

SELECTING GENES WITH DISSIMILAR DISCRIMINATION STRENGTH FOR SAMPLE CLASS PREDICTION

Zhipeng Cai, Randy Goebel, Mohammad R. Salavatipour, Yi Shi, Lizhe Xu, and Guohui Lin*

*Department of Computing Science, University of Alberta
Edmonton, Alberta T6G 2E8, Canada
Email: zhipeng, goebel, mreza, ys3, ghlin@cs.ualberta.ca*

One of the main applications of microarray technology is to determine the gene expression profiles of diseases and disease treatments. This is typically done by selecting a small number of genes from amongst thousands to tens of thousands, whose expression values are collectively used as classification profiles. This gene selection process is notoriously challenging because microarray data normally contains only a very small number of samples, but range over thousands to tens of thousands of genes. Most existing gene selection methods carefully define a function to score the differential levels of gene expression under a variety of conditions, in order to identify top-ranked genes. Such single gene scoring methods suffer because some selected genes have very similar expression patterns so using them all in classification is largely redundant. Furthermore, these selected genes can prevent the consideration of other individually-less but collectively-more differentially expressed genes. We propose to cluster genes in terms of their class discrimination strength and to limit the number of selected genes per cluster. By combining this idea with several existing single gene scoring methods, we show by experiments on two cancer microarray datasets that our methods identify gene subsets which collectively have significantly higher classification accuracies.

Keywords: Microarray data analysis; gene clustering; gene selection; classification

1. Introduction

DNA microarrays provide the opportunity to measure the expression levels of thousands of genes simultaneously. This novel technology supplies us a large volume of data to systematically understand various gene regulations under different conditions. As one of the main applications, it is very important to determine the gene expression profiles of diseases and disease treatments. Among the thousands of genes in the arrays, many of them do not have their expression values distinguishably changed across different condition, e.g., so-called “house keeping” genes. These genes certainly would not be very useful in profiling since they do not contribute much to disease or treatment class recognition. In practice, a small number, typically in tens, of genes that are highly differentially expressed across different conditions are to be selected to compose profiles for the purpose of class prediction. This process is known as *gene selection*; there are many existing methods, which typically define a function to score the level of how differentially expressed a gene is, under different conditions, and identify those top ranked genes.¹⁻⁵ Such single gene scoring methods typically suffer the problem that some selected genes have very simi-

lar expression patterns, therefore using them all in classification is largely redundant, and those selected genes prevent other individually-less but collectively-more differentially expressed genes from being selected.

Several other gene selection methods have recognized the problem with the redundancy of some highly expressed genes, and look for a subset of genes that collectively maximize the classification accuracy. For example, Xiong *et al.* define a function to measure the classification accuracy for individual genes and select a subset of genes through *Sequential Forward [Floating] Selection* (SF[F]S),⁶ which was developed decades ago for general feature selection. Guyon *et al.* propose another method that uses *Support Vector Machines* (SVMs) and *Recursive Feature Elimination* (RFE).⁷ In terms of effectiveness, these gene selection methods perform much better than those single gene scoring methods, since they measure the classification strength of the whole set of selected genes. Computationally, they are essentially heuristics which replace exhaustive enumeration of an optimal subset of genes which typically take much longer to return a solution. The inefficiency of these methods actually prevent them from being used in

*To whom correspondence should be addressed.

practice. Nevertheless, there are alternative implementations of the key idea, which is to exclude a gene when there is already a similar gene been selected.

We propose another implementation which first clusters genes according to their class discrimination strength, namely, two genes that have very close class discrimination strength are placed in a common cluster; we then limit the number of genes per cluster to be selected. This provides a more efficient clustering process which, when combined with a single gene scoring method, leads to an efficient and effective gene selection algorithm. We call our method an *EEGS-based gene selection method*. In the next section, we present the details of a novel measure of class discrimination strength difference between two genes, using their expression values. With this distance measure, we briefly explain how to adopt the *k-means* algorithm⁸ to cluster genes. We also briefly introduce three single gene scoring methods, namely F-test³, Cho⁴ and GS,⁵ and two classifiers, namely, a linear kernel SVM classifier⁷ and a *k* Nearest Neighbor (KNN) classifier.³ Finally, we outline a complete high level description of the EEGS-based gene selection methods. In Section 3, we briefly introduce our performance measurements, followed by the dataset descriptions, and our experimental results. Section 4 discusses parameter selection, the effects of variety within data sets, classifiers and the performance measurements, and finally, the overall results compared to single gene scoring methods. Section 5 summarizes our main contributions, our conclusions on the suitable datasets for the EEGS-based methods, and some plans for future work.

2. The EEGS-Based Gene Selection Methods

There are two challenges in microarray data classification. One is class discovery to define previously unrecognized classes. The other is to assign individual samples to already-defined classes, which is the focus here.

2.1. The Performance Measurements

The genes selected by a method are evaluated by their class discrimination strength, measured by the classification accuracy, defined as follows. For gene selection purposes, a number of microarray samples with known class labels are provided, which form a

training dataset. The selected genes are then used for building a classifier, which can take a new microarray sample and assign it a class label. The set of such samples for testing purpose is referred to as the *testing dataset*, and the percentage of the correctly labeled samples is defined as the *classification accuracy* of the method (on this particular testing dataset). Note that we have to have the class labels for the samples in the testing dataset in order to calculate the classification accuracy. For computational convenience, given a microarray dataset whose samples all have known class labels, only a portion of it is used to form the training dataset; the rest of the samples have their class labels removed and are used to form the testing dataset. There are two popular cross validation schemes adopted in the literature to evaluate a method, which are *l-Fold* and *Leave One Out* (LOO). In the *l-Fold* cross validation, the whole dataset is (randomly) partitioned into *l* equal parts and, at one time, one part is used as testing dataset and the other *l* - 1 parts are used as training dataset. The process is repeated for each part and the average classification accuracy over these *l* ones is taken as the final classification accuracy. Here set *l* = 5 and repeat the process for 20 iterations. Therefore, the final classification accuracy is the average over 100 values. For each of the LOO cross validation experiments, only one sample is used as a testing dataset, and the others form the training dataset. The process is repeated for every sample and the final classification accuracy is defined as the percentage of the correctly labeled samples. We report both the LOO classification accuracies and the 5-Fold classification accuracies for all the six tested gene selection methods in Section 3.

2.2. The Classifiers

We adopt two classifiers in our study. One is a linear kernel SVM classifier that has been used in Guyon *et al.*⁷ and the other is a KNN classifier that has been used in Dudoit *et al.*³ Essentially, with a given set of selected genes determined by some gene selection method, the SVM classifier, which contains multiple SVMs, finds decision planes to best separate the labeled samples based on the expression values of these selected genes. Subsequently, it uses this set of decision planes to predict the class label of a test sample. For a more detailed explanation of how the

decision planes are constructed, the readers are referred to Guyon *et al.*⁷ The KNN classifier predicts the label of a testing sample in a different way. Using the expression values of (only) the selected genes, the classifier identifies the k most similar samples in the training dataset. It then uses the class labels of these k similar samples through a majority vote. In our experiments, we set the value of k to be 5 as default, after testing for several values in the range 4 to 10.

2.3. The Single Gene Scoring Methods

As mentioned in the introduction, many of the existing gene selection methods are single gene scoring methods that define a function to approximate the class discrimination strength of a gene.¹⁻⁵ Typically, an F-test gene selection method^{2, 9} is presented, which basically captures the variance of the class variances of the gene expression values in the dataset. A bigger variance indicates that a gene is more differentially expressed and thus ranked higher. Because class sizes might differ a lot, Cho *et al.*⁴ proposed a weighted variant, which was further refined by Yang *et al.*⁵ We denote these three single gene scoring methods as F-test, Cho and GS, respectively, and combine the EEGS idea with them to have the EEGS-based methods, denoted as EEGS-F-test, EEGS-Cho and EEGS-GS, respectively.

2.4. Gene Clustering

Gene clustering in microarray data analysis is an independent research subject, in which genes having a similar expression pattern are clustered for certain applications. In our work here, we are particularly interested in the class discrimination strength of the genes, since we do not want to select too many genes that have similar class discrimination strength. Note that genes having a similar expression pattern would certainly have similar class discrimination strength, but the other way around is not necessarily true. Therefore, we define a new measure trying to better capture the difference in the class discrimination strength between two genes.

Assume there are p genes and n samples in the microarray training dataset, and these n samples belong to L distinct classes. Let a_{ij} denote the expression value of gene i in sample j . This way, the

training dataset can be represented as a matrix

$$A_{p \times n} = (a_{ij})_{p \times n}.$$

Let C_1, C_2, \dots, C_L denote the L classes, and $n_q = |C_q|$, for $q = 1, 2, \dots, L$. Let \bar{a}_{iq} be the mean expression value of gene i in class C_q : $\bar{a}_{iq} = \frac{1}{n_q} \sum_{j \in C_q} a_{ij}$, for $q = 1, 2, \dots, L$. The centroid matrix is thus

$$\bar{A}_{p \times L} = (\bar{a}_{iq})_{p \times L}.$$

The *discrimination strength vector* of gene i is defined as

$$v_i = \langle |\bar{a}_{iq_1} - \bar{a}_{iq_2}| \mid 1 \leq q_1 < q_2 \leq L \rangle,$$

where the order of $\frac{1}{2}L(L-1)$ vector entries is fixed the same for all genes, for example the lexicographical order. After all the discrimination strength vectors have been calculated, the k -means algorithm⁸ is applied to cluster these p genes into k clusters using their discrimination strength vectors. Essentially, k -means is a centroid-based clustering algorithm that partitions the genes based on their pairwise distances. We adopt both the Euclidean distance and the Pearson correlation coefficient in our experiments. Again, we have tested several values of k in the k -means algorithm (cf. Section 4.1) and we have set it to 100 as default.

The following small example is cut from the CAR dataset (cf. Section 3.1 and Section 4.3), including 4 genes and their expression values in 9 samples: $G_1 = 40794_at$, $G_2 = 660_at$, $G_3 = 32200_at$, and $G_4 = 41238_s_at$. These 9 samples belong to three classes, C_1, C_2, C_3 , and $C_1 = \{S_1, S_2\}$, $C_2 = \{S_3, S_4, S_5\}$, and $C_3 = \{S_6, S_7, S_8, S_9\}$. $A_{4 \times 9} =$

$$\left(\begin{array}{c|cccc|cccc} & S_1 & S_2 & S_3 & S_4 & S_5 & S_6 & S_7 & S_8 & S_9 \\ \hline G_1 & 1.301 & 1.301 & 1.301 & 1.301 & 1.301 & 4.212 & 4.123 & 4.005 & 4.301 \\ G_2 & 1.301 & 1.301 & 1.301 & 1.301 & 1.301 & 3.914 & 3.954 & 3.835 & 4.112 \\ G_3 & 1.531 & 1.491 & 1.301 & 1.301 & 1.301 & 3.809 & 4.032 & 3.684 & 3.842 \\ G_4 & 1.301 & 1.301 & 3.869 & 4.270 & 4.459 & 1.301 & 1.301 & 1.301 & 1.301 \end{array} \right).$$

It follows that the centroid matrix $\bar{A}_{4 \times 3}$ is

$$\bar{A}_{4 \times 3} = \left(\begin{array}{c|ccc} & C_1 & C_2 & C_3 \\ \hline G_1 & 1.301 & 1.301 & 4.160 \\ G_2 & 1.301 & 1.301 & 3.954 \\ G_3 & 1.511 & 1.301 & 3.842 \\ G_4 & 1.301 & 4.199 & 1.301 \end{array} \right),$$

and the discrimination strength vectors are

$$\begin{aligned} v_1 &= \langle 0.000, 2.859, 2.859 \rangle, \\ v_2 &= \langle 0.000, 2.653, 2.653 \rangle, \\ v_3 &= \langle 0.210, 2.331, 2.541 \rangle, \\ v_4 &= \langle 2.898, 0.000, 2.898 \rangle. \end{aligned}$$

Therefore, the pairwise Euclidean distances are $d(1, 2) = 0.291$, $d(1, 3) = 0.651$, $d(2, 3) = 0.400$, $d(1, 4) = 4.071$, $d(2, 4) = 3.937$, and $d(3, 4) = 3.576$. These distances indicate that the first three genes have very similar discriminatory strength, while gene *41238_s_at* has different discriminatory strength (cf Section 4.3).

2.5. The Complete EEGS-Based Gene Selection Methods

Given a microarray training dataset containing p genes and n samples in L classes, an EEGS-based gene selection method first calls the k -means algorithm (with $k = 100$) to cluster genes. Next, depending on the detailed single gene scoring method integrated in the method, which is one of F-test, Cho and GS, it calls the single gene scoring method to score all the genes and sort them into non-increasing order. Using this gene order and the gene cluster information, the EEGS-based method selects a pre-specified number, x , of top ranked genes with the constraint that there are at most T genes per cluster can be selected. In more details, it scans through the gene order and picks up a gene only if there are less than T genes from the same cluster been selected. These x selected genes are then fed to classifier construction, either the SVM classifier or the KNN classifier. In our experiments, we have tested x ranging from 1 to 80 and several values for T (cf. Section 3.2). We have set $T = 1$ as default (cf. Section 4.1).

Depending on the single gene scoring method integrated into the EEGS-based gene selection method, which is one of F-test, Cho and GS, the method is referred to as EEGS-F-test, EEGS-Cho and EEGS-GS, respectively.

3. Experimental Results

We compare the three EEGS-based gene selection methods with the three ordinary gene selection methods, measured by the classification accuracy. For ease of presentation, we separate the 5-Fold cross

validation results and the LOO cross validation results. Note that we have adopted two distance measures in the k -means clustering algorithm. We have a broader collection of experimental results, but here report only those based on the Euclidean distance, as there is essentially no difference between the results based on the Pearson correlation coefficient (cf. Section 4.1). Note also that we have adopted two classifiers, a linear kernel SVM classifier and a KNN classifier. We choose to plot their classification accuracies together labeled by different notations, for instance, EEGS-Cho-KNN labels the accuracies of the KNN classifier. The experiments are done on two real cancer microarray datasets, CAR¹⁰ and LUNG,⁹ whose details are described in the following subsection.

3.1. Dataset Descriptions

The CAR dataset contains 174 samples in eleven classes: *prostate*, *bladder/ureter*, *breast*, *colorectal*, *gastroesophagus*, *kidney*, *liver*, *ovary*, *pancreas*, *lung adenocarcinomas*, and *lung squamous cell carcinoma*, which have 26, 8, 26, 23, 12, 11, 7, 27, 6, 14, and 14 samples, respectively.¹⁰ Each sample originally contained 12,533 genes. We preprocessed the dataset as described in Su *et al.*¹⁰ to include only those probe sets whose maximum hybridization intensity in at least one sample is ≥ 200 ; Subsequently, all hybridization intensity values ≤ 20 , including negative hybridization intensity values, were raised to 20, and the values were log transformed. After preprocessing, we obtained a dataset of 9,182 genes.

The LUNG dataset⁹ contains in total 203 samples in five classes: *adenocarcinomas*, *squamous cell lung carcinomas*, *pulmonary carcinoids*, *small-cell lung carcinomas* and *normal lung*, which have 139, 21, 20, 6, 17 samples, respectively. Each sample originally had 12,600 genes. A preprocessing step which removed genes with standard deviations smaller than 50 expression units, produced a dataset with 3,312 genes.⁹

3.2. Cross Validation Classification Accuracies

The classification accuracies reported here were obtained under the default setting which uses Euclidean distance, $k = 100$ in the k -means clustering algorithm, and at most $T = 1$ gene per cluster could be

selected. On each of the two datasets, all six gene selection methods, F-test, Cho, GS, EEGS-F-test, EEGS-Cho, and EEGS-GS, were run and both the 5-Fold and LOO cross validation classification accuracies were collected and plotted in Figures 1 – 4.

Obviously, these plots show that regardless of which cross validation scheme and which classifier were used, the classification accuracies of the EEGS-based gene selection methods were significantly higher than that of their non-EEGS-based counterparts. Typically, on the CAR dataset, the classification accuracies of the EEGS-based methods were unbelievably higher — though the difference between the classification accuracies became smaller with the increasing number of selected genes, it remained to be more than 10%. From Figures 1 and 2, among the single gene scoring methods, another observation is that the GS method performed better than the Cho method and the F-test method. The EEGS-based methods had the same performance tendency on the CAR dataset. On the LUNG dataset, similar results were obtained, although the performance differences between the EEGS-based methods and the non-EEGS-based methods were smaller than those on the CAR dataset.

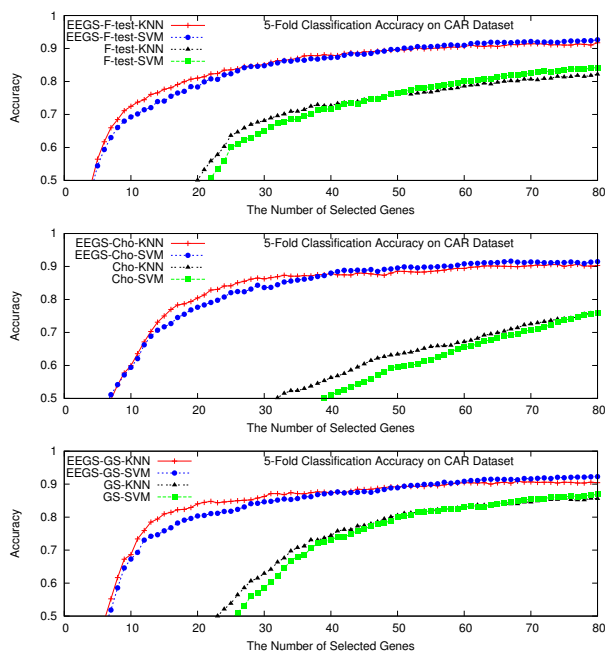


Fig. 1. The 5-Fold cross validation classification accuracies of the six gene selection methods on the CAR dataset.

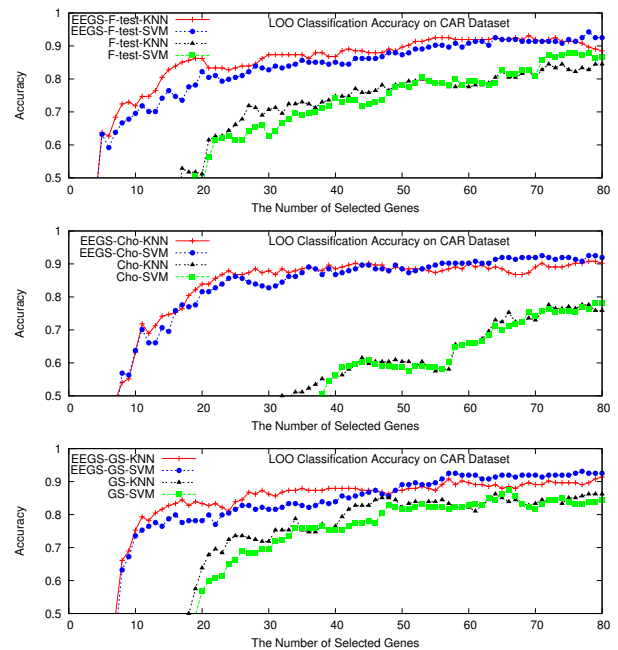


Fig. 2. The LOO cross validation classification accuracies of the six gene selection methods on the CAR dataset.

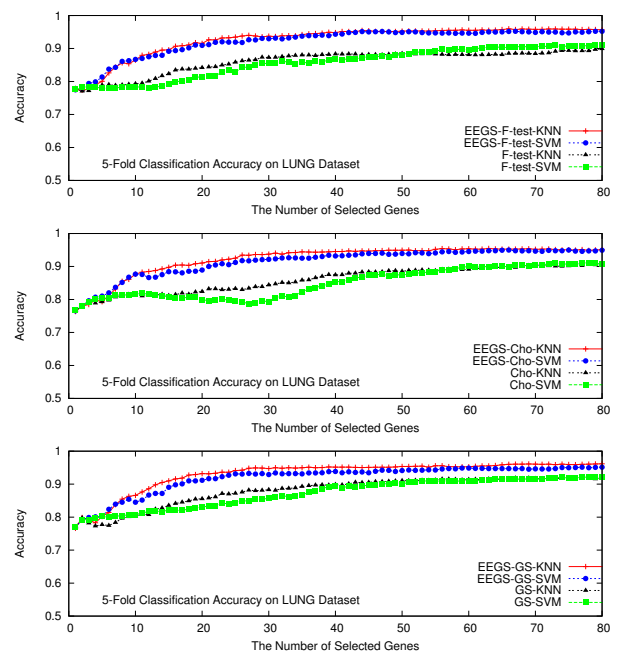


Fig. 3. The 5-Fold cross validation classification accuracies of the six gene selection methods on the LUNG dataset.

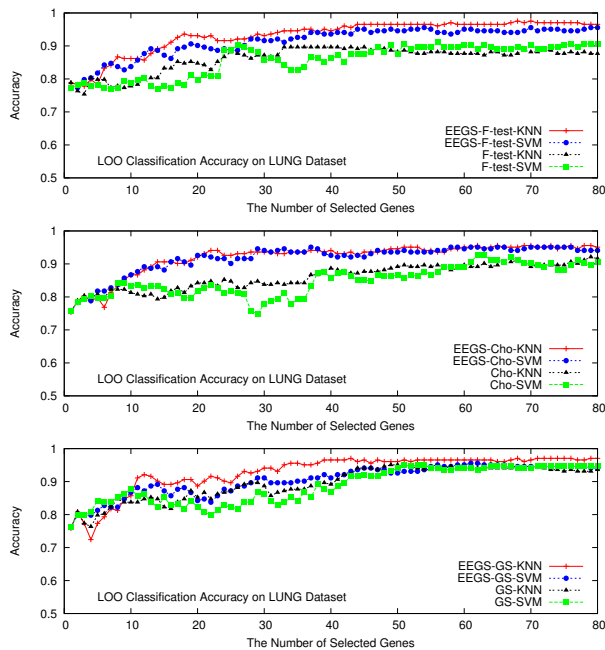


Fig. 4. The LOO cross validation classification accuracies of the six gene selection methods on the LUNG dataset.

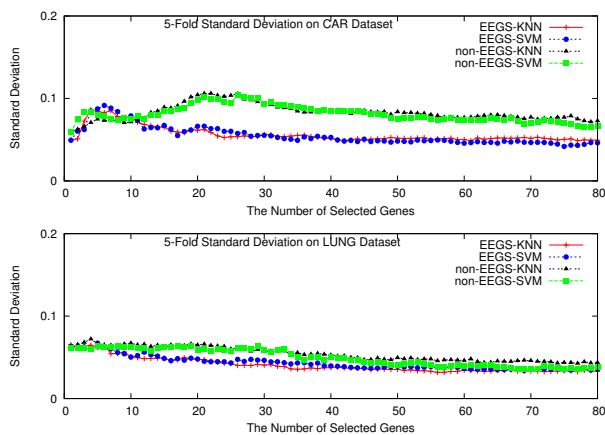


Fig. 5. Plots of average standard deviations of the 5-Fold classification accuracies of the EEGS-based and the non-EEGS-based gene selection methods, combined with the KNN and the SVM classifiers, on the CAR dataset (top) and LUNG dataset (bottom), respectively.

We have also calculated the standard deviations of the 5-Fold cross validation classification accuracies. Note that the accuracies plotted in Figures 1 and 3 were averages over 100. Figure 5 plots the av-

erage standard deviations of the EEGS-based methods and the non-EEGS-based methods, on the CAR and the LUNG datasets, respectively. Namely, the EEGS-KNN plot records the average standard deviations of the three EEGS-based methods (EEGS-F-test, EEGS-Cho and EEGS-GS) combined with the KNN classifier, and the non-EEGS-SVM plot records the average standard deviations of the three non-EEGS-based methods (F-test, Cho and GS) combined with the SVM classifier, and so on. These results show that the standard deviations of the classification accuracies of the EEGS-based methods were even smaller than those of the non-EEGS-based methods, indicating that the EEGS-based methods performed more consistently.

Having observed the performance differences between the EEGS-based and the non-EEGS-based methods, we wanted to tell if they were statistically significant. Analysis of variance (ANOVA) is a statistical procedure to test hypotheses about differences between two or more means. We chose the ANOVA tools of the SPSS system (<http://spss.com/>). In addition to other parameters, ANOVA gives a significance, or p value, of the hypothesis that the two methods under comparison have no difference, at a specified confidence level, which was set to 95% in our experiments. Low p values, normally less than 0.001, indicate that the hypothesis is false, that is, the two methods do have difference. ANOVA also gives a confidence interval for the difference of the accuracy means associated with the two methods, and if the interval contains 0, then the two methods did not have much difference; Otherwise, the hypothesis is false. We collected all the detailed ANOVA values for each EEGS-based gene selection method versus the non-EEGS-based method, combined with a classifier. The values on the CAR dataset are in Table 1, for four different numbers of selected genes. The values on the LUNG dataset are not included because of space, but they are available upon request. Clearly, these values all show that the EEGS-based methods were significant better than their non-EEGS-based counterparts.

4. Discussion

4.1. Gene Clustering

We adopted the k -means algorithm for gene clustering, in which k , the number of expected clusters, has

Table 1. The ANOVA analysis of performance differences between the EEGS-based gene selection methods and their non-EEGS-based counterparts, respectively. A lower than 0.001 p -value and an interval $[L, U]$ excluding 0 both indicate that the EEGS-based gene selection method and its non-EEGS-based counterpart performed significantly different. The dataset is CAR and the numbers of selected genes are 20, 40, 60 and 80.

	#Genes = 20			#Genes = 40			#Genes = 60			#Genes = 80		
	p	L	U	p	L	U	p	L	U	p	L	U
Cho-KNN	0.000	0.3823	0.4236	0.000	0.2989	0.3400	0.000	0.1929	0.2285	0.000	0.1279	0.1650
Cho-SVM	0.000	0.3606	0.4069	0.000	0.3499	0.3893	0.000	0.2329	0.2725	0.000	0.1414	0.1761
F-test-KNN	0.000	0.2835	0.3366	0.000	0.1225	0.1638	0.000	0.0969	0.1139	0.000	0.0560	0.0967
F-test-SVM	0.000	0.3098	0.3654	0.000	0.1391	0.1774	0.000	0.0882	0.1245	0.000	0.0625	0.0949
GS-KNN	0.000	0.3472	0.4030	0.000	0.1048	0.1571	0.000	0.0450	0.0842	0.000	0.0228	0.0540
GS-SVM	0.000	0.3780	0.4290	0.000	0.1005	0.1512	0.000	0.0608	0.1005	0.000	0.0281	0.0583

to be set beforehand. Obviously, the value of k will affect the sizes of resultant clusters, and therefore will affect T ultimately, which is the maximum number of genes per cluster to be selected. We chose to empirically determine these two values. To this end, we experimented with 15 values for k : from 10 to 150 (in tens), and five values for T : 1, 2, 3, 4 and 5. All three EEGS-based methods combined with two classifiers were tested on the CAR dataset, under the 5-Fold cross validation, for each combination of k and T . Associated with each value of T , a classification accuracy is defined as the mean value of $100 \times 3 \times 2 \times 15 = 9,000$ ones, where there are 100 runs in the 5-Fold cross validation, 3 EEGS-based methods, 2 classifiers, and 15 values of k in the test. These classification accuracies, with respect to the number of selected genes, are plotted in Figure 6 (top), where $T = 1$ clearly performed the best.

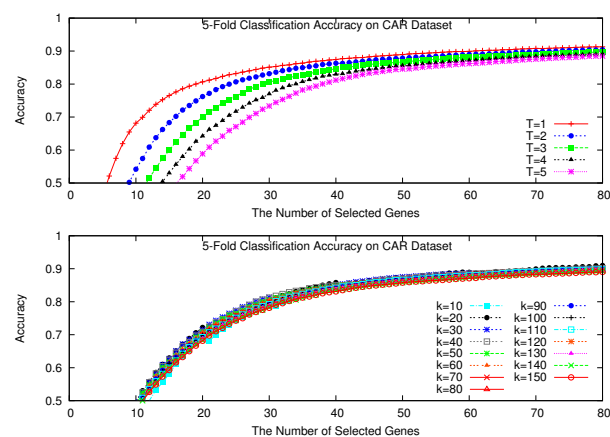


Fig. 6. The effects of the number of clusters in gene clustering and the maximum number of genes per cluster can be selected.

Similarly, associated with each value of k , a classification accuracy is defined as the mean value of $100 \times 3 \times 2 \times 5 = 3,000$ ones, where there are 5 values of T in the test. Again, these classification accuracies, with respect to the number of selected genes, are plotted in Figure 6 (bottom), where we can see that the value of k didn't really affect the performance. Since we decided to set the maximum number of selected genes to be 80, we determined to set $k = 100$ and $T = 1$ as default.

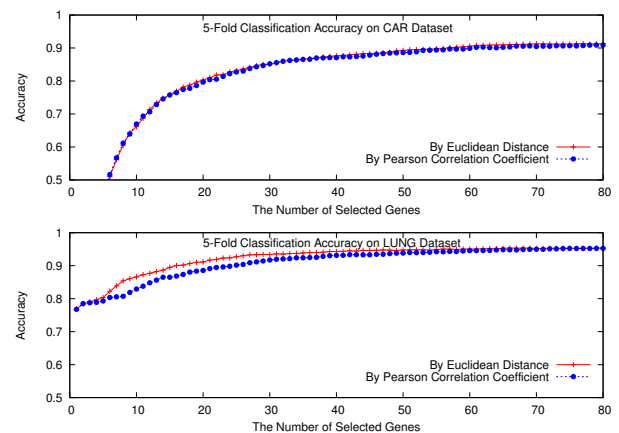


Fig. 7. The effects of the Euclidean distance and the Pearson correlation coefficient in gene clustering.

Another important factor in gene clustering that might affect its performance is the distance measure, for which the Euclidean distance and the Pearson correlation coefficient are two most commonly adopted ones. We have experimented with both of them in the k -means clustering algorithm on the CAR and the LUNG datasets. With the default setting for k and T , we collected the 5-Fold classification accuracies which are the mean values of

$100 \times 3 \times 2 = 600$ ones and plotted them in Figure 7. Clearly seen, the detailed distance measure did not seem to affect the overall performance of the EEGS-based methods, in terms of their classification accuracy. Therefore, we chose the Euclidean distance as our default setting.

4.2. Datasets

Note that in the EEGS-based gene selection methods, a discrimination strength vector is computed for every gene, and genes are clustered using the Euclidean distance defined on their discrimination strength vectors. The main intention for such clustering is to limit the number of genes that have very similar class discrimination strength to be selected, and thus to provide space for other individually-less but collectively-more differentially expressed genes to participate in the class prediction. This goal would not be achieved when there are only two classes in the dataset (binary classification), which would mean that the discrimination strength vectors have only one entry and the EEGS-based method reduces to its component basic gene selection method. For similar reasons, we suspect that the EEGS-based gene selection methods would work well when the number of classes in the dataset is three. The CAR and the LUNG datasets contain eleven and five classes, respectively, and therefore the discrimination strength vectors have 55 and 10 entries, respectively. The EEGS-based gene selection methods all performed excellent on them.

For various reasons, microarray datasets are often unbalanced, that is, the sizes of the classes are highly variable. For example, in the LUNG dataset, the maximum class size is 139 while the minimum class size is only 6. Since it is possible that during the 5-Fold cross validation the random partition produces a training dataset containing only a few samples, or maybe even none, from a small class, the testing would make mistakes on the samples from the same class. To verify how much the dataset unbalance would affect the performance of a gene selection method, we removed the classes of sizes smaller than 10 from the CAR and the LUNG datasets to produce two reduced but more balanced datasets, denoted as CAR^r and $LUNG^r$, respectively. Consequently, the CAR^r and the $LUNG^r$ datasets contain 153 samples in 8 classes and 197 samples in 4 classes, respectively.

We then ran all six methods combined with both the KNN classifier and the SVM classifier on the full and the reduced datasets, and plotted the average classification accuracies (each over three methods with two classifiers, i.e., six values) in Figure 8. In the figure, one can see that the performance of the EEGS-based methods did not change a lot on the reduced CAR^r and the $LUNG^r$ datasets, compared with their performance on the full datasets. Interestingly, for the non-EEGS-based methods, their performance increased significantly on the CAR^r dataset, but not on the $LUNG^r$ dataset. Nevertheless, these results show that the EEGS-based methods performed more stable (and better) than the non-EEGS-based methods on unbalanced datasets. One of the possible reasons is that the EEGS-based methods might be able to select some genes that are signatures of the samples in the small classes, for which further studies are needed to understand better.

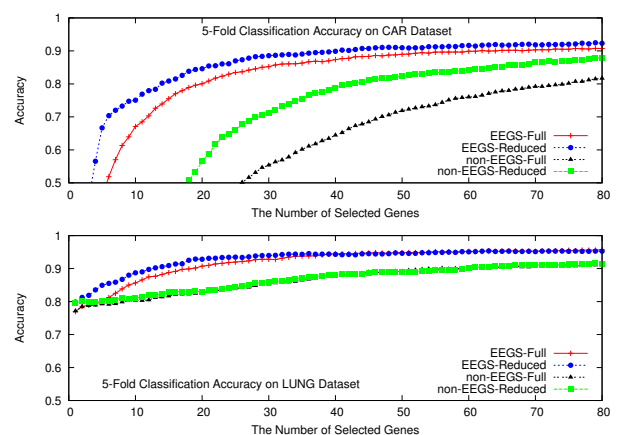


Fig. 8. Classification accuracies of the EEGS-based and the non-EEGS-based methods on the full and reduced datasets, where EEGS-Full plots the average classification accuracies of the EEGS-based methods on the full dataset.

To further understand these two datasets, we refined the statistics to collect the classification *precision* and *recall* for each class in the 5-Fold cross validation. We fixed the number of selected genes to be 20. For each class, the classification precision is defined as the number of correctly labeled samples divided by the total number of samples predicted to be in the class; and the classification recall is defined as the number of correctly labeled samples divided by the size of the class. Again, we took the pre-

cision and recall of the EEGS-based methods to be the average over six values, and similarly for the non-EEGS-based methods. All six methods were run on both the full and the reduced datasets and the precision and recall on the CAR dataset are plotted in Figure 9, and on the LUNG dataset are plotted in Figure 10, in which the x -axis represents the class.

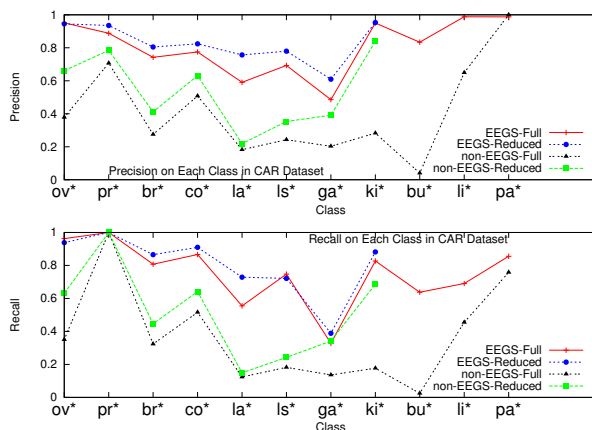


Fig. 9. The classification precision and recall on the classes in the full and reduced CAR dataset. Class label abbreviations, where the numbers are the class sizes: ov*: ovary (27), pr*: prostate (26), br*: breast (26), co*: colorectal (23), la*: lung adenocarcinomas (14), ls*: lung squamous cell carcinoma (14), ga*: gastroesophagus (12), ki*: kidney (11), bu*: bladder/ureter (8), li*: liver (7), pa*: pancreas (6).

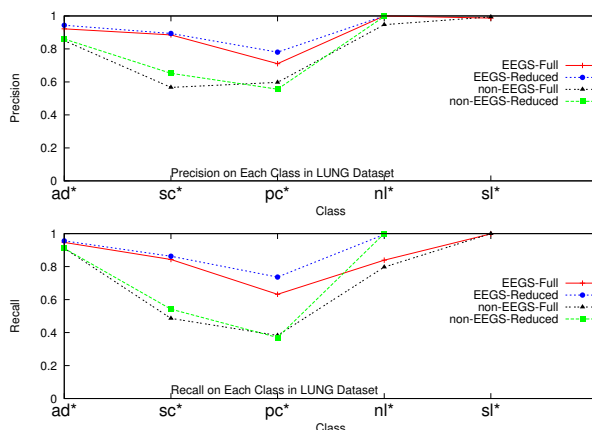


Fig. 10. The classification precision and recall on the classes in the full and reduced LUNG dataset. Class label abbreviations, where the numbers are the class sizes: ad*: adenocarcinomas (139), sc*: squamous cell lung carcinomas (21), pu*: pulmonary carcinoids (20), nl*: normal lung (17), sl*: small-cell lung carcinomas (6).

Surprisingly, the performance of either of the EEGS-based methods or the non-EEGS-based methods was not directly related to the class size, for example, both the precision and recall were high on the smallest class in each dataset. On the other hand, both precision and recall for the non-EEGS-based methods were low on the class *bladder/ureter* in the CAR dataset, while those of the EEGS-based methods were significantly higher. This demonstrates the extreme class prediction power of the genes selected by the EEGS-based methods, besides the fact that the EEGS-based methods performed consistently better than the non-EEGS-based methods on every class.

4.3. A Case Study on the CAR Dataset

The EEGS-based gene selection methods are designed to not select too many genes having similar class discrimination strength, so as to consider individually-less but collectively-more differentially expressed genes. In this sense, some selected genes might be lower in the gene order, but have strength to discriminate some classes on which its preceding genes might not do well. To examine if this indeed happened, we collected more detailed results on the first 10 genes selected by F-test and EEGS-F-test, respectively. The 5-Fold cross validation classification accuracies of the SVM classifier built on the first x genes were also collected, for $x = 1, 2, \dots, 10$. We summarized the results in Table 2, in which column ‘Probe Set’ records the probe set (gene) id in the CAR dataset, column ‘R’ records the rank of the gene in the gene order by F-test, and column ‘Accuracy’ records the classification accuracy of the gene subset up to the gene at the row.

Note that the third gene (probe set) selected by EEGS-F-test, *765_s_at*, has a rank 17, which was thus not selected by F-test. The classification accuracy of the top 10 genes selected by F-test was only 30.63%, while adding the third gene *765_s_at* in EEGS-F-test lifted the classification accuracy to 42.18%, already significantly higher than 30.63%. On average, the contribution of each gene, except the first, selected by EEGS-F-test was 6.10% in terms of classification accuracy; the contribution of each gene, except the first, selected by F-test was only 1.23%. These figures suggested that when the number of selected

genes was fixed, the genes selected by EEGS-F-test had much higher class discrimination strength compared to the genes selected by F-test.

Table 2. The first 10 genes selected by the EEGS-F-test and the F-test methods on the CAR dataset, respectively, and the respective 5-Fold cross validation classification accuracies of the SVM classifiers built on the genes. Column 'R' records the rank of the gene in the gene order by F-test, and column 'Accuracy' records the classification accuracy of the gene subset up to the gene at the row.

EEGS-F-test-SVM			F-test-SVM		
Probe Set	R	Accuracy	Probe Set	R	Accuracy
40794_at	1	19.54%	40794_at	1	19.54%
41238_s_at	4	28.91%	660_at	2	25.46%
765_s_at	17	42.18%	32200_at	3	25.40%
1500_at	21	60.52%	41238_s_at	4	30.00%
35220_at	24	64.54%	34941_at	5	30.69%
32771_at	27	68.62%	41468_at	6	30.52%
34797_at	38	70.75%	36141_at	7	30.12%
35194_at	50	73.56%	617_at	8	30.35%
36806_at	52	73.16%	37812_at	9	30.29%
40511_at	63	74.48%	217_at	10	30.63%

5. Conclusions and Future Works

We have proposed to differentiate the class discrimination strength between two genes by defining a novel distance between their discrimination strength vectors, based on their expression values. With such a distance measurement, we clustered genes and only allowed a limited number of genes per cluster to be selected for class prediction purpose. Through combining this gene cluster information with a single gene scoring method, we were able to efficiently select a subset of genes that collectively, improve the determination of class membership. The experimental results on two real microarray datasets confirmed our expectations. These results also demonstrated that the EEGS-based gene selection methods were more stable and had significantly higher classification precision and recall on smaller classes, where the non-EEGS-based methods normally performed poorly. The EEGS-based methods performed particularly well when there are many classes in the dataset. One immediate way to extend this work is to closely examine whether or not the selected genes by the EEGS-based methods, but not by the non-EEGS-based methods, are biologically meaningful.

Acknowledgments

This research is supported in part by AHFMR (to LX), AICML (to RG), CFI (to GL), iCore (to YS

and RG), NSERC (to RG, MS and GL) and the University of Alberta (to MS).

References

1. T. R. Golub, D. K. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, J. P. Mesirov, H. Coller, M. L. Loh, J. R. Downing, M. A. Caligiuri, C. D. Bloomfield, and E. S. Lander. Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring. *Science*, 286:531–537, 1999.
2. P. Baldi and A. D. Long. A Bayesian framework for the analysis of microarray expression data: Regularized t-test and statistical inferences of gene changes. *Bioinformatics*, 17:509–519, 2001.
3. S. Dudoit, J. Fridlyand, and T. P. Speed. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American Statistical Association*, 97:77–87, 2002.
4. J. Cho, D. Lee, J. H. Park, and I. B. Lee. New gene selection for classification of cancer subtype considering within-class variation. *FEBS Letters*, 551:3–7, 2003.
5. K. Yang, Z. Cai, J. Li, and G.-H. Lin. A stable gene selection in microarray data analysis. *BMC Bioinformatics*, 7:228, 2006.
6. M. Xiong, X. Fang, and J. Zhao. Biomarker identification by feature wrappers. *Genome Research*, 11:1878–1887, 2001.
7. I. Guyon, J. Weston, S. Barnhill, and V. Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46:389–422, 2002.
8. F. Blanchot-Jossic, A. Jarry, D. Masson, K. Bach-Ngohou, J. Paineau, M. G. Denis, C. L. Laboisse, and J. F. Mosnier. Up-regulated expression of ADAM17 in human colon carcinoma: co-expression with EGFR in neoplastic and endothelial cells. *Oncogene*, 207:156–163, 2005.
9. A. Bhattacharjee, W. G. Richards, J. Staunton, C. Li, S. Monti, P. Vasa, C. Ladd, J. Beheshti, R. Bueno, M. Gillette, M. Loda, G. Weber, E. J. Mark, E. S. Lander, W. Wong, B. E. Johnson, T. R. Golub, D. J. Sugarbaker, and M. Meyerson. Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses. *Proceedings of National Academy of Sciences of the United States of America*, 98:13790–13795, 2001.
10. A. I. Su, J. B. Welsh, L. M. Sapinoso, S. G. Kern, P. Dimitrov, H. Lapp, P. G. Schultz, S. M. Powell, C. A. Moskaluk, H. F. Frierson, Jr., and G. M. Hampton. Molecular classification of human carcinomas by use of gene expression signatures. *Cancer Research*, 61:7388–7393, 2001.