

A BAYESIAN NONPARAMETRIC MIXTURE MODEL FOR SELECTING GENES AND GENE SUBNETWORKS

BY YIZE ZHAO¹, JIAN KANG^{1,2} AND TIANWEI YU³

Emory University

It is very challenging to select informative features from tens of thousands of measured features in high-throughput data analysis. Recently, several parametric/regression models have been developed utilizing the gene network information to select genes or pathways strongly associated with a clinical/biological outcome. Alternatively, in this paper, we propose a non-parametric Bayesian model for gene selection incorporating network information. In addition to identifying genes that have a strong association with a clinical outcome, our model can select genes with particular expressional behavior, in which case the regression models are not directly applicable. We show that our proposed model is equivalent to an infinity mixture model for which we develop a posterior computation algorithm based on Markov chain Monte Carlo (MCMC) methods. We also propose two fast computing algorithms that approximate the posterior simulation with good accuracy but relatively low computational cost. We illustrate our methods on simulation studies and the analysis of Spellman yeast cell cycle microarray data.

1. Introduction. In high-throughput data analysis, selecting informative features from tens of thousands of measured features is a difficult problem. Incorporating pathway or network information into the analysis has been a promising approach. Generally the setup of the problem contains two pieces of information. The first is the measurements of the features in multiple samples, typically with a clinical outcome associated with each sample. The second piece of information is a network depicting the biological relationship between the features, which is based on existing biological knowledge. The network could contain information such as protein interaction, transcriptional regulation, enzymatic reaction and signal transduction, etc. [Cerami et al. (2011)].

Some methods are developed using the available network topology for high-throughput data analysis. These methods incorporate the gene-pathway relationships or gene network information into a parametric/regression model. The primary goal is to identify either the important pathways or the genes that are strongly

Received May 2013; revised November 2013.

¹Equally contributed.

²Supported in part by the National Center for Advancing Translational Sciences of the National Institutes of Health under Award Number UL1TR000454.

³Supported in part by the National Institutes of Health grants P20 HL113451 and U19 AI090023.

Key words and phrases. Dirichlet process mixture, ising priors, density estimation, feature selection, microarray data.

associated with clinical outcomes of interest. For example, there are a series of works [Wei and Li (2007, 2008); Wei and Pan (2010)] that model the gene network using a Discrete or Gaussian Markov random field (DMRF or GMRF). Li and Li (2008) and Pan, Xie and Shen (2010) used the gene network to build penalties in a regression model for gene pathway selection. Ma et al. (2010) incorporated the gene co-expression network in identification of cancer prognosis markers using a survival model. Li and Zhang (2010) and Stingo et al. (2011) developed Bayesian linear regression models using MRF priors or Ising priors that capture the dependent structure of transcription factors or the gene network/pathway. Recently, Jacob, Neuvial and Dudoit (2012) proposed a powerful graph-structured two-sample test to detect differentially expressed genes.

Although regression models are widely used for the selection of the gene sub-network that is associated with an outcome variable, in some situations the question of interest is to study the expressional behavior of genes, for example, periodicity, without an outcome variable. In other situations, the experimental design is more complex than simple case-control. For example, some gene expression studies involve longitudinal/functional measurements for which the parametric models [Leng and Müller (2006); Zhou et al. (2010); Breeze et al. (2011)] or the multivariate testing procedure [Jacob, Neuvial and Dudoit (2012)] may not be applicable without a major modification. A straightforward approach to this problem is to perform large-scale simultaneous hypothesis testing on gene behavior. A set of genes can be selected based on the testing statistics or p -values, where a correct choice of a null distribution for those correlated testing statistics [Efron (2004, 2010)] should be used. However, this approach ignores the gene network information that is useful to identify the subnetwork of genes with the particular expressional behavior. Due to the diverse behavior of neighboring genes on the network, it is generally believed that genes in close proximity on a network are likely to have joint effects on biological/medical outcomes or have similar expressional behavior. This motivates the needs of analyzing the large-scale testing statistics or statistical estimates incorporating the network information. Another motivation is that a linear regression or parametric model of gene expression levels might not be suitable in some cases. For example, we may be interested in finding subnetworks of genes that have nonlinear relations with an outcome without specifying a parametric form. To address these problems, a simple framework can be adopted. First, a certain statistic is computed for each feature without considering the network structure. The statistic can come from a test of nonlinear association, a test of periodic behavior or a certain regression model. After obtaining the feature-level statistics, a mixture model that takes into account the network structure can be used to select interesting features/subnetworks. More recently, Qu, Nettleton and Dekkers (2012) developed a Bayesian semiparametric model to take into account dependencies across genes by extending a mixture model to a regression model over the generated pseudo-covariates. This method could be sensitive to the choices of the

pseudo-covariates. [Wei and Pan \(2012\)](#) proposed a Bayesian joint model of multiple gene networks using a two-component Gaussian mixture model with a MRF prior. This approach assumes the Gaussian distribution for each component which might not fit the data very well in other applications.

To mitigate problems of the current methods, we propose a Bayesian nonparametric mixture model for large-scale statistics incorporating network information. Specifically, the gene specific statistics are assumed to fall into two classes: “unselected” and “selected,” corresponding to whether the statistics are generated from a null distribution, with prior probabilities p_0 and $p_1 = 1 - p_0$. A statistic has density either $f_0(r)$ or $f_1(r)$ depending on its class, where $f_0(r)$ represents “unselected” density and $f_1(r)$ represents “selected” density. Thus, without knowing the classes, the statistics follow a mixture distribution:

$$(1.1) \quad p_0 f_0(r) + p_1 f_1(r).$$

As suggested by [Efron \(2010\)](#), it is reasonable to assume statistics are normally distributed. This justifies the use of a Dirichlet process mixture (DPM) of normal distributions to estimate both $f_0(x)$ and $f_1(x)$. Note that different from [Wei and Pan \(2012\)](#), our model does not assume that f_0 and f_1 directly take the form of a normal density function. The DPM model has been discussed extensively and widely used in Bayesian statistics [[Antoniak \(1974\)](#); [Escobar \(1994\)](#); [Escobar and West \(1995\)](#); [Müller and Quintana \(2004\)](#); [Dunson \(2010\)](#)], due to the availability of efficient computational techniques [[Neal \(2000\)](#); [Ishwaran and James \(2001\)](#); [Wang and Dunson \(2011\)](#)] and the nonparametric nature with good performance on density estimation. The DPM has been extended to make inference for differential gene expression [[Do, Müller and Tang \(2005\)](#)] and estimate positive false discovery rates [[Tang, Ghosal and Roy \(2007\)](#)] but without incorporating the network information. In our model, we assign an Ising prior [[Li and Zhang \(2010\)](#)] to class labels of all genes according to the dependent structure of the network. As discussed previously, the class label only takes two values, “selected” and “unselected,” while a DPM model is equivalent to an infinity mixture model [[Neal \(2000\)](#); [Ishwaran and James \(2001, 2002\)](#)], based on which we develop a posterior computation algorithm. Our method selects genes and gene subnetworks automatically during the model fitting. To reduce the computational cost, we propose two fast computation algorithms that approximate the posterior distribution either using finite mixture models or guided by a standard DPM model fitting, for which we develop a hierarchical ordered distribution clustering (HODC) algorithm. It essentially performs clustering on ordered density functions. The fast computation algorithms can be tailored from any routine algorithms for the standard DPM model and combined with the HODC algorithm. Also, we suggest two approaches to choosing the hyperparameters in the model.

To the best of our knowledge, our work is among the very first to extend the DPM model to incorporate the gene network for gene selections. Our method has

the following features: (1) It provides a general framework for gene selection based on large-scale statistics using the network information. It can be used for detecting a particular expressional/functional behavior, as well as the association with a clinical/biological outcome. (2) It produces good uncertainty estimates of gene selection from a Bayesian perspective, taking into account the variability from many sources. (3) It introduces more flexibility on model fitting adaptive to data in light of the advantages of Bayesian nonparametric modeling. It is more robust than a parametric model (e.g., two-component Gaussian mixture model) which is sensitive to model assumptions. (4) The fast computational algorithms have been developed for the posterior inference approximation. From our experience, achieving a similar accuracy, it can be 50–150 times (depending on the number of genes in the analysis) faster than the standard Markov chain Monte Carlo (MCMC) algorithm. For a data set more than 2000 genes, the analysis can be done within half an hour using a typical personal computer. (5) Compared with the standard DPM model, our model achieves much better selection accuracy in the simulation studies and provides much more interpretable and biologically meaningful results in the analysis of Spellman yeast cell cycle microarray data. One potential issue of our method is that it only allows the borrowing of information based on the network vicinity, without considering possible compensatory effects between neighboring genes. Such issues can be addressed by downstream analyses after a small number of genes/subnetworks are selected.

The rest of paper is organized as follows. In Section 2 we describe the proposed model and an equivalent model representation. We discuss the choice of priors and the details of the posterior computation algorithms for gene selection. In Section 2.4 we introduce fast computational algorithms for approximating posterior computation. In Section 3 we analyze an example data set, the Spellman yeast cell cycle microarray data. We evaluate the performance of our model via simulation studies in Section 4, where we compare our results with a standard DPM model ignoring the network information. We conclude the paper with discussions and future directions in Section 5.

2. The model. Let n be the total number of genes in our analysis. For $i = 1, \dots, n$, let r_i denote a statistic for gene i . It represents either a functional behavior or the association with a clinical outcome. For the association analysis, it is common to have an outcome Y and a gene expression profile X_i for each gene, i . As an alternative to a regression model, we can produce statistics for each gene, that is, $r_i = s(X_i, Y)$, where $s(\cdot, \cdot)$ can be a covariance function or other dependence test statistics. For a large-scale testing problem, we usually obtain p -values, p_1, \dots, p_n , which can be transformed to normally distributed statistics, that is, $r_i = -\Phi^{-1}(p_i)$, where $\Phi(r)$ denotes the cumulative distribution function for the standard normal distribution. This transformation is a monotone transformation and it ensures the “selected” genes have a larger value of r_i . Let z_i be the

class label for gene selection, where $z_i = 1$ if gene i is selected and $z_i = 0$ is unselected. For $i, j = 1, \dots, n$, let c_{ij} denote the gene network configuration, where $c_{ij} = 1$ if gene i and j are connected, $c_{ij} = 0$ otherwise. Write $\mathbf{r} = (r_1, \dots, r_n)'$, $\mathbf{z} = (z_1, \dots, z_n)'$ and $\mathbf{C} = (c_{ij})$. In our model, \mathbf{r} and \mathbf{C} are observed data, and \mathbf{z} is a latent vector of our primary interest.

2.1. *A network based DPM model for gene selection.* As suggested by Efron (2010), we assume r_i 's are normally distributed. Let $N(\mu, \sigma^2)$ denote a normal distribution with mean μ and standard deviation σ . Let $\mathcal{DP}(G, \alpha)$ represent a Dirichlet process with base measure G and scalar precision α . Given the class label \mathbf{z} , we consider the following DPM model: for $i = 1, 2, \dots, n$ and $k = 0, 1$,

$$\begin{aligned}
 & [r_i \mid \mu_i, \sigma_i^2] \sim N(\mu_i, \sigma_i^2), \\
 (2.1) \quad & [(\mu_i, \sigma_i^2) \mid z_i = k, G_k] \sim G_k, \\
 & G_k \sim \mathcal{DP}[G_{0k}, \tau_k],
 \end{aligned}$$

where μ_i and σ_i^2 are latent mean and variance parameters for each r_i . The random measure G_k and the base measure G_{0k} are both defined on $(-\infty, +\infty) \times (0, +\infty)$. We specify $G_{0k} = N(\gamma_k, \xi_k^2) \times \text{IG}(\alpha_k, \beta_k)$, where $\text{IG}(\alpha, \beta)$ denotes an inverse gamma distribution with shape α and scale β . Note that given latent parameters μ_i, σ_i^2 , the statistic r_i is conditionally independent of z_i . By integrating out (μ_i, σ_i^2) , we build the conditional density of r_i given $z_i = k$ in (1.1), that is,

$$(2.2) \quad f_k(r) = \int \pi(r \mid \boldsymbol{\theta}) dG_k(\boldsymbol{\theta}), \quad \pi(r \mid \boldsymbol{\theta}) = \frac{1}{\sigma} \phi\left(\frac{r - \mu}{\sigma}\right),$$

where $\boldsymbol{\theta} = (\mu, \sigma^2)$ and $\phi(r)$ is the standard Gaussian density function. This provides a Bayesian nonparametric construction of $f_k(r)$.

To incorporate the network structure, we assign a weighted Ising prior to \mathbf{z} :

$$(2.3) \quad \pi(\mathbf{z} \mid \boldsymbol{\pi}, \boldsymbol{\varrho}, \boldsymbol{\omega}, \mathbf{C}) \propto \exp\left[\sum_{i=1}^n \left(\tilde{\omega}_i \log(\pi_{z_i}) + \varrho_{z_i} \sum_{j \neq i} \omega_j c_{ij} I[z_i = z_j]\right)\right],$$

where $\boldsymbol{\pi} = (\pi_0, \pi_1)$ with $0 < \pi_1 = 1 - \pi_0 < 1$, $\boldsymbol{\varrho} = (\varrho_0, \varrho_1)$ with $\varrho_k > 0$ for $k = 0, 1$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)'$ with $\omega_i > 0$ for $i = 1, \dots, n$, and $\tilde{\omega}_i = \sum_{j=1}^n c_{ij} \omega_j / \sum_{j=1}^n c_{ij}$. The indicator function $I[\mathcal{A}] = 1$ if event \mathcal{A} is true, $I[\mathcal{A}] = 0$ otherwise. The parameter $\boldsymbol{\pi}$ controls the sparsity of \mathbf{z} , and the parameter $\boldsymbol{\varrho}$ characterizes the smoothness of \mathbf{z} over the network. For each gene i , a weight ω_i is introduced to control the information inflow to gene i from other connected genes, which can adjust the prior distribution of z_i based on biologically meaningful knowledge, if any. The term $\tilde{\omega}_i$ is introduced to balance the contribution from $\boldsymbol{\pi}$ and $\boldsymbol{\varrho}$ to the prior probability of \mathbf{z} . When $\boldsymbol{\varrho} = (0, 0)$ and $\boldsymbol{\omega} = (1, \dots, 1)'$, the latent class labels z_i 's are independent identically distributed as Bernoulli with parameter π_1 .

2.2. *Model representations.* As discussed by Neal (2000), the DPM models can also be obtained by taking the limit as the number of components goes to infinity. With a similar fashion, we construct an equivalent model representation of (2.1) for efficient posterior computations. Let $\text{Discrete}(\mathbf{a}, \mathbf{b})$ denote a discrete distribution taking values in $\mathbf{a} = (a_1, \dots, a_L)'$ with probability $\mathbf{b} = (b_1, \dots, b_L)'$, that is, if $\xi \sim \text{Discrete}(\mathbf{a}, \mathbf{b})$, then $\Pr(\xi = a_l) = b_l$, for $l = 1, \dots, L$. Let $\text{Dirichlet}(\boldsymbol{\alpha})$ denote a Dirichlet distribution with parameter $\boldsymbol{\alpha}$. Let L_k , for $k = 0, 1$, represent the number of components for density $f_k(r)$. We define the index sets $\mathbf{a}_0 = (-L_0 + 1, -L_0 + 2, \dots, 0)$ and $\mathbf{a}_1 = (1, 2, \dots, L_1)$. Let $\mathbf{q}_0 = (q_{-L_0+1}, q_{-L_0+2}, \dots, q_0)$ and $\mathbf{q}_1 = (q_1, \dots, q_{L_1})$ with $\sum_{g \in \mathbf{a}_k} q_g = 1$. Let $\mathbf{1}_n = (\underbrace{1, \dots, 1}_n)$. Then model (2.1) is equivalent to the following model, as $L_0 \rightarrow \infty$ and $L_1 \xrightarrow{n} \infty$:

$$\begin{aligned}
 (2.4) \quad & [r_i \mid g_i, \tilde{\boldsymbol{\theta}}] \stackrel{\text{i.i.d.}}{\sim} N(\tilde{\mu}_{g_i}, \tilde{\sigma}_{g_i}^2), \\
 & [g_i \mid z_i = k, \mathbf{q}_k] \stackrel{\text{i.i.d.}}{\sim} \text{Discrete}(\mathbf{a}_k, \mathbf{q}_k), \\
 & \tilde{\boldsymbol{\theta}}_g \sim G_{0k} \quad \text{for } g \in \mathbf{a}_k, \\
 & \mathbf{q}_k \sim \text{Dirichlet}(\tau_k \mathbf{1}_{L_k} / L_k),
 \end{aligned}$$

where $\tilde{\boldsymbol{\theta}} = \{\tilde{\boldsymbol{\theta}}_g\}_{g \in \mathbf{a}_0 \cup \mathbf{a}_1}$ and $\tilde{\boldsymbol{\theta}}_g = (\tilde{\mu}_g, \tilde{\sigma}_g^2)$. The index g_i indicates the latent class associated with each data point r_i . Write $\mathbf{g} = (g_1, \dots, g_n)$ and $\mathbf{z} = (z_1, \dots, z_n)$. For each class, g , the parameter $\tilde{\boldsymbol{\theta}}_g$ determines the distribution of r_i from that class. The conditional distributions of g_i and $\tilde{\boldsymbol{\theta}}_{g_i}$ given $z_i = 0$ and $z_i = 1$ are different. Based on model (2.4), the conditional density of $f_k(r)$ in (2.2) becomes

$$(2.5) \quad f_k(r) = \sum_{g \in \mathbf{a}_k} \frac{q_g}{\tilde{\sigma}_g} \phi\left(\frac{r - \tilde{\mu}_g}{\tilde{\sigma}_g}\right).$$

This further implies that given L_0 and L_1 , the marginal distribution of r_i also has a form of finite mixture normals, that is,

$$(2.6) \quad \pi(r) = \sum_{k=0}^1 p_k f_k(r) = \sum_{g=-L_0+1}^{L_1} \frac{\tilde{q}_g}{\tilde{\sigma}_g} \phi\left(\frac{r - \tilde{\mu}_g}{\tilde{\sigma}_g}\right),$$

where $\tilde{q}_g = p_0 q_g$ if $g \leq 0$, $\tilde{q}_g = p_1 q_g$ otherwise.

Model (2.4) is not identifiable for z_i in the sense that if we switch the gene selection class label “0” and “1,” the marginal distribution of r_i (2.6) is unchanged. Without loss of generality, we assume that the “selected gene” should be more likely to have large statistics compared to the “unselected genes.” Thus, we impose an order restriction on the parameter $\tilde{\boldsymbol{\theta}}$, for $g = -L_0 + 1, \dots, L_1$,

$$(2.7) \quad \tilde{\mu}_g < \tilde{\mu}_{g+1}.$$

This also sorts out the nonidentifiability of parameter $\tilde{\boldsymbol{\theta}}$. In many cases, the functional behaviors of some genes are strongly evident from prior biological knowledge. Whether or not those genes are selected is not necessarily determined by

other genes in the network. Those genes are likely to be the hubs of the networks, thus, the determination of the status of these genes might help select genes in their neighborhood. This suggests that it is reasonable to preselect a small amount of genes that can be surely elicited by biologists from their experience and knowledge. We refer to them as “surely selected” (SS) genes. These genes are usually associated with very large statistics. We evaluate the performance via the simulation studies in Section 4.2.

2.3. *Posterior computation.* In model (2.4), given L_0 and L_1 , we have the full conditional distribution of $g_i = g$ and $z_i = k$ given $\mathbf{g}_{-i} = (g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_n)$, $\mathbf{z}_{-i} = (z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n)$ and data \mathbf{r} :

$$\begin{aligned}
 (2.8) \quad & \pi(g_i = g, z_i = k \mid \mathbf{g}_{-i}, \mathbf{z}_{-i}, \mathbf{r}, \tilde{\boldsymbol{\theta}}) \\
 & \propto \frac{1}{\tilde{\sigma}_g} \phi\left(\frac{r_i - \tilde{\mu}_g}{\tilde{\sigma}_g}\right) \frac{n_{-ig} + \tau_k/L_k}{\tau_k + m_k - 1} \\
 & \quad \times \exp\left(\tilde{\omega}_i \log(\pi_k) + \varrho_k \sum_{j \neq i} \omega_j c_{ij} I[z_j = k]\right),
 \end{aligned}$$

where $m_k = \sum_{i=1}^n I[z_i = k]$ is the number of genes in class k and $n_{-ig} = \sum_{j \neq i} I[g_j = g]$ represents the number of g_j for $j \neq i$ that are equal to g .

As $L_0 \rightarrow \infty$ and $L_1 \rightarrow \infty$, if $(g, k) = (g_j, z_j)$ for some $j \neq i$, then

$$\begin{aligned}
 (2.9) \quad & \pi(g_i = g, z_i = k \mid \mathbf{g}_{-i}, \mathbf{z}_{-i}, \mathbf{r}, \tilde{\boldsymbol{\theta}}) \\
 & \propto \frac{n_{-ig}}{\tau_k + m_k - 1} \exp\left(\tilde{\omega}_i \log(\pi_k) + \varrho_k \sum_{j \neq i} \omega_j c_{ij} I[z_j = k]\right) \\
 & \quad \times \frac{1}{\tilde{\sigma}_g} \phi\left(\frac{r_i - \tilde{\mu}_g}{\tilde{\sigma}_g}\right)
 \end{aligned}$$

and

$$\begin{aligned}
 (2.10) \quad & \pi(g_i \neq g_j, z_i \neq z_j \text{ for all } j \neq i \mid \mathbf{g}_{-i}, \mathbf{z}_{-i}, \mathbf{r}, \tilde{\boldsymbol{\theta}}) \\
 & \propto \frac{\tau_k}{\tau_k + m_k - 1} \exp\left(\tilde{\omega}_i \log(\pi_k) + \varrho_k \sum_{j \neq i} \omega_j c_{ij} I[z_j = k]\right) \\
 & \quad \times \frac{\Gamma(\alpha_k + 1/2) \beta_k^{\alpha_k}}{\sqrt{2\pi} \Gamma(\alpha_k) \xi_k} \int \phi\left(\frac{\mu - \gamma_k}{\xi_k}\right) \left(\beta_k + \frac{1}{2}(r_i - \mu)^2\right)^{-(\alpha_k + 1/2)} d\mu,
 \end{aligned}$$

where the integral can be efficiently computed by the Gaussian quadrature method in practice. See Section A in the supplemental article [Zhao, Kang and Yu (2014)] for the derivations of equations (2.9) and (2.10).

The full conditionals of $\tilde{\mu}_g$ and $\tilde{\sigma}_g^2$ for $g \in \{g_1, \dots, g_n\}$ are given by

$$(2.11) \quad [\tilde{\mu}_g \mid \tilde{\sigma}_g^2, \mathbf{r}] \sim \text{N}\left(\frac{\tilde{\sigma}_g^2 \gamma_k + \xi_k^2 \sum_{i:g_i=g} r_i}{\tilde{\sigma}_g^2 + \xi_k^2 n_g}, \frac{\tilde{\sigma}_g^2 \xi_k^2}{\tilde{\sigma}_g^2 + \xi_k^2 n_g}\right),$$

$$(2.12) \quad [\tilde{\sigma}_g^2 \mid \tilde{\mu}_g, \mathbf{r}] \sim \text{IG}\left(\alpha_k + \frac{n_g}{2}, \beta_k + \frac{1}{2} \sum_{i:g_i=g} (r_i - \tilde{\mu}_g)^2\right),$$

where $k = I[g > 0]$ and $n_g = \sum_{i=1}^n I[g_i = g]$. We summarize this algorithm in Section B.1 in the supplemental article [Zhao, Kang and Yu (2014)] and refer to it as NET-DPM-1. It is computationally intensive when n is very large. To mitigate this problem, we propose two fast algorithms to fit finite mixture models (FMM) with appropriate choices of the number of components.

2.4. Fast computation algorithms.

2.4.1. *FMM approximation.* When L_1 and L_0 fit the data well, we can accurately approximate the infinite mixture model (2.1) by the FMM (2.4). Given a fixed L_0 and L_1 , it is straightforward to perform posterior computation for model (2.4) based on (2.8). We refer to this algorithm as NET-DPM-2 (see Section B.2 in the supplemental article [Zhao, Kang and Yu (2014)] for details). This algorithm does not change the dimension of $\tilde{\theta}$ over iterations. In this sense, it simplifies the computation. Also, in order to keep computation efficient, we search for smaller values of L_0 and L_1 which fit the data well. This can be achieved under the guidance of a DPM density fitting for which we introduce an algorithm in the next section.

2.4.2. *Hierarchical ordered density clustering.* Without using the network information, a DPM model fitting on data \mathbf{r} provides an approximation to the marginal density (2.6). It generates posterior samples for mixture densities, where the mean number of components should be close to $L_0 + L_1$. Let us focus on one sample. Suppose $L_0 + L_1$ is equal to the number of components in this sample. To further obtain an estimate of L_0 and L_1 for this sample, we need to partition the $L_0 + L_1$ components into two classes. Thus, we propose an algorithm to cluster a set of ordered densities. We call it hierarchical ordered density clustering (HODC). Here, the density order is determined by the mean location of that density. For example, a set of Gaussian density functions are sorted according to their mean parameters. Similar to the classical hierarchical clustering analysis, we define a distance metric of density functions:

$$(2.13) \quad d(f, f') = \int_{-\infty}^{+\infty} [f(x) - f'(x)]^2 dx,$$

where f and f' are two univariate density functions. Let $\mathcal{P} = \{(\hat{\mu}_g, \hat{\sigma}_g^2, \hat{p}_g)\}_{g=1}^{L_0+L_1}$ denote parameters for $L_0 + L_1$ Gaussian densities, where $\hat{\mu}_g < \hat{\mu}_{g+1}$, $g =$

1, 2, . . . , L₀ + L₁ - 1. This is the input data to the HODC algorithm totally consisting of L₀ + L₁ - 2 steps. At the m step, there are L₀ + L₁ - m clusters of densities and let $\mathbf{s}_l^{(m)}$, for l = 1, . . . , L₀ + L₁ - m, denote the density indices in cluster l. To simplify, we define

$$(2.14) \quad \tilde{\phi}(r; \mathbf{s}, \mathcal{P}) = \sum_{g \in \mathbf{s}} \frac{\hat{p}_g}{\hat{\sigma}_g} \phi\left(\frac{r - \hat{\mu}_g}{\hat{\sigma}_g}\right) / \sum_{g \in \mathbf{s}} \hat{p}_g,$$

which represents a mixture of Gaussian densities, where the components indexed by \mathbf{s} are a subset of $\{\phi[(r - \hat{\mu}_g)/\hat{\sigma}_g]/\hat{\sigma}_g\}_{g=1}^{L_0+L_1}$.

HODC:

Input: Parameters for a mixture of Gaussian densities, that is, \mathcal{P} .

Initialization: Set m = 0 and $\mathbf{s}_l^{(0)} = \{l\}$, for l = 1, 2, . . . , L₀ + L₁;

Repeat the following steps until m = L₀ + L₁ - 2:

Step 1: Find

$$l^{(m)} = \arg \min_l d(\tilde{\phi}(\cdot; \mathbf{s}_l^{(m)}, \mathcal{P}), \tilde{\phi}(\cdot; \mathbf{s}_{l+1}^{(m)}, \mathcal{P})).$$

Step 2: For l = 1, 2, . . . , L₀ + L₁ - m - 1, set

$$\mathbf{s}_l^{(m+1)} = \begin{cases} \mathbf{s}_l^{(m)} & \text{if } l < l^{(m)}, \\ \mathbf{s}_l^{(m)} \cup \mathbf{s}_{l+1}^{(m)} & \text{if } l = l^{(m)}, \\ \mathbf{s}_{l+1}^{(m)} & \text{if } l > l^{(m)}. \end{cases}$$

Step 3: Set m = m + 1.

Output: $\{\mathbf{s}_l^{(m)}\}_{l=1}^{L_0+L_1-m}$ for m = 1, 2, . . . , L₀ + L₁ - 2.

Figure 1 illustrates the HODC algorithm. The algorithm stops when m = L₀ + L₁ - 2, where the ordered density components are partitioned into two classes indexed by $\mathbf{s}_1^{(m)}$ and $\mathbf{s}_2^{(m)}$. This suggests that the number of indices in $\mathbf{s}_{k+1}^{(m)}$, denoted by $|\mathbf{s}_{k+1}^{(m)}|$, is an estimate for L_k in model (2.4). By running the HODC, we can obtain one L_k estimate for each posterior sample generated from a DPM fitting. We take the average of L_k estimates over all the posterior samples as the input of NET-DPM-2. The HODC also provides an approximation to f_k(r) in (2.5), that is, $\tilde{\phi}(r; \mathbf{s}_{k+1}^{(m)}, \mathcal{P})$. This implies that we can further simplify the computation with the algorithm in the following section.

2.4.3. *FMM guided by a DPM model fitting.* From a DPM model fitting, we obtain V posterior samples of the parameters for the marginal density of \mathbf{r} . We denote them as $\mathcal{P}_v = \{(\hat{\mu}_{vg}, \hat{\sigma}_{vg}^2, \hat{p}_{vg})\}_{g=1}^{L_{v0}+L_{v1}}$, for v = 1, 2, . . . , V. For each \mathcal{P}_v , the HODC algorithm partitions L_{v0} + L_{v1} components into two classes, where the

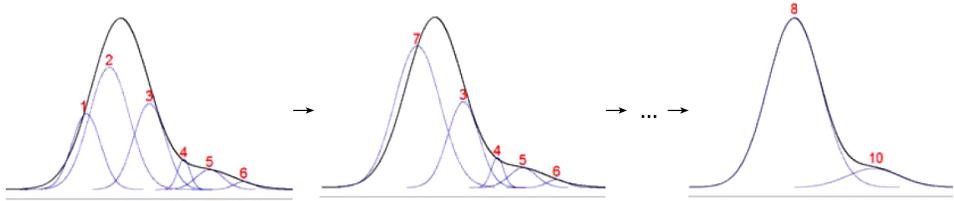


FIG. 1. An illustration of the HODC algorithm for six density components: the HODC starts with clustering densities 1 and 2 as a mixture density labeled as 7, since the “distance” between 1 and 2 is shorter than all other adjacent density pairs. Then the HODC computes the “distance” between densities 3 and 7, densities 3 and 4, . . . , to proceed the clustering. Following this procedure, the HODC ends up with clustering densities 1, 2, 3 as a mixture density (labeled as 8) and 4, 5, 6 as another mixture density (labeled as 10).

class-specific components are indexed by $\mathbf{a}_{v,0}$ and $\mathbf{a}_{v,1}$. This leads to V approximations of $f_k(r)$, that is, $\tilde{\phi}(r; \mathbf{a}_{v,k}, \mathcal{P}_v)$. Given $f_k(r)$, our proposed gene selection model reduces to

$$(2.15) \quad [r_i | z_i = k] \stackrel{\text{i.i.d.}}{\sim} f_k(r)$$

for $i = 1, 2, \dots, n$ and $k = 0, 1$, and \mathbf{z} follows (2.3). To make inference on the posterior distribution of \mathbf{z} by combining all V approximations of $f_k(r)$, we consider

$$(2.16) \quad \pi(\mathbf{z} | \mathbf{r}) \approx \frac{1}{V} \sum_{v=1}^V \pi(\mathbf{z} | \mathbf{r}, \tilde{\phi}_v),$$

where $\tilde{\phi}_v = \{\tilde{\phi}(r; \mathbf{a}_{v,0}, \mathcal{P}_v), \tilde{\phi}(r; \mathbf{a}_{v,1}, \mathcal{P}_v)\}$. For each v , the full conditional of z_i is given by

$$(2.17) \quad \begin{aligned} &\pi(z_i = k | \mathbf{z}_{-i}, \mathbf{r}, \tilde{\phi}_v) \\ &\propto \tilde{\phi}(r_i; \mathbf{a}_{v,k}, \mathcal{P}_v) \exp\left(\tilde{\omega}_i \log(\pi_k) + \varrho_k \sum_{j \neq i} \omega_j c_{ij} I[z_j = k]\right). \end{aligned}$$

We refer to this algorithm as NET-DPM-3 (see Section B.3 in the supplemental article [Zhao, Kang and Yu (2014)] for details). It is extremely fast with a moderate V . Since the marginal density is estimated without using the network information, it might introduce bias on the distribution of z_i and underestimate the variability of z_i . From our experience, those issues do not affect the selection accuracy much. Some examples are provided in Section 4.

2.5. *The choice of hyperparameters.* To proceed NET-DPMs, we need to specify the hyperparameters $\boldsymbol{\pi}$, $\boldsymbol{\varrho}$ and $\boldsymbol{\omega}$ in (2.3). We assume that $\boldsymbol{\omega}$ is prespecified according to biological information. In this paper, we choose equal weight, that is, $\boldsymbol{\omega} = \mathbf{1}_n$ without incorporating any biological prior knowledge. We suggest two approaches to choosing $\boldsymbol{\pi}$ and $\boldsymbol{\varrho}$: (1) we assign hyperpriors on $\boldsymbol{\pi}$ and $\boldsymbol{\varrho}$ and

make posterior inference; (2) for a set of possible choices of π and ϱ , we employ the Bayesian model averaging. The details are provided in Section C in the supplemental article [Zhao, Kang and Yu (2014)].

3. Application. To demonstrate the behavior of our method, we apply the proposed method to the analysis of the Spellman yeast cell cycle microarray data set [Spellman et al. (1998)]. The data set is intended to detect genes with periodic behavior along the procession of the cell cycle. It has been extensively used in the development of computational methods. The network is summarized from the Database of Interacting Proteins (DIP) [Xenarios et al. (2002)]. We use the high-confidence connections between yeast proteins from the DIP. Eventually, the network contains 2031 genes, where the mean, median, maximum and minimum edges per gene are 3.948, 2, 57 and 1, respectively.

There is no outcome variable in the cell-cycle data set. In this demonstration we focus on the selection of genes with periodic behavior in light of the network. It is known that such genes show different phase shifts along the cell cycle and may not be correlated with each other [Yu (2010)]. We first perform the Fisher's exact G test for periodicity [Wichert, Fokianos and Strimmer (2004)] for each gene. We then transform the p -values to normal quantiles, $r_i = -\Phi^{-1}(p_i)$ for gene i . We apply the fully Bayesian inference (NET-DPM-1), one fast computation approach (NET-DPM-3) and the standard DPM model fitting (STD-DPM) to this data set. For the NET-DPM-1, set $\tau_0 = 10$, $\tau_1 = 2$; following the results by STD-DPM, set $\gamma_k = \bar{\mu}_k$, $\xi_k^2 = \bar{\sigma}_k^2$, $\beta_k = 10$, $\alpha_k = \bar{\sigma}_k^2/\xi_k^2 + 1$ with $k = 0, 1$, where $\{\bar{\mu}_k\}$ and $\{\bar{\sigma}_k^2\}$ are preliminary estimations by the STD-DPM. We also conduct a sensitivity analysis for the hyperparameters specification (see the details in Section E in the supplemental article [Zhao, Kang and Yu (2014)]) to verify the robustness of the proposed methods. For both methods, the choices of π_0 and ϱ for the model averaging algorithm are (0.75, 0.8, 0.85, 0.9) and (0.5, 1, 5, 10, 15) \times (0.5, 1, 5, 10, 15) with restriction $\varrho_0 < \varrho_1$. We run all the algorithms 5000 iterations with 2000 burn-in. In this article, the standard DPM fitting is obtained by an R package: `DPpackage` and all the proposed algorithms are implemented in R.

Table 1 presents the gene selection results based on three methods in a two-by-two table format. The number of the "selected" genes by the NET-DPM-1, the NET-DPM-3 and the STD-DPM are 201, 216 and 114, respectively. The summation of the diagonal elements of the table comparing the NET-DPM-3 and the NET-DPM-1 is larger than that for NET-DPM-3 and the STD-DPM. This indicates a stronger agreement between the two algorithms for NET-DPM.

We focus our discussion on the NET-DPM-3 results. After removing all unselected genes, as well as selected genes not connected to any other selected genes, 163 of the 216 genes fall into 11 subnetworks. Of the 11 subnetworks, 10 are very small, each containing 5 or less genes. The remaining subnetwork contains 135 genes. Considering the purpose of the study is to find genes with periodic behavior, and most such genes are functionally related and regulated by the cell cycle

TABLE 1
Gene selection results by the three methods for the cell cycle data set

	NET-DPM-1		STD-DPM	
	Selected	Unselected	Selected	Unselected
NET-DPM-3				
Selected	170	46	100	116
Unselected	31	1784	14	1801

process, this result is expected. We present the subnetwork in Figure 2. Sixty-one of the 135 genes belong to the mitotic cell cycle process based on gene ontology [Ashburner et al. (2000)]. The yeast mitotic cell cycle can be roughly divided into the M phase and the interphase, which contains S and G phases [Ashburner et al. (2000)]. We do not further divide the interphase because the number of genes annotated to its descendant nodes are small. Among the 135 genes, 45 are annotated to the M phase, and 21 are annotated to the interphase. By coloring the M phase genes in red, the interphase genes in blue and the genes annotated to both phases in green, we see that the majority of the selected M phase genes are clustered on

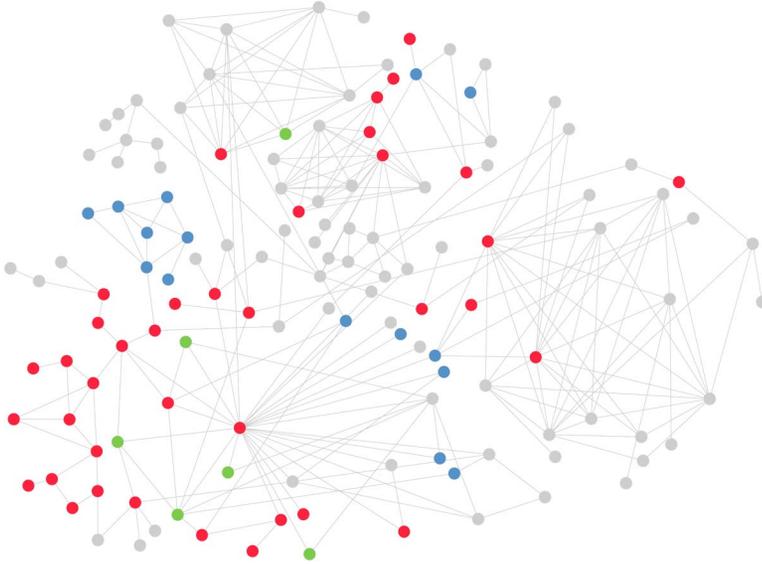


FIG. 2. A subnetwork composed of genes with periodic behavior. The subnetwork consists of 135 genes. Red nodes: genes functionally involved in the M-phase of cell cycle; blue node: genes functionally involved in the interphase of cell cycle; green nodes: genes functionally involved in both M and interphase of cell cycle.

the subnetwork, while the selected interphase genes are somewhat scattered, with 7 falling into a small but tight cluster.

We show part of the subnetwork detected by the NET-DPM-3 with the corresponding one under the STD-DPM in Figure 3, where the genes that are linked by a dashed line are connected to other genes that are not shown in the figure. In this subnetwork, the gene selection results by the NET-DPM-1 agree with the NET-DPM-3 except for only one gene “YML064” for which the NET-DPM-1 does not select it with probability 0.478, while the NET-DPM-3 selects it with probability 0.687. This implies that both methods provide large uncertainty on this gene. Comparing the top panel (our method, NET-DPM-3) and bottom panel (STD-DPM), we observe a number of genes selected by NET-DPM but not by STD-DPM, and almost all such genes are cell cycle-related (denoted by a star by the ORF name). Examples include YAL041W (CLS4), which is required for the establishment and maintenance of polarity and critical in bud formation [Chenevert, Valtz and Herskowitz (1994); Cherry et al. (2012)]. The gene only shows moderate periodic behavior, as denoted by the color of the node. However, due to its links to other genes that have strong periodic behavior, it is selected by our method as an interesting gene. Another example is YFL008W (SMC1). It is a subunit of the cohesion complex, which is essential in sister chromatid cohesion of mitosis and meiosis. The complex is also involved in double-strand DNA break repair [Strunnikov and Jessberger (1999); Cherry et al. (2012)]. Similar to CLS4, the periodic behavior of SMC1 is not strong enough. It is only selected when the information is borrowed from linked genes that are functionally related and show strong periodic behavior. A number of other cell cycle-related genes in Figure 3 are in a similar situation, for example, YBR106W, YDR052C, YJL157C, YGL003C and YMR076C. These examples clearly show the benefit of utilizing the biological information stored in the network structure.

To assess the functional relevance of the selected genes globally, we resort to mapping the genes onto gene ontology biological processes [Ashburner et al. (2000)]. We limit our search to the GO Slim terms using the mapper of the Saccharomyces Genome Database [Cherry et al. (2012)]. The full result is listed in the supplementary file. Clearly, the overrepresented GO Slim terms are centered around cell cycle. Here we discuss some GO terms that are nonredundant. Among the 216 selected genes, 70 (32.4%, compared to 4.5% among all genes) belong to the process response to DNA damage stimulus (GO:0006974). The term shares a large portion of its genes with DNA recombination (GO:0006310) and DNA replication (GO:0006260) processes, which are integral to the cell cycle. Sixty-seven of the selected genes (31.0%, compared to 4.7% among all genes) belong to the process mitotic cell cycle (GO:0000278). Twenty-six of the 67 genes are shared with response to DNA damage stimulus (GO:0006974). Forty-one of the selected genes (19.0%, compared to 3.0% among all genes) belong to the process regulation of cell cycle (GO:0051726), among which 29 also belong to mitotic cell cycle (GO:0000278). Thirty-one of the selected genes (14.4%, compared to 2.6% among

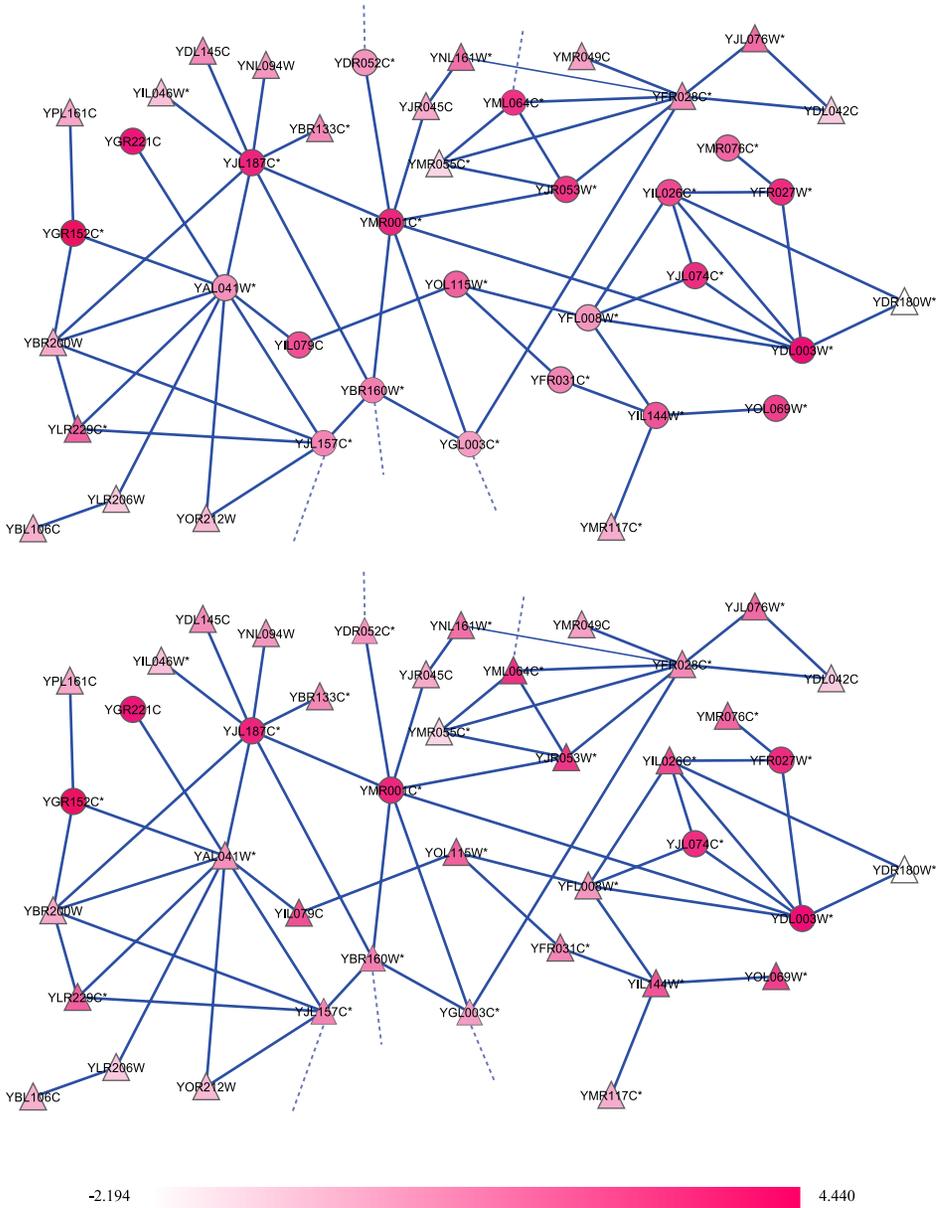


FIG. 3. A portion of the subnetwork shown in Figure 2, together with the immediate neighbors of the selected genes. Upper panel: NET-DPM-3 results; lower panel: STD-DPM results. The node labels indicate the gene name; circles and triangles represent “selected” and “unselected” genes; colors denote the value of the normal quantiles; a star in superscript represents the genes functionally annotated to the cell cycle process. Dash lines denote connections to genes not shown in the figure.

all genes) belong to the process meiotic cell cycle (GO:0051321), among which 12 are shared with mitotic cell cycle (GO:0000278). Other major enriched terms include chromatin organization (12.5%, compared to 3.5% overall), cytoskeleton organization (12.5%, compared to 3.4% overall), regulation of organelle organization (9.7%, compared to 2.4% overall) and cytokinesis (7.9%, compared to 1.7% overall). These terms clearly show strong relations with the yeast cell cycle.

4. Simulation studies. In this section we illustrate the performance of our methods (NET-DPMs) using simulation studies with various network structures and data settings compared with other methods. In Simulation 1 we study the similarity between the fully computational algorithm NET-DPM-1 and two fast computation approaches NET-DPM- x , $x = 2, 3$, in terms of gene selection accuracy and uncertainty estimations. Each of the three algorithms can be used along with one of the two methods for choosing hyperparameters: the posterior inference and model averaging. In Simulation 2 we focus on the gene network selection under a particular network structure and two types of simulated data to demonstrate the flexibility of the proposed methods. In both simulations we compare the NET-DPMs with a STD-DPM combined with the HODC algorithm without using any network information. In Section D in the supplemental article [Zhao, Kang and Yu (2014)], we also demonstrate the flexibility of the proposed methods by conducting a simulation on the selection of genes that are strongly associated with an outcome variable, and compare our NET-DPMs with a network based Bayesian variable selection (NET-BVS) proposed by Li and Zhang (2010).

4.1. *Simulation 1.* In this simulation we investigate the performance of the proposed algorithms using a simulated data set that mimic the real data in Section 3. We generate a scale-free network with 1000 genes based on the rich-get-rich algorithm [Barabási and Albert (1999)], that is, $n = 1000$. Two hub genes with 64 and 69 connections to other genes are in this network; the mean and median edges per gene are 1.998 and 1. The partial network structure with the two hub genes included is shown in Figure 4. From the network structure, we generate \mathbf{z} from the Ising model (2.3) with the sparsity parameter $\pi_0 = 0.8$ and smoothness parameters $\boldsymbol{\varrho} = (\varrho_0, \varrho_1) = (5, 10)$. For $i = 1, \dots, n$, in light of the results in Section 3, we simulate data r_i given z_i from the empirical distributions (Figure 5) of

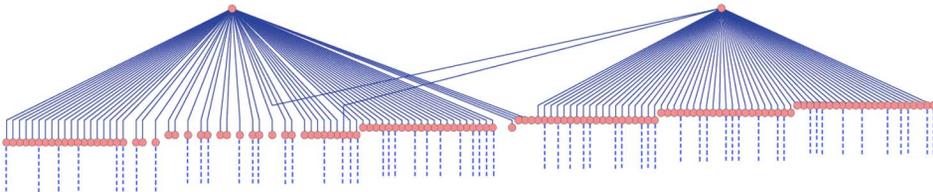


FIG. 4. Partial network structure with the dash lines representing connections to other nodes not shown in the figure.

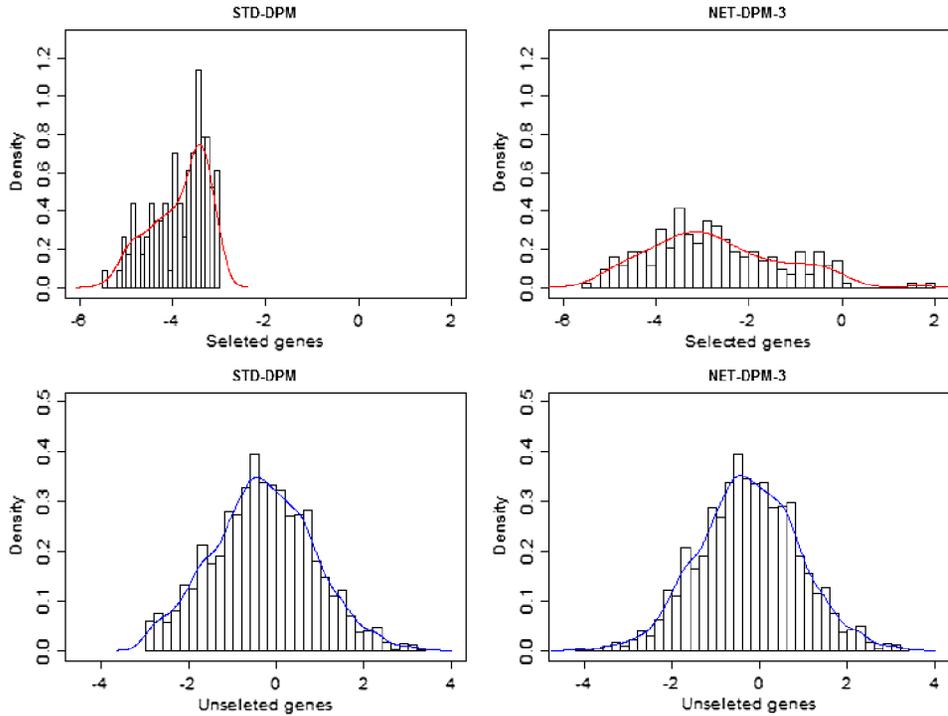


FIG. 5. Empirical distributions of “selected” genes (upper panel) and “unselected” genes (lower panel) in the Spellman yeast cell cycle data estimated by the NET-DPM-3 (right panel) and the STD-DPM (left panel).

the test statistics for “selected” and “unselected” genes in the Spellman yeast cell cycle microarray data. As shown in Section 3, the NET-DPM-3 (Scenario 1) and the STD-DPM (Scenario 2) provide different gene selections results. We set both scenarios as the truth to simulate data.

We apply the NET-DPM- x , for $x = 1, 2, 3$, and the STD-DPM to the simulated data set. To choose the sparsity and smoothness parameters, the NET-DPM-1 and the NET-DPM-3 are both combined with model averaging, where the possible choices of π_0 and $\boldsymbol{\varrho}$ are $(0.75, 0.8, 0.85, 0.9)$ and $(1, 5, 10, 20, 50) \times (1, 5, 10, 20, 50)$, while the NET-DPM-2 is combined with the posterior inference on $(\pi_0, \boldsymbol{\varrho})$. As for other hyperparameters, we specify $\tau_k, \xi_k, \gamma_k, \beta_k, \alpha_k; k = 0, 1$ the same way as in the data application for the NET-DPM- x , for $x = 1, 2$. With random starting values, each algorithm is run 10 times under 10,000 iterations with 2000 burn-in. For each gene i , the mode of the marginal posterior probability of z_i is taken to determine whether gene i is selected or not. The selection performance for each method based on the average of the 10 runs is presented in Table 2. We also compare the posterior probability estimates of \mathbf{z} between different algorithms under Scenario 1 in Figure 6.

TABLE 2
Gene selection accuracy in Simulation 1

	STD-DPM	NET-DPM-1	NET-DPM-2	NET-DPM-3
Scenario 1				
True positive rate	0.893	0.973	0.920	0.920
False positive rate	0.292	0.001	0.000	0.006
False discovery rate	0.801	0.014	0.000	0.080
Scenario 2				
True positive rate	1.000	1.000	1.000	1.000
False positive rate	0.232	0.000	0.000	0.007
False discovery rate	0.741	0.000	0.000	0.085
Typical computation time (hrs)	0.100	8.500	2.800	0.150

From Table 2, it is clear that the NET-DPMs achieve a better selection performance than the STD-DPM method under both scenarios. The STD-DPM without using the gene network information provides an extreme high false discovery rate in each scenario. This implies that it is critical to incorporate the gene network information to control FDR. Table 2 also suggests the NET-DPM-2 and the NET-DPM-3 approximate the NET-DPM-1 very well in terms of the gene selection accuracy with a substantial lower computational cost (3.4 GHz CPU, 8 GB Memory, Windows System). In addition, a comparison between the NET-DPM-2 and the NET-DPM-3 shows that the Bayesian model averaging over hy-

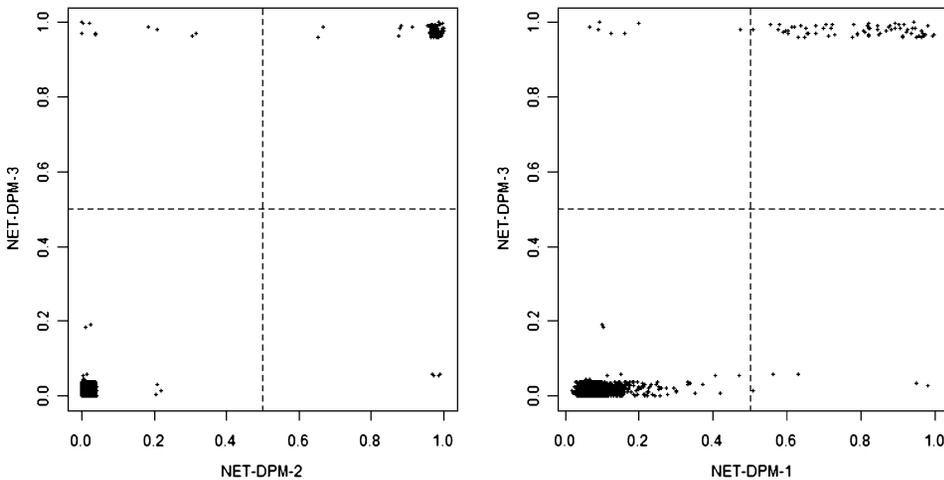


FIG. 6. Marginal posterior probabilities of the class labels of all 1000 genes by the different methods: NET-DPM-3 vs. NET-DPM-2 (left panel) and NET-DPM-2 vs. NET-DPM-1 (right panel). The probability values are jittered by tiny random noises for better presenting.

perparameters (π_0, ϱ) provides an efficient alternative to the standard Bayesian posterior inference procedure. For the posterior probability estimates, the NET-DPM-2 and the NET-DPM-3 achieve a good agreement, as shown in the left panel of Figure 6. However, in the right panel of Figure 6, compared with the NET-DPM-1, the NET-DPM-3 tends to provide larger probability estimates for the “selected” genes, but smaller probability estimates for “unselected” genes. This implies the fast computation approaches underestimate the uncertainty of gene selection.

4.2. *Simulation 2.* In this simulation we demonstrate the flexibility of the proposed methods and their ability to identify subnetworks of interest. We consider a 94-gene network which consists of an 11-gene subnetwork by design and an 83-gene scale-free network simulated from the rich-get-rich algorithm. The mean and median edges per node for the whole network are 2.02 and 1. Figure 7 shows the designed 11-gene subnetwork, where genes 5, 6 and 11 are connected with three other genes from the 83 gene scale-free network. Rather than simulating from priors, we directly specify the class label \mathbf{z} as $z_i = 1$ for $i \in \{1, 2, 3, 4, 5, 8, 9, 10\}$, $z_i = 0$ otherwise. In Figure 7 the blue nodes represent the “selected” genes and the red nodes are “unselected” genes. In addition, all other genes in the scale-free network (not shown in the figure) are “unselected.” The gene subnetwork of interest includes genes 1, 2, 3, 4 and 5, which are encircled by a rectangle frame in Figure 7. The null distribution for “unselected” r_i is specified as a standard normal distribution: $[r_i | z_i = 0] \sim N(0, 1)$. For the distribution of “selected” genes, we

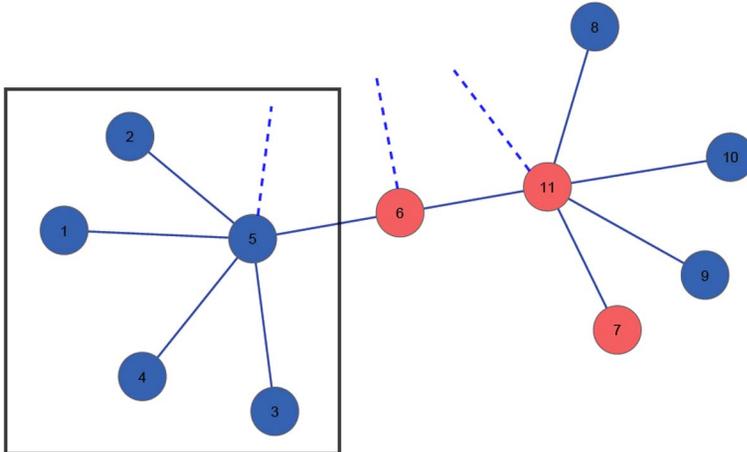


FIG. 7. *Partial simulated gene network structure: the blue nodes represent “selected” genes and the red nodes represent “unselected” genes. Dash lines denote connections to genes not shown in the figure. A subnetwork of interest includes nodes 1, 2, 3, 4 and 5, which are encircled by a rectangle frame.*

TABLE 3

Selection accuracy of gene subnetwork* by TPR (true positive rate), FPR (false positive rate) and FDR (false discovery rate) in Simulation 2

Method	Gaussian data			Non-Gaussian data		
	TPR	FPR	FDR	TPR	FPR	FDR
NET-DPM-3	63%	11%	15%	60%	5%	8%
STD-DPM	15%	33%	69%	17%	26%	60%

*For gene subnetwork selection, the TPR is defined as the percentage of exactly selecting the correct network. The FPR is the percentage of selecting a larger network containing the correct network and at least one more other gene that has connection to the network. The FDR is the proportion of falsely selecting a larger network among all the network discoveries (selecting a correct or larger network).

consider two settings:

$$\text{Gaussian data: } [r_i | z_i = 1] \sim 0.4 \times N(3, 1) + 0.6 \times N(2, 0.5),$$

$$\text{Non-Gaussian data: } [r_i | z_i = 1] \sim 0.4 \times G(5, 2) + 0.6 \times G(6, 3),$$

where $G(a, b)$ denotes a gamma distribution with shape a and rate b . According to the above procedure, we simulate 100 data sets for each type of data. We apply the NET-DPM-3 and the STD-DPM to each data set. We utilize the model averaging for choosing hyperparameters and a set of possible choices are given by $\{1, 2, 5, 10, 15\}$ for both ϱ_0 and ϱ_1 , and $\{0.8, 0.85, 0.9, 0.95\}$ for π_0 . We run 10,000 iterations with 2000 burn-in on each data set for both methods. In each simulated data set, we predetermine one gene as a “sure selected” gene. It has the largest number of connections with the “selected” genes estimated by the STD-DPM model.

Table 3 summarizes the selection accuracy of the gene subnetwork based on the 100 simulated data sets for each type of data. It is clear that the NET-DPM-3 provides much higher accuracy of the subnetwork selection than the STD-DPM. The NET-DPM-3 achieves a more than 60% accuracy rate in correctly identifying the subnetwork with an additional low false positive and false negative occurrences regardless of the type of data. This verifies the overall better performance of NET-DPM-3 than the STD-DPM in terms of identifying the gene subnetwork and the robustness of the proposed methods on different types of data.

5. Discussion. In the article we propose a Bayesian nonparametric mixture model for gene/gene subnetwork selection. Our model extends the standard DPM model incorporating the gene network information to significantly improve the accuracy of the gene selections and reduce the false discovery rate. We demonstrate that the proposed method has the ability to identify the subnetworks of genes and individual genes with a particular expressional behavior. We also show

that it is able to select genes which are strongly associated with clinical variables. We develop a posterior computation algorithm along with two fast approximation approaches. The posterior inference can produce more accurate uncertainty estimates of gene selection, while the fast computing algorithms can achieve a similar gene selection accuracy. Due to the nonparametric nature, our method has the flexibility to fit various data types and has robustness to model assumptions.

When we observe gene expression data along with measurements of a clinical outcome, we need to create statistics to perform the selection of genes that are strongly associated with the clinical outcome. The choice of the statistics is crucial to the performance of our methods. To model the relationship between the clinical outcome and gene expression data, much literature suggests a linear regression model [Li and Li (2008); Pan, Xie and Shen (2010); Li and Zhang (2010); Stingo et al. (2011)], from which we produce testing statistics or coefficient estimates as the candidates. For instance, as we suggest in Section D in the supplemental article [Zhao, Kang and Yu (2014)], the most straightforward approach is to fit simple linear regression on each gene and use the t -statistics as the input data to our methods. However, there is no scientific evidence that the relationship between gene expression profiles and the clinical outcome should follow a linear regression model. Without making this assumption, we may test the independence between each gene expression profile and the clinical outcome via a nonparametric model suggested by Einmahl and Van Keilegom (2008) and use our model to fit the testing statistics. Other potential choices of statistics for the nonlinear problems include mutual information statistics [Peng, Long and Ding (2005)] and maximal information coefficient (MIC) statistics [Reshef et al. (2011)].

Although the development of our method is motivated by gene selection problems, our method can conduct variable selection for a general purpose and it has broad applications. For example, functional neuroimaging studies (e.g., fMRI and PET) usually produce large-scale statistics, one for each voxel in the brain. Those statistics are used to localize the brain activity regions related to particular brain functions. This essentially is a voxel selection problem to which our method is applicable, where the networks may be defined according to the spatial locations of the voxels. In addition to this, we discuss two future directions:

- (1) It is common that we have multiple hypothesis tests for each gene, and we have interest in jointly analyzing these statistics. This motivates an extension of the current NET-DPM model from one dimension to multiple dimensions for multivariate large-scale statistics.
- (2) The selection of one gene might be affected by not only the genes that are directly connected to it, but also the genes close to it over the network. It

would be interesting to extend the prior specifications of the class label by incorporating a network distance. This should provide more biologically meaningful results.

Acknowledgments. The authors would like to thank the Editor, the Associate Editor and two referees for their helpful suggestions and constructive comments that substantially improved this manuscript.

SUPPLEMENTARY MATERIAL

Supplement to “A Bayesian nonparametric mixture model for selecting genes and gene subnetworks” (DOI: [10.1214/14-AOAS719SUPP](https://doi.org/10.1214/14-AOAS719SUPP); .pdf). In this online supplemental article we provide (A) derivations of the proposed methods, (B) details of the main algorithms for posterior computations, (C) details of posterior inference for hyperparameters, (D) additional simulation studies and (E) sensitivity analysis.

REFERENCES

- ANTONIAK, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Ann. Statist.* **2** 1152–1174. [MR0365969](#)
- ASHBURNER, M., BALL, C. A., BLAKE, J. A., BOTSTEIN, D., BUTLER, H., CHERRY, J. M., DAVIS, A. P., DOLINSKI, K., DWIGHT, S. S., EPPIG, J. T., HARRIS, M. A., HILL, D. P., ISSEL-TARVER, L., KASARSKIS, A., LEWIS, S., MATESE, J. C., RICHARDSON, J. E., RINGWALD, M., RUBIN, G. M. and SHERLOCK, G. (2000). Gene ontology: Tool for the unification of biology. The Gene Ontology Consortium. *Nat. Genet.* **25** 25–29.
- BARABÁSI, A.-L. and ALBERT, R. (1999). Emergence of scaling in random networks. *Science* **286** 509–512. [MR2091634](#)
- BREEZE, E., HARRISON, E., MCHATTIE, S., HUGHES, L., HICKMAN, R., HILL, C., KIDDLE, S., KIM, Y., PENFOLD, C., JENKINS, D. et al. (2011). High-resolution temporal profiling of transcripts during arabidopsis leaf senescence reveals a distinct chronology of processes and regulation. *The Plant Cell Online* **23** 873–894.
- CERAMI, E., GROSS, B., DEMIR, E., RODCHENKOV, I., BABUR, Ö., ANWAR, N., SCHULTZ, N., BADER, G. and SANDER, C. (2011). Pathway commons, a web resource for biological pathway data. *Nucleic Acids Res.* **39** D685–D690.
- CHENEVERT, J., VALTZ, N. and HERSKOWITZ, I. (1994). Identification of genes required for normal pheromone-induced cell polarization in *saccharomyces cerevisiae*. *Genetics* **136** 1287–1297.
- CHERRY, J. M., HONG, E. L., AMUNDSEN, C., BALAKRISHNAN, R., BINKLEY, G., CHAN, E. T., CHRISTIE, K. R., COSTANZO, M. C., DWIGHT, S. S., ENGEL, S. R., FISK, D. G., HIRSCHMAN, J. E., HITZ, B. C., KARRA, K., KRIEGER, C. J., MIYASATO, S. R., NASH, R. S., PARK, J., SKRZYPEK, M. S., SIMISON, M., WENG, S. and WONG, E. D. (2012). Saccharomyces Genome Database: The genomics resource of budding yeast. *Nucleic Acids Res.* **40** D700–D705.
- DO, K.-A., MÜLLER, P. and TANG, F. (2005). A Bayesian mixture model for differential gene expression. *J. Roy. Statist. Soc. Ser. C* **54** 627–644. [MR2137258](#)
- DUNSON, D. B. (2010). Nonparametric Bayes applications to biostatistics. In *Bayesian Nonparametrics* (N. L. Hjort, C. Holmes, P. Müller and S. G. Walker, eds.) 223–273. Cambridge Univ. Press, Cambridge. [MR2730665](#)

- EFRON, B. (2004). Large-scale simultaneous hypothesis testing: The choice of a null hypothesis. *J. Amer. Statist. Assoc.* **99** 96–104. [MR2054289](#)
- EFRON, B. (2010). Correlated z -values and the accuracy of large-scale statistical estimates. *J. Amer. Statist. Assoc.* **105** 1042–1055. [MR2752597](#)
- EINMAHL, J. H. J. and VAN KEILEGOM, I. (2008). Tests for independence in nonparametric regression. *Statist. Sinica* **18** 601–615. [MR2411617](#)
- ESCOBAR, M. D. (1994). Estimating normal means with a Dirichlet process prior. *J. Amer. Statist. Assoc.* **89** 268–277. [MR1266299](#)
- ESCOBAR, M. D. and WEST, M. (1995). Bayesian density estimation and inference using mixtures. *J. Amer. Statist. Assoc.* **90** 577–588. [MR1340510](#)
- ISHWARAN, H. and JAMES, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *J. Amer. Statist. Assoc.* **96** 161–173. [MR1952729](#)
- ISHWARAN, H. and JAMES, L. F. (2002). Approximate Dirichlet process computing in finite normal mixtures: Smoothing and prior information. *J. Comput. Graph. Statist.* **11** 508–532. [MR1938445](#)
- JACOB, L., NEUVIAL, P. and DUDOIT, S. (2012). More power via graph-structured tests for differential expression of gene networks. *Ann. Appl. Stat.* **6** 561–600. [MR2976483](#)
- LENG, X. and MÜLLER, H.-G. (2006). Classification using functional data analysis for temporal gene expression data. *Bioinformatics* **22** 68–76.
- LI, C. and LI, H. (2008). Network-constrained regularization and variable selection for analysis of genomic data. *Bioinformatics* **24** 1175–1182.
- LI, F. and ZHANG, N. R. (2010). Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics. *J. Amer. Statist. Assoc.* **105** 1202–1214. [MR2752615](#)
- MA, S., SHI, M., LI, Y., YI, D. and SHIA, B.-C. (2010). Incorporating gene co-expression network in identification of cancer prognosis markers. *BMC Bioinformatics* **11** 271.
- MÜLLER, P. and QUINTANA, F. A. (2004). Nonparametric Bayesian data analysis. *Statist. Sci.* **19** 95–110. [MR2082149](#)
- NEAL, R. M. (2000). Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Statist.* **9** 249–265. [MR1823804](#)
- PAN, W., XIE, B. and SHEN, X. (2010). Incorporating predictor network in penalized regression with application to microarray data. *Biometrics* **66** 474–484. [MR2758827](#)
- PENG, H., LONG, F. and DING, C. (2005). Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Analysis and Machine Intelligence* **27** 1226–1238.
- QU, L., NETTLETON, D. and DEKKERS, J. C. M. (2012). A hierarchical semiparametric model for incorporating intergene information for analysis of genomic data. *Biometrics* **68** 1168–1177. [MR3040023](#)
- RESHEF, D. N., RESHEF, Y. A., FINUCANE, H. K., GROSSMAN, S. R., MCVEAN, G., TURNBAUGH, P. J., LANDER, E. S., MITZENMACHER, M. and SABETI, P. C. (2011). Detecting novel associations in large data sets. *Science* **334** 1518–1524.
- SPELLMAN, P. T., SHERLOCK, G., ZHANG, M. Q., IYER, V. R., ANDERS, K., EISEN, M. B., BROWN, P. O., BOTSTEIN, D. and FUTCHER, B. (1998). Comprehensive identification of cell cycle-regulated genes of the yeast *Saccharomyces cerevisiae* by microarray hybridization. *Mol. Biol. Cell* **9** 3273–3297.
- STINGO, F. C., CHEN, Y. A., TADESSE, M. G. and VANNUCCI, M. (2011). Incorporating biological information into linear models: A Bayesian approach to the selection of pathways and genes. *Ann. Appl. Stat.* **5** 1978–2002. [MR2884929](#)
- STRUNNIKOV, A. V. and JESSBERGER, R. (1999). Structural maintenance of chromosomes (SMC) proteins: Conserved molecular properties for multiple biological functions. *Eur. J. Biochem.* **263** 6–13.
- TANG, Y., GHOSAL, S. and ROY, A. (2007). Nonparametric Bayesian estimation of positive false discovery rates. *Biometrics* **63** 1126–1134, 1312. [MR2414590](#)

- WANG, L. and DUNSON, D. B. (2011). Fast Bayesian inference in Dirichlet process mixture models. *J. Comput. Graph. Statist.* **20** 196–216. Supplementary material available online. [MR2816545](#)
- WEI, Z. and LI, H. (2007). A Markov random field model for network-based analysis of genomic data. *Bioinformatics* **23** 1537–1544.
- WEI, Z. and LI, H. (2008). A hidden spatial-temporal Markov random field model for network-based analysis of time course gene expression data. *Ann. Appl. Stat.* **2** 408–429. [MR2415609](#)
- WEI, P. and PAN, W. (2010). Network-based genomic discovery: Application and comparison of Markov random-field models. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **59** 105–125. [MR2750134](#)
- WEI, P. and PAN, W. (2012). Bayesian joint modeling of multiple gene networks and diverse genomic data to identify target genes of a transcription factor. *Ann. Appl. Stat.* **6** 334–355. [MR2951540](#)
- WICHERT, S., FOKIANOS, K. and STRIMMER, K. (2004). Identifying periodically expressed transcripts in microarray time series data. *Bioinformatics* **20** 5–20.
- XENARIOS, I., SALWIŃSKI, L., DUAN, X. J., HIGNEY, P., KIM, S.-M. and EISENBERG, D. (2002). DIP, the Database of Interacting Proteins: A research tool for studying cellular networks of protein interactions. *Nucleic Acids Res.* **30** 303–305.
- YU, T. (2010). An exploratory data analysis method to reveal modular latent structures in high-throughput data. *BMC Bioinformatics* **11** 440.
- ZHAO, Y., KANG, J. and YU, T. (2014) Supplement to “A Bayesian nonparametric mixture model for selecting genes and gene subnetworks.” DOI:[10.1214/14-AOAS719SUPP](#).
- ZHOU, B., XU, W., HERNDON, D., TOMPKINS, R., DAVIS, R., XIAO, W., WONG, W., TONER, M., WARREN, H., SCHOENFELD, D. et al. (2010). Analysis of factorial time-course microarrays with application to a clinical study of burn injury. *Proc. Natl. Acad. Sci. USA* **107** 9923.

DEPARTMENT OF BIOSTATISTICS AND
BIOINFORMATICS
EMORY UNIVERSITY
1518 CLIFTON RD.
ATLANTA, GEORGIA 30322
USA
E-MAIL: yize.zhao@emory.edu
jian.kang@emory.edu
tianwei.yu@emory.edu