

Bayesian clustering in decomposable graphs

Luke Bornn* and François Caron†

Abstract. In this paper we propose a class of prior distributions on decomposable graphs, allowing for improved modeling flexibility. While existing methods solely penalize the number of edges, the proposed work empowers practitioners to control clustering, level of separation, and other features of the graph. Emphasis is placed on a particular prior distribution which derives its motivation from the class of product partition models; the properties of this prior relative to existing priors are examined through theory and simulation. We then demonstrate the use of graphical models in the field of agriculture, showing how the proposed prior distribution alleviates the inflexibility of previous approaches in properly modeling the interactions between the yield of different crop varieties. Lastly, we explore American voting data, comparing the voting patterns amongst the states over the last century.

Keywords: decomposable graphs, Bayesian analysis, product partition models, agriculture, clustering, American voting

1 Introduction

This paper is concerned with the inference of the conditional independence graph \mathcal{G} of a multivariate random vector Y of dimension n , a problem sometimes referred to as structure learning. We focus here on undirected decomposable graphs, whose popularity is mainly due to the tractable factorization they allow for the likelihood (Dawid and Lauritzen 1993; Lauritzen 1996); related work for directed graphical models can be found in Koller and Friedman (2009). Learning the conditional independence graph \mathcal{G} is an onerous task due to the large number of graphs on a set of n nodes, or variables. It is possible using optimization methods to find the graph which best fits the data according to some metric (Meinshausen and Buhlmann 2006; Yuan and Lin 2007; Friedman et al. 2008); alternatively Bayesian model averaging may be used to accommodate uncertainty in the estimated graph, or maximum a posteriori estimation may be used to select a given model from the posterior over graphs. Such an approach relies on a prior distribution $\pi(\mathcal{G})$ over the set of decomposable graphs of a given size; through Bayes theorem, this prior is updated based on the data to give an a posteriori estimate of the distribution over graphs.

Current approaches have been limited in their ability to accommodate varying forms of prior information on the graph. For instance, in an effort to encourage interpretable graphs, the standard approach has been to penalize the number of edges (conditional dependencies) in the graph. However, many situations exist where one might expect

*Department of Statistics, University of British Columbia, Vancouver, BC, Canada
l.bornn@stat.ubc.ca

†INRIA Bordeaux Sud-Ouest, Institut de Mathématiques de Bordeaux, University of Bordeaux, Bordeaux, France, Francois.Caron@inria.fr

variables to be clustered together and the graph to exhibit block structure. At the moment no such prior distribution exists to handle this problem. Our contribution in this article is to propose a class of prior distributions motivated from the class of product partition models which will allow improved flexibility in the specification of prior information on the graph.

The field of agriculture is particularly suitable to the application of graphical models. Due to large spatial domains as well as multifarious crop varieties, it is valuable to have models which both handle the complexity of the biophysical process as well as allow straightforward interpretation. In particular, one might examine the set of zero/non-zero correlations between crop varieties' yields, using the presence or absence of edges to make decisions regarding crop management, marketing, and insurance policies. In addition, due to small sample sizes in many agricultural applications, the choice of prior distribution becomes particularly important.

2 Bayesian Inference on Decomposable Graphs

We begin with a brief overview of graphical models, following the exposition in [Dawid and Lauritzen \(1993\)](#); see also [Lauritzen \(1996\)](#) for further details on graphical models. Let $\mathcal{G} = (V, E)$ be a graphical model with vertices $V = \{1, \dots, n\}$ and pairwise edges E . The pair of nodes $\{i, j\} \in V$ are adjacent if $(i, j) \in E$, and a subset $C \subset V$ is said to be complete if all its elements are adjacent to each other. A complete subgraph that is maximal (i.e. not contained within another complete subgraph) is called a clique. An ordering of the cliques of an undirected graph, (C_1, \dots, C_{n_c}) is said to be perfect if the vertices of each clique C_i also contained in any previous clique C_1, \dots, C_{i-1} are all members of one previous clique; that is, for $i = 2, 3, \dots, n_c$

$$H_i = C_i \cap \bigcup_{j=1}^{i-1} C_j \subseteq C_h$$

for some $h \in \{1, 2, \dots, i-1\}$. The sets H_i , $i = 1, \dots, n_c - 1$ are called separators. We write S_1, \dots, S_{n_s} the non-empty separators (some might appear multiple times). If an undirected graph admits a perfect ordering it is said to be decomposable.

We associate to each vertex i a random variable Y_i . For $A \subseteq V$, let $Y_A = \{Y_i | i \in A\}$. A distribution P over V is Markov with respect to \mathcal{G} if, for any decomposition (A, B) of \mathcal{G} , X_A is independent of X_B given $X_{A \cap B}$. The widespread use of decomposable models is due to the resulting factorization of densities. Specifically, if P satisfies the conditional independencies implied by a decomposable graph \mathcal{G} , then the likelihood of the graphical model specified by P can be factorized according to the graph's cliques and separators

$$p(y|\mathcal{G}, \theta) = \frac{\prod_{i=1}^{n_c} p(y_{C_i} | \theta_{C_i})}{\prod_{j=1}^{n_s} p(y_{S_j} | \theta_{S_j})} \quad (1)$$

where θ is a quantity parameterizing the graphical model P over the graph \mathcal{G} and satisfying some consistency conditions with respect to \mathcal{G} ([Dawid and Lauritzen 1993](#)).

Traditionally, focus has been on Gaussian graphical models, also known as covariance selection models ([Dempster 1972](#)) where $P = N_n(\mu, \Sigma)$ is an n -dimensional multivariate

Gaussian distribution and θ is the $n \times n$ covariance matrix Σ . Conditional independence structure is represented by the precision matrix Σ^{-1} . If the edge $(i, j) \notin E$, then the variables Y_i and Y_j are conditionally independent given the remaining variables, and $\Sigma_{(i,j)}^{-1} = \Sigma_{(j,i)}^{-1} = 0$. As such, the Gaussian graphical model may be factorized as (1) with the covariance Σ replacing θ , and the corresponding likelihood terms written as

$$p(y_B | \Sigma_B) = (2\pi)^{-|B|/2} \det(\Sigma_B)^{-|B|/2} \exp\left[-\frac{1}{2} \text{tr}(S_B(\Sigma_B)^{-1})\right] \quad (2)$$

for each complete set B , where $|B|$ denotes the cardinality of B and S_B is the empirical covariance matrix of y_B .

From a Bayesian perspective, we are interested in the posterior distribution $p(\theta, \mathcal{G} | y) \propto p(y | \theta, \mathcal{G}) p(\theta | \mathcal{G}) \pi(\mathcal{G})$. Much work has been dedicated to specifying proper priors $p(\theta | \mathcal{G})$, (see [Giudici and Green 1999](#); [Dawid and Lauritzen 1993](#)). The main focus of this paper is the specification of a prior distribution $\pi(\mathcal{G})$ over the space of decomposable graphs. As this space is very large compared to the number of observations, it is crucial to add as much prior information as possible on the structure of the unknown graph \mathcal{G} . Moreover, we are generally interested in obtaining sparse graph estimates for needs of interpretation and prediction. Up until now, the specification of $\pi(\mathcal{G})$ has been limited to the uniform distribution, or priors which penalize the complexity as measured by the number of edges. This brings us to the focus of this work, namely a class of prior distributions $\pi(\mathcal{G})$ which subsumes control over the structure and features of \mathcal{G} .

3 Priors on Decomposable Graphs

3.1 Previous work

While early work on inference in decomposable models often assumed a uniform prior over graphs ([Giudici and Green 1999](#)), such priors put considerable mass on models of intermediate size. In an effort to put more weight on smaller graphs, several authors have proposed using a binomial prior distribution with parameter ρ on the number of edges r in the graph. This yields priors of the type ([Dobra et al. 2004](#); [Jones et al. 2005](#))

$$\pi(\mathcal{G}) \propto \rho^r (1 - \rho)^{m-r} \quad (3)$$

where $m = \frac{n(n-1)}{2}$ is the maximal number of possible edges on n nodes. When $\rho = 1/2$, it reduces to the aforementioned uniform prior over graphs. [Jones et al. \(2005\)](#) suggest the use of $\rho = 2/(n-1)$, motivated from the resulting density's peak at n edges in the unconstrained graph. Some authors also consider adding a hierarchical Beta prior $\rho \sim \text{Be}(a, b)$ ([Carvalho and Scott 2009](#)), giving the marginal prior on the graph as

$$\pi(\mathcal{G}) = \int_0^1 \pi(\mathcal{G} | \rho) \pi(\rho) d\rho \propto \frac{\beta(a+k, b+m-k)}{\beta(a, b)}$$

where $\beta(\cdot, \cdot)$ is the beta function. [Carvalho and Scott \(2009\)](#) suggest a default choice of $a = b = 1$, implying a uniform prior on ρ . Interestingly, the resulting prior on \mathcal{G} is

$$\pi(\mathcal{G}) = \frac{1}{m+1} \binom{m}{r}^{-1},$$

which penalizes medium-sized graphs as desired. Such a prior weights each graph according to the number of graphs in the unrestricted space with the same number of edges. However, as shown by [Armstrong et al. \(2009\)](#), the space of decomposable graphs can be considerably different than the unrestricted space. To address this, [Armstrong et al. \(2009\)](#) have proposed a uniform prior on decomposable graphs given the number of edges. However, calculating the number of decomposable graphs of a given size is an arduous task: there exists no list in the literature of decomposable graphs and their breakdown in terms of number of edges, nor are there straightforward ways of computing such quantities. As a result, [Armstrong et al. \(2009\)](#) propose an MCMC estimation scheme, testing its accuracy up to 12 nodes, although such a scheme will likely become prohibitive in higher dimensions.

While the priors in the above references allow one to control the size of the resulting graphs through the number of edges, often doing so results in undesirable graph structures, namely those with a high number of separators and long strings of nodes. [Figure 1](#)(top) shows random samples from a binomial prior over 20-node graphs with $\rho = 0.1$ (closely echoing the choice of [Jones et al. \(2005\)](#), namely $\rho = 2/(n-1) \approx 0.1$) and $\rho = 0.5$ (the uniform prior). We see from this plot that there is no clustering of the cliques, making interpretation difficult. In addition, the long strings/trees seen for $\rho = 0.1$ do not mesh with reality in most cases. Clearly such a class of priors is not suitable if one suspects clustering amongst the variables, clique sizes to be upper (or lower) bounded, or nearly full separation between cliques. Our focus therefore is on moving beyond priors which focus on the number of edges to priors which focus on graph (clique and separator) structure.

3.2 A new prior distribution on decomposable graphs

Motivated from the class of product partition models ([Hartigan 1990](#); [Barry and Hartigan 1992, 1993](#)), we consider prior distributions of the form

$$\pi(\mathcal{G}) \propto \frac{\prod_{j=1}^{n_c} \psi_C(C_j)}{\prod_{j=1}^{n_s} \psi_S(S_j)} \quad (4)$$

where ψ_C and ψ_S are respectively called the clique/separator cohesion functions, with the convention that $\psi_S(\emptyset) = 1$. Evidently one could choose to penalize only cliques or separators by setting ψ_C or ψ_S to constant values. Alternatively, one could simply penalize clique sizes by setting $\psi_C = a^{|C|}$. Motivated from the class of product partition models, consider the cohesion functions $\psi_C(B) = a(|B|-1)!$ and $\psi_S(B) = \frac{1}{b}(|B|-1)!$, $a > 0$, $b > 0$, hence

$$\pi(\mathcal{G}) \propto a^{n_c} b^{n_s} \frac{\prod_{j=1}^{n_c} (|C_j|-1)!}{\prod_{j=1}^{n_s} (|S_j|-1)!}. \quad (5)$$

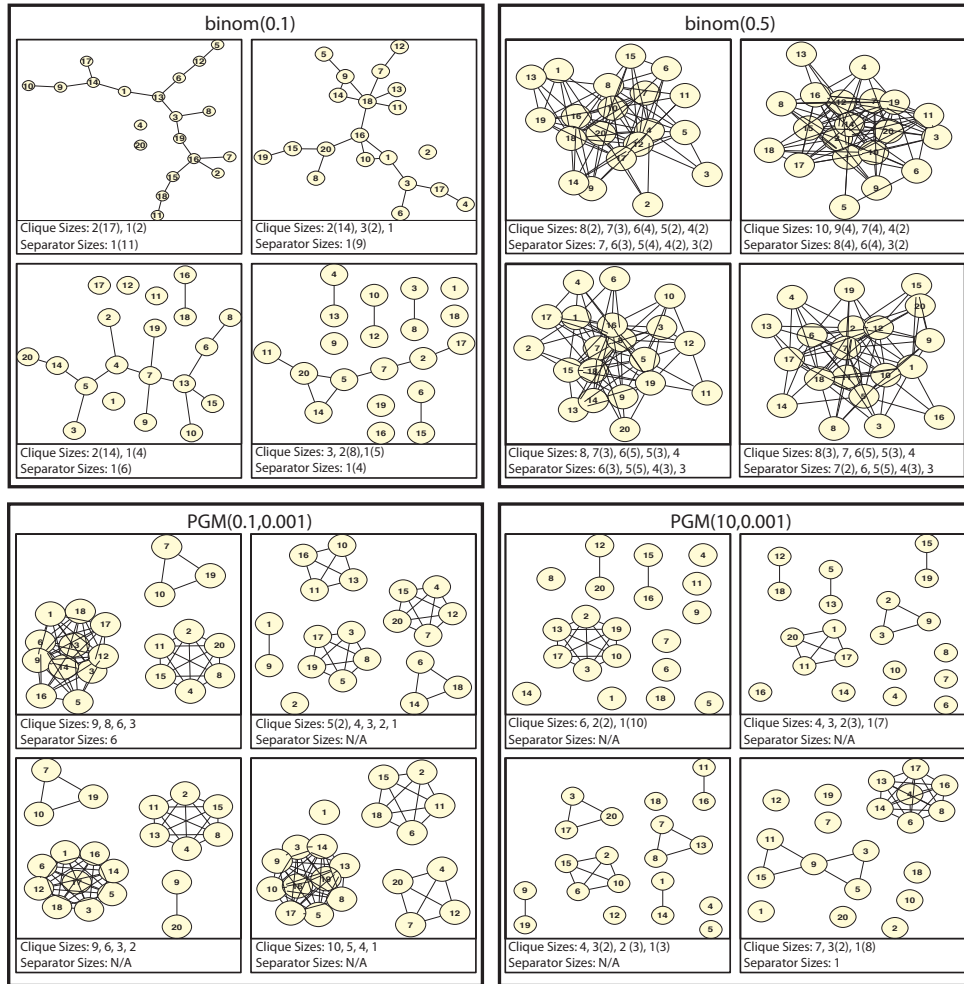


Figure 1: Four random samples from binomial and product graphical model (PGM) priors. Clique and separator sizes for each graph are also shown (“Clique Sizes: 2(3)” implies 3 cliques of size 2). 4 million samples were generated using Markov chain Monte Carlo, and every millionth is shown. While the binomial is characterized by large strings and many separators, the product graphical model allows one to induce clustering by setting b small.

The factorial terms result in predilection towards large cliques and small separators – a desirable trait in terms of interpretability of the resulting graph. For instance, even if $a = b = 1$ with 20 nodes, the completely connected graph would be preferred over the complete independence graph by a factor of $20!$. The parameters a and b respectively tune the number of cliques and separators in the decomposable graph. For a small,

the prior will favour a small number of large cliques. Likewise for b , with small values favouring fewer separators. Figure 1 (bottom) shows samples from this prior. Because of its relation to product partition models (described later), we term this prior the product graphical model prior. To clearly demonstrate the control the product graphical model prior (5) gives relative to the binomial prior, we set $b = 1/1000$, highly penalizing the number of separators and hence resulting in highly separated cliques. In addition, we look at two different values for a ; $a = 0.1$, resulting in fewer and larger cliques, and $a = 10$, resulting in more (but smaller) cliques. Fig. 1 demonstrates the ability of the prior to induce clustering of the cliques, and therefore sparsity in correlation.

We have seen some general properties of the prior (5), namely the ability to control the number of cliques and separators. Figure 2 shows $\log_{10}(\pi(\mathcal{G})/\pi(\mathcal{G}'))$ for different

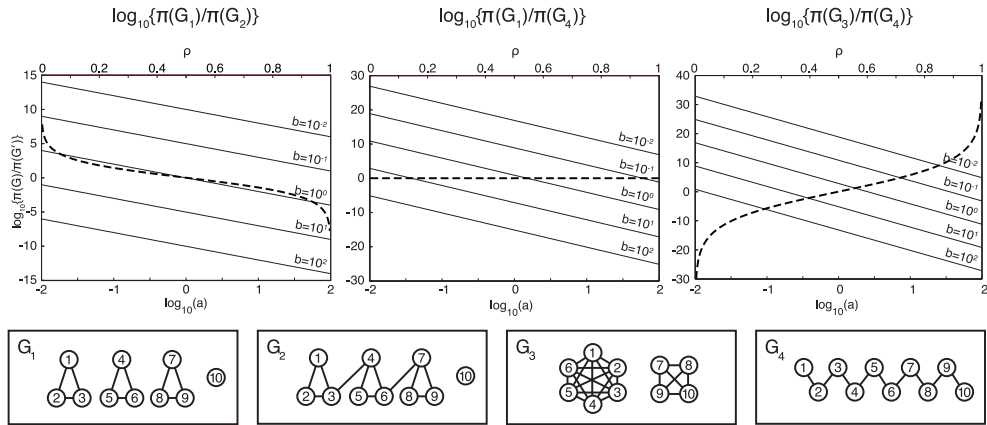


Figure 2: Log ratio of priors over two graphs for product graphical model prior for various a , b (solid, bottom axis) and binomial prior for various ρ (dashed, top axis). While the binomial prior allows one to control the number of edges, for instance choosing \mathcal{G}_1 over \mathcal{G}_2 , the same parameter would seldom choose \mathcal{G}_3 over \mathcal{G}_4 , despite \mathcal{G}_3 having a sparse covariance matrix, and \mathcal{G}_4 having a saturated covariance matrix.

graphs $\mathcal{G}, \mathcal{G}'$. Specifically, decreases in b result in increased prior probability on models with few separators; in addition we see that as a is increased, more mass is put on models with many cliques. In contrast, we also plot the same ratio for the binomial prior (3). From this one can see the limited control such a prior gives, favouring small models in terms of number of edges, but putting very little mass on models, for example, which feature clusters of fully-connected nodes (as in \mathcal{G}_3) and therefore have sparse covariance matrices.

Selecting the appropriate cohesion functions in equation (4) is a difficult problem, but one for which we may gain insight from the existing literature on product partition models (Crowley 1997; Quintana and Iglesias 2003; Quintana 2006). For instance, one may use Figure 2 to select a and b to best fit with prior intuition regarding the features of the graph, then verify the choice through generation of Monte Carlo samples from

the prior as in Figure 1. Alternatively, cross-validation or related methods may be used to select a and b ; due to the potential computational cost of such methods, sequential Monte Carlo approaches may be used to speed up prior distribution selection (Bornn et al. 2010).

Given that the likelihood decomposes as (1) and the prior is of the form (4), the posterior will also be of the form (4) with cohesions $\psi_C(C_j)p(y_{C_j})$ and $\psi_S(S_j)p(y_{S_j})$. The prior admits several other attractive properties and connections with well-known clustering methods as well. If $\psi_S(S_j) \rightarrow \infty$ for all $S_j \neq \emptyset$, then Equation (4) reduces to

$$\pi(\mathcal{G}) \propto \prod_{j=1}^{n_c} \psi_C(C_j)$$

if $n_s = 0$ and 0 otherwise. The resulting prior puts only positive mass on graphs with no separators. It has been introduced as a prior over partitions by Hartigan (1990) and Barry and Hartigan (1992, 1993) under the name of *product partition models*. In the particular case of (4) with $b \rightarrow 0$, the prior over \mathcal{G} reduces to

$$\pi(\mathcal{G}) = \frac{a^{n_c} \Gamma(a)}{\Gamma(a+n)} \prod_{j=1}^{n_c} (|C_j| - 1)!$$

As shown by Quintana and Iglesias (2003) (see also Quintana 2006), this is the distribution over partitions induced by a Dirichlet Process (Ferguson 1973; Antoniak 1974). We also have

$$\mathbb{E}(n_c) = \sum_{i=0}^{n-1} \frac{a}{a+i} \simeq a \log(1+n/a) + \gamma, \quad \text{var}(n_c) = \sum_{i=1}^{n-1} \frac{ai}{(a+i)^2}$$

where γ is Euler's constant and

$$\text{pr}(n_c = k) = s(n, k) a^k \Gamma(a) / \Gamma(a+n)$$

where the coefficients $s(n, k)$ are the absolute values of Stirling numbers of the first kind (Antoniak 1974). In this limiting case, the number of cliques increases logarithmically with the number of nodes.

3.3 Extensions

Motivated by the larger class of exchangeable partition functions (Pitman 1995; Lau and Green 2007), we can also consider four-parameters models, allowing more control over the relative sizes of the cliques/separators

$$\pi(\mathcal{G}) \propto \frac{\prod_{j=1}^{n_c} (a_2 + a_1(j-1)) \frac{\Gamma(|C_j| - a_1)}{\Gamma(1 - a_1)}}{\prod_{j=1}^{n_s} (b_2 + b_1(j-1)) \frac{\Gamma(|S_j| - b_1)}{\Gamma(1 - b_1)}}$$

where $a_2 > -a_1, 0 \leq a_1 < 1$, likewise for b_1, b_2 . The above model reduces to (5) when $a_1 = b_1 = 0$. We can also consider models that control the maximal number of cliques/separators

$$\pi(\mathcal{G}) \propto \frac{\prod_{j=1}^{n_c} (c_1 - j + 1) \frac{\Gamma(c_2 + |C_j|)}{\Gamma(c_2)}}{\prod_{j=1}^{n_s} (d_1 - j + 1) \frac{\Gamma(d_2 + |C_j|)}{\Gamma(d_2)}}$$

where $c_1, c_2, d_1, d_2 > 0$, and $c_1 > d_1$ are the maximal number of cliques/separators. These two models respectively admit as limiting cases the distribution over partitions induced by the two-parameter Poisson-Dirichlet distribution and the finite Dirichlet-multinomial distribution, (see [Lau and Green \(2007\)](#) for further details on these distributions). Using such extensions, one is able to both extend the product graphical model prior to control relative sizes and the maximal number of cliques and separators, as well as borrow from the wealth of literature on Dirichlet and related distributions to gain insight into the prior distribution's characteristics.

4 Example: Modeling Agricultural Output of Different Species

Determining agricultural policies to govern crop production, harvesting, and export is a challenge fraught with high variability both temporally and spatially. Enabling effective crop management, handling, and marketing thus requires accurate understanding of crop yield, which accounts for and explains these variations. While much effort has been made in developing models for predicting single crops ([Potgieter et al. 2006](#); [Bornn and Zidek 2010](#)), little effort has been made in understanding statistically the relationship between crop yield of different crop varieties.

Understanding the connection between yields of different crop varieties is valuable for a multitude of reasons. Firstly, because certain crops are planted and harvested at different times, the management of one crop might benefit from knowledge obtained from harvesting a similar crop earlier in the year. Additionally, by accounting for correlation between different crops, insurers might better cover themselves against extreme events and better control insurance rates for farmers. Lastly, farmers themselves might wish to ensure some level of stability in their income, and therefore might prefer to plant crops which are uncorrelated in yield. Through such a practice, a farmer would be proactive in preventing disasters across his entire crop portfolio. Simply by looking at the resulting undirected graph, a farmer could select two crops which do not have a path connecting them, and are therefore uncorrelated.

We examine the total production (in thousands of bushels) of 24 crops in the state of California from the years 1990 to 2009 (20 years). The data is compiled from the U.S Department of Agriculture website, where a considerable database is available for viewing and analysis. The 24 crops include, for example, several varieties of wheat, rice, and beans. We use the now-standard Gaussian hyper-inverse Wishart model: the likelihood of yield is given in (1) and (2), and the prior for the covariance matrix Σ is hyper-inverse Wishart, which factorizes similarly to (1), as a ratio of inverse Wishart

distributions over cliques and separators (Giudici and Tarantola 1996). See Carvalho and Scott (2009) for some alternative marginal likelihoods based on fractional Bayes factors which can help to induce parsimony. The parameters chosen for the hyper-inverse Wishart distribution are as described in Jones et al. (2005); we focus on the specification of $\pi(\mathcal{G})$. Looking at the list of crops, one would expect that there will be clustering of the yields according to crop characteristics. For instance, it would be reasonable to expect the yield of beans to be correlated with each other. We also seek an interpretable graph, namely one with small complexity (in terms of number of edges and/or separators). The first such prior we examine is the binomial prior of Jones et al. (2005) with $\rho = 2/(n - 1)$, chosen due to its prevalence in the literature. While such a prior allows for penalization on the number of edges, no control is available over clustering. In contrast, by using the prior (5), we can set $b = .01$ to put strong penalization on the number of separators (and hence induce separation of the cliques and therefore sparsity in the correlation matrix), and set $a = .01$ to encourage a small number of cliques in the pursuit of simplicity in the resulting graph.

We run MCMC of length 10 million over the space of decomposable graphs (Giudici and Green 1999) for both the binomial and product graphical model priors, thinning to every 100 samples. With both priors, one may save computational resources by making local moves, merging and splitting cliques within the Markov chain. As a result, one need not re-determine the structure of the entire graph at each move. Figure 3 shows the 4 graphs with highest posterior probabilities from each prior. The product graphical model prior results in the top 4 graphs having posterior density values in the range 0.11 to 0.49, whereas for the binomial the range is 0.04 to 0.06, indicating that the binomial prior spreads mass much more evenly across distributions relative to the product graphical model prior with $a = b = 0.01$. Immediately evident from the figure is the different forms resulting from each prior. Specifically, the binomial prior induces long strings of nodes with many separators, whereas the product graphical model posterior reflects our prior beliefs that variables will cluster together, resulting in sparsity in the correlations between variables. A commercial farmer desiring to plant two plots with uncorrelated crops to minimize the risk of loss might reach quite different conclusions from each prior. Specifically, the large strings of nodes from the binomial prior suggest correlation between the majority of crops. The farmer might not plant winter wheat (planted in late fall) and a strain of beans (harvested in early fall) on his two plots, despite their very different growing seasons, due to their connection in two of the highest posterior probability graphs in Figure 3. In contrast, the separation of cliques from the product graphical model prior (5) would allow these crops to be planted together. Such decisions could be made from the highest posterior graph, or by conducting Bayesian model averaging to obtain the expected utility of a given decision.

To gain an understanding of the product graphical model prior’s prediction performance, we split the data into a training set (first 12 years) and testing set (last 8 years). After simulating from the posterior distribution arising from the binomial and product graphical model priors, we use Bayesian model averaging via the marginal likelihood evaluated on the test data to judge the model’s prediction performance. We evaluate the resulting posterior predictive evaluated on the test set in Table 1; indeed, the product

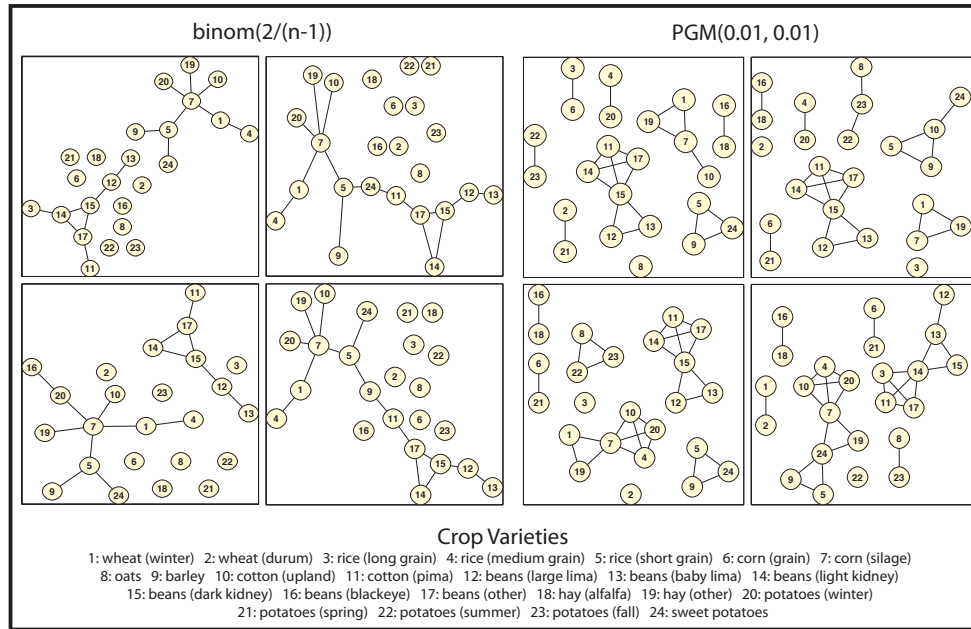


Figure 3: Four samples with highest posterior probability from crop yield model using binomial and product graphical model (PGM) priors. We see the bean yields (nodes 12 through 17) seem to cluster together, as do summer and fall potatoes (nodes 22 and 23). We also observe that the product graphical model prior induces separated cliques, whereas the binomial prior results in long strings and trees of connected variables. As a result, the product graphical model prior will induce sparsity in the resulting posterior covariance.

graphical model prior provides better prediction in this example, even over a variety of parameter choices. We also show the number of edges for each model, indicating that sparsity in terms of edges alone is not responsible for the improved prediction.

5 Example: Modeling 20th Century American Voting Patterns

In an effort to demonstrate the product graphical model prior in higher dimensions, we now turn to the modeling of American voting data by state. For each federal election from 1904 to 1976, occurring every four years, we measure the proportion of votes for the Republican party in each of the 50 states (Carr 2005). Our goal is to model and visualize correlation in voting pattern changes over the last century. Some immediate questions come to mind: “Do certain states have an important role in determining election outcomes?”, “Are there groups of states which vote together, operating independently from

Table 1: Log predictive density evaluated on test data using various priors

Distribution:	Binomial		PGM		
Parameters:	$2/(24 - 1)$	0.5	(0.01, 0.01)	(0.1, 0.1)	(1, 1)
Avg. Log Predictive:	-688	-707	-675	-677	-686
Avg. Number of Edges:	17.8	29.6	18.1	16.4	20.3

the US as a whole?”

We proceed by exploring the posterior distribution resulting from the binomial prior with edge probability 0.1, and the product graphical model prior with parameters $a = 10, b = 10^{-3}$, in an effort to make the overall number of edges resulting from each model comparable. Figure 4 shows the two graphs with highest posterior density from each model. As expected, the binomial graphs contain long strings of variables, while the product graphical model prior demonstrates clustering and grouping of variables. While the binomial prior results in similar variables placed along the same string, the grouping from the product graphical model allows for clearer interpretation. For instance, we immediately observe that the southern states (SC, MS, LA, AL, GA, TX, VA, FL) generally vote in a group. Other patterns of interest also arise, including a close connection between AR, NC, and TN. Also, notice that NY and KS are consistently the single node connecting clusters of variables. As such, these states might be considered key indicators of voting behavior.

6 Discussion

While we have focused on the Bayesian approach to covariance selection, significant work has also been done in a non-Bayesian framework. A common approach involves placing an ℓ_1 penalization on the precision matrix Σ^{-1} , which leads to sparse estimates (Meinshausen and Bühlmann 2006; Yuan and Lin 2007; Friedman et al. 2008). Closer to the heart of this paper, Marlin and Murphy (2009) examine the case of estimating \mathcal{G} when clustering is expected, and therefore Σ^{-1} exhibits block structure. However, these models are neither decomposable nor generative.

While we have focused in this article on Gaussian graphical models, the prior defined in this article is far more general and can be used with any type of model for handling discrete or mixed data (see Madigan and York 1995; Lauritzen 1996). We have also considered the hyperparameters a and b to be known constants. Estimating them within the MCMC sampler would require one to compute the normalizing constant in (5), which is in general not tractable. An exception of interest is the case $b \rightarrow 0$, where we can assign a gamma prior to a and use the data augmentation algorithm described in West (1992) to update a given the other variables.

In conclusion, the proposed product graphical model prior improves flexibility in modeling decomposable graphical models and borrows strength from the immense literature on product partition and related models. The product graphical model prior

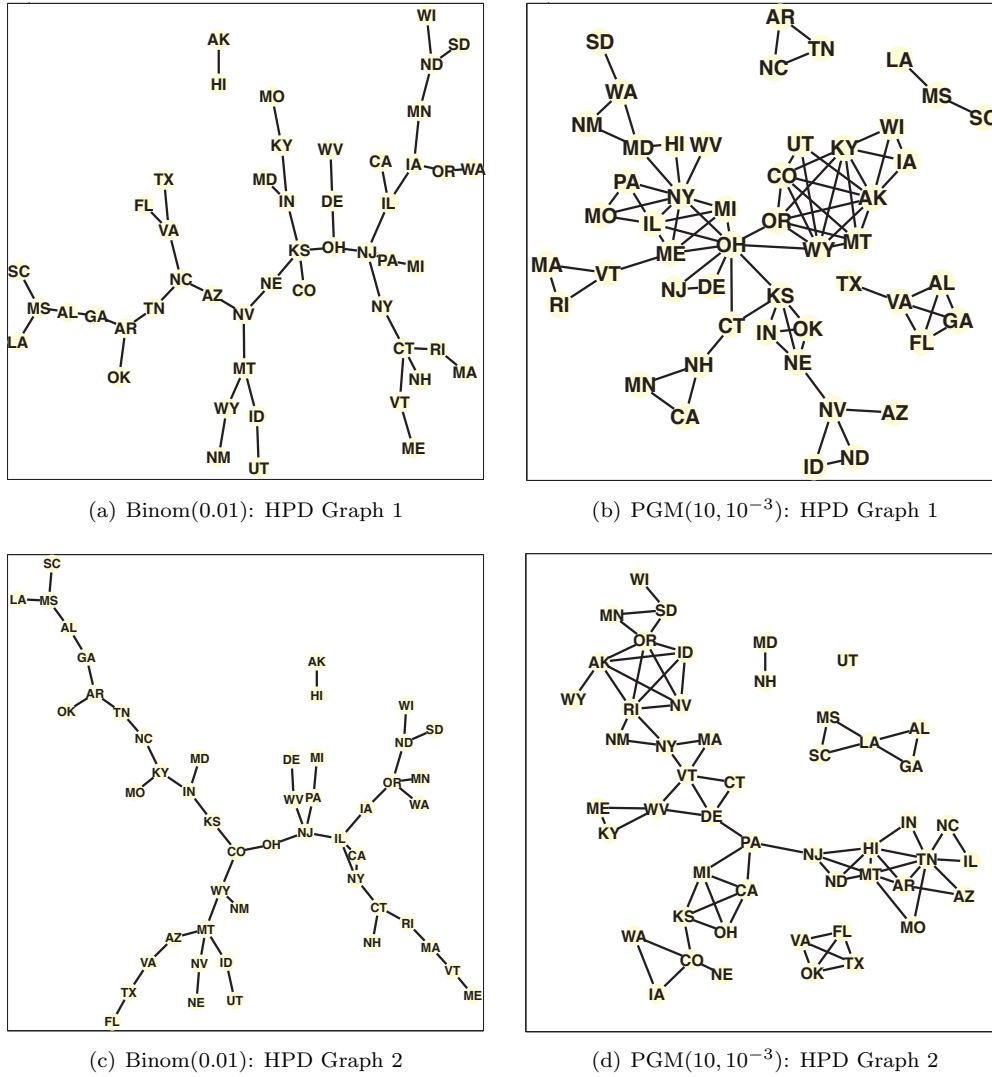


Figure 4: Voting example: two graphs with highest posterior density (HPD) from binomial and product graphical model priors.

allows one to encourage (or discourage) clustering of the graphs, and therefore can induce sparsity in the correlation matrix through clique separation; consequently, the product graphical model empowers practitioners to encapsulate their true prior beliefs to build a model more attuned to the problem at hand.

References

- Antoniak, C. (1974). “Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems.” *Annals of Statistics*, 2: 1152–1174. 7
- Armstrong, H., Carter, C., Wong, K., and Kohn, R. (2009). “Bayesian covariance matrix estimation using a mixture of decomposable graphical models.” *Statistics and Computing*, 19(3): 303–316. 4
- Barry, D. and Hartigan, J. (1992). “Product Partition Models for Change Point Problems.” *Annals of Statistics*, 20: 260–279. 4, 7
- (1993). “A Bayesian analysis for Change Point Problems.” *Journal of the American Statistical Association*, 88: 309–319. 4, 7
- Bornn, L., Doucet, A., and Gottardo, R. (2010). “An Efficient Computational Approach for Prior Sensitivity Analysis and Cross-Validation.” *Canadian Journal of Statistics*, 38(1): 47–64. 7
- Bornn, L. and Zidek, J. (2010). “Efficient Stabilization of Crop Yield Prediction in the Canadian Prairies.” Technical Report 258, Department of Statistics, University of British Columbia. 8
- Carr, A. (2005). “Presidential Elections 1789-2000.”
URL <http://psephos.adam-carr.net/countries/u/usa/pres.shtml> 10
- Carvalho, C. and Scott, J. (2009). “Objective Bayesian model selection in Gaussian graphical models.” *Biometrika*, 96: 1–16. 3, 4, 9
- Crowley, E. (1997). “Product Partition Models for Normal Means.” *Journal of the American Statistical Association*, 92(437): 192–198. 6
- Dawid, A. and Lauritzen, S. (1993). “Hyper-Markov Laws in the Statistical Analysis of Decomposable Graphical Models.” *Annals of Statistics*, 3: 1272–1317. 1, 2, 3
- Dempster, A. (1972). “Covariance selection.” *Biometrics*, 28(1): 157–175. 2
- Dobra, A., Hans, C., Jones, B., Nevins, J., Yao, G., and West, M. (2004). “Sparse graphical models for exploring gene expression data.” *Journal of Multivariate Analysis*, 90(1): 196–212. 3
- Ferguson, T. (1973). “A Bayesian analysis of some nonparametric problems.” *Annals of Statistics*, 1: 209–230. 7
- Friedman, J., Hastie, T., and Tibshirani, R. (2008). “Sparse inverse covariance estimation with the graphical lasso.” *Biostatistics*, 9(3): 432–441. 1, 11
- Giudici, P. and Green, P. J. (1999). “Decomposable graphical Gaussian model determination.” *Biometrika*, 84(4): 785–801. 3, 9

- Giudici, P. and Tarantola, C. (1996). “Global prior distributions for the analysis of discrete graphical models.” *Statistical Methods and Applications*, 5(1): 129–147. 9
- Hartigan, J. (1990). “Partition models.” *Communications in statistics. Theory and methods*, 19: 2745 – 2756. 4, 7
- Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C., and West, M. (2005). “Experiments in stochastic computation for high-dimensional graphical models.” *Statistical Science*, 20(4): 388–400. 3, 4, 9
- Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press. 1
- Lau, J. and Green, P. (2007). “Bayesian model-based clustering procedures.” *Journal of Computational and Graphical Statistics*, 16: 526–558. 7, 8
- Lauritzen, S. (1996). *Graphical Models*. Oxford University Press. 1, 2, 11
- Madigan, D. and York, J. (1995). “Bayesian Graphical Models for discrete data.” *International Statistical Review*, 63: 215–232. 11
- Marlin, B. and Murphy, K. (2009). “Sparse Gaussian graphical models with unknown block structure.” In *Proceedings of the 26th Annual International Conference on Machine Learning*. 11
- Meinshausen, N. and Bühlmann, P. (2006). “High-dimensional graphs and variable selection with the Lasso.” *Annals of Statistics*, 34(3): 1436–1462. 1, 11
- Pitman, J. (1995). “Exchangeable and partially exchangeable random partitions.” *Probability Theory and Related Fields*, 102: 145–158. 7
- Potgieter, A., Hammer, G., and Doherty, A. (2006). “Oz-wheat: a regional-scale crop yield simulation model for Australian wheat.” *Queensland Department of Primary Industries & Fisheries, Information Series No. QI06033, Brisbane, Qld (ISSN 0727-6273)*. 8
- Quintana, F. (2006). “A predictive view of Bayesian clustering.” *Journal of Statistical Planning and Inference*, 136: 2407–2429. 6, 7
- Quintana, F. and Iglesias, P. (2003). “Bayesian clustering and product partition models.” *Journal of the Royal Statistical Society B*, 65: 557–574. 6, 7
- West, M. (1992). “Hyperparameter estimation in Dirichlet process mixture models.” ISDS discussion paper no. 92-03, Duke University, Durham, NC. 11
- Yuan, M. and Lin, Y. (2007). “Model selection and estimation in the Gaussian graphical model.” *Biometrika*, 94: 19–35. 1, 11

Acknowledgments

The authors wish to thank Arnaud Doucet, the associate editor, and several referees for helpful comments which improved the manuscript.

