Statistical Methods in Medical Research 2009; **18**: 565–575

Gene set enrichment analysis made simple

Rafael A Irizarry Department of Biostatistics, Johns Hopkins School of Public Health, 615 N. Wolfe St. E3620, Baltimore, MD 21205, USA, **Chi Wang** Statistics Department, University of California, Riverside, 900 University Avenue, 2626 Statistics Computer Bldg. Riverside, CA 92521-0122, USA, **Yun Zhou** Gilead Sciences, Inc. 333 Lakeside Dr., Foster City, CA 94404, USA and **Terence P Speed** Department of Statistics, 367 Evans Hall, #3860, University of California, Berkeley, CA 94720-3860, USA

Among the many applications of microarray technology, one of the most popular is the identification of genes that are differentially expressed in two conditions. A common statistical approach is to quantify the interest of each gene with a *p*-value, adjust these *p*-values for multiple comparisons, choose an appropriate cut-off, and create a list of *candidate genes*. This approach has been criticised for ignoring biological knowledge regarding how genes work together. Recently a series of methods, that do incorporate biological knowledge, have been proposed. However, the most popular method, gene set enrichment analysis (GSEA), seems overly complicated. Furthermore, GSEA is based on a statistical test known for its lack of sensitivity. In this article we compare the performance of a simple alternative to GSEA. We find that this simple solution clearly outperforms GSEA. We demonstrate this with eight different microarray datasets.

1 Introduction

The problem of identifying genes that are differentially expressed in two conditions has received much attention from the statistical community and data analysts in general. Most of the work has focused on designing appropriate test statistics[1,2] and developing procedures to account for multiple comparisons.[3,4] Most approaches follow a similar recipe: decide on a null hypothesis, test this hypothesis for each gene, produce a *p*-value, and attach a significance level that accounts for multiplicity. At the end, each gene receives a score, which we use to decide if it is in our final list of significant genes. Those on this final list are typically called *candidate genes* because further validation tests are commonly performed. In this article, we refer to this as the *marginal* approach. A limitation of this approach is that genes that are known to be biologically associated are scored independently. Although many important discoveries have been made with this approach, the resulting gene lists do not always provide useful biological insights.

Recently, various approaches have been proposed to incorporate biological knowledge into the analysis. The vast majority of these have relied on the results from the marginal approach instead of starting from the original expression data. Because many of these marginal procedures have been useful and given the complicated nature of microarray data, we view this as a correct first approach. In this article we do not discuss nor propose methods that start from scratch.

Address for correspondence: Terence P Speed, Department of Statistics, 367 Evans Hall, #3860, University of California, Berkeley, CA 94720-3860, USA. E-mail: terry@stat.berkeley.edu

10.1177/0962280209351908

There are currently two major types of procedure for incorporating biological knowledge into differential expression analysis. We will refer to these as the *over-representation* and the *aggregate score* approaches. In both, gene categories or *gene sets* are formed prior to the statistical analysis. The sets are formed by, for example, grouping genes that are part of the same cellular components, are essential for a biological process, or have the same molecular function. In many cases the gene sets target the condition that is being studied. However, it is more common to use category definitions from the Gene Ontology project.[5] The Gene Ontology project provides a controlled vocabulary to describe gene and gene product attributes in any organism.[6]

Over-representation analysis can be summarised as follows: first, form a list of candidate genes using the marginal approach. Then, for each gene set, we create a two-by-two table comparing the number of candidate genes that are members of the category to those that are not members. The significance of over-representation can be assessed, for example, using the hypergeometric distribution or its binomial approximation. More elaborate approaches exist and a large number of over-representation methods have been published. Many of these have been implemented as web-tools. A comprehensive list can be found at `http://www.geneontology.org/GO.tools.microarray.shtml`

A limitation of the over-representation approach is that it ignores all the genes that did not make the list of candidate genes. Therefore, the results will be highly dependent on the cutoff used in constructing this list. In fact, examples can be found where very few, or even none, of the genes in functional groups known to behave differently in the two conditions survive the typical filters, and therefore the groups are not detected as interesting. Mootha *et al.*[7] describe a particularly interesting example. The *aggregate score* approach, does not have this limitation. The basic idea is to assign scores to each gene set based on all the gene-specific scores for that gene set. There are various ways to calculate these aggregate scores[7–14]. In this article we focus on the aggregate score method rather than the over-representation approach.

Of these methods gene set enrichment analysis (GSEA)[7,13] is by far the most popular. Surprisingly, GSEA is based on the Kolmogorov–Smirnov (K–S) test, which is well known for its lack of sensitivity and limited practical use. Subramanian *et al.*[13] seem to have realised this and developed an *ad hoc* modification of the K–S test. A further limitation of the K–S test and its modified versions, is that the null distribution of the score is hard to compute. Tian *et al.*[14] propose the use of the standard statistical approach for detecting shifts in centre: a one sample z-test. Tian *et al.* propose the use of permutation tests for assessing the significance of the z-test. However, they do not explore the performance of the standard parametric approach. We find that using the one sample t-test along with a standard multiple comparison adjustment[15] of the normal distribution p-value works well in practice. This procedure is extremely simple in comparison to GSEA and requires practically no computation time.

A possible advantage of GSEA, i.e. the K–S test, over the one sample z-test is that the latter is specifically designed to identify gene sets with mean shifts and the K–S test is designed to find general difference in the cumulative distribution. In principle, we want to be able to detect gene sets for which some members are up-regulated and others are down-regulated. The z-test is not sensitive to this change as there is no shift in mean. We therefore, propose the use of another standard statistical test useful for detecting changes in scale: the χ^2-test.

In this article, we compare GSEA to the one sample z-test and χ^2-test using all the datasets described in Mootha *et al.*[7] and Subramanian *et al.*[13]. In Section 2 we briefly describe the methods in question. In Section 3 we present the results from the comparison. Finally, in Section 4 we discuss these results, describe some current work that we expect to improve upon our proposed method, and give concluding remarks.

2 Methods

Most aggregate score approaches start with the results from a marginal analysis. For example, we may start with a t-statistic t_i for each gene $i = 1, \ldots, N$. We then identify gene set g with a subset $A_g \subset \{1, \ldots, N\}$. We want our score, say E_g (E for enrichment), to quantify how *different* the $t_i, i \in A_g$ are from the $t_i, i \notin A_g$. A second task is to assign a level of significance to each E_g. Most methods take the approach of defining a null hypothesis, calculating the null distribution, and assigning a level of significance. Because the scores for dozens of gene sets are considered, the significance levels are adjusted for multiple comparisons. The competing methods differ in the way that *different* is quantified and the null hypothesis defined and calculated. Notice, that the t_is need not be a t-statistic. In fact the GSEA article uses another statistics that summarised the signal-to-noise ratio (SNR) for each gene. Because the resulting values are very similar to a t-statistic we refer to the t_i as signal-to-noise value and t-statistic interchangeably.

Mootha *et al.*[7] use a version of the K–S statistic to test for differences in the distributions of the t-statistics related to members of a gene set compared to t-statistics from the rest of the genes. Because they were interested in comparing these scores across gene sets of different sizes, and then null distribution of the K–S statistic depends heavily on this size, Mootha *et al.* defined a normalised K–S statistics as their score E^{GSEA}. To assess the significance of these scores a permutation test was performed. Specially, they permuted the sample labels and re-computed E_g^{GSEA} 1000 times. In each permutation the maximum enrichment score was recorded. These 1000 values defined the null distribution and used to assign p-values.

Subramanian *et al.*[13] seem to have noticed the lack of power of the K–S test, a well-known fact, and proposed an *ad hoc* modification to improve this. Furthermore, in the original version of GSEA, an adjusted p-value was calculated only for the enrichment score of the top ranking set. In Subramanian *et al.*, after normalising the test statistic for each gene set, the FDR q-value for each gene set is calculated and used to select candidate gene sets. The end result is a rather complicated method that takes minutes to run on a typical laptop computer.

Determining if two sets of numbers have different distributions is certainly not a new problem. Many solutions exist. The K–S test is one that has not been used in many (or any) other applications, so why use it here? Let us start with the most basic statistical approach: test for a shift in centre/mean as proposed by Tian *et al.*[14] If, under the null hypothesis, the t_i are normally distributed with mean 0 and standard deviation 1, inference can be done with a one sample z-test. For a robust version we could use a Wilcoxon test. When enough replicates are available in each condition we expect the t-statistics to follow a standard normal distribution under the null-hypothesis of no

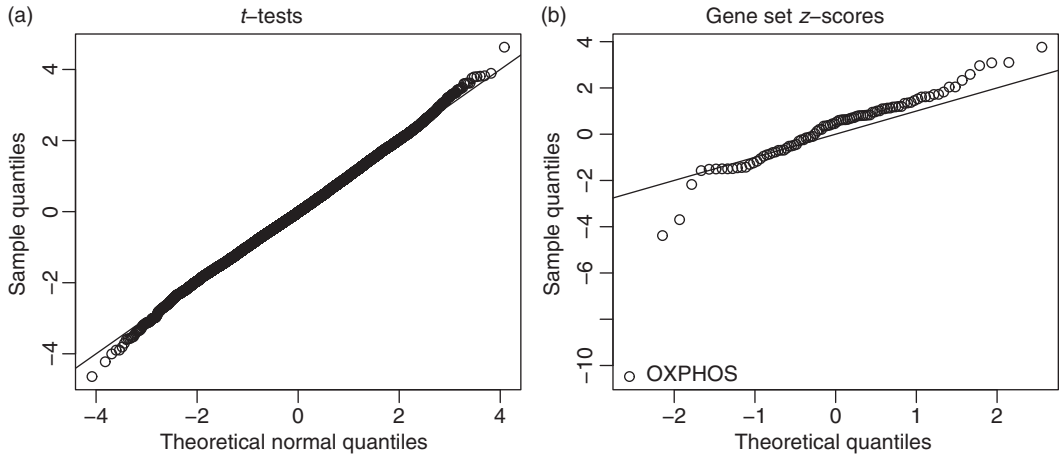

Figure 1 Quantile–quantile plots. (a) For the diabetes data presented in Mootha *et al.*[7] we plot the quantiles of the observed t-statistics vs the theoretical quantiles of the standard normal distribution. The identity line is shown. (b) For the same data we show the enrichment score based on the z-test for the gene sets presented by Mootha *et al.* The score for the OXPHOS gene set is highlighted.

difference between the conditions. The data presented by Mootha *et al.*[7] seem to satisfy this assumption. Figure 1(a) shows a quantile–quantile plot comparing the t-tests used in Mootha *et al.* to a standard normal distribution. Figure 2 shows this quantile–quantile plot for all datasets in Subramanian *et al.*[13] Barring a few outliers, which are likely associated to differentially expressed genes, the assumption appears appropriate in all datasets. If we assume that these tests are independent (under the null) then for any given gene set the z-score:

$$E_g^z = \sqrt{N_g}\,\bar{t}, \quad \text{with } \bar{t} = \frac{1}{N_g}\sum_{i \in A_g} t_i, \tag{1}$$

with N_g the number of genes in A_g, also follows a standard normal distribution. This implies that we can easily obtain a p-value.

With appropriate p-values calculated we have numerous multiple comparison adjustment methods to choose from and do not need to perform permutation tests. Tian *et al.*[14] argue that the normality assumption is not appropriate because we expect the t_i to be correlated even under the null hypothesis. However, they do not appear to have tested this empirically. We find that assuming the E_g^z are normally distributed under the null hypothesis is in fact a useful approximation for all the examples we examined. For example, Figure 1(b) shows the z-score for the dataset presented in Mootha *et al.* for the same gene sets they considered. Notice that the obvious outlier in Figure 1(b), is the OXPHOS gene set discovered to be important by Mootha *et al.* Thus, the discovery that merited their publication would have been made with a statistical method that could be explained in one paragraph instead of several pages.

A possible limitation of the one sample z-test is that it will not detect changes in scale. A gene set where half the gene sets are up-regulated and the other half are down-regulated

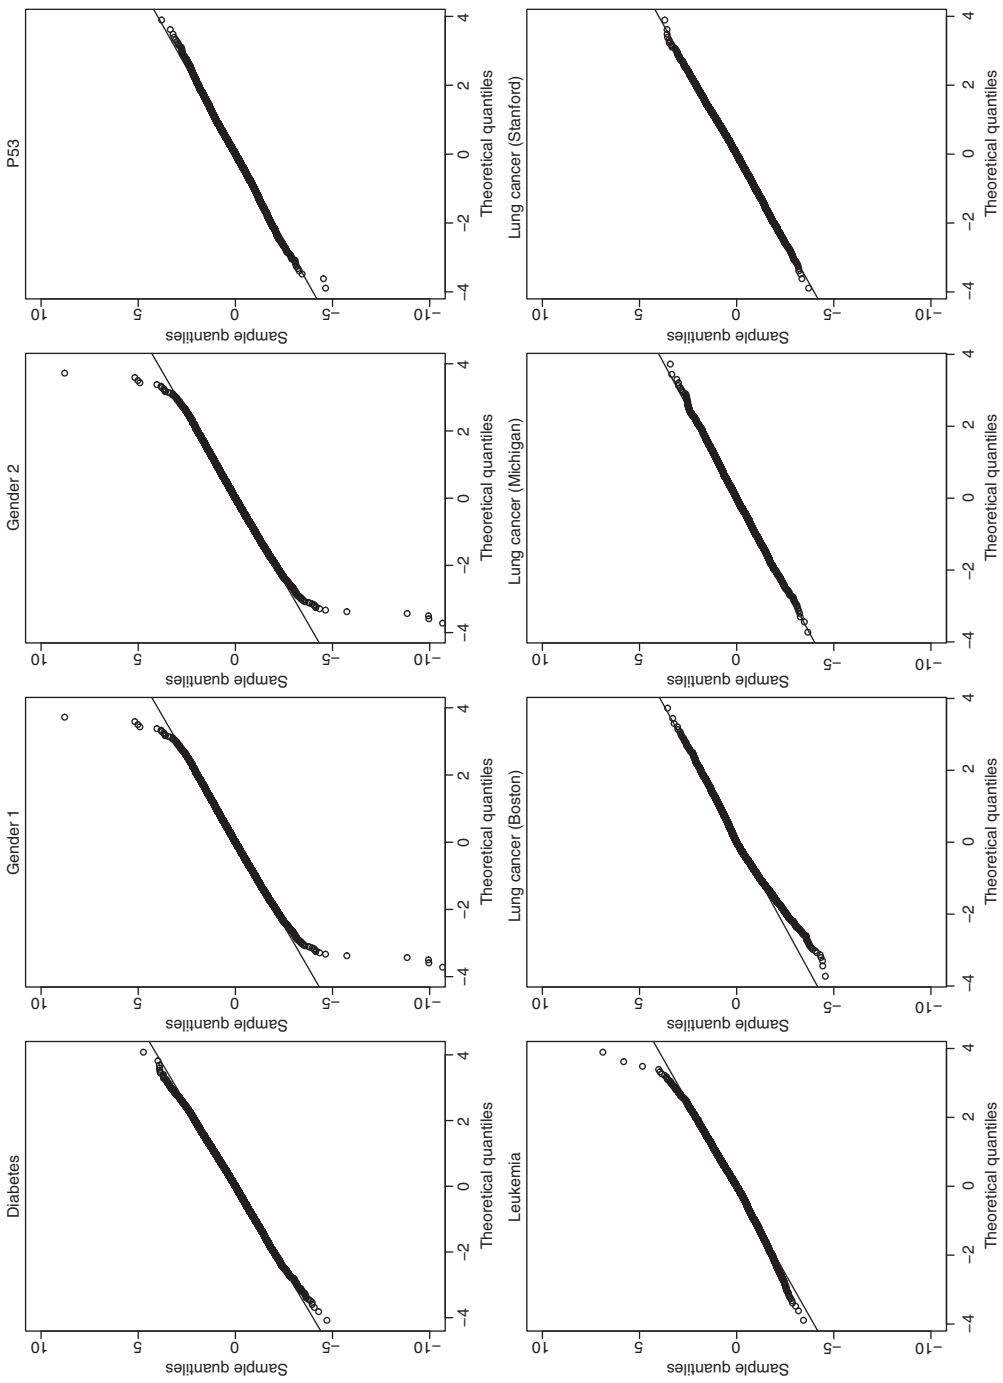

Figure 2 As for Figure 1(a) but for all the datasets presented in Mootha *et al.*[7] and Subramanian *et al.*[13] The identity line is shown.

may have no mean shift but is certainly interesting from a biological standpoint. The standard test for scale change, i.e. the χ^2-test, is useful for this. We define a standardised χ^2-test that permits us to compare gene sets of different sizes and different mean shifts:

$$E_g^{\chi^2} = \frac{\sum_{i \in A_g} (t_i - \bar{t})^2 - (N_g - 1)}{2(N_g - 1)}. \tag{2}$$

For gene sets that are large enough, say > 20, $E_g^{\chi^2}$ follows a standard normal distribution as well. Thus, computing p-values and adjusting these is just as straightforward as for the z-test.

3 Results

We computed the z-score and normalised χ^2 for all gene sets and all datasets presented in Mootha *et al*. and Subramanian *et al*. We used the latest version of GSEA. We adjusted for multiple comparisons using Storey's q-value.[15] We compared these to the q-values computed using GSEA. Table 1 shows all the gene sets achieving a GSEA q-value of less than 0.25, as done by Subramanian *et al*. With the exception of only three cases out of 4139, all gene sets found by GSEA to have q-values < 0.025 were either in the top 10 gene sets or had a q-value less than 0.05 for either the z-test or the χ^2-test. The three

Table 1 For each of the eight datasets studies by Mootha *et al*. and Subramanian *et al*. we found the gene sets for which GSEA reports a q-value of 0.25 or less. Note that the Stanford dataset had no gene sets passing this requirement. For the rest we show the q-values obtained for these same gene sets when using the z-test and the χ^2-test. The ranks of the gene sets obtained with each of these three methods, within the dataset, are also shown. There are only three examples for which the q-value was larger than 0.05 and the rank was larger than 10 in both the z-test and the χ^2- test. These are shown in bold

Study	Gene set	Size	GSEA		z-test		χ^2-test	
			q-value	Rank	q-value	Rank	q-value	Rank
Diabetes	MAP00360 Phenylalanine metabolism	23	0.06	2	0.07	9	0.6	46
Diabetes	MAP00910 Nitrogen metabolism	30	0.3	3	<0.01	6	0.6	43
Diabetes	OXPHOS HG-U133A probes	114	0.04	1	<0.001	1	0.6	66
Gender 1	chrY	40	<0.001	1	<0.001	1.5	<0.001	2.5
Gender 1	chrYp11	18	<0.001	3	<0.001	3	<0.001	2.5
Gender 1	chrYq11	16	<0.001	2	<0.001	1.5	<0.001	2.5
Gender 2	XINACT MERGED	20	<0.001	1	<0.001	6	<0.001	2
Gender 2	GNF FEMALE GENES	85	0.05	3	<0.001	7	<0.001	2
Gender 2	TESTIS GENES	73	0.02	2	<0.001	2.5	<0.001	2
P53	rasPathway	22	0.2	6	<0.01	5	0.9	123
P53	p53hypoxiaPathway	20	<0.001	2	0.03	22	<0.001	1
P53	hsp27Pathway	15	<0.001	2	0.01	14	0.4	40
P53	p53Pathway	16	<0.001	2	<0.01	4	<0.001	2

(*continued*)

Table 1 Continued

Study	Gene set	Size	GSEA		*z*-test		χ^2-test	
			q-value	Rank	*q*-value	Rank	*q*-value	Rank
P53	P53 UP	40	0.01	4	<0.001	2	<0.001	6
P53	radiation sensitivity	26	0.08	5	0.02	16	<0.001	3
Leukemia	chr6q21	31	0.01	1	<0.001	2	0.8	23
Leukemia	chr5q31	59	0.05	2	0.03	7	0.1	86
Leukemia	chr13q14	31	0.06	3	0.2	16	0.4	7
Leukemia	chr14q32	64	0.08	5	<0.01	3	<0.01	2
Leukemia	chr17q23	39	0.07	4	<0.01	4	0.7	18
Boston	p53hypoxiaPathway	19	0.05	1	<0.001	13	<0.01	18
Boston	Aminoacyl tRNA biosynthesis	15	0.1	5	<0.001	12	0.2	63
Boston	INSULIN 2F UP	113	0.1	2	<0.001	2.5	<0.01	22
Boston	tRNA Synthetases	16	0.2	7	<0.001	9	0.3	91
Boston	LEU DOWN	124	0.1	4	<0.001	2.5	<0.01	27
Boston	HTERT UP	104	0.1	3	<0.001	5	0.05	38
Boston	GLUT DOWN	199	0.2	6	<0.001	2.5	<0.001	8
Boston	cell cycle checkpoint	19	0.2	8	<0.001	16	0.3	98
Michigan	amiPathway	22	0.01	3.5	<0.001	6.5	1	208.5
Michigan	cskPathway	22	0.01	3.5	<0.001	6.5	1	208.5
Michigan	badPathway	19	<0.01	2	0.03	29	0.9	151
Michigan	Il12Pathway	22	0.05	6	0.01	23	0.9	79
Michigan	no2il12Pathway	16	0.08	7	0.02	25	1	246
Michigan	**GO ROS**	18	0.09	8	0.06	54	0.9	156
Michigan	**tob1Pathway**	18	0.2	17	0.06	53	0.9	69
Michigan	HEMO TF LIST JP	66	0.2	13	<0.01	18	1	245
Michigan	ctla4Pathway	16	0.2	20	<0.01	10	0.9	26
Michigan	**ST G alpha i Pathway**	29	0.2	16	0.05	50	0.9	68
Michigan	MAP00010 Glycolysis Gluconeogenesis	45	<0.01	1	<0.001	8	0.9	30
Michigan	vegfPathway	21	0.03	5	<0.01	17	1	173
Michigan	INSULIN 2F UP	113	0.2	9	<0.001	2	0.9	65
Michigan	insulin signalling	77	0.2	10	0.04	39	0.9	8
Michigan	HTERT UP	104	0.2	12	<0.001	5	0.3	4
Michigan	MAP00251 Glutamate metabolism	18	0.2	14	0.01	21	0.9	19
Michigan	ceramidePathway	18	0.2	15	<0.01	19	0.9	111
Michigan	p53 signalling	65	0.2	11	<0.01	11	0.9	60
Michigan	tRNA Synthetases	16	0.2	18	<0.01	14	0.9	55
Michigan	MAP00970 Aminoacyl tRNA biosynthesis	15	0.2	19	<0.01	16	0.9	73

cases are highlighted with bold letters in Table 1. Notice that all three were found in the Michigan Lung Cancer dataset.

Figure 3 shows two gene sets: the GO ROS group in the Michigan Lung Cancer dataset and the GLUT DOWN gene set in one of the Gender datasets. GO ROS would be considered interesting in the Michigan Lung Cancer study by GSEA but not by the simpler methods. GLUT DOWN would be considered interesting in the Gender dataset

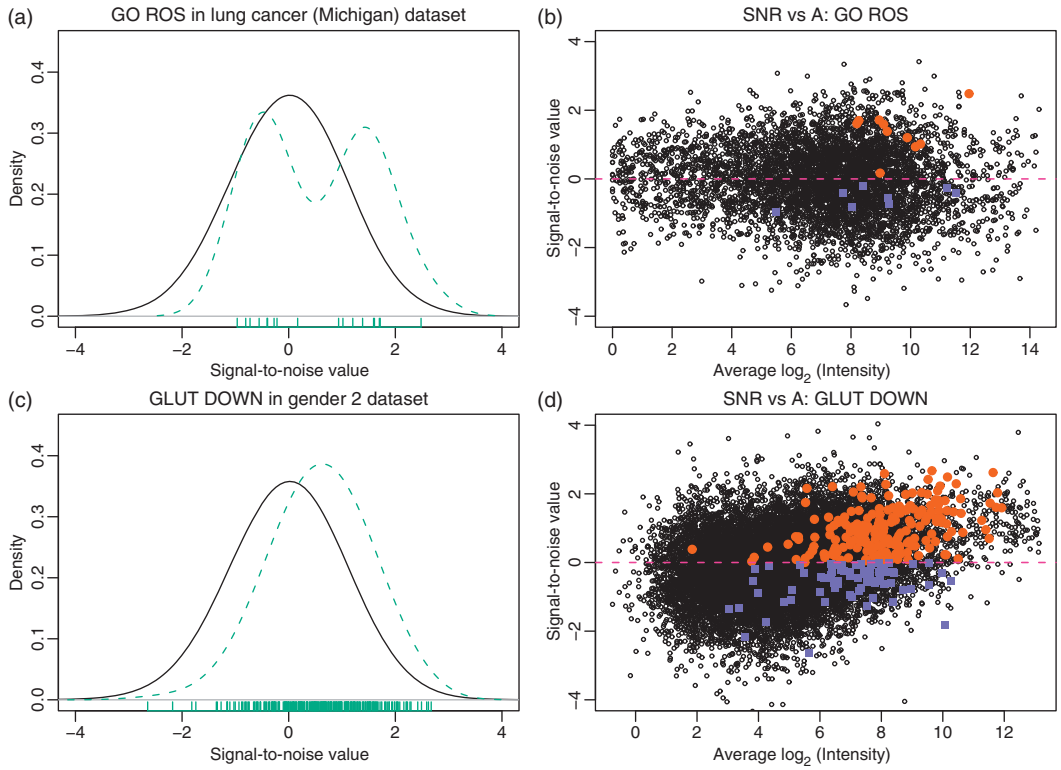

Figure 3 Gene sets showing disagreement between GSEA and the z-test. (a) Empirical density estimate of the the signal to noise values for the GO ROS group (dashed lines) and the rest of the genes (solid line). The ticks on the x-axis show the actual observations. This particular group had a small GSEA q-value but a z-test and $\chi^2 > 0.25$. (b) For each gene, signal-to-noise values plotted against the average intensity for the same dataset as in (a). The values for the GO ROS gene set are highlighted. Circles denote the up-regulated genes in the gene set and squares denote the down-regulated genes. (c) As (a) but for the GLUT DOWN gene set in the gender dataset. The z-test approach results in a very small q-value (< 0.001) for this gene set but a GSEA q-value larger than 0.25. (d) As (b) but for the data described in (c).

by the z-test but not by GSEA. The only interesting feature of the GO ROS group is a gap (no observations of t_i) between 0.5 and 1. We do not consider this to be interesting enough to merit detection. On the other hand, the GLUT DOWN has a clear shift in mean. Figure 4 shows the other two gene sets found by GSEA and not by the other methods. They do not appear interesting in any way.

Subramanian *et al.* point out that there is very little agreement in the results obtained from the three lung cancer datasets they studied. They demonstrate the advantages of GSEA over the marginal approach by showing better agreement between aggregate scores as compared to marginal ones. We created lists of the top gene sets for these three studies using four different approaches: the top 30 gene sets (lowest q-values) in each group as found by GSEA and the z-test, all the gene sets with FDR < 0.25 for GSEA, and all the gene sets with FDR < 0.05 for the z-test. Figure 5 shows Venn diagrams for the results. It is clear that much better agreement is found with the z-tests than with GSEA.

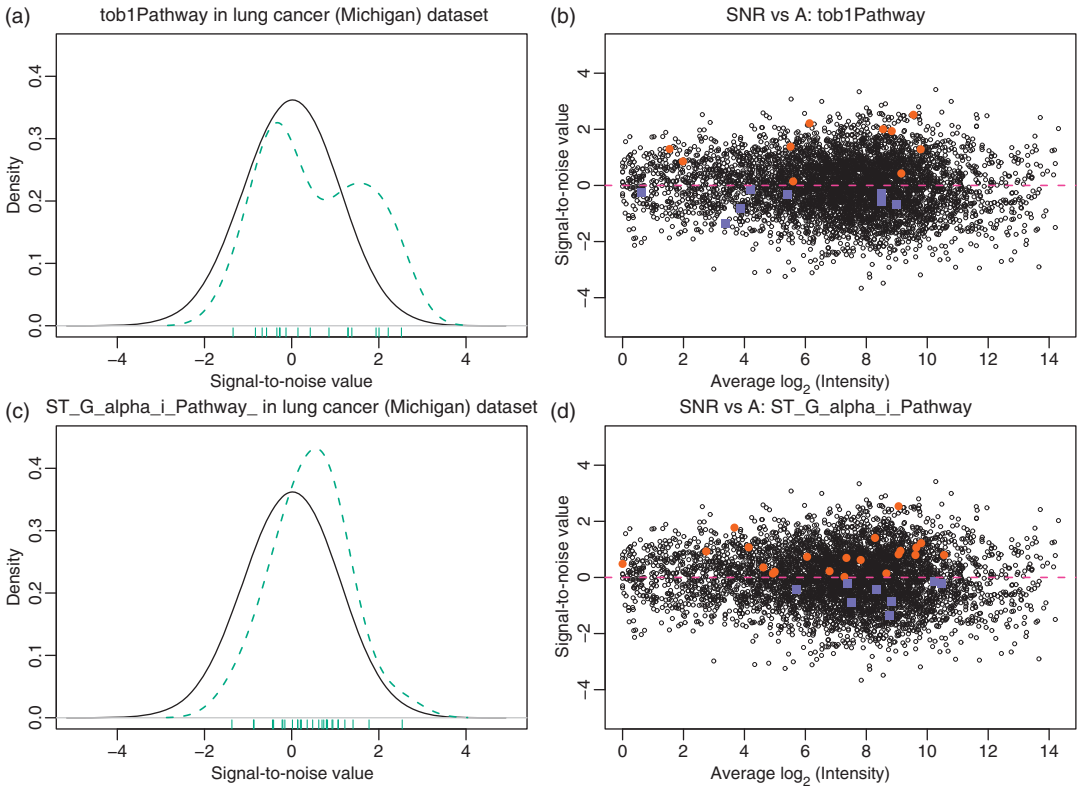

Figure 4 As Figure 3 but for the two other gene sets found by GSEA and not by the z-test or χ^2-test.

4 Discussion

We have compared GSEA to two very simple procedures based on standard statistical approaches: the one-sided z-test and the χ^2-test. We found that the simpler methods outperformed GSEA in assessments based on the eight datasets used in the GSEA papers and a simulation study. The great majority of gene sets found by GSEA to be interesting are also found by the z-test. Notice that if we expect gene sets to be interesting due to mean shifts then it is no surprise that the z-test outperforms GSEA since statistical theory predicts this test to be much more powerful than the K–S test. In fact, this is one reason we use the 0.05 cut-off, instead of 0.25, for the z-test q-value. An argument for GSEA could be that some gene sets are interesting for reasons other than mean shifts, such as scale changes. For many of these cases the χ^2-test was able to identify them as interesting. The only three gene sets not found by either the z-test or χ^2-test are shown in Table 1, Figures 3 and 4. For all three it is hard to argue that they are interesting in any way. We notice that all three gene sets are small in size as compared to other gene sets and have unexpected gaps in the observations of the signal-to-noise values. It is possible that the *ad hoc* modification of the K–S test is biased in favour of small gene sets.

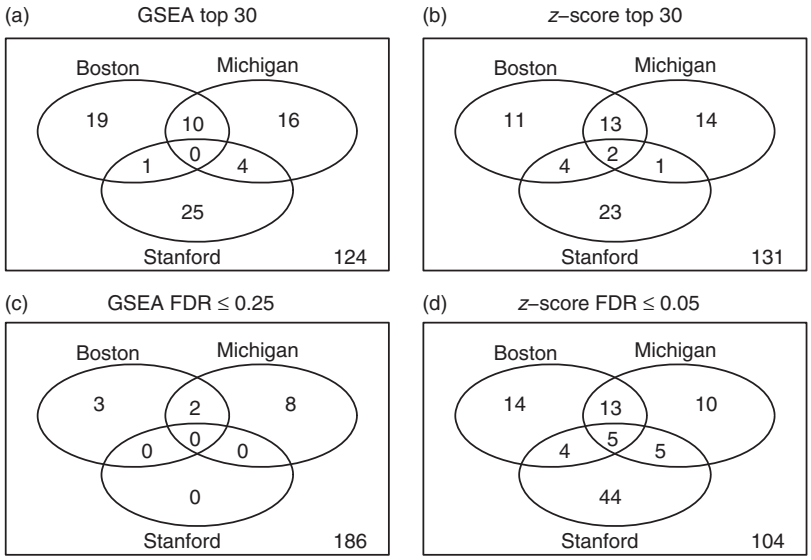

Figure 5 Gene set agreement, shown with Venn diagram, in the Boston, Michigan and Stanford lung cancer dataset. The numbers in the lower right corners are the number of gene sets that were not in any list. (a) Agreement among top 30 gene sets ranked by their GSEA q-value. (b) As (a) but for the z-test. (c) Agreement among gene sets achieving a GSEA q-value smaller than 0.25. (d) As (c) but for gene sets achieving a q-value smaller than 0.05 with the z-test.

Another advantage of the method presented here is that it can be easily extended to application other than the comparison of two conditions. There is no need for the statistics used to compute the enrichment scores described here, Equations (1) and (2), to be *t*-statistics. Any statistics that we expect to follow a standard normal distribution can be used. For example, another common application of microarrays examines cancer survival data. In these cases the summary statistic is commonly a parameter estimate from a survival model. The standard normal approximation is a common approximation of the standardised versions of these estimates. Tian *et al.*[14] argue against the use of the normal approximation for the averaged *t*-tests and propose the use of permutation-based tests. A disadvantage of their proposed permutation tests is that they are not easily extended to cases other than comparison of two conditions. Tian *et al.* correctly point out that if the *t*-statistics are correlated under the null hypothesis, the assumption that the *z*-score is normal with standard deviation 1 is incorrect. We did not find this to be a problem in practice. Furthermore, we find that the the average correlation in gene sets is of the order of 0.1 (data not shown), which only corresponds to a 5% inflation of the score. A correction factor can easily be inserted at the appropriate place.

An entirely parametric approach, as the one described here, has been previously proposed by Kim and Volsky.[12] Their approach, referred to as PAGE, ignores the marginal *t*-tests, and computes a *t*-test based on the effect sizes (log fold changes) within each gene set. A limitation of this approach is that it does not take into account the gene-specific variances. This is problematic because different genes are known to result in measurements with different variances.[16] Furthermore, PAGE is restricted to applications of comparing two conditions. However, we expect PAGE to outperform GSEA as well.

We have made an argument against the use of GSEA. Methods that are much simpler, require hardly any computation time, and can be easily implemented in any data analysis package, have been demonstrated to outperform GSEA. However, we do not think the methods we have described here are a final solution. We describe them here because they are an obvious first step that has been ignored. Efron and Tibshirani[17] propose an approach that includes a statistic that specifically targets gene sets with only a fraction of the genes differentially expressed and a novel permutation approach. Falcon and Gentleman[18] describe methodology that takes into account the fact that overlap exists between the different gene sets. These approaches certainly seem promising.

References

1 Tusher VG, Tibshirani R, Chu G. Significance analysis of microarrays applied to the ionizing radiation response. *Proceedings of the National Academy of Sciences* 2001; **98**(18): 10515.

2 Smyth GK. Linear models and empirical bayes methods for assessing differential expression in microarray experiments. *Statistical Applications in Genetics and Molecular Biology* 2004; **3**(1): 1–25.

3 Storey JD, Tibshirani R. Statistical significance for genomewide studies. *Proceedings of the National Academy of Sceinces* 2003; **100**(16): 9440–5.

4 Dudoit S, Shaffer JP, Boldrick JC. Multiple hypothesis testing in microarray experiments. *Statistical Science* 2003; **18**: 71–103.

5 Lee HK, Braynen W, Keshav K, Pavlidis P. Erminej: Tool for functional analysis of gene expression data sets. *BMC Bioinformatics* 2005; **6**: 269.

6 The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nature Genetics* 2000; **25**: 25–29.

7 Mootha VK, Lindgren CM, Eriksson KF *et al*. Pgc-1alpha-responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nature Genetics* 2003; **34**(3): 267–73.

8 Pavlidis P, Lewis DP, Noble WS. Exploring gene expression data with class scores. In *Pacific Symposium on Biocomputing* 2002: 474–85.

9 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. *Neurochemical Research* 2004; **29**: 1213–22.

10 Goeman J, van de Geer S, de Kort F, van Houwelingen H. A global test for groups of genes: testing association with a clinical outcome. *Bioinformatics* 2004; **20**: 93–99.

11 Goeman JJ, Oosting J, Cleton-Jansen AM, Anninga JK, van Houwelingen HC. Testing association of a pathway with survival using gene expression data. *Bioinformatics* 2005; **21**(9): 1950–7.

12 Kim SY, Volsky DJ. PAGE: Parametric analysis of gene set enrichment. *BMC Bioinformatics* 2005; **6**(144): 1–12.

13 Subramanian A, Tamayo P, Mootha VK *et al*. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of National Academy of Sciences* 2005; **102**(43): 15545–50.

14 Tian L, Greenberg SA, Kong SW, Altschuler J, Kohane IS, Park PJ. Discovering statistically significant pathways in expression profiling studies. *Proceedings of National Academy of Sciences* 2005; **102**(38): 13544–9.

15 Storey JD. A direct approach to false discovery rates. *Journal of the Royal Statistical Society B* 2002; **64**: 479–98.

16 Kendziorski C, Irizarry RA, Chen KS, Haag JD, Gould MN. On the utility of pooling biological samples in microarray experiments. *Proceedings of National Academy of Sciences* 2005; **102**(12): 4252–7.

17 Efron B, Tibshirani R. On testing the significance of sets of genes. *Annals of Applied Statistics* 2007; **1**(1): 107–29.

18 Falcon S, Gentleman R. Using gostats to test gene lists for go term association. *Bioinformatics* 2006; **567**(1): 257–8.