

A BGe score for tied-covariance mixtures of Gaussian Bayesian networks

Marco Grzegorzcyk

Bernoulli Institute (BI)
Rijksuniversiteit Groningen
E-mail: m.a.grzegorzcyk@rug.nl

Abstract

Mixtures of Gaussian Bayesian networks have previously been studied under full-covariance assumptions, where each mixture component has its own covariance matrix. We propose a mixture model with tied-covariance, in which all components share a common covariance matrix. Our main contribution is the derivation of its marginal likelihood, which remains analytic. Unlike in the full-covariance case, however, the marginal likelihood no longer factorizes into component-specific terms. We refer to the new likelihood as the BGe scoring metric for tied-covariance mixtures of Gaussian Bayesian networks. For model inference, we implement MCMC schemes combining structure MCMC with a fast Gibbs sampler for mixtures, and we empirically compare the tied- and full-covariance mixtures of Gaussian Bayesian networks on simulated and benchmark data.

Keywords

Gaussian Bayesian networks, mixture model, tied-covariance, Markov chain Monte Carlo

1 Introduction

Bayesian networks (BNs) are popular graphical models for representing conditional (in-) dependencies among random variables using directed acyclic graphs (DAGs). Each node in a DAG corresponds to a variable, and directed edges encode probabilistic relationships. This structure allows joint distributions with many conditional independencies to be expressed compactly. BNs are widely used for learning under uncertainty, making them important tools in statistics, machine learning, and artificial intelligence. Thorough introductions to Bayesian networks are available in standard textbooks (Pearl, 1988; Neapolitan, 1989; Koller and Friedman, 2009; Scutari and Denis, 2014). Structure learning algorithms aim to infer DAGs from data, but since the search space grows super-exponentially with the number of nodes (Chickering, 1996), the problem is computationally challenging. For efficient structure learning algorithms, see Kuipers and Moffa (2017); Scutari et al. (2018, 2019); Scutari (2020); Kitson et al. (2023). In the case of sparse data, there is typically considerable uncertainty about the DAG. This uncertainty can be addressed by Bayesian structure learning algorithms, by sampling DAGs from the posterior distribution and averaging over the sampled DAGs. This approach is commonly referred to as Bayesian model averaging (Constantinou et al., 2022). For Bayesian BN structure learning algorithms, see Madigan and York (1995); Giudici and Castelo (2003); Friedman and Koller (2003)

or Grzegorzczuk and Husmeier (2008).

Beyond the standard BN structure learning literature, tailored approaches have been developed for incomplete (Friedman, 1997; Ruggieri et al., 2020; Bodewes and Scutari, 2021), ordinal (Luo et al., 2021; F., 2024; Scauda et al., 2025), and hybrid data (Chickering et al., 1995; Scutari, 2020). We have recently contributed to this line of work by proposing Bayesian counterparts for these important special cases, addressing missing (Grzegorzczuk, 2023), ordinal (Grzegorzczuk, 2024), and hybrid (Grzegorzczuk, 2025) data. In the present paper, we extend this trilogy (Grzegorzczuk, 2023, 2024, 2025) by introducing a Bayesian approach for tied-covariance mixtures of Gaussian Bayesian networks. A mixture model assumes that the data stem from several subpopulations, each represented by its own distribution. Mixture models are used in many frameworks, and we are not the first to propose mixture Bayesian networks: both non-Bayesian (Ko et al., 2007) and Bayesian (Grzegorzczuk et al., 2008) formulations have been introduced, each assuming that all subpopulations share a common DAG but differ in their network parameters. Because they use component-specific covariance matrices, these earlier models are full-covariance mixtures of Gaussian Bayesian networks, which we refer to as M1. The Bayesian M1 model from Grzegorzczuk et al. (2008) employs the structure MCMC sampler (Madigan and York, 1995; Giudici and Castelo, 2003) to infer the DAGs and the allocation sampler (Nobile and Fearnside, 2007) to assign the data to K mixture components. Each mixture component $k = 1, \dots, K$ has a weight $\pi_k \in [0, 1]$, a mean vector $\boldsymbol{\mu}_k$, and a covariance matrix $\boldsymbol{\Sigma}_k$. The resulting density is

$$p(\mathbf{x}|\text{DAG}, \mathcal{M}_1) = \sum_{k=1}^K \pi_k \cdot \mathcal{N}_n(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k),$$

where $\mathcal{N}_n(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the density of the n -variate normal distribution $\mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ at $\mathbf{x}$, and each $\boldsymbol{\Sigma}_k$ must be consistent with the DAG, that is, each $\boldsymbol{\Sigma}_k$ must encode the conditional independence relations implied by the DAG. For a dataset $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ with given class labels and under suitable prior assumptions,¹ the marginal likelihood of the mixture BN model factorizes into the product of the component-specific marginal likelihoods:

$$p(\mathcal{D}|\text{DAG}, \mathcal{M}_1) = \prod_{k=1}^K p(\mathcal{D}_k|\text{DAG}),$$

where $\mathcal{D}_k$ denotes the subset of observations with class label k , and each $p(\mathcal{D}_k|\text{DAG})$ can be computed analytically.

In this paper, we introduce the tied-covariance counterpart of the M1 model, which we refer to as the M2 mixture model. In the M2 model the K components share a common covariance matrix $\boldsymbol{\Sigma}$:

$$p(\mathbf{x}|\text{DAG}, \mathcal{M}_2) = \sum_{k=1}^K \pi_k \cdot \mathcal{N}_n(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}).$$

The main contribution of this paper is the derivation of the marginal likelihood $p(\mathcal{D}|\text{DAG}, \mathcal{M}_2)$ for the tied-covariance mixture model (M2). We show that it admits an analytic solution, but unlike in the M1 case, it no longer decomposes into K local marginal likelihoods. We

¹In each component k , an independent Normal-Wishart prior is used (Geiger and Heckerman, 1994, 2002; Kuipers et al., 2014).

refer to the new marginal likelihood as the BGe scoring metric for tied-covariance mixtures of Gaussian Bayesian networks. We implement both types of mixture models and compare their performances on simulated and benchmark data. Our Markov chain Monte Carlo (MCMC) sampling algorithm combines structure MCMC sampling of DAGs (Madigan and York, 1995; Giudici and Castelo, 2003) with a recently proposed fast Gibbs sampler for mixture models (Newman, 2025).

2 Statistical methods

2.1 Bayesian networks

Bayesian networks (BNs) are graphical models that use directed acyclic graphs (DAGs) to represent conditional independencies among a set of random variables, which we denote by $X_1, \dots, X_n$. The variables correspond to the n nodes of a DAG, and directed edges between nodes encode direct conditional dependencies. Node X_j is called a parent of X_i if there is an edge $X_j \rightarrow X_i$, and it is called an ancestor of X_i if X_i can be reached from X_j by following a sequence of one or more directed edges along their orientations, symbolically $X_j \rightarrow \dots \rightarrow X_i$. By definition, a node cannot be its own ancestor, so cyclic paths of directed edges, such as $X_i \rightarrow \dots \rightarrow X_i$, are excluded. The set of parent nodes of X_i is denoted by Π_i . A DAG can be identified with its n parent sets $\Pi_1, \dots, \Pi_n$. Given $\mathcal{G}$, i.e. the conditional independence relations implied by $\mathcal{G}$, the joint density factorizes as

$$p(X_1, \dots, X_n | \mathcal{G}) = \prod_{i=1}^n p(X_i | \Pi_i). \quad (1)$$

Two DAGs, say $\mathcal{G}_\star$ and $\mathcal{G}_\bullet$, that imply the same conditional independencies are called equivalent, and they yield the same factorization:

$$p(X_1, \dots, X_n | \mathcal{G}_\star) = \prod_{i=1}^n p(X_i | \Pi_i^\star) = \prod_{i=1}^n p(X_i | \Pi_i^\bullet) = p(X_1, \dots, X_n | \mathcal{G}_\bullet),$$

where $\Pi_i^\star$ and $\Pi_i^\bullet$ denote the parent sets of X_i in $\mathcal{G}_\star$ and $\mathcal{G}_\bullet$, respectively. Equivalent DAGs belong to the same equivalence class. Chickering (2002) shows that DAG equivalence classes can be represented by completed partially directed acyclic graphs (CPDAGs), and that two DAGs $\mathcal{G}_\star$ and $\mathcal{G}_\bullet$ are equivalent if and only if they share the same skeleton and the same v-structures. The skeleton of a DAG is obtained by replacing each directed edge with an undirected edge. A v-structure is a configuration in which a node X_i has two parent nodes with no edge between them, for example $X_k \rightarrow X_i \leftarrow X_l$ with no edge between X_k and X_l . CPDAGs contain a mixture of directed and undirected edges. A directed edge $X_j \rightarrow X_i$ in a CPDAG indicates that all DAGs in the equivalence class share this directed edge. An undirected edge $X_j - X_i$ indicates that all DAGs in the equivalence class contain an edge between X_j and X_i , but not all agree on its orientation. In particular, there is at least one DAG with the edge $X_j \rightarrow X_i$ and at least one DAG with the edge $X_j \leftarrow X_i$.

The random variables $X_1, \dots, X_n$ can be stacked into a random vector $\mathbf{X} = (X_1, \dots, X_n)^\top$. A dataset $\mathcal{D}$ consists of m independent realizations of $\mathbf{X}$:

$$\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\},$$

where $\mathbf{x}_i$ denotes the i -th observation of $\mathbf{X}$. The posterior distribution of DAG $\mathcal{G}$ is defined as:

$$p(\mathcal{G}|\mathcal{D}) := \frac{p(\mathcal{D}|\mathcal{G})p(\mathcal{G})}{p(\mathcal{D})} \propto p(\mathcal{D}|\mathcal{G})p(\mathcal{G}), \quad (2)$$

where $p(\mathcal{D}|\mathcal{G})$ is the marginal likelihood of $\mathcal{G}$, $p(\mathcal{G})$ is the prior probability of $\mathcal{G}$, and $p(\mathcal{D})$ is a normalizing constant. Using Markov chain Monte Carlo (MCMC) sampling, DAGs can be drawn from the posterior distribution, yielding a posterior sample $\mathcal{G}_1, \dots, \mathcal{G}_R$.

In the BN literature, the marginal likelihood $p(\mathcal{D}|\mathcal{G})$ is often referred to as the score of the DAG $\mathcal{G}$. Two BN scoring metrics allow the marginal likelihood to be computed analytically:

- The *BDe scoring metric* (Heckerman et al., 1995; Chickering et al., 1995; Heckerman and Geiger, 1995) is used to compute the marginal likelihood of DAGs over discrete (categorical) random variables. BDe stands for “Bayesian scoring metric for discrete networks having score equivalence.”
- The *BGe scoring metric* (Geiger and Heckerman, 2002; Kuipers et al., 2014; Heckerman and Geiger, 1995; Geiger and Heckerman, 1994) is used to compute the marginal likelihood of DAGs over Gaussian random variables. BGe stands for “Bayesian scoring metric for Gaussian networks having score equivalence.”

These two scoring metrics not only allow the marginal likelihood to be computed analytically, but also ensure score equivalence, i.e. equivalent DAGs receive the same marginal likelihood value. This property is crucial for BNs, since equivalent graphs describe the same conditional dependencies and therefore should not be assigned different scores.

2.2 Gaussian Bayesian networks with the BGe score

In Gaussian Bayesian networks (GBNs), the random vector $\mathbf{X} \in \mathbb{R}^n$ follows an n -dimensional Gaussian distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. For a dataset $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ we then have

$$p(\mathcal{D}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \prod_{i=1}^m \mathcal{N}_n(\mathbf{x}_i|\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (3)$$

where $\mathcal{N}_n(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the density of the n -variate normal distribution $\mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ at $\mathbf{x}$. A fully conjugate prior for the mean vector $\boldsymbol{\mu}$ and the precision matrix $\mathbf{W} = \boldsymbol{\Sigma}^{-1}$ is given by the Normal–Wishart distribution:

$$\begin{aligned} \mathbf{W} &\sim \mathcal{W}_n(\alpha_w, \mathbf{T}_\dagger), \\ \boldsymbol{\mu}|\mathbf{W} &\sim \mathcal{N}_n(\boldsymbol{\nu}, (\alpha_\mu \mathbf{W})^{-1}), \end{aligned} \quad (4)$$

where $\mathbf{T}_\dagger \in \mathbb{R}^{n \times n}$, $\boldsymbol{\nu} \in \mathbb{R}^n$, $\alpha_w > n - 1$, and $\alpha_\mu > 0$ are hyperparameters. $\mathbf{T}_\dagger$ is the positive definite parametric matrix of the Wishart prior, $\boldsymbol{\nu}$ is the mean vector of the Gaussian prior, and α_w and α_μ can be interpreted as prior equivalent sample sizes (Geiger and Heckerman, 2002). Since the prior is fully conjugate, the marginal likelihood

$$p(\mathcal{D}) = \int \int p(\mathcal{D}|\boldsymbol{\mu}, \mathbf{W}^{-1}) \cdot p(\boldsymbol{\mu}|\mathbf{W}) \cdot p(\mathbf{W}) d\boldsymbol{\mu} d\mathbf{W} \quad (5)$$

can be computed analytically. However, the model specified in Eqs. (3–4) does not impose any conditional independence relations among the variables. That is, every pair of variables is dependent and remains dependent even when conditioning on any subset of the others. In terms of BNs, Eq. (5) therefore corresponds to the marginal likelihood of a complete DAG $\mathcal{G}_C$, symbolically $p(\mathcal{D}) = p(\mathcal{D}|\mathcal{G}_C)$. A complete DAG $\mathcal{G}_C$ has the maximal number of edges, implying that no pairwise independencies exist. For general DAGs, by contrast, the marginal likelihood incorporates the conditional independencies implied by the parent sets Π_i through the factorization in Eq. (1).

In their seminal works on GBNs (Geiger and Heckerman, 2002; Kuipers et al., 2014; Heckerman and Geiger, 1995; Geiger and Heckerman, 1994), Geiger and Heckerman extend the model in Eqs. (3–4) to non-complete DAGs. A DAG $\mathcal{G}$ specifies conditional independence relations among the variables, and it must be ensured that the precision matrix $\mathbf{W}$ is consistent with $\mathcal{G}$, i.e. that $\mathcal{G}$ and $\mathbf{W}$ imply the same conditional independencies. Geiger and Heckerman also show how to compute the marginal likelihood $p(\mathcal{D}|\mathcal{G})$ for any given DAG $\mathcal{G}$, and they refer to this marginal likelihood as the BGe score of $\mathcal{G}$.

2.3 Mixture Gaussian Bayesian networks

In this section we assume that the data arise from a mixture of GBNs. In Section 2.3.1 we consider the full-covariance mixture model, in which each component has its own Gaussian distribution with a separate covariance matrix. In Section 2.3.2 we introduce the tied-covariance mixture model, in which all components share a common covariance matrix. The full-covariance model yields a marginal likelihood that factorizes into component-specific BGe scores. For the tied-covariance model we derive the marginal likelihood, which we refer to as the BGe score for tied-covariance mixture GBNs.

In both sections we let $z_i = k$ indicate that observation vector $\mathbf{x}_i$ is assigned to mixture component k . We assume there are K components, and define the assignment vector $\mathbf{z} = (z_1, \dots, z_m)$, which allocates m_k observations $\mathbf{x}_{k,l}$ ($l = 1, \dots, m_k$) to component k . Accordingly, we can write

$$\mathcal{D} = \{\mathbf{x}_{k,l} | k = 1, \dots, K; l = 1, \dots, m_k\}, \text{ where } \sum_{k=1}^K m_k = m.$$

For each component k , we have the empirical mean vector and the scatter matrix (cross-product matrix):

$$\bar{\mathbf{x}}_{[k]} := \frac{1}{m_k} \sum_{l=1}^{m_k} \mathbf{x}_{k,l}, \text{ and } \mathbf{S}_{[k]} := \sum_{l=1}^{m_k} (\mathbf{x}_{k,l} - \bar{\mathbf{x}}_{[k]})(\mathbf{x}_{k,l} - \bar{\mathbf{x}}_{[k]})^\top. \quad (6)$$

2.3.1 Full-covariance mixture Gaussian Bayesian networks

Suppose each observation vector $\mathbf{x}_i$ ($i = 1, \dots, m$) arises from a Gaussian mixture distribution with K components:

$$p(\mathbf{x}_i | \Theta) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (7)$$

where $\Theta := \{(\pi_1, \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1), \dots, (\pi_K, \boldsymbol{\mu}_K, \boldsymbol{\Sigma}_K)\}$ is the set of model parameters, $\pi_k > 0$ is the weight of component k , and $\sum_{k=1}^K \pi_k = 1$. On each pair $(\boldsymbol{\mu}_k, \mathbf{W}_k)$, where $\mathbf{W}_k = \boldsymbol{\Sigma}_k^{-1}$, we

impose the fully conjugate Normal–Wishart prior from Eq. (4):

$$\begin{aligned}\mathbf{W}_k &\sim \mathcal{W}_n(\alpha_{w,k}, \mathbf{T}_{\dagger,k}), \\ \boldsymbol{\mu}_k | \mathbf{W}_k &\sim \mathcal{N}_n(\boldsymbol{\nu}_k, (\alpha_{\mu,k} \mathbf{W}_k)^{-1}),\end{aligned}\tag{8}$$

where $\mathbf{T}_{\dagger,k} \in \mathbb{R}^{n \times n}$ is positive definite, $\boldsymbol{\nu}_k \in \mathbb{R}^n$, $\alpha_{w,k} > n - 1$, and $\alpha_{\mu,k} > 0$ are hyperparameters. As before, $\alpha_{w,k}$ and $\alpha_{\mu,k}$ can be interpreted as prior equivalent sample sizes (Geiger and Heckerman, 2002).

For a given assignment vector $\mathbf{z}$ and a complete DAG $\mathcal{G}_C$, the component-specific posterior distributions are

$$\begin{aligned}\mathbf{W}_k | (\mathcal{D}, \mathbf{z}, \mathcal{G}_C) &\sim \mathcal{W}_n(\alpha_{w,k} + m_k, \mathbf{T}_{\dagger,k}), \\ \boldsymbol{\mu}_k | (\mathcal{D}, \mathbf{z}, \mathcal{G}_C, \mathbf{W}_k) &\sim \mathcal{N}_n(\boldsymbol{\mu}_k^\diamond, ((\alpha_{\mu,k} + m_k) \mathbf{W}_k)^{-1}),\end{aligned}\tag{9}$$

with

$$\begin{aligned}\mathbf{T}_{\dagger,k} &:= \mathbf{T}_{\dagger,k} + \mathbf{S}_{[k]} + \frac{\alpha_{\mu,k} \cdot m_k}{\alpha_{\mu,k} + m_k} (\boldsymbol{\nu}_k - \bar{\mathbf{x}}_{[k]})(\boldsymbol{\nu}_k - \bar{\mathbf{x}}_{[k]})^\top, \\ \boldsymbol{\mu}_k^\diamond &:= \frac{\alpha_{\mu,k} \boldsymbol{\nu}_k + m_k \bar{\mathbf{x}}_{[k]}}{\alpha_{\mu,k} + m_k}.\end{aligned}\tag{10}$$

For the marginal likelihood we obtain:

$$p(\mathcal{D} | \mathbf{z}, \mathcal{G}_C) = \prod_{k=1}^K \left(\pi^{-\frac{n \cdot m_k}{2}} \cdot \left(\frac{\alpha_{\mu,k}}{\alpha_{\mu,k} + m_k} \right)^{\frac{n}{2}} \cdot \frac{\Gamma_n(\frac{\alpha_{w,k} + m_k}{2})}{\Gamma_n(\frac{\alpha_{w,k}}{2})} \cdot \frac{\det(\mathbf{T}_{\dagger,k})^{\alpha_{w,k}/2}}{\det(\mathbf{T}_{\dagger,k})^{(\alpha_{w,k} + m_k)/2}} \right),\tag{11}$$

where $\Gamma_n(\cdot)$ denotes the n -dimensional multivariate gamma function. From Eq. (11) it can be seen that the marginal likelihood factorizes into K terms. For any l -dimensional variable subset $L \subset \{X_1, \dots, X_n\}$ we have

$$p(\mathcal{D}^L | \mathbf{z}, \mathcal{G}_C^L) = \prod_{k=1}^K \left(\pi^{-\frac{l \cdot m_k}{2}} \cdot \left(\frac{\alpha_{\mu,k}}{\alpha_{\mu,k} + m_k} \right)^{\frac{l}{2}} \cdot \frac{\Gamma_l(\frac{\alpha_{w,k} - n + l + m_k}{2})}{\Gamma_l(\frac{\alpha_{w,k} - n + l}{2})} \cdot \frac{\det(\mathbf{T}_{\dagger,k}^{L,L})^{\alpha_{w,k}/2}}{\det(\mathbf{T}_{\dagger,k}^{L,L})^{(\alpha_{w,k} + m_k)/2}} \right),\tag{12}$$

where $\mathcal{D}^L$ denotes the data restricted to the l variables in L , $\mathcal{G}_C^L$ is the complete DAG on those l nodes, and $\mathbf{T}_{\dagger,k}^{L,L}$ and $\mathbf{T}_{\dagger,k}^{L,L}$ are the submatrices of $\mathbf{T}_{\dagger,k}$ and $\mathbf{T}_{\dagger,k}$ corresponding to the variables in L .

2.3.2 Tied-covariance mixture Gaussian Bayesian networks

In the tied-covariance mixture model, in contrast to the full-covariance case in Eq. (7), we assume that all components share the same covariance matrix $\boldsymbol{\Sigma}$, while only the mean vectors $\boldsymbol{\mu}_k$ are component-specific:

$$p(\mathbf{x}_i | \boldsymbol{\Theta}) = \sum_{k=1}^K \pi_k \mathcal{N}_n(\mathbf{x}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}).\tag{13}$$

We impose a Wishart prior on the precision matrix $\mathbf{W} = \boldsymbol{\Sigma}^{-1}$. Conditional on $\mathbf{W}$, we then assign independent Gaussian priors to the component-specific mean vectors $\boldsymbol{\mu}_k$ ($k = 1, \dots, K$):

$$\mathbf{W} \sim \mathcal{W}_n(\alpha_w, \mathbf{T}_{\dagger}), \quad \boldsymbol{\mu}_k | \mathbf{W} \sim \mathcal{N}_n(\boldsymbol{\nu}_k, (\alpha_{\mu,k} \mathbf{W})^{-1}),\tag{14}$$

where $\mathbf{T}_\dagger \in \mathbb{R}^{n \times n}$ is positive definite, $\boldsymbol{\nu}_k \in \mathbb{R}^n$, $\alpha_w > n - 1$, and $\alpha_{\mu,k} > 0$ are hyperparameters. As before, α_w and $\alpha_{\mu,k}$ can be interpreted as prior equivalent sample sizes (Geiger and Heckerman, 2002).

The shared precision matrix $\mathbf{W}$ links the mixture components, so that we no longer have K independent component-specific Gaussian models. Nevertheless, the posterior distribution and the marginal likelihood can still be computed analytically. Since the derivations are lengthy and technical, they are presented in **Part A of the supplementary material**; here we report only the results. For a complete DAG $\mathcal{G}_C$ and a given assignment vector $\mathbf{z}$, the tied-covariance model yields the following posterior distribution:

$$\begin{aligned} \mathbf{W} | (\mathcal{D}, \mathbf{z}, \mathcal{G}_C) &\sim \mathcal{W}_n(\alpha_w + m, \mathbf{T}_\dagger), \\ \boldsymbol{\mu}_k | (\mathcal{D}, \mathbf{z}, \mathcal{G}_C, \mathbf{W}) &\sim \mathcal{N}(\boldsymbol{\mu}_k^\diamond, ((\alpha_{\mu,k} + m_k)\mathbf{W})^{-1}), \end{aligned} \quad (15)$$

with

$$\begin{aligned} \mathbf{T}_\dagger &:= \mathbf{T}_\dagger + \sum_{k=1}^K \mathbf{T}_{[k]}, \text{ and } \boldsymbol{\mu}_k^\diamond := \frac{\alpha_{\mu,k} \boldsymbol{\nu}_k + m_k \bar{\mathbf{x}}_{[k]}}{\alpha_{\mu,k} + m_k}, \\ \text{where } \mathbf{T}_{[k]} &:= \mathbf{S}_{[k]} + \frac{\alpha_{\mu,k} \cdot m_k}{\alpha_{\mu,k} + m_k} (\boldsymbol{\nu}_k - \bar{\mathbf{x}}_{[k]})(\boldsymbol{\nu}_k - \bar{\mathbf{x}}_{[k]})^\top. \end{aligned} \quad (16)$$

For a complete DAG $\mathcal{G}_C$, the marginal likelihood is

$$p(\mathcal{D} | \mathbf{z}, \mathcal{G}_C) = \pi^{-\frac{n \cdot m}{2}} \cdot \prod_{k=1}^K \left(\frac{\alpha_{\mu,k}}{\alpha_{\mu,k} + m_k} \right)^{\frac{n}{2}} \cdot \frac{\Gamma_n(\frac{\alpha_w + m}{2})}{\Gamma_n(\frac{\alpha_w}{2})} \cdot \frac{\det(\mathbf{T}_\dagger)^{\alpha_w/2}}{\det(\mathbf{T}_\dagger)^{(\alpha_w + m)/2}}. \quad (17)$$

For an l -dimensional subset $L \subset \{X_1, \dots, X_n\}$, this implies

$$p(\mathcal{D}^L | \mathbf{z}, \mathcal{G}_C^L) = \pi^{-\frac{l \cdot m}{2}} \cdot \prod_{k=1}^K \left(\frac{\alpha_{\mu,k}}{\alpha_{\mu,k} + m_k} \right)^{\frac{l}{2}} \cdot \frac{\Gamma_l(\frac{\alpha_w - n + l + m}{2})}{\Gamma_l(\frac{\alpha_w - n + l}{2})} \cdot \frac{\det(\mathbf{T}_\dagger^{L,L})^{\alpha_w/2}}{\det(\mathbf{T}_\dagger^{L,L})^{(\alpha_w + m)/2}}, \quad (18)$$

where $\mathcal{D}^L$ is the data restricted to the variables in L , $\mathcal{G}_C^L$ is the complete DAG on L , and $\mathbf{T}_\dagger^{L,L}$ and $\mathbf{T}_\dagger^{L,L}$ are the corresponding submatrices of $\mathbf{T}_\dagger$ and $\mathbf{T}_\dagger$.

2.3.3 Comparison of full-covariance and tied-covariance models

To highlight the differences between the two marginal likelihoods, we set $\mathbf{T}_{\dagger,k} = \mathbf{T}_\dagger$, $\alpha_{w,k} = \alpha_w$, and $\alpha_{\mu,k} = \alpha_\mu$. For complete DAGs $\mathcal{G}_C$, the marginal likelihoods for the full-covariance (M1) and tied-covariance (M2) mixture GBNs are then given by

$$\begin{aligned} p(\mathcal{D} | \mathbf{z}, \mathcal{G}_C, \mathcal{M}_1) &= \tilde{c} \cdot \prod_{k=1}^K \frac{\Gamma_n(\frac{\alpha_w + m_k}{2})}{\Gamma_n(\frac{\alpha_w}{2})} \cdot \prod_{k=1}^K \frac{\det(\mathbf{T}_\dagger)^{\alpha_w/2}}{\det(\mathbf{T}_\dagger + \mathbf{T}_{[k]})^{(\alpha_w + m_k)/2}}, \\ p(\mathcal{D} | \mathbf{z}, \mathcal{G}_C, \mathcal{M}_2) &= \tilde{c} \cdot \frac{\Gamma_n(\frac{\alpha_w + \sum_k m_k}{2})}{\Gamma_n(\frac{\alpha_w}{2})} \cdot \frac{\det(\mathbf{T}_\dagger)^{\alpha_w/2}}{\det(\mathbf{T}_\dagger + \sum_k \mathbf{T}_{[k]})^{(\alpha_w + \sum_k m_k)/2}}, \end{aligned}$$

where $\tilde{c} := \pi^{-\frac{n \cdot m}{2}} \cdot \prod_{k=1}^K \left(\frac{\alpha_\mu}{\alpha_\mu + m_k} \right)^{\frac{n}{2}}$ is a constant. For $K = 1$, both marginal likelihoods reduce to the standard marginal likelihood of a Gaussian model with a Normal–Wishart prior, i.e. the BGe score of a homogeneous GBN.

2.3.4 Marginal likelihood and posterior distribution of any DAG

For both mixture models, the marginal likelihood for any DAG $\mathcal{G}$ can be computed using the factorization of Geiger and Heckerman (1994, 2002); Kuipers et al. (2014):

$$p(\mathcal{D}|\mathbf{z}, \mathcal{G}) = \prod_{i=1}^n \frac{p(\mathcal{D}^{\{X_i, \mathbf{\Pi}_i\}}|\mathbf{z}, \mathcal{G}_C^{\{X_i, \mathbf{\Pi}_i\}})}{p(\mathcal{D}^{\{\mathbf{\Pi}_i\}}|\mathbf{z}, \mathcal{G}_C^{\{\mathbf{\Pi}_i\}})}, \quad (19)$$

where $\mathbf{\Pi}_i$ denotes the parent node set of X_i implied by $\mathcal{G}$, and $p(\mathcal{D}^{\{X_i, \mathbf{\Pi}_i\}}|\mathbf{z}, \mathcal{G}_C^{\{X_i, \mathbf{\Pi}_i\}})$ and $p(\mathcal{D}^{\{\mathbf{\Pi}_i\}}|\mathbf{z}, \mathcal{G}_C^{\{\mathbf{\Pi}_i\}})$ are marginal likelihoods for the subsets $L_{i,1} := \{X_i, \mathbf{\Pi}_i\}$ and $L_{i,2} := \{\mathbf{\Pi}_i\}$. These marginal likelihoods for variable subsets can be computed with either Eq. (12) or Eq. (18), depending on whether the full-covariance or the tied-covariance approach is applied.

Covariance matrices $\mathbf{\Sigma}_k = \mathbf{W}_k^{-1}$ sampled from Eq. (9) or $\mathbf{\Sigma} = \mathbf{W}^{-1}$ sampled from Eq. (15) correspond to complete DAGs $\mathcal{G}_C$, since they do not encode any conditional independence relations. To obtain posterior samples of covariance matrices that are consistent with the conditional independencies implied by an arbitrary DAG $\mathcal{G}$, we apply the algorithm of Grzegorzcyk (2023).

2.4 Joint posterior inference of class assignments and DAGs

In Section 2.3 we assumed that the number of mixture components K and the assignment vector $\mathbf{z}$ were fixed and known. To infer both from the data, we place a Dirichlet prior with parameters $\alpha_1 = \dots = \alpha_K = 1$ on the mixture weights $\pi_1, \dots, \pi_K$, and, following Newman (2025), we exclude empty components.² For any assignment vector $\mathbf{z} = (z_1, \dots, z_m)$ with component counts $(m_1, \dots, m_K)$ such that $m_k > 0$ for all $k = 1, \dots, K$, the induced distribution on assignments (obtained by integrating out the mixture weights under a Dirichlet(1, ..., 1) prior) is the Dirichlet–multinomial distribution conditioned on non-empty components (Newman, 2025):

$$p(\mathbf{z}) = \binom{m-1}{K-1}^{-1} \cdot \frac{\prod_{k=1}^K m_k!}{m!}. \quad (20)$$

Conditional on $\mathbf{z}$, the posterior distribution of the mixture weights is

$$p(\pi_1, \dots, \pi_K|\mathbf{z}) = \frac{\Gamma(m+K)}{\prod_{k=1}^K \Gamma(m_k+1)} \cdot \prod_{k=1}^K \pi_k^{m_k}, \quad (21)$$

which is the Dirichlet density with parameters $\tilde{\alpha}_k = m_k + 1$ ($k = 1, \dots, K$).

Furthermore, since empty components are ruled out, the number of mixture components K is fully determined by the assignment vector $\mathbf{z}$. To encourage sparser models, we place a Poisson prior with rate parameter $\lambda > 0$ on K

$$p(K) = \frac{\lambda^K \cdot e^{-\lambda}}{K!}.$$

For the joint posterior distribution of a DAG $\mathcal{G}$ and the assignment vector $\mathbf{z}$, we have:

$$p(\mathcal{G}, \mathbf{z}|\mathcal{D}) \propto p(\mathcal{D}|\mathcal{G}, \mathbf{z}) \cdot p(\mathcal{G}) \cdot p(\mathbf{z}) \cdot p(K), \quad (22)$$

where $p(\mathcal{D}|\mathcal{G}, \mathbf{z})$ is the marginal likelihood from Eq. (19), $p(\mathcal{G})$ is the prior on DAGs, $p(\mathbf{z})$ is the Dirichlet–multinomial distribution over assignments, and the number of mixture components K is determined by $\mathbf{z}$.

²That is, we exclude all assignment vectors $\mathbf{z}$ for which $m_k = 0$ for some $k \in \{1, \dots, K\}$.

2.5 MCMC sampling scheme for mixture GBNs

To generate samples from the joint posterior distribution in Eq. (22), we combine the classical structure MCMC sampler for Bayesian networks (Madigan and York, 1995; Giudici and Castelo, 2003) with the recently proposed fast collapsed Gibbs sampler for mixture models by Newman (2025).³ Given the current state $[\mathcal{G}, \mathbf{z}]$, where $\mathcal{G}$ is a DAG on the n variables and $\mathbf{z} = (z_1, \dots, z_m)$ is the assignment vector of the mixture model, our MCMC scheme alternates between two sampling steps:

Step 1: Metropolis-Hastings move on the DAG:

Given the current state $[\mathcal{G}, \mathbf{z}]$, let $N(\mathcal{G})$ denote the neighborhood of the DAG $\mathcal{G}$, i.e. the set of all DAGs that can be reached from $\mathcal{G}$ by adding, deleting or reversing one single edge. We propose a candidate $\mathcal{G}^* \in N(\mathcal{G})$ by selecting a neighbor at random. The acceptance probability of the move is:

$$A([\mathcal{G}, \mathbf{z}] \rightarrow [\mathcal{G}^*, \mathbf{z}]) = \min \left\{ 1, \frac{p(\mathcal{D}|\mathcal{G}^*, \mathbf{z})}{p(\mathcal{D}|\mathcal{G}, \mathbf{z})} \cdot \frac{p(\mathcal{G}^*)}{p(\mathcal{G})} \cdot \text{HR}(\mathcal{G}, \mathcal{G}^*) \right\} \quad (23)$$

with the Hastings ratio $\text{HR}(\mathcal{G}, \mathcal{G}^*) = \frac{|N(\mathcal{G})|}{|N(\mathcal{G}^*)|}$, where $|\cdot|$ denotes the cardinality of the corresponding neighborhood. If the move is accepted, the state is updated to $[\mathcal{G}^*, \mathbf{z}]$; otherwise it remains at $[\mathcal{G}, \mathbf{z}]$.

Step 2: Collapsed Gibbs move on the assignment vector:

Given the current state $[\mathcal{G}, \mathbf{z}]$, the assignment vector $\mathbf{z} = (z_1, \dots, z_m)$ implies K mixture components. Let $\mathbf{z}_{-i}$ denote the assignment vector with the i -th observation removed, i.e. $\mathbf{z}_{-i} = (z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_m)$.

- First, choose one of the K components uniformly at random. From this component k , select one of its m_k observations uniformly at random. This yields an observation index $i \in \{1, \dots, m\}$ assigned to component k . Remove observation i from k . If $m_k = 1$, the removal leaves an empty component; in this case, delete component k and decrement all subsequent component indices $k+1, \dots, K$ by one.
- After removal (and possible deletion), let $\tilde{K}$ be the current number of components. For $s = 1, \dots, \tilde{K} + 1$, compute the unnormalized weights:

$$w_s = \begin{cases} \frac{m - \tilde{K}}{\tilde{K}} \cdot p(\tilde{K}) \cdot p(\mathcal{D}|\mathcal{G}, \mathbf{z}_{-i}, z_i = s), & s = 1, \dots, \tilde{K}, \\ \tilde{K} \cdot p(\tilde{K} + 1) \cdot p(\mathcal{D}|\mathcal{G}, \mathbf{z}_{-i}, z_i = \tilde{K} + 1), & s = \tilde{K} + 1. \end{cases}$$

and the corresponding probabilities

$$p_s = \frac{w_s}{\sum_{\nu=1}^{\tilde{K}+1} w_\nu}, \quad s = 1, \dots, \tilde{K} + 1.$$

- Finally, re-allocate observation i to one of the components $1, \dots, \tilde{K} + 1$, choosing component s with probability p_s . The new state is then $[\mathcal{G}, \mathbf{z}^*]$, where $\mathbf{z}^* = (z_1, \dots, z_{i-1}, s, z_{i+1}, \dots, z_m)$.

³The fast Gibbs sampler of Newman (2025) applies to mixture models with a categorical-Dirichlet prior and analytically computable marginal likelihoods.

Repeatedly performing the two sampling steps yields a Markov chain whose stationary distribution is the joint posterior in Eq. (22). We run the algorithm for T iterations. To account for burn-in and to reduce autocorrelation, we discard the first 50% of the output and thin the remaining 50% by a factor of ξ . This results in a posterior sample of size $R = \lfloor \frac{T}{2\xi} \rfloor$:

$$[\mathcal{G}^{(1)}, \mathbf{z}^{(1)}], \dots, [\mathcal{G}^{(R)}, \mathbf{z}^{(R)}]. \quad (24)$$

2.6 Evaluation of the posterior sample

We use the posterior sample from Eq. (24) for a range of purposes:

- From the posterior samples of the assignment vectors $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(R)}$ we compute the **pair-wise co-allocation probabilities**, defined as the probability that two observations i and j are assigned to the same mixture component. These probabilities can be visualized in a heatmap, where block structures indicate groups of observations that are consistently co-assigned across the posterior sample.
- From the sampled DAGs, we compute the **marginal edge posterior probability** of an edge $X_i \rightarrow X_j$ as

$$\hat{p}_{i,j} := \frac{1}{R} \sum_{r=1}^R \mathcal{I}_{X_i \rightarrow X_j}(\mathcal{G}^{(r)}),$$

where $\mathcal{I}_{X_i \rightarrow X_j}(\mathcal{G}) = 1$ if the CPDAG of $\mathcal{G}$ contains either the directed edge $X_i \rightarrow X_j$ or the undirected edge $X_i - X_j$, and $\mathcal{I}_{X_i \rightarrow X_j}(\mathcal{G}) = 0$ otherwise.⁴

- Imposing a threshold $\psi \in [0, 1]$ on the edge probabilities $\hat{p}_{i,j}$ yields the **network prediction** $\hat{\mathcal{G}}_\psi$, which contains only those edges with posterior probability exceeding ψ .
- To assess the **network reconstruction accuracy** beyond a fixed threshold ψ , we evaluate the **area under the precision–recall curve (AUC)**. Given the true CPDAG, each edge posterior probability $\hat{p}_{i,j}$ serves as a score for predicting whether the edge is present. Undirected edges in the true CPDAG are treated as undirected, meaning that either orientation $X_i \rightarrow X_j$ or $X_j \rightarrow X_i$ is counted as a correct recovery; this interpretation is consistent with the edge score computation. Varying the threshold ψ from 1 down to 0 yields a sequence of precision–recall pairs, where precision is the proportion of predicted edges that are true, and recall is the proportion of true edges that are recovered. The $\text{AUC} \in [0, 1]$ summarizes this curve into a single measure, with higher values indicating better agreement between the inferred edge probabilities and the true CPDAG.
- For every posterior sample $[\mathcal{G}^{(r)}, \mathbf{z}^{(r)}]$ we can draw model parameters with covariance matrices coherent with $\mathcal{G}^{(r)}$ (Grzegorzczuk, 2023):

$$\Theta^{(r)} = \begin{cases} \{(\pi_1^{(r)}, \boldsymbol{\mu}_1^{(r)}, \boldsymbol{\Sigma}_1^{(r)}), \dots, (\pi_{K_r}^{(r)}, \boldsymbol{\mu}_{K_r}^{(r)}, \boldsymbol{\Sigma}_{K_r}^{(r)})\}, & \text{(full-covariance),} \\ \{(\pi_1^{(r)}, \boldsymbol{\mu}_1^{(r)}), \dots, (\pi_{K_r}^{(r)}, \boldsymbol{\mu}_{K_r}^{(r)})\}, \boldsymbol{\Sigma}^{(r)}, & \text{(tied-covariance),} \end{cases}$$

where the number of components K_r is determined by $\mathbf{z}^{(r)}$.

⁴Because DAGs fall into equivalence classes, each DAG is first replaced by its CPDAG. If the CPDAG has the undirected edge $X_i - X_j$, we interpret it as a bidirectional edge $X_i \leftrightarrow X_j$. Under this interpretation, the undirected edge implies that $X_i \rightarrow X_j$ is present.

- For a new dataset with m^* observations $\mathcal{D}^* := \{\mathbf{x}_1^*, \dots, \mathbf{x}_{m^*}^*\}$, we compute the **geometric mean predictive probability** using Monte Carlo integration. The likelihood of the new data is averaged over the posterior samples of the parameters obtained from $\mathcal{D}$ and then normalized by the m^* -th root:

$$p(\mathcal{D}^* | \mathcal{D}) := \left[\frac{1}{R} \sum_{r=1}^R \prod_{i=1}^{m^*} p(\mathbf{x}_i^* | \boldsymbol{\Theta}^{(r)}) \right]^{1/m^*},$$

where

$$p(\mathbf{x}_i^* | \boldsymbol{\Theta}^{(r)}) = \begin{cases} \sum_{k=1}^{K_r} \pi_k^{(r)} \mathcal{N}_n(\mathbf{x}_i^* | \boldsymbol{\mu}_k^{(r)}, \boldsymbol{\Sigma}_k^{(r)}), & \text{(full-covariance),} \\ \sum_{k=1}^{K_r} \pi_k^{(r)} \mathcal{N}_n(\mathbf{x}_i^* | \boldsymbol{\mu}_k^{(r)}, \boldsymbol{\Sigma}^{(r)}), & \text{(tied-covariance).} \end{cases}$$

2.7 Implementation details

If no prior knowledge about the network is available, we assume all DAGs $\mathcal{G}$ to be equally likely. The hyperparameters of the Normal–Wishart priors are set to yield uninformative (weak) priors. Specifically, we take $\mathbf{T}_{\dagger,k} = \mathbf{T}_{\dagger} = \mathbf{I}$ and $\boldsymbol{\nu}_k = \boldsymbol{\nu} = \mathbf{0}$, where $\mathbf{I}$ is the identity matrix and $\mathbf{0}$ the zero vector. For the equivalent sample size parameters we choose very small values: $\alpha_{w,k} = \alpha_w = n + 1$ (recall that $\alpha_w > n - 1$ is required) and $\alpha_{\mu,k} = \alpha_{\mu} = 1$. This makes the Normal–Wishart priors weakly informative. For example, in the tied-covariance model the priors are: $\mathbf{W} \sim \mathcal{W}_n(n + 1, \mathbf{I})$ and for each component k : $\boldsymbol{\mu}_k | \mathbf{W} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{W}^{-1})$. Finally, we set the rate parameter of the Poisson prior on the number of mixture components to $\lambda = 1$.

To assess convergence, we selected a few representative datasets and ran independent MCMC simulations on each of them. Standard convergence diagnostics were applied, such as trace plots and scatter plots of the edge scores (Grzegorzcyk, 2023, 2024), to compare the outputs of the independent runs. We determined the required number of iterations T as the point at which the trace plots reached stable plateaus and the independent runs produced nearly identical marginal posterior probabilities.

The results reported in Sections 4.1 and 4.2 are based on 250,000 iterations, while those in Section 4.3 are based on 1,000,000 iterations. In each case, the first $T/2$ iterations were discarded as burn-in, and a thinning factor ξ was applied so that the final posterior sample consisted of $R = 500$ draws from the second half of the Markov chain.

Software: Upon acceptance, we will share the R code at <https://github.com/MarcoAndreas/MIX-GBN>

3 Data

We compare model performance on both simulated and real-world datasets. For the simulated data, the true networks (CPDAGs) are known, allowing us to evaluate network reconstruction accuracy using AUC and rSHD scores. For the real-world datasets, where the true network structures are unknown, we instead assess performance through predictive probabilities.

3.1 Simulated network data

We assume that the network domain consists of n variables $X_1, \dots, X_n$ arranged in topological order. This yields $\zeta = n(n-1)/2$ possible directed edges $X_j \rightarrow X_i$ with $j < i$. Each edge is assumed to be present independently with probability p , and we set $p = 20/\zeta$ so that the expected number of edges equals 20. In a Gaussian Bayesian network with zero mean vector, $\boldsymbol{\mu} = \mathbf{0}$, each observation v satisfies

$$x_{i,v} = \sum_{j \in \boldsymbol{\Pi}_i} b_{i,j} x_{j,v} + \epsilon_{i,v}, \quad (25)$$

where $x_{i,v}$ is the value of variable X_i , $\boldsymbol{\Pi}_i \subset \{X_1, \dots, X_{i-1}\}$ is its parent set, and $\epsilon_{i,v} \sim \mathcal{N}(0, \sigma_i^2)$ denotes Gaussian noise. Eq. (25) involves the parameters $\{b_{i,j} : j \in \boldsymbol{\Pi}_i\}$ and σ_i . We sample these parameters from uniform distributions on $[1, 2]$ and assign a random sign to each $b_{i,j}$. Finally, we rescale the parameters so that their Euclidean norm equals 1:

$$\sigma_i^2 + \sum_{j \in \boldsymbol{\Pi}_i} b_{i,j}^2 = 1.$$

The joint distribution of $\mathbf{X} = (X_1, \dots, X_n)^\top$ is Gaussian with covariance matrix $\boldsymbol{\Sigma}$ determined by the rescaled parameters $\{b_{i,j} : j \in \boldsymbol{\Pi}_i\}$ and σ_i ($i = 1, \dots, n$) (Shachter and Kenley, 1989). We generate m observations

$$\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_m \sim \mathcal{N}_n(\mathbf{0}, \boldsymbol{\Sigma}).$$

We assume K mixture components, each associated with a mean vector $\boldsymbol{\mu}_k$. The mean vectors are sampled from standard Gaussian distributions:

$$\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K \sim \mathcal{N}_n(\mathbf{0}, \mathbf{I}).$$

Each observation $\tilde{\mathbf{x}}_v$ is then randomly assigned to a component k , with equal probabilities $\pi_k = 1/K$. If $\tilde{\mathbf{x}}_v$ is assigned to k , we add the corresponding mean vector, so that $\mathbf{x}_v \leftarrow \tilde{\mathbf{x}}_v + \boldsymbol{\mu}_k$. Thus, each $\mathbf{x}_v$ is drawn from the tied-covariance mixture GBN in Eq. (13), with equal mixture weights.

3.2 Benchmark datasets

We use four benchmark datasets that are widely used in the statistics and machine learning literature. Within each dataset, all n continuous variables are z -score standardized to have mean 0 and variance 1.⁵

3.2.1 Fish

The Fish Market dataset (Aung Pyae, 2019) contains $m = 159$ fish samples from a market in Finland, belonging to $K = 7$ species: *Bream*, *Roach*, *Whitefish*, *Parkki*, *Perch*, *Pike*, and *Smelt*. The dataset has $n = 6$ variables: weight, vertical length, diagonal length, cross length, height, and diagonal width. Although often presented as real-world data, it remains unclear whether these fish were actually caught in the wild.

⁵The data can, for instance, be obtained from <https://www.kaggle.com/datasets>.

3.2.2 Iris

The classic Iris dataset (Fisher, 1936) contains $m_i = 50$ samples from each of $K = 3$ species: *I. setosa*, *I. versicolor*, and *I. virginica*. It has $n = 4$ variables: sepal length, sepal width, petal length, and petal width.

3.2.3 Wine

The Wine dataset (Forina et al., 1986) contains chemical analyses of $m = 178$ wines from $K = 3$ cultivars of the same Italian region. It includes $n = 13$ variables: alcohol, malic acid, ash, alkalinity of ash, magnesium, total phenols, flavonoids, nonflavonoid phenols, proanthocyanins, color intensity, hue, OD280/OD315 (of diluted wines), and proline.

3.2.4 Seeds

The Wheat Seeds Kernel dataset (Charytanowicz et al., 2010) contains $m_i = 70$ kernel samples from each of $K = 3$ wheat varieties: *Kama*, *Rosa*, and *Canadian*. It has $n = 7$ variables: area, perimeter, compactness, kernel length, kernel width, asymmetry coefficient, and kernel groove length.

4 Results

In Section 4.1 we use simulated data to compare model performance in terms of network reconstruction accuracy. The simulation study consists of two parts: first, we assume that the class labels are known; second, we treat the class labels as latent and infer them from the data. In Section 4.2 we evaluate the models on the four benchmark datasets introduced in Section 3.2. Since the true networks are unknown, performance is assessed using predictive probabilities. In Section 4.3 we evaluate the four benchmark datasets with the new tied-covariance mixture GBN model (M2) and present the empirical results.

4.1 Simulated data

We generate data as described in Section 3.1. We cross four sample sizes, $m \in \{100, 200, 400, 800\}$, with four numbers of mixture components, $K \in \{2, 4, 8, 16\}$. For each (m, K) combination we generate 25 independent replicates, resulting in a total of 400 simulated datasets.

Part 1: First, we assume that the class labels are known, i.e. the m observations are correctly allocated to the K mixture components. In this setting, we compare the performance of five models.

- **H:** The homogeneous Gaussian Bayesian network (GBN), using the BGe score. All m observations are merged into a single component.
- **M1:** The *full-covariance* mixture model, using the product of component-specific BGe scores from Section 2.3.1.
- **M2:** The *tied-covariance* mixture model, using the newly derived modified BGe score from Section 2.3.2.

- **clG**: The conditional linear Gaussian model for hybrid data, using the implementation in the **R** package *bnlearn* (Scutari and Denis, 2014; Scutari, 2020; Bodewes and Scutari, 2021).
- **mBGe**: The mean-adjusted Gaussian Bayesian network for hybrid data, using the approach from Grzegorzczak (2025).

The clG and mBGe models are designed for hybrid data, i.e., datasets that contain both discrete and continuous variables. Both models allow edges from discrete to continuous variables. In principle, we could treat the component allocation as a single discrete variable with K categories; however, if this variable were made a parent of all continuous nodes, clG and mBGe would effectively reduce to the full-covariance (M1) and tied-covariance (M2) mixture models with known class labels, respectively. To avoid this, we instead represent the class membership by $\log_2(K)$ binary indicator variables and include them alongside the continuous variables. Each value combination of the binary indicators corresponds to a specific k , and clG and mBGe are then applied in the usual way, i.e., structure learning infers which indicators serve as parents of which continuous variables. This setup is consistent with the study reported in Section 5.1 of Grzegorzczak (2025).⁶

Figure 1 shows the average AUC scores for each combination of m and K . Overall, the tied-covariance mixture model (M2) performs best and yields the highest AUC scores. H (BGe) and clG are consistently outperformed by M2. The models M1, M2, and mBGe perform similarly for $K = 2$, but M2 achieves higher AUCs for $K \geq 4$. We also compared model performance in terms of relative structural Hamming distance (rSHD). The rSHD results, shown in Figure 1 of **Part B of the supplementary material**, are in good agreement with the AUC findings.

Part 2: Second, we assume the class labels to be latent (i.e., unknown) and focus on the performance of the two mixture GBN models, M1 (full-covariance) and M2 (tied-covariance). For details we refer to Sections 2.4 and 2.5.⁷ We also include the homogeneous Gaussian Bayesian network (H) as a reference model. However, since H ignores the heterogeneity of the observations and merges them into a single component, its performance does not differ from **Part 1**.

Figure 2 shows the method-specific average AUC scores with latent class labels. The AUC scores of M2 decrease only slightly, indicating that the class labels can indeed be inferred from the data. Again, M2 consistently achieves the highest AUC scores. As expected, the AUC scores of both mixture models decrease with the number of mixture components K , but increase with the sample size m . The relative structural Hamming distance (rSHD) results are shown in Figure 2 of **Part B of the supplementary material** and are in good agreement with the AUC findings.

⁶Note that the notation differs from the present paper: in Grzegorzczak (2025) the models were denoted BGe (H), hBe (M1), clG, and mBGe, and were applied to learn Bayesian networks from hybrid data. For hybrid data, the hBe (M1) model assumes that each of the K value combinations of the discrete variables corresponds to a mixture component with separate network parameters. In Grzegorzczak (2025), this assumption introduced many redundant components and thus weakened the performance of the hBe (M1) model. In our study, however, this assumption is beneficial, as it aligns with the ground truth.

⁷Latent class labels disqualify the two hybrid data models (clG and mBGe). In the absence of discrete variables, they reduce to homogeneous GBNs.




Figure 1: **Areas under the Precision-Recall Curve (AUCs) for Simulated Data with Known Component Labels.** Each panel shows a bar chart of the average AUC scores with error bars indicating 95% confidence intervals. Each of the five bars refers to a specific method (see legend). The rows and columns refer to sample sizes m and numbers of mixture components K , respectively. The models clG and mBGe required inferring the data allocation from a set of binary variables.




Figure 2: **Areas under the Precision-Recall Curve (AUCs) for Simulated Data with Latent (Unknown) Component Labels.** Each panel shows a bar chart of the average AUC scores with error bars indicating 95% confidence intervals. The rows refer to four sample sizes m , and the columns refer to four numbers of mixture components K .

4.2 Predictive probabilities for benchmark data

For the four benchmark datasets from Section 3.2 the true networks (CPDAGs) are unknown, and therefore we evaluate model performance using predictive probabilities. Although the class labels are known, we treat them as latent to demonstrate that they can be inferred from the data.

When a benchmark dataset comprises m_{all} samples, we split it into a training set $\mathcal{D}$ and a validation set $\mathcal{D}^*$. For the training set, we randomly select $m \in \{25, 50, 100, 150\}$ samples from the m_{all} available, and use the remaining $m^* = m_{\text{all}} - m$ observations for validation.⁸ To ensure comparability across validation sets of different sizes m^* , we compute the geometric mean predictive probabilities, as described in Section 2.6.

Figure 3 shows boxplots of the geometric mean predictive probabilities. As expected, predictive performance improves with increasing sample size m , while relative differences between models tend to diminish. Overall, the new M2 tied-covariance mixture GBN model yields the highest predictive probabilities, whereas the homogeneous GBN (H) model consistently yields the lowest. For the four benchmark datasets, we observe the following trends:

- **Fish:** The tied-covariance mixture GBN (M2) performs best, followed by the full-covariance mixture GBN (M1), with the homogeneous GBN (H) performing worst.
- **Iris:** The homogeneous GBN (H) performs poorly. Between the mixture models, M2 is slightly better only at $m = 25$; for larger m , their performances are similar.
- **Wine:** M2 again performs best overall. At $m = 25$, the homogeneous GBN (H) is competitive, while at $m \geq 100$, M1 approaches M2’s performance.
- **Seeds:** As with the Fish data, M2 clearly outperforms the other models. Both H and M1 perform poorly, suggesting that M2 strikes the right balance between the under-complex H and the over-complex M1.

4.3 Applications to benchmark data

In Section 4.2, the tied-covariance mixture GBN (M2) achieved the highest predictive probabilities overall. We therefore adopt this model for analyzing the four benchmark datasets. The purpose of this application study is twofold: (i) to demonstrate that the underlying class allocations can indeed be learned from the data, and (ii) to predict the unknown network structures (CPDAGs). Although these benchmark datasets are well known and widely used, to the best of our knowledge no network structures have been reported for them in the literature.

Figure 4 shows heatmaps of the pairwise co-allocation probabilities. The four heatmaps reveal the following trends:

- **Fish:** Approximately $K = 7$ groups are inferred. These groups correspond broadly, but not one-to-one, to the fish species. The three largest groups are: *Bream* (1–35); *Roach* together with most of the *Perch* (36–55 and 76–106); and the remaining *Perch* with three *Whitefish* (107–128 and 59–61). In addition, *Parkki* (62–72), *Pike* (129–141), and *Smelt* together with three *Perch* (146–159 and 73–75) each form smaller groups.

⁸The Iris dataset has only $m_{\text{all}} = 150$ samples. We therefore restrict the maximum training size to $m = 145$.




Figure 3: **Geometric Mean Predictive Probabilities (PPs) for Benchmark Data.** Each figure shows three boxplots of the distribution of log predictive probabilities (PPs) computed from 100 random data subsets. The boxplots refer to the models H (white), M1 (light grey), and M2 (dark grey). The columns refer to four different datasets, and the rows refer to four sample sizes m .




Figure 4: **Heatmaps of pairwise co-allocation probabilities for the four benchmark datasets.** All panels show probabilities under the tied-covariance mixture GBN (M2) model. Axis tick marks indicate the positions of the true transitions between mixture components.

- **Iris:** Three clear groups are inferred. *I. setosa* (1–50) is completely separated from the other two species, while *I. versicolor* and *I. virginica* show partial overlap. A small subset of *I. virginica* flowers (101–150) appears separated from the rest.
- **Wine:** Two large groups are inferred, corresponding to *Cultivar 1* (1–59) and much of *Cultivar 3* (131–178). In contrast, the wines from *Cultivar 2* (36–55) do not form a distinct cluster but are spread across several smaller groups.
- **Seeds:** The three wheat varieties are almost perfectly separated, indicating that the class allocation can be recovered very accurately from the data.

Overall, the co-allocation heatmaps demonstrate that the model is able to recover meaningful group structures from the data. For Seeds and Iris the inferred groups align almost perfectly with the true classes, while for Fish and Wine the model uncovers finer substructures beyond the original labels, highlighting additional heterogeneity. These results confirm that the underlying data allocation can indeed be learned directly from the data.

Figures 5–6 show the four inferred networks (CPDAGs). For the **Iris**, **Fish**, and **Seeds** datasets, we display all edges with posterior scores above the threshold $\psi = 0.75$. For **Wine**, however, we relaxed the threshold to $\psi = 0.5$, since otherwise the network would have contained only four edges. We represent an undirected edge $X_i - X_j$ whenever one orientation ($X_i \rightarrow X_j$ or $X_j \rightarrow X_i$) exceeded the threshold ψ , and the opposite orientation had a comparably high score.⁹ Unlike the three smaller networks (**Iris**, **Seeds**, **Fish**) with $n \in \{4, 6, 7\}$ nodes, the **Wine** network with $n = 13$ nodes contains many directed edges. A detailed substantive interpretation of the inferred structures is beyond the scope of this paper. Nevertheless, we hope that these exemplary applications may encourage the use of these four widely known datasets also as benchmarks for graphical model analysis (cf. Section 5).

5 Conclusion and discussion

In this paper, we proposed an extension of the BGe scoring metric originally introduced by Geiger and Heckerman. For mixture data consisting of distinct components, the standard extension of the BGe score assumes that each mixture component has its own network parameters, namely a separate mean vector and covariance matrix. In contrast, our extension assumes component-specific mean vectors while allowing the covariance matrix to be shared across components. This construction is directly analogous to classical mixture models with full-covariance and tied-covariance, and we therefore referred to the two mixture GBMs as the full-covariance mixture model (M1) and the tied-covariance mixture model (M2).

Our empirical results on both simulated and benchmark data demonstrate that the M2 mixture model can yield significant improvements over the M1 mixture model. In particular, consistent with classical findings for mixture models in statistics, we observed that the tied-covariance model (M2) performs especially well at smaller sample sizes, while this advantage diminishes or may even reverse as sample sizes increase. These results underline the practical value of the new scoring metric as a complementary alternative to the standard full-covariance approach (M1). We have applied both mixture models to four well-known and widely used datasets (Fish, Iris, Wine, Seeds). Although these datasets have traditionally been employed primarily to evaluate

⁹We considered $X_i \rightarrow X_j$ and $X_j \rightarrow X_i$ to have similar scores $e_{i,j}$ and $e_{j,i}$ if $|e_{i,j} - e_{j,i}| < 0.25$.

classification accuracy, we believe their heterogeneity offers considerable potential for graphical modeling. We therefore hope that our study contributes to establishing them as datasets for Bayesian network analyses. While our focus was on mixtures of Gaussian Bayesian networks, these datasets also appear well suited for assessing other Bayesian network approaches, particularly models designed to analyze heterogeneous data arising from different experimental conditions, as, for example, proposed by Scutari et al. (2022).




Figure 5: **Inferred networks (CPDAGs) for the Iris and Seeds data.** The class labels were assumed to be unknown and the tied-covariance mixture model (M2) was applied to all data. Shown are the edges whose scores exceeded the threshold $\psi = 0.75$.




Figure 6: **Inferred networks (CPDAGs) for the Fish and Wine data.** The class labels were assumed to be unknown and the new tied-covariance mixture model (M2) was applied to all data. Shown are the edges whose scores exceeded the threshold $\psi = 0.75$ (**Fish**) and $\psi = 0.5$ (**Wine**), respectively.

References

- J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann, San Francisco, CA, USA, 1988.
- R. Neapolitan, Probabilistic Reasoning in Expert Systems: Theory and Algorithms, CreateSpace, Scotts Valley CA, USA, 1989.
- D. Koller, N. Friedman, Probabilistic Graphical Models: Principles and Techniques, Adaptive Computation and Machine Learning series, MIT Press, Cambridge, MA, USA, 2009. ISBN: 9780262013192.
- M. Scutari, J.-B. Denis, Bayesian Networks: With Examples in R, Chapman & Hall, Boca Raton, FL, 2014.
- D. M. Chickering, Learning Bayesian networks is NP-complete, in: D. Fisher, H. J. Lenz (Eds.), Learning from Data: Artificial Intelligence and Statistics, volume 5, Springer, New York, 1996, pp. 121–130.
- J. Kuipers, G. Moffa, Partition MCMC for inference on acyclic digraphs, Journal of the American Statistical Association 112 (2017) 282–299.
- M. Scutari, C. E. Graafland, J. M. Gutiérrez, Who learns better Bayesian network structures: Constraint-based, score-based or hybrid algorithms?, in: International Conference on Probabilistic Graphical Models, PMLR, 2018, pp. 416–427.
- M. Scutari, C. Graafland, J. Gutiérrez, Who learns better Bayesian network structures: Accuracy and speed of structure learning algorithms, International Journal of Approximate Reasoning 115 (2019) 235–253.
- M. Scutari, Bayesian network models for incomplete and dynamic data, Statistica Neerlandica 74 (2020) 397–419.
- N. K. Kitson, A. C. Constantinou, Z. Guo, Y. Liu, K. Chobtham, A survey of Bayesian network structure learning, Artificial Intelligence Review 56 (2023) 8721–8814.
- A. C. Constantinou, Y. Liu, N. K. Kitson, K. Chobtham, Z. Guo, Effective and efficient structure learning with pruning and model averaging strategies, International Journal of Approximate Reasoning 151 (2022) 292–321.
- D. Madigan, J. York, Bayesian graphical models for discrete data, International Statistical Review 63 (1995) 215–232.
- P. Giudici, R. Castelo, Improving Markov chain Monte Carlo model search for data mining, Machine Learning 50 (2003) 127–158.
- N. Friedman, D. Koller, Being Bayesian about network structure: A Bayesian approach to structure discovery in Bayesian networks, Machine Learning 50 (2003) 95–126.
- M. Grzegorzcyk, D. Husmeier, Improving the structure MCMC sampler for Bayesian networks by introducing a new edge reversal move, Machine Learning 71 (2008) 265–305.

- N. Friedman, Learning belief networks in the presence of missing values and hidden variables, in: D. H. Fisher (Ed.), *Proceedings of the Fourteenth International Conference on Machine Learning (ICML)*, Morgan Kaufmann, Nashville, Tennessee, USA, 1997, pp. 125–133.
- A. Ruggieri, F. Stranieri, F. Stella, M. Scutari, Hard and soft EM in Bayesian network learning from incomplete data, *Algorithms* 13 (2020).
- T. Bodewes, M. Scutari, Learning Bayesian networks from incomplete data with the node-average likelihood, *International Journal of Approximate Reasoning* 138 (2021) 145–160.
- X. Luo, G. Moffa, J. Kuipers, Learning Bayesian networks from ordinal data, *Journal of Machine Learning Research* 22 (2021) 1–44.
- C. F., Learning Bayesian networks: A Copula approach for mixed-type data., *Psychometrika* 89 (2024) 658–686.
- M. Scauda, J. Kuipers, G. Moffa, A latent causal inference framework for ordinal variables, 2025. [arXiv:2502.10276](https://arxiv.org/abs/2502.10276).
- D. Chickering, D. Geiger, D. Heckerman, Learning Bayesian networks: Search methods and experimental results, *Machine Learning* 20 (1995) 112–128.
- M. Grzegorzczuk, Being Bayesian about learning Gaussian Bayesian networks from incomplete data, *International Journal of Approximate Reasoning* 160 (2023). Article 108954.
- M. Grzegorzczuk, Being Bayesian about learning Bayesian networks from ordinal data, *International Journal of Approximate Reasoning* 170 (2024). Article 109205.
- M. Grzegorzczuk, Being Bayesian about learning Bayesian networks from hybrid data, *International Journal of Approximate Reasoning* 187 (2025). Article 109549.
- Y. Ko, C. Zhai, S. Rodriguez-Zas, Inference of gene pathways using Gaussian mixture models, in: *BIBM International Conference on Bioinformatics and Biomedicine*, Fremont, CA, 2007, pp. 362–367.
- M. Grzegorzczuk, D. Husmeier, K. Edwards, P. Ghazal, A. Millar, Modelling non-stationary gene regulatory processes with a non-homogeneous Bayesian network and the allocation sampler, *Bioinformatics* 24 (2008) 2071–2078.
- A. Nobile, A. Fearnside, Bayesian finite mixtures with an unknown number of components: The allocation sampler, *Statistics and Computing* 17 (2007) 147–162.
- D. Geiger, D. Heckerman, Learning Gaussian networks, in: *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, CA., 1994, pp. 235–243.
- D. Geiger, D. Heckerman, Parameter priors for directed acyclic graphical models and the characterization of several probability distributions, *The Annals of Statistics* 30 (2002) 1412–1440.
- J. Kuipers, G. Moffa, D. Heckerman, Addendum on the scoring of Gaussian directed acyclic graphical models, *The Annals of Statistics* 42 (2014) 1689–1691.

- M. E. J. Newman, Fast sampling and model selection for Bayesian mixture models, 2025. [arXiv:2501.07668](https://arxiv.org/abs/2501.07668), preprint.
- D. M. Chickering, Learning Equivalence Classes of Bayesian-Network Structures, *Journal of Machine Learning Research* 2 (2002) 445–498.
- D. Heckerman, D. Geiger, D. M. Chickering, Learning Bayesian networks: The combination of knowledge and statistical data, *Machine Learning* 20 (1995) 245–274.
- D. Heckerman, D. Geiger, Learning Bayesian networks: A unification for discrete and Gaussian domains, in: *Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, Morgan Kaufmann, San Francisco, CA, 1995, pp. 274–82.
- R. D. Shachter, R. Kenley, Gaussian influence diagrams, *Management Science* 35 (1989) 527–550.
- A. P. Aung Pyae, Fish Market Dataset, Technical Report, Kaggle, 2019. URL: <https://www.kaggle.com/aungpyaeap/fish-market>, accessed: 2025-08-21.
- R. Fisher, The use of multiple measurements in taxonomic problems, *Annals of Eugenics* 7 (1936) 179–188.
- M. Forina, C. Armanino, M. Castino, M. Ubigli, Multivariate data analysis as a discriminating method of the origin of wines, *Vitis* 25 (1986) 189–201.
- M. Charytanowicz, J. Niewczas, P. Kulczycki, P. A. Kowalski, S. Łukasik, S. Żak, Complete gradient clustering algorithm for features analysis of x-ray images, in: E. Pietka, J. Kawa (Eds.), *Information Technologies in Biomedicine*, volume 69 of *Advances in Soft Computing*, Springer, Berlin, Heidelberg, 2010, pp. 15–24.
- M. Scutari, C. Marquis, L. Azzimonti, Using mixed-effects models to learn bayesian networks from related data sets, in: *Proceedings of The 11th International Conference on Probabilistic Graphical Models*, volume 186 of *Proceedings of Machine Learning Research*, PMLR, 2022, pp. 73–84.