



Available at

www.ElsevierMathematics.com

POWERED BY SCIENCE @ DIRECT®

Journal of Statistical Planning and
Inference 124 (2004) 379–398

journal of
statistical planning
and inference

www.elsevier.com/locate/jspi

Controlling the number of false discoveries: application to high-dimensional genomic data

Edward L. Korn^{a,*}, James F. Troendle^b, Lisa M. McShane^a,
Richard Simon^a

^a*Biometric Research Branch, EPN-8128, National Cancer Institute, Bethesda, MD 20892, USA*

^b*Biometry and Mathematical Statistics Branch, National Institute of Child Health and Human
Development, Bethesda, MD 20892, USA*

Received 8 November 2001; accepted 6 April 2003

Abstract

Researchers conducting gene expression microarray experiments often are interested in identifying genes that are differentially expressed between two groups of specimens. A straightforward approach to the identification of such “differentially expressed” genes is to perform a univariate analysis of group mean differences for each gene, and then identify those genes that are most statistically significant. However, with the large number of genes typically represented on a microarray, using nominal significance levels (unadjusted for the multiple comparisons) will lead to the identification of many genes that truly are not differentially expressed, “false discoveries.” A reasonable strategy in many situations is to allow a small number of false discoveries, or a small proportion of the identified genes to be false discoveries. Although previous work has considered control for the *expected* proportion of false discoveries (commonly known as the false discovery rate), we show that these methods may be inadequate. We propose two stepwise permutation-based procedures to control with specified confidence the *actual* number of false discoveries and approximately the *actual* proportion of false discoveries. Limited simulation studies demonstrate substantial gain in sensitivity to detect truly differentially expressed genes even when allowing as few as one or two false discoveries. We apply these new methods to analyze a microarray data set consisting of measurements on approximately 9000 genes in paired tumor specimens, collected both before and after chemotherapy on 20 breast cancer patients. The methods described are broadly applicable to the problem of identifying which variables of any large set of measured variables differ between pre-specified groups.

Published by Elsevier B.V

Keywords: False discovery; Multiple comparisons; Stepwise procedure; Permutation method; Gene expression; Microarray

* Corresponding author. Tel.: +1-301-4020635; fax: +1-301-4020560.

E-mail address: korne@ctep.nci.nih.gov (E.L. Korn).

1. Introduction

Technological advances have made possible detailed genetic characterization of biological specimens. For example, cDNA microarray technology (Schena et al., 1995) permits the simultaneous evaluation of expression levels of thousands of genes on a single specimen, generating a gene expression “profile” for that specimen. The cDNA microarray technology has now been applied to profile numerous human cancer specimens; it is hoped that gene expression profiles of tumors might aid in distinguishing aggressive from indolent tumors and might guide choice of therapies. Another exciting opportunity comes with the completion of the initial sequencing and analysis of the human genome. More than 1.4 million single nucleotide polymorphisms (SNPs) in the human DNA sequence have been now identified (International Human Genome Sequencing Consortium, 2001). Differing patterns of SNPs might be related to risk of developing disease or predict response to, or toxicity from, drug therapies. A typical experimental approach in these settings would be to select specimens from two or more groups it was desired to compare, and then to measure a large number of characteristics on each specimen. Often one would want to identify characteristics that univariately are significantly different between the groups. The issue we address in this paper is how one can identify individual characteristics that are significantly different between groups of specimens while maintaining control over spurious findings amid the potentially enormous number of comparisons being made.

The particular example we consider in this paper consists of gene expression profiles obtained by cDNA microarray analysis of approximately 9000 genes for 40 paired breast tumor specimens. The specimens were collected on 20 breast cancer patients, before and after chemotherapy. Our interest is in identifying genes whose expression levels differed significantly after chemotherapy as compared to before. The example data are continuous, log-transformed expression ratios that measure the relative abundance of each gene’s mRNA in the test specimen compared to a reference sample using a two-color fluorescent probe hybridization system (Schena et al., 1995). We note, however, that the general approaches described in this paper are very broadly applicable to both continuous and discrete data, and to censored data.

If one were to simply conduct univariate tests of characteristics, for example gene expression levels, using conventional significance levels, there would be an enormous multiple comparisons problem. Many genes would likely be claimed significantly differentially expressed when, in truth, they were not differentially expressed, i.e., a type I error. Such false claims of significance are often called “false discoveries.” One can use a procedure to account for these multiple comparisons and control the probability of any false discovery. This overall probability of any error is usually referred to as the familywise error (FWE) rate. A Bonferroni adjustment to the p -values, or preferably less conservative stepwise procedures such as those described by Westfall and Young (1993, pp. 72–74) and Troendle (1996) can be used to control the familywise error rate. For example, Callow et al. (2000) have applied the Westfall and Young method to microarray data. These procedures will guarantee that the probability of any false discovery is less than the designated significance level, e.g., 0.05. However, the criterion of not making any false discovery is too stringent for most microarray

Table 1
 Simulated distribution of the number of false discoveries for 10,000 hypotheses tested using univariate nominal significance levels of 0.001

Correlation, ρ	Mean	Percentile				
		10th	25th	50th	75th	90th
0	10.0	6	8	10	12	14
0.5	9.9	3	5	8	13	18
0.8	9.9	0	1	4	12	27

Note: Test statistics are simulated under a global null as standard normal with block diagonal correlation matrix, block size 100 and pairwise correlation ρ within a block. The number of simulated data sets is 10,000.

investigations, in which the identification of these genes will be followed by further study of them for biological relevance. On the other hand, making no adjustment for multiple comparisons could generate many false leads.

A reasonable compromise is to use a procedure that will allow some false discoveries, but not too many. A simple procedure is to lower the nominal significance level and appeal to Bonferroni, e.g., using a significance level of 0.001 would ensure in expectation at most 10 false discoveries with 10,000 variables. A slightly more complex procedure attempts to control for the expected proportion of discoveries (identified genes) that are false discoveries (with the proportion set to 0 when no genes are identified): Order the univariate p -values from the k variables, $P_{(1)} < P_{(2)} < \dots < P_{(k)}$. To keep the expected false discovery proportion less than γ (e.g., $\gamma = 0.10$), identify as differentially expressed those genes that are associated with the indices $1, 2, \dots, i$, where i is the largest index satisfying $P_{(i)}k < i\gamma$. This procedure is attributed to Eklund by Seeger (1968) and was studied by Benjamini and Hochberg (1995). Tusher et al. (2001) present a procedure they call significance analysis of microarrays (SAM) for estimating a false discovery rate from data. Their procedure is based on estimating the expected number of false positives from a complete null permutation distribution; they do not discuss the statistical properties of their procedure. Procedures targeting control of the expected number or proportion of false discoveries rather than the actual number or proportion can give a false sense of security. This is demonstrated in Tables 1 and 2 for simulated data with 10,000 variables. In Table 1, we consider using a univariate nominal significance level of 0.001. The expected number of false discoveries is less than or equal to 10, but the spread of the distribution of the actual number of false discoveries becomes quite large when the correlation between the variables increases. For example, there is a 10% chance of having 18 or more false discoveries with block correlation 0.5. The same problem arises when controlling the expected false discovery proportion. Table 2 displays the distribution of the actual false discovery proportion when using the simple procedure described above to control the expected false discovery proportion to be less than 0.10. Even with no correlation the results here are troubling: 10% of the time the false discovery proportion will be 0.29 or more.

Table 2

Simulated distribution of the false discovery proportion for 10,000 hypotheses tested using a simple step-up procedure to control the expected false discovery rate to be no more than 10%

Correlation, ρ	Mean	Percentile				
		10th	25th	50th	75th	90th
0	0.098	0	0	0	0.18	0.29
0.5	0.090	0	0	0	0.13	0.30
0.8	0.055	0	0	0	0	0.17

Note: Simulated test statistics are normally distributed with block diagonal correlation matrix, block size 100, and pairwise correlation ρ within a block. The first 10 test statistics have standardized mean of 4, and the remaining have mean zero. The number of simulated data sets is 10,000. The step-up procedure applied is that described and studied by [Benjamini and Hochberg \(1995\)](#).

In this paper, we discuss methods for controlling the actual (rather than expected) number and proportion of false discoveries. We prove that our procedure for the former guarantees control with specified confidence, and provide evidence that our procedure for the latter achieves approximate control. That is, application of these methods will allow statements such as “with 95% confidence, the number of false discoveries does not exceed 2” or “with approximate 95% confidence, the proportion of false discoveries does not exceed 0.10”. In Section 2, we describe our procedures for controlling the number and proportion of false discoveries and provide justifications for the algorithms which appear in Appendix A. The methods are applied to the analysis of the pre-post chemotherapy breast cancer specimens in Section 3. In Section 4, we describe some limited simulation studies to assess our procedures and to compare them to procedures designed to control the FWE rate. We end with a discussion in Section 5.

2. Controlling false discoveries

We consider two probabilistic frameworks for inference. The first, a population model, specifies that for each group the k -dimensional vector of variables associated with a specimen is independent and identically distributed with a multivariate distribution that can depend on the group. For unpaired data, the multidimensional distribution of the subvector of null hypothesis variables is assumed to be the same in the different groups. For paired data, the pairs represent population strata, and the multivariate distribution of this subvector is assumed to be the same within each pair regardless of group membership. The other framework we consider is a multivariate permutation model. In this model, every permutation of the subvector of null hypothesis variables into the different groups is assumed equally likely. For example, in an unpaired two-group comparison with n_1 and n_2 observations in the groups, all $(n_1 + n_2)! / (n_1! n_2!)$ assignments of group labels of subvectors of null hypothesis variables are equally likely. In applications we have in mind the population model is probably more relevant. However, the null hypothesis population model implies the null hypothesis permutation model ([Lehmann, 1975](#), pp. 55–65), and therefore results concerning type I error control

using the permutation model will be more general than using the population model. We use the permutation model for the proofs below, and probability statements refer to this model. For the simulations in Section 4 we use the population model.

We assume that for each variable we have performed an appropriate univariate statistical test of the null hypothesis that the distribution of that variable is the same across the groups, and we have obtained a p -value associated with that test. For example, if the data consist of continuous gene expression levels, t -tests or Wilcoxon tests might be performed for each gene. If the data consist of binary variables, each indicating the presence or absence of a particular marker, then chi-squared tests or Fisher’s exact tests may be used. The tests may be for paired data (as in our breast cancer example) or unpaired data depending on the design. Let

$$P_{(1)} < P_{(2)} < \dots < P_{(k)} \tag{1}$$

be the ordered p -values from the univariate tests based on the k variables, and let $H_{(1)}, H_{(2)}, \dots, H_{(k)}$ denote the null hypotheses in the corresponding order. For now, we assume that there are no ties in the p -values (as would be the case if the variables were continuous and parametric hypothesis tests were used). For any subset $T = \{t_1, t_2, \dots, t_j\}$ of $K = \{1, 2, \dots, k\}$, consider the multivariate permutation distribution of p -values $\{P'_{t_1}, P'_{t_2}, \dots, P'_{t_j}\}$ that would be generated by permuting the group labels on the specimens and calculating the test statistics on the variables with indices in T using the permuted data. For paired data, these permutations are obtained by switching the characteristic profiles within each pair, yielding 2^n possible permutations for n pairs of specimens. For the unpaired or multi-group case, permutations are performed by shuffling the group membership labels. Note that in each case, the characteristic profiles measured on any given specimen remain intact so as to preserve the correlation among the measured characteristics. With a small number of specimens, it may be possible to enumerate all possible permutations. However, typically the number of permutations is very large, so they are randomly sampled. For example, for our paired breast tumor example, permutations are performed by switching with probability $\frac{1}{2}$ the before and after gene expression profiles within each pair.

We now present two procedures. Procedure A is designed to control the number of false discoveries, and Procedure B is designed to control the proportion of false discoveries. The control of the error rates is “strong” in the sense that the error rate is controlled regardless of which variables satisfy the null hypothesis (Westfall and Young, 1993, p. 10). Both Procedures A and B described below are “step-down” permutation methods. That is, they proceed from smallest p -value to largest, which is equivalent to “stepping down” from the largest test statistic to the smallest.

2.1. Controlling the number of false discoveries

Procedure A. Suppose we wish to be $1 - \alpha$ confident that the number of false discoveries is $\leq u$ ($u > 0$). For any $T = \{t_1, t_2, \dots, t_j\} \subseteq \{1, 2, \dots, k\}$, let $y_{T,u}^\alpha$ denote the largest number such that

$$\Pr_{00}((u + 1)\text{st smallest of the } \{P'_{t_1}, P'_{t_2}, \dots, P'_{t_j}\} < y_{T,u}^\alpha) \leq \alpha,$$

where Pr_{00} refers to the multivariate permutation distribution under the global null hypothesis, i.e., all variables satisfy the null hypothesis. (In the proof below, the probabilities are not under this global-null permutation distribution. It is used here only as a means of defining the quantity $y_{T,u}^\alpha$.) We reject null hypotheses sequentially, in order of smallest to largest p -value as follows:

Automatically reject $H_{(1)}, H_{(2)}, \dots, H_{(u)}$. For $r > u$, having rejected $H_{(r-1)}$, reject $H_{(r)}$ if $P_{(r)} < y_{r,u}^\alpha$ where

$$y_{r,u}^\alpha = \min(y_{\{(i_1), (i_2), \dots, (i_u), (r), (r+1), \dots, (k)\}, u}^\alpha | \{i_1 < i_2 < \dots < i_u\} \subseteq \{1, 2, \dots, r-1\}) \tag{2}$$

and subscript (i_j) denotes the index associated with the i_j th ordered p -value from among the full set of observed p -values. Once a hypothesis is not rejected, all further hypotheses are not rejected.

For $u = 0$, we define Procedure A to be the usual stepwise procedure for controlling the familywise error (Westfall and Young, 1993, p. 66). That is, reject $H_{(1)}$ if $P_{(1)} < y_{K_0,0}^\alpha$, and having rejected $H_{(r-1)}$, reject $H_{(r)}$ if $P_{(r)} < y_{\{(r), (r+1), \dots, (k)\}, 0}^\alpha$.

Proposition A. Under the multivariate null-hypothesis permutation model described at the beginning of this section, Procedure A controls the number of false discoveries to be less than or equal to u with confidence $1 - \alpha$.

Proof of Proposition A. The proof follows closely along the lines of the proof of the control of the FWE rate for the standard stepwise FWE method (Westfall and Young, 1993, pp. 72–74). Let s be the rank of the p -value in the ordered list (1) corresponding to $(u+1)$ st smallest p -value associated with a variable that satisfies the null hypothesis. (For example, if $u = 1$ and the variables associated with $H_{(2)}$ and $H_{(4)}$ satisfy the null hypothesis, and the variables associated with $H_{(1)}$ and $H_{(3)}$ do not, then s would equal 4.) Let K_0 be the set of indices corresponding to the variables that satisfy the null hypothesis. In particular, the variable associated with $H_{(s)}$ is in K_0 and there are u other variables in K_0 associated with u of the $H_{(1)}, H_{(2)}, \dots, H_{(s-1)}$. Then

$\text{Pr}(\text{more than } u \text{ true null hypotheses are rejected})$

$$\begin{aligned} &= \text{Pr}(H_{(1)}, H_{(2)}, \dots, H_{(s)} \text{ are rejected}) \\ &\leq \text{Pr}(P_{(s)} < y_{s,u}^\alpha) \\ &\leq \text{Pr}(P_{(s)} < y_{K_0,u}^\alpha) \\ &\leq \alpha. \end{aligned}$$

The penultimate inequality is true because $y_{s,u}^\alpha \leq y_{K_0,u}^\alpha$ which follows because $K_0 \subseteq \{(i_1), (i_2), \dots, (i_u), (s), (s+1), \dots, (k)\}$ for a choice of $i_1 < i_2 < \dots < i_u$. Since $P_{(s)}$ is the $(u+1)$ st smallest p -value corresponding to variables with indices in the set K_0 , its distribution only involves p -values associated with variables satisfying the null hypothesis. Therefore, the permutation distribution of the $(u+1)$ st smallest of the $\{P'_t, t \in K_0\}$ can be used as the reference distribution for $P_{(s)}$, verifying the last inequality.

Note that in the above proof the control of the type I error is strong because there can be variables that do not satisfy the null hypothesis, i.e., those variables not in

K_0 . Subset pivotality, as defined by Westfall and Young (1993, p. 42), is sometimes needed for proofs of multiple testing results. However, since we only have one test per variable, it does not enter here (Westfall and Young, 1993, p. 115). (A proof of a result similar to Proposition A but involving a more general set of hypotheses would require subset pivotality.)

If the variables or p -values are discrete, there can be ties in the p -values given in (1), but this does not present a problem. Regardless of the ordering of the tied variables in (1), if the hypothesis associated with the first variable in the order is rejected, then the hypotheses associated with the other tied variables will also be rejected because the minimization (2) will be over smaller sets for the other variables. In addition, which of the tied variables is considered first for rejection will not matter, as the permutation distribution will include all of them when considering the first rejection. Also, if the first of the tied variables fails to reject, the procedure ceases and no further hypotheses are rejected, so that the situation in which the first tied variable fails to reject, and the later tied variables do reject, need not be considered.

We have used univariate p -values in (1) and (2) upon which to base our permutation tests. Any univariate statistic could be used here to distinguish between the groups, e.g., the absolute difference between two group means or the absolute difference between two proportions for binary data. The only requirement is that the same statistic is used for the $\{P'_t, t \in K\}$ generated on the permuted data. Obviously, using a statistic that captures well the evidence of group differences should result in more power to make true discoveries.

The computational burden implicit in (2) can be large if $\binom{r}{u}$ is not small. A conservative procedure in these situations would be to set $y_{r,u}^\alpha = y_{K,u}^\alpha$. When $r \ll k$ and the number of possible permutations is not very small, we expect the conservatism of this less computationally intensive procedure to be minimal. We examine this in Section 4 with some limited simulations.

As an alternative to fixing α in advance, one can obtain an “adjusted p -value” for each hypothesis. The adjusted p -value associated with hypothesis $H_{(r)}$ represents the smallest significance level at which $H_{(r)}$ would be rejected under the procedure. Appendix A contains an algorithm for calculating the adjusted p -values using Procedure A and the conservative version of Procedure A.

2.2. Controlling the proportion of false discoveries

Intuitively, Procedure A could be applied in an iterative way to control the false discovery proportion (FDP) which we define as the number of null hypotheses rejected divided by the total number of hypotheses rejected. If no hypotheses are rejected, we define $\text{FDP} = 0$. Note that the expected value of the FDP is the false discovery rate (FDR) that has been previously studied by Benjamini and Hochberg (1995). Suppose that the goal is to have $1 - \alpha$ confidence that the false discovery proportion is less than some value γ . Then, if the r th-ordered hypothesis is the last one to be rejected by the procedure, we want to be $1 - \alpha$ confident that no more than $\lceil r\gamma \rceil$ false discoveries are among those r rejected hypotheses, where the notation $\lceil x \rceil$ denotes the greatest integer less than or equal to x .

Procedure B. Suppose we wish to be $1 - \alpha$ confident that the proportion of false discoveries is no more than γ ($\gamma > 0$). Reject $H_{(1)}$ if $P_{(1)} < y_{K,0}^\alpha$. Having rejected $H_{(r-1)}$, reject $H_{(r)}$ if either $|\lceil r\gamma \rceil| > |\lceil (r - 1)\gamma \rceil|$ (“automatic rejection”) or if $P_{(r)} < y_{r,|\lceil r\gamma \rceil|}^\alpha$.

Conjecture. Procedure B approximately controls the proportion of false discoveries to be less than or equal to γ with confidence $1 - \alpha$.

Heuristic justification of the conjecture. Let S be the rank of the p -value in the ordered list (1) corresponding to the first time that if hypotheses $1, 2, \dots, S$ were rejected, then the FDP would be larger than γ at that point. S is a random variable that depends on the complete set of k tests. Note that the variable associated with $H_{(S)}$ satisfies the null hypothesis (i.e., is in K_0), and that there are $|\lceil S\gamma \rceil|$ other variables in K_0 associated with $|\lceil S\gamma \rceil|$ of the $H_{(1)}, H_{(2)}, \dots, H_{(S-1)}$. Also note that $H_{(S)}$ cannot be an “automatic rejection” because if $|\lceil S\gamma \rceil| > |\lceil (S - 1)\gamma \rceil|$, then the FDP would be larger than γ for hypothesis $(S - 1)$. The conjecture is justified by the following:

$$\begin{aligned} \Pr[\text{FDP} > \gamma] &\leq \Pr[H_{(1)}, H_{(2)}, \dots, H_{(S)} \text{ are rejected}] \\ &\leq \Pr[P_{(S)} < y_{S,|\lceil S\gamma \rceil|}^\alpha] \\ &\leq \Pr[P_{(S)} < y_{K_0,|\lceil S\gamma \rceil|}^\alpha] \end{aligned}$$

The first inequality is not an equality as it was in the proof of Proposition A because the procedure may continue rejecting beyond $H_{(S)}$ leading to the possibility that the false discovery proportion may subsequently fall below γ . If $|\lceil S\gamma \rceil|$ were a constant rather than a random quantity, then the last probability would be less than α . However, it is not a constant, and it is possible to construct counterexamples with the last probability exceeding α . However, in the counterexamples we have been able to construct, the exceedance of the probability over α is very small. This, along with the fact that there are other inequalities in going from $\Pr[\text{FDP} > \gamma]$ to $\Pr[P_{(S)} < y_{K_0,|\lceil S\gamma \rceil|}^\alpha]$, suggests that if it is possible for $\Pr[\text{FDP} > \gamma]$ to be greater than α then it will not exceed α by very much. We examine the properties of Procedure B in some limited simulation studies in Section 4.

The issue of tied p -values is more subtle here than for Procedure A because of the automatic rejections. If the p -value associated with an automatically rejected hypothesis is tied with the next p -value to be considered, then depending upon the arbitrary order of the variables with the tied p -values, the procedure might or might not reject the hypothesis immediately following the automatic rejection. We recommend in this situation that all possible ordering of the variables with tied p -values be considered, and the one that leads to the most rejections be used. If we have two tied p -values and consider both orderings, then we have two possible values of $y_{S,|\lceil S\gamma \rceil|}^\alpha$, say $y_{S,|\lceil S\gamma \rceil|}^{\alpha,1}$ and $y_{S,|\lceil S\gamma \rceil|}^{\alpha,2}$. Therefore, in the heuristic justification of the conjecture, we can replace the line $\Pr[P_{(S)} < y_{S,|\lceil S\gamma \rceil|}^\alpha]$ with $\Pr[P_{(S)} < y_{S,|\lceil S\gamma \rceil|}^{\alpha,1} \text{ or } P_{(S)} < y_{S,|\lceil S\gamma \rceil|}^{\alpha,2}]$. However, since $y_{S,|\lceil S\gamma \rceil|}^{\alpha,1}$ and $y_{S,|\lceil S\gamma \rceil|}^{\alpha,2}$ are each less than $y_{K_0,|\lceil S\gamma \rceil|}^\alpha$, the next line in the justification follows.

We can obtain an adjusted p -value for each hypothesis using a method similar to that employed for Procedure A; see Appendix A. This adjusted p -value associated with hypothesis $H_{(r)}$ represents the smallest significance level at which $H_{(r)}$ would be rejected under the procedure.

3. Application to microarray data from 40 paired breast tumor specimens

We demonstrate the methods on a subset of a previously published data set involving gene expression from cDNA microarrays using specimens from 65 breast tumors from 42 individuals (Perou et al., 2000). Our analysis is based on data from 20 individuals with specimens taken both before and after a 16-week course of doxorubicin chemotherapy. The microarray analyses generate two gene expression profiles for each specimen, one before, and one after chemotherapy. Each profile consists of log expression ratios measured on approximately 9000 genes. Based on a cluster analysis of these profiles, Perou et al. (2000, p. 747) note “Gene expression patterns in two tumor samples from the same individual were almost always more similar to each other than either was to any other sample.” This similarity does not eliminate the equally interesting possibility of finding large and statistically significant differences in gene expression in the 20 paired pre vs. post chemotherapy specimens. Genes showing different expression before compared to after chemotherapy will henceforth be referred to as “differentially expressed,” and the goal of our analysis will be to identify these differentially expressed genes.

The primary data were obtained at < http://genome-www.stanford.edu/breast_cancer/molecularp portraits/download.shtml >. Fluorescent intensities for the two labeled samples (test and reference samples) are recorded in two channels. Typically, a red and green image is produced to display the intensities in the two channels. For each channel and each spot (gene) on the array, both foreground (feature) and background intensity measures are given, and a “signal” can be calculated as the foreground intensity minus the background intensity. An important first step in the analysis of microarray data is careful examination of the red and green images of individual spots on the array to determine the quality of the fluorescence measurements. Owing to experimental artifacts such as dust specs, scratches, bubbles, or poor adherence of the DNA to the glass slide, some spots will need to be flagged and not used in the data analysis. This spot flagging had already been performed by the original investigators, and the flagging data was supplied. Data from spots flagged and not used by the original investigators were also not used here. Spots labeled as “EMPTY” were not used in any analyses.

A second data quality issue is the problem of low intensity spots. In particular, when foreground intensity is close to background intensity, and consequently the signal is low, it is known that the measurements tend to be unreliable. We have found it helpful to examine plots of signal in channel 1 vs. signal in channel 2 for each array. Fig. 1 shows such a plot for one of the arrays. A “fanning” of the point scatter in the lower left corner indicates the range in which signal is lost amid the noise. This plot and the many others examined (not shown) suggested that signal measurements less than 100 (2 on the log base 10 scale) should not be considered reliable. We chose to

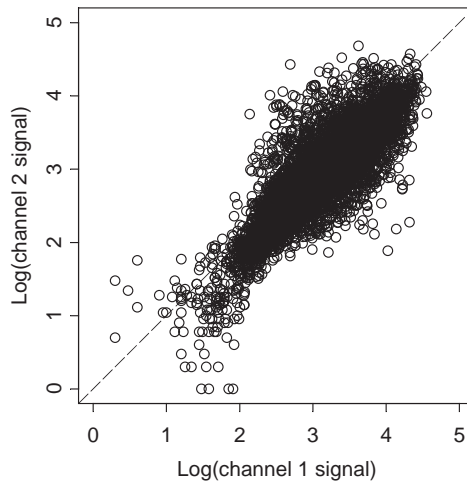


Fig. 1. Logarithm (base 10) of signal in channel 2 plotted vs. logarithm of signal in channel 1 for the genes having measurements in both channels in microarray experiment #25 in the Perou data set. Signal is computed as foreground intensity minus background intensity. Dashed line represents the 45° line.

exclude from the analysis spots for which signal was less than 100 in both channels. If signal was less than 100 in only one channel, the spot was used with the signal in that channel set to 100.

Also noticeable in Fig. 1 is a shift of the scatter of points to slightly below the 45° line. This shift typically results from differences in signal intensities in the two channels due to nuisance factors such as different physical properties of the two fluorescent dyes (for example, different efficiencies of dye incorporation or different rates of degradation), different photo-multiplier tube settings for scanning, or differences in the starting amounts of the two fluorescently labeled samples; it indicates the need for data normalization. An expression ratio for each gene was formed as channel 2 divided by channel 1 signal for that gene. Ratios were then median normalized within each array by dividing the ratios by the median of the ratios from nonexcluded spots for that array. We note that recent statistical work has described more complex alternatives to median normalization to normalize gene expression data (Yang et al., 2001; Chen et al., 1997). However, in our experience with applications involving considerable between-patient variability, the simple median normalization has worked well.

Genes for which data were missing from more than half of the 20 pairs of tumor specimens were eliminated from consideration. This left 8029 genes for analysis. All p -values were calculated on log transformed median-normalized expression ratios.

Table 3 shows the genes with the 28 smallest unadjusted paired t -test p -values for testing the null hypothesis that the mean pre- and post-chemotherapy expression of the gene is the same. (Nominal univariate parametric paired t -test p -values are used throughout this application.) Geometric means are also provided along with GenBank accession numbers. Information on these genes can be found by searching on the

Table 3

Genes identified when controlling false discovery proportion to be no more than approximately 10% with approximately 95% confidence

Gene rank	Accession number	Number of patients	Gene expression ratio (geometric mean)			Unadjusted <i>p</i> -value ^a
			Pre-chemo	Post-chemo	Ratio of post/pre	
1	AA478553	19	0.74	2.21	2.98	0.0000002
2	N23941	20	1.27	2.18	1.72	0.0000005
3	W96134	20	1.66	3.08	1.85	0.0000006
4	N95402	20	1.19	2.07	1.74	0.0000011
5	AA040944	20	0.42	1.79	4.21	0.0000011
6	AA442853	20	1.53	2.61	1.71	0.0000012
7	AA134757	20	2.32	4.69	2.02	0.0000021
8	AA418077	20	1.75	3.38	1.94	0.0000027
9	R12840	20	0.50	1.77	3.55	0.0000052
10	AI831083	20	1.24	2.48	2.00	0.0000057
11	AA044993	20	0.65	1.37	2.09	0.0000057
12	AA031596	20	1.19	2.06	1.73	0.0000067
13	AA454868	18	2.95	5.05	1.71	0.0000067
14	AA598794	20	0.72	1.52	2.10	0.0000077
15	T74141	19	8.28	16.44	1.98	0.0000097
16	H21041	20	0.95	1.83	1.93	0.0000110
17	AA133129	20	0.68	1.37	2.01	0.0000119
18	AA485377	20	0.52	1.33	2.56	0.0000221
19	AA167222	20	4.21	9.62	2.29	0.0000293
20	AA287695	19	1.27	2.18	1.72	0.0000325
21	AA489234	20	1.31	1.95	1.49	0.0000582
22	AA293362	20	1.71	2.94	1.72	0.0000600
23	N94487	20	1.19	2.39	2.01	0.0000658
24	H86754	19	2.00	3.79	1.89	0.0000974
25	AA430629	19	1.82	1.25	0.69	0.0001050
26	AA485743	20	1.20	0.92	0.77	0.0001155
27	H82948	19	0.84	0.72	0.86	0.0001209
28	H05099	19	0.97	0.77	0.80	0.0001309

^aBased on a paired Student's *t*-test with degrees of freedom equal to one less than the number of patients with paired data available for that gene.

accession numbers at <<http://www.ncbi.nlm.nih.gov/UniGene/>>. It is interesting to note that the first 24 genes identified have more expression post-chemotherapy than pre-chemotherapy, and all of the specimens for 4 of the genes showed more expression after the chemotherapy (Fig. 2). This list of 28 genes is the set identified by application of our Procedure B controlling the proportion of false discoveries to be no more than $\gamma = 0.10$ with approximate 95% confidence ($\alpha = 0.05$). Table 4 provides a comparison of the lists of genes identified under several methods: Bonferroni, step-down FWE control (equivalent to our Procedure A with $u = 0$ and $\alpha = 0.05$), Procedure A with $u = 1$ and $\alpha = 0.05$, Procedure A with $u = 2$ and $\alpha = 0.05$, and Procedure B with $\gamma = 0.10$ and $\alpha = 0.05$. Permutations in this paired data setting consist of switching




Fig. 2. For each of the genes given in Table 3, plotted points are the ratios of the post-chemotherapy to pre-chemotherapy gene expression for each of 18–20 patients, and arrows are the geometric means of the post/pre ratios. (a) The 17 most significant genes, comprising the set identified as differentially expressed by the level 0.05 step-down FWE procedure, or equivalently Procedure A allowing no false discoveries with 95% confidence. (b) The 18th through 28th most significant genes, which together with the first 17 comprise the set identified as differentially expressed by Procedure B allowing no more than 10% false discoveries with approximate 95% confidence.

the before- and after-chemotherapy expression profiles. Thus, there are 2^{20} possible permutations of the profile data. We used Monte-Carlo sampling to randomly select 19,999 of these permutations. Adjusted p -values are given under each method, with




Fig. 2. (Continued).

values in bold type denoting rejected hypotheses ($\alpha < 0.05$). Setting the confidence level at 95%, the number of rejections ranges from 11 to 28, with the Bonferroni method yielding the least and Procedure B, tied with Procedure A with $u = 2$, yielding the most. By exploiting the correlations among genes, the step-down FWE procedure was able to reject 6 more genes than the Bonferroni procedure. (These two procedures are directly comparable since they are controlling the same FWE rate.) Allowing just one false discovery allowed identification of an additional 6 genes, and allowing 2 or 10% false discoveries allowed identification of an additional 5 genes, for a total of 28. The results presented in Table 4 all used the less computational conservative ($y_{r,u}^z = y_{K,u}^z$) versions of Procedures A and B. Calculation of all results in Table 4 took approximately 18 min on a 866 MHz pentium III processor with 1 GB of memory. The fully computational method was also run and yielded very similar p -values with the same numbers of hypotheses being rejected.

4. Simulation studies

We conducted several simulations to assess the performance of Procedures A and B and to compare them to the procedure to control the FWE rate. Specifically, we evaluated the same procedures that we applied to the breast tumor data in the previous

Table 4
Adjusted *p*-values for genes identified in Table 3

Gene rank	Accession number	Adjusted <i>P</i> -value				
		Bonferroni	Step-down FWE control (Procedure A with $u = 0$)	Procedure A: number of false discoveries not to exceed 1 ($u = 1$)	Procedure A: number of false discoveries not to exceed 2 ($u = 2$)	Procedure B: false discovery proportion not to exceed 0.10 ($\gamma = 0.10$)
1	AA478553	0.0016	0.0010	0.0000	0.0000	0.0010
2	N23941	0.0040	0.0019	0.0002	0.0000	0.0019
3	W96134	0.0048	0.0022	0.0002	0.0001	0.0022
4	N95402	0.0088	0.0035	0.0002	0.0001	0.0035
5	AA040944	0.0088	0.0035	0.0002	0.0001	0.0035
6	AA442853	0.0096	0.0038	0.0002	0.0001	0.0038
7	AA134757	0.0169	0.0068	0.0002	0.0001	0.0068
8	AA418077	0.0217	0.0088	0.0003	0.0001	0.0088
9	R12840	0.0418	0.0176	0.0005	0.0002	0.0176
10	AI831083	0.0458	0.0189	0.0007	0.0002	0.0176
11	AA044993	0.0458	0.0189	0.0007	0.0002	0.0176
12	AA031596	0.0538	0.0223	0.0008	0.0002	0.0176
13	AA454868	0.0538	0.0223	0.0008	0.0002	0.0176
14	AA598794	0.0618	0.0254	0.0010	0.0002	0.0176
15	T74141	0.0779	0.0318	0.0016	0.0003	0.0176
16	H21041	0.0883	0.0375	0.0021	0.0003	0.0176
17	AA133129	0.0955	0.0416	0.0023	0.0003	0.0176
18	AA485377	0.1774	0.0753	0.0070	0.0014	0.0176
19	AA167222	0.2352	0.1019	0.0118	0.0022	0.0176
20	AA287695	0.2609	0.1137	0.0143	0.0026	0.0176
21	AA489234	0.4673	0.2039	0.0386	0.0094	0.0176
22	AA293362	0.4817	0.2090	0.0409	0.0099	0.0176
23	N94487	0.5283	0.2267	0.0472	0.0121	0.0176
24	H86754	0.7820	0.3158	0.0879	0.0276	0.0276
25	AA430629	0.8430	0.3352	0.0985	0.0318	0.0318
26	AA485743	0.9273	0.3591	0.1127	0.0382	0.0382
27	H82948	0.9707	0.3724	0.1210	0.0423	0.0423
28	H05099	1.00	0.3959	0.1352	0.0492	0.0492

Note: All permutation procedures used 19,999 resampled permutations. The computationally conservative versions of Procedures A and B, i.e., $y_{r,u}^z = y_{K,u}^z$, were used. Univariate *p*-values upon which all procedures were based were computed using a paired Student's *t*-test with degrees of freedom equal to one less than the number of patients with paired data available for that gene. *P*-values in bold type correspond to tests that would be rejected under the given procedure at level 0.05.

section. Similar to the paired breast tumor example, we considered cases in which 8000 hypotheses were to be tested, and we assumed 40 paired specimens. The paired difference data were generated as multivariate Gaussian with a variety of block correlation structures. The blocks of correlated variables can be viewed, for example, as representing genes that are in the same pathway or that are co-regulated. Variables are independent between blocks. Both positive and negative correlations were considered.

Table 5
Simulation results applying Procedures A and B at level 0.05 to sets of 40 paired profiles generated under block-exchangeable correlation structure

	Step-down FWE (Proc. A, $u = 0$)	Procedure A: $FD^a \leq 1$ ($u = 1$)	Procedure A: $FD \leq 2$ ($u = 2$)	Procedure B: $FDP^b \leq 0.10$ ($\gamma = 0.10$)
$\rho = 0$				
Sensitivity (in %) to detect nonnull	66.09	87.03	92.71	93.49
% of simulations with $FD > 0$	5.39	28.80	54.51	63.59
% of simulations with $FD > 1$	0.17	5.22	19.22	30.82
% of simulations with $FD > 2$	0.02	0.71	4.89	12.12
% of simulations with $FDP > 0.10$	0.03	0.28	1.79	4.43
$\rho = 0.5$				
Sensitivity (in %) to detect nonnull	66.91	85.82	91.27	87.29
% of simulations with $FD > 0$	4.78	22.59	40.50	46.52
% of simulations with $FD > 1$	0.37	4.62	13.55	19.05
% of simulations with $FD > 2$	0.06	1.20	4.53	8.15
% of simulations with $FDP > 0.10$	0.52	1.21	2.93	4.32
$\rho = 0.9$				
Sensitivity (in %) to detect nonnull	82.43	88.63	91.16	87.95
% of simulations with $FD > 0$	5.00	9.25	12.03	13.01
% of simulations with $FD > 1$	2.68	5.07	7.19	7.91
% of simulations with $FD > 2$	1.86	3.63	5.05	5.55
% of simulations with $FDP > 0.10$	1.92	3.59	4.69	4.89

Note: Paired difference data were generated as 8000 Gaussian variables, each with variance 1, in 80 blocks of 100 correlated variables each. Pairwise correlation between variables was ρ within a block and 0 between blocks. Thirty of the variables out of the 8000 tested had standardized mean equal to 1.5, and the remaining had mean zero. Percentages in bold type have nominal value 5.00%.

^aFD=number of false discoveries.

^bFDP=false discovery proportion.

We varied the number of differentially expressed genes (nonnull variables), and the patterns of correlations among those differentially expressed genes. Each simulation was repeated 10,000 times. For purposes of computational feasibility, 99 resamplings were performed at each stage, and the conservative method of critical value calculation ($y_{r,u}^z = y_{K,u}^z$) was used except where noted otherwise. The effect of using 99 resamplings was investigated in some limited simulation studies, and the results were found not to differ substantially from using 199 resamplings.

In Table 5, performance of the methods is assessed under a block exchangeable correlation structure with 30 nonnull variables out of 8000. All procedures satisfied the targeted 95% confidence, with little conservatism. The increased sensitivity to detect nonnull variables afforded by allowing even one or two false discoveries is striking, for example an absolute increase of over 20% as compared to the procedure controlling the FWE when the within-block correlation was 0 or 0.5. Several variants of these cases were also examined and are now described (results not shown). For the results

Table 6

Simulation results applying Procedures A and B at level 0.05 to sets of 40 paired profiles generated under correlation structure with mixed positive and negative correlations

	Step-down FWE (Proc. A, $u = 0$)	Procedure A: $FD^a \leq 1$ ($u = 1$)	Procedure A: $FD \leq 2$ ($u = 2$)	Procedure B: $FDP^b \leq 0.10$ ($\gamma = 0.10$)
Sensitivity (in %) to detect nonnull	67.60	85.18	90.61	90.21
% of simulations with $FD > 0$	4.85	20.57	35.29	42.10
% of simulations with $FD > 1$	0.54	4.79	12.24	16.84
% of simulations with $FD > 2$	0.15	1.59	5.14	8.09
% of simulations with $FD > 0.10$	0.24	1.13	3.16	4.92

Note: Paired difference data were generated as 8000 Gaussian variables in 80 blocks of 100 correlated variables each. Within each block, the set of variables is divided into thirds (33|33|34). Any pair of variables within the same third have positive correlation equal to $\frac{2}{3}$, and between thirds, variables have negative correlation $-\frac{1}{3}$. Between blocks, variables are independent. Of the 100 variables in the first block, 10 variables in the first third, 10 in the second third, and 10 in the last third each had standardized mean equal to 1.5. All remaining variables had mean zero. This resulted in 30 correlated, nonnull variables, each with standardized mean equal to 1.5. Percentages in bold type have nominal value 5.00%.

^aFD=number of false discoveries.

^bFDP=false discovery proportion.

in Table 5, the 30 nonnull variables were all placed within the same block, that is, they were correlated. For the case with correlation $\rho = 0.5$, when we distributed the nonnull variables one per block (uncorrelated), the results were very similar. The sensitivities and levels were also very similar when a block size of 20 rather than 100 was used. An interesting result was obtained when the number of nonnull variables was reduced to 5. In this case, the performance of Procedure A to control for no more than 1 or 2 errors was essentially the same as before, but the sensitivity advantage of Procedure B was lost, as it became nearly equivalent to the FWE procedure; even one false rejection would result in 20% false discoveries, which exceeds the allowable 10%. An additional simulation under the complete null (i.e., no nonnull variables) with block correlation 0.5 verified that stated confidence levels were also maintained by each procedure. In terms of the mean numbers of false discoveries, the observed mean numbers were 0.05, 1.07, and 2.09 for the procedures controlling for 0, 1, and 2 false discoveries, respectively.

A case of negatively correlated variables in combination with positively correlated variables was also considered. As in Table 5, there were 30 nonnull variables out of 8000. Results are presented in Table 6, and they are very similar to those given in Table 5. Confidence levels are maintained, and substantial sensitivity is gained by using the procedures that allow a few false discoveries.

We also assessed the difference between the conservative procedure and the fully computational procedure for the case of block correlation structure with correlation 0.5

and 30 nonnull variables. We found that the two computational methods had essentially the same performance in terms of sensitivity to detect nonnull variables.

5. Discussion

This paper has considered both the control of the absolute number of false discoveries as well as the control of the false discovery proportion. Investigators may prefer one or the other of these types of control. For example, if an investigator considers 10 false discoveries out of 100 discoveries acceptable but not 10 false discoveries out of 12 discoveries, then he or she would be more interested in controlling the false discovery proportion. On the other hand, if following up 10 false discoveries is considered acceptable to find even one true discovery, then control of the number of false discoveries would be more appropriate and more sensitive. Also, when it is expected that there are few truly nonnull hypotheses, the discreteness of the false discovery proportion should be kept in mind when setting the bound γ . For example, setting $\gamma=0.10$ with five nonnull variables is essentially equivalent to allowing no false discoveries.

The sensitivity gains from allowing even one or two false discoveries or a small proportion of false discoveries can be very large. This was clearly evident in Table 5 in the cases of low to moderate correlation where the sensitivity to detect nonnull variables increased from approximately 65% to about 85% by allowing just one false discovery, and to over 90% if two false discoveries were allowed. In the breast cancer example, the number of identified genes allowing either 2 or 10% false discoveries was more than 1.5 times the number identified when controlling the FWE and more than 2.5 times the number identified by the Bonferroni procedure.

The main potential downside of these procedures is the computational burden, but we do not view this as a serious disadvantage. In Section 2.1, we suggest the computational simplification of setting $y_{r,u}^\alpha = y_{k,u}^\alpha$ when $r \ll k$. In our breast tumor example, the slight conservatism that this might have introduced did not alter the number of differentially expressed genes we were able to identify compared to using the fully computational procedure. Also, our simulations which all used the conservative simplification did not suggest any substantial conservatism in the confidence levels or significant reduction in sensitivity for the cases we considered. The degree of conservatism would be most severe in a situation in which the number of nonnull variables is large and a moderately large number or proportion of false discoveries is allowed. Such situations are not likely to be the norm, but further work might be needed to develop appropriate computational simplifications or computationally more efficient algorithms for those cases. For example, for Procedure A, a potentially useful hybrid approach might be to use $y_{r,u}^\alpha$ up to a certain value of r , say $r^* = 20$ or 30, and then switch to the fixed critical value $y_{r^*,u}^\alpha$ thereafter. In any case, undoubtedly a substantial amount of time has been invested to measure the large number of variables, so even if a statistical procedure requires a few hours to complete, the time seems well spent to improve the sensitivity for detecting important variables.

Procedures A and B are immediately generalizable to a wide range of situations. In our example and simulations, we based the permutation procedures on parametric

paired t -test p -values. However, the procedures did not depend on any of the parametric assumptions. Any orderable univariate statistics can form the basis for the permutation procedure. With appropriate choice of univariate statistics, these methods apply to discrete data or even to censored data provided that the censoring distributions are the same in the groups. Unpaired or multi-group sampling designs are handled by appropriately modifying the permutation method.

Acknowledgements

This study utilized the high-performance computational capabilities of the Biowulf/LoBoS3 cluster at the National Institutes of Health, Bethesda, MD. We thank the referees and editorial board member for their helpful comments.

Appendix A. Algorithms for calculating adjusted p -values

There are various algorithms one could use to implement the procedures, especially the computationally intensive ones. We describe algorithms for the computationally intensive procedures that would be appropriate when there are large number of variables but we are interested in only those with small adjusted p -values. For all the algorithms, let the p -values of the observed data be (p_1, p_2, \dots, p_k) , and let the ordered p -values for the observed data be $p_{v_1} \leq p_{v_2} \leq \dots \leq p_{v_k}$. The sequence $\{v_j\}$ is fixed throughout the algorithms. Denote the adjusted p -values by $(\hat{p}_{v_1}, \hat{p}_{v_2}, \dots, \hat{p}_{v_k})$.

Algorithm A. This algorithm determines the adjusted p -values for Procedure A to allow u false discoveries. The adjusted p -values are calculated recursively, starting by setting $\hat{p}_{v_1} = \hat{p}_{v_2} = \dots = \hat{p}_{v_u} = 0$. The algorithm is used to calculate \hat{p}_{v_r} , the adjusted p -value for variable v_r ($r > u$) after having found $\hat{p}_{v_1}, \hat{p}_{v_2}, \dots, \hat{p}_{v_{r-1}}$.

Step 1. Choose B random permutations of the data vectors consistent with the experimental design. Denote the univariate p -values for the variables from the j th permutation by $(p_{v_1}^{*(j)}, p_{v_2}^{*(j)}, \dots, p_{v_k}^{*(j)})$ for $j = 1, 2, \dots, B$.

Step 2. Initialize $COUNT = 1$. Let $W = \{v_{w_1}, v_{w_2}, \dots, v_{w_u}\}$ be a collection of distinct indices from $\{v_1, v_2, \dots, v_{r-1}\}$. Define the set $R = \{v_{w_1}, v_{w_2}, \dots, v_{w_u}\} \cup \{v_r, v_{r+1}, \dots, v_k\}$. For each $j = 1, 2, \dots, B$, find ϕ_j , the $(u + 1)$ st smallest p -value for the permuted data over the subset R , i.e.,

$$\phi_j = [\{p_i^{*(j)} : i \in R\}]_{(u+1)},$$

where the notation $[A]_{(s)}$ is defined to refer to the s th smallest of the elements of the finite set A .

Step 3. For $j = 1, 2, \dots, B$, if $p_{v_r} \geq \phi_j$ then $COUNT \leftarrow COUNT + 1$.

Step 4. $p_{v_r}^*(W) \leftarrow COUNT/(B + 1)$.

Step 5. For each W repeat steps 2–4. There will be $\binom{r-1}{u}$ such sets W . (Note that when $u = 0$, W can only be empty, and in that case $R = \{v_r, v_{r+1}, \dots, v_k\}$.)

- Step 6. $p_{v_r}^* \leftarrow \max_W p_{v_r}^*(W)$.
- Step 7. $\hat{p}_{v_r} \leftarrow \max\{\hat{p}_{v_1}, \hat{p}_{v_2}, \dots, \hat{p}_{v_{r-1}}, \hat{p}_{v_r}^*\}$.

If the sample sizes are small enough so that all permutations of the data vectors consistent with the experimental design can be enumerated, then $j = 1, \dots, B$ in Step 1 should run over all permutations except the one corresponding to the observed data.

Algorithm B. This algorithm determines the adjusted p -values for Procedure B to allow the false discovery proportion $\leq \gamma$. The adjusted p -values are calculated recursively, starting with \hat{p}_{v_1} . The algorithm below is used to calculate \hat{p}_{v_r} , the adjusted p -value for variable v_r after having found $\hat{p}_{v_1}, \hat{p}_{v_2}, \dots, \hat{p}_{v_{r-1}}$. The algorithm is identical to Algorithm A except rather than u being fixed in advance, there is a Step 0:

Step 0. Let $u = \lceil r\gamma \rceil$. If $u > \lceil (r-1)\gamma \rceil$, then an automatic rejection is allowed and set $p_{v_r}^* = 0$ and skip to Step 7. Otherwise, continue with Step 1.

Algorithm A*. This algorithm determines the adjusted p -values for the conservative version of Procedure A to allow u false discoveries. First, one starts by setting $\hat{p}_{v_1} = \hat{p}_{v_2} = \dots = \hat{p}_{v_u} = 0$. The algorithm then calculates all of the remaining adjusted p -values.

Step 0. Initialize the counters $\text{COUNT}_i = 1$ for $i = u + 1, \dots, k$.

Step 1. Choose a random permutation of the data vectors consistent with the experimental design. Denote the univariate p -values for the variables from this permutation by $(p_1^*, p_2^*, \dots, p_k^*)$.

Step 2. Let $q^* = [\{p_i^* : i = 1, \dots, k\}]_{(u+1)}$.

Step 3. If $p_{v_i} \geq q^*$ then $\text{COUNT}_i \leftarrow \text{COUNT}_i + 1$ for $i = u + 1, \dots, k$.

Step 4. Repeat Steps 1–3 B times.

Step 5. $\hat{p}_{v_i} \leftarrow \text{COUNT}_i / (B + 1)$ for $i = u + 1, \dots, k$.

As with Algorithm A, if the sample sizes are small enough so that all permutations of the data vectors consistent with the experimental design can be enumerated, then all permutations except the one corresponding to the observed data should be used in Steps 1–3.

Algorithm B*. This algorithm determines the adjusted p -values for the conservative version of Procedure B to allow the false discovery proportion $\leq \gamma$.

Step 0. Initialize the counters $\text{COUNT}_r = 1$ for $r = 1, \dots, k$.

Step 1. If $\lceil r\gamma \rceil > \lceil (r-1)\gamma \rceil$, then $\tilde{p}_{v_r} = 0$ for $r = 1, \dots, k$.

Step 2. Choose a random permutation of the data vectors consistent with the experimental design. Denote the univariate p -values for the variables from this permutation by $(p_1^*, p_2^*, \dots, p_k^*)$. (Note that the ordering does not matter for this step.)

Step 3. If $\lceil r\gamma \rceil = \lceil (r-1)\gamma \rceil$ and $p_{v_r} \geq [\{p_i^* : i = 1, \dots, k\}]_{(\lceil r\gamma \rceil + 1)}$ then $\text{COUNT}_r \leftarrow \text{COUNT}_r + 1$ for $r = 1, \dots, k$.

Step 4. Repeat Steps 2–3 B times.

Step 5. If $\lceil r\gamma \rceil = \lceil (r-1)\gamma \rceil$ then $\tilde{p}_{v_r} \leftarrow \text{COUNT}_r / (B + 1)$ for $r = 1, \dots, k$.

Step 6. $\hat{p}_{v_r} \leftarrow \max(\tilde{p}_{v_1}, \dots, \tilde{p}_{v_r})$ for $r = 1, \dots, k$.

References

- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. Ser. B* 57, 289–300.
- Callow, M.J., Dudoit, S., Gong, E.L., Speed, T.P., Rubin, E.M., 2000. Microarray expression profiling identifies genes with altered expression in HDL-deficient mice. *Genome Res.* 10, 2022–2029.
- Chen, Y., Dougherty, E.R., Bittner, M.L., 1997. Ratio-based decisions and the quantitative analysis of cDNA microarray images. *J. Biomed. Opt.* 2, 364–374.
- International Human Genome Sequencing Consortium, 2001. Initial sequencing and analysis of the human genome. *Nature* 409, 860–921.
- Lehmann, E.L., 1975. *Nonparametrics*. Holden-Day, New York.
- Perou, C.M., Sorlie, T., Eisen, M.B., van de Rijn, M., Jeffrey, S.S., Rees, C.A., Pollack, J.R., Ross, D.T., Johnsen, H., Akslen, L.A., Fluge, O., Pergamenschikov, A., Williams, C., Zhu, S.X., Lonning, P.E., Borresen-Dale, A., Brown, P.O., Botstein, D., 2000. Molecular portraits of human breast tumors. *Nature* 406, 747–752.
- Schena, M., Shalon, D., Davis, R.W., Brown, P.O., 1995. Quantitative monitoring of gene expression patterns with a complementary DNA microarray. *Science* 270, 467–470.
- Seeger, P., 1968. A note on a method for the analysis of significance en masse. *Technometrics* 10, 586–593.
- Troendle, J.F., 1996. A permutational step-up method of testing multiple outcomes. *Biometrics* 52, 846–859.
- Tusher, V., Tibshirani, R., Chu, G., 2001. Significance analysis of microarrays applied to transcriptional responses to ionizing radiation. *Proc. Natl. Acad. Sci.* 98, 5116–5121.
- Westfall, P.H., Young, S.S., 1993. *Resampling-Based Multiple Testing*. Wiley, New York; pp. 72–74.
- Yang, Y.H., Dudoit, S., Luu, P., Speed, T.P., 2001. Normalization for cDNA microarray data. Technical Report 589, Department of Statistics, UC Berkeley, CA.