Published in final edited form as:

Wiley Interdiscip Rev Comput Stat. 2013 ; 5(4): . doi:10.1002/wics.1260.

Identification of significant features in DNA microarray data

Eric Bair^{1,2,*}

¹Department of Endodontics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

²Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

Abstract

DNA microarrays are a relatively new technology that can simultaneously measure the expression level of thousands of genes. They have become an important tool for a wide variety of biological experiments. One of the most common goals of DNA microarray experiments is to identify genes associated with biological processes of interest. Conventional statistical tests often produce poor results when applied to microarray data owing to small sample sizes, noisy data, and correlation among the expression levels of the genes. Thus, novel statistical methods are needed to identify significant genes in DNA microarray experiments. This article discusses the challenges inherent in DNA microarray analysis and describes a series of statistical techniques that can be used to overcome these challenges. The problem of multiple hypothesis testing and its relation to microarray studies are also considered, along with several possible solutions.

Keywords

microarray; genetics; feature selection; multiple testing

INTRODUCTION

High-dimensional biological data sets have become increasingly common in recent years. Examples include data collected from DNA microarrays, comparative genome hybridization experiments, mass spectrometry, genome-wide association studies, and DNA/RNA sequencing. These new technologies have revolutionized our understanding of the genetics of human disease and numerous other biological processes. However, statistical analysis of such data sets is challenging for several reasons. These data sets are high dimensional, and the sample sizes are often small. Moreover, many of these data sets tend to be “noisy,” and the correlation between the features that are measured can be complex. For these reasons conventional statistical methods often produce unsatisfactory results when applied to modern high-dimensional biological data.

This study focuses on one of the most common problems in the analysis of high-dimensional biological data, which is the identification of significant genes in DNA microarray studies. This is one of the best-studied problems in the analysis of high-dimensional biological data

© 2013 Wiley Periodicals, Inc.

*Correspondence to: ebair@email.unc.edu.

Conflict of interest: The authors have declared no conflicts of interest for this article.

FURTHER READING Cui and Churchill¹³⁵ and Allison et al.³⁰ provide good overviews of feature selection methods for microarray data. Jeffery et al.⁴⁷ describe several of the most commonly used feature selection methods for microarrays and compare the performance of these methods on 9 publicly available data sets. There are also numerous books containing information on feature selection and other aspects of microarray data analysis not considered in this review. Good references include Causton et al.,¹³⁶ Parmigiani et al.,¹³⁷ Speed,¹³⁸ Wit and McClure,¹³⁹ McLachlan et al.,¹⁴⁰ Do et al.,¹⁴¹ and Draghici.¹⁴²

sets, and many of the methods that are applied to this problem may also be applied to other types of high-dimensional biological data. In a typical microarray study, one may wish to identify genes that are associated with a disease or some other biological process of interest. For example, one might attempt to identify genes associated with a disease by collecting a set of biological samples from diseased patients and another set of samples from healthy patients. Genes whose expression levels differ between the diseased and the control samples may be associated with the disease of interest. Alternatively, one might wish to identify genes that may be used to predict the prognosis of patients with a specific type of cancer. One might identify such genes by collecting tumor samples from a cohort of cancer patients and searching for genes whose expression levels are associated with the survival times of the patients. Ultimately, this information may be used for personalized treatment of cancer and other diseases. If the gene expression profile of a tumor indicates that the risk of metastasis is high, then the cancer should be treated more aggressively than another tumor whose gene expression suggests a low risk of metastasis.

This article consists of three main sections. In the first section, we will briefly describe DNA microarray technology and how DNA microarray data are collected. In the second section, we will provide a brief overview of some of the methods that have been used to identify significant genes in DNA microarray experiments. Numerous methods have been proposed in recent years, and space does not permit a detailed discussion of all possible methods. We have attempted to focus on several of the most commonly used approaches, along with an overview of some of the common principles and techniques used in these methods. We also briefly describe a few more recent methods for combining information across genes. In the final section, we discuss the problem of multiple hypothesis testing, which inevitably arises when identifying significant features in high-dimensional data sets.

DNA MICROARRAY DATA

Overview of Molecular Biology

Each organism's genetic information is contained in a molecule called deoxyribonucleic acid, more commonly known as DNA. DNA is a double-stranded molecule that is a chain of four possible nucleotides, namely adenine (A), cytosine (C), guanine (G), and thymine (T). The two strands of DNA are joined to one another by hydrogen bonds between nucleotides on the opposite strands. A always pairs with T, and G always pairs with C. Thus, if the sequence of one strand of DNA is known, then the sequence of the other strand is also known. Each such pair of bonded nucleotides is known as a base pair.

There are approximately 3.2 billion base pairs in the human genome (i.e., the entire sequence of DNA in a given human cell).¹ Different segments of DNA perform different functions, and much of the DNA performs no known function. The DNA segments of primary interest in most studies are the segments which contain instructions for building proteins. These segments are known as genes, and they comprise about 1.5% of the DNA sequence in humans.² Proteins perform most of the important functions in cells, including metabolism, DNA replication and repair, and communication with other cells.

The information contained in DNA is converted to proteins in a two-step process: In the first step, known as transcription, a given sequence of DNA is transcribed into an intermediary called messenger ribonucleic acid (mRNA), which is a single-stranded molecule that contains a copy of the complements of the base sequence of the DNA. The one difference is that thymine is replaced by uracil (U). In the second step, known as translation, the sequence of base pairs in the mRNA is translated into a protein, which is composed of a sequence of amino acids. Each set of three base pairs in the mRNA corresponds to one of 20 amino acids, a relationship that is known as the genetic code. This process by which the

information in the sequence of DNA is converted to mRNA and then to proteins is known as the fundamental dogma of molecular biology. See Dudoit et al.³ for a discussion of this process and its relationship to DNA microarray data.

DNA Microarray Technology

An important implication of the fundamental dogma of molecular biology is that there should be a strong association between the presence of a given protein in a cell and the presence of the mRNA sequence that is transcribed to build that protein. If a protein is active in a given cell, there should be a large number of copies of the mRNA sequence corresponding to that protein. Conversely, if a protein is not active in a cell, there should be few copies of the corresponding mRNA sequence. Thus, DNA microarrays attempt to evaluate the presence or absence of proteins in a cell and their relative abundance by measuring the relative abundance of the corresponding mRNA sequences.

DNA microarrays measure the relative abundance of mRNA sequences in the cells in a sample by taking advantage of complementary base pairing. Recall that in a DNA (or RNA) sequence, C always pairs with G and A always pairs with T (or U). A DNA microarray is typically constructed by placing an array of probes on a glass microscope slide. Each probe consists of a sequence of nucleotides that is complementary to the nucleotide sequence of a specific mRNA or its corresponding DNA sequence. Thus, one can measure the expression level of a given gene by measuring the amount of mRNA that hybridizes to the spot on the microarray corresponding to the gene. Different forms of DNA microarrays exist, such as oligonucleotide microarrays⁴ and cDNA microarrays^{5,6} but all of the most commonly used microarrays are based on this principle.

Figure 1 shows a typical (cDNA) microarray experiment. Two samples are collected, namely an experimental sample and a control sample. For example, the experimental sample may contain tissue from a cancerous tumor, and the control sample may contain noncancerous tissue from the same location in the body. First, mRNA is extracted from both samples. The extracted mRNA is treated with an enzyme called reverse transcriptase to convert it to a complementary DNA (or cDNA) sequence. Then each sample is treated with a fluorescent dye. Typically, the red dye Cy5 and the green dye Cy3 are used. Equal amounts of the two samples are then hybridized onto an array. To determine which genes are expressed at a high (or a low) level in the experimental group compared with the control group, one may measure the ratio of Cy5 to Cy3 at the probe on the array corresponding to that gene. For example, suppose that the experimental sample was treated with the red dye and the control sample was treated with the green dye. Then a red spot on the array indicates that there was more mRNA for that particular gene produced by the experimental group than the control group, indicating that this gene is expressed at a higher level in the control group. An image of a DNA microarray slide is shown in Figure 2.

Before microarray data is analyzed, the ratio of red dye to green dye at each spot on the array is measured using an appropriate scanner. This ratio is stored in a large data matrix. Typically, each row of the data matrix contains all the measured expression levels for a given gene, and each column of the data matrix corresponds to a particular sample. Such a data set is often visualized in the form of a heat map, as shown in Figure 3. The task of the microarray data analyst is to answer the biological question(s) of interest using this data matrix.⁷

It is important to attempt to remove extraneous variation in microarray data prior to data analysis. Variations in design of the arrays, sample preparation and scanner reading can produce “batch effects” where some subset of samples exhibit systematic differences in gene expression that are unrelated to the biological process of interest. Failure to account for such

batch effects can result in spurious findings.^{8,9} Thus, normalization is often necessary to remove batch effects. Numerous methods have been proposed to normalize microarray data.^{10–18} A detailed description of these normalization methods is beyond the scope of this review; see the aforementioned references for more information.

METHODS FOR IDENTIFYING SIGNIFICANT FEATURES IN DNA MICROARRAY DATA

Perhaps the most common objective of microarray experiments is to identify genes that are associated with a biological process of interest. For example, one may wish to identify genes associated with a disease of interest by comparing the expression levels of genes in diseased samples to the corresponding expression levels in healthy samples. Other outcomes of interest are also possible. For example, in cancer studies, one frequently wishes to identify genes associated with the survival time of cancer patients. The motivation is that genes that are associated with lower survival are likely to be associated with more serious forms of cancer that require more aggressive treatment.

In statistical terms, one has a large number of features (genes) and an outcome variable (e.g., disease versus control, survival time). The objective is to identify genes that are associated with the outcome variable. In principle, this objective can be accomplished using conventional statistical methods. To compare the expression of genes between two groups, one may calculate a *t*-test statistic for each gene. If there are three or more groups, an ANOVA *F*-test statistic may be used. To find genes associated with a continuous outcome variable, one may calculate a standardized regression coefficient, and to find genes associated with a survival outcome, one may calculate the Cox score for each gene.^{19,20}

However, these conventional methods often perform poorly on microarray data sets for several reasons, which will be discussed in more detail below. DNA microarray data sets are frequently noisy, and sample sizes are often small. Moreover, the gene expression levels are often highly correlated with one another, and failing to account for this fact may result in a loss of power. Also, one will typically perform several thousand hypothesis tests in a microarray experiment, so specialized methods are needed to control for type I error.

Throughout the remainder of this section, we will assume that one is comparing two different conditions using a *t*-test or a variation of the conventional *t*-test. However, the methods discussed below are easily generalized to other test statistics, such as ANOVA *F*-tests, standardized regression coefficients, and Cox scores.

Fold-Change Methods

One simple method for identifying differentially expressed features is to compute the average value of each feature under each condition and then compute the ratio of these averages. If the ratio exceeds some arbitrary cutoff, then the difference is called “significant.” For example, a gene may be called “significant” if the average expression level of a gene is more than twice as large (or less than half as large) in one condition compared with the other.

This approach has the benefit of simplicity, and it has been used in previous microarray studies.^{21,22} However, this method has some serious shortcomings. It is not based on a formal statistical test, so there is no simple way to calculate a *P*-value or confidence interval or other measure of the statistical validity of the association. Moreover, it is easy to see that this fold change has higher variance for genes expressed at lower levels, which is true of the majority of genes in microarray studies.^{23–25} For these reasons fold-change methods are

generally accepted to be inferior to other methods for identifying differentially expressed features.^{26–30}

t-Tests

An alternative approach is to identify significant genes based on a two-sample t -test of the null hypothesis that the mean expression level of the gene is the same under both conditions. This approach has also been used in microarray studies,³¹ and it has several advantages over fold-change methods. It is straightforward to calculate P -values and confidence intervals using t -tests, and for large samples the distribution of the t -statistic is independent of the overall expression level of the gene. In contrast, fold-change statistics have higher variance for genes expressed at low levels.

Unfortunately, using t -tests to identify differentially expressed genes can be problematic when the sample size is small, which is commonly the case in microarray experiments. It can be difficult to obtain accurate estimates of the variances of each group when the sample sizes are small. In particular, if the estimated variance of a gene is small, which occurs frequently when a gene is expressed at low levels,²⁵ then the gene may have a large t -statistic even if the fold change is small.

Alternative Versions of t-Tests

Given the shortcomings of t -tests described above, numerous authors have proposed alternative versions of t -tests for identifying significant features in gene expression data. Typically these methods combine data from all the genes to obtain a regularized estimator of the variance of a particular gene. In general, such variance estimates are biased. However, since the usual estimator of the variance has high variance when the sample size is low, a biased estimator of the variance may have lower prediction error than an unbiased estimator, since these biased estimators have lower variance than the unbiased estimator. This is especially true when the sample size is small. See Hastie et al.³² for a more detailed discussion of this phenomenon, which is commonly referred to as the “bias-variance trade-off.”

An example of the bias-variance trade-off is shown in Figure 4. Suppose the objective is to predict y based on x given the data in the figure. If one predicts y using a linear regression estimator based on x , the variance will be relatively low, but the bias will be high, since it cannot model the nonlinear relationship between x and y . At the other extreme, if one predicts y by interpolating the data with a smoothing spline, the bias will be 0, since the interpolation function can model any arbitrary relationship between x and y . However, the variance will be high, since the predicted values of y may change drastically if such a model is fit to a new data set.

Figure 5 shows how the bias and variance of a series of models varies as the complexity of the model increases. Each model in this figure represents a smoothing spline³³ fit to the data from Figure 4. As the complexity of the model increases, the variance of the model increases and the bias of the model decreases. One attempts to choose the model complexity that minimizes the expected prediction error or mean squared error (MSE), which can be shown to be equal to the sum of the variance, the square of the bias, and an irreducible error term due to unexplainable variance in y .³²

These figures illustrate why a regularized (i.e., biased) estimator of the variance of the genes may produce better results when identifying significant features based on microarray data. By regularizing the estimates of the variance, the complexity of each individual model is reduced, increasing the bias of the model but decreasing the variance. If the decrease in

variance is sufficient to offset the increase in bias, the accuracy of the overall model may be increased.

One possible approach is to estimate the variance of each gene by using the pooled estimator of the variance of all genes. Although this method has been used for several microarray studies,^{34–36} it also has some serious shortcomings. This obviously assumes that the variances of the expression levels of all genes are approximately the same, which is unlikely to be true in most situations. More importantly, since the denominator of the t -test will be the same for all genes, this method is essentially equivalent to the fold-change method, since it selects the genes with the largest mean differences without regard for the variance of an individual gene and thus suffers from the same drawbacks as fold-change methods. In terms of the bias-variance trade-off, this pooled variance estimate has low variance but high bias.

An alternative approach is to combine the variance estimator of each gene with some sort of pooled estimator of the variance across the genes. This avoids the high variance that results from estimating the variance of each gene individually as well as the high bias that results from relying entirely on a pooled variance estimate. For example, the “Significance Analysis of Microarrays” (SAM) procedure of Tusher et al.²⁵ uses the following test statistic:

$$t_i = \frac{\bar{X}_i - \bar{Y}_i}{s_i + s_0} \quad (1)$$

Here t_i represents the t -statistic for the i th gene, and $\bar{X}_i$ and $\bar{Y}_i$ represent the mean expression level of the gene under each experimental condition. The variance is estimated by summing the estimated variance of the i th gene (denoted by s_i) and a normalizing constant s_0 . This normalizing constant reduces the variance of the estimator of the variance and hence reduces the likelihood of obtaining false positive findings as a result of genes whose estimated variance is small. Typically, s_0 is chosen to be some quantile (such as the median) of the s_i 's across all of the genes. The SAM software is publicly available as an add-in for Microsoft Excel (<http://www-stat.stanford.edu/tibs/SAM/>). It is also implemented in the “samr” R package.

Other normalized estimators of the variance of microarray samples have been proposed. For example, Huber et al.³⁷ apply an arsinh transformation to the gene expression data that is designed to produce stable variance estimates irrespective of the gene's overall expression level. This method is implemented in the “vsn” R package (available through the Bioconductor project at <http://www.bioconductor.org>). Cui et al.³⁸ combine gene-specific and between-gene variance estimates using a James-Stein estimator.³⁹ R code for implementing this method is available at <http://www.stjude-research.org/depts/biostats/documents/cui-Fstat.R>. In general, any method to reduce the variance in the estimates of the variances of the individual genes can produce more accurate results when the sample size is small.

Bayesian Methods

Bayesian methods can also be used to combine information across genes to avoid inaccurate variance estimates as a result of small sample sizes. Typically, these methods impose some type of Bayesian prior distribution on the gene expression data and estimate the posterior distribution for each gene by combining information across all of the genes. For example, Baldi and Long⁴⁰ impose a prior distribution on the variances of the genes to obtain the following regularized t -test:

$$t_i = \frac{\bar{X}_i - \bar{Y}_i}{\sqrt{\frac{\nu_0 \sigma_0^2 + (n-1)s_i^2}{\nu_0 + n - 2}}} \quad (2)$$

In this expression $\bar{X}_i$, $\bar{Y}_i$, and s_i are defined as they were in Eq. (1). The parameter σ_0^2 is an estimator of the pooled variance across genes, which is calculated using data from all the genes, and ν_0 is a tuning parameter that controls the relative contributions of the gene-specific variance estimate and the global variance estimate. R code for implementing this method is available at <http://molgen51.biol.rug.nl/cybert/help/index.html>.

Note that Eq. (2) is similar to Eq. (1) in that the denominator of the t -statistic consists of a linear combination of an estimator of the variance of gene i plus a pooled estimate of the variance of all the genes. The similarity between the two expressions is not surprising. In general Bayesian methods tend to produce biased parameter estimates, but these estimators may have lower variance or mean-squared error than unbiased estimators, which is the same motivation for considering the regularized variance estimators discussed previously. Indeed, in some situations regularized frequentist parameter estimators can be shown to be Bayesian estimators with the appropriate choice of prior.⁴¹

Other similar Bayesian approaches have also been proposed for different types of microarray problems.^{42–46} In particular, the “limma” method of Smyth⁴⁵ uses an empirical Bayes test statistic that consistently performed well in a recent study comparing feature selection method for microarray data.⁴⁷

Calculating P -Values

If a t -test (or other conventional parametric test, such as ANOVA or regression) is used to test the null hypothesis of no association between the expression level of a given gene and an outcome, then calculating the P -value for this null hypothesis is straightforward if the assumptions of the test are satisfied. However, it may be dangerous to assume that these test statistics are normally distributed when the sample size is small. Moreover, as discussed previously, in many situations it is preferable to use biased estimators of the variance of a gene’s expression level. When a biased estimator of the variance is used, a t -statistic may no longer have a t distribution. Thus, alternative approaches may be needed to compute P -values in these situations.

One possible alternative is to calculate P -values based on the permutation distribution of the test statistic. Let t_j denote the t -statistic (or other test statistic) associated with gene j . Suppose the sample labels are then permuted K times, and let $t_{j,k}$ denote the test statistic associated with gene j for the k th permuted data set. Then one can estimate the P -value for gene j (denoted by P_j) as follows:

$$P_j = \frac{1}{K} \sum_{k=1}^K I(|t_{j,k}| > |t_j|) \quad (3)$$

Here $I(x)$ denotes an indicator function that is equal to 1 if the condition is true and 0 otherwise. In other words, the P -value is estimated by counting the number of times that the permuted version of the test statistic is “more extreme” than the original (unpermuted) version of the test statistic. A very large (or very small) test statistic is unlikely to occur by chance, so very few permuted data sets will produce a larger test statistic and the P -value will be small. This approach is used by the “SAM” software package²⁵ to calculate P -values.

This approach requires a choice of the number of permutations K . For small data sets, one may simply evaluate all possible permutations. In the case where one wishes to compare n_1 samples from one condition to n_2 samples from another condition using a t -test (or variant

thereof), there are a total of $\binom{n_1+n_2}{n_1}$ possible permutations. However, this would be computationally intractable for larger data sets, so it is common to arbitrarily select a value of $K = 1000$ or an even larger number if more precision is desired.

One possible problem with calculating P -values using Eq. (3) is that it can be difficult to estimate P -values that are close to 0. If $|t_j| > |t_{j,k}|$ for all k , then Eq. (3) implies that $P_j = 0$, when in reality all that can be inferred is that $P_j < 1/K$. This is problematic because certain methods for adjusting for multiple hypotheses testing in microarray experiments require precise estimation of P -values that are very close to 0. See below for more details.

There are a few possible solutions to this problem. The simplest approach is to increase the value of K . This will solve the problem given sufficient computing power, but it can be computationally intractable for large data sets. Another possibility is to pool the results of all the genes when calculating the permutation P -values. Suppose there are a total of N genes in the experiment. Then we estimate p_j as follows:

$$P_j = \frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K I(|t_{i,k}| > |t_j|) \quad (4)$$

In other words, rather than simply counting the number of times that the permuted test statistic for gene j is more extreme than the unpermuted test statistic for gene j , one counts the number of times that the permuted test statistic for any gene is greater than the permuted test statistic for gene j . This can increase the precision of the estimates of P_j without increasing the computational burden. See Hastie et al.³² for a complete discussion of calculating P -values based on the permutation distribution of the test statistics.

Methods for Combining Information Across Genes

The methods discussed thus far assume that hypothesis tests will be performed on each gene one at a time, and that the results of a hypothesis test on a given gene will not be affected by the hypothesis tests performed on other genes. This strategy may be inefficient on DNA microarray studies. Genes often act in pathways, meaning that several genes may be involved with the same biological process and hence be activated and deactivated simultaneously. If several related genes show evidence of differential expression at the same time, that is stronger evidence that the differential expression represents biological signal than if such a pattern were observed for a single gene. Several methods have been proposed for combining information across genes when searching for differentially expressed genes in microarray studies, which will be discussed below.

Biologically Motivated Methods—One approach for combining information across genes is to utilize known biological relationships among the genes. Typically genes are classified into groups using biological databases such as Gene Ontology.⁴⁸ Each group represents a set of biologically similar genes. The most commonly used methods compare the number of significant features in each group to the number expected if the genes in the group are not differentially expressed. If there are an unusually high number of significant features in a given group, suggesting that the pathway corresponding to the group is differentially expressed.

One strategy for identifying pathways containing differentially expressed genes is known as overrepresentation analysis (ORA). ORA first identifies a list of “significant” genes using any of the previously described methods for detecting differentially expressed genes. The M “most significant” genes are selected, which are typically the genes with the smallest P -values. Then for each group of genes, Fisher’s exact test (or some approximation thereof) is used to test the null hypothesis that the number of genes called significant in each group does not exceed the number of genes expected to be called significant due to chance. Various implementations of ORA have been proposed in the literature,^{49–54} and it has been used in some microarray experiments.⁵⁵

Despite the popularity of ORA, it has several shortcomings. Typically only the top M genes are used to compute the ORA statistics, resulting in the loss of any information available from genes not among the M most significant genes. The choice of M is often arbitrary as well. Moreover, all of the top M genes are treated equally, meaning that genes with extremely small univariate P -values are given the same weight as genes whose univariate P -values are much larger. Finally, ORA considers the gene to be the unit of analysis rather than the subject, which is inappropriate in virtually all real-world situations. Among other issues, it implies that the gene sets should be independent of one another, which is almost certainly not true in practice. See Pavlidis et al.,⁵⁶ Tian et al.,⁵⁷ or Allison et al.³⁰ for more information on the shortcomings of ORA.

An alternative strategy that avoids the problems associated with ORA is gene set enrichment analysis (GSEA). GSEA functions as follows: First, the genes are ordered according to their t -statistics or P -values or some other measure of univariate statistical significance. Then for each group of genes, the distribution of the t -statistics of the genes in the group is compared to the distribution of the genes not in the group using a one-sided Kolmogorov–Smirnov statistic or some other similar statistic. The idea is that if a group of genes is differentially expressed, then the distribution of the t -statistics among that set of genes should be different than the distribution of the t -statistics among the remaining genes. A P -value can be calculated for each set of genes by permuting the data multiple times and using Eq. (3) or (4) or alternative methods. Various implementations of GSEA have been proposed in the literature.^{56–65} There are also several software implementations of GSEA. For example, the Broad Institute offers software to perform GSEA in both Java and R (<http://www.broadinstitute.org/gsea/index.jsp>) and a variant of GSEA is implemented in the “SAM” software package.

The main shortcoming of GSEA is the fact that it tests a “competitive null hypothesis.” Suppose we have two groups of genes, which we will call gene group 1 and gene group 2. Then a smaller P -value for testing the null hypothesis of no differential expression in gene group 1 implies a larger P -value for testing this null hypothesis in gene group 2 even if the expression levels in gene group 2 remain unchanged. This occurs because the P -value for gene group 2 is calculated by comparing the test statistics of the genes in gene group 2 to the test statistics of all genes not in gene group 2, including the test statistics in gene group 1. Thus, if extreme test statistics are observed in gene group 1, this decreases the significance of gene group 2. See Damian and Gorfine⁶⁶ or Allison et al.³⁰ for a more detailed discussion of this phenomena. The development of methods for identifying groups of genes associated with an outcome of interest that avoids the shortcomings of ORA and GSEA is an active research area.

Statistically Based Methods—Other strategies for combining information across genes use novel statistical methods that do not require any knowledge of the biological relationship between the genes. We have previously discussed one possible statistical strategy for combining information across genes, namely regularized or Bayesian estimators of the

variance of individual genes. By using information about the variance of other genes to estimate the variance of a specific gene, the variance of the test statistic is greatly decreased, and hence the risk of false positives and false negatives is also decreased. However, in recent years several more advanced methods have been proposed for combining information across genes that we will briefly describe below.

One strategy is known as the optimal discovery procedure (ODP).^{67,68} The motivation for ODP is similar to the motivation for the pathway-based methods discussed previously. Since genes function in pathways, we expect that genes in the same pathway are likely to be co-expressed. Thus, if a gene shows evidence of differential expression, one can be more confident that the differential expression is not due to chance if other genes show a similar expression pattern. See Figure 6 for an illustration of this idea. The difference between ODP and pathway-based methods is that pathway-based methods require one to know in advance which genes are expected to be co-expressed based on previously collected biological data whereas ODP does not.

The ODP is a generalization of the Neyman–Pearson lemma⁶⁹ that states the most powerful test of a given null hypothesis against a given alternative hypothesis rejects the null hypothesis when the ratio

$$\frac{\text{Probability of the observed data under the alternative hypothesis}}{\text{Probability of the observed data under the null hypothesis}} \quad (5)$$

is large. The ODP generalizes the Neyman–Pearson lemma to situations where multiple hypotheses are tested by rejecting the null hypothesis that gene i is not differentially expressed when the ratio

$$\frac{\text{Sum of the probabilities of observing data } i \text{ under each alternative hypothesis}}{\text{Sum of the probabilities of observing data } i \text{ under each null hypothesis}} \quad (6)$$

is large. Thus, if a set of genes with similar expression patterns all show evidence of differential expression, then Eq. (6) will be larger than Eq. (5) for a given gene in the set, meaning that the null hypothesis of no differential expression is more likely to be rejected under ODP than under the traditional Neyman–Pearson paradigm.

In practice, Eq. (6) cannot be computed exactly and must be approximated.^{67,68} Software for computing the ODP is publicly available.⁷⁰

An alternative strategy for combining information across genes without any biological information about the relationship between the genes is the lassoed principal components (LPC) method of Witten and Tibshirani.⁷¹ The motivation for LPC is similar to the motivation for ODP: A gene is more likely to be differentially expressed if there are other genes with similar expression patterns than it is if there are no such similar genes. However, LPC uses a different strategy to determine if there are other genes with similar expression patterns. The idea behind LPC is that if a group of genes are co-regulated, then it is likely that a principal component of the gene expression matrix (sometimes called an eigenarray⁷²) will capture the variance in this group of genes. Thus, the LPC algorithm attempts to identify an eigenarray or group of eigenarrays that are associated with the biological process of interest and projects the t -statistics (or other relevant test statistics) onto this group of eigenarrays. This method can be shown to significantly reduce the false discovery rate in a variety of situations.⁷¹ This method is implemented in the “lpc” R package.

Clustering and Prediction Methods

Identifying features associated with an outcome of interest is not the only objective of microarray studies. One may also wish to partition the data into homogeneous subgroups and/or use the data to predict an outcome of interest. Clustering methods and prediction methods are useful in this situation.

There is a vast literature devoted to methods for clustering or predicting an outcome based on microarray data. A full description of such methods is beyond this scope of this review (which focuses on feature selection). However, it is noteworthy that there are methods for clustering^{73–85} and prediction^{20,86–95} that also perform feature selection. These methods generally do not evaluate whether a selected gene is “statistically significant” nor do they indicate which genes are the “most significant.” Also, the user of these methods often has limited control over the number of features selected. Thus, these methods have serious disadvantages if feature selection is the primary goal of the analysis. Nevertheless, these methods can identify a list of genes for further study, particularly in cases where clustering and prediction are important goals of the experiment.

Comparison of Feature Selection Methods

Numerous methods have been proposed for identifying significant features in DNA microarray data. However, the question of which methods produce the best results (i.e., maximize power while controlling type I error) has not been studied extensively. In practice researchers often choose feature selection methods based on the ease of implementing the method rather than the performance of the method. The “SAM” software package has become a popular tool for microarray analysis largely due to the fact that it is available as an Excel add-in and does not require the use of R or command-line programs.

Limited research indicates that the “limma” method of Smyth⁴⁵ performs well for a wide variety of problems, although other methods may perform better in specific situations.^{47,71,96} Limma is implemented in the “limma” R package, which is available from Bioconductor. Determining which feature selection method is likely to produce the best results on a given data set is an important area for future research.

ISSUES RELATED TO MULTIPLE HYPOTHESIS TESTING

Identifying significant genes in microarray studies requires performing a large number of hypothesis tests, which presents statistical challenges. When performing a single hypothesis test, it is conventional to choose a significance level such that the probability of rejecting the null hypothesis when it is true is equal to . However, when multiple hypothesis tests are performed, the probability of at least one false positive test will be much larger than . Thus, methods are needed to control the number of false positive tests while maintaining sufficient power to identify truly significant genes.

The Family-Wise Error Rate

One possible solution is to control the family-wise error rate (FWER) at a specified level. The FWER is defined to be the probability of rejecting at least one null hypothesis that is true. The most common way to control the FWER at a specified level is to use a Bonferroni correction. Each individual null hypothesis is rejected if and only if $P < \alpha/N$, where P is the P -value for the test and N is the total number of tests. It is easy to show that the probability of at least one type I error is no greater than using this procedure.

Although the Bonferroni correction controls the number of false positive tests, it is a very stringent criterion that typically results in a substantial loss of power. In experiments with

small sample sizes it is common for no tests to satisfy the Bonferroni criteria. Thus, most microarray analysts prefer less stringent approaches.³⁰ Methods exist for controlling the FWER using more permissive criteria than the Bonferroni correction,^{3,97} but these methods also suffer from lower power and are not commonly used.

The False Discovery Rate

The false discovery rate (FDR) is defined to be the expected proportion of false positives among the set of genes that are called significant. One may also adjust for multiple comparisons by controlling the FDR rather than the FWER. This approach typically yields greater power than FWER-based methods and hence is generally regarded as preferable.³⁰

The FDR was first proposed by Benjamini and Hochberg.⁹⁸ To control the FDR at a given level α , they proposed the following procedure: Let $P_{(1)} \leq P_{(2)} \leq \dots \leq P_{(N)}$ be the ordered P -values, and let $H_{(1)}, H_{(2)}, \dots, H_{(N)}$ be the corresponding null hypotheses. Then reject $H_{(1)}, H_{(2)}, \dots, H_{(j)}$, where

$$j = \max_i \{i : P_{(i)} \leq \alpha i / N\} \quad (7)$$

Benjamini and Hochberg⁹⁸ prove that the FDR of this procedure is at most α . This procedure is always valid if the P -values are independent. It remains valid in some cases even when dependency exists among the P -values, and methods exist for estimating the FDR where any type of dependency exists.^{99–111}

Rather than choosing a specific FDR in advance, one may wish to estimate the FDR when the top m genes are called significant. This is easy to do using the methodology of Benjamini and Hochberg⁹⁸: If we let

$$\hat{\alpha} = P_{(m)} N / m \quad (8)$$

Then Eq. (7) implies that the FDR should be approximately $\hat{\alpha}$.

If one estimates the null distribution of the test statistics using permutations of the data as in Eqs (3) and (4), then an alternative estimator of the false discovery rate may be used. Once again, let t_j denote the t -statistic (or other test statistic) associated with gene j , and let $t_{j,k}$ denote the test statistic associated with gene j for the k th permuted data set. Also, let $t_{(1)} \leq t_{(2)} \leq \dots \leq t_{(N)}$ be the order statistics of the absolute values of the t_j 's. Then one may estimate the FDR $\hat{\alpha}$ when the top m genes are called significant as follows:

$$\hat{\alpha} = \frac{1}{mK} \sum_{i=1}^N \sum_{k=1}^K I(|t_{i,k}| > t_{(N-m)}) \quad (9)$$

In other words, one estimates the FDR by dividing the average number of genes called significant over K permuted data sets by the number of genes called significant in the unpermuted data set (which is m , since the top m genes were called significant). It can be shown that $\hat{\alpha}$ in Eq. (9) is a consistent estimator of the FDR.^{100,112} Also, it can be shown that estimators (8) and (9) are equivalent.³² There are several R packages which will compute the FDR using this methodology (such as the “multtest” package, which is available from CRAN, and the “fdr” package, which is available from Bioconductor). This methodology is also implemented in the “SAM” software package.

The q -Value

In multiple testing problems, the q -value¹¹³ of a given test statistic t is defined to be the smallest possible FDR that can occur among all possible rejection regions that reject the null hypothesis when $T = t$. For example, if a t -statistic is calculated for each gene and the j th such t -statistic is t_j and $|t_j| = C$, then the q -value for the j th hypothesis test is the FDR for the rejection region $|t_j| \geq C$. In other words, the q -value is the FDR that results when one calls gene j significant along with all other genes that have a more extreme test statistic than gene j . Obviously genes with more extreme test statistics will have smaller q -values. The q -value may be estimated using Eq. (8) or (9), although other approaches are possible (see below). The q -value may be calculated using the “qvalue” R package (available from Bioconductor) as well as the “SAM” software package.

A Bayesian interpretation of the q -value is possible.^{112,113,114} Suppose that each gene comes from one of two populations, one of which consists of genes that are differentially expressed, and the other which consists of genes that are not differentially expressed. Under this assumption, the test statistic for each gene may be modeled using a mixture model. Define a set of random variables Z_j such that $Z_j = 0$ if gene j is not differentially expressed and $Z_j = 1$ if gene j is differentially expressed. Also let $|t_j| = C$, and let $q(t_j)$ be the q -value corresponding to gene j . Then one can show^{112,115} that

$$q(t_j) = P(Z_j = 0 | |t_j| \geq C) \quad (10)$$

In other words, under this mixture model, the q -value is the posterior probability that the j th null hypothesis is true given the test statistic for gene j . Although we have assumed a rejection region of the form $|t_j| \geq C$ in Eq. (10), this result holds under more general rejection regions.

Note: Equation (10) is only true if we calculate the P -value based on the positive false discovery rate (pFDR) as defined by Storey¹¹² rather than the traditional FDR proposed by Benjamini and Hochberg.⁹⁸ The pFDR is defined to be the expected proportion of false positives among the set of genes that are called significant conditional on the fact that at least one gene is called significant.

Methods exist for directly estimating the mixture distribution under the model described above and thereby estimating the pFDR/ q -value corresponding to individual genes using Eq. (10) or similar procedures.^{116–119} Limited research suggests that many of these methods produce comparable results.^{30,120} Such mixture models may also be used for omnibus testing.^{121,122}

CONCLUSION

Microarrays have been important for a variety of biological applications for over a decade. However, technology for generating high-throughput biological data is improving at a rapid pace, and technology that is commonly used today may be replaced in the near future. Indeed, some have suggested that newer technologies such as RNA-seq may soon replace microarrays¹²³ just as microarrays have largely replaced older techniques such as Northern blotting.¹²⁴ As technology advances, new methods will be necessary for analyzing data sets generated by the new techniques, and some methods for analyzing microarray data may no longer be useful in the future if microarrays are replaced by newer methods. Indeed, methods for identifying differentially expressed genes based on RNA-seq data is currently an active research area.^{125–130}

Despite these changing technologies, we feel that a discussion of methods for analyzing microarray data is still relevant and timely. DNA microarrays are still cheaper than RNA-seq assays, and RNA-seq gene expression measurements can be unreliable for genes expressed at lower levels.¹³¹ More importantly, however, many of the statistical techniques that have been developed for analyzing microarray data can also be applied to data produced by other high-throughput biological assays. For example, using normalized or Bayesian estimators of the variance of an estimator is useful for performing feature selection in any situation where the number of features is large and the number of observations is small. Similarly, use of the FDR and pFDR to control type I error is useful for a wide variety of multiple testing problems, which arise in the analysis of nearly all types of modern high-throughput biological data. For example, the “SAM” software package for DNA microarray analysis was recently upgraded to analyze RNA-seq data in addition to DNA microarray data. The new method continues to use resampling-based approaches to estimate the null distribution of each test statistic which is then used to estimate the FDR. See Li and Tibshirani¹³⁰ for details. Likewise, GSEA and other pathway-based methods for feature selection have been applied to genome-wide association studies.^{132–134} Thus, we see that the methods developed for DNA microarray analysis will be useful for many years in the future even as technology changes.

Acknowledgments

This work was partially supported by NIEHS grant P30ES010126 and NCATS grant UL1RR025747. We thank the four anonymous reviewers for their helpful suggestions.

REFERENCES

1. International Human Genome Sequencing Consortium. Finishing the euchromatic sequence of the human genome. *Nature*. 2004; 431:931–945. [PubMed: 15496913]
2. Lander ES, Linton LM, Birren B, Nusbaum C, Zody MC, Baldwin J, Devon K, Dewar K, Doyle M, FitzHugh W, et al. Initial sequencing and analysis of the human genome. *Nature*. 2001; 409:860–921. [PubMed: 11237011]
3. Dudoit S, Yang Y, Callow MJ, Speed TP. Statistical methods for identifying differentially expressed genes in replicated cDNA microarray experiments. *Stat Sin*. 2002; 12:111–139.
4. Lockhart DJ, Dong H, Byrne MC, Follettie MT, Gallo MV, Chee MS, Mittmann M, Wang C, Kobayashi M, Horton H, et al. Expression monitoring by hybridization to high-density oligonucleotide arrays. *Nat Biotechnol*. 1996; 14:1675–1680. [PubMed: 9634850]
5. DeRisi J, Penland L, Brown PO, Bittner ML, Meltzer PS, Ray M, Chen Y, Su YA, Trent JM. Use of a cDNA microarray to analyse gene expression patterns in human cancer. *Nat Genet*. 1996; 14:457–460. [PubMed: 8944026]
6. Hughes TR, Mao M, Jones AR, Burchard J, Marton MJ, Shannon KW, Lefkowitz SM, Ziman M, Schelter JM, Meyer MR, et al. Expression profiling using microarrays fabricated by an ink-jet oligonucleotide synthesizer. *Nat Biotechnol*. 2001; 19:342–347. [PubMed: 11283592]
7. Bullinger L, Döhner K, Bair E, Fröhling S, Schlenk RF, Tibshirani R, Döhner H, Pollack JR. Use of gene-expression profiling to identify prognostic subclasses in adult acute myeloid leukemia. *N Eng J Med*. 2004; 350:1605–1616. doi: 10.1056/NEJMoa031046.
8. Baggerly KA, Coombes KR, Neeley ES. Run batch effects potentially compromise the usefulness of genomic signatures for ovarian cancer. *J Clin Oncol*. 2008; 26:1186–1187. doi: 10.1200/JCO.2007.15.1951. [PubMed: 18309960]
9. Baggerly KA, Coombes KR. Deriving chemosensitivity from cell lines: forensic bioinformatics and reproducible research in high-throughput biology. *Ann Appl Stat*. 2009; 3:1309–1334.
10. Tseng GC, Oh MK, Rohlin L, Liao JC, Wong WH. Issues in cDNA microarray analysis: quality filtering, channel normalization, models of variations and assessment of gene effects. *Nucleic Acids Res*. 2001; 29:2549–2557. doi:10.1093/nar/29.12.2549. [PubMed: 11410663]

11. Yang YH, Dudoit S, Luu P, Lin DM, Peng V, Ngai J, Speed TP. Normalization for cDNA microarray data: a robust composite method addressing single and multiple slide systematic variation. *Nucleic Acids Res.* 2002; 30:e15. doi:10.1093/nar/30.4.e15. [PubMed: 11842121]
12. Quackenbush J. Microarray data normalization and transformation. *Nat Genet.* 2002; 32:496–501. [PubMed: 12454644]
13. Smyth GK, Speed T. Normalization of cDNA microarray data. *Methods.* 2003; 31:265–273. [PubMed: 14597310]
14. Bolstad B, Irizarry R, Åstrand M, Speed T. A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics.* 2003; 19:185–193. doi: 10.1093/bioinformatics/19.2.185. [PubMed: 12538238]
15. Irizarry RA, Warren D, Spencer F, Kim IF, Biswal S, Frank BC, Gabrielson E, Garcia JG, Geoghegan J, Germino G, et al. Multiple-laboratory comparison of microarray platforms. *Nat Methods.* 2005; 2:345–350. [PubMed: 15846361]
16. Brettschneider J, Collin F, Bolstad BM, Speed TP. Quality assessment for short oligonucleotide microarray data. *Technometrics.* 2008; 50:241–264. doi:10.1198/004017008000000334.
17. Stafford, P. *Methods in Microarray Normalization.* Chapman & Hall/CRC; Boca Raton, FL: 2008. Drug Discovery Series
18. Leek JT, Scharpf RB, Bravo HC, Simcha D, Langmead B, Johnson WE, Geman D, Baggerly K, Irizarry RA. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat Rev Genet.* 2010; 11:733–739. [PubMed: 20838408]
19. Beer DG, Kardia SL, Huang CC, Giordano TJ, Levin AM, Misek DE, Lin L, Chen G, Gharib TG, Thomas DG, et al. Gene-expression profiles predict survival of patients with lung adenocarcinoma. *Nat Med.* 2002; 8:816–824. [PubMed: 12118244]
20. Bair E, Tibshirani R. Semi-supervised methods to predict patient survival from gene expression data. *PLoS Biol.* 2004; 2:e108. doi:10.1371/journal.pbio.0020108. [PubMed: 15094809]
21. Schena M, Shalon D, Heller R, Chai A, Brown PO, Davis RW. Parallel human genome analysis: microarray-based expression monitoring of 1000 genes. *Proc Natl Acad Sci.* 1996; 93:10614–10619. [PubMed: 8855227]
22. DeRisi JL, Iyer VR, Brown PO. Exploring the metabolic and genetic control of gene expression on a genomic scale. *Science.* 1997; 278:680–686. [PubMed: 9381177]
23. Rocke DM, Durbin B. A model for measurement error for gene expression arrays. *J Comput Biol.* 2001; 8:557–569. [PubMed: 11747612]
24. Newton MA, Kendzierski CM, Richmond CS, Blattner FR, Tsui KW. On differential variability of expression ratios: improving statistical inference about gene expression changes from microarray data. *J Comput Biol.* 2001; 8:37–52. [PubMed: 11339905]
25. Tusher VG, Tibshirani R, Chu G. Significance analysis of microarrays applied to the ionizing radiation response. *Proc Natl Acad Sci.* 2001; 98:5116–5121. [PubMed: 11309499]
26. Chen Y, Dougherty ER, Bittner ML. Ratio-based decisions and the quantitative analysis of cDNA microarray images. *J Biomed Opt.* 1997; 2:364–374. doi: 10.1117/12.281504. [PubMed: 23014960]
27. Miller RA, Galecki A, Shmookler-Reis RJ. Interpretation, design, and analysis of gene array expression experiments. *J Gerontol A Biol Sci Med Sci.* 2001; 56:B52–B57. [PubMed: 11213267]
28. Budhraja V, Spitznagel E, Schaiff WT, Sadovsky Y. Incorporation of gene-specific variability improves expression analysis using high-density DNA microarrays. *BMC Biol.* 2003; 1:1. [PubMed: 14641937]
29. Hsiao A, Worrall DS, Olefsky JM, Subramaniam S. Variance-modeled posterior inference of microarray data: detecting gene-expression changes in 3 T3-L1 adipocytes. *Bioinformatics.* 2004; 20:3108–3127. [PubMed: 15217816]
30. Allison DB, Cui X, Page GP, Sabripour M. Microarray data analysis: from disarray to consolidation and consensus. *Nat Rev Genet.* 2006; 7:55–65. [PubMed: 16369572]
31. Callow MJ, Dudoit S, Gong EL, Speed TP, Rubin EM. Microarray expression profiling identifies genes with altered expression in HDL-deficient mice. *Genome Res.* 2000; 10:2022–2029. [PubMed: 11116096]

32. Hastie, T.; Tibshirani, R.; Friedman, J. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. 2 ed. Springer; New York: 2009. Springer Series in Statistics
33. Hastie, T.; Tibshirani, R. Generalized Additive Models. Chapman & Hall/CRC; Boca Raton, FL: 1990. Monographs on Statistics and Applied Probability Series
34. Tanaka TS, Jaradat SA, Lim MK, Kargul GJ, Wang X, Grahovac MJ, Pantano S, Sano Y, Piao Y, Nagaraja R, et al. Genome-wide expression profiling of midgestation placenta and embryo using a 15,000 mouse developmental cDNA microarray. *Proc Natl Acad Sci*. 2000; 97:9127–9132. [PubMed: 10922068]
35. Arfin SM, Long AD, Ito ET, Toller L, Riehle MM, Paegle ES, Hatfield GW. Global gene expression profiling in *Escherichia coli* K12. The effects of integration host factor. *J Biol Chem*. 2000; 275:29672–29684. [PubMed: 10871608]
36. Kerr MK, Martin M, Churchill GA. Analysis of variance for gene expression microarray data. *J Comput Biol*. 2000; 7:819–837. [PubMed: 11382364]
37. Huber W, von Heydebreck A, Sültmann H, Poustka A, Vingron M. Variance stabilization applied to microarray data calibration and to the quantification of differential expression. *Bioinformatics*. 2002; 18(suppl 1):S96–S104. doi:10.1093/bioinformatics/18.suppl_1.S96. [PubMed: 12169536]
38. Cui X, Hwang JTG, Qiu J, Blades NJ, Churchill GA. Improved statistical tests for differential gene expression by shrinking variance components estimates. *Biostatistics*. 2005; 6:59–75. doi:10.1093/biostatistics/kxh018. [PubMed: 15618528]
39. Stein CM. Confidence sets for the mean of a multivariate normal distribution. *J R Stat Soc B*. 1962; 24:265–296.
40. 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; 17:509–519. doi:10.1093/bioinformatics/17.6.509. [PubMed: 11395427]
41. Goldstein M. Bayesian analysis of regression problems. *Biometrika*. 1976; 63:51–58. doi:10.1093/biomet/63.1.51.
42. Lönnstedt I, Speed T. Replicated microarray data. *Stat Sin*. 2002; 12:31–46.
43. Kendziora CM, Newton MA, Lan H, Gould MN. On parametric empirical Bayes methods for comparing multiple groups using replicated gene expression profiles. *Stat Med*. 2003; 22:3899–3914. doi: 10.1002/sim.1548. [PubMed: 14673946]
44. Wright GW, Simon RM. A random variance model for detection of differential gene expression in small microarray experiments. *Bioinformatics*. 2003; 19:2448–2455. doi: 10.1093/bioinformatics/btg345. [PubMed: 14668230]
45. Smyth GK. Linear models and empirical Bayes methods for assessing differential expression in microarray experiments. *Stat Appl Genet Mol Biol*. 2004; 3 doi: 10.2202/1544-6115.1027.
46. Newton MA, Noueiry A, Sarkar D, Ahlquist P. Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*. 2004; 5:155–176. doi:10.1093/biostatistics/5.2.155. [PubMed: 15054023]
47. Jeffery I, Higgins D, Culhane A. Comparison and evaluation of methods for generating differentially expressed gene lists from microarray data. *BMC Bioinformatics*. 2006; 7:359. doi: 10.1186/1471-2105-7-359. [PubMed: 16872483]
48. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, et al. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat Genet*. 2000; 25:25–29. [PubMed: 10802651]
49. Dahlquist KD, Salomonis N, Vranizan K, Lawlor SC, Conklin BR. GenMAPP, a new tool for viewing and analyzing microarray data on biological pathways. *Nat Genet*. 2002; 31:19–20. [PubMed: 11984561]
50. Doniger SW, Salomonis N, Dahlquist KD, Vranizan K, Lawlor SC, Conklin BR. MAPPFinder: using gene ontology and GenMAPP to create a global gene-expression profile from microarray data. *Genome Biol*. 2003; 4:R7. [PubMed: 12540299]
51. Zeeberg BR, Feng W, Wang G, Wang MD, Fojo AT, Sunshine M, Narasimhan S, Kane DW, Reinhold WC, Lababidi S, et al. GoMiner: a resource for biological interpretation of genomic and proteomic data. *Genome Biol*. 2003; 4:R28. [PubMed: 12702209]

52. Draghici S, Khatri P, Bhavsar P, Shah A, Krawetz SA, Tainsky MA. Onto-tools, the toolkit of the modern biologist: onto-express, onto-compare, onto-design and onto-translate. *Nucleic Acids Res.* 2003; 31:3775–3781. [PubMed: 12824416]
53. Zhong S, Li C, Wong WH. ChipInfo: software for extracting gene annotation and gene ontology information for microarray analysis. *Nucleic Acids Res.* 2003; 31:3483–3486. [PubMed: 12824349]
54. Berriz GF, King OD, Bryant B, Sander C, Roth FP. Characterizing gene sets with FuncAssociate. *Bioinformatics.* 2003; 19:2502–2504. [PubMed: 14668247]
55. Blalock EM, Chen KC, Sharrow K, Herman JP, Porter NM, Foster TC, Landfield PW. Gene microarrays in hippocampal aging: statistical profiling identifies novel processes correlated with cognitive impairment. *J Neurosci.* 2003; 23:3807–3819. [PubMed: 12736351]
56. Pavlidis P, Qin J, Arango V, Mann JJ, Sibille E. Using the gene ontology for microarray data mining: a comparison of methods and application to age effects in human prefrontal cortex. *Neurochem Res.* 2004; 29:1213–1222. [PubMed: 15176478]
57. Tian L, Greenberg SA, Kong SW, Altschuler J, Kohane IS, Park PJ. Discovering statistically significant pathways in expression profiling studies. *Proc Natl Acad Sci U S A.* 2005; 102:13544–13549. [PubMed: 16174746]
58. Mootha VK, Lindgren CM, Eriksson KF, Subramanian A, Sihag S, Lehar J, Puigserver P, Carlsson E, Ridderstrale M, Laurila E, et al. PGC-1 α -responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nat Genet.* 2003; 34:267–273. [PubMed: 12808457]
59. Breitling R, Amtmann A, Herzyk P. Iterative Group Analysis (iGA): a simple tool to enhance sensitivity and facilitate interpretation of microarray experiments. *BMC Bioinformatics.* 2004; 5:34. [PubMed: 15050037]
60. Rahnenfuhrer J, Domingues FS, Maydt J, Lengauer T. Calculating the statistical significance of changes in pathway activity from gene expression data. *Stat Appl Genet Mol Biol.* 2004; 3 Article16.
61. Barry WT, Nobel AB, Wright FA. Significance analysis of functional categories in gene expression studies: a structured permutation approach. *Bioinformatics.* 2005; 21:1943–1949. [PubMed: 15647293]
62. Subramanian A, Tamayo P, Mootha VK, Mukherjee S, Ebert BL, Gillette MA, Paulovich A, Pomeroy SL, Golub TR, Lander ES, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci U S A.* 2005; 102:15545–15550. [PubMed: 16199517]
63. Zahn JM, Sonu R, Vogel H, Crane E, Mazan-Mamczarz K, Rabkin R, Davis RW, Becker KG, Owen AB, Kim SK. Transcriptional profiling of aging in human muscle reveals a common aging signature. *PLoS Genet.* 2006; 2:e115. [PubMed: 16789832]
64. Newton MA, Quintana FA, Boon JA, Sengupta S, Ahlquist P. Random-set methods identify distinct aspects of the enrichment signal in gene-set analysis. *Ann Appl Stat.* 2007; 1:85–106.
65. Efron B, Tibshirani R. On testing the significance of sets of genes. *Ann Appl Stat.* 2007; 1:107–129.
66. Damian D, Gorfine M. Statistical concerns about the GSEA procedure. *Nat Genet.* 2004; 36:663. [PubMed: 15226741]
67. Storey JD. The optimal discovery procedure: a new approach to simultaneous significance testing. *J R Stat Soc B.* 2007; 69:347–368. doi: 10.1111/j.1467-9868.2007.005592.x.
68. Storey JD, Dai JY, Leek JT. The optimal discovery procedure for large-scale significance testing, with applications to comparative microarray experiments. *Biostatistics.* 2007; 8:414–432. [PubMed: 16928955]
69. Neyman J, Pearson ES. On the problem of the most efficient tests of statistical hypotheses. *Philos Trans R Soc Lond A.* 1933; 231:289–337. doi:10.1098/rsta.1933.0009.
70. Leek JT, Monsen E, Dabney AR, Storey JD. EDGE: extraction and analysis of differential gene expression. *Bioinformatics.* 2006; 22:507–508. [PubMed: 16357033]
71. Witten DM, Tibshirani R. Testing significance of features by lassoed principal components. *Ann Appl Stat.* 2008; 2:986–1012. [PubMed: 19756232]

72. Alter O, Brown PO, Botstein D. Singular value decomposition for genome-wide expression data processing and modeling. *Proc Natl Acad Sci U S A*. 2000; 97:10101–10106. [PubMed: 10963673]
73. Xing EP, Karp RM. CLIFF: clustering of high-dimensional microarray data via iterative feature filtering using normalized cuts. *Bioinformatics*. 2001; 17:S306–S315. doi: 10.1093/bioinformatics/17.suppl_1.S306. [PubMed: 11473022]
74. Wang J, Bo T, Jonassen I, Myklebost O, Hovig E. Tumor classification and marker gene prediction by feature selection and fuzzy c-means clustering using microarray data. *BMC Bioinformatics*. 2003; 4:60. doi:10.1186/1471-2105-4-60. [PubMed: 14651757]
75. Tadesse MG, Sha N, Vannucci M. Bayesian variable selection in clustering high-dimensional data. *J Am Stat Assoc*. 2005; 100:602–617. doi:10.1198/016214504000001565.
76. Raftery AE, Dean N. Variable selection for model-based clustering. *J Am Stat Assoc*. 2006; 101:168–178. doi: 10.1198/016214506000000113.
77. Kim S, Tadesse MG, Vannucci M. Variable selection in clustering via Dirichlet process mixture models. *Biometrika*. 2006; 93:877–893. doi:10.1093/biomet/93.4.877.
78. Pan W, Shen X, Jiang A, Hebbel RP. Semi-supervised learning via penalized mixture model with application to microarray sample classification. *Bioinformatics*. 2006; 22:2388–2395. doi:10.1093/bioinformatics/btl393. [PubMed: 16870935]
79. Pan W, Shen X. Penalized model-based clustering with application to variable selection. *J Mach Learn Res*. 2007; 8:1145–1164.
80. Bondell HD, Reich BJ. Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. *Biometrics*. 2008; 64:115–123. doi:10.1111/j.1541-0420.2007.00843.x. [PubMed: 17608783]
81. Swartz MD, Mo Q, Murphy ME, Lupton JR, Turner ND, Hong M, Vannucci M. Bayesian variable selection in clustering high-dimensional data with substructure. *J Agric Biol Environ Stat*. 2008; 13:407–423. doi:10.1198/108571108X378317.
82. Wang S, Zhu J. Variable selection for model-based high-dimensional clustering and its application to microarray data. *Biometrics*. 2008; 64:440–448. doi:10.1111/j.1541-0420.2007.00922.x. [PubMed: 17970821]
83. Maugis C, Celeux G, Martin-Magniette ML. Variable selection for clustering with Gaussian mixture models. *Biometrics*. 2009; 65:701–709. doi:10.1111/j.1541-0420.2008.01160.x. [PubMed: 19210744]
84. Koestler DC, Marsit CJ, Christensen BC, Karagas MR, Bueno R, Sugarbaker DJ, Kelsey KT, Houseman EA. Semi-supervised recursively partitioned mixture models for identifying cancer subtypes. *Bioinformatics*. 2010; 26:2578–2585. doi:10.1093/bioinformatics/btq470. [PubMed: 20834038]
85. Witten DM, Tibshirani R. A framework for feature selection in clustering. *J Am Stat Assoc*. 2010; 105:713–726. doi:10.1198/jasa.2010.tm09415. [PubMed: 20811510]
86. Tibshirani R, Hastie T, Narasimhan B, Chu G. Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proc Natl Acad Sci*. 2002; 99:6567–6572. doi:10.1073/pnas.082099299. [PubMed: 12011421]
87. Sha N, Vannucci M, Tadesse MG, Brown PJ, Dragoni I, Davies N, Roberts TC, Contestabile A, Salmon M, Buckley C, et al. Bayesian variable selection in multinomial probit models to identify molecular signatures of disease stage. *Biometrics*. 2004; 60:812–819. doi:10.1111/j.0006-341X.2004.00233.x. [PubMed: 15339306]
88. Bair E, Hastie T, Paul D, Tibshirani R. Prediction by supervised principal components. *J Am Stat Assoc*. 2006; 101:119–137. doi:10.1198/016214505 000000628.
89. Wu B. Differential gene expression detection and sample classification using penalized linear regression models. *Bioinformatics*. 2006; 22:472–476. doi:10.1093/bioinformatics/bti827. [PubMed: 16352654]
90. Tai F, Pan W. Incorporating prior knowledge of gene functional groups into regularized discriminant analysis of microarray data. *Bioinformatics*. 2007; 23:3170–3177. doi:10.1093/bioinformatics/btm488. [PubMed: 17933851]

91. Wang S, Zhu J. Improved centroids estimation for the nearest shrunken centroid classifier. *Bioinformatics*. 2007; 23:972–979. doi:10.1093/bioinformatics/btm046. [PubMed: 17384429]
92. Guo Y, Hastie T, Tibshirani R. Regularized linear discriminant analysis and its application in microarrays. *Biostatistics*. 2007; 8:86–100. doi:10.1093/biostatistics/kxj035. [PubMed: 16603682]
93. Paul D, Bair E, Hastie T, Tibshirani R. “Preconditioning” for feature selection and regression in high-dimensional problems. *Ann Stat*. 2008; 36:1595–1618.
94. Guo J. Simultaneous variable selection and class fusion for high-dimensional linear discriminant analysis. *Biostatistics*. 2010; 11:599–608. doi:10.1093/biostatistics/kxq023. [PubMed: 20505058]
95. Stingo FC, Vannucci M. Variable selection for discriminant analysis with markov random field priors for the analysis of microarray data. *Bioinformatics*. 2011; 27:495–501. doi:10.1093/bioinformatics/btq690. [PubMed: 21159623]
96. Murie C, Woody O, Lee A, Nadon R. Comparison of small n statistical tests of differential expression applied to microarrays. *BMC Bioinformatics*. 2009; 10:45. doi: 10.1186/1471-2105-10-45. [PubMed: 19192265]
97. Wu, H.; Kerr, M.; Cui, X.; Churchill, G. MAANOVA: A software package for the analysis of spotted cDNA microarray experiments. In: Parmigiani, G.; Garrett, E.; Irizarry, R.; Zeger, S., editors. *The Analysis of Gene Expression Data. Statistics for Biology and Health*. Springer; London: 2003. p. 313–341.
98. Benjamini Y, Hochberg Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J R Stat Soc B*. 1995; 57:289–300.
99. Benjamini Y, Yekutieli D. The control of the false discovery rate in multiple testing under dependency. *Ann Stat*. 2001; 29:1165–1188.
100. Storey JD, Taylor JE, Siegmund D. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *J R Stat Soc B*. 2004; 66:187–205. doi:10.1111/j.1467-9868.2004.00439.x.
101. Farcomeni A. More powerful control of the false discovery rate under dependence. *Stat Methods Appl*. 2006; 15:43–73.
102. Meinshausen N. False discovery control for multiple tests of association under general dependence. *Scand J Stat*. 2006; 33:227–237. doi:10.1111/j.1467-9469.2005.00488.x.
103. Pawitan Y, Calza S, Ploner A. Estimation of false discovery proportion under general dependence. *Bioinformatics*. 2006; 22:3025–3031. doi:10.1093/bioinformatics/btl527. [PubMed: 17046978]
104. Efron B. Correlation and large-scale simultaneous significance testing. *J Am Stat Assoc*. 2007; 102:93–103.
105. Finner H, Dickhaus T, Roters M. Dependency and false discovery rate: Asymptotics. *Ann Stat*. 2007; 35:1432–1455.
106. Leek JT, Storey JD. A general framework for multiple testing dependence. *Proc Natl Acad Sci*. 2008; 105:18718–18723. doi:10.1073/pnas.0808709105. [PubMed: 19033188]
107. Romano JP, Shaikh AM, Wolf M. Control of the false discovery rate under dependence using the bootstrap and subsampling. *Test*. 2008; 17:417–442.
108. Sun W, Tony Cai T. Large-scale multiple testing under dependence. *J R Stat Soc B*. 2009; 71:393–424. doi:10.1111/j.1467-9868.2008.00694.x.
109. Clarke S, Hall P. Robustness of multiple testing procedures against dependence. *Ann Stat*. 2009; 37:332–358.
110. Friguet C, Kloareg M, Causeur D. A factor model approach to multiple testing under dependence. *J Am Stat Assoc*. 2009; 104:1406–1415.
111. Fan J, Han X, Gu W. Estimating false discovery proportion under arbitrary covariance dependence. *J Am Stat Assoc*. 2012; 107:1019–1035. doi:10.1080/01621459.2012.720478.
112. Storey JD. A direct approach to false discovery rates. *J R Stat Soc B*. 2002; 64:479–498. doi: 10.1111/1467-9868.00346.
113. Storey JD. The positive false discovery rate: a Bayesian interpretation and the q-value. *Ann Stat*. 2003; 31:2013–2035.

114. Efron B, Tibshirani R. Empirical Bayes methods and false discovery rates for microarrays. *Genet Epidemiol.* 2002; 23:70–86. doi:10.1002/gepi.1124. [PubMed: 12112249]
115. Efron B, Tibshirani R, Storey JD, Tusher V. Empirical Bayes analysis of a microarray experiment. *J Am Stat Assoc.* 2001; 96:1151–1160. doi:10.1198/016214501753382129.
116. Allison DB, Gadbury GL, Heo M, Fernández JR, Lee CK, Prolla TA, Weindrich R. A mixture model approach for the analysis of microarray gene expression data. *Comput Stat Data Anal.* 2002; 39:1–20. doi:10.1016/S0167-9473(01)00046-9.
117. Pounds S, Morris SW. Estimating the occurrence of false positives and false negatives in microarray studies by approximating and partitioning the empirical distribution of p-values. *Bioinformatics.* 2003; 19:1236–1242. doi:10.1093/bioinformatics/btg148. [PubMed: 12835267]
118. Do KA, Müller P, Tang F. A Bayesian mixture model for differential gene expression. *J R Stat Soc C.* 2005; 54:627–644. doi:10.1111/j.1467-9876.2005.05593.x.
119. Efron B. Microarrays, empirical Bayes and the two-groups model. *Stat Sci.* 2008; 23:1–22.
120. Datta S, Datta S. Empirical Bayes screening of many p-values with applications to microarray studies. *Bioinformatics.* 2005; 21:1987–1994. doi:10.1093/bioinformatics/bti301. [PubMed: 15691856]
121. Dai H, Charnigo R. Omnibus testing and gene filtration in microarray data analysis. *J Appl Stat.* 2008; 35:31–47. doi:10.1080/02664760701683528.
122. Dai H, Charnigo R. Contaminated normal modeling with application to microarray data analysis. *Can J Stat.* 2010; 38:315–332. doi:10.1002/cjs.10053.
123. Shendure J. The beginning of the end for microarrays? *Nat Methods.* 2008; 5:585–587. [PubMed: 18587314]
124. Taniguchi M, Miura K, Iwao H, Yamanaka S. Quantitative assessment of DNA microarrays—comparison with Northern blot analyses. *Genomics.* 2001; 71:34–39. [PubMed: 11161795]
125. Wang L, Feng Z, Wang X, Wang X, Zhang X. DEGseq: an R package for identifying differentially expressed genes from RNA-seq data. *Bioinformatics.* 2010; 26:136–138. [PubMed: 19855105]
126. Robinson MD, Oshlack A. A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol.* 2010; 11:R25. [PubMed: 20196867]
127. Hardcastle TJ, Kelly KA. baySeq: empirical Bayesian methods for identifying differential expression in sequence count data. *BMC Bioinformatics.* 2010; 11:422. [PubMed: 20698981]
128. Anders S, Huber W. Differential expression analysis for sequence count data. *Genome Biol.* 2010; 11:R106. [PubMed: 20979621]
129. Bullard JH, Purdom E, Hansen KD, Dudoit S. Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments. *BMC Bioinformatics.* 2010; 11:94. [PubMed: 20167110]
130. Li J, Tibshirani R. Finding consistent patterns: a nonparametric approach for identifying differential expression in RNA-Seq data. *Stat Methods Med Res.* 2011
131. Łabaj PP, Leparć GG, Linggi BE, Markillie LM, Wiley HS, Kreil DP. Characterization and improvement of RNA-Seq precision in quantitative transcript expression profiling. *Bioinformatics.* 2011; 27:i383–i391. [PubMed: 21685096]
132. Medina I, Montaner D, Bonifaci N, Pujana MA, Carbonell J, Tarraga J, Al-Shahrour F, Dopazo J. Gene set-based analysis of polymorphisms: finding pathways or biological processes associated to traits in genome-wide association studies. *Nucleic Acids Res.* 2009; 37:W340–W344. [PubMed: 19502494]
133. Zhang K, Cui S, Chang S, Zhang L, Wang J. i-GSEA4GWAS: a web server for identification of pathways/gene sets associated with traits by applying an improved gene set enrichment analysis to genome-wide association study. *Nucleic Acids Res.* 2010; 38:W90–W95. [PubMed: 20435672]
134. Nam D, Kim J, Kim SY, Kim S. GSA-SNP: a general approach for gene set analysis of polymorphisms. *Nucleic Acids Res.* 2010; 38:W749–W754. [PubMed: 20501604]
135. Cui X, Churchill G. Statistical tests for differential expression in cDNA microarray experiments. *Genome Biol.* 2003; 4:210. doi:10.1186/gb-2003-4-4-210. [PubMed: 12702200]

136. Causton, HC.; Quackenbush, J.; Brazma, A. Microarray Gene Expression Data Analysis: A Beginner's Guide. Blackwell Publishing; Malden, MA: 2003.
137. Parmigiani, G.; Garrett, ES.; Irizarry, RA.; Zeger, SL. The Analysis of Gene Expression Data: Methods and Software. Springer; New York: 2003.
138. Speed, T. Statistical Analysis of Gene Expression Microarray Data. Interdisciplinary Statistics. Chapman & Hall/CRC; Boca Raton, FL: 2003.
139. Wit, E.; McClure, J. Statistics for Microarrays: Design, Analysis and Inference. John Wiley & Sons; Chichester: 2004.
140. McLachlan, G.; Do, K.; Ambrose, C. Analyzing Microarray Gene Expression Data. John Wiley & Sons; Hoboken, NJ: 2005. Wiley Series in Probability and Statistics
141. Do, K.; Müller, P.; Vannucci, M. Bayesian Inference for Gene Expression and Proteomics. Cambridge University Press; New York: 2006.
142. Draghici, S. Statistics and Data Analysis for Microarrays Using R and Bioconductor. 2nd ed. Chapman & Hall/CRC; Boca Raton, FL: 2011. Mathematical and Computational Biology Series

FIGURE 1.

Illustration of a typical microarray experiment (using cDNA technology). First, mRNA is extracted from two groups of cells, namely an experimental sample of interest and a control sample. Each sample is labeled with a different color of fluorescent dye. The samples are then combined and hybridized onto an array. The relative abundance of the mRNA corresponding to a particular gene can be measured by calculating the ratio of red dye to green dye at the appropriate spot on the array.

FIGURE 2.

Image of a DNA microarray slide. One may measure the relative gene expression of each gene by comparing the ratio of the amount of red dye to the amount of green dye at each probe on the array.

FIGURE 3.

Heat map of the leukemia microarray data of Bullinger et al.⁷ Each colored square on the map corresponds to the expression level of a given gene for a given patient. In the above figure, each row represents a gene and each column represents a patient. The brighter the color of a given square, the higher (or lower) the expression level of the corresponding gene. Usually hierarchical clustering is performed on the rows and columns of the data set prior to drawing the heat map.

FIGURE 4.

Illustration of the bias-variance trade-off. The above figure shows a regression problem where the objective is to predict y given a value of x . The dotted line shows the true relationship between x and y . The linear regression estimator (shown in blue) has high bias and low variance, and the interpolation estimator (shown in orange) has low bias and high variance.

FIGURE 5.

Illustration of the association between the complexity of a model and the bias/variance of the model. In general, as the complexity of a model increases, the variance of the model increases and the bias of the model decreases.

FIGURE 6.

Illustration of the optimal discovery procedure (ODP). Suppose that the test statistic for the null hypothesis of no differential expression is $t = -2$ for one gene and $t = 2$ for a second gene. Suppose further that there are several other genes with similar expression patterns to the second gene for which $t \approx 2$. Using traditional hypothesis testing procedures, one would be equally likely to reject the null hypothesis of no differential expression for both of the two genes. Using ODP, one would be more likely to reject the null hypothesis for the gene where $t = 2$, since the existence of several genes with similar expression patterns increases one's confidence that the result is not due to chance.