

# Bayesian Bi-clustering Methods with Applications in Computational Biology

Han Yan<sup>1\*</sup>      Jiéxing Wu<sup>2\*</sup>      Yang Li<sup>3\*†</sup>      Jun S. Liu<sup>1</sup>

<sup>1</sup>Department of Statistics, Harvard University

<sup>2</sup>Google

<sup>3</sup>Citadel Securities

## Abstract

. Bi-clustering is a useful approach in analyzing biological data when observations come from heterogeneous groups and have a large number of features. We outline a general Bayesian approach in tackling bi-clustering problems in moderate to high dimensions, and propose three Bayesian bi-clustering models on categorical data, which increase in complexities in their modeling of the distributions of features across bi-clusters. Our proposed methods apply to a wide range of scenarios: from situations where data are cluster-distinguishable only among a small subset of features but masked by a large amount of noise, to situations where different groups of data are identified by different sets of features or data exhibit hierarchical structures. Through simulation studies, we show that our methods outperform existing (bi-)clustering methods in both identifying clusters and recovering feature distributional patterns across bi-clusters. We apply our methods to two genetic datasets, though the area of application of our methods is even broader.

## 1 Introduction

Over the past few decades, the biology community has witnessed an explosion of data due to technology advances. For example, DNA microarrays and the next-generation RNA sequencing technologies enable scientists to measure expression levels of tens of thousands of genes at once; and the single-cell technologies can even extend such measurements to the individual cell level. Another example is the Human Genome Project, which collected the whole genome sequence information, i.e., the information over three billion nucleotide base pairs that make up human DNA, from many thousands of individuals. Data collected in these cases are of high dimension in nature. Analyzing such data can reveal important biological functions, identify mutations that are responsible for certain types of diseases, shed light on human evolutionary history, etc. A useful first attempt is to cluster data into homogeneous groups. Traditional clustering methods, such as the hierarchical clustering-based UPGMA (Sokal, 1958) and the distance-based K-means method (MacQueen et al., 1967; Hartigan and Wong, 1979) have been successfully applied to biological data to identify groups of genes or samples that are of biological significance. However, the underlying biological mechanisms and the high dimensional nature of the data both suggest a further examination of the behavior of features while performing clustering. For example, when

---

\*These authors contributed equally to the manuscript.

†This presentation reflects the analysis and views of Yang Li. No recipient should interpret this presentation to represent the general views of Citadel Securities or its personnel. Facts, analysis, and views presented in this presentation have not been reviewed by, and may not reflect information known to, other Citadel Securities professionals.

analyzing gene expression data across cell types in the same tissue/organ, only a small fraction of the genes may be highly expressed and show different expression levels in different cell types. To cluster cells by cell types, it is advised to find out the differentially expressed genes and cluster cells only based on the behavior of those genes. Additionally, different biological functions or conditions activate different sets of genes. A gene may be highly expressed for a subset of experiments, while inactive in the rest of conditions. It is of great interest to identify samples corresponding to different biological conditions or functions and, at the same time, identify the set of genes that are active under each condition or function. This motivates the need of bi-clustering, i.e., methods that cluster both samples and features simultaneously.

## 1.1 A case study

In an effort to detect population structures encoded by human single nucleotide polymorphisms (SNPs), we downloaded the human SNP data from the international HapMap project website, which made available the SNP information for thousands of individuals across the world. A SNP is a single nucleotide base-pair variation among populations, and is the most common type of sequence variations found in human DNA. The raw data consist of 1,397 individuals from 11 populations across the world, and 1,457,897 SNPs for each individual. Acknowledging that not all SNPs are population or continent specific we take a bi-clustering approach. Before formally introducing our methods and analyzing the real data, we demonstrate using a semi-synthetic dataset created from the human genetic data introduced above potential advantages of bi-clustering over traditional clustering methods.

We randomly select 50 individuals from each of the following four populations: two populations with Africa ancestry: Yoruba in Ibadan, Nigeria (YRI) and Maasai in Kinyawa, Kenya (MKK), one population in east Asia: Han Chinese in Beijing, China (CHB), and one population from Europe: Toscans in Italy (TSI). We keep 50 original SNPs (features) and add 450 noise features (SNPs generated randomly) with each generated from a probability vector following the Dirichlet distribution  $Dir(1, 1, 1)$ . We apply to this dataset our biclustering method to be introduced in Section 3.2, the K-means method based on the first three principal components (PCs) (Jackson, 2005), and STRUCTURE (Pritchard et al., 2000), a clustering method that treats each SNP as cluster-differentiating. Clustering results are shown in Table 1. With only 50 informative SNPs and a signal to noise ratio of  $\frac{1}{9}$ , our bi-clustering method achieves the highest adjusted Rand Index (ARI, (Hubert and Arabie, 1985)) of 0.68, followed by 0.41 of STRUCTURE, and 0.40 of the PCA+Kmeans method. ARI is a measure of similarity between two data clustering results after correcting for chance. Its maximum value is 1 and occurs when two clustering results match perfectly. All three methods perform reasonably well in separating populations by continents. However our bi-clustering method is least affected by noise as fewer samples from each underlying population is misclassified into other clusters. All methods have difficulties in distinguishing the two African populations, as they are geographically and genetically close to each other. PCA+Kmeans creates two clusters for the two populations, but with both populations almost equally mixed in both clusters. STRUCTURE puts almost all individuals from both African populations into one cluster. In contrast, our bi-clustering method can still put the individuals into two clusters with each cluster dominated by one underlying population.

## 1.2 Literature on bi-clustering methods

The idea of bi-clustering was first introduced in Hartigan (1972). Cheng and Church (2000) pioneer the application of bi-clustering to study gene expression data. They define a bi-cluster as a submatrix of the expression data matrix (possibly transformed and nor-

Methods		Bi-clustering				PCA+Kmeans				STRUCTURE			
TrueLabel	EstLabel	C1	C2	C3	C4	C1	C2	C3	C4	C1	C2	C3	C4
		47	1	2	0	36	3	4	7	39	2	9	0
CHB		1	49	0	0	1	43	0	6	4	40	3	3
TSI		1	1	34	14	1	1	27	21	2	4	42	2
MKK		0	0	11	39	2	0	27	21	1	1	48	0
YRI													

Table 1: Clustering results by three methods using the semi-synthetic dataset.

malized) that have a high similarity score, where their similarity score reflects how well the submatrix can be fitted by a two-way ANOVA model. They use a greedy algorithm to search for bi-clusters that have high scores. In [Bergmann et al. \(2003\)](#), bi-clusters are termed as “transcription modules” that contain coregulated genes in a set of relevant experimental conditions. Mathematically, a transcription module contains genes whose normalized expression levels exceed a threshold under all conditions that belong to the bi-cluster. Statistical-Algorithmic Method for Bi-cluster Analysis (SAMBA) is a graph based approach proposed by [Tanay et al. \(2002\)](#). They define a bipartite graph whose two stable sets correspond to genes and conditions, with edges for significant expression changes. Bi-clusters correspond to bipartite subgraphs. They assign weights to vertex pairs so that significant bi-clusters have large weights. [Ben-Dor et al. \(2003\)](#) work with the relative ordering of gene expression values, and define bi-clusters to be order-preserving sub-matrices (OPSMs). A submatrix is order preserving if there is a permutation of its columns under which the sequence of values in every row is strictly increasing.

Another set of approaches utilizes singular value decomposition (SVD) to discover bi-clusters. A rank- $k$  data matrix  $\mathbf{Y}$  with dimensions  $n$  and  $p$  can be decomposed into the sum of  $k$  rank-1 matrices:  $\mathbf{Y} = \mathbf{U}\Lambda\mathbf{V}^T = \sum_{i=1}^k s_i \mathbf{u}_i \mathbf{v}_i^T$ , where  $\Lambda = \text{diag}(s_1, \dots, s_k)$ ,  $s_1 \geq \dots \geq s_k > 0$  are  $k$  singular values,  $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_k)$  are orthonormal left singular vectors and  $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_k)$  are orthonormal right singular vectors. [Asgarian and Greiner \(2007\)](#) use the rank-1 matrix  $s_1 \mathbf{u}_1 \mathbf{v}_1^T$  to approximate  $\mathbf{Y}$ . Applying a hinge function, a bi-cluster is extracted by taking out rows and columns whose corresponding entries in  $\mathbf{u}_1$  and  $\mathbf{v}_1$  have large absolute values. Once a bi-cluster is found, the same procedure is applied to the remaining matrix  $\mathbf{Y} - s_1 \mathbf{u}_1 \mathbf{v}_1^T$ . Instead of using a hinge function, [Lee et al. \(2010\)](#) directly impose a sparsity constraint on the top singular vector and define a bi-cluster to be rows and columns that have non-zero entries in  $\mathbf{u}_1$  and  $\mathbf{v}_1$ .

A useful feature of SVD is that if two columns of  $\mathbf{Y}$  are the same, the corresponding columns in  $\mathbf{V}^T$  must also be equal. A similar observation applies to the rows of  $\mathbf{Y}$ . This property can be used to find bi-clusters that have constant values. [Yang et al. \(2007\)](#) combine this idea with hierarchical clustering to cluster left and right singular vectors. Rows and columns in the same cluster define a constant bi-cluster of the original data matrix.

The plaid model introduced by [Lazzeroni and Owen \(2002\)](#) treats each expression value in a bi-cluster as a sum of the main effect, the gene effect and the condition effect. If a gene-condition pair is in multiple bi-clusters, its expression value is the sum of module effects from all the bi-clusters it falls into. Bi-clusters are found by minimizing mean squared error through an iterative algorithm. Inspired by the plaid model, two Bayesian bi-clustering algorithms were developed about the same time: [Gu and Liu \(2008\)](#) and [Caldas and Kaski \(2008\)](#). The main difference between the two methods is that [Gu and Liu \(2008\)](#) assume that the expression value of a gene in a bi-cluster is centered around the sum of three

terms: the bi-cluster-specific main effect, the gene effect, and the condition effect. Those genes that do not fall into any bi-cluster have their expression values following a normal distribution with mean 0. In contrast, in [Caldas and Kaski \(2008\)](#), the cluster-specific main effect term is replaced by a global main effect term that is shared by every observation regardless of whether the data point belongs to any cluster or not.

All aforementioned bi-clustering methods are designed for continuous observations. We also find a few methods that target or may be modified to target categorical data. [Neuwald et al. \(2003\)](#) introduce a Bayesian Partition with Pattern Selection model to simultaneously cluster aligned protein sequences in a super family into sub-families and discover important positions in the aligned sequences that underlie these subfamilies. [Patrikainen and Mannila \(2004\)](#) formulate a bi-clustering task on binary data through a finite mixture model and term it “subspace clustering”. Positing that the Bernoulli parameter for each feature equal either a feature-specific background value or a cluster-specific value if the feature is relevant for a bi-cluster  $k$ , they develop an EM-like algorithm to fit the model and use Chi-squared tests and Bayesian Information Criterion (BIC) ([Schwarz et al., 1978](#)) to select features for each cluster and the total number of clusters. [Hoff \(2005\)](#) puts a Bayesian spin on the model of [Patrikainen and Mannila \(2004\)](#) by imposing a Dirichlet process mixture model. A version on continuous data is discussed in [Hoff et al. \(2006\)](#). [Labiod and Nadif \(2011\)](#) propose the SpecCo algorithm to find a pre-specified  $K$  number of clusters to partition the input binary data matrix based on a generalized modularity measure.

An idea similar to bi-clustering is co-clustering or block clustering, which partitions both rows and columns of a dataset. For example, [Keribin et al. \(2015\)](#) partition both rows and columns into an unknown number of clusters. Each row-cluster and column-cluster pair form a data block. Within each block, elements are homogeneous, i.e., following the same distribution. One underlying assumption of block clustering is that all features are of the same nature, which greatly limits the applicability of block clustering.

Some non-probabilistic criterion-based algorithms for bi-clustering of categorical data have also been proposed. [Pensa et al. \(2005\)](#) analyze the Boolean context of a data matrix. SUBCAD in [Gan and Wu \(2004\)](#) finds bi-clusters iteratively by optimizing a combination of “compactness” within a bi-cluster and “separation” outside bi-clusters. [de França \(2016\)](#) uses Locality Sensitive Hashing (LSH) to find an initial set of seed clusters and then uses the InClose ([Andrews, 2011](#)) algorithm from Formal Concept Analysis to find bi-clusters.

### 1.3 Main themes of the article

We here use the term “bi-clustering” broadly and do not distinguish between subspace clustering and bi-clustering, thus including both unsupervised learning (aka clustering) with variable selection and clustering with cluster-specific feature selection. For example, BBC2 introduced in Section 3.2 assumes for each feature a feature-specific background distribution and a set of cluster-specific foreground distributions depending on whether the feature is selected by any cluster. Some Bayesian methods for clustering with variable selection have been developed in the early 2000s ([Liu et al., 2003; Neuwald et al., 2003; Tadesse et al., 2005; Raftery and Dean, 2006](#)). They are mostly model-based and aim to use a small subset of features to partition objects into different clusters. Each selected feature is usually assumed to follow the same distribution within a cluster but different distributions for different clusters. In contrast, when we allow for cluster-specific feature selection, the goal is to identify subsets of objects and features simultaneously. This strategy may give us more freedom in modeling features and greater interpretability along both dimensions.

In some studies, data from different sources collected under different conditions need to be integrated. In these cases, bi-clustering may be more appropriate given that certain biological functions are gene-specific and condition-specific. In some other studies, samples are more similar to each other: e.g., experiments conducted under the same condition and

from the same tissue. The subset of genes that are active are likely to be the same across groups. Then, one may want to perform a variable selection to identify this subset. BBC1 to be introduced in Section 3.1 falls under clustering with variable selection.

The model we develop in Section 2 shares some similarity with the multiple partitions model in [Galimberti and Soffritti \(2007\)](#) and [Marbac and Vandewalle \(2019\)](#). The multiple partitions mixture model assumes that features can be grouped into an unknown  $B$  number of independent blocks. Within each block, observations over the set of features follow a mixture model with an unknown number of mixture components. The emphasis of multi-partitions clustering is that no single partition of the objects can fully describe the behavior of the data over all features. Two samples can be in the same cluster over a set of features, but in two different clusters over another set of features. Our bi-clustering model can also be viewed as a special multi-partitions clustering model if in each feature block we view the partition at row-cluster level, but not individual sample level. Our model also restricts how row-clusters interact in each feature block. A mixture component either represents a single row-cluster or a group of row-clusters.

We propose a general framework for bi-clustering problems in the Bayesian context in Section 2. Although our framework works for various data types, we present in Section 3 detailed treatments of a few special twists aiming at different tasks in biology. Specifically, three methods with increasing complexity are introduced to tackle categorical data. The first method (BBC1) works on binary data, and performs a Bayesian clustering with variable selections. We then model a more complicated situation by allowing a feature to have cluster-specific distributions in only some of the clusters (BBC2). Compared with BBC1, BBC2 is able to identify both features that are unique to a cluster, and clusters that share certain features. We further extend BBC2 to build a hierarchy of clusters (HBBC), which is able to detect finer differences in the distributions of features among clusters compared with BBC1 and BBC2. For continuous data, we describe a general framework for detecting bi-clusters based on independent correlation matrices of the features. This enables us to integrate datasets from different sources and identify consistent patterns among datasets. Applications of the proposed methods on a few genetic and genomic datasets are presented in Section 4. We conclude with a short discussion in Section 5. More technical details are provided in the Appendix.

## 2 A general framework for Bayesian bi-clustering

### 2.1 A statistical model

Let  $\mathbf{Y}$  denote the  $n \times p$  data matrix. For consistency, we call rows of  $\mathbf{Y}$  objects, and columns features. The goal is to assign objects into clusters, and for each cluster, find a subset of features that supports the cluster assignment, i.e., features that follow cluster-specific distributions. The parameters of interest are: the total number of clusters  $K$ , the latent cluster assignment vector  $\mathbf{C} = (C_1, \dots, C_n)^T$ , and the  $K \times p$  binary feature selection matrix  $\mathbf{S} = (S_{k,j})_{k=1, j=1}^{K, p}$ . Note that  $\mathbf{C} \in \{0, 1, \dots, K\}^n$ , where 0 represents a null group hosting objects that are not assigned to any cluster and  $k > 0$  represents cluster  $k$ .  $S_{k,j} = 1$  means that feature  $j$  is included in cluster  $k$ . For each  $k > 0$ , let  $I_k = \{i : i \in \{1, \dots, n\}, C_i = k\}$  and  $J_k = \{j : j \in \{1, \dots, p\}, S_{k,j} = 1\}$ . The submatrix  $\{\mathbf{Y}_{i,j}\}_{i \in I_k, j \in J_k}$  form a bi-cluster. Conditioning on all the parameters, each element of  $\mathbf{Y}$  follows a posited parametric model  $F$ , and we use  $\boldsymbol{\theta}_{k,j}$  to denote model parameters for feature  $j$  under cluster  $k$ :

$$Y_{i,j} \mid \mathbf{C}, \mathbf{S}, \boldsymbol{\Theta}, K \sim F[\boldsymbol{\theta}_{c_i,j} \mathbb{I}(S_{c_i,j} = 1) + \boldsymbol{\theta}_{0,j} \mathbb{I}(S_{c_i,j} = 0)].$$

where  $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_{0,j}, \boldsymbol{\theta}_{1,j}, \dots, \boldsymbol{\theta}_{K,j}\}_{j=1}^p$ . This model assumes that the distribution of  $Y_{i,j}$  depends on its cluster assignment  $C_i$  and whether feature  $j$  is selected as a distinguishing

feature for cluster  $C_i$ . Parameter vectors  $\boldsymbol{\theta}_{c_i,j}$  and  $\boldsymbol{\theta}_{0,j}$  represent the cluster-specific and feature-specific distributions of feature  $j$ , respectively. Conjugate priors for elements of  $\boldsymbol{\Theta}$  are particularly convenient to work with in this setting since the posterior distributions are still from the same distribution family, and the normalizing constants can be calculated explicitly. The  $C_i$ 's are often assumed to be independently and identically distributed (*iid*) as:

$$P(C_i = 0 | K) = \gamma_0, \quad \text{and} \quad P(C_i = k | K) = (1 - \gamma_0) / K, \quad k = 1, \dots, K, \quad (1)$$

where  $\gamma_0$  is the prior probability for an object to belong to the null cluster, and can be set to 0 if we wish for a partition of all objects. As a binary variable,  $S_{k,j}$  is assumed to be independent Bernoulli *a priori*:

$$S_{k,j} | K \sim \text{Bernoulli}(\pi_S), \quad j = 1, \dots, p, \quad k = 1, \dots, K. \quad (2)$$

In (2), all the  $S_{k,j}$ 's share the same prior distribution. This is reasonable if one does not have any prior belief that a particular subset of features are more likely to be differently distributed across clusters. If one wants to incorporate prior knowledge on feature importance, one can replace  $\pi_S$  with feature specific parameters  $\pi_{S_j}$ , where  $\pi_{S_j}$  is set larger for more importance features. The prior in (2) gives rise to all feature distribution patterns across clusters. In the special case of clustering with variable selection, only two patterns are considered for each feature  $j$ :  $S_{k,j} = 0$  for all  $k$  and  $S_{k,j} = 1$  for all  $k$ . In this case  $\mathbf{S}$  degenerates to a  $p$ -dimensional binary vector.

Similar discrete-data bi-clustering problems and models have also been studied in [Neuwald et al. \(2003\)](#), [Guo \(2013\)](#), and [Wu \(2017\)](#), although in more specific ways. Our general framework here can be viewed as an extension of their frameworks.

## 2.2 Posterior inference via Gibbs sampler

Our parameters of interest are  $\mathbf{C}$ ,  $\mathbf{S}$ , and  $K$ . The dimension of  $\mathbf{S}$  depends on  $K$ . The prior distribution of  $\mathbf{C}$  also depends on  $K$ . Theoretically we can use a reversible jump MCMC algorithm ([Green, 1995](#)) for posterior computation, but it is computationally expensive with poor convergence behaviors. We therefore opt for a less principled but computationally more friendly two-step approach: for a range of  $K$  values, we make inference on  $\mathbf{C}$  and  $\mathbf{S}$  conditional on  $K$ ; and then we choose the optimal  $K$  according to a certain criterion, such as maximizing an approximation of  $P(K | \mathbf{Y})$ . In the current subsection, we focus on the inference of  $\mathbf{C}$  and  $\mathbf{S}$ , and will discuss how to find the optimal  $K$  in the next subsection.

It is straightforward to write down the joint posterior distribution of  $\mathbf{C}$  and  $\mathbf{S}$  conditioning on  $K$ :

$$\begin{aligned} P(\mathbf{C}, \mathbf{S} | \mathbf{Y}, K) &= \frac{P(\mathbf{Y} | \mathbf{C}, \mathbf{S}, K)P(\mathbf{C} | K)P(\mathbf{S} | K)}{P(\mathbf{Y} | K)} \\ &\propto P(\mathbf{Y} | \mathbf{C}, \mathbf{S}, K)P(\mathbf{C} | K)P(\mathbf{S} | K). \end{aligned} \quad (3)$$

We will mainly focus on finding the posterior modal estimates of the model parameters because of the multimodality of the posterior distribution resulting from near-nonidentifiability of typical clustering and bi-clustering models. Due to high dimensionality and the complexity of the parameters and latent variables involved in the model, it is infeasible to obtain either an analytical or a guaranteed numerical optimal solution. Instead, we can apply Markov Chain Monte Carlo (MCMC) methods to draw samples from the posterior distribution to identify an approximate mode of (3).

Conditioning on  $K$ , we run a Gibbs sampler to iteratively draw elements in  $\mathbf{C}$  and  $\mathbf{S}$  conditional on the rest. By the Bayes rule, the posterior probability of selecting feature  $j$

for cluster  $k$  is,

$$\begin{aligned} & P(S_{k,j} = 1 \mid \mathbf{C}, \mathbf{S}_{[-(k,j)]}, \mathbf{Y}, K) \\ &= \frac{\pi_S P(\mathbf{Y} \mid \mathbf{C}, S_{k,j} = 1, \mathbf{S}_{[-(k,j)]}, K)}{\pi_S P(\mathbf{Y} \mid \mathbf{C}, S_{k,j} = 1, \mathbf{S}_{[-(k,j)]}, K) + (1 - \pi_S) P(\mathbf{Y} \mid \mathbf{C}, S_{k,j} = 0, \mathbf{S}_{[-(k,j)]}, K)}, \end{aligned} \quad (4)$$

where  $\mathbf{S}_{[-(k,j)]}$  denotes the feature selection matrix with the  $(k, j)$ -entry omitted. When  $K$  is not too large, we can also calculate  $P(\mathbf{S}_j \mid \mathbf{C}, \mathbf{S}_{[-j]}, Y, K)$ , where  $\mathbf{S}_{[-j]}$  denotes the sub-matrix of  $\mathbf{S}$  with the  $j$ -th column removed, for all configurations of  $\mathbf{S}_j = (S_{1,j}, \dots, S_{K,j})^T$ , and sample from it. The conditional distribution of each element of  $\mathbf{C}$  is categorical with probability:

$$P(C_i = k \mid \mathbf{C}_{[-i]}, \mathbf{S}, \mathbf{Y}, K) \propto P(\mathbf{Y} \mid C_i = k, \mathbf{S}, \mathbf{C}_{[-i]}, K) P(C_i = k \mid K), \quad (5)$$

for  $k = 0, \dots, K$ , where  $\mathbf{C}_{[-i]} = (C_1, C_2, \dots, C_{i-1}, C_{i+1}, \dots, C_n)^T$ .

We obtain the maximum *a posteriori* (MAP) estimate of  $(\mathbf{C}, \mathbf{S})$  from the MCMC samples. For any given  $K$ , running the Gibbs sampler  $M$  times after burn-in gives us  $(\mathbf{C}^1(K), \mathbf{S}^1(K)), \dots, (\mathbf{C}^M(K), \mathbf{S}^M(K))$ . The MAP estimate is defined as

$$\begin{aligned} & (\hat{\mathbf{C}}(K), \hat{\mathbf{S}}(K)) \\ &= \arg \max_{m=1, \dots, M} P(\mathbf{C}^m(K), \mathbf{S}^m(K) \mid \mathbf{Y}, K) \\ &= \arg \max_{m=1, \dots, M} P(\mathbf{Y} \mid \mathbf{C}^m(K), \mathbf{S}^m(K), K) P(\mathbf{C}^m(K) \mid K) P(\mathbf{S}^m(K) \mid K). \end{aligned} \quad (6)$$

We sometimes can further integrate out  $\mathbf{S}$  to obtain the marginal posterior distribution of  $\mathbf{C}$  and  $K$  up to a normalizing factor, if  $K$  is not too large. This marginalization can further improve the efficiency of the Gibbs sampler as suggested in [Liu \(1994\)](#). Since the features are independent *a priori* and all entries in  $\mathbf{S}$  are binary, we have

$$\begin{aligned} P(\mathbf{Y} \mid \mathbf{C}, K) &= \sum_{\mathbf{S}} P(\mathbf{Y} \mid \mathbf{C}, \mathbf{S}, K) P(\mathbf{S} \mid K) \\ &= \prod_{j=1}^p \sum_{\mathbf{S}_j} P(\mathbf{Y}_{\cdot,j} \mid \mathbf{C}, \mathbf{S}_j, K) P(\mathbf{S}_j \mid K), \end{aligned} \quad (7)$$

where  $\mathbf{Y}_{\cdot,j} = (Y_{1,j}, Y_{2,j}, \dots, Y_{n,j})^T$ . This leads to the marginal MAP estimate of  $\mathbf{C}$ , which can also be approximated using MCMC samples:

$$\tilde{\mathbf{C}}(K) = \arg \max_{m=1, \dots, M} P(\mathbf{Y} \mid \mathbf{C}^m(K), K) P(\mathbf{C}^m(K) \mid K). \quad (8)$$

After obtaining  $\tilde{\mathbf{C}}(K)$ , we can proceed to get an estimate of  $\mathbf{S}$ :

$$\tilde{\mathbf{S}}(K) = \arg \max_{m=1, \dots, M} P(\mathbf{Y} \mid \mathbf{S}^m(K), \tilde{\mathbf{C}}(K), K) P(\mathbf{S}^m(K) \mid K), \quad (9)$$

or get the “exact” posterior mode of  $\mathbf{S}$  by evaluating  $P(\mathbf{Y} \mid \mathbf{S}(K), \tilde{\mathbf{C}}(K), K) P(\mathbf{S}(K) \mid K)$  for all configurations of  $\mathbf{S}(K)$ .

In our bi-clustering model, we do not assume a mixture membership for each object. Since cluster labels are of categorical nature without a natural ordering, calculating their posterior means or medians is not appropriate here. Hence we use the MAP estimate in (6) or (8) as a point estimate of the latent cluster membership. Although it may induce biases for estimating the model parameters if we fix the latent membership indicators at their MAP estimates, the biases are typically not a major concern because the posterior

distribution is mostly highly concentrated at the MAP in such cases. For example, for the human genetic data we analyzed in Section 4.1, 1,195 out of 1,198 (99.75%) latent cluster assignment variables are concentrated at their respective posterior mode throughout the sampling steps after the burn-in stage. For the single cell RNA sequence data we analyzed in Section 4.3, each  $C_i$  takes the value of its posterior mode in all MCMC steps after burn-in. These real data examples imply that the posterior probability for the latent cluster assignment variable  $C_i$  to be at its posterior mode is very close to 1. Furthermore, the parameters of primary interest to us are  $\mathbf{C}$  and  $\mathbf{S}$  but not  $\boldsymbol{\Theta}$ , and we integrate out  $\boldsymbol{\Theta}$  in our posterior computation and inference.

### 2.3 Inference of number of clusters $K$

To infer the number of clusters  $K$ , a natural approach is to choose  $K$  that maximizes the marginal posterior probability  $P(K | \mathbf{Y})$ , i.e.,

$$\hat{K} = \arg \max_K P(\mathbf{Y} | K) P(K),$$

which requires the marginalization computation:

$$P(\mathbf{Y} | K) = \sum_{\mathbf{C}} \sum_{\mathbf{S}} P(\mathbf{Y} | \mathbf{C}, \mathbf{S}, K) P(\mathbf{C} | K) P(\mathbf{S} | K). \quad (10)$$

Equation (10) generally does not have an analytical solution. We can approximate it using Chib (1995)'s method, which will be briefly reviewed below, with its specific application to our task deferred to Section 3.2.

Using  $\boldsymbol{\theta}_K$  to denote the set of parameters in the model, whose dimension may be related to  $K$ , and letting  $\mathbf{Z}$  be the latent variable (or missing data) in the model, we can write  $P(\mathbf{Y} | K)$  as

$$P(\mathbf{Y} | K) = \frac{P(\mathbf{Y} | \boldsymbol{\theta}_K, K) P(\boldsymbol{\theta}_K | K)}{P(\boldsymbol{\theta}_K | \mathbf{Y}, K)}. \quad (11)$$

where  $P(\boldsymbol{\theta}_K | K)$  is the prior density of  $\boldsymbol{\theta}_K$ , and  $P(\mathbf{Y} | \boldsymbol{\theta}_K, K) = \int P(\mathbf{Y}, \mathbf{Z} | \boldsymbol{\theta}_K, K) d\mathbf{Z}$ . Equation (11) holds for any  $\boldsymbol{\theta}_K$ . However, to estimate  $P(\boldsymbol{\theta}_K | \mathbf{Y}, K)$  in the denominator accurately, it is advised to choose a  $\boldsymbol{\theta}_K$  with high posterior density, such as the posterior mean or mode computed from the Gibbs outputs. We denote the selected  $\boldsymbol{\theta}_K$  value as  $\boldsymbol{\theta}_K^*$ .

Suppose Gibbs sampling can be applied to sample  $\boldsymbol{\theta}_K$  and  $\mathbf{Z}$  from their conditional densities:  $P(\boldsymbol{\theta}_K | \mathbf{Y}, \mathbf{Z}, K)$  and  $P(\mathbf{Z} | \mathbf{Y}, \boldsymbol{\theta}_K, K)$ . Let  $\{\boldsymbol{\theta}_K^{(m)}, \mathbf{Z}^{(m)}\}_{m=1}^M$  be  $M$  such samples. We obtain the approximation

$$\hat{P}(\boldsymbol{\theta}_K^* | \mathbf{Y}, K) = \frac{1}{M} \sum_{m=1}^M P(\boldsymbol{\theta}_K^* | \mathbf{Z}^{(m)}, \mathbf{Y}, K). \quad (12)$$

Note that the likelihood function in the numerator of Equation (11) can be written as

$$P(\mathbf{Y} | \boldsymbol{\theta}_K, K) = \frac{P(\mathbf{Y}, \mathbf{Z} | \boldsymbol{\theta}_K, K)}{P(\mathbf{Z} | \mathbf{Y}, \boldsymbol{\theta}_K, K)}.$$

In many examples, such as when the missing data distribution  $P(\mathbf{Z} | \mathbf{Y}, \boldsymbol{\theta}_K, K)$  and the complete-data likelihood  $P(\mathbf{Y}, \mathbf{Z} | \boldsymbol{\theta}_K, K)$  are analytically available, the numerator of Equation (11) can be evaluated exactly and, thus, (11) can be approximated by using  $\hat{P}(\boldsymbol{\theta}_K^* | \mathbf{Y}, K)$  in the denominator. When  $P(\mathbf{Z} | \mathbf{Y}, \boldsymbol{\theta}_K, K)$  is not available analytically, we will need to find some way to approximate it (Chib, 1995).

In practice, computing marginal likelihoods for all possible values that  $K$  can take, i.e., from 1 to  $N$ , and choose the best among them is computationally demanding and

not realistic. Narrowing down the potential values to a smaller range is the first step. To achieve this, we can first run HBBC, a model proposed in Section 3.3 that can automatically infer the number of clusters, denoted as  $\hat{K}_{HBBC}$ . With  $\hat{K}_{HBBC}$ , we can run other bi-clustering algorithms proposed in Section 3 for a sequence of consecutive  $K$  values around  $\hat{K}_{HBBC}$ , and select the best model using the marginal likelihood as the criterion. Without resorting to HBBC, assuming we know *a priori*  $K \in [K_{min}, K_{max}]$ , we can first compute the approximate posterior  $P(K | \mathbf{Y})$  for a small candidate set of ordered  $K$  values (possibly equally spaced) in between  $K_{min}$  and  $K_{max}$ . Identify the value  $K = K_1$  whose approximate  $P(K | \mathbf{Y})$  is the largest. Let  $K_2$  and  $K_3$  be the two values in the candidate set that are immediately before and after  $K_1$ . Replace  $K_{min}$  by  $K_2$  and  $K_{max}$  by  $K_3$ . Repeat the same procedure till we narrow down the search to a small, manageable set of candidate values. In many examples we encounter,  $P(K | \mathbf{Y})$  increases with  $K$  first, reaching its maximum and then decreases. Hence the above approach ensures that the optimal  $K$  is always inside the narrowed searching interval. If there are concerns that  $P(K | \mathbf{Y})$  is not unimodal, We may also need to consider the neighborhood of those  $K$ s whose  $P(K | \mathbf{Y})$ s are not significantly smaller than the maximum  $P(K | \mathbf{Y})$ . In practice, we may take  $K_{min} = 1$  and  $K_{max} = [n/20]$ , for example.

### 3 Bayesian bi-clustering with various data types

Besides the massive amount of gene expression data, there is also an abundance of discrete data available in genetic and genomic studies. For example, the Human Genome Project has made available complete or partial DNA sequences and genetic variations of many individuals of different population origins, making it possible to study genetic bases of various diseases and the human migration history. Protein sequences and structural data for multiple species available from the Protein Data Bank make it possible to analyze similarities and differences among proteins in a superfamily, revealing insights on their functionalities. In this section we introduce three Bayesian bi-clustering methods with increasing complexity that target categorical datasets, and a novel bi-clustering model for data integration.

#### 3.1 Bayesian clustering with variable selection for binary data

Binary data matrices arise in many applications and are also ideal prototypes for developing relevant theories. For example, in text analyses, columns correspond to words (or concepts) and rows correspond to articles. Bi-clustering can provide a complementary view to the classic analysis based on topic models (Blei et al., 2003). Such data can also be generated from single-cell RNA sequencing technologies, in which a “1” means that the gene is highly expressed in the cell, and “0” otherwise. Since experimented single cells are often from the same tissue, we usually expect only a small fraction of genes to show different expressed/repressed patterns among cell sub-types. The clustering problem then becomes a bi-clustering one: identifying differently expressed genes and clustering cells based on these genes.

To make our presentation more targeted, we refer to rows of  $\mathbf{Y}$  as “cells”, columns as “genes”, and selected columns as “biomarkers”. A simplifying assumption that we adopt here is that biomarkers have cluster-specific distributions across all clusters and non-biomarkers assume only feature-specific distributions. Thus, the feature selection matrix  $\mathbf{S}$  in the general framework is simplified to a  $p$ -dimensional vector,  $\mathbf{S} = (S_1, \dots, S_p)^T \in \{0, 1\}^p$ , with  $S_j = 1$  indicating that gene  $j$  is selected as a biomarker, and 0 otherwise. Biomarkers are assumed to have different probabilities of being in the expressed state in each cluster, while each non-biomarker follows its feature-specific background distribution across all

clusters. The model, termed as ‘‘BBC1’’, is thus:

$$\text{BBC1: } Y_{i,j} \mid \mathbf{C}, \mathbf{S}, \boldsymbol{\Theta}, K \sim \text{Bernoulli}(\boldsymbol{\theta}_{c_i,j} \mathbb{I}(S_j = 1) + \boldsymbol{\theta}_{0,j} \mathbb{I}(S_j = 0)). \quad (13)$$

The priors for  $C_i$  and  $S_j$  follow (1) and (2) and the priors for the  $\theta$ ’s are i.i.d. Beta ( $\alpha_{\theta_1}, \alpha_{\theta_2}$ ) for biomarkers, and i.i.d. Beta ( $\alpha_{w_1}, \alpha_{w_2}$ ) for non-biomarkers. The marginal likelihood for BBC1 after integrating out  $\boldsymbol{\Theta}$  is,

$$P(\mathbf{Y} \mid \mathbf{C}, \mathbf{S}, K) = \left[ \prod_{j: S_j=0} \frac{B(\alpha_{w_1} + n_{j,1}, \alpha_{w_2} + n_{j,0})}{B(\alpha_{w_1}, \alpha_{w_2})} \right] \left[ \prod_{j: S_j=1} \prod_{k=1}^K \frac{B(\alpha_{\theta_1} + n_{k,j,1}, \alpha_{\theta_2} + n_{k,j,0})}{B(\alpha_{\theta_1}, \alpha_{\theta_2})} \right],$$

where  $n_{j,0} = \sum_i \mathbb{I}\{Y_{i,j} = 0\}$ ,  $n_{j,1} = \sum_i \mathbb{I}\{Y_{i,j} = 1\}$ ,  $n_{k,j,0} = \sum_i \mathbb{I}\{C_i = k, Y_{i,j} = 0\}$ ,  $n_{k,j,1} = \sum_i \mathbb{I}\{C_i = k, Y_{i,j} = 1\}$  and  $B(\cdot, \cdot)$  is the beta function. After incorporating the priors, we sample  $\mathbf{C}$  and  $\mathbf{S}$  from their joint posterior distribution via Gibbs sampling. To accelerate the convergence of the Gibbs sampler, we can further integrate out  $\mathbf{S}$  if either clustering is the major concern, or we want to get the marginal MAP estimate of  $\mathbf{C}$  first, and infer  $\mathbf{S}$  conditioning on it:

$$\begin{aligned} & P(\mathbf{Y} \mid \mathbf{C}, K) \\ &= \sum_{S_1 \in \{0,1\}} \cdots \sum_{S_p \in \{0,1\}} P(\mathbf{Y} \mid \mathbf{C}, \mathbf{S}, K) P(\mathbf{S} \mid K) \\ &= \prod_{j=1}^p \left[ (1 - \pi_S) \frac{B(\alpha_{w_1} + n_{j,1}, \alpha_{w_2} + n_{j,0})}{B(\alpha_{w_1}, \alpha_{w_2})} + \pi_S \prod_{k=1}^K \frac{B(\alpha_{\theta_1} + n_{k,j,1}, \alpha_{\theta_2} + n_{k,j,0})}{B(\alpha_{\theta_1}, \alpha_{\theta_2})} \right]. \end{aligned} \quad (14)$$

Combining Equation (14) with the prior in (1), we have the posterior distribution for  $\mathbf{C}$ , from which we can sample using Gibbs sampling. If a single optimal solution for  $\mathbf{C}$  is preferred, as it is easier to interpret and proceed to further experimental investigations, the sample MAP estimator (i.e., the one obtained via MCMC samples) can be used. After inferring  $\mathbf{C}$ , we proceed to derive the posterior distribution of  $\mathbf{S}$ :

$$S_j \mid \mathbf{C}, \mathbf{Y}, K \sim \text{Bernoulli} \left( \frac{\pi_S \prod_{k=1}^K \frac{B(\alpha_{\theta_1} + n_{k,j,1}, \alpha_{\theta_2} + n_{k,j,0})}{B(\alpha_{\theta_1}, \alpha_{\theta_2})}}{\pi_S \prod_{k=1}^K \frac{B(\alpha_{\theta_1} + n_{k,j,1}, \alpha_{\theta_2} + n_{k,j,0})}{B(\alpha_{\theta_1}, \alpha_{\theta_2})} + (1 - \pi_S) \frac{B(\alpha_{w_1} + n_{j,1}, \alpha_{w_2} + n_{j,0})}{B(\alpha_{w_1}, \alpha_{w_2})}} \right),$$

for  $j = 1, \dots, p$ .

To infer the number of cell sub-types,  $K$ , we select the number that maximizes the marginal likelihood  $P(\mathbf{Y} \mid K)$ , which can be written as:

$$P(\mathbf{Y} \mid K) = \frac{P(\mathbf{Y} \mid \mathbf{C}, K) P(\mathbf{C} \mid K)}{P(\mathbf{C} \mid \mathbf{Y}, K)}. \quad (15)$$

The two terms in the numerator are readily available from (1) and (14) for any realization of  $\mathbf{C}$ . The denominator can be approximated by the empirical frequency of the posterior samples at  $\mathbf{C}$  given  $K$ . To have a good approximation to the denominator, we evaluate the term at the sample posterior mode  $\mathbf{C}^*$ .

For any  $\mathbf{C}$ , a permutation of the cluster indices  $1, 2, \dots, K$  generates an equivalent bi-cluster pattern. Denoting the equivalent class as  $\tilde{\mathbf{C}}$ , we have  $P(\mathbf{C} \mid \mathbf{Y}, K) = \frac{1}{K!} P(\tilde{\mathbf{C}} \mid \mathbf{Y}, K)$ . Each equivalent pattern in  $\tilde{\mathbf{C}}$  is a local mode for the Gibbs sampler. As a result, our sampling result may serve as a representation not for the posterior distribution of  $\mathbf{C}$  but for that of  $\tilde{\mathbf{C}}$ . Let  $(\tilde{\mathbf{C}}^m)_{m=1}^M$  be outputs from the Gibbs sampler, and  $\tilde{\mathbf{C}}^*$  represents

the sample posterior mode and its equivalent patterns. We approximate  $P(\tilde{\mathbf{C}} \mid \mathbf{Y}, K)$  at  $\tilde{\mathbf{C}}^*$  by its frequency among the  $M$  samples.

The model is not strictly identifiable as permuting the labels in  $\mathbf{C}$  does not change the model likelihood. [Allman et al. \(2009\)](#) proved that finite mixture of  $K$  different Bernoulli products are “generically identifiable”, i.e., the set of non-identifiable parameters is of Lebesgue measure zero, up to label switching if the number of Bernoulli components  $p$  satisfies  $p \geq 2\lceil \log_2 K \rceil + 1$ . Due to the presence of non-biomarkers, we cannot directly apply this result here. However, it is not difficult to show directly using Theorem 4 and Corollary 5 of [Allman et al. \(2009\)](#) that if the number of biomarkers  $p_1 \equiv \sum_j S_j \geq 2\lceil \log_2 K \rceil + 1$ , BBC1 is generically identifiable. Conditioning on  $K$  and  $\mathbf{S}$ , model parameters for biomarkers are generically identifiable. For non-biomarkers, all clusters share the same set of model parameters, which are strictly identifiable. Combining the two together, BBC1 is generically identifiable, up to label switching. When both  $\mathbf{S}$  and  $\Theta$  are unknown, the model is also generically identifiable. Two different  $\mathbf{S}$  vectors give rise to two different distributions since the set of features that are independent of each other is different under different configurations of  $\mathbf{S}$ .

To test the method, we simulate datasets from model (13) with  $n = 200$ ,  $p = 1,000$ ,  $K = 5$ , and the number of biomarkers  $N_s = 10, 20, 30$ , and  $40$ . Each setting is replicated 20 times. In each dataset, we randomly pick  $N_s$  biomarkers and assign each object to one of the  $K$  clusters with equal probability. For a non-biomarker  $j$ , we draw  $\theta_{0,j} \sim \text{Beta}(1, 1)$ , and for a biomarker  $j$ , we draw  $\theta_{k,j} \sim \text{Beta}(0.2, 0.2)$ ,  $k = 1, \dots, 5$ .

We compare performances of BBC1, FABIA: factor analysis for bicluster acquisition ([Hochreiter et al., 2010](#)), and two Latent Class Analysis with variable selection approaches in [Marbac and Sedki \(2017a\)](#), termed as “VarSelLCM-BIC” and “VarSelLCM-MICL” respectively, implemented in the R package *VarSelLCM* ([Marbac and Sedki, 2017b](#)). FABIA is a popular biclustering method for continuous data. Adding in a comparison with FABIA, we demonstrate that biclustering methods that are designed for continuous data may not perform well on categorical data. Conditioning on model parameters and the number of clusters, both approaches in *VarSelLCM* assume a mixture model: biomarkers follow cluster-specific distributions and non-biomarkers follow feature-specific distributions in all clusters. “VarSelLCM-BIC” employs a non-Bayesian approach, using the EM algorithm ([Dempster et al., 1977](#)) to obtain the MLE and selecting the best model based on the BIC. “VarSelLCM-MICL” is a Bayesian approach similar to BBC1. Model selection in this case is based on maximizing  $P(\mathbf{Y}, \mathbf{C} \mid \mathbf{S}, K)$ .

We implement BBC1 with the following parameter specifications.  $\theta_{k,j} \sim \text{Beta}(1, 1)$  for  $k = 0, 1, \dots, K$ , and  $j = 1, \dots, p$ . The prior belief on the proportion of biomarkers is  $\pi_S = 0.1$ . We initialize  $\mathbf{C}$  and  $\mathbf{S}$  at random, run the MCMC chain for 900 steps, and discard the first 200 steps as burn-in. We let  $K = 2, 3, \dots, 9$  and choose the best  $K$ . The running time depends on the signal level. When the number of biomarkers is 10 or 20, it takes 4 to 6 minutes to finish analyzing one dataset. When the number of biomarkers is 30, the average running time is 30 seconds. The times further reduce to less than 15 seconds when the number of biomarkers is 40.

Results are shown in Table 2. The clustering error (CE) rate refers to the misclassification error rate after optimal alignment between the inferred and true cluster assignments. BBC1 outperforms the other methods. The false positive rate and false negative rate refer to those for features. BBC1 can always find the actual number of clusters, with no or very low clustering errors. Most biomarkers and non-biomarkers can be correctly classified. VarSelLCM-BIC almost always estimates the number of clusters to be 2. In many cases, the two estimated clusters are mixed with objects from all 5 underlying groups. The high false negative rates indicate that it cannot find biomarkers. The performance for VarSelLCM-MICL increases significantly as we increase the number of biomarkers. The

model for VarSelLCM-MICL is very similar to BBC1 except for some prior and hyper-parameter specifications. However, it employs an iterative *ad hoc* optimization strategy to estimate  $\mathbf{C}$  and  $\mathbf{S}$ . Briefly, with a fixed  $K$  it finds  $\mathbf{S}$  to maximize  $P(\mathbf{S} \mid \mathbf{Y}, K, \mathbf{C})$  conditional on  $\mathbf{C}$ , and then conditional on  $\mathbf{S}$ , it randomly selects some components of  $\mathbf{C}$  and updates them by maximizing  $P(\mathbf{Y}, \mathbf{C} \mid \mathbf{S}, K)$ . However, without a complete sweep, some elements of  $\mathbf{C}$  may seldom get updated, resulting in clustering errors even when  $N_s$  is large. When  $N_s$  is too small, the algorithm is also easily stuck in a local mode. FABIA performs the worst among all methods. Under the model setup for FABIA, all data entries that do not belong to any bi-cluster follow the same distribution. Hence we treat such data as being in a null cluster. The estimated number of clusters for FABIA is the number of bi-clusters identified plus a null cluster. Most of the time, it cannot find any bi-cluster. Its high false negative rates indicate that it misses out many signal features. Its high false positive rates implies that many noise features are wrongly picked up to infer bi-clusters.

Number of biomarkers	Method	Number of Clusters	CE Rate	ARI	False Positive Rate	False Negative Rate
Ns = 10	BBC1	5	3.45%	0.93	0.46%	12%
	VarSelLCM-BIC	2	72.53%	0.04	2.95%	78.5%
	VarSelLCM-MICL	4.3	59.4%	0.22	20%	38%
	FABIA	1.7	75.93%	0	35.85%	80%
Ns = 20	BBC1	5	0.025%	1	0.17%	2.5%
	VarSelLCM-BIC	2	71.38%	0.06	2.44%	90.75%
	VarSelLCM-MICL	3.6	31.48%	0.62	2.90%	9%
	FABIA	1.4	76.48%	0	35.34%	74.75%
Ns = 30	BBC1	5	0	1	0.14%	2.18%
	VarSelLCM-BIC	2	65.33%	0.17	2.53%	95.5%
	VarSelLCM-MICL	5.1	15.98%	0.83	1.47%	2.34%
	FABIA	1.65	76.4%	0	35.69%	75.79%
Ns = 40	BBC1	5	0	1	0.13%	2.88%
	VarSelLCM-BIC	2	64.83%	0.18	2.26%	99.88%
	VarSelLCM-MICL	5.05	11.78%	0.87	0.21%	0.89%
	FABIA	1.6	74.95%	0.01	36.18%	73.35%

Table 2: Simulation results to testing out the effectiveness of Bayesian bi-clustering based on BBC1 for binary data.

### 3.2 Bayesian bi-clustering on categorical data

The model in the previous section assumes that a feature follows either a feature-specific background distribution (the same for all clusters), or a cluster-and-feature-specific distribution, which is different for different clusters. However, in some scenarios a feature may follow cluster-specific distributions only in some clusters, and follow the feature-specific background distribution in the remaining clusters. To deal with such scenarios, we introduce BBC2, which is largely based on the framework detailed in [Wu \(2017\)](#).

Let  $\mathbf{Y}$  be an  $n \times p$  matrix with entry  $Y_{i,j}$  taking value from  $\{1, 2, \dots, L\}$ , where  $L$  is the total number of categories. The goal is to partition the  $n$  objects into an unknown  $K$  number of clusters and identify, for each cluster, a subset of features that follow cluster-specific distributions. Given the model parameters, each data entry is assumed to follow a

categorical distribution:

$$\text{BBC2: } Y_{i,j} \mid \mathbf{C}, \mathbf{S}, \boldsymbol{\Theta}, K = \begin{cases} \text{Categorical}(\boldsymbol{\theta}_{c_i,j}) & \text{if } S_{c_i,j} = 1 \\ \text{Categorical}(\boldsymbol{\theta}_{0,j}) & \text{if } S_{c_i,j} = 0. \end{cases} \quad (16)$$

We assume that  $K - 1$  follows a truncated Poisson prior with parameter  $\alpha$ :

$$P(K) \propto \frac{\alpha^{K-1}}{(K-1)!} \exp(-\alpha), \text{ for } K = 1, 2, \dots, n.$$

Conditioning on  $K$ ,  $C_i$  follows the prior distribution (1) with  $\gamma_0 = 0$  and  $S_{k,j}$  follows (2). Compared to SVD-based approaches, an attractive feature of this bi-cluster modeling approach is that its discovered features are directly interpretable.

The binary vector  $\mathbf{S}_j$  has  $2^K$  possible configurations. However, the realization  $\mathbf{S}_j = (1, 1, \dots, 1)^T$  is indistinguishable from any of the  $K$  realizations of  $\mathbf{S}_j$  that have exactly one entry being 0 and the rest being 1s, since feature  $j$  in cluster  $k$  with  $S_{k,j} = 0$  follows a distribution not shared by any other cluster. Using  $(1, 1, \dots, 1)^T$  to represent the  $K + 1$  all cluster-specific configurations, we get a total of  $2^K - K$  configurations of  $\mathbf{S}_j$  that may lead to different likelihood values. The prior for  $\mathbf{S}_j$  conditioning on  $K$  is:

$$P(\mathbf{S}_j \mid K) = \begin{cases} \pi_S^{\sum_{k=1}^K S_{k,j}} (1 - \pi_S)^{K - \sum_{k=1}^K S_{k,j}}, & \text{if } \sum_{k=1}^K S_{k,j} < K \\ \pi_S^K + K\pi_S^{K-1}(1 - \pi_S), & \text{if otherwise.} \end{cases}$$

With Dirichlet priors assigned to  $\boldsymbol{\theta}_{k,j}, \boldsymbol{\theta}_{0,j}$  in (16), i.e.,

$$\theta_{k,j}, \theta_{0,j} \sim \text{Dirichlet}(\boldsymbol{\gamma}), \quad \boldsymbol{\gamma} = \underbrace{(\gamma_1, \gamma_2, \dots, \gamma_L)}_L^T, \quad (17)$$

we integrate out  $\boldsymbol{\Theta} = \{(\boldsymbol{\theta}_{0,j})_{j=1:p}, (\boldsymbol{\theta}_{k,j})_{k=1:K, j=1:p}\}$  to obtain the posterior distribution:

$$P(\mathbf{C}, \mathbf{S}, K \mid \mathbf{Y}) \propto P(K)P(\mathbf{C} \mid K)P(\mathbf{S} \mid K) \times \frac{1}{B(\boldsymbol{\gamma})^{p + \sum_{j=1}^p \sum_{k=1}^K S_{k,j}}} \prod_{j=1}^p \left[ B \left( \sum_{k: S_{k,j}=0} \mathbf{n}_{k,j} + \boldsymbol{\gamma} \right) \cdot \prod_{k: S_{k,j}=1} B(\mathbf{n}_{k,j} + \boldsymbol{\gamma}) \right]. \quad (18)$$

where  $\mathbf{n}_{k,j} = (n_{k,j,1}, \dots, n_{k,j,L})^T$  with  $n_{k,j,l} = \sum_{i=1}^n \mathbb{I}(C_i = k, Y_{i,j} = l)$ , and  $B(x_1, \dots, x_L) = \prod_{l=1}^L \Gamma(x_l) / \Gamma(\sum_{l=1}^L x_l)$ .

Given  $K$ , the Gibbs sampler can be implemented to iteratively draw samples from the conditional posterior distributions of  $C_i$  and  $S_{k,j}$  respectively. The choice of  $K$  is based on its marginal posterior probability  $P(K \mid \mathbf{Y}) \propto P(K)P(\mathbf{Y} \mid K)$ . To approximate the marginal likelihood  $P(\mathbf{Y} \mid K)$ , we consider a variant of Chib's method by making use of the posterior distribution of  $\{\mathbf{C}, \mathbf{S}, K\}$  in Equation (18). Specifically, we make use of the identity

$$P(\mathbf{Y} \mid K) = \frac{P(\mathbf{Y} \mid \mathbf{C}, \mathbf{S}, K)P(\mathbf{C} \mid K)P(\mathbf{S} \mid K)}{P(\mathbf{C} \mid \mathbf{Y}, K)P(\mathbf{S} \mid \mathbf{Y}, \mathbf{C}, K)}. \quad (19)$$

All the terms in the numerator and  $P(\mathbf{S} \mid \mathbf{Y}, \mathbf{C}, K)$  in the denominator can be computed exactly, while  $P(\mathbf{C} \mid \mathbf{Y}, K)$  is intractable. We can write  $P(\mathbf{C} \mid \mathbf{Y}, K)$  as:

$$P(\mathbf{C} \mid \mathbf{Y}, K) = \int \int P(\mathbf{C} \mid \mathbf{Y}, \boldsymbol{\Theta}, \mathbf{S}, K)P(\boldsymbol{\Theta}, \mathbf{S} \mid \mathbf{Y}, K)d\boldsymbol{\Theta}d\mathbf{S}. \quad (20)$$

Similar to Section 3.1, we consider the equivalent class of  $(\mathbf{C}, \mathbf{S})$ , denoted as  $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$ . We have  $P(\mathbf{C} | \mathbf{Y}, K) = \frac{1}{K!} P(\tilde{\mathbf{C}} | \mathbf{Y}, K)$ . Let  $(\tilde{\mathbf{C}}^m, \tilde{\mathbf{S}}^m, \tilde{\Theta}^m)_{m=1}^M$  be outputs from the Gibbs sampler. We approximate  $P(\tilde{\mathbf{C}} | \mathbf{Y}, K)$  at point  $\mathbf{C}^*$  (including its equivalent representations) by:

$$\hat{P}(\tilde{\mathbf{C}} | \mathbf{Y}, K) = \frac{1}{M} \sum_{m=1}^M P(\mathbf{C}^* | \mathbf{Y}, \tilde{\Theta}^m, \tilde{\mathbf{S}}^m, K).$$

To demonstrate the strength of BBC2, we conduct a simulation study to compare the performances of BBC2 and a few existing methods, including STRUCTURE (Pritchard et al., 2000), which assumes that all features follow cluster-specific distributions in all clusters, the Latent Block Model (Keribin et al., 2015) (LBM), SUBCAD, and FABIA. More specifically, we generate a dataset with  $n = 300$  objects,  $p = 3000$  features,  $K = 3$ , and  $\pi_S = 0.15$ . Parameter  $\theta$  for each categorical distribution is drawn independently from  $\text{Dirichlet}(1, 1, 1)$ , and the objects are equally likely to be in any cluster. The experiment is replicated 20 times independently.

When running BBC2, we let  $\alpha = 0.05$ ,  $\pi_S = 0.1$  and  $\gamma = (1, 1)$ . We start each MCMC chain at random values of  $\mathbf{C}$  and  $\mathbf{S}$  and run for 500 steps, with the first 200 steps as burn-in. The results are presented in Table 3, showing that BBC2 can correctly estimate  $K$  and  $\mathbf{C}$  in all experiments. In terms of its ability to infer  $\mathbf{S}$ , we calculate its feature recovery accuracy, which is the proportion of  $\mathbf{S}_{js}$  that are correctly estimated. The feature recovery accuracy for BBC2 is as high as 91%. Feature recovery accuracy is not applicable to the other methods.

STRUCTURE predicts more clusters than the true number, and tends to output small clusters of sizes smaller than 10. This overestimation is likely caused by treating all features as following different distributions in all clusters, which brings in noise to the model. The comparison between BBC2 and STRUCTURE illustrates the importance of allowing feature homogeneity among some clusters by having a background distribution for each feature. SUBCAD also achieves 0 clustering error. However, it requires the user to input the number of clusters, which is usually unknown in advance. LBM performs slightly worse than BBC2, but is still highly accurate. LBM tends to overestimate the number of (row-)clusters by separating a single or a few objects from the other objects in the same underlying group. This is observed in three experiments. Computationally, to get the optimal number of (bi-)clusters, LBM requires more model comparisons than BBC2. For BBC2, We run the algorithm for a number of  $K$  values. For each  $K$ , the algorithm automatically infers the feature matrix  $\mathbf{S}$ . However, LBM requires both the number of row-clusters and the number of column-clusters as input. Each row-cluster number and column-cluster number combination leads to a different model. For more than half of the datasets, FABIA cannot find any bi-cluster; and for those cases that it can find some bi-clusters, the identified clusters are of very small sizes, hence resulting in a very high clustering error rate.

We also compare BBC2 with Hoff (2005)'s subspace clustering model on simulated binary outcome data (since the latter cannot be applied to other types of data). We run two experiments, each with 20 replicates. In the first experiment, we fix  $n = 300$ ,  $p = 3000$  and choose  $K = 2, 4, 6$ . In the second experiment, we fix  $n = 300$ ,  $K = 4$  and set  $p = 2000, 3000, 4000$ .  $\pi_S = 0.25$  for all features. Running time increases linearly with the number of features since we treat each feature independently. At  $K = 2$ , it takes 40 seconds to finish the algorithm when  $p = 2000$ . The time increases to 60 seconds when  $p = 3000$ , and 80 seconds when  $p = 4000$ . Since the total number of feature configurations is  $O(2^K)$ , the running time increases exponentially with  $K$ . With  $p = 2000$  features, the running time for  $K = 2, 5, 7$  is 40, 120, 360 seconds, respectively.

Method	Number of Clusters	CE Rate	ARI
BBC2	3	0	1
STRUCTURE	4.9	5.63%	0.93
LBM	3.25	0.43%	0.99
SUBCAD	3*	0	1
FABIA	1.55	64.1%	0

Table 3: Simulation results to compare BBC2 with other (Bi-)clustering methods.

\*For SUBCAD, we pre-specify  $K = 3$ .

Clustering results for both experiments are presented in Table 4. The feature recovery accuracy and true negative rates, i.e., specificity for BBC2 and Hoff’s method are shown in Table 5 for both experiments. For experiment 1, both BBC2 and Hoff’s method are robust to the increase of  $K$ , incurring zero clustering error. However, as shown in Table 5, the feature recovery rate and true negative rate decrease with  $K$  for both methods. Increasing the total number of clusters increases the chance of letting a feature to follow a cluster-specific distribution in any cluster, hence specificity decreases. Increasing  $K$  also increases the total number of configurations for each  $\mathbf{S}_j$  exponentially, resulting in more difficulties in fully recovering  $\mathbf{S}_j$ . However, BBC2 maintains a higher feature recovery accuracy (by 50% or more) and a higher true negative rate (by 30% or more) than Hoff’s method in all  $K$  values. When we increase the total number of features in experiment 2, performances for both methods stay unchanged with both retaining the zero clustering error rate. The feature recovery rate stays at 69% for BBC2, and 45% for Hoff’s. BBC2 also maintains a much higher true negative rate (91%) than Hoff’s method (51%).

SUBCAD also achieves zero clustering error in all settings, and it depends on knowing the true number of clusters upfront. LBM overestimates the number of clusters. The clustering error rate increases with  $K$ s. In many cases, it outputs small clusters having only a few objects. There are also cases that objects from the same underlying cluster are split into two clusters of roughly equal size, and cases that two underlying clusters are estimated to be one. These cases can significantly increase the error rate. For experiment 1, FABIA can not find any bi-cluster when  $K = 2$ , and has a much lower ARI than other methods. Its performance is also quite variable: in some datasets, it can recover the underlying clustering patterns, but in others, it incurs significant clustering errors.

Method	Vary K, Fix p	Number of Clusters	CE Rate	ARI	Vary p, Fix K	Number of Clusters	CE Rate	ARI
BBC2	K=2 (p=3000)	2	0	1	p=2000 (K=4)	4	0	1
Hoff		2	0	1		4	0	1
LBM		2.6	0.83%	0.98		4.95	1.37%	0.98
SUBCAD		2*	0	1		4*	0	1
FABIA		1	47.47%	0		3.65	20.68%	0.69
BBC2	K=4 (p = 3000)	4	0	1	p=3000 (K=4)	4	0	1
Hoff		4	0	1		4	0	1
LBM		5.45	3.15%	0.97		4.7	4.97%	0.94
SUBCAD		4*	0	1		4*	0	1
FABIA		2.55	22.47%	0.67		3.8	18.6%	0.70
BBC2	K=6 (p = 3000)	6	0	1	p=4000 (K=4)	4	0	1
Hoff		6	0	1		4	0	1
LBM		7.7	10.85%	0.89		6.2	2.97%	0.96
SUBCAD		6*	0	1		4*	0	1
FABIA		3.45	27.88%	0.65		3.4	21.2%	0.69

Table 4: Comparing BBC2 with other Bi-clustering methods for binary data under different simulation settings.

\*For SUBCAD, we input the true number of clusters for all experiments.

Simulation Settings		Change number of clusters K (p=3000)			Change number of features p (K=4)		
Method	K(p)	K=2	K=4	K=6	p=2000	p=3000	p=4000
BBC2		86.3 / 96.5	68.8 / 91.0	57.7 / 85.6	69.0 / 91.0	68.8 / 90.9	68.6 / 91.0
Hoff		50.5 / 73.3	45.3 / 52.0	33.6 / 32.9	45.5 / 51.6	45.5 / 50.8	45.2 / 51.0

Table 5: Feature recovery accuracy (RA) and true negative rate (TNR) of BBC2 and Hoff’s method. Numbers in each cell are in percentage, with the first representing the RA and the second the TNR.

### 3.3 Hierarchical Bayesian bi-clustering

Philosophically, the task of “clustering” (aka *unsupervised learning*) is a very subjective matter and can result in drastically different findings if one focuses on different things, or even different levels of details. For example, given a group of people, one may cluster them into rough ethnicity groups, e.g., Asian, Caucasian, African, etc. However, further dividing them into a finer level of ethnicity groups (e.g., Chinese, Japanese, etc) will need some professional training. Indeed, as shown in the genetics example in Section 4, BBC2 can well separate distant populations, but tends to group intra-continent sub-populations into the same cluster because it has a hard time re-focusing its attention to more refined details. This observation argues that a unified clustering model might not be able to simultaneously account for both large and subtle differences. A natural solution similar to human’s learning strategy is to do clustering recursively or hierarchically (Sokal, 1958; Johnson, 1967). As attempted in Wu (2017), we extend the classical hierarchical clustering strategy to “hierarchical Bayesian bi-clustering” (HBBC, henceforth), in which we build a

tree based on BBC2. As the tree grows, one leaf node is allowed to be further split into two child nodes, and the choice of whether and which leaf node to split are governed by a Bayesian criterion.

HBBC starts with assigning all objects to one cluster - the root node (step 0). Each subsequent step splits an existing cluster and increases the number of clusters by one until the tree stops growing. At the beginning of step  $t$ ,  $t$  clusters have been formed with corresponding data blocks  $\{\mathbf{Y}^r\}_{r=1}^t$ , a row-wise partition of the whole dataset  $\mathbf{Y}$ . We assume that the data of all clusters are mutually independent conditioning on the corresponding hierarchical structure. Denote the current hierarchical structure as model  $H_0$ , and its “descendant” structure as  $H_r$  ( $r = 1, 2, \dots, t$ ) if node-splitting takes place at leaf node  $r$ . Note that  $H_0$  corresponds to the case that the tree terminates at step  $t$ . We assign a prior probability of  $q/t$  to each  $H_r$  model, and  $1 - q$  to model  $H_0$ , where  $0 < q < 1$  is a hyper-parameter. As the tree grows deeper, it is harder to split further at any leaf node *a priori*.

To determine whether and where the tree should get split, we consider the potential bi-clustering for data block  $\mathbf{Y}^r$  at step  $t$ . Let  $K_r$  be the number of clusters it can split into. Under  $H_0$ ,  $K_r = 1$ , and under  $H_r$ ,  $K_r = 2$ . The model for  $\mathbf{Y}^r$  and the priors for  $K_r$ ,  $\mathbf{C}^r$  and  $\mathbf{S}^r$  all follow BBC2’s setup. We define  $w_r$  to be a step-scaled ratio of model posterior probabilities under  $H_r$  and  $H_0$ :

$$\begin{aligned} w_r &:= t \cdot \frac{P(H_r \mid \mathbf{Y})}{P(H_0 \mid \mathbf{Y})} = \frac{q}{1 - q} \frac{P(\mathbf{Y} \mid H_r)}{P(\mathbf{Y} \mid H_0)} = \frac{q}{1 - q} \frac{\prod_{l=1}^t P(\mathbf{Y}^l \mid H_r)}{\prod_{l=1}^t P(\mathbf{Y}^l \mid H_0)} \\ &= \frac{q}{1 - q} \frac{P(\mathbf{Y}^r \mid H_r)}{P(\mathbf{Y}^r \mid H_0)} = \frac{q}{1 - q} \frac{P(\mathbf{Y}^r \mid K_r = 2)}{P(\mathbf{Y}^r \mid K_r = 1)}, \end{aligned}$$

and  $w_0 = t$  by default.

To calculate  $w_r$ ,  $r = 1, \dots, t$ , the only difficult part is  $P(\mathbf{Y}_r \mid K_r = 2)$ , which corresponds to Equation (19). It has been dealt with in the previous subsection with a variant of Chib’s method.  $P(\mathbf{Y}_r \mid K_r = 1)$  corresponds to the case that all objects are the same cluster. Based on the model defined in (16) and (17), it can be easily computed as

$$\begin{aligned} P(\mathbf{Y}^r \mid K_r = 1) &= \int P(\mathbf{Y}^r \mid \boldsymbol{\Theta}, K_r = 1) P(\boldsymbol{\Theta} \mid K_r = 1) d\boldsymbol{\Theta} \\ &= \prod_{j=1}^p P(\mathbf{Y}_{\cdot,j}^r \mid \boldsymbol{\theta}_j, K_r = 1) P(\boldsymbol{\theta}_j \mid K_r = 1) d\boldsymbol{\theta}_j = \prod_{j=1}^p \frac{B(\mathbf{n}_{r,j} + \boldsymbol{\gamma})}{B(\boldsymbol{\gamma})}. \end{aligned}$$

At each step, we identify the largest  $w_r$ ,  $r = 1, \dots, t$ . If it is greater than  $w_0 = t$ , we split the corresponding node, and go to the next step. Otherwise, the tree stops at level  $t$ .

### 3.4 Bayesian bi-clustering for data integration

Although some Normal distribution-based Bayesian bi-clustering methods have been developed earlier as reviewed in Section 1.2, they are mostly used for analyzing “homogeneous” data types (such as gene expression data produced by the same technology from the same lab) and are not flexible enough to handle complex data structures, such as those encountered in integrative genomics tasks. For example, a typical integrative genomic problem is to combine many thousands of published mRNA expression datasets (such as those in GEO and TCGA) to gain deeper biological insights and suggest new biological discoveries. It is inappropriate to impose a simple Bayesian bi-clustering structure developed in Section 2 to the combined dataset since different labs and technologies typically induce different types of error distributions and normalization issues. Instead, one can work on the gene correlation matrices resulting from different datasets, which can be regarded as

data summaries and are more uniform and comparable with each other. A correlation matrix can also be a reasonable summary of discrete data, especially in high dimensions. The ‘‘CLIC’’ algorithm developed in [Li et al. \(2017\)](#) is based on this observation, of which we here describe a generalized formulation.

Let  $\mathbf{X}_1, \dots, \mathbf{X}_p$  denote  $p$  datasets (e.g., microarrays), each with  $N$  rows (corresponding to  $N$  genes), and let  $\mathbf{X}_d$  have  $l_d$  columns (corresponding to  $l_d$  experimental samples). For robustness, we can screen out those datasets with  $l_d < 10$ . We first convert dataset  $\mathbf{X}_d$  into a transformed  $N \times N$  gene correlation matrix,  $\mathbf{Y}_d$ , whose  $(i, j)$ -th entry is  $Y_{d,i,j} = \frac{1}{2} \log \frac{1+r_{d,i,j}}{1-r_{d,i,j}}$ , with  $r_{d,i,j}$  being the sample correlation coefficient between gene  $i$  and  $j$  in dataset  $d$ . Suppose we have a query list  $Q_n$  of  $n$  genes (a subset of the  $N$  genes), and are interested in inferring gene co-expression modules, i.e., ‘‘clusters,’’ among genes in  $Q_n$  and datasets in which these modules are active. Each co-expression module with the relevant datasets forms a bi-cluster. After identifying these bi-clusters, we may fish out more genes not in  $Q_n$  for each module based on their co-expression patterns with the inferred module in the selected datasets.

Again, we let  $\mathbf{C} = (C_1, \dots, C_n)^T$  be the vector of cluster assignment as before, with each  $C_i$  taking values in  $\{0, 1, \dots, K\}$ , with  $k = 0$  representing a null cluster. Each non-null cluster  $k$  is assumed to be active in only a subset of datasets. Let  $S_{k,d} \in \{0, 1\}$  indicates whether dataset  $d$  is selected or not for cluster  $k$ . If dataset  $d$  is selected into cluster  $k$ , genes in cluster  $k$  co-express over the set of experimental conditions in dataset  $d$ , and we assume their transformed sample correlations follow the same (cluster-specific) normal distribution with mean  $\theta_{k,d}$  and variance  $\sigma_{k,d}^2$ . If dataset  $d$  is not selected ( $S_{k,d} = 0$ ) or genes  $i$  and  $j$  are not in the same non-null cluster, the corresponding  $Y_{d,i,j}$  is normally distributed with dataset-specific mean  $\theta_{0,d}$  and variance  $\sigma_{0,d}^2$ . In summary, for  $k = 1, \dots, K$ ,

$$\begin{aligned} [Y_{d,i,j} \mid S_{k,d} = 1, C_i = C_j = k] &\sim N(\theta_{k,d}, \sigma_{k,d}^2), \\ [Y_{d,i,j} \mid S_{k,d} = 0 \text{ or } C_i \neq C_j \text{ or } C_i = C_j = 0] &\sim N(\theta_{0,d}, \sigma_{0,d}^2). \end{aligned}$$

For computational ease, we adopt the conjugate Normal-Inv-Gamma prior for  $(\theta_{k,d}, \sigma_{d,k}^2)$ . Priors for  $\mathbf{C}$  and  $S_{k,d}$  follow (1) and (2). Since the number of genes of a typical gene module is no more than a few hundreds, in most practical cases, the number of gene pairs in which both genes are within the same co-expression module is very small compared to the total number of gene pairs  $N(N - 1)/2$ , we thus fix the dataset-specific background parameters  $\theta_{0,d}$  and  $\sigma_{0,d}^2$  for two random genes not in the same module at their total-sample estimates for dataset  $d$ , i.e.,

$$\theta_{0,d} = \frac{2}{N(N - 1)} \sum_{1 \leq i < j \leq N} Y_{d,i,j}, \quad \sigma_{0,d}^2 = \frac{2}{N(N - 1)} \sum_{1 \leq i < j \leq N} (Y_{d,i,j} - \theta_{0,d})^2.$$

Furthermore, we can integrate out  $\boldsymbol{\theta}$ ,  $\boldsymbol{\sigma}$  and  $\mathbf{S}$  to obtain the marginal likelihood  $P(\mathbf{Y} \mid \mathbf{C}, K)$  and implement a Gibbs sampler on the space of  $(\mathbf{C}, K)$  (see [Appendix E](#) for details).

In the above formulation, although bi-clusters correspond to gene-dataset combinations, the actual data we work with are matrices of the transformed gene pairwise correlations obtained from different datasets. To connect back to our general framework, we can reformat the  $p \times n \times n$  array  $\mathbf{Y}$  to a  $n(n - 1)/2 \times p$  matrix  $\tilde{\mathbf{Y}}$ , where each row corresponds to a pair of query genes (recording their correlation), and each column a dataset. Then, we can map the clustering of the original set of genes to the clustering of the rows of  $\tilde{\mathbf{Y}}$ : a row of  $\tilde{\mathbf{Y}}$  is assigned to cluster  $k \in \{1, \dots, K\}$  if the two genes correspond to this row are both in cluster  $k$ , otherwise the row is assigned to cluster 0. In other words, this formulation differs from the typical bi-clustering settings in that the clustering of the rows of  $\tilde{\mathbf{Y}}$  has an additional restrictive structure, which can be reflected in its prior. The rest follows the general framework of [Section 2](#).

What we just described can be viewed as the second stage of a two-stage procedure. In the first stage we obtain a summary for each dataset in the form of a correlation matrix of the objects involved – the individual features of the objects are no longer used in the second stage although the derivation of this summary matrix relies on the features and may involve feature selections. Our bi-clustering model is designed only for the second stage, i.e., the bi-clustering of objects and datasets. Generalized Biclustering (GBC) developed in [Li et al. \(2020\)](#) is also designed to integrate information from data of different types. GBC concatenates all data matrices that share a set of common objects, retaining all features without any data reduction. It models this concatenated matrix jointly through an exponential family-based factor model and conducts bi-clustering on all the features and all the objects. While GBC is more delicate and principled, our approach is more robust and much more scalable (e.g., can be applied to combine information from thousands of datasets each with hundreds to many thousands features). In principle, our method can be applied to integrate datasets of various types, provided that information of each dataset can be approximated summarized as a correlation matrix for certain characteristics (e.g., mRNA expressions, or mutation loads) of the objects (e.g., genes) common to all datasets.

## 4 Real-data applications of Bayesian bi-clustering

### 4.1 Bi-clustering analysis of genetic data by BBC2

Let us reconsider the human genetic data introduced in Section [1.1](#). BBC2 assumes that features are independent of each other. To obtain a set of SNPs that are not in linkage disequilibrium (i.e. nearly independent), we use PLINK to remove correlated SNPs by setting the threshold for the squared correlation  $r^2$  to be  $10^{-6}$  over a sliding window of 200 variants and an increasing step of 20. We have also tried other thresholds such as  $10^{-4}$  and  $10^{-2}$ , and find the clustering result relatively robust (see Appendix [B](#) for details). All children from trio-families are dropped, resulting in unrelated individuals. The two filtering steps lead to a dataset of 4,217 SNPs and 1,198 individuals from 11 populations based on population origin. A description of the dataset can be found in Appendix [A](#).

To run BBC2, we let  $\alpha = 0.05$ ,  $\beta = 0.05$ , and  $\gamma = (1, 1, 1)$ . A small  $\alpha$  represents a prior belief that samples come from a small number of clusters. A small  $\beta$  expresses the prior belief that most features are background ones shared across clusters. The hyper-parameter of the Dirichlet prior for the genotype frequency  $\theta$  of each SNP is set at  $\gamma = (1, 1, 1)$ . The total number of MCMC steps is 600 with the first 200 steps as burn-in.

BBC2’s clustering result is shown in Table [6](#), with an ARI of 0.56. The method estimates the number of clusters to be 6. Individuals from the ASW population are split into two clusters, mixing with the other populations with African ancestry. For the rest of the populations, almost all individuals from the same population are grouped together. Populations in mixed groups are close to each other geographically, i.e., from the same continent. Out of the 4,217 SNPs, 96.42% show certain extent of diversity cross populations. Table [7](#) summarizes the number of SNPs that follow cluster-specific distributions in the specified numbers of clusters. For example, only 151 SNPs are background features following a common distribution, and only 11 SNPs are all-cluster-specific, meaning that their distributions are different in different clusters. These numbers suggest neither clustering nor clustering with feature selection is adequate in modeling the heterogeneity in the distributions of features.

When we specify the number of clusters to be 11, the true number of populations, BBC2 still outputs 6 nonempty clusters, and the cluster assignments for the 1,198 individuals are exactly the same as Table [6](#). This suggests that more refined (weak) signals do not seem to help in global clustering. To further investigate this phenomenon, we randomly select  $R\%$  of the features (SNPs) and randomly permute their observations (so that these features

are made to be noninformative). We rerun the algorithm with  $R$  ranging from 10 to 90. Clustering error rate stays almost the same when  $R$  is below 75, but increases to about 54% after that. This observation indicates that a majority of the SNPs carry some amount of information, which supports the previous finding that over 95% of the SNPs are “important”. It also indicates BBC2 is robust to the presence of many noninformative features. Among the randomized features, the algorithm correctly estimates 90-97% of them to follow background distributions, i.e., 3-10% false positive rate. More details can be found in Appendix A.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	31	22	0	0	0	0
LWK	110	0	0	0	0	0
MKK	4	152	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	53	5	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 6: BBC2’s clustering result on the human genetic dataset: rows correspond to population origins and columns correspond to estimated clusters.

Number of Cluster-Specific Distributions	0	1	2	3	4	6
Number of SNPs	151	611	1448	1553	443	11

Table 7: Numbers of SNPs (second row) that are estimated by BBC2 to have the specified numbers (first row) of cluster-specific distributions for the human genetic dataset.

## 4.2 Analyzing the genetic data with HBBC

In the previous subsection, the grouping of intra-continent sub-populations into the same cluster by BBC2 indicates that intra-continent minor variations across sub-populations have been missed out. We demonstrate here that HBBC, with its hierarchical nature, is able to pick up these weak signals. From the modeling perspective, focusing on one node at a time enables us to detect minor differences between closely related sub-populations. It can also capture the scenario that a feature has several background distributions, corresponding to different higher level groups, e.g., continents. The feature can have sub-population-specific distributions in certain continents, and continent-specific background distributions in the rest. From the evolutionary perspective, human population evolves roughly in a hierarchical fashion (the out-of-Africa theory): *Homo sapiens* developed first in Africa and then spread to other continents about 50-70,000 years ago (Maca-Meyer et al., 2001; Posth et al., 2016). Then, within a continent, sub-populations emerged at different times. Hence, sub-populations within a continent are genetically closer to each other, and are on a smaller sub-tree before joining other sub-populations in other continents to form a larger (sub)-tree.

We apply HBBC with hyper-parameter  $q = 0.05$ , reflecting a fairly conservative prior on developing new clusters. To further alleviate excessive learning, we restrict the minimum size of any node to be 50, which corresponds to around 5% of the total number of individuals. We have also tested different minimum node sizes: 5, 20, 30 and 40, and found the results to be robust (see Appendix C for details). Other parameters are set to the same values as in BBC2. The result is shown in Figure 1, with the leaf nodes colored in orange. Each blue internal node is labelled with the number of individuals assigned to it. The numbers inside each orange node are numbers of individuals by population. At the first split, populations with African ancestry are separated from the rest; at the second split, the East Asians are separated from individuals with European (CEU, TSI) or European-mixed (MEX, GIH) ancestries<sup>1</sup>, followed by the third and fifth splits that further separate populations in the latter three ancestry groups. The sixth step identifies subpopulations in the East Asian group: the JPT samples are separated from the two Chinese populations. Among the populations with African ancestry, the MKK population and a small fraction of the ASW population are the first to be separated from the rest. The second major cluster with African ancestry consists of the YIH population and another subset of the ASW population. The LWK population makes up the third major cluster. The ARI of the clustering result is equal to 0.75.

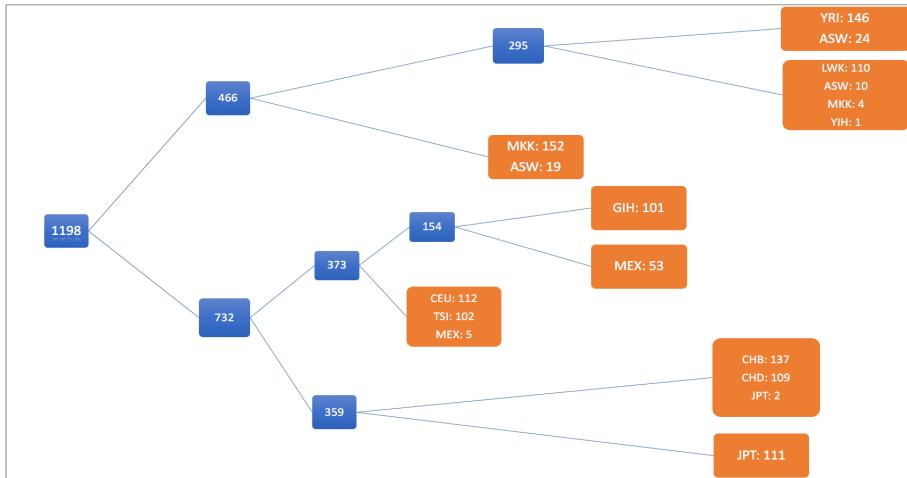


Figure 1: HBBC’s estimation result on the human genetic data, minimum node size=50. The number in each blue node is the number of individuals assigned to the node. In each orange node, i.e., a leaf node: each population name-number pair indicates for that population, the number of individuals that are assigned to this node. Reading from the left to the right of the figure gives the order of the node splitting process.

The advantage of HBBC over BBC2 has been demonstrated in this example through its ability to distinguish the Japanese from the Chinese, and the Luhya samples from the Yoruba samples. The order of the node splitting process is likely to imply the relative genetic proximities of different populations. For example, each of the first three steps separates populations by continents. While intra-continent populations are identified after step 3. In terms of selected features, the number of SNPs that have different distributions between the Chinese and the Japanese populations is estimated to be 354 (8.39%), while this number is 1,053 (24.97%) between the GIH population and the MEX population. Geographically, GIH and MEX are more distant to each other than JPT and CHB+CHD.

<sup>1</sup>Some recent studies show that there is a strong European influence on the genetic composition of the Indian population (Reich et al., 2009; Narasimhan et al., 2019). Raghavan et al. (2014) and Sans (2000) also document the genetic proximity of European populations and the native American, and Latin American populations.

Among the African populations, the MKK-dominate group is most different from the rest of the African populations as it is the first to be split out and is distinguished on 1,167 (27.67%) SNPs, while the number is only 317 (7.52%) between the YIH-dominate group and LWK-dominate group.

Benefits of identifying a set of SNPs that have different distributions in different ethnicity groups are not limited to clustering of individuals. For example, if we want to estimate polygenic risk scores of a certain disease for a group of individuals, the developed bi-clustering methods (BBC2 and HBBC) help account for hidden population structures and estimate the baseline distributions of SNPs more accurately, which may result in a better risk score estimate. The SNPs selected in each step of HBBC are also potentially related to human phenotypes that show variations among different populations. We take human hair and hair color related traits as an example. The differences in the traits are larger for populations in different ethnicity groups across continents, while it is more subtle when comparing geographically closer populations. So we should expect to see the traits-related SNPs becoming less likely to be selected in later steps of HBBC. The analysis of hair-related SNPs in Appendix D provides some support.

### 4.3 Clustering single-cell RNA sequencing data using BBC1

We apply BBC1 to the dataset in [Chu et al. \(2016\)](#). The dataset contains the expression levels of 19,097 genes in 1,018 cell samples from 7 human cell types: DE cells (endoderm derivatives), endothelial cells (EC), neuronal progenitor cells (NPC), trophoblast-like cells, single undifferentiated H1 and H9 ES cells, and foreskin fibroblasts (HFF). We binarize the data such that 1 represents a high expression level and 0 otherwise before applying BBC1. The prior probability of a gene being a biomarker is set to be  $\pi_S = 0.1$ . Flat priors are given to  $\theta$ s: i.e.,  $\alpha_{\theta_1} = \alpha_{\theta_2} = \alpha_{w_1} = \alpha_{w_2} = 1$ . The number of MCMC steps is 700 with the first 200 as burn-in. We compare BBC1’s clustering result to those produced by SC3 ([Kiselev et al., 2017](#)) and Seurat ([Stuart et al., 2019](#)). The original, un-binarized, data are used for both SC3 and Seurat. For Seurat, 20 principal components are used for its *FindNeighbors* function and the resolution is set to 0.5 for the function *FindClusters*.

The clustering results from the three algorithms are shown in Tables 8-10. SC3 gives the highest ARI of 0.93. BBC1 is the second best with an ARI of 0.74, closely followed by Seurat with an ARI of 0.73. BBC1 and Seurat have very similar clustering results, both having difficulties in separating H1 and H9 cells. Samples from these two cell types are mixed together over the first 2 principal component directions. The higher ARI of SC3 is entirely attributed to its ability to distinguish H1 and H9 cells. SC3 performs slightly worse in clustering DE cells compared to BBC1 and Seurat. Though all algorithms separate NPC into two clusters, BBC1 is able to put the majority of NPCs into one cluster, and the rest to a minor cluster. SC3 and Seurat tend to separate NPCs into two equally sized clusters. BBC1 also put all EC in one cluster (i.e., Cluster 2 in Table 8), while SC3 and Seurat both have one EC mixed with DE cells. There are 9,837 genes whose posterior probabilities of being biomarkers is greater than 0.99 according to BBC1. The number of marker genes predicted by SC3 is 8,499, and 5,548 by Seurat. BBC1 picks up 87% and 51% of the marker genes from SC3 and Seurat, respectively, as biomarkers.

Comparing the performance of BBC1 with SC3 and Seurat, we observe that binarization does not result in much loss of information, but rather reduces noise and improves robustness. In certain extremely noisy cases and/or when model assumptions are severely violated, we believe that BBC1 is more likely to cluster samples from the same cell type together than existing approaches.

Cell Type	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7
DEC	138	0	0	0	0	0	0
EC	0	105	0	0	0	0	0
H1	0	0	207	0	0	0	5
H9	0	0	162	0	0	0	0
HFF	0	0	0	0	159	0	0
NPC	0	0	0	43	0	130	0
TB	0	0	0	0	0	0	69

Table 8: BBC1’s estimation result on the scRNA-seq data. Rows correspond to true cell types, and columns correspond to estimated clusters.

Cell Type	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7
DEC	137	1	0	0	0	0	0
EC	1	104	0	0	0	0	0
H1	0	0	212	0	0	0	0
H9	0	0	162	0	0	0	0
HFF	0	0	0	0	159	0	0
NPC	0	0	0	79	0	94	0
TB	0	0	0	0	0	0	69

Table 9: Seurat’s clustering result on the scRNA-seq data. Rows correspond to true cell types, and columns correspond to estimated clusters.

Cell Type	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7	Cluster 8	Cluster 9
DEC	127	0	0	0	0	0	0	0	11
EC	0	104	0	0	0	0	0	0	1
H1	0	0	212	0	0	0	0	0	0
H9	0	0	0	162	0	0	0	0	0
HFF	0	0	0	0	159	0	0	0	0
NPC	0	0	0	0	0	89	0	84	0
TB	0	0	0	0	0	0	69	0	0

Table 10: SC3’s clustering result on the scRNA-seq data. Rows correspond to true cell types, and columns correspond to estimated clusters.

#### 4.4 Gene clustering based on co-expression patterns

To demonstrate the bi-clustering approach described in Section 3.4 for integrative genomics, we amass a collection of 1,774 mRNA datasets from GEO for the mouse Affymetrix chip Mouse430\_v2 platform (Li et al., 2017), and apply the correlation-matrix-based bi-clustering algorithm to a list of 58 genes. The ground truth, which is unknown to the algorithm, is that the genes are from two pathways: the DNA replication pathway with 22 genes, and the mitochondrial respiratory chain complex I with 36 genes.

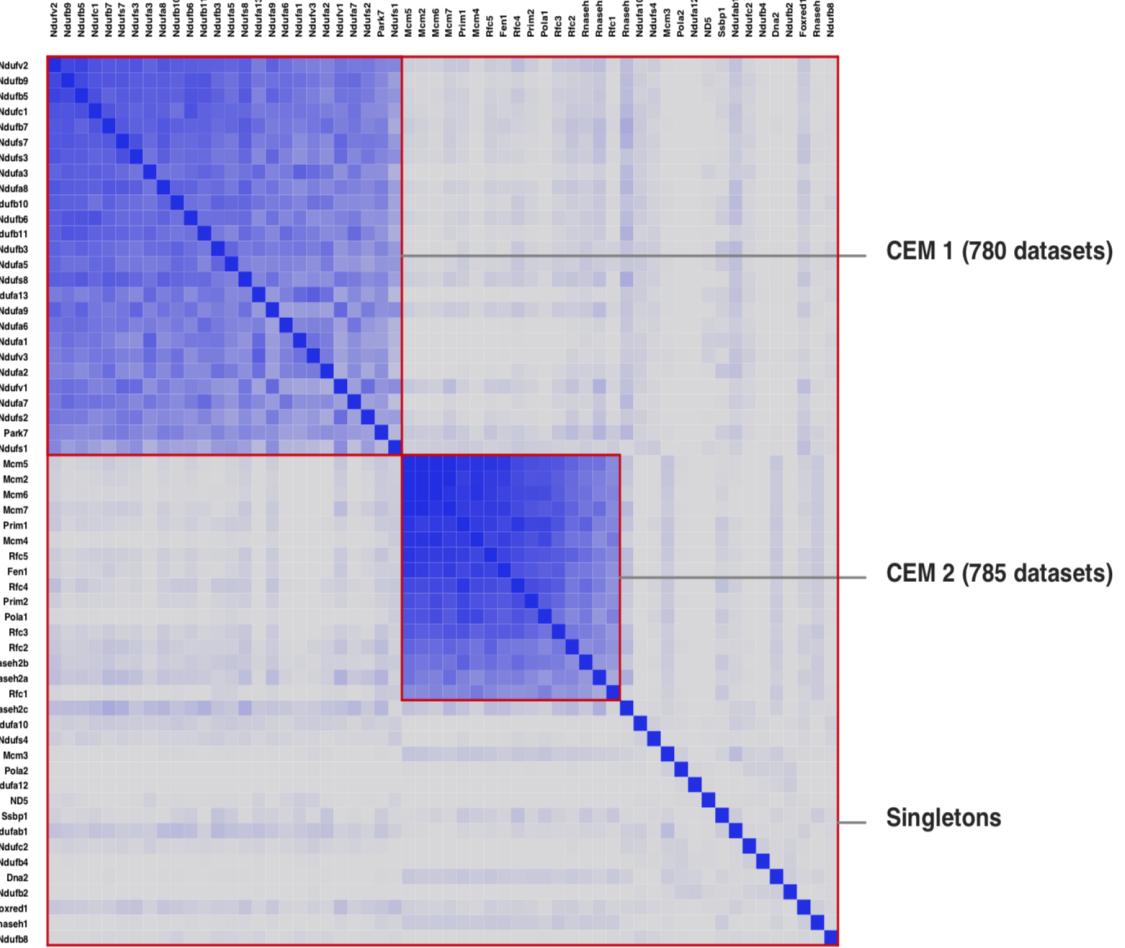


Figure 2: CLIC’s clustering result on the input two pathways. Both rows and columns correspond to the input genes. CEM1 consists of 26 genes from one pathway and is supported by 780 datasets. CEM2 consists of 16 genes from the other pathway and is supported by 785 datasets. Genes that are in neither CEM are identified as singletons.

Figure 2 displays the clustering result, indicating that the algorithm identifies two nontrivial gene co-expression clusters, CEM1 and CEM2. CEM1 consists of 26 genes from the pathway for mitochondrial respiratory chain complex I (missing 10 genes) and CEM2 consists of 16 genes from the DNA replication pathway (missing 6 genes). No gene from one pathway is assigned to the cluster of the other pathway. The 16 genes not included in the two clusters are identified as singletons.

Genes in CEM1 are estimated to be strongly co-expressed in 780 datasets, whereas those in CEM2 are supported by 785 datasets. For datasets that support both clusters, different co-expression patterns can occur in the two clusters. As an illustration, we select one such dataset with 59 experimental samples, and show its gene expression patterns in Figure 3. In the first 19 samples, genes in CEM1 are all highly expressed, while all genes in CEM2 have low expression values. The pattern is reversed in the next 21 samples. In the last 19 samples, genes in CEM1 show low expression levels, while moderate expression levels are observed among genes in CEM2. The different co-expression patterns across samples in one dataset for one cluster is due to the fact that different samples correspond to different experimental conditions. Analyzing the expression patterns in the two clusters across datasets tells us when the two pathways will be both activated and work together. Examining the selected datasets helps characterize conditions under which genes in the cluster co-express and reveals insights on biological functions of the gene clusters. For

example, among the top 10 datasets selected by CEM1, 5 of them are related to inflammatory processes. Two are related to experiments studying rheumatoid arthritis and another two are related to inflammation in liver.

## 5 Discussion

Modern high-throughput biomedical technologies have generated massive amount of data. The huge volume of data, complicated by the heterogeneity of data and the complexity of biological systems, posts increasing challenges for discovering hidden patterns and understanding biological functions. Bi-clustering is one way to deal with such challenges by identifying and clustering objects on a subset of features. The goal of bi-clustering, when compared to clustering, is not only to cluster objects, but also to pinpoint important features and model them effectively for each cluster, which reduces the impact of noise and helps unwind hidden structures in the data.

We have outlined a general Bayesian framework for tackling bi-clustering problems and described a few variants that target different types of data and attempt for different goals. For categorical data, we start with the basic model that clusters objects on a subset of features, which are assumed to have cluster-specific distributions in every cluster. The method can detect rare signals and has accurate clustering performance. This approach works well when data are homogeneous, such as cell sub-types within one tissue/organ. We then extend the model to allow for different features to be used by different clusters. In this way, we are able to identify features that are differently distributed in a subset of clusters, but each follows a common feature-specific distribution in the remaining clusters. We then apply this general approach in a hierarchical fashion to identify bi-clusters in a step-wise manner. The sequence of node-splitting and the hierarchical structure show the relative heterogeneity among data and enable us to discover cluster-level relationships. As a useful and nontrivial twist of the standard bi-clustering framework, we lastly introduce a data-integration method that imposes a Bayesian bi-clustering structure on the correlation matrices derived from multiple datasets. We show that this approach can effectively extract important gene-module information by combining many publicly available gene expression datasets.

The true number of clusters in a dataset is often unknown. In our general framework, we have discussed the inference of  $K$  based on its marginal posterior distribution. Another way of simultaneously inferring bi-cluster memberships and the total number of clusters is to employ a Dirichlet process mixture model. As  $K$  changes throughout MCMC iterations, algorithms such as reversible jump MCMC and split-merge MCMC can be used. Readers can refer to methods in Hoff (2005); Hoff et al. (2006); Kim et al. (2006) for a detailed study on bi-clustering/clustering using Dirichlet process mixture models.

The general framework can be easily extended to deal with data of mixed types, which often occur when one tries to combine data from different sources measuring different aspects of a common set of objects. The assumption of independence among features in the general framework makes the analysis straightforward. However, in many cases, feature dependency cannot and should not be ignored. For example, genes in a pathway co-express in many biological functions. We have discussed one way to identify bi-clusters based on gene expression correlations. Another way to extend the general framework is to directly impose dependency structures on the observed data. For example, Raftery and Dean (2006) discussed the design and parameterization of the covariance matrix in Gaussian model-based clustering. A potential issue with this more idealistic approach is that the model can be too complex to allow for meaningful inferential results. More research along this direction is certainly warranted.

We have implemented the three Bayesian bi-clustering algorithms in R and C++, and

it is available at <https://github.com/HanY-H/BayesianBiClustering>.

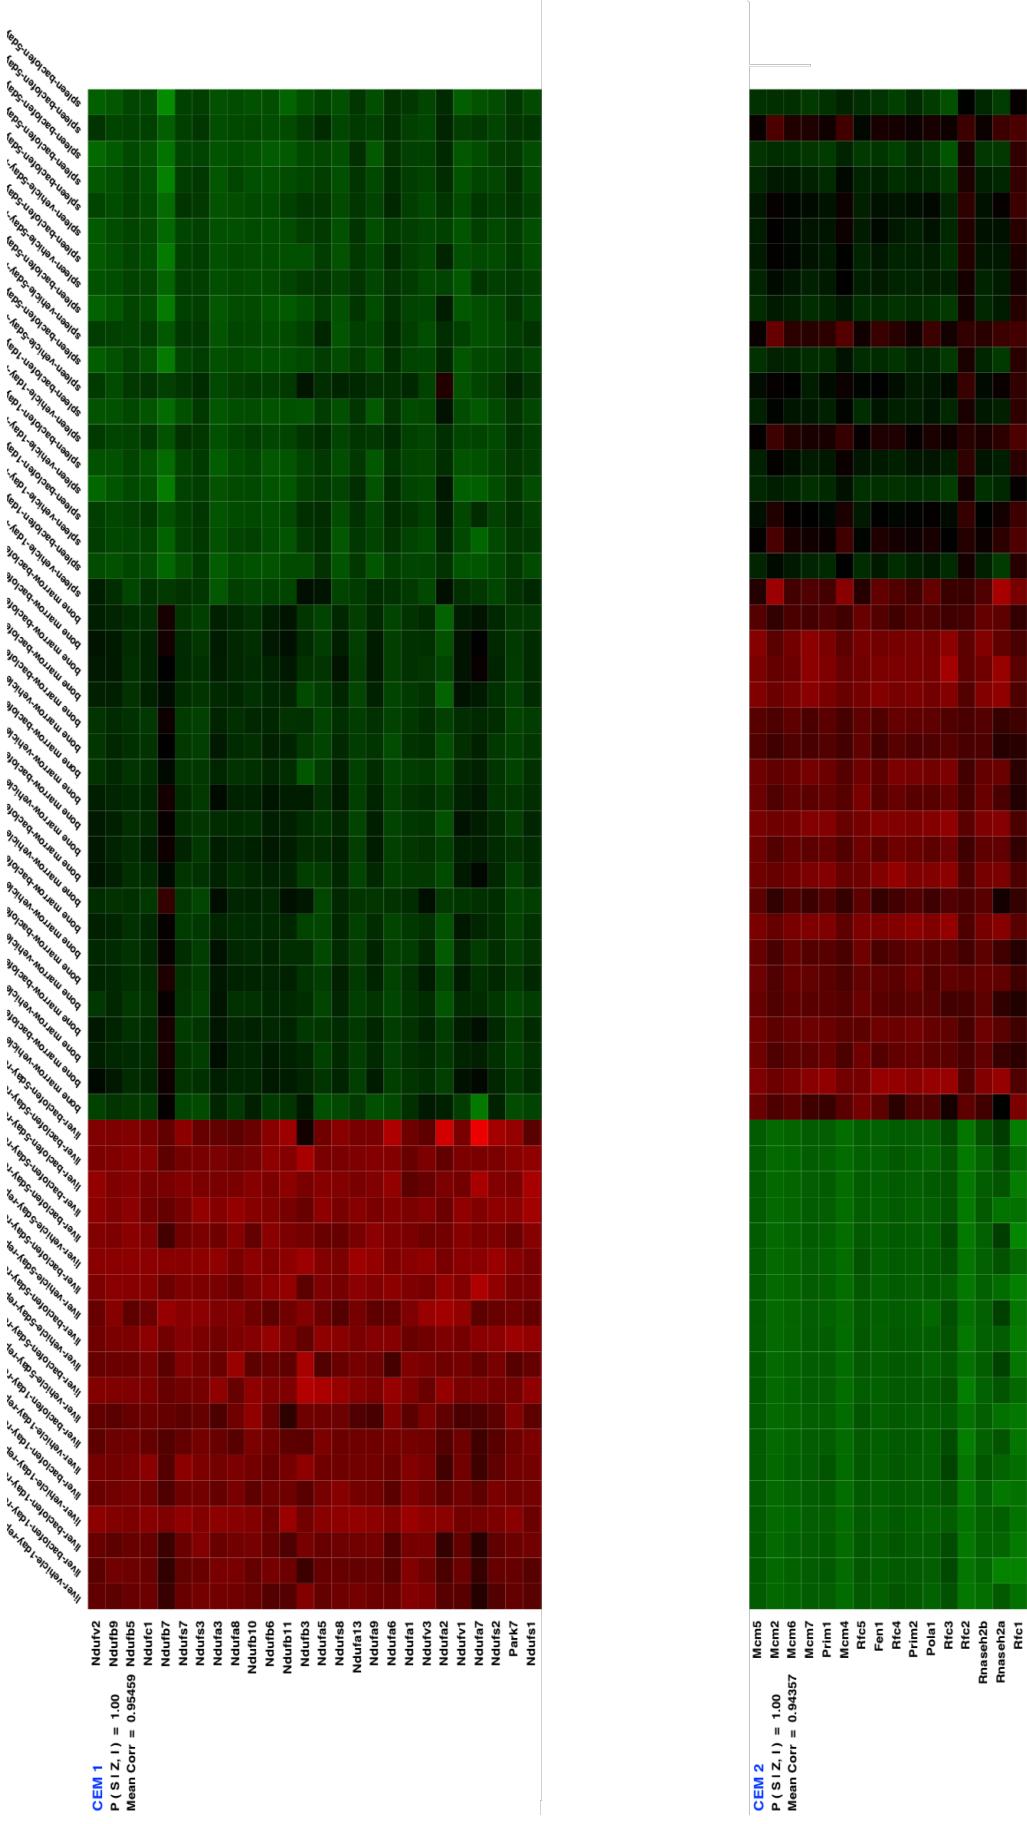


Figure 3: Expression levels of genes in CEM1 and CEM2 over samples in one dataset that is selected by both CEMs using CLJC. The top 26 rows correspond to genes in CEM1, and the bottom 16 rows correspond to genes in CEM2. Columns correspond to 49 experimental samples in the dataset. The expression level is in terms of Z scores. It ranges from -3 (green) to 3 (red).

## A Human genetic data description and BBC2’s performance on the same data with shuffled columns

Table 11 summarizes the 11 populations in the human genetic data set analyzed in Sections 4.1 and 4.2.

Population Label	Population Description	Number of Individuals
ASW	African ancestry in Southwest USA	53
LWK	Luhya in Webuye, Kenya	110
MKK	Maasai in Kinyawa, Kenya	156
YRI	Yoruba in Ibadan, Nigeria	147
GIH	Gujarati Indians in Houston, Texas	101
MEX	Mexican ancestry in LA, California	58
CHB	Han Chinese in Beijing, China	137
CHD	Chinese in Metropolitan Denver, Colorado	109
JPT	Japanese in Tokyo, Japan	113
CEU	Utah residents with Northern and Western European ancestry	112
TSI	Toscani in Italia	102

Table 11: A summary of the human genetic data used in Section 4.1 and 4.2.

In the following we present some results for the analysis of the shuffled data using BBC2. At each percentage of features (columns) to be shuffled (permuted), we repeat the shuffling procedure 10 times, i.e., for each level, 10 datasets are generated. Each dataset contains a randomly selected subset of specified percentage of features being randomly permuted. Table 12 shows the optimal number of clusters and clustering error rates of BBC2 at each pre-specified shuffling level for one shuffled dataset. Results for the rest 9 shuffled datasets at each level are very similar to Table 12. Clustering results on the shuffled data at each pre-specified shuffling level are shown in Tables 13 to 29.

Percentage of shuffled columns	Number of Clusters	CE rate
10%	6	41.40%
15%	6	41.57%
20%	6	41.24%
25%	6	41.32%
30%	6	41.57%
35%	6	41.57%
40%	6	41.74%
45%	6	41.40%
50%	6	41.74%
55%	6	41.74%
60%	6	41.65%
65%	6	41.65%
70%	6	41.57%
75%	5	53.51%
80%	7	54.09%
85%	4	57.76%
90%	6	53.84%

Table 12: A summary of BBC2’s performance on the shuffled human genetic data at different shuffling levels.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	32	21	0	0	0	0
LWK	110	0	0	0	0	0
MKK	3	153	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	52	6	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 13: BBC2’s clustering result on the shuffled data when 10% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	35	18	0	0	0	0
LWK	110	0	0	0	0	0
MKK	5	151	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	52	6	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 14: BBC2's clustering result on the shuffled data when 15% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	32	21	0	0	0	0
LWK	110	0	0	0	0	0
MKK	2	154	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	53	5	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 15: BBC2's clustering result on the shuffled data when 20% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	31	22	0	0	0	0
LWK	110	0	0	0	0	0
MKK	4	152	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	54	4	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 16: BBC2's clustering result on the shuffled data when 25% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	33	20	0	0	0	0
LWK	110	0	0	0	0	0
MKK	6	150	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	53	3	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 17: BBC2's clustering result on the shuffled data when 30% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	32	21	0	0	0	0
LWK	110	0	0	0	0	0
MKK	4	152	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	51	7	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 18: BBC2's clustering result on the shuffled data when 35% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	33	20	0	0	0	0
LWK	110	0	0	0	0	0
MKK	5	151	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	50	8	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 19: BBC2's clustering result on the shuffled data when 40% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	31	22	0	0	0	0
LWK	110	0	0	0	0	0
MKK	3	153	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	52	6	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 20: BBC2’s clustering result on the shuffled data when 45% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	33	20	0	0	0	0
LWK	110	0	0	0	0	0
MKK	9	147	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	54	4	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 21: BBC2’s clustering result on the shuffled data when 50% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	32	21	0	0	0	0
LWK	110	0	0	0	0	0
MKK	7	149	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	52	6	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 22: BBC2's clustering result on the shuffled data when 55% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	36	17	0	0	0	0
LWK	110	0	0	0	0	0
MKK	8	148	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	54	4	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 23: BBC2's clustering result on the shuffled data when 60% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	32	21	0	0	0	0
LWK	110	0	0	0	0	0
MKK	4	152	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	50	8	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 24: BBC2's clustering result on the shuffled data when 65% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	36	17	0	0	0	0
LWK	110	0	0	0	0	0
MKK	5	151	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	52	6	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 25: BBC2's clustering result on the shuffled data when 70% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
ASW	53	0	0	0	0
LWK	110	0	0	0	0
MKK	156	0	0	0	0
YRI	147	0	0	0	0
GIH	0	100	1	0	0
MEX	0	0	6	52	0
CHB	0	0	0	0	137
CHD	0	0	0	0	109
JPT	0	0	0	0	113
CEU	0	0	112	0	0
TSI	0	0	102	0	0

Table 26: BBC2’s clustering result on the shuffled data when 75% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7
ASW	53	0	0	0	0	0	0
LWK	110	0	0	0	0	0	0
MKK	156	0	0	0	0	0	0
YRI	147	0	0	0	0	0	0
GIH	0	101	0	0	0	0	0
MEX	0	0	44	10	2	2	0
CHB	0	0	0	0	0	0	137
CHD	0	0	0	0	0	0	109
JPT	0	0	0	0	0	0	113
CEU	0	0	0	112	0	0	0
TSI	0	0	0	102	0	0	0

Table 27: BBC2’s clustering result on the shuffled data when 80% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4
ASW	53	0	0	0
LWK	110	0	0	0
MKK	156	0	0	0
YRI	147	0	0	0
GIH	0	101	0	0
MEX	0	50	8	0
CHB	0	0	0	137
CHD	0	0	0	109
JPT	0	0	0	113
CEU	0	0	112	0
TSI	0	0	102	0

Table 28: BBC2’s clustering result on the shuffled data when 85% of the columns are shuffled.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	53	0	0	0	0	0
LWK	110	0	0	0	0	0
MKK	156	0	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	101	0	0	0	0
MEX	0	0	47	5	6	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 29: BBC2’s clustering result on the shuffled data when 90% of the columns are shuffled.

Table 30 records false positive rates for shuffled features at different shuffling levels computed based on BBC2’s outputs. Since for each randomly selected feature that has its observations randomly shuffled, we expect the feature to be a pure noise column and not to be selected by any clusters. Low false positive rates indicate that BBC2 is powerful to detect such noises and not use them for clustering.

Shuffling Level	False Positive Rate	Shuffling Level	False Positive Rate	Shuffling Level	False Positive Rate
10%	3.55%	40%	3.79%	70%	3.15%
15%	4.42%	45%	3.53%	75%	6.35%
20%	3.32%	50%	3.22%	80%	10.23%
25%	3.22%	55%	6.21%	85%	9.93%
30%	3.63%	60%	3.16%	90%	7.22%
35%	3.39%	65%	6.02%		

Table 30: False positive rates among shuffled columns at different shuffling levels computed from BBC2’s outputs.

## B BBC2's clustering results on the human genetic data with different thresholds on pairwise correlation

By setting the upper bound for  $r^2$  to be  $10^{-4}$ , 4,518 SNPs are included in the analysis. BBC2 still selects the optimal number of clusters to be 6, and its clustering result is shown in Table 31. A summary of feature heterogeneity is shown in Table 32.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	32	21	0	0	0	0
LWK	110	0	0	0	0	0
MKK	4	152	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	53	5	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 31: BBC2's clustering result when the upper bound for  $r^2$  is  $10^{-4}$ .

Number of Cluster-Specific Distributions	0	1	2	3	4	6
Number of SNPs	179	689	1553	1608	480	9

Table 32: BBC2's estimate on numbers of SNPs that have the specified numbers of cluster-specific distributions for the human genetic dataset with  $r^2 = 10^{-4}$ .

By setting the upper bound for  $r^2$  to be  $10^{-2}$ , 14,840 SNPs are included in the analysis, which is more than 3 times the number of SNPs included when the upper bound is  $10^{-4}$ . These additional SNPs seem to help in better clustering ASW samples, with most of them being in cluster 1 and only 5 individuals in cluster 2. However, it also introduces additional noises, resulting in more MEX individuals being separated from its main cluster, i.e., cluster 4, and being mixed with the European cluster. The clustering and feature selection results are shown in Tables 33 and 34.

Population Abbreviation	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
ASW	48	5	0	0	0	0
LWK	110	0	0	0	0	0
MKK	3	153	0	0	0	0
YRI	147	0	0	0	0	0
GIH	0	0	101	0	0	0
MEX	0	0	0	46	12	0
CHB	0	0	0	0	0	137
CHD	0	0	0	0	0	109
JPT	0	0	0	0	0	113
CEU	0	0	0	0	112	0
TSI	0	0	0	0	102	0

Table 33: BBC2’s clustering result when the upper bound for  $r^2$  is  $10^{-2}$ .

Number of Cluster-Specific Distributions	0	1	2	3	4	6
Number of SNPs	4991	2554	3549	3029	706	11

Table 34: BBC2’s estimate on numbers of SNPs that have the specified numbers of cluster-specific distributions for the human genetic dataset with  $r^2 = 10^{-2}$ .

## C HBBC’s results on the human genetic data with different minimum node sizes

Figures 4–6 show HBBC’s results on the human genetic dataset analyzed in Section 4.2 with the minimum node size decreasing from 40 to 5. HBBC is quite robust to the minimum node size specification.

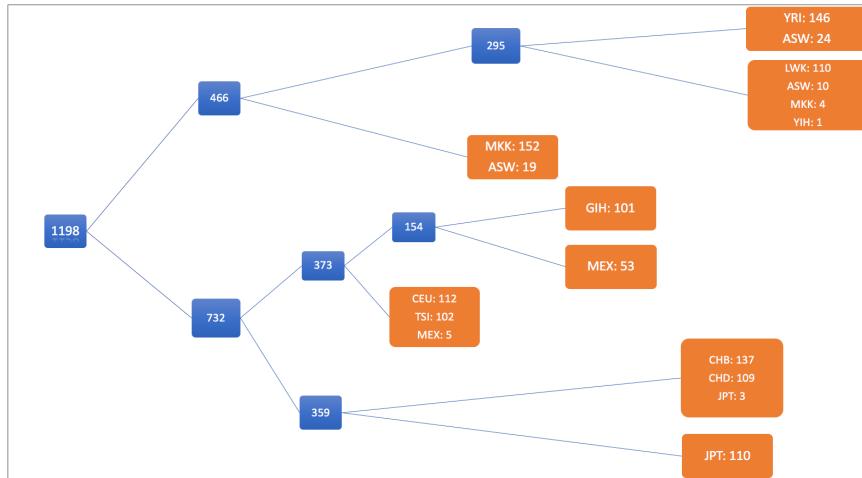


Figure 4: HBBC’s estimation result on the human genetic data, minimum node size=40. Tree and node interpretations follow from Section 4.2.

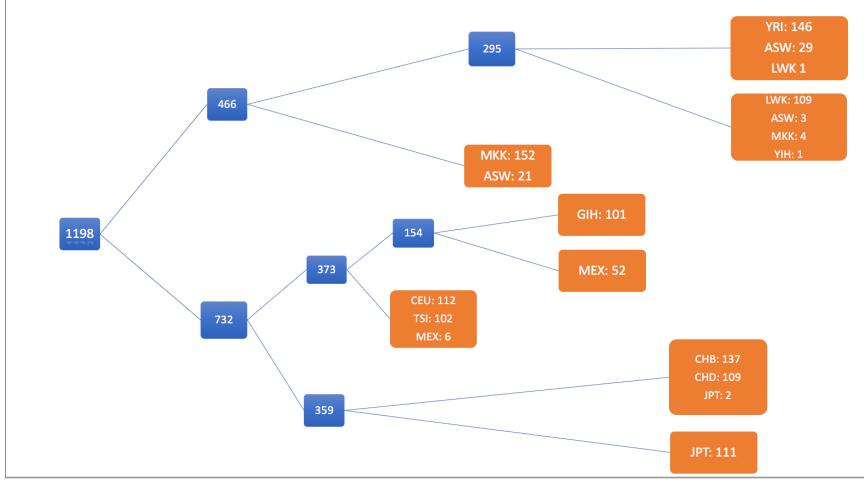


Figure 5: HBBC's estimation result on the human genetic data, minimum node size=30. Tree and node interpretations follow from Section 4.2.

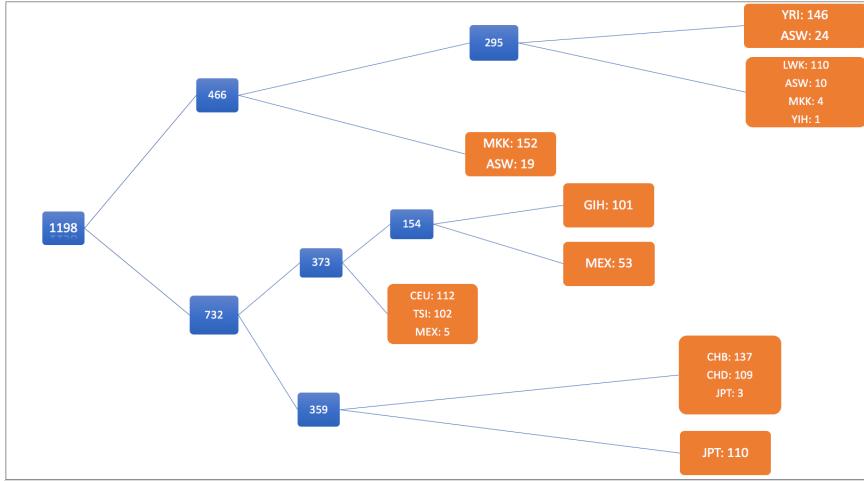


Figure 6: HBBC's estimation result on the human genetic data, minimum node size=20. Tree and node interpretations follow from Section 4.2.

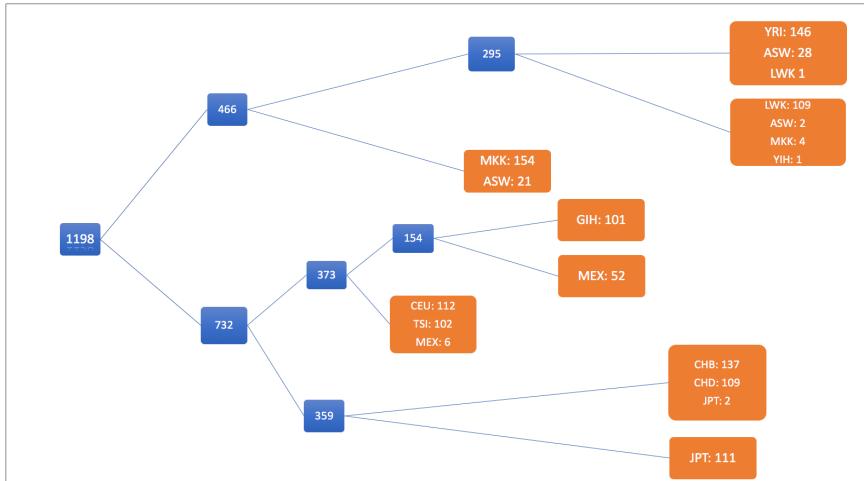


Figure 7: HBBC's estimation result on the human genetic data, minimum node size=5. Tree and node interpretations follow from Section 4.2.

## D Analysis of hair-related SNPs

We first identify a larger pool of SNPs from our dataset. This pool includes both the 4,217 SNPs in the final dataset fed to HBBC and SNPs that are close to and have high correlations with the 4,217 SNPs. The resulting pool contains 34,052 SNPs. In each step, we partition the 34,052 SNPs into two sets: a target set that includes SNPs that are either selected in this step or are highly correlated and near a selected SNP, and a non-target set that contains SNPs that are either not selected in this step or are highly correlated and near a not-selected SNP. Any overlap between the two sets are removed from both sets. We obtain a list of 93 hair and hair color related SNPs from Phenotype-Genotype Integrator website. 4 of such SNPs are in the pool of 34,052 SNPs. We calculate the odds ratio of the following two: being a traits-related SNP and in the target set over being a traits-related SNP but in the non-target set, and being in the target set over being in the non-target set. In step 1, i.e., separating the African populations from the rest of the populations, all 4 SNPs are in the target set. The odds ratios are 2.92, 1.83 and 1.29 in step 5, step 3 and step 2, respectively. All of these steps involve separating populations in different continents. While in step 4, separating the MKK population from the other African sub-populations, the odds ratio is 0.82, indicating hair and hair color become less distinguishable at this stage. The odds ratios are 0, i.e., none of the traits-related SNPs is selected in step 6 and 7: separating sub-populations in Africa and sub-populations in east Asia. This may suggest that specific aspects in hair and hair color related traits that are associated with the 4 SNPs are not distinguishable between these closely related sub-populations.

## E More details on BBC for data integration

We conjugate a Normal-Inverse-Gamma prior for  $\theta_{d,k}$  and  $\sigma_{d,k}^2$  in Section 3.4

$$\begin{aligned}\theta_{d,k} &\sim N(\mu_\theta, \sigma_{d,k}^2/\kappa_\theta), \quad k = 1, \dots, K, \\ \sigma_{d,k}^2 &\sim \text{Inv-Gamma}(\alpha_\sigma, \beta_\sigma), \quad k = 1, \dots, K,\end{aligned}$$

where  $\mu_\theta$ ,  $\kappa_\theta$ ,  $\alpha_\sigma$  and  $\beta_\sigma$  are hyper-parameters.

To enable the Gibbs sampler, we integrate out  $\boldsymbol{\theta}$  and  $\boldsymbol{\sigma}$  to obtain the marginal likelihood function for  $\mathbf{C}$  and  $\mathbf{S}$ ,

$$\begin{aligned}P(\mathbf{Y} | \mathbf{C}, \mathbf{S}, K) &= \int \int P(\mathbf{Y} | \boldsymbol{\theta}, \boldsymbol{\sigma}, \mathbf{C}, \mathbf{S}, K) P(\boldsymbol{\theta} | \mathbf{C}, K) P(\boldsymbol{\sigma} | \mathbf{C}, K) d\boldsymbol{\theta} d\boldsymbol{\sigma} \\ &= \prod_{d=1}^p \left[ \prod_{k=1}^K \left\{ \frac{\left\{ \beta_\sigma^{\alpha_\sigma} (2\pi)^{-\frac{n_k}{2}} / \sqrt{1 + \kappa_\theta^{-1} n_k} \right\} \cdot \left\{ \Gamma(\alpha_\sigma + \frac{n_k}{2}) / \Gamma(\alpha_\sigma) \right\}}{\left( \beta_\sigma + \frac{1}{2} \left( \sum_{i < j: c_i = c_j = k} Y_{d,i,j}^2 + \kappa_\theta \mu_\theta^2 - \frac{\left( \kappa_\theta \mu_\theta + \sum_{i < j: c_i = c_j = k} Y_{d,i,j} \right)^2}{\kappa_\theta + n_k} \right) \right)^{\alpha_\sigma + \frac{n_k}{2}}} \right\}^{S_{d,k}} \right. \\ &\quad \left. \cdot \prod_{d=1}^p \left[ \prod_{i < j: c_i \neq c_j} \frac{\exp \left\{ - \frac{(Y_{d,i,j} - \theta_{d,0})^2}{2\sigma_{d,0}^2} \right\}}{\sqrt{2\pi\sigma_{d,0}^2}} \right]^{S_{d,k}} \cdot \prod_{d=1}^p \left[ \prod_{i < j} \frac{\exp \left\{ - \frac{(Y_{d,i,j} - \theta_{d,0})^2}{2\sigma_{d,0}^2} \right\}}{\sqrt{2\pi\sigma_{d,0}^2}} \right]^{1-S_{d,k}} \right],\end{aligned}$$

where  $n_k$  denotes the number of gene pairs in cluster  $k$ . We can also integrate out  $\mathbf{S}$  from the likelihood function  $P(\mathbf{Y} | \mathbf{C}, \mathbf{S}, K)$  to obtain the marginal distribution of  $\mathbf{C}$ , so that the Gibbs sampler only cycles through  $C_1, \dots, C_n$ :

$$\begin{aligned}
& P(\mathbf{Y} | \mathbf{C}, K) \\
&= \sum_{\mathbf{S}_1 \in \{0,1\}^K} P(\mathbf{S}_1 | \mathbf{C}, K) \cdots \sum_{\mathbf{S}_p \in \{0,1\}^K} P(\mathbf{S}_p | \mathbf{C}, K) P(\mathbf{Y} | \mathbf{C}, \mathbf{S}, K) \\
&= \prod_{d=1}^p \left[ \prod_{k=1}^K (\pi_S) \left( \frac{\left\{ \beta_\sigma^{\alpha_\sigma} (2\pi)^{-\frac{n_k}{2}} / \sqrt{1 + \kappa_\theta^{-1} n_k} \right\} \cdot \left\{ \Gamma(\alpha_\sigma + \frac{n_k}{2}) / \Gamma(\alpha_\sigma) \right\}}{\left( \beta_\sigma + \frac{1}{2} \left( \sum_{i < j: c_i = c_j = k} Y_{d,i,j}^2 + \kappa_\theta \mu_\theta^2 - \frac{\left( \kappa_\theta \mu_\theta + \sum_{i < j: c_i = c_j = k} Y_{d,i,j} \right)^2}{\kappa_\theta + n_k} \right) \right)^{\alpha_\sigma + \frac{n_k}{2}}} \right) \right. \\
&\quad \left. + (1 - \pi_S) \left( \prod_{i < j: c_i = c_j = k} \frac{\exp \left\{ - \frac{(Y_{d,i,j} - \theta_{d,0})^2}{2\sigma_{d,0}^2} \right\}}{\sqrt{2\pi\sigma_{d,0}^2}} \right) \right] \cdot \left( \prod_{i < j: c_i \neq c_j} \frac{\exp \left\{ - \frac{(Y_{d,i,j} - \theta_{d,0})^2}{2\sigma_{d,0}^2} \right\}}{\sqrt{2\pi\sigma_{d,0}^2}} \right).
\end{aligned}$$

We run the sampler for a range of cluster numbers,  $K = 1, \dots, \mathcal{K}$ , and obtain the point estimator of number of clusters  $\hat{K}$  and cluster assignment  $\hat{\mathbf{C}}$  by the MAP estimator. Let  $\mathbf{C}_{(K)}^{(1)}, \dots, \mathbf{C}_{(K)}^{(M)}$  denote MCMC samples of  $\mathbf{C}$  with the number of clusters  $K$ . The MAP estimator for  $K$  and  $\mathbf{C}$  is defined as

$$(\hat{K}, \hat{\mathbf{C}}) = \arg \max_{K, \mathbf{C}_{(K)}^{(m)}: K=1, \dots, \mathcal{K}; m=1, \dots, M} P(\mathbf{C}_{(K)}^{(m)} | \mathbf{Y}). \quad (21)$$

## Acknowledgements

This research was supported in part by NSF Grants DMS-1903139 and DMS-2015411.

## References

Allman, E. S., C. Matias, J. A. Rhodes, et al. (2009). Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics* 37(6A), 3099–3132.

Andrews, S. (2011). In-close2, a high performance formal concept miner. In *International Conference on Conceptual Structures*, pp. 50–62. Springer.

Asgarian, N. and R. Greiner (2007). Using rank-one biclusters to classify microarray data. Technical report, Technical Report CPDC-TR-9906-010, University of Alberta.

Ben-Dor, A., B. Chor, R. Karp, and Z. Yakhini (2003). Discovering local structure in gene expression data: the order-preserving submatrix problem. *Journal of computational biology* 10(3-4), 373–384.

Bergmann, S., J. Ihmels, and N. Barkai (2003). Iterative signature algorithm for the analysis of large-scale gene expression data. *Physical review E* 67(3), 031902.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent dirichlet allocation. *Journal of machine Learning research* 3(Jan), 993–1022.

Caldas, J. and S. Kaski (2008). Bayesian biclustering with the plaid model. In *2008 IEEE Workshop on Machine Learning for Signal Processing*, pp. 291–296. IEEE.

Cheng, Y. and G. M. Church (2000). Biclustering of expression data. In *Ismb*, Volume 8, pp. 93–103.

Chib, S. (1995). Marginal likelihood from the gibbs output. *Journal of the american statistical association* 90(432), 1313–1321.

Chu, L.-F., N. Leng, J. Zhang, Z. Hou, D. Mamott, D. T. Vereide, J. Choi, C. Kendziorski, R. Stewart, and J. A. Thomson (2016). Single-cell rna-seq reveals novel regulators of human embryonic stem cell differentiation to definitive endoderm. *Genome biology* 17(1), 173.

de França, F. O. (2016). A hash-based co-clustering algorithm for categorical data. *Expert Systems with Applications* 64, 24–35.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* 39(1), 1–22.

Galimberti, G. and G. Soffritti (2007). Model-based methods to identify multiple cluster structures in a data set. *Computational statistics & data analysis* 52(1), 520–536.

Gan, G. and J. Wu (2004). Subspace clustering for high dimensional categorical data. *ACM SIGKDD Explorations Newsletter* 6(2), 87–94.

Green, P. J. (1995). Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika* 82(4), 711–732.

Gu, J. and J. S. Liu (2008). Bayesian biclustering of gene expression data. *BMC genomics* 9(1), S4.

Guo, L. (2013). Bayesian biclustering on discrete data: Variable selection methods. PhD Thesis, Department of Statistics, Harvard University.

Hartigan, J. A. (1972). Direct clustering of a data matrix. *Journal of the american statistical association* 67(337), 123–129.

Hartigan, J. A. and M. A. Wong (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28(1), 100–108.

Hochreiter, S., U. Bodenhofer, M. Heusel, A. Mayr, A. Mitterecker, A. Kasim, T. Khamiakova, S. Van Sanden, D. Lin, W. Talloen, et al. (2010). Fabia: factor analysis for bicluster acquisition. *Bioinformatics* 26(12), 1520–1527.

Hoff, P. D. (2005). Subset clustering of binary sequences, with an application to genomic abnormality data. *Biometrics* 61(4), 1027–1036.

Hoff, P. D. et al. (2006). Model-based subspace clustering. *Bayesian Analysis* 1(2), 321–344.

Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of classification* 2(1), 193–218.

Jackson, J. E. (2005). *A user's guide to principal components*, Volume 587. John Wiley & Sons.

Johnson, S. C. (1967). Hierarchical clustering schemes. *Psychometrika* 32(3), 241–254.

Keribin, C., V. Brault, G. Celeux, and G. Govaert (2015). Estimation and selection for the latent block model on categorical data. *Statistics and Computing* 25(6), 1201–1216.

Kim, S., M. G. Tadesse, and M. Vannucci (2006). Variable selection in clustering via dirichlet process mixture models. *Biometrika* 93(4), 877–893.

Kiselev, V. Y., K. Kirschner, M. T. Schaub, T. Andrews, A. Yiu, T. Chandra, K. N. Natarajan, W. Reik, M. Barahona, A. R. Green, et al. (2017). Sc3: consensus clustering of single-cell rna-seq data. *Nature methods* 14(5), 483–486.

Labiod, L. and M. Nadif (2011). Co-clustering for binary and categorical data with maximum modularity. In *2011 IEEE 11th International Conference on Data Mining*, pp. 1140–1145. IEEE.

Lazzeroni, L. and A. Owen (2002). Plaid models for gene expression data. *Statistica sinica*, 61–86.

Lee, M., H. Shen, J. Z. Huang, and J. Marron (2010). Biclustering via sparse singular value decomposition. *Biometrics* 66(4), 1087–1095.

Li, Y., A. A. Jourdain, S. E. Calvo, J. S. Liu, and V. K. Mootha (2017). Clic, a tool for expanding biological pathways based on co-expression across thousands of datasets. *PLoS computational biology* 13(7), e1005653.

Li, Z., C. Chang, S. Kundu, and Q. Long (2020). Bayesian generalized biclustering analysis via adaptive structured shrinkage. *Biostatistics* 21(3), 610–624.

Liu, J. S. (1994). The collapsed gibbs sampler in bayesian computations with applications to a gene regulation problem. *Journal of the American Statistical Association* 89(427), 958–966.

Liu, J. S., J. L. Zhang, M. J. Palumbo, and C. E. Lawrence (2003). Bayesian clustering with variable and transformation selections. *Bayesian statistics* 7, 249–275.

Macá-Meyer, N., A. M. González, J. M. Larruga, C. Flores, and V. M. Cabrera (2001). Major genomic mitochondrial lineages delineate early human expansions. *BMC genetics* 2(1), 1–8.

MacQueen, J. et al. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, Volume 1, pp. 281–297. Oakland, CA, USA.

Marbac, M. and M. Sedki (2017a). Variable selection for mixed data clustering: a model-based approach. *arXiv preprint arXiv:1703.02293*.

Marbac, M. and M. Sedki (2017b). Variable selection for model-based clustering using the integrated complete-data likelihood. *Statistics and Computing* 27(4), 1049–1063.

Marbac, M. and V. Vandewalle (2019). A tractable multi-partitions clustering. *Computational Statistics & Data Analysis* 132, 167–179.

Narasimhan, V. M., N. Patterson, P. Moorjani, N. Rohland, R. Bernardo, S. Mallick, I. Lazaridis, N. Nakatsuka, I. Olalde, M. Lipson, et al. (2019). The formation of human populations in south and central asia. *Science* 365(6457).

Neuwald, A. F., N. Kannan, A. Poleksic, N. Hata, and J. S. Liu (2003). Ran's c-terminal, basic patch, and nucleotide exchange mechanisms in light of a canonical structure for rab, rho, ras, and ran gtpases. *Genome research* 13(4), 673–692.

Patrikainen, A. and H. Mannila (2004). Subspace clustering of high-dimensional binary data—a probabilistic approach. In *Proc. Workshop on Clustering High Dimensional Data, SIAM International Conference on Data Mining*, pp. 57–65. Citeseer.

Pensa, R. G., C. Robardet, and J.-F. Boulicaut (2005). A bi-clustering framework for categorical data. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pp. 643–650. Springer.

Posth, C., G. Renaud, A. Mittnik, D. G. Drucker, H. Rougier, C. Cupillard, F. Valentin, C. Thevenet, A. Furtwängler, C. Wißing, et al. (2016). Pleistocene mitochondrial genomes suggest a single major dispersal of non-africans and a late glacial population turnover in europe. *Current Biology* 26(6), 827–833.

Pritchard, J. K., M. Stephens, and P. Donnelly (2000). Inference of population structure using multilocus genotype data. *Genetics* 155(2), 945–959.

Raftery, A. E. and N. Dean (2006). Variable selection for model-based clustering. *Journal of the American Statistical Association* 101(473), 168–178.

Raghavan, M., P. Skoglund, K. E. Graf, M. Metspalu, A. Albrechtsen, I. Moltke, S. Rasmussen, T. W. Stafford Jr, L. Orlando, E. Metspalu, et al. (2014). Upper palaeolithic siberian genome reveals dual ancestry of native americans. *Nature* 505(7481), 87–91.

Reich, D., K. Thangaraj, N. Patterson, A. L. Price, and L. Singh (2009). Reconstructing indian population history. *Nature* 461(7263), 489–494.

Sans, M. (2000). Admixture studies in latin america: from the 20th to the 21st century. *Human biology*, 155–177.

Schwarz, G. et al. (1978). Estimating the dimension of a model. *The annals of statistics* 6(2), 461–464.

Sokal, R. R. (1958). A statistical method for evaluating systematic relationships. *Univ. Kansas, Sci. Bull.* 38, 1409–1438.

Stuart, T., A. Butler, P. Hoffman, C. Hafemeister, E. Papalexi, W. M. Mauck III, Y. Hao, M. Stoeckius, P. Smibert, and R. Satija (2019). Comprehensive integration of single-cell data. *Cell* 177(7), 1888–1902.

Tadesse, M. G., N. Sha, and M. Vannucci (2005). Bayesian variable selection in clustering high-dimensional data. *Journal of the American Statistical Association* 100(470), 602–617.

Tanay, A., R. Sharan, and R. Shamir (2002). Discovering statistically significant biclusters in gene expression data. *Bioinformatics* 18(suppl\_1), S136–S144.

Wu, J. (2017). Bayesian inference of relationships. PhD Thesis, Department of Statistics, Harvard University.

Yang, W.-H., D.-Q. Dai, and H. Yan (2007). Bioclustering of microarray data based on singular value decomposition. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 194–205. Springer.