SPARSE INTEGRATIVE CLUSTERING OF MULTIPLE OMICS DATA SETS

BY RONGLAI SHEN[1], SIJIAN WANG AND QIANXING MO

Memorial Sloan-Kettering Cancer Center, University of Wisconsin-Madison and Baylor College of Medicine

High resolution microarrays and second-generation sequencing platforms are powerful tools to investigate genome-wide alterations in DNA copy number, methylation and gene expression associated with a disease. An integrated genomic profiling approach measures multiple omics data types simultaneously in the same set of biological samples. Such approach renders an integrated data resolution that would not be available with any single data type. In this study, we use penalized latent variable regression methods for joint modeling of multiple omics data types to identify common latent variables that can be used to cluster patient samples into biologically and clinically relevant disease subtypes. We consider lasso [*J. Roy. Statist. Soc. Ser. B* **58** (1996) 267–288], elastic net [*J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** (2005) 301–320] and fused lasso [*J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** (2005) 91–108] methods to induce sparsity in the coefficient vectors, revealing important genomic features that have significant contributions to the latent variables. An iterative ridge regression is used to compute the sparse coefficient vectors. In model selection, a uniform design [*Monographs on Statistics and Applied Probability* (1994) Chapman & Hall] is used to seek "experimental" points that scattered uniformly across the search domain for efficient sampling of tuning parameter combinations. We compared our method to sparse singular value decomposition (SVD) and penalized Gaussian mixture model (GMM) using both real and simulated data sets. The proposed method is applied to integrate genomic, epigenomic and transcriptomic data for subtype analysis in breast and lung cancer data sets.

1. Introduction. Clustering analysis is an unsupervised learning method that aims to group data into distinct clusters based on a certain measure of similarity among the data points. Clustering analysis has many applications in a wide variety of fields including pattern recognition, image processing and bioinformatics. In gene expression microarray studies, clustering cancer samples based on their gene expression profile has revealed molecular subgroups associated with histopathological categories, drug response and patient survival differences [Perou et al. (1999), Alizadeh et al. (2000), Sorlie et al. (2001), Lapointe et al. (2003), Hoshida et al. (2003)].

In the past few years, *integrative genomic studies* are emerging at a fast pace where in addition to gene expression data, genome-wide data sets capturing somatic mutation patterns, DNA copy number alterations and DNA methylation changes are simultaneously obtained in the same biological samples. A fundamental challenge in translating cancer genomic findings into clinical application lies in the ability to find "driver" genetic and genomic alterations that contribute to tumor initiation, progression and metastasis [Chin and Gray (2008), Simon (2010)]. As integrated genomic studies have emerged, it has become increasingly clear that true oncogenic mechanisms are more visible when combining evidence across patterns of alterations in DNA copy number, methylation, gene expression and mutational profiles [Cancer Genome Atlas Research Network (2008), TCGA Network (2011)]. Integrative analysis of multiple "omic" data types can help the search for potential "drivers" by uncovering genomic features that tend to be dysregulated by multiple mechanisms [Chin and Gray (2008)]. A well-known example is the *HER2* oncogene which can be activated through DNA amplification and mRNA over-expression. We will discuss the *HER2* example further in our motivating example.

In this paper, we focus on the class discovery problem given multiple omics data sets (multidimensional data) for tumor subtype discovery. A major challenge in subtype discovery based on gene expression microarray data is that the clinical and therapeutic implications for most existing molecular subtypes of cancer are largely unknown. A confounding factor is that expression changes may be related to cellular activities independent of tumorigenesis, and therefore leading to subtypes that may not be directly relevant for diagnostic and prognostic purposes. By contrast, as we have shown in our previous work [Shen, Olshen and Ladanyi (2009)], a joint analysis of multiple omics data types offers a new paradigm to gain additional insights. Individually, none of the genomic-wide data type alone can completely capture the complexity of the cancer genome or fully explain the underlying disease mechanism. Collectively, however, true oncogenic mechanisms may emerge as a result of joint analysis of multiple genomic data types.

Somatic DNA copy number alterations are key characteristics of cancer [Beroukhim et al. (2010)]. Copy number gain or amplification may lead to activation of oncogenes (e.g., *HER2* in Figure 1). Tumor suppressor genes can be inactivated by copy number loss. High-resolution array-based comparative genomic hybridization (aCGH) and SNP arrays have become dominant platforms for generating genome-wide copy number profiles. The measurement typical of aCGH platforms is a log-ratio of normalized intensities of genomic DNA in experimental versus control samples. For SNP arrays, copy number measures are represented by log of total copy number (logR) and parent-specific copy number as captured by a B-allele frequency (BAF) [Chen, Xing and Zhang (2011), Olshen et al. (2011)]. Both platforms generate contiguous copy number measures along ordered chromosomal locations (an example is given in Figure 6). Spatial smoothing methods are desirable for modeling copy number data.

In addition to copy number aberrations, there are widespread DNA methylation changes at CpG dinucleotide sites (regions of DNA where a Cytocine nucleotide occurs next to a Guanine nucleotide) in the cancer genome. DNA methylation is the most studied epigenetic event in cancer [Feinberg and Vogelstein (1983), Holliday (1979), Laird (2003, 2010)]. Tumor suppressor genes are frequently inactivated by hypermethylation (increased methylation of CpG sites in the promoter region of the gene), and oncogenes can be activated through the promoter hypomethylation. DNA methylation arrays measure the intensities of methylated probes relative to unmethylated probes for tens of thousands of CpG sites located at promoter regions of protein coding genes. M-values are calculated by taking log-ratios of methylated and unmethylated probe intensities [Irizarry et al. (2008)], similar to the M-values used for gene expression microarrays which quantify the relative expression level (abundance of a gene's mRNA transcript) in cancer samples compared to a normal control.

In this paper, we focus on the class discovery problem given multiple omics data sets for tumor subtype discovery. Suppose $t = 1, \ldots, T$ different genome-scale data types (DNA copy number, methylation, mRNA expression, etc.) are obtained in $j = 1, \ldots, n$ tumor samples. Let $\mathbf{X}_t$ be the $p_t \times n$ data matrix where $\mathbf{x}_{it}$ denote the ith row and $\mathbf{x}_{jt}$ the jth column of $\mathbf{X}_t$. Rows are genomic features and columns are samples. Here we use the term *genomic feature* and the corresponding feature index i in the equations throughout the paper to refer to either a protein-coding gene (typically for expression and methylation data) or ordered genomic elements that do not necessarily have a one-to-one mapping to a specific gene (copy number measure along chromosomal positions) depending on the data type.

Let $\mathbf{Z}$ be a $g \times n$ matrix where rows are latent variables and columns are samples, and g is the number of latent variables. Latent variables can be interpreted as "fundamental" variables that determine the values of the original p variables [Jolliffe (2002)]. In our context, we use latent variables to represent disease driving factors (underlying the wide spectrum of genomic alterations of various types) that determine biologically and clinically relevant subtypes of the disease. Typically, $g \ll \sum_t p_t$, providing a low-dimension latent subspace to the original genomic feature space. Following a similar argument for reduced-rank linear discriminant analysis in Hastie, Tibshirani and Friedman (2009), a rank-g approximation where $g \leq K - 1$ is sufficient for separating K clusters among the n data points. For the rest of the paper, we assume the dimension of $\mathbf{Z}$ is $(K - 1) \times n$ with mean zero and identity covariance matrix. A joint latent variable model expressed in matrix form is

$$\mathbf{X}_t = \mathbf{W}_t \mathbf{Z} + \mathbf{E}_t, \qquad t = 1, \ldots, T. \tag{1}$$

In the above, $\mathbf{W}_t$ is a $p_t \times (K - 1)$ coefficient (or loading) matrix relating $\mathbf{X}_t$ and $\mathbf{Z}$ with $\mathbf{w}_{jt}$ being the jth row and $\mathbf{w}_{kt}$ the kth column of $\mathbf{W}_t$, and $\mathbf{E}_t$ is a $p_t \times n$ matrix where the column vectors $\mathbf{e}_j$, $j = 1, \ldots, n$, represent uncorrelated error terms

that follow a multivariate distribution with mean zero and a diagonal covariance matrix $\boldsymbol{\Psi}_t = (\sigma_1^2, \ldots, \sigma_{p_t}^2)$. Each data matrix is row-centered so no intercept term is presented in equation (1).

Equation (1) provides an effective integration framework in which the latent variables $\mathbf{Z} = (\mathbf{z}_1, \ldots, \mathbf{z}_{K-1})$ are common for all data types, representing a probabilistic low-rank approximation simultaneously to the T original data matrices. In Section 3.2 we point out its connection and differences from singular value decomposition (SVD). In Sections 6 and 7 we illustrate that applying SVD to the combined data matrix broadly fails to achieve an effective integration of various data types.

Equation (1) is the basis of our initial work [Shen, Olshen and Ladanyi (2009)] in which we introduced an integrative model called iCluster. We considered a soft-thresholding estimate of $\mathbf{W}_t$ that continuously shrinks the coefficients for noninformative features toward zero. The motivation for sparse coefficient vectors is clearly indicated by Figure 1 panels (D) and (E). A basic sparsity-inducing approach is to use a lasso penalty [Tibshirani (1996)]. Nevertheless, different data types call for appropriate penalty terms such that each $\mathbf{W}_t$ is sparse with a specific sparsity structure. In particular, copy number aberrations tend to occur in contiguous regions along chromosomal positions (Figure 6), for which the fused lasso penalty [Tibshirani et al. (2005)] is appropriate. In gene expression data where groups of genes involved in the same biological pathway are co-regulated and thus highly correlated in their expression levels, the elastic net penalty [Zou and Hastie (2005)] is useful to encourage a grouping effect by selecting strongly correlated features together. In this paper, we present a sparse iCluster framework that employs different penalty terms for the estimation of $\mathbf{W}_t$ associated with different data types.

In Section 3 we present the methodological details of the latent variable regression combined with lasso, elastic net and fused lasso penalty terms. To determine the optimal combination of the penalty parameter values, a very large search space needs to be covered, which presents a computational challenge. An exhaustive grid search is ineffective. We use a uniform design by Fang and Wang (1994) that seeks "experimental" points that scattered uniformly across the search domain has superior convergence rates over the conventional grid search (Section 3.3). Section 4 presents an EM algorithm for maximizing the penalized data log-likelihood. The number of clusters K is unknown and must be estimated. Section 5 discusses the estimation of K based on a cross-validation approach. Section 6 presents results from simulation studies. Section 7 presents results from real data applications. In particular, Section 7.1 presents an integrative analysis of epigenomic and transcriptomic profiling data using a breast cancer data set [Holm et al. (2010)]. In Section 7.2 we illustrate our proposed method to construct a genome-wide portrait of copy number induced gene expression changes using a lung cancer data set [Chitale et al. (2009)]. We conclude the paper with a brief summary in Section 8.

2. Motivating example. In this section we show an example where an integrated analysis of multiple omics data sets is far more insightful than separate analyses. Pollack et al. (2002) used customized microarrays to generate measurements of DNA copy number and mRNA expression in parallel for 37 primary breast cancer and 4 breast cancer cell line samples. Here the number of data types $T = 2$. In the mRNA expression data matrix $\mathbf{X}_1$, the individual element x_{ij1} refers to the observed expression of the ith gene in the jth tumor. In the DNA copy number data matrix $\mathbf{X}_2$, the individual element x_{ij2} refers to the observed log-ratio of tumor versus normal copy number of the ith gene in the jth tumor. In this example, both data types have gene-centric measurement by design.

A heatmap of the genomic features on chromosome 17 is plotted in Figure 1. In the heatmap, rows are genes ordered by their genomic position and columns are samples ordered by hierarchical clustering [panels (A)] or by the lasso iCluster method [panels (B)]. There are two main subclasses in the 41 samples: the cell line subclass (samples labeled in red) and the *HER2* tumor subclass (samples labeled in green). It is clear in Figure 1(A) that these subclasses cannot be distinguished well from separate hierarchical clustering analyses.

Separate clustering followed by manual integration as depicted in Figure 1(A) remains the most frequently applied approach to analyze multiple omics data sets in the current literature due to its simplicity and the lack of a truly integrative approach. However, Figure 1(A) clearly shows its lack of a unified system for cluster assignment and poor correlation of the outcome with biological and clinical annotation. As we will illustrate in the simulation study in Section 7, separate clustering can fail drastically in estimating the true number of clusters, classifying samples to the correct clusters and selecting cluster-associated features. Several limitations of this common approach are responsible for its poor performance:

- Correlation between data sets is not utilized to inform the clustering analysis, ignoring an important piece of information that plays a key role for identifying "driver" features of biological importance.
- Separate analysis of *paired* genomic data sets is an inefficient use of the available information.
- It is not straightforward to integrate the multiple sets of cluster assignments that are data-type dependent without extensive prior biological information.
- The standard clustering method includes all genomic features regardless of their relevance to clustering.

Our method aims to overcome these obstacles by formulating a joint analysis across multiple omics data sets. The heatmap in Figure 1(B) demonstrates the superiority of our working model in correctly identifying the subgroups (vertically divided by solid black lines). From left to right, cluster 1 (samples labeled in red) corresponds to the breast cancer cell line subgroup, distinguishing cell line samples from tumor samples. Cluster 2 corresponds to the *HER2* tumor subtype (samples
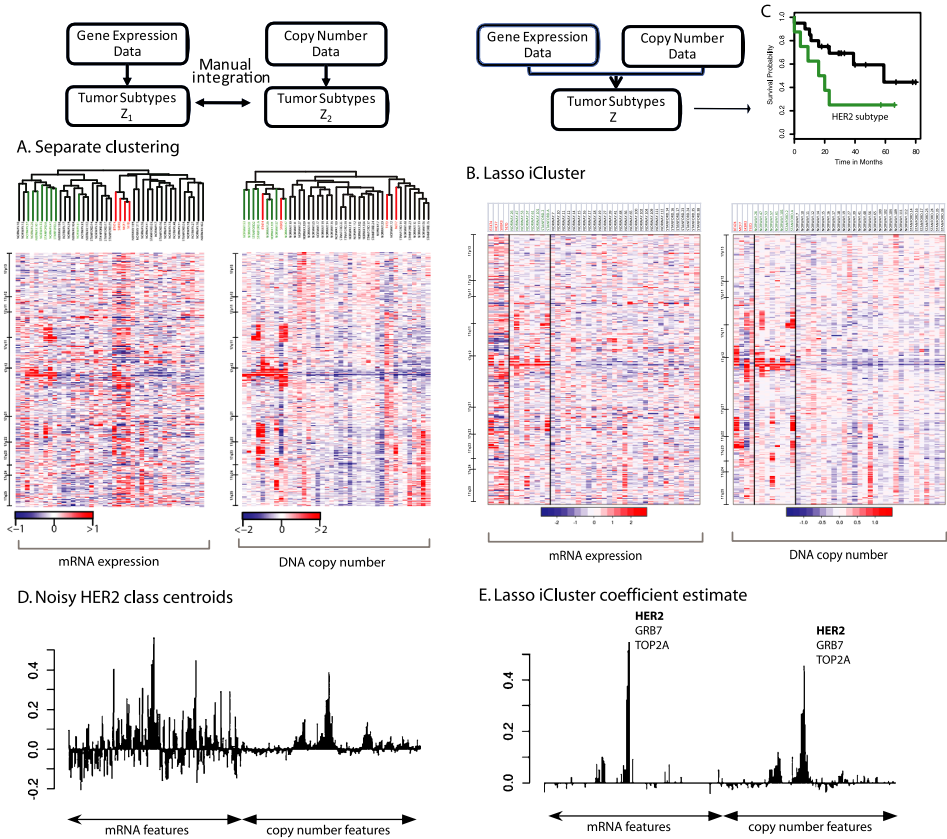
FIG. 1.   *A motivating example using the Pollack data set to demonstrate that a joint analysis using the lasso iCluster method outperforms the separate clustering approach in subtype analysis given DNA copy number and mRNA expression data.* (A) *Heatmap with samples ordered by separate hier-archical clustering. Rows are genes and samples are columns. Samples labeled in red are breast can-cer cell line samples. Samples labeled in green are HER2 breast tumors.* (B) *Heatmap with samples ordered by integrative clustering using the lasso iCluster method.* (C) *Kaplan–Meier plot indicates the HER2 subtype has poor survival outcome.* (D) *Standard cluster centroid estimates.* (E) *Sparse coefficient estimates under the lasso iCluster model.*

labeled in green), showing concordant amplification in the DNA and overexpression in mRNA at the *HER2* locus (chr 17q12). This subtype is associated with poor survival as shown in Figure 1(C). Cluster 3 (samples labeled in black) did not show any distinct patterns, though a pattern may have emerged if there were additional data types such as DNA methylation.

The motivation for sparseness in the coefficient estimates is illustrated by Figure 1(E). It clearly reveals the *HER2*-subtype specific genes (including *HER2*, *GRB7*, *TOP2A*). By contrast, the standard cluster centroid estimation is flooded

with noise [Figure 1(D)], revealing an inherent problem with clustering methods without regularization.

The copy number data example in Figure 1 depicts a narrow (focal) DNA amplification event on a single chromosome involving only a few genes (including *HER2*). Nevertheless, copy number is more frequently altered across long contiguous regions. In the lung cancer data example we will discuss in Section 7.2, chromosome arm-level copy number gains (log-ratio > 0) and losses (log-ratio < 0) as illustrated in Figure 6 are frequently observed, motivating the use of a fused lasso penalty to account for such structural dependencies. In the next section we discuss methodological details on lasso, fused lasso and elastic net in the latent variable regression.

3. Method. Assuming Gaussian error terms, equation (1) implies the following conditional distribution

$$(2) \qquad \mathbf{X}_t | \mathbf{Z} \sim N(\mathbf{W}_t \mathbf{Z}, \boldsymbol{\Psi}_t), \qquad t = 1, \dots, T.$$

Further assuming $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{I})$, the marginal distribution for the observed data is then

$$(3) \qquad \mathbf{X}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma}_t),$$

where $\boldsymbol{\Sigma}_t = \mathbf{W}_t \mathbf{W}_t' + \boldsymbol{\Psi}_t$. Direct maximization of the marginal data log-likelihood is difficult. We consider an expectation–maximization (EM) algorithm [Dempster, Laird and Rubin (1977)]. In the EM framework, the latent variables are considered "missing data." Therefore, the "complete" data log-likelihood that consists of these latent variables is

$$(4) \qquad \begin{aligned} \ell_c = {} & -\frac{n}{2} \sum_{t=1}^{T} \log |\boldsymbol{\Psi}_t| - \frac{1}{2} \sum_{t=1}^{T} \operatorname{tr}\big((\mathbf{X}_t - \mathbf{W}_t \mathbf{Z})' \boldsymbol{\Psi}_t^{-1} (\mathbf{X}_t - \mathbf{W}_t \mathbf{Z})\big) \\ & - \frac{1}{2} \operatorname{tr}(\mathbf{Z}' \mathbf{Z}). \end{aligned}$$

The constant term in ℓ_c has been omitted. In the next section we discuss a penalized complete data log-likelihood to induce sparsity in $\mathbf{W}_t$.

3.1. *Penalized likelihood approach.* As mentioned earlier, sparsity in $\mathbf{W}_t$ directly impacts the interpretability of the latent variables. A zero entry in the ith row and kth column ($w_{ikt} = 0$) means that the ith genomic feature has no weight on the kth latent variable in data type t. If the entire row $\mathbf{w}_{it} = 0$, then this genomic feature has no contribution to the latent variables and is considered noninformative. We use a penalized complete-data log-likelihood as follows to enforce desired sparsity in the estimated $\mathbf{W}_t$:

$$(5) \qquad \ell_{c,p}(\{\mathbf{W}_t\}_{t=1}^{T}, \{\boldsymbol{\Psi}\}_{t=1}^{T}) = \ell_c - \sum_{t=1}^{T} J_{\lambda_t}(\mathbf{W}_t),$$

where ℓ_c is the complete-data log-likelihood function defined in (4) which controls the fitness of the model; $J_{\lambda_t}(\mathbf{W}_t)$ is a penalty function which controls the complexity of the model; and λ_t is a nonnegative tuning parameter that determines the balance between the two. The subscript p in $\ell_{c,p}$ stands for penalized.

Different data types call for different penalty functions. We introduce three types of penalties in the iCluster model: lasso, elastic net, and fused lasso. Both lasso and elastic net regression methods have been applied to gene expression data [Zhao and Simon (2010), Barretina et al. (2012)]. For feature selection, the elastic net may have an additional advantage by shrinking coefficients of correlated features toward each other, and thus encourages a grouping effect toward selecting highly correlated features together. Copy number aberrations tend to occur in contiguous regions along chromosomal positions, motivating the use of fused lasso.

3.1.1. *The lasso penalty.* The lasso penalty is a basic sparsity-inducing that takes the form

$$(6) \qquad J_{\lambda_t}(\mathbf{W}_t) = \lambda_t \sum_{k=1}^{K-1} \sum_{i=1}^{p_t} |w_{ikt}|,$$

where w_{ikt} is the element in the ith row and kth column of $\mathbf{W}_t$. The ℓ_1-penalty continuously shrinks the coefficients toward zero and thereby yields a substantial decrease in the variance of the coefficient estimates. Owing to the singularity of the ℓ_1-penalty at the origin ($w_{ikt} = 0$), some estimated $\hat{w}_{ikt}$ will be *exactly* zero. The degree of sparseness is controlled by the tuning parameter λ_t.

3.1.2. *The fused lasso penalty.* To account for the strong spatial dependence along genomic ordering typical in DNA copy number data, we consider the fused lasso penalty [Tibshirani et al. (2005)], which takes the following form:

$$(7) \qquad J_{\lambda_t}(\mathbf{W}_t) = \lambda_{1t} \sum_{k=1}^{K-1} \sum_{i=1}^{p_t} |w_{ikt}| + \lambda_{2t} \sum_{k=1}^{K-1} \sum_{i=2}^{p_t} |w_{ikt} - w_{(i-1)kt}|,$$

where λ_{1t} and λ_{2t} are two nonnegative tuning parameters. The first penalty encourages sparseness while the second encourages smoothness along index i. The fused lasso penalty is particularly suitable for DNA copy number data where contiguous regions of a chromosome tend to be altered in the same fashion [Tibshirani and Wang (2008)].

3.1.3. *The elastic net penalty.* The elastic net penalty [Zou and Hastie (2005)] takes the form

$$(8) \qquad J_{\lambda_t}(\mathbf{W}_t) = \lambda_{1t} \sum_{k=1}^{K-1} \sum_{i=1}^{p_t} |w_{ikt}| + \lambda_{2t} \sum_{k=1}^{K-1} \sum_{i=1}^{p_t} w_{ikt}^2,$$
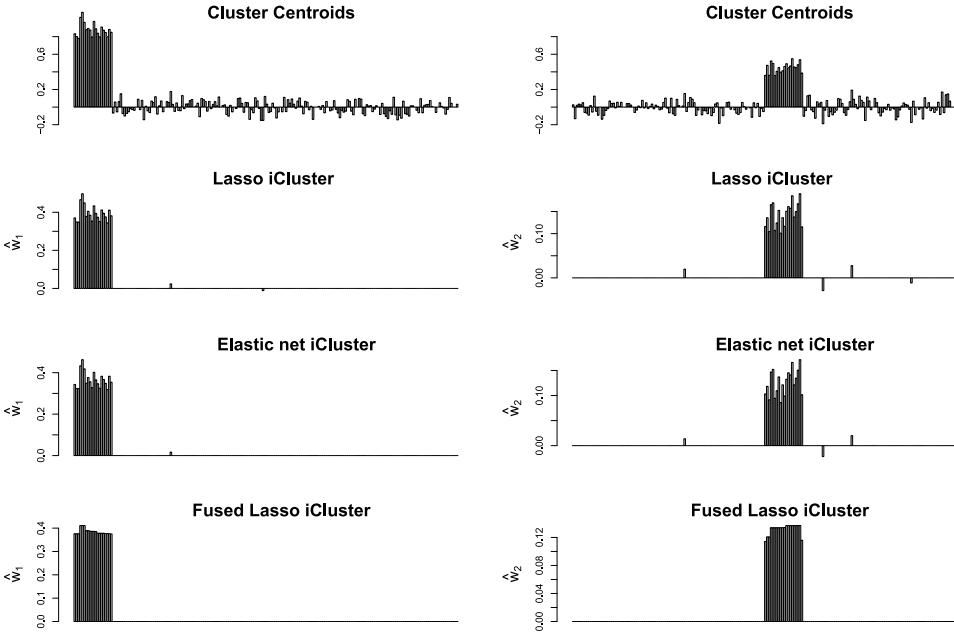
FIG. 2. *A simulated pair of data sets each with* 100 *subjects* ($n = 100$) *and* 200 *features* ($p_t = 200, t = 1, 2$), *and* 2 *subgroups* ($K = 2$). *Top panel plots the cluster centroids in data set* 1 (*left*) *and in data set* 2 (*right*). *Estimated sparse iCluster coefficients are plotted below.*

where λ_{1t} and λ_{2t} are two nonnegative tuning parameters. Zou and Hastie (2005) showed that the elastic net penalty tends to select or remove highly correlated predictors together in a linear regression setting by enforcing their estimated co-efficients to be similar. In our experience, the elastic net penalty tends to be more numerically stable than the lasso penalty in our model.

Figure 2 shows the effectiveness of sparse iCluster using a simulated pair of data sets ($T = 2$). We simulated a single length-n latent variable $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})$ where $n = 100$. The coefficient matrix $\mathbf{W}_1$ consists of a single column $\mathbf{w}_1$ of length $p_1 = 200$ with the first 20 elements set to 1.5 and the remaining elements set to 0, that is, $w_{i1} = 1.5$ for $i = 1, \ldots, 20$ and 0 elsewhere. The coefficient matrix $\mathbf{W}_2$ consists of a single column $\mathbf{w}_2$ of length $p_2 = 200$ and set to have $w_{i2} = 1.5$ for $i = 101, \ldots, 120$ and 0 elsewhere. The lasso, elastic net and fused lasso coefficient estimates are plotted to contrast the noisy cluster centroids estimated separately in data type 1 (left) and in data type 2 (right) in the top panel of Figure 2. The algo-rithm for computing these sparse estimates will be discussed in Section 4.

3.2. *Relationship to singular value decomposition* (*SVD*). An SVD/PCA on the concatenated data matrix $\mathbf{X} = (\mathbf{X}'_1, \ldots, \mathbf{X}'_T)'$ is a special case of equation (1) that requires a common covariance matrix across data types. Specifically, it can be shown that when $\mathbf{\Psi}_1 = \cdots = \mathbf{\Psi}_T = \sigma^2 \mathbf{I}$, equation (1) reduces to a "probabilistic

SVD/PCA" on the concatenated data matrix $\mathbf{X}$. Following similar derivation in Tipping and Bishop (1999), the maximum likelihood estimates of $\mathbf{W}$, where $\mathbf{W} = (\mathbf{W}'_1, \ldots, \mathbf{W}'_T)'$ is the concatenated coefficient matrix, coincide with the first $K - 1$ eigenvectors of the sample covariance matrix $\mathbf{XX}'$ or the right singular vector of the concatenated data matrix $\mathbf{X}$. The MLE of σ^2 is the average of the remaining $n - K + 1$ eigenvalues, capturing the residual variation averaged over the "lost" dimensions.

The major assumption is the requirement that all features have the same variance. The genomic data types, however, are fundamentally different and the method we propose primarily aims to deal with heteroscedasticity among genomic features of various types. The common covariance assumption that leads to SVD is therefore not suitable for integrating omics data types. It is worth mentioning that feature scaling may not necessarily yield $\sigma_1^2 = \cdots = \sigma_{p_t}^2$. In our modeling framework, σ_i^2 is the conditional variance of x_{ij} given $\mathbf{z}_j$. Standardization on x_{ij} will yield the same marginal variance across features, but the conditional variances of features are not necessarily the same after standardization.

Our method aims to identify common influences across data types through the latent component $\mathbf{Z}$. The independent error terms $\mathbf{E}_t, t = 1, \ldots, T$ capture the remaining variances unique to each data type after accounting for the common variance. In SVD, however, the unique variances are absorbed in the term $\mathbf{WZ}$ by enforcing $\boldsymbol{\Psi}_1 = \cdots = \boldsymbol{\Psi}_T = \sigma^2 \mathbf{I}$. As a result, common and unique variations are no longer separable. This is in fact one of the fundamental differences between the factor analysis model and PCA, which has practical importance in integrative modeling.

In Sections 6 and 7 we illustrate that SVD on the concatenated data matrix broadly fails to achieve an effective integration in both simulated and real data sets. By contrast, our method can more effectively deal with heteroscedasticity among genomic features of various types. The contrast with a sparse SVD method lies in the fact that our framework allows block-wise sparse constraints to the coefficient matrix.

3.3. *Uniform sampling.* An exhaustive grid search for the optimal combination of the penalty parameters that maximizes a certain criterion (the optimization criterion will be discussed in Section 5) is inefficient and computationally prohibitive. We use the uniform design (UD) of Fang and Wang (1994) to generate good lattice points from the search domain, a similar strategy adopted by Wang et al. (2008). A key theoretical advantage of UD over the traditional grid search is the uniform space filling property that avoids wasteful computation at close-by points. Let D be the search region. Using the concept of discrepancy that measures uniformity on $D \subset R^d$ with arbitrary dimension d, which is basically the Kolmogorov statistic for a uniform distribution on D, Fang and Wang (1994) point out that the discrepancy of the good lattice point set from a uniform design converges to zero with a rate of $O(n^{-1}(\log n)^d)$, where n (a prime number) denotes

the number of generated points on D. They also point out that the sequence of equi-lattice points on D has a rate of $O(n^{-1/d})$ and the sequence of uniformly distributed random numbers on D has a rate of $O(n^{-1/2}(\log\log n)^{1/2})$. Thus, the uniform design has an optimal rate for $d \geq 2$.

4. Algorithm. We now discuss the details of our algorithm for parameter estimation in sparse iCluster. The latent variables (columns of $\mathbf{Z}$) are considered to be "missing" data. The algorithm therefore iterates between an E-step for imputing $\mathbf{Z}$ and a penalized maximization step (M-step) that updates the estimates of $\mathbf{W}_t$ and $\mathbf{\Psi}_t$ for all t. Given the latent variables, the data types are conditionally independent and, thus, the integrative omics problem can be decomposed into solving T independent subproblems with suitable penalty terms. The penalized estimation procedures are therefore "decoupled" for each data type given the latent variables $\mathbf{Z}$. When convergence is reached, cluster membership will be assigned for each tumor based on the posterior mean of the latent variable $\mathbf{Z}$.

E-step. In the E-step, we take the expectation of the penalized complete-data log-likelihood $\ell_{c,p}$ as defined in equations (4) and (5), which primarily involves computing two conditional expectations given the current parameter estimates:

$$(9) \qquad E[\mathbf{Z}|\mathbf{X}] = \mathbf{W}'\mathbf{\Sigma}^{-1}\mathbf{X}$$

$$(10) \qquad E[\mathbf{Z}\mathbf{Z}'|\mathbf{X}] = \mathbf{I} - \mathbf{W}'\mathbf{\Sigma}^{-1}\mathbf{W} + E[\mathbf{Z}|\mathbf{X}]E[\mathbf{Z}|\mathbf{X}]',$$

where $\mathbf{\Sigma} = \mathbf{W}\mathbf{W}' + \mathbf{\Psi}$ and $\mathbf{\Psi} = \mathrm{diag}(\mathbf{\Psi}_1, \ldots, \mathbf{\Psi}_T)$. Here, the posterior mean in (9) effectively provides a simultaneous rank-$(K-1)$ approximation to the original data matrices $\mathbf{X}$.

M-step. In the M-step, given the quantities in equations (9) and (10), we maximize the penalized complete-data log-likelihood to update the estimates of $\mathbf{W}_t$ and $\mathbf{\Psi}_t$.

(1) *Sparse estimates of* $\mathbf{W}_t$. For $t = 1, \ldots, T$, we obtain the penalized estimates by

$$
\mathbf{W}_t \leftarrow \underset{\mathbf{W}_t}{\mathrm{argmin}}\, \frac{1}{2}\sum_{t=1}^{T} E\big[\mathrm{tr}((\mathbf{X}_t - \mathbf{W}_t\mathbf{Z})'\mathbf{\Psi}_t^{-1}(\mathbf{X}_t - \mathbf{W}_t\mathbf{Z}))|\hat{\mathbf{W}}_t, \hat{\mathbf{\Psi}}_t\big]
$$

(11)

$$
+ J_{\lambda_t}(\mathbf{W}_t),
$$

where $\hat{\mathbf{W}}_t$ and $\hat{\mathbf{\Psi}}_t$ denote the parameter estimates in the last EM iteration. We apply a local quadratic approximation [Fan and Li (2001)] to the ℓ_1 term involved in the penalty function $J_{\lambda_t}(\mathbf{W}_t)$. Using the fact $|\alpha| = \alpha^2/|\alpha|$ when $\alpha \neq 0$, we consider the following quadratic approximation to the ℓ_1 term:

$$(12) \qquad \lambda_t \sum_{k=1}^{K-1}\sum_{i=1}^{p_t} \frac{w_{ikt}^2}{|\hat{w}_{ikt}|}.$$

Due to the uncorrelated error terms (diagonal $\boldsymbol{\Psi}_t$) and "noncoupling" structure of the lasso and elastic net penalty terms, the estimation of $\mathbf{W}_t$ can then be computed feature-by-feature by taking derivatives with respect to each row $\mathbf{w}_{it}$ for $i = 1, \ldots, p_t$. The solution for (11) under various penalty terms can then be obtained by iteratively computing the following ridge regression estimates:

(a) *Lasso estimates.* For $i = 1, \ldots, p_t$,

(13)
$$\mathbf{w}_{it} = \mathbf{x}_{it} E[\mathbf{Z}' | \mathbf{X}_t, \hat{\mathbf{W}}_t, \hat{\boldsymbol{\Psi}}_t] (E[\mathbf{Z}\mathbf{Z}' | \mathbf{X}_t, \hat{\mathbf{W}}_t, \hat{\boldsymbol{\Psi}}_t] + \mathbf{A}_i)^{-1},$$

where $\mathbf{A}_i = 2\sigma_i^2 \lambda_t \, \mathrm{diag}\{1/|\hat{w}_{i1t}|, \ldots, 1/|\hat{w}_{i(K-1)t}|\}$. Computing (13) only requires the inversion of a $(K-1) \times (K-1)$ matrix in the latent subspace.

(b) *Elastic net estimates.* Similarly, we consider a quadratic approximation to the ℓ_1 term in the elastic net penalty and obtain the solution for (11) by iteratively computing a ridge regression estimate similar to (13) but with $\mathbf{A}_i = 2\sigma_i^2(\lambda_{1t} \, \mathrm{diag}\{1/|\hat{w}_{i1t}|, \ldots, 1/|\hat{w}_{i(K-1)t}|\} + \lambda_{2t}\mathbf{I})$.

(c) *Fused lasso estimates.* For fused lasso penalty terms, we consider the following approximation:

(14)
$$\lambda_{1t} \sum_{k=1}^{K-1} \sum_{i=1}^{p} \frac{w_{ikt}^2}{|\hat{w}_{ikt}|} + \lambda_{2t} \sum_{k=1}^{K-1} \sum_{i=2}^{p} \frac{(w_{ikt} - w_{(i-1)kt})^2}{|\hat{w}_{ikt} - \hat{w}_{(i-1)kt}|}.$$

In the fused lasso scenario, the parameters are coupled together and the estimation of $\mathbf{w}_i$ are no longer separable. However, we circumvent the problem by expressing the estimating equation in terms of a vectorized form $\tilde{\mathbf{w}}_t = \mathrm{vec}(\mathbf{W}_t') = (\mathbf{w}_1, \ldots, \mathbf{w}_{K-1})'$, a column vector of dimension $s = p_t \cdot (K-1)$ by concatenating the columns of $\mathbf{W}_t$. Then (14) can be expressed in the following form:

$$\lambda_{1t} \tilde{\mathbf{w}}_t' \mathbf{A} \tilde{\mathbf{w}}_t + \lambda_{2t} \tilde{\mathbf{w}}_t' \mathbf{L} \tilde{\mathbf{w}}_t,$$

where

$\mathbf{A} = \mathrm{diag}\{1/|\hat{w}_1|, \ldots, 1/|\hat{w}_s|\},$

$\mathbf{L} = \mathbf{D} - \mathbf{M},$

$\mathbf{M} = \begin{cases} 1/|\hat{w}_i - \hat{w}_j|, & |i - j| = K - 1, \\ 0, & \text{otherwise} \end{cases}$ ($s \times s$ dimension),

$\mathbf{D} = \mathrm{diag}\{d_1, \ldots, d_s\}$ where d_j is the summation of the jth row of $\mathbf{M}$.

Letting $\mathbf{C} = \mathbf{X}_t E[\mathbf{Z}' | \mathbf{X}_t, \hat{\mathbf{W}}_t, \hat{\boldsymbol{\Psi}}_t]$ and $\mathbf{Q} = E[\mathbf{Z}\mathbf{Z}' | \mathbf{X}_t, \hat{\mathbf{W}}_t, \hat{\boldsymbol{\Psi}}_t]$, the corresponding estimating equation is then

(15)
$$\frac{\partial}{\partial \tilde{\mathbf{w}}} J(\tilde{\mathbf{w}}) + \tilde{\mathbf{Q}}\tilde{\mathbf{w}} = \tilde{\mathbf{C}},$$

where

(16)
$$\tilde{\mathbf{Q}} = \begin{pmatrix} \sigma_1^{-2}\mathbf{Q} & & \\ & \ddots & \\ & & \sigma_{p_t}^{-2}\mathbf{Q} \end{pmatrix}, \qquad \tilde{\mathbf{C}} = \begin{pmatrix} \sigma_1^{-2}\mathbf{c}_1' \\ \vdots \\ \sigma_{p_t}^{-2}\mathbf{c}_{p_t}' \end{pmatrix},$$

where $\mathbf{c}_j$ is the jth row of $\mathbf{C}$. The solution for (11) under the fused lasso penalty is then computed by iteratively computing

(17)
$$\tilde{\mathbf{w}}_t = (\tilde{\mathbf{Q}} + 2\lambda_{1t}\mathbf{A} + 2\lambda_{2t}\mathbf{L})^{-1}\tilde{\mathbf{C}}.$$

(2) *Estimates of* $\boldsymbol{\Psi}_t$. Finally, for $t = 1, \ldots, T$, we update $\boldsymbol{\Psi}_t$ in the M-step as follows:

(18)
$$\boldsymbol{\Psi}_t = \frac{1}{n} \operatorname{diag}\big(\mathbf{X}_t\mathbf{X}_t' - \hat{\mathbf{W}}_t E\big[\mathbf{Z}|\{\mathbf{X}_t\}_{t=1}^T, \{\hat{\mathbf{W}}_t\}_{t=1}^T, \{\hat{\boldsymbol{\Psi}}_t\}_{t=1}^T\big]\mathbf{X}_t'\big).$$

The algorithm iterates between the E-step and the M-step as described above until convergence. Cluster membership will then be assigned by applying a standard K-means clustering on the posterior mean $E[\mathbf{Z}|\mathbf{X}]$. In other words, cluster partition in the final step is performed in the integrated latent variable subspace of dimension $n \times (K - 1)$. Applying k-means on latent variables to obtain discrete cluster assignment is commonly used in spectral clustering methods [Ng, Jordan and Weiss (2002), Rohe, Chatterjee and Yu (2010)].

5. Choice of tuning parameters. We use a resampling-based criterion for selecting the penalty parameters and the number of clusters. The procedure entails repeatedly partitioning the data set into a learning and a test set. In each iteration, sparse iCluster (for a given K and tuning parameter values) will be applied to the learning set to obtain a classifier and subsequently predict the cluster membership for the test set samples. In particular, we first obtain parameter estimates from the learning set. For new observations in the test data $\mathbf{X}^*$, we then compute the posterior mean of the latent variables $E[\mathbf{Z}|\mathbf{X}^*] = \hat{\mathbf{W}}_\ell' \hat{\boldsymbol{\Sigma}}_\ell^{-1} \mathbf{X}^*$, where $\hat{\mathbf{W}}_\ell, \hat{\boldsymbol{\Sigma}}_\ell^{-1}$ denote parameter estimates from the learning set. A K-means clustering is then applied to $E[\mathbf{Z}|\mathbf{X}^*]$ to partition the test set samples into K clusters. Denote this as partition C_1. In parallel, the procedure applies an independent sparse iCluster with the same penalty parameter values to the test set to obtain a second partition C_2, giving the "observed" test sample cluster labels. Under the true model, the predicted C_1 and the "observed" C_2 (regarded as the "truth") would have good agreement by measures such as the adjusted Rand index. We therefore define a reproducibility index (RI) as the median adjusted Rand index across all repetitions. Values of RI close to 1 indicate perfect cluster reproducibility and values of RI close to 0 indicate poor cluster reproducibility. In this framework, the concepts of bias, variance and prediction error that typically apply to classification analysis where the true cluster labels are known now become relevant for clustering. The idea is similar to the "Clest" method proposed by Dudoit

and Fridlyand (2002), the prediction strength measure proposed by Tibshirani and Walther (2005) and the in-group proportion (IGP) proposed by Kapp and Tibshirani (2007).

6. Simulation. In this section we present results from two simulation studies. In the first simulation setup, we simulate a single length-n latent variable $\mathbf{z} \sim N(0, 1)$ where $n = 100$. Subject j, $j = 1, \ldots, n$, belongs to cluster 1 if $z_j > 0$ and cluster 2 otherwise. For simplicity, the pair of coefficient matrices $(\mathbf{W}_1, \mathbf{W}_2)$ are of the same dimension 200×1 ($p_1 = p_2 = 200$), with $w_{it} = 3$ for $i = 1, \ldots, 20$ for both data types ($t = 1, 2$) and zero elsewhere. Next we obtain the data matrices $(\mathbf{X}_1, \mathbf{X}_2)$ with each element generated according to equation (1) with standard normal error terms. This simulation represents a scenario where an effective joint analysis of two data sets should be expected to enhance the signal strength and thus improve clustering performance.

Table 1 summarizes the performances of each method in terms of the ability to choose the correct number of clusters, cross-validated error rates and cluster reproducibility. In Table 1 separate K-means methods perform poorly in terms of the ability to choose the correct number of clusters, cluster reproducibility and the cross-validation error rates (with respect to the true simulated cluster membership). K-means on concatenated data performs even worse, likely due to noise accumulation. For sparse SVD, a cluster assignment step is needed. We took a similar approach of applying K-means on the first $K - 1$ right singular vectors of the data matrix. Sparse SVD performs better than simple K-means, though data

TABLE 1
Clustering performance summarized over 50 simulated data sets under setup 1 ($K = 2$). Separate clustering methods have two sets of numbers associated with model fit to each individual data type. Numbers in parentheses are the standard deviations over 50 simulations

Method	Percent of times choosing the correct K	Cross-validation error rate	Cluster reproducibility
Separate K-means	58	0.08 (0.04)	0.67 (0.17)
	62	0.08 (0.04)	0.70 (0.19)
Concatenated K-means	50	0.06 (0.04)	0.66 (0.19)
Separate sparse SVD	74	0.07 (0.06)	0.71 (0.13)
	76	0.07 (0.07)	0.72 (0.12)
Concatenated sparse SVD	78	0.07 (0.08)	0.70 (0.12)
Separate AHP-GMM	38	0.06 (0.04)	0.72 (0.15)
	40	0.05 (0.04)	0.74 (0.14)
Concatenated AHP-GMM	46	0.06 (0.04)	0.75 (0.13)
Lasso iCluster	90	0.04 (0.02)	0.81 (0.08)
Enet iCluster	94	0.03 (0.02)	0.85 (0.07)
Fused lasso iCluster	94	0.03 (0.02)	0.83 (0.08)

TABLE 2
Feature selection performance summarized over 50 simulated data sets for K = 2. There are a total of 20 true features simulated to distinguish the two sample clusters

Method	Data 1		Data 2	
	True positives	False positives	True positives	False positives
Separate K-means	–	–	–	–
Concatenated K-means	–	–	–	–
Separate sparse SVD	18.7 (3.2)	21.5 (37.7)	18.8 (2.9)	27.4 (43.6)
Concatenated sparse SVD	14.0 (5.3)	22.5 (16.1)	13.7 (5.2)	22.8 (16.4)
Separate AHP-GMM	19.6 (2.1)	0.02 (0.16)	19.1 (3.1)	0 (0)
Concatenated AHP-GMM	18.8 (3.6)	0.02 (0.15)	18.6 (4.0)	0.02 (0.15)
Lasso iCluster	20 (0)	0.07 (0.3)	20 (0)	0.07 (0.3)
Enet iCluster	20 (0)	0.1 (0.3)	20 (0)	0.02 (0.1)
Fused lasso iCluster	20 (0)	0 (0)	20 (0)	0 (0)

concatenation does not seem to offer much advantage. In this simulation scenario, AHP-GMM models show good performance in feature selection (Table 2), but appear to have a low frequency of choosing the correct $K = 2$. A common theme in this simulation is that a data concatenation approach is generally ineffective regardless of the clustering methods used. By contrast, sparse iCluster methods achieved an effective integrative outcome across all performance criteria.

Table 2 summarizes the associated feature selection performance. No numbers are shown for the standard K-means methods, as they do not have an inherent feature selection method. Among the methods, sparse iCluster methods perform the best in identifying the true positive features while keeping the number of false positives close to 0.

In the second simulation, we vary the setup as follows. We simulate 150 subjects belonging to three clusters ($K = 3$). Subjects $j = 1, \ldots, 50$ belong to cluster 1, subjects $j = 51, \ldots, 100$ belong to cluster 2, and subjects $j = 101, \ldots, 150$ belong to cluster 3. A total of $T = 2$ data types ($\mathbf{X}_1, \mathbf{X}_2$) are simulated. Each has $p_1 = p_2 = 500$ features. Here each data type alone only defines two clusters out of the three. In data set 1, $x_{ij1} \sim N(2, 1)$ for $i = 1, \ldots, 10$ and $j = 1, \ldots, 50$, $x_{ij1} \sim N(1.5, 1)$ for $i = 491, \ldots, 500$ and $j = 51, \ldots, 100$, and $x_{ij1} \sim N(0, 1)$ for the rest. In data set 2, $x_{ij2} = 0.5 * x_{ij1} + e$ where $e \sim N(0, 1)$ for $i = 1, \ldots, 10$ and $j = 1, \ldots, 50$, $x_{ij2} \sim N(2, 1)$ for $i = 491, \ldots, 500$ and $j = 101, \ldots, 150$, and $x_{ij2} \sim N(0, 1)$ for the rest. The first 10 features are correlated between the two data types. In Tables 3 and 4, the sparse iCluster methods consistently outperform the other methods in clustering and feature selection.

The core iCluster EM iterations are implemented in C. Table 5 shows some typical computation times for problems of various dimensions on a 3.2 GHz Xeon Linux computer.

TABLE 3
Clustering performance summarized over 50 *simulated data sets under setup* 2 *($K = 3$)*

Method	Frequency of choosing the correct K	Cross-validation error rate	Cluster reproducibility
Separate K-means	2	0.33 (0.001)	0.54 (0.07)
	0	0.33 (0.002)	0.47 (0.04)
Concatenated K-means	100	0.01 (0.07)	0.96 (0.03)
Separate sparse SVD	0	0.28 (0.10)	0.45 (0.03)
	0	0.31 (0.07)	0.44 (0.04)
Concatenated sparse SVD	16	0.01 (0.002)	0.59 (0.05)
Separate AHP-GMM	0	0.07 (0.13)	0.63 (0.05)
	0	0.32 (0.02)	0.54 (0.06)
Concatenated AHP-GMM	100	0.01 (0.07)	0.98 (0.03)
Lasso iCluster	100	0.0003 (0.001)	0.98 (0.01)
Enet iCluster	100	0.0003 (0.001)	0.97 (0.02)
Fused lasso iCluster	100	0 (0)	0.94 (0.05)

7. Results. In this section we present details of two real data applications.

7.1. *Integration of epigenomic and transcriptomic profiling data in the Holm breast cancer study.* In Section 2 we discussed a motivating example using the Pollack et al. (2002) data set. In this section we present our first real data application which involves integrative analysis of DNA methylation and gene expression data from the Holm et al. (2010) study. In this data set, methylation profiling in 189 breast cancer samples using Illumina methylation arrays for 1452 CpG sites (cor-

TABLE 4
Feature selection performance summarized over 50 *simulated data sets under $K = 3$*

Method	Data 1		Data 2	
	True positives	False positives	True positives	False positives
Separate K-means	–	–	–	–
Concatenated K-means	–	–	–	–
Separate sparse SVD	19.8 (0.7)	349.6 (167.1)	19.9 (0.3)	347.5 (142.5)
Concatenated sparse SVD	20 (0)	396.6 (128.7)	19.6 (1.6)	395.4 (128.3)
Separate AHP-GMM	15.8 (5.0)	239.9 (245.5)	15.5 (5.5)	269.9 (246)
Concatenated AHP-GMM	19.2 (1.7)	0.33 (0.64)	14.4 (4.0)	0.21 (0.66)
Lasso iCluster	20 (0)	1.5 (1.4)	19.9 (0.2)	1.9 (1.5)
Enet iCluster	20 (0)	0.5 (0.6)	19.8 (0.5)	0.7 (1.0)
Fused lasso iCluster	20 (0)	0 (0)	20 (0)	0 (0)

TABLE 5
Computing time (*in seconds*) *for typical runs of sparse iCluster under various dimensions*

		Time (in seconds)		
p	N	Lasso iCluster	Elastic net iCluster	Fused lasso iCluster
200	100	0.10	0.11	0.37
500	100	0.50	0.36	3.56
1000	100	1.40	1.45	25.05
2000	100	6.49	5.90	76.40
5000	100	18.93	18.94	33 (min)

responding to 803 cancer-related genes) is available. The original study performed a hierarchical clustering on the methylation data alone. Through manual integration, the authors then correlated the methylation status with gene expression levels for 511 oligonucleotide probes for genes with CpG sites on the methylation assays in the same sample set. Here we compare clustering of individual data types to various integration approaches. We included the most variable 288 CpG sites (following a similar procedure taken in the Holm study) in the methylation data.

We applied sparse iCluster for a joint analysis of the methylation ($p_1 = 288$) and gene expression ($p_2 = 511$) data using different penalty combinations. In Figure 3(A) the first two latent variables separated the samples into three distinct clusters. By associating the cluster membership with clinical variables, it becomes clear that tumors in cluster 1 are predominantly estrogen receptor (ER)-negative and associated with the basal-like breast cancer subtype (Figure 4). Among the rest of the samples, sparse iCluster further identifies a subclass (cluster 3) that highly expresses platelet-derived growth factor receptors (*PDGFRA/B*), which have been associated with breast cancer progression [Carvalho et al. (2005)].

In Section 3.2 we discussed an SVD approach on a combined data matrix as a special case of our model. Here we present results from SVD and a sparse SVD algorithm proposed by Witten, Tibshirani and Hastie (2009) on the concatenated data matrix. Figures 3(B) and 3(C) indicate that SVD applied to each data type alone can only separate one out of the three clusters. Figures 3(D) and 3(E) indicate that data concatenation does not perform any better in this analysis than separate analyses of each data type alone.

In Table 6 the results from sparse iCluster with two different sets of penalty combinations are presented: the combination of (lasso, lasso) and the combination of (lasso, elastic net) for methylation and gene expression data, respectively (Table 6 top panel). The reproducibility index (RI) is computed for various K's and penalty parameters are sampled based on a uniform design described in Section 3.3. As described in Section 5, RI (ranges between 0 and 1) measures the agreement between the predicted cluster membership and the "observed" cluster membership using a 10-fold cross-validation.
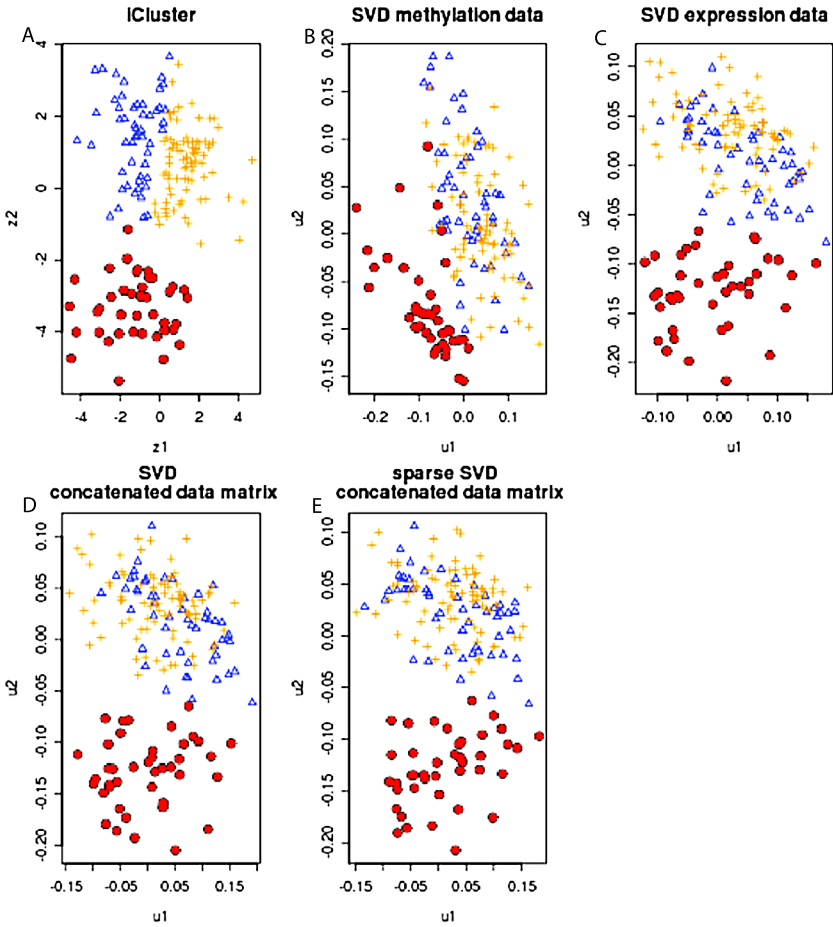
FIG. 3. *Separation of the data points by* (A) *latent variables from sparse iCluster,* (B) *right singular vectors from SVD of the methylation data alone,* (C) *right singular vectors from SVD of the expression data alone,* (D) *SVD on the concatenated data matrix, and* (E) *sparse SVD on the concatenated data matrix. Red dots indicate samples belonging to cluster* 1, *blue open triangles indicate samples belonging to cluster* 2, *and orange pluses indicate samples belonging to cluster* 3.
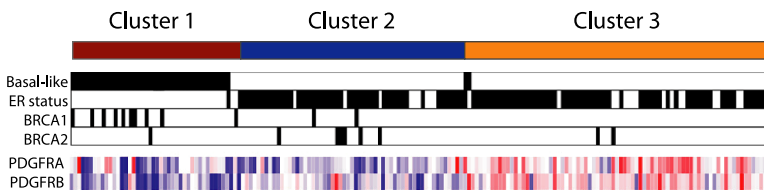

FIG. 4. *Integrative clustering of the Holm study DNA methylation and gene expression data revealed three clusters with a cross-validated reproducibility of* 0.7 *and distinct clinical and molecular characteristics.*

TABLE 6

Cluster reproducibility and number of genomic features selected using sparse iCluster, sparse SVD on concatenated data matrix and Adaptive Hierarchically Penalized Gaussian Mixture Model (AHP-GMM) on concatenated data matrix. Two variations of the sparse iCluster method were presented: iCluster (lasso, lasso) implements lasso penalty for both data types, and iCluster (lasso, elastic net) implements lasso penalty for the methylation data and elastic net penalty for the gene expression data. K: the number of clusters. RI: reproducibility index

K	RI	Selected methylation features	Selected expression features	RI	Selected methylation features	Selected expression features
		iCluster (lasso, lasso)			iCluster (lasso, elastic net)	
2	0.68	138	151	0.70	183	353
3	0.46	150	204	0.70	273	182
4	0.42	183	398	0.48	273	182
5	0.42	205	454	0.47	282	223
		sparse SVD			AHP-GMM	
2	0.78	1	105	0.93	9	63
3	0.34	1	134	0.42	28	105
4	0.27	288	511	0.49	116	368
5	0.22	273	504	0.43	42	243

Both methods identified a 2-cluster solution with an RI around 0.70, distinguishing the ER-negative, Basal-like subtype from the rest of the tumor samples (Figures 3 and 4, samples labeled in red). The iCluster (lasso, elastic net) method adds an ℓ_2 penalty term to encourage grouped selection of highly correlated genes in the expression data. This approach further identified a 3-cluster solution with high reproducibility (RI = 0.70). The additional division finds a subgroup that highly expresses platelet-derived growth factor receptors (Figure 4).

Figure 5 displays heatmaps of the methylation and expression data. Columns are samples ordered by the integrated cluster assignment. Rows are cluster-discriminating genes (with nonzero coefficient estimates) grouped into gene clusters by hierarchical clustering. In total, there are 273 differentially methylated genes and 182 differentially expressed genes. Several cancer genes including *MUC1*, *SERPINA5*, *RARA*, *MECP2* and *RAD50* are hypermethylated and show concordant underexpression in cluster 1. On the other hand, hypomethylation of cancer genes including *ETS1*, *HDAC1*, *FANCE*, *RAB32* and *JAK3* are observed and, correspondingly, these genes show increased expression levels.

To compare with other methods, we implemented the sparse SVD method by Witten, Tibshirani and Hastie (2009) and an adaptive hierarchical penalized Gaussian mixture model (AHP-GMM) by Wang and Zhu (2008) on the concatenated data matrix. None of these methods generated additional insights beyond separating the ER-negative and basal-like tumors from the others (Figure 3 and Table 6).
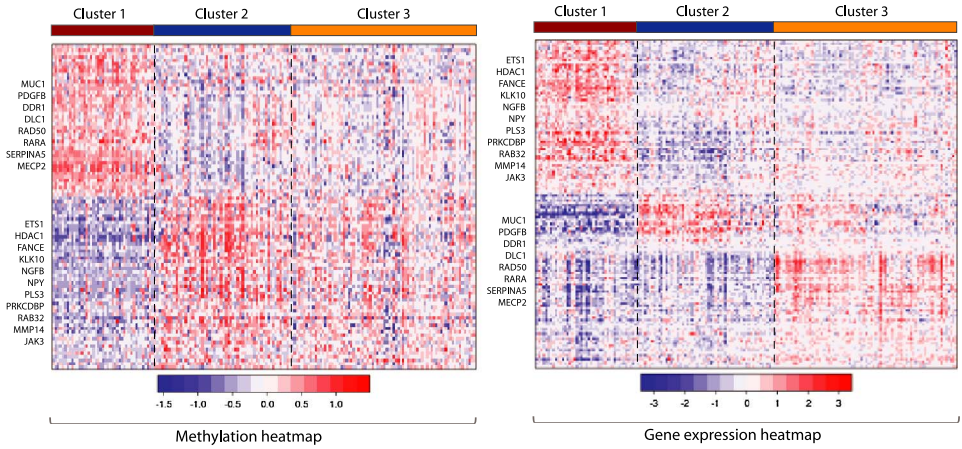
FIG. 5. *Integrative clustering of the Holm study DNA methylation and gene expression data revealed three clusters with a cross-validated reproducibility of* 0.7. *Selected genes with negatively correlated methylation and expression changes are indicated to the left of the heatmap.*

Feature selection is predominantly "biased" toward gene expression features when directly applying sparse SVD on the combined data matrix (bottom panel of Table 6), likely due to the larger between-cluster variances observed in the gene expression data.

7.2. *Constructing a genome-wide portrait of concordant copy number and gene expression pattern in a lung cancer data set.*   We applied the proposed method to integrate DNA copy number (aCGH data) and mRNA expression data in a set of 193 lung adenocarcinoma samples [Chitale et al. (2009)]. Figure 6 displays an example of the probe-level data (log-ratios of tumor versus normal copy number) on chromosomes 3 and 8 in one tumor sample. Many samples in this data set display similar chr 3p whole-arm loss and chr 3q whole-arm gain.
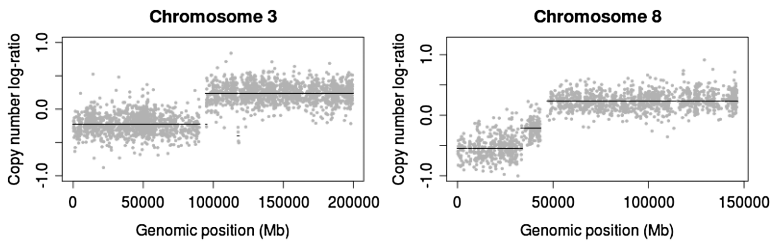

FIG. 6. *Illustration of copy number probe-level data from a lung tumor sample* [Chitale et al. (2009)]. *Log-ratios of copy number* (tumor versus normal) *on chromosomes* 3 *and* 8 *are displayed. Log-ratio greater than zero indicates copy number gain and log-ratio below zero indicates loss. Black line indicates the segmented value using the circular binary segmentation method* [Olshen et al. (2004), Venkatraman and Olshen (2007)].

Arm-length copy number aberrations are surprisingly common in cancer [Beroukhim et al. (2010)], affecting up to thousands of genes within the region of alteration. A broader challenge is thus to pinpoint the "driver" genes that have functional roles in tumor development from those that are functionally neutral ("passengers"). To that end, an integrative analysis with gene expression data could provide additional insights. Genes that show concordant copy number and transcriptional activities are more likely to have functional roles.

In the search for copy number-associated gene expression patterns, we fit a sparse iCluster model for each of the 22 chromosomes using (fused lasso, lasso) a penalty combination for joint analysis of copy number and gene expression data. To facilitate comparison, we compute a 2-cluster solution with a single latent variable vector $\mathbf{z}$ (instead of estimating K) to extract the major pattern for each chromosome. Penalty parameter tuning is performed as described before. In Figure 7 we plot the 22 pairs of the sparse coefficient vectors ordered by chromosomal position. The coefficients can be interpreted as the difference between the two cluster means. Positive and negative coefficient values in Figure 7(A) thus indicate copy number gains and losses in one cluster relative to the other. Similarly, in Figure 7(B), coefficient signs indicate over- or under-expression in one cluster
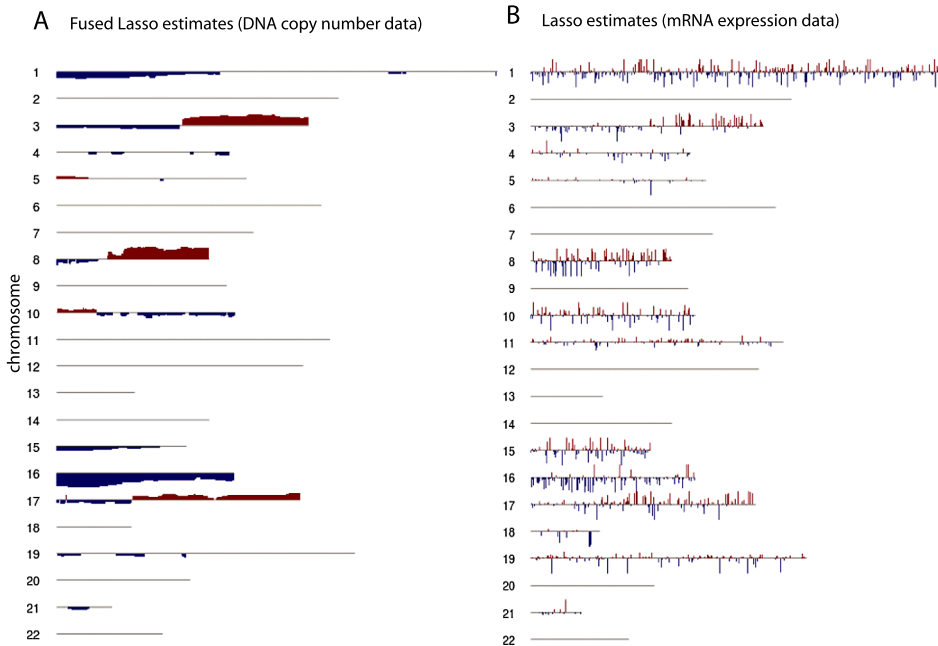

FIG. 7. *Penalized coefficient vector estimates arranged by chromosomes* 1 *to* 22 *derived by iCluster* (*fused lasso*, *lasso*) *applied to the Chitale et al. lung cancer data set. A single latent variable vector is used to identify the major pattern of each chromosome.*

relative to the other. Concordant copy number and gene expression changes can thus be directly visualized from Figure 7.

Several chromosomes (1, 3, 8, 10, 15 and 16) show contiguous regions of gains or losses spanning whole chromosome arms. As discussed before, arm-length aberrations can affect up to thousands of genes within the region of alteration. A great challenge is thus to pinpoint the "driver" genes that have important roles in tumor development from those that are functionally neutral ("passengers"). To that end, an integrative analysis could provide additional insights for identifying potential drivers by revealing genes with concordant copy number and transcriptional activities. Figure 7 shows that the application of the proposed method can unveil a genome-wide pattern of such concordant changes, providing a rapid way for identifying candidate genes of biological significance. Several arm-level copy number alterations (chromosomes 3, 8, 10, 16) exhibit concerted influence on the expression of a small subset of the genes within the broad regions of gains and losses.

8. Discussion. Integrative genomics is a new area of research accelerated by large-scale cancer genome efforts including the Cancer Genome Atlas Project. New integrative analysis methods are emerging in this field. van Wieringen and van de Wiel (2009) proposed a nonparametric testing procedure for DNA copy number induced differential mRNA gene expression. Peng et al. (2010) and Vaske et al. (2010) considered pathway and network analysis using multiple genomic data sources. A number of others [Waaijenborg, Verselewel de Witt Hamer and Zwinderman (2008), Parkhomenko, Tritchler and Beyene (2009), Le Cao, Martin and Robert-Granie (2009), Witten, Tibshirani and Hastie (2009), Witten and Tibshirani (2009), Soneson et al. (2010)] suggested using canonical correlation analysis (CCA) to quantify the correlation between two data sets (e.g., gene expression and copy number data). Most of this previous work focused on integrating copy number and gene expression data, and none of these methods were specifically designed for tumor subtype analysis.

We have formulated a penalized latent variable model for integrating multiple genomic data sources. The latent variables can be interpreted as a set of distinct underlying cancer driving factors that explain the molecular phenotype manifested in the vast landscape of alterations in the cancer genome, epigenome and transcriptome. Lasso, elastic net and fused lasso penalty terms are used to induce sparsity in the feature space. We derived an efficient and unified algorithm. The implementation scales well for increasing data dimension.

A future extension on group-structured penalty terms is to incorporate a grouping structure defined a priori. Two types of group structures are relevant for our application. One is to treat the $w_{i1}, \ldots, w_{i(K-1)}$ as a group since they are associated with the same feature. Yuan and Lin's group lasso penalty [Yuan and Lin (2006)] can be applied directly. Similar to our current algorithm, by using Fan and Li's local quadratic approximation, the problem reduces to a ridge-type regression in each iteration. The other extension is to incorporate the grouping structure

among features to boost the signal to noise ratio, for example, to treat the genes within a pathway as a group. We can consider a hierarchical lasso penalty [Wang et al. (2009)] to achieve sparsity at both the group level and the individual variable level.

REFERENCES

ALIZADEH, A. A., EISEN, M. B., DAVIS, E. E. et al. (2000). Distinct types of diffuse large B-cell lymphoma identified by gene expression profiling. *Nature* **403** 503–511.

BARRETINA, J., CAPONIGRO, G., STRANSKY, N., VENKATESAN, K., MARGOLIN, A. A., KIM, S., WILSON, C. J., LEHAR, J., KRYUKOV, G. V., SONKIN, D., REDDY, A., LIU, M., MURRAY, L., BERGER, M. F., MONAHAN, J. E., MORAIS, P., MELTZER, J., KOREJWA, A., JANE-VALBUENA, J., MAPA, F. A., THIBAULT, J., BRIC-FURLONG, E., RAMAN, P., SHIPWAY, A., ENGELS, I. H., CHENG, J., YU, G. K., YU, J., ASPESI, P., DE SILVA, M., JAGTAP, K., JONES, M. D., WANG, L., HATTON, C., PALESCANDOLO, E., GUPTA, S., MAHAN, S., SOUGNEZ, C., ONOFRIO, R. C., LIEFELD, T., MACCONAILL, L., WINCKLER, W., REICH, M., LI, N., MESIROV, J. P., GABRIEL, S. B., GETZ, G., ARDLIE, K., CHAN, V., MYER, V. E. and WEBER, B. L. (2012). The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature* **483** 603–307.

BEROUKHIM, R., MERMEL, C. H., PORTER, D., WEI, G., RAYCHAUDHURI, S., DONOVAN, J., BARRETINA, J., BOEHM, J. S., DOBSON, J., URASHIMA, M., MC HENRY, K. T., PINCHBACK, R. M., LIGON, A. H., CHO, Y. J., HAERY, L., GREULICH, H., REICH, M., WINCKLER, W., LAWRENCE, M. S., WEIR, B. A., TANAKA, K. E., CHIANG, D. Y., BASS, A. J., LOO, A., HOFFMAN, C., PRENSNER, J., LIEFELD, T., GAO, Q., YECIES, D., SIGNORETTI, S., MAHER, E., KAYE, F. J., SASAKI, H., TEPPER, J. E., FLETCHER, J. A., TABERNERO, J., BASELGA, J., TSAO, M., DEMICHELIS, F., RUBIN, M. A., JANNE, P. A., DALY, M. J., NUCERA, C., LEVINE, R. L., EBERT, B. L., GABRIEL, S., RUSTGI, A., ANTONESCU, C. R., LADANYI, M., LETAI, A., GARRAWAY, L., LODA, M., BEER, D., TRUE, L. D., OKAMOTO, A., POMEROY, S. L., SINGER, S., GOLUB, T. R., LANDER, E. S., GETZ, G. and SELLERS, W. R. (2010). The landscape of somatic copy-number alteration across human cancers. *Nature* **463** 899–905.

CANCER GENOME ATLAS RESEARCH NETWORK (2008). Comprehensive genomic characterization defines human glioblastoma genes and core pathways. *Nature* **455** 1061–1068.

CARVALHO, I., MILANEZI, F., MARTINS, A., REIS, R. M. and SCHMITT, F. (2005). Overexpression of platelet-derived growth factor receptor alpha in breast cancer is associated with tumour progression. *Breast Cancer Res.* **7** R788–R795.

CHEN, H., XING, H. and ZHANG, N. R. (2011). Estimation of parent specific DNA copy number in tumors using high-density genotyping arrays. *PLoS Comput. Biol.* **7** e1001060, 15. MR2776334

CHIN, L. and GRAY, J. W. (2008). Translating insights from the cancer genome into clinical practice. *Nature* **452** 553–563.

CHITALE, D., GONG, Y., TAYLOR, B. S., BRODERICK, S., BRENNAN, C., SOMWAR, R., GOLAS, B., WANG, L., MOTOI, N., SZOKE, J., REINERSMAN, J. M., MAJOR, J., SANDER, C., SESHAN, V. E., ZAKOWSKI, M. F., RUSCH, V., PAO, W., GERALD, W. and LADANYI, M. (2009). An integrated genomic analysis of lung cancer reveals loss of DUSP4 in EGFR-mutant tumors. *Nature* **28** 2773–2783.

DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **39** 1–38. MR0501537

DUDOIT, S. and FRIDLYAND, J. (2002). A prediction-based resampling method for estimating the number of clusters in a dataset. *Genome Biology* **3** 1–21.

FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. MR1946581

FANG, K. T. and WANG, Y. (1994). *Number-Theoretic Methods in Statistics. Monographs on Statistics and Applied Probability* **51**. Chapman & Hall, London. MR1284470

FEINBERG, A. P. and VOGELSTEIN, B. (1983). Hypomethylation distinguishes genes of some human cancers from their normal counterparts. *Nature* **301** 89–92.

HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning*: *Data Mining*, *Inference*, *and Prediction*, 2nd ed. Springer, New York. MR2722294

HOLLIDAY, R. (1979). A new theory of carcinogenesis. *Br. J. Cancer* **40** 513–522.

HOLM, K., HEGARDT, C., STAAF, J. et al. (2010). Molecular subtypes of breast cancer are associated with characteristic DNA methylation patterns. *Breast Cancer Research* **12** R36.

HOSHIDA, Y., NIJMAN, S. M., KOBAYASHI, M., CHAN, J. A., BRUNET, J. P., CHIANG, D. Y., VILLANUEVA, A., NEWELL, P., IKEDA, K., HASHIMOTO, M., WATANABE, G., GABRIEL, S., FRIEDMAN, S. L., KUMADA, H., LLOVET, J. M. and GOLUB, T. R. (2003). Integrative transcriptome analysis reveals common molecular subclasses of human hepatocellular carcinoma. *Cancer Research* **69** 7385–7392.

IRIZARRY, R. A., LADD-ACOSTA, C., CARVALHO, B., WU, H., BRANDENBURG, S. A., JEDDELOH, J. A., WEN, B. and FEINBERG, A. P. (2008). Comprehensive high-throughput arrays for relative methylation (CHARM). *Genome Res.* **18** 780–790.

JOLLIFFE, I. T. (2002). *Principal Component Analysis*, 2nd ed. Springer, New York. MR2036084

KAPP, A. V. and TIBSHIRANI, R. (2007). Are clusters found in one dataset present in another dataset? *Biostatistics* **8** 9–31.

LAIRD, P. W. (2003). The power and the promise of DNA methylation markers. *Nat. Rev. Cancer* **3** 253–266.

LAIRD, P. W. (2010). Principles and challenges of genome-wide DNA methylation analysis. *Nat. Rev. Genet.* **11** 191–203.

LAPOINTE, J., LI, C., HIGGINS, J. P., VAN DE RIJN, M., BAIR, E., MONTGOMERY, K., FERRARI, M., EGEVAD, L., RAYFORD, W., BERGERHEIM, U., EKMAN, P., DEMARZO, A. M., TIBSHIRANI, R., BOTSTEIN, D., BROWN, P. O., BROOKS, J. D. and POLLACK, J. R. (2003). Gene expression profiling identifies clinically relevant subtypes of prostate cancer. *Proc. Natl. Acad. Sci. USA* **101** 811–816.

LE CAO, K. A., MARTIN, P. G., ROBERT-GRANIE, P. and BESSE, P. (2009). Sparse canonical methods for biological data integration: Application to a cross-platform study. *BMC Bioinformatics* **26** 34.

NG, A. Y., JORDAN, M. I. and WEISS, Y. (2002). On spectral clustering: Analysis and an algorithm. *Adv. Neural Inf. Process. Syst.* **2** 849–856.

OLSHEN, A. B., VENKATRAMAN, E. S., LUCITO, R. and WIGLER, M. (2004). Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics* **5** 557–572.

OLSHEN, A. B., BENGTSSON, H., NEUVIAL, P., SPELLMAN, P. T., OLSHEN, R. A. and SESHAN, V. E. (2011). Parent-specific copy number in paired tumor-normal studies using circular binary segmentation. *Bioinformatics* **27** 2038–2046.

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., ZHU, J., BERGAMASCHI, A., HAN, W., NOH, D.-Y., POLLACK, J. R. and WANG, P. (2010). Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *Ann. Appl. Stat.* **4** 53–77. MR2758084

PEROU, C. M., JEFFREY, S. S., VAN DE RIJN, M. et al. (1999). Distinctive gene expression patterns in human mammary epithelial cells and breast cancers. *Proc. Natl. Acad. Sci. USA* **96** 9212–9217.

POLLACK, J. R., SØRLIE, T., PEROU, C. M. et al. (2002). Microarray analysis reveals a major direct role of DNA copy number alteration in the transcriptional program of human breast tumors. *Proc. Natl. Acad. Sci. USA* **99** 12963–12968.

ROHE, K., CHATTERJEE, S. and YU, B. (2010). Spectral clustering and the high-dimensional stochastic block model. Available at arXiv:1007.1684.

SHEN, R., OLSHEN, A. B. and LADANYI, M. (2009). Integrative clustering of multiple genomic data types using a joint latent variable model with application to breast and lung cancer subtype analysis. *Bioinformatics* **25** 2906–2912.

SIMON, R. (2010). Translational research in oncology: Key bottlenecks and new paradigms. *Expert Reviews Molecular Medicine* **12** e32.

SONESON, C., LILLJEBJÖRN, H., FIORETOS, T. and FONTES, M. (2010). Integrative analysis of gene expression and copy number alterations using canonical correlation analysis. *BMC Bioinformatics* **11** 191.

SORLIE, T., PEROU, C. M., TIBSHIRANI, R. et al. (2001). Gene expression patterns of breast carcinomas distinguish tumor subclasses with clinical implications. *Proc. Natl. Acad. Sci. USA* **98** 10869–10874.

TCGA NETWORK (2011). Integrated genomic analyses of ovarian carcinoma. *Nature* **474** 609–615.

TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. MR1379242

TIBSHIRANI, R. and WALTHER, G. (2005). Cluster validation by prediction strength. *J. Comput. Graph. Statist.* **14** 511–528. MR2170199

TIBSHIRANI, R. and WANG, P. (2008). Spatial smoothing and hot spot detection for CGH data using the fused lasso. *Biostatistics* **9** 18–29.

TIBSHIRANI, R., SAUNDERS, M., ROSSET, S., ZHU, J. and KNIGHT, K. (2005). Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** 91–108. MR2136641

TIPPING, M. E. and BISHOP, C. M. (1999). Probabilistic principal component analysis. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **61** 611–622. MR1707864

VAN WIERINGEN, W. N. and VAN DE WIEL, M. A. (2009). Nonparametric testing for DNA copy number induced differential mRNA gene expression. *Biometrics* **65** 19–29. MR2665842

VASKE, C. J., BENZ, S. C., SANBORN, J. Z., EARL, D., SZETO, C., ZHU, J., HAUSSLER, D. and STUART, J. M. (2010). Inference of patient-specific pathway activities from multi-dimensional cancer genomics data using PARADIGM. *Bioinformatics* **26** 237–245.

VENKATRAMAN, E. S. and OLSHEN, A. B. (2007). A faster circular binary segmentation algorithm for the analysis of array CGH data. *Bioinformatics* **23** 657–663.

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 correlaton analysis. *Stat. Appl. Genet. Mol. Biol.* **7** Art. 3, 29. MR2386320

WANG, S. and ZHU, J. (2008). Variable selection for model-based high-dimensional clustering and its application to microarray data. *Biometrics* **64** 440–448, 666. MR2432414

WANG, S., NAN, B., ZHU, J. and BEER, D. G. (2008). Doubly penalized Buckley–James method for survival data with high-dimensional covariates. *Biometrics* **64** 132–140, 323. MR2422827

WANG, S., NAN, B., ZHOU, N. and ZHU, J. (2009). Hierarchically penalized Cox regression with grouped variables. *Biometrika* **96** 307–322. MR2507145

WITTEN, D. M. and TIBSHIRANI, R. J. (2009). Extensions of sparse canonical correlation analysis with applications to genomic data. *Stat. Appl. Genet. Mol. Biol.* **8** Art. 28, 29. MR2533636

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.

YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68** 49–67. MR2212574

ZHAO, Y. and SIMON, R. (2010). Development and validation of predictive indices for a continuous outcome using gene expression profiles. *Cancer Inform.* **9** 105–114.

ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** 301–320. MR2137327

R. SHEN
DEPARTMENT OF EPIDEMIOLOGY
  AND BIOSTATISTICS
MEMORIAL SLOAN-KETTERING CANCER CENTER
NEW YORK, NEW YORK 10065
USA
E-MAIL: shenr@mskcc.org

S. WANG
DEPARTMENT OF BIOSTATISTICS
  AND MEDICAL INFORMATICS
DEPARTMENT OF STATISTICS
UNIVERSITY OF WISCONSIN-MADISON
MADISON, WISCONSIN 53792-4675
USA
E-MAIL: swang@biostat.wisc.edu

Q. MO
DIVISION OF BIOSTATISTICS
DAN L. DUNCAN CANCER CENTER
BAYLOR COLLEGE OF MEDICINE
HOUSTON, TEXAS, 77030
USA
E-MAIL: qmo@bcm.edu