_i}}^f$ centred around $\boldsymbol{\alpha}_{0,i}/s_{0,i}$. For example, if the $\boldsymbol{\alpha}_{0,i}$ are symmetrical vectors (i.e. $\alpha_{0,ijk} = \alpha_{0,i}$), then large values of $s_{0,i}$ provide $\theta_{X_i, \Pi_{X_i}}^f$ with symmetrical expected value.

In the following we will drop both hyper-parameters from the notation for brevity.

The marginal posterior distribution for $\theta_{X_i, \Pi_{X_i}}^f$ is not analytically tractable, as noted in [17]. However, the posterior average can be compactly expressed as:

$$[\widehat{\theta}_{X_i, \Pi_{X_i}}^f]_{jk} = \frac{E[\alpha_{ijk}] + n_{ijk}^f}{s_i + n_i^f}, \quad \text{where } n_i^f = \sum_{jk} n_{ijk}^f. \quad (7)$$

$E[\alpha_{ijk}]$ represents the posterior average of α_{ijk} ; it cannot be written in closed form but can be approximated using variational inference [25,26]. The resulting $\widehat{\theta}_{X_i, \Pi_{X_i}}^f$ are data-set-specific but depend on all the available data via the *partial pooling* [10] of the information present in the $|F|$ related data sets, thanks to the shared $E[\alpha_{ijk}]$ term. On the one hand, this produces more reliable estimates for sparse data and for related data sets with unbalanced sample sizes [27]. On the other hand, the

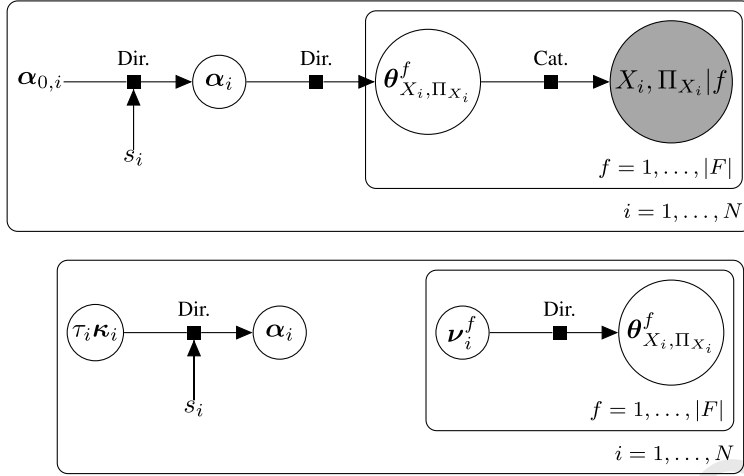


Fig. 1. Directed factor graphs representing hierarchical Multinomial-Dirichlet model for related data sets (top panel) and its variational approximation (bottom panel). Cat. and Dir. represent respectively Categorical and Dirichlet distributions.

prior in (6) violates the parameter independence assumption, leading to a marginal likelihood that does not decompose over parent configurations and that is not score-equivalent. The prior is specified on (X_i, Π_{X_i}) , as opposed to $X_i | \Pi_{X_i}$, which is only later computed from the joint distribution. As a result, the distribution of $(X_i, \Pi_{X_i} | F)$ is different from the product of the distributions of $(X_i | \Pi_{X_i}, F)$ and $(\Pi_{X_i} | F)$ because $(X_i, \Pi_{X_i} | F)$ and $(\Pi_{X_i} | F)$ are estimated by applying the hierarchical model separately to two different sets of variables, thus pooling the available information differently.

3. Structure learning from related data sets

In this section we derive the marginal likelihood score associated with the hierarchical model in (6) to implement structure learning from related data sets containing the same variables. As the hierarchical model is not analytically tractable, we approximate the associated posterior distribution with the product of two independent distributions by means of variational inference. The approximate variational model, shown in the bottom panel of Fig. 1, is the following:

$$\begin{aligned} \alpha_i | s_i, \tau_i, \kappa_i &\sim s_i \cdot \text{Dirichlet}(\tau_i \kappa_i), \\ \theta_{X_i, \Pi_{X_i}}^f | \nu_i^f &\sim \text{Dirichlet}(\nu_i^f) \quad f = 1, \dots, |F|, \end{aligned} \quad (8)$$

where $\nu_i^f = \{\nu_{ijk}^f\}$, $\kappa_i = \{\kappa_{ijk}\}$ with $i = 1, \dots, N$; $j = 1, \dots, |\Pi_{X_i}|$; $k = 1, \dots, |X_i|$ and $f = 1, \dots, |F|$; $\sum_{jk} \kappa_{ijk} = 1$ and $\tau_i \in \mathbb{R}^+$ for $i = 1 \dots N$. These parameters are estimated from the available data by minimising the Kullback-Leibler divergence between the exact posterior distribution p and its variational approximation q , as described in [17]. The algorithm used to estimate the variational parameters is summarised in Appendix A.

Since F is assumed to be the parent of any node in the network and to be always observed, we treat it as an input variable in a conditional Bayesian network [11, Section 5.6] and we do not explicitly assign it a distribution. Therefore, the auxiliary variable F will not influence the score.

The variational model (8) is similar to the original hierarchical MD model (6), but it removes the dependence between $\theta_{X_i, \Pi_{X_i}}^f$ and α_i thus making it possible to derive in closed form the variational approximation of the marginal likelihood $P(\mathcal{D} | F, \mathcal{G})$.

Lemma 1. Given $|F|$ complete and related data sets $\mathcal{D} = \{\mathcal{D}_f, f = 1, \dots, |F|\}$, under the assumption that the related data sets have the same dependence structure \mathcal{G} and that each local distribution follows the hierarchical MD (6) with positive parameters, the variational approximation of the marginal likelihood of the data $P(\mathcal{D} | F, \mathcal{G})$ is

$$q(\mathcal{D} | F, \mathcal{G}) = \prod_{i=1}^N \prod_{f=1}^{|F|} \prod_{j=1}^{|\Pi_{X_i}|} \left[\frac{\Gamma(s_i \widehat{\kappa}_{ij})}{\Gamma(s_i \widehat{\kappa}_{ij} + n_{ij}^f)} \prod_{k=1}^{|X_i|} \frac{\Gamma(s_i \widehat{\kappa}_{ijk} + n_{ijk}^f)}{\Gamma(s_i \widehat{\kappa}_{ijk})} \right], \quad (9)$$

where $n_{ij}^f = \sum_k n_{ijk}^f$, $\widehat{\kappa}_{ij} = \sum_k \widehat{\kappa}_{ijk}$ and $s_i \widehat{\kappa}_{ijk}$ represents the posterior average of α_{ijk} under the variational model (8).

Proof. Under the hierarchical model (6), the conditional distribution of X_i given $\Pi_{X_i} = j$ and $F = f$ is a categorical distribution with parameters $\theta_{X_i|j}^f$ whose k th element is $[\theta_{X_i, \Pi_{X_i}}^f]_{jk} / \sum_{\tilde{k}} [\theta_{X_i, \Pi_{X_i}}^f]_{j\tilde{k}}$. The distribution of $\theta_{X_i|j}^f$ is a Dirichlet distribution with parameter $\alpha_{i|j}$, whose k -th element is $[\alpha_{i|j}]_k = \alpha_{ijk} / \sum_{\tilde{k}} \alpha_{ij\tilde{k}}$.

Under the variational model (8), the approximate posterior distribution of $\theta_{X_i|j}^f$ is a Dirichlet distribution with parameters v_{ij}^f whose k th element is v_{ijk}^f . Since the parameter v_{ijk}^f is estimated *a posteriori* as $\widehat{v}_{ijk}^f = s_i \widehat{\kappa}_{ijk} + n_{ijk}^f$ (see [17] for a detailed derivation), we can approximate the prior distribution of $\theta_{X_i|j}^f$ with $q(\theta_{X_i|j}^f | \mathcal{G}) = \text{Dirichlet}(s_i \widehat{\kappa}_{ij})$ because of conjugacy.

The variational approximation of the conditional distribution satisfies *independence between related data sets*. Moreover, given a data set $F = f$, both *parameter modularity* and *parameter independence* are satisfied. Thus,

$$q(\mathcal{D} | F, \mathcal{G}) = \iint \prod_{f=1}^{|F|} \prod_{i=1}^N \prod_{j=1}^{|\Pi_{X_i}|} q(X_i | \Pi_{X_i} = j, F = f, \theta_{X_i|j}^f, \alpha_{i|j}, \mathcal{G}) q(\theta_{X_i|j}^f, \alpha_{i|j} | \mathcal{G}) d\theta_{X_i|j}^f d\alpha_{i|j}.$$

Thanks to the independence between $\theta_{X_i|j}^f$ and $\alpha_{i|j}$ induced by the variational model and the fact that $\int q(\alpha_{i|j} | \mathcal{G}) d\alpha_{i|j} = 1$, we obtain

$$q(\mathcal{D} | F, \mathcal{G}) = \prod_{f=1}^{|F|} \prod_{i=1}^N \prod_{j=1}^{|\Pi_{X_i}|} \int q(X_i | \Pi_{X_i} = j, F = f, \theta_{X_i|j}^f, \mathcal{G}) q(\theta_{X_i|j}^f | \mathcal{G}) d\theta_{X_i|j}^f,$$

which has the same form as the marginal likelihood of the classic Multinomial-Dirichlet model but with $s_i \widehat{\kappa}_{ijk}$ as the parameter of the Dirichlet distribution. The approximate marginal likelihood can be thus written as (9).

Note that the marginal likelihood (9) has the same form as the classic BD score (5), with α_{ijk} replaced by $s_i \widehat{\kappa}_{ijk}$, which represents the posterior average of α_{ijk} under the hierarchical variational model. The posterior average is shared between different related data sets, thus inducing a pooling effect that makes $\theta_{X_i, \Pi_{X_i}}^f$ and α_i dependent once more.

From Lemma 1, we define the approximated Bayesian hierarchical Dirichlet score as

$$\text{BHD}(\mathcal{G}, \mathcal{D} | F) = q(\mathcal{D} | F, \mathcal{G}).$$

The proposed BHD score can be factorised over the nodes, *i.e.*,

$$\text{BHD}(\mathcal{G}, \mathcal{D} | F) = \prod_{i=1}^N \text{BHD}(X_i | \Pi_{X_i}, F),$$

and can be used to learn a common structure for all related data sets, taking into account potential differences in the probabilistic relationships between variables.

4. Computational complexity

Estimating the BHD score in (9) is more complex than estimating the classic BD score in (5) because the latter is available in closed-form but the former is not. In this section we will assess the computational complexity of $\text{BHD}(X_i | \Pi_{X_i}, F)$ and $\text{BD}(X_i | \Pi_{X_i}, \alpha_i)$.

In the case of $\text{BD}(\text{eu})$, the score in (5) is a closed-form function of the counts n_{ijk} which are tallied from $\{X_i, \Pi_{X_i}\}$ in $O(n(1 + N_{\Pi_{X_i}}))$ time, where $N_{\Pi_{X_i}}$ is the dimension of the parent set Π_{X_i} . Assuming that each variable takes at most l values, there are $l^{1+N_{\Pi_{X_i}}}$ counts. Hence both computing the marginal counts n_{ij} and multiplying/summing up all the terms in (5) take $O(l^{1+N_{\Pi_{X_i}}})$ time. The overall computational complexity of computing $\text{BD}(X_i | \Pi_{X_i}; \alpha_i)$ then is

$$O\left(n(1 + N_{\Pi_{X_i}}) + l^{1+N_{\Pi_{X_i}}}\right). \quad (10)$$

As for BHD, the counts n_{ijk}^f are tallied from $\{X_i, \Pi_{X_i}, F\}$ in $O(n(2 + N_{\Pi_{X_i}}))$ time because of the auxiliary variable F . Computing the marginal counts n_{ij}^f and multiplying/summing up all the terms in (9) takes $O(|F|l^{1+N_{\Pi_{X_i}}})$ time.

The increased complexity of $\text{BHD}(X_i | \Pi_{X_i}, F)$, however, comes from the algorithm used to estimate the variational parameters τ_i , v_{ijk}^f and κ_{ijk} . The variational algorithm is derived in [17] and is reproduced for convenience as Algorithm 1 in Appendix A.

The update of \widehat{v}_{ijk}^f in step 1 of Algorithm 1 takes $O(|F|l^{1+N_{\Pi X_i}})$ since there are as many \widehat{v}_{ijk}^f as there are n_{ijk}^f .

Each update of $\widehat{\tau}_i$ in step 3a requires the computation of $\partial \mathcal{L} / \partial \tau_i$ and $\partial^2 \mathcal{L} / \partial \tau_i^2$: both are closed-form functions that sum over the indices j and k of $\widehat{\kappa}_{ijk}$. Updating the parameter $\widehat{\tau}_i$ by means of (A.2) thus takes $O(l^{1+N_{\Pi X_i}})$.

Each update of $\widehat{\kappa}_{ijk}$ in step 3b requires the computation of $\partial \mathcal{L} / \partial \kappa_{ijk}$ and $\partial^2 \mathcal{L} / \partial \kappa_{ijk}^2$, which scale respectively as $|F|l^{1+N_{\Pi X_i}}$ and as $l^{1+N_{\Pi X_i}}$. Given the partial derivatives, the cost of updating the parameter $\widehat{\kappa}_{ijk}$ by means of (A.3) scales as the number of elements, that is, $l^{1+N_{\Pi X_i}}$. Thus, since there are $l^{1+N_{\Pi X_i}}$ terms $\widehat{\kappa}_{ijk}$, the computational cost of step 3b is $O((|F| + 3)l^{1+N_{\Pi X_i}})$.

Once we take into account the number of iterations m_1 and m_2 , we have that the overall computational complexity of computing $\text{BHD}(X_i | \Pi_{X_i}, F)$ is

$$\begin{aligned} & \underbrace{O(n(2 + N_{\Pi X_i}) + |F|l^{1+N_{\Pi X_i}})}_{\text{formula in (9)}} + \underbrace{O(m_1|F|l^{1+N_{\Pi X_i}})}_{\text{step 1}} + \\ & \quad + \underbrace{O(m_1m_2l^{1+N_{\Pi X_i}})}_{\text{step 3a}} + \underbrace{O(m_1m_2(|F| + 3)l^{1+N_{\Pi X_i}})}_{\text{step 3b}} = \\ & \quad = O(m_1m_2|F|l^{1+N_{\Pi X_i}}) \end{aligned} \quad (11)$$

If we compare (11) with (10), we can see that they are both linear in $l^{1+|\Pi X_i|}$, but the former also depends on the number of related data sets $|F|$. The computational cost of each step of the iterative procedure for computing BHD is thus comparable to the computational cost associated with learning a network by means of BDeu.

4.1. Empirical evaluation

In order to confirm the derived computational complexity of BHD, we evaluate the time needed to compute the BHD score for a single node as different parameters vary. We consider in particular:

- $N_{\Pi X_i} \in \{2, 3, 4, 6, 8, 10, 12, 14, 16\}$, where $N_{\Pi X_i}$ represents the size of the parent set;
- $|F| \in \{5, 10, 20, 40\}$, where $|F|$ represents the number of related data sets;
- $l = |X_i| \in \{2, 3, 4, 5\}$, where $|X_i|$ represents the number of states for each variable.

For each parameter combination, we sample 10 different parameter sets with the following methods:

- hier:** the parameters associated with each of the related data sets are sampled from a hierarchical Dirichlet distribution with imaginary sample size equal to 10 and with a parameter α_i sampled from a Dirichlet distribution with all $\alpha_{0,ijk} = 1$;
- iid:** the parameters associated with each of the related data sets are independently sampled from the same Dirichlet distribution with imaginary sample size equal to 10 and uniform α_i .

For each parameter set, we sample $|F|$ related data sets comprising the same number of observations $n_f \in \{5000, 10000, 20000\}$.

Fig. 2 shows how the logarithm of the computational time varies as a function of the dimension of the parent set $N_{\Pi X_i}$ for different values of $|X_i|$. The logarithm of the computational time scales linearly in the dimension of the parent set for each value of $|X_i|$, with a slope that is proportional to the value $|X_i|$. A deviation from this behaviour is visible for small values of $N_{\Pi X_i}$ in the **iid** case, where the time needed to estimate the parameters is negligible compared to fixed computational costs like memory allocation.

The effect of $|F|$ on the computational times is weaker than that of $N_{\Pi X_i}$ and $|X_i|$, while the effect of the number of observations n_f is negligible.

To estimate the effect of $N_{\Pi X_i}$, $|X_i|$ and $|F|$ on the computational time, we estimated the parameters of the linear model

$$\log(\text{Time}) = \beta_0 + \beta_1 \mathbb{I}_{\text{iid}} + \beta_2 \log(|F|) + \beta_3 (1 + N_{\Pi X_i}) \log(|X_i|) + \epsilon,$$

corresponding to

$$\text{Time} = e^{\beta_0 + \beta_1 \mathbb{I}_{\text{iid}}} |F|^{\beta_2} |X_i|^{\beta_3 (1 + N_{\Pi X_i})} e^{\epsilon},$$

where ϵ represents the measurement error. All the parameter estimates are significantly different from zero (the associated p-values are smaller than 10^{-15}) and the model fits the computational times well ($R^2 = 0.95$). Moreover, the estimated parameters $\beta_0 = -6.20$, $\beta_1 = -0.84$, $\beta_2 = 0.71$ and $\beta_3 = 1.15$ are consistent with the theoretical computational complexity derived in (11).




Fig. 2. Boxplots of the logarithm of the computational time with different dimensions of the parent set $N_{\Pi_{X_i}}$ and number of states $|X_i|$, with parameters sampled with the hierarchical (top panel) or i.i.d (bottom panel) approach.

Table 1

Scores associated with each node and overall score obtained in the numerical example hill-climbing optimisation. Step 0 corresponds to the starting empty graph, step 5 corresponds to the estimated network. The scores updated at each step due to arc addition are highlighted in bold.

Step	X_1	X_2	X_3	X_4	X_5	Overall
0	-466.205	-528.450	-695.655	-642.506	-672.477	-3005.293
1	-466.205	-463.584	-695.655	-642.506	-672.477	-2940.427
2	-466.205	-463.584	-651.734	-642.506	-672.477	-2896.506
3	-466.205	-463.584	-651.734	-642.506	-666.052	-2890.081
4	-466.205	-463.584	-651.734	-642.506	-650.698	-2874.727
5	-466.205	-463.584	-651.734	-638.011	-650.698	-2870.232

5. Numerical example

We consider a simple example to illustrate the steps involved in learning a network structure with BHD and in estimating the associated parameters.

We consider in this example the set of variables X_1, \dots, X_5 , with $|X_i| = 2$ for $i = 1, \dots, 5$, and $|F| = 2$ related data sets. We assume that the true underlying structure for both the related data sets is that shown in the top panel of Fig. 3, and that the parameters for the two related data sets are those summarised in Tables B.2-B.6 in Appendix B. These parameters have been sampled from a hierarchical Dirichlet distribution with imaginary sample size equal to 10 and with a parameter α that is sampled from a Dirichlet distribution with all $\alpha_{0,ijk} = 1$.

We learn the structure from $|F|$ related data sets, each containing $n_f = 1000$ observations sampled from the true underlying distribution, with the hill-climbing implementation in bnlearn [28] with the BHD score with imaginary sample $s_i = s = 1$.

The true underlying network is recovered after five optimisation steps, shown in Fig. 3. The scores associated with each node and the increasing overall score obtained during the 5 steps of hill-climbing are summarised in Table 1. Tables B.2-B.6 show the parameters estimated by means of the method described in [17], associated with the estimated network. The average absolute error in parameter estimation is 0.023. The average absolute error decreases to 0.005 when the number of observations increases to $n_f = 10000$, thus showing the consistency of both the BHD score and the parameter estimation method as the size of the data set increases. In contrast, even with $n_f = 10000$ both BDeu and BIC are unable to learn the true underlying network from the pooled data sets due to the differences between their distributions. Notice, e.g., the parameters associated to the node X_2 summarised in Table B.3, which represent different associations between X_1 and X_2

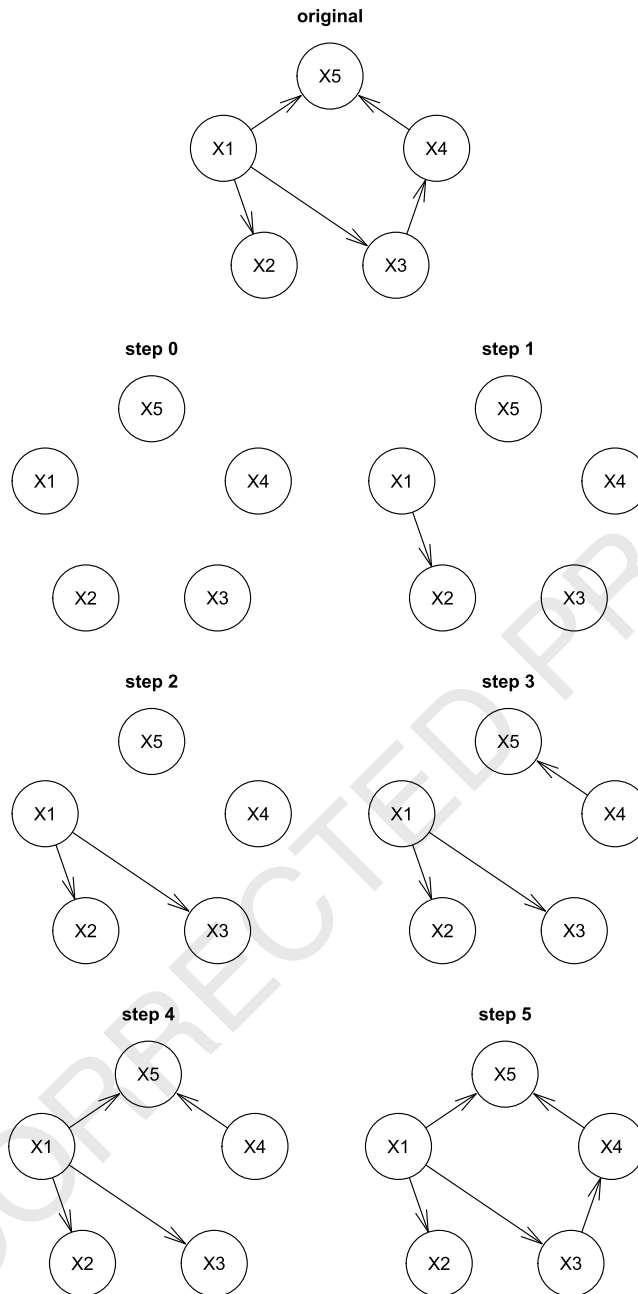


Fig. 3. Original underlying network (top central panel) and networks estimated by means of BHD score during the 5 hill-climbing steps.

across the two data sets. BHD is able to take into account different associations across data sets and to properly estimate the associated parameters (see Table B.3). Both BDeu and BIC estimate instead no association between X_1 and X_2 because they pool the two data sets.

5.1. Sensitivity to the hyperparameters

We repeat the experiment for different values of the hyperparameters $s_i = s$ and $s_{0,i} = s_0$ in the set $\{0.1, 0.5, 1, 10, 100\}$. The results are summarised in Fig. 4. As s increases the learned networks become more connected, similarly to what usually happens with BDeu (see [29] for a more detailed discussion). However, given a value of s , the same graph is learned for all the values of s_0 . In future works, it may be interesting to further study how to choose a suitable value for s for the BHD score or to model it as an hidden random variable with its own prior distribution.




Fig. 4. Network learned by BHD with $s \in \{0.1, 0.5, 1, 10, 100\}$ and $s_0 = 1$. For a given value of s , the learned graph does not change as s_0 varies.

6. Simulation studies

We now perform some simulation studies to compare the empirical performance of BHD to that of BDeu and BIC. For brevity, we will not discuss the results for BIC in detail since they are fundamentally the same as those for BDeu. We are interested in structure learning in the following two scenarios:

- the true underlying network is the same for all the related data sets;
- the true underlying network is the same for all the related data sets, apart from N_F data sets in which N_A randomly selected arcs have been removed.

For each scenario, we generate synthetic data following three different models for the local distributions of each node:

- hier**: the parameters associated with each of the related data sets are sampled from a hierarchical Dirichlet distribution with imaginary sample size equal to 10 and with a parameter α_i sampled from a Dirichlet distribution with all $\alpha_{0,ijk} = 1$;
- iid**: the parameters associated with each of the related data sets are independently sampled from the same Dirichlet distribution with imaginary sample size equal to 10 and uniform α_i vector;
- id**: the parameters are identical for all data sets.

The first approach follows the distributional assumptions of the hierarchical model underlying BHD and may favour the proposed score. The last approach may favour methods that do not take into account that data may comprise related data sets, thus pooling all the data and assuming that all observations are generated from the same distribution. The second approach is a middle ground between the first and the third, since parameters associated with the related data sets are different but they are not generated from the hierarchical model.

We then perform structure learning on the simulated data using the hill-climbing implementation in bnlearn [28] with the BHD and BDeu scores, both with imaginary sample $s_i = s = 1$. In the case of BDeu we pool all the available data from different related data sets. We evaluate the accuracy of network reconstruction with the Structural Hamming Distance (SHD) [30] between the estimated and the true underlying structure, True Positive (TP), False Positive (FP) and False Negative (FN) arcs.

6.1. Simulation study 1

The aim of this simulation study is to evaluate the performance of BHD as different networks parameters vary. Specifically, we consider different number of nodes, related data sets and states for each variable.

We first sample 3 network structures for each of three different levels of sparsity, such that they contain $\{1, 1.2, 1.5\} \cdot N$ arcs, and each combination of:

- $N \in \{5, 10\}$, where N represents the number of nodes;
- $|F| \in \{2, 5, 10\}$, where $|F|$ represents the number of related data sets;
- $|X_i| \in \{2, 5\}$, where $|X_i|$ represents the number of states for each variable.

Then, for both scenario (a) and (b), we replicate the same structure for all the $|F|$ related data sets. In scenario (b), for each of the N_F data sets differing from the others, we randomly remove N_A arcs from the network, with $N_F \in \{1, 2\}$ and $N_A \in \{1, 2\}$. Thus, in scenario (b) we deal with N_F structures that differ from one another and from the main structure by N_A arcs.

Once the network structures have been generated, we sample 10 different parameter sets for each of **hier**, **iid** and **id**. Then, for each of these parameter sets, we sample $|F|$ related data sets, each containing $n_f \in \{10, 100, 200, 500, 1000\}$ observations.

The difference between BDeu and BHD in terms of SHD for scenarios (a) and (b) is shown in Fig. 5, respectively in the left and right panel. Positive values favour the proposed BHD score. When parameters are sampled from a hierarchical distribution (**hier**), BHD outperforms BDeu in both scenarios, with a larger improvement in scenario (b). In the **iid** case, BHD




Fig. 5. Simulation study 1: Boxplots of SHD difference between BHD and BDeu score for scenario (a) (left panel) and (b) (right panel). Positive values favour the hierarchical score.

is competitive with BDeu when the underlying network structures are homogeneous, and it outperforms BDeu when the underlying network structures are different. On the other hand, in the **id** case BDeu has better accuracy than BHD because it correctly assumes that all the data are generated from the same distribution, while BHD has a large number of redundant parameters that would model the non-existing related data sets.

Fig. 6 shows how the difference in SHD between BHD and BDeu varies for different simulation parameters in scenario (a). Specifically, the differences between BHD and BDeu (positive for **hier** and **iid**, negative for **id**) become increasingly large in magnitude as the number of variables N or the number of related data sets $|F|$ increase. On the other hand, the differences between BHD and BDeu gradually decrease as the number of states $|X_i|$ increases. As for the sample size, BHD increasingly outperforms BDeu in both **hier** and **iid** as n_f increases. In the **id** case we expect the two scores to be asymptotically equivalent, but the values we consider for n_f are not large enough to clearly show it empirically.

Fig. 7 shows the relationship between the difference in SHD and some key simulation parameters in scenario (b). The effect of both the number of related data sets $|F|$ and the number of observations n_f is more marked than in scenario (a). For the same $|F|$ and n_f , BHD outperforms BDeu by a larger margin when some network structures are different (scenario (b)) compared to when they are all identical (scenario (a)).

Fig. 8 shows the difference in TP (left), FP (centre) and FN (right panel) between BDeu and BHD for scenario (b). Positive values favour the proposed BHD score. While the two methods perform similarly in terms of TP and FN, BHD outperforms BDeu in terms of FP in the **hier** case. The structures learned by BHD are thus sparser and more interpretable than those learned by BDeu.

We also perform some experiments with different values $s \in \{1, 2, 5, 10\}$ of the imaginary sample size. As s increases, BHD achieves marginally lower SHDs. However, its average SHD is not significantly different from that of BDeu for the same value of s .

6.2. Simulation study 2

Given the results of the first simulation study, we now focus on the effect of specific parameters on the performance of BHD. In particular, the aim of this simulation study is to evaluate the behaviour of the proposed score as the number of related data sets increases.

Similarly to the first simulation study, we sample one network structure for each of three different levels of sparsity ($\{1, 1.2, 1.5\} \cdot N$ arcs as before) and each of $|F| \in \{2, 5, 10, 25, 50, 100\}$ related data sets. We treat both the number of nodes ($N = 10$) and the number of states for each variable ($|X_i| = 2$) as fixed.

Then, for both scenario (a) and (b), we replicate the same structure for all the $|F|$ related data sets. In scenario (b), for each of the N_F data sets differing from the others, we randomly remove N_A arcs from the network, with $N_F \in \{1, 2, |F|\}$ and $N_A \sim \text{Bin}(N_T, p)$, where N_T is the total number of arcs and $p \in \{0.01, 0.1\}$. Thus, in scenario (b) we deal with N_F structures that differ from one another and from the main structure by N_A arcs.

Once the network structures have been generated, we sample 10 different parameter sets for each structure and for each of **hier** and **iid**; for each of these parameter sets, we sample $|F|$ related data sets composed of $n_f \in \{10, 100, 200, 500, 1000, 2000, 5000, 10000\}$ observations each. In this simulation study we disregard the **id** case and we focus instead on the more interesting **hier** and **iid** cases.

Figs. 9 and 10 show how the SHD difference between BHD and BDeu varies as the number of related data sets $|F|$ and the number of observations n_f increase. Results obtained in scenario (a) and (b) are presented together for brevity. Moreover, we did not notice any practical difference as the number of removed arcs N_A or the number of networks with a reduced number of arcs N_F vary.

For small values of $|F|$ the improvement of BHD with respect to BDeu increases with $|F|$, while for large values of $|F|$ the difference between BHD and BDeu reaches a plateau. As expected, the gain is much larger in **hier**. In both **hier** and **iid** BHD increasingly outperforms BDeu as n_f increases. The difference between the two scores is particularly clear for large values of n_f .

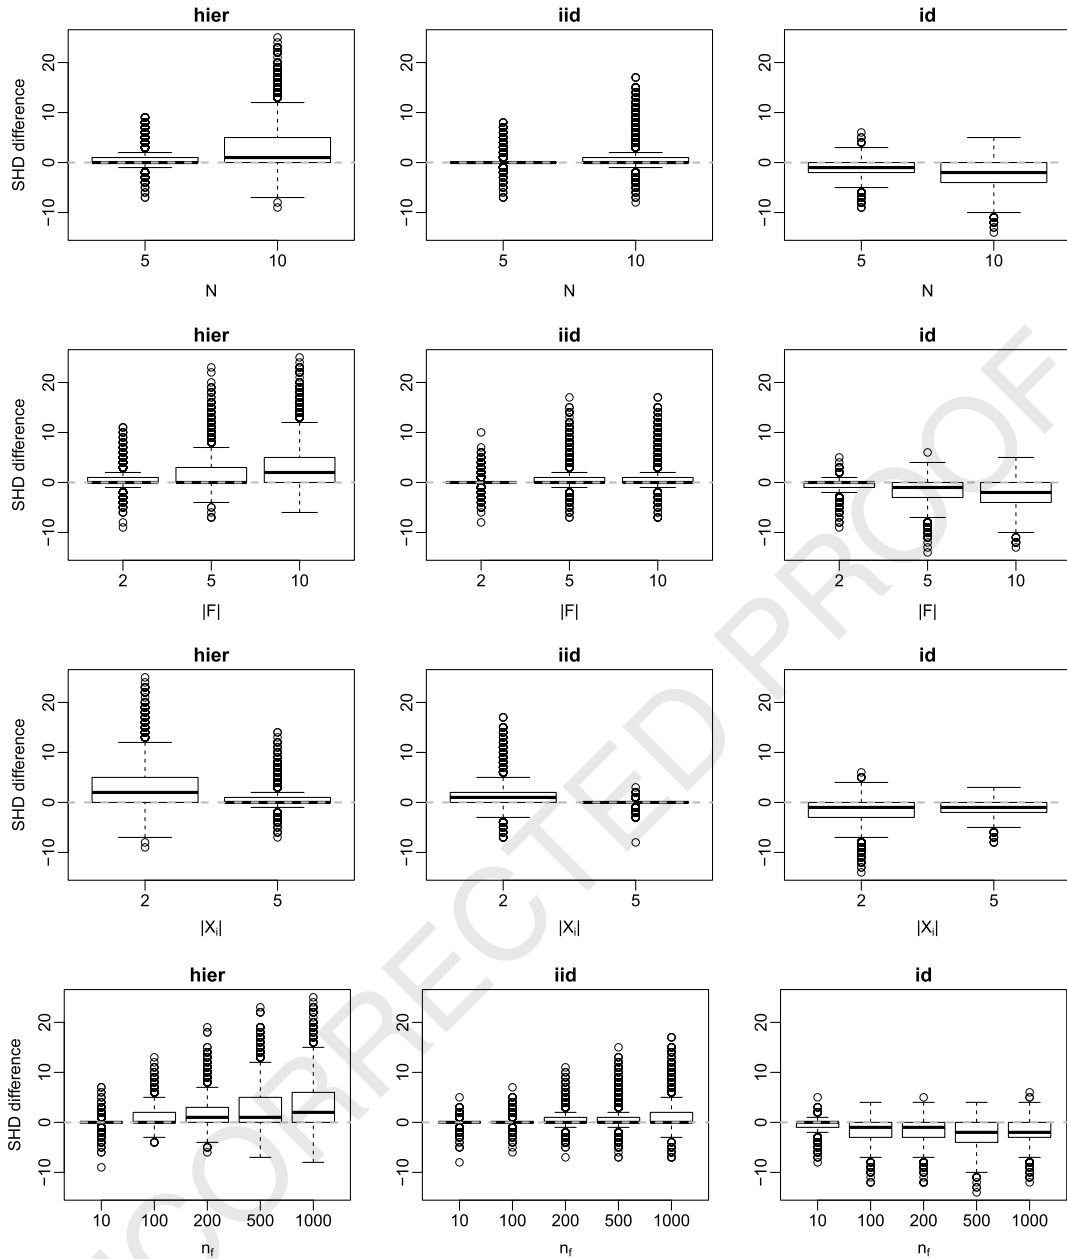


Fig. 6. Simulation study 1: Boxplots of SHD difference between BHD and BDeu score for scenario (a) (equal structures) with different values of number of variables N , number of related data sets $|F|$, number of states $|X_i|$ and number of observations n_f , with parameters sampled with the hierarchical (left panels), i.i.d (central panels) or identical distribution (right panels) approach. Positive values favour the hierarchical score.

The improvement of BHD with respect to BDeu in terms of SHD is the result of a lower number of FP arcs, as in the first simulation study. However, in this study BHD outperforms BDeu also in terms of TP and FN for large values of n_f .

6.3. Simulation study 3

Following up from the second simulation study, we now evaluate the performance of BHD as the number of nodes varies. We sample a network structure for each of three different levels of sparsity (the same as in the first two studies) and each of $N \in \{5, 10, 25, 50, 100\}$ nodes. We treat the number of related data sets ($|F| = 10$) and the number of states for each variable ($|X_i| = 2$) as fixed.

The simulations are performed following the same steps as in the previous simulation study. For both scenario (a) and (b), we replicate the same structure for all the $|F|$ related data sets. We then we randomly remove N_A arcs from the

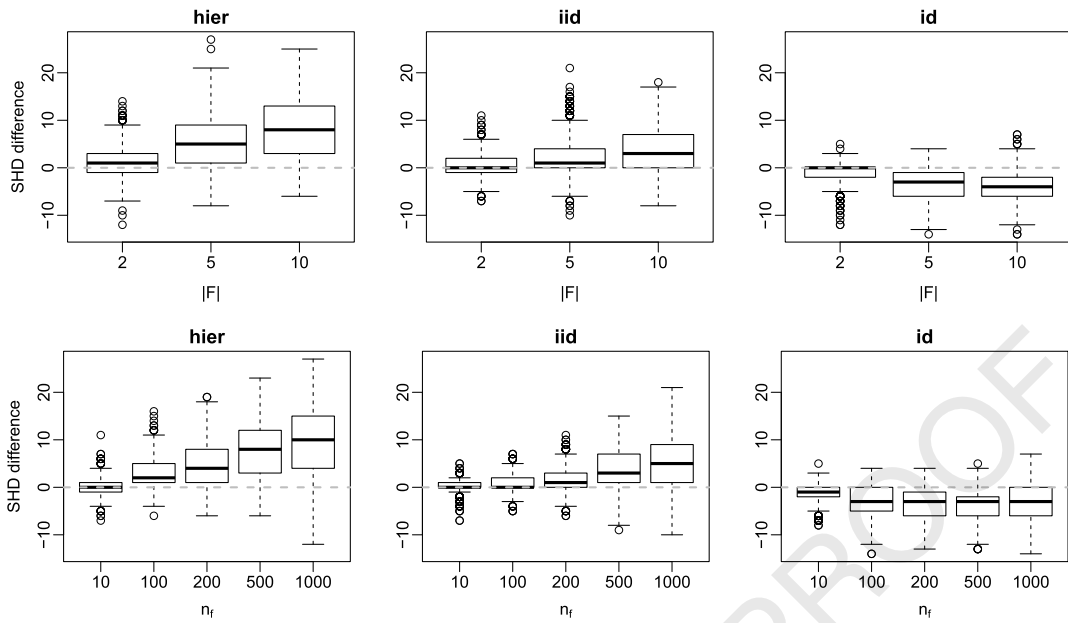


Fig. 7. Simulation study 1: Boxplots of SHD difference between BHD and BDeu score for scenario (b) (different structures) with different values of number of related data sets $|F|$ and number of observations n_f , with parameters sampled with the hierarchical (left panels), i.i.d (central panels) or identical distribution (right panels) approach. Positive values favour the hierarchical score.




Fig. 8. Simulation study 1: Boxplots of TP (left panel), FP (central panel) and FN (right panel) difference between BHD and BDeu score for scenario (b) (different structures). Positive values always favour the hierarchical score.




Fig. 9. Simulation study 2: Boxplots of SHD difference between BHD and BDeu score with different values of number of related data sets $|F|$, with parameters sampled with the hierarchical (left panel) or i.i.d (right panel) approach. Positive values favour the hierarchical score.

networks for each of the N_F data sets differing from the others, as in simulation study 2; then we sample 10 parameter sets and, for each of them, $|F|$ related data sets with $n_f \in \{10, 100, 200, 500, 1000, 2000, 5000, 10000\}$ observations.

Fig. 11 shows the SHD difference between BHD and BDeu as a function of the number of nodes N . Results obtained in scenario (a) and (b) are presented together for brevity as in simulation study 2. The boxplots show that the bigger the

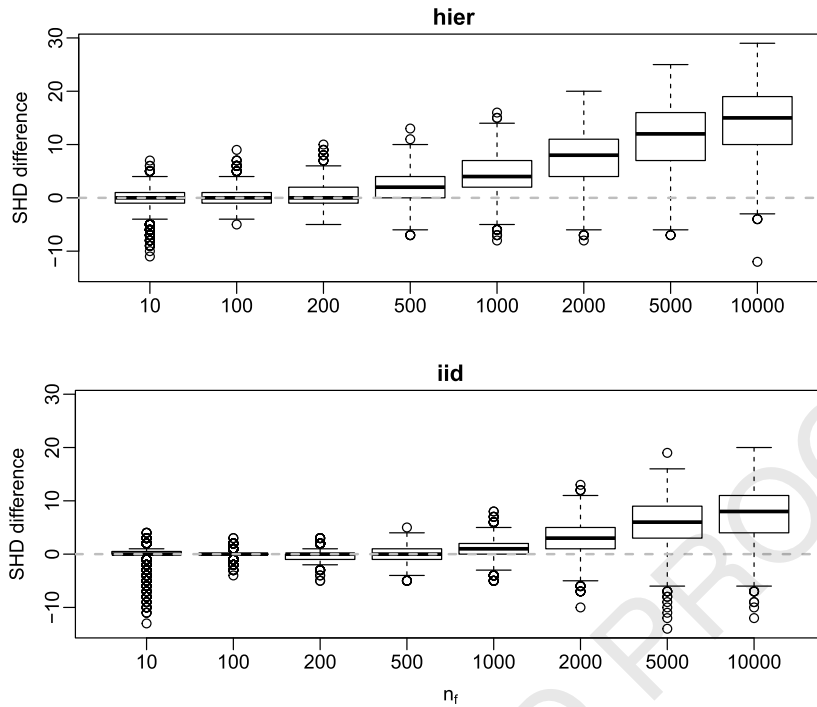


Fig. 10. Simulation study 2: Boxplots of SHD difference between BHD and BDeu score with different values of number of observations n_f , with parameters sampled with the hierarchical (top panel) or i.i.d (bottom panel) approach. Positive values favour the hierarchical score.




Fig. 11. Simulation study 3: Boxplots of SHD difference between BHD and BDeu score with different values of number of variables N , with parameters sampled with the hierarchical (left panel) or i.i.d (right panel) approach. Positive values favour the hierarchical score.

number of nodes N , the larger the improvement of BHD with respect to BDeu. As expected, the gain is much larger in **hier** compared to **iid** for the same N .

As in the previous simulation study, the improvement of BHD with respect to BDeu can be attributed to a lower number of FP arcs. However, in this simulation study BHD outperforms BDeu also in terms of TP and FN for large values of N .

6.4. Simulation study 4

The aim of this simulation study is to evaluate the performance of BHD when each related data set is composed by a different number of observations. The simulation proceeds as in the previous studies, while treating the number of related data sets ($|F| = 10$), the number of nodes ($N = 10$) and the number of states for each variable ($|X_i| = 2$) as fixed.

Only for scenario (a), we replicate the same structure for all the $|F|$ related data sets, we sample 10 different parameter sets for each structure and for each of **hier** and **iid**. Then, for each of these parameter sets, we sample $|F|$ related data sets, each of them composed of a number of observations $n_f \in \{100, 500, 1000\}$ which is the same for all the $|F|$ related data sets apart from N_F data sets that are reduced to $p \cdot n_f$ observations, where $p \in \{0.1, 0.25, 0.5\}$.

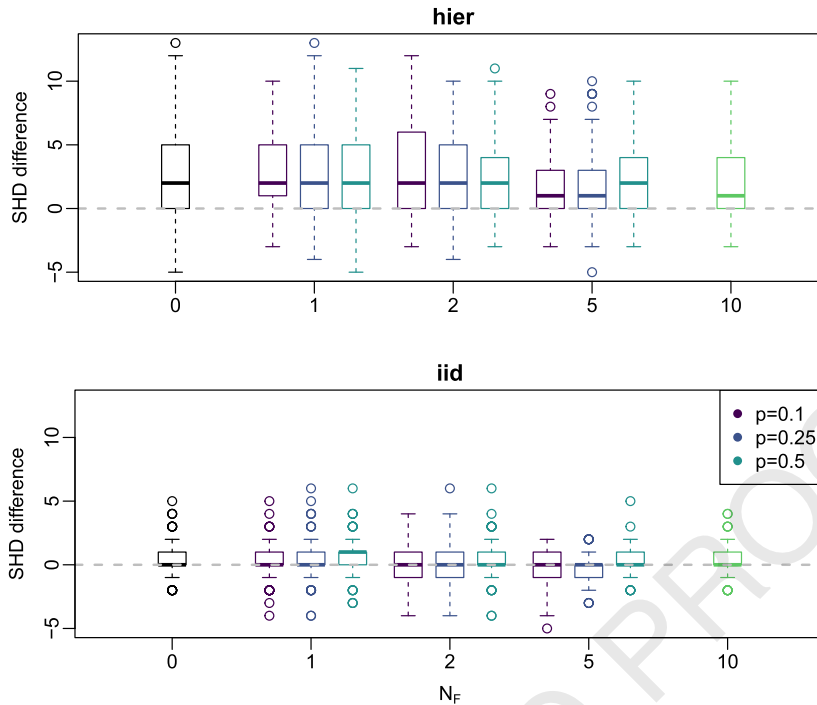


Fig. 12. Simulation study 4: Boxplots of SHD difference between BHD and BDeu score with different number of sub-sampled data sets N_F and proportions of sub-sampling p , with parameters sampled with the hierarchical (top panel) or i.i.d (bottom panel) approach. $N_F = 0$ corresponds to the case where all the data sets have the same number of observations, $N_F = 10$ corresponds to the case where the samples have dimensions $(0.25n_f, 0.25n_f, 0.5n_f, 0.5n_f, 0.75n_f, 0.75n_f, 0.75n_f, n_f, n_f, n_f)$. Positive values favour the hierarchical score.

Furthermore, we consider an additional case with a different number of observations for each related data set. Specifically we consider for the 10 related data sets a number of samples equal to $(0.25n_f, 0.25n_f, 0.5n_f, 0.5n_f, 0.75n_f, 0.75n_f, 0.75n_f, n_f, n_f, n_f)$. This composite case will be identified by means of $N_F = 10$. We compare all these cases using the standard case where all the related data sets are composed by the same number of observations n_f as a baseline. This case will be identified by means of $N_F = 0$.

Fig. 12 shows the SHD difference between BHD and BDeu for different values of N_F and p . The performance of BHD in all the sub-sampled scenarios is similar to the standard case ($N_F = 0$). Also in the composite scenario, with a different number of observations for each related data set ($N_F = 10$), the performance of BHD do not decrease significantly with respect to the equal-number-of-observation scenario. Analogously to the previous simulation studies, the gain is much larger in **hier**.

7. Conclusions and future work

In this work we propose a new Bayesian score, BHD, to learn a common BN structure from related data sets. BHD assumes that their joint distribution in each node of the network follows a mixture of Dirichlet distributions, thus pooling information between the data sets. The joint distribution in each node is approximated by means of a variational method. We found that the resulting computational complexity is linear in the number of related data sets, and that otherwise it is in the same class as BDeu. We showed with a comprehensive set of simulation studies that BHD outperforms both BDeu and BIC when applied to data that comprise related data sets; and that it has comparable performance to BDeu and BIC when the data are a single, homogeneous set of observations. Moreover, the larger is the number of the nodes in the network or the larger is the number of observations, the larger is the improvement with respect to both BDeu and BIC.

Learning a common BN structure with BHD builds on and complements our previous work on parameter learning from related data sets, described in [17]. We can use the latter to learn the parameters associated with a network structure learned using BHD, thus obtaining different BNs (one for each related data set) with the same structure and related parameters, as shown in the numerical example. Combining the two approaches may increase the performance of the BN models such as BN classifiers when dealing with related data sets. Future applications of the combined approach include, e.g., meta-analysis studies, which aim at combining information from several data sources [10, Sec. 5.6].

The assumptions underlying BHD can be relaxed in several ways to extend its applicability to more complex scenarios. For instance, relaxing the assumption that related data sets share the same dependence structure may allow to detect independencies that hold only in certain contexts, as in [31]. Such context-specific independences would be directly modelled by learning different but related network structures for each data set. Another interesting development would be to derive

Algorithm 1 Variational Estimator.

For each $i = 1, \dots, N$, while $\text{iter}_1 < m_1$ and $\text{tol}_1 > t_1$:

1. Update \hat{v}_i by means of (A.1)
2. Fix the starting values of $\hat{\tau}_i, \hat{\kappa}_i$.
3. While $\text{iter}_2 < m_2$ and $\text{tol}_2 > t_2$:
 - (a) update $\hat{\tau}_i$ given $\hat{\kappa}_i$ and \hat{v}_i by means of (A.2);
 - (b) update $\hat{\kappa}_i$ given $\hat{\tau}_i$ and \hat{v}_i by means of (A.3);
 - (c) increase the iterator iter_2 .
4. Update $\hat{\tau}_i, \hat{\kappa}_i$ with the values estimated at the end of step 3.
5. Increase the iterator iter_1 .

a conditional independence test from BHD to learn BNs from related data sets with constraint-based algorithm similarly to, e.g., [32].

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We would like to acknowledge support for this project from the Swiss National Science Foundation (NSF, Grant No. IZKSZ2_162188).

Appendix A. Estimation of variational parameters

For each node X_i , with $i = 1, \dots, N$, we estimate the variational parameters of model (8) by maximising a lower bound of the evidence lower bound (ELBO) for the marginal log-likelihood for the node X_i . The derivation of the bound and the algorithm used to estimate the variational parameters are described in detail in [17]. Here we summarise the parameter estimation method in Algorithm 1 and we adapt it to match the notation used in this paper.

The lower bound of the ELBO for the node X_i is the functional \mathcal{L}_i , which is defined as

$$\begin{aligned}
 \mathcal{L}_i = & \sum_{f=1}^{|F|} \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} (n_{ijk}^f - v_{ijk}^f + s_i \kappa_{ijk}) (\psi(v_{ijk}^f) - \psi(v_{i..}^f)) + \\
 & + \sum_{f=1}^{|F|} \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \log \Gamma(v_{ijk}^f) - |F| \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \log \Gamma(s_i \kappa_{ijk}) + \\
 & + |F| \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} (s_i \kappa_{ijk} - 1) (\log(\kappa_{ijk}) - \psi(\tau_i \kappa_{ijk}) + \psi(\tau_i)) + \\
 & + \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} ([\alpha_0]_{ijk} - \tau_i \kappa_{ijk}) (\psi(\tau_i \kappa_{ijk}) - \psi(\tau_i)) + \\
 & + \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \log \Gamma(\tau_i \kappa_{ijk}) - \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \log \Gamma([\alpha_0]_{ijk}) - \sum_{f=1}^{|F|} \log \Gamma(v_{i..}^f) + \\
 & + |F| \log \Gamma(s_i) - \frac{S_i}{\tau_i} |F| (|X_i| |\Pi_{X_i}| - 1) + \log \Gamma(s_{0,i}) - \log \Gamma(\tau_i),
 \end{aligned}$$

where $v_{i..}^f = \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} v_{ijk}^f$ and $\psi(\cdot)$ is the digamma function, the derivative of the $\log \Gamma(\cdot)$ function.

For each data set $F = f$, with $f = 1, \dots, |F|$, we can first estimate the quantity v_{ijk}^f , associated with the configuration k of X_i and j of Π_{X_i} , by maximising \mathcal{L}_i with respect to v_{ijk}^f and by assuming κ_{ijk} to be fixed. By setting the partial derivative of \mathcal{L}_i with respect to v_{ijk}^f to zero, we obtain:

$$\hat{v}_{ijk}^f = n_{ijk}^f + s_i \kappa_{ijk}. \quad (\text{A.1})$$

We can then estimate $\hat{\tau}_i$ and $\hat{\kappa}_{ijk}$ given the value of v_{ijk}^f by means of a fixed-point method, which alternates the optimisation of \mathcal{L}_i with respect to τ_i and κ_{ijk} . Since no analytical solution is available, we perform this optimisation by means of a Newton algorithm.

We obtain the Newton update for τ_i by treating κ_i as fixed to a vector whose elements are the quantities κ_{ijk} with $j = 1, \dots, |\Pi_{X_i}|$ and $k = 1, \dots, |X_i|$. If we define

$$g_{\tau_i}(\tau_i, \kappa_i) = \partial \mathcal{L}_i / \partial \tau_i = \frac{s_i}{\tau_i^2} |F| (|X_i| |\Pi_{X_i}| - 1) + \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} (\psi'(\tau_i \kappa_{ijk}) \kappa_{ijk} - \psi'(\tau_i)) ([\alpha_0]_{ijk} - \tau_i \kappa_{ijk} - |F| (s_i \kappa_{ijk} - 1))$$

and

$$h_{\tau_i}(\tau_i, \kappa_i) = \partial^2 \mathcal{L}_i / \partial \tau_i^2 = -\frac{2s_i}{\tau_i^3} |F| (|X_i| |\Pi_{X_i}| - 1) - \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \psi''(\tau_i \kappa_{ijk}) \kappa_{ijk}^2 + \psi''(\tau_i) + \sum_{k=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} (\psi''(\tau_i \kappa_{ijk}) \kappa_{ijk}^2 - \psi''(\tau_i)) ([\alpha_0]_{ijk} - \tau_i \kappa_{ijk} - |F| (s_i \kappa_{ijk} - 1)),$$

the Newton update for the parameter τ_i becomes

$$\hat{\tau}_i = \hat{\tau}_i^{\text{old}} \exp \left(-\frac{g_{\tau_i}(\hat{\tau}_i^{\text{old}}, \kappa_i)}{h_{\tau_i}(\hat{\tau}_i^{\text{old}}, \kappa_i) \hat{\tau}_i^{\text{old}} + g_{\tau_i}(\hat{\tau}_i^{\text{old}}, \kappa_i)} \right), \quad (\text{A.2})$$

where $\hat{\tau}_i^{\text{old}}$ is the estimate for τ_i in the previous iteration of the Newton algorithm.

We obtain the Newton update for the parameter vector κ_i by treating τ_i and \mathbf{v}_i as fixed: their elements are the quantities v_{ijk}^f with $f = 1, \dots, |F|$, $j = 1, \dots, |\Pi_{X_i}|$ and $k = 1, \dots, |X_i|$. If we define

$$g_{\kappa_{ijk}}(\kappa_i, \tau, \mathbf{v}_i) = \partial \mathcal{L}_i / \partial \kappa_{ijk} = \sum_{f=1}^{|F|} s_i (\psi(v_{ijk}^f) - \psi(v_{i..}^f)) + \tau_i \psi'(\tau_i \kappa_{ijk}) ([\alpha_0]_{ijk} + \tau_i \kappa_{ijk} - |F| (s_i \kappa_{ijk} - 1)) + \tau_i \psi(\tau_i) + s_i |F| (\psi(\tau_i) - \psi(\tau_i \kappa_{ijk}) - \psi(s_i \kappa_{ijk})) + \log(\kappa_{ijk} + 1) - \frac{|F|}{\kappa_{ijk}}$$

and

$$h_{\kappa_{ijk}}(\kappa_i, \tau, \mathbf{v}_i) = \partial^2 \mathcal{L}_i / \partial \kappa_{ijk}^2 = \tau_i^2 \psi''(\tau_i \kappa_{ijk}) ([\alpha_0]_{ijk} - \tau_i \kappa_{ijk} + |F| (s_i \kappa_{ijk} - 1)) - \tau_i \psi'(\tau_i \kappa_{ijk}) (\tau_i + 2s_i |F|) - s_i^2 |F| \psi'(s_i \kappa_{ijk}) + \frac{s_i |F|}{\kappa_{ijk}} + \frac{|F|}{\kappa_{ijk}^2},$$

we can write the Newton update for κ_{ijk} as

$$\hat{\kappa}_{ijk} = \hat{\kappa}_{ijk}^{\text{old}} + \frac{\sum_{l=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \frac{g_{\kappa_{ijl}}(\kappa_i^{\text{old}}, \tau_i, \mathbf{v}_i)}{h_{\kappa_{ijl}}(\kappa_i^{\text{old}}, \tau_i, \mathbf{v}_i)}}{\sum_{l=1}^{|X_i|} \sum_{j=1}^{|\Pi_{X_i}|} \frac{h_{\kappa_{ijl}}(\kappa_i^{\text{old}}, \tau_i, \mathbf{v}_i)}{h_{\kappa_{ijk}}(\kappa_i^{\text{old}}, \tau_i, \mathbf{v}_i)}} - \frac{g_{\kappa_{ijk}}(\kappa_i^{\text{old}}, \tau_i, \mathbf{v}_i)}{h_{\kappa_{ijk}}(\kappa_i^{\text{old}}, \tau_i, \mathbf{v}_i)}, \quad (\text{A.3})$$

where $\kappa_{ijk}^{\text{old}}$ is the estimate for κ_{ijk} in the previous iteration of the Newton algorithm.

Appendix B. Parameters of numerical example

Tables B.2-B.6 show the original and the estimated parameters of the network used in the numerical example presented in Section 5.

Table B.2
Parameters associated with X_1 .

$X_1 F$	Original		Estimated	
	1	2	1	2
1	0.42	0.18	0.41	0.17
2	0.58	0.82	0.59	0.83

Table B.3
Parameters associated with X_2 .

F	Original				Estimated			
	1		2		1		2	
$X_2 X_1$	1	2	1	2	1	2	1	2
1	0.31	0.65	0.56	0.16	0.33	0.64	0.58	0.14
2	0.69	0.35	0.44	0.84	0.67	0.36	0.42	0.86

Table B.4
Parameters associated with X_3 .

F	Original				Estimated			
	1		2		1		2	
$X_3 X_1$	1	2	1	2	1	2	1	2
1	0.61	0.42	0.15	0.53	0.64	0.39	0.15	0.54
2	0.39	0.58	0.85	0.47	0.36	0.61	0.85	0.46

Table B.5
Parameters associated with X_4 .

F	Original				Estimated			
	1		2		1		2	
$X_4 X_3$	1	2	1	2	1	2	1	2
1	0.47	0.51	0.38	0.28	0.45	0.51	0.40	0.28
2	0.53	0.49	0.62	0.72	0.55	0.49	0.60	0.72

Table B.6
Parameters associated with X_5 .

F	Original				Estimated			
	1		2		1		2	
X_4	1		2		1		2	
$X_5 X_1$	1	2	1	2	1	2	1	2
1	0.46	0.56	0.57	0.52	0.65	0.43	0.22	0.38
2	0.54	0.44	0.43	0.48	0.35	0.57	0.78	0.62

F	Estimated				Estimated			
	1		2		1		2	
X_4	1		2		1		2	
$X_5 X_1$	1	2	1	2	1	2	1	2
1	0.41	0.58	0.53	0.53	0.77	0.44	0.14	0.38
2	0.59	0.42	0.47	0.47	0.23	0.56	0.86	0.62

References

[1] R.J. Gray, A Bayesian analysis of institutional effects in a multicenter cancer clinical trial, *Biometrics* 50 (1994) 244–253.
 [2] D.J. Spiegelhalter, K.R. Abrams, J.P. Myles, *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*, Wiley, 2004.
 [3] M.E. Goddard, Genomic selection: prediction of accuracy and maximisation of long term response, *Genetica* 136 (2009) 245–257.
 [4] Y.C.J. Wientjes, P. Bijma, R.F. Veerkamp, M.P.L. Calus, An equation to predict the accuracy of genomic values by combining data from multiple traits, populations, or environments, *Genetics* 202 (2016) 799–823.
 [5] R. Makowsky, N.M. Pajewski, Y.C. Klimentidis, A.I. Vazquez, C.W. Duarte, D.B. Allison, G. de los Campos, Beyond missing heritability: prediction of complex traits, *PLoS Genet.* 7 (2011) e1002051.
 [6] G. de los Campos, A.I. Vazquez, R.L. Fernando, Y.C. Klimentidis, D. Sorensen, Prediction of complex human traits using the genomic best linear unbiased predictor, *PLoS Genet.* 9 (2013) e1003608.

- [7] A. Russell, M. Ghalaieny, B. Gazdiyeva, S. Zhumabayeva, A. Kurmanbayeva, K.K. Akhmetov, Y. Mukanov, M. McCann, M. Ali, A. Tucker, C. Vitolo, A. Althonayan, A spatial survey of environmental indicators for Kazakhstan: an examination of current conditions and future needs, *Int. J. Environ. Res.* 12 (2018) 735–748.
- [8] C. Vitolo, M. Scutari, A. Tucker, A. Russell, Modelling air pollution, climate and health data using Bayesian networks: a case study of the English regions, *Earth Space Sci.* 5 (2018) 76–88.
- [9] S.S. Qian, T.F. Cuffney, I. Alameddine, G. McMahon, K.H. Reckhow, On the application of multilevel modeling in environmental and ecological studies, *Ecology* 91 (2010) 355–361.
- [10] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, D.B. Rubin, *Bayesian Data Analysis*, 3rd ed., CRC Press, 2014.
- [11] D. Koller, N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, 2009.
- [12] A. Niculescu-Mizil, R. Caruana, Inductive transfer for Bayesian network structure learning, in: *Proceedings of Artificial Intelligence and Statistics, 2007*, pp. 339–346.
- [13] C.J. Oates, J.Q. Smith, S. Mukherjee, J. Cussens, Exact estimation of multiple directed acyclic graphs, *Stat. Comput.* 26 (2016) 797–811.
- [14] F. De Michelis, P. Magni, P. Piergiorgi, M.A. Rubin, R. Bellazzi, A hierarchical naive Bayes model for handling sample heterogeneity in classification problems: an application to tissue microarrays, *BMC Bioinform.* 7 (2006) 514.
- [15] A. Malovini, N. Barbarini, R. Bellazzi, F. De Michelis, Hierarchical naive Bayes for genetic association studies, *BMC Bioinform.* 13 (2012) S6.
- [16] E. Bareinboim, J. Pearl, Causal inference and the data-fusion problem, *Proc. Natl. Acad. Sci.* 113 (2016) 7345–7352.
- [17] L. Azzimonti, G. Corani, M. Zaffalon, Hierarchical estimation of parameters in Bayesian networks, *Comput. Stat. Data Anal.* 137 (2019) 67–91.
- [18] D. Heckerman, D. Geiger, D.M. Chickering, Learning Bayesian networks: the combination of knowledge and statistical data, *Mach. Learn.* 20 (1995) 197–243. Available as, Technical Report MSR-TR-94-09.
- [19] S.L. Lauritzen, N. Wermuth, Graphical models for associations between variables, some of which are qualitative and some quantitative, *Ann. Stat.* 17 (1989) 31–57.
- [20] G. Schwarz, Estimating the dimension of a model, *Ann. Stat.* 6 (1978) 461–464.
- [21] M. Ueno, Learning networks determined by the ratio of prior and data, in: *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence, 2010*, pp. 598–605.
- [22] M. Scutari, An empirical-Bayes score for discrete Bayesian networks, in: *Proceedings Track, PGM 2016, J. Mach. Learn. Res.* 52 (2016) 438–448.
- [23] M. Scutari, Dirichlet Bayesian network scores and the maximum relative entropy principle, *Behaviormetrika* 45 (2018) 337–362.
- [24] D.M. Chickering, A transformational characterization of equivalent Bayesian network structures, in: *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence, UAI '95, 1995*, pp. 87–98.
- [25] M.I. Jordan, Z. Ghahramani, T.S. Jaakkola, L.K. Saul, An introduction to variational methods for graphical models, *Mach. Learn.* 37 (1999) 183–233.
- [26] M.J. Wainwright, M.I. Jordan, Graphical models, exponential families, and variational inference, *Found. Trends Mach. Learn.* 1 (2008) 1–305.
- [27] G. Casella, E. Moreno, Assessing robustness of intrinsic tests of independence in two-way contingency tables, *J. Am. Stat. Assoc.* 104 (2009) 1261–1271.
- [28] M. Scutari, Learning Bayesian networks with the bnlearn R package, *J. Stat. Softw.* 35 (2010) 1–22.
- [29] T. Silander, P. Kontkanen, P. Myllymäki, On sensitivity of the map bayesian network structure to the equivalent sample size parameter, in: *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence, 2007*, pp. 360–367.
- [30] I. Tsamardinos, L.E. Brown, C.F. Aliferis, The max-min hill-climbing Bayesian network structure learning algorithm, *Mach. Learn.* 65 (2006) 31–78.
- [31] C. Boutilier, N. Friedman, M. Goldszmidt, D. Koller, Context-specific independence in Bayesian networks, in: *Proceedings of the 12th International Conference on Uncertainty in Artificial Intelligence, UAI '97, 1996*, pp. 115–123.
- [32] R.E. Tillman, Structure learning with independent non-identically distributed data, in: *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09, 2009*, pp. 1041–1048.