

Bayesian Clustering of Transcription Factor Binding Motifs

Shane T. JENSEN and Jun S. LIU

Genes are often regulated in living cells by proteins called transcription factors that bind directly to short segments of DNA in close proximity to specific genes. These binding sites have a conserved nucleotide appearance, which is called a motif. Several recent studies of transcriptional regulation require the reduction of a large collection of motifs into clusters based on the similarity of their nucleotide composition. We present a principled approach to this clustering problem based on a Bayesian hierarchical model that accounts for both within- and between-motif variability. We use a Dirichlet process prior distribution that allows the number of clusters to vary and we also present a novel generalization that allows the core width of each motif to vary. This clustering model is implemented, using a Gibbs sampling strategy, on several collections of transcription factor motif matrices. Our stochastic implementation allows us to examine the variability of our results in addition to focusing on a set of best clusters. Our clustering results identify several motif clusters that suggest that several transcription factor protein families are actually mixtures of several smaller groups of highly similar motifs, which provide substantially more refined information compared with the full set of motifs in the family. Our clusters provide a means by which to organize transcription factors based on binding motif similarities and can be used to reduce motif redundancy within large databases such as JASPAR and TRANSFAC, which aids the use of these databases for further motif discovery. Finally, our clustering procedure has been used in combination with discovery of evolutionarily conserved motifs to predict co-regulated genes. An alternative to our Dirichlet process prior distribution is presented that differs substantially in terms of a priori clustering characteristics, but shows no substantive difference in the clustering results for our dataset. Despite our specific application to transcription factor binding motifs, our Bayesian clustering model based on the Dirichlet process has several advantages over traditional clustering methods that could make our procedure appropriate and useful for many clustering applications.

KEY WORDS: Bayesian hierarchical modeling; Dirichlet process; Gibbs sampling; Motif clustering.

1. INTRODUCTION

1.1 Transcription Factor Binding Motifs

The complete information that defines the characteristics of living cells within an organism is encoded in the form of a moderately simple molecule, deoxyribonucleic acid (DNA). The building blocks of DNA are four nucleotides, abbreviated by their attached organic bases as A, C, G, and T. The pairs A–T and C–G are complementary bases between which hydrogen bonds can form. A DNA molecule consists of two long chains of nucleotides that are complementary to each other and joined by hydrogen bonds twisted into a double helix. This structure gives rise to the term *base pair* when describing a DNA sequence. A long DNA chain in a living cell is called a *chromosome*. For example, every human cell has 23 pairs of chromosomes with lengths ranging from 47 million to 245 million base pairs. The specific ordering of the four types of nucleotides along these chains is the means by which the information that completely defines all functions within a cell is stored. The term *gene* refers to sequence segments along a chromosome that are used to code the information for making *proteins*, the fundamental action molecules of the cell. Surprisingly, only about 1–2% of the entire human genome (set of chromosomes) corresponds to gene regions. It is believed that much of the mechanism for controlling when, where, and how much protein will be produced is located in the “noncoding” region located *upstream* (i.e., directly before) the gene sequence.

Transcribing or activating a gene requires not only the DNA sequence in the upstream region, but also many proteins called

transcription factors (TFs). When these TFs are present, they bind to specific DNA patterns in the upstream sequence of genes, and either induce or repress the transcription of these genes by recruiting other necessary proteins (Lodish et al. 1995). A particular transcription factor protein is able to bind and regulate only certain target genes by recognizing a short (6–20 base pair long) sequence of nucleotides called a transcription factor binding site (or, more simply, a *site*). Different binding sites (located near different genes) of the same transcription factor protein show a substantial sequence conservation, which we call a *motif*, but some variability is also present.

1.2 Statistical Formulation of Motifs

Each motif is mathematically formulated as a *motif matrix*, which measures the desirability of each base at each position of the motif. The simplest matrix is an alignment or count matrix \mathbf{N}_{jk} , which records the occurrence of base k at position j of all the sites for this motif. Table 1 shows the count matrix for motif MA0011 from the database we will use in Section 3. Also shown in Table 1 is the corresponding frequency matrix ($f_{jk} = \mathbf{N}_{jk}/N$, where N is the number of motif sites) for motif MA0011.

Schneider and Stephens (1990) used the motif matrix to construct a *sequence logo* as a means by which to visualize the appearance of the motif. Figure 1 gives the sequence logo for the same motif MA0011 as constructed by the program WebLogo (Crooks, Hon, Chandonia, and Brenner 2004). The height of each position is equal to its *information content* ($\sum_k f_{jk} \log[f_{jk}/\theta_{0k}]$, where θ_{0k} is the proportion of base k in the nonmotif background positions) and the size of each letter is proportional to that letter’s relative frequency f_{jk} .

Shane T. Jensen is Assistant Professor, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104 (E-mail: stjensen@wharton.upenn.edu). Jun S. Liu is Professor, Department of Statistics, Harvard University, Cambridge, MA 02138 (E-mail: jliu@stat.harvard.edu). The authors thank Geetu Tuteja for help with the JASPAR and TRANSFAC databases. Liu was supported in part by NSF DMS-02-04674, NSF 10228102, and NIH R01-GM078990, funded through the NSF–NIH joint program on mathematical biology.

Table 1. Matrix representations of the motif MA0011

Count matrix									Frequency matrix								
Pos	1	2	3	4	5	6	7	8	Pos	1	2	3	4	5	6	7	8
A	3	5	0	0	12	1	2	1	A	.25	.42	.00	.00	1.00	.08	.17	.08
C	1	2	10	1	0	1	0	2	C	.08	.17	.83	.08	.00	.08	.00	.17
G	1	1	0	0	0	2	1	1	G	.08	.08	.00	.00	.00	.17	.08	.08
T	7	4	2	11	0	8	9	8	T	.58	.33	.17	.92	.00	.67	.75	.67

1.3 Discovery of Motifs

As reviewed by Jensen, Liu, Zhou, and Liu (2004), Bayesian motif discovery models are the foundation of many popular programs for discovering conserved binding sites in large sequence datasets. These models are usually implemented by an iterative strategy, such as the expectation-maximization (EM) algorithm (Dempster, Laird, and Rubin 1977) or the Gibbs sampler (Geman and Geman 1984), that utilizes the ease with which the motif frequency matrix can be estimated if the binding site locations are known and the corresponding ease with which the binding site locations can be estimated if the motif appearance is known. Jensen et al. (2004) also discussed several constraints of these motif discovery procedures, such as assumptions of known motif width (the number of columns in the motif frequency matrix) and known abundance of binding sites. Jensen and Liu (2004) demonstrated that allowing the motif width to vary leads to more accurate results in several motif discovery applications.

Although the discovery and characterization of a single motif is often the goal of a particular biological investigation (see, e.g., Eichenberger et al. 2003), it is common for scientists to be interested in examining the similarities and differences between an entire collection of discovered TF motifs. Large collections of discovered motifs have been utilized in various applications, such as the phylogenetic discovery of co-regulated genes (Qin et al. 2003) and the prediction of synergistic relationships between transcription factors (Hannenhalli and Levy 2002). These applications each represent a highly specialized approach to the utilization of a collection of motifs and do not address the issue of a general statistical approach for the sharing of information between motifs.

1.4 Modeling Motif Similarity by Clustering

Figure 2 shows the sequence logos for four different motifs from the JASPAR database, which will be analyzed in Section 3. It is clear that these motifs all show differences in width

and appearance, but there exists some similarity or common structure within this set. One could certainly argue for grouping MA0031 and MA0032 together based on their similar appearances, although they do differ in both width and composition, but this decision is based on ad hoc personal judgment. The statistical problem of interest here is to model the common structure between these different motifs and find a principled means by which to group motifs together based on their similarity.

There are several traditional statistical techniques for clustering observations together that were reviewed by Hartigan (1975). *Hierarchical tree clustering* joins observations together into successively larger clusters based on some sort of similarity measure, whereas *K-means clustering* groups observations into a predetermined number of clusters by minimizing a within-cluster distance measure. Each of these techniques has elements that are not ideally suited for our desired goal of motif clustering. Hierarchical tree clustering requires the user to specify a distance measure between the observations (in this case, motif matrices) and for comparing motifs, it is not clear what type of simple distance metric should be used. In addition, the result of this algorithm is a tree that joins all observations together, and it is not clear where the tree should be “cut” to produce a set of clusters. Kielbasa, Gonze, and Herzel (2005) proposed two different distance metrics for the clustering of motif matrices, but needed to impose arbitrary thresholds on those distances to produce a set of clusters. Similarly, Schones, Sumazin, and Zhang (2005) presented several distance metrics for comparing motif matrices, and grouped together pairs of matrices below a *p*-value-based threshold. Situations where the number of clusters is known a priori are ideally suited for *K*-means clustering. When the number of clusters is unknown, *K*-means clustering becomes more difficult, and usually a cross-validation strategy is employed to estimate the number of clusters. For our motif clustering applications, there is very little prior intuition as to how many motifs might cluster together in a particular collection of motifs, so we seek a model that easily allows for an unknown number of clusters.




Figure 1. Sequence logo of the motif MA0011.

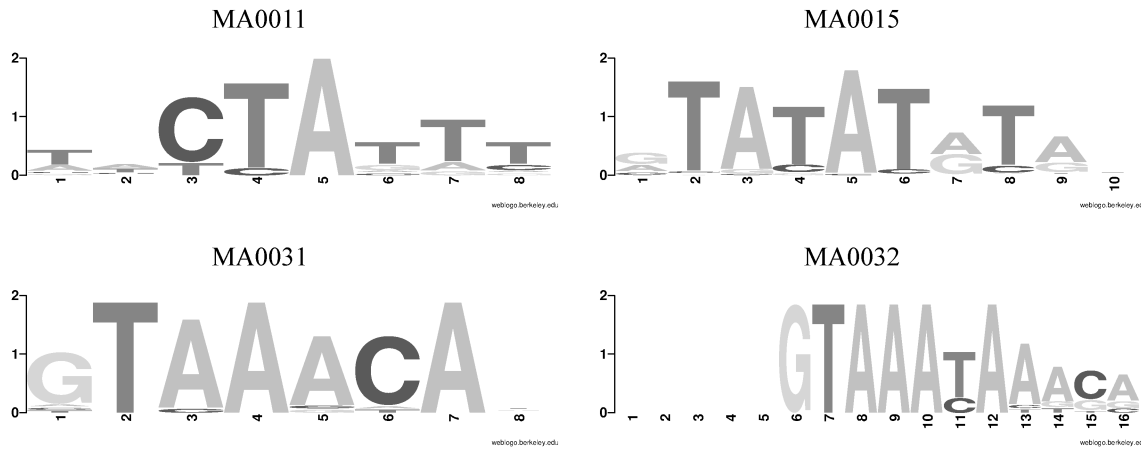


Figure 2. Four different motifs from the JASPAR database.

In addition, these techniques consider the observations themselves to be fixed and known, which is not the case for our applications, where each motif is only an estimate generated by a prior motif discovery procedure. Recognizing that our discovered motifs themselves are estimated quantities, we need to model both within-motif and between-motif variability. In Section 2 we outline a Bayesian hierarchical clustering model that encompasses both levels of variability and does not require prior knowledge of the number of clusters. As discussed briefly in Section 1.3, most motif discovery procedures assume a fixed and known motif width, but in reality the width is often unknown for many motifs and can vary substantially between different motifs. In Section 2.5 we extend our Bayesian motif clustering model to allow each motif width to be an unknown variable that will be estimated by our procedure. Our model is implemented stochastically by a Gibbs sampling algorithm, which allows us to examine not only “best estimates” of motif clusters, but also the variability within our clustering results. In Section 3 we present various techniques for summarizing and understanding the results from our clustering procedure, within the context of an application to the JASPAR database of transcription factor binding motifs.

These clusters provide a way to organize transcription factor proteins based on the similarity of their binding motifs. We utilize this clustering to analyze large databases of transcription factor matrices—TRANSFAC (Wingender et al. 2000) and JASPAR (Sandelin, Alkema, Engstrom, Wasserman, and Lenhard 2004)—for motifs with high similarity (Sec. 4) and discuss how this clustering information can be used to refine the search for additional transcription factor binding sites. In addition, we briefly discuss combining our clustering procedure with motif discovery in sequences that have been conserved by evolution, so as to predict genes that share similar transcription factor binding motifs and thus are possibly co-regulated (Sec. 5). Software and supplemental materials for our procedure are available at <http://stat.wharton.upenn.edu/~stjensen/research/cluster.html>.

2. BAYESIAN MOTIF CLUSTERING MODEL

We use a Bayesian hierarchical model to infer common structure, in the form of clusters, within a collection of discovered

motifs. The data for each discovered motif constitute a count matrix \mathbf{N}_i , which can have different widths and number of counts (i.e., number of binding sites included in the matrix) compared to other TF motifs. For now, we assume that our clustering will be based on motif matrices with a fixed and known width w , so we assume each of these n raw motif matrices \mathbf{N}_i should contain a submatrix \mathbf{Y}_i of dimension $w \times 4$ that will be considered the *core* upon which the clustering will be based. We will later extend our model to allow the core width within each cluster of motifs to vary.

2.1 Hierarchical Framework

Hierarchical models are useful in a variety of scientific problems when the structure of the data suggests multiple levels of uncertainty. We want to include components for both within-motif and between-motif variability of the nucleotide counts Y_{ijk} , where i indexes the motif, j indexes the w columns within each motif core, and k indexes the four possible nucleotides ($k = a, c, g, \text{ or } t$) within each column. Our model on the within-motif variability between different binding sites for a count motif \mathbf{Y}_i is a product multinomial model. We assume that each position (column) of the core count matrix \mathbf{Y}_i follows an independent multinomial distribution parameterized by the corresponding column of an unknown frequency matrix Θ_i , that is, $p(\mathbf{Y}_i | \Theta_i) = \prod_{j=1}^w p(\mathbf{Y}_{ij} | \theta_{ij})$, where

$$\mathbf{Y}_{ij} = (Y_{ija}, Y_{ijc}, Y_{ijg}, Y_{ijt})$$

$$\sim \text{Multinomial}(n_i, \theta_{ij} = (\theta_{ija}, \theta_{ijc}, \theta_{ijg}, \theta_{ijt})).$$

For our between-motif variability, we simply assume that each motif frequency matrix Θ_i in our collection shares a common but completely unknown distribution, denoted $F(\cdot)$, that is,

$$\Theta_i = (\theta_{i1}, \dots, \theta_{iw}) \sim F(\cdot),$$

where $F(\cdot)$ is an unknown distribution with w dimensions for the columns and four dimensions for the nucleotides (constrained to sum to 1). This unknown distribution $F(\cdot)$ represents the common structure between the different motifs in the dataset. Estimation of this unknown distribution is complicated by the fact that our frequency matrices Θ_i are unknown; only the count matrices \mathbf{Y}_i are observed. A popular Bayesian approach to nonparametric problems is to give the unknown distribution $F(\cdot)$ a Dirichlet process prior $\mathcal{D}(\gamma)$, where γ is a

finite (nonnegative) measure, that typically is smooth (Ferguson 1974). Here, because we have a multidimensional $F(\cdot)$, we use a Dirichlet process prior $\mathcal{D}(\gamma_1 \times \cdots \times \gamma_w)$, where each smooth measure γ_j is four dimensional, taking the form of $\gamma_j = b \times \text{Dirichlet}(\alpha, \dots, \alpha)$ for $j = 1, \dots, w$. The parameter b is a weighting factor that characterizes how close the unknown distribution F is to the shape of γ and the smoothness of F .

2.2 Clustering of Observations

An important consequence of our model is that it enables similar motifs to be clustered together into a group modeled by one common frequency matrix. As explained by Ferguson (1974), if $\Theta_1, \dots, \Theta_n$ are n iid observations from the probability function F whose prior distribution is the Dirichlet process $\mathcal{D}(\gamma)$, where γ is a finite measure on the domain, then

$$F(\cdot) | \Theta_1, \dots, \Theta_n \sim \mathcal{D}(\gamma^*) = \mathcal{D}\left(\gamma + \sum_{j=1}^n \delta_{\Theta_j}\right).$$

Thus, the posterior mean of $F(\cdot)$, or the predictive distribution of a new observation, is proportional to $\gamma + \sum_{j=1}^n \delta_{\Theta_j}$. If the Θ 's only take on C distinctive values, then we have a mixture of the smooth measure γ and the C point masses (with potentially different weights). These point mass components allow for the clustering of similar observations. If we were to draw an additional $(n+1)$ th observation Θ^* from this distribution $\mathcal{D}(\gamma^*)$, that new observation would either come from the smooth measure γ or would take on a value exactly equal to one of the current Θ_j 's, say Θ_c , in which case Θ_c and Θ^* are defined as being in the same cluster. The conditional distribution $p(\Theta_i | \Theta_{-i})$ of one current observation Θ_i , given all other observations Θ_{-i} , is also a mixture between the smooth measure and the C point masses at each of the Θ_{-i} that represent the unique values within Θ_{-i} . Any observations Θ_m and Θ_n that have the same value are defined as being in the same cluster. This conditional distribution allows us to implement our model via Gibbs sampling (Geman and Geman 1984). The Dirichlet process has been used as a prior distribution in nonparametric Bayesian analyses, such as MacEachern (1994) and Escobar (1994) for the estimation of normal means and Liu (1996) in a binomial hierarchical setting. Green and Richardson (2001) discussed the use of the Dirichlet process as a flexible model for clustering observations, and presented an extended class of Dirichlet multinomial allocations for which the Dirichlet process is a limiting case. Medvedovic and Sivaganesan (2002) used the clustering properties of the Dirichlet process prior as part of a hierarchical model for gene expression profiles from microarray data.

2.3 Gibbs Sampling Implementation

For our motif clustering model, a Gibbs sampler could intuitively be based on $p(\Theta_i | \Theta_{-i})$. However, because our Θ_i 's are actually unknown, a more efficient clustering procedure involves drawing values of the clustering indicators directly, as in MacEachern (1994), without dealing with drawing a frequency matrix Θ_i for each motif i at each iteration. We denote our clustering indicators $\mathbf{z} = (z_1, \dots, z_n)$, which simply defines a partition of $\{1, \dots, n\}$. Algorithmically, we let $z_i = c$ if Θ_i belongs to the c th cluster or let $z_i = 0$ if Θ_i is drawn

from the prior measure γ and hence forms a new cluster. Our ‘‘collapsed’’ Gibbs sampler (Liu 1994) iteratively samples from $p(z_i | \mathbf{z}_{-i}, \mathbf{Y})$, where we again use the notation \mathbf{z}_{-i} or Θ_{-i} to mean all the \mathbf{z} or Θ parameters except the i th one. As mentioned earlier, no matter what the current Θ_{-i} is, as long as it corresponds to the same indicator vector \mathbf{z}_{-i} , we have the same (almost surely) conditional prior distribution $p(z_i | \Theta_{-i})$. Thus, we can write that

$$\begin{aligned} p(z_i = 0 | \mathbf{z}_{-i}) &= \frac{b}{b + n - 1}, \\ p(z_i = c | \mathbf{z}_{-i}) &= \frac{n_c}{b + n - 1}, \end{aligned} \quad (1)$$

where n_c is the size of cluster c (i.e., the number of z 's in \mathbf{z}_{-i} that are equal to c) and b is the weight parameter for forming a new cluster. Thus, under this model the prior probability for joining a particular cluster increases as the number of observations in that cluster increases, implying that the Dirichlet process prior favors unequal allocations of observations. With observations \mathbf{Y} , we have the posterior conditional distribution

$$p(z_i | \mathbf{z}_{-i}, \mathbf{Y}) \propto p(\mathbf{z}, \mathbf{Y}) \propto p(\mathbf{Y} | \mathbf{z}) p(z_i | \mathbf{z}_{-i}).$$

It is evident from our model that $p(\mathbf{Y} | \mathbf{z})$ is straightforward to derive and can be written as the product of the normalizing constants of C product multinomial distributions, where C is the total number of distinctive z 's. After some simplifications, we have the probability that observation Y_i forms a new cluster is

$$p(z_i = 0 | \mathbf{z}_{-i}, \mathbf{Y}) \propto \frac{b}{b + n - 1} \prod_{j=1}^w \frac{\prod_k \Gamma(Y_{ijk} + \alpha)}{\Gamma(\sum_k Y_{ijk} + 4\alpha)} \frac{\Gamma(4\alpha)}{\Gamma(\alpha)^4}, \quad (2)$$

where b is the weighting factor as defined at the end of Section 2.1. For the case where $z_i = c \neq 0$ (i.e., joining an existing cluster that already has a count matrix $\tilde{\mathbf{Y}}_c$), we have

$$\begin{aligned} p(z_i = c | \mathbf{z}_{-i}, \mathbf{Y}) &\propto \frac{n_c}{b + n - 1} \prod_{j=1}^w \frac{\prod_k \Gamma(Y_{ijk} + \tilde{Y}_{cjk} + \alpha)}{\Gamma(\sum_k Y_{ijk} + \tilde{Y}_{cjk} + 4\alpha)} \\ &\quad \times \frac{\Gamma(\sum_k \tilde{Y}_{cjk} + 4\alpha)}{\prod_k \Gamma(\tilde{Y}_{cjk} + \alpha)}. \end{aligned} \quad (3)$$

A complete iteration of our Gibbs sampling algorithm results in a complete sample \mathbf{z} of our clustering indicators, which also represents a complete *partition* of our motif matrices.

2.4 Motif Alignment

An additional component of our model addresses the fact that we do not necessarily know which *core* \mathbf{Y}_i of width w to use within the raw alignment matrix of width $n_i > w$ for motif i . We use the notation $a_i = j$ to mean that we are using the columns $(j, j+1, \dots, j+w-1)$ of our raw motif matrix \mathbf{N}_i as our core \mathbf{Y}_i . For example, if our clustering algorithm is based on a fixed width $w = 6$ and our i th raw motif matrix \mathbf{N}_i has eight positions, then we have three possible choices for our core motif: $a_i = 1$ ($\mathbf{Y}_i =$ columns 1–6 of \mathbf{N}_i), $a_i = 2$ ($\mathbf{Y}_i =$ columns 2–7 of \mathbf{N}_i), or $a_i = 3$ ($\mathbf{Y}_i =$ columns 3–8 of \mathbf{N}_i). We thus need an additional step where, for each raw data matrix, the best location of the central motif a_i is drawn conditional the other motifs \mathbf{Y}_{-i} and clustering indicators \mathbf{z}_{-i} for the other motifs. Let $\mathbf{Y}_i^{a_i}$

denote the core \mathbf{Y}_i that corresponds to the choice of a particular a_i . Then the posterior probability $p(a_i) = p(a_i|\mathbf{z}_{-i}, \mathbf{Y}_{-i})$ of a_i is

$$\begin{aligned}
 p(a_i) &\propto \int p(\mathbf{Y}_i^{a_i}|\Theta_i)p(\Theta_i|\mathbf{z}_{-i}, \mathbf{Y}_{-i})d\Theta_i \\
 &= \int \Theta_i^{Y_i^{a_i}} \frac{b}{b+n-1} \left[\frac{\Gamma(4\alpha)}{\Gamma\alpha^4} \right]^w \Theta_i^\alpha d\Theta_i \\
 &\quad + \sum_{c=1}^C \int \Theta_i^{Y_i^{a_i}} \frac{n_c}{b+n-1} \delta_{(\Theta_i=\tilde{\Theta}_c)} \frac{p(\tilde{\mathbf{Y}}_c, \tilde{\Theta}_c|\mathbf{z}_{-i})}{p(\tilde{\mathbf{Y}}_c|\mathbf{z}_{-i})} d\Theta_i \\
 &= \frac{b}{b+n-1} \left[\frac{\Gamma(4\alpha)}{\Gamma\alpha^4} \right]^w \frac{\prod_k \Gamma(Y_{ijk}^{a_i} + \alpha)}{\Gamma(\sum_k Y_{ijk}^{a_i} + 4\alpha)} \\
 &\quad + \sum_{c=1}^C \frac{n_c}{b+n-1} \prod_{j=1}^w \frac{\prod_k \Gamma(Y_{ijk}^{a_i} + \tilde{Y}_{cjk} + \alpha)}{\Gamma(\sum_k Y_{ijk}^{a_i} + \tilde{Y}_{cjk} + 4\alpha)} \\
 &\quad \times \frac{\Gamma(\sum_k \tilde{Y}_{cjk} + 4\alpha)}{\prod_k \Gamma(\tilde{Y}_{cjk} + \alpha)}. \tag{4}
 \end{aligned}$$

This alignment procedure is performed every tenth iteration of the collapsed Gibbs sampler described in the previous section.

2.5 Allowing Cluster Width to Vary

We now extend our model to allow the core motif widths within each cluster, $\mathbf{w} = (w_1, \dots, w_C)$, where C is the current number of clusters, to be unknown variables. Each cluster width w_c is modeled as being independent with prior distribution $w_c \sim \text{Poisson}(\lambda)$, where λ is the expected a priori width of the motif in each cluster, which we assume is fixed and specified. We let B_{ik} be the ‘‘background’’ counts of nucleotide k over all columns of the raw matrix \mathbf{N}_i that are not included in the core matrix \mathbf{Y}_i , which now must be taken into account by our model because each motif width is allowed to vary. These background columns represent the edges of the raw motif matrices that are not well conserved. We assume that the background counts, $\mathbf{B}_i = (B_{ia}, \dots, B_{it})$, are a multinomial realization from an underlying vector of background nucleotide frequencies $\theta_0 = (\theta_{0a}, \dots, \theta_{0t})$. With these added distributions, the posterior probability for the core width of a particular cluster c , conditional on the current members of that cluster \mathbf{z}_c and their core alignments \mathbf{a}_c , is

$$\begin{aligned}
 p(w_c|\mathbf{z}_c, \mathbf{a}_c) &\propto \prod_{j=1}^{w_c} \frac{\prod_k \Gamma(\tilde{Y}_{cjk} + \alpha)}{\Gamma(\sum_k \tilde{Y}_{cjk} + 4\alpha)} \frac{\Gamma(4\alpha)}{\Gamma\alpha^4} \prod_k \theta_{0k}^{B_{ck}} \frac{\lambda^{w_c} e^{-\lambda}}{\Gamma(w_c)}, \tag{5}
 \end{aligned}$$

where \tilde{Y}_{cjk} and B_{ck} are, respectively, the core and background nucleotide counts in cluster c . We implement this added component of our model by an additional step in our Gibbs sampling algorithm that, for each cluster c , samples a new value of w_c from the conditional posterior distribution (5). Some methods, such as Schones et al. (2005), use a fixed number of columns to calculate the similarity between matrices and these core widths are then held fixed during the clustering procedure. In contrast, our Gibbs sampling implementation allows the core width of each motif to change, depending on the current set of clusters. This means that each core width in our model is being estimated

using additional information from other motifs in the dataset, rather than independently estimating each core width using only information from each matrix separately. Other methods, such as Kielbasa et al. (2005), cluster motifs on the basis of a pairwise distance and also estimate core width based on a pairwise comparison. This is a slight improvement compared to treating each matrix individually, but still does not use as much information as our model, which makes width decisions on the level of an entire cluster, not just for pairs of motifs.

2.6 Alternative Clustering Priors

As mentioned in Section 2.3, the Dirichlet process prior favors unequal allocation of observations, meaning that each new observation has a greater prior probability of being placed in a cluster that already has many observations. An alternative is a uniform clustering prior that favors equal allocations of observations, that is, the prior probability that a new observation is placed in any one of the existing clusters is uniform. If we already have n observations divided into clusters $c = 1, \dots, C$ with n_1, \dots, n_C members, then the two prior distributions are

Dirichlet process prior:

$$P(z_{n+1} = c|\mathbf{z}, \text{DP}) = \frac{n_c}{b+n},$$

$$P(z_{n+1} = \text{new}|\mathbf{z}, \text{DP}) = \frac{b}{b+n};$$

Uniform clustering prior:

$$P(z_{n+1} = c|\mathbf{z}, \text{unif}) = \frac{1}{b+C},$$

$$P(z_{n+1} = \text{new}|\mathbf{z}, \text{unif}) = \frac{b}{b+C},$$

where b is the weight given to forming a new cluster. In fact, we can consider both the Dirichlet process and the uniform clustering specifications as particular cases of a more general clustering prior distribution, where

$$\begin{aligned}
 P(z_{n+1} = \text{new}|\mathbf{z}) &\propto b, & P(z_{n+1} = c|\mathbf{z}) &\propto f(n_c), \\
 & & c &= 1, \dots, C. \tag{6}
 \end{aligned}$$

This general clustering model reduces to the Dirichlet process when $f(n_c) = n_c$ and the uniform clustering prior when $f(n_c) = 1$, but more general functions may be desirable in particular situations. The prior density of a partition \mathbf{z} with C clusters under either our Dirichlet process or our uniform clustering model can be calculated recursively. With the Dirichlet process prior, we have prior density

$$p(\mathbf{z}|\text{DP}) = \frac{b^C \cdot \prod_{c=1}^C (n_c - 1)!}{\prod_{i=1}^n (b+i-1)}. \tag{7}$$

With the uniform clustering prior, we have prior density

$$p(\mathbf{z}|\text{unif}) = \frac{b^{C-1} \cdot (b+C)}{\prod_{c=1}^C (b+c)^{n_c}}. \tag{8}$$

It is worth noting that the Dirichlet process prior density (7) is not affected by the ordering of the clusters, whereas the uniform clustering prior density (8) is. In other words, different partitions with the same cluster sizes are exchangeable under

the Dirichlet process model, but we will get different values of the uniform clustering density for different, but exchangeable orderings of unequally sized clusters. As suggested by Green and Richardson (2001), to ensure exchangeability of our uniform clustering model, we need to make our prior density $p(\mathbf{z}|\text{unif})$ a function of a “signature” of the partition that is identical for exchangeable partitions. For example, if we let $p^*(\mathbf{z}|\text{unif}) = k \cdot p(\mathbf{z}'|\text{unif})$, where \mathbf{z}' is \mathbf{z} with the z_i 's arranged in order from the largest cluster to the smallest, then the calculation of (8) for \mathbf{z}' will be the same for all exchangeable values of \mathbf{z} . All of these complications are avoided in the Dirichlet process model, which automatically gives the same prior density value for exchangeable partitions.

We compared the behavior of these two clustering prior specifications for a simple simulation study, where 1,000 complete partitions $\mathbf{z} = (z_1, \dots, z_n)$ with $n = 106$ and $b = 1$ were generated under both prior distributions. This particular sample size n was chosen to be comparable with the real data in our first application in Section 3. Figure 3 displays the distributions of the number of clusters as well as the size of the multiple-member ($n_c > 1$) clusters over all of our simulated partitions. As expected, the number of clusters (with multiple members) is larger under the uniform prior and the sizes of some clusters from the Dirichlet process are larger than any generated from the uniform prior specification.

We now discuss several applications of our Bayesian motif clustering model to collections of transcription factor matrices. In Section 3 we apply our method to JASPAR (Sandelin et al. 2004), a small but heavily curated database of transcription factor matrices. In Section 4 we apply our method to a collection that consists of all matrices from JASPAR and the TRANSFAC database (Wingender et al. 2000), which is a larger but less curated database of transcription factor matrices. Finally, in Section 5 we discuss a cross-species application where discovered motifs are used to infer co-regulated genes in bacteria.

3. APPLICATION TO A SINGLE DATABASE: JASPAR

We use the JASPAR (Sandelin et al. 2004) database as an example to illustrate different strategies for visualizing and analyzing results from our clustering model. This database contains 111 nucleotide-count matrices that differ substantially in appearance, number of counts, and motif width. Between different motifs, the number of binding sites used to construct the motif matrix (the total number of multinomial counts) varies from 6 to 389, with an average of around 35 counts per matrix. The range of matrix widths was from 4 to 30 base pairs (bp), although the matrices were generally short, with an average width of approximately 11 bp. From this database, we also have species information for almost all motifs, as well as a classification into a particular “protein family” based on the common physical structure of each motif’s DNA-binding domains. For example, one family of transcription factors is the helix–loop–helix family, which has two DNA-binding helix domains that bind directly to the DNA strand and are joined by a loop domain. We will use this extra protein family and species information when we examine the results produced by our clustering model.

For our Dirichlet process prior, we chose prior parameters α and prior weight b both to be equal to 1. We chose a small prior weight b for simplicity, because we had no prior knowledge on the number of clusters to expect and our clustering results did not change substantially when we used larger values of b . We also assumed our background frequencies were uniform across nucleotides, that is, $\theta_{0k} = .25$ for $k = 1, \dots, 4$. Although we allowed the motif widths to vary, we restricted the core width of each motif to be at least 6 bp to reflect biological reality, with an a priori expected core width $\lambda = 8$ bp. This restriction reduced our dataset from 111 matrices down to 106 matrices. As

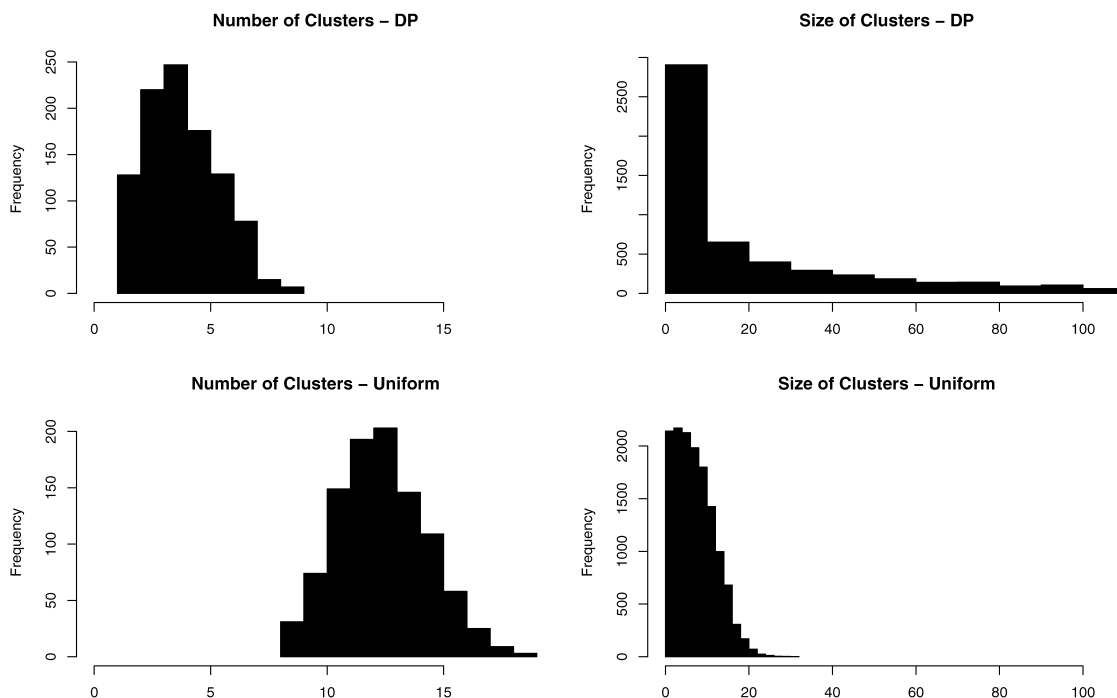


Figure 3. Comparison of clustering statistics between DP and uniform clustering priors.

described in Sections 2.3–2.5, our Bayesian hierarchical clustering model was implemented using a Gibbs sampling algorithm. Details of our evaluation of convergence are given in the supplemental materials.

3.1 Tree of Pairwise Clustering Probabilities

An intuitive means for examining our overall clustering results is the posterior probability p_{ij} that a particular pair of motifs i and j is in the same cluster. The value of p_{ij} for any two motifs i and j can be estimated by the proportion of iterations that have motif i and j in the same cluster, which is the Monte Carlo estimate of the posterior mean of the indicator variable for motif i and j being in the same cluster. The full list of pairwise clustering probabilities is given in the supplemental materials. Based on these pairwise clustering probabilities p_{ij} , a pairwise distance measure can be calculated between each pair of motifs in the dataset, $d_{ij} = 1 - p_{ij}$. Medvedovic and Sivaganesan (2002) used the same distance measure in their gene expression application. Our distance matrix was then converted into a tree diagram by an average-linkage hierarchical algorithm, which is shown in Figure 4.

This clustering tree shows several strong relationships, such as the group of NUCLEAR motifs and the group of bHLH motifs on the left side of Figure 4. Many weaker relationships are also present, implying that many motifs have a low but nonzero probability of being grouped together. Although many of the stronger clusters of motifs belong to the same protein family and same species, there are several interesting exceptions. The NUCLEAR and bHLH groups previously mentioned contain mostly motifs from *Homo sapiens*, but the NUCLEAR group also contains a motif from fruit fly *Drosophila melanogaster* and the bHLH group includes motifs from mouse (*Mus musculus*).

3.2 Best Clustering Partition and Cluster Strength

Although Figure 4 allows us to examine the clustering structure of the entire dataset, the tree is not ideal for deducing the “best partition” or best set of clusters in the dataset, because these pairwise probabilities are calculated across many different partitions, similar to the problem for hierarchical tree clustering mentioned in Section 1.4. One could “cut the tree” at any number of different threshold distances and thereby produce any number of possible partitions, but a less arbitrary alternative is to estimate the posterior mode of our clusters from our Markov Chain Monte Carlo simulation. We estimate this posterior mode by calculating the posterior probability of the partition \mathbf{z} at the end of each iteration of our sampler and retaining the partition $\hat{\mathbf{z}}$ with the highest posterior probability as our best estimate of the mode. The probability $p(\mathbf{z}|\mathbf{Y}, \mathbf{B})$ of \mathbf{z} is calculated as the product of the likelihood value of our cluster matrices $\tilde{\mathbf{Y}}_c$ (the sum of all $w_c \times 4$ count matrices in cluster c) and total background counts B_k ,

$$p(\mathbf{Y}, \mathbf{B}|\mathbf{z}) \propto \prod_{c=1}^C \int p(\tilde{\mathbf{Y}}_c|\Theta_c) p(\Theta_c|\mathbf{z}) p(\mathbf{B}|\theta_0) d\Theta_c$$

$$\propto \prod_{c=1}^C \prod_{j=1}^{w_c} \frac{\prod_k \Gamma(\tilde{Y}_{cjk} + \alpha)}{\Gamma(\sum_k \tilde{Y}_{cjk} + 4\alpha)} \prod_k \theta_{0k}^{B_k},$$

and the prior densities of our clusters $p(\mathbf{z})$ and variable motif widths $p(\mathbf{w})$. Although this best partition can reduce our dataset to a list of interesting clusters, we have lost information about the variability of these clusters by focusing on a point estimate. To retain some measure of variability, we incorporate cluster-level and observation-level clustering characteristics within our “best clusters.” We can measure the strength of each cluster by calculating the logarithm of the *Bayes factor* (Kass and Raftery 1995) for the current cluster c , with members $\mathbf{z} = (z_1, z_2, \dots, z_{n_c})$, versus each member of the cluster forming its own cluster,

$$\text{strength}(\text{cluster } c) = \log \left[\frac{P(\mathbf{Y}|\mathbf{z} \text{ all same})}{P(\mathbf{Y}|\mathbf{z} \text{ all different})} \times \frac{P(\mathbf{z} \text{ all same})}{P(\mathbf{z} \text{ all different})} \right].$$

For a cluster of motifs $(\mathbf{Y}_1, \dots, \mathbf{Y}_m)$ and clustering indicators $\mathbf{z} = (z_1, \dots, z_m)$,

$$\text{strength}(\text{cluster } c) = \log \left[\frac{\int \tilde{\Theta}^{\tilde{\mathbf{Y}}+\alpha-1} d\tilde{\Theta}}{\prod_{i=1}^m \int \Theta_i^{\mathbf{Y}_i+\alpha-1} d\Theta_i} \times \frac{(m-1)!}{b^{m-1}} \right]$$

$$= \log \left[\frac{\prod_{j=1}^{w_c} ((\prod_k \Gamma(\tilde{Y}_{jk} + \alpha)) / (\sum_k \Gamma(\tilde{Y}_{jk} + 4\alpha)))}{\prod_{i=1}^m \prod_{j=1}^{w_c} ((\prod_k \Gamma(Y_{ijk} + \alpha)) / (\sum_k \Gamma(Y_{ijk} + 4\alpha)))} \times \frac{(m-1)!}{b^{m-1}} \right],$$

where $\tilde{\mathbf{Y}}$ and $\tilde{\Theta}$ again denote the count and frequency matrices for the entire cluster together. The clusters within our best partition can then be ranked by this measure of cluster strength, giving us an extra measure of confidence/uncertainty about inference based on a specific cluster. We can also measure clustering strength at the level of individual motifs within our best partition by calculating, for each motif, the posterior probability that it should belong to that cluster, as opposed to any of the other existing clusters or being its own cluster. For each motif i , this posterior probability $p(z_i|\mathbf{z}_{-i}, \mathbf{Y})$ is the same calculation that is performed during each iteration of our Gibbs sampling algorithm, but in this case we are conditioning on the best partition [i.e., $p(\hat{\mathbf{z}}|\hat{\mathbf{z}}_{-i}, \mathbf{Y})$].

For our dataset, the partition $\hat{\mathbf{z}}$ with the highest posterior value consisted of 26 multiple-member clusters containing 69 out of 106 total motifs. The strongest 6 of these 26 clusters (using our cluster strength measure) are listed in Table 2, along with their cluster size, width, and consensus sequence. The consensus sequence is a representation of the total count matrix for the cluster, giving the nucleotide with the highest count in each position. A nucleotide is capitalized only if its nucleotide frequency is greater than .75 in that position. It is clear from the table that this measure of cluster strength is quite dependent on the size of the cluster: larger clusters tend to have a higher value of cluster strength.

Most of our strongest clusters contain motifs from within a single TF protein family, although there are exceptions, such as the fifth cluster that contains both ZN-FINGER and FORK-HEAD motifs. Clearly, this cluster would not have been de-




Figure 4. Tree of pairwise clustering probabilities. The motifs are labeled by the motif name as well as the protein family and species associated with that motif. Motifs that did not cluster with any other motifs are not shown. The length of each tree “branch” (vertical line shared by a group of motifs) is proportional to the average distance d_{ij} between that group of motifs and the next closest group of motifs. The arbitrary ordering of the branches is chosen by the hierarchical algorithm.

Table 2. Strongest clusters from the best partition of JASPAR matrices

Clus	Size	Strength	Width	Consensus	Protein Families	Species	Motifs
1	6	186.5	6	aGGTCA	NUCLEAR	<i>H. sapiens</i> <i>D. melanogaster</i>	MA0065 MA0066 MA0071 MA0072 MA0074 MA0016
2	5	179.5	6	CACGTG	bHLH-ZIP bHLH	<i>H. sapiens</i> <i>M. musculus</i> <i>M. musculus</i>	MA0058 MA0059 MA0093 MA0104 MA0004
3	3	82.2	7	aTGACGT	bZIP	<i>A. majus</i> <i>H. sapiens</i>	MA0096 MA0097 MA0018
4	3	72.7	7	CCGGAAG	ETS	<i>H. sapiens</i> <i>D. melanogaster</i>	MA0028 MA0076 MA0026
5	3	72.5	8	GTAACAa	FORKHEAD ZN-FINGER	<i>H. sapiens</i> <i>D. melanogaster</i>	MA0030 MA0031 MA0013
6	2	40.0	8	cAATtATT	HOMEO-ZIP	<i>A. thaliana</i> <i>A. thaliana</i>	MA0008 MA0110

tected if motifs were grouped together based only on TF family. It is also interesting to note that most of the larger clusters contain similar motifs from different species, with motifs from human (*H. sapiens*) being grouped with motifs from mouse (*M. musculus*), fruit fly (*D. melanogaster*), and snapdragon flowers (*Antirrhinum majus*). All of the motifs in Table 2 have individual clustering probabilities close to 1, but many of the weaker clusters in the best partition contain motifs that have individual probabilities substantially less than 1. Our entire best partition of JASPAR matrices is given in the supplemental materials.

3.3 Core Width Variability in JASPAR

A key component of our clustering procedure is that the width of the core motif within each raw matrix is not considered to be fixed and known, but is instead allowed to vary by cluster. We examine the variability of these core widths in Figure 5, which is a plot of 95% posterior intervals for the core width of each motif in our JASPAR dataset.

The 95% posterior intervals are quite different between motifs: some motifs have a wide interval, while other motifs have an interval that consists only of the minimum motif width of 6 bp, which in many cases is because the raw data matrix is




Figure 5. The 95% posterior intervals of each motif width in JASPAR matrices.

only 6 bp wide. In several motifs, the 95% posterior interval for the motif width does not even include the a priori expected motif width of 8 bp. The wide intervals for several motifs suggest that considering core widths as fixed and known, as in Cartharius et al. (2005) and Schones et al. (2005), would result in a substantial loss of information.

3.4 Effect of Prior Specification on JASPAR Results

In Section 2.6 we discussed the a priori differences between the Dirichlet process prior compared to the uniform clustering prior. We now investigate whether these differences are also apparent in the posterior JASPAR results. The distribution (over all partitions produced by the Gibbs sampler) of the number of multiple-member clusters and the average size of our JASPAR clusters is given in Figure 6.

The Dirichlet process prior and uniform prior give dramatically different clustering results based on prior simulation alone (Fig. 3), but show very slight differences in the posterior clustering results (Fig. 6). Only minor differences were observed between our Dirichlet process and uniform clustering models in terms of the clustering trees (Sec. 3.1) and best partitions (Sec. 3.2), which indicates that our choice of a Dirichlet process prior distribution was not very influential on our posterior clustering results, at least in comparison to a uniform clustering prior alternative. However, other datasets may show a larger influence of the prior specification on the posterior clustering results, and our software allows the user to specify the use of either prior distribution. Green and Richardson (2001) demonstrated with several datasets that the unequal allocations favored by the Dirichlet process priors can persist in the posterior distribution.

4. APPLICATION TO COMBINED DATABASES: JASPAR AND TRANSFAC

Our motivations for applying a clustering model to established databases, such as JASPAR or TRANSFAC, are to eliminate “redundant” matrices within these databases and to understand the similarities as well as differences among the different transcription factors. If two (or more) motifs have nearly identical core matrices, we can consider them to be redundant because they do not contain unique information in terms of appearance. The presence of redundant matrices complicates the use of these databases for further motif discovery, because searches that involve several nearly identical matrices will lead to an excessive number of predicted sites (Kielbasa et al. 2005), and it also affects how we evaluate computational results. As mentioned in Sections 2.4 and 2.5, an additional complication in these databases is that the motif core has unknown width and location within each matrix, which is also addressed by our clustering model. In Section 3 we already observed substantial redundancy within the JASPAR database alone. Now, we want to combine our JASPAR matrices with the larger TRANSFAC database (Wingender et al. 2000), which contains 714 transcription factor matrices. In this combined collection of 825 matrices, we expect substantial redundancy, which we will address with our clustering procedure. Similar to our JASPAR-only application, we chose prior parameters α and b both equal to 1, expected core width $\lambda = 8$ bp, and background frequencies $\theta_{0k} = .25$ for $k = 1, \dots, 4$. We again restricted the core width of each motif to be at least 6 bp, which reduced our dataset from 825 matrices down to 817 matrices. Evaluation of Gibbs sampling convergence for this combined dataset is given in the supplemental materials.

The best partition produced by our Gibbs sampling implementation consisted of 165 multiple-member clusters containing 746 out of the 817 matrices in our combined collection. The

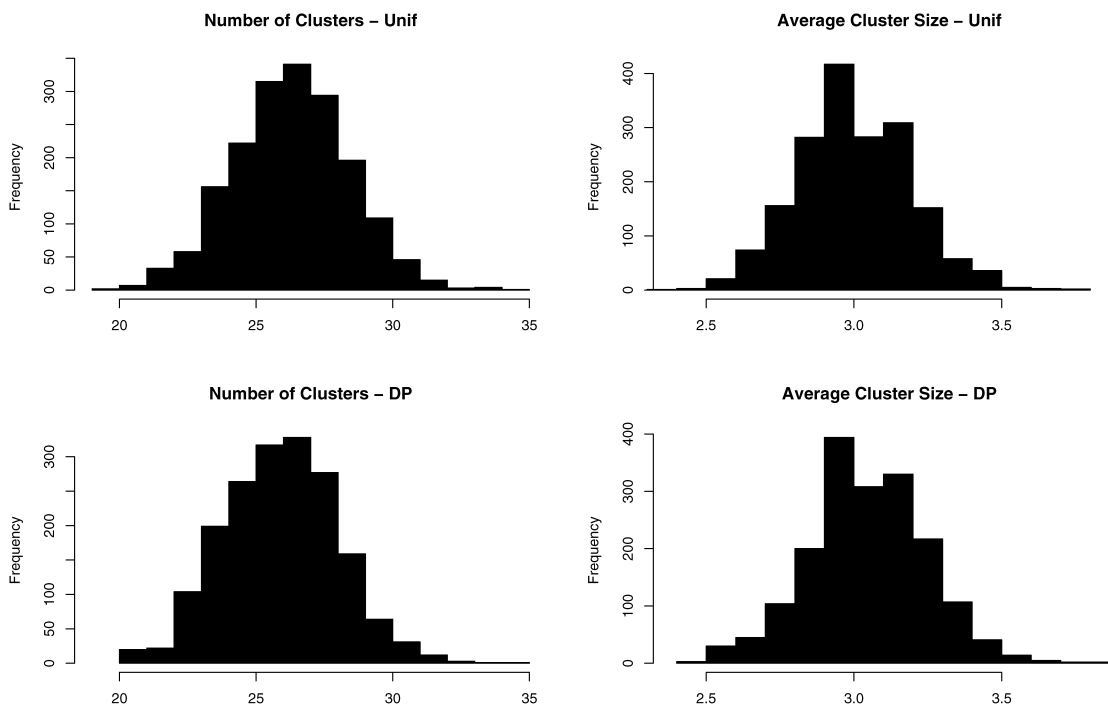


Figure 6. Clustering statistics between uniform and DP models.

Table 3. Strongest clusters from the best partition of TRANSFAC and JASPAR matrices

Cluster	Size	Strength	Width	Consensus	Number of members	
					JASPAR	TRANSFAC
1	24	1,015.3	6	CACGTG	1	23
2	25	974.9	6	TGACGT	2	23
3	15	697.5	8	TTTcGCGC	1	14
4	18	550.1	6	aGATAa	1	17
5	16	548.4	6	CACGTG	4	12
6	17	486.5	6	cGGAAG	3	14
7	7	358.6	6	TGTTCT	0	7
8	11	338.7	7	tCACGTG	0	11
9	14	336.7	6	TGAcCt	1	13
10	11	319.3	6	AAAGcG	4	7

clusters in our best partition are given in the supplemental materials, along with a full list of the pairwise clustering probabilities. We also examined this combined dataset using our model with the alternative uniform clustering prior, and observed very little difference in the posterior clustering results (details given in supplemental materials). Our expectations of high levels of redundancy in these databases seems to be confirmed by the fact that approximately 90% of the matrices in our combined collection are partitioned into multiple-member clusters. The strongest 10 of these 165 clusters (using our cluster strength measure) are listed in Table 3, along with their cluster size, width, and consensus sequence (as defined in our previous application). For brevity, the individual members of each cluster are not given in Table 3. Instead, the total number of members from either JASPAR or TRANSFAC in each cluster is given.

From Table 3, we can see that there is substantial redundancy both between databases and within each database. The redundancy seems to be more substantial within the TRANSFAC database, even taking into account the fact that the JASPAR matrices are a small component (13%) of the combined collection. This is not surprising, because there was a substantial amount of manual curation involved in the creation of the JASPAR database. We also attempted additional merge moves to combine our strong clusters with similar appearance, but these moves did not actually increase our posterior density. Just as in our smaller JASPAR application (Sec. 3), we again observed that a substantial number of motifs in our combined dataset also had substantial variability in their core widths and alignments: 87% of the motifs in the combined dataset had 95% posterior intervals for core width that covered more than a single fixed value (details given in the supplemental materials). It is also worth noting that a majority of the remaining 13% of motifs had no variability simply due to the fact that the raw motif width was equal to the minimum core width. Clearly, these results suggest that the assumption of fixed and known motif widths in Cartharius et al. (2005) and Schones et al. (2005) is a tenuous and results in the loss of substantial information.

Several existing motif-finding programs, such as MatInspector (Cartharius et al. 2005), attempt to utilize nonredundant collections of matrices to scan large sets of genomic sequences for transcription factor binding sites. MatInspector uses a collection of matrix clusters created by manual curation, whereas

we obviously prefer an automated method to generate clusters of matrices. Schones et al. (2005) presented a clustering method for motif matrices that allows matrices to belong to multiple clusters, which means their method does not produce a *partition* of matrices that can be used to reduce redundancy. Our model implementation produces a best partition of clusters, alignments, and core widths, all of which are needed to create a *supermatrix* for each cluster, which is the sum of the aligned core matrices in that cluster. The set of super-matrices can then be used as inputs for a sequence scanning algorithm, such as that by Huang, Kao, Zhou, Liu, and Wong (2004). When looking for matches within genomic sequences, using these super-matrices reduces the number of comparisons that are needed compared to using each matrix in a database individually. This reduction of redundant comparisons not only eases the computational burden of sequence scanning, but also helps somewhat with the usual problem of evaluating the statistical significance of good matches in this multiple comparison setting, as mentioned in Kielbasa et al. (2005).

5. APPLICATION TO CROSS-SPECIES CONSERVATION

In Jensen, Shen, and Liu (2005), we presented an additional application of our motif clustering procedure where *phylogenetically* discovered motifs are clustered to infer co-regulated genes. *Phylogenetic* motif discovery (McCue et al. 2001) searches for conserved motifs in upstream sequences from different, but related, species under the assumption that transcription factor binding motifs are likely to be conserved by evolution. If the motifs found upstream of several *Bacillus subtilis* genes are similar enough to be clustered together, then it is possible that the same TF (recognizing that common motif) is targeting each of the genes in that cluster. Thus, by combining statistical techniques for both motif discovery and motif clustering, one can infer clusters of potentially co-regulated genes. Qin et al. (2003) inferred co-regulated genes by applying a clustering algorithm to previously discovered motifs in *E. coli*. Instead of using a Dirichlet process prior, Qin et al. (2003) used the same uniform clustering prior presented in Section 2.6. However, their clustering technique was simpler than our model in that the core motif width was considered to be fixed and known, which we suggest in Sections 3 and 4 is quite a restrictive assumption.

We used phylogenetic motif discovery on genomic sequences from the bacteria *Bacillus subtilis* and six related bacterial species (Jensen et al. 2005). Our Bayesian hierarchical clustering model was then used to create clusters of highly similar motifs within our collection of discovered motif matrices. Because each motif contains sites discovered in close proximity to a particular *B. subtilis* gene, our motif clusters also can be interpreted as clusters of possibly co-regulated genes. These gene clusters show substantial evidence of co-regulation when compared to several sources of external information about *B. subtilis* genes, such as microarray gene expression data and functional ontology. In addition to using the best partition of clusters, we also were able to incorporate our clustering variability into this application. Our cluster strength measure was again used to rank clusters, and several of the strongest clusters also were shown to be strong in terms of external evidence of co-regulation. Our individual motif strengths $p(\hat{z}_i | \hat{z}_{-i}, \mathbf{Y})$ were also used to filter out weakly associated motifs from our clustering results.

6. DISCUSSION

Although we have presented our clustering procedure in the context of several specific applications to motif matrices, these advantages of our Bayesian clustering model are not specific to this particular type of data. Bayesian hierarchical clustering models based on a Dirichlet process prior distribution should be considered an attractive approach to many clustering problems. Our hierarchical framework lets us account for uncertainty in the count matrices that represents each TF motif by assuming a product multinomial distribution, whereas most clustering programs assume the count matrices are fixed and known without error. The second advantage is that our clustering strategy uses a model-based similarity measure rather than some ad hoc distance measure to compare motifs. At each iteration of the Gibbs sampling algorithm, the decision to cluster a particular observation is determined by the conditional distribution of z_i given all other information (\mathbf{z}_{-i} , \mathbf{Y}_i). Thus, our distance metric is exactly equal to the conditional posterior distribution under our full Bayesian clustering model. A third advantage of our clustering model allows not only the clusters themselves to vary (in terms of which motifs are members of which clusters), but also the number of clusters is allowed to vary. This is a key improvement over a clustering technique that requires the number of clusters to be fixed (such as K -means clustering), because in this situation, we have very little idea a priori how many motifs might be similar to each other. Standard hierarchical tree clustering is also less ideal in this situation, because an arbitrary threshold must be used to produce a set of clusters (e.g., Kielbasa et al. 2005 and Schones et al. 2005). Another general advantage of our procedure is that our posterior sampling implementation gives us an idea of the variability of our clustering results, whereas traditional clustering methods typically give only a point estimate. We explore several approaches to of this variability, including tree structures that summarize the pairwise clustering probability of our motifs, as well as measures of strength for entire clusters and individual motifs within clusters. However, further research into effective techniques for analyzing stochastic clustering results is needed, because the usual procedure of averaging across iterations is not appropriate when both cluster sizes and individual memberships within clusters vary between iterations.

We also presented a novel extension of our model that allows the motif width within each cluster to vary. Our results indicate that many motifs have substantial motif width variability. Previous methods such as Qin et al. (2003) may be ignoring important information by considering motif core widths to be fixed and known a priori. Our model also addresses the alignment issue that, within each raw motif matrix, it is not obvious where the core motif is located. Our model allows us to condition on the motif core in all other raw matrices within the current cluster when we calculate the most likely location of the motif core within a particular matrix. In many cases, other matrices may show very similar compositions to the matrix in question (especially matrices within the same cluster), in which case the conditioning provides a substantial amount of information pertaining to the motif core location. This extra information is ignored by methods (e.g., Schones et al. 2005) that use fixed widths during their clustering procedure and is only partially

captured by methods (e.g., Kielbasa et al. 2005) that use widths estimated by pairwise comparisons of matrices.

As mentioned in Section 4, our clustering results eliminate the redundancy within current matrix databases, which will benefit future motif discovery by reducing the number of redundant hits when scanning sequences for known transcription factors. However, it may be possible to utilize our clustering model for motif discovery in a more sophisticated way. A Bayesian framework for motif discovery presented by Jensen et al. (2004) is based on a motif model where very little is known a priori about the appearance of an unknown motif. However, once a set of motifs has been discovered (and clustered), we should incorporate this information directly into our motif discovery procedure. One proposal would be to use the posterior predictive distribution from our motif clustering model as the scoring function for motif discovery, which would increase the ability of our motif-finding algorithms to detect a motif that is similar to motifs that have already been discovered elsewhere. In addition to using our best partition to aid the discovery of new transcription factor binding sites, this strategy would also allow us to utilize the clustering uncertainty and variability in motif core widths and alignments that are estimated by our model.

[Received December 2004. Revised February 2006.]

REFERENCES

- Cartharius, K., Frech, K., Grote, K., Klocke, B., Haltmeier, M., Klingenhoff, A., Frisch, M., Bayerlein, M., and Werner, T. (2005), "MatInspector and Beyond: Promoter Analysis Based on Transcription Factor Binding Sites," *Bioinformatics*, 21, 2933–2942.
- Crooks, G., Hon, G., Chandonia, J., and Brenner, S. (2004), "WebLogo: A Sequence Logo Generator," *Genome Research*, 14, 1188–1190.
- Dempster, A., Laird, N., and Rubin, D. (1977), "Maximum Likelihood From Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society, Ser. B*, 39, 1–38.
- Eichenberger, P., Jensen, S., Conlon, E., van Ooij, C., Silvaggi, J., Gonzalez-Pastor, J., Fujita, M., Ben-Yehuda, S., Stragier, P., Liu, J., and Losick, R. (2003), "The σ^E Regulon and the Identification of Additional Sporulation Genes in *Bacillus subtilis*," *Journal of Molecular Biology*, 327, 945–972.
- Escobar, M. (1994), "Estimating Normal Means With a Dirichlet Process Prior," *Journal of the American Statistical Association*, 89, 268–277.
- Ferguson, T. (1974), "Prior Distributions on Spaces of Probability Measures," *The Annals of Statistics*, 2, 615–629.
- Geman, S., and Geman, D. (1984), "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6, 721–741.
- Green, P., and Richardson, S. (2001), "Modelling Heterogeneity With and Without the Dirichlet Process," *Scandinavian Journal of Statistics*, 28, 355–375.
- Hannenhalli, S., and Levy, S. (2002), "Predicting Transcription Factor Synergism," *Nucleic Acids Research*, 30, 4278–4284.
- Hartigan, J. (1975), *Clustering Algorithms*, New York: Wiley.
- Huang, H., Kao, M., Zhou, X., Liu, J., and Wong, W. (2004), "Determination of Local Statistical Significance of Patterns in Markov Sequences With Application to Promoter Element Identification," *Journal of Computational Biology*, 11, 1–14.
- Jensen, S., and Liu, J. (2004), "BioOptimizer: A Bayesian Scoring Function Approach to Motif Discovery," *Bioinformatics*, 20, 1557–1564.
- Jensen, S., Liu, X., Zhou, Q., and Liu, J. (2004), "Computational Discovery of Gene Regulatory Binding Motifs: A Bayesian Perspective," *Statistical Science*, 19, 188–204.
- Jensen, S., Shen, L., and Liu, J. (2005), "Combining Phylogenetic Motif Discovery and Motif Clustering to Predict Co-Regulated Genes," *Bioinformatics*, 21, 3832–3839.
- Kass, R., and Raftery, A. (1995), "Bayes Factors," *Journal of the American Statistical Association*, 90, 773–795.
- Kielbasa, S., Gonze, D., and Herzog, H. (2005), "Measuring Similarities Between Transcription Factor Binding Sites," *BMC Bioinformatics*, 6, 237.
- Liu, J. (1994), "The Collapsed Gibbs Sampler in Bayesian Computations With Applications to a Gene Regulation Problem," *Journal of the American Statistical Association*, 94, 958–966.

- (1996), "Nonparametric Hierarchical Bayes via Sequential Imputations," *The Annals of Statistics*, 24, 911–930.
- Lodish, H., Baltimore, D., Berk, A., Zipursky, S., Matsudaira, P., and Darnell, J. (1995), "Regulation of Transcription Initiation," in *Molecular Cell Biology* (4th ed.), New York: Scientific American, pp. 405–481.
- MacEachern, S. (1994), "Estimating Normal Means With a Conjugate Style Dirichlet Process Prior," *Communication in Statistics, Simulation and Computation*, 23, 727–741.
- McCue, L., Thompson, W., Carmack, C., Ryan, M., Liu, J., Derbyshire, V., and Lawrence, C. E. (2001), "Phylogenetic Footprinting of Transcription Factor Binding Sites in Proteobacterial Genomes," *Nucleic Acids Research*, 29, 774–782.
- Medvedovic, M., and Sivaganesan, S. (2002), "Bayesian Infinite Mixture Models Based Clustering of Gene Expression Profiles," *Bioinformatics*, 18, 1194–1206.
- Qin, Z. S., McCue, L. A., Thompson, W., Mayerhofer, L., Lawrence, C. E., and Liu, J. S. (2003), "Identification of Co-Regulated Genes Through Bayesian Clustering of Predicted Regulatory Binding Sites," *Nature Biotechnology*, 21, 435–439.
- Sandelin, A., Alkema, W., Engstrom, P., Wasserman, W., and Lenhard, B. (2004), "JASPAR: An Open Access Database for Eukaryotic Transcription Factor Binding Profiles," *Nucleic Acids Research*, 32, ??–??.
- Schneider, T. D., and Stephens, R. M. (1990), "Sequence Logos: A New Way to Display Consensus Sequences," *Nucleic Acids Research*, 18, 6097–6100.
- Schones, D., Sumazin, P., and Zhang, M. (2005), "Similarity of Position Frequency Matrices for Transcription Factor Binding Sites," *Bioinformatics*, 21, 307–313.
- Wingender, E., Chen, X., Hehl, R., Karas, H., Liebich, I., Matys, V., Meinhardt, T., Pruss, M., Reuter, I., and Schacherer, F. (2000), "TRANSFAC: An Integrated System for Gene Expression Regulation," *Nucleic Acids Research*, 28, 316–319.