

INFERRING GENE–GENE INTERACTIONS AND FUNCTIONAL MODULES USING SPARSE CANONICAL CORRELATION ANALYSIS¹

BY Y. X. RACHEL WANG, KENI JIANG, LEWIS J. FELDMAN,
PETER J. BICKEL AND HAIYAN HUANG

University of California, Berkeley

Networks pervade many disciplines of science for analyzing complex systems with interacting components. In particular, this concept is commonly used to model interactions between genes and identify closely associated genes forming functional modules. In this paper, we focus on gene group interactions and infer these interactions using appropriate partial correlations between genes, that is, the conditional dependencies between genes after removing the influences of a set of other functionally related genes. We introduce a new method for estimating group interactions using sparse canonical correlation analysis (SCCA) coupled with repeated random partition and subsampling of the gene expression data set. By considering different subsets of genes and ways of grouping them, our interaction measure can be viewed as an aggregated estimate of partial correlations of different orders. Our approach is unique in evaluating conditional dependencies when the correct dependent sets are unknown or only partially known. As a result, a gene network can be constructed using the interaction measures as edge weights and gene functional groups can be inferred as tightly connected communities from the network. Comparisons with several popular approaches using simulated and real data show our procedure improves both the statistical significance and biological interpretability of the results. In addition to achieving considerably lower false positive rates, our procedure shows better performance in detecting important biological pathways.

1. Introduction. Many complex systems in science and nature are composed of interacting parts. Such parts can be modeled as nodes and their relationships as edges in a network. Network modeling has found numerous applications [Newman (2010)]. Gene association networks is one such example, with genes modeled as nodes and their interactions as edges. One important application of gene networks is the identification of communities corresponding to genes with related functional groupings. Many of these functional groups encode biological pathways. A major task in understanding biological processes is to identify these pathway genes and

Received September 2014.

¹Supported in part by Grants NIH EY019094, NIH U01 HG007031, NSF DMS-06-36667 and NSF DMS-11-60319.

Key words and phrases. Gene association networks, community structure, sparse canonical correlation analysis (SCCA), partial correlation.

elucidate the relationships between them. We focus in this paper on modeling gene interactions. As a result, a gene network can be constructed using the interaction measures as edge weights and gene functional groups can be inferred as tightly connected communities in the network.

In gene networks, direct observation of gene relationships by experimental approaches is extremely cost-prohibitive given that the typical size of the networks is in the tens of thousands. The gene expression levels, on the other hand, are easier to measure and can be regarded as sets of covariates associated with the nodes. Constructing gene networks using expression data has remained a challenging unsupervised learning problem in the statistics literature due to the complexity of data structure and the difficulty of finding an appropriate measure for characterizing gene relationships. A review of existing methods can be found in [Wang and Huang \(2014\)](#).

Most methods for inferring edges in gene networks are based on the notion of measuring expression profile similarity or co-expression, which aims to estimate marginal relationships between pairs of genes. Widely used co-expression measures include the Euclidean distance or the angle between vectors of observed expression levels or, most commonly, the marginal covariance or correlation. Measures detecting general statistical dependence such as mutual information (MI) are also explored. MI offers the advantage of being able to detect nonlinear correlations [[Daub et al. \(2004\)](#)], but some empirical studies [[Steuer et al. \(2002\)](#)] also show it yields almost identical results as the Pearson correlation. Recently, a new measure named the maximal information coefficient (MIC) was proposed by [Reshef et al. \(2011\)](#) based on normalized estimates of MI. [Kinney and Atwal \(2014\)](#) offer some criticisms and discussions of MIC.

The above measures for estimating marginal dependencies only consider pairwise relationships. However, in a real biological pathway, a gene can interact with a group of genes but their marginal relationships may remain weak. Such higher-level interactions (i.e., gene group interactions) are better modeled by Gaussian graphical models (GGM) due to its interpretation in terms of conditional correlations. Under the assumption of multivariate normality of gene expression vectors, the GGM uses the inverse of the gene covariance matrix (or precision matrix) as a measure for gene associations. This approach is closely related to the concept of partial correlations: the (i, j) th element in the precision matrix is proportional to the partial correlation between gene i and j conditional on the rest of the genes. To address the “curse of dimensionality” (the number of genes being much larger than the number of samples) in estimating the precision matrix, one can exploit the belief that gene networks are inherently sparse and reframe the problem of estimating partial correlations in a penalized regression setting [[Meinshausen and Bühlmann \(2006\)](#), [Peng, Zhou and Zhu \(2009\)](#)]. More studies on estimating the sparse precision matrix in high-dimensional GGMs can be found in, for example, [Schäfer and Strimmer \(2005\)](#), [Friedman, Hastie and Tibshirani \(2008\)](#) and [Zhou et al. \(2011\)](#).

Despite their attractive theoretical properties, these partial correlation-based methods still have limitations in their estimation methods. In the current literature, partial correlation is usually calculated conditioned on either all of the available genes or a more or less arbitrary subset of them that may contain noisy (biologically unrelated) genes. [de la Fuente et al. \(2004\)](#) reported that conditioning on all genes simultaneously can introduce spurious dependencies which are not from a direct causal or common ancestors effect. To alleviate this concern, there are alternative approaches using lower order partial correlations [[de la Fuente et al. \(2004\)](#), [Li \(2002\)](#), [Magwene and Kim \(2004\)](#), [Wille and Bühlmann \(2006\)](#), [Wille et al. \(2004\)](#)] which condition on one or two other genes. However, these methods come at a cost of lowering the sensitivity for inferring higher level gene associations and do not necessarily eliminate the effect of noisy genes. [Kim et al. \(2012\)](#) proposed to minimize the impact of noisy genes by conditioning on a small set (3–5 genes) of “seed genes” (i.e., known pathway genes). However, such prior biological information is not always available, especially in exploratory studies.

In this paper we tackle the problem of estimating gene relationships when the correct conditional set for partial correlation is unknown. We introduce a new method of inferring the strength of gene group interactions using sparse canonical correlation analysis (SCCA) with repeated random partition and subsampling of the gene expression data set. There has been a growing interest in applying SCCA to genomic data sets [[Lee et al. \(2011\)](#), [Parkhomenko, Tritchler and Beyene \(2009\)](#), [Waaaijenborg, Verselewele de Witt Hamer and Zwinderman \(2008\)](#), [Witten and Tibshirani \(2009\)](#)] in the context of studying relationships between two or more sets of variables, such as gene expression levels, copy numbers and other phenotype variations, with measurements taken from the same sample. One novelty of our method lies in the application of SCCA to a single data set facilitated by a random partition scheme. By randomly separating the genes into two groups, SCCA searches for a strong linear relationship between a small set of genes, for example, 5–20 genes, from both groups of genes (e.g., 500–2000 genes in total). Through multiple rounds of random partition, this SCCA approach, reframed in a linear regression setting, gives estimates proportional to partial correlations conditioned on different sets of signal genes (with noisy genes eliminated through sparsity). The subsampling procedure analyzes different subsets of the genes at a time and enables simultaneous identification of multiple interacting groups with different signal strengths. Using this construction, we build an edge weight matrix for the whole gene network whose interaction measure reflects an aggregated estimate of partial correlations of different orders. Our approach is flexible and can be adapted to work with or without prior biological knowledge.

The rest of the paper is organized as follows. In [Section 2](#) we discuss in detail the motivations behind our new scheme of computing edge weights in a gene network by assessing gene group interactions and provide an outline of the full procedure. To identify densely connected communities as potential gene functional modules in the constructed network, we implemented two well-known methods in the network literature, the stochastic block model (SBM) and hierarchical

clustering (HC). In Section 3 comparisons are made between our procedure and correlation-based methods. We demonstrate that our procedure in general achieves a significant reduction in the rate of false positives. To test its performance in real data applications, our procedure is applied to an *Arabidopsis thaliana* microarray data set obtained under oxidation stress. Finally, in Section 4 we discuss the advantages and potential extensions of the present method.

2. Methods. As mentioned in Section 1, the conditional correlation interpretation of partial correlation suggests it is a more appropriate framework for modeling higher level interactions in gene networks, provided the conditional computation is carried out properly. In this section, we discuss some of the limitations of the partial correlation approach that arise due to its reliance on the correct selection of conditional sets of genes and how our SCCA-based approach circumvents this difficulty. We then give a detailed description of our new method of estimating an edge weight matrix using SCCA with subsampling.

2.1. *Method motivation.* Recall that when the gene expression levels follow a multivariate normal distribution, for a set of genes W , the partial correlation between genes i and j can be expressed as

$$(2.1) \quad \rho_{ij} = \text{cor}(i, j | W \setminus \{i, j\}) = \begin{cases} -\frac{\omega_{ij}}{\sqrt{\omega_{ii}\omega_{jj}}}, & i \neq j, \\ 1, & i = j, \end{cases}$$

where ω_{ij} are elements in the precision matrix $(\Sigma^G)^{-1}$ with Σ^G being the gene covariance matrix of the set W [see, e.g., Edwards (2000)]. Genes i and j being conditionally independent is equivalent to the corresponding partial correlation and element in the precision matrix being zero.

As pointed out in de la Fuente et al. (2004) and Kim et al. (2012), the selection of a proper set of genes on which the correlation in (2.1) is conditioned determines the effectiveness of using partial correlation to measure gene interactions. The inclusion of noisy (biologically unrelated) genes in the set $W \setminus \{i, j\}$ may introduce spurious dependencies and, consequently, false edges in the estimated network. The use of partial correlation may also prove problematic when W contains multiple pathways. Here is a minimal example: suppose the set W has two pathways $\{x, y, z\}$ and $\{u, v\}$ and two independent noisy genes p and q , with expression relationships

$$(2.2) \quad z = x + y + \varepsilon_1 u + \varepsilon_2 v + \varepsilon_3 p, \quad u = \delta_1 x + \delta_2 y + \delta_3 z + \delta_4 q + v,$$

where ε_i and δ_j are small constants so that the dependencies between the two pathways are negligible, and gene v is independent of genes x and y . Computing the partial correlations, we have the desired dependencies:

$$\begin{aligned} \text{cor}(z, x | W \setminus \{z, x\}) &= \text{cor}(z, y | W \setminus \{z, y\}) = 1, \\ \text{cor}(u, v | W \setminus \{u, v\}) &= 1, \end{aligned}$$

but also some spurious ones:

$$\text{cor}(u, x|W \setminus \{u, x\}) = \text{cor}(u, y|W \setminus \{u, y\}) = \text{cor}(u, z|W \setminus \{u, z\}) = 1.$$

Using these partial correlations to construct an edge weight matrix would imply the two pathways are fully connected. The proper calculation should condition only on genes in the same pathway, but such information is usually hard to obtain in practice. Alternatively, a more appropriate edge weight measure can take into account the magnitude of the linear coefficients in (2.2) so that it reflects the amount of contribution each gene makes to a pathway and the two-block nature of the network. Recall that in a regression setting, the regression coefficients are multiplicative functions of the corresponding partial correlations. In this sense, the coefficients encompass more information and provide a better resolution on gene relationships than the partial correlations alone.

Motivated by these observations, we propose a new way to assess gene group interactions. In particular, we aim to identify strong linear relationships possessed by a small subset of the candidate genes. We make direct use of the linear coefficients found by SCCA when applied to two randomly partitioned gene groups. With repeated random partition on subsampled gene sets, an edge weight matrix built by the average SCCA coefficients over iterations reflects an aggregated level of direct or partial gene interactions. More discussion on how CCA coefficients relate to partial correlations can be found in Section 4 of the supplementary information [Wang et al. (2015)]. Sparsity is imposed to reduce dimensionality and, in particular in the example above, ensures the mixing of the two pathways is negligible on average.

2.2. Review of sparse canonical correlation analysis and its implementation.

Let $\mathbf{X} \in \mathbb{R}^{n \times q_1}$ be a matrix comprised of n observations on q_1 variables, and $\mathbf{Y} \in \mathbb{R}^{n \times q_2}$ a matrix comprised of n observations on q_2 variables. CCA introduced by Hotelling (1936) involves finding maximally correlated linear combinations between the two sets of variables. More explicitly, one finds $\boldsymbol{\alpha} \in \mathbb{R}^{q_2}$ and $\boldsymbol{\beta} \in \mathbb{R}^{q_1}$ that solve the optimization problem

$$(2.3) \quad \max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \boldsymbol{\alpha}^T \Sigma_{YX} \boldsymbol{\beta} \quad \text{subject to } \boldsymbol{\alpha}^T \Sigma_{YY} \boldsymbol{\alpha} = 1, \boldsymbol{\beta}^T \Sigma_{XX} \boldsymbol{\beta} = 1,$$

where $\Sigma_{(\cdot, \cdot)}$ represent the correlation matrices. Note that provided the variables in \mathbf{X} and \mathbf{Y} have nonzero variances, this is equivalent to the usual CCA formulation in terms of covariance matrices.

In practice, the population correlations are replaced with their sample counterparts. That is, $S_{YX} = \mathbf{Y}^T \mathbf{X} / (n - 1)$, $S_{XX} = \mathbf{X}^T \mathbf{X} / (n - 1)$ and $S_{YY} = \mathbf{Y}^T \mathbf{Y} / (n - 1)$, assuming the columns of \mathbf{X} and \mathbf{Y} have been centered and scaled. Let \mathbf{a} and \mathbf{b} be the weight vectors solving the optimization problem

$$(2.4) \quad \max_{\mathbf{a}, \mathbf{b}} \mathbf{a}^T S_{YX} \mathbf{b} \quad \text{subject to } \mathbf{a}^T S_{YY} \mathbf{a} = 1, \mathbf{b}^T S_{XX} \mathbf{b} = 1$$

for sample correlations.

For high throughput biological data, q_1 and q_2 are typically much larger than n . It is thus natural to impose sparsity on \mathbf{a} and \mathbf{b} , and this can be done by including (typically convex) penalty functions in (2.4). A number of studies [Parkhomenko, Tritchler and Beyene (2009), Waaijenborg, Verselewe de Witt Hamer and Zwinderman (2008), Witten, Tibshirani and Hastie (2009)] have proposed various methods for formulating the penalized optimization problem and obtaining sparse solutions. Here we adopt the diagonal penalized CCA criterion given by Witten, Tibshirani and Hastie (2009), which treats the covariance matrices in (2.4) as diagonal and relaxes the equality constraints for convexity:

$$(2.5) \quad \max_{\mathbf{a}, \mathbf{b}} \mathbf{a}^T \mathbf{Y}^T \mathbf{X} \mathbf{b} \quad \text{subject to } \mathbf{a}^T \mathbf{a} \leq 1, \mathbf{b}^T \mathbf{b} \leq 1, p_1(\mathbf{a}) \leq c_1, p_2(\mathbf{b}) \leq c_2,$$

where p_1 and p_2 are convex penalty functions. In this paper, we consider an L_1 penalty and solve the above optimization using the modified NIPALS algorithm proposed by Lee et al. (2011), which is reported to yield better empirical performance than Witten, Tibshirani and Hastie's (2009) algorithm. The modified NIPALS algorithm performs penalized regressions iteratively on \mathbf{X} and \mathbf{Y} with the penalty functions $p_{\lambda_1}(\cdot) = \lambda_1 \|\cdot\|_1$ and $p_{\lambda_2}(\cdot) = \lambda_2 \|\cdot\|_1$. This is an equivalent formulation to iteratively optimizing (2.5) using the bounded constraints.

It is important to note that one more complication arises when SCCA is applied to gene expression data. In CCA, the estimation of the correlation matrix using sample correlations requires the data matrices \mathbf{X} and \mathbf{Y} have independent rows. However, given a gene expression matrix with genes in columns and experiments in rows, it is often the case that row-wise and column-wise dependencies co-exist. Row-wise dependencies, or experiment dependencies, can be defined as the dependencies in gene expression between experiments due to the similar or related cellular states induced by the experiments [Teng and Huang (2009)]. When unaccounted for, they can introduce redundancies that overwhelm the important signals and lead to inaccurate estimates of the gene correlation matrix. To decouple the effect of experiment dependencies from the estimation of gene correlations, we apply the *Knorm* procedure from Teng and Huang (2009). The *Knorm* model assumes a multiplicative structure for the gene-experiment interactions, and iteratively estimates the gene covariance matrix and experiment covariance matrix through a weighted correlation formula. In addition, row subsampling and covariance shrinkage are used to ensure robust estimation.

2.3. Constructing an edge weight matrix by SCCA with repeated random partition and subsampling. Suppose an observed data set contains measurements of the expression levels of p genes in n experiments, where each experiment has a small number of replicates. We next describe our new procedure of computing edge weights that reflect gene group interactions in the gene network.

Summary of procedure:

Step (i): Data normalization by Knorm. A gene expression matrix \mathbf{Z}_b of dimension $n \times p$ can be generated from the full data set by sampling one replicate from each experiment. Using the *Knorm* model in [Teng and Huang \(2009\)](#), we normalize \mathbf{Z}_b as

$$(2.6) \quad \mathbf{Z}_b^* = (\hat{\Sigma}^E)^{-1/2} (\mathbf{Z}_b - \hat{\mathbf{M}}),$$

where $\hat{\mathbf{M}}$ is the estimated mean matrix and $\hat{\Sigma}^E$ is the estimated experiment correlation matrix.

Step (ii): Subsampling. For each normalized expression matrix \mathbf{Z}_b^* , sample (without replacement) a fixed fraction s , say, 70%, of the genes to obtain an $n \times sp$ submatrix $\mathbf{Z}_b^{\text{sub}}$.

Step (iii): SCCA with random partition on the subsampled matrix. For each partition t , randomly split the columns (genes) of $\mathbf{Z}_b^{\text{sub}}$ into two groups of equal size (more explanation given in the remarks below) to form $\mathbf{X}_{b,t}^{\text{sub}}$ and $\mathbf{Y}_{b,t}^{\text{sub}}$. Run SCCA on $\mathbf{X}_{b,t}^{\text{sub}}$ and $\mathbf{Y}_{b,t}^{\text{sub}}$: find sparse weight vectors $\mathbf{a}_{b,t}^{\text{sub}}$ and $\mathbf{b}_{b,t}^{\text{sub}}$ using the modified NIPALS algorithm [[Lee et al. \(2011\)](#)] with the L_1 penalty and tuning parameters $\lambda = (\lambda_1, \lambda_2)$, the choice of which will be discussed in Section 3.

Let $\mathbf{c}_{b,t}$ be the list of the absolute values $|\mathbf{a}_{b,t}^{\text{sub}}|$ and $|\mathbf{b}_{b,t}^{\text{sub}}|$ ordered according to the gene list. For the genes not included in the subsampled matrix, the corresponding values in $\mathbf{c}_{b,t}$ are set to 0. Average over all the partitions to obtain the average weights $\bar{\mathbf{c}}_b$. Define edge weight matrix $\mathbf{A}_b = \bar{\mathbf{c}}_b \bar{\mathbf{c}}_b^T$, setting $\text{diag}(\mathbf{A}_b) = 0$ to exclude self loops.

Step (iv): Repeat steps (ii) and (iii) B times. Define $\bar{\mathbf{A}} = 1/B \sum_{b=1}^B \mathbf{A}_b$ and normalize by the maximum value in $\bar{\mathbf{A}}$.

As will be demonstrated in Section 3.1, $\bar{\mathbf{A}}$ defined above exhibits a natural block structure when there is one or multiple functional groups. Here are more remarks on our procedure to construct $\bar{\mathbf{A}}$:

1. Step (i) can be skipped when dependencies between experimental conditions are weak and not of concern.

2. Step (ii) subsampling is necessary if we aim to identify multiple functional groups (that may overlap) simultaneously. As there will be multiple groups with strong interactions, not all of them can be detected unless different subsets of genes are considered. For more discussion about the subsampling step and the choice of subsampling levels, we refer to Section 5 in the supplementary information [[Wang et al. \(2015\)](#)].

3. During the random partition in step (iii), the two sets of genes do not have to be exactly equal in size, but they need to be comparable in order to maximize the chance of separating any gene functional group of interest into two sets.

4. Through multiple rounds of random partition, SCCA gives estimates in a regression setting proportional to partial correlations conditioned on different sets

of signal genes. Overall, subsampling and random partition enable us to consider different subsets of the genes and ways to group them. Thus, the elements in $\bar{\mathbf{A}}$ can be interpreted as an aggregated measure of partial correlations of different orders as the algorithm steps through different conditional sets of genes.

5. As we search through different subsets of genes, different signal groups are identified depending on the strengths of linear associations in the subset. As will be shown empirically in Section 3.1, the averaged result leads to the formation of a distinct block structure with different connectivities in the matrix.

Our procedure is flexible and can be modified easily to incorporate the following variants:

1. If prior knowledge is available on a pathway of interest, for example, it is known in advance that some genes are actively involved in that pathway, one may focus on the identification of the gene group related to this pathway first and incorporate the prior knowledge by lowering the penalties associated with those known pathway genes in the SCCA algorithm. Examples involving using prior knowledge of pathway genes can be found in Section 5 of the supplementary information [Wang et al. (2015)].

2. If the interest is to identify disjoint gene groups and running time is not a concern, we can run the whole procedure iteratively with no subsampling, each time identifying one dominating signal group and removing it from the subsequent analysis.

Asymptotic behavior of our procedure. Here we first show asymptotically the validity of our procedure by considering a simple case where there exists only one functional group and all the other genes are uncorrelated. Due to this simplification, no subsampling is needed, and the use of CCA without sparsity suffices since in the asymptotics we consider the regime of n (number of experiments) going to infinity with p (number of genes) fixed. Without loss of generality, in the entire gene set $G = \{1, 2, \dots, p\}$ let the first k genes $K = \{1, 2, \dots, k\}$ form one pathway.

For every partition t , let \mathbf{a}_t and \mathbf{b}_t be the solutions to (2.4) and \mathbf{c}_t be the list of the absolute values $|\mathbf{a}_t|$ and $|\mathbf{b}_t|$ ordered according to the gene list. Assuming \mathbf{Z} follows a multivariate normal distribution and the inverse covariance matrix has a diagonal block structure (detailed assumptions are presented in Section 2 of the supplementary information [Wang et al. (2015)]), we have the following proposition regarding the asymptotic difference between the values of $\{c_{i,t}, i \in K\}$ and $\{c_{j,t}, j \notin K\}$ averaged over t . For convenience suppose p is even and denote $q = p/2$.

PROPOSITION 2.1. Let $\bar{\mathbf{c}} = \sum_{t=1}^N \mathbf{c}_t / N$, where N is the number of partitions, then given $1 < k < q$,

$$(2.7) \quad \lim_{N \rightarrow \infty} \lim_{n \rightarrow \infty} \left(\min_{i \in K} \bar{c}_i - \max_{j \notin K} \bar{c}_j \right) = D$$

for some positive constant D .

In Section 2 of the supplementary information [Wang et al. (2015)], we give the proof of Proposition 2.1 with a lower bound on D that quantifies the asymptotic difference in the assigned weights between functional group genes and noisy genes. The separation in $\bar{\mathbf{c}}$ implies the genes in the graph characterized by the edge weight matrix $\bar{\mathbf{A}} = \bar{\mathbf{c}}\bar{\mathbf{c}}^T$ can be grouped into different clusters based on their connectivity.

To further understand the asymptotic behavior of our procedure in general cases when multiple functional groups exist, we present an example that consists of two (disjoint) groups of interacting genes and other unrelated genes in supplementary information Section 2 [Wang et al. (2015)]. We show a theoretical derivation of $\bar{\mathbf{A}} = 1/B \sum_{b=1}^B \mathbf{A}_b = 1/B \sum_{b=1}^B \bar{\mathbf{c}}_b \bar{\mathbf{c}}_b^T$ for this example in detail to highlight and explain the role of subsampling. We can see that with subsampling, the limiting $\bar{\mathbf{A}}$ (when $n \rightarrow \infty$) exhibits a natural block structure corresponding to the two gene groups, thus extending the validity of Proposition 2.1. The ideas underlying the analytical derivation in this simple example are straightforward and directly applicable to general cases, though the computations involved would be very tedious. Note that the analytical computations look tedious even in this small example.

2.4. Identify community structures given the edge weight matrix $\bar{\mathbf{A}}$. To demonstrate that $\bar{\mathbf{A}}$ possesses advantages over traditional approaches in identifying gene functional modules, subsequent analysis of $\bar{\mathbf{A}}$ based on community detection tools is needed. Many methods are available in this field. In particular, clustering has been a popular and well-studied technique. Jain, Murty and Flynn (1999), Kaufman and Rousseeuw (2009), Theodoridis and Koutroumbas (2005) provide general reviews of various clustering techniques, and reviews with more specific focus on gene expression data can be found in D'haeseleer, Liang and Somogyi (2000), Jiang, Tang and Zhang (2004), Kerr et al. (2008). Variants of spectral clustering are also widely explored for detecting communities in sparse networks [Ramesh et al. (2010)]. Viewing gene relationships as edges in a graph, a natural approach is to consider functional modules as tightly connected subgraphs. Genes with related functionalities are expected to have dense connections, whereas biologically unrelated (noisy) genes may be only sparsely connected. The Stochastic Block Model (SBM) builds a general probabilistic graph model based on such an assumption that nodes (genes) have different connectivities depending on their block memberships.

Below we introduce two popular community detection tools, SBM and hierarchical clustering (HC), which we will use in later simulation and real data analysis to dissect gene interaction groups from $\bar{\mathbf{A}}$. As we have mentioned, there are many other choices for performing this task. The structure of $\bar{\mathbf{A}}$ itself may also imply some methods are more suitable than others. In this paper, it is not our intention

to suggest or evaluate the best community detection tools that should be applied to $\bar{\mathbf{A}}$. Here we are presenting SBM and HC just as two illustrative approaches.

The SBM, formally introduced by Holland, Laskey and Leinhardt (1983), generalizes the Erdős–Rényi model and defines a family of probability distributions for a graph. Here is a detailed model definition.

DEFINITION 2.2. A SBM is a family of probability distributions for a graph with node set $\{1, 2, \dots, p\}$ and Q node blocks defined as follows:

1. Let $\mathbf{C} = (C_1, C_2, \dots, C_p)$ denote the set of labels such that $C_i = k$ if the node i belongs to block k :

$$\mathbf{C} \stackrel{\text{i.i.d.}}{\sim} \text{Multinomial}(\boldsymbol{\gamma}),$$

where $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots, \gamma_Q)$ is the vector of proportions.

2. Let $\boldsymbol{\pi} = (\pi_{lk})_{1 \leq l, k \leq Q}$ be a symmetric matrix of a block dependent edge probability matrix and \mathbf{A} be the adjacency matrix. Conditioned on the block labels \mathbf{C} , (\mathbf{A}_{ij}) for $i < j$ are independent, and

$$P(\mathbf{A}_{ij} | \mathbf{C}) = P(\mathbf{A}_{ij} = 1 | C_i = l, C_j = k) = \pi_{lk}.$$

Discretizing $\bar{\mathbf{A}}$ defined in Section 2.3 into a 0–1 matrix, the class labels and the parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\pi}$ are estimated using the psuedo-likelihood algorithm by Amini et al. (2013). The unconditional version of the algorithm fits the conventional SBM above, while the conditional version takes into account the variability of node degrees within blocks [Karrer and Newman (2011)]. Potential functional groups are identified as classes having large diagonal entries in $\boldsymbol{\pi}$.

Agglomerative HC is another widely used nonmodel-based technique for extracting communities, especially in the study of social networks [Scott and Peter (2011)]. In our application, we adopt Ward’s distance [Ward (1963)] for the computation of merging costs. Let g_i be the nodes, the distance between two clusters M_1, M_2 defined as

$$\begin{aligned} d(M_1, M_2) &= \frac{n_1 n_2}{n_1 + n_2} \|m_1 - m_2\|^2 \\ &= \frac{1}{2(n_1 + n_2)} \sum_{i, j \in M_1 \cup M_2} \|g_i - g_j\|^2 - \frac{1}{2n_1} \sum_{i, j \in M_1} \|g_i - g_j\|^2 \\ &\quad - \frac{1}{2n_2} \sum_{i, j \in M_2} \|g_i - g_j\|^2, \end{aligned}$$

where n_1 and n_2 denote the sizes of M_1 and M_2 , and m_1 and m_2 are the cluster centers of M_1 and M_2 , respectively. A natural way to define the square of the pairwise distance is $\|g_i - g_j\|^2 = 1 - \bar{\mathbf{A}}_{ij}$ for $i \neq j$, and zero otherwise. Since Ward’s method minimizes the increase in the within group sum of squares at each

merging and tends to merge clusters that are close to each other and small in size, a small cluster that manages to survive a long distance before coalescing is likely to be a tight cluster, indicating the genes it contains have high connectivity with each other. Thus, at an appropriately chosen cutoff level Q , we identify the smallest few clusters as potential functional groups.

Both SBM and HC require a priori knowledge of the number of clusters Q , and the proper selection of Q remains an open problem in the literature. For SBM, we refer to some discussions in Daudin, Picard and Robin (2008) and Channarond, Daudin and Robin (2012). For HC, a common way to choose the cutoff Q is to set it as the number just before the merging cost starts to rise sharply. Due to the scale and complexity of a typical gene expression data set, this criterion is not very applicable. In this paper, for the HC approach we choose Q empirically based on the sizes of the clusters each Q produces. That is, Q is increased incrementally until small clusters start to emerge. A comparison between SBM and HC can be found in Section 3.1.

2.5. Flow chart summarizing the whole procedure. A comprehensive summary of the whole procedure, including the tuning parameters needed in constructing $\bar{\mathbf{A}}$ and illustrative subsequent analysis of $\bar{\mathbf{A}}$, is provided in Figure 1. The choices of the parameters are explained in the paper and summarized again in Section 3 of the supplementary information [Wang et al. (2015)].

3. Results. In this section we evaluate the performance of the proposed method and other approaches using simulated and real microarray data sets. In particular, we compare the quality of the estimated gene functional groups, resulting from different ways of computing edge weights, and the two methods of community detection (SBM and HC) discussed in Section 2.4. We use *precision* and *recall*, defined as $precision = TP/(TP + FP)$ and $recall = TP/(TP + FN)$, as measures for evaluating classification performance. Here TP is the number of true positive findings of functional group genes, FP is the number of false positives and FN is the number of false negatives. In the context of this study, they can be regarded as a measure of exactness and completeness of our search results, respectively. The problems of choosing the appropriate proportion of subsampling and λ for sparsity are also discussed. For detailed analysis of the effects of the tuning parameters, we refer to Section 3.3 and Section 7 of the supplementary information [Wang et al. (2015)].

3.1. Simulation.

3.1.1. Generation of simulation data sets. We simulate a microarray data set consisting of $p = 150, 300$ or 500 genes and $n = 30$ experiments, with 5 replicates for each experiment. To make the data more realistic, we introduce experiment dependencies, multiple functional groups and random noise. The simulation parameters are generated as follows:

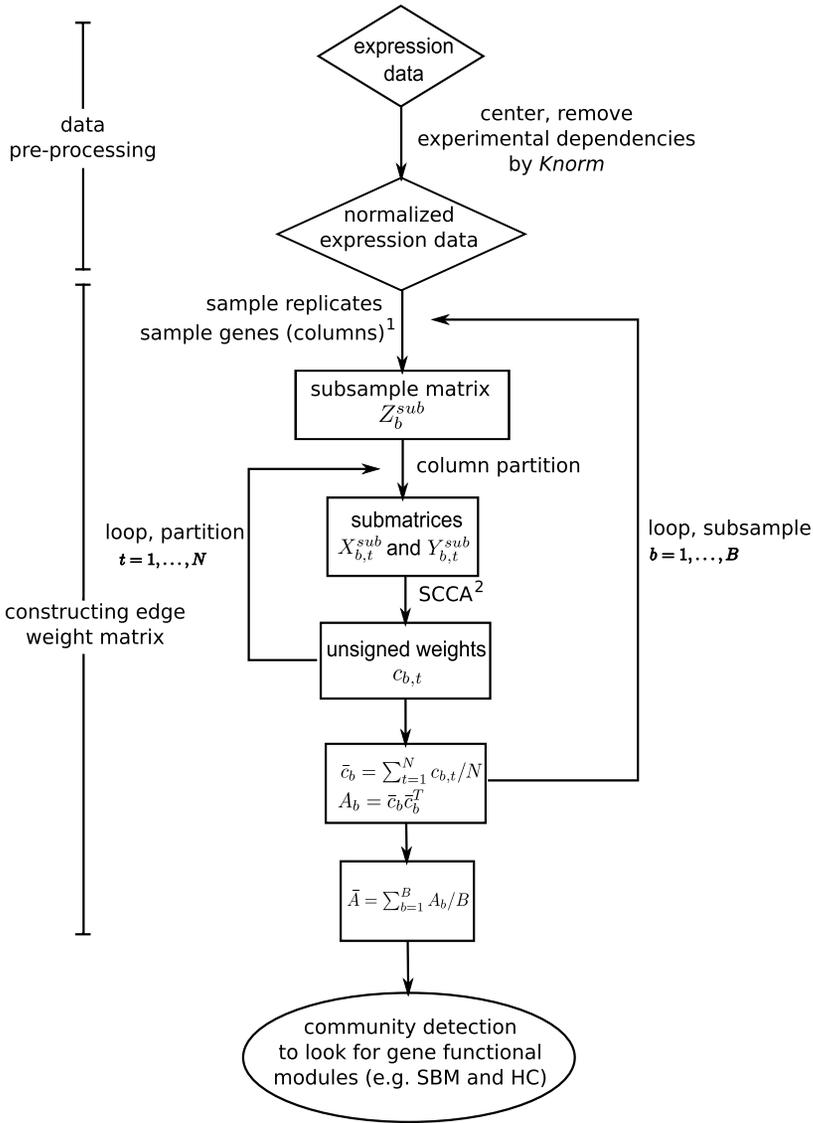


FIG. 1. Flow chart summarizing the whole procedure. Each numeric superscript in the diagram indicates the need for tuning parameters: 1. Subsampling level, 2. Penalty parameter λ .

(i) Experiment correlation matrix, Σ^E . For illustrative purpose, we set the experiment correlation matrix to have 0, 33 and 67% dependencies. In the case of a 33% dependency, for example, 33% of the experiments have high dependencies (correlation between 0.5 and 0.6) while the remaining experiments are uncorrelated with one another.

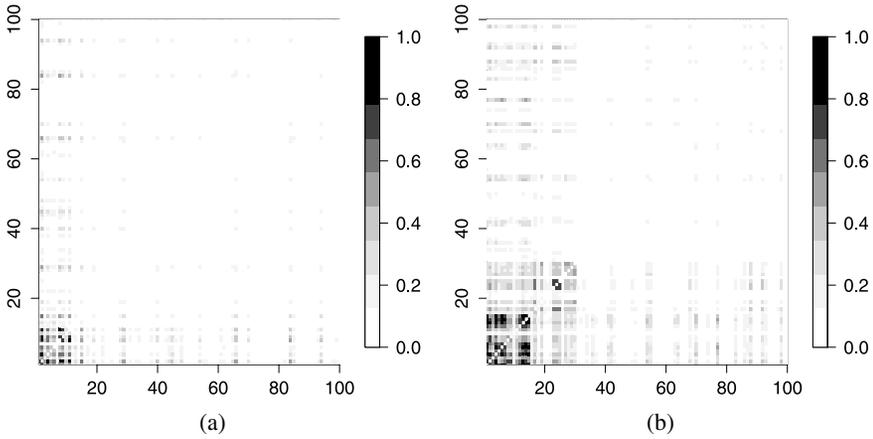


FIG. 2. Heatmaps of the matrix $\bar{\mathbf{A}}$ using data sets with (a) $p = 150$, 0% experiment dependency, one functional group, subsampling level 70% and $(\lambda_1, \lambda_2) = (9, 9)$; (b) $p = 300$, 0% experiment dependency, two functional groups, subsampling level 70% and $(\lambda_1, \lambda_2) = (9, 15)$. For clarity, only the first 100×100 entries are shown and the functional groups are placed at positions 1–15 and 16–30, respectively.

(ii) Gene correlation matrix, Σ^G . In each data set, we introduce one or two functional groups with 15 genes in each. Genes in the same group are correlated, having either high correlations (0.5–0.6) or low correlations (0.1–0.2) with the other genes, and otherwise they are not.

Using the above parameters, we generate the simulation data as follows. First, we generate a 30×500 gene expression matrix \mathbf{Z} , with $\text{vec}(\mathbf{Z}^T)$, from a multivariate normal distribution with mean zero and a covariance matrix $\Sigma^G \otimes \Sigma^E$. To introduce linear relationships, within each group we take linear combinations of some genes to replace their original values. Using the final 30×500 gene expression matrix, we add random noise with a small SD (e.g., 0.01) to each row to generate the 5 replicates for each experiment.

3.1.2. *Estimated $\bar{\mathbf{A}}$ and tuning parameter selection.* Figure 2 shows the heatmaps of the matrix $\bar{\mathbf{A}}$ for two data sets with different numbers of functional groups. For visual clarity, the genes are ordered according to their true group memberships. In both cases, the matrix demonstrates a clear block structure. In particular, in the two-group case both pathways are visible, although the first one is more prominent. We remark here that the difference in signal strength between the two pathways is introduced by chance variation during data generation and the use of subsampling is necessary for the identification of the weaker group. Although we present results obtained with a subsampling level of 70%, a range of reasonable subsampling levels can be chosen without significantly affecting the final results (supplementary information Section 7 [Wang et al. (2015)]). The other

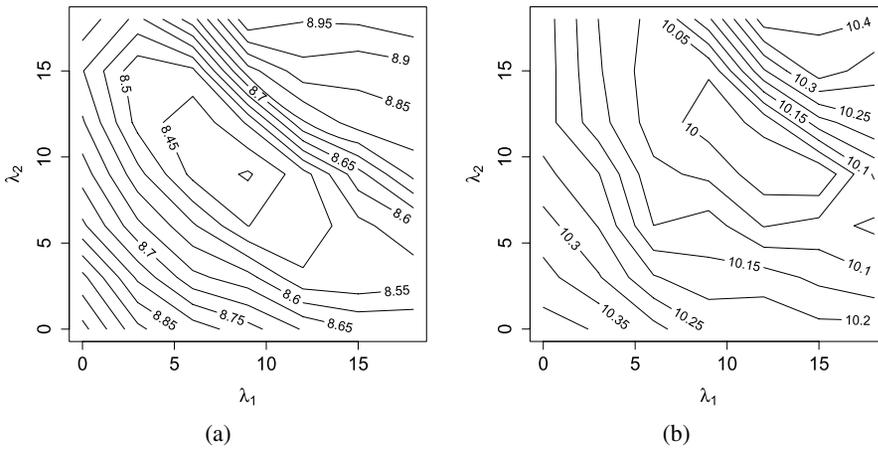


FIG. 3. Contour plots of the entropy of the upper triangular entries of $\bar{\mathbf{A}}$ on the grid $(\lambda_1, \lambda_2) \in \{0, 3, \dots, 18\}^2$ using data sets with (a) $p = 150$, 0% experiment dependency, one functional group and subsampling level 70%; (b) $p = 300$, 0% experiment dependency, two functional groups and subsampling level 70%.

tuning parameter λ is chosen such that the matrix $\bar{\mathbf{A}}$ displays optimal contrast between the pathway and nonpathway groups, and we shall use this as guidance for assessing the quality of $\bar{\mathbf{A}}$ and selecting λ .

Among the common approaches for the selection of optimal tuning parameters, cross-validation-based methods are used in Waaijenborg, Verselewele de Witt Hamer and Zwinderman (2008), Parkhomenko, Trichler and Beyene (2009) and Lee et al. (2011). However, all of their methods involve dividing a sample into multiple sets, which is impractical for data sets with only a few tens of observations. Witten and Tibshirani (2009) proposed an alternative permutation-based method which estimates the p -value of the maximal correlation found by performing SCCA on permuted samples. Due to the large number of partitions and subsamplings required in our method, this approach would be very computationally expensive. Instead we measure the effectiveness of λ using the entropy of $\bar{\mathbf{A}}$, defined as

$$(3.1) \quad H(\mathbf{A}) = - \sum_{i < j, \mathbf{A}_{ij} > 0} (\mathbf{A}_{ij} / S_{\mathbf{A}}) \log(\mathbf{A}_{ij} / S_{\mathbf{A}}),$$

where $S_{\mathbf{A}} = \sum_{i < j} \mathbf{A}_{ij}$. The entropy quantifies the sharpness of its distribution and thus is indicative of the signal intensity. Figure 3 plots the contours of $H(\bar{\mathbf{A}})$ for the same two data sets used in Figure 2. Regions with low entropy correspond to λ , leading to a matrix with better signal intensity.

3.1.3. *Performance comparison.* Figure 4 compares the classification performance of our methods, *scca.sbm* and *scca.hc*, with four correlation-based methods, *pearson.hc*, *pearson.sbm*, *module.dynamic* and *module.hybrid*. The methods

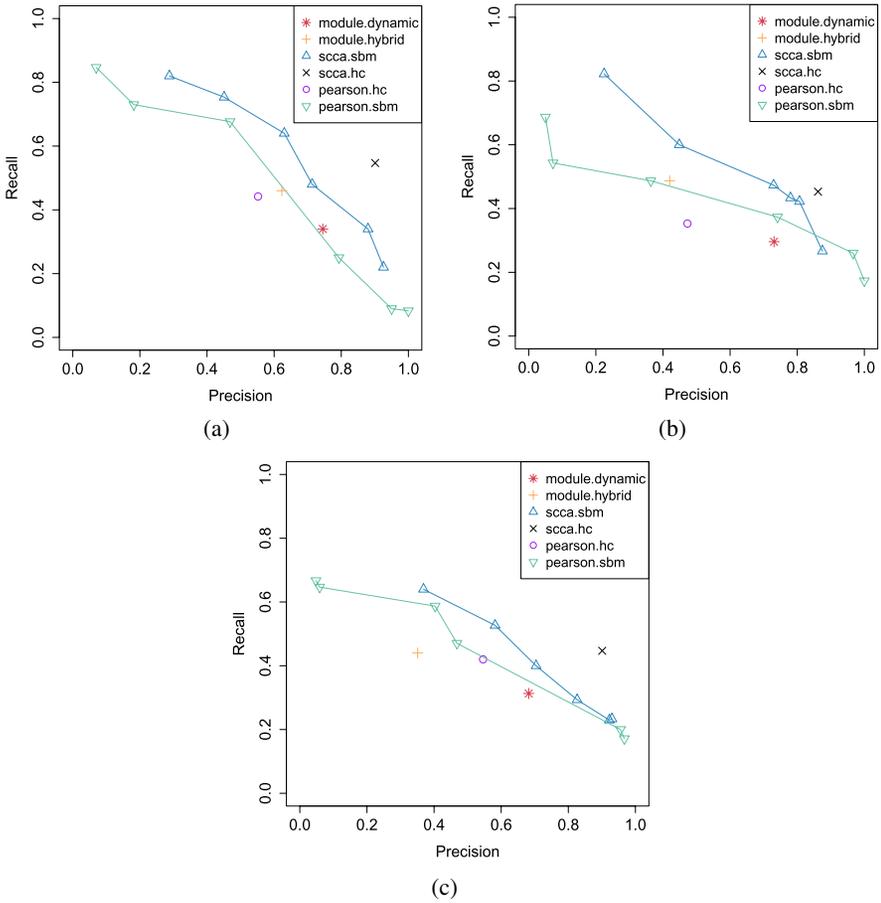


FIG. 4. Classification performance of different methods using data sets with $p = 500$, one pathway group, subsampling level 70%, and (a) 0%, (b) 33% and (c) 67% of experiment dependency. *pearson.sbm* and *scca.sbm* are applied to matrices at discretization levels $\{0.3, 0.4, \dots, 0.8\}$ (from left to right on the curve).

are named by cross-mixing the following to allow for comparisons in the two-stage procedure:

scca: Calculate $\bar{\mathbf{A}}$'s with $\lambda \in \{9, 12, \dots, 27\}^2$ and select 10 of these with the smallest entropy values. The final cluster membership (after community detection) is decided by a majority vote based on the selected $\bar{\mathbf{A}}$'s, so only stable clusters and cluster members are chosen.

pearson: Pearson's correlation matrix after the data is normalized using equation (2.6) and $Knorm$ estimates.

module: Transformed Pearson's correlation matrix used in Langfelder and Horvath (2007).

TABLE 1

Classification performance of different methods using data sets with $p = 500$, two pathway groups, subsampling level 70%, and various levels (0%, 33% and 67%) of experiment dependency

	0%		33%		67%	
	Precision	Recall	Precision	Recall	Precision	Recall
	Pathway 1					
<i>scca.hc</i>	0.861	0.533	0.831	0.441	0.811	0.433
<i>pearson.hc</i>	0.238	0.233	0.497	0.427	0.471	0.393
<i>module.dynamic</i>	0.718	0.3	0.742	0.333	0.764	0.38
<i>module.hybrid</i>	0.439	0.407	0.544	0.447	0.453	0.385
	Pathway 2					
<i>scca.hc</i>	0.808	0.487	0.890	0.489	0.833	0.420
<i>pearson.hc</i>	0.438	0.387	0.323	0.307	0.460	0.273
<i>module.dynamic</i>	0.758	0.4	0.808	0.347	0.8	0.4
<i>module.hybrid</i>	0.565	0.473	0.529	0.387	0.455	0.46

sbm: Fit a SBM on a discretized edge weight matrix (at level $\{0.3, 0.4, \dots, 0.8\}$) using the unconditional pseudo-likelihood algorithm in Amini et al. (2013) with $Q = 2$ (or 3) initialized by spectral clustering with perturbation. Select the cluster with the highest internal connectivity based on the estimates.

hc: HC with Ward's distance and cut the dendrogram when clusters of size less than 25 start to appear as the number of clusters Q increases. The choice of this upper bound is based on the size of the cluster selected in *scca.sbm*, and a range of reasonable numbers can be used without affecting the final results.

dynamic, hybrid: HC with dendrogram cutting methods in the R package `dynamicTreeCut` [Langfelder, Zhang and Horvath (2008)].

Figure 4 plots the average *precision* and *recall* of the above six methods calculated on 10 simulation data sets for each level of experiment dependency. It can be seen that using our SCCA approach to compute edge weights in general leads to higher *precision* across all experiment dependency levels. Of the two ways of community identification, *scca.hc* produces higher *precision* than *scca.sbm* at comparable *recall* levels.

Table 1 shows the same performance measures obtained from data sets containing two independent functional groups for *scca.hc*, *pearson.hc*, *module.dynamic* and *module.hybrid*. The numbers are averages from 10 simulation data sets for each level of experiment dependency. Similar to the one-group case, we choose the smallest Q that produces two clusters of size less than 25 as the cutoff in HC. We remark here that when multiple groups are present, *scca.sbm* tends to detect only the strongest signal group while failing to pick up the weaker one. This can be explained by considering the within-class homogeneity assumption in the SBM model and noting that the degree distribution is often less homogeneous in the weaker signal group (see, e.g., Figure 2). The conditional pseudo-likelihood

algorithm in Amini et al. (2013) is also not sensitive enough to detect the finer distinctions. Results from *pearson.sbm* are also omitted as they are very noisy. In all the cases, *scca.hc* demonstrates the best *precision* at comparable, if not better, *recall*.

3.2. Application to real data. We tested the performance of our procedure by applying it to *Arabidopsis thaliana* microarray expression data retrieved from AtGenExpress (http://www.arabidopsis.org/servlets/TairObject?type=expression_set&id=1007966941). The analyzed data set included expression measurements collected from shoot tissues subject to oxidation stress for 22,810 genes under 13 experiment conditions with two replicates for each experiment. In these experiments, the plants were treated with methyl viologen (MV), which led to the formation of reactive oxygen species (ROS). Various studies have shown that depending on the type of ROS, a different biological response is provoked. Thus, by focusing on the ROS induced by MV, we were able to show and validate that the results of our pathway gene search were supported, in part, by other already published ROS-related microarray experiments.

A subset of all 22,810 genes was selected for analysis based on the following criteria. (i) The experiment variance of the gene exceeds 0.1. An unvarying expression profile suggests the gene has an activity level unaltered by the particular stress condition, and hence is unlikely to be part of any stress-induced pathway. The inclusion of such genes may cause problems in covariance estimation as well. We also removed genes with a suspiciously high experiment variance, as it could suggest inaccuracy in measurements. (ii) The discrepancy between the two replicates is smaller than 2 for each experiment. This ensures only genes with consistent measurements are included in our analysis. (iii) The minimum expression level exceeds 7. More active genes are likely to possess stronger signals, making our search easier. This requirement further trims down the data set to a smaller size more desirable for our procedure. We note here that the inclusion of (iii) is optional—if running time is not a concern, the minimum expression level could be either lowered or entirely removed. The final subset for analysis contained 2718 genes.

Potential functional groups were found by *scca.hc*. Due to the complexity and noise level of the data set, we did not expect the entropy (3.1) to have a clean-cut unimodal distribution. Furthermore, the presence of many groups with varying signal strengths implies each may need a different optimal λ for detection. For example, strong groups are likely to require more regularization or, in other words, larger λ . For this reason, we performed our search in multiple stages starting from large λ for stronger groups to smaller λ for weaker ones. At every stage, the groups found were removed from the original set before proceeding to the next stage. The upper bound on λ was found by increasing λ until the entropy stabilized. Searching down from this upper bound, we chose λ from three grids: $\{90, 100, 110\}^2$,

TABLE 2
GO enrichment of groups

Group ID	Enriched GO term	Number of genes with enriched terms	P-values
1	Chloroplast organellar gene	10 out of 15 ¹	1.10×10^{-4}
2	Phenylpropanoid-flavonoid biosynthesis	3 out of 4	6.65×10^{-7}
3	Glucosinolate biosynthesis	7 out of 7	1.95×10^{-14}
4	Chloroplast organellar gene	3 out of 3	7.83×10^{-3}
5	Ribosome	10 out of 15	7.20×10^{-13}
8	Ribosome	5 out of 6	8.31×10^{-8}
10	Photosystem I or II	8 out of 10	2.87×10^{-14}
12	Endomembrane system	3 out of 4	2.35×10^{-3}

¹4 out of the 10 chloroplast genes are mitochondrial organellar genes.

{60, 70, 80}² and {30, 40, 50}². The cutoff level Q in HC was increased incrementally until at least five clusters of size less than 30 appeared. A reasonable range of numbers can be used to choose the cutoff and our results are not very sensitive to the choice of this number. The full procedure produced 13 groups of genes, the full list of which, including annotations, can be found in Section 6 of the supplementary information [Wang et al. (2015)].

To test the biological significance of all 13 groups found (i.e., whether there is a functional relationship between genes within the various groups), we first examined for enrichment of gene product properties, collectively designated gene ontology (GO) annotations, within each group using information available at The Arabidopsis Information Resource (<http://www.arabidopsis.org/tools/bulk/index.jsp>). We determined that 8 out of 13 groups were highly enriched with genes having the same GO annotation and calculated their p -values using Fisher's exact test to compare with the counts obtained from the full analyzed data set (Table 2).

In addition to the GO enrichment approach for validating the groups, and in order to support the biological significance of the groups found, we also evaluated other forms of evidence. We were able to determine that for several groups the genes placed in the groups encode for known pathways. For example, group 2 genes encode steps in the phenylpropanoid-flavonoid (FB) biosynthesis pathway, and group 3 genes encode for steps in the glucosinolate (GSL) biosynthesis pathway. Both are well-studied secondary metabolic pathways. Flavonoids are compounds of diverse biological activities such as anti-oxidants, functioning in UV protection, in defense, in auxin transport inhibition and in flower coloring [Gachon et al. (2005), Naoumkina et al. (2010), Taylor and Grotewold (2005), Woo, Jeong and Hawes (2005)], and GSLs are sulfur-rich amino acid-containing compounds which become active in response to tissue damage and are believed to offer a protective function [Sønderby, Geu-Flores and Halkier (2010), Verkerk et al. (2009),

TABLE 3
GO enrichment of groups—first cut

Group ID	Enriched GO term	Number of genes with enriched terms	<i>P</i> -values
9	Cell wall	16 out of 81	4.46×10^{-6}
10	Defense response	29 out of 78	1.58×10^{-2}
11	Phenylpropanoid-flavonoid biosynthesis	11 out of 76	5.42×10^{-12}

Yan and Chen (2007)]. A considerable number of genes in both pathways are induced by broad environmental stresses and regulated at the transcriptional level. Based on the lists of genes associated with these two pathways reported in Kim et al. (2012), our analyzed data set contained 13 FB pathway genes and 26 GSL pathway genes. The precisions of our search are 75% and 100%, respectively.

In order to assess the likelihood that genes in the remaining groups could also encode steps within specific pathways, we reviewed microarray data from plants subjected to other forms of oxidative stress (these experiments are similar to the experiment from which our data set using MV was obtained). Using this approach we found that genes in each of the additional seven groups (1, 4, 5, 8, 9, 11, 12) were strongly associated in these independent experiments (supplementary information Section 6 [Wang et al. (2015)]).

Of all the groups found, groups 6, 7 and 13 remain uncharacterized in the literature. Nonetheless, using CoExSearch [part of the ATTD-II database (http://atted.jp/top_search.shtml #CoexVersion)], all four genes in group 7 were correlated to some degree with abiotic stress conditions. We also found these genes were common anoxia-repressed genes [Loreti et al. (2005)]. The lack of complete characterization for these groups in the current literature leaves potential scope for further biological examination.

For comparison we applied *pearson.hc*, *module.dynamic* and *module.hybrid* to the same data. As the simulation study suggests the latter two methods in general have better performance than *pearson.hc*, particularly in the multi-group case, we will present the results from these two methods and refer to Section 6 in the supplementary information [Wang et al. (2015)] for *pearson.hc*-based results. In order to compare with our results, we chose two cuts of the dendrogram such that the first cut produced the same number of groups as our method, and the second one led to groups with sizes comparable to ours. The first cut resulted in 13 groups with sizes ranging from 60 to 293. We picked the three most promising groups based on their annotations and the GO analysis is summarized in Table 3. Although all of them have statistically significant *p*-values, their precisions are quite low. In particular, group 11 contains our group 2 as a subset and includes 11 genes (out of 76) in the FB pathway and 5 genes are in the isoprenoid biosynthesis pathway. These two pathways are derived from different initial precursors and are known

TABLE 4
GO enrichment of groups—second cut

Group ID	Enriched GO term	Number of genes with enriched terms
62	NA	0 out of 6
63	Chloroplast	4 out of 6
64	Located in plasma membrane	2 out of 5
65	Located in plasma membrane	3 out of 5
66	Pyridoxine biosynthetic process	2 out of 5

to be unrelated. We note here that at this cut level, the GSL pathway cannot be identified by the method. The second cut produces 66 groups with sizes from 5 to 81. We picked five small groups for analysis and only one group with genes localized in chloroplast has significant GO enrichment (Table 4). Even so, these genes are unlikely to be functionally related. The comparison suggests our method can achieve better precision and lead to more biologically meaningful groupings of genes.

3.3. Effects of tuning parameters. To systematically study the effects of different tuning parameters on the identification of gene functional groups, we perform sensitivity analysis for different choices of subsampling levels and penalty parameter λ using both the simulated and real data discussed above. For the sake of completeness, we also compare tuning parameters from the HC and SBM procedures. Overall, our results are reasonably stable for a range of λ values. Further stability can be achieved by pooling results from different λ . As expected, the choice of subsampling level is more important when there exist multiple functional groups. Our results suggest levels between 50% and 80% can all be considered in practice. For community detection, HC is more robust than SBM in the sense that the classification results are not sensitive to the cutoff chosen. The results are summarized in Section 7 of the supplementary information [Wang et al. (2015)].

4. Discussion. In this paper we focus on the problem of estimating gene group interactions in gene networks, where data are given in the form of nodes and their associated covariates and estimation of the true network is a challenging task. We propose a new method to construct an edge weight matrix for the full network by applying SCCA to sampled subsets of genes with random partitioning. To evaluate the quality of the constructed network, subsequent analysis of the community structures is applied to identify potential gene functional groups. Although the work is presented under the setting of gene networks, we believe our approach can be generally applicable to answer similar questions in other biochemical networks and even networks in other fields that are sparse and have similar covariate features.

Compared to other popular ways of measuring gene interactions, our SCCA approach is more conceptually appealing. By seeking maximally correlated sets of genes among randomly sampled subsets, this approach provides an aggregated measure of gene partial correlations when the correct conditional set is unknown, and thus gives us a better chance of capturing group interactions. As demonstrated in both simulation and real data applications, one of the main attractions of our procedure is its high *precision*. Although it does not seem to greatly improve *recall*, this is not a huge drawback in light of the search algorithm by Kim et al. (2012). Given the accuracy of our search results in general, one can use these identified genes as “seed genes” to initiate a more complete search and expand on the current lists.

Our approach can be modified to handle other practical situations. When it is known in advance that some genes operate in the same functional group, one may incorporate the prior knowledge by lowering the penalties associated with those genes in the SCCA algorithm. Although we have focused on the case with disjoint functional groups, our method of constructing an edge weight matrix is still applicable to the overlapping case as long as the shared genes possess strong direct or partial interactions with all the other functional genes (supplementary information Section 5 [Wang et al. (2015)]). However, a different community detection method [e.g., mixed membership SBM; Airoldi et al. (2008)] should be applied to identify the overlapping structures.

The core of our procedure consists of an implementation of SCCA by LASSO regression, and this naturally opens room for further investigation. For example, it would be interesting to find out if using other penalty functions yields different results, more importantly, whether SCCA can be implemented using a different optimization criterion or a more efficient algorithm to lessen the computational cost of our procedure. In the theoretical aspect, it would be desirable to incorporate sparsity into our asymptotic analysis.

On the community detection side, although we used SBM and HC as examples, there are many other available methods to be further explored, especially their properties in relation to the edge weight matrix $\bar{\mathbf{A}}$. The use of SBM and HC also gives rise to other interesting extensions. As noted in Section 3.1, conventional SBM does not perform well when there are multiple groups, which is mainly caused by the heterogeneity of node degrees. However, fitting a degree-corrected model using the conditional pseudo-likelihood algorithm does not seem to offer significant improvement. It would be desirable to carry out further study on the theoretical properties of the degree-corrected SBM and characterize its identifiability problem. Another possible extension is to modify these algorithms to take weighted adjacency matrices without discretization. Developing a practical but more systematic way of choosing the cutoff level for HC also invites future study.

SUPPLEMENTARY MATERIAL

Supplementary information (DOI: [10.1214/14-AOAS792SUPP](https://doi.org/10.1214/14-AOAS792SUPP); .pdf). Asymptotic analysis and additional explanations of the procedure, additional simulation and real data results. The code for estimating the edge weight matrix can be requested from hhuang@stat.berkeley.edu.

REFERENCES

- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9** 1981–2014.
- AMINI, A. A., CHEN, A., BICKEL, P. J. and LEVINA, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. *Ann. Statist.* **41** 2097–2122. [MR3127859](#)
- CHANNAROND, A., DAUDIN, J.-J. and ROBIN, S. (2012). Classification and estimation in the stochastic blockmodel based on the empirical degrees. *Electron. J. Stat.* **6** 2574–2601. [MR3020277](#)
- D’HAESELEER, P., LIANG, S. and SOMOGYI, R. (2000). Genetic network inference: From co-expression clustering to reverse engineering. *Bioinformatics* **16** 707–726.
- DAUB, C. O., STEUER, R., SELBIG, J. and KLOSKA, S. (2004). Estimating mutual information using B-spline functions—An improved similarity measure for analysing gene expression data. *BMC Bioinformatics* **5** 118.
- DAUDIN, J.-J., PICARD, F. and ROBIN, S. (2008). A mixture model for random graphs. *Stat. Comput.* **18** 173–183. [MR2390817](#)
- DE LA FUENTE, A., BING, N., HOESCHELE, I. and MENDES, P. (2004). Discovery of meaningful associations in genomic data using partial correlation coefficients. *Bioinformatics* **20** 3565–3574.
- EDWARDS, D. (2000). *Introduction to Graphical Modelling*, 2nd ed. Springer, New York. [MR1880319](#)
- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- GACHON, C. M. M., LANGLOIS-MEURINNE, M., HENRY, Y. and SAINDRENAN, P. (2005). Transcriptional co-regulation of secondary metabolism enzymes in Arabidopsis: Functional and evolutionary implications. *Plant Mol. Biol.* **58** 229–245.
- HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5** 109–137. [MR0718088](#)
- HOTELLING, H. (1936). Relations between two sets of variates. *Biometrika* **28** 321–377.
- JAIN, A. K., MURTY, M. N. and FLYNN, P. J. (1999). Data clustering: A review. *ACM Computing Surveys* **31** 264–323.
- JIANG, D., TANG, C. and ZHANG, A. (2004). Cluster analysis for gene expression data: A survey. *IEEE Transactions on Knowledge and Data Engineering* **16** 1370–1386.
- KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Phys. Rev. E* (3) **83** 016107, 10. [MR2788206](#)
- KAUFMAN, L. and ROUSSEEUW, P. J. (2009). *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley, New York.
- KERR, G., RUSKIN, H. J., CRANE, M. and DOOLAN, P. (2008). Techniques for clustering gene expression data. *Comput. Biol. Med.* **38** 283–293.
- KIM, K., JIANG, K., TENG, S. M., FELDMAN, L. J. and HUANG, H. (2012). Using biologically interrelated experiments to identify pathway genes in Arabidopsis. *Bioinformatics* **28** 815–822.
- KINNEY, J. B. and ATWAL, G. S. (2014). Equitability, mutual information, and the maximal information coefficient. *Proc. Natl. Acad. Sci. USA* **111** 3354–3359. [MR3200177](#)

- LANGFELDER, P. and HORVATH, S. (2007). Eigengene networks for studying the relationships between co-expression modules. *BMC Syst. Biol.* **1** 54.
- LANGFELDER, P., ZHANG, B. and HORVATH, S. (2008). Defining clusters from a hierarchical cluster tree: The Dynamic Tree Cut package for R. *Bioinformatics* **24** 719–720.
- LEE, W., LEE, D., LEE, Y. and PAWITAN, Y. (2011). Sparse canonical covariance analysis for high-throughput data. *Stat. Appl. Genet. Mol. Biol.* **10** Art. 30, 26. [MR2823516](#)
- LI, K.-C. (2002). Genome-wide coexpression dynamics: Theory and application. *Proc. Natl. Acad. Sci. USA* **99** 16875–16880.
- LORETI, E., POGGI, A., NOVI, G., ALPI, A. and PERATA, P. (2005). A genome-wide analysis of the effects of sucrose on gene expression in Arabidopsis seedlings under anoxia. *Plant Physiol.* **137** 1130–1138.
- MAGWENE, P. and KIM, J. (2004). Estimating genomic coexpression networks using first-order conditional independence. *Genome Biology* **5** R100.
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the Lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363](#)
- NAOUMKINA, M. A., ZHAO, Q., GALLEGÓ-GIRALDO, L., DAI, X., ZHAO, P. X. and DIXON, R. A. (2010). Genome-wide analysis of phenylpropanoid defence pathways. *Mol. Plant Pathol.* **11** 829–846.
- NEWMAN, M. E. J. (2010). *Networks: An Introduction*. Oxford Univ. Press, Oxford. [MR2676073](#)
- PARKHOMENKO, E., TRITCHLER, D. and BEYENE, J. (2009). Sparse canonical correlation analysis with application to genomic data integration. *Stat. Appl. Genet. Mol. Biol.* **8** Art. 1, 36. [MR2471148](#)
- PENG, J., ZHOU, N. and ZHU, J. (2009). Partial correlation estimation by joint sparse regression models. *J. Amer. Statist. Assoc.* **104** 735–746. [MR2541591](#)
- RAMESH, A., TREVINO, R., VON HOFF, D. D. and KIM, S. (2010). Clustering context-specific gene regulatory networks. In *Pacific Symposium on Biocomputing* 444–455.
- 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.
- SCHÄFER, J. and STRIMMER, K. (2005). An empirical Bayes approach to inferring large-scale gene association networks. *Bioinformatics* **21** 754–764.
- SCOTT, J. and PETER, J. C. (2011). *The SAGE Handbook of Social Network Analysis*. SAGE Publications, London.
- SØNDERBY, I. E., GEU-FLORES, F. and HALKIER, B. A. (2010). Biosynthesis of glucosinolates—Gene discovery and beyond. *Trends in Plant Science* **15** 283–290.
- STEUER, R., KURTHS, J., DAUB, C. O., WEISE, J. and SELBIG, J. (2002). The mutual information: Detecting and evaluating dependencies between variables. *Bioinformatics* **18** S231–S240.
- TAYLOR, L. P. and GROTEWOLD, E. (2005). Flavonoids as developmental regulators. *Curr. Opin. Plant Biol.* **8** 317–323.
- TENG, S. L. and HUANG, H. (2009). A statistical framework to inter functional gene relationships from biologically interrelated microarray experiments. *J. Amer. Statist. Assoc.* **104** 465–473. [MR2751431](#)
- THEODORIDIS, S. and KOUTROUMBAS, K. (2005). *Pattern Recognition*, 4th ed. Academic Press, Burlington, MA.
- VERKERK, R., SCHREINER, M., KRUMBEIN, A., CISKA, E., HOLST, B., ROWLAND, I., SCHRIJVER, R. D., HANSEN, M., GERHÄUSER, C., MITHEN, R. and DEKKER, M. (2009). Glucosinolates in Brassica vegetables: The influence of the food supply chain on intake, bioavailability and human health. *Mol. Nutr. Food Res.* **53** Suppl 2 S219.
- WAAIJENBORG, S., VERSELEWEL DE WITT HAMER, P. C. and ZWINDERMAN, A. H. (2008). Quantifying the association between gene expressions and DNA-markers by penalized canonical correlation analysis. *Stat. Appl. Genet. Mol. Biol.* **7** Art. 3, 29. [MR2386320](#)

- WANG, Y. X. R., JIANG, K., FELDMAN, L. J., BICKEL, P. J. and HUANG, H. (2015). Supplement to “Inferring gene–gene interactions and functional modules using sparse canonical correlation analysis.” DOI:10.1214/14-AOAS792SUPP.
- WANG, Y. X. R. and HUANG, H. (2014). Review on statistical methods for gene network reconstruction using expression data. *J. Theoret. Biol.* **362** 53–61.
- WARD, J. H. JR. (1963). Hierarchical grouping to optimize an objective function. *J. Amer. Statist. Assoc.* **58** 236–244. MR0148188
- WILLE, A. and BÜHLMANN, P. (2006). Low-order conditional independence graphs for inferring genetic networks. *Stat. Appl. Genet. Mol. Biol.* **5** Art. 1, 34 pp. (electronic). MR2221304
- WILLE, A., ZIMMERMANN, P., VRANOVA, E., FÜRHOLZ, A., LAULE, O., BLEULER, S., HENNIG, L., PRELIC, A., VON ROHR, P., THIELE, L., ZITZLER, E., GRUISSEM, W. and BÜHLMANN, P. (2004). Sparse graphical Gaussian modeling of the isoprenoid gene network in *Arabidopsis thaliana*. *Genome Biology* **5** 1–13.
- WITTEN, D. M. and TIBSHIRANI, R. (2009). Extensions of sparse canonical correlation analysis with applications to genomic data. *Stat. Appl. Genet. Mol. Biol.* **8** 1–27.
- WITTEN, D. M., TIBSHIRANI, R. and HASTIE, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics* **10** 515–534.
- WOO, H.-H., JEONG, B. R. and HAWES, M. C. (2005). Flavonoids: From cell cycle regulation to biotechnology. *Biotechnol. Lett.* **27** 365–374.
- YAN, X. and CHEN, S. (2007). Regulation of plant glucosinolate metabolism. *Planta* **226** 1343–1352.
- ZHOU, S., RÜTIMANN, P., XU, M. and BÜHLMANN, P. (2011). High-dimensional covariance estimation based on Gaussian graphical models. *J. Mach. Learn. Res.* **12** 2975–3026. MR2854354

Y. X. R. WANG
P. J. BICKEL
H. HUANG
DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: rachelwang@stat.berkeley.edu
bickel@stat.berkeley.edu
hhuang@stat.berkeley.edu

K. JIANG
L. J. FELDMAN
DEPARTMENT OF PLANT AND MICROBIAL BIOLOGY
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: kenij@berkeley.edu
ljfeldman@berkeley.edu