

# Automated Extraction of Mutual Independence Patterns Using Bayesian Comparison of Partition Models

Guillaume Marrelec<sup>1</sup> and Alain Giron<sup>2</sup>

**Abstract**—Mutual independence is a key concept in statistics that characterizes the structural relationships between variables. Existing methods to investigate mutual independence rely on the definition of two competing models, one being nested into the other and used to generate a null distribution for a statistic of interest, usually under the asymptotic assumption of large sample size. As such, these methods have a very restricted scope of application. In this article, we propose to change the investigation of mutual independence from a hypothesis-driven task that can only be applied in very specific cases to a blind and automated search within patterns of mutual independence. To this end, we treat the issue as one of model comparison that we solve in a Bayesian framework. We show the relationship between such an approach and existing methods in the case of multivariate normal distributions as well as cross-classified multinomial distributions. We propose a general Markov chain Monte Carlo (MCMC) algorithm to numerically approximate the posterior distribution on the space of all patterns of mutual independence. The relevance of the method is demonstrated on synthetic data as well as two real datasets, showing the unique insight provided by this approach.

**Index Terms**—Mutual independence, Bayesian analysis, model comparison, likelihood ratio criterion, minimum discrimination information statistic, Markov chain Monte Carlo, Gibbs sampling, parallel tempering

## 1 INTRODUCTION

MUTUAL independence is a key concept in statistics whose goal is to characterize the structural relationships between variables. As a consequence, a fundamental problem is to be able to determine from a sample of finite size whether some subsets of variables are mutually independent or not. To this end, several approaches have been proposed in the literature. For two- and three-way contingency tables, one can use the minimum discrimination information statistic against the null hypothesis of independence [1, Chap. 8, Sections 2 and 3.1]; in the multidimensional case, there is the chi-squared test for independence [2, Section 23.8]. For multivariate normal distributions, one can resort to the likelihood ratio criterion [3, Chap. 9] or, again, the minimum discrimination information statistic [1, pp. 306–307] against the null hypothesis of independence. Such approaches, however, have a major drawback, in that one needs two and only two competing models, one being nested into the other one and used to generate a null distribution for a statistic of interest, usually under the asymptotic assumption of large sample size; if the statistic is beyond a certain threshold, the null assumption is rejected,

but usually with no hint about the validity of the second model, not to mention the true underlying structure of dependence. In the case of several competing models, or when there is no hint regarding the underlying structure of mutual independence, the above-mentioned methods do not provide any way to approach the problem. As a consequence, these methods have a very restricted scope of application.

In the present study, we propose to change the investigation of mutual independence from a hypothesis-driven task that can only be applied in very specific cases to a blind and automated search within patterns of mutual independence. To this end, we develop a general framework for a data-driven investigation of all potential patterns of mutual independence. More specifically, we propose to treat the issue as one of model comparison that we solve in a Bayesian framework. A first step in this direction was proposed by [4], who used Bayesian model comparison to quantify the probability for two discrete variables to be independent. A second step was performed by [5], who showed that a Bayes factor, comparing two models with and without independence, provided a relevant measure of similarity for agglomerative hierarchical clustering in the case of two subvectors of a multivariate normal distribution. We here propose to go further and perform a full probabilistic exploration of the independence pattern underlying the data. Such an approach heavily relies on the one-to-one mapping that exists between a pattern of mutual independence between  $D$  variables and a partition of  $\{1, \dots, D\}$ . Comparing models of mutual independence is then equivalent to comparing partitions, which can be seen as a clustering problem and solved in the general framework of

• The authors are with the Laboratoire d'imagerie biomédicale (LIB), Sorbonne Université, CNRS, INSERM, F-75006 Paris, France, and also with the Centre de recherches et d'études en sciences des interactions (CRÉSI) — Center for Interaction Science (CIS), F-75006 Paris, France.  
E-mail: {guillaume.marrelec, alain.giron}@inserm.fr.

Manuscript received 25 June 2019; revised 13 Jan. 2020; accepted 14 Jan. 2020. Date of publication 22 Jan. 2020; date of current version 3 June 2021.  
(Corresponding author: Guillaume Marrelec.)

Recommended for acceptance by M. Sugiyama.

Digital Object Identifier no. 10.1109/TPAMI.2020.2968065

Bayesian model-based clustering [6], [7], [8], [9]. The present problem, however, can be solved by neither of the two main classes of clustering methods, namely the clustering of mixture models, e.g., [10], [11], [12], [13], [14], and the clustering of functional data or curves, e.g., [15], [16], [17], [18], [19], [20], but requires a class of its own. After specification of a model of dependence, the Bayesian machinery makes it possible, from a theoretical perspective, to calculate the posterior probability of any partition (corresponding to any pattern of mutual independence) given some data. From a practical point of view, it allows the resulting posterior distribution on the set of all partitions (corresponding to the set of all patterns of mutual independence) to be explicitly computed (if possible), or otherwise approximated through a numerical sampling scheme.

The outline of the article is the following. We first expose the problem and provide a theoretical treatment in the form of Bayesian model comparison. We then investigate the particular cases of multivariate normal and discrete distributions, providing asymptotic expressions for the log posterior distributions which turn out to be compatible with the Bayes information criterion [21], the likelihood ratio criterion [3] and the minimum discrimination information statistic [1]. We propose a general Markov chain Monte Carlo (MCMC) algorithm to numerically approximate the posterior distribution on the space of all patterns of mutual independence. We then demonstrate the interest of the method on synthetic data as well as two real datasets, showing the unique insight provided by this approach. The discussion sums it up and rises some outstanding issues.

## 2 METHOD

We start by proposing a quick description of the problem (Section 2.1) and its interpretation in a Bayesian framework in terms of model comparison (Section 2.2). We consider the special cases of multivariate normal distributions (Section 2.3) and discrete distributions (Section 2.4). In Section 2.5, we propose an efficient sampling scheme to explore the set of all partitions.

### 2.1 Mutual Independence and Partitions

Let  $X$  be a  $D$ -dimensional variate following a distribution  $g$  with parameter  $\theta$ . By definition, if  $X$  can be decomposed into  $K$  mutually independent subvectors  $X_1, \dots, X_K$ , then  $g(X|\theta)$  can be decomposed as the product

$$g(X|\theta) = \prod_{k=1}^K g_k(X_k|\theta). \quad (1)$$

The definition of a decomposition of  $X$  into mutually independent subvectors is equivalent to the choice of a partition  $\mathcal{B} = \{B_1, \dots, B_K\}$  of  $[D] = \{1, \dots, D\}$ , that is, disjoint and non-empty subsets (or blocks)  $B_k$ 's of  $[D]$  whose union is equal to  $[D]$ . A partition is denoted by concatenating the subsets composing it separated by the sign “|”, e.g., 12|3 for the partition of [3] into  $\{\{1, 2\}, \{3\}\}$ . Let  $\mathcal{C}_D$  be the set of all partitions of  $[D]$ . For instance,  $\mathcal{C}_2 = \{12, 1|2\}$  has 2 elements, while  $\mathcal{C}_3 = \{1|2|3, 12|3, 13|2, 23|1, 123\}$  has 5. More generally, the total number of partitions of a set with  $D$  elements is given by the  $D$ th Bell number  $\varpi_D$  [22], [23], [24], [25], [26].

### 2.1.1 A Note on Notations

Consider the partition  $\mathcal{B} = \{B_1, \dots, B_K\}$  of  $[D]$  into  $K$  blocks. Using the language of partitions, we should denote by  $X_A$  the  $\text{card}(A)$ -dimensional subvector of  $X$  defined by  $(X_a)_{a \in A}$ , and, if  $X_{B_1}, \dots, X_{B_K}$  are mutually independent, then  $g(X|\theta)$  should be decomposed as the product

$$g(X|\theta) = \prod_{k=1}^K g_{B_k}(X_{B_k}|\theta). \quad (2)$$

For the sake of simplicity, we will stick to  $g_k$  and  $X_k$  instead of  $g_{B_k}$  and  $X_{B_k}$ , respectively. Note however that these notations implicitly refer to a partition  $\mathcal{B} = \{B_1, \dots, B_K\}$ .

### 2.2 Model Comparison

Consider a partition  $\mathcal{B}$  that decomposes  $X$  into  $K$  mutually independent subvectors  $X_1, \dots, X_K$ ,  $X_k$  being of size  $D_k$ . Given a realization  $x$  of  $X$ , we quantify the relevance of  $\mathcal{B}$  by its posterior probability  $\Pr(\mathcal{B}|x)$ . According to Bayes' updating rule, this quantity yields

$$\Pr(\mathcal{B}|x) \propto \Pr(\mathcal{B}) p(x|\mathcal{B}). \quad (3)$$

$\Pr(\mathcal{B})$  is the prior probability for  $\mathcal{B}$ ; it characterizes what is known about  $\mathcal{B}$  before the data are available. In the following, it will be set as uniform on  $\mathcal{C}_D$ ; issues related to selecting the prior are discussed in Section 6.4.  $p(x|\mathcal{B})$  is the likelihood. It expresses how the data are related to the model. In the present case, it is also called marginal model likelihood, as it is obtained after removing the effect of the model parameters (see below). Finally,  $\Pr(\mathcal{B}|x)$  is the posterior probability of  $\mathcal{B}$ . It summarizes all the information that is available regarding  $\mathcal{B}$  after acquisition of the data.

#### 2.2.1 Marginal Model Likelihood

As the distribution  $g$  is assumed to be a function of some parameter  $\theta$ , the marginal model likelihood can be obtained by marginalization of the usual model likelihood, yielding in the assumption where  $\mathcal{B}$  holds

$$p(x|\mathcal{B}) = \int p(x, \theta|\mathcal{B}) d\theta \quad (4)$$

$$= \int p(x|\mathcal{B}, \theta) p(\theta|\mathcal{B}) d\theta. \quad (5)$$

In this expression,  $p(x|\mathcal{B}, \theta)$  is equal to

$$p(x|\mathcal{B}, \theta) = g(x|\theta) \quad (6)$$

$$= \prod_{k=1}^K g_k(x_k|\theta), \quad (7)$$

leading to

$$p(x|\mathcal{B}) = \int \prod_{k=1}^K g_k(x_k|\theta) p(\theta|\mathcal{B}) d\theta. \quad (8)$$

If  $\theta$  can itself be partitioned into subvectors,  $\theta = (\theta_k)_{k=1, \dots, K}$ , each subvector  $\theta_k$  characterizing the parameters of  $g_k$ , and assuming prior independence of these parameters

$$p(\theta|\mathcal{B}) = \prod_{k=1}^K p(\theta_k|\mathcal{B}), \tag{9}$$

the marginal likelihood reads

$$p(x|\mathcal{B}) = \prod_{k=1}^K \int g_k(x_k|\theta_k) p(\theta_k|\mathcal{B}) d\theta_k. \tag{10}$$

Often we are under the assumption that the data are composed of  $N$  independent and identically distributed (i.i.d.) realizations of  $X$ , that is,  $\{x_1, \dots, x_N\}$ . While this assumption does not simplify much the theoretical expression of the posterior distribution, it usually provides key simplifications in practical cases, as we will see in the next two particular cases of multivariate normal distributions and multivariate discrete distributions.

### 2.3 Multivariate Normal Distributions

We here consider the specific case where the data are multivariate normal. More specifically, let  $X$  be a  $D$ -dimensional multivariate normal variable with known mean  $\mu$  and unknown covariance matrix  $\Sigma$ . Under  $\mathcal{B}$ ,  $\Sigma$  is block-diagonal, each block  $\Sigma_k$  corresponding to a subset  $X_k$  of size  $D_k$ . Given a dataset  $y = \{y_1, \dots, y_N\}$  of i.i.d. realizations of  $X$  and  $S$  the corresponding sample sum-of-square matrix

$$S = \sum_{n=1}^N (y_n - \mu)(y_n - \mu)^t, \tag{11}$$

and with conjugate (i.e., inverse-Wishart) prior for the covariance matrix,  $p(S|\mathcal{B})$  can be calculated in closed form and yields [5]

$$p(S|\mathcal{B}) = \frac{|S|^{\frac{N-D-1}{2}}}{Z(D, N)} \prod_{k=1}^K \frac{Z(D_k, N + \nu_k)}{Z(D_k, \nu_k)} \frac{|\Lambda_k|^{\frac{\nu_k}{2}}}{|\Lambda_k + S_k|^{\frac{N+\nu_k}{2}}}, \tag{12}$$

where  $\Lambda$  is the prior (diagonal) scale matrix,  $\Lambda_k$  its  $k$ th block,  $|\cdot|$  the determinant of a matrix,  $\nu \geq D + 1$  the prior degree of freedom,  $\nu_k = \nu - D + D_k$ , and  $Z(d, n)$  the inverse of a normalization constant

$$Z(d, n) = 2^{\frac{nd}{2}} \pi^{\frac{d(d-1)}{4}} \prod_{d'=1}^d \Gamma\left(\frac{n+1-d'}{2}\right). \tag{13}$$

Incorporating the expression of  $p(S|\mathcal{B})$  from Eq. (12) into Bayes' rule, Eq. (3), and taking into account the fact that  $|S|^{\frac{N-D-1}{2}}/Z(D, N)$  does not depend on the model of dependence, and so is part of the normalization constant, we obtain for the posterior distribution

$$\Pr(\mathcal{B}|S) \propto \Pr(\mathcal{B}) \prod_{k=1}^K \frac{Z(D_k, N + \nu_k)}{Z(D_k, \nu_k)} \frac{|\Lambda_k|^{\frac{\nu_k}{2}}}{|\Lambda_k + S_k|^{\frac{N+\nu_k}{2}}}. \tag{14}$$

Note that if the mean is unknown, this calculation is still valid, with  $\mu$  replaced by the sample mean

$$m = \frac{1}{N} \sum_{n=1}^N y_n, \tag{15}$$

and the degree of freedom  $N$  by  $N - 1$ .

### 2.3.1 Asymptotic Form

Let  $\widehat{\Sigma}_k = S_k/N$  be the sample covariance matrix corresponding to block  $k$ . Asymptotically ( $N \rightarrow \infty$ ), the log of the marginal likelihood  $\ln p(S|\mathcal{B})$  can be expressed as [5]

$$-\frac{N}{2} \sum_{k=1}^K \ln |\widehat{\Sigma}_k| - \left[ \sum_{k=1}^K \frac{D_k(D_k + 1)}{4} \right] \ln N, \tag{16}$$

plus a term that is proportional to  $DN$  and, hence, does not depend on  $\mathcal{B}$ , and terms that are  $O(1)$ . In this equation, the first term corresponds to the part of the maximum-likelihood that does depend on  $\mathcal{B}$  (see Section 1.1 of online supplement, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2020.2968065>), while the second term is the BIC penalization for the dimension of the problem [21], i.e., of the form

$$-\frac{\#\text{model parameters}}{2} \ln N. \tag{17}$$

In the case where we compare a model  $\mathcal{B}_1$  of full dependence with another nested model  $\mathcal{B}_0$  with mutually independent subvectors  $X_1, \dots, X_K$ , the log-ratio of marginal model likelihoods yields

$$\begin{aligned} \ln \frac{p(S|\mathcal{B}_1)}{p(S|\mathcal{B}_0)} &= \frac{N}{2} \ln \frac{\prod_{k=1}^K |\widehat{\Sigma}_k|}{|\widehat{\Sigma}|} \\ &\quad - \frac{D(D+1)}{4} \ln N \\ &\quad + \left[ \sum_{k=1}^K \frac{D_k(D_k + 1)}{4} \right] \ln N + O(1). \end{aligned} \tag{18}$$

The first term of the right-hand side corresponds to the minimum discrimination information statistic against the null hypothesis of independence in the case of a multivariate normal distribution [1, Chap. 12, Section 3.6] or, equivalently, to the log of the likelihood ratio criterion for testing independence between sets of variates [3, Chap. 9]. This result can easily be generalized to any pair of nested models.

### 2.4 Cross-Classified Multinomial Distributions

We now consider the case of a  $D$ -dimensional discrete distribution, also coined cross-classified multinomial distribution [27, Section 7.1]. To this aim, let  $X = (X_1, \dots, X_D)$  be a  $D$ -dimensional discrete multivariate variable, such that each  $X_d$  takes values in set  $E_d$  with cardinality  $I_d$ . For each block  $B_k$ , we also define the set  $E_{B_k} = \times_{d \in B_k} E_d$  with cardinality  $I_{B_k} = \prod_{d \in B_k} I_d$ . Under  $\mathcal{B}$ ,  $X$  can be decomposed into  $K$  mutually independent variables  $(X_1, \dots, X_K)$ , the model is parameterized by  $K$  multidimensional parameters  $\theta_k = (\theta_{x_k})_{x_k \in E_{B_k}}$ , and the likelihood reads

$$\Pr(x|\theta_1, \dots, \theta_K) = \Pr(x_1, \dots, x_K|\theta_1, \dots, \theta_K) \tag{19}$$

$$= \prod_{k=1}^K \Pr(x_k|\theta_k) \tag{20}$$

$$= \prod_{k=1}^K \theta_{x_k}, \quad (21)$$

where  $\theta_{x_k}$  is the probability to have  $X_k = x_k$ . Given a dataset  $\mathbf{y}$  of  $N$  i.i.d. realizations of  $X$  and under the usual assumption of a Dirichlet prior with parameter  $(a_{x_k})_{x_k \in E_{B_k}}$  for each  $\theta_k$ , the marginal model likelihood has a simple expression (see Section 2.1 of online supplement, available online), leading to a posterior probability of

$$\Pr(\mathcal{B}|\mathbf{y}) \propto \Pr(\mathcal{B}) \prod_{k=1}^K \frac{\Gamma(\sum_{x_k \in E_{B_k}} a_{x_k})}{\prod_{x_k \in E_{B_k}} \Gamma(a_{x_k})} \times \frac{\prod_{x_k \in E_{B_k}} \Gamma(N_{x_k} + a_{x_k})}{\Gamma(\sum_{x_k \in E_{B_k}} N_{x_k} + a_{x_k})}, \quad (22)$$

where  $N_{x_k}$  is the number of times that we observe  $X_k = x_k$  for  $x_k \in E_{B_k}$ , and  $\Gamma$  the usual Gamma function.

#### 2.4.1 Asymptotic Form

For  $k = 1, \dots, K$ , let  $f_{x_k} = N_{x_k}/N$ , so that  $\sum_{x_k \in E_{B_k}} f_{x_k} = 1$ . Set also  $\mathbf{f}_k = (f_{x_k})_{x_k \in E_{B_k}}$ . Then the log posterior can be asymptotically expressed as (see Section 2.2 of online supplement, available online)

$$\ln \Pr(\mathcal{B}|\mathbf{y}) = \sum_{k=1}^K \left[ -NH(\mathbf{f}_k) - \frac{I_{B_k} - 1}{2} \ln N \right] + O(1), \quad (23)$$

where  $H(\mathbf{f}_k)$  is the classical entropy function associated with  $\mathbf{f}_k$ , that is

$$H(\mathbf{f}_k) = - \sum_{x_k \in E_{B_k}} f_{x_k} \ln(f_{x_k}). \quad (24)$$

The first term in the asymptotic expression of  $\ln \Pr(\mathcal{B}|\mathbf{y})$ , Eq. (23), corresponds to the maximum-likelihood estimate, while the second term is the BIC penalization for a model with  $I_{B_k} - 1$  parameters (corresponding to the  $I_{B_k}$ -dimensional parameter  $\theta_k$  together with the additional constraint  $\sum_{x_k \in E_{B_k}} \theta_{x_k} = 1$ ).

As we did for the multivariate normal distributions, consider the case where we compare a model  $\mathcal{B}_1$  of full dependence with another nested model  $\mathcal{B}_0$  with mutually independent subvectors  $X_1, \dots, X_K$ . The log-ratio of marginal model likelihoods then yields

$$\begin{aligned} \ln \frac{\Pr(\mathbf{y}|\mathcal{B}_1)}{\Pr(\mathbf{y}|\mathcal{B}_0)} &= N \sum_{x \in \times_{d \in [D]} E_d} f_x \ln \frac{f_x}{\prod_{k=1}^K f_{x_k}} \\ &\quad - \frac{\sum_{d \in [D]} I_d - 1}{2} \ln N \\ &\quad + \sum_{k=1}^K \left[ \frac{I_{B_k} - 1}{2} \ln N \right] \\ &\quad + O(1). \end{aligned} \quad (25)$$

The first term of the right-hand side generalizes the minimum discrimination information statistic against the null hypothesis of independence in the case of discrete two- and

three- way contingency tables [1, Chap. 8, Sections 2 and 3.1]. This result can easily be generalized to any pair of nested models.

## 2.5 Sampling Scheme

As mentioned above, the total number of partitions of a set with  $D$  elements is given by the  $D$ th Bell number  $\varpi_D$ . The first six Bell numbers are 1, 2, 5, 15, 52, and 203. The growth rate of  $\varpi_D$  is given by (see Section 3.1 of online supplement, available online)

$$\varpi_D = O\left[\left(\frac{D}{\ln D}\right)^D\right], \quad (26)$$

which is faster than exponential and slower than factorial. For instance, we have  $\varpi_{10} = 115\,975$ , while  $\varpi_{20}$  is larger than  $5.17 \times 10^{13}$ . As a consequence, exhaustive examination of all potential partitions quickly becomes intractable. To circumvent this issue, it is possible to resort to Markov chain Monte Carlo (MCMC) sampling, a powerful tool that is widely used in Monte Carlo integration but also in Bayesian data analysis to generate samples that, under certain conditions, will approximate a distributions of interest [28, Chap. 5]. In our case, we use it to generate a sample of partitions from our posterior distribution. Following [29, p. 322], we proceed according to the following sampling scheme:

- Generate  $M$  partitions from the uniform distribution [30].
- From these  $M$  partitions, use importance resampling to draw  $C$  partitions (that is, sample  $C$  partitions from the  $M$  partitions without replacement with a probability of sampling each partition proportional to its posterior probability, see [29, Section 10.5]).
- Use these  $C$  partitions as starting points to run  $C$  independent parallel chains of  $J$  samples.

For the sampling itself, [31] proposed a Gibbs sampling approach (henceforth coined Gibbs) that sequentially scans all the elements one by one and considers moves from the current partition to any other partition differing from the current one in only that one element. Such an algorithm has the advantage of considering only a limited number of potential partitions at each step. Each step scans through the  $D$  variables and, for each variable, there are as many options as there are partitions, which is limited by the number of variables; it is therefore  $O(D)$ , and the number of partitions considered is  $O(D^2)$ . However, this algorithm, by only moving one element at a time, is expected to generate highly correlated states and preclude large changes. To improve convergence, we also considered parallel tempering (PT) and a sampling scheme that can be conceptualized as an implementation of 2-way stochastic hierarchical clustering (2wSHC).

#### 2.5.1 Parallel Tempering

To allow for an exploration of the hypothesis space that is less likely to get trapped in (or around) local maxima, we can resort to parallel tempering [28, Section 10.4], see also [32] or [33]. Parallel tempering is a sampling scheme that runs several dependent sequences with different target probabilities and allows state swaps between sequences




Fig. 1. Example of 2w-SHC. Set  $D = 6$  and assume that the current state is partition  $12|356|4$ . From this partition, an agglomerative clustering algorithm would consider  $3 \times 2/2 = 3$  potential states ( $1|2|356|4$ ,  $1|2|435|6$ , and  $1|2|345|6$ ), and a divisive clustering algorithm  $\binom{2}{2} + \binom{3}{2} = 4$  potential states ( $1|2|356|4$  from the division of 12;  $12|3|4|56$ ,  $12|36|4|5$ , and  $12|35|4|6$  from the division of 356). In this particular case, 2w-SHC would then compare the posterior probabilities of 8 states.

with a certain probability. In our case, if the posterior probability for a model involving partition  $\mathcal{B}$  is denoted by  $\phi(\mathcal{B})$ , we set the target probabilities of the  $L$  sequences as

$$\pi_l(\mathcal{B}) = \exp\left[\frac{\ln \phi(\mathcal{B})}{T_l}\right], \quad l = 1, \dots, L. \quad (27)$$

We set  $T_1 = 1 < T_2 < \dots < T_L$ , so that  $\pi_1$  corresponds to the original distribution, while the  $\pi_l$ 's,  $l > 1$ , correspond to increasingly flattened versions of it. In that sense, parallel tempering is similar to simulated annealing, but with the advantage of respecting the detailed balance equation and, therefore, ensuring convergence towards the target distributions. At each step  $j \in \{1, \dots, J-1\}$ , the algorithm chooses between swapping (with probability  $\alpha_1$ ) and updating (with probability  $1 - \alpha_1$ ). Swapping is proposed between  $\mathcal{B}_l^{[j]}$ , the current state of sequence  $l$ , and  $\mathcal{B}_{l+1}^{[j]}$ , the current state of sequence  $l+1$ , uniformly on  $\{1, \dots, L-1\}$  and accepted with probability

$$\min\left[1, \frac{\pi_l(\mathcal{B}_{l+1}^{[j]})}{\pi_l(\mathcal{B}_l^{[j]})} \frac{\pi_{l+1}(\mathcal{B}_l^{[j]})}{\pi_{l+1}(\mathcal{B}_{l+1}^{[j]})}\right] \\ = \min\left(1, \exp\left\{\left[\ln \phi(\mathcal{B}_{l+1}^{[j]}) - \ln \phi(\mathcal{B}_l^{[j]})\right] \left(\frac{1}{T_l} - \frac{1}{T_{l+1}}\right)\right\}\right). \quad (28)$$

### 2.5.2 2-Way Stochastic Hierarchical Clustering

An alternative approach to [31] is to resort to a method that also relies on MCMC but, at each step, considers as potential new states the current partition as well as all partitions that are obtained by either the merging of two blocks of the current partition or the division of one block of the current partition into two blocks. For a partition of  $[D]$  into  $K$  blocks  $\{B_1, \dots, B_K\}$ , there are  $K(K-1)/2$  partitions that can be obtained by merging, and

$$\sum_{k: \#B_k \geq 2} \left\{ \begin{matrix} \#B_k \\ 2 \end{matrix} \right\} \quad (29)$$

that can be obtained by division, where  $\left\{ \begin{matrix} a \\ b \end{matrix} \right\}$  is the Stirling number of the second kind, i.e., the number of partitions of a set with  $a$  elements in  $b$  blocks. It can furthermore be shown that  $\left\{ \begin{matrix} a \\ 2 \end{matrix} \right\} = 2^{a-1} - 1$  (see Section 3.2 of online supplement, available online). Such an algorithm can be seen as the stochastic exploration of a hierarchy through a simultaneous combination of agglomerative (bottom-up) and divisive (top-down) hierarchical clustering, by considering moving up the current state (through merging) or down (through division)—whence the term “2-way”, see also

Fig. 1. From an algorithmic perspective, all merged partitions can be obtained in a straightforward manner, while divided partitions can be obtained using an algorithm proposed by [34], see also [26, Section 7.2.1.5]. This algorithm has the advantage of allowing for larger moves compared to [31]. However, it also has two downsides. First, the structure of the discrete space might still make it hard to escape local maxima, all the more that their probabilities can become very large with increasing  $D$  and  $N$ . Besides, the number of potential partitions considered at each step quickly increases with the number of variables and may considerably slow down the sampling scheme.

### 2.5.3 Our Approach

Practically, we implemented a sampler based on parallel tempering where swapping states is performed with probability  $\alpha_1$ , element-wise Gibbs sampling with probability  $\alpha_2$ , and sampling with the 2-way stochastic hierarchical clustering with probability  $1 - \alpha_1 - \alpha_2$ . See Fig. 2 for an algorithmic description of the sampling scheme. As specified in Table 1, this approach includes as special cases Gibbs ( $L=1$ ,  $\alpha_1=0$  and  $\alpha_2=1$ ), 2wSHC ( $L=1$ ,  $\alpha_1=0$  and  $\alpha_2=0$ ), Gibbs+2wSHC ( $L=1$ ,  $\alpha_1=0$  and  $\alpha_2 < 1$ ), Gibbs+PT ( $L > 1$ ,  $\alpha_1 > 0$  and  $\alpha_2 = 1 - \alpha_1$ ), 2wSHC+PT ( $L > 1$ ,

- 1: Start with initial step  $(\mathcal{B}_1^{[1]}, \dots, \mathcal{B}_L^{[1]})$
- 2: **for**  $j = 1$  to  $J-1$  **do** {Iterate on steps}
- 3:  $u \leftarrow \text{rand}()$
- 4: **if**  $u < \alpha_1$  **then** {Swapping step}
- 5: Uniformly choose  $l_0 \in \{1, \dots, L-1\}$
- 6:  $r \leftarrow \exp\left\{\left[\ln \phi(\mathcal{B}_{l_0+1}^{[j]}) - \ln \phi(\mathcal{B}_{l_0}^{[j]})\right] \left(\frac{1}{T_{l_0}} - \frac{1}{T_{l_0+1}}\right)\right\}$
- 7:  $v \leftarrow \text{rand}()$
- 8: **if**  $v < \min(1, r)$  **then** {Swap}
- 9:  $\mathcal{B}_{l_0}^{[j+1]} \leftarrow \mathcal{B}_{l_0+1}^{[j]}$
- 10:  $\mathcal{B}_{l_0+1}^{[j+1]} \leftarrow \mathcal{B}_{l_0}^{[j]}$
- 11:  $\mathcal{B}_l^{[j+1]} \leftarrow \mathcal{B}_l^{[j]}$  for  $l \notin \{l_0, l_0+1\}$
- 12: **else** {No change}
- 13:  $\mathcal{B}_l^{[j+1]} \leftarrow \mathcal{B}_l^{[j]}$  for all  $l \in \{1, \dots, L\}$
- 14: **end if**
- 15: **else if**  $u < \alpha_1 + \alpha_2$  **then** {Parallel step, Gibbs}
- 16:  $(\mathcal{B}_1^{[j+1]}, \dots, \mathcal{B}_L^{[j+1]}) \leftarrow$  updated state using Gibbs
- 17: **else** {Parallel step, 2wSHC}
- 18:  $(\mathcal{B}_1^{[j+1]}, \dots, \mathcal{B}_L^{[j+1]}) \leftarrow$  updated state using 2wSHC
- 19: **end if**
- 20: **end for**

Fig. 2. Sampling scheme. General description of the MCMC sampling scheme. Behavior of each of the  $C$  independent chains.

TABLE 1  
Sampling Scheme

	$L$	$\alpha_1$	$\alpha_2$
Gibbs	1	0	1
2wSHC	1	0	0
Gibbs+2wSHC	1	0	$< 1$
Gibbs+PT	$> 1$	$> 0$	$1 - \alpha_1$
2wSHC+PT	$> 1$	$> 0$	0
Gibbs+2wSHC+PT	$> 1$	$> 0$	$< 1 - \alpha_1$

Parameter values corresponding to specific sampling families.  $L$  is the number of dependent sequences run in the parallel tempering;  $\alpha_1$  is the probability to swap states;  $\alpha_2$  the probability to use a step of component-wise Gibbs sampling.

$\alpha_1 > 0$  and  $\alpha_2 = 0$ ), and Gibbs+2wSHC+PT ( $L > 1$ ,  $\alpha_1 > 0$  and  $\alpha_2 < 1 - \alpha_1$ ). Note that, due to the parallel tempering algorithm, each of the  $C$  independent parallel chain is itself composed of  $L$  dependent sequences.

### 3 SIMULATION STUDY

We now consider a simulation study with  $D = 6$  variables, corresponding to  $\varpi_6 = 203$  potential partitions. The low dimension of the problem allows for an exhaustive calculation of all posterior probabilities on the solution space.

#### 3.1 Data

For  $D = 6$ , we considered partitions with an increasing number of blocks  $K$  ( $1 \leq K \leq 6$ ). For a given value of  $K$ , we performed 500 simulations. For each simulation, the 6 variables were randomly partitioned into  $K$  clusters, all partitions having equal probability of occurrence [35, Chap. 12], [36]. For a given partition  $\mathcal{B} = \{B_1, \dots, B_K\}$  of [6], we generated 300 i.i.d. samples following a distribution  $g(x)$  structured as in Eq. (1).

In a first batch of simulations (“Gaussian data”), each  $g_k$  was set to either a univariate (if the size  $D_k$  of  $B_k$  was equal to 1) or multivariate (if  $D_k > 1$ ) normal distribution with mean 0 and covariance matrix  $\Sigma_k$  sampled according to an inverse-Wishart distribution with  $D_k + 1$  degrees of freedom and scale matrix the identity matrix and then rescaled to a correlation matrix. Such a sampling scheme on  $\Sigma_k$  generated correlation matrices with uniform marginal distributions for all correlation coefficients [37].

For the second batch of simulations (“non-Gaussian data”), we kept all  $\Sigma_k$ ’s generated in the first batch but set each  $g_k$  to a Student- $t$  distribution with  $\zeta$  degrees of freedom, location parameter 0 and scale matrix  $\Sigma_k$  [38].  $\zeta$  was set in  $\{1, 3, 5\}$  (the normal case would correspond to  $\zeta \rightarrow \infty$ ).

#### 3.2 Analysis

We have to deal with multivariate normal distributions, and the corresponding resolution method involves two hyperparameters, namely a degree of freedom  $\nu$  and a scale matrix  $\Lambda$  (see Section 2.3). We here considered three alternative approaches to set these hyperparameters. The first approach, `BayesOptim`, sets the degree of freedom to the lowest value that still corresponds to a well-defined distribution, that is  $\nu = D$ , and a diagonal scale matrix that optimizes the marginal model likelihood of Eq. (12) in the case of 6 mutually independent variables [5]. An alternative approach, `BayesCorr`,

works with the sample correlation matrix instead of the sample covariance matrix. One can then set the number of degrees of freedom to  $\nu = D + 1$  and the scale matrix to the identity matrix. As mentioned in the previous paragraph, the corresponding prior distribution yields uniform marginal distributions for the correlation coefficients [37]. A third approach, `Bic`, computes the posterior probability using the BIC approximation, which does not involve hyperparameters and is also insensitive to the fact that the input is the covariance matrix or the correlation matrix.

To assess the quality of the inference process, we considered three quantities: the posterior probability of the true underlying partition, the ratio between this posterior probability and the probability of the maximum a posteriori (MAP) partition, and the entropy of the posterior distribution in  $\log_{203}$ . The posterior probability of the true underlying partition gives an absolute sense of the quality of the inference process as a consequence of both the information contained in the data and the quality of the method used. The two other quantities help to disentangle the relative contribution of the two factors (information contained in the data and method used). The entropy of the posterior distribution in  $\log_{203}$  yields values ranging from 0 for a degenerate posterior distribution (one partition has a posterior probability of 1, while all other partitions have a posterior probability of 0) to 1 for a uniform posterior distribution (all partitions have posterior probability of  $1/203$ ). It is an indicator of the (lack) of information contained in the data. As to the ratio between the posterior probability of the true model and the probability of the MAP partition, it gives a sense of how far away the inference process is from picking the true model as the preferred model.

### 3.3 Results

#### 3.3.1 Gaussian Data

We first compared the probability distributions obtained for the three methods applied to the Gaussian data (see Fig. 3 or Section 4.1 of online supplement, available online). There was a relative correspondence between values, with a relationship that varied depending on the number of blocks ( $K$ ) in the synthetic data.

The posterior probability of the true model increased with the number of samples and decreased with the number of clusters (Fig. 4, top left panel). Globally, the inference process tended to detect the true model as one of the most probable ones (Fig. 4, top right and bottom right panels). The difference lay in the entropy of the posterior distributions (Fig. 4, bottom left panel). For  $K = 1$  cluster, it quickly tended toward 0, indicating a very localized distribution and, hence, a precise inference. By contrast, increasing  $K$  led to both an increase in entropy for a given sample size, but also a decrease in the speed at which entropy decreased with the sample size.

All in all, we found that the number of clusters in the simulated data had a dramatic influence on the inference process, in that it was all the harder to confirm the existence of a specific model that the given model had many blocks of mutually independent variables. Such a behavior in disfavor of models with sparse covariance matrices is further discussed in Section 6.6.

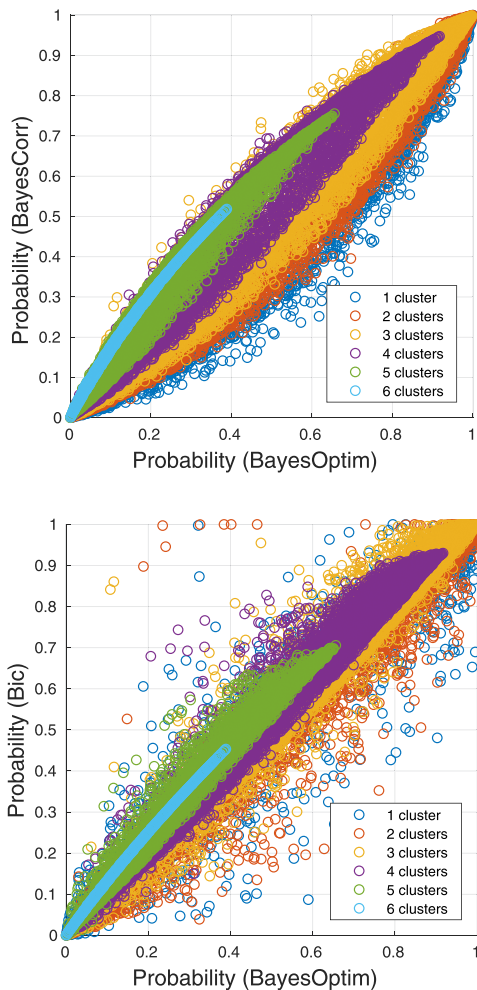


Fig. 3. Simulation study. Comparison of probability obtained for BayesOptim and either BayesCorr (top) or Bic (bottom) as a function of the number of clusters  $K$  in the simulated data.

### 3.3.2 Non-Gaussian Data

Analysis of the non-Gaussian data with the same methods led to several changes (see Section 4.2 of online supplement, available online). First, the relationship between probabilities computed with BayesOptim and BayesCorr in the non-Gaussian case remained similar to the Gaussian case regardless of the degrees of freedom. By contrast, Bic behaved quite differently from BayesOptim, with a difference that decreased when the degrees of freedom increased to become quite similar to the Gaussian case for  $\zeta = 5$ .

Besides, the medians of the four quantities used to assess the behavior of our approach were found to be quite similar to the Gaussian case. However, they also exhibited much more variability than in the Gaussian case, with again a variability that decreased with increasing  $\zeta$  to become quite similar to the Gaussian case for  $\zeta = 5$ .

## 4 HIV STUDY DATA

In this section, we consider a toy example of mutual independence extraction. The problem was already analyzed elsewhere [5], [39], [40]. As in the case of the simulation study, its low dimension ( $D = 6$  variables and  $\varpi_6 = 203$  potential partitions) allows for an exhaustive calculation of all posterior probabilities on the solution space.

## 4.1 Data

The data originates from a study investigating early diagnosis of HIV infection in children from HIV positive mothers [39]. The variables are related to various measures on blood and its components:  $X_1$  and  $X_2$  immunoglobulin G and A, respectively;  $X_4$  the platelet count;  $X_3$ ,  $X_5$  lymphocyte B and T4, respectively; and  $X_6$  the T4/T8 lymphocyte ratio. The observed correlation matrix is given in Table 2. According to [39], discussion with experts suggested the existence of a strong association between variables  $X_1$  and  $X_2$  as well as between variables  $X_3$ ,  $X_5$ , and  $X_6$ . Using conditional independence graphs, [39] found that the values of partial correlation between  $X_4$  and other variables had probability around zero, which led him to hypothesize that his original model was over-parameterized. Still with conditional independence graphs, [40] found that the links between  $X_4$  and the five other variables had low probability of existence and that no single graphical model was able to accurately account for the data, hinting that models of conditional independence graphs might be too refined for that specific dataset. [5] found that various hierarchical clustering methods tended to cluster  $X_3$  and  $X_5$  as well as  $X_1$  and  $X_2$ ; variable  $X_6$  tended to be cluster with  $(X_1, X_2)$  or  $(X_3, X_5)$  depending on the clustering method.

## 4.2 Analysis

We first computed the exact probability distribution of all potential partitions using the three variants BayesOptim, BayesCorr, and Bic mentioned in the simulation study. We then focused on BayesOptim. Regarding the quantities of interest, we first computed the probabilities for all potential partitions. From there, we computed the relevances associated to all subsets of  $[D]$  [41]. For a subset  $B$  of  $[D]$ , the relevance of  $B$  is the probability to find  $B$  as a block in the partitioning of  $[D]$ ; it is calculated as the sum of all probabilities associated with partitions for which  $B$  is a block. Finally, we computed the probability corresponding to the following two expert statements:

- “There is a strong association between  $X_1$  and  $X_2$ ”: posterior probability for  $X_1$  and  $X_2$  to be partitioned in the same block regardless of the rest.
- “There is a strong association between  $X_3$ ,  $X_5$ , and  $X_6$ ”: posterior probability for  $X_3$ ,  $X_5$ , and  $X_6$  to be partitioned in the same block regardless of the rest.

## 4.3 Results

The three methods (BayesOptim, BayesCorr, and Bic) yielded similar results. The four most probable patterns of mutual independence were found to be the same in the same order (12356|4, 12|356|4, 126|35|4, and 124|356) with a good agreement as to the weight of these models (see Table 3). These four models accounted for more than 99 percent of the probability distribution. Importantly, these four models are *not* all nested in one another (e.g., 12356|4 and 124|356; 12|356|4 and 126|35|4). The fact that some of them were found to be nested in one another (12|356|4 nested in 124|356 and 12356|4; 126|35|4 nested in 12356|4) was extracted by the analysis and not imposed a priori.

Relevances were computed for all subsets of  $[D]$  in the case of BayesOptim (see Section 5 of online supplement,

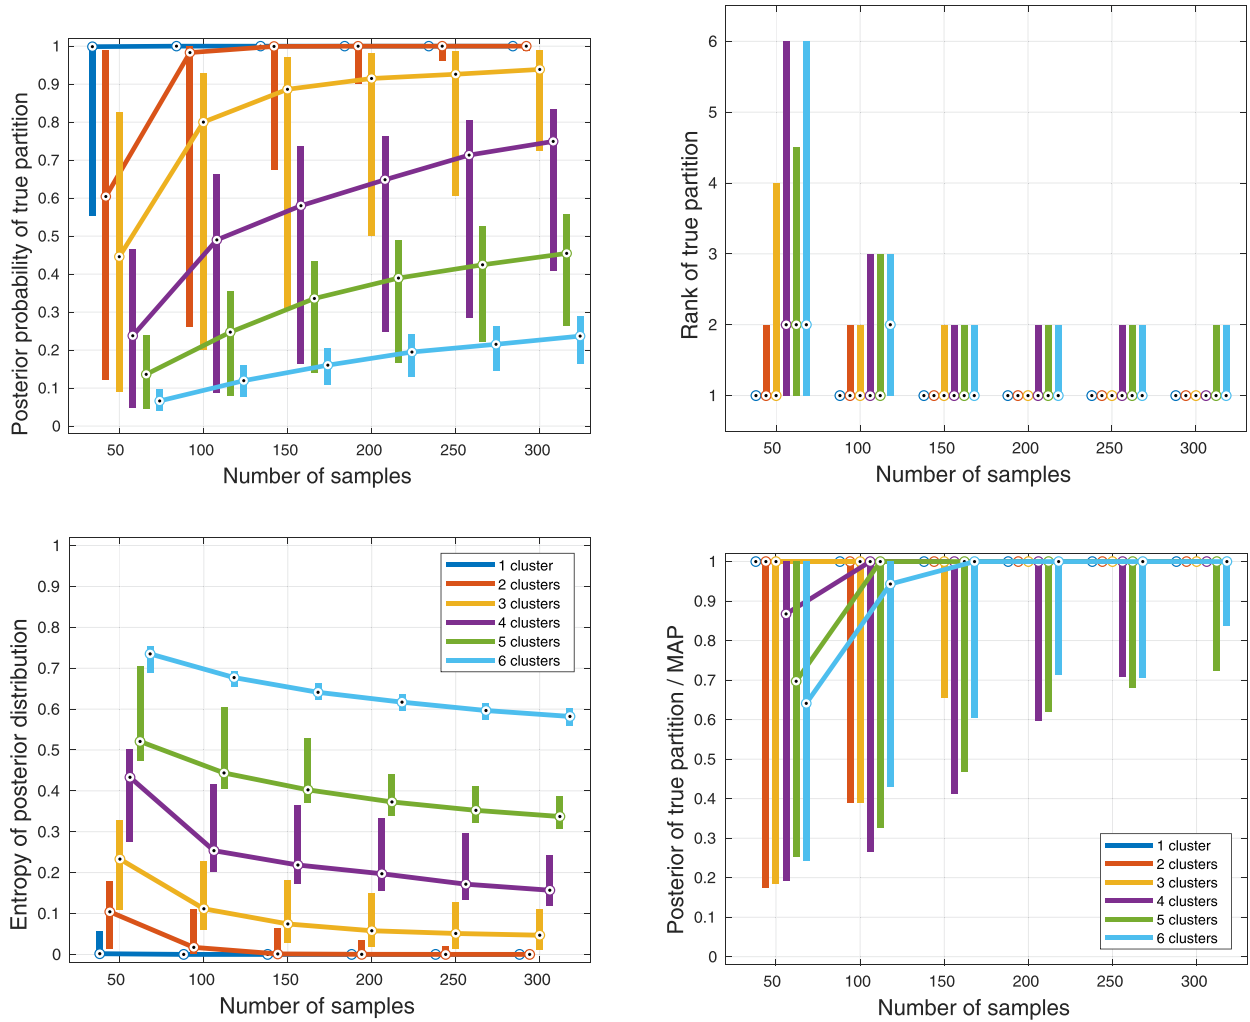


Fig. 4. Simulation study. For BayesOptim, boxplot (median and [25%, 75%] probability interval) of posterior probability for the true model (top left), entropy of posterior distribution (bottom left), rank of true model when ranking potential models by decreasing posterior probability (top right), and ratio of posterior probability of true model to posterior probability of maximum a posteriori (bottom right).

available online). The most important features were that; 4 had a relevance of 0.994 and 12|356 a relevance of 0.852. This means that the probability to find 4 in a block all by itself was equal to 0.994, while the probability to find block 12|356 was equal to 0.852. The fact that the relevance of 4 was found to be larger than that of 12|356 means that in some cases we found a partition with a block containing 4 alone but where 12|356 was split, e.g., 12|356|4 and 126|35|4.

TABLE 2  
HIV Study Data

	$X_1$	$X_2$	$X_3$	$X_4$	$X_5$	$X_6$
$X_1$	<b>8.8374</b>	0.479	-0.040	-0.033	0.356	-0.236
$X_2$	0.483	<b>0.1919</b>	0.068	-0.084	-0.224	-0.110
$X_3$	0.220	0.057	<b>8924231.9</b>	0.085	0.552	-0.330
$X_4$	-0.040	-0.133	0.149	<b>20392.4</b>	0.091	0.013
$X_5$	0.253	-0.124	0.523	0.179	<b>1952795.2</b>	0.384
$X_6$	-0.276	-0.314	-0.183	0.064	0.213	<b>1.378</b>

Summary statistics for the HIV data. Sample variances (main diagonal, bold), correlations (lower triangle) and partial correlations (upper triangle, italic). Data from [39].

TABLE 3  
HIV Study Data

Rank	Model	Posterior probability		
		BayesOptim	BayesCorr	Bic
# 1	12356 4	0.852	0.648	0.912
# 2	12 356 4	0.132	0.320	$7.90 \times 10^{-2}$
# 3	126 35 4	$8.21 \times 10^{-3}$	$1.94 \times 10^{-2}$	$4.51 \times 10^{-3}$
# 4	124 356	$3.80 \times 10^{-3}$	$4.77 \times 10^{-3}$	$2.00 \times 10^{-3}$
Total		0.996	0.992	0.998

Comparison of variants. Posterior probabilities for the four most probable patterns of mutual independence as computed using BayesOptim, BayesCorr, and Bic.



## 5 fMRI FUNCTIONAL CONNECTIVITY ANALYSIS

### 5.1 Materials and Methods

#### 5.1.1 Data

Functional magnetic resonance imaging (fMRI) is an imaging modality that makes it possible to dynamically and non-invasively follow metabolic and hemodynamic consequences of brain activity. fMRI functional connectivity analysis is a field that investigates the organization of brain networks in resting-state fMRI data by extracting clusters of brain regions whose spontaneous activities are highly correlated [42], [43], [44]. We applied our approach to a real resting-state fMRI data composed of 205 time samples recorded for 82 brain regions in a young healthy subject (subject #1 of [5]). Since both assumptions of variables following a multivariate normal distribution and of i.i.d. realizations are quite common in the field, we resorted to the results of Section 2.3; to speed up the process, we used the BIC approximation of Eq. (16). We ran the analyses with Matlab on a desktop PC with Intel Xeon Silver 4114, 2.2 GHz, 10 cores and 64 GB RAM.

#### 5.1.2 Numerical Sampling

A system with  $D = 82$  variables can be associated with  $\varpi_{82} \approx 6.2439 \times 10^{89}$  different partitions, preventing explicit computation of all the corresponding posterior probabilities. We therefore resorted to the approximate sampling scheme discussed in Section 2.5. More specifically, we ran 5 distinct sampling schemes: Gibbs, 2wSHC, Gibbs+2wSHC (with  $\alpha_2 = 0.8$ ), Gibbs+PT (with  $L = 7$  and  $\alpha_1 = 0.5$ ), 2wSHC+PT (with  $L = 7$  and  $\alpha_1 = 0.5$ ), and Gibbs+2wSHC+PT (with  $L = 7$ ,  $\alpha_1 = 0.5$  and  $\alpha_2 = 0.4$ ). All sampling schemes were run with  $M = 10^4$  initial samples according to a uniform distribution,  $C = 4$  parallel chains, and  $J_1 = 10^5$  samples for each chain. To account for burn-in, we discarded the first half of the samples and computed our quantities of interest on the second half.

We observed that 2wSHC tended to visit a limited number of states, but each visited state was associated with many proposal states. To speed up the calculation, we kept a database of states associated with more than  $10^3$  proposal states that were visited in the last 200 steps. At each iteration, the procedure checked if the current state was in the database; if so, it used the values already computed instead of computing them over. This significantly sped up the process, but also increased the memory load.

To assess algorithmic variability, all sampling schemes were run 10 times, for a total of 50 runs.

#### 5.1.3 Comparing Algorithms

We assessed convergence of each of the 50 runs by quantifying between-chain heterogeneity with the average  $L_1$  distance between estimated probabilities. More specifically, assume that  $Q$  partitions appeared during the sampling regardless of the chain. Let  $f_{cq}$  be the frequency of the  $q$ th partition in chain  $c$ ,  $f_q$  the frequency of the  $q$ th partition when all chains are pooled,  $\mathbf{f}_c = (f_{c1}, \dots, f_{cQ})^t$  the vector of frequency estimates for chain  $c$ , and  $\mathbf{f} = (f_{.1}, \dots, f_{.Q})^t$  the vector of frequency estimates from the whole sample. Between-chain heterogeneity was measured as the average




Fig. 5. Real data. Comparison of total computation time (top) and time per step (bottom) for Gibbs, 2wSHC, Gibbs+2wSHC, Gibbs+PT, and Gibbs+2wSHC+PT. Computational time was obtained by running  $J_1 = 10^5$  samples for Gibbs, 2wSHC, Gibbs+2wSHC, and Gibbs+PT, and  $J_2 = 10^4$  samples for Gibbs+2wSHC+PT.

of the  $L_1$  distances between the frequencies observed in each chain and the frequencies observed in the whole sample, that is

$$\frac{1}{C} \sum_{c=1}^C \|\mathbf{f}_c - \mathbf{f}\|_{L_1} = \frac{1}{C} \sum_{c=1}^C \sum_{q=1}^Q |f_{cq} - f_q|. \quad (30)$$

This quantity is greater than 0, and equal to 0 only when all chains yield the same estimates. To assess within- and between-algorithm variability, we also computed the  $L_1$  distance between probability estimates obtained from the 50 runs.

## 5.2 Results

All algorithms were quite computationally demanding. Gibbs, Gibbs+2wSHC, and Gibbs+PT ran smoothly for all 10 repetitions with the above parameters. 2wSHC failed twice out of ten repetitions. Gibbs+2wSHC+PT failed 10 times out of 10 with the original parameters; with a reduced chain length of  $J_2 = 10^4$ , it failed once out of 10. As to 2wSHC+PT, we were not able to run it, even for a chain length as low as  $J_3 = 10^3$ . All failures were due to insufficient memory to store the database of previous states visited and corresponding probabilities (see Section 5.1.2), either because the current state was related to too many potential states through 2wSHC, or because the global database was too large.

Computational times are summarized in Fig. 5. Gibbs was the fastest algorithm. Allowing for steps of 2wSHC added an extra burden. Running parallel tempering had two opposite effects. First, 7 parallel steps were run for each chain, adding computational burden in terms of time and memory load. However, about half of the time, a step only consisted of testing if two states could be swapped, which is faster than computing several probabilities. The two effects made that Gibbs+PT had the fastest steps and Gibbs+2wSHC+PT the longest steps.

The level of convergence varied quite widely between algorithms, as summarized in Fig. 6. Gibbs and 2wSHC performed quite badly, as all chains of a given algorithm converged to different states (characterized by a heterogeneity of 3/2). Convergence of Gibbs was improved by adding

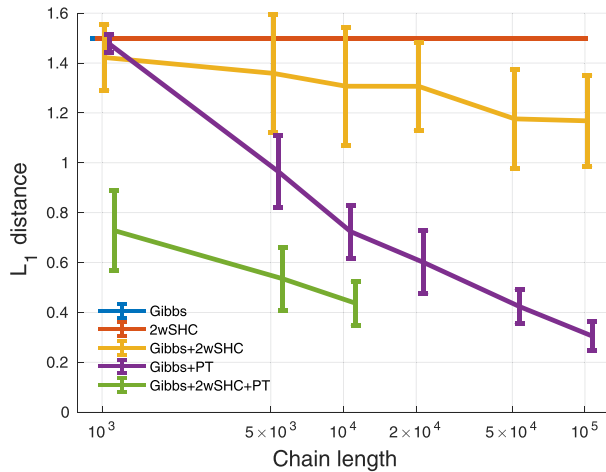


Fig. 6. Real data. Comparison of convergence (mean  $\pm$  standard deviation of  $L_1$  distance) as a function of chain length. When the four chains of a given algorithm converged to four different states, the corresponding heterogeneity was  $3/2$ . The curves are slightly shifted along the  $x$ -axis to avoid superposition. Gibbs and 2wSHC obtained very similar results and cannot be distinguished on the plot.

steps of 2wSHC and parallel tempering. After  $J_2 = 10^4$  samples, Gibbs+2wSHC+PT was the algorithm that provided chains with the best convergence level. However, taking time into account, Gibbs+PT was able to provide a better convergence level with  $J_1 = 10^5$  in less time.

These differences in convergence were confirmed when considering the probabilities estimated from the 50 runs (see Fig. 7 and Table 4). The 10 repetitions of Gibbs and 2wSHC essentially converged to different partitions, leading to distance values close to 2. The estimates given by Gibbs+2wSHC were less dissimilar to one another. Gibbs+PT and Gibbs+2wSHC+PT generated estimates that were consistent across repetitions, and relatively similar between algorithms.

## 6 DISCUSSION

In the present manuscript, we first exposed the problem of extracting patterns of mutual independence, translated it in terms of comparing partitions, and proposed a theoretical treatment in the form of Bayesian model comparison. We investigated two particular cases: multivariate normal distributions and cross-classified multinomial distributions, showing that the Bayesian solutions relate to the likelihood ratio criterion, the minimum discrimination information criterion and the Bayes information criterion in the asymptotic case of large sample size. We proposed a general sampling scheme on the set of all partitions which combines Gibbs sampling, a stochastic 2-way exploration of a hierarchy, and

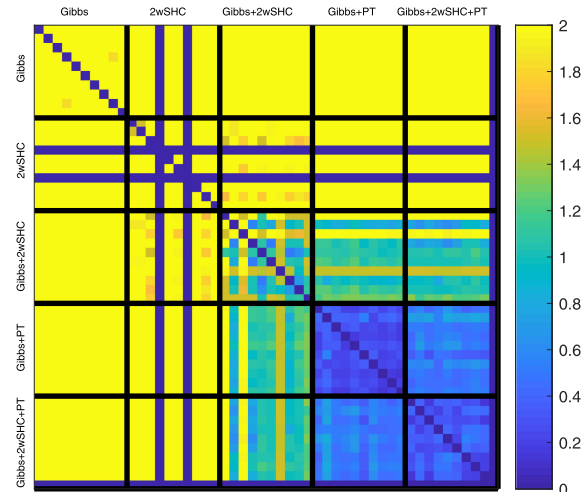


Fig. 7. Real data. Comparison of  $L_1$  distances between estimated probabilities. Two sets of estimated probabilities with different supports have a distance of 2.

parallel tempering. We finally demonstrated the interest of the method on synthetic data as well as two real datasets, showing the unique insights provided by this approach.

From a theoretical perspective, we believe that the present approach opens a new avenue for the investigation of mutual independence. Compared to existing tools (likelihood ratio criterion, minimum discrimination information statistic), the main advantage of our method is that it does not require to restate the problem in the form of two competing nested models to be compared in the asymptotic case of large sample size, but allows for a full data-driven exploration of the patterns of mutual independence. It also allows for a variety of questions to be formulated regarding any feature induced by the potential pattern of mutual independence, such as the number of blocks, the specific relationship between two variables, or the relationship of one variable to all other variables.

### 6.1 Mutual Independence Extraction and Clustering

As mentioned in the introduction, the problem faced in the present study belongs to the general area of clustering, but does not fall into either of two of the main subareas that are the clustering of mixture models and the clustering of functional data and curves. For instance, a classical way to perform cluster analysis of the HIV data would be to consider the children as objects and the six variables as features. One could then want to cluster the 107 objects (children) into subgroups of objects that share similar profiles in terms of the 6 features. In this context, a common assumption is that

TABLE 4  
Real Data

	Gibbs	2wSHC	Gibbs+2wSHC	Gibbs+PT	Gibbs+2wSHC+PT
Gibbs	$1.996 \pm 0.024$	$2 \pm 0$	$2 \pm 0$	$2 \pm 0$	$2 \pm 0$
2wSHC		$1.982 \pm 0.094$	$1.948 \pm 0.112$	$2 \pm 0$	$2 \pm 0$
Gibbs+2wSHC			$1.323 \pm 0.436$	$1.300 \pm 0.387$	$1.239 \pm 0.429$
Gibbs+PT				$0.275 \pm 0.086$	$0.521 \pm 0.107$
Gibbs+2wSHC+PT					$0.335 \pm 0.103$

Comparison of  $L_1$  distances between estimated probabilities from different runs grouped by algorithm (average  $\pm$  standard deviation).

(i) each observation in the sample may arise from any of a small number of different distributions (one per cluster), (ii) the objects to be classified (children) are independent from one another given the cluster they belong to, and (iii) the distribution corresponding to each cluster is characterized by a set of parameters whose dimension is fixed (and, in particular, does not depend on the number of elements in the corresponding class). This is not the issue that the current study tried to address. Instead, it considered the 6 variables as objects, each object having 107 features (the realizations over children), and tried to classify them according to their *relationships* with one another. In this respect, our problem shares similarities with the clustering of functional data or curves, e.g., [15], [16], [17], [18], [19], [20]. However, while curve clustering usually assumes a certain form of temporal dependence, we here assumed independent and identically distributed realizations of a variable. This is the reason why the clustering of variables based on their pattern of independence belongs to neither class mentioned above and defines a class of its own.

Another possibility would be to consider biclustering [45]. Biclustering has a natural translation in terms of partitions and Bayesian model comparison, but these models usually relate the values taken by the variables across examples. We are therefore not quite sure how it would apply to models of mutual independence.

## 6.2 Bayesian Analysis, Asymptotic Approximation and Model Selection Criteria

In two common cases—multivariate normal distributions (Section 2.3) and cross-classified multinomial distributions (Section 2.4)—we showed that the log of the posterior distribution could be asymptotically approximated by a criterion that is the sum of a maximum-likelihood ratio and a BIC correction for model complexity. In the simulation study (Section 3), we showed that the method based on the asymptotic approximation (coined *Bic*) gave results that were similar to the posterior distribution (see in particular Fig. 3, bottom). Similar results were found in the HIV study data (Section 4; see in particular Table 3). Retrospectively, one could have thought of the BIC criterion as a valid approach to perform blind extraction of mutual independence patterns. Yet, we are not aware of any reference advocating such an approach. It is only after the full Bayesian analysis we conducted in a general case, its application to multivariate normal distributions, and an investigation of the asymptotic behavior that the BIC appeared as a potential solution to the problem. As a matter of fact, we believe that considering extraction of mutual independence patterns as one of model comparison is a major contribution of the present manuscript.

## 6.3 Sampling Scheme

From a practical point of view, the sampling scheme is a key factor for a full exploration of all potential patterns of mutual independence. In the present study, we combined a Gibbs sampling scheme that sequentially scans elements (*Gibbs*), a stochastic 2-way algorithm for hierarchical clustering (*2wSHC*), and parallel tempering (*PT*). We observed that an algorithm solely based on *Gibbs* or *2wSHC*

generated chains that were slow to converge and had very limited mixing. By contrast, mixing *Gibbs* and *2wSHC* as well as introduction of *PT* allowed for a better exploration of the state space (quantified in terms of between-chain heterogeneity).

Our algorithm could be modified by completing the initial (uniform) random sampling in the set of all partitions with a step consisting of running the agglomerative hierarchical clustering on the data [5] and then basing the importance resampling step on the visited states. This potential improvement was not carried through in the present study for the following reason. We used mixing of the chains as a measure of convergence. The validity of such a monitoring could be challenged in the case of chains starting from seeds that are close to one another from the algorithm's perspective. Indeed, by construction, the various states of the hierarchical clustering are rather close to one another from the perspective of *2wSHC* (at most  $D$  steps). As a consequence, using seeds from the agglomerative hierarchical clustering might have artificially accelerated convergence for *2wSHC* and biased our convergence analysis in favor of *2wSHC*. By contrast, we hoped that feeding the chains with seeds sampled uniformly would make between-chain similarity a good indicator of convergence.

Besides the sampling algorithm proposed here, another approach might be to use the generalized Swendsen–Wang sampler (*SWC*) or generalized Gibbs sampler [46], even though such approaches are expected to work best on relatively sparse connectivity graphs—by contrast, in our case, the graph would be fully connected.

In the cases that we considered, the problem was simplified by the fact that all parameters could be integrated out and the marginal likelihood computed explicitly. When this is not possible, one has to deal with a joint posterior distribution of the partition and corresponding parameters, whose dimensions vary depending on the partition. To perform numerical sampling of such a distribution, one could consider resorting to reversible jump MCMC (*RJ-MCMC*) and, more precisely, split-and-merge [47], [48], [49], which would have to be adapted to deal with the specificity of the problem at hand.

## 6.4 Prior Distributions

Setting prior distributions is a key and touchy issue that one usually has to deal with when performing a Bayesian analysis. Here, we faced the problem at two different levels: for the competing models of mutual independence,  $p(\mathcal{B})$ , and for the model parameters, e.g.,  $p(\Sigma|\mathcal{B})$  for the multivariate normal distribution and  $p(\theta|\mathcal{B})$  for the cross-classified multinomial distribution.

In the present manuscript, the priors for the parameters were set as conjugate priors for the sake of convenience. While this choice is not fully satisfying, it has the advantage of allowing for closed form expressions for the marginal model likelihoods. The choice of a prior for  $\Sigma$  is further discussed in [5]. More generally, consistency requires that parameters found in different models be assigned the same prior. One way to do this is to define the parameter's prior distribution for the model with no mutual independence and then derive all other distributions through marginalization (see again discussion of

[5]). In any case, the importance of the prior distributions decreases when the data size increases. In the limiting case of large sample size, the asymptotic expressions do not depend on the priors set for the model parameters, as we found for multivariate normal distributions and cross-classified multinomial distributions.

Regarding the prior probabilities for the competing models of mutual independence, we set them to a uniform prior. The consequences of setting such a uniform prior on the set of all partitions need to be investigated, for the structure of this set and its size might induce undesired properties. For instance, in the simulation study and the HIV study data, setting a uniform prior entailed that the probability to have a partition with 1, 2, 3, 4, 5, or 6 blocks was given by  $4.93 \times 10^{-3}$ , 0.153, 0.443, 0.3202,  $7.39 \times 10^{-2}$ , and  $4.93 \times 10^{-3}$ , respectively. With increasing  $D$ , the difference in probability increases dramatically. For  $D = 100$ , the prior probability for a partition to have a number of blocks in the range [21, 40] is given by  $1 - 1.12 \times 10^{-4}$ , see also [26].

Now, depending on the information available for a given problem, one might wish to set different priors. Various models have been proposed for partitions. A general family of priors is the so-called product partition models [31], [41]. Also, specific features might be desirable for the prior distribution. For instance, in the case where the assignment of the labels  $[D]$  to the  $D$  variables is arbitrary, it would make sense to require the prior distribution to be exchangeable [50, Section 2]—see also Section 3.3 of online supplement, available online. By contrast, other features might be rejected, such as consistency as defined in [13]—see also Section 3.4 of online supplement, available online.

Finally, it might be of interest to incorporate expert knowledge into the analysis. In a Bayesian framework, expert statements can easily be incorporated in the form of prior information, that is, by selecting a prior distribution over potential partition models that respects the prior knowledge. Information regarding the number of blocks as well as preferred or forbidden partitions may be relatively easily to model. In other cases (e.g., preferential connectivity patterns in the case of brain networks observed in fMRI, Section 5), translating specific information into prior probability might turn out to be a real challenge.

## 6.5 Dimension of the Problem

The previous discussion on priors shows that the complexity of the model, and, hence, the ease with which it can be solved, is influenced by two factors: the set of all potential patterns of mutual independence and the model parameters corresponding to each pattern of mutual independence. For  $D$  variables, the set of all potential partitions of a given set of variables into mutually independent components is of size  $\varpi_D$ , which is a function of  $D$  only. As to the dimension of the model parameter space, it depends on the underlying model of mutual independence. For instance, in the case of multivariate normal distributions, a covariance matrix is fully specified by  $D(D+1)/2$  parameters; of these, one could argue that only  $D(D-1)/2$  (correlation coefficients) have a decisive influence on the pattern of mutual independence. Patterns of mutual independence, by setting some correlation coefficients to 0, reduce this number. In the case

of a cross-classified multinomial distribution, the maximal number of parameters is given by  $\prod_{d=1}^D I_d - 1$ . For instance, for binary variables ( $I_d = 2$ ), we have a maximum of  $2^D - 1$  parameters, which is  $O(e^D)$ , but grows slower than  $\varpi_D$  according to Eq. (26). For a given data size, the number of model parameters has an influence on how well these parameters can be estimated and, as a consequence, our capacity to discriminate between models.

## 6.6 The Difficulty to Extract Sparse Models

In the simulation study (Section 3), we observed that the number of clusters underlying the data had a dramatic influence on the inference process, in that it was all the harder to confirm the existence of a specific model that the given model had many groups of mutually independent variables. An explanation for such a behavior, which can be related to the common issue of overfitting in statistics and machine learning, can be seen in the BIC approximation. Consider for instance the multivariate normal case (Section 2.3). The Bayes factor between a “reference” partition  $\mathcal{B}_1$  and another model  $\mathcal{B}_0$  associated with a sparser covariance matrix (henceforth shortened as “sparser model”), Eq. (18), can be expressed as

$$\ln \frac{p(S|\mathcal{B}_1)}{p(S|\mathcal{B}_0)} \approx \frac{N}{2} \hat{I}(\mathcal{B}_1 : \mathcal{B}_0) - \frac{\#\text{param}(\mathcal{B}_1) - \#\text{param}(\mathcal{B}_0)}{2} \ln N. \quad (31)$$

$\hat{I}(\mathcal{B}_1 : \mathcal{B}_0)$  is positive and tends to its theoretical counterpart  $\delta$ , which is a multivariate generalization of mutual information [51], [52]. When the data originate from  $\mathcal{B}_1$ ,  $\delta > 0$  and the Bayes factor is positive, tending to  $+\infty$  as  $O(N)$ , a clear indication that  $\mathcal{B}_1$  is to be preferred over  $\mathcal{B}_0$ . By contrast, when the data originate from  $\mathcal{B}_0$ ,  $\hat{I}(\mathcal{B}_1 : \mathcal{B}_0)$  will tend to  $\delta = 0$ . In the most favorable case, where  $\hat{I}(\mathcal{B}_1 : \mathcal{B}_0)$  is indeed equal to 0 (which is actually of probability 0 on real data), the Bayes factor will be negative, but going to  $-\infty$  at the much slower rate of  $O(\ln N)$ . We suspect that this behavior might not be specific to the simulation performed but more general, making extraction of sparse models particularly challenging.

## 6.7 Mutually Exclusive and Exhaustive Models

In a Bayesian analysis, a key strategy when comparing competing hypotheses is the use of mutually exclusive and exhaustive propositions or models [53, Section 2.4]. This is the case for our competing patterns of mutual independence. In particular, the set of all partitions with the finer-than relationship forming a lattice [54, p. 15], we have to explicitly state that selection of one model precludes the selection of another, nested model to ensure mutual exclusion. This should be contrasted to the classical definition of mutual independence where such exclusion is not mentioned. For instance, for the simulation study and the HIV study data analyzed in the present manuscript with  $D = 6$ , assuming a pattern of mutual independence of 12356|4 means that the distribution for  $(X_1, X_2, X_3, X_4, X_5, X_6)$  can be decomposed as

$$g(X_1, X_2, X_3, X_4, X_5, X_6) = g(X_1, X_2, X_3, X_5, X_6) g(X_4), \quad (32)$$

but does not specify anything regarding  $g(X_1, X_2, X_3, X_5, X_6)$ . In particular, we could further have

$$g(X_1, X_2, X_3, X_5, X_6) = g(X_1, X_2)g(X_3, X_5, X_6), \quad (33)$$

which would entail a pattern of mutual independence of  $12|356|4$ . This standard definition of mutual independence is in line with the prior distributions set on the model parameters, which theoretically do not prevent such situations to occur. In the example above, the prior probability distribution for all correlation coefficients between  $(X_1, X_2)$  and  $(X_3, X_5, X_6)$  to be equal to 0 is not equal to 0. To achieve mutual exclusion of models, we should exclude for each prior distribution parameter values that are compatible with a more refined model. Such an approach is rarely carried out. Instead, many analyses using Bayesian model comparison (e.g., polynomial approximations of increasing degree) rely on parameter spaces that are embedded in one another. A theoretical justification of this approach can be found in [55]. Practically, the consequence of this issue is rather limited. In the above example, the probability to obtain  $12|356|4$  with a full correlation matrix on  $(X_1, X_2, X_3, X_5, X_6)$  is equal to 0. Therefore, a model with a full correlation matrix exhibits almost surely no pattern of mutual independence.

## 6.8 Covariance Matrix versus Full Dataset

In the special case of multivariate normal distributions (Section 2.3), we performed our inference based solely on the sample covariance matrix. The rationale for this is that, from a theoretical point of view, mutual independence is a property of the covariance matrix, which has to be block diagonal. As a consequence, it is often thought that the sole statistic of interest is the sample covariance matrix (see for instance the HIV study, Section 4), and we wanted our approach to be applicable to such cases. Note however that, to do so, we used a formula that is only valid when the number of samples is larger than the dimension of the problem,  $N \geq D$ . When this is not the case, we need to use the whole data set  $y$  and perform Bayesian inference on  $p(\mathcal{B}|y)$  instead of  $p(\mathcal{B}|S)$  as was done here. This implies, for instance, using non-degenerate prior distributions for  $\mu$  and  $\Sigma$ , e.g., conjugate priors [29, Section 3.6]. The corresponding analysis was performed in Section 1.2 of online supplement, available online, yielding a posterior probability that tends to the one found in Section 2.3 when a hyperparameter tends to 0.

## 6.9 Mutual Independence Extraction and Graphical Models

A point to discuss is the relationship between the extraction of patterns of independence and graphical models. Graphical models are graphical representations of variables that conveniently encode relations of dependence and independence between variables. There are for instance conditional independence graphs [27] and covariance graphs [56]. Conditional independence graphs are usually considered in the context of multivariate normal distributions (graphical Gaussian models) or discrete distributions (graphical log-linear models), while covariance graphs are mostly used in the context of multivariate normal distributions. Conditional independence graphs are more refined than models

of mutual independence. For instance, for  $D$  variables, there are  $\varpi_D$  potential models of mutual independence but  $2^{D(D-1)/2}$  models of conditional independence graphs, with  $\ln \varpi_D = O(D \ln D)$  and  $\ln 2^{D(D-1)/2} = O(D^2)$ . In conditional independence graphs, mutual independence has a natural translation: mutually independent clusters of variables are disjoint maximal connected components of the graph. But conditional independence graphs can code for much more than connected components. As such, conditional independence graphs are much more flexible models. One could therefore contemplate using tools from the field of graphical models to estimate the patterns of mutual independence. However, we do not expect this direction to provide efficient solutions. First, the problem of inferring the pattern of a graphical model is itself challenging, generally requiring specific assumptions and numerical approximations. For instance, in the simplest case of multivariate normal distributions, the structure of conditional independence is strongly related to that of the precision (or concentration) matrix, that is, the inverse of the covariance matrix. Such a matrix is particularly complex to estimate, as it involves a matrix inversion. The maximum-likelihood is complex to obtain under constraint of a given conditional independence graph (e.g., using the iterative proportional fitting algorithm); nondecomposable graphical models are particularly hard to deal with. Then, once performed, such an analysis provides information regarding the underlying pattern of conditional independence, a large part of which is irrelevant for mutual independence and would have to be discarded by marginalization. We therefore expect mutual independence extraction to be more robust than graphical model procedures. Also, one could think of cases (e.g., multivariate Student-t distributions) where models of mutual independence could be compared, whereas graphical models would be more problematic to define. Finally, from a practical point of view, [5] showed that some such tools (spectral clustering and graphical lasso) performed worse than Bayesian model-based hierarchical clustering. Also, the HIV example was analyzed using conditional independence graphs in [39] and [40]. The results are rather complex to interpret in terms of mutual independence. [39] found that the values of partial correlation between  $X_4$  and other variables had probability around zero, which led him to hypothesize that his original model was over-parameterized. [40] found that links between  $X_4$  and other variables had low probability of existence. The (approximate) multivariate analysis found that the data could only poorly be accounted for by a single graph, as no graph had a probability over  $6.60 \times 10^{-2}$ . Among the 8 most probable graphs, ranging in posterior probability from  $6.60 \times 10^{-2}$  to  $2.00 \times 10^{-2}$ , the first and third most probable graphs (with posterior probabilities of  $6.60 \times 10^{-2}$  and  $3.20 \times 10^{-2}$ , respectively, for a total of  $9.80 \times 10^{-2}$ ) found  $X_4$  isolated from the other 5 variables, while the 6 others found  $X_4$  connected (with posterior probabilities of  $4.73 \times 10^{-2}$ ,  $2.93 \times 10^{-2}$ ,  $2.73 \times 10^{-2}$ ,  $2.66 \times 10^{-2}$ ,  $2.20 \times 10^{-2}$ ,  $2.00 \times 10^{-2}$ , and a total probability of 0.173). These 8 models accounted for only 27 percent of the total probability, and 73 percent remained to be explained. This is a typical example where the level of refinement carried by conditional independence graphs is counter-productive when one is only looking for the simpler information of mutual independence.

## 6.10 Dealing With Non-Gaussian Data

To deal with continuous data, we assumed an underlying model of multivariate normal distributions. Such a model has the advantage of simplicity, as it makes it possible to derive a closed form expression for the marginal model likelihood (Section 2.3). In the simulation study (Section 3), we investigated how this assumption fared when dealing with data generated according to independent multivariate Student- $t$  distributions. Multivariate Student- $t$  distributions have heavier tails than multivariate normal distributions and provide a general yet simple and flexible parametric framework to assess the robustness of our approach to non-normal data. Unlike what happens for multivariate normal distributions, the product of multivariate Student- $t$  distributions is in general not a multivariate Student- $t$  distribution, and a block-diagonal scale matrix in a multivariate Student- $t$  distribution is not a sign of mutually independent variables. Therefore, a product of multivariate Student- $t$  distributions has no reason to be a multivariate Student- $t$  distribution itself. We found that our approach fared decently with such data, with two limits: an increased variability, and a BIC approximation diverging from the other two methods BayesOptim and BayesCorr. How this result can be extended to variables of larger dimension remains to be investigated. In this perspective, a particular question that is not clear to us is the relationship between the dimensions of the independent variables, the degrees of freedom of their respective distributions, and how non-Gaussian the resulting data are.

Going beyond the multivariate normal model and explicitly dealing with multivariate Student- $t$  data would definitely broaden the scope of our approach and allow for a more robust inference of mutual independence patterns. Obtaining the closed form expression for the marginal model likelihood in this case is likely to be a real challenge. Still, one could maybe consider using the BIC approximation of the correct model, with a maximum likelihood that would be computed numerically and a correction for the number of parameters that would be straightforward to obtain. Whether such a method has any practical interest has yet to be tested.

## 7 CONCLUSION AND FUTURE WORK

In the present manuscript, we advocated that the problem of extracting patterns of mutual independence could efficiently be considered as a Bayesian model comparison. For multivariate normal distributions and cross-classified multinomial distributions, we showed that the Bayesian solution provides a principled yet efficient and flexible generalization of existing approaches. We proposed a general sampling scheme to perform a blind exploration of the posterior distribution on the set of all partitions. Finally, we demonstrated the interest of the method, showing the unique insights provided by this approach.

### 7.1 Effective Selection of Sparse Models

In the simulation study, we found that sparse models tended to require large datasets to be successfully extracted. This result first need to be confirmed in other situations. If this finding happens to be general, more research would be

warranted to find ways to counter this trend. Directions of research may include incorporating prior information favoring sparse models.

### 7.2 Further Investigation of the Relationship Between Our Method and Existing Approaches

In the manuscript, we emphasized the relationship between the Bayes factor and minimum discrimination information in the case of multivariate normal distributions, Eq. (16), and cross-classified multinomial distributions, Eq. (25). This result was also used to understand the difficulty to extract simple models that we observed on simulated data (see Sections 3 and 6.6). We believe that this relationship goes much deeper and deserves an in-depth investigation. We hope to be able to propose a first step in that direction in the near future.

### 7.3 Efficient Representation of Results

A key issue in the investigation of patterns of mutual independence is our (in)ability to represent the results of probabilistic inference in a synthetic manner. Indeed, while a pattern of mutual independence between  $D$  variables can be summarized as a partition of  $[D]$  (as was done in the present manuscript), we are not aware of any method to represent a probability distribution over the set of all patterns of mutual independence. For small  $D$  (practically,  $D \leq 4$ ), one could think of representing the lattice structure of all such patterns together with a color coding for the probability of each pattern. But the growth rate of  $\varpi_D$  prevents to use this method for even moderate values of  $D$ . For subsets of cardinality 1 or 2, relevances could be represented, e.g., in the form of graphs. For statistics of order 3, we would need a 3d space, and statistics of higher order could not be represented. We believe that being able to efficiently represent patterns of mutual independence would greatly help in the investigation of mutual independence.

## ACKNOWLEDGMENTS

The authors would like to thank the three anonymous reviewers, whose comments helped significantly improve the quality of this article.

## REFERENCES

- [1] S. Kullback, *Information Theory and Statistics*. Mineola, NY, USA: Dover, 1968.
- [2] J. H. Zar, *Biostatistical Analysis*, 5th ed. Upper Saddle River, NJ, USA: Prentice Hall, 2010.
- [3] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*. New York, NY, USA: Wiley, 1958.
- [4] D. R. Wolf, "Mutual information as a Bayesian measure of independence," 1994. [Online]. Available: <https://arxiv.org/abs/comp-gas/9511002>
- [5] G. Marrelec, A. Messé, and P. Bellec, "A Bayesian alternative to mutual information for the hierarchical clustering of dependent random variables," *PLoS One*, vol. 10, no. 9, 2015, Art. no. e0137278.
- [6] J. D. Banfield and A. E. Raftery, "Model-based Gaussian and non-Gaussian clustering," *Biometrics*, vol. 49, no. 3, pp. 803–821, 1993.
- [7] C. Fraley and A. E. Raftery, "Model-based clustering, discriminant analysis, and density estimation," *J. Amer. Statist. Assoc.*, vol. 97, no. 458, pp. 611–631, 2002.
- [8] J. W. Lau and P. J. Green, "Bayesian model-based clustering procedures," *J. Comput. Graphical Statist.*, vol. 16, no. 3, pp. 526–558, 2007.

- [9] A. K. Jain, "Data clustering: 50 years beyond K-means," *Pattern Recognit. Lett.*, vol. 31, no. 8, pp. 651–666, 2010.
- [10] C. E. Rasmussen, "The infinite Gaussian mixture model," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2000, pp. 554–560.
- [11] K. A. Heller and Z. Ghahramani, "Bayesian hierarchical clustering," Gatsby Unit Tech. Rep. GCNU-TR 2005-002, 2005. Online. [Available]: <http://mlg.eng.cam.ac.uk/zoubin/clustering.html>
- [12] C. Fraley and A. E. Raftery, "Bayesian regularization for normal mixture estimation and model-based clustering," *J. Classification*, vol. 24, pp. 155–181, 2007.
- [13] J. G. Booth, G. Casella, and J. P. Hobert, "Clustering using objective functions and stochastic search," *J. Roy. Statist. Soc. Ser. Statist. Methodol.*, vol. 70, pp. 119–139, 2008.
- [14] K. Sirinukunwattana, R. S. Savage, M. F. Bari, D. R. J. Snead, and N. M. Rajpoot, "Bayesian hierarchical clustering for studying cancer gene expression data with unknown statistics," *PLoS One*, vol. 8, no. 10, 2013, Art. no. e75748.
- [15] J. C. Wakefield, C. Zhou, and S. G. Self, "Modelling gene expression data over time: Curve clustering with informative prior distributions," *Bayesian Anal.*, vol. 7, pp. 721–732, 2003.
- [16] J. O. Ramsay and B. W. Silverman, *Functional Data Analysis*, 2nd ed. New York, NY, USA: Springer, 2005.
- [17] N. Serban and L. Wasserman, "CATS: Clustering after transformation and smoothing," *J. Amer. Statist. Assoc.*, vol. 100, no. 471, pp. 990–999, 2005.
- [18] N. A. Heard, C. C. Holmes, and D. A. Stephens, "A quantitative study of gene regulation involved in the immune response of anopheline mosquitoes: An application of Bayesian hierarchical clustering of curves," *J. Amer. Statist. Assoc.*, vol. 101, pp. 18–29, 2006.
- [19] L. Ferreira and D. B. Hitchcock, "A comparison of hierarchical methods for clustering functional data," *Commun. Statist. – Simul. Comput.*, vol. 38, no. 9, pp. 1935–1949, 2009.
- [20] J. Jacques and C. Preda, "Functional data clustering: A survey," Inria, Rocquencourt, France, Tech. Rep. 8198, 2013.
- [21] G. Schwarz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, 1978.
- [22] E. T. Bell, "Exponential numbers," *Amer. Math. Monthly*, vol. 41, no. 7, pp. 411–419, 1934.
- [23] E. T. Bell, "Exponential polynomials," *Ann. Math.*, vol. 35, no. 2, pp. 258–277, 1934.
- [24] E. T. Bell, "The iterated exponential integers," *Ann. Math.*, vol. 39, no. 3, pp. 539–557, 1938.
- [25] G.-C. Rota, "The number of partitions of a set," *Amer. Math. Monthly*, vol. 71, pp. 498–504, 1964.
- [26] D. E. Knuth, *The Art of Computer Programming*, vol. 4, Fascicle 3B. Boston, MA, USA: Addison-Wesley, 2005.
- [27] J. Whittaker, *Graphical Models in Applied Multivariate Statistics*. Hoboken, NJ, USA: Wiley, 1990.
- [28] J. Liu, *Monte Carlo Strategies in Scientific Computing*. New York, NY, USA: Springer, 2002.
- [29] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*. London, U.K.: Chapman & Hall, 1998.
- [30] J. Pitman, "Some probabilistic aspects of set partitions," *Amer. Math. Monthly*, vol. 104, pp. 201–209, 1997.
- [31] E. M. Crowley, "Product partition models for normal means," *J. Amer. Statist. Assoc.*, vol. 92, no. 437, pp. 192–198, 1997.
- [32] R. M. Neal, "Sampling from multimodal distributions using tempered transitions," *Statist. Comput.*, vol. 6, no. 4, pp. 353–366, 1996.
- [33] D. J. Earl and M. W. Deem, "Parallel tempering: Theory, applications, and new perspectives," *Phys. Chem. Chem. Phys.*, vol. 7, 2005, Art. no. 3910.
- [34] F. Ruskey, "Simple combinatorial gray codes constructed by reversing sublists," in *Proc. 4th Int. Symp. Algorithms Comput.*, 1993, pp. 201–208.
- [35] A. Nijenhuis and H. Wilf, *Combinatorial Algorithms for Computers and Calculators*, 2nd ed. Orlando, FL, USA: Academic, 1978.
- [36] H. S. Wilf, "East side, west side," 1999. [Online]. Available: <http://www.math.upenn.edu/wilf/lecnotes.html>
- [37] J. Barnard, R. McCulloch, and X.-L. Meng, "Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage," *Statistica Sinica*, vol. 10, no. 4, pp. 1281–1311, 2000.
- [38] S. Kotz and S. Nadarajah, *Multivariate t Distributions and Their Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [39] A. Roverato, "Asymptotic prior to posterior analysis for graphical gaussian models," in *Classification and Data Analysis*, M. Vichi and O. Opitz, Eds. Berlin, Germany: Springer, 1999, pp. 335–342.
- [40] G. Marrelec and H. Benali, "Asymptotic Bayesian structure learning using graph supports for Gaussian graphical models," *J. Multivariate Anal.*, vol. 97, pp. 1451–1466, 2006.
- [41] J. A. Hartigan, "Partition models," *Commun. Statist. Theory Methods*, vol. 19, pp. 2745–2756, 1990.
- [42] B. T. T. Yeo *et al.*, "The organization of the human cerebral cortex estimated by intrinsic functional connectivity," *J. Neurophysiol.*, vol. 106, pp. 1125–1165, 2011.
- [43] C. Kelly *et al.*, "A convergent functional architecture of the insula emerges across imaging modalities," *NeuroImage*, vol. 61, no. 4, pp. 1129–1142, 2012.
- [44] O. Sporns, "Cerebral cartography and connectomics," *Philosoph. Trans. Roy. Soc. London Ser. Biol. Sci.*, vol. 370, no. 1668, 2015, Art. no. 20140173.
- [45] J. A. Hartigan, "Direct clustering of a data matrix," *J. Amer. Statist. Assoc.*, vol. 67, no. 337, pp. 123–129, 1972.
- [46] A. Barbu and S.-C. Zhu, "Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1239–1253, Aug. 2005.
- [47] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.
- [48] S. Richardson and P. J. Green, "On Bayesian analysis of mixtures with unknown number of components," *J. Roy. Statist. Soc. Ser. Statist. Methodol.*, vol. 59, pp. 731–792, 1997.
- [49] D. I. Hastie and P. J. Green, "Model choice using reversible jump Markov chain Monte Carlo," *Statistica Neerlandica*, vol. 66, no. 3, pp. 309–338, 2012.
- [50] J. Pitman, "Combinatorial stochastic processes," Dept. Statist., Univ. California, CA, USA. Tech. Rep. 621, 2002. [Online]. Available: <https://statistics.berkeley.edu/tech-reports?page=6>
- [51] H. Joe, "Relative entropy measures of multivariate dependence," *J. Amer. Statist. Assoc.*, vol. 84, pp. 157–164, 1989.
- [52] M. Studeny, "Complexity of structural models," in *Proc. Joint Session 6th Prague Conf. Asymptotic Statist. 13th Prague Conf. Inf. Theory Statist. Decis. Functions Random Processes*, 1998, pp. 23–28.
- [53] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge, U.K.: Cambridge Univ. Press, 2003.
- [54] G. Birkhoff, *Lattice Theory*, 3rd ed. Providence, RI, USA: American Mathematical Society, 1973.
- [55] J. O. Berger and M. Delampady, "Testing precise hypotheses," *Statist. Sci.*, vol. 2, no. 3, pp. 317–335, 1987.
- [56] D. Edwards, *Introduction to Graphical Modelling*, 2nd ed. New York, NY, USA: Springer, 2000.

**Guillaume Marrelec** received the MS degree in engineering and applied mathematics jointly from the École Centrale Paris, Gif-sur-Yvette, France and the Universität Stuttgart, Stuttgart, Germany, in 1999, and the PhD degree in medical imaging from the Université Paris XI, Orsay, France, in 2003. He is currently a senior research scientist with the Institut National de la Santé et de la Recherche Médicale (INSERM) and works with the Laboratoire d'imagerie Biomédicale (LIB). His research interests include brain imaging, functional magnetic resonance imaging, functional connectivity, information theory, Bayesian analysis, and interactions.

**Alain Giron** received the PhD degree in biomedical engineering from the Université de Technologie de Compiègne (UTC), Compiègne, France, in 1990. He is currently a senior research engineer with the Institut National de la Santé et de la Recherche Médicale (INSERM) and works with the Laboratoire d'imagerie Biomédicale (LIB). His research interests include medical imaging, statistics, and artificial intelligence.

► For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/csdl](http://www.computer.org/csdl).