

# Block clustering with collapsed latent block models

Jason Wyse · Nial Friel

Received: 13 April 2010 / Accepted: 17 January 2011 / Published online: 5 May 2011  
© Springer Science+Business Media, LLC 2011

**Abstract** We introduce a Bayesian extension of the latent block model for model-based block clustering of data matrices. Our approach considers a block model where block parameters may be integrated out. The result is a posterior defined over the number of clusters in rows and columns and cluster memberships. The number of row and column clusters need not be known in advance as these are sampled along with cluster memberships using Markov chain Monte Carlo. This differs from existing work on latent block models, where the number of clusters is assumed known or is chosen using some information criteria. We analyze both simulated and real data to validate the technique.

**Keywords** Block clustering · Latent block model · Bayesian model choice · Collapsed model

## 1 Introduction

Many data sets arise as a result of a number of features or variables being observed for a collection of objects. As examples, shoppers and the items which they do or do not buy; whether a document contains specific words or not; the expression levels of a gene under a series of conditions in a DNA experiment. Such data will be recorded in a matrix, say, with rows indexing objects and columns indexing features or variables. Often interest will focus on clustering rows and further, clustering the features which distinguish these row clusters. We refer to this task as block clustering

although it is also known as block modelling, biclustering, co-clustering and two-mode clustering.

One of the first approaches to block clustering was suggested by Hartigan (1972) and since then, many have been proposed. Much recent work in block clustering and related areas has been either on the analysis of microarray data (Tibshirani et al. 1999; Cheng and Church 2000; Getz et al. 2000; Lazzeroni and Owen 2002; Kluger et al. 2003) or document classification (Hofmann 2001; Blei et al. 2003; Griffiths and Steyvers 2004). Approaches vary in whether they allow clusters to overlap or not. In our case, the problem can be thought of as permuting the rows and columns of the data matrix to make a “chessboard” of blocks of data having similar value.

While many approaches to document classification are model-based i.e. a parametric underlying model is assumed when clustering data, this is often not the case in microarray analysis. Some exceptions are Lazzeroni and Owen (2002) who assume a Gaussian error model for gene expression with additive effects for gene and condition clusters and Sheng et al. (2003) who assume a multinomial model for expression level in a discretized microarray. It is common to use two-way hierarchical clustering for this data or other partitioning methods (for example Getz et al. 2000). One drawback of these methods is the lack of probabilistic justification as noted by Wit and McClure (2004) (Chapter 7, page 171). A model-based approach allows explicit modelling of noise in the data. This can be an advantage in data, such as microarrays, which is particularly prone to noise, incorporating uncertainty in cluster membership.

In this paper we consider an extension of the latent block model (LBM) approach of Govaert and Nadif (2008). The LBM was developed as an intuitive extension of the finite mixture model used in model-based clustering (Fraley and Raftery 2002) to allow for clustering of objects and features.

---

J. Wyse (✉) · N. Friel  
University College London, London, UK  
e-mail: [jason@stats.ucl.ac.uk](mailto:jason@stats.ucl.ac.uk)

N. Friel  
e-mail: [nial.friel@ucl.ie](mailto:nial.friel@ucl.ie)

We propose a Bayesian LBM. This has been considered previously by van Dijk et al. (2009). In their approach the number of clusters in objects and features is assumed known and Gibbs sampling is used to find clusterings. They choose the number of clusters using an information criterion based on maximum likelihood.

We show that it is possible to sample the number of clusters and the cluster membership jointly using simple Markov chain Monte Carlo (MCMC) on a collapsed model, so that uncertainty in the number of clusters is naturally incorporated as part of our Bayesian LBM. The collapsed model is obtained by integrating out block parameters analytically. This is possible using standard prior assumptions. There is no need to resort to a trans-dimensional sampler, such as the reversible jump sampler of Green (1995). Our idea extends the allocation sampler of Nobile and Fearnside (2007) to two directions, with slight modifications. We discuss differences and similarities of our approach to block clustering with those which are most comparable qualitatively. The sampler is applied to both simulated and real datasets to gauge performance.

The remainder of the paper is organized as follows. Section 2 reviews the LBM and introduces the collapsed Bayesian LBM. Section 3 gives the MCMC sampler which we use. The differences between our sampler and reversible jump samplers are also discussed. Label switching is mentioned and the section concludes with approaches to summarize the output of the sampler. Section 4 applies the approach to simulated data. In Sect. 5 we analyze voting records data from the U.S. congress and compare this to a maximum likelihood analysis. In Sect. 6 we discuss the analysis of microarray data, and use our sampler to analyze data from a yeast microarray experiment. The paper concludes with a discussion.

## 2 Models

The data is  $Y = (y_{ij})$ , an  $n \times m$  matrix. It is assumed rows and columns may be reordered so that the matrix can be represented as  $K \times G$  blocks with data in blocks modelled by the same density, where  $K$  and  $G$  are the number of row and column clusters respectively. This could be imagined as a “chessboard” effect, with  $K - 1$  divisions in the direction of the rows and  $G - 1$  in the direction of the columns. The parameters of the data density are conditional on the block and  $\theta_{kg}$  denotes the parameters for block  $(k, g)$ , with  $\Theta$  denoting the collection of these. We now give a review of the LBM of Govaert and Nadif (2008).

### 2.1 Latent block models

Conditional on  $K$  and  $G$ , let  $\mathcal{U}$  be a latent space indexing the set of all possible clusterings of rows and columns. Then

the distribution of the data  $Y$  can be written

$$p(Y|K, G, \Theta, \phi) = \sum_{\mathbf{u} \in \mathcal{U}} p(\mathbf{u}|K, G, \phi) p(Y|K, G, \mathbf{u}, \Theta)$$

where  $\phi$  are parameters for the distribution of  $\mathbf{u}$ . Here  $\mathbf{u}$  is a multinomial random variable taking values over each possible allocation of rows and columns to clusters. In this regard, suppose these allocations are indexed  $1, \dots, K^n G^m$ ; then

$$\mathbf{u}|K, G, \phi \sim \text{Multinomial}(\phi_1, \dots, \phi_{K^n G^m})$$

with  $\sum_{l=1}^{K^n G^m} \phi_l = 1$ . Govaert and Nadif (2008) make the assumption that row and column clusterings are independent *a priori*, so that  $p(\mathbf{u}|K, G, \phi) = p(\mathbf{z}|K, \omega) p(\mathbf{w}|G, \rho)$  where  $z_i = k$  if row  $i$  is in cluster  $k$  and  $w_j = g$  if column  $j$  is in cluster  $g$ . The probability of a row belonging to cluster  $k$  is  $\omega_k$  and  $\rho_g$  denotes the probability that a column belongs to cluster  $g$ . The LBM is then

$$\begin{aligned} p(Y|K, G, \Theta, \omega, \rho) \\ = \sum_{(\mathbf{z}, \mathbf{w}) \in \mathcal{Z} \times \mathcal{W}} p(\mathbf{z}|K, \omega) p(\mathbf{w}|G, \rho) p(Y|K, G, \mathbf{z}, \mathbf{w}, \Theta) \end{aligned} \quad (1)$$

where  $\mathcal{Z}$  and  $\mathcal{W}$  denote the latent spaces of all row and column clusterings respectively. When constructing the data likelihood given the latent allocations, we make the assumption of local independence. That is, within a block, data are independent. This gives data likelihood conditional on  $\mathbf{z}, \mathbf{w}$ ,

$$p(Y|K, G, \mathbf{z}, \mathbf{w}, \Theta) = \prod_{k=1}^K \prod_{g=1}^G \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij}|\theta_{kg}).$$

As  $|\mathcal{Z} \times \mathcal{W}| = K^n G^m$ , it is not feasible to calculate (1). We now review a way to fit this model using a method based on Expectation-Maximization (EM) (Dempster et al. 1977) due to Govaert and Nadif (2008).

#### 2.1.1 Estimation using BEM2

Here we outline the BEM2 algorithm of Govaert and Nadif (2008) which we will compare our approach with later (Sect. 5). Let the random variables  $r_{ik} = \mathbb{I}\{z_i = k\}$  and  $c_{jg} = \mathbb{I}\{w_j = g\}$ . Then

$$\mathbf{r}_i \sim \text{Multinomial}(\omega_1, \dots, \omega_K)$$

$$\mathbf{c}_j \sim \text{Multinomial}(\rho_1, \dots, \rho_G).$$

The complete (or classification) log-likelihood associated with the LBM (1) is

$$\begin{aligned} \mathcal{L}(\mathbf{r}, \mathbf{c}, \boldsymbol{\omega}, \boldsymbol{\rho}, \Theta) &= \sum_{i=1}^n \sum_{k=1}^K r_{ik} \log \omega_k + \sum_{j=1}^m \sum_{g=1}^G c_{jg} \log \rho_g \\ &+ \sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^K \sum_{g=1}^G r_{ik} c_{jg} \log p(y_{ij} | \theta_{kg}). \end{aligned} \quad (2)$$

The E step using this log-likelihood directly is intractable due to the dependence structure among the rows and columns. Govaert and Nadif (2008) suggest a variational approximation to the joint distribution of the latent  $\mathbf{r}, \mathbf{c}$  which leaves  $\mathbf{r}$  and  $\mathbf{c}$  independent. Then using the interpretation of EM due to Neal and Hinton (1998) this leads to a new “fuzzy” criterion for block clustering

$$\begin{aligned} \mathcal{G}(\mathbf{s}, \mathbf{t}, \boldsymbol{\omega}, \boldsymbol{\rho}, \Theta) &= \mathcal{L}(\mathbf{s}, \mathbf{t}, \boldsymbol{\omega}, \boldsymbol{\rho}, \Theta) - \sum_{i=1}^n \sum_{k=1}^K s_{ik} \log s_{ik} \\ &- \sum_{j=1}^m \sum_{g=1}^G t_{jg} \log t_{jg} \end{aligned}$$

which can be alternately maximized with respect to  $\mathbf{s}, \mathbf{t}$  and  $\boldsymbol{\omega}, \boldsymbol{\rho}, \Theta$  where  $s_{ik} = \Pr(z_i = k)$  and  $t_{jg} = \Pr(w_j = g)$ .

The possible ways in which this criterion may be alternately maximized determines different algorithms. The BEM2 algorithm maximizes it as follows.

1. Initialize the unknowns  $\mathbf{s}, \mathbf{t}, \boldsymbol{\omega}, \boldsymbol{\rho}, \Theta$  at some sensible value.
2. Maximize  $\mathcal{G}$  with respect to  $\mathbf{s}, \boldsymbol{\omega}$  and  $\Theta$  keeping  $\mathbf{t}$  and  $\boldsymbol{\rho}$  fixed.
3. Maximize  $\mathcal{G}$  with respect to  $\mathbf{t}, \boldsymbol{\rho}$  and  $\Theta$  keeping  $\mathbf{s}$  and  $\boldsymbol{\omega}$  fixed.
4. Iterate steps 2-3 until convergence.

It is noted that each sweep of BEM2 has two maximizations of  $\Theta$ . This maximization procedure is reported to have outperformed the other schemes considered in Govaert and Nadif (2008), so we use it here to compare with our Bayesian approach.

### 2.1.2 Choosing $K$ and $G$ when using BEM2

In (1), it is assumed that  $K$  and  $G$  are known. The number of clusters assumed can have a considerable effect on the output of clustering algorithms. Usually, many runs, each with a different number of clusters, are necessary. These are then compared to find the best clustering, either based on some information criterion or visual inspection of plots.

Since the LBM is defined in terms of the latent allocation vectors  $\mathbf{z}$  and  $\mathbf{w}$ , it is not clear how one could use a standard information criterion (e.g. BIC, Schwarz 1978) here

to choose the number of components best supported by the data. One approach may be to use the maximized complete log-likelihood treating the row and column allocations as unknown parameters,  $\mathcal{L}(\hat{\mathbf{r}}, \hat{\mathbf{c}}, \hat{\boldsymbol{\omega}}, \hat{\boldsymbol{\rho}}, \hat{\Theta})$ . van Dijk et al. (2009) have used this approach for LBMs when using AIC-3 (Bozdogan 1994) to choose  $K$  and  $G$ . In this case the number of parameters to be estimated is  $n(K-1) + m(G-1) + dKG + (K-1) + (G-1)$  where  $d$  is the dimension of any  $\theta_{kg}$ . A separate model estimation is required for each  $K$  and  $G$  combination over a grid of plausible models. Adopting this type of approach crucially involves replacing the maximum log likelihood with a maximized complete log likelihood and also the issue of selection of a particular information criterion, and could therefore be criticized for both reasons.

The Bayesian LBM we propose seeks to incorporate uncertainty in  $K$  and  $G$  into the model. This is so that the clustering task is also one of cluster model determination. The model determination task and the allocation task are dealt with simultaneously through a fully Bayesian approach. This has analogy with some other block clustering strategies, which undertake greedy searches to find new row and column clusters. See for example Hartigan (2000). An advantage here is that the search has a probabilistic justification based on a posterior distribution for  $K$  and  $G$ . In the next section we introduce the Bayesian LBM which is at the core of the clustering procedures we discuss.

## 2.2 Bayesian latent block models

The Bayesian LBM is formed by taking prior densities on  $K, G, \Theta, \boldsymbol{\omega}$  and  $\boldsymbol{\rho}$ . Let  $\pi(\cdot)$  denote prior and posterior densities. Then we may write down the posterior of the number of clusters and latent cluster allocations from Bayes' theorem

$$\begin{aligned} \pi(K, G, \mathbf{z}, \mathbf{w}, \boldsymbol{\omega}, \boldsymbol{\rho}, \Theta | Y) &\propto p(\mathbf{z} | K, \boldsymbol{\omega}) p(\mathbf{w} | G, \boldsymbol{\rho}) p(Y | K, G, \mathbf{z}, \mathbf{w}, \Theta) \\ &\times \pi(\Theta | K, G) \pi(\boldsymbol{\omega} | K) \pi(\boldsymbol{\rho} | G) \pi(K) \pi(G). \end{aligned} \quad (3)$$

Adopting a conjugate prior for  $\boldsymbol{\omega}, \boldsymbol{\rho}$  and each  $\theta_{kg}$  allows one to integrate these from the posterior analytically. We call this collapsing. Doing this allows us to obtain the marginal posterior  $\pi(K, G, \mathbf{z}, \mathbf{w} | Y)$ . Samples can be generated from this posterior using the MCMC sampler of Sect. 3. This is similar to the general approach of Nobile and Fearnside (2007). The idea of collapsing has been used by Sheng et al. (2003) in the analysis of a discretized microarray and by Griffiths and Steyvers (2004) in latent Dirichlet analysis for document classification. It would be possible to estimate this model without integrating out parameters by using reversible jump MCMC (RJMCMC) (Green 1995). We discuss this further in Sect. 3.2.

We choose a standard conjugate prior for each of the parameters to be integrated out. For example,  $\omega \sim \text{Dirichlet}(\alpha, \dots, \alpha)$  and  $\rho \sim \text{Dirichlet}(\beta, \dots, \beta)$  *a priori*. For the examples considered in this paper, we take the non-informative values  $\alpha = 1, \beta = 1$ . The prior on  $\theta_{kg}$  will depend on the distribution assumed for the data. For the most widely used models, a standard conjugate prior will be available. The  $\theta_{kg}$  are assumed independent *a priori*. This leads to the posterior

$$\pi(K, G, \mathbf{z}, \mathbf{w} | Y) \propto \pi(K)\pi(G) \frac{\Gamma\{\alpha K\} \prod_{k=1}^K \Gamma\{n_k + \alpha\}}{\Gamma\{\alpha\}^K \Gamma\{n + \alpha K\}} \times \frac{\Gamma\{\beta G\} \prod_{g=1}^G \Gamma\{m_g + \beta\}}{\Gamma\{\beta\}^G \Gamma\{m + \beta G\}} \prod_{k=1}^K \prod_{g=1}^G M_{kg}, \tag{4}$$

where

$$M_{kg} = \int \pi(\theta_{kg}) \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij} | \theta_{kg}) d\theta_{kg},$$

where  $n_k$  is the number of rows in cluster  $k$  and  $m_g$  is the number of columns in cluster  $g$ . An outline of the calculation of the posterior is given in Appendix A.

The priors for the number of clusters,  $\pi(K)$  and  $\pi(G)$  are taken to be truncated Poisson(1) over the ranges  $1, \dots, K_{\max}$  and  $1, \dots, G_{\max}$ . Examples of Poisson priors being adopted for the number of components include Phillips and Smith (1996) and Stephens (2000). The use of a truncated Poisson(1) prior has been justified in Nobile (2005). We did experiment with a uniform prior on the number of clusters. However, we found that this gave unnecessary empty clusters in some situations.

We now give the  $M_{kg}$  for two useful data models which we will use in examples later.

### 2.2.1 Bernoulli model for binary data

Assume that  $\Pr(y_{ij} = 1 | z_i = k, w_j = g) = \theta_{kg}$ . We take a Beta( $\gamma, \delta$ ) prior on  $\theta_{kg}$ . Then

$$M_{kg} = \frac{\Gamma\{\gamma + \delta\} \Gamma\{s_{kg} + \gamma\} \Gamma\{n_k m_g - s_{kg} + \delta\}}{\Gamma\{\gamma\} \Gamma\{\delta\} \Gamma\{n_k m_g + \gamma + \delta\}}$$

where  $s_{kg}$  are the block sufficient statistics given by  $s_{kg} = \sum_{i:z_i=k} \sum_{j:w_j=g} y_{ij}$ . Further detail on the calculation is given in Appendix B.

### 2.2.2 Gaussian model for continuous data

Assume  $y_{ij} | z_i = k, w_j = g \sim N(\mu_{kg}, \sigma_{kg}^2)$ . Take the priors  $\mu_{kg} \sim N(\xi, \tau^2 \sigma_{kg}^2)$  and  $\sigma_{kg}^2 \sim \text{IG}(\delta/2, \gamma/2)$  where  $\text{IG}(a, b)$

is the Inverse-Gamma distribution:  $p(x) = \frac{b^a}{\Gamma(a)} x^{-(a+1)} \times \exp\{-b/x\}$ . Then

$$M_{kg} = \frac{\gamma^{\delta/2} \Gamma\{(n_k m_g + \delta)/2\}}{\pi^{n_k m_g/2} \Gamma\{\delta/2\} (n_k m_g \tau^2 + 1)^{1/2}} \times \left( ss_{kg} - \frac{\tau^2 (s_{kg} + \xi/\tau^2)^2}{n_k m_g \tau^2 + 1} + \frac{\xi^2}{\tau^2} + \gamma \right)^{-(n_k m_g + \delta)/2}$$

where  $s_{kg} = \sum_{i:z_i=k} \sum_{j:w_j=g} y_{ij}$  and  $ss_{kg} = \sum_{i:z_i=k} \sum_{j:w_j=g} y_{ij}^2$ . Further details on calculating  $M_{kg}$  are given in Appendix B.

## 3 MCMC sampling of clusterings

The sampler which we propose consists of four different moves. The first is just a standard Gibbs update for the row/column label. The second proposes to reallocate collections of rows and columns. The final two moves propose to add or remove clusters. We describe the moves for rows, but they apply to columns analogously. When running the algorithm, the moves are each applied to the rows and columns in a single sweep. Since the LBM will be invariant to cluster labellings, we will encounter the label switching problem. We outline how to deal with this as well as discussing how to summarize the output from the sampler.

### 3.1 MCMC moves

#### 3.1.1 Gibbs sampling to update the allocation of one row

Suppose row  $i$  is currently in cluster  $k$ . We then sample its new allocation,  $\tilde{z}_i$  from the distribution

$$p(\tilde{z}_i = k' | Y, K, G, \mathbf{z}_{-i}, \mathbf{w}) \propto \frac{n_{k'} + \alpha}{n_k - 1 + \alpha} \prod_{g=1}^G \frac{M_{k'g}^{(+i)} M_{kg}^{(-i)}}{M_{k'g} M_{kg}}, \quad k' \neq k \tag{5}$$

and  $p(\tilde{z}_i = k | Y, K, G, \mathbf{z}_{-i}, \mathbf{w}) \propto 1$  where  $M_{kg}^{(-i)}$  and  $M_{k'g}^{(+i)}$  are obtained respectively by removing row  $i$  from cluster  $k$  and adding it to cluster  $k'$  within column cluster  $g$ . The total computational effort required for the Gibbs sweep on rows and columns is  $O((n + m)KG)$  which may be prohibitive for large  $K, G, n$  or  $m$ . It is possible to move one row and column between clusters using a Metropolis-Hastings move. This could be alternated with a Gibbs update to reduce computational overhead or some mixture of the two moves could be used.

### 3.1.2 Move to update the allocation of more than one row

This move is similar to move M3 in Nobile and Fearnside (2007). Its role is to move more than one row at a time. The way in which new row allocations are proposed should isolate clusters more quickly than just performing one row Gibbs updates. The procedure is as follows. Choose two row clusters  $k$  and  $k'$  at random. Let  $S$  be the index set of rows currently belonging to clusters  $k$  and  $k'$ . The members of  $S$  are randomly reordered. Imagining clusters  $k$  and  $k'$  to be empty initially and  $S$  to be full, we sequentially take each row from  $S$  and allocate it to  $k$  or  $k'$ . This allocation is done using the probability that the current clusters  $k$  or  $k'$  generated that row conditioning on rows that have already been reallocated to  $k$  or  $k'$ . For row  $i$  in  $S$  these probabilities are denoted by  $p_k^{(i)}$  and  $p_{k'}^{(i)}$  with  $p_k^{(i)} + p_{k'}^{(i)} = 1$ . To write down the proposal probability of this move we use  $\bar{M}_{kg}$ ,  $g = 1, \dots, G$  to represent the integrated likelihood of the members placed in cluster  $k$  before member  $i$  has been processed. Similarly  $\bar{n}_k$  represents the number of rows in cluster  $k$  before  $i$  has been processed. Then using similar notation to the Gibbs move it can be shown (see Appendix A.2, Nobile and Fearnside 2007) that

$$\frac{p_{k'}^{(i)}}{p_k^{(i)}} = \frac{\bar{n}_{k'} + \alpha}{\bar{n}_k + \alpha} \prod_{g=1}^G \frac{\bar{M}_{k'g}^{(+i)} \bar{M}_{kg}}{\bar{M}_{k'g} \bar{M}_{kg}^{(+i)}}$$

Using  $p_k^{(i)} + p_{k'}^{(i)} = 1$ , the above can be solved for  $p_{k'}^{(i)}$ . The proposed allocation of row  $i$ ,  $\tilde{z}_i$  may then be sampled. Once the quantities  $\bar{n}_{k'}$ ,  $\bar{n}_k$ ,  $\bar{M}_{k'g}$  and  $\bar{M}_{kg}$  have been updated based on  $\tilde{z}_i$ , the next row in  $S$  can be dealt with.

When all members of  $S$  have been processed the proposal probability of moving from  $\mathbf{z}$  to  $\tilde{\mathbf{z}}$  is

$$\frac{1}{K(K-1)} \prod_{i \in S} p_{\tilde{z}_i}^{(i)}$$

For the reverse move the proposal probability is

$$\frac{1}{K(K-1)} \prod_{i \in S} p_{z_i}^{(i)}$$

The new allocation  $\tilde{\mathbf{z}}$  is then accepted with probability  $\min(1, A)$  where

$$A = \frac{\Gamma\{\bar{n}_k + \alpha\} \Gamma\{\bar{n}_{k'} + \alpha\}}{\Gamma\{n_k + \alpha\} \Gamma\{n_{k'} + \alpha\}} \prod_{g=1}^G \frac{\bar{M}_{k'g} \bar{M}_{kg}}{\bar{M}_{k'g}^{(+i)} \bar{M}_{kg}^{(+i)}} \times \prod_{i \in S} \frac{p_{z_i}^{(i)}}{p_{\tilde{z}_i}^{(i)}}$$

and  $\bar{n}_k$ ,  $\bar{n}_{k'}$ ,  $\bar{M}_{k'g}$ ,  $\bar{M}_{kg}$  are the proposed cluster sizes and integrated block likelihoods when all the members of  $S$  have been processed.

### 3.1.3 Moves to split or combine clusters

To add a cluster we first randomly propose a cluster,  $k$ , to “split”. The new cluster will be labelled  $K + 1$  if the current number of clusters is  $K$ . In the same way as the move to reallocate more than one row (Sect. 3.1.2), the probability of a row proposed as being in cluster  $k$  or  $K + 1$  is given by the conditional probability it was generated by that cluster, the rows being processed sequentially. Clearly the order in which rows are processed is important. Thus for the split and combine moves we place an ordering on the members of cluster  $k$ , that is, the order in which the members are arranged in cluster  $k$  is important. As well as taking members out from cluster  $k$  and placing them in cluster  $K + 1$ , this is important when we place all members back into cluster  $k$  in the combine move. It is possible to propose a label swap of  $K + 1$  with any other label selected at random (itself included), say  $k'$ . This then would split cluster  $k$  into clusters  $k$  and some  $\{1, \dots, K + 1\} \setminus k$ .

Let  $S$  denote the index set of rows currently belonging to cluster  $k$ . We choose a split move with probability  $p_S^K$ . For the split move, the denominator in the proposal ratio will be

$$p(\mathbf{z} \rightarrow \tilde{\mathbf{z}}) = p_S^K \frac{1}{K(K+1)} \frac{1}{n_k!} \prod_{i \in S} p_{z_i}^{(i)}$$

where the second term accounts for selecting the cluster to split, and then the cluster to swap labels with, the third term accounts for the number of ways in which members may be arranged (processed), and the fourth term is the product of conditional probabilities (see Sect. 3.1.2).

For the combine move, two clusters are selected at random, say  $k$  and  $k'$  from the  $K + 1$  available. Then all members of cluster  $k'$  are proposed to be placed back in  $k$ . Thus the numerator in the proposal probability for the split move is

$$p(\tilde{\mathbf{z}} \rightarrow \mathbf{z}) = (1 - p_S^{K+1}) \frac{1}{K(K+1)} \frac{1}{n_k!}$$

where the first term is the probability of proposing a combine move, the second accounts for the clusters selected, and the third accounts for the number of ways in which the members of cluster  $k$  may be arranged.

The acceptance probability for the split move is then  $\min(1, A)$  where

$$A = \frac{\pi(K+1)}{\pi(K)} \frac{\Gamma\{n + \alpha K\}}{\Gamma\{n + \alpha(K+1)\}} \times \frac{\Gamma\{\alpha(K+1)\} \Gamma\{\bar{n}_k + \alpha\} \Gamma\{\bar{n}_{k'} + \alpha\}}{\Gamma\{\alpha\} \Gamma\{\alpha K\} \Gamma\{n_k + \alpha\}} \times \prod_{g=1}^G \frac{\bar{M}_{kg} \bar{M}_{k'g}}{M_{kg}} \times \frac{1 - p_S^{K+1}}{p_S^K} \left( \prod_{i \in S} p_{z_i}^{(i)} \right)^{-1}$$

and  $\tilde{n}_k, \tilde{n}_{k'}, \tilde{M}_{kg}, \tilde{M}_{k'g}$  give the proposed sizes and integrated block likelihoods of the proposed clusters.

The acceptance probability for the combine move is  $\min(1, A^{-1})$ . These moves are similar to the “split and combine” moves discussed by Richardson and Green (1997). Our experiments showed that they gave satisfactory mixing and higher acceptance rates than proposing empty clusters.

### 3.2 Form of reversible jump sampler

As noted, the moves discussed in the previous section resemble moves used in the RJMCMC sampler of Richardson and Green (1997) and other RJ samplers for related classification problems (for example Robert et al. 2000). The difference with our sampler is that the space we sample from is of fixed dimension. This is due to collapsing. Performing an equivalent RJ analysis to that presented here would be challenging for LBMs. This would mean extending the Gibbs sampler of van Dijk et al. (2009) to include variable dimension moves for splitting or combining clusters. The construction of proposal densities for variable dimensional moves in RJ samplers can be crucial to their performance. Work has been done in this area (Brooks et al. 2003; Green 2003), but for many applications construction of proposals is case specific. The reason for entertaining a RJ analysis here is that we are not only concerned with finding a cluster allocation for a specified LBM. We are also interested in exploring different cluster models, and so the task also becomes one of model determination as discussed in Sect. 2.1.2.

Consider splitting a row cluster  $k$  into  $k$  and  $k'$  in a typical RJ approach. This is more difficult than the component splitting case in Richardson and Green (1997), since splitting each row cluster gives rise to  $d(G+1)$  new parameters where  $d$  is the dimension of any  $\theta_{kg}$ . Finding a proposal that will mix well may require lots of trial and error, especially if  $d > 1$  or  $G$  is even moderately large. Moreover, computational time would increase dramatically with respect to the collapsed LBM in these situations.

Using a collapsed model, is, in a sense, a form of variance reduction for this model. We reduce variability in sampling of allocations, by integrating out  $\omega, \rho$  and  $\Theta$ . This should give better sampling of the high probability clusterings of the data, since uncertainty due to parameter values has vanished.

### 3.3 Label switching

The joint posterior of cluster models and allocations or labels (4) is invariant to label switching, that is, the labels are not identifiable. If there is one labelling  $1, \dots, K$  of the rows, then any permutation of this, say,  $\sigma(1), \dots, \sigma(K)$ ,

gives exactly the same information about clustering relationships. The posterior on row labels has  $K!$  indistinguishable modes. Generally as the Markov chain progresses, we will observe switches between these equivalent modes; the well known label switching phenomenon. In our case label switching can occur for row and column labels independently. There are many approaches for dealing with the label switching problem (Stephens 2000; Celeux et al. 2000). The approach we adopt here is due to Nobile and Fearnside (2007). It is ideal for our purposes since it does not involve loss functions based on sampled model parameters (which are no longer in our model). It just requires the samples of  $\mathbf{z}$  and  $\mathbf{w}$ .

We now outline the procedure we use to deal with label switching. Some more details are given in Appendix C. We post-process row and column allocation vectors separately. The re-labelled data can then be used to compute posterior probability of cluster memberships and other quantities of interest. To post process the label vectors  $\mathbf{z}_1, \mathbf{z}_2, \dots$  output from MCMC we begin by arranging these in order of increasing number of non-empty components. This gives the ordering  $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \dots$ , where for  $s < t$ ,  $\mathbf{z}^{(s)}$  uses either the same number of components as  $\mathbf{z}^{(t)}$  in total, or less. For example, with  $K = 4$ ,  $\mathbf{z}^{(s)} = (3, 3, 2, 2, 1)$  would come before  $\mathbf{z}^{(t)} = (4, 4, 3, 3, 1, 2)$ . Suppose we have processed and re-labelled the vectors  $\mathbf{z}^{(t)}$  up to time  $T-1$ , and there are  $K_{T-1}$  non-empty components in  $\mathbf{z}^{(T-1)}$ . Compute a cost matrix with general element

$$C(k_1, k_2) = \sum_{t=1}^{T-1} \sum_{i=1}^n \mathbf{I} \left\{ z_i^{(t)} \neq k_1, z_i^{(T)} = k_2 \right\}.$$

Then the more  $\mathbf{z}^{(T)}$  disagrees with the vectors already processed, the higher this cost will be (see Appendix C). The square assignment algorithm of Carpaneto and Toth (1980) returns the permutation  $\sigma(\cdot)$  of the labels in  $\mathbf{z}^{(T)}$  which minimizes the total cost  $\sum_{k=1}^{K_{T-1}} C(k, \sigma(k))$ . We then relabel  $\mathbf{z}^{(T)}$  by permuting the labels according to  $\sigma(\cdot)$ . This approach will become computationally more expensive as the number of clusters  $K$  increases (each call to the square assignment algorithm is at worst  $O(K^4)$  for each MCMC iterate).

### 3.4 Summarizing MCMC output

Having sampled both the number of clusters and cluster memberships, it will be of interest to give a summary of the sampling. As different  $(K, G)$  cluster models are structurally different, it is not possible to give an “average” of cluster membership. We suggest two summaries.

#### 3.4.1 Using the modal cluster model

The first summary focuses on using the modal, or most visited model from the MCMC output. It takes the series of

$(K, G)$  visited models and chooses the pair which appear most often. Call this pair  $(\hat{K}, \hat{G})$ . Suppose this pair has occurred  $N$  times in the post burn-in sample. We extract the  $N$  pairs of label vectors  $\mathbf{z}$  and  $\mathbf{w}$  corresponding to these occurrences. We then post process these label vectors using the procedure to undo label switching outlined in the previous section. This will be necessary to compute posterior distributions of row and column cluster membership. After computing the posterior distributions of row membership, row  $i$  has distribution  $(q_{i1}, \dots, q_{i\hat{K}})$  where  $q_{ik}$  is the estimated posterior probability row  $i$  belongs to cluster  $k$  in the  $(\hat{K}, \hat{G})$  cluster model. For the summary we assign  $i$  to cluster  $\arg \max_k q_{ik}$ . The columns are given the same treatment.

### 3.4.2 Using the MAP

Since we are sampling from the fixed dimensional posterior  $\pi(K, G, \mathbf{z}, \mathbf{w}|Y)$ , the maximum *a posteriori* (MAP) cluster model and cluster membership  $(K, G, \mathbf{z}, \mathbf{w})_{\text{MAP}}$  is also a useful summary of the MCMC output. The MAP gives the visited  $(K, G, \mathbf{z}, \mathbf{w})$  having highest probability *a posteriori* from the samples obtained.


## 4 Simulation experiment

To see how the sampler discriminates between different cluster models, it was run on some simulated data. We gen-

erated three  $200 \times 200$  binary matrices with 4 and 4, 2 and 5 and 1 and 4 row and column clusters respectively. In each case, the block parameter  $\theta_{kg}$  was drawn uniformly from  $[0, 1]$ . The blocks were then generated using  $\text{Bernoulli}(\theta_{kg})$  random variates. This is shown in the left of Fig. 1. Clusters were made less distinguishable by transforming the generated  $\theta_{kg}$  to the intervals  $[0.2, 0.8]$  and  $[0.3, 0.7]$  using  $\theta_{kg}^{[a,b]} = a + \theta_{kg}(b - a)$  and generating two further matrices. The rows and columns of the resulting matrices were then randomly reordered, disguising the data structure. The chain was run for 1000 burn-in iterations and a further 16 000 iterations on each data set. We assumed a  $\text{Beta}(1, 1)$  prior for  $\theta_{kg}$  in all cases. The priors for  $\omega$  and  $\rho$  are as in Sect. 2.2.

We looked at two performance diagnostics of the sampler. The first was the posterior model probability (PMP) of the model used to generate the data and the second was the integrated autocorrelation time (IAT) of sampled cluster models. Computing the PMP just amounts to counting the number of times the model in question was visited and dividing by the total number of samples. For the IAT we identify cluster models  $(K, G)$  by a model index  $R = 1, \dots, K_{\text{max}}G_{\text{max}}$ . Then we estimate the quantity  $\tau = 1 + 2 \sum_{t=1}^{\infty} \rho_R(t)$ , where  $\rho_R(t)$  is the autocorrelation of the series of post burn-in samples  $R_1, R_2, \dots$  at lag  $t$ . The series here refers to the cluster models sampled from the posterior (4). Lower values of the IAT indicate better mixing and better performance of the MCMC sampler. The IAT can thus be

**Fig. 1** Simulated data with decreasing distinguishability



used as a measure of efficiency for MCMC algorithms. See for example Liu (2004, Chapter 5) and Roberts (1996).

The results are shown in Table 1. The  $\theta_{kg}$  column is coded A for  $\theta_{kg} \sim \text{Uniform}[0, 1]$ , B for transformation to  $[0.2, 0.8]$  and C for  $[0.3, 0.7]$ . As the noise in the data increases, the ability to identify the model which generated the data de-

**Table 1** Results of simulation experiment. The PMP gives the posterior model probability of the generating model.  $\hat{\tau}$  is the estimated IAT


$(K, G)$	$\theta_{kg}$	PMP	$\hat{\tau}$
(4, 4)	A	0.9550	8.79
	B	0.9463	10.57
	C	0.9014	17.43
(2, 5)	A	0.9343	4.55
	B	0.8886	9.79
	C	0.8369	13.66
(1, 4)	A	0.8035	7.86
	B	0.3000	8.97
	C	0.1494	4.61

creases. This is to be expected. The estimated IAT indicates that we get less efficient sampling as the noise increases, with the exception of the (1, 4) cluster model. This is a particularly challenging situation, since two of the clusters are very similar (see Fig. 1). The transformed  $\theta_{11}, \theta_{13}$  were for  $[0.2, 0.8] : 0.288, 0.320$  and for  $[0.3, 0.7] : 0.36, 0.38$ . The fact that these two clusters are practically indistinguishable would make the sampler choose a (1, 3) model as the best model after scrambling of the data. In fact the most visited model in both these cases had 1 row cluster and 3 column clusters (62.37% and 82.66% of the posterior probability). In this situation, the best cluster model was not the same as the generating model. This is an artifact of the simulation process but shows that a sensible clustering can be achieved.

## 5 Congressional voting in US senate

We apply the sampler to the UCI Congressional Voting data assuming the Bernoulli model of Sect. 2.2.1. The data records whether 435 members of the 98th congress (267

**Fig. 2** (Color online) Voting data. Color key: white = “nay”, black = “yay”. Left panel: Raw data. Right panels: summary cluster membership from the modal 7 row and 12 column cluster model and the cluster membership obtained from BEM2. Row clusters are numbered 1–7, top–bottom. Column clusters are numbered left–right. The red lines divide clusters of congressmen, and the blue lines divide the issue clusters





**Table 2** Distribution of cluster models for voting data

Rows	Columns							
	9	10	11	12	13	14	15	16
5	0.0000	0.0002	0.0013	0.0008	0.0005	0.0000	0.0002	0.0000
6	0.0001	0.0021	0.0401	0.1238	0.1491	0.0946	0.0348	0.0029
7	0.0001	0.0041	0.0543	0.1675	0.1614	0.0697	0.0188	0.0012
8	0.0000	0.0009	0.0101	0.0270	0.0191	0.0075	0.0029	0.0005
9	0.0000	0.0000	0.0010	0.0013	0.0013	0.0007	0.0001	0.0000

democrats, 168 republicans) voted “yay”, “nay”, abstained or were absent in votes on 16 different key issues. Here the members of congress are represented by rows, and the issues are represented by columns. The data is available from <http://archive.ics.uci.edu/ml/datasets/Congressional+Voting+Records> and is shown in the left panel of Fig. 2. The aim was to see whether the sampler could discover any clustering by party and issue. For example, one may expect that democrats voted differently to republicans on certain issues. It was thought best to ignore absent and abstain votes. Here, this is equivalent to treating these votes as a “nay”, since our focus is on clustering rows and columns. The only sample sizes entering into our calculations are the number of rows and columns in each cluster. For the Bernoulli model, the block sufficient statistic is the sum of the data. This is not affected by a missing data point. We do point out however that missing data could be easily imputed by inclusion of a Gibbs step to sample from the full conditional of any missing point.

The sampler was run for 110 000 iterations with 10 000 as a burn-in initialized at the no cluster model. To reduce correlation in samples, we took every 10th sample after burn-in. The move of Sect. 3.1.2 had a 16% acceptance rate for rows and 51% for columns. The cluster split and combine moves had about a 1.5% acceptance rate for rows and 8% for columns. The run took just over an hour on a 2.5 GHz processor.

Table 2 shows the distribution of the number of row and column clusters. It can be seen that about 60% of the posterior probability is placed on 6/7 row clusters and 12/13 column clusters. We extracted the samples with 7 row clusters and 12 column clusters to construct an estimated clustering following Section 3.4.1. The estimated clustering is shown in the middle panel of Fig. 2. The red horizontal lines here divide the clusters of congressmen, and the blue vertical lines divide the issue clusters. When referencing Fig. 2 we say that the congressman clusters (rows) are numbered 1 to 7 top-bottom, and the issue clusters (columns) are numbered 1 to 12 left-right.

Issues have only three non-singleton clusters. The first contains “anti-satellite-test-ban”, “aid-to-nicaraguan-contras” and “mx-missile” (column cluster 1). The second

**Table 3** Party distribution over row clusters from collapsed sampling

Cluster	Democrat	Republican
1 (131)	8	123
2 (125)	125	0
3 (77)	71	6
4 (38)	37	1
5 (36)	3	33
6 (23)	21	2
7 (5)	2	3

**Table 4** Party distribution over row clusters from BEM2 algorithm

Cluster	Democrat	Republican
1 (131)	8	123
2 (104)	104	0
3 (62)	61	1
4 (60)	50	10
5 (35)	5	30
6 (30)	26	4
7 (13)	13	0

contains “physician-fee-freeze” and “education-spending” (column cluster 3) and the third has “handicapped-infants” and “duty-free-exports” (column cluster 6). Row cluster composition by party is shown in Table 3. The majority democrat party roughly splits into four clusters, while the republican party splits into two. The main discrepancy between the two large democrat clusters, 2 and 3, appear to be the issues “religious-groups-in-schools” and “crime” in issue clusters 9 and 12. Row cluster 6 which is also mainly democrat appears to vote similarly to the republican cluster 1. Row clusters 4 (democrat) and 5 (republican) appear to deviate from their core party vote.

We compared the results obtained from our algorithm with those obtained from the BEM2 algorithm of Govaert and Nadif (2008), reviewed in Sect. 2.1.1. The algorithm was run using 7 row clusters and 12 column clusters. It should be noted that BEM2 requires the number of row and column clusters to be assumed known in advance. To obtain


an estimated clustering, we took the cluster with the maximum probability of membership. The composition by party of the row clusters from BEM2 is shown in Table 4. Row cluster 1 is similar in both, but there are some differences in the other clusters. The BEM2 clustering only used 10 column clusters of the 12 available. The right panel of Fig. 2 shows the clustering from BEM2. For comparison purposes with the collapsed LBM clustering, the columns have been arranged in the same order. The collapsed LBM appears to identify more small clusters, leading to a marginally more homogeneous blocking of the data.

There is an advantage here over an EM approach to fitting a cluster model in that the number of clusters need not be assumed known in advance. Here, model uncertainty is naturally inbuilt into the approach and it is dealt with automatically by the sampler. There is no user intervention to choose the cluster model. User intervention is only in the choice of prior hyperparameters and priors on the number of clusters. In our experience, standard non informative priors for the data model parameters, and the truncated Poisson(1) prior on the number of clusters (as argued in Nobile 2005), seem to perform well. The computations for the collapsed LBM are also numerically stable if clusters empty out. In our experimentation with EM algorithms this caused instability. Empty clusters could easily occur, say, if the chosen cluster model is not well supported by the data.

## 6 Microarray experiments

A DNA microarray experiment records expression levels of a large number of genes over a number of conditions or samples. The number of conditions or samples is usually less than 100, while the number of genes could be in the thousands. Discovering which genes behave similarly and under which subgroups of conditions is the aim of analysis. One way to do this is to group together genes with similar expression levels. Methods differ in whether they allow clusters to overlap or not. Here we will not allow clusters to overlap due to the form of the LBM.

Analyzing DNA experiments can be challenging, due to the large row dimension and the general uncertainty in how many clusters may be present in the data. We apply our sampler to data from DNA experiments on the budding yeast *Saccharomyces Cerevisiae*. The microarray contains 419 genes and records the expression level of these under 70 conditions. It was obtained from the R package `bi-clust` (Kaiser et al. 2009). Expression levels lay between  $-6$  and  $+7$ . The aim is to see how much structure the sampler can uncover, so the rows and columns of the microarray were randomly reordered (Fig. 3(a)). In our application the rows represent the genes and the columns represent the conditions.



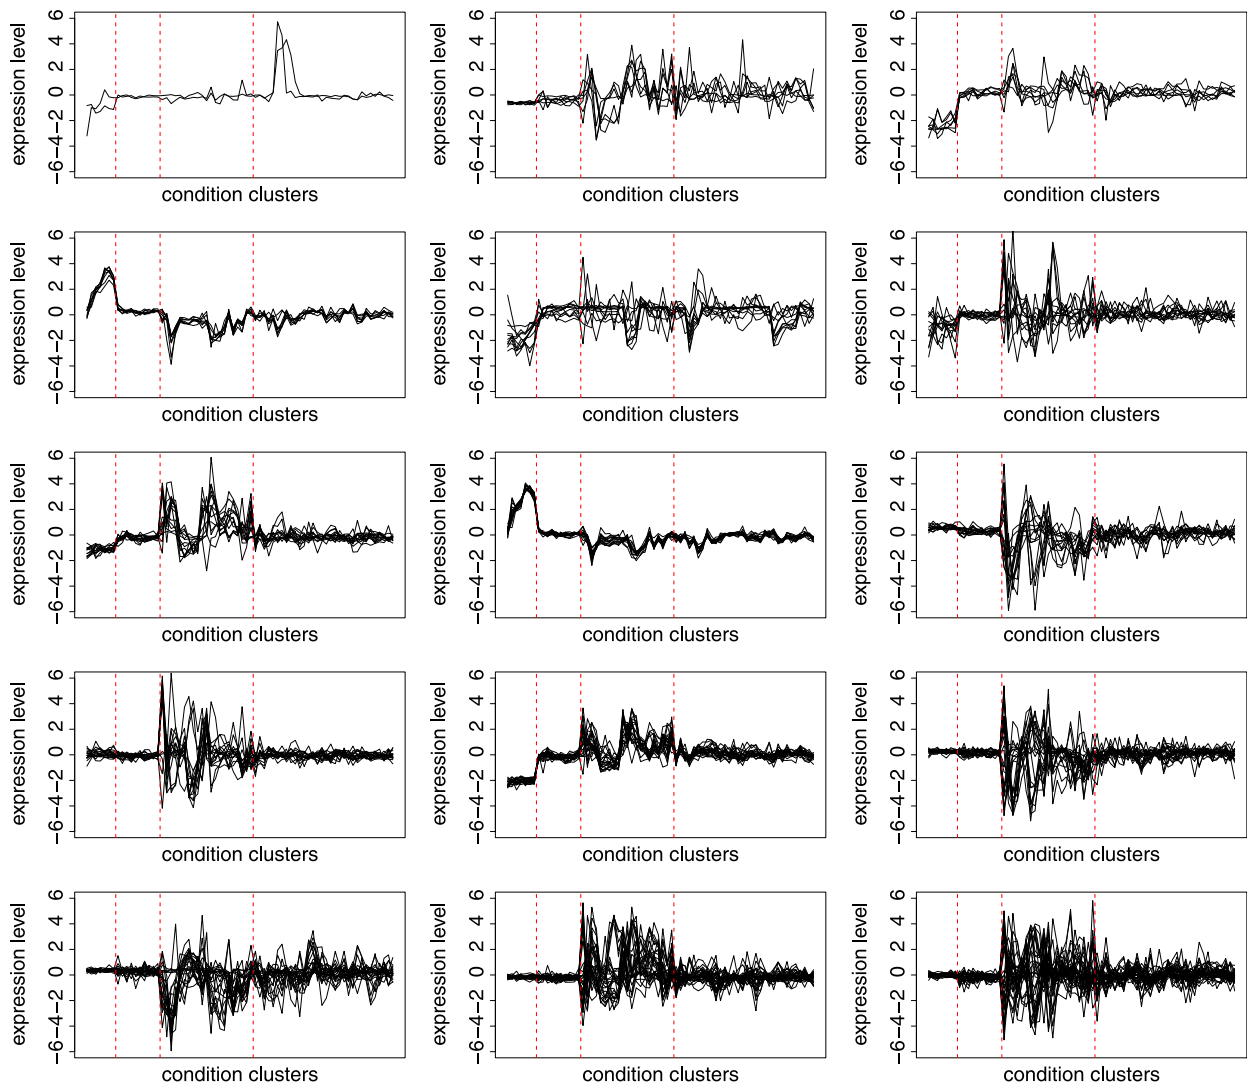
**Fig. 3** Yeast data. (a) Original microarray (b) MAP clustering from sampler

We use the Gaussian model of Sect. 2.2.2 for expression level. This model requires specification of four hyperparameters. Two of these ( $\gamma$  and  $\delta$ ) are for the prior on the block error variances and two are for the prior on the block means ( $\xi$  and  $\tau^2$ ). We choose  $\gamma = \delta = 0.02$  and  $\xi = 0$ ,  $\tau^2 = 100$ . This choice of  $\gamma$  and  $\delta$  gives a proper density on the error variance which is non-informative (see for example Spiegelhalter et al. 1996). Similarly, choosing  $\xi = 0$  is a reasonable non-informative choice given the range of the data. Setting  $\tau^2 = 100$  says that the prior information on a block mean is equal to 1% of the information in the observed expression level of one gene under one condition within that block. This is also non-informative. The prior on the number of row and column clusters was again taken to be truncated Poisson(1) with  $K_{\max} = 200$  and  $G_{\max} = 70$ .

The sampler was run for 220 000 iterations with 20 000 taken as burn-in. We stored every 20th iteration thereafter. The run was time consuming, taking approximately 3 hours. This said, the large gene dimension of such an array does pose a challenge when searching for two way clusters. The initial cluster model assumed had 1 row and column cluster i.e. no cluster structure. Acceptance rates for the move of Sect. 3.1.2 were 25% for rows and 18% for columns. Split and combine acceptances were about 0.5% for rows and about 25% each for columns. The low acceptance rates of split and combine moves for rows would be expected since finding clusters will be more difficult in a larger dimension.

**Table 5** Posterior distribution of cluster models for the microarray data

Rows	Columns											
	3	4	5	6	7	8	9	10	11	12	13	
18	0.0000	0.0001	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
19	0.0005	0.0002	0.0005	0.0001	0.0019	0.0044	0.0013	0.0011	0.0006	0.0000	0.0000	0.0000
20	0.0005	0.0000	0.0002	0.0032	0.0110	0.0192	0.0121	0.0057	0.0012	0.0004	0.0000	0.0000
21	0.0029	0.0027	0.0019	0.0023	0.0128	0.0147	0.0090	0.0028	0.0008	0.0003	0.0001	0.0001
22	0.0076	0.0119	0.0074	0.0053	0.0109	0.0129	0.0085	0.0033	0.0007	0.0000	0.0001	0.0001
23	0.0368	0.0409	0.0242	0.0094	0.0089	0.0041	0.0030	0.0015	0.0002	0.0000	0.0000	0.0000
24	0.0643	0.0706	0.0420	0.0144	0.0073	0.0049	0.0023	0.0005	0.0002	0.0000	0.0000	0.0000
25	0.1016	0.1198	0.0702	0.0272	0.0082	0.0018	0.0011	0.0003	0.0002	0.0000	0.0000	0.0000
26	0.0373	0.0457	0.0228	0.0095	0.0020	0.0008	0.0003	0.0000	0.0000	0.0000	0.0000	0.0000
27	0.0090	0.0109	0.0058	0.0021	0.0008	0.0002	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
28	0.0008	0.0015	0.0010	0.0002	0.0002	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000



**Fig. 4** (Color online) Selection of row clusters from the MAP clustering. Each plot corresponds to a different row cluster. Each profile (black line) gives the gene expression level for over all conditions.

Conditions are arranged by condition cluster membership. The condition clusters are separated by the red dashed line

Table 5 gives the PMP of the visited models from the MCMC output. The model space visited by the sampler is large. The modal model (25 row clusters and 4 column clusters) gave a posterior probability of 11.98%. There is posterior support for anything from 3 to 5 column clusters, for 23 to 26 row clusters (PMP > 0.02 in all cases). This is an example with considerable model uncertainty and it may be difficult to know the models to include using an information criterion over a grid of possible models as discussed in Sect. 2.1.2 and adopted by van Dijk et al. (2009). Our approach has the obvious advantage of exploring the uncertainty in the posterior model space and attaching a probability to each model.

Instead of constructing a summary based on the modal model, we took the MAP clustering here (shown in Fig. 3(b)). The MAP had 26 row clusters and 4 column clusters. To have a closer look at the row clusters, we plot a selection of these in Fig. 4. The plots show the gene expression profiles for genes in the same row cluster over conditions arranged by condition cluster. It can be seen that in certain cases, there is a clear clustering of genes with similar profiles. The gene clusters shown are arranged by size (left–right, top–bottom). Some of the larger clusters appear quite noisy, while some follow a common trend closely. Auxiliary runs of the sampler on the subsets of row clusters could be performed to try and isolate further cluster structures.

### 7 Conclusion

We have considered a collapsed Bayesian extension of the Latent Block Model of Govaert and Nadif (2008). We showed how an MCMC sampler could be used to sample both the cluster model and the cluster memberships when clustering a data matrix into blocks. The approach was demonstrated on simulated data and two real data examples. The application to simulated data suggested that the sampler’s performance deteriorates as clusters become less distinguishable. We applied the sampler to Congressional voting records from the U.S. senate. It was shown to perform well in isolating clusters of congressmen and “yay”, “nay” votes in the data. In the second real data example, we used the sampler for analysis of a DNA microarray experiment. This demonstrated that there can be considerable uncertainty in the number of clusters in certain situations. Knowing even the range of possible models may be difficult. Results from the microarray experiment demonstrated that clusters could be found by a search strategy with a probabilistic basis using the collapsed LBM sampler. Overall, the approach seems to be a robust way to block cluster a data matrix. The user need only specify prior hyperparameters and priors on the number of clusters. Code implementing the sampler written in the C language is available at

[www.ucd.ie/statdept/jwyse](http://www.ucd.ie/statdept/jwyse). This code can be easily modified to experiment with different priors on the number of clusters.

**Acknowledgements** The authors would like to thank Brendan Murphy for helpful advice on aspects of this paper. Nial Friel’s research was supported by a Science Foundation Ireland Research Frontiers Program grant, 09/RFP/MTH2199.

### Appendix A: Calculation of the posterior

Writing out all posterior terms longhand, assuming  $\omega \sim \text{Dirichlet}(\alpha, \dots, \alpha)$  and  $\rho \sim \text{Dirichlet}(\beta, \dots, \beta)$  a priori gives

$$\begin{aligned} \pi(K, G, \mathbf{z}, \mathbf{w}, \omega, \rho, \Theta|Y) &\propto \pi(K)\pi(G) \frac{\Gamma\{\alpha K\} \Gamma\{\beta G\}}{\Gamma\{\alpha\}^K \Gamma\{\beta\}^G} \\ &\times \prod_{k=1}^K \omega_k^{n_k} \prod_{g=1}^G \rho_g^{m_g} \prod_{k=1}^K \prod_{g=1}^G \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij}|\theta_{kg}) \\ &\times \prod_{k=1}^K \prod_{g=1}^G \pi(\theta_{kg}) \prod_{k=1}^K \omega_k^{\alpha-1} \prod_{g=1}^G \rho_g^{\beta-1} \\ &= \pi(K)\pi(G) \frac{\Gamma\{\alpha K\} \Gamma\{\beta G\}}{\Gamma\{\alpha\}^K \Gamma\{\beta\}^G} \prod_{k=1}^K \omega_k^{n_k+\alpha-1} \prod_{g=1}^G \rho_g^{m_g+\beta-1} \\ &\times \prod_{k=1}^K \prod_{g=1}^G \pi(\theta_{kg}) \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij}|\theta_{kg}). \end{aligned}$$

Integrating the left and right hand sides of the above with respect to  $\omega, \rho$  and  $\Theta$  gives

$$\begin{aligned} \pi(K, G, \mathbf{z}, \mathbf{w}|Y) &\propto \pi(K)\pi(G) \frac{\Gamma\{\alpha K\} \prod_{k=1}^K \Gamma\{n_k + \alpha\}}{\Gamma\{\alpha\}^K \Gamma\{n + \alpha K\}} \\ &\times \frac{\Gamma\{\beta G\} \prod_{g=1}^G \Gamma\{m_g + \beta\}}{\Gamma\{\beta\}^G \Gamma\{m + \beta G\}} \prod_{k=1}^K \prod_{g=1}^G M_{kg} \end{aligned}$$

where

$$M_{kg} = \int \pi(\theta_{kg}) \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij}|\theta_{kg}) d\theta_{kg}.$$

**Appendix B: Calculation of  $M_{kg}$**

**B.1 Bernoulli model for binary data**

Assume that  $\Pr(y_{ij} = 1|z_i = k, w_j = g) = \theta_{kg}$ . We take a Beta( $\gamma, \delta$ ) prior on  $\theta_{kg}$ . Then

$$\begin{aligned} \pi(\theta_{kg}) & \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij}|\theta_{kg}) \\ &= \frac{\Gamma\{\gamma + \delta\}}{\Gamma\{\gamma\}\Gamma\{\delta\}} \theta_{kg}^{\gamma-1} (1 - \theta_{kg})^{\delta-1} \\ & \times \prod_{i:z_i=k} \prod_{j:w_j=g} \theta_{kg}^{y_{ij}} (1 - \theta_{kg})^{1-y_{ij}} \\ &= \frac{\Gamma\{\gamma + \delta\}}{\Gamma\{\gamma\}\Gamma\{\delta\}} \theta_{kg}^{s_{kg} + \gamma - 1} (1 - \theta_{kg})^{n_k m_g - s_{kg} + \delta - 1}. \end{aligned}$$

Integrating the left and right hand side of this with respect to  $\theta_{kg}$  gives

$$M_{kg} = \frac{\Gamma\{\gamma + \delta\}}{\Gamma\{\gamma\}\Gamma\{\delta\}} \frac{\Gamma\{s_{kg} + \gamma\} \Gamma\{n_k m_g - s_{kg} + \delta\}}{\Gamma\{n_k m_g + \gamma + \delta\}}.$$

**B.2 Gaussian model for continuous data**

Assume  $y_{ij}|z_i = k, w_j = g \sim N(\mu_{kg}, \sigma_{kg}^2)$ . Take the priors  $\mu_{kg} \sim N(\xi, \tau^2 \sigma_{kg}^2)$  and  $\sigma_{kg}^2 \sim \text{IG}(\delta/2, \gamma/2)$  where  $\text{IG}(a, b)$  is the Inverse-Gamma distribution:  $p(x) = \frac{b^a}{\Gamma(a)} x^{-(a+1)} \times \exp\{-b/x\}$ . Then

$$\begin{aligned} \pi(\mu_{kg})\pi(\sigma_{kg}^2) & \prod_{i:z_i=k} \prod_{j:w_j=g} p(y_{ij}|\mu_{kg}, \sigma_{kg}^2) \\ &= \frac{(\gamma/2)^{\delta/2}}{\Gamma\{\delta/2\}} \sigma_{kg}^{-2(\delta/2+1)} \exp\{-\gamma/2\sigma_{kg}^2\} \\ & \times (2\pi \tau^2 \sigma_{kg}^2)^{-1/2} \exp\{-(\mu_{kg} - \xi)^2/2\tau^2 \sigma_{kg}^2\} \\ & \times (2\pi \sigma_{kg}^2)^{-n_k m_g/2} \\ & \times \exp\left\{-\frac{(s_{kg} - 2\mu_{kg} s_{kg} + n_k m_g \mu_{kg}^2)/2\sigma_{kg}^2}{\sigma_{kg}^2}\right\}. \end{aligned}$$

Completing the square on  $\mu_{kg}$  and integrating with respect to it gives

$$\begin{aligned} (2\pi)^{-n_k m_g/2} & \frac{(\gamma/2)^{\delta/2}}{\Gamma\{\delta/2\}} \sigma^{-2((n_k m_g + \delta)/2 + 1)} (n_k m_g \tau^2 + 1)^{-1/2} \\ & \times \exp\left\{-\frac{1}{2\sigma_{kg}^2} \left( s_{kg} - \frac{\tau^2(s_{kg} + \xi/\tau^2)^2}{n_k m_g \tau^2 + 1} + \frac{\xi^2}{\tau^2} + \gamma \right)\right\}. \end{aligned}$$

Finally, integrating with respect to  $\sigma_{kg}^2$  and tidying up gives

$$\begin{aligned} M_{kg} &= \frac{\gamma^{\delta/2} \Gamma\{(n_k m_g + \delta)/2\}}{\pi^{n_k m_g/2} \Gamma\{\delta/2\} (n_k m_g \tau^2 + 1)^{1/2}} \\ & \times \left( s_{kg} - \frac{\tau^2(s_{kg} + \xi/\tau^2)^2}{n_k m_g \tau^2 + 1} + \frac{\xi^2}{\tau^2} + \gamma \right)^{-(n_k m_g + \delta)/2}. \end{aligned}$$

**Appendix C: Cost matrix for undoing label switching**

The cost matrix for processing the vector  $\mathbf{z}^{(T)}$  is

$$C(k_1, k_2) = \sum_{t=1}^{T-1} \sum_{i=1}^n \mathbf{I}\{z_i^{(t)} \neq k_1, z_i^{(T)} = k_2\}.$$

The more  $\mathbf{z}^{(T)}$  disagrees with the vectors already processed, the higher this cost will be. This is made clearer by rewriting the general entry of the cost matrix:

$$\begin{aligned} C(k_1, k_2) &= \sum_{t=1}^{T-1} \sum_{i=1}^n \left( 1 - \mathbf{I}\{z_i^{(t)} = k_1, z_i^{(T)} = k_2\} \right) \\ &= n(T-1) - \sum_{t=1}^{T-1} \sum_{i=1}^n \mathbf{I}\{z_i^{(t)} = k_1, z_i^{(T)} = k_2\} \\ &= n(T-1) - \sum_{i=1}^n N_i(T-1, k_1) \mathbf{I}\{z_i^{(T)} = k_2\} \end{aligned} \tag{C.1}$$

where  $N_i(T-1, k_1)$  gives the number of processed samples up to  $z^{(T-1)}$  which have given label  $k_1$  to row  $i$ . For the sake of discussion, consider processing a sample where no label switching has occurred,  $K$  is fixed and there are no changes in labels from one MCMC sample to the next. In this case the costs will be

$$\begin{aligned} C(k, k) &= n(T-1) - \sum_{i=1}^n N_i(T-1, k) \mathbf{I}\{z_i^{(T)} = k\} \\ &= n(T-1) - n_k(T-1) \\ &= (n - n_k)(T-1) \end{aligned}$$

and for  $k' \neq k$

$$\begin{aligned} C(k, k') &= n(T-1) - \sum_{i=1}^n N_i(T-1, k) \mathbf{I}\{z_i^{(T)} = k'\} \\ &= n(T-1) - 0 \\ &= n(T-1). \end{aligned}$$

A cost of 0 could only be obtained when all rows have the same label, that is, when there is no clustering. Of course this discussion simplifies the problem somewhat. The key is in finding a permutation of the labels to minimize all costs.

This permutation is found by solving the square assignment problem using the algorithm of Carpaneto and Toth (1980) in our case. Carpaneto et al. (1988) give an extensive review of the assignment problem and algorithms for its solution. Finally, we note that (C.1) can be exploited to give an online post processing procedure. Define the  $K \times n$  matrix  $S^{(T-1)}$  with general entry  $S_{ki}^{(T-1)} = \sum_{t=1}^{T-1} I\{z_i^{(t)} = k\}$ . Then we have  $C(k, k') = n(T-1) - \sum_{i=1}^n S_{ki}^{(T-1)} I\{z_i^{(T)} = k'\}$ . After calling the square assignment algorithm and permuting the labels  $\mathbf{z}^{(T)}$  according to its solution, we can update  $S$ , using  $S_{ki}^{(T)} = S_{ki}^{(T-1)} + I\{z_i^{(T)} = k\}$ .

## References

- Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022 (2003)
- Bozdogan, H.: Mixture-model cluster analysis using model selection criteria and a new information measure of complexity. In: Bozdogan, H. (ed.) *Proceedings of the first US/Japan Conference on the Frontiers of Statistical Modeling: An Informational Approach*, vol. 2, pp. 69–113. Kluwer Academic, Boston (1994)
- Brooks, S.P., Giudici, P., Roberts, G.O.: Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions (with discussion). *J. R. Stat. Soc., Ser. B, Stat. Methodol.* **65**, 3–39 (2003)
- Carpaneto, G., Martello, S., Toth, P.: Algorithms and codes for the assignment problem. *Ann. Oper. Res.* **13**, 193–223 (1988)
- Carpaneto, G., Toth, P.: Algorithm 548: Solution of the assignment problem. *ACM Trans. Math. Softw.* **6**, 104–111 (1980)
- Celeux, G., Hurn, M., Robert, C.P.: Computational and inferential difficulties with mixtures posterior distribution. *J. Am. Stat. Assoc.* **95**, 957–979 (2000)
- Cheng, Y., Church, G.M.: Biclustering of expression data. In: *ISMB 2000 Proceedings*, pp. 93–103 (2000)
- Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm (with discussion). *J. R. Stat. Soc., Ser. B, Stat. Methodol.* **39**, 1–38 (1977)
- Fraley, C., Raftery, A.E.: Model-based clustering, discriminant analysis and density estimation. *J. Am. Stat. Assoc.* **97**, 611–631 (2002)
- Getz, G., Levine, E., Domany, E.: Coupled two-way clustering analysis of gene microarray data. *Proc. Natl. Acad. Sci. USA* **97**, 12079–12084 (2000)
- Govaert, G., Nadif, M.: Block clustering with Bernoulli mixture models: Comparison of different approaches. *Comput. Stat. Data Anal.* **52**, 3233–3245 (2008)
- Green, P.: Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, 711–732 (1995)
- Green, P.J.: Trans-Dimensional Markov chain Monte Carlo. In: Green, P.J., Hjord, N.L., Richardson, S. (eds.) *Highly Structured Stochastic Systems*, pp. 179–198. Oxford University Press, Oxford (2003)
- Griffiths, T.L., Steyvers, M.: Finding scientific topics. *Proc. Natl. Acad. Sci. USA* **101**, 5228–5235 (2004)
- Hartigan, J.A.: Direct clustering of a data matrix. *J. Am. Stat. Assoc.* **67**, 123–129 (1972)
- Hartigan, J.A.: Bloc voting in the United States senate. *J. Classif.* **17**, 29–49 (2000)
- Hofmann, T.: Unsupervised learning by probabilistic latent semantic analysis. *Mach. Learn.* **42**, 177–196 (2001)
- Kaiser, S., Santamaria, R., Sill, M., Theron, R., Quintales, L., Leisch, F.: Biclust: BiCluster Algorithms. R package version 0.9.1. (2009)
- Kluger, Y., Basri, R., Chang, J.T., Gerstein, M.: Spectral Biclustering of microarray data: Co-clustering genes and conditions. *Genome Res.* **13**, 703–716 (2003)
- Lazzeroni, L., Owen, A.: Plaid models for gene expression data. *Stat. Sin.* **12**, 61–86 (2002)
- Liu, J.S.: *Monte Carlo Strategies in Scientific Computing*. Springer, New York (2004)
- Neal, R.M., Hinton, G.E.: A view of the EM algorithm that justifies incremental, sparse, and other variants. In: Jordan, M.I. (ed.) *Learning in Graphical Models*, pp. 355–358. Kluwer Academic, Dordrecht (1998)
- Nobile, A.: Bayesian finite mixtures: a note on prior specification and posterior computation. Technical report. Department of Statistics, University of Glasgow (2005)
- Nobile, A., Fearnside, A.T.: Bayesian finite mixtures with an unknown number of components: The allocation sampler. *Stat. Comput.* **17**, 147–162 (2007)
- Phillips, D.B., Smith, A.F.M.: Bayesian model comparison via jump diffusions. In: Gilks, W.R., Richardson, S., Spiegelhalter, D.J. (eds.) *Markov Chain Monte Carlo in Practice*, pp. 215–239. Chapman & Hall, London (1996)
- Richardson, S., Green, P.J.: On Bayesian analysis of mixtures with an unknown number of components (with discussion). *J. R. Stat. Soc. B* **59**, 731–792 (1997)
- Robert, C.P., Rydén, T., Titterton, D.M.: Bayesian inference in hidden Markov models through the reversible jump Markov chain Monte Carlo method. *J. R. Stat. Soc. B* **62**, 57–76 (2000)
- Roberts, G.O.: Markov chain concepts related to sampling algorithms. In: Gilks, W.R., Richardson, S., Spiegelhalter, D.J. (eds.) *Markov Chain Monte Carlo in Practice*, pp. 45–58. Chapman & Hall, London (1996)
- Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978)
- Sheng, Q., Moreau, Y., Moor, B.D.: Biclustering microarray data by Gibbs sampling. *Bioinformatics* **19**, 196–205 (2003)
- Spiegelhalter, D.J., Best, N.G., Gilks, W.R., Inskip, H.: Hepatitis B: a case study in MCMC methods. In: Gilks, W.R., Richardson, S., Spiegelhalter, D.J. (eds.) *Markov Chain Monte Carlo in Practice*, pp. 21–44. Chapman & Hall, London (1996)
- Stephens, M.: Bayesian analysis of mixture models with an unknown number of components- an alternative to reversible jump methods. *Ann. Stat.* **28**, 40–74 (2000)
- Tibshirani, R., Hastie, T., Eisen, M., Ross, D., Botstein, D., Brown, P.: Clustering methods for the analysis of DNA microarray data. Technical report, Stanford University (1999)
- van Dijk, B., van Rosmalen, J., Paap, R.: A Bayesian approach to two-mode clustering. Technical report, Econometric Institute Report, Erasmus University Rotterdam (2009)
- Wit, E., McClure, J.: *Statistics for Microarrays: Design, Analysis and Inference*. Wiley, New York (2004)