

Data Mining in Genomics

Jae K. Lee, PhD*, Paul D. Williams, PhD,
Sooyoung Cheon, PhD

*Division of Biostatistics and Epidemiology, Department of Public Health Sciences,
Box 800717, University of Virginia, Charlottesville, VA 22908, USA*

A great explosion of genomic data has occurred recently, because of the advances in various high-throughput biotechnologies, such as RNA gene expression microarrays. These large genomic data sets are information-rich and often contain much more information than the researchers generating the data may have anticipated. Such an enormous data volume enables new types of analyses but also makes research questions difficult to answer using traditional methods. Analysis of these massive genomic data has several unprecedented challenges.

Challenge 1: multiple comparisons issue

Analysis of high-throughput genomic data requires handling an astronomical number of candidate targets, most of which are false-positives [1,2]. For example, a traditional statistical testing criterion of 5% significance level would result in an average of 500 false-positive genes from a 10K microarray study comparing two biological conditions for which no real biological differences appeared in gene regulation. If a small number of 100 genes, for example, are differentially regulated, the *real* differentially expressed genes will be mixed with the 500 false-positives without any a priori information to discriminate the groups of genes. Confidence on the 600 targets identified with this statistical test is low, and further investigation of these candidates will have a poor yield. Simply tightening this statistical criterion, such as with a 1% or lower significance level, will result in a high false-negative error rate with failure to identify many important real biological targets. This pitfall, the so-called “multiple comparisons issue,” becomes even more serious when trying to find novel biological mechanisms

This study is supported by NIH grant 1R01HL081690 of JKL.

* Corresponding author.

E-mail address: jaeklee@virginia.edu (J.K. Lee).

and biomarker prediction models that involve multiple interacting targets and genes, because the number of candidate pathways or interaction mechanisms grows exponentially. Thus, data mining techniques must effectively minimize both false-positive and false-negative error rates in these genome-wide investigations.

Challenge 2: high-dimensional biological data

The second challenge is the high-dimensional nature of biological data in many genomic studies [3]. In genomic data analysis, many gene targets are investigated simultaneously, yielding dramatically sparse data points in the corresponding high-dimensional data space. Mathematical and computational approaches often fail to capture these high-dimensional phenomena accurately. For example, many search algorithms cannot freely move between local maxima in a high-dimensional space. Furthermore, inference based on the combination of several lower-dimensional observations may not provide a correct understanding of the real phenomenon in their joint high-dimensional space. Consequently, unless appropriate statistical dimension reduction techniques are used to convert high-dimensional data problems into lower-dimensional ones, important variation and information in the biological data may be obscured.

Challenge 3: small n and large p problem

The third challenge is the so-called “small n and large p ” problem [2]. Desired performance of conventional statistical methods is achieved when the sample size of the data, namely n —the number of independent observations and subjects—is much larger than the number of candidate prediction parameters and targets, namely p . In many genomic data analyses this situation is often completely reversed. For example, in a microarray study, tens of thousands of gene expression patterns may become the candidate prediction factors for a biological phenomenon of interest (eg, response versus resistance to a chemotherapeutic regimen), but the number of independent observations (eg, different patients or samples) is often a few tens or hundreds at most. Because of the experimental costs and limited availability of biological materials, the number of independent samples may be even smaller, sometimes only a few. Traditional statistical methods are not designed for these circumstances and often perform very poorly; furthermore, statistical power must be strengthened using all sources of information in large-screening genomic data.

Challenge 4: computational limitation

No matter how powerful a computer system becomes, solving many genomic data mining problems through exhaustive combinatorial search and comparisons is often prohibitive [4]. In fact, many current problems

in genomic data analysis have been theoretically proven to be of nonpolynomial-hard complexity, implying that no computational algorithm can search for all possible candidate solutions. Thus, heuristic—most frequently statistical—algorithms that effectively search and investigate a very small portion of all possible solutions are often sought for genomic data mining problems. The success of many bioinformatics studies critically depends on the construction and use of effective and efficient heuristic algorithms, most of which are based on the careful application of probabilistic modeling and statistical inference techniques.

Challenge 5: noisy high-throughput biological data

The next challenge derives from the fact that high-throughput biotechnical data and large biological databases are inevitably noisy because biological information and signals of interest are often observed with many other random or confounding factors. Furthermore, a one-size-fits-all experimental design for high-throughput biotechniques can introduce bias and error for many candidate targets. Therefore, many investigations in bioinformatics can be performed successfully only when the variability of genomics data is well understood. In particular, the distributional characteristics of each data set must be analyzed using statistical and quality control techniques on initial data sets so that relevant statistical approaches may be applied appropriately. This preprocessing step is critical for all subsequent bioinformatics analyses, and reconciling dramatically different results that may stem from slightly different preprocessing procedures can sometimes be difficult. Although this issue has no easy answer, consistent preprocessing procedures within each and across different analyses, with good documentation of procedures, must be used.

Challenge 6: integration of multiple, heterogeneous biological data for translational bioinformatics research

The last challenge is the integration of genomic data with heterogeneous biological data and associated metadata, such as gene function, biological subjects' phenotypes, and patient clinical parameters. For example, multiple heterogeneous data sets, including gene expression data, biological responses, clinical findings, and outcomes data, may need to be combined to discover genomic biomarkers and gene networks that are relevant to disease and predictive of clinical outcomes, such as cancer progression and chemosensitivity to an anticancer compound. Some data sets exist in different formats and may require combined preprocessing, mapping between data elements, or other preparatory steps before correlative analysis, depending on their biological characteristics and data distributions. Effective combination and use of the information from these heterogeneous genomic, clinical, and other data resources remain a significant challenge.

This article reviews novel concepts and techniques for tackling various genomic data mining problems. In particular, because DNA microarrays and GeneChips techniques have become an important tool in biological and biomedical investigations, this article focuses on statistical approaches that have been applied to various microarray data analyses to overcome some challenges mentioned earlier.

A new concept of statistical significance: false discovery rate

To avoid a large number of false-positive findings, the family-wise error rate (FWER) has been classically controlled for the random chance of multiple hypotheses (or candidates) through evaluating the probability that at most one false-positive is included at a cutoff level of a test statistic among all candidates. However, FWER has been found to be very conservative in microarray studies, resulting in a high false-negative error rate, often very close to 100% [1]. To avoid this pitfall, a novel concept of statistical significance, the so-called “false discovery rate” (FDR) and its refinement, *Q* value, have been suggested [2,5] (*Q*-value package, www.bioconductor.org).

To illustrate what the false discovery rate is, suppose *M* candidates are available for simultaneously testing to reject the null hypothesis of no biological significance. Assume *M*₀ among *M* to be the number of true negative candidates and *M*₁ (= *M* − *M*₀) to be the number of true positive candidates. At a cutoff value of a test statistic or data mining tool, let *R* denote the number of all positives (or significantly identified candidates), *V* the number of false-positives, and *S* the number of false-negatives (Table 1).

Then, the FDR is defined as *V*/*R* if *R* is greater than 0, the ratio between false-positives (*V*) and all positive findings (*R* = *V* + *S*). Note that FDR is thus derived based on the null (no significance) and alternative (significant target) distributions. In contrast, the classical *P* value (or type 1 error), here *V*/*M*₀, and the statistical power (1, type 2 error), or *S*/*M*₁, are based only on one of the null and alternative distributions. Therefore, the FDR criterion can simultaneously balance between false-positives and false-negatives, whereas the classical *P* value and power can address only one of the two errors.

Table 1
Classification of the candidate hypotheses

	Null accept	Null reject	Total
Null true	U	V	M ₀
Alternative true	T	S	M ₁
Total	W	R	M

Abbreviations: S, true positive; T, false negative; U, true negative; V, false positive; W, U + T; R, V + S; M₀, U + V; M₁, T + S; M, M₀ + M₁.

The FDR evaluation has been rapidly adopted for microarray data analysis, including the widely used significance analysis of microarrays (SAM) and other approaches [1,6]. Many different methods have been suggested for estimating FDR directly from test statistics, or indirectly from classical P values of these statistics. The latter methods are convenient, because standard P values can be simply converted into their corresponding FDR values [5,7] and Q value, especially the latter based on a resampling technique. More careful FDR assessment can also be found in many other recent studies [7].

Pairwise statistical tests for genomic data

The differential expression pattern of each gene in a microarray experiment is usually assessed with (typically pairwise) contrasts of mean expression values among experimental conditions. These comparisons have been routinely measured as fold changes, whereby genes with greater than two- or threefold changes are selected for further investigation. Results have frequently found that a gene that shows a high fold change between comparison conditions might also exhibit high variability in general, and hence its differential expression may not be significant. Similarly, a modest change in gene expression may be significant if its differential expression pattern is highly reproducible. Several authors have indicated this fundamental flaw in the fold change-based approach [1]. Thus, the emerging standard approach is based on statistical significance and hypothesis testing, with careful attention to reliability of variance estimates and multiple comparison issues.

The classical t test and other traditional test statistics have been initially used for testing the differential expression of each gene [6]. These classical testing procedures, however, rely on reasonable estimates of reproducibility or within-gene error, requiring a large number of replicated arrays. When a small number of replicates are available per condition (eg, duplicate or triplicate), the use of within-gene estimates of variability does not provide a reliable hypothesis-testing framework. For example, a gene may have similar differential expression values in duplicate experiments through chance alone. Furthermore, comparing means can be misled by outliers with dramatically smaller or larger expression intensities than other replicates, and therefore error estimates constructed solely within genes may result in underpowered tests for differential expression comparisons and also in large numbers of false-positives. Several approaches to improving estimates of variability and statistical tests of differential expression have thus recently emerged [8–10].

Significance analysis of microarrays

SAM has been proposed to improve the unstable error estimation in the 2-tailed t test through adding a variance stabilization factor that minimizes

the variance variability across different intensity ranges [1]. Based on the observation that the signal-to-noise ratio varies with different gene expression intensities, SAM tries to stabilize gene-specific fluctuations and is defined based on the ratio of change in gene expression to the standard deviation in the data for that gene. The relative difference $d(i)$ in gene expression is defined as:

$$d(i) = (x_I(i) - x_U(i)) / (s(i) + s_0)$$

where $x_I(i)$ and $x_U(i)$ are the average expression values of gene i in states I and U , respectively. The gene-specific scatter $s(i)$ is the standard pooled deviation of replicated expression values of the gene in the two states. To compare values of $d(i)$ across all genes, the distribution of $d(i)$ is assumed to be independent of the level of gene expression. However, at low expression levels, variability in $d(i)$ can be high because of small values of $s(i)$. To ensure that the variance of $d(i)$ is independent of gene expression, a positive constant s_0 is added to the denominator. The value for s_0 is chosen to minimize the coefficient of variation, where the coefficient of variability of $d(i)$ is computed as a function of $s(i)$ in moving windows across all the genes.

Local pooled error

Based on a more careful error-pooling technique, the so-called “local-pooled-error” (LPE) test was also introduced. This testing technique is particularly useful when the sample size is very small (eg, two or three per condition). LPE variance estimates for genes are formed by pooling variance estimates for genes with similar expression intensities from replicated arrays within experimental conditions [6]. The LPE approach leverages the observations that genes with similar expression intensity values often show similar array-experimental variability within experimental conditions; and that variance of individual gene expression measurements within experimental conditions typically decreases as a (nonlinear) function of intensity. LPE has been introduced specifically for analyzing small-sample microarray data, whereby error variance estimates for genes are formed through pooling variance estimates for genes with similar expression intensities from replicated arrays within experimental conditions (LPE package, www.bioconductor.org). The LPE approach is possible because common background noise can often be found within each local intensity region of the microarray data. At high levels of expression intensity, this background noise is dominated by the expression intensity, whereas at low levels the background noise is a larger component of the observed expression intensity, which can be easily observed in the so-called “ M versus A log-intensity scatter plot” of two replicated chips among three different immune conditions (Fig. 1) [6]. The LPE approach controls the situation when a gene with low expression may have very low variance by chance and the resulting signal-to-noise ratio is unrealistically large.




Fig. 1. Log-intensity ratio (M) as a function of average gene expression between replicated chips (A). Top panels represent the estimated error distributions (based on a non-parametric regression) for (A) naive, (B) 48 hour activated, and (C) T-cell clone D4 conditions in the mouse immune response microarray study.

Statistical significance of the LPE-based test is evaluated through first calculating each gene's medians m_1 and m_2 under the two compared conditions to avoid artifacts from outliers. The LPE statistic for the median (log-intensity) difference Z is then calculated as:

$$Z = (m_1 - m_2) / S_{\text{LPE}_{\text{pooled}}}$$

where $S_{\text{LPE}_{\text{pooled}}}$ is the pooled standard error from the LPE-estimated baseline variances from the two conditions. The LPE approach shows a significantly better performance than 2-tailed t test, SAM, and Westfall-Young's permutation tests, especially when the number of replicates is smaller than 10 [6].

Statistical modeling on genomic data

Genomic expression profiling studies are also frequently performed for comparing complex, multiple biological conditions and pathways. Several linear modeling approaches have been introduced for analyzing microarray data with multiple conditions. For example, an analysis of variance (ANOVA) model approach was considered to capture the effects of dye,

array, gene, condition, array–gene interaction, and condition–gene interaction separately on complementary DNA (cDNA) microarray data [11], and a two-stage mixed model was proposed first to model cDNA microarray data with the effects of array, condition, and condition–array interaction and then fit the residuals with the effects of gene, gene–condition interaction, and gene–array interaction [12]. Several approaches have also been developed using the Bayesian paradigm for analyzing microarray data, including Bayesian parametric modeling [13], Bayesian regularized t test [8], Bayesian hierarchical modeling with a multivariate normal prior [14], and Bayesian heterogeneous error model (HEM) with two error components [15].

Analysis of variance modeling

The use of ANOVA models has been suggested to estimate relative gene expression and to account for other sources of variation in microarray data [16]. Although the exact form of the ANOVA model depends on the particular data set, a typical ANOVA model for two-color–based cDNA microarray data can be defined as

$$y_{ikg} = \mu + A_i + D_j + V_k + G_g + AD_{ij} + AG_{ig} + DG_{ig} + VG_{kg} + \epsilon_{ijk}$$

where y_{ikg} is the measured intensity from array i , dye j , variety k , and gene g on an appropriate scale (typically the log scale). The generic term *variety* is often used to refer to the mRNA samples studied, such as treatment and control samples; cancer and normal cells; or time points of a biological process. The terms A , D , and AD account for the overall effects that are not gene-specific. The gene effects G_g capture the average levels of expression for genes and the array-by-gene interactions AG_{ig} capture differences caused by varying sizes of spots on arrays. The dye-by-gene interactions DG_{ig} represent gene-specific dye effects. None of these effects are of biological interest but amount to a normalization of the data for ancillary sources of variation. The effects of primary interest are the interactions between genes and varieties, VG_g . These terms capture differences from overall averages that are attributable to the specific combination of variety k and gene g . Differences among these variety-by-gene interactions provide the estimates for the relative expression of gene g in varieties 1 and 2 through $VG_{1g} - VG_{2g}$. Note that AV , DV , and other higher-order interaction terms are typically assumed to be negligible and are considered together with the error terms. The error terms ϵ_{ijk} are often assumed to be independent and normal with mean zero and a common variance. However, such a global ANOVA model is difficult to implement in practice because of its computational restriction. Instead, one often considers gene-by-gene ANOVA models such as

$$y_{ikg} = \mu_g + A_i + D_j + V_k + AD_{ij} + VG_{kg} + \epsilon_{ijk}$$

Alternatively, a two-stage ANOVA model may be used [12]. The first layer is for main effects non-specific to the gene effects:

$$y_{ijk} = \mu_g + A_i + D_j + V_k + AD_{ij} + AG_{ig} + \epsilon_{ijk}$$

Let r_{ijk} be the residuals from this first ANOVA fit. Then, the second-layer ANOVA model for gene-specific effects is considered as

$$r_{ijk} = G_g + AG_{ig} + DG_{ig} + VG_{kg} + v_{ijk}$$

Excepting the main effects of G and V and their interaction effects, the other terms A , D , (AD) , (AG) , and (DG) can be considered as random effects. These within-gene ANOVA models can be implemented using most standard statistical packages, such as R, SAS, or SPSS.

Heterogeneous error model

Similar to the statistical tests for comparing two sample conditions, the within-gene ANOVA modeling methods are underpowered and have inaccurate error estimation in microarray data with limited replication. HEM has been suggested as an alternative (HEM package at www.bioconductor.org). It is based on Bayesian hierarchical modeling and LPE error-pooling-based prior constructions, with two layers of error that decompose the total error variability into the technical and biological error components in microarray data [15]. The first layer is constructed to capture the array technical variation caused by many experimental error components, such as sample preparation, labeling, hybridization, and image processing: $y_{ijkl} = x_{ijk} + \epsilon_{ijkl}$, where ϵ_{ijkl} is approximately *iid* Normal $[0, \sigma^2(x_{ijk})]$, where $i = 1, 2, \dots, G$; $j = 1, 2, \dots, C$; $k = 1, 2, \dots, m_{ij}$; $l = 1, 2, \dots, n_{ijk}$. The second layer is then hierarchically constructed to capture the biological error component: $x_{ijk} = \mu + g_i + c_j + r_{ij} + b_{ijk}$, where b_{ijk} is approximately *iid* Normal $[0, \sigma^2_b(ij)]$. Here, the genetic parameters are for the grand mean (shift or scaling) constant, gene, cell, interaction effects, and the biological error; the last error term varies and is heterogeneous for each combination of different genes and conditions. The biological variability is individually assessed for discovery of biologically relevant expression patterns. The HEM approach shows a significantly better performance than standard ANOVA methods, especially when the number of replicates is small (Fig. 2).

Unsupervised learning: clustering

Clustering analysis is widely applied to search for the groups (clusters) in microarray data, because these techniques can effectively reduce the high-dimensional gene expression data into a two-dimensional dendrogram organized by each gene's expression association patterns (Fig. 3). Currently,




Fig. 2. Receiver operating characteristic curves from heterogeneous error model (*solid lines*) and analysis of variance (*dotted lines*) models with two and five replicated arrays. The horizontal axis is 1-false-positive error rate (FPR) and the vertical axis is 1-false-negative error rate (FNR).

clustering analysis is one of the most frequently used techniques for genomic data mining in biomedical studies [17–19]. Some technical aspects of these approaches are summarized. A clustering approach first must be defined through a measure or distance index of similarity or dissimilarity, such as

- Euclidean: $d(x, y) = \Sigma(x_k - y_k)^2$
- Manhattan: $d(x, y) = \Sigma|x_k - y_k|$
- Correlation: $d(x, y) = 1 - r(x, y)$, where $r(x, y)$ is a correlation coefficient

Next, an allocation algorithm must be defined based on one of these distance metrics. Two classes of clustering algorithms have been used in genomic data analysis: hierarchical and partitioning allocation algorithms.

Hierarchical algorithms that allocate each subject to its nearest subject or group include:

- Agglomerative methods: average linkage based on group average distance, single linkage based on minimum nearest distance, and complete linkage based on maximum furthest distance;




Fig. 3. Dendrogram (*top panel*) and heatmap (*bottom panel*) of hierarchical clustering analysis for the concordant complementary DNA (cDNA) and oligo array expression patterns on the NCI-60 cancer cell lines. A region of heatmap occupied by melanoma genes are shown from the combined set of 3297 oligo and cDNA transcripts. Each gene expression pattern is designated as coming from the cDNA or oligo array transcripts. The concordant oligo and cDNA microarray expressions are marked with blue bars.

- Probabilistic methods: Bayes factor, posterior probabilities of subclusters;
- Divisive methods: monothetic variable division, polythetic division.

Partitioning algorithms divide the data into a prespecified number of subsets, including

- Self-organizing map: division into a geometrically preset grid structure of subclusters;
- K-means: iterative relocation into a predefined number of subclusters;
- Partitioning around medoids: similar to but more robust than K-means clustering;
- Clara: division of fixed-size subdatasets for applications to large data sets;
- Fuzzy algorithm: probabilistic fractions of membership rather than deterministic allocations.

One of the most difficult aspects of using these clustering analyses is the interpretation of their heuristic, often unstable, clustering results. To overcome this shortcoming, several refined clustering approaches have been suggested. For example, the use of *bootstrapping* was suggested to evaluate the consistency and confidence of each gene's membership to particular cluster groups [11]. The *gene shaving* approach was suggested to find the clusters directly relevant to major variance directions of an array data set [3]. Recently, *tight clustering*, a refined bootstrap-based hierarchical clustering, was proposed to formally assess and identify the groups of genes that are most tightly clustered with each other [20].

Supervised learning: classification

Supervised classification learning on genomic data is often performed to obtain genomic prediction models for different groups of biological subjects (eg, macrophage cells under different immunologic conditions) and patients (different subclasses of patients who have cancer, such as those who have acute lymphoblastic leukemia versus acute myeloid leukemia). Prediction based on genomic expression signatures has received considerable attention in many challenging classification problems in biomedical research [21,22]. For example, these analyses have been conducted in cancer research as alternative diagnostic techniques to the traditional ones, such as classification through the origin of cancer tissues or microscopic appearance, which can be problematic for predicting many critical human disease subtypes [23]. Several different approaches to microarray classification modeling have been proposed, including gene voting [21], support vector machines (SVMs) [24], Bayesian regression models [22], partial least squares [25], and GA/KNN [26]. The following discussion considers strategies to evaluate and compare the performance of these different classification methods.

Measures for classification model performance

Microarray data often have tens of thousands of genes on each chip, whereas only a few tens of samples or replicated arrays are available in a microarray study. In the classification modeling on genomic data, avoiding overfitting is essential, as is finding an optimal subset of the thousands of genes for constructing classification rules and models that are robust in different choices of training samples and consistent in prediction performance on future samples. Typically, in this kind of supervised learning, separate training and test sets (of subjects with known classes) are used, the former to fit classification prediction models, and the latter, which is independent of the former set, for rigorous model validation. Evaluation of prediction performance should then be carefully conducted among the extremely large number of competing models, especially in using appropriate performance selection criteria and in using the whole data for model training and

evaluation. Several different measures are currently used to evaluate performance of classification models: classification error rate, area under the receiver operating characteristic curve (area under the curve [AUC]), and the product of posterior classification probabilities [27,28].

When a large number of candidate models (eg, approximately 10^8 two-gene models on 10K array data) are compared in their performance, these measures are often saturated—their maximum performance levels are achieved using many competing models—so that identification of the best (most robust) prediction model among them is extremely difficult. Furthermore, these measures cannot capture an important aspect of classification model performance as follows: suppose three samples are classified using two classification models (or rules); one model provides the correct posterior classification probabilities 0.8, 0.9, and 0.4, and the other 0.8, 0.8, and 0.4 for the three samples. Assuming these were unbiased estimates of classification error probabilities (on future data), the former model would be preferred because this model will perform better in terms of the expected number of correctly classified samples in future data.

Note that the two models provide the same misclassification error rate, one third. This aspect of classification performance cannot be captured through evaluating the commonly used error rate or AUC criteria, which simply add one count for each correctly classified sample, ignoring its degree of classification error probability.

To overcome this limitation, the so-called “misclassification-penalized posterior” (MiPP) criterion has been suggested recently [4]. This measure is the sum of the correct-classification (posterior) probabilities of correctly classified samples subtracted by the sum of the misclassification (posterior) probabilities of misclassified samples. Suppose there are m classes π_k , $i = 1, \dots, m$, from a population of N samples. Let $X_j, j = 1, \dots, n_i$, be the j th sample. Under a particular prediction model (eg, one- or two-gene model) from a classification rule, such as linear discriminant analysis or SVMs, MiPP is then defined as:

$$\Lambda = \sum_{\text{correct}} p_k(X_j) - \sum_{\text{wrong}} (1 - p_k(X_j))$$

where $p_k(X_j)$ is the posterior classification probability of sample X_j into the k th class. Here *correct* and *wrong* correspond to the samples that are correctly and incorrectly classified. In the two-class problem, *correct* simply means $p_k(X_j)$ is more than 0.5, but in general it occurs when $p_k(X_j) = \max_{i=1, \dots, m} (p_i(X_j))$. MiPP can also be shown to be the sum of the posterior probabilities of correct classification penalized by the number of misclassified samples (N_M): $\Lambda = \sum p_k(X_j) - N_M$. Thus, MiPP is a continuous measure (compared with the discrete error rate) of classification performance that accounts for the degree of classification certainty and the error rate, and is sensitive enough to distinguish subtle differences in prediction performance among many competing models.

Classification modeling

Several classification modeling approaches are currently widely used in genomic data analysis.

Gene voting

Gene voting [21] is an intuitively-derived technique that aggregates the weighted votes from all modeling gene signatures; the advantage of this technique is that it can be easily implemented without complicated computing and statistical arguments. It has been proposed for predicting subclasses of patients who have acute leukemia observed with microarray gene expression data [21]. This method gains accuracy through aggregating predictors built from a learning set and casting their voting weights. For binary classification, each gene casts a vote for class 1 or 2 among p samples, and the votes are aggregated over genes. For gene g_j , the vote is $v_j = a_j (g_j - b_j)$, where $a_j = (m_1 - m_2)/(s_1 + s_2)$ and $b_j = (m_1 + m_2)/2$ for sample means m_1 and m_2 and sample standard deviations s_1 and s_2 . Using this method based on 50 gene predictors, 36 of 38 patients in an independent validation set have been correctly classified between acute myeloid leukemia and acute lymphoblastic leukemia.

Linear and quadratic discriminant analysis

Linear or quadratic discriminate analysis is a classical statistical classification technique based on the multivariate normal distribution assumption. This technique is frequently found to be robust and powerful for many different applications, despite the distributional assumption; the gene voting technique can be considered as a variant of linear discriminant analysis. Linear discriminant analysis can be applied with leave-one-out classification, assuming each class follows a multivariate normal distribution. Each sample will then be allocated to group k , to which its classification probability is maximized. The quadratic discriminate analysis can be similarly performed, except that the covariance matrix of the multivariate normal distribution (for each of m classes) is now considered differently among m classes. Differences between linear discriminant analysis and quadratic discriminant analysis are typically small, especially if polynomial factors are considered in linear discriminant analysis. In general, quadratic discriminant analysis requires more observations to estimate each variance–covariance matrix for each class. Linear and quadratic discriminant analysis have consistently shown high performance, not because the data is likely derived from Gaussian distributions, but more likely because the data support only a simple boundaries, such as linear or quadratic [28].

Logistic regression

The logistic regression classification technique is based on the regression fit on probabilistic odds among comparing conditions. This technique

requires no specific distribution assumption but is often found to be less sensitive than other approaches. Logistic regression methods simply maximize the conditional likelihood $Pr(G = k|X)$, typically by a Newton-Raphson algorithm [29]. The allocation decision on a sample is based on the logistic regression fit:

$$\text{Logit}(p_i) = \log(p_i/(1 - p_i)) \sim \beta^T x$$

where β is the logistic regression estimated coefficient vector for the microarray data. Logistic regression discriminant analysis is often used because of its flexible assumption about the underlying distribution, but if it is actually from Gaussian distribution, logistic regression shows a loss of 30% efficiency in the (misclassification) error rate compared with linear discriminant analysis.

Support vector machines

Conceptually similar to gene voting, SVMs are one of the recent machine-learning classification techniques based on the data projection to high-dimensional kernel space. This technique also does not require distributional assumption, yet can perform better than other approaches in some complicated cases. However, it often requires large numbers of samples and predictor gene signatures for optimal performance. SVMs separate a given set of binary labeled training data with a hyperplane that is maximally distant from them, known as the *maximal margin hyperplane* [24]. Based on a kernel, such as a polynomial of dot products, the current data space will be embedded in a higher dimensional space. Commonly used kernels include

- Radial basis function kernel: $K(x, y) = \exp(-|x - y|^2/2s^2)$
- Polynomial kernel: $K(x, y) = \langle x, y \rangle^d$ or $K(x, y) = (\langle x, y \rangle + c)^d$, where $\langle, \rangle$ denotes the inner product.

Comparison of classification methods

These classification techniques must be carefully applied in prediction model training on genomic data. In particular, if all the samples are used for model search/training and evaluation in a large screening search for classification models, a serious selection bias is inevitably introduced [30]. To avoid this pitfall, a stepwise (leave-one-out) cross-validated discriminant procedure has been suggested that gradually adds genes to the training set [4,28]. The prediction performance is typically found to be continuously improved (or not decreased) through adding more features into the model. This result is again caused by a sequential search-and-selection strategy against an astronomically large number of candidate models; some of them can show overoptimistic prediction performance for a particular

training set by chance. Furthermore, even though a leave-one-out or similar cross-validation strategy is used in this search, the number of candidate models is too big to eliminate many random ones that survive by chance from cross-validation. Thus, test data should be completely independent from the training data to obtain an unbiased estimate of each model's performance. To address these pitfalls, the stepwise cross-validated discriminant (SCVD) procedure sequentially adds one gene at a time to identify the most optimal prediction model based both on *n*-fold modeling and train-test validation strategies. SCVD can be used with any of the aforementioned classification methods.

Linear discriminant analysis, quadratic discriminant analysis, logistic regression, and SVMs with linear or radial basis function kernels have been compared using the SCVD approach [28]. The leukemia microarray data in the study by Golub and colleagues [21] had a training set of 27 acute lymphoblastic leukemia and 11 acute myeloid leukemia samples and an independent test set of 20 acute lymphoblastic leukemia and 14 acute myeloid leukemia samples. Because two distinct data sets exist, the model is constructed based on the training data and evaluated on the test data set. Each rule identified a somewhat different subset of features that showed the best performance within each classification method (Table 2). In terms of error rate, the SVM with a linear kernel seems to be the most accurate rule. However, linear discriminant analysis only misclassified one sample and the SVM with the radial basis function kernel and quadratic discriminant analysis misclassified two samples on the independent test data. Logistic regression does not seem to perform as well as the other rules, misclassifying 4 of 34 samples. Comparing the rules based on MiPP is somewhat complicated for SVMs, because the estimated probabilities of correct classification from SVMs are based on how far samples are from a decision boundary. Therefore, these are not true probabilities, as is the case with linear discriminant analysis, quadratic discriminant analysis, and logistic

Table 2
Classification results of the classification rules and the corresponding gene model

Method	Gene model	Error rate on training data	MiPP on training data	Error rate on test data	MiPP on test data
LDA	1882 + 1144	0%	37.91	2.9%	31.46
QDA	4847 + 5062	0%	37.96	5.8%	29.81
Logistic	1807 + 4211 + 575	0%	37.99	11.8%	25.64
SVM K =	2020 + 4377 + 1882	0%	35.16	0%	29.26
linear					
SVM K =	4847 + 3867 + 6281	0%	32.52	5.9%	21.71
RBF					

Abbreviations: K, kernel; LDA, linear discriminant analysis; MiPP, misclassification-penalized posterior; QDA, quadratic discriminant analysis; RBF, radial basis function; SVM, support vector machines.

regression. In an application to a different microarray study on colon cancer, the radial basis function–kernel SVM model with three genes was found to perform best among these classification techniques.

The MiPP-based SCVD procedure was the most robust classification model and could accurately classify samples with a very small number of features—only two or three genes for the two well-known microarray data sets, outperforming many previous models with 50 to 100 features—although different classification methods may perform differently in different data sets. These data are consistent with the notion that many correlated genes share more or less similar information and may discriminate similarly among different subtypes of a particular disease, and that multiple small-feature models may perform well in terms of the construction of a classification model. As shown, the prediction performance on the training set is quickly saturated with a 0% error rate and very close to the maximum MiPP value of 38 (total sample size). However, error rates and MiPP values vary greatly on the independent test set. The error rates were also found to be misleading and less informative than MiPP.

Genomic pathway modeling

Many recent pathway-modeling studies for transcriptional regulation and gene functional networks have been performed based on genomic expression data. These approaches can be largely divided into three categories, qualitative, quantitative, and integrative pathway modeling, based on the different types of genomic data used. This article introduces several pathway-modeling approaches with references. Pathway modeling is one of the most active research fields in current genomic sciences, and substantial additional information can be found in the references.

Qualitative pathway modeling

Pathway modeling has been performed using functional and annotation information from several genomic databases. For example, computationally predicting genomewide transcription units based on pathway–genome databases (PGDBs) and other organizational (ie, protein complexes) annotation improved transcription unit organization information in *Escherichia coli* and *Bacillus subtilis* [31]. A classification of transcription factors is also proposed to organize pathways that connect extracellular signaling to the regulation of transcription and constitutively activate nuclear factors in eukaryotic cells. This latter classification was performed based on known cellular characteristics that describe the roles of these factors within regulatory circuits to identify many downstream functional mechanisms, such as serine and tyrosine phosphorylation; Rel/nuclear factor- κ B family, Ci/Gli, Wnt and Notch pathway, nuclear factor of activated T cells transcription factor activation; and Ca²⁺ increase [32].

Quantitative pathway modeling

Gene regulation networks have also been explored based on quantitative genomic expression data. For example, Bayesian network modeling was used for capturing regulatory interactions between genes based on genome-wide expression measurements on yeast [33].

Probabilistic models for context-specific regulatory relationships were also proposed to capture complex expression patterns of many genes in various biological conditions, accounting for known variable factors, such as experimental settings, putative binding sites, or functional information in yeast stress and compendium data [34]. These quantitative pathway models have been found to effectively characterize both relationships and magnitude of relevant genes' expression patterns, and have been used extensively in recent pathway modeling in various microarray studies [33–35].

Integrative pathway modeling

Integration of qualitative and quantitative gene network information has been attempted in recent pathway modeling studies. For example, a comprehensive genomic module map was constructed through combining gene expression and known functional and transcriptional information, wherein each module represents a distinctive set of directly regulated and associated genes that act in concert to perform a specific function. In this study, different expression activities in tumors were described in terms of the behavior of these modules [35]. Regression on transcription motifs is proposed for discovering candidate genes' upstream sequences that undergo expression changes in various biological conditions. This method combines the known motif structural information and gene expression patterns based on an integrated regression analysis [36].

Genomic biomarkers for disease progression and chemosensitivity

Recently, genomic data have been used to predict the outcome and tumor chemosensitivity in patients who have cancer. The prognosis of these patients is frequently uncertain, and histologically based prognostic indicators are sometimes inaccurate because of the inherent complexities involved. Deregulation of tumor cells leads to uncontrolled division, invasion, and metastasis; the specific patterns of deregulation in a particular tumor and which genetic pathways are altered likely affect the course of the disease. Based on these observations, genomic predictors have been developed for the progression of metastatic disease after removal of breast cancer tumors [37]. In another study, genomic predictors were generated using tumor samples from 78 patients, of whom 34 developed metastasis within 5 years of surgical resection [38]. The first step in the development of the predictor was to determine, using microarray analysis, which genes were significantly

differentially expressed compared with a pool of all specimens. Of 24,479 genes on the chip, 4968 were found to have at least a twofold change in expression and a P value less than 0.01 in at least five tumors. To determine which genes could be used to predict metastasis, the correlation coefficients between gene expression values and metastasis development were calculated and 231 genes were found with an absolute correlation coefficient greater than 0.3. Leave-one-out cross-validation on sequentially larger subsets of genes determined that a group of 70 genes yielded the least-erroneous classifier. Testing the classifier on a 19-patient dataset resulted in two incorrect predictions, and this was later expanded to predict survival times for patients who had breast cancer, accurately predicting disease-free long-term survivors against patients who had poor prognosis [39].

A combination of in vitro assays and genomic data was used to predict the prognosis of patients who had non-small cell lung cancer based on Bayesian computational classification techniques [40]. In this study, microarray profiling analysis was performed to determine gene expression profiles of the various tumors, using Affymetrix HG-U133 plus 2.0 chips. First filtering found approximately 2070 genes highly correlated to lung cancer recurrence. The k -means method was used to generate gene clusters, which were then analyzed using singular value decomposition to generate a metagene, the dominant average expression pattern of the cluster. These metagenes were used in a binary regression tree to partition tumor samples into subsets on which predictions of recurrence could be made; the prediction accuracy for 5-year disease-free survival was more than 90% among the predicted long-term survivors compared with less than 40% among the patients who had predicted poor prognosis.

Summary

The authors believe that several reasons exist why these genomic data mining approaches have been successful and represent a promising direction for future work. First, gene networks inferred from genomic expression signatures have been found to be highly relevant to patient prognosis and chemotherapeutic responses [41–43]. Even though many individual gene expression values in these networks are often variable and noisy, the entire gene networks have been found to be consistent in their overall expression patterns [44–46]. Second, genomewide RNA expression profiling techniques such as microarrays and GeneChips have been dramatically improved recently, so that the expression patterns of the entire human genome can now be accurately and cost-effectively measured on patient samples. In fact, microarray RNA profiling is one of the most accurately quantifiable and comprehensive profiling biotechnologies among all current high-throughput biotechniques, including comparative genomic hybridization, spectral karyotyping, serial analysis of gene expression, two-dimensional gel electrophoresis, mass spectrometry, or protein arrays [47,48]. Third, bioinformatics analysis methods

and techniques for these microarray data have been significantly improved, especially in testing (eg, SAM, LPE, false discovery rate), clustering (eg, hierarchical, self-organizing map, K-means, response projected clustering), classification (linear discriminant analysis, SVMs, logistic regression, random forest), and pathway analysis (Gene Map Annotator And Pathway Profiler, ingenuity pathway analysis) for investigating the complex and extensive information in massive genomic data sets effectively and efficiently [1,4,24]. Finally, most importantly, based on significant efforts by the National Institutes of Health (Gene Expression Omnibus [GEO]) and the European Bioinformatics Institute (ArrayExpress), many precious microarray data sets of cancer—cell lines and patients—have been archived for public access. For example, GEO currently archived more than 5550 microarray data sets on more than 150,000 different biomedical samples and human patients with more than 1500 sets for cancer alone. Furthermore, despite their technical differences, microarray data sets from different time points, laboratories, and even platforms contain consistent information for many gene expression patterns, so that investigations can be performed successfully across those different genomic data sets. This large and rapidly increasing compendium of data demands data mining approaches and ensures that genomic data mining will continue to be a necessary and highly productive field.

References

- [1] Tusher VG, Tibshirani R, Chu G. Significance analysis of microarrays applied to the ionizing radiation response. *Proc Natl Acad Sci U S A* 2001;98(9):5116–21.
- [2] Storey JD, Tibshirani R. Statistical significance for genomewide studies. *Proc Natl Acad Sci U S A* 2003;100(16):9440–5.
- [3] Hastie T, Tibshirani R, Eisen MB, et al. ‘Gene shaving’ as a method for identifying distinct sets of genes with similar expression patterns. *Genome Biol* 2000;1(2):p. RESEARCH0003.
- [4] Soukup M, Cho H, Lee JK. Robust classification modeling on microarray data using misclassification penalized posterior. *Bioinformatics* 2005;21(Suppl 1):i423–30.
- [5] Benjamini Y, Drai D, Elmer G, et al. Controlling the false discovery rate in behavior genetics research. *Behav Brain Res* 2001;125(1–2):279–84.
- [6] Jain N, Thattai J, Braciale T, et al. Local-pooled-error test for identifying differentially expressed genes with a small number of replicated microarrays. *Bioinformatics* 2003;19(15):1945–51.
- [7] Jain N, Cho H, O’Connell N, et al. Rank-invariant resampling based estimation of false discovery rate for analysis of small sample microarray data. *BMC Bioinformatics* 2005;6:187.
- [8] Baldi P, Long AD. A Bayesian framework for the analysis of microarray expression data: regularized t-test and statistical inferences of gene changes. *Bioinformatics* 2001;17(6):509–19.
- [9] Efron B, Tibshirani R. Empirical Bayes methods and false discovery rates for microarrays. *Genet Epidemiol* 2002;23(1):70–86.
- [10] Kerr MK, Martin M, Churchill GA. Analysis of variance for gene expression microarray data. *J Comput Biol* 2000;7(6):819–37.
- [11] Kerr MK, Churchill GA. Bootstrapping cluster analysis: assessing the reliability of conclusions from microarray experiments. *Proc Natl Acad Sci U S A* 2001;98(16):8961–5.
- [12] Wolfinger RD, Gibson G, Wolfinger ED, et al. Assessing gene significance from cDNA microarray expression data via mixed models. *J Comput Biol* 2001;8(6):625–37.

- [13] Newton MA, Kendziorski CM, Richmond CS, et al. On differential variability of expression ratios: improving statistical inference about gene expression changes from microarray data. *J Comput Biol* 2001;8(1):37–52.
- [14] Ibrahim JGaC, M.-H., Gray RJ. Bayesian models for gene expression with DNA microarray Data. *J Am Stat Assoc* 2002;97:88–99.
- [15] Cho H, Lee JK. Bayesian hierarchical error model for analysis of gene expression data. *Bioinformatics* 2004;20(13):2016–25.
- [16] Kerr MK, Churchill GA. Statistical design and the analysis of gene expression microarray data. *Genet Res* 2001;77(2):123–8.
- [17] Lee JK, Bussey KJ, Gwadry FG, et al. Comparing cDNA and oligonucleotide array data: concordance of gene expression across platforms for the NCI-60 cancer cells. *Genome Biol* 2003;4(12):R82.
- [18] Scherf U, Ross DT, Waltham M, et al. A gene expression database for the molecular pharmacology of cancer. *Nat Genet* 2000;24(3):236–44.
- [19] Weinstein JN, Scherf U, Lee JK, et al. The bioinformatics of microarray gene expression profiling. *Cytometry* 2002;47(1):46–9.
- [20] Tseng GC, Wong WH. Tight clustering: a resampling-based approach for identifying stable and tight patterns in data. *Biometrics* 2005;61(1):10–6.
- [21] Golub TR, Slonim DK, Tamayo P, et al. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* 1999;286(5439):531–7.
- [22] West M, Blanchette C, Dressman H, et al. Predicting the clinical status of human breast cancer by using gene expression profiles. *Proc Natl Acad Sci U S A* 2001;98(20):11462–7.
- [23] Su AI, Welsh JB, Sapinoso LM, et al. Molecular classification of human carcinomas by use of gene expression signatures. *Cancer Res* 2001;61(20):7388–93.
- [24] Furey TS, Cristianini N, Duffy N, et al. Support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics* 2000;16(10):906–14.
- [25] Nguyen DV, Rocke DM. Partial least squares proportional hazard regression for application to DNA microarray survival data. *Bioinformatics* 2002;18(12):1625–32.
- [26] Li L, Darden TA, Weinberg CR, et al. Gene assessment and sample classification for gene expression data using a genetic algorithm/k-nearest neighbor method. *Comb Chem High Throughput Screen* 2001;4(8):727–39.
- [27] Hand DJ. Construction and assessment of classification rules. Chichester: John Wiley & Sons; 1997.
- [28] Soukup M, Lee JK. Developing optimal prediction models for cancer classification using gene expression data. *J Bioinform Comput Biol* 2004;1(4):681–94.
- [29] Pampel FC. Logistic regression: a primer. Sage University Papers Series on Quantitative Applications of the Social Sciences; 2000.
- [30] Ambroise C, McLachlan GJ. Selection bias in gene extraction on the basis of microarray gene-expression data. *Proc Natl Acad Sci U S A* 2002;99(10):6562–6.
- [31] Romero PR, Karp PD. Using functional and organizational information to improve genome-wide computational prediction of transcription units on pathway-genome databases. *Bioinformatics* 2004;20(5):709–17.
- [32] Brivanlou AH, Darnell JE Jr. Signal transduction and the control of gene expression. *Science* 2002;295(5556):813–8.
- [33] Friedman N, Linial M, Nachman I, et al. Using Bayesian networks to analyze expression data. *J Comput Biol* 2000;7(3–4):601–20.
- [34] Segal E, Taskar B, Gasch A, et al. Rich probabilistic models for gene expression. *Bioinformatics* 2001;17(Suppl 1):S243–52.
- [35] Segal E, Friedman L, Koller D, et al. A module map showing conditional activity of expression modules in cancer. *Nat Genet* 2004;36(10):1090–8.
- [36] Conlon EM, Liu XS, Lieb JD, et al. Integrating regulatory motif discovery and genome-wide expression analysis. *Proc Natl Acad Sci U S A* 2003;100(6):3339–44.

- [37] van 't Veer LJ, Dai H, van de Vijver MJ, et al. Gene expression profiling predicts clinical outcome of breast cancer. *Nature* 2002;415(6871):530–6.
- [38] van 't Veer LJ, Dai H, van de Vijver MJ, et al. Expression profiling predicts outcome in breast cancer. *Breast Cancer Res* 2003;5(1):57–8.
- [39] Dressman HK, Hans C, Bild A, et al. Gene expression profiles of multiple breast cancer phenotypes and response to neoadjuvant chemotherapy. *Clin Cancer Res* 2006;12(3 Pt 1): 819–26.
- [40] Potti A, Mukherjee S, Petersen R, et al. A genomic strategy to refine prognosis in early-stage non-small-cell lung cancer. *N Engl J Med* 2006;355(6):570–80.
- [41] Miller LD, Smeds J, George J, et al. An expression signature for p53 status in human breast cancer predicts mutation status, transcriptional effects, and patient survival. *Proc Natl Acad Sci U S A* 2005;102(38):13550–5.
- [42] Havaleshko DM, Cho H, Conaway M, et al. Prediction of drug combination chemosensitivity in human bladder cancer. *Mol Cancer Ther* 2007;6(2):578–86.
- [43] Paik S, Shak S, Tang G, et al. A multigene assay to predict recurrence of tamoxifen-treated, node-negative breast cancer. *N Engl J Med* 2004;351(27):2817–26.
- [44] Horvath S, Zhang B, Carlson M, et al. Analysis of oncogenic signaling networks in glioblastoma identifies ASPM as a molecular target. *Proc Natl Acad Sci U S A* 2006;103(46): 17402–7.
- [45] Bild AH, Yao G, Chang JT, et al. Oncogenic pathway signatures in human cancers as a guide to targeted therapies. *Nature* 2006;439(7074):353–7.
- [46] Potti A, Yao G, Chang JT, et al. Genomic signatures to guide the use of chemotherapeutics. *Nat Med* 2006;12(11):1294–300.
- [47] Ma XJ, Patel R, Wang X, et al. Molecular classification of human cancers using a 92-gene real-time quantitative polymerase chain reaction assay. *Arch Pathol Lab Med* 2006; 130(4):465–73.
- [48] Puskas LG, Juhasz F, Zarva A, et al. Gene profiling identifies genes specific for well-differentiated epithelial thyroid tumors. *Cell Mol Biol (Noisy-le-grand)* 2005;51(2):177–86.