

RESEARCH ARTICLE

A Bayesian Alternative to Mutual Information for the Hierarchical Clustering of Dependent Random Variables

Guillaume Marrelec^{1*}, Arnaud Messé², Pierre Bellec³

1 Sorbonne Universités, UPMC Univ Paris 06, CNRS, INSERM, Laboratoire d'imagerie biomédicale (LIB), F-75013, Paris, France, **2** Department of Computational Neuroscience, University Medical Center Hamburg-Eppendorf, Hamburg University, Hamburg, Germany, **3** Département d'informatique et recherche opérationnelle, Centre de recherche de l'institut universitaire de gériatrie de Montréal, Université de Montréal, Montréal, Qc, Canada

* marrelec@lib.upmc.fr



OPEN ACCESS

Citation: Marrelec G, Messé A, Bellec P (2015) A Bayesian Alternative to Mutual Information for the Hierarchical Clustering of Dependent Random Variables. PLoS ONE 10(9): e0137278. doi:10.1371/journal.pone.0137278

Editor: Magnus Rattray, University of Manchester, UNITED KINGDOM

Received: January 19, 2015

Accepted: August 16, 2015

Published: September 25, 2015

Copyright: © 2015 Marrelec et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data have been deposited to Figshare: <http://dx.doi.org/10.6084/m9.figshare.1521155>.

Funding: A.M. is supported by Deutsche Forschungsgemeinschaft (DFG) grant SFB 936/Z3. P.B. is supported by the Fonds de Recherche du Québec (Santé (FRQS)). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: PB offers consulting in brain image analysis for two Contract Research Organizations, NeuroRX and Biospective, Montreal,

Abstract

The use of mutual information as a similarity measure in agglomerative hierarchical clustering (AHC) raises an important issue: some correction needs to be applied for the dimensionality of variables. In this work, we formulate the decision of merging dependent multivariate normal variables in an AHC procedure as a Bayesian model comparison. We found that the Bayesian formulation naturally shrinks the empirical covariance matrix towards a matrix set a priori (e.g., the identity), provides an automated stopping rule, and corrects for dimensionality using a term that scales up the measure as a function of the dimensionality of the variables. Also, the resulting log Bayes factor is asymptotically proportional to the plug-in estimate of mutual information, with an additive correction for dimensionality in agreement with the Bayesian information criterion. We investigated the behavior of these Bayesian alternatives (in exact and asymptotic forms) to mutual information on simulated and real data. An encouraging result was first derived on simulations: the hierarchical clustering based on the log Bayes factor outperformed off-the-shelf clustering techniques as well as raw and normalized mutual information in terms of classification accuracy. On a toy example, we found that the Bayesian approaches led to results that were similar to those of mutual information clustering techniques, with the advantage of an automated thresholding. On real functional magnetic resonance imaging (fMRI) datasets measuring brain activity, it identified clusters consistent with the established outcome of standard procedures. On this application, normalized mutual information had a highly atypical behavior, in the sense that it systematically favored very large clusters. These initial experiments suggest that the proposed Bayesian alternatives to mutual information are a useful new tool for hierarchical clustering.

Canada. This does not alter the authors' adherence to PLOS ONE policies on sharing data and materials.

Introduction

Cluster analysis aims at uncovering natural groups of objects in a multivariate dataset (see [1] for a review). In the vast variety of methods used in cluster analysis, an agglomerative hierarchical clustering (AHC) is a generic procedure that sequentially merges pairs of clusters that are most similar according to an arbitrary function called similarity measure, thereby generating a nested set of partitions, also called hierarchy [2]. The choice of the similarity measure indirectly defines the shape of the clusters, and thus plays a critical role in the clustering process. While this choice is guided by the features of the problem at hand, it is also often restricted to a limited number of commonly used measures, such as the Euclidean distance or Pearson correlation coefficient [3]. In the present work, we focus on the clustering of random variables based on their mutual information, which has recently gained in popularity in cluster analysis, notably in the field of genomics [4–7] and in functional magnetic resonance imaging (fMRI) data analysis [8–10]. Mutual information is a general measure of statistical dependency derived from information theory [11–13]. A key feature of mutual information is its ability to capture nonlinear interactions for any type of random variables [14]; also of interest, it indifferently applies to univariate or multivariate variables and can thus be applied to clusters of arbitrary size. Yet, mutual information is an extensive measure that increases with variable dimensionality. In addition, \hat{I} , the finite-sample estimator of mutual information, suffers from a dimensionality-dependent bias (see SA of S1 File). Several authors have proposed to correct mutual information for dimensionality by using a “normalized” version of mutual information [15–17]. In the clustering literature, normalized mutual information is routinely used. However, the impact of such correction procedure has not been extensively evaluated so far.

In the present paper, we consider Bayesian model-based clustering [1, 18–20] as an alternative to mutual information for the hierarchical clustering of dependent multivariate normal variables. Specifically, we derive a similarity measure by comparing two models: \mathcal{M}_I where X_i and X_j are independent (i.e., the covariance between any element of X_i and any element of X_j is equal to zero), against \mathcal{M}_D where the covariance between X_i and X_j can be set to any admissible value. The proposed similarity measure is then the log Bayes factor in favor of \mathcal{M}_D against \mathcal{M}_I [21]. With appropriate priors on the model parameters, we show that the similarity measure $s(X_i, X_j)$ between X_i and X_j can be expressed in closed form. As will be shown below, the Bayesian formulation naturally (1) allows for clustering even when the sample covariance matrix is ill-defined; (2) provides for an automated stopping rule when the clustering reached has $s(X_i, X_j) \leq 0$ for any pair of remaining clusters; (3) corrects for dimensionality using a term that scales up the measure as a function of the dimensionality of the variables; and (4) provides for a local and global measure of similarity, in that it can be used to decide which pair of variables to cluster at each step (local level) as well as to compare different levels of the resulting hierarchy (global level). Asymptotically (i.e., when the number of samples $N \rightarrow \infty$), the similarity measure is a linear function of mutual information, with a penalization factor that is in agreement with the Bayesian information criterion (BIC) [22]. In this sense, the present paper makes an explicit connection between Bayesian model comparison for the clustering of dependent random variables and mutual information. The code corresponding to the Bayesian approach is freely available online (<https://github.com/SIMEXP/arXiv-1501.05194/releases/tag/1.0>)

We evaluated an AHC procedure based on this approach with synthetic datasets. The experiment aimed to evaluate how it behaved under both its exact and asymptotic forms compared to other approaches, including raw and normalized mutual information. We finally tested the new measures on two real datasets: a toy dataset and functional magnetic resonance imaging (fMRI) data.

Analysis

In the following, we develop a Bayesian solution to the problem of clustering detailed above. We first introduce the model together with the Bayesian framework and a general expression for the similarity measure. In subsequent subsections, we derive a closed-form expression for the marginal model likelihoods under both assumptions of dependence and independence as well as exact and asymptotic expressions for the similarity measure. We then provide a description of the hierarchical agglomerative clustering algorithm resulting from the present development. We examine how the same framework can be conveniently used to compare nested partitions, that is, different levels of a hierarchy. We also deal with the issue of setting the hyperparameters. Finally, we show how the Bayesian solution can naturally provide for an automatic stopping rule.

Bayesian model comparison

Let \mathbf{X} be a D -dimensional multivariate normal variable with known mean $\boldsymbol{\mu}$ and unknown covariance matrix $\boldsymbol{\Sigma}$. Define \mathbf{X}_i and \mathbf{X}_j as two disjoint subvectors of \mathbf{X} (of size D_i and D_j , respectively), and $\mathbf{X}_{i \cup j}$ as their union (of size $D_{i \cup j} = D_i + D_j$). Assume that we have to decide whether we should cluster \mathbf{X}_i and \mathbf{X}_j based on their level of dependence. To this end, consider two competing models \mathcal{M}_I and \mathcal{M}_D . In \mathcal{M}_I , \mathbf{X}_i and \mathbf{X}_j are independent and the distribution of $\mathbf{X}_{i \cup j}$ can be decomposed as the product of the marginal distributions of \mathbf{X}_i and \mathbf{X}_j . Under such condition, $\boldsymbol{\Sigma}_{i \cup j}$, the restriction of $\boldsymbol{\Sigma}$ to $\mathbf{X}_{i \cup j}$, is block diagonal with blocks $\boldsymbol{\Sigma}_i$ and $\boldsymbol{\Sigma}_j$, the restrictions of $\boldsymbol{\Sigma}$ to \mathbf{X}_i and \mathbf{X}_j , respectively. In \mathcal{M}_D , by contrast, \mathbf{X}_i and \mathbf{X}_j are dependent. Given a dataset $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ of N independent and identically distributed (i.i.d.) realizations of \mathbf{X} and \mathbf{S} the corresponding sample sum-of-square matrix

$$\mathbf{S} = \sum_{n=1}^N (\mathbf{x}_n - \boldsymbol{\mu})(\mathbf{x}_n - \boldsymbol{\mu})^\dagger,$$

we propose to quantify the similarity between \mathbf{X}_i and \mathbf{X}_j as the log Bayes factor, that is, the log ratio of the marginal model likelihoods of \mathcal{M}_D versus \mathcal{M}_I

$$s(\mathbf{X}_i, \mathbf{X}_j) = \ln \frac{p(\mathbf{S}_{i \cup j} | \mathcal{M}_D)}{p(\mathbf{S}_{i \cup j} | \mathcal{M}_I)}. \tag{1}$$

Each marginal model likelihood can be expressed as an integral over the model parameters as described below.

Note that we assumed a known mean in the following theoretical development for the sake of simplicity. If the mean is unknown (as will be the case in the simulation and real data sections), this development is still valid, with N replaced by $N - 1$ and $\boldsymbol{\mu}$ by the sample mean

$$\mathbf{m} = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n.$$

in the expression of the sample sum-of-square matrix.

Marginal model likelihood under the hypothesis of dependence

Let us first calculate $P(\mathbf{S}_{i \cup j} | \mathcal{M}_D)$, the marginal model likelihood under the hypothesis of dependence. Expressing this quantity as a function of the model parameters yields

$$p(\mathbf{S}_{i \cup j} | \mathcal{M}_D) = \int p(\mathbf{S}_{i \cup j} | \mathcal{M}_D, \boldsymbol{\Sigma}_{i \cup j}) p(\boldsymbol{\Sigma}_{i \cup j} | \mathcal{M}_D) d\boldsymbol{\Sigma}_{i \cup j}. \tag{2}$$

Calculation of the integral requires to know the likelihood $p(\mathbf{S}_{i \cup j} | \mathcal{M}_D, \Sigma_{i \cup j})$ and the prior distribution $p(\Sigma_{i \cup j} | \mathcal{M}_D)$ of the covariance matrix under \mathcal{M}_D . With multivariate normal data, $\mathbf{S}_{i \cup j}$ given $\Sigma_{i \cup j}$ is Wishart distributed with N degrees of freedom and scale matrix $\Sigma_{i \cup j}$ ([23] Corollary 7.2.2), leading to the following likelihood

$$p(\mathbf{S}_{i \cup j} | \mathcal{M}_D, \Sigma_{i \cup j}) = \frac{|\mathbf{S}_{i \cup j}|^{\frac{N - D_{i \cup j} - 1}{2}}}{Z(D_{i \cup j}, N)} |\Sigma_{i \cup j}|^{-\frac{N}{2}} \exp \left[-\frac{1}{2} \text{tr}(\Sigma_{i \cup j}^{-1} \mathbf{S}_{i \cup j}) \right], \tag{3}$$

where $Z(d, n)$ is the inverse of the normalization constant

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

As to the prior distribution, this quantity is here set as a conjugate prior, namely an inverse-Wishart distribution with $v_{i \cup j}$ degrees of freedom and inverse scale matrix $\Lambda_{i \cup j}$ ([24] §3.6)

$$p(\Sigma_{i \cup j} | \mathcal{M}_D) = \frac{|\Lambda_{i \cup j}|^{\frac{v_{i \cup j}}{2}}}{Z(D_{i \cup j}, v_{i \cup j})} |\Sigma_{i \cup j}|^{-\frac{v_{i \cup j} + D_{i \cup j} + 1}{2}} \exp \left[-\frac{1}{2} \text{tr}(\Sigma_{i \cup j}^{-1} \Lambda_{i \cup j}) \right]. \tag{4}$$

Bringing Eqs (3) and (4) together into Eq (2) yields for the marginal model likelihood

$$p(\mathbf{S}_{i \cup j} | \mathcal{M}_D) = \frac{|\Lambda_{i \cup j}|^{\frac{v_{i \cup j}}{2}} |\mathbf{S}_{i \cup j}|^{\frac{N - D_{i \cup j} - 1}{2}}}{Z(D_{i \cup j}, N) Z(D_{i \cup j}, v_{i \cup j})} \times \int |\Sigma_{i \cup j}|^{-\frac{N + v_{i \cup j} + D_{i \cup j} + 1}{2}} \exp \left\{ -\frac{1}{2} \text{tr} \left[(\Lambda_{i \cup j} + \mathbf{S}_{i \cup j}) \Sigma_{i \cup j}^{-1} \right] \right\} d\Sigma_{i \cup j}.$$

The integrand is proportional to an inverse-Wishart distribution with $N + v_{i \cup j}$ degrees of freedom and scale matrix $\Lambda_{i \cup j} + \mathbf{S}_{i \cup j}$, leading to

$$p(\mathbf{S}_{i \cup j} | \mathcal{M}_D) = \frac{|\mathbf{S}_{i \cup j}|^{\frac{N - D_{i \cup j} - 1}{2}}}{Z(D_{i \cup j}, N)} \frac{Z(D_{i \cup j}, N + v_{i \cup j})}{Z(D_{i \cup j}, v_{i \cup j})} \frac{|\Lambda_{i \cup j}|^{\frac{v_{i \cup j}}{2}}}{|\mathbf{S}_{i \cup j} + \Lambda_{i \cup j}|^{\frac{N + v_{i \cup j}}{2}}}. \tag{5}$$

Marginal model likelihood under the hypothesis of independence

We can now calculate $P(\mathbf{S}_{i \cup j} | \mathcal{M}_I)$, the marginal model likelihood under the hypothesis of independence. If \mathcal{M}_I holds, then $\Sigma_{i \cup j}$ is block-diagonal with two blocks Σ_i and Σ_j the submatrix restrictions of $\Sigma_{i \cup j}$ to \mathbf{X}_i and \mathbf{X}_j , respectively. Introduction of the model parameters therefore yields for the marginal likelihood

$$p(\mathbf{S}_{i \cup j} | \mathcal{M}_I) = \int p(\mathbf{S}_{i \cup j} | \mathcal{M}_I, \Sigma_i, \Sigma_j) p(\Sigma_i, \Sigma_j | \mathcal{M}_I) d\Sigma_i d\Sigma_j. \tag{6}$$

To calculate this integral, we again need to know the likelihood $p(\mathbf{S}_{i \cup j} | \mathcal{M}_I, \Sigma_i, \Sigma_j)$ and the prior distribution $p(\Sigma_i, \Sigma_j | \mathcal{M}_D)$ of the two blocks of the covariance matrix under \mathcal{M}_I . The likelihood is the same as for \mathcal{M}_D and has the form of Eq (3). Furthermore, since $\Sigma_{i \cup j}$ is here block diagonal, we have $|\Sigma_{i \cup j}| = |\Sigma_i| |\Sigma_j|$ and $\text{tr}(\Sigma_{i \cup j}^{-1} \mathbf{S}_{i \cup j}) = \text{tr}(\Sigma_i^{-1} \mathbf{S}_i) + \text{tr}(\Sigma_j^{-1} \mathbf{S}_j)$, where \mathbf{S}_i and \mathbf{S}_j are the restrictions of \mathbf{S} to \mathbf{X}_i and \mathbf{X}_j , respectively. Consequently, the likelihood can be further

expanded as

$$p(\mathbf{S}_{i|j} | \mathcal{M}_I, \boldsymbol{\Sigma}_i, \boldsymbol{\Sigma}_j) = \frac{|\mathbf{S}_{i|j}|^{\frac{N - D_{i|j} - 1}{2}}}{Z(D_{i|j}, N)} \prod_{k \in \{i,j\}} |\boldsymbol{\Sigma}_k|^{-\frac{N}{2}} \exp \left[-\frac{1}{2} \text{tr}(\boldsymbol{\Sigma}_k^{-1} \mathbf{S}_k) \right]. \tag{7}$$

As to the prior distribution, assuming no prior dependence between $\boldsymbol{\Sigma}_i$ and $\boldsymbol{\Sigma}_j$ yields

$$p(\boldsymbol{\Sigma}_i, \boldsymbol{\Sigma}_j | \mathcal{M}_I) = p(\boldsymbol{\Sigma}_i | \mathcal{M}_I) p(\boldsymbol{\Sigma}_j | \mathcal{M}_I). \tag{8}$$

For the sake of consistency, $p(\boldsymbol{\Sigma}_i | \mathcal{M}_I)$ and $p(\boldsymbol{\Sigma}_j | \mathcal{M}_I)$ are set equal to $p(\boldsymbol{\Sigma}_i | \mathcal{M}_D)$ and $p(\boldsymbol{\Sigma}_j | \mathcal{M}_D)$, respectively, which are in turn obtained by marginalization of $p(\boldsymbol{\Sigma}_{i|j} | \mathcal{M}_D)$ as given by Eq (4). For $k \in \{i, j\}$, $p(\boldsymbol{\Sigma}_k | \mathcal{M}_I)$ can be found to have an inverse-Wishart distribution with $\nu_k = \nu_{i|j} - D_{i|j} + D_k$ degrees of freedom and inverse scale matrix $\boldsymbol{\Lambda}_k$ the restriction of $\boldsymbol{\Lambda}_{i|j}$ to \mathbf{X}_k ([25] §5.2)

$$p(\boldsymbol{\Sigma}_k | \mathcal{M}_I) = \frac{|\boldsymbol{\Lambda}_k|^{\frac{\nu_k}{2}}}{Z(D_k, \nu_k)} |\boldsymbol{\Sigma}_k|^{-\frac{\nu_k + D_k + 1}{2}} \exp \left[-\frac{1}{2} \text{tr}(\boldsymbol{\Lambda}_k \boldsymbol{\Sigma}_k^{-1}) \right]. \tag{9}$$

Incorporating Eqs (7), (8), and (9) into Eq (6) yields

$$p(\mathbf{S}_{i|j} | \mathcal{M}_I) = \frac{|\mathbf{S}_{i|j}|^{\frac{N - D_{i|j} - 1}{2}}}{Z(D_{i|j}, N)} \prod_{k \in \{i,j\}} \frac{|\boldsymbol{\Lambda}_k|^{\frac{\nu_k}{2}}}{Z(D_k, \nu_k)} \times \int |\boldsymbol{\Sigma}_k|^{-\frac{N + \nu_k + D_k + 1}{2}} \exp \left\{ -\frac{1}{2} \text{tr}[(\mathbf{S}_k + \boldsymbol{\Lambda}_k) \boldsymbol{\Sigma}_k^{-1}] \right\}.$$

Each integrand is proportional to an inverse-Wishart distribution with $N + \nu_k$ degrees of freedom and scale matrix $\mathbf{S}_k + \boldsymbol{\Lambda}_k$, leading to

$$p(\mathbf{S}_{i|j} | \mathcal{M}_I) = \frac{|\mathbf{S}_{i|j}|^{\frac{N - D_{i|j} - 1}{2}}}{Z(D_{i|j}, N)} \prod_{k \in \{i,j\}} \frac{Z(D_k, N + \nu_k)}{Z(D_k, \nu_k)} \frac{|\boldsymbol{\Lambda}_k|^{\frac{\nu_k}{2}}}{|\mathbf{S}_k + \boldsymbol{\Lambda}_k|^{\frac{N + \nu_k}{2}}}. \tag{10}$$

Log Bayes factor of dependence versus independence

Let us now express the Bayesian similarity measure by incorporating Eqs (5) and (10) into Eq (1), yielding

$$s(\mathbf{X}_i, \mathbf{X}_j) = \left[\sum_{k \in \{i,j\}} \frac{N + \nu_k}{2} \ln |\boldsymbol{\Lambda}_k + \mathbf{S}_k| \right] - \frac{N + \nu_{i|j}}{2} \ln |\boldsymbol{\Lambda}_{i|j} + \mathbf{S}_{i|j}| + \text{cst} \tag{11}$$

with

$$\begin{aligned} \text{cst} = & \frac{\nu_{i|j}}{2} \ln |\boldsymbol{\Lambda}_{i|j}| + \sum_{d=1}^{D_{i|j}} \left[\ln \Gamma \left(\frac{N + \nu_{i|j} + 1 - d}{2} \right) - \ln \Gamma \left(\frac{\nu_{i|j} + 1 - d}{2} \right) \right] \\ & - \sum_{k \in \{i,j\}} \left\{ \frac{\nu_k}{2} \ln |\boldsymbol{\Lambda}_k| + \sum_{d=1}^{D_k} \left[\ln \Gamma \left(\frac{N + \nu_k + 1 - d}{2} \right) + \ln \Gamma \left(\frac{\nu_k + 1 - d}{2} \right) \right] \right\}. \end{aligned} \tag{12}$$

Another form for $s(\mathbf{X}_i, \mathbf{X}_j)$ is

$$s(\mathbf{X}_i, \mathbf{X}_j) = \Delta \phi_{i|j} - \Delta \phi_i - \Delta \phi_j, \tag{13}$$

with

$$\Delta\phi_k = \phi(N + v_k, \mathbf{\Lambda}_k + \mathbf{S}_k) - \phi(v_k, \mathbf{\Lambda}_k), \quad k \in \{i, j, i \cup j\}, \tag{14}$$

and

$$\phi(n, \mathbf{A}) = -\frac{n}{2} \ln |\mathbf{A}| + \sum_{d=1}^{\dim(\mathbf{A})} \ln \Gamma\left(\frac{n+1-d}{2}\right). \tag{15}$$

$\Delta\phi_k$ quantifies, up to a constant that cancels out in $s(\mathbf{X}_i, \mathbf{X}_j)$, the amount by which the data support a model of multivariate normal distributions with unrestricted covariance matrix for \mathbf{X}_k .

Asymptotic form of the log Bayes factor

We can now provide an asymptotic expression for $s(\mathbf{X}_i, \mathbf{X}_j)$. Define $\hat{\mathbf{S}}_k$ as the standard sample covariance matrix, i.e., $\mathbf{S}_k = N\hat{\mathbf{S}}_k$. When $N \rightarrow \infty$, we can use Stirling approximation ([26] p. 257) to expand the expression ϕ of Eq (15), leading to (see §B of S1 File)

$$\begin{aligned} s(\mathbf{X}_i, \mathbf{X}_j) &= \frac{N}{2} \ln \frac{|\hat{\mathbf{S}}_i| |\hat{\mathbf{S}}_j|}{|\hat{\mathbf{S}}_{i \cup j}|} - \frac{1}{2} \left[\frac{D_{i \cup j}(D_{i \cup j} + 1)}{2} - \sum_{k \in \{i, j\}} \frac{D_k(D_k + 1)}{2} \right] \ln N + O(1) \\ &= N \hat{I}(\mathbf{X}_i, \mathbf{X}_j) - \frac{D_i D_j}{2} \ln N + O(1), \end{aligned} \tag{16}$$

where

$$\hat{I}(\mathbf{X}_i, \mathbf{X}_j) = \frac{1}{2} \ln \frac{|\hat{\mathbf{S}}_i| |\hat{\mathbf{S}}_j|}{|\hat{\mathbf{S}}_{i \cup j}|}$$

is the plug-in estimator of mutual information for a multivariate normal distribution. Alternatively, $N \hat{I}(\mathbf{X}_i, \mathbf{X}_j)$ can be seen as the minimum discrimination information for the independence of \mathbf{X}_i and \mathbf{X}_j ([12] Chap. 12, §3.6).

Hierarchical agglomerative clustering

A hierarchy on a set of D variables is a nested set of partitions $(C_l)_{l=0}^D$, where C_l is a partition of $\{1, \dots, D\}$ into $D - l$ clusters [2]. A hierarchical agglomerative clustering (AHC) is a generic procedure to generate such a hierarchy, outlined in pseudo-code in Algorithm 1. The main steps of the algorithm are: (1) Initialize the partition with singletons (line 2); (2) derive a matrix s_l where each element represents the similarity between two clusters \mathbf{X}_i and \mathbf{X}_j of C_l , based on an arbitrary function $s(\mathbf{X}_i, \mathbf{X}_j)$ (line 4); (3) identify the two clusters that are most similar (line 5); (4) form a new partition identical to the previous one, except that the two most similar clusters are now merged (lines 6–7); (5) iterate Steps 2–4 until the partition has only one single element which covers the whole set of variables (line 3). In the case of the methods proposed here, the similarity measure is given by either Eq (11) for the exact formulation or Eq (16) for the asymptotic BIC approximation.

Note that for the selection of the pair of clusters that are most similar (step 3), there may be more than one pair of clusters which maximize the similarity function. Most implementations of AHC proceed by selecting arbitrarily one such pair (e.g., the first one to be detected). In the in-house implementation we used, the pair was selected randomly amongst all these pairs. This was done to properly capture the instability of the algorithm. In such a form, AHC may not be

deterministic anymore, in that two runs of the same algorithm on the same dataset may result in different hierarchies.

```

1: return Hierarchy ( $\mathcal{C}_l$ ) $l=0$  $D$ 
2:  $\mathcal{C}_0 \leftarrow \{\{X_1\}, \dots, \{X_D\}\}$ 
3: for  $l = 0, \dots, D - 2$  do
4:    $\mathbf{s}_l \leftarrow [s(\mathbf{X}_i, \mathbf{X}_j)]_{\mathbf{x}_i, \mathbf{x}_j \in \mathcal{C}_l}$ 
5:    $(i^*, j^*) \leftarrow \arg \max_{i, j} \mathbf{s}_l(i, j)$ 
6:    $\mathcal{C}_{l+1} \leftarrow \mathcal{C}_l \setminus \{i^*, j^*\}$ 
7:    $\mathcal{C}_{l+1} \leftarrow \mathcal{C}_l \cup \{i^* \cup j^*\}$ 
8: end for

```

Algorithm 1 General description of the hierarchial agglomerative clustering algorithm.

Comparing distinct levels of the hierarchy

The last development aims at providing a way to compare nested partitions, i.e., different levels of the hierarchy. Once the hierarchical clustering has been performed, each step is associated with a partition of \mathbf{X} . Assume that, at level l , the partition reads $\{\mathbf{X}_1, \dots, \mathbf{X}_K\}$ and that, at step $l + 1$, \mathbf{X}_i and \mathbf{X}_j are clustered. Denote by \mathcal{C}_l the assumption that the partition at level l is the correct partition of \mathbf{X} . The global improvement brought by the clustering of \mathbf{X}_i and \mathbf{X}_j between steps l and $l + 1$ can be quantified by the log ratio between the marginal model likelihoods

$$\ln \frac{p(\mathcal{S}|\mathcal{C}_{l+1})}{p(\mathcal{S}|\mathcal{C}_l)},$$

where both quantities can be computed in a manner similar to what was done for the similarity measure. For instance, if \mathcal{C}_l is true, then Σ is block-diagonal with K blocks Σ_k 's, the submatrix restrictions of Σ to \mathbf{X}_k . Introducing the model parameters then yields

$$p(\mathcal{S}|\mathcal{C}_l) = \int p(\mathcal{S}|\mathcal{C}_l, \Sigma_1, \dots, \Sigma_K) p(\Sigma_1, \dots, \Sigma_K | \mathcal{C}_l) \prod_{k=1}^K d\Sigma_k. \tag{17}$$

In a way similar to what was done previously, the likelihood $p(\mathcal{S}|\mathcal{C}_l, \Sigma_1, \dots, \Sigma_K)$ can be expanded as

$$p(\mathcal{S}|\mathcal{C}_l, \Sigma_1, \dots, \Sigma_K) = \frac{|\mathcal{S}|^{\frac{N-D-1}{2}}}{Z(D, N)} \prod_{k=1}^K |\Sigma_k|^{-\frac{N}{2}} \exp \left[-\frac{1}{2} \text{tr}(\Sigma_k^{-1} \mathcal{S}_k) \right]. \tag{18}$$

Turning to the prior distribution $p(\Sigma_1, \dots, \Sigma_K | \mathcal{C}_l)$ and assuming no prior dependence between the Σ_k 's, we can set

$$p(\Sigma_1, \dots, \Sigma_K | \mathcal{C}_l) = \prod_{k=1}^K p(\Sigma_k | \mathcal{C}_l). \tag{19}$$

Each $p(\Sigma_k | \mathcal{C}_l)$ can be obtained by the marginalization of $p(\Sigma | \mathcal{C}_l)$, which is here taken as an inverse-Wishart distribution with ν degrees of freedom and inverse scale matrix the D -by- D diagonal matrix Λ . Note that such a prior on Σ is compatible with the prior used earlier for $\Sigma_{i \cup j}$ if one sets $\nu_{i \cup j} = \nu - D + D_{i \cup j}$ and if $\Lambda_{i \cup j}$ is the restriction of Λ to $\mathbf{X}_{i \cup j}$ ([25] §5.2). We then have

$$p(\Sigma_k | \mathcal{C}_l) = \frac{|\Lambda_k|^{-\frac{\nu_k}{2}}}{Z(D_k, \nu_k)} |\Sigma_k|^{-\frac{\nu_k + D_k + 1}{2}} \exp \left[-\frac{1}{2} \text{tr}(\Lambda_k \Sigma_k^{-1}) \right]. \tag{20}$$

Incorporating Eqs (18), (19), and (20) into Eq (17) and integrating leads to

$$p(\mathcal{S}|\mathcal{C}_l) = \frac{|\mathcal{S}|^{\frac{N-D-1}{2}}}{Z(D, N)} \prod_{k=1}^K \frac{Z(D_k, N + v_k)}{Z(D_k, v_k)} \frac{|\Lambda_k|^{\frac{v_k}{2}}}{|\Lambda_k + \mathcal{S}_k|^{\frac{N+v_k}{2}}}. \tag{21}$$

The same calculation can be performed for $p(\mathcal{S}|\mathcal{C}_{l+1})$. The result is the same as in Eq (21), except that the product is composed of $K - 1$ terms. Of these terms, $K - 2$ correspond to clusters that are unchanged from \mathcal{C}_l to \mathcal{C}_{l+1} and, as a consequence, are identical to those of Eq (21). The $(K - 1)$ th term corresponds to the cluster obtained by the merging of \mathbf{X}_i and \mathbf{X}_j . As a consequence, the log Bayes factor reads

$$\begin{aligned} \ln \frac{p(\mathcal{S}|\mathcal{C}_{l+1})}{p(\mathcal{S}|\mathcal{C}_l)} &= \ln \left[\frac{Z(D_{i,j}, N + v_{i,j})}{Z(D_{i,j}, v_{i,j})} \frac{|\Lambda_{i,j}|^{\frac{v_{i,j}}{2}}}{|\Lambda_{i,j} + \mathcal{S}_{i,j}|^{\frac{N+v_{i,j}}{2}}} \right] \\ &\quad - \sum_{k \in \{i,j\}} \ln \left[\frac{Z(D_k, N + v_k)}{Z(D_k, v_k)} \frac{|\Lambda_k|^{\frac{v_k}{2}}}{|\Lambda_k + \mathcal{S}_k|^{\frac{N+v_k}{2}}} \right]. \end{aligned}$$

But this quantity is nothing else than $s(\mathbf{X}_i, \mathbf{X}_j)$. In other words, we proved that

$$\ln \frac{p(\mathcal{S}|\mathcal{C}_{l+1})}{p(\mathcal{S}|\mathcal{C}_l)} = s(\mathbf{X}_i, \mathbf{X}_j), \tag{22}$$

i.e., $s(\mathbf{X}_i, \mathbf{X}_j)$, the *local* measure of similarity between \mathbf{X}_i and \mathbf{X}_j , can be used to compute the *global* measure of relative probability between two successive levels of the hierarchy.

Setting the hyperparameters

Hyperparameter selection is often a thorny issue in Bayesian analysis. We here considered two approaches. The first approach (coined BayesCov) is to set 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 (21) before any clustering (that is, with $K = D$ clusters and $D_k = 1$ for all k), yielding (see §C of S1 File)

$$\Lambda_{dd} = \frac{\nu - D + 1}{N} S_{dd},$$

where Λ_{dd} and S_{dd} are the diagonal elements of the prior inverse scale matrix Λ and sum-of-square matrix \mathcal{S} , respectively. An alternative approach (coined BayesCorr) is to work 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. The corresponding prior distribution yields uniform marginal distributions for the correlation coefficients [27]. Note that clustering with the asymptotic form of Eq (16) (coined Bic) does not involve hyperparameters; it is also insensitive to the fact that the input is the covariance matrix or the correlation matrix.

Automatic stopping rule

An advantage of the Bayesian clustering scheme proposed here and its BIC approximation is that they come naturally with an automatic stopping rule. By definition of s in Eq (1), the fact that $s(\mathbf{X}_{i_0}, \mathbf{X}_{j_0}) > 0$ for the pair that is selected for clustering also means that the marginal model likelihood for \mathcal{M}_D is larger than that for \mathcal{M}_I . As such, \mathbf{X}_{i_0} and \mathbf{X}_{j_0} are more likely to belong to the same cluster than not and, as a consequence, it indeed makes sense to cluster them. By

contrast, if we have $s(X_{i_0}, X_{j_0}) < 0$ for the same pair, the marginal model likelihood for \mathcal{M}_D is smaller than that for \mathcal{M}_I . X_{i_0} and X_{j_0} are therefore more likely to belong to different clusters. If, at a given step of the clustering, the pair with highest similarity measure has a negative similarity measure, then all pairs do, meaning that all pairs of clusters tested more probably belong to different clusters. It therefore makes sense to stop the clustering procedure at that point. This shows that an automatic stopping rule can simply be implemented into the clustering algorithm: Stop the clustering if the pair (X_{i_0}, X_{j_0}) selected for clustering at a given step has $s(X_{i_0}, X_{j_0}) < 0$. Note that, according to Eq (22), the resulting clustering corresponds to the one in the hierarchy that has largest marginal likelihood. We will refer to BayesCovAuto, BayesCorrAuto and BicAuto when applying the clustering schemes with this automated stopping rule.

Results

Validation on synthetic data

To assess the behavior of the method expounded here, we examined how it fared on synthetic data. We used the two variants of the Bayes factor mentioned above (BayesCov and BayesCorr), Bic, as well as the same methods with automatic stopping rule (BayesCovAuto, BayesCorrAuto and BicAuto). As a mean of comparison, we also used the following methods—

- A random hierarchical clustering scheme, where variables were clustered uniformly at random at each step. This category contains only one algorithm: Random, which was implemented for the purpose of the present study.
- Hierarchical clustering schemes with similarity measures given by either Pearson correlation or absolute Pearson correlation, and a merging rule based on either the single, average, or complete linkage, or using Ward criterion. This category contains 8 algorithms: Single, Average, Complete, Ward, SingleAbs, AverageAbs, CompleteAbs, WardAbs. We used the implementations of these methods proposed in NIAK (<https://github.com/SIMEXP/niak>)
- Hierarchical clustering with a similarity measure given by mutual information, with and without normalization. This category contains 2 algorithms: Infomut and InfomutNorm. These methods were implemented for the purpose of the present study. Note that neither algorithm can run on small samples.
- An approach where the clusters are estimated as the blocks of the precision (i.e., inverse covariance) matrix estimated with the graphical lasso—essentially a maximum-likelihood with L_1 -norm penalization [28]. The penalization parameter λ was set in $[0, 1]$ by step of 0.1, then to 5, 10, 20, 50, 100, 200 [29]. A version that optimizes λ with a BIC criterion was also used [30]. Since the graphical lasso is not invariant by transformation of a covariance matrix into a correlation matrix, we used either the covariance matrix or the correlation matrix as input. Note that this approach automatically determined a number of clusters. Also, for $\lambda = 0$ (unconstrained case), the algorithm cannot run on small samples. This category contains 34 algorithms: 16 algorithms gLassoCov-x and 16 algorithms gLassoCorr-x, where x is the value of λ , and 2 algorithms gLassoCovOpt and gLassoCorrOpt. For this category of algorithms, we used a package freely available (<http://www.cs.ubc.ca/~schmidtm/Software/L1precision.html>) and already used in [29].
- An approach based on the spectral clustering [31] of either the raw value or the absolute value of either the correlation or the partial correlation matrix. This approach required the number of clusters as input. Since this clustering has a step of k -means, which is stochastic by nature, we considered 2 variants: one with 1 step of k -means and the other one with 10 repetitions of k -means and selection of the best clustering in terms of inertia. The similarity

measures were defined so that the range would be the same (namely in $[0, 1]$) when using the signed or the absolute value of correlation: $0.5(1+r)$ and $|r|$, respectively. This category contains 8 algorithms: 2 algorithms SpecCorr-x, 2 algorithms SpecCorrAbs-x, 2 algorithms SpecCorrpar-x and 2 algorithms SpecCorrparAbs-x, where x is the number of times that k -means is performed. The spectral clustering algorithms of this category were coded for the purpose of the present study, while we used the implementations of the k -means algorithm proposed in NIAK.

All in all, 59 variants were tested.

Data description. In order to assess the performance of the Bayesian approach, we performed the following set of simulations. For each value of D in $\{6, 10, 20, 40\}$, we considered partitions with increasing number of clusters C ($1 \leq C \leq D$). For a given value of C , we performed 500 simulations. For each simulation, the D variables were randomly partitioned into C clusters, all partitions having equal probability of occurrence ([32] chap. 12); [33]. For a given partition $\{X_1, \dots, X_K\}$ of X , we generated data according to

$$f(\mathbf{x}) = \prod_{k=1}^K f_k(\mathbf{x}_k), \tag{23}$$

where all f_k 's were taken either as multivariate normal distributions (parameters: mean μ_k and covariance matrix Σ_k) or multivariate Student- t distributions (parameters: degrees of freedom ν , location parameter μ_k and scale matrix Σ_k). In both case, the μ_k 's were set to $\mathbf{0}$ and the Σ_k 's were first sampled according to a Wishart distribution with $D_k + 1$ degrees of freedom and scale matrix the identity matrix and then rescaled to a correlation matrix. The sampling scheme on Σ_k generated correlation matrices with uniform marginal distributions for all correlation coefficients [27]. For the multivariate Student- t distributions, ν was set in $\{1, 3, 5\}$. Eq (23) was used to generate synthetic datasets of length N varying from 10 to 300 by increment of 40. Each dataset was summarized by its sample correlation matrix and hierarchical clustering was performed using the methods mentioned above. All simulations were implemented using the Pipeline System for Octave and Matlab, PSOM (<https://github.com/SIMEXP/psom>) [34] under Matlab 7.2 (The MathWorks, Inc.) and run on a 24-core server.

To assess the efficiency of the various methods, we thresholded each clustering at the true number of clusters, except for BayesCovAuto, BayesCorrAuto, BicAuto and gLasso for which we used the clustering obtained by the method. We then quantified the quality of the resulting partition using the proportion of correct classifications as well as the adjusted Rand index, which computes the fraction of variable pairs that are correctly clustered corrected for chance [35, 36]. Results corresponding to a given dimension D and a given method were then pooled across numbers of clusters C , lengths N and distributions (multivariate normal and Student). We classified the methods from best to worst based on these global results using the following indices (in this order): median of adjusted Rand index, 25% percentile value of adjusted Rand index, 5% percentile value of adjusted Rand index, smallest value of adjusted Rand index, and proportion of correct classifications. Note that some algorithms (Infomut, InfomutNorm and SpecCorrpar) require the sample covariance matrix to be definite positive. As a consequence, these algorithms could not run on small samples. We therefore restrained our evaluation to cases where all algorithms were operational. Finally, we performed a Bayesian ANOVA-like regression analysis ([24] §15.6), where we explained the adjusted Rand index of nine algorithms (BayesCov, BayesCovAuto, BayesCorr, BayesCorrAuto, Bic, BicAuto, Infomut, InfomutNorm, and AverageAbs) with the following effects: clustering algorithm (9 levels), number of variables D (4 levels: $D \in \{6, 10, 20, 40\}$), type of distribution (4 levels: multivariate Gaussian or multivariate Student- t with 1, 3, or 5 degrees of freedom), number of samples N (8 levels: N

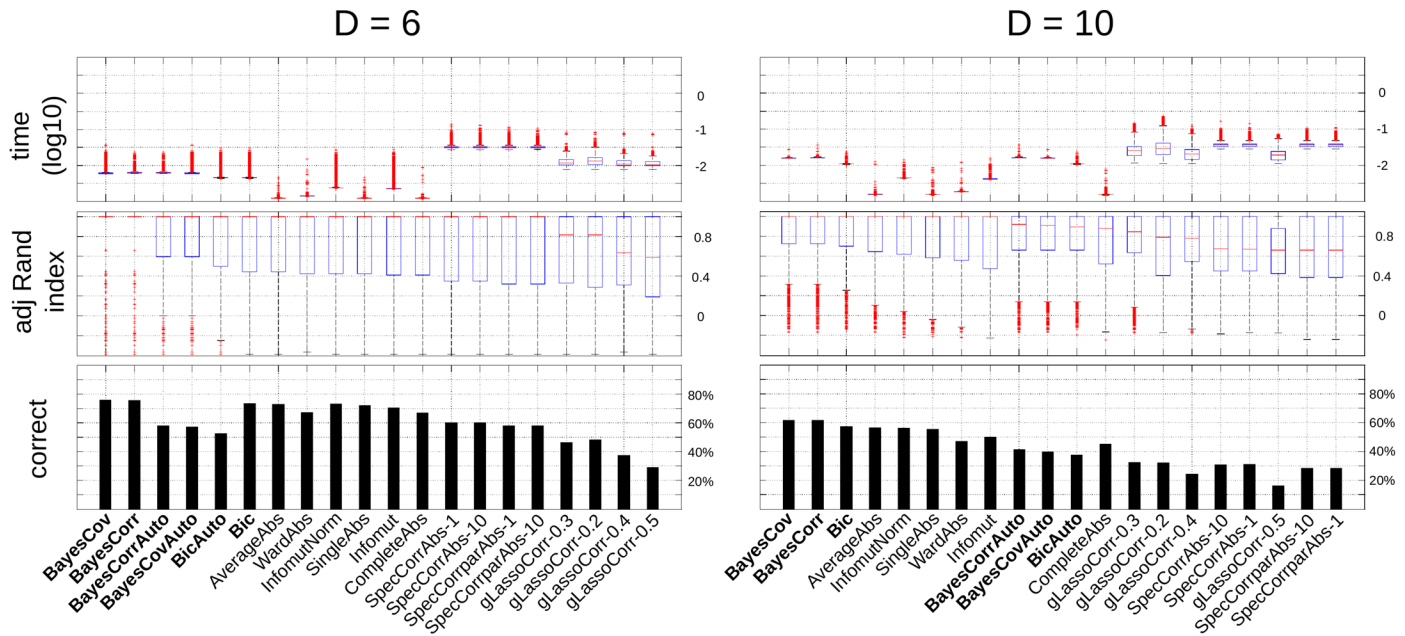


Fig 1. Simulation study. Computational time (top), adjusted Rand index (middle) and proportion of correct classifications (bottom) for $D = 6$ (left) and $D = 10$ (right).

doi:10.1371/journal.pone.0137278.g001

$\in \{10,50,90,130,170,210,250,290\}$), and number of clusters C (68 levels: from 2 to $D - 1$ clusters for each D). In other words, the model considered was of the form

$$\text{adj Rand index}(\text{algo}, D, \text{distr}, N, C) = \beta_0 + \beta_{\text{algo}} + \beta_D + \beta_{\text{distr}} + \beta_N + \beta_{D,C} + \epsilon. \quad (24)$$

Interactions between effects, while potentially relevant, were not considered to keep the analysis tractable. The posterior distribution for the various regression parameters were estimated using Gibbs sampling.

Results. The results corresponding to the adjusted Rand index and fraction of correct classification are summarized in Figs 1 and 2 for the 20 best methods. Globally, and as expected, all methods were adversely affected by an increase in the number of variables. In all cases, the variants proposed in the present paper performed very well compared to other methods. BayesCov and BayesCovAuto were always classified as the two best algorithms and Bic was never outperformed by a method already published. The methods with automatic thresholding of the hierarchy performed surprisingly well, considered that they were compared against methods with oracle. In particular, they clearly outperformed all variants of gLasso, the only method that was able to automatically determine the number of clusters. Of note, all variants of gLasso proved too slow to be applied to our simulation data for $D \in \{20,40\}$.

The results of the regression analysis are represented in Fig 3. The 9 algorithms selected included the ones proposed in the present manuscript (BayesCov, BayesCovAuto, BayesCov, BayesCovAuto, Bic, and BicAuto), the best-performing classic algorithm in the previous analysis (AverageAbs) as well as the algorithms based on mutual information (Infomut and InfomutNorm). We found that increasing the dataset size (N) increased the performance of the algorithm, while increasing the dimensionality of the problem (D) and the number of clusters

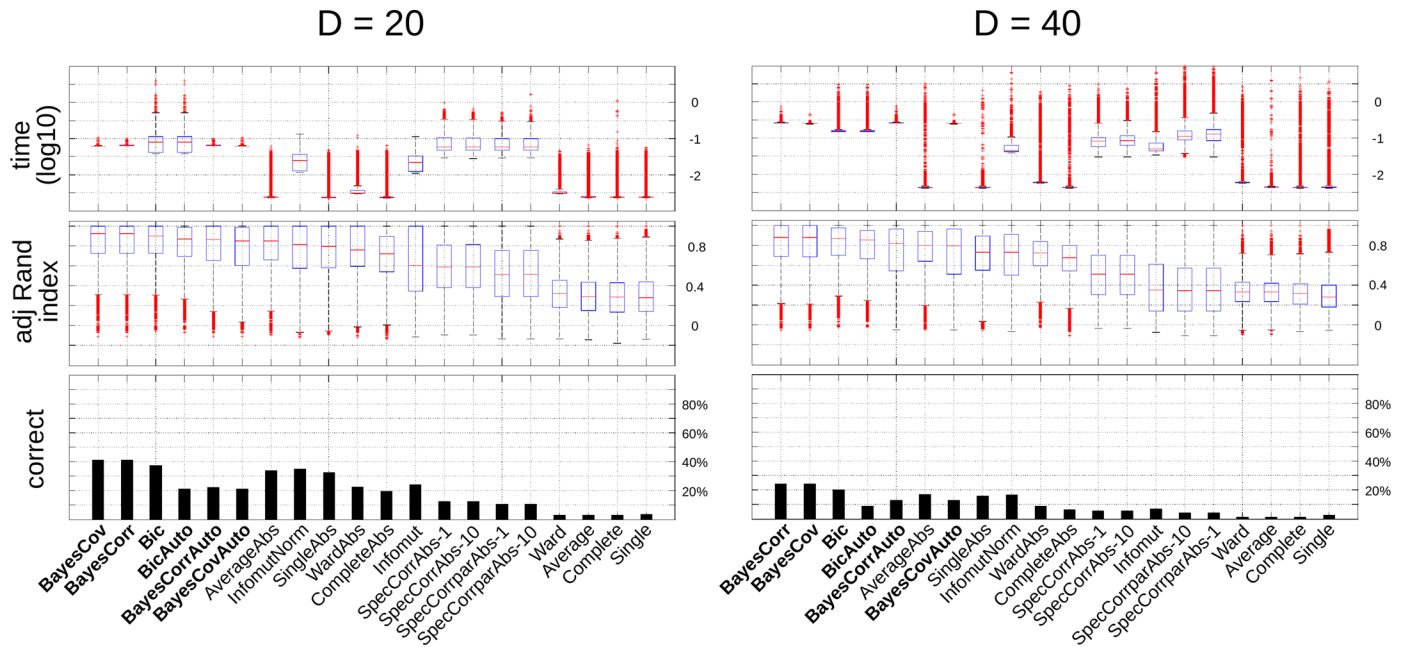


Fig 2. Simulation study. Computational time (top), adjusted Rand index (middle) and proportion of correct classifications (bottom) for $D = 20$ (left) and $D = 40$ (right).

doi:10.1371/journal.pone.0137278.g002

(C) decreased it. Note that dimension was found to have a negative influence on the adjusted Rand index, even though this index was partly proposed as a modification from the raw Rand index to overcome this limitation. Finally, this analysis confirmed the superior behavior of the methods proposed here, even when the method included the automatic stopping rule.

Toy example

This data set was first used in [37] and later re-analyzed in [38] in the context of conditional independence graphs. It originates from a study investigating early diagnosis of HIV infection in children from HIV positive mothers. 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 sample summary statistics are given in Table 1. [37] found that the correlations between X_4 and any other variable had strong probability around zero and hypothesized that the model was overparameterized. Based on the simulation study, we performed clustering of the data with the following methods: BayesCov(Auto), BayesCov(Auto), Bic(Auto), Infomut, InfomutNorm, SingleAbs, AverageAbs, CompleteAbs, WardAbs, SpecCorrAbs and SpecCorrparAbs. For spectral clustering, we used either 1 or 10 repetitions of the k -means step; since the results obtained for 1 step of k -means were highly variable for 3, 4, and 5 clusters, we discarded these results.

The resulting clusterings are given in Fig 4 and Table 2. All hierarchical clusterings started by clustering X_3 and X_5 (lymphocyte). This was confirmed by SpecCorrAbs-10 when required to provide 5 clusters. All hierarchical clustering methods then clustered X_1 and X_2 (immunoglobulin). This result was in agreement with both SpecCorrAbs-10 and SpecCorrparAbs-10

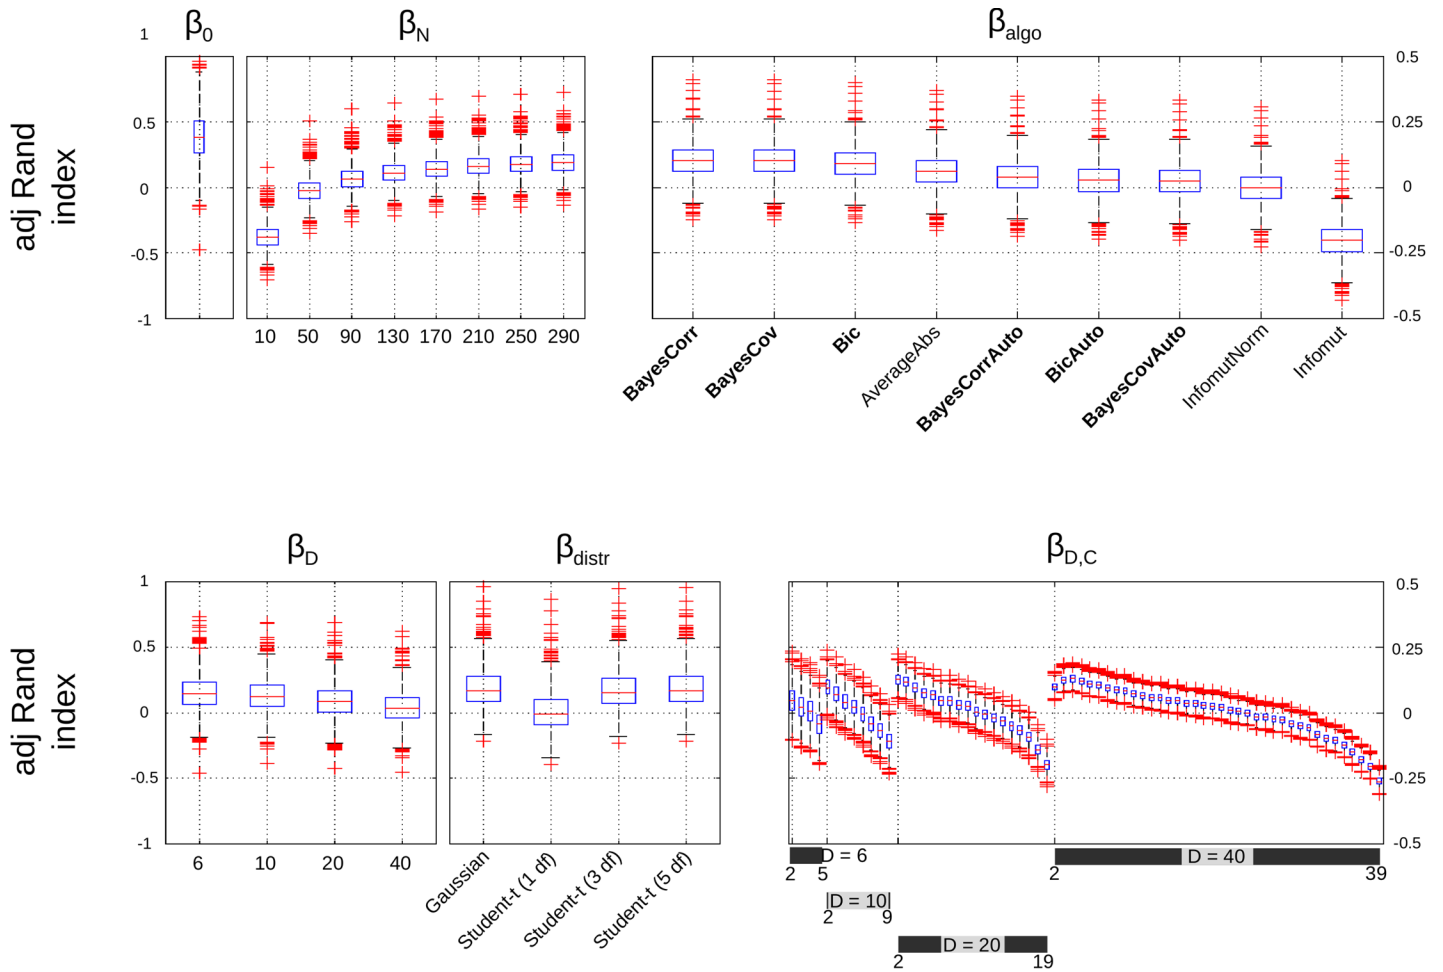


Fig 3. Simulation study. Result of the regression analysis. Posterior distribution for the different regression coefficients: β_0 (global effect), β_N (dataset size N), β_{algo} (algorithm), β_D (number of variables D), β_{distr} (type of distribution), and $\beta_{D,C}$ (number of clusters) [see Eq (24)].


doi:10.1371/journal.pone.0137278.g003

Table 1. Toy example. Summary statistics for the HIV data: sample variances (main diagonal), correlations (lower triangle) and partial correlations (upper triangle) (from [37]).


x_1	8.8374	0.479	-0.043	-0.033	0.356	-0.236
x_2	0.483	0.1919	0.068	-0.084	-0.224	-0.110
x_3	0.220	0.057	8924231.9	0.085	0.552	-0.330
x_4	-0.040	-0.133	0.149	20392.4	0.091	0.013
x_5	0.253	-0.124	0.523	0.179	1952795.2	0.384
x_6	-0.276	-0.314	-0.183	0.064	0.213	1.378
	x_1	x_2	x_3	x_4	x_5	x_6

doi:10.1371/journal.pone.0137278.t001


(G1): BayesCov, BayesCorr, Infomut, InfomutNorm



(G2): SingleAbs, AverageAbs



(G3): Bic



(G4): CompleteAbs, WardAbs




Fig 4. Toy example. Result of clustering. Algorithms in the top row clustered X_4 at the last step, while it was clustered at the before the last step for algorithms in the bottom row. Algorithms in the left column clustered X_6 with $\{X_3, X_5\}$, while X_6 was clustered with $\{X_1, X_2\}$ for the algorithms in the right column. Parts in grey correspond to clustering steps that were not performed by BayesCovAuto or BayesCorrAuto in (G1), or Bic in (G3).

doi:10.1371/journal.pone.0137278.g004

Table 2. Toy example. Result of spectral clustering with increasing number of clusters.

	SpecCorrAbs-10	SpecCorrparAbs-10
2 clusters		12356 4
3 clusters	126 35 4	12 356 4
4 clusters		12 35 4 6
5 clusters	1 2 35 4 6	not reproducible

doi:10.1371/journal.pone.0137278.t002

when required to provide 4 clusters. After the second step, we observed four behaviors for the AHC algorithms and classified them accordingly:

- (G1) BayesCov, BayesCorr, Infomut and >InfomutNorm;
- (G2) SingleAbs and AverageAbs;
- (G3) Bic;
- (G4) CompleteAbs and WardAbs.

While not a hierarchical clustering, SpecCorrAbs-10 provided successive clusterings that were in agreement with methods in (G2). Algorithms in (G1) and (G3) clustered X_6 with $\{X_3, X_5\}$, creating a cluster of variables related to lymphocyte. Algorithms in (G1) and (G2) (and SpecCorrAbs-10) agreed that a partitioning of the variables in two clusters should lead to $\{X_1, X_2, X_3, X_5, X_6\}$ on the one hand and $\{X_4\}$ on the other hand. This clustering was also found by SpecCorrpar-10 when constrained to extract two clusters from the data. It was also considered the best clustering for BayesCov and BayesCorr. For Bic, the optimal partitioning was composed of three clusters, namely, $\{X_1, X_3, X_6\}$, $\{X_1, X_2\}$, and $\{X_4\}$, which is in agreement with what would methods in (G1) yield for three clusters; furthermore, it still kept X_4 separated from the other variables.

In Fig 5, we represented the evolution of the \log_{10} Bayes factor during hierarchical clustering for BayesCov, BayesCorr and Bic. Note that, while the clustering steps are identical for BayesCov and BayesCorr, the log Bayes factors are similar but not identical. Likewise, while the first two clustering steps of Bic is identical to those of BayesCov and BayesCorr, one can see that, unlike BayesCov and BayesCorr, Bic considered the merging of $\{X_3, X_5\}$ with X_6 almost as likely as that of X_1 with X_2 . Also, while the successive clusterings of X_3 with X_5 and then X_6 as well as that of X_1 with X_2 strongly increased the Bayes factor for BayesCov and BayesCorr, the improvement brought by the clustering of $\{X_3, X_5, X_6\}$ with $\{X_1, X_2\}$ in these methods was less important.

All in all, this analysis led us to the following conclusion: it is very likely that variables X_1 and X_2 belong to the same cluster of dependent variables, and similarly for variables X_3 and X_5 . Also, there is very strong evidence in favor of the fact that X_4 is independent from the other variables. Finally, we suspect that X_3, X_5 and X_6 could belong to the same cluster of variables.

fMRI data

Cluster analysis is a popular tool to study the organization of brain networks in resting-state fMRI [39, 40], by identifying clusters of brain regions with highly correlated spontaneous activity. We applied the 13 methods that were found to have good performance on simulations (see Fig 6) to a collection of resting-state time series. The time series had 205 time samples and were recorded for 82 brain regions in 19 young healthy subjects. See §D of S1 File for details on




Fig 5. Toy example. Result of clustering for BayesCorr, BayesCov and Bic. The values on the y axis correspond to the \log_{10} Bayes factor in favor of the global clustering obtained at each step compared to a model where all variables are independent (step 0 of hierarchical clustering). The dotted lines correspond to clustering steps that were not performed with the automatic stopping rule.

doi:10.1371/journal.pone.0137278.g005

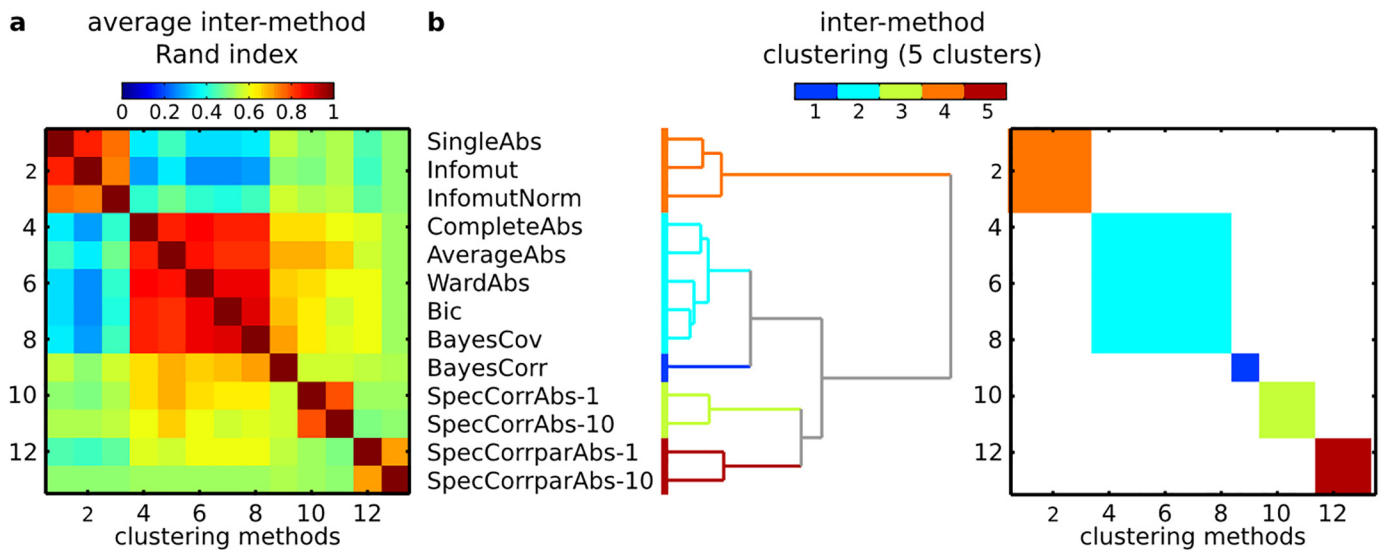


Fig 6. Real resting-state fMRI data—between-method similarity. Panel a: Rand indices between individual partitions generated with different methods, averaged across all subjects and scales (number of clusters). Panel b: Hierarchical clustering using matrix shown in Panel a as a similarity measure and Ward’s criterion.

doi:10.1371/journal.pone.0137278.g006

data collection and preparation. The data are available online (http://figshare.com/articles/Atlanta_resting_state_fmri_time_series_preprocessed_using_the_AAL_template/1521155).

We first aimed at establishing which clustering algorithms yielded similar results on these datasets. We more specifically investigated a 7-cluster solution, as this level of decomposition has been examined several times in the literature [39, 41, 42]. Each clustering algorithm was applied to the time series of each subject independently. For a given pair of methods, the similarity between the cluster solutions generated by both methods on the same subject was

evaluated with the Rand. Note that the raw, unnormalized Rand index was used here, as we did not compare cluster solutions with different numbers of clusters, which is the main motivation of the normalization. The unnormalized Rand index has a more intuitive interpretation than its adjusted counterpart. The Rand indices were averaged across subjects, hence resulting into a method-by-method matrix capturing the (average) similarity of clustering outputs for each pair of methods (see Fig 6). An AHC with Ward's criterion was applied to this matrix in order to identify clusters of methods with similar cluster outputs. We visually identified five clusters of methods that had high (> 0.7) average within-cluster Rand index. The largest cluster included classical AHCs such as CompleteAbs, AverageAbs, WardAbs, as well as the Bic and BayesCov methods proposed here. It should be noted that this class of algorithms generated solutions for this problem that were very close to one another (average within-cluster Rand index > 0.8). The BayesCorr method was also close to that large group of methods, but not quite as much as the aforementioned methods (average Rand index of about 0.7), and was thus singled out as a separate cluster. The spectral methods were split into two clusters, depending on whether they were based on correlation (SpecCorrAbs-1 and SpecCorrAbs-10) or partial correlation (SpecCorrparAbs-1 and SpecCorrparAbs-10). Finally, the two variants of mutual information (Infomut and InfomutNorm) generated solutions that were highly similar to SingleAbs. It was reassuring that the variants of Bayes methods proposed here performed similarly to algorithms known to produce physiologically plausible solutions, such as Ward [43, 44]. While we found some analogy between BayesCorr, BayesCov, Infomut and InfomutNorm, it was intriguing that the variants of mutual information tested seemed to generate markedly different classes of solutions from the Bayes methods. We decided to examine the cluster solutions of these methods in more details.

As a reference, we examined the cluster solutions generated by WardAbs, in addition to two variants of Bayes clustering that yielded slightly different solutions (BayesCov and BayesCorr), and normalized mutual information, InfomutNorm. To represent the typical behavior of each method across subjects, we generated a "group" consensus clustering summarizing the stable features of the ensemble of individual cluster solutions. This consensus clustering was generated by the evidence accumulation algorithm [45] outlined below. First, each partition of each subject was represented as a binary 81-by-81 adjacency matrix $A = (A_{ij})$, where for each pair (i, j) of brain regions, $A_{ij} = 1$ if areas i and j were in the same cluster, and $A_{ij} = 0$ otherwise. The adjacency matrices were then averaged across subjects and selected methods, yielding a 81-by-81 stability matrix $C = (C_{ij})$ where each element C_{ij} coded for the frequency at which brain areas i and j fell in the same cluster. Finally this stability matrix was used as a similarity matrix in a AHC with Ward's criterion to generate one consensus partition. The brain regions were grouped into the same consensus cluster if they had a high probability of falling into the same cluster on average across subjects and methods, hence the name consensus clustering.

Fig 7 represents the stability matrices and consensus clusters, for the four methods of interest. As expected based on our first experiment on the similarity of cluster outputs, the WardAbs and BayesCov methods were associated with highly similar stability matrices and almost identical consensus clusters. Many areas of high consensus could be identified (with values close of 0 or 1), illustrating the very good agreement of the cluster solutions across subjects. The outline of the consensus clusters as well as a volumetric representation of the brain parcellation are presented in Fig 7b. Some of these clusters closely matched those reported in previous studies: cluster 7 can be recognized as being the visual network, cluster 2 the sensorimotor network, and clusters 6 and 3 the anterior and posterior parts of the default-mode network, respectively [41, 46, 47]. By contrast with WardAbs and BayesCorr, the InfomutNorm tended to generate very large clusters, which was apparent both on the stability matrix and the consensus clusters. The BayesCorr method was intermediate between BayesCov and InfomutNorm in

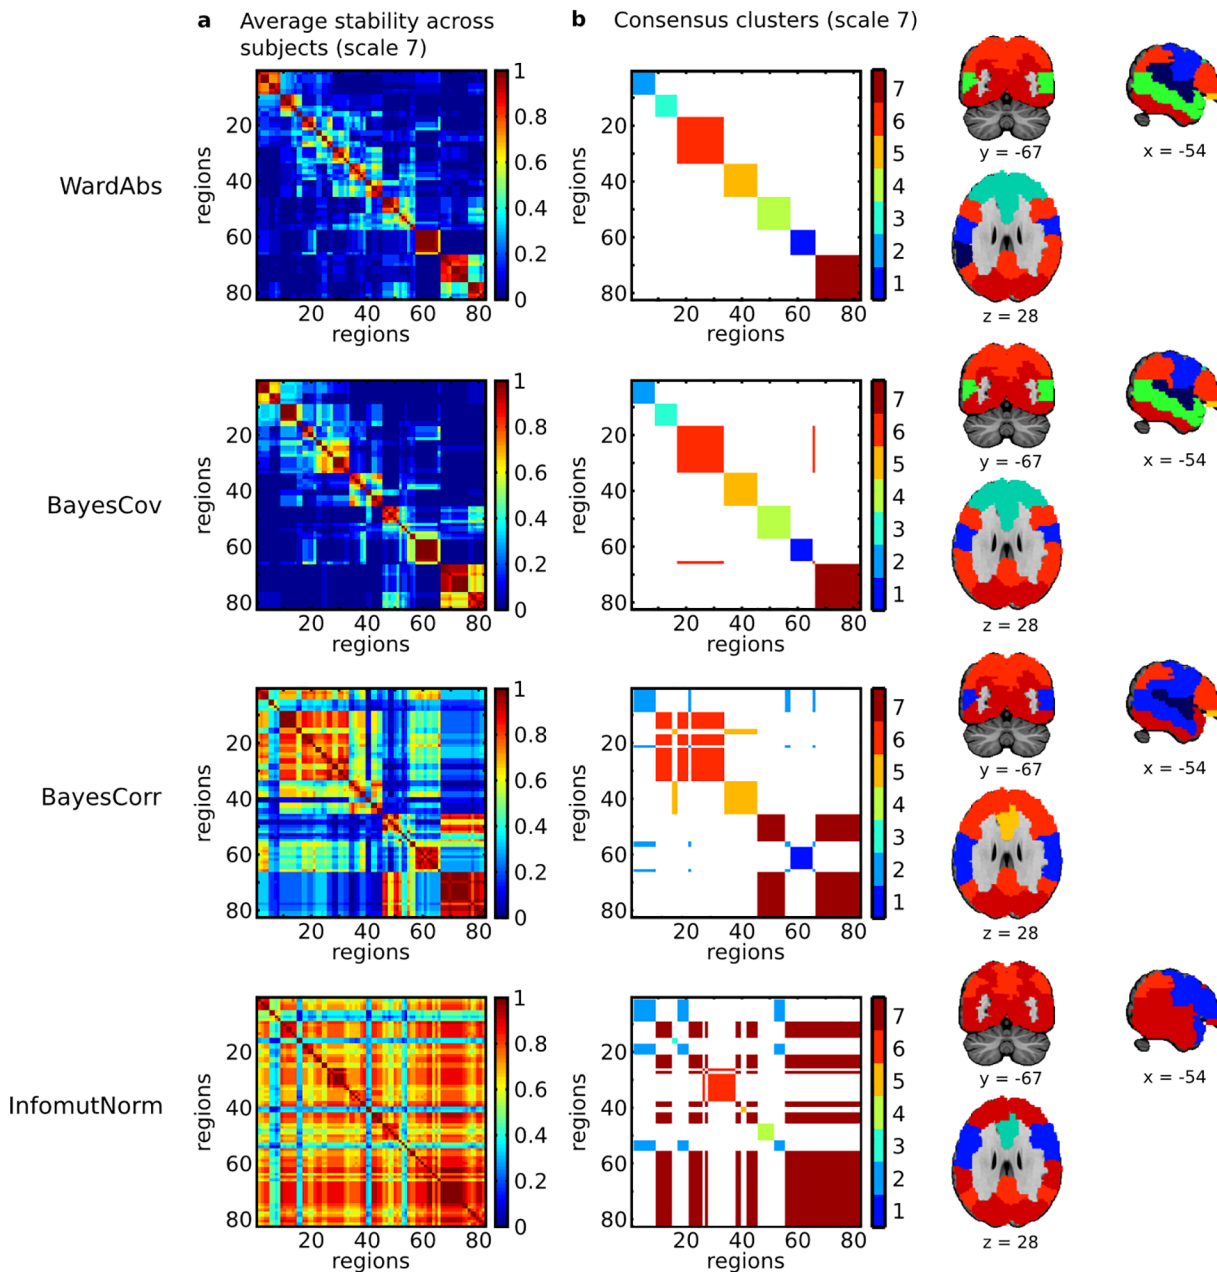


Fig 7. Real data—consensus clustering. A consensus clustering was generated based on the average adjacency matrices across all subjects (Panel a). The (weighted) adjacency matrix associated with the consensus clustering is represented along with a volumetric brain parcellation (Panel b). The weights in the adjacency matrix were added to establish a visual correspondence with the volumetric representation. Note that the brain regions have been order based on the hierarchical clustering generated with WardAbs.

doi:10.1371/journal.pone.0137278.g007

terms of cluster size. These decompositions into very large clusters do not fit current views on the organization of resting-state networks.

Overall, our analysis on real fMRI data led to the following conclusions. Three big subsets of methods emerged: spectral methods, mutual information (with SingleAbs), and finally all the other methods. Application of this last group of methods, which included the Bayes variants proposed here, resulted in a plausible decomposition into resting-state networks. In the

Table 3. Real resting-state fMRI data—computational cost. Time required by each method to cluster one dataset.

method	minimum	median	maximum
CompleteAbs	10.9 ms	11.5 ms	24.5 ms
AverageAbs	11.0 ms	11.6 ms	25.2 ms
SingleAbs	11.1 ms	11.7 ms	49.5 ms
WardAbs	14.1 ms	14.8 ms	18.9 ms
InfomutNorm	159 ms	170 ms	261 ms
InfoMut	299 ms	322 ms	416 ms
Bic	666 ms	673 ms	810 ms
BayesCov	1.083 s	1.094 s	1.263 s
BayesCorr	1.108 s	1.151 s	1.133 s
SpecCorrparAbs-1*	1.928 s	2.065 s	2.291 s
SpecCorrAbs-1*	1.992 s	2.176 s	2.417 s
SpecCorrparAbs-10*	13.251 s	13.939 s	14.301 s
SpecCorrAbs-10*	14.456 s	15.381 s	16.101 s

* For nonhierarchical methods, we summed the times used to perform clustering at each scale.

doi:10.1371/journal.pone.0137278.t003

absence of ground truth, it is not possible to further comment on the relevance of the differences in cluster solutions identified by the three groups of methods. We still noted that the methods based on mutual information led to large clusters that were difficult to interpret. Our interpretation is that the strategies implemented in Infomut and InfomutNorm did not behave well for these datasets.

As a final computational note, the time required by the different methods to cluster data is summarized in Table 3. Although the differences in execution time may reflect the quality of the implementation, the methods proposed here were the slowest of the hierarchical methods, but were still faster than spectral clustering.

Discussion

Contributions

Summary. We here proposed some novel similarity measures well suited for the agglomerative hierarchical clustering of dependent variables. These measures rely on a Bayesian model comparison for multivariate normal random variables. On synthetic data with a known (ground truth) partition, hierarchical clustering based on the Bayesian measures was found to outperform several standard clustering procedures in terms of adjusted Rand index and classification accuracy. On the toy example, the Bayesian approaches led to result similar to those of mutual information clustering techniques, with the advantage of an automated thresholding. On the real fMRI data, the Bayesian measures led to results consistent with standard clustering methods, in contrast to methods based on mutual information, which exhibited a highly atypical behavior.

Bayesian normalization of mutual information. A key feature of the Bayesian approach is its ability to take the dimension of the clusters into account. Dimensionality is an important issue in two respects (see §A of S1 File for an illustration). First, mutual information is an

extensive measure that depends on the dimension of the variables. This has motivated the introduction of a normalization factor in the application of mutual information to hierarchical clustering [16, 17]. A second issue is the existence of a bias in the estimation of mutual information. This bias mechanically increases with the dimensionality of the variables. Because of the two issues described above, hierarchical clustering based on mutual information will tend to cluster unrelated but large variables rather than correlated variables of lower dimensions. As demonstrated on real fMRI data, the heuristic proposed by [16] does not provide a general solution to the issue of dimensionality. Furthermore, it removes some interesting features of mutual information, such as the additivity of the clustering measure. By contrast, the Bayesian approach takes the dimensionality into account in a principled way, providing a quantitative version of Occam’s factor ([48] Chap. 20). The Bayesian normalization is additive rather than multiplicative, thus preserving the additive properties of mutual information.

From similarity measure to log Bayes factor. We defined the similarity measure $s(X_i, X_j)$ between any two pairs of variables X_i and X_j as a log Bayes factor. At each step, the pair (X_{i_0}, X_{j_0}) that had the largest similarity measure was merged. Taking into account the unique features of s as the log Bayes factor defined in Eq (1) allowed us to have access to a global measure of fit as defined in Eq (22) as well as to provide an automatic stopping rule that behaved very well on simulated data. Going from a similarity measure to a log Bayes factor has other advantages that could take the clustering proposed here even further (see below).

Practical value of the Bayes/Bic clustering in fMRI. The new alternatives to mutual information introduced in this paper (i.e., Bayes and Bic) proved useful for the analysis of resting-state fMRI. The benefits were particularly clear when compared to InfomutNorm, which tended to create large, inhomogeneous clusters. By contrast, both Bayes and Bic had a behavior similar to standard hierarchical clustering based on Pearson’s linear correlation. The possible benefits of Bayes and Bic over those canonical methods are still substantial. The mutual information first provides a multivariate measure of interaction that is well adapted to hierarchical brain decomposition [49, 50] and which has a clear interpretation in information theory [51–53]. For these reasons, the mutual information for Gaussian variables is more appealing than a simple average of pairwise correlation coefficients across clusters. Because mutual information is measured between clusters, it is natural to build the clusters themselves based on this metric. A second benefit of the proposed approach is that Bic proved to be the most stable of all tested methods in the range of 5–15 clusters on real fMRI datasets. The increase in stability over Ward’s was modest, but significant. This advantage may become even more substantial if the clustering is performed in higher dimension, i.e., with smaller areas than the AAL brain parcellation or even at the voxel level.

Similarity vs. distance. Clustering techniques are based on either a similarity measure or a distance measure. While the description of the present manuscript mostly relied on the notion of similarity, going from one concept to the other one can generally be done with minor changes. For instance, standard hierarchical procedures which rely on the minimization of a distance to perform clustering (e.g., single, complete and average linkage) can be applied to cases where closeness is quantified by a measure of similarity, simply by using the opposite of the similarity matrix as a distance matrix. Although the resulting measure may not define a mathematically valid distance, it is not required for the procedure to work. Similarly, in a Bayesian framework, switching from a similarity measure to a distance measure tantamounts to switching from

$$s(X_i, X_j) = \ln \frac{p(\mathbf{S}_{i,j} | \mathcal{M}_D)}{p(\mathbf{S}_{i,j} | \mathcal{M}_I)}$$

to

$$d(\mathbf{X}_i, \mathbf{X}_j) = -s(\mathbf{X}_i, \mathbf{X}_j) = \ln \frac{p(\mathcal{S}_{i,j} | \mathcal{M}_I)}{p(\mathcal{S}_{i,j} | \mathcal{M}_D)},$$

that is, from the log ratio of the marginal model likelihoods in favor of dependent variables to the log ratio of the marginal model likelihoods in favor of independent variables.

Modeling choices

Choice of priors. Any Bayesian analysis requires the introduction of prior distributions. In the present study, we needed the prior distribution for the covariance matrix associated to any clustering of \mathbf{X} . Our choices were guided by one assumption, one rule of consistency, and one rule of simplicity. First, our assumption was to not assume a priori any sort of dependence between covariance matrices associated to different clusters. This allowed for the decomposition of any prior as a product of independent priors, see Eqs (8) and (19). The rule of consistency was to consider that the prior for a given covariance matrix should not be contradictory at different levels of the hierarchy. This is why we set the prior distribution for the global covariance matrix Σ as an inverse-Wishart distribution and then derived the prior for any covariance matrix Σ_k as the prior distribution for Σ integrated over all parameters that do not appear in Σ_k ; using such an approach, the resulting prior turned out to be an inverse-Wishart distribution as well. Last, the rule of simplicity is the one that dictated the use of inverse-Wishart distributions as prior distributions for the covariance matrices. Such a family of priors had the twofold advantage of being closed by marginalization and allowing for closed-form expressions of the quantities of interest. An inverse-Wishart distribution is characterized by two parameters: the degrees of freedom ν and the inverse scale matrix Λ . If Σ is a D -by- D matrix, we must have $\nu > D - 1$ for the distribution to be normalized. Also, Λ must be positive definite. From there, we had two strategies. The first one was to set the degree of freedom to the lowest value that still corresponded to a well defined distribution ($\nu = D$) and a diagonal scale matrix that optimized the marginal model likelihood. An alternative approach was to work with the sample correlation matrix, set $\nu = D + 1$ and equal Λ with the D -by- D identity matrix I , since this choice corresponds to a distribution that is associated with uniform marginal distributions of the correlation coefficients [27]. While we believe that our assumption and the rule of consistency are sensible choices, we must admit that we are not quite as content with the choice of inverse-Wishart distributions for priors. The major issue with such a family of priors is that they simultaneously constrain the structure of correlation and the variances. By contrast, it seems intuitive that clustering should depend on the correlation structure only, not on the variances. As such, the prior on the variances should ideally be set independently from the correlation structure. Priors that separate variance and correlation have already been proposed [27]. Unfortunately, the use of such priors would make it impossible to provide a closed form for the marginal model likelihood. While numeric schemes could be implemented to circumvent this issue, it would render the procedure proposed much more complex and computationally burdensome. By contrast, the algorithm detailed here is rather straightforward. Also, the influence of the prior vanishes when the sample size increases. From a practical perspective, the three methods proposed here (two, BayesCov and BayesCorr, with different priors and one, Bic, not influenced by the prior) exhibited similar behaviors and still outperformed other existing methods in the simulation study. We take it as a proof of the robustness of the method to the choice of prior.

Covariance vs. precision matrix modeling. The presence of clusters of variables that are mutually independent is equivalent to having a covariance matrix Σ that is block diagonal, which is itself equivalent to having a precision (or concentration) matrix $\Upsilon = \Sigma^{-1}$ that is also

block diagonal. One could therefore solve the problem using precision matrices instead of covariance matrices. The corresponding calculations can be found in §E of [S1 File](#). The main difference between the two approaches stems from the fact that a submatrix of the inverse covariance matrix is not equal to the inverse of the corresponding submatrix of the covariance matrix, that is, $(\Lambda^{-1})_k \neq (\Lambda_k)^{-1}$, unless Λ is block diagonal; also, Wishart and inverse-Wishart distributions do not marginalize the same way. These differences are to be related to the fact that a submatrix of a covariance matrix is better estimated than the whole covariance matrix, while the same does not hold for a precision matrix. From there, we could expect the covariance-based approach to perform better than the precision-based approach, and the difference to increase with increasing D . This was confirmed on our synthetic data, where the precision-based approach behaved as well as BayesCov and BayesCorr for $D = 6$ but had worse results than these two approaches for $D \in \{10, 20, 40\}$. Besides, performance of the automated stopping rule was much more efficient with BayesCov and BayesCorr than it was with the precision-based approach. As a final note, basing the method on concentration matrices yielded a slower algorithm, arguably because of the matrix inversions that are required.

Sample covariance matrix vs. full dataset. Intuitively, the structure of dependence of a multivariate normal distribution is included in its covariance matrix. All existing algorithms that we used here do not need the full dataset but only the sample covariance (or correlation) matrix. This is why we started with a likelihood model that only considers the covariance matrix [see Eqs (3) and (7)]. Rigorously, this model is only valid for $N \geq D$; for $N < D$, one should resort to a model of the full data as being multivariate normal with a mean μ and covariance matrix Σ . Nonetheless, we kept the ‘intuitive’ approach as it has fewer hyperparameters, is easier to deal with, and leads to formulas that are simpler to interpret. Also, the resulting similarity measure [Eq (11)] is not restricted to $N \geq D$, but is well defined for $N < D$ as well. From a practical perspective, the ‘intuitive’ algorithm only requires the sample covariance matrix as input, while a full model would also require the sample mean. Finally, this simpler model exhibited good behavior on our synthetic data, even for small datasets.

Directions for future work

Computational costs. Regarding the computational cost, measures derived from mutual information or a Bayesian approach are more demanding than standard methods such as Average or Ward. The derivation of the similarity matrix and the search for the most similar pairs of clusters are the two critical operations that can be optimized to decrease computation time. It is always possible to speed up these two steps by taking advantage of the fact that the similarity matrices of two successive iterations of the algorithm have many elements in common, as all but one element of a partition at a given iteration are identical to the elements of the partition of the previous iteration. Critically, in the case of Average and Ward methods, it is in addition possible to derive the similarity matrix at every iteration only based on the similarity matrix at initialization through successive updates using the Lance-Jambu-Williams recursive formula [54]. By contrast, other measures, including BayesCov, BayesCorr, Bic, Infomut, and InfomutNorm, have to be re-evaluated independently at every step, which means that the determinant of a covariance matrix with increasing size has to be computed. Finding an update formula analogous to Lance-Jambu-Williams for clustering methods based on variants of the mutual information would substantially accelerate these algorithms.

From deterministic to stochastic clustering. Another extension would be to replace the deterministic rule of selecting the pair with largest similarity measure for merging by a probabilistic rule where the probability to cluster a given pair is given by the posterior probability of the resulting clustering.

Group analysis. A last extension that could be easily implemented in the present framework is the generalization of the method to account for E different entities (e.g., subjects in fMRI) sharing the same structure but with potentially different covariance matrices for that structure. If each entity e is characterized by a variable $\mathbf{X}^{[e]}$ and corresponding sample sum-of-square matrix $\mathbf{S}^{[e]}$, one can perform AHC on each $\mathbf{X}^{[e]}$ using $\mathbf{S}^{[e]}$ and the corresponding similarity measure $s^{[e]}$. However, with a straightforward modification of the present method, one could also perform global AHC of all E covariance matrices considered simultaneously. Assuming that the covariance matrices of the different elements are independent given the common underlying structure, then the resulting similarity measure is the sum of all individual similarity measures $s^{[e]}$'s.

Generalization to other types of distribution. Altogether, the Bayesian framework that we used provides a principled way to generalize our approach to distributions other than multivariate normal ones. Such generalization would potentially account for a wide variety of situations, such as nonlinear dependencies or discrete distributions. This would widen the scope of possible applications of the technique, e.g., genetics [4–7]. The issues related to this generalization are twofold. First, one needs a model of dependence. In the discrete case, one could think of multinomial distributions, originating from categorical, i.e., generalized Bernoulli distributions ([55] §3.4). In the continuous case, multivariate normal distributions are a first choice model beyond which it is not clear what to use. Multivariate Student- t distributions could be considered, even though our results on simulated data would tend to hint that the difference with multivariate normal distributions might not be that large. One could also consider using models where dependence is controlled independently of the marginal distributions, such as multivariate copulas [56, 57]. Another issue is the possibility to express the marginal posterior likelihood of the data given the model selected. For multivariate discrete variables, we expect it to be feasible, albeit computationally very challenging and sensitive to the type of prior distribution. For other distributions, obtaining a closed form might be out of reach. Nonetheless, the marginal posterior likelihood could be approximated using various criteria, such as the AIC [58] or variants thereof—AICc [59]; ([60] §2.3.1) or AICu ([60] §2.4.1)—, or the BIC [22], which naturally appeared in the present derivation. In any case, any approach beyond multivariate normal distributions would drastically increase the complexity of our approach, both in terms of inference and computation.

Application to truly hierarchical data. In the present manuscript, we used a hierarchical algorithm as a way to extract an underlying structure of dependence from data. Our assumption was that there was one such structure. Such an approach provided a simple and efficient clustering algorithm with an interesting connection to mutual information. However, the method as presented here is not able to deal with data that are truly organized hierarchically. Extending it to deal with such data would improve the scope of the algorithm. One way to do would be to use Dirichlet process mixtures [20, 61–65], together with a model of dependent variables.

Conclusion

In this paper, we proposed a procedure based on Bayesian model comparison to decide whether or not to merge Gaussian multivariate variables in an agglomerative hierarchical clustering procedure. The resulting similarity measure was found to be closely related to the standard mutual information, with some additional corrections for the dimensionality of the datasets. These new Bayesian alternatives to mutual information turned out to be beneficial to hierarchical clustering on simulations and real datasets alike. Because of the simplicity of its implementation, its good practical performance and the potential generalizations to other types of random variables, we believe that the approach presented here is a useful new tool in the context of hierarchical clustering.

Supporting Information

S1 File. Supporting information.
(PDF)

Acknowledgments

The authors are grateful to the members of the 1000 functional connectome consortium for publicly releasing the ‘Atlanta’ data sample, to the UNF (Unité de neuroimagerie fonctionnelle, Université de Montréal, Montréal, Qc, Canada) for providing them with computational resources, to Mathieu Desrosiers (UNF) for his expertise and support running the simulations on the grid engine. Part of this study was performed while G.M. was at the UNF.

Author Contributions

Conceived and designed the experiments: GM PB. Performed the experiments: GM PB. Analyzed the data: GM PB. Contributed reagents/materials/analysis tools: GM AM PB. Wrote the paper: GM AM PB. Conceived the theoretical model: GM.

References

1. Jain AK. Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*. 2010; 31(8):651–666. doi: [10.1016/j.patrec.2009.09.011](https://doi.org/10.1016/j.patrec.2009.09.011)
2. Duda RO, Hart PE, Stork DG. *Pattern Classification*. 2nd ed. Wiley-Interscience; 2000.
3. D’haeseleer P. How does gene expression clustering work? *Nature Biotechnology*. 2005; 23(12):1499–1501. doi: [10.1038/nbt1205-1499](https://doi.org/10.1038/nbt1205-1499) PMID: [16333293](https://pubmed.ncbi.nlm.nih.gov/16333293/)
4. Butte AJ, Kohane IS. Mutual information relevance networks: functional genomic clustering using pairwise entropy measurements. In: *Proceedings of the 5th Pacific Symposium on Biocomputing*. vol. 5; 2000. p. 415–426.
5. Zhou X, Wang X, Dougherty ER, Russ D, Suh E. Gene clustering based on clusterwise mutual information. *Journal of Computational Biology*. 2004; 11(1):147–161. doi: [10.1089/106652704773416939](https://doi.org/10.1089/106652704773416939) PMID: [15072693](https://pubmed.ncbi.nlm.nih.gov/15072693/)
6. Dawy Z, Goebel B, Hagenauer J, Andreoli C, Meitinger T, Mueller JC. Gene mapping and marker clustering using Shannon’s mutual information. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*. 2006; 3(1):47–56. doi: [10.1109/TCBB.2006.9](https://doi.org/10.1109/TCBB.2006.9) PMID: [17048392](https://pubmed.ncbi.nlm.nih.gov/17048392/)
7. Priness I, Maimon O, Ben-Gal I. Evaluation of gene-expression clustering via mutual information distance measure. *BMC Bioinformatics*. 2007; 8:111. doi: [10.1186/1471-2105-8-111](https://doi.org/10.1186/1471-2105-8-111) PMID: [17397530](https://pubmed.ncbi.nlm.nih.gov/17397530/)
8. Stausberg S, Elger CE, Lehnertz K. Hierarchical mutual information clustering for an improved classification of fMRI data. *Clinical Neurophysiology*. 2009; 120(1):e33. doi: [10.1016/j.clinph.2008.07.072](https://doi.org/10.1016/j.clinph.2008.07.072)
9. Benjaminsson S, Fransson P, Lansner A. A novel model-free data analysis technique based on clustering in a mutual information space: application to resting-state fMRI. *Frontiers in Systems Neuroscience*. 2010; 4:34. doi: [10.3389/fnsys.2010.00034](https://doi.org/10.3389/fnsys.2010.00034) PMID: [20721313](https://pubmed.ncbi.nlm.nih.gov/20721313/)
10. Kolchinsky A, van den Heuvel MP, Griffa A, Hagmann P, Rocha LM, Sporns O, et al. Multi-scale integration and predictability in resting state brain activity. *Frontiers in Neuroinformatics*. 2014; 8:66. doi: [10.3389/fninf.2014.00066](https://doi.org/10.3389/fninf.2014.00066) PMID: [25104933](https://pubmed.ncbi.nlm.nih.gov/25104933/)
11. Shannon CE. A mathematical theory of communication. *The Bell System Technical Journal*. 1948; 27:379–423, 623–656. doi: [10.1002/j.1538-7305.1948.tb01338.x](https://doi.org/10.1002/j.1538-7305.1948.tb01338.x)
12. Kullback S. *Information Theory and Statistics*. Dover, Mineola, NY; 1968.
13. Cover TM, Thomas JA. *Elements of Information Theory*. Wiley Series in Telecommunications and Signal Processing. Wiley; 1991.
14. Steuer R, Kurths J, Daub CO, Weise J, Selbig J. The mutual information: Detecting and evaluating dependencies between variables. *Bioinformatics*. 2002; 18(suppl. 2):S231–S240. doi: [10.1093/bioinformatics/18.suppl_2.S231](https://doi.org/10.1093/bioinformatics/18.suppl_2.S231) PMID: [12386007](https://pubmed.ncbi.nlm.nih.gov/12386007/)
15. Li M, Badger JJ, Chen X, Kwong S, Kearney P, Zhang H. An information-based sequence distance and its application to whole mitochondrial genome phylogeny. *Bioinformatics*. 2001; 17(2):149–154. doi: [10.1093/bioinformatics/17.2.149](https://doi.org/10.1093/bioinformatics/17.2.149) PMID: [11238070](https://pubmed.ncbi.nlm.nih.gov/11238070/)

16. Kraskov A, Stögbauer H, Andrzejak RG, Grassberger P. Hierarchical clustering using mutual information. *Europhysics Letters*. 2005; 70(2):278–284. doi: [10.1209/epl/i2004-10483-y](https://doi.org/10.1209/epl/i2004-10483-y)
17. Kraskov A, Grassberger P. MIC: Mutual information based hierarchical clustering. In: Emmert-Streib F, Dehmer M, editors. *Information Theory and Statistical Learning*. Springer; 2009. p. 101–123.
18. Scott AJ, Symons MJ. Clustering methods based on likelihood ratio criteria. *Biometrics*. 1971; 27(2):387–397. doi: [10.2307/2529003](https://doi.org/10.2307/2529003)
19. Binder DA. Approximations to Bayesian cluster analysis. *Biometrika*. 1981; 68:275–285. doi: [10.1093/biomet/68.1.275](https://doi.org/10.1093/biomet/68.1.275)
20. Heller KA, Ghahramani Z. Bayesian hierarchical clustering. *Gatsby Unit*; 2005.
21. Kass RE, Raftery AE. Bayes factors. *Journal of the American Statistical Association*. 1995; 90(430):773–795. doi: [10.1080/01621459.1995.10476572](https://doi.org/10.1080/01621459.1995.10476572)
22. Schwarz G. Estimating the dimension of a model. *The Annals of Statistics*. 1978; 6(2):461–464. doi: [10.1214/aos/1176344136](https://doi.org/10.1214/aos/1176344136)
23. Anderson TW. *An Introduction to Multivariate Statistical Analysis*. 3rd ed. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, New York; 2003.
24. Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis*. 2nd ed. Texts in Statistical Science. Chapman & Hall, London; 2004.
25. Press SJ. *Applied Multivariate Analysis. Using Bayesian and Frequentist Methods of Inference*. 2nd ed. Dover, Mineola; 2005.
26. Abramowitz M, Stegun IA, editors. *Handbook of Mathematical Functions*. No. 55 in Applied Math.. National Bureau of Standards; 1972.
27. Barnard J, McCulloch R, Meng XL. Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*. 2000; 10(4):1281–1311.
28. Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*. 2008; 9(3):432–441. doi: [10.1093/biostatistics/kxm045](https://doi.org/10.1093/biostatistics/kxm045) PMID: [18079126](https://pubmed.ncbi.nlm.nih.gov/18079126/)
29. Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, et al. Network modeling methods for fMRI. *NeuroImage*. 2011; 54(2):875–891. doi: [10.1016/j.neuroimage.2010.08.063](https://doi.org/10.1016/j.neuroimage.2010.08.063) PMID: [20817103](https://pubmed.ncbi.nlm.nih.gov/20817103/)
30. Lian H. Shrinkage tuning parameter selection in precision matrices estimation. *Journal of Statistical Planning and Inference*. 2011; 141(8):2839–2848. doi: [10.1016/j.jspi.2011.03.008](https://doi.org/10.1016/j.jspi.2011.03.008)
31. von Luxburg U. A tutorial on spectral clustering. *Max-Planck-Institut für biologische Kybernetik*; 2006. TR-149.
32. Nijenhuis A, Wilf H. *Combinatorial Algorithms for Computers and Calculators*. 2nd ed. Academic Press, Orlando, FL, USA; 1978.
33. Wilf HS. *East Side, West Side*; 1999. Available from: <http://www.math.upenn.edu/~wilf/lecnotes.html>
34. Bellec P, Lavoie-Courchesne S, Dickinson P, Lerch JP, Zijdenbos AP, Evans AC. The pipeline system for Octave and Matlab (PSOM): a lightweight scripting framework and execution engine for scientific workflows. *Frontiers in Neuroinformatics*. 2012; 6:7. doi: [10.3389/fninf.2012.00007](https://doi.org/10.3389/fninf.2012.00007) PMID: [22493575](https://pubmed.ncbi.nlm.nih.gov/22493575/)
35. Rand WM. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*. 1971; 66(336):846–850. doi: [10.1080/01621459.1971.10482356](https://doi.org/10.1080/01621459.1971.10482356)
36. Hubert L, Arabie P. Comparing partitions. *Journal of Classification*. 1985; 2(1):193–218. doi: [10.1007/BF01908075](https://doi.org/10.1007/BF01908075)
37. Roverato A. Asymptotic prior to posterior analysis for graphical Gaussian models. In: Vichi M, Opitz O, editors. *Classification and Data Analysis*. Springer; 1999. p. 335–342.
38. Marrelec G, Benali H. Asymptotic Bayesian structure learning using graph supports for Gaussian graphical models. *Journal of Multivariate Analysis*. 2006; 97:1451–1466. doi: [10.1016/j.jmva.2005.08.008](https://doi.org/10.1016/j.jmva.2005.08.008)
39. Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*. 2011; 106:1125–1165. doi: [10.1152/jn.00338.2011](https://doi.org/10.1152/jn.00338.2011) PMID: [21653723](https://pubmed.ncbi.nlm.nih.gov/21653723/)
40. Kelly C, Toro R, Di Martino A, Cox CL, Bellec P, Castellanos FX, et al. A convergent functional architecture of the insula emerges across imaging modalities. *NeuroImage*. 2012; 61(4):1129–1142. doi: [10.1016/j.neuroimage.2012.03.021](https://doi.org/10.1016/j.neuroimage.2012.03.021) PMID: [22440648](https://pubmed.ncbi.nlm.nih.gov/22440648/)
41. Bellec P, Rosa-Neto P, Lyttelton OC, Benali H, Evans AC. Multi-level bootstrap analysis of stable clusters in resting-state fMRI. *NeuroImage*. 2010; 51(3):1126–1139. doi: [10.1016/j.neuroimage.2010.02.082](https://doi.org/10.1016/j.neuroimage.2010.02.082) PMID: [20226257](https://pubmed.ncbi.nlm.nih.gov/20226257/)

42. Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, et al. Functional network organization of the human brain. *Neuron*. 2011; 72(4):665–678. doi: [10.1016/j.neuron.2011.09.006](https://doi.org/10.1016/j.neuron.2011.09.006) PMID: [22099467](https://pubmed.ncbi.nlm.nih.gov/22099467/)
43. Thirion B, Varoquaux G, Dohmatob E, Poline JB. Which fMRI clustering gives good brain parcellations? *Frontiers in Neuroscience*. 2014; 8:167. doi: [10.3389/fnins.2014.00167](https://doi.org/10.3389/fnins.2014.00167) PMID: [25071425](https://pubmed.ncbi.nlm.nih.gov/25071425/)
44. Orban P, Doyon J, Petrides M, Mennes M, Hoge R, Bellec P. The richness of task-evoked hemodynamic responses defines a pseudohierarchy of functionally meaningful brain networks. *Cerebral Cortex*. in press;
45. Fred ALN, Jain AK. Combining multiple clusterings using evidence accumulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2005; 27(6):835–850. doi: [10.1109/TPAMI.2005.113](https://doi.org/10.1109/TPAMI.2005.113) PMID: [15943417](https://pubmed.ncbi.nlm.nih.gov/15943417/)
46. Salvador R, Suckling J, Coleman M, Pickard JD, Menon D, Bullmore E. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex*. 2005; 34(4):387–413.
47. van den Heuvel M, Mandl R, Hulshoff Pol H. Normalized cut group clustering of resting-state fMRI data. *PLoS ONE*. 2008; 3(4):e2001. doi: [10.1371/journal.pone.0002001](https://doi.org/10.1371/journal.pone.0002001) PMID: [18431486](https://pubmed.ncbi.nlm.nih.gov/18431486/)
48. Jaynes ET. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge; 2003.
49. Tononi G, Sporns O, Edelman GM. A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences of the USA*. 1994; 91(11):5033–5037. doi: [10.1073/pnas.91.11.5033](https://doi.org/10.1073/pnas.91.11.5033) PMID: [8197179](https://pubmed.ncbi.nlm.nih.gov/8197179/)
50. Marrelec G, Bellec P, Krainik A, Duffau H, Pélégrini-Issac M, Lehericy S, et al. Regions, systems, and the brain: hierarchical measures of functional integration in fMRI. *Medical Image Analysis*. 2008; 12(4):484–496. doi: [10.1016/j.media.2008.02.002](https://doi.org/10.1016/j.media.2008.02.002) PMID: [18396441](https://pubmed.ncbi.nlm.nih.gov/18396441/)
51. Watanabe S. Information theoretical analysis of multivariate correlation. *IBM Journal of Research and Development*. 1960; 4(1):66–82. doi: [10.1147/rd.41.0066](https://doi.org/10.1147/rd.41.0066)
52. Joe H. Relative entropy measures of multivariate dependence. *Journal of the American Statistical Association*. 1989; 84:157–164. doi: [10.1080/01621459.1989.10478751](https://doi.org/10.1080/01621459.1989.10478751)
53. Studený M, Vejnarová J. The multiinformation function as a tool for measuring stochastic dependence. In: Jordan MI, editor. *Proceedings of the NATO Advanced Study Institute on Learning in Graphical Models*; 1998. p. 261–298.
54. Batagelj V. Generalized Ward and related clustering problems. In: Bock HH, editor. *Classification and Related Methods of Data Analysis*. North-Holland, Amsterdam; 1988. p. 67–74.
55. Papoulis A. *Probability, Random Variables, and Stochastic Processes*. International student edition ed. McGraw-Hill Series in Systems Science. McGraw-Hill Kogakusha, Tokyo; 1965.
56. Nelsen RB. *An Introduction to Copulas*. Springer, New York; 1999.
57. Fischer M. Multivariate copulae. In: Kurowicka D, Joe H, editors. *Dependence Modeling: Vine Copula Handbook*. World Scientific; 2011. p. 19–36.
58. Akaike H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*. 1974; 19(6):716–723. doi: [10.1109/TAC.1974.1100705](https://doi.org/10.1109/TAC.1974.1100705)
59. Burnham KP, Anderson DR. Multimodel inference: understanding AIC and BIC in model selection. In: *Amsterdam Workshop on Model Selection*; 2004. p. 261–304.
60. McQuarrie ADR, Tsai CL. *Regression and Time Series Model Selection*. World Scientific; 1998.
61. Heller KA. *Efficient Bayesian methods for clustering*. Gatsby Computational Neuroscience Unit, University College London; 2007.
62. Savage RS, Heller K, Xu Y, Ghahramani Z, Truman WM, Grant M, et al. R/BHC: fast Bayesian hierarchical clustering for microarray data. *BMC Bioinformatics*. 2009; 10:242. doi: [10.1186/1471-2105-10-242](https://doi.org/10.1186/1471-2105-10-242) PMID: [19660130](https://pubmed.ncbi.nlm.nih.gov/19660130/)
63. Cooke EJ, Savage RS, Kirk PDW, Darkins R, Wild DL. Bayesian hierarchical clustering for microarray time series data with replicates and outlier measurements. *BMC Bioinformatics*. 2011; 12:399. doi: [10.1186/1471-2105-12-399](https://doi.org/10.1186/1471-2105-12-399) PMID: [21995452](https://pubmed.ncbi.nlm.nih.gov/21995452/)
64. Darkins R, Cooke EJ, Ghahramani Z, Kirk PDW, Wild DL, Savage RS. Accelerating Bayesian hierarchical clustering of time series data with a randomised algorithm. *PLoS ONE*. 2013; 8(4):e59795. doi: [10.1371/journal.pone.0059795](https://doi.org/10.1371/journal.pone.0059795) PMID: [23565168](https://pubmed.ncbi.nlm.nih.gov/23565168/)
65. Sirinukunwattana K, Savage RS, Bari MF, Snead DRJ, Rajpoot NM. Bayesian hierarchical clustering for studying cancer gene expression data with unknown statistics. *PLoS ONE*. 2013; 8(10):e75748. doi: [10.1371/journal.pone.0075748](https://doi.org/10.1371/journal.pone.0075748) PMID: [24194826](https://pubmed.ncbi.nlm.nih.gov/24194826/)