



Modeling Dependent Gene Expression

The Harvard community has made this article openly available. [Please share](#) how this access benefits you. Your story matters

Citation	Telesca, Donatello, Peter Müller, Giovanni Parmigiani, and Ralph S. Freedman. 2012. "Modeling Dependent Gene Expression." <i>The Annals of Applied Statistics</i> 6 (2): 542–60. https://doi.org/10.1214/11-aos525 .
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:41552044
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

MODELING DEPENDENT GENE EXPRESSION

BY DONATELLO TELESCA, PETER MÜLLER,
GIOVANNI PARMIGIANI AND RALPH S. FREEDMAN

UCLA School of Public Health, University of Texas M.D. Anderson Cancer Center, Dana Farber Cancer Institute and Harvard School of Public Health and University of Texas M.D. Anderson Cancer Center

In this paper we propose a Bayesian approach for inference about dependence of high throughput gene expression. Our goals are to use prior knowledge about pathways to anchor inference about dependence among genes; to account for this dependence while making inferences about differences in mean expression across phenotypes; and to explore differences in the dependence itself across phenotypes. Useful features of the proposed approach are a model-based parsimonious representation of expression as an ordinal outcome, a novel and flexible representation of prior information on the nature of dependencies, and the use of a coherent probability model over both the structure and strength of the dependencies of interest. We evaluate our approach through simulations and in the analysis of data on expression of genes in the Complement and Coagulation Cascade pathway in ovarian cancer.

1. Introduction. Inferring patterns of dependence from high throughput geneomic data poses significant challenges. Statistically, the problem is one of learning about dependence structures in high dimension, with relatively low signal. A promising direction for strengthening this inference is the explicit consideration of information from known “pathways”—biochemical processes described in terms of a series of relationships among genes and their products.

In this paper we take this perspective, and propose a Bayesian approach to achieve three related goals in the context of gene expression analysis: to use prior knowledge about pathways to anchor inference about dependence among genes; to account for this dependence while making inferences

Received February 2010; revised October 2011.

Key words and phrases. Conditional independence, microarray data, probability of expression, probit models, reciprocal graphs, reversible jumps MCMC.

<p>This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in <i>The Annals of Applied Statistics</i>, 2012, Vol. 6, No. 2, 542–560. This reprint differs from the original in pagination and typographic detail.</p>

about differences in mean expression across phenotypes; and to explore differences in the dependence itself across phenotypes. The proposed model builds on the POE model (Parmigiani et al. 2002) and integrates inference about probability of differential expression with inference about dependence between genes through the formulation of a coherent probability model. Our proposed inferences are local in the sense that the model is centered around a specific pathway. Formally, variable selection is used to remove and add structure relative to the centering pathway. This is in contrast to approaches aimed at learning dependence structures de novo from expression data, without guidance by a prior pathway structure.

Some of the existing approaches for probabilistic modeling of dependence structures attempt to explore the space of all possible graphical models, often restricted to Directed Acyclic Graphs (DAGs) or Bayesian networks (BN) (Lauritzen 1996) and decomposable models (Dawid and Lauritzen 1993). A comprehensive review of statistical methodology for network data is provided in Kolaczyk (2009). Recent literature includes the application of BN and dynamic BN to microarray data (Murphy and Mian 1999, Friedman et al. 2000), with applications and extensions of this methodology reported in Ong, Glasner and Page (2002) and Beal et al. (2005), among others. Although appealing, these techniques have computational and methodological limitations related to modeling conditional independence under the “large p , small n ” paradigm and the difficult specification of consistent prior models across dimensions (Dobra et al. 2004). Other authors (Scott and Carvalho 2008, Jones et al. 2005) have reported difficulties with the performance of standard trans-dimensional MCMC methods (Giudici and Green 1999) in the exploration of the model space, and suggested alternative stochastic search schemes. For a decision theoretic perspective on graphical model selection see Sebastiani and Ramoni (2005).

To overcome these problems, we focus on variations of a baseline model that represents known dependence structures. The centering anchors the model space to a prior path diagram elicited from sets of molecular interactions derived by previous experiments.

Our idea is similar to the modeling approaches described in Wei and Li (2007) and Wei and Li (2008), who introduced conditional independence between genes, via a Markov random field defined over binary hidden states of differential expression. These authors propose to consider a fixed Markov random field mirroring exactly the topology of a prior pathway and ignoring the directionality of the edges. The construction of dependence patterns based on hidden Markov random fields had also been previously explored by Broët and Richardson (2006) in the analysis of CGH microarrays. We contrast the approach of Wei and Li (2007, 2008) in three fundamental ways. First, we provide an alternative interpretation of the connections encoded

into a prior pathway. We develop a prior model for the dependence structure, that is, based on the reciprocal graphs (Koster 1996). This class of graphical models takes account of the directionality of the edges included in the pathway and allows for the Markovian characterization of cycles, which often arise in biological depictions of genetic interactions. Also, recognizing that a known pathway is often summarizing results obtained under different experimental conditions, we allow for significant deviations from the prior dependence structure. This extension requires explicit consideration of a model determination strategy, but enables inference on the model parameters as well as inference on the dependence structure between genes. Finally, our focus is on identifying significant interactions between genes in a prior pathway, as opposed to identifying differentially expressed genes in a given pathway.

The proposed methodology finds motivation in the analysis of gene expression of Epithelial Ovarian Cancer (EOC) patients. In this setting, the complement and coagulation cascade pathway (Figure 2) represents a key study target, as disease progression is thought to be highly linked to inflammation and vascularization processes [Wang, Wang and Kavanagh (2005)].

The rest of this article is organized as follows. In Section 2 we introduce the proposed model. Section 3 discusses estimation and inferential details associated with the proposed model. We validate our approach with a simulation study in Section 4. Section 5 employs the model for the analysis of epithelial ovarian cancer expression data, to derive inference about active genetic interactions. In the example, a well-known molecular pathway provides prior information on the dependence structure. A final section concludes with a critical discussion of limitations and possible extensions.

2. Dependent probability of expression. In Section 2.1, we discuss graphical models and notation, and in Section 2.2, we review the POE (Probability of Expression) model Parmigiani et al. (2002), which defines biological events via latent three-way indicators of relevant biological states. The original POE model assumes independence across genes, conditional on hyperparameters. We extend the original model by formalizing more complex relationships among variables via a cascade of conditional dependences, guided by a predefined interaction map. The predefined interaction map is formally represented as a graph. In Section 2.3 we introduce a prior probability model on this graph.

2.1. Representing dependence through graphical models. Networks of relationships among expression levels can be represented as graphs that describe how genes influence each other [for an example in ovarian cancer see Wang, Wang and Kavanagh (2005)]. More formally, a graph is often characterized as an algebraic structure $\mathcal{G} = \{V, E\}$, composed of a set of nodes V , in

our case genes, and a set of edges $E \subseteq \{(v_{i_1}, v_{i_2}), v_i \in V\} \cup \{\{v_{i_1}, v_{i_2}\}, v_i \in V\}$. Here (v_i, v_j) denotes a directed edge from v_i to v_j , and $\{v_i, v_j\}$ denotes an undirected edge. A graph \mathcal{G} defines the Markov properties of a statistical model in a graphical fashion, via the specification of a set of conditional dependencies.

Biochemical networks often include the presence of cycles and feedback relationships. This requires some care when trying to characterize a coherent probabilistic model that accurately portrays prior biochemical knowledge. For this purpose, we focus on a class of graphical models known as reciprocal graphs [Koster (1996)]. Reciprocal graphs are defined as a natural generalization of other well-known classes, including directed acyclic graphs (DAG) and Markov random fields, among others. Reciprocal graphs are defined through the coherent probabilistic interpretation of directed $a(\rightarrow b)$, undirected $(a - b)$ and reciprocal edges $(a \rightleftarrows b)$. Here, for simplicity, we consider a subset of the reciprocal graph family excluding undirected edges. The restriction to only directed edges will later be important to facilitate the mapping of \mathcal{G} to a simultaneous equation model.

The proposed model and inference is based on the directed graph \mathcal{G} . But sometimes it is of interest to describe the implied conditional independence structure, that is, the Markov properties. When desired, the Markov properties of our model are defined in terms of an undirected graph $\mathcal{G}^m = \{V, E^m\}$ elicited via moralization [Koster (1996), Lauritzen (1996)] of a graph \mathcal{G} . In substance, the moralization procedure consists in adding an undirected edge between parents of a common child and replacing the remaining directed edges with undirected ones. In \mathcal{G}^m , standard Markov field properties hold, in the sense that two genes i and j are disconnected when they are conditionally independent, given the rest of the genes [Besag (1974)]. For example, consider the reciprocal graph \mathcal{G} represented in Figure 1. The class of Markov equivalent models spanned by \mathcal{G} may be represented with the moral (undirected) graph \mathcal{G}^m , for which the following Markov property holds: $1 \perp 2|3, 4$, that is, $P(1, 3|2, 4) = P(1|2, 4)P(3|2, 4)$. The correspondence between \mathcal{G} and \mathcal{G}^m is not 1-to-1 as \mathcal{G}^m could arise from the moralization of an entire class of Markov equivalent reciprocal graphs. Further details about moralization in reciprocal graphs and covariance equivalence

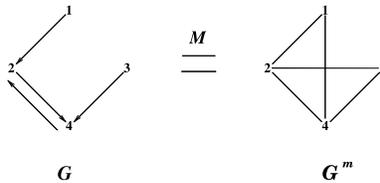


FIG. 1. (Example) moralization of a reciprocal graph.

are discussed in Koster (1996) and Spirtes et al. (1998). Here, our inference will be based on \mathcal{G} only, and the directionality will be based on prior knowledge. The undirected graph \mathcal{G}^m provides a convenient summary of the conditional independence structure if desired.

2.2. Dependent gene expression and hidden systems of simultaneous equations. Following Parmigiani et al. (2002), we consider data in the form of an $(p \times n)$ expression matrix \mathbf{Y} , with the generic element y_{ij} denoting the observed gene expression for gene i in sample j , $i = 1, \dots, p$ and $j = 1, \dots, n$. We introduce latent variables $e_{ij} \in \{-1, 0, 1\}$ indexing three possible expression categories for each entry in \mathbf{Y} . For example, if \mathbf{Y} represents ratios of expression level relative to a normal reference, they can be interpreted as high, normal and low. Given e_{ij} , for each gene i and each sample j we assume a mixture parameterized with $\theta = (\alpha_j, \mu_i, \kappa_i^-, \kappa_i^+)$ as follows:

$$(1) \quad p(y_{ij} - (\alpha_j + \mu_i) | e_{ij}) = f_{ij}(y_{ij}) \quad \text{with} \quad \begin{cases} f_{-1i} = U(-\kappa_i^-, 0), \\ f_{0i} = N(0, \sigma_i^2), \\ f_{1i} = U(0, \kappa_i^+). \end{cases}$$

In words, we assume that the observed expressions arise from a mixture of a Gaussian distribution and two uniform distributions designed to capture a broad range of departures relative to the Gaussian. The interpretation of the Gaussian component varies depending on the experimental design and sampling scheme, and can be trained in a supervised way if data are available (Garrett and Parmigiani 2004). When the technology used for measuring expression has an internal reference, as in Section 5, the high (low) class can be interpreted as over- (under) expression compared to the reference. The upcoming definition of a dependence structure will focus on the latent e_{ij} and define dependence at the level of these indicators. In other words, the proposed model could be characterized as a boolean network on the latent e_{ij} .

In (1), α_j is a sample-specific effect, included to adjust for systematic variation across samples; μ_i is a gene-specific effect, modeling the overall abundance of each gene, and κ_i^- and κ_i^+ parameterize the limits of variation in the tails. Finally, σ_i^2 is the variance of the normal component of the distribution of gene i . We follow Parmigiani et al. (2002) in defining a conditionally conjugate prior for μ_i , σ_i^2 and κ_i^- and κ_i^+ . Let $\mathcal{G}a(a, b)$ denote a Gamma distribution with expectation a/b :

$$\begin{aligned} p(\mu_i | m_\mu, \tau_\mu) &= N(m_\mu, \tau_\mu), & p(1/\sigma_i^2 | \gamma_\sigma, \lambda_\sigma) &= \mathcal{G}a(\gamma_\sigma, \lambda_\sigma), \\ p(1/\kappa_i^- | \gamma_\kappa^-, \lambda_\kappa^-) &= \mathcal{G}a(\gamma_\kappa^-, \lambda_\kappa^-), & p(1/\kappa_i^+ | \gamma_\kappa^+, \lambda_\kappa^+) &= \mathcal{G}a(\gamma_\kappa^+, \lambda_\kappa^+); \end{aligned}$$

where $\min(\kappa_i^+, \kappa_i^-) > \kappa_0 \sigma_i$ and $\kappa_0 = 5$. The restriction on κ_i^- and κ_i^+ ensures that the gene-specific mixture distribution has heavier tails than its normal

component, preserving interpretability of the three-way latent classes. For the sample-specific effect α_j , we impose an identifiability constraint $\alpha_j \sim N(0, \tau_\alpha^2)$ with $\sum_{j=0}^n \alpha_t = 0$.

Specifying a prior model for e_{ij} , we deviate from Parmigiani et al. (2002), defining the model in terms of latent normal variables [Albert and Chib (1993)]. For each gene and sample we introduce a latent Gaussian variable z_{ij} , and define

$$(2) \quad e_{ij} = \begin{cases} 1, & \text{if } z_{ij} > 1 \text{ high expression,} \\ 0, & \text{if } -1 < z_{ij} \leq 1 \text{ normal expression,} \\ -1, & \text{if } z_{ij} \leq -1 \text{ low expression,} \end{cases}$$

where the distribution of z_{ij} is defined by the following simultaneous equations model (SEM):

$$(3) \quad z_{ij} = m_{ij} + \sum_{k \neq i} \beta_{ik}(z_{kj} - m_{kj}) + \varepsilon_{ij}, \quad i = 1, \dots, p, j = 1, \dots, n,$$

with $\varepsilon_{ij} \sim N(0, s_i^2)$. Let $\mathbf{Z}_j = (z_{1j}, \dots, z_{pj})'$ denote the p -dimensional vector of latent probit scores associated with sample j . Also, let \mathbf{B} be the $(p \times p)$ matrix whose diagonal elements are unity and whose off-diagonal (i, k) components are $-\beta_{ik}$. Provided \mathbf{B} is nonsingular, the process above defines a proper joint probability density function [Besag (1974)]. More precisely, defining the marginal precision matrix $\mathbf{H}_z = \text{diag}(1/s_1, \dots, 1/s_p)$ and $\mathbf{\Omega} = \mathbf{B}'\mathbf{H}_z\mathbf{B}$, we have

$$(4) \quad P(\mathbf{Z}_j | \mathbf{m}_j, \mathbf{\Omega}) = \frac{|\mathbf{\Omega}|^{1/2}}{(2\pi)^{1/2}} \exp\left\{-\frac{1}{2}(\mathbf{Z}_j - \mathbf{m}_j)'\mathbf{\Omega}(\mathbf{Z}_j - \mathbf{m}_j)\right\},$$

where $\mathbf{m}_j = (m_{1j}, \dots, m_{pj})'$.

If $\mathbf{e}_j = (e_{1j}, \dots, e_{pj})'$, the implied probabilities for the indicators e_{ij} are

$$(5) \quad P(\mathbf{e}_j | \mathbf{m}_j, \mathbf{\Omega}) = \int_{A_{pj}} \dots \int_{A_{1j}} P(\mathbf{Z}_j | \mathbf{m}_j, \mathbf{\Omega}) d\mathbf{Z}_j,$$

where A_{ij} is the interval $(-\infty, -1]$ if $e_{ij} = -1$, $(-1, 1]$ if $e_{ij} = 0$ and $(1, \infty)$ if $e_{ij} = 1$. We use notation $\pi_{ij}^+ = p(z_{ij} > 1 | y)$, $\pi_{ij}^- = p(z_{ij} < -1 | y)$ and $p_{ij}^* = \pi_{ij}^+ - \pi_{ij}^-$.

In the context of this SEM, we propose to use a reciprocal graph, $\mathcal{G} = \{V, E\}$, to describe a dependence structure among the three-way indicators e_{ij} that reflects a priori knowledge about a pathway. Relationships between genes are captured via a set of conditional independences over the joint distribution of the classes $\mathbf{e}_j = (e_{ij}, i = 1, \dots, p)$. This is implemented by structuring the matrix \mathbf{B} so that the off-diagonal element (i, k) is null ($\beta_{ik} = 0$), if and only if the edge $k \rightarrow i$ is not in $\{E\}$ [Spirtes et al. (1998)]. The resulting concentration matrix $\mathbf{\Omega} = \mathbf{B}'\mathbf{H}_z\mathbf{B}$ will have zero off-diagonal

elements ($\omega_{ik} = 0$) structured compatibly with the Markov properties encoded in the moral graph $\mathcal{G}^m = \{V, E^m\}$ [Koster (1996)]. In summary, we use the SEM to define a probability model that matches the conditional independence structure given by \mathcal{G} . The coefficients \mathbf{B} of the SEM index a family of probability models that adhere to a given independence structure \mathcal{G} , including an interpretation of the edge directions.

For each gene and sample, the mean m_{ij} may be modeled as a linear function ($m_{ij} = \mathbf{x}_j' b_i$) of, say, a design vector \mathbf{x}_j . This allows for comparisons across groups. For example, if $\mathbf{x}_j = 1$ and -1 for samples under two different biologic conditions, then the posterior distribution for b_i formalizes inference on the differential expression of gene i under the two conditions, adjusting for the dependence among the genes.

Finally, the autoregressive scheme in (3) implicitly assumes that genetic interactions are invariant across all the cross-sample biological variation represented in the study. Relaxing this assumption is important and can be achieved by including an interaction term relating the covariate or phenotype information in \mathbf{x}_t with the neighboring probit scores \mathbf{z}_{kj} in (3).

In summary, we assume a mixture model for the observed gene expressions y_{ij} . The noisy data y_{ij} is reduced to latent trinary indicators which are used to define the dependence structure. Because of the nonlinear shrinkage induced by the mixture model, the y_{ij} do not come from a multivariate normal, and the patterns of dependence could be more complex.

2.3. Priors over graphical structures and dependence parameters. We define a prior probability model for the dependence structure \mathcal{G} . In words, the prior is based on a pathway diagram that summarizes substantive prior information about the pathway of interest. We interpret the pathway as a reciprocal graph $\mathcal{G}_0 = \{V, E_0\}$ (see example in Section 2.1). The prior on \mathcal{G} is defined on the set of all graphs that can be obtained by deleting edges from \mathcal{G}_0 . More formally, we define the model space generated by \mathcal{G}_0 as $M(\mathcal{G}_0) = \{\mathcal{G} = (V, E) : E \subset E_0\}$. If E_0 comprises a total number of K edges, then $M(\mathcal{G}_0)$ includes $D = 2^K$ possible models.

The definition of the the prior $p(\mathcal{G})$ can be seen as stating joint probabilities for the multiple hypothesis testing problem implicitly defined by inclusion versus exclusion of all possible edges. Following the standard Bayesian variable selection scheme [George and McCulloch (1993), Brown, Vannucci and Fearn (1998), Dobra et al. (2004)], we can consider edge inclusions as exchangeable Bernoulli trials with common inclusion probability φ . If $k_{\mathcal{G}}$ is the number of edges included in \mathcal{G} , it follows that $P(\mathcal{G}|\varphi) = \varphi^{k_{\mathcal{G}}}(1 - \varphi)^{K - k_{\mathcal{G}}}$. When the inclusion probability φ comes from the Beta family ($\varphi \sim \mathcal{B}(a_{\varphi}, b_{\varphi})$), Scott and Berger (2006) and Carvalho and Scott (2009) show that this class of prior model probabilities yield a strong control over the number of “false” edges included in \mathcal{G} . The associated marginal prior on \mathcal{G} becomes $p(\mathcal{G}) = \Gamma(\kappa_{\mathcal{G}} + a_{\varphi})\Gamma(K + b_{\varphi} - \kappa_{\mathcal{G}})/\Gamma(K + b_{\varphi} + \kappa_{\mathcal{G}})$.

A key feature of the proposed prior is the restriction to subsets of \mathcal{G}_0 . Inference under the proposed model populates existing pathways with probabilistic information associated with a biological system at a temporal cross section of its dynamic. The restriction to $M(\mathcal{G}_0)$ is important to keep MCMC posterior simulation across the model space practicable. For global searches, without restriction to a focused set of models, trans-dimensional MCMC becomes impracticable. Local focus does not preclude some extensions beyond $M(\mathcal{G}_0)$ to facilitate discovery of previously unknown interactions. For example, consider an arbitrary graph \mathcal{G} , without restriction to $M(\mathcal{G}_0)$, and let $m_{\mathcal{G}}$ denote the number of deleted *and* added edges relative to \mathcal{G}_0 . One could replace $k_{\mathcal{G}}$ in the prior by $m_{\mathcal{G}}$ and allow for graphs beyond \mathcal{G}_0 . Little would change in the proposed inference. But centering on models close to \mathcal{G}_0 is important. See also related comments in Section 6.

Our model is completed defining priors over the nonzero parameters $\beta_{ij} \sim N(0, \sigma_{\beta}^2)$ ($i, j = 1, \dots, p$). This defines a conjugate prior for the normal model (3). This formulation is derived as a natural characterization of the SEM in (3).

We recognize that assuming full exchangeability over the edges does not make active use of potential prior information on inter-gene relationships, possibly available through public data-bases like KEGG or Gene Ontology. We note, however, that fine scale prior information on individual interactions is easily included in the proposed inferential framework defining partially exchangeable or independent Bernoulli trials with interaction-specific inclusion probabilities, say, φ_{ij} . If desired, the model can be extended with Beta(a_{ij}, c_{ij}) for φ_{ij} , with hyperparameters chosen to reflect interaction summaries perhaps elicited through available tools like the R package GOSim. These elicitation processes are, however, still the subject of active research [Fröhlich et al. (2007), Mistry and Pavlidis (2008)]. We therefore limit our analysis to purely structural priors.

3. Estimation and inference.

3.1. *Model determination via RJ-MCMC.* Let $\boldsymbol{\theta}$ denote all population parameters and unknown quantities directly associated with the sampling model introduced in (1). We implement posterior inference for $(\boldsymbol{\theta}, \mathbf{B}, \mathcal{G})$ by setting up posterior MCMC simulation. We define the current state $x = (\boldsymbol{\theta}, \mathbf{B}, \mathcal{G})$ as the complete set of unknowns and write $\pi(dx)$ short for the target posterior distribution $p(\boldsymbol{\theta}, \mathbf{B}, \mathcal{G} | \mathbf{Y})$.

The MCMC is defined by the following transition probabilities: (a) Update the parameter vector $(\boldsymbol{\theta}, \mathbf{B})$; and (b) Update \mathcal{G} , ensuring that the proposed graph \mathcal{G}' is in the set $M\{\mathcal{G}_0\}$. This move usually involves changes to \mathbf{B} as well.

The updates in (a) follow the usual Metropolis-within-Gibbs scheme. We sample components of $\boldsymbol{\theta}$ directly from their conditional posterior distributions (Gibbs sampling details are reported in the [Appendix](#)). We update the matrix \mathbf{B} by row via multivariate random-walk Metropolis–Hastings transition probabilities. Let $pa(i)$ denote the parent nodes of node i in the directed graph \mathcal{G} . We define the i th row of \mathbf{B} as $\boldsymbol{\beta}_i$ and propose a new state $\boldsymbol{\beta}'_i | \boldsymbol{\beta}_{pa(i)} \sim N_{|pa(i)|}(\boldsymbol{\beta}_i | \boldsymbol{\beta}_{pa(i)}; V_{\boldsymbol{\beta},i}^*)$, where $V_{\boldsymbol{\beta},i}^* = c(1/s_i^2 W_i^T W_i + 1/\sigma_\beta^2 I)^{-1}$. Here W_i is an $(n \times |pa(i)|)$ design matrix including all mean adjusted probit scores for parents of gene i and c is a Metropolis–Hastings tuning parameter. For each row, this proposal scheme changes \mathbf{B} to \mathbf{B}' , defining a local approximation of a reciprocal graph by a directed acyclical graph. Letting \tilde{Z} denote a $p \times n$ matrix of mean-adjusted probit scores, the proposed transition is accepted with probability

$$R(\mathbf{B}, \mathbf{B}') = \min \left\{ 1; \frac{|\mathbf{B}'|^n}{|\mathbf{B}|^n} \text{etr} \left[-\frac{1}{2} \tilde{Z}^T (\mathbf{B}' - \mathbf{B})^T H_z (\mathbf{B}' - \mathbf{B}) \tilde{Z} \right] \right\}.$$

Some care is needed for the updates in (b), as they involve adding or deleting an edge in \mathcal{G} , therefore changing the dimension of the parameter space. We implement a reversible jumps MCMC (RJ) [Green (1995)]:

(i) Draw an edge $(k \rightarrow i)$ at random from E_0 . If in the current state \mathcal{G} , $(k \rightarrow i) \notin E$, propose the birth of the new edge $k \rightarrow i$. If $(k \rightarrow i) \in E$, propose the death of $k \rightarrow i$.

(ii) If we propose the birth $k \rightarrow i$, the structural matrix \mathbf{B} gets populated with a new element $\beta'_{ik} = u$, where $u \sim q(u)$. If we propose the death of edge $k \rightarrow i$, we simply set $\beta'_{ik} = 0$.

Steps (i) and (ii) generate a candidate $x' = (\mathbf{B}', \mathcal{G}')$. Let $m = \text{index the move proposed in step (i)}$, and let m' index the reverse move. The acceptance probability is [Green (1995)]

$$(6) \quad R(x, x') = \min \left\{ 1, \frac{\pi(dx')}{\pi(dx)} \frac{q(m'|x')}{q(m|x)q(u)} \right\},$$

where $q(m|x)$ is the probability of proposing move m when the chain is in state x , and $q(u)$ is the density function of u . In general, $R(x, x')$ might include an additional factor involving the Jacobian of a possible (deterministic) transformation of (x, u) to define x' . The described RJ involves no such transformation. The move m is generated in step (i) by a uniform draw from E_0 , implying $q(m|x) = q(m'|x')$. Finally, $q(u)$ is the proposal p.d.f. The acceptance probability of a birth R_b is then defined as

$$\begin{aligned} R_b &= \min \left\{ 1; \frac{p(x'|\mathbf{Y})}{p(x|\mathbf{Y})} q(u)^{-1} \right\} \\ &= \min \left\{ 1; \frac{|\mathbf{B}'|^n}{|\mathbf{B}|^n} \text{etr} \left[-\frac{1}{2} \tilde{Z}^T (\mathbf{B}' - \mathbf{B})^T H_z (\mathbf{B}' - \mathbf{B}) \tilde{Z} \right] \frac{\varphi}{(1 - \varphi)q(u)} \right\}. \end{aligned}$$

Even though nonsingular matrices define a dense open set in \mathbb{R}^p , if the proposed element β'_{ik} of \mathbf{B}' defines a numerically singular matrix, $\mathbf{\Omega}$ will not be positive definite and we reject move m' setting R_b to zero. Given this sampling scheme, the probability of a deletion is simply defined as $R_d = 1/R_b$, with the roles of x' and x as currently imputed and proposed state reversed.

3.2. Graphical model selection. The posterior probability $p(\mathcal{G}, \mathbf{B} | \mathbf{Y})$ and the corresponding MCMC posterior simulation characterize our knowledge about the pathway in the light of the data. Based on this posterior probability, we may be interested in selecting a representative graph \mathcal{G} . The posterior only summarizes the evidence for each \mathcal{G} . It does not yet tell us which \mathcal{G} s we should finally report.

This model selection problem has been discussed by different authors. Drton and Perlman (2007) discuss graphical model selection from the frequentist perspective, under the assumption that $n \geq p + 1$, while Jones et al. (2005) or Meinshausen and Bühlmann (2006) describe selection techniques for problems where the sample size n is small when compared to the number of variables p . From a Bayesian perspective, Carvalho and Scott (2009) provide a comprehensive discussion of Objective Bayesian model selection in Gaussian Graphical Models.

In the context of the model described in Section 2, graphical model selection can be defined by removing elements $(k \rightarrow i) \in E_0$ specified by the prior graph $\mathcal{G}_0 = \{V, E_0\}$. This is equivalent to the vanishing of the structural parameters β_{ik} in the matrix \mathbf{B} , characterizing the joint distribution of latent probit scores \mathbf{Z} [Ronning and Kukuk (1996)]. If the edge set E_0 has size $|E_0| = Q$, graphical model selection involves testing Q hypothesis

$$H_q^0 : \beta_{(q)} = 0 \quad \text{vs.} \quad H_q^1 : \beta_{(q)} \neq 0 \quad \text{for } q = 1, \dots, Q.$$

When testing a large number of hypotheses it is important to address possible multiplicity problems by controlling some predefined error rate. A popular choice is to control the False Discovery Rate (FDR) [Benjamini and Hochberg (1995)]. Several authors [Carvalho and Scott (2009), Scott and Carvalho (2008), Scott and Berger (2006)] consider the shrinkage prior defined in Section 3.1 and report how including edges with inclusion probability $P(\beta_{ik} \neq 0) > 0.5$ (median model) yields strong control over the number of false positives.

4. Simulation study. We validate and illustrate the proposed method with a simulation study with $p = 50$ genes from $n = 30$ samples. We define \mathbf{Y} as the $(p \times n)$ matrix of simulated mRNA intensities and consider a balanced design where 15 columns of \mathbf{Y} are from “normal” samples and 15 columns of \mathbf{Y} are associated with “tumor” samples. Thus, $\mathbf{x}_{ij} = (1, 0)'$ if y_{ij} is a normal sample and $\mathbf{x}_{ij} = (1, 1)'$ if y_{ij} is a tumor sample.

We generate simulated data \mathbf{Y} as follows. Given a set of latent scores $\mathbf{W} \sim \mathcal{MN}(\mathbf{0}, \boldsymbol{\Sigma}_z, \mathbf{I}_T)$, where $\boldsymbol{\Omega}_z = \boldsymbol{\Sigma}_z^{-1}$ encodes a known conditional dependence structure, and covariate effects $\mathbf{b}_i \sim N_2(\mathbf{m}_i, \sigma_b^2 \mathbf{I}_2)$, we define $z_{ij} = w_{ij} + \mathbf{x}'_{ij} \mathbf{b}_i$. We then generate the intensity matrix \mathbf{Y} from a three-way mixture of Gaussian distributions:

$$(7) \quad \begin{aligned} y_{ij} | z_{ij} \leq -1 &\sim N(-4, 2^2), \\ y_{ij} | z_{ij} > 3 &\sim N(4, 2^2), \\ y_{ij} | -1 < z_{ij} \leq 3 &\sim N(0, 1). \end{aligned}$$

The precision matrix $\boldsymbol{\Omega}_z$ is defined as follows. First we obtain the $p \times p$ matrix \mathbf{B} by defining $\gamma_{ij} \stackrel{d}{=} Ga(2, 1)$, $c_{ij} = \{-1, 1\}$ with $P(c_{ij} = 1) = 0.5$ and δ_0 a Dirac mass at 0, so that $\mathbf{B}_{ii} = 1$ (for $i = 1, \dots, p$), and the off-diagonal elements $\mathbf{B}_{ij} = \pi_0 \delta_0 + (1 - \pi_0) c_{ij} \gamma_{ij}$. The simulation truth is deliberately chosen differently from the assumed analysis model (1).

We then generate $\boldsymbol{\Omega}_z$ by rescaling $\mathbf{B}'\mathbf{B}$ to a correlation matrix. The simulation model (7) is deliberately different from the assumed analysis model, but still includes a meaningful notion of true dependence structure and strength.

We use a prior Graph $\mathcal{G}_0 = \{V, E_0\}$ spanned by the set of edges $E = E^* \cup \tilde{E}$, with E^* spanning the simulation truth of nonzero elements in \mathbf{B}_{ij} (in our example $|E^*| = 50$) and \tilde{E} serving as a random mispecification set including false edges (in our example $|\tilde{E}| = 87$).

In Figure 3 we display the classification results for the expression measurements generated under the dependence schemes just described. We calculated posterior probabilities of over- and under-expression from 50,000 posterior samples (thinned by 10), obtained after conservatively discarding 50,000 iterations. Our C++ implementation of the algorithm, described in Section 3.1, performed this simulation in about 6 hours on a standard desktop (2.94 GHz processor).

Figure 3 (*left panel*) shows the simulation truths as indicators (e_{ij}) of over- (white), normal- (grey) and under-expression (black). The right panel reports a unidimensional summary of the probabilities of over- or under-expression ($p_{ij}^* = \pi_{ij}^+ - \pi_{ij}^-$). The elements p_{ij}^* are defined in the $[-1, 1]$ scale and may be compared directly with the three-way indicators e_{gt} . We note that the p^* scale provides improved resolution over genes with signal and recovers well the generating truth.

Posterior inference includes a posterior distribution on the dependence structure. In Figure 4 (*left panel*) we report the number of edges included in the model by MCMC iteration, for two chains starting at opposite sides of the model saturation spectrum. Despite the size of the mispecification set \tilde{E} , the trans-dimensional Markov chains converge fairly rapidly toward models

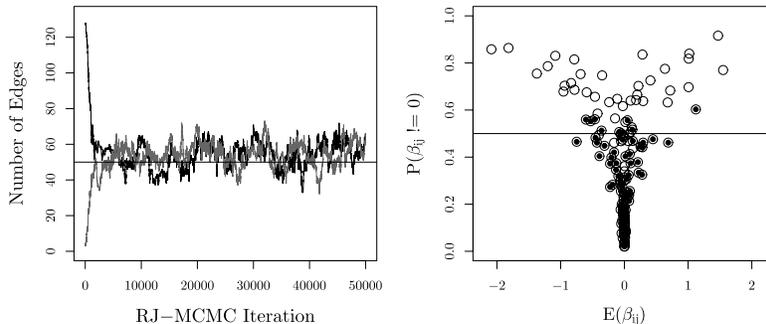


FIG. 4. *Simulation study: (Left panel) Number of edges included in the model by MCMC iteration, for two chains with starting points at the two extremes of the saturation spectrum. (Right panel) Posterior expected SEM coefficients $E(\beta_{ik}|\mathbf{Y})$ vs. posterior inclusion probabilities $P(\beta_{ik} \neq 0)$. False edges are represented with a solid circle.*

probability scheme not only penalizes for model complexity, but effectively controls the number of false discoveries, allowing for a genuine recovery of the generating conditional dependence structure.

In our simulation experiments we found that selecting edges with posterior inclusion probabilities greater than 0.5 tends to control the false discovery rate at level 0.01. We compared our model to the (independence) PoE model of Parmigiani et al. (2002) and found that including network inference as a new inferential goal does not diminish the classification accuracy of under- and over-expressed samples. For details see the Web-based supplement [Telesca et al. (2011)]. Furthermore, comparison with standard global search algorithms based on dynamic shrinkage of partial correlation estimates point to substantial inferential gains associated with the proposed methodology (see Web-based supplementary material, Section 3).

5. Case study. Wang, Wang and Kavanagh (2005) report a study of epithelial ovarian cancer (EOC). The goal of the study is to characterize the role of the tumor microenvironment in favoring the intra-peritoneal spread of EOC. To this end, the investigators collected tissue samples from patients with benign (b) and malignant (m) ovarian pathology. Specimens were collected, among other sites, from peritoneum adjacent to the primary tumor. RNA was co-hybridized with reference RNA to a custom made cDNA microarray including combination of the Research genetics RG_HsKG_031901 8k clone set and 9000 clones selected from RG_Hs_seq_ver_070700. A complete list of genes is available at http://nciarray.nci.nih.gov/gal_files/index.shtml, “custom printings.” See the array labeled Hs-CCDTM-17.5k-1px.

In the following discussion we focus on the comparison of 10 peritoneal samples from patients with benign ovarian pathology (bPT) versus 14 sam-

ples from patients with malignant ovarian pathology (mPT). The raw data was processed using BRB ArrayTool (<http://linus.nci.nih.gov/BRB-ArrayTools.html>). In particular, spots with minimum intensity less than 300 in both fluorescence channels were excluded from further analysis. See Wang, Wang and Kavanagh (2005) for a detailed description.

One subset of genes reported on the NIH custom microarray are 61 genes in the coagulation and complement pathway from KEGG (<http://www.genome.ad.jp>), shown in Figure 2. Genes on this pathway are of interest for their role in the inflammatory process. The arches in the pathway are interpreted as prior judgement about (approximate) conditional dependence (Section 2.1). However, recognizing that the pathway represents a protein system rather than gene expression, we allow for significant deviation from this structure, explicitly including model determination in our analysis.

We fit the model presented in Section 2 to this set of 61 genes. The prior set of conditional dependences between genes is represented as a reciprocal graph in Figure 2 and includes a set of 148 possible edges. Reported inference is based on 50,000 MCMC samples, thinned by 10, after discarding 50,000 observation for burn-in.

Recording the number of times the sampler visits a particular edge, we calculate the posterior probability $v_{ik} = P(\beta_{ik} | \mathbf{Y})$, for each edge ($k \rightarrow i$) in the prior graph \mathcal{G}_0 . In Figure 5 (Panel b) we show the set of selected genetic interactions when we consider edges with inclusion probabilities greater than 0.5 (median model). Edge directionality is inherited from \mathcal{G}_0 (Figure 2).

The posterior distribution on \mathbf{e}_g provides inference on differential expression, appropriately adjusted for dependence. Starting from the Complement and Coagulation Cascade pathway, we identify a set of 24 genes exhibiting patterns of dependence in their differential expression profiles across healthy and tumor tissues.

To interpret our findings, we searched the scientific literature using the Information Hyperlinked Over Protein (IHOP) tool implemented by Hoffman [Hoffman and Valencia (2004)], available at: www.ihop-net.org. For example, our study confirms the centrality of the peptide IL8 (Interleukin-8) in the regulation of the chemokine (CXC and CC motifs) genes. The protein encoded by this gene has been reported by several authors to play an important role in the response to inflammatory stimuli, resistance to apoptosis and tumoral angiogenesis. See Terranova and Rice (1997) or Brat, Bellail and Erwin (2005) for comprehensive discussions on IL8 and its receptors. One other example is the finding of dependent expression profiles associated with the Thrombin pathway ($F2 \rightarrow F2R$ and $F2 \rightarrow THBD$). This pathway plays a central role in the coagulation cascade and has been reported as a potential mediator of cellular function in the ovarian follicle [Roach et al. (2002)].

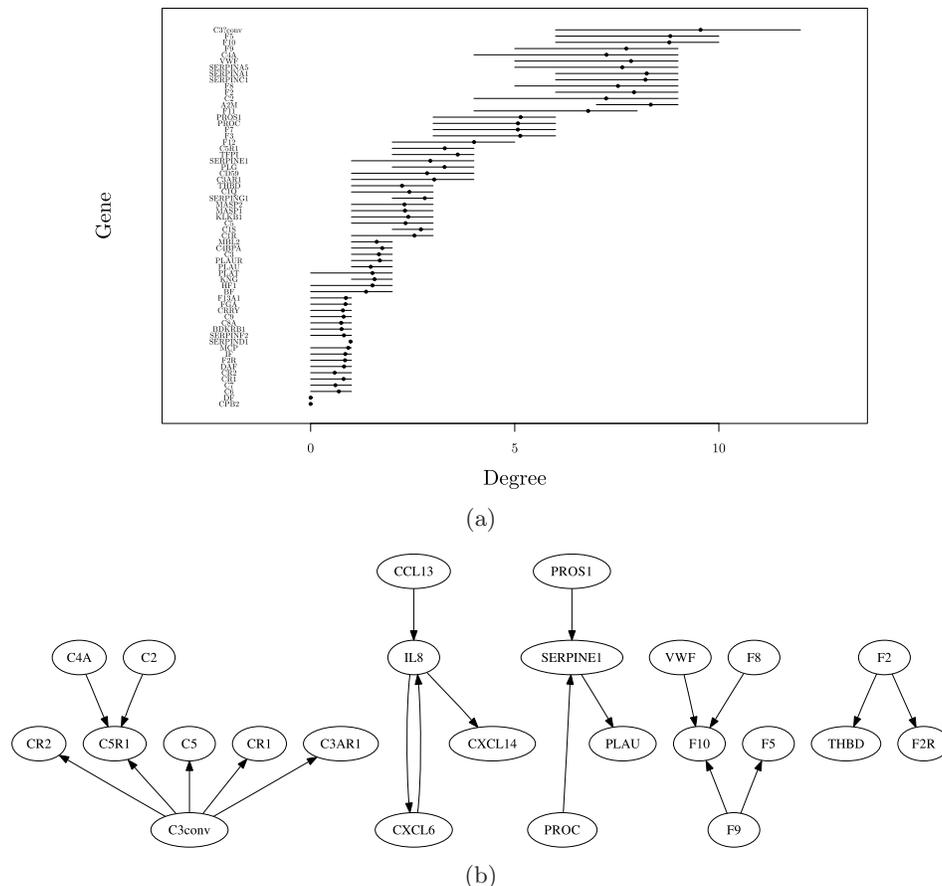


FIG. 5. Case study. Panel (a): Posterior mean degree and associated 95% credible intervals by gene. Panel (b): Posterior pathway obtained selecting edges with inclusion probabilities greater than 0.5 (Median model).

Posterior edge inclusion probabilities allow for the calculation of networks’ summaries at the gene level, which summarize the role played by individual genes in the prior pathway. In Figure 5 (Panel a), we report the posterior distribution of the degree of the node associated with each gene. This quantity is simply defined as $d_i = p(|ne(i)||\mathbf{Y})$, the posterior distribution of the number of neighbors associated with each gene. This measure is often used in social science as a way to summarize an individual’s centrality in a relational network [Sabidussi (1966)]. From a molecular biology perspective, genes with a high degree may be interpreted as playing active roles in the regulation of the pathway under study, in association with the biological process of interest. Our analysis of the degree distribution in Figure 5 (Panel a) confirms what we observe in the selected posterior set of genetic

interactions (Panel b) and identifies important active components of the complement and coagulation pathway. For example, we confirm the central role of C3-convertase in the promotion and progression of malignant ovarian cancer in humans, often reported as a key activation component in mouse studies [Markiewski and Lambris (2009)].

6. Discussion. We propose a probability model for the analysis of dependent gene expression data. Dependence between genes is modeled via the explicit consideration of prior information from pathways representing known biochemical processes. We characterize a biochemical pathway as a reciprocal graph depicting a coherent set of conditional dependence relationships between three-way classes of gene under-, normal- and over-expression. Modeling dependence between latent indicators of class membership is likely to represent a more sensible approach for this kind of data, when compared with methods that model correlations between observables directly. Acknowledging that a known pathway represents only prior information, we seek posterior inference for the model parameters as well as for the pathway itself via an RJ-MCMC scheme. We showed, through simulation studies, that our model enables the recovery of the true dependence structure, even under a misspecified prior pathway.

Our model of mRNA abundance relies on the Probability of Expression (POE) Model of Parmigiani et al. (2002), and assumes that the variability of expression across tissue samples can be fully characterized by heavy tailed mixtures of Normal and Uniform random variables. While this is a simplification of reality, it contributes to denoising data and is likely to provide useful summaries, allowing for the investigation of the many aspects associated with expression data analysis, from data normalization, to DE analysis, to the characterization of molecular profiles. The general framework presented in this article is also adaptable to other models of gene expression analysis.

In the construction of the dependent probability model, it is important to acknowledge the limitations of the information provided in a biochemical network. In fact, a pathway may not necessarily describe relations among transcript levels, although it carries some information about it. The proposed methodology is currently restricted to known biochemical pathways. Nevertheless, structural restrictions to one or more pathways of interest substantially simplifies computational tractability. The proposed model complexity is, in fact, only linear in the number of genes and interactions included in the prior graph. In our simulation example this provided substantially higher power in the detection of meaningful interactions, when compared to standard global search strategies. Computational scalability of the proposed methodology could, however, be an issue, when considering highly saturated pathways including a large number of genes. In these cases, methods based

on simplifying assumptions and approximate inference may indeed prove more feasible as exploratory analytical tools [Dobra et al. (2004)].

Our model could be extended to discover novel genetic interactions, by allowing the adding of new edges between nodes in the prior graph \mathcal{G}_0 . This extension would, however, come at a substantial computational cost and would require a challenging reformulation of the prior over graphs $p(\mathcal{G})$, to penalize for model complexity and, at the same time, to favor models closer to the structure of the prior pathway \mathcal{G}_0 . Initial progress in this direction was reported by Braun, Cope and Parmigiani (2008) and, in the context of Bayesian Networks, by Mukherjee and Speed (2008) and it is the subject of active research.

In this article we model dependence between three-way variables as dependence between latent Gaussian quantities. This probability scheme is only a convenient restriction on the possible shapes of dependence characterizing a matrix of ordinal random variables. Extensions of our model considering a richer class of dependence structures are, in principle, appealing. However, these changes would require a higher level of complexity and possible ad hoc limitations on the clique size contributing to the joint distribution of the three-way indicators.

APPENDIX: FULL CONDITIONAL DISTRIBUTIONS

Sample-specific means α_t . From Section 2.2 we have that $p(\alpha_j|\tau_\alpha^2) \propto \exp\{-\alpha_j^2/(2\tau_\alpha^2)\}$. Using standard conjugate analysis, it is easy to show that $\alpha_j|\mathbf{Y}_j, \boldsymbol{\theta}_{\setminus\alpha_j} \sim N(\alpha_j^*, v_{\alpha_j}^*)I(l_j^* < \alpha_j < u_j^*)$, where $v_{\alpha_j}^* = \{\tau_\alpha^{-2} + \sum_i \sigma_g^{-2} I(e_{ij} = 0)\}^{-1}$, $\alpha_j^* = v_{\alpha_j}^* \sum_i \{(y_{ij} - \mu_i)/\sigma_i^2 I(e_{ij} = 0)\}$, $l_j^* = \max_{\{g:e_{ij}=1\}}(y_{ij} - \mu_i - k_i^+)$ and $u_j^* = \min_{\{i:e_{ij}=-1\}}(y_{ij} - \mu_i + k_i^-)$.

Gene-specific means μ_i . From Section 2.2 we have that $p(\mu_i|m_\mu, \tau_\mu^2) \propto \exp\{-(\mu_i^2 - 2m_\mu\mu_i)/(2\tau_\mu^2)\}$. Using standard conjugate analysis, it is easy to show that $\mu_i|\mathbf{Y}_i, \boldsymbol{\theta}_{\setminus\mu_i} \sim N(\mu_i^*, v_i^*)I(l_i^* < \mu_i < u_i^*)$, where $v_i^* = \{\tau_\mu^{-2} + \sigma_i^{-2} \sum_j I(e_{ij} = 0)\}^{-1}$, $\mu_i^* = v_i^* \{m_\mu/\tau_\mu^2 + \sigma_i^{-2} \sum_j (y_{ij} - \alpha_j) I(e_{ij}=0)\}$, $l_i^* = \max_{\{j:e_{ij}=1\}}(y_{ij} - \alpha_j - k_i^+)$ and $u_i^* = \min_{\{j:e_{ij}=-1\}}(y_{ij} - \alpha_j + k_i^-)$.

Gene-specific variances σ_i^2 . We introduced a conditionally conjugate Inverse Gamma prior for σ_i^2 in Section 2.2. For ease of notation we define $h_i = 1/\sigma_i^2$, and $n_{0i} = \sum_j I(e_{ij} = 0)$. It is easy to verify that $h_i|\mathbf{Y}_i, \boldsymbol{\theta}_{\setminus h_i} \sim \text{Gamma}(a_i^*, b_i^*)I\{h_i \geq (\kappa_0/\min(k_i^-, k_i^+))^2\}$, where $a_i^* = \gamma_\sigma + n_{0i}/2$ and $b_i^* = \lambda_\sigma + \sum_j I(e_{ij} = 0)(y_{ij} - \mu_i - \alpha_j)^2/2$.

Uniform bounds k_i^- and k_i^+ . For ease of notation we define $\nu_{i0} = 1/k_i^-$ and $\nu_{i1} = 1/k_i^+$. We have $p(\nu_{i\ell}) \propto \nu_{i\ell}^{\gamma_k - 1} e^{-\lambda_k \nu_{i\ell}}$, ($\ell = 0, 1$). The conditional posterior distribution of these parameters is defined as $\nu_{i\ell}|\mathbf{Y}_i, \boldsymbol{\theta}_{\setminus\nu_{i\ell}} \sim \text{Gamma}(a_{i\ell}^*, b_{i\ell}^*) \times I(S_{i\ell})$, where $a_{i\ell}^* = \gamma_k + \sum_j I(e_{ij} = 2\ell - 1)$, $b_{i\ell}^* = \lambda_k$ and $S_{i\ell} = \{\nu_{i\ell} : \nu_{i\ell} \leq \min[\min_{\{j:e_{ij}=2\ell-1\}}(y_{ij} - \mu_i - \alpha_j), (\kappa_0 \sigma_i)^{-1}]\}$.

Probit score precisions $1/s_i^2$. For ease of notation we define $h_{si} = 1/s_i^2$. Taking advantage of conditional conjugacy with the distribution of probit scores, we define $p(h_{si}|a_s, b_s) \propto h_{si}^{a_s-1} \exp\{-b_s h_{si}\}$. Let $\tilde{z}_{ij} = z_{ij} - m_{ij} - \sum_{k \neq i} \beta_{ik}(z_{kj} - m_{kj})'$. It is easy to show that the conditional posterior density of h_{si} is then Gamma with $p(h_{si}|\mathbf{Y}_i, \boldsymbol{\theta}_{\setminus h_{si}}) \propto h_{si}^{n/2+a_s-1} \exp\{-h_{si} \sum_j (\tilde{z}_{ij}^2)/2\}$.

SUPPLEMENTARY MATERIAL

Convergence diagnostics and model comparisons

(DOI: [10.1214/11-AOAS525SUPP](https://doi.org/10.1214/11-AOAS525SUPP); .pdf). We provide an extended discussion of some aspects associated with the proposed model. In particular, we compare our results to the PoE model of Parmigiani et al. (2002) as well as some current methods used to infer networks.

REFERENCES

- ALBERT, J. H. and CHIB, S. (1993). Bayesian analysis of binary and polychotomous response data. *J. Amer. Statist. Assoc.* **88** 669–679. [MR1224394](#)
- BEAL, M., FALCIANI, F., GHAHRAMANI, Z., RANGEL, C. and WILD, D. (2005). A Bayesian approach to reconstructing genetic regulatory networks with hidden factors. *Bioinformatics* **21** 349–356.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *J. Roy. Statist. Soc. Ser. B* **36** 192–236. [MR0373208](#)
- BRAT, D. J., BELLAIL, A. C. and ERWIN, G. V. M. (2005). The role of interleukin-8 and its receptors in gliomagenesis and tumoral angiogenesis. *Neuro-Oncology* **7** 122–133.
- BRAUN, R., COPE, L. and PARMIGIANI, G. (2008). Identifying differential correlation in gene/pathway combinations. *BMC Bioinformatics* **9** 488.
- BROËT, P. and RICHARDSON, S. (2006). Detection of gene copy number changes in CGH microarrays using a spatially correlated mixture model. *Bioinformatics* **22** 911–918.
- BROWN, P. J., VANNUCCI, M. and FEARN, T. (1998). Multivariate Bayesian variable selection and prediction. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **60** 627–641. [MR1626005](#)
- CARVALHO, C. M. and SCOTT, J. G. (2009). Objective Bayesian model selection in Gaussian graphical models. *Biometrika* **96** 497–512. [MR2538753](#)
- DAWID, A. P. and LAURITZEN, S. L. (1993). Hyper-Markov laws in the statistical analysis of decomposable graphical models. *Ann. Statist.* **21** 1272–1317. [MR1241267](#)
- DOBRA, A., HANS, C., JONES, B., NEVINS, J. R., YAO, G. and WEST, M. (2004). Sparse graphical models for exploring gene expression data. *J. Multivariate Anal.* **90** 196–212. [MR2064941](#)
- DRTON, M. and PERLMAN, M. D. (2007). Multiple testing and error control in Gaussian graphical model selection. *Statist. Sci.* **22** 430–449. [MR2416818](#)
- FRIEDMAN, N., LINIAL, M., NACHMAN, I. and PE'ER, D. (2000). Using Bayesian networks to analyze expression data. *J. Comput. Biol.* **7** 601–620.
- FRÖHLICH, H., SPEER, N., POUTSKA, A. and BEIBART, T. (2007). GOSim—an R-package for computation of theoretic GO similarities between terms and ene products. *BMC Bioinformatics* **8** 166.

- GARRETT, E. S. and PARMIGIANI, G. (2004). A nested unsupervised approach to identifying novel molecular subtypes. *Bernoulli* **10** 951–969. [MR2108038](#)
- GEORGE, E. I. and McCULLOCH, R. E. (1993). Variable selection via Gibbs sampling. *J. Amer. Statist. Assoc.* **88** 881–889.
- GIUDICI, P. and GREEN, P. J. (1999). Decomposable graphical Gaussian model determination. *Biometrika* **86** 785–801. [MR1741977](#)
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82** 711–732. [MR1380810](#)
- HOFFMAN, R. and VALENCIA, A. (2004). A gene network for navigating the literature. *Nature Genetics* **36** 664–664.
- JONES, B., CARVALHO, C., DOBRA, A., HANS, C., CARTER, C. and WEST, M. (2005). Experiments in stochastic computation for high-dimensional graphical models. *Statist. Sci.* **20** 388–400. [MR2210226](#)
- KOLACZYK, E. D. (2009). *Statistical Analysis of Network Data: Methods and Models*. Springer, New York. [MR2724362](#)
- KOSTER, J. T. A. (1996). Markov properties of nonrecursive causal models. *Ann. Statist.* **24** 2148–2177. [MR1421166](#)
- LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series* **17**. Clarendon Press, New York. [MR1419991](#)
- MARKIEWSKI, M. M. and LAMBRIS, J. D. (2009). Unwelcome complement. *Cancer Research* **69** 6367.
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363](#)
- MISTRY, M. and PAVLIDIS, P. (2008). Gene ontology term overlap as a measure of gene functional similarity. *BMC Bioinformatics* **9** 327.
- MUKHERJEE, S. and SPEED, T. P. (2008). Network inference using informative priors. *PNAS* **105** 14133–14318.
- MURPHY, K. and MIAN, S. (1999). Modeling gene expression data using dynamic Bayesian networks. Technical report, Computer Science Division, Univ. California, Berkeley.
- ONG, I., GLASNER, J. and PAGE, D. (2002). Modelling regulatory pathways in e.coli from time series expression profiles. *Bioinformatics* **18** S241–S248.
- PARMIGIANI, G., GARRETT, E. S., ANBAZHAGAN, R. and GABRIELSON, E. (2002). A statistical framework for expression-based molecular classification in cancer. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 717–736. [MR1979385](#)
- ROACH, L. E., PETRIK, J. J., PLANTE, L. and LAMARRE, J. (2002). Thrombin generation and presence of thrombin in ovarian follicles. *Biology of Reproduction* **66** 1350–1358.
- RONNING, G. and KUKUK, M. (1996). Efficient estimation of ordered probit models. *J. Amer. Statist. Assoc.* **91** 1120–1129. [MR1424612](#)
- SABIDUSSI, G. (1966). The centrality index of a graph. *Psychometrika* **31** 581–603. [MR0205879](#)
- SCOTT, J. G. and BERGER, J. O. (2006). An exploration of aspects of Bayesian multiple testing. *J. Statist. Plann. Inference* **136** 2144–2162. [MR2235051](#)
- SCOTT, J. G. and CARVALHO, C. M. (2008). Feature-inclusion stochastic search for Gaussian graphical models. *J. Comput. Graph. Statist.* **17** 790–808. [MR2649067](#)
- SEBASTIANI, P. and RAMONI, M. (2005). Normative selection of Bayesian networks. *J. Multivariate Anal.* **93** 340–357. [MR2162642](#)
- SPIRTEs, P., RICHARDSON, T. S., MEEK, C., SCHEINES, R. and GLYMOUR, C. (1998). Using path diagrams as a structural equation modeling tool. *Sociol. Methods Res.* **27** 182–225.

- TELESCA, D., MÜLLER, P., PARMIGIANI, G. and FREEDMAN, R. S. (2011). Supplement to “Modeling dependent gene expression.” DOI:[10.1214/11-AOAS525SUPP](https://doi.org/10.1214/11-AOAS525SUPP).
- TERRANOVA, P. F. and RICE, V. M. (1997). Review: Cytokine involvement in ovarian processes. *American Journal of Reproductive Immunology* **37** 50–63.
- WANG, X., WANG, E. and KAVANAGH, J. (2005). Ovarian cancer, the coagulation pathway, and inflammation. *Journal of Translational Medicine* **3** 25.
- WEI, Z. and LI, H. (2007). A Markov random field model for network-based analysis of genomic data. *Bioinformatics* **23** 1357–1544.
- 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](https://doi.org/10.1214/08-BA156)

D. TELESCA
DEPARTMENT OF BIostatISTICS
UCLA SCHOOL OF PUBLIC HEALTH
LOS ANGELES, CALIFORNIA 90095–1772
USA
E-MAIL: donatello.telesca@gmail.com

G. PARMIGIANI
DEPARTMENT OF BIostatISTICS
HARVARD SCHOOL OF PUBLIC HEALTH
BOSTON, MASSACHUSETTS 02115
USA

P. MÜLLER
UNIVERSITY OF TEXAS AUSTIN
DEPARTMENT OF MATHEMATICS
AUSTIN, TEXAS 78712
USA

R. S. FREEDMAN
UNIVERSITY OF TEXAS
M.D. ANDERSON CANCER CENTER
DEPARTMENT OF GYNECOLOGIC ONCOLOGY
HOUSTON, TEXAS 7030
USA